Electronic Colloquium on Computational Complexity, Report No. 329 (2026)

A PROOF OF THE ACHLIOPTAS–FRIEDGUT CONVERGENCE CONJECTURE

GAIA CARENINI

Abstract. For every fixed integer $q \geq 3$, we show that the q-colourability transition in a sparse random graph takes place at a limiting edge density, thus proving the Achlioptas-Friedgut convergence conjecture. The proof adapts the argument we recently introduced to settle the satisfiability conjecture.

1. Introduction

A proper (vertex) q-colouring of a graph is an assignment of colours from the set $[q] = \{1, \dots, q\}$ to its vertices such that adjacent vertices receive different colours. For a fixed integer $q \geq 3$, a central question in probabilistic graph theory is to determine when a sparse random graph admits such a colouring. Denote by $G(n, p)$ the random graph on n vertices in which each edge is added independently with probability p. As the average degree d increases, $G(n, d/n)$ passes from being q-colourable with high probability to being non-q-colourable with high probability. One way to approach this question is through the chromatic number $\chi(G)$, that is the least number of colours needed for a proper colouring. The first concentration results showed that this parameter fluctuates remarkably little. Shamir and Spencer [20] established concentration on the scale $\sqrt{n}$ for general p, and on a bounded number of values in sufficiently sparse graphs. Subsequent work of Łuczak [17] and Alon and Krivelevich [3] reduced this to two consecutive values, the latter throughout the range $p = n^{-\alpha}$ with fixed $\alpha > 1/2$. For graphs of bounded average degree, the problem was therefore to locate these two values and to decide when one of them occurs with high probability. Achlioptas and Naor [4] made this concentration explicit. Let k_d be the smallest integer k such that $d < 2k \log k$, they proved that $\mathbb{P}(\chi(G(n, d/n)) \in \{k_d, k_d + 1\}) \longrightarrow 1$. Their argument also determines the chromatic number exactly when $d \in [(2k-1)\log k, 2k \log k)$; it is then $k+1$ with high probability. These intervals form a set of asymptotic density $1/2$ in the positive real line. The underlying colourability bound is obtained by applying the second moment method to the number of balanced colourings. In particular, the correlations between colourings become central to locating the transition.

Statistical physics offers a broader description of these correlations. The cavity method studies how the organisation of the space of proper colourings changes as edges are added [18, 16, 15, 21]. In the predicted clustered phase, colourings are organised into many groups, with colourings in different groups separated by changes on a large number of vertices. At the condensation transition, a small collection of large clusters carries a substantial proportion of all colourings. This is a change in the distribution of solutions while proper colourings still exist; the colourability transition occurs only when the solution space becomes empty. The distinction matters for counting arguments, since pairs of colourings from the same cluster can contribute disproportionately to a second moment. The physics predictions , therefore, suggest both a location for the colorability threshold and an explanation of the difficulties encountered in proving it.

Recent advancements have increasingly exploited the structure of the solution space highlighted in the aforementioned studies. The bound of Achlioptas and Naor gives colourability for $d < 2(q-1)\log(q-1)$, whereas the first moment method gives non-colourability for $d > 2q \log q - \log q$. Coja-Oghlan and Vilenchik [11] refined the second moment argument to prove a lower bound of $2q \log q - \log q - 2 \log 2 + o_q(1)$. This is the asymptotic location of condensation, whose existence

Date: November 12, 2026.

ISSN 1433-8092

and predicted location were established for all sufficiently large q by Bapst, Coja-Oghlan, Hetterich, Raßmann and Vilenchik [6]. On the other side, Coja-Oghlan [10] obtained the upper bound $2q\log q - \log q - 1 + o_q(1)$, agreeing with the cavity prediction for the colourability threshold to within $o_q(1)$. The resulting bounds leave an interval of width $2\log 2 - 1 + o_q(1)$ for the transition as $q \to \infty$. They also allow Coja-Oghlan and Vilenchik [11] to determine the chromatic number for every fixed average degree in an explicit set of asymptotic density one, extending the earlier result of Achlioptas and Naor.

These advances locate the transition very precisely for large q. A separate question concerns its behaviour as n grows with q fixed: does the transition approach a single limiting density? A fundamental result of Achlioptas and Friedgut [2] establishes that the transition is sharp. For each fixed $q \geq 3$, there exists a sequence $d_q(n)$ such that, for every fixed $\varepsilon > 0$, $G(n, (d_q(n) - \varepsilon)/n)$ is q-colourable with high probability, whereas $G(n, (d_q(n) + \varepsilon)/n)$ is not. Thus, the transition occurs in an interval whose width tends to zero on the scale of average degree. Its location could nevertheless "oscillate" as n grows. The colourability threshold conjecture, also known as the Achlioptas–Friedgut convergence conjecture, asserts that this oscillation does not occur.

Conjecture 1.1 (Achlioptas–Friedgut convergence conjecture)**.** *For every integer $q \geq 3$, there exists a constant $d_q \in (0, \infty)$ such that, for every fixed $d \geq 0$ with $d \neq d_q$,*

$$\lim_{n\to\infty} \mathbb{P}\big(G(n, d/n) \text{ is } q\text{-colourable}\big) = \begin{cases} 1, & d < d_q, \\ 0, & d > d_q. \end{cases}$$

Proving Conjecture 1.1 faces the same difficulty arose in proving the satisfiability conjecture, asserting that, for every fixed $k \geq 3$, there is a limiting clause density separating satisfiability from unsatisfiability of k-CNFs with high probability. Friedgut [14] proved the existence of a sharp threshold sequence, leaving open whether its location converges. Ding, Sly and Sun [13] proved the conjecture for all sufficiently large k and identified the limiting threshold. Recently, in [9], we resolved the conjecture for every fixed $k \geq 3$ by proving a polynomial bound on the width of the satisfiability transition and combining it with an interpolation criterion of Abbe and Montanari [1].

In this paper, we adapt the technique introduced in [9] to prove Conjecture 1.1.

Theorem 1.2. *Achlioptas–Friedgut convergence conjecture (Conjecture 1.1) holds.*

1.1. **Proof overview.** In what follows, we fix $q \geq 3$ and call a graph *colourable* if it admits a proper q-colouring. We write $\alpha_q = d_q/2$ for the corresponding number of edges per vertex. We work with an edge process in which pairs are sampled independently, allowing repetitions. Let $e_1, e_2, \ldots$ be independent uniform unordered pairs of distinct vertices in $[n]$, and let G_m consist of the first m pairs, with multiplicities. We let $H_n = \min\{m : G_m \text{ is not colourable}\}$ and $a_n = \mathbb{E}H_n$. Repeated edges do not affect colourability. Indeed, among $O(n)$ samples their number is $o(n)$ in probability, while the edge count in $G(n, d/n)$ is concentrated around $dn/2$. A coupling therefore transfers a limiting threshold for this process to the independent-edge model.

We explain the successive reductions in our proof of the satisfiability conjecture [9], and how each is adapted to colouring. The common strategy is to pass from the existence of a limiting threshold to a scaling window estimate, then to a variance bound, and finally to a bound on the remaining lifetime of the last instance that still admits a solution.

1.1.1. *From convergence to a scaling window bound.* The first step of the SAT proof uses the criterion of Abbe and Montanari [1], saying that a scaling window of width $O(n/(\log n)^{1+\delta})$ in the number of clauses, for any fixed $\delta > 0$, suffices to establish a limiting threshold (see also [19, Section 4.5]). Our polynomial window bound in [9] supplies the required concentration. For colouring, we establish the corresponding reduction in Section 2, using an interpolation technique due to Bayati, Gamarnik and Tetali [7].

