Two-round Ramsey games on random graphs

Yahav Alon* Patrick Morris† Wojciech Samotij‡

May 3, 2023

Abstract

Motivated by the investigation of sharpness of thresholds for Ramsey properties in random graphs, Friedgut, Kohayakawa, Rödl, Ruciński and Tetali introduced two variants of a single-player game whose goal is to colour the edges of a random graph, in an online fashion, so as not to create a monochromatic triangle. In the two-round variant of the game, the player is first asked to find a triangle-free colouring of the edges of a random graph G_1 and then extend this colouring to a triangle-free colouring of the union of G_1 and another (independent) random graph G_2, which is disclosed to the player only after they have coloured G_1. Friedgut et al. analysed this variant of the online Ramsey game in two instances: when G_1 has $\Theta(n^{4/3})$ edges and when the number of edges of G_1 is just below the threshold above which a random graph typically no longer admits a triangle-free colouring, which is located at $\Theta(n^{3/2})$.

The two-round Ramsey game has been recently revisited by Conlon, Das, Lee and Mészáros, who generalised the result of Friedgut at al. from triangles to all strictly 2-balanced graphs. We extend the work of Friedgut et al. in an orthogonal direction and analyse the triangle case of the two-round Ramsey game at all intermediate densities. More precisely, for every $n^{-4/3} \ll p \ll n^{-1/2}$, with the exception of $p = \Theta(n^{-3/5})$, we determine the threshold density q at which it becomes impossible to extend any triangle-free colouring of a typical $G_1 \sim G_{n,p}$ to a triangle-free colouring of the union of G_1 and $G_2 \sim G_{n,q}$. An interesting aspect of our result is that this threshold density q 'jumps' by a polynomial quantity as p crosses a 'critical' window around $n^{-3/5}$.

1 Introduction

Given graphs G and H, we say that G is *H-Ramsey* if any red/blue-colouring of the edges of G results in a monochromatic copy of H. The classical theorem of Ramsey [19], from which the term *Ramsey theory* stems, implies that K_n is H-Ramsey for all n large enough in terms H. It is natural to ask what other graphs G are also H-Ramsey and a prominent theme has been to explore the existence of Ramsey graphs G that are *sparse*; see, for example, [17] and the references therein. One famous example is the work of Frankl and Rödl [7], who constructed K_3-Ramsey graphs that are K_4-free by considering sparse random graphs. This prompted Łuczak, Ruciński and Voigt [13] to initiate the systematic study of thresholds for Ramsey properties in random graphs, which has since become a prominent theme in probabilistic combinatorics. In particular, this pair of papers established the following. (Here and throughout, we denote by $G_{n,p}$ the binomial random graph with n vertices

*School of Mathematical Sciences, Tel Aviv University, Tel Aviv 6997801, Israel. Email: yahavalo@tauex.tau.ac.il.

†Departament de Matemàtiques, Universitat Politècnica de Catalunya (UPC), Carrer de Pau Gargallo 14, 08028 Barcelona, Spain. Email: pmorrismaths@gmail.com. Research supported in part by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) Walter Benjamin program - project number 504502205.

‡School of Mathematical Sciences, Tel Aviv University, Tel Aviv 6997801, Israel. Email: samotij@tauex.tau.ac.il. Research supported in part by the European Research Council Consolidator Grant 101044123 (RandomHypGra) and by the Israel Science Foundation grant 2110/22.

and edge probability p and, for a graph property P, we say that P holds *asymptotically almost surely* (a.a.s. for short) in $G_{n,p}$ if the probability that P holds tends to 1 as n tends to infinity.)

Theorem 1.1 ([7, 13])**.** *There exist constants* $0 < c < C$ *such that:*

(a) If $p_0 \leq cn^{-1/2}$*, then a.a.s.* G_{n,p_0} *is* not K_3*-Ramsey.*

(b) If $p_1 \geq Cn^{-1/2}$*, then a.a.s.* G_{n,p_1} is K_3*-Ramsey.*

The 1-statement (b) was implicit in the work of Frankl and Rödl [7]; Łuczak, Ruciński, and Voigt [13] proved the 0-statement (a) and provided an alternative proof of (b). The study of Ramsey properties of random graphs culminated in a seminal series of papers by Rödl and Ruciński [20, 21, 22], who greatly generalised Theorem 1.1, establishing thresholds for any graph H and also any number of colours $2 \leq r \in \mathbb{N}$. Recently, Nenadov and Steger [16] provided a short proof of (b) (and the analogous 1-statements for all H and number of colours) by using the theory of hypergraph containers [3, 23]; see Section 2.3 for more on this.

1.1 Ramsey games on random graphs

The pioneering work [7, 13, 20, 21, 22] on Ramsey properties of random graphs has been highly influential, with many extensions and variations being studied. In this paper, we will focus on *Ramsey games* played on random graphs, a viewpoint introduced in by Friedgut, Kohayakawa, Rödl, Ruciński and Tetali [8]. The starting point of their work was to view Theorem 1.1 as a single-player game played against a random source. In this game, a player is presented with a set of M random edges of K_n, for some $1 \leq M \leq \binom{n}{2}$, and asked to colour the edges, trying to avoid a monochromatic H. Standard results on the asymptotic equivalence of random graph models (see, for example [12, Section 1.4]), along with Theorem 1.1, show that, in the case that $H = K_3$, the player will asymptotically almost surely fail when $M \geq Cn^{3/2}$ and succeed when $M \leq cn^{3/2}$, for appropriately chosen $c, C > 0$. In [8], the authors introduced the following two variants on this game:

The online game. The player is presented with edges of K_n one at a time, according to a uniformly random permutation. Upon seeing each edge, the player must colour the edge with only the knowledge of the previous (already coloured) edges. The game ends when a monochromatic H occurs and the aim of the player is to last as long as possible.

The two-round game. Here, the player is given two random graphs: a uniformly random n-vertex graph G_1 with M_1 edges and a second, *independent* random graph G_2 with M_2 edges. The player must colour the first random graph, avoiding monochromatic copies of H and with no knowledge of the second random graph. The player is then presented with the second random graph and asked to extend their colouring of G_1 to an H-free colouring of $G_1 \cup G_2$.

The work of Friedgut et al. [8] formalised these games, studied the case where $H = K_3$, and drew interesting connections between the games, the original random Ramsey problem and other related research directions. In particular, the two-round game arose naturally in another work of a subset of the authors [10] that established that the property of being K_3-Ramsey has a sharp threshold in $G_{n,p}$.

With regards to the online game, the authors of [8] gave a simple argument showing that, when played with two colours, the game typically ends with a graph having $\Theta(n^{4/3})$ edges. More precisely, for any[1] $M = M(n) \ll n^{4/3}$, a.a.s. the player has a strategy to colour the edges avoiding a monochromatic K_3. Moreover for any $M \gg n^{4/3}$, a.a.s. the player will be forced to create a monochromatic K_3 while colouring the first M edges.

[1] Here, and throughout, for real-valued functions $f = f(n)$ and $g = (n)$, we write $f \ll g$ (or $g \gg f$) to denote that f/g tends to 0 as n tends to infinity.

Interestingly, as discussed in more detail below, the two-round game is used to prove this upper bound on the running time of the online game. There has been a considerable amount of work [2, 14, 15, 18] generalising this result and establishing the expected running time of the online game under optimal play. However, a full understanding for all graphs H and number of colours remains elusive; see [18] and the references therein for the currently best known bounds.

In the context of the two-round game with respect to K_3 and with two colours, Theorem 1.1 implies that, when $M_1 \geq Cn^{3/2}$ for a large enough C, a.a.s. the player will not be able to survive even the first round, as any colouring of G_1 will induce a monochromatic K_3. Friedgut, Kohayakawa, Rödl, Ruciński and Tetali [8] explored what happens near this extreme, when $M_1 = cn^{3/2}$ for some small constant $c > 0$, as well as when $M_1 = \Theta(n^{4/3})$. They established the following results.

Theorem 1.2 ([8])**.** *Let G_1 be a uniformly random graph on n vertices with M_1 edges and G_2 an independent random graph on n vertices with M_2 edges. Suppose either*

(a) $M_1 = cn^{3/2}$ for some $c > 0$ and $M_2 \gg 1$; or

(b) $M_1 = cn^{4/3}$ for some $c > 0$ and $M_2 \gg n^{4/3}$.

Then a.a.s. no G_1-measurable K_3-free red/blue-colouring of G_1 can be extended to a K_3-free red/blue-colouring of $G_1 \cup G_2$.

Theorem 1.2 (a) shows that just below the threshold for the K_3-Ramsey property, the random graph G is a.a.s. very close to being K_3-Ramsey in the sense that, although G does admit a K_3-free colouring, no such colouring can be extended after one adds to G some $\omega(1)$ random edges.

Theorem 1.2 (b) highlights the connection between the online game and the two-round game, as it implies that a.a.s. the online game cannot last $M \gg n^{4/3}$ steps. Indeed, even if we allow the player of the online game a 'grace period' and do not ask for any colouring until $n^{4/3}$ edges are revealed, a.a.s. no matter how the player chooses to colour these, they will not be able to extend to the next $M - n^{4/3} \gg n^{4/3}$ random edges.

In contrast to the online game, the two-round game has not been further explored until the recent work of Conlon, Das, Lee and Mészáros [6]. They investigated to what extent Theorem 1.2 (a) can be extended to two-round games with respect to other graphs H. Answering a question from [8], they showed that a large family of graphs H (namely strictly 2-balanced graphs, see [6] for a definition) exhibit the same behaviour as K_3 in that just below the threshold for the H-Ramsey property a.a.s. all H-free colourings can be killed by adding a super-constant more random edges. They also showed that this is not the case for all graphs H and posed the interesting question as to what properties of H determine this behaviour. Finally, we mention that both [8] and [6] explore the two-round game with three colours near the Ramsey threshold.

1.2 Our results

The aim of our work here is to expand on the work of [8] on the two-round game in a different direction. We keep our focus on $H = K_3$ and two colours and investigate the outcome of the two-round game as the number of edges in each round is varied. We will state and prove our results in the setting of binomial random graphs, as opposed to uniform graphs with a fixed number of edges. It is well-known that these models are asymptotically equivalent [12, Section 1.4], but the independence of the binomial model makes it more convenient to work with. In order to systematically study two-round games on random graphs at different densities, we introduce the following notion of a *Ramsey completion threshold*, which captures the critical density of the second graph at which the probability that the player succeeds in extending their H-free colouring from the first to the second graph jumps from $1 - o(1)$ to $o(1)$.

Definition 1.3. *Given some probability $p = p(n) \in [0, 1]$ and a graph H, we say that $q = q(n)$ is a* Ramsey completion threshold *for H with respect to p if the following holds a.a.s. for $G \sim G_{n,p}$:*

(a) There exists a G-measurable H-free red/blue-colouring φ of G such that, if $q_0 \ll q$, then a.a.s., φ can be extended to an H-free colouring of $G \cup G_{n,q_0}$.

(b) For every G-measurable H-free red/blue-colouring φ of G, if $q_1 \gg q$, then a.a.s., φ cannot be extended to an H-free colouring of $G \cup G_{n,q_1}$.

If such a completion threshold exists, we denote it by[2] *$q(n; H, p)$. If all red/blue-colourings of $G_{n,p}$ contain a monochromatic H with probability $\Omega(1)$, we set $q(n; H, p) = 0$.*

We refer to (a) as the 0-*statement* of the definition and (b) as the 1-*statement*. Observe that Theorem 1.1 gives that there exists $C > 0$ such that $q(n; K_3, p) = 0$ for all $p \geq Cn^{-1/2}$. Moreover, Theorem 1.2 (a) gives for *any* constant $c > 0$, if $p = cn^{-1/2}$ and $q(n; K_3, p) \neq 0$, then $q(n; K_3, p) = n^{-2}$ whereas Theorem 1.2 (b) gives that[3] $q(n; K_3, p) = n^{-2/3}$ when $p = \Theta(n^{-2/3})$. Here, we complete the picture for almost all intermediate values of p.

Theorem 1.4. *We have that*

$$q(n; K_3, p) = \begin{cases} n^{-6}p^{-8} & \text{if } n^{-3/5} \ll p \ll n^{-1/2}; \qquad \text{(upper range)} \\ n^{-3}p^{-7/2} & \text{if } n^{-2/3} \ll p \ll n^{-3/5}. \qquad \text{(lower range)} \end{cases}$$

An interesting aspect of this result is that there is a 'jump' in the completion threshold at around $n^{-3/5}$. Indeed, when $p = n^{-3/5}$, then $n^{-3}p^{-7/2} = n^{-9/10}$ whereas $n^{-6}p^{-8} = n^{-6/5}$. We refer to the values of p larger than $n^{-3/5}$ as the *upper range* and those smaller than $n^{-3/5}$ as the *lower range*. Determining the behaviour of $q(n; K_3, p)$ when $p = \Theta(n^{-3/5})$ remains an intriguing open question, as does exploring the behaviour of the two-round game for different H as the densities of the two random graphs vary. In particular, it is far from clear what properties of H could determine this behaviour; Theorem 1.4 does not suggest any obvious conjecture.

Our proof of Theorem 1.4 incorporates several different approaches to capture the different behaviour occurring at different densities. One particular feature that we would like to highlight is a novel use of the 'discharging method' to prove the existence of a desired colouring of the first graph in the lower range (see the proof of Lemma 3.13). This method is inspired by an argument in recent work of Friedgut, Kuperwasser, Schacht and the third author [9] that establishes sufficient conditions for sharpness of thresholds for various Ramsey properties. Here, we build on this general idea of using discharging to find 'easily colourable configurations' in graphs of small density, but we adopt a more involved discharging scheme catered to our purposes. We believe that this method may find further applications in the study of Ramsey properties of graphs and other discrete structures and our work here demonstrates its flexibility.

Organisation. The proof of Theorem 1.4 naturally splits into four parts. In Section 3, we establish the 0-statement, that is, the lower bound on $q(n; K_3, p)$. The (shorter) argument for the upper range is presented in Section 3.1 and the (longer) argument for the lower range – in Section 3.2. In Section 4, we establish the 1-statement. The 1-statement—the upper bound on $q(n; K_3, p)$—for the (easier) lower range will be proved in Section 4.2 and for the (harder) upper range – in Section 4.3. Before embarking on this, we collect the relevant notation and tools in Section 2.

[2] As usual, we abuse notation here as such a threshold will not be determined uniquely, but rather up to constants.

[3] In fact they only prove the 1-statement (b) of Definition 1.3 but the corresponding 0-statement also holds and essentially follows from the analysis of the online game in [8], see Remark 3.7.

Acknowledgement. The research on this project was initiated during a joint research workshop of Tel Aviv University and the Freie Universität Berlin on Ramsey Theory, held in Tel Aviv in March 2020, and partially supported by the GIF grant G-1347-304.6/2016. We would like to thank the German–Israeli Foundation (GIF) and both institutions for their support. We also thank Shagnik Das for suggesting this problem.

2 Preliminaries

In this section, we present the notation and tools that we will use in our proofs.

2.1 Notation

We use standard probabilistic and graph theory notation throughout. For a graph G and a subgraph F, we let $N_F(G)$ denote the number of copies of F in G. For vertex subsets $A, B \subseteq V(G)$ of a graph G, $G[A]$ denotes the graph induced by G on A and $e(A, B)$ denotes the number of edges of G with one endpoint in A and the other in B (edges in $G[A \cap B]$ are counted *once* here). For a vertex $v \in V(G)$, we let $d_G(v)$ denote the degree of v in G and $\Delta(G) \coloneqq \max_{v \in V(G)} d_G(v)$ denote the maximum degree.

Given a hypergraph $\mathcal{H}$, we denote the numbers of its edges and vertices by $e(\mathcal{H})$ and $v(\mathcal{H})$, respectively. Further, for a vertex subset $T \subset V(G)$, $d_{\mathcal{H}}(T)$ denotes the number of edges of $\mathcal{H}$ containing T. For an integer $\ell \geq 0$, we write $\Delta_\ell(\mathcal{H}) \coloneqq \max\{d_{\mathcal{H}}(T) : T \subset V(\mathcal{H}), |T| = \ell\}$ for the maximum degree of a vertex set of size ℓ in $\mathcal{H}$. A set of edges $M \subseteq E(\mathcal{H})$ in a hypergraph $\mathcal{H}$ is a *matching* if $e \cap f = \emptyset$ for all $e \neq f \in M$. The *matching number* $\nu(\mathcal{H}) \coloneqq \max\{|M| : M \text{ is a matching in } \mathcal{H}\}$ is the size of the largest matching in a hypergraph $\mathcal{H}$.