The concentration of H_n describes the width of the colourability transition in the number of added edges, while convergence of a_n/n identifies its limiting location. We show that a sufficiently

narrow window again rules out oscillation once it is combined with interpolation; in particular, we prove that:

$$\mathbb{P}\left(|H_n - a_n| > \frac{n}{(\log n)^2}\right) \longrightarrow 0 \quad \Longrightarrow \quad \frac{a_n}{n} \text{ converges.}$$

Therefore, the problem has been reduced to showing the following.

Theorem 1.3 (Scaling window)**.** *For every fixed integer $q \geq 3$,*

$$\mathbb{P}\left(|H_n - a_n| > \frac{n}{(\log n)^2}\right) \longrightarrow 0. \tag{1.1}$$

1.1.2. *From a scaling window to remaining lifetime.* In the study of k-SAT, we added independent uniform clauses one at a time, and let T_n be the first time the resulting formula is unsatisfiable. The identity $\mathbb{P}(T_n > m) = \mathbb{P}(F_{n,m} \text{ is satisfiable})$ expressed the satisfiability transition as the distribution of a single random variable. Chebyshev's inequality therefore reduced a scaling window bound to an upper bound on $\mathrm{Var}(T_n)$. We apply the same reduction for colouring. The process defined above adds independent uniform pairs, and H_n is its first uncolourable time. Since $\mathbb{P}(H_n > m) = \mathbb{P}(G_m \text{ is colourable})$, Chebyshev gives

$$\mathbb{P}\left(|H_n - a_n| > \frac{n}{(\log n)^2}\right) \leq \frac{(\log n)^4}{n^2}\,\mathrm{Var}(H_n).$$

It is consequently enough to prove $\mathrm{Var}(H_n) = o(n^2/(\log n)^4)$.

The next step of the SAT proof regarded the clause process as a covering process on the Boolean assignments: a clause covers the assignments that violate it, and T_n is the time at which every assignment has been covered. Aldous's covering-time inequality [5] bounds the variance of this time in terms of the expected remaining lifetime of the last satisfiable formula, under an independent continuation by fresh clauses. For colouring, the assignment space is $[q]^{[n]}$, and an edge uv covers the maps σ for which $\sigma(u) = \sigma(v)$. The uncovered maps are exactly the proper colourings, so the same inequality applies. Let $\Gamma_n = G_{H_n-1}$ be the last colourable graph, and let $r(G)$ be the expected number of fresh independent uniform pairs needed to destroy the colourability of G, with $r(G) = 0$ when G is already uncolourable. Aldous's inequality gives

$$\mathrm{Var}(H_n) \leq C a_n \mathbb{E} r(\Gamma_n) \leq C n \mathbb{E} r(\Gamma_n),$$

where the second inequality uses the elementary first moment bound $a_n = O(n)$. The independent continuation is important. Indeed, the observed terminal edge destroys Γ_n, but it has been selected by that property, whereas a fresh edge need not do so. We prove that, for every fixed $B > 0$,

$$\mathbb{E} r(\Gamma_n) = O_{q,B}\left(\frac{n}{(\log n)^B}\right).$$

Taking $B > 4$ gives the variance estimate required above.

1.1.3. *From remaining lifetime to forced statements.* In SAT, the remaining-lifetime problem leads to the study of *forced* variables, i.e., variables that have the same value in every satisfying assignment. A new clause destroys satisfiability precisely when all its literals are forced to be false. If a satisfiable formula F has $\ell(F)$ forced variables, then a uniform clause on k distinct variables with independent uniform signs is fatal with probability $h_k(F) = \binom{\ell(F)}{k}/2^k\binom{n}{k}$. In the setting of SAT, the proof estimates the occurrence of these forced variables along the process, using deletion of a variable and the two possible values with which it can be restored. These estimates control the remaining lifetime of the last satisfiable formula.

For colouring, permuting the colours shows that no vertex can have an absolute forced colour. The analogous forced statement is a relation between two vertices. Culberson and Gent [12] introduced *frozen pairs* in this global sense and related them to SAT backbones. We write $u \sim_G v$ if every proper colouring of G gives u and v the same colour, and denote the equivalence class of u

by $C_G(u)$. A fresh edge is fatal exactly when its endpoints form a frozen pair. Its failure probability is therefore

$$\kappa(G) = \frac{\sum_u (|C_G(u)| - 1)}{n(n-1)}.$$

This choice of relation is dictated by the failure event: to identify the forced statement relevant to a new constraint, we ask exactly which surviving solutions that constraint excludes. Both persist as constraints are added while a solution remains. In particular,

$$r(G)\kappa(G) \le 1, \qquad \sum_{m \ge 0} \mathbb{E}\kappa(G_m) = 1,$$

where κ is set to zero on uncolourable graphs. The first inequality follows by waiting for an edge inside a current frozen class. In the second identity, the summand is the probability that step $m+1$ is the first failure. These observations connect frozen pairs along the process with the remaining lifetime at its terminal state.

1.1.4. *Estimating frozen pairs.* To obtain the quantitative bound, we use two deletion operations. First, deleting an edge that creates a frozen equality reveals a choice that was forced before the edge was added. If adding ua forces $u \sim v$, then every previous colouring satisfies $\sigma(u) = \sigma(a)$ or $\sigma(u) = \sigma(v)$. Adding an edge that excludes one alternative resolves the choice and enlarges a frozen class. We measure this growth through $\Phi_u(G) = \log |C_G(u)|$. The logarithm accounts for the fact that any vertex in the current class can impose the same restriction. Its total possible logarithmic growth is at most $\log n$. A drift estimate for this potential bounds the time for which unresolved choices can persist, and hence controls the degree of an endpoint of the terminal edge.

Second, we adapt the variable-deletion step of the proof for satisfiability. Deleting a Boolean variable leaves two possible extensions of each assignment of the remaining variables. Deleting a vertex u leaves the possible colours $[q] \setminus \sigma(N_G(u))$ for each colouring σ of $G - u$. If u has a frozen partner v, every extendable colouring leaves exactly one available colour, namely $\sigma(v)$. In particular, no colouring of $G - u$ can make all neighbours of u monochromatic since this would leave at least $q - 1 \ge 2$ colours for u. Identifying its neighbours must therefore prevent colourability. Poissonisation and a comparison of the resulting incident-edge distributions turn this obstruction into a remaining-lifetime estimate with a weight depending on the degree of u.

More precisely, let U_n be a uniformly chosen endpoint of the terminal edge and let $\Delta_n = \deg_{\Gamma_n}(U_n)$. The two arguments give, for suitable constants $A, \Lambda, C > 0$,

$$\mathbb{E}[2^{\Delta_n}; H_n \le \lfloor \Lambda n \rfloor] \le C \log n, \qquad \mathbb{E}[r(\Gamma_n) e^{-A\Delta_n^3}] \le C n^{5/6}.$$

The first estimate controls the cost appearing in the second. A change of probability measure, followed by Jensen's and Hölder's inequalities, removes the weight and gives the required bound on $\mathbb{E}r(\Gamma_n)$.

2. From a narrow window to a limiting threshold

We show that a sufficiently narrow scaling window implies convergence of the threshold location, using the Poisson interpolation method of Bayati, Gamarnik and Tetali [7].

Proposition 2.1. *Suppose that*

$$\mathbb{P}\left(|H_n - a_n| > \frac{n}{(\log n)^2} \right) \longrightarrow 0. \tag{2.1}$$

Then a_n/n converges to a constant $\alpha_q \in (0, \infty)$, $H_n/n \to \alpha_q$ in probability, and Theorem 1.2 holds with $d_q = 2\alpha_q$.

For a colourable graph, call two vertices equivalent if they have the same colour in every proper colouring. These equivalence classes are its *frozen classes*. An added edge destroys colourability exactly when its endpoints belong to the same class. This is the only fact about frozen classes needed for the interpolation.

Let $\mathcal{A}_n(c)$ consist of a Pois(cn) number of independent uniform *ordered* pairs from $[n]^2$, regarded as edges. Let $\rho_n(c) = \mathbb{P}(\mathcal{A}_n(c) \text{ is colourable})$. The number of loops is Pois(c) and is independent of all non-loop pairs. Conditional on their number, the non-loop pairs are uniform unordered distinct-vertex pairs. Thus, for $n \geq 2$,

$$\begin{cases} \rho_n(c) = e^{-c}\mathbb{P}(H_n > \Pi_{c(n-1)}) \\ \rho_1(c) = e^{-c} \end{cases} \tag{2.2}$$

where the Poisson variable is independent of the edge process.

Lemma 2.2. *For positive integers n_1, n_2 and every $c \geq 0$, $\rho_{n_1+n_2}(c) \geq \rho_{n_1}(c)\rho_{n_2}(c)$.*

Proof. Let $N = n_1 + n_2$ and partition $[N]$ into parts V_1, V_2 of the indicated sizes. For $0 \leq t \leq 1$, take independently a Pois(tcN) number of uniform ordered pairs on $[N]^2$ and, for each i, a Pois$((1-t)cn_i)$ number on V_i^2. Let $f(t)$ be the probability that their union is colourable.

Differentiating with respect to the Poisson means expresses $f'(t)$ in terms of the effect of adding one pair. In a colourable graph, let its frozen classes have a_j vertices in V_1 and b_j in V_2. An added ordered pair destroys colourability exactly when its endpoints lie in one frozen class, including the case of a loop. Therefore

$$\begin{aligned} f'(t) &= c\mathbb{E}\left[\mathbf{1}_{\{G \text{ colourable}\}} \sum_j \left(\frac{a_j^2}{n_1} + \frac{b_j^2}{n_2} - \frac{(a_j+b_j)^2}{N}\right)\right] \\ &= c\mathbb{E}\left[\mathbf{1}_{\{G \text{ colourable}\}} \sum_j \frac{(n_2 a_j - n_1 b_j)^2}{n_1 n_2 N}\right] \geq 0. \end{aligned}$$

An uncolourable graph contributes zero. At $t = 0$, the two parts are independent, so $f(0) = \rho_{n_1}(c)\rho_{n_2}(c)$, while at $t = 1$, $f(1) = \rho_N(c)$. Continuity at the endpoints and integration conclude the proof. □

Let $x_n = a_n/n$ and $h_n = 4/(\log n)^2$ for $n \geq 3$. By (3.1) and (3.2), there exist constants $0 < a < A$ with $a \leq x_n \leq A$ for all sufficiently large n. Fix a value $L > A + 1$ and choose $0 < \eta < e^{-2L}/16$. The assumed window estimate (2.1) implies, uniformly for $0 \leq c \leq L$ and all sufficiently large n,

$$\begin{cases} c \leq x_n - h_n & \Longrightarrow \quad \rho_n(c) \geq e^{-c}(1-\eta) \\ c \geq x_n + h_n & \Longrightarrow \quad \rho_n(c) \leq e^{-c}\eta \leq \eta. \end{cases}$$

Let $w_n = n/(\log n)^2$ and $\mu = c(n-1)$. Chebyshev gives

$$\mathbb{P}(|\Pi_\mu - \mu| > w_n) \leq \frac{Ln}{w_n^2} = \frac{L(\log n)^4}{n} = o(1),$$

uniformly in c. If $c \leq x_n - h_n$, then $\mu \leq a_n - 4w_n$, and on the two concentration events $\Pi_\mu \leq a_n - 3w_n < H_n$. If $c \geq x_n + h_n$, then $\mu \geq a_n + 4w_n - L \geq a_n + 3w_n$ for large n, and on those events $\Pi_\mu \geq a_n + 2w_n > H_n$. Their complementary probabilities tend to zero by (2.1). Equation (2.2) proves the claims.

For n, m sufficiently large, let $\ell = \min(n, m)$ and $c = \min(x_n, x_m) - h_\ell$. Taking the lower bound on n, m large enough ensures $0 < c < L$. Both factors in the inequality of Lemma 2.2 are at least $e^{-L}(1-\eta)$, so $\rho_{n+m}(c) \geq e^{-2L}(1-\eta)^2 > \eta$. The fact that $c \geq x_n + h_n$ implies $\rho_n(c) \leq e^{-c}\eta \leq \eta$ forces $x_{n+m} > c - h_{n+m}$. Since $h_{n+m} \leq h_\ell$,

$$x_{n+m} \geq \min(x_n, x_m) - \frac{8}{(\log \min(n,m))^2}. \tag{2.3}$$

The neighbouring-size comparison follows in the same way, using the one-vertex model explicitly. Let $c = x_n - h_n$. Then, $\rho_{n+1}(c) \geq \rho_n(c)\rho_1(c) \geq e^{-2L}(1-\eta) > \eta$, and hence

$$x_{n+1} \geq x_n - \frac{8}{(\log n)^2} \tag{2.4}$$

for all sufficiently large n.

The following deterministic lemma completes the convergence argument.

Lemma 2.3. *Let (x_n) be a bounded real sequence, and consider $g(t) = C/(\log t)^2$. Suppose that for some C and all sufficiently large integers n, m,*

$$\begin{cases} x_{n+m} \geq \min(x_n, x_m) - g(\min(n, m)) \\ x_{n+1} \geq x_n - g(n) \end{cases}$$

Then (x_n) converges.

Proof. Fix a sufficiently large integer k. For every $N \geq k^2$, write $N = tk + r$, with $0 \leq r < k \leq t$. Thus, N is the sum of t integers, each equal to k or $k+1$. We take r of the latter and $t - r$ of the former. Arrange these summands as the leaves of a binary tree, splitting the list of leaves as evenly as possible at each vertex and assigning to each internal vertex the sum of its descendant leaves. Each child of an internal vertex has at least one quarter, and at most three quarters, of its parent's total size. Indeed, for $s \geq 2$ leaves the smaller child has at least $\lfloor s/2 \rfloor$ leaves. In particular, its size relative to the parent is at least

$$\frac{\lfloor s/2 \rfloor k}{s(k+1)} \geq \frac{k}{3(k+1)} \geq \frac{1}{4},$$

whenever $k \geq 3$. Both children satisfy this lower bound, giving the upper bound as well.