For a set W, we let $\mathcal{P}(W) \coloneqq \{U : U \subseteq W\}$ denote the power set of W, the set of all subsets of W. For a Boolean statement A, we denote by $\mathbb{1}[A]$, the indicator function which evaluates to 1 if A holds and 0 if A does not hold.

2.2 Concentration inequalities

We will frequently use concentration inequalities for two families of random variables. The first inequality, often attributed to Chernoff [5] (see also [12, Theorem 2.1]), deals with the case of binomial random variables.

Lemma 2.1 (**Chernoff's inequality**). *Let $X \sim \mathrm{Bin}(n, p)$ be binomially distributed and let $\mu = \mathbb{E}[X] = np$. Then for any $k \geq 0$, we have that*

$$\Pr[X \geq \mu + k] \leq \exp\left(-\frac{k^2}{2(\mu + k/3)}\right).$$

The following inequality, known as Janson's inequality [11] (see also [12, Theorem 2.14]) provides an exponential bound for the lower tail of the number of edges induced by a hypergraph on a random subset of its vertices.

Lemma 2.2 (**Janson's inequality**). *Let Γ be a finite set, let $p\colon \Gamma \to [0, 1]$, and let Γ_p be a random subset such that every element $a \in \Gamma$ is in Γ_p with probability $p(a)$, independently of all other elements.*

Suppose that $A_1, \ldots, A_m$ is a sequence of nonempty subsets of Γ. For each $i \in \{1, \ldots, m\}$, denote by $I_i \coloneqq \mathbb{1}[A_i \subseteq \Gamma_p]$ the indicator random variable for the event $A_i \subseteq \Gamma_p$. Finally, denote

$$X \coloneqq \sum_{i=1}^{m} I_i, \qquad \mu \coloneqq \mathbb{E}[X], \qquad \text{and} \qquad \Delta \coloneqq \sum_{i,j=1}^{m} \mathbb{1}[A_i \cap A_j \neq \emptyset] \cdot \mathbb{E}[I_i I_j].$$

Then, for every $0 \leq k \leq \mu$,

$$\Pr[X \leq \mu - k] \leq \exp\left(-\frac{k^2}{2\Delta}\right).$$

In the setting of Janson's inequality (Lemma 2.2), we will also be interested in showing that a.a.s., there is a large collection of *disjoint* subsets A_i which all appear in Γ_p. Our next result provides an upper bound on the probability that this does not happen, and can be derived from Lemma 2.2. We provide the details of this derivation in Appendix A. We recall that for a hypergraph $\mathcal{H}$, $\nu(\mathcal{H})$ denotes the matching number of $\mathcal{H}$, that is, the size of that largest matching in $\mathcal{H}$.

Corollary 2.3 (**Maximal disjoint families**). *Let Γ, p, $A_1, \dots, A_m$, μ, and Δ be as in the statement of Lemma 2.2 and set $D := \frac{\mu^2}{800\Delta}$. Writing $\mathcal{A}$ for the hypergraph with vertex set Γ and edge set $\{A_1, \dots, A_m\}$, we have*

$$\Pr\left[\nu\big(\mathcal{A}[\Gamma_p]\big) \le D\right] \le \exp(-D).$$

2.3 Containers

We will appeal to the method of hypergraph containers, developed by Balogh, Morris and the third author of the present paper [3], and independently, by Saxton and Thomason [23]. The key idea underlying this method is that, given a uniform hypergraph whose edge set is evenly distributed, one can distribute its independent sets into a well-behaved collection of *containers*. In more detail, these containers are vertex subsets that are almost independent (in that they induce few edges of the hypergraph), every independent set of the hypergraph lies in some container and, crucially, we have a bound on the number of containers. As there are many fewer containers than independent sets in the hypergraph, reasoning about containers rather than independent sets leads to more efficient arguments and this technique has proven to be extremely powerful. Indeed, the setting of independent sets in hypergraphs can be used to encode a wide range of problems in combinatorics and the method of hypergraph containers has been successfully exploited in a multitude of different settings, see [4]. Particularly relevant to our work here are the applications of the method in sparse Ramsey theory, a program which was initiated by Nenadov and Steger [16], who reproved the 1-statement of Theorem 1.1 utilising containers.

We state the container lemma in the following form, which follows from a general container lemma of Saxton and Thomason [23]. We show how to derive this form of the container theorem in Appendix B. We remark that the following version of the container theorem (and indeed the version of Saxton and Thomason [23]) gives slightly more than we promised in the discussion of the general method given above: The following theorem allows us to conclude not only that all independent sets lie in containers, but also sets that are very close to being independent (see condition (i) of the theorem). Also, the theorem posits that every container is of the form $f(S_1, \dots, S_t)$ for some family of small sets S_i, that is, each container is determined by a (constant-sized) collection of small subsets of the container. This in turn will allow us to run various union bounds over such collections of small subsets.

Theorem 2.4. *For every positive integer $2 \le k \in \mathbb{N}$ and all $\varepsilon \in (0,1)$ and $1 \le K \in \mathbb{N}$, there exist $t \in \mathbb{N}$ and $\delta > 0$ such that the following holds. Suppose that a nonempty k-uniform hypergraph $\mathcal{H}$ with vertex set V and $\tau \in (0, 1/t)$ satisfy*

$$\Delta_\ell(\mathcal{H}) \le K\tau^{\ell-1} \cdot \frac{e(\mathcal{H})}{v(\mathcal{H})}$$

for every $\ell \in \{2, \dots, k\}$. Then, there exists a function $f\colon \mathcal{P}(V)^t \to \mathcal{P}(V)$ with the following properties:

(i) For every set $I \subseteq V$ satisfying $e(\mathcal{H}[I]) \le \delta\tau^k e(\mathcal{H})$, there are $S_1, \dots, S_t \subseteq I$, each of size at most $\tau v(\mathcal{H})$ and such that $I \subseteq f(S_1, \dots, S_t)$.

(ii) For every $S_1, \dots, S_t \subseteq V$, the set $f(S_1, \dots, S_t)$ induces fewer than $\varepsilon e(\mathcal{H})$ edges in $\mathcal{H}$.

2.4 Typical properties of $G_{n,p}$

In this section, we derive some properties of $G_{n,p}$ that a.a.s. hold. These will be useful throughout the paper. We recall from Section 2.1, that $N_F(G)$ denotes the number of copies of F in G. We also recall the standard notion of density of a graph F, which is $m(F) \coloneqq \max\{e_J/v_J : J \subseteq F\}$.

Lemma 2.5. *Suppose that $G \sim G_{n,p}$ with $n^{-2/3} \ll p \ll n^{-1/2}$. Then a.a.s., G has the following properties:*

(P1) $\Delta(G) \le 2np$*;*

(P2) $N_F(G) \le 2n^{v_F}p^{e_F}$ *for all F that satisfy both $v_F \le 8$ and $m(F) \le 3/2$; moreover, $N_F(G) \le n^{v_F}p^{e_F}\log n$ for all F that only satisfy $v_F \le 8$.*

(P3) $N_{K_{2,10}}(G) \le n^{11}p^{18}$*;*

(P4) *There is a constant $\theta > 0$ such that, for all $U \subseteq V(G)$ with $|U| \ge \frac{n}{2}$, the graph $G[U]$ contains a set of $\theta|U|^3p^3$ edge-disjoint triangles;*

(P5) *For all integer-valued $a = a(n), b = b(n)$ such that $(\log n)^7/p \ll a, b \le n$, the following holds for any $A, B \subseteq V$ with $|A| \le a, |B| \le b$:*

$$e(A,B) \le |A| \cdot |B| \cdot p + \frac{a \cdot b \cdot p}{\log^3 n}.$$

Proof. Properties **(P1)** and **(P2)** are standard properties of the distribution of the edges and small subgraphs in $G_{n,p}$ that can be proved using the second moment method and union bounds, see for example [1, Chapter 4]. Property **(P3)** follows from Markov's inequality. Indeed, we have that

$$\mathbb{E}[N_{K_{2,10}}] \le n^{12}p^{20} = \frac{n^{11}p^{18}}{np^2},$$

and $np^2 \ll 1$ due to the fact that $p \ll n^{-1/2}$.