Apply the first assumed inequality at each internal vertex. At any leaf, the value is at least $x_k - g(k)$, by the second inequality. Along a path from a leaf towards the root, the ancestor at distance j has size at least $k(4/3)^j$. We have that the smaller of its two children therefore has size at least $(k/4)(4/3)^j$. Since g decreases,

$$x_N \geq x_k - g(k) - \sum_{j \geq 1} g\left(\frac{k}{4}\left(\frac{4}{3}\right)^j\right).$$

All sizes are in the range of the assumptions if k is large enough. The sum is bounded by

$$C \sum_{j \geq 1} \frac{1}{(\log(k/4) + j \log(4/3))^2} \leq \frac{C}{\log(4/3)\log(k/4)}.$$

Thus,

$$\liminf_{N \to \infty} x_N \geq x_k - O(1/\log k).$$

Let k tend to infinity along a subsequence realising $\limsup x_k$. This gives $\liminf x_n \geq \limsup x_n$. □

We apply Lemma 2.3 to (2.3)–(2.4), and prove that $a_n/n \to \alpha_q$, with $\alpha_q \geq c_0 > 0$ by (3.2) and $\alpha_q < \infty$ by (3.1). Together with the hypothesis (2.1), this gives $H_n/n \to \alpha_q$ in probability.

To complete the proof of Proposition 2.1, we read the pairs in order, retaining only the first occurrence of each pair. The resulting sequence is a uniform ordering of all $\binom{n}{2}$ pairs. Let D_m be the number of distinct pairs among the first m occurrences. Then, for $m = O(n)$, we have that

$$\mathbb{E}(m - D_m) \leq \frac{\binom{m}{2}}{\binom{n}{2}} = O(1),$$

and thus $m - D_m = o_{\mathbb{P}}(n)$.

Independently of this ordering, take $B \sim \mathrm{Bin}(\binom{n}{2}, d/n)$, and retain its first B pairs. The resulting graph has distribution $G(n, d/n)$, since conditional on B its edge set is uniform among sets of size B. Moreover, $B/n \to d/2$ in probability, because $\mathbb{E}B = d(n-1)/2$ and $\mathrm{Var}(B) \leq dn/2$. For fixed $d > 0$ and $0 < \varepsilon < d/2$, the preceding estimates therefore give, with probability tending to one,

$$G_{\lfloor (d/2-\varepsilon)n \rfloor} \subseteq G(n, d/n) \subseteq G_{\lfloor (d/2+\varepsilon)n \rfloor}.$$

If $d \neq 2\alpha_q$, choose ε so that both bounding edge densities lie on the same side of α_q. Monotonicity and $H_n/n \to \alpha_q$ then prove Theorem 1.2 with $d_q = 2\alpha_q$. The case $d = 0$ is immediate. This completes the proof of Proposition 2.1.

3. The edge process and the remaining-lifetime reduction

We now recall a few elementary estimates for the edge process, describe its frozen pairs, and reduce the window bound to an estimate on $r(\Gamma_n)$. We begin with uniform linear bounds on the expected lifetime. These provide the boundedness and positivity used in Section 2, as well as the tail estimate needed below. For the upper bound, each fixed assignment is destroyed with a probability bounded away from zero at every step. For the lower bound, a sufficiently sparse graph has no subgraph of minimum degree three and hence is 3-colourable.

For $n \geq 2q$ and any map $\sigma : [n] \to [q]$, let $n_i = |\sigma^{-1}(i)|$. The probability that a uniform pair is monochromatic is

$$\frac{\sum_i n_i(n_i-1)}{n(n-1)} \geq \frac{n-q}{q(n-1)} \geq \frac{1}{2q}.$$

A union bound over the at most q^n maps gives $\mathbb{P}(H_n > m) \leq q^n e^{-m/(2q)}$. The same bound applies to the survival time starting from any fixed colourable G, since adding its existing edges can only reduce the set of colourings. Summing the bound after a sufficiently large constant times n proves

$$r(G) \leq Cn \quad \text{, and} \quad a_n \leq Cn. \tag{3.1}$$

It also shows that H_n has all finite moments. Fix henceforth a constant $\Lambda > 1$, depending only on q, sufficiently large that $\mathbb{P}(H_n > \lfloor \Lambda n \rfloor) \leq e^{-bn}$ for some $b > 0$ and all sufficiently large n.

We also need that a_n is bounded below by a positive multiple of n. This requires only a "rough" estimate. Let $c_0 = 10^{-3}$ and $m = \lfloor c_0 n \rfloor$. If a simple graph is not 3-colourable, it has a nonempty subgraph of minimum degree at least 3. Thus, such a graph has a vertex set of size $v \geq 4$ spanning at least $s = \lceil 3v/2 \rceil$ distinct edges.

For a specified set of s pairs, the probability that all occur among $e_1, \ldots, e_m$ is at most $(m/\binom{n}{2})^s$, by assigning a distinct occurrence to each specified pair and taking a union bound. The probability that some v-set spans s distinct pairs is therefore at most

$$\binom{n}{v}\binom{\binom{v}{2}}{s}\left(\frac{m}{\binom{n}{2}}\right)^s \leq \left(\frac{en}{v}\right)^v \left(\frac{4ec_0 v}{3n}\right)^s \leq \left(K\sqrt{\frac{v}{n}}\right)^v,$$

where $K = e\left(\frac{4ec_0}{3}\right)^{3/2} < \frac{1}{2}$.

If $s > \binom{v}{2}$, the probability is zero. The last inequality uses $4ec_0 v/(3n) < 1$. Summing for $4 \leq v \leq \sqrt{n}$ gives $O(n^{-1})$, and summing for $v > \sqrt{n}$ gives $O(2^{-\sqrt{n}})$. Hence

$$\mathbb{P}(H_n > \lfloor c_0 n \rfloor) \longrightarrow 1, \qquad \liminf_{n\to\infty} \frac{a_n}{n} \geq c_0. \tag{3.2}$$

For a colourable graph G, write $u \sim_G v$ if every proper colouring of G gives u and v the same colour. This is an equivalence relation. Let $C_G(u)$ be the class of u. A pair of distinct vertices in one class is called *frozen*. Adding a pair to a colourable graph destroys colourability exactly when that pair is frozen. Recall that $\kappa(G)$ denotes the fraction of frozen pairs, setting it to zero on uncolourable graphs. Frozen equalities persist as long as the graph remains colourable. Therefore, a fresh occurrence of any currently frozen pair destroys colourability, and $r(G)\kappa(G) \leq 1$. In particular, since $\kappa(G) \geq s(s-1)/(n(n-1)) \geq (s-1)^2/n^2$, for a class of size s, the following holds:

$$(s-1)\sqrt{r(G)} \leq n. \tag{3.3}$$

The last edge is selected by its success in preventing colourability. The next identity removes that selection by summing over all possible times and frozen pairs. For any nonnegative function $f(G, e)$, understood to be zero unless G is colourable and e is frozen, independence of the next pair gives

$$\mathbb{E}f(\Gamma_n, e_{H_n}) = \sum_{m\geq 0} \mathbb{E}\left[\frac{1}{\binom{n}{2}} \sum_e f(G_m, e)\right]. \tag{3.4}$$

In particular, $\sum_m \mathbb{E}\kappa(G_m) = 1$.

Let Π_t be a rate-one Poisson counting process, independent of the pairs, and let $\mathcal{G}_n(t) = G_{\Pi_t}$. By Poisson splitting, the pair multiplicities at time t are independent $\mathrm{Pois}(t/\binom{n}{2})$ variables. For every nonnegative function f,

$$\int_0^\infty \mathbb{E}f(\mathcal{G}_n(t))\,\mathrm{d}t = \sum_{m\geq 0} \mathbb{E}f(G_m). \tag{3.5}$$

If the process is run at rate θ, the integral is divided by θ. For any fixed assignment, the occurrences on its monochromatic pairs form a Poisson variable of mean at least $t/(2q)$. A union bound therefore gives the first-moment estimate $\mathbb{P}(\mathcal{G}_n(t) \text{ is colourable}) \leq q^n e^{-t/(2q)}$.

We use Aldous's covering-time inequality [5]. Let T be the time needed to cover a finite set S by independent copies of a random subset, with each point covered with positive probability. Let U be the last uncovered set, and let $c(U)$ be its expected covering time with fresh independent subsets. Aldous's inequality states that, for an absolute constant C_0, $\mathrm{Var}(T) \leq C_0(\mathbb{E}T)\mathbb{E}c(U)$.

Apply this to $S = [q]^{[n]}$, the set of all assignments. A pair $\{u, v\}$ covers the assignments with $\sigma(u) = \sigma(v)$. Every assignment has a monochromatic pair when $n > q$, so every point of S is covered with positive probability. The cover time is H_n, and the last uncovered set consists exactly of the proper colourings of Γ_n. Thus $c(U) = r(\Gamma_n)$, and (3.1) gives $\mathrm{Var}(H_n) \leq C_0 a_n \mathbb{E}r(\Gamma_n) \leq Cn\mathbb{E}r(\Gamma_n)$. Consequently, Chebyshev's inequality yields, for every $t > 0$,

$$\mathbb{P}(|H_n - a_n| > t) \leq \frac{Cn\mathbb{E}r(\Gamma_n)}{t^2}. \tag{3.6}$$

4. Forced choices and growth of frozen classes

To understand how a frozen equality is created, delete the edge that creates it. Suppose that $G + \{u, a\}$ is colourable and that u and v have the same colour in every colouring of this enlarged graph. A colouring of G either violates the deleted edge, in which case $\sigma(u) = \sigma(a)$, or survives its addition, in which case $\sigma(u) = \sigma(v)$. Conversely, this disjunction ensures the frozen equality after the edge is added. Thus, deleting a constraint leads directly to a forced choice between two equalities. We now keep trace of these choices so that we can count their resolutions.

Fix a colourable graph G and a root u. A nonempty set S is a *forced choice for* u if every proper colouring σ of G satisfies $\sigma(u) \in \{\sigma(v) : v \in S\}$. A singleton $\{v\}$ is a forced choice exactly when $v \in C_G(u)$. We shall only need choices of size two with neither alternative already in $C_G(u)$. Let $C = C_G(u)$, $s = |C|$, and $W = [n] \setminus C$. We keep trace of these two-element choices as the symmetric relation

$$R_G(u) = \{(a, b) \in W^2 : \sigma(u) \in \{\sigma(a), \sigma(b)\} \text{ for every proper colouring } \sigma \text{ of } G\}. \tag{4.1}$$

There are no diagonal pairs. Indeed, $(a, a) \in R_G(u)$ would force a in C. Let $R(a) = \{b : (a, b) \in R_G(u)\}$ and $\beta_G(u) = |R_G(u)|/n^2$, setting $\beta_G(u) = 0$ on uncolourable graphs. We count ordered pairs in $R_G(u)$, so each two-element choice is counted twice.

There are two ways an edge can resolve these choices. First, if $a \in W$, adding ua preserves colourability and gives exactly

$$C_{G+\{u,a\}}(u) = C \cup R(a). \tag{4.2}$$

To see this, a vertex $b \in W$ joins the class precisely when every colouring with $\sigma(u) \neq \sigma(a)$ has $\sigma(u) = \sigma(b)$. This is equivalent to the disjunction defining $(a, b) \in R_G(u)$. The same identity holds when the endpoint u in the added edge is replaced by any member of C.

Second, for distinct $a, b \in W$, adding ab, if colourability survives, forces all of $R(a) \cap R(b)$ into C. Indeed, if $x \in R(a) \cap R(b)$ had a colour different from u, the two forced choices would require both a and b to have the colour of u, contrary to the new edge. Thus, edges within W also contribute to the growth of the root's frozen class.

Resolving a two-element choice enlarges $C_G(u)$. If this class has size s, each of its s vertices can impose the same restriction, so we measure relative growth through $\log s$. An increase by d

contributes $\log(1 + d/s)$, and the total possible increase is at most $\log n$. Define

$$\Phi_u(G) = \begin{cases} \log |C_G(u)|, & G \text{ colourable}, \\ \log n, & G \text{ uncolourable}. \end{cases}$$

This potential lies in $[0, \log n]$ and never decreases when an edge is added. The following drift bound uses both ways of resolving choices.

Lemma 4.1. *For a fresh uniform pair e, $\beta_G(u) \le \mathbb{E}_e[\Phi_u(G + e) - \Phi_u(G)]$.*

Proof. The assertion is immediate when G is uncolourable. Otherwise use the notation above, let $d_a = |R(a)|$, and let $k_{ab} = |R(a) \cap R(b)|$. We count the uniform pair as an ordered pair of distinct vertices. An edge from C to $a \in W$ increases the potential by $\log(1+d_a/s)$, by (4.2). An edge $\{a, b\}$ within W increases it by at least $\log(1 + k_{ab}/s)$. This remains true if the new graph is uncolourable, since $s + k_{ab} \le n$ and its potential is then $\log n$. Using $\log(1 + x) \ge x/(1 + x)$ and discarding the other nonnegative increments, then including and subtracting the diagonal terms $k_{aa} = d_a$, gives

$$n(n-1)\mathbb{E}_e[\Phi_u(G + e) - \Phi_u(G)] \ge \sum_{a,b \in W} \frac{k_{ab}}{s + k_{ab}} + (2s - 1) \sum_{a \in W} \frac{d_a}{s + d_a}. \tag{4.3}$$

Since $k_{ab} \le d_a$, double-counting common neighbours in R shows

$$\sum_{a,b \in W} \frac{k_{ab}}{s + k_{ab}} \ge \sum_{a \in W} \frac{\sum_b k_{ab}}{s + d_a} = \sum_{(a,y) \in R_G(u)} \frac{d_y}{s + d_a}.$$

As $s \ge 1$, the right side of (4.3) is at least

$$\sum_{(a,y) \in R_G(u)} \frac{d_y + 2s - 1}{s + d_a} \ge \sum_{(a,y) \in R_G(u)} \frac{s + d_y}{s + d_a} \ge |R_G(u)|.$$

For the last inequality, pair (a, y) with (y, a) and use $x/x' + x'/x \ge 2$ for positive x, x'. Dividing by $n(n - 1)$, and using $n(n - 1) \le n^2$, proves the lemma. □

We now count frozen pairs according to the degree of their first endpoint. Deleting an incident occurrence either preserves the equality or reveals a two-element choice. Counting every possible deleted occurrence will give a factor equal to that degree. For $j \ge 0$, define the functions

$$\kappa_j(G) = \frac{1}{n(n-1)} \sum_{u \ne v} \mathbf{1}_{\{G \text{ colourable}, \ u \sim_G v, \deg_G(u) = j\}}.$$

Thus $\sum_j \kappa_j(G) = \kappa(G)$. All these functions are zero on uncolourable graphs.

Lemma 4.2. *For $m \ge 1$, $j \ge 1$, and any fixed root u,*

$$j\mathbb{E}\kappa_j(G_m) \le \frac{2m}{n}\mathbb{E}\kappa_{j-1}(G_{m-1}) + \frac{2mn}{(n-1)^2}\mathbb{E}\big[\mathbf{1}_{\{\deg_{G_{m-1}}(u) = j-1\}}\beta_{G_{m-1}}(u)\big]. \tag{4.4}$$

Proof. For each ordered frozen pair (u, v) counted by $\kappa_j(G_m)$, count all j occurrences incident with u. Sum this count over the m possible occurrence indices before taking expectations. For each fixed index, the other $m - 1$ occurrences form a graph F with law G_{m-1}, independently of the deleted uniform pair. Exchangeability makes the contributions of the indices equal. On a counted event, the root has degree $j - 1$ in F.