We will establish property **(P4)** using Corollary 2.3. To this end, fix any $U \subseteq V(G)$ with $|U| \ge \frac{n}{2}$, let $X \coloneqq N_{K_3}(G[U])$ and note that $\mu \coloneqq \mathbb{E}[N_{K_3}(G[U])] = \binom{|U|}{3}p^3$ for n large. Denoting by I_T, for every triangle T in U, the indicator random variable of the event that T appears in G, we have

$$\Delta \coloneqq \sum \{I_T I_{T'} : T, T' \text{ triangles in } K_n[U], E(T) \cap E(T') \neq \emptyset\} \le \mu \cdot (1 + np^2) \le 2\mu,$$

where μ accounts for a choice of T, the first summand in the bracket accounts for a choice of T' such that $|E(T') \cap E(T)| \ge 2$ (and hence $T = T'$) and the second summand accounts for a choice of T' such that $|E(T) \cap E(T')| = 1$; we also used that $p \ll n^{-1/2}$ in the last inequality. By Corollary 2.3, with $\Gamma = K_n[U]$ and $\mathcal{A}$ the collection of all triangles in Γ, the probability that every collection of pairwise edge-disjoint triangles in $G[U]$ is smaller than

$$D \coloneqq \frac{\mu^2}{800\Delta} \ge \frac{\mu}{1600} \ge \frac{|U|^3p^3}{10000}$$

is at most e^{-D}. Property **(P4)** then follows from a union bound over the (less than 2^n) choices of vertex subset U, using that $D \gg n$ for all such U because $p \gg n^{-2/3}$.

Finally, in order to prove that G has **(P5)**, note that, for any choice of $a, b \gg (\log n)^7/p$ and $A, B \subseteq V$ such that $|A| \le a, |B| \le b$, we have that $\mu \coloneqq \mathbb{E}[e(A,B)] = \left(|A||B| - \binom{|A\cap B|}{2}\right)p \le |A||B|p$. Moreover $e(A,B)$ is binomially distributed and, setting $k \coloneqq abp/\log^3 n$, we have that $\mu + k/3 \le 2abp$. Hence, by Lemma 2.1, we have that

$$\Pr\left[e(A,B) \ge |A||B|p + k\right] \le \Pr\left[e(A,B) \ge \mu + k\right] \le \exp\left(-\frac{k^2}{2(\mu + k/3)}\right)$$
$$\le \exp\left(-\frac{(abp)^2}{\log^6 n \cdot 4abp}\right) \le \exp\left(-\frac{abp}{4\log^6 n}\right).$$

Therefore, applying a union bound over the choice of a, b, A, B, we get:

$$\begin{aligned}\Pr[G \notin \textbf{(P5)}] &\leq \sum_{a,b} \sum_{\substack{1\leq\alpha\leq a\\1\leq\beta\leq b}} \binom{n}{\alpha}\cdot\binom{n}{\beta}\cdot\exp\left(-\frac{abp}{4\log^6 n}\right)\\ &\leq \sum_{a,b} \sum_{\substack{1\leq\alpha\leq a\\1\leq\beta\leq b}} n^\alpha\cdot n^\beta\cdot\exp\left(-\frac{abp}{4\log^6 n}\right)\\ &\leq n^2\cdot\sum_{a,b} n^a\cdot n^b\cdot\exp\left(-\frac{abp}{4\log^6 n}\right)\\ &\leq n^2\cdot\sum_{a,b}\exp\left(a\log n + b\log n - \frac{abp}{4\log^6 n}\right),\end{aligned}$$

where the (outer) sum goes over all choices of $a = a(n) \in \mathbb{N}$ and $b = b(n) \in \mathbb{N}$ such that $(\log n)^7/p \ll a, b \leq n$. For all such a, b, we have that $\frac{abp}{4\log^6 n} \gg a\log n, b\log n$, and so

$$\Pr[G \notin \textbf{(P5)}] \leq n^2\cdot\sum_{a,b} n^{-g_{a,b}(n)} \ll 1,$$

where $g_{a,b}(n) = \frac{abp}{4\log^7 n} - a - b \gg 1$ for all a, b. □

3 Proof of the 0-statements

In this section, we prove our 0-statements, establishing the lower bounds on $q(n; K_3, p)$ in Theorem 1.4. Our proofs in the lower and upper ranges follow the same scheme. First, we will show the existence of a *good* colouring of $G_1 \sim G_{n,p}$, i.e., a colouring with certain desirable properties specified in Definition 3.1 below. Second, we will show that any such good colouring of G_1 can be a.a.s. extended to the second independent random graph $G_2 \sim G_{n,q}$ when q is chosen appropriately. While the arguments showing existence of a good colouring will be different in the lower and the upper ranges, extendability of good colourings, stated as Proposition 3.2 below, is proved in the full range of interest.

Definition 3.1. *For a colouring $\varphi\colon E(G) \to \{\text{red}, \text{blue}\}$ of the edges of a graph G, we define the coloured graph C_{rrbb} to be a 4-cycle with two adjacent red edges and two adjacent blue edges. Then for $t \geq 0$, we say a colouring $\varphi\colon E(G) \to \{\text{red}, \text{blue}\}$ is* t-good *if it has the following properties:*

1. *φ has no monochromatic triangles;*
2. *every edge of G that is not in a triangle is coloured blue;*
3. *the number of C_{rrbb} in G coloured by φ is less than t.*

Moreover, if the colouring is 0-good, we we will refer to it as being very good.

We remark that conditions 1 and 2 will be easy to impose on a colouring of $G_1 \sim G_{n,p}$. Indeed, since p is always below the K_3-Ramsey threshold of Theorem 1.1, a K_3-free colouring exists and one can recolour edges not in triangles so that condition 2 is also satisfied. Thus, the critical condition in the definition is condition 3. The motivation for considering copies of C_{rrbb} is that they pose a direct threat to being able to extend a colouring to G_2. Indeed, consider an edge that forms a triangle with both the red edges and the blue edges of a copy of of C_{rrbb} in G_1. If this edge appears in G_2, then clearly there is no way to colour the edge without creating a monochromatic triangle.

Proposition 3.2. *Suppose that $n^{-2/3} \ll p \ll n^{-1/2}$, $t > 0$ and $0 < q \ll \min\{t^{-1}, n^{-3}p^{-7/2}\}$ and let $G_1 \sim G_{n,p}$ and $G_2 \sim G_{n,q}$ be independent. Then a.a.s. any G_1-measurable colouring $\varphi\colon E(G_1) \to \{\text{red}, \text{blue}\}$ which is t-good can be extended to a monochromatic triangle-free colouring of $G_1 \cup G_2$.*

In the proof of Proposition 3.2 and the further proofs of the 0-statements, we will consider copies of certain uncoloured subgraphs, which are defined in Figure 1. Our next lemma explains why the graphs F_0, F_1, K_4, and their edge-deleted companions F_0^-, F_1^-, K_4^-, appear naturally when one looks for colourings with few copies of C_{rrbb}.

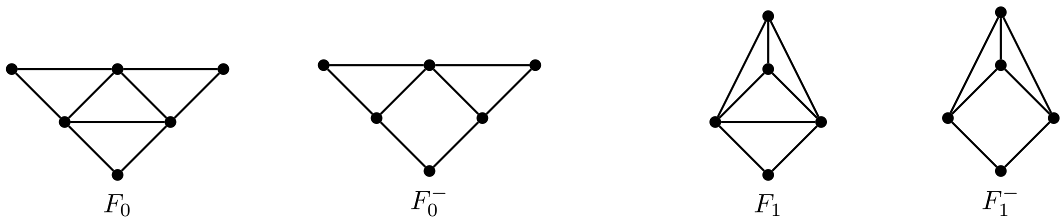

Figure 1: The graphs F_0, F_0^-, F_1 and F_1^-.