If $u \sim_F v$ already, discard the condition that adding the pair preserves colourability. The probability that the fresh pair is incident with u is $2/n$, giving the first term on the right. If the equality is new, write the added pair as $\{u, a\}$. The enlarged graph is colourable, so $a \notin C_F(u)$, and (4.2) says precisely that $(a, v) \in R_F(u)$. For a fixed F, u, summing over the uniform choice of $v \ne u$ and the fresh unordered pair gives

$$\frac{|R_F(u)|}{(n-1)\binom{n}{2}} = \frac{2n}{(n-1)^2}\beta_F(u).$$

Averaging over the root gives the same expectation as a fixed root, by vertex symmetry. Multiplying by m proves (4.4). Repeated pairs cause no difficulty because we counted and deleted occurrences, not distinct edges. □

The recurrence sums over time because the total drift of the potential is bounded. A generating function then keeps all degrees in the same calculation. Let $M = \lfloor \Lambda n \rfloor$ and define

$$b_j = \sum_{m=0}^{M-1} \mathbb{E}\kappa_j(G_m), \qquad d_j = \sum_{m=0}^{M-1} \mathbb{E}\big[\mathbf{1}_{\{\deg_{G_m}(u)=j\}}\beta_{G_m}(u)\big].$$

Here u is any fixed vertex. Since G_0 has no frozen distinct pair, summing (4.4) and enlarging the time sums on the right gives a constant K, depending only on Λ, such that

$$jb_j \le K(b_{j-1} + d_{j-1}) \quad (j \ge 1), \qquad \sum_{j\ge 0} b_j \le 1. \tag{4.5}$$

The last inequality is the terminal identity.

Let $\mathcal{B}(z) = \sum_j b_j z^j$ and $\mathcal{D}(z) = \sum_j d_j z^j$. These are polynomials, since no degree exceeds M. For $z \ge 1$, Lemma 4.1 gives

$$\begin{aligned}\mathcal{D}(z) &\le \mathbb{E}\sum_{m=0}^{M-1} z^{\deg_{G_m}(u)}\big(\Phi_u(G_{m+1}) - \Phi_u(G_m)\big)\\ &\le (\log n)\mathbb{E}z^{\deg_{G_M}(u)} \le (\log n)e^{2\Lambda(z-1)}.\end{aligned}$$

For the second inequality, degrees and the potential are increasing, so on each trajectory the weighted sum is at most $z^{\deg_{G_M}(u)}\log n$. For the last, $\deg_{G_M}(u) \sim \mathrm{Bin}(M, 2/n)$, whose generating function is $(1 + 2(z-1)/n)^M \le e^{2\Lambda(z-1)}$.

Multiplying (4.5) by z^{j-1} and summing gives $\mathcal{B}'(z) \le K\mathcal{B}(z) + K\mathcal{D}(z)$. Integrating this differential inequality from 1 to 2, with $\mathcal{B}(1) \le 1$, yields

$$\mathcal{B}(2) \le e^K + K\int_1^2 e^{K(2-z)}\mathcal{D}(z)\,\mathrm{d}z \le C\log n.$$

The terminal identity identifies b_j with $\mathbb{P}(H_n \le M, \Delta_n = j)$. We have therefore proved

$$\mathbb{E}\big[2^{\Delta_n}; H_n \le M\big] \le C\log n.$$

In particular, for every positive integer p, the inequality $x^{3p} \le (3p)!(\log 2)^{-3p}2^x$ gives

$$\mathbb{E}\big[\Delta_n^{3p}; H_n \le M\big] \le C\frac{(3p)!}{(\log 2)^{3p}}\log n. \tag{4.6}$$

5. Deleting a vertex and identifying its neighbours

We next bound the occupation of nontrivial frozen classes, with a weight depending on the vertex degree. We begin by examining which colourings extend after a deleted vertex is restored.

For a colouring σ of $G - u$, the available colours at u are $[q] \setminus \sigma(N_G(u))$. If u has a frozen partner $v \ne u$, every extendable colouring leaves exactly one available colour, namely $\sigma(v)$. Now suppose that some colouring of $G - u$ makes all neighbours of u monochromatic. There are then at least $q - 1 \ge 2$ available colours for u, so one can extend that colouring with a colour different from $\sigma(v)$, a contradiction. An empty neighbourhood likewise leaves all q colours available.

Identifying the neighbours in $G - u$ therefore destroys colourability. To bound the occupation of this event, we first estimate the effect of adding an ordinary vertex, then compare its star with that of an identified set.

A star that destroys colourability has a first fatal edge. For a Poisson number of star edges, its probability can be compared with the probability that the next edge is fatal. This converts failure of extension into the occupation of frozen pairs.

Lemma 5.1. *Let V have $v \geq 2$ vertices and let $o \notin V$. Fix $\lambda > 0$. At time t, let B have independent* $\mathrm{Pois}(\lambda t)$ *occurrences on each pair of V. We define a graph F as follows. Independently choose a sequence of uniform vertices of V, and add a star from o to its first J entries, where $J \sim \mathrm{Pois}(z)$ and $z = \lambda v t$. Then,*

$$\mathbb{P}(B \text{ is colourable, } F \text{ is not colourable}) \leq (e^z - 1)\mathbb{E}\kappa(F). \tag{5.1}$$

Moreover, as t varies with the indicated laws,

$$\int_0^\infty \mathbb{E}\kappa(F)\,\mathrm{d}t \leq \frac{1}{\lambda\binom{v+1}{2}}.$$

Proof. Let I be the first star length making the graph uncolourable, with $I = \infty$ if none does. On $\{B \text{ colourable}\}$, $I \geq 1$, and I is independent of J. For $i \geq 1$,

$$\frac{\mathbb{P}(J \geq i)}{\mathbb{P}(J = i-1)} = \sum_{\ell \geq 1} \frac{z^\ell (i-1)!}{(i-1+\ell)!} \leq e^z - 1$$

when $z > 0$; the desired inequality at $z = 0$ is immediate. Summing over $I = i$ bounds the left side of (5.1) by $(e^z - 1)\mathbb{P}(B \text{ colourable}, I = J + 1)$. The last probability is the chance that F is colourable and the next uniform star edge is frozen. By Poisson splitting, every pair of $V \cup \{o\}$ has independent $\mathrm{Pois}(\lambda t)$ multiplicity, so F is invariant under all vertex permutations. Averaging its root over the vertices shows that this chance equals $\mathbb{E}\kappa(F)$. Finally, F has the law of a uniform pair process run at rate $\lambda\binom{v+1}{2}$. In its discrete version, the expected fraction of frozen pairs at step m is the probability of first losing colourability at step $m+1$. These events are disjoint, so their probabilities sum to at most one, even if the graph never becomes uncolourable. Equation (3.5) concludes the proof of the lemma. □

For a vertex set Z, let G/Z be the multigraph obtained by identifying all vertices of Z into one vertex, keeping all pairs and hence allowing loops. A loop makes a graph uncolourable. Colourings of G/Z correspond exactly to colourings of G that are constant on Z.

Identifying D vertices multiplies the mean degree at the new vertex by D. A change of measure compares this larger star with an ordinary one. The following estimate records the resulting cost.

Lemma 5.2. *Let $1 \leq D \leq N/4$, let Z be a uniform D-subset of $[N]$ independent of the pair process, and define $A_N(D) = \sum_{m \geq 0} \mathbb{P}(G_m \text{ colourable, } G_m/Z \text{ uncolourable})$. For all sufficiently large N,*

$$A_N(D) \leq C(D^2 + 1 + N^{1/3} e^{CD^3}).$$

Proof. Use Poisson time and, by symmetry, fix Z. Let $\lambda = \binom{N}{2}^{-1}$ and $t_0 = 4q(\log q)N$. By the first-moment estimate, the part of the occupation integral after t_0 is at most $2q\,q^{-N}$.

Before t_0, the probability of any occurrence inside Z is at most $\lambda\binom{D}{2}t$. Its integral is $O(D^2)$. In the absence of such an occurrence, the quotient consists of the graph B on $V = [N] \setminus Z$ and a star at the identified vertex. Let $v = N - D$ and $z = \lambda v t$. The graph B has independent $\mathrm{Pois}(\lambda t)$ multiplicities, and the star has $\mathrm{Pois}(Dz)$ uniform independent endpoints in V, independently of B. If the original graph is colourable, then B is colourable.

Let $\delta_D(t)$ be the probability that B is colourable but this star makes it uncolourable, and define $\delta_1(t)$ with star mean z. Relative to $J \sim \mathrm{Pois}(z)$, the density of $\mathrm{Pois}(Dz)$ is

$$L_D(J) = e^{-(D-1)z} D^J, \qquad \mathbb{E}L_D(J)^3 = e^{z(D^3 - 3D + 2)}.$$

The formula is interpreted by continuity at $z = 0$. Hölder's inequality, applied to the indicator of the event defining δ_1, gives

$$\delta_D(t) \leq e^{z(D^3-3D+2)/3} \delta_1(t)^{2/3} \leq e^{CD^3} \delta_1(t)^{2/3} \qquad (0 \leq t \leq t_0),$$

because z is bounded by a constant depending only on q. Lemma 5.1 gives

$$\int_0^{t_0} \delta_1(t)\,\mathrm{d}t \leq \frac{e^{\lambda v t_0} - 1}{\lambda\binom{v+1}{2}} \leq C,$$

where $D \le N/4$ keeps the denominator bounded away from zero. A second application of Hölder yields

$$\int_0^{t_0} \delta_D(t)\,\mathrm{d}t \le e^{CD^3} t_0^{1/3} \left(\int_0^{t_0} \delta_1(t)\,\mathrm{d}t \right)^{2/3} \le CN^{1/3} e^{CD^3}.$$

Together with the internal-pair and large-time terms, this proves the lemma. □

5.1. **Occupation of vertices with a frozen partner.** For a colourable graph G on $[n]$, let $B_j(G)$ be the fraction of all vertices which have degree exactly j and belong to a nontrivial frozen class. Set $B_j(G) = 0$ if G is uncolourable. We first show that, for a constant C_1 depending only on q,

$$\sum_{m \ge 0} \mathbb{E} B_j(G_m) \le C(j^2 + 1 + n^{1/3} e^{C_1 j^3}) \qquad (j \ge 1). \tag{5.2}$$

There are no frozen isolated vertices, so $B_0 = 0$.

Fix a root u, and expose the Poisson graph after deleting u. Call the remaining graph B. Its process runs at rate $\theta = \binom{n-1}{2}/\binom{n}{2} = (n-2)/n$. The degree of u and its list of independent uniform neighbours are independent of B. If u lies in a nontrivial frozen class, every colouring of B uses at least two colours on that list. Indeed, if a colouring of B were constant on the neighbours, there would be at least $q - 1 \ge 2$ available colours for u. At least one differs from the colour of its alleged frozen partner, a contradiction.

For $1 \le j \le (n-1)/4$, condition on degree j and extend the set of distinct neighbours to a j-set Z, choosing the additional vertices uniformly from its complement. Conditional also on the number of distinct neighbours, the marginal law of Z is uniform and independent of B. Identifying Z destroys colourability whenever the root belongs to a nontrivial frozen class. Dropping the factor $\mathbb{P}(\deg(u) = j) \le 1$ therefore gives, at every time t,

$$\mathbb{P}(u \text{ counted by } B_j(\mathcal{G}_n(t))) \le \mathbb{P}(\mathcal{G}_{n-1}(\theta t) \text{ colourable}, \mathcal{G}_{n-1}(\theta t)/Z \text{ uncolourable}).$$

Integrating and using vertex symmetry proves $\sum_m \mathbb{E} B_j(G_m) \le \theta^{-1} A_{n-1}(j)$. This gives (5.2) by Lemma 5.2. For the remaining integers j, the same displayed bound is immediate from $\sum_m \mathbb{E} B_j(G_m) \le \mathbb{E} H_n \le Cn$: its right side already exceeds Cn. Thus (5.2) holds for every $j \ge 1$, including degrees larger than the number of vertices. Choose a fixed $A > C_1$ and let $w(j) = e^{-Aj^3}$. Summing (5.2) with these weights gives

$$\sum_{m \ge 0} \mathbb{E} \left[\sum_{j \ge 0} w(j) B_j(G_m) \right] \le Cn^{1/3}. \tag{5.3}$$

Both numerical series $\sum_{j \ge 1} (j^2 + 1) e^{-Aj^3}$ and $\sum_{j \ge 1} e^{-(A - C_1)j^3}$ converge.

Recall that Δ_n is the degree of a uniformly chosen endpoint of the terminal edge. To pass from vertex occupation to the law of this endpoint, define the weighted frozen-pair fraction by

$$\kappa_w(G) = \frac{1}{n(n-1)} \sum_{u=1}^{n} (|C_G(u)| - 1) w(\deg_G(u))$$

when G is colourable, and set it to zero otherwise. The terminal identity, with the two endpoints averaged, gives

$$\mathbb{E}[r(\Gamma_n) w(\Delta_n)] = \sum_{m \ge 0} \mathbb{E}[r(G_m) \kappa_w(G_m)].$$

By (3.3) and (3.1),

$$r(G)\kappa_w(G) \le \frac{n}{n-1} \sqrt{r(G)} \sum_{j \ge 0} w(j) B_j(G) \le C\sqrt{n} \sum_{j \ge 0} w(j) B_j(G).$$

Consequently, (5.3) proves

$$\mathbb{E}[r(\Gamma_n) e^{-A\Delta_n^3}] \le Cn^{5/6}. \tag{5.4}$$

The factor $n^{1/3}$ comes from the identification estimate, and the factor $n^{1/2}$ from the bound on frozen-class size. This proves the second of the two estimates announced in the introduction.

The following lemma combines the weighted remaining-lifetime estimate with the degree moments by a change of measure and Jensen's inequality.