Lemma 3.3. *Suppose that a 4-cycle C in a graph G has two adjacent edges e_1, e_2 such that, for $i = 1, 2$, e_i is contained in a triangle of G that does not contain e_{3-i}. Then C is contained in a copy of F^- in G, for some $F \in \{F_0, F_1, K_4\}$. Moreover, in each case, adding the edge that forms a triangle with e_1 and e_2 completes this copy of F^- to a copy of F.*

Proof. Label the vertices of C as x, y, w, z so that $e_1 = xy$ and $e_2 = xw$. Further, for $i = 1, 2$, let u_i be the vertex which forms the triangle with e_i given by the statement of the lemma; we therefore have that $u_1 \neq w$ and $u_2 \neq y$. We consider the following cases. Firstly, if one of the u_i is equal to z, then the vertices of C host a K_4^- in G and we are done. So we can assume that neither u_i is equal to z. If $u_1 = u_2$, then we get a copy of F_1^- that contains C, whilst if $u_1 \neq u_2$, we get a copy of F_0^- containing C. The moreover statement can also be easily checked in each case. □

The following simple consequence of Lemma 3.3 is more easily applicable in some of our proofs.

Corollary 3.4. *Let G be a graph coloured by some $\varphi\colon E(G) \to \{\text{red}, \text{blue}\}$ which satisfies conditions 1 and 2 of Definition 3.1. Then any copy of C_{rrbb} in G is contained in some copy of F^- in G for some $F \in \{F_0, F_1, K_4\}$.*

Proof. Let e_1 and e_2 be the red edges in the copy of C_{rrbb}. Condition 2 in Definition 3.1 implies that each of these edges is in a triangle of G. Moreover, G does not contain the edge e_3 that forms a triangle with e_1 and e_2, as otherwise there would be no way to colour e_3 without creating a monochromatic triangle, contradicting the fact that φ is triangle-free, which is condition 1 in Definition 3.1. □

We are now in a position to prove Proposition 3.2.

Proof of Proposition 3.2. We first show the following claim.

Claim 3.5. *A.a.s. every copy of F_0, F_1 and K_4 in $G_1 \cup G_2$ has at most one edge in G_2.*

Proof. Let X_{F_0} count the number of copies of F_0 in $G_1 \cup G_2$ with at least two edges in G_2. There are at most n^6 copies of F_0 in K_n and the probability that each such copy appears in $G_1 \cup G_2$ with at least two edges in G_2 is at most $\binom{9}{2}(p+q)^7 q^2$. Since $q \ll p$, by our assumptions on q and p, we can bound the expectation of X_{F_0} as follows:

$$\mathbb{E}[X_{F_0}] \leq \binom{9}{2} n^6 (p+q)^7 q^2 \leq 2\binom{9}{2} n^6 p^7 q^2 \ll 1,$$

using that $q \ll n^{-3}p^{-7/2}$ in the last inequality here. Similarly, defining X_F to be the number of copies F in $G_1 \cup G_2$ with at least two edges in G_2 for $F \in \{F_1, K_4\}$, we have

$$\mathbb{E}[X_{F_1}] \leq 2^8 n^5 p^6 q^2 \ll n^{-1}p^{-1} \ll 1, \qquad \mathbb{E}[X_{K_4}] \leq 2^6 n^4 p^4 q^2 \ll n^{-2}p^{-3} \ll 1.$$

The assertion of the claim now follows by Markov's inequality. ◀

Given G_1 and a red/blue-colouring of its edges, we call an edge $uv \in K_n \setminus G_1$ *dangerous* if there is a copy of C_{rrbb} in G_1 with edges uw, wv, vx, xu for some vertices $w, x \notin \{u, v\}$ such that uw, wv are coloured red and vx, xu are coloured blue.

Claim 3.6. *A.a.s. G_1 has the following properties:*

(i) A.a.s. (with respect to G_2), every copy of $F \in \{F_0, F_1, K_4\}$ in $G_1 \cup G_2$ has at most one edge in G_2,

(ii) For every t-good colouring φ of G_1, a.a.s. G_2 contains no dangerous edges.

Proof. Property (i) follows from Claim 3.5 and Fubini's theorem. Further, each t-good colouring φ of G_1 contains at most t copies of C_{rrbb} and hence at most t dangerous edges. Since each such edge appears in G_2 with probability $q \ll t^{-1}$, the expected number of dangerous edges that appear in G_2 is $o(1)$ and (ii) follows from Markov's inequality. ◀

In view of Claim 3.6, it suffices to show that, for every G_1 that has the two properties described in the claim, every t-good colouring φ of G_1 can be extended to a K_3-free colouring of $G_1 \cup G_2$ for every graph G_2 that satisfies the events described in both (i) and (ii), which a.a.s. hold in G_2 (for fixed G_1 and φ). To this end, consider an arbitrary ordering of the edges of $G_2 \setminus G_1$ and colour them one-by-one according to the following rule. We colour $e \in G_2 \setminus G_1$ blue *unless* e forms a blue triangle with previously coloured edges (of $G_1 \cup G_2$), in which case we colour e red. We claim that the resulting colouring of $G_1 \cup G_2$ is K_3-free.

Note that the only possible monochromatic triangles that can occur in our colouring of $G_1 \cup G_2$ must be red and contain an edge of $G_2 \setminus G_1$, as φ is good and thus triangle-free. Suppose that e is the *last* edge of $G_2 \setminus G_1$ that completes a red triangle; denote this triangle by T and its two remaining edges by f_r and g_r. Note that e is also contained in a triangle with two blue edges, say f_b and g_b, already coloured in $G_1 \cup G_2$, as otherwise our rule would colour e blue. We claim that f_r and g_r are both contained in triangles other than T in $G_1 \cup G_2$. Indeed, let $h \in \{f_r, g_r\}$ be arbitrary. If $h \in G_1$, then h must be contained in a (non-monochromatic) triangle of G_1 due to the fact that our colouring of G_1 was good and the fact that h is coloured red. If $h \in G_2 \setminus G_1$, then it must be contained in a triangle whose remaining two edges are blue, as otherwise we would have coloured it blue. Consequently, by Lemma 3.3, we have that the 4-cycle $C := \{f_r, g_r, f_b, g_b\}$ is contained in a copy of F^- in $G_1 \cup G_2$, for some $F \in \{F_0, F_1, K_4\}$, and the edge e completes this copy of F^- to a copy of F. However, as we have assumed that G_2 contains no dangerous edges, one of the edges of the C_{rrbb}-copy $C := \{f_r, g_r, f_b, g_b\}$ must belong to G_2. This contradicts the assumed conclusion of (i), as the copy of F containing C has at least two edges in G_2, namely e and one of the edges in C. This shows that no red triangle can occur, which concludes the proof of the proposition. □

Remark 3.7. Our proof of the 0-statements adopts a greedy strategy to colour the second random graph. In fact, our colouring is identical to the colouring used in [8] to prove that the online game a.a.s. lasts $\Omega(n^{4/3})$ rounds under optimal play. Indeed, they also colour each edge blue as it appears unless it creates a blue triangle, in which case they colour it red. The authors of [8] show that this colouring will only fail to avoid monochromatic triangles if a copy of F_0 (Figure 1) or K_4 appears, which a.a.s. does not happen with $\ll n^{4/3}$ rounds/edges. Similarly, one can adjust our proof presented above to show that Theorem 1.2 (b) is tight in the following sense:

When $M_1 = cn^{4/3}$ for some $c > 0$ and $M_2 \ll n^{4/3}$, then a.a.s. there is a red/blue-colouring of G_1 that can be extended to the edges of G_2 avoiding monochromatic copies of K_3. Indeed, as in our proof of Proposition 3.2, one can colour G_1 avoiding monochromatic triangles such that every edge not in a triangle is blue. Colouring the second random graph according to the online greedy approach, as in [8], the player will only fail if a copy of F_0 or K_4 appears in $G_1 \cup G_2$, with at least one edge of G_2. Such copies a.a.s. do not exist when $M_2 \ll n^{4/3}$ and so the player a.a.s. succeeds.

3.1 The 0-statement in the upper range

In this section, we prove the lower bound on $q(n; K_3, p)$ in the upper range $n^{-3/5} \ll p \ll n^{-1/2}$ of Theorem 1.4. In this case, the proof follows easily from Proposition 3.2.

Theorem 3.8. *Suppose that $n^{-2/3} \ll p \ll n^{-1/2}$ and $q \ll n^{-6}p^{-8}$ and let $G_1 \sim G_{n,p}$ and $G_2 \sim G_{n,q}$ be independent. Then a.a.s. there is a G_1-measurable colouring $\varphi\colon E(G_1) \to \{\text{red}, \text{blue}\}$ that can be a.a.s. extended to a triangle-free colouring of $G_1 \cup G_2$.*

Proof. By Proposition 3.2, it suffices to show that a.a.s., $G_1 \sim G_{n,p}$ admits a t-good colouring φ with $t = 150n^6p^8$. Firstly, we claim that a.a.s. G_1 contains at most n^6p^8 copies of F^-, for each $F \in \{F_0, F_1, K_4\}$. Indeed, this follows from Lemma 2.5 **(P2)**, which applies since $m(F_0^-) = 4/3$, $m(F_1^-) = 7/5$, and $m(K_4^-) = 5/4$, and the fact that $n^4p^5, n^5p^7 \ll n^6p^8$ for $p \gg n^{-2/3}$. We also have that a.a.s. G_1 is *not* K_3-Ramsey due to Theorem 1.1 (a). It is therefore enough to show that G_1 admits a t-good colouring φ under the assumption that these two asymptotically-almost-sure events occur.

We define φ as follows. Take any triangle-free colouring of G_1 and recolour any edge not in a triangle blue. As we only changed the colour of edges not in triangles, it is clear that φ remains triangle-free; it only remains to show that φ induces at most t copies of C_{rrbb}. This follows from Corollary 3.4, as each copy of C_{rrbb} is contained in some copy of F^- for some $F \in \{F_0, F_1, K_4\}$. Each such copy of some F^- with $F \in \{F_0, F_1, K_4\}$, hosts at most[4] 50 copies of C_{rrbb} and so the number of C_{rrbb} is at most 50 times the number of copies of some F^- with $F \in \{F_0, F_1, K_4\}$. This completes the proof due to our upper bounds on the number of these copies above. □

3.2 The 0-statement in the lower range

In this section, we prove the lower bound on $q(n; K_3, p)$ in the lower range $n^{-2/3} \ll p \ll n^{-3/5}$ of Theorem 1.4. We will again appeal to Proposition 3.2, but now we will be able to show the existence of a t-good colouring of G_1 for the much larger value $t = n^3p^{7/2}$.

Theorem 3.9. *Suppose that $n^{-2/3} \ll p \ll n^{-3/5}$ and $q \ll n^{-3}p^{-7/2}$ and let $G_1 \sim G_{n,p}$ and $G_2 \sim G_{n,q}$ be independent. Then a.a.s., there is a G_1-measurable colouring $\varphi\colon E(G_1) \to \{\text{red}, \text{blue}\}$ that can be a.a.s. extended to a triangle-free colouring of $G_1 \cup G_2$.*

Recall that in the proof of Theorem 3.8, we showed that every copy of C_{rrbb} is contained in a copy of F^-, for some $F \in \{F_0, F_1, K_4\}$, and we used simple upper bounds on the number of copies of F^-. Here, such simple bounds will no longer suffice and we will have to explore how the copies of these fixed graphs interact. Clearly, any singular copy of such an F^- can be coloured so that it avoids both monochromatic triangles *and* copies of C_{rrbb}. Therefore, we are only forced to create copies of C_{rrbb} if these copies of some F^- and the copies of

[4] We state this loose upper bound for simplicity. It is easy to check that there are at most 45 copies of C_4 in any graph with at most 6 vertices, but of course our specific F^- contain many fewer 4-cycles.

triangles interact in certain ways. The following definition captures the subgraphs of K_n that correspond to a collection of interacting copies of K_3, F_0^- and F_1^- (we exclude K_4^- from this list, as it is composed of two interacting triangles).

Definition 3.10. *Let H be the hypergraph with vertex set $E(K_n)$ whose hyperedges are all (3-, 7- or 8-element) sets of edges that form a copy of K_3, F_0^- or F_1^- in K_n. We will call a graph $C \subseteq K_n$ a* collage *if C induces a connected subhypergraph of H. We will denote the collection of collages in K_n by $\mathcal{C}$.*

We also define collages that are *well-behaved* as follows.

Definition 3.11. *We say a collage $C \in \mathcal{C}$ is well-behaved if*

(i) $v(C) \leq \log n$;

(ii) For any subgraph $C' \subseteq C$ such that $C' \in \mathcal{C}$, we have that $e(C')/v(C') < 5/3$.

Moreover, we say that $C \in \mathcal{C}$ is very well-behaved *if, in addition to (i) and (ii), C satisfies the following further condition:*

(iii) C contains no copies of a graph F with $(v_F, e_F) \in \{(4,6),(5,7),(8,12)\}$.

We will reduce Theorem 3.9 to two key lemmas. The first shows that our random graph $G_1 \sim G_{n,p}$ will a.a.s. only contain well-behaved collages.

Lemma 3.12. *Suppose that $n^{-2/3} \ll p \ll n^{-3/5}$ and let $G_1 \sim G_{n,p}$. Then a.a.s. every collage $C \in \mathcal{C}$ such that $C \subseteq G_{n,p}$ is well-behaved.*

Our second lemma asserts that *very* well-behaved collages can be coloured avoiding any copies of C_{rrbb}.

Lemma 3.13. *Every very well-behaved collage $C \in \mathcal{C}$ admits a* very *good colouring.*

We remark that condition (i) of Definition 3.11 is in fact irrelevant here and we will prove that the conclusion of Lemma 3.13 holds for all collages that satisfy conditions (ii) and (iii). Before proving these lemmas, let us see how they imply Theorem 3.9.

Proof of Theorem 3.9. By Proposition 3.2, it suffices to show that a.a.s. G_1 admits a t-good colouring with $t := n^3p^{7/2}$. Let us assume the asymptotically-almost-sure conclusions of Lemma 3.12 and Lemma 2.5 **(P2)** and also that G_1 is not K_3-Ramsey, which happens a.a.s. due to Theorem 1.1 (a).

We colour the edges of G_1 according to the following scheme, where we define $\mathcal{C}(G_1)$ to be the collection of collages $C \in \mathcal{C}$ such that $C \subseteq G_1$:

1. Colour all maximal subgraphs $C \in \mathcal{C}(G_1)$ which are very well-behaved with a very good colouring (this is possible, due to Lemma 3.13).

2. Colour all the other maximal subgraphs in $\mathcal{C}(G_1)$ in a triangle-free way, such that all edges not in a triangle are blue (this is possible as G_1 is not K_3-Ramsey).

We claim that the resulting colouring is good. Firstly, note that all edges of G_1 are indeed coloured as every edge of G_1 lies in some maximal collage contained in G_1 (the collage may just consist of the single edge). Clearly, all edges not in a triangle are coloured blue. Moreover, as every triangle of G_1 lies in some maximal collage, the resulting colouring contains no monochromatic triangles. It remains to show that there are at most t copies of C_{rrbb}.

Corollary 3.4 implies that any copy of C_{rrbb} must lie in some copy of F_1^-, F_0^-, or K_4^-. However, each of these graphs is contained in some maximal collage (they all induce connected subhypergraphs in H) and thus all copies of C_{rrbb} are contained in maximal collages that are not very well-behaved. It thus suffices to show that the total number of 4-cycles that lie in such collages is at most t. By the assumed conclusion of Lemma 3.12, every $C \in \mathcal{C}(G_1)$ is well-behaved, and so if C is not very well-behaved, then it contains some copy of a subgraph F with v_F vertices and e_F edges such that $(v_F, e_F) \in \{(4,6),(5,7),(8,12)\}$. However, the assumed conclusion of Lemma 2.5 **(P2)** implies that there are at most

$$n^4p^6 \log n + \binom{\binom{5}{2}}{7} n^5p^7 \log n + \binom{\binom{8}{2}}{12} n^8p^{12} \log n \leq 2^{10} n^5 p^7 \log n$$

copies of such an F in G_1 (using that $n^4p^6 \ll n^8p^{12} \ll n^5p^7$ here). Therefore, there are at most $2^{10}n^5p^7 \log n$ maximal collages $C \in \mathcal{C}(G)$ that are not very well-behaved. A collage $C \in \mathcal{C}(G)$ contains at most $v(C)^4$ copies of 4-cycles, and each collage $C \in \mathcal{C}(G)$ has at most $\log n$ vertices on account of it being well-behaved. So in total, there are at most $3n^5p^7 + 2^{10}n^5p^7 \cdot \log^5 n \ll n^3p^{7/2}$ copies of C_{rrbb} in our colouring of G, finishing our proof. □

It remains to prove Lemmas 3.12 and 3.13. We begin by proving Lemma 3.12.

Proof of Lemma 3.12. Let $\mathcal{C}_{\text{bad}}$ be the collection of all $C \in \mathcal{C}$ such that either $e(C)/v(C) \geq 5/3$ or $v(C) \geq \log n$. It suffices to show that a.a.s. $G_1 \sim G_{n,p}$ does not contain any subgraphs in $\mathcal{C}_{\text{bad}}$. In fact, we will focus on another family $\mathcal{C}^*$ of (nonempty) subgraphs of K_n with the following properties:

(a) Every set in $\mathcal{C}_{\text{bad}}$ contains some element of $\mathcal{C}^*$.

(b) Every $C^* \in \mathcal{C}^*$ satisfies $e(C^*) \geq 5v(C^*)/3$ or both $v(C^*) \geq \log n$ and $e(C^*) \geq 5v(C^*)/3 - 3$.

(c) For every $5 \leq k \leq n$, there are at most $(2k)^{150}(8n)^k$ graphs $C^* \in \mathcal{C}^*$ with $v(C^*) = k$.

Assuming we can find such a family $\mathcal{C}^*$ of subgraphs, we claim that we are done. Indeed, by (a),

$$\Pr[\exists C \in \mathcal{C}_{\text{bad}} : C \subseteq G_1] \leq \Pr[\exists C^* \in \mathcal{C}^* : C^* \subseteq G_1] \leq \sum_{C^* \in \mathcal{C}^*} p^{e(C^*)}.$$

Moreover, by (b) and (c),

$$\sum_{C^* \in \mathcal{C}^*} p^{e(C^*)} = \sum_{k=5}^{n} \sum_{\substack{C^* \in \mathcal{C}^* \\ v(C^*)=k}} p^{e(C^*)} \leq \sum_{k=5}^{\log n} (2k)^{150}(8n)^k p^{5k/3} + \sum_{k=\log n}^{n} (2k)^{150}(8n)^k p^{5k/3-3} \ll 1,$$

where the last inequality follows from the assumption that $p \ll n^{-3/5}$. We also used that (b) easily implies that there are no $C^* \in \mathcal{C}^*$ with less than 5 vertices.

It remains to define a family $\mathcal{C}^*$ of subgraphs of K_n satisfying conditions (a)–(c) above. Given a collage $C \in \mathcal{C}_{\text{bad}}$, we construct the 'core' of C algorithmically as follows. We fix some order σ on $E(K_n)$ and initiate our algorithm with logs L_V, L_E, L_O and L_D all being empty. Throughout the algorithm, we will have that L_V is a sequence of distinct vertices in $V(C) \subseteq V(K_n)$, L_E is a sequence of distinct edges in $E(C) \subseteq E(K_n)$, L_O is a sequence of positive natural numbers and L_D is a sequence whose each entry indicates a time step $i \geq 0$ and some set of edges $F \subseteq E(C) \subseteq E(K_n)$. We will maintain, at the end of every time step $i \geq 0$ of the algorithm, that the set of vertices in L_V and the set of edges in L_E define a subgraph of C, which we denote as C_i (and so C_0 is the empty graph).

Now, in the first step of the algorithm, we choose some edge $e_1 \in C$, add its endpoints (in an arbitrary order) to L_V and add e_1 to L_E (so that C_1 is the one-edge graph e_1). In every subsequent step $i \geq 2$, we do the following:

- We terminate and output $C^* = C_{i-1}$ if one of the following is true:

$$|L_D| = 7, \qquad |L_V| > \log n \qquad \text{or} \qquad C_{i-1} = C.$$

- Otherwise, since $C_{i-1} \neq C$ and C is a collage, there must be a copy of K_3, F_0^- or F_1^- in C that intersects C_{i-1} in at least one edge but is not fully contained in C_{i-1}. Call such a copy *regular* if it is a copy of F_0^- and it intersects C in a triangle; otherwise, call it *degenerate*. We say that a regular copy of F_0^- is *rooted* at e if e belongs to its intersection with C_{i-1} and it is the edge of the triangle that also participates in the (unique) 4-cycle of F_0^-.
  - If there exist regular copies of F_0^-, then to each copy associate a number $x \in \mathbb{N}$ which is the position in L_E of the edge that the copy is rooted at. Take a copy of F_0^- that minimises this position and add this minimum x to L_O. Update L_V and L_E by appending to it the vertices and the edges of our copy of F_0^- that do not lie in C_{i-1}: the five such edges are added according to their relative order in σ and the three vertices in some canonical order. Note that this defines C_i with $e(C_i) = e(C_{i-1}) + 5$ and $v(C_i) = v(C_{i-1}) + 3$.
  - If there are no regular copies of F_0^-, add some degenerate copy of K_3, F_0^- or F_1^-, append to L_V and to L_E the vertices and the edges of this copy that do not lie in C_{i-1}: the edges are added according their relative order in σ and the vertices (if there are any) in an arbitrary order; this again defines C_i. Finally, detail this degenerate step i by logging it in L_D along with the set of edges in $E(C_i) \setminus E(C_{i-1})$.

Since each step increases L_E by at least one, this algorithm terminates for any $C \in \mathcal{C}_{\text{bad}}$. We may thus define $\mathcal{C}^*$ as the set of all its outputs, that is, $\mathcal{C}^* := \{C^* : C \in \mathcal{C}_{\text{bad}}\}$. This definition guarantees that $\mathcal{C}^*$ satisfies (a) above; we will show that it also satisfies (b) and (c). First though we establish the following key estimate that bounds the distance of $e(C_i)$ from $5v(C_i)/3$ in terms of the number of degenerate steps (equivalently, the size of $|L_D|$) at the end of step i, which we denote by $d(i)$.

Claim 3.14. *For all $i \geq 1$, we have that $d(i) \leq 3e(C_i) - 5v(C_i) + 7 \leq 21d(i)$.*

Proof. Both inequalities hold with equality when $i = 1$ since $e(C_1) = 1$, $v(C_1) = 2$, and $d(1) = 0$. Suppose that $i \geq 2$ and the claim holds for $i-1$. If the ith step is regular, the claim continues to hold since $e(C_i) = e(C_{i-1})+5$, $v(C_i) = v(C_{i-1}) + 3$, and $d(i) = d(i-1)$. If the ith step is degenerate, then there is some $H' = H \cap C_{i-1}$ such that H is a copy of K_3, F_0^- or F_1^- and H' is a proper subgraph of H that contains at least one edge and $(H, H') \neq (F_0^-, K_3)$. In this case we have that $e(C_i) = e(C_{i-1})+e$, $v(C_i) = v(C_{i-1})+v$, where $e := e(H)-e(H')$ and $v := v(H) - v(H')$. For every such H, H', we have $1 \leq 3e - 5v \leq 21$, and so the claim will hold also for i. Indeed the upper bound of $3e - 5v \leq 21$ follows from the fact that $e \leq 7$ as H has at most 8 edges and H' is non-empty. The lower bound $3e - 5v \geq 1$ follows from a simple case analysis considering $v = 1, 2, 3, 4$ and noting that one can assume that H' is an induced subgraph of H. We leave the details to the reader. ◀

Property (b) now follows from Claim 3.14. Indeed, if $C^* = C$ then certainly $e(C^*) \geq 5v(C^*)/3$ as $C \in \mathcal{C}_{\text{bad}}$. If $C^* \neq C$ and $v(C^*) < \log n$, then at the time τ at which the algorithm terminates we have that $C^* = C_{\tau-1}$ and $d(\tau - 1) = 7$ and so $0 \leq 3e(C^*) - 5v(C^*)$ as required. Finally if $v(C^*) > \log n$, the lower bound on $e(C^*)$ also follows from Claim 3.14, using that, trivially, 0 is a lower bound on the number of degenerate steps taken when the algorithm terminates.

It remains to prove property (c) of the collection $\mathcal{C}^*$, so let us fix some $5 \leq k \leq n$. We bound the number of $C^* \in \mathcal{C}^*$ with $v(C^*) = k$ as follows. Firstly, we note that C^* can be completely determined by the logs L_V, L_O

and L_D when the algorithm terminates. That is, we can recover L_E from these logs. Indeed, the first two vertices in L_V determine the first edge in L_E. Now, suppose we have recovered L_E up to time $i-1$ and consider time step i. If the step is not degenerate (which we know from L_D), the new edges of the regular copy of F_0^- are completely determined by the edge that the copy is rooted at, which we know from L_O and our recovery of L_E so far, and the next three vertices, which appear (in a canonical order) in L_V. Hence, we can add the new edges (in order according to the order σ on $E(K_n)$) to L_E and recover L_E up to time i. If, on the other hand, the step i is degenerate, then L_D will signify this and indicate the new edges that need to be added to L_E. Again the order they are added is determined by σ.

This shows that in order to bound the number of C^* with $v(C^*) = k$, it suffices to bound the number of possibilities for the logs L_V, L_O and L_D output by the algorithm with $|L_V| = k$. For the logs L_V, we use the simple upper bound that there are at most n^k choices. For the logs L_O, note first that Claim 3.14 implies that any C^* with $v(C^*) = k$ has

$$e(C^*) \le \frac{5k + 21 \cdot 7}{3} \le 2k + 49,$$

using here that there are at most 7 degenerate steps before the algorithm terminates. Now in each instance of the algorithm, note that L_O is a nondecreasing sequence of numbers bounded by $2k+49$. Moreover, the length of the sequence is the number of regular steps which is certainly less than k and we can append repeated entries with value $2k+50$ to make all the sequences length k. Hence we can bound the number of possible L_O by the number of nondecreasing sequences of length k with elements in $\{1, \ldots, 2k+50\}$, which is $\binom{k+(2k+50)-1}{k} \le 2^{3k+50}$. Finally, each of the at most 7 entries of L_D indicates a step (at most k choices) and a selection of at most 7 edges which lie on vertices in L_V (at most k^2 choices for each edge). Hence there at most $\sum_{j=0}^{7}(k(\sum_{h=1}^{7}(k^2)^h))^j \le k^{140}$ choices for L_D. Combining our estimates of the number of choices of L_V, L_0 and L_D gives (c) and completes the proof of Lemma 3.12. □

Finally, it remains to prove Lemma 3.13, which is the subject of the rest of this section. In order to prove this, we use a novel 'discharging' method, similar in spirit to the method used in the recent work of Friedgut, Kuperwasser, Schacht and the third author [9] in proving sharp thresholds for Ramsey properties.

Proof of Lemma 3.13. Assume the contrary and let $C \in \mathcal{C}$ be a smallest counterexample. We claim that every proper subgraph $D \subsetneq C$ admits a very good colouring. Indeed, let $D = D_1 \cup \cdots \cup D_t$ be the partition of D into maximal collages and note each D_i is also very well-behaved. Since C is a smallest counterexample, each D_i admits a very good colouring. We claim that the union of these colourings is a very good colouring of D. Indeed, every triangle in D is contained in some D_i (as it is a maximal collage) and thus it is not monochromatic. It is clear that every edge not in a triangle is coloured blue. If there was a copy of C_{rrbb} in D, it would lie in some copy of F_0^-, F_1^- or K_4^-, by Corollary 3.4, and thus it would lie in some D_i (as each of F_0^-, F_1^- and K_4^- induces a connected subhypergraph of H), a contradiction.

Our aim is now to remove from C a carefully chosen selection of edges and show that any very good colouring of the remaining subgraph (which exists, from above) can be extended to the removed edges while remaining very good, thus contradicting our assumption that C is a counterexample. In order to find such a *removable* set of edges, we define a discharging procedure which assigns weights to small subgraphs of C that we call a *blocks*.

To define our blocks, notice that condition (iii) of Definition 3.11 implies that C does not contain copies of K_4 and the graphs F_2 and F_3 depicted in Figure 2. This in turn implies that every triangle in C shares edges with at most one other triangle in C. We define our *blocks* $\mathcal{B} = \mathcal{B}(C)$ to be all copies of K_4^- in C and all triangles in C that are not contained in a K_4^-; this definition guarantees that blocks are pairwise edge-disjoint. Fix an ordering σ of $\mathcal{B}$ such that every triangle in $\mathcal{B}$ precedes every copy of K_4^- and assign weights to the blocks in $\mathcal{B}$ as follows:

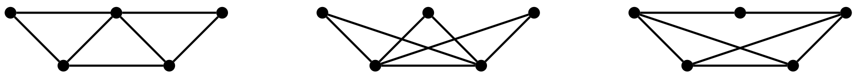

F_2 F_3 F_4

Figure 2: The three graphs constructed from K_4^- by connecting a pair of its vertices by a path of length two.

1. Assign weight 5 to each vertex of C and weight -3 to each edge.
2. For every $v \in V(C)$ contained in exactly one block, send its weight to this block.
3. For every vertex $v \in V(C)$ contained in more than one block, redistribute its weight equally to the two smallest blocks that contain v according to the ordering σ on $\mathcal{B}$. Note that in this step, if v is a vertex in $i \geq 0$ triangles in $\mathcal{B}$, then $i' := \min\{i, 2\}$ triangles containing v and $2 - i'$ copies of K_4^- containing v increase their weight by $5/2$.
4. For every $e \in E(C)$ contained in a block, redistribute its weight to the block containing it (recall that blocks are pairwise edge-disjoint).
5. For every $v \in V(C)$ not contained in a block, redistribute its weight equally to the edges incident to v (note that these edges were not yet handled, since they are also not in a block).
6. For every $e \in E(C)$ not contained in a block, redistribute its weight in the following way: Since e belongs to a copy of F_0^- or F_1^-, at least one of its endpoints must also be a vertex in a block.
   (a) If only one of the endpoints belongs to some block, distribute e's weight equally among all the blocks it belongs to.
   (b) Otherwise, split e's weight equally among its two endpoints and, for each of the endpoints, distribute the weight equally among all its blocks.

Note that by the end of this process, the total weight of all the edges and vertices in C has been redistributed to $\mathcal{B}$, and the total weight remains unchanged. By our assumption that $e(C)/v(C) < 5/3$, the total weight of the vertices and edges before redistribution to blocks was positive, and therefore so is the total weight of all the blocks. Therefore, C must contain (at least) one positive-weight block $X \in \mathcal{B}$. We will split the argument into two cases, depending on whether X is a copy of K_3 or K_4^-. In each case, we will find a removable set of edges. Before embarking on this, we prove a technical claim that allows us to reason about the weight of X by inspecting the graph C locally, without knowledge of the whole of C.

Claim 3.15. *Suppose that $C' \subseteq C$ satisfy $\mathcal{B}(C') = \mathcal{B}(C) =: \mathcal{B}$ and let $w_C, w_{C'} : \mathcal{B} \to \mathbb{R}$ be the weight assignments defined by the above process on C and C', respectively (with the same order σ on $\mathcal{B}$). Then $w_C(X) \leq w_{C'}(X)$ for each $X \in \mathcal{B}$.*

Proof. Since stages 1–4 depend only on the set of blocks, by the end of stage 4, and $\mathcal{B}(C) = \mathcal{B}(C')$, all the vertices, edges and blocks in C and C' have the same weight. Now let J be the graph spanning all the edges of C that do not lie in a triangle (and so have not been dealt with by the end of Stage 4 of the process). It is enough to show that, after stage 5, the C'-weight of each edge of $J \cap C'$ is at least as large as its C-weight and that the C-weight of each edge of J is at most $-1/2$. This implies the assertion of the claim, as in stage 6, the change in weight of every block depends only on the edges of J and their C-weights are negative and never larger than their C'-weights.

Pick an arbitrary $e \in J$. If both endpoints of e lie in blocks, its weight is -3, in both processes. Otherwise, exactly one endpoint of e does not lie in a block. If we denote this endpoint by v, then, for both $H \in \{C, C'\}$, the H-weight of e at the end of stage 5 is $-3 + 5/d_H(v)$. Since $d_{C'}(v) \leq d_C(v)$ for all $v \in V(C')$ and $d_C(v) \geq 2$ for all $v \in V(C)$, the C-weight of e is at most $-1/2$ and not larger than its C'-weight. ◀

Claim 3.16. *If X is a positive-weight copy of K_3, then one of its edges $e \in E(X)$ is not in a 4-cycle in a copy of F_0^- in C.*

Claim 3.17. *Suppose X is a positive-weight copy of K_4^- with edges e_1, f_1, e_2, f_2, g such that e_i, f_i and g form a triangle for $i = 1, 2$. Then there exists an $i \in \{1, 2\}$ such that neither e_i nor f_i belong to a 4-cycle in a copy