Lemma 5.3. *Let $0 \le X \le 1$, $Y \ge 0$, $A > 0$, and $0 < \varepsilon < 1$. If $\mathbb{E}[Xe^{-AY^3}] \le \varepsilon$, then, for every integer $p \ge 1$ with $\mathbb{E}Y^{3p} < \infty$,*

$$\mathbb{E}X \le \sqrt{\varepsilon} + \left(\frac{2A}{\log(1/\varepsilon)}\right)^p \mathbb{E}Y^{3p}. \tag{5.5}$$

Proof. Write $x = \mathbb{E}X$. There is nothing to prove if $x \le \sqrt{\varepsilon}$. Otherwise define a probability measure ν by $\mathrm{d}\nu = (X/x)\,\mathrm{d}\mathbb{P}$. Jensen's inequality for the convex function e^{-At} gives $\frac{\varepsilon}{x} \ge \mathbb{E}_\nu e^{-AY^3} \ge e^{-A\mathbb{E}_\nu Y^3}$. Hölder's inequality and $0 \le X \le 1$ give

$$\log(x/\varepsilon) \le \frac{A}{x}\mathbb{E}[XY^3] \le Ax^{-1/p}(\mathbb{E}Y^{3p})^{1/p}.$$

For $p = 1$ the last inequality follows directly from $X \le 1$; for $p > 1$, use $\mathbb{E}X^{p/(p-1)} \le x$. Since $x > \sqrt{\varepsilon}$, $\log(x/\varepsilon) > \frac{1}{2}\log(1/\varepsilon)$. Raising the resulting inequality to the power p proves (5.5). □

Choose C_0 so that $r(G) \le C_0 n$ for every graph on $[n]$, and apply the lemma with

$$X = \frac{r(\Gamma_n)}{C_0 n}\mathbf{1}_{\{H_n \le M\}}, \qquad Y = \Delta_n \mathbf{1}_{\{H_n \le M\}}, \qquad \varepsilon = C_2 n^{-1/6}.$$

Choose C_2 large enough for (5.4) to give $\mathbb{E}[Xe^{-AY^3}] \le \varepsilon$. For all sufficiently large n, $\varepsilon < 1$ and $\log(1/\varepsilon) \ge (\log n)/7$. Equation (4.6) and Lemma 5.3 then yield, for every fixed integer $p \ge 1$,

$$\mathbb{E}X \le Cn^{-1/12} + \frac{C_p}{(\log n)^{p-1}}.$$

The omitted times have probability at most e^{-bn} by $\mathbb{P}(H_n > \lfloor \Lambda n \rfloor) \le e^{-bn}$, and $r(\Gamma_n) \le C_0 n$ there. For any fixed $B > 0$, take an integer $p > B + 1$. We conclude that

$$\mathbb{E}r(\Gamma_n) = O_{q,B}\left(\frac{n}{(\log n)^B}\right). \tag{5.6}$$

The constants may depend on B. The estimate holds for every fixed B because the exponential moment supplies every fixed polynomial moment of Δ_n.

Taking $B = 8$ in (3.6), with $t = n/(\log n)^2$, yields

$$\mathbb{P}\left(|H_n - a_n| > \frac{n}{(\log n)^2}\right) \le \frac{C}{(\log n)^4} \longrightarrow 0. \tag{5.7}$$

This proves Theorem 1.3 and establishes the hypothesis of Proposition 2.1. That proposition gives convergence of a_n/n and transfers the result to the independent-edge model, completing the proof of Theorem 1.2.

Statement regarding artificial intelligence. None of the mathematical arguments in this paper was generated by a large language model.

Acknowledgments. The author is grateful to Imre Leader for his guidance and support. This work was supported by the CB European PhD Studentship funded by Trinity College, Cambridge. The author would also like to thank the Isaac Newton Institute for Mathematical Sciences, Cambridge, for support and hospitality during the program Logical Foundations of Complexity Theory, where part of the work on this paper was undertaken. This work was supported by EPSRC grant EP/Z000580/1.

References

[1] E. Abbe and A. Montanari, *On the concentration of the number of solutions of random satisfiability formulas*, Random Structures Algorithms **45** (2014), no. 3, 362–382.

[2] D. Achlioptas and E. Friedgut, *A sharp threshold for k-colorability*, Random Structures Algorithms **14** (1999), no. 1, 63–70.

[3] N. Alon and M. Krivelevich, *The concentration of the chromatic number of random graphs*, Combinatorica **17** (1997), no. 3, 303–313.

[4] D. Achlioptas and A. Naor, *The two possible values of the chromatic number of a random graph*, Annals of Mathematics (2) **162** (2005), no. 3, 1335–1351. Preliminary version in STOC 2004.

[5] D. J. Aldous, *Threshold limits for cover times*, Journal of Theoretical Probability **4** (1991), no. 1, 197–211.

[6] V. Bapst, A. Coja-Oghlan, S. Hetterich, F. Raßmann and D. Vilenchik, *The condensation phase transition in random graph coloring*, Communications in Mathematical Physics **341** (2016), 543–606.

[7] M. Bayati, D. Gamarnik and P. Tetali, *Combinatorial approach to the interpolation method and scaling limits in sparse random graphs*, Annals of Probability **41** (2013), no. 6, 4080–4115.

[8] G. Carenini, *The scaling window of random k-SAT*, arXiv:2609.26222 (2026); also Electronic Colloquium on Computational Complexity, Report TR26-145.

[9] G. Carenini, *A polynomial scaling window for random k-SAT and a proof of the satisfiability conjecture*, Electronic Colloquium on Computational Complexity, Report TR26-229 (2026).

[10] A. Coja-Oghlan, *Upper-bounding the k-colorability threshold by counting covers*, Electronic Journal of Combinatorics **20** (2013), no. 3, Paper 32.

[11] A. Coja-Oghlan and D. Vilenchik, *The chromatic number of random graphs for most average degrees*, International Mathematics Research Notices **2016**, no. 19, 5801–5859. Preliminary version in FOCS 2013.

[12] J. Culberson and I. Gent, *Frozen development in graph coloring*, Theoretical Computer Science **265** (2001), no. 1–2, 227–264.

[13] J. Ding, A. Sly and N. Sun, *Proof of the satisfiability conjecture for large k*, Annals of Mathematics (2) **196** (2022), no. 1, 1–388.

[14] E. Friedgut, *Sharp thresholds of graph properties, and the k-SAT problem* (with an appendix by J. Bourgain), Journal of the American Mathematical Society **12** (1999), no. 4, 1017–1054.

[15] F. Krzakala, A. Montanari, F. Ricci-Tersenghi, G. Semerjian and L. Zdeborová, *Gibbs states and the set of solutions of random constraint satisfaction problems*, Proceedings of the National Academy of Sciences **104** (2007), no. 25, 10318–10323.

[16] F. Krzakala, A. Pagnani and M. Weigt, *Threshold values, stability analysis, and high-q asymptotics for the coloring problem on random graphs*, Phys. Rev. E **70** (2004), 046705.

[17] T. Łuczak, *The chromatic number of random graphs*, Combinatorica **11** (1991), no. 1, 45–54.

[18] R. Mulet, A. Pagnani, M. Weigt and R. Zecchina, *Coloring random graphs*, Phys. Rev. Lett. **89** (2002), 268701.

[19] W. Perkins, *Searching for (sharp) thresholds in random structures: where are we now?*, Bulletin of the American Mathematical Society **62** (2025), no. 1, 113–143.

[20] E. Shamir and J. Spencer, *Sharp concentration of the chromatic number on random graphs $G_{n,p}$*, Combinatorica **7** (1987), no. 1, 121–129.

[21] L. Zdeborová and F. Krzakala, *Phase transitions in the coloring of random graphs*, Phys. Rev. E **76** (2007), 031131.

Department of Pure Mathematics and Mathematical Statistics, University of Cambridge, Cambridge, United Kingdom

Email address: gc645@cam.ac.uk

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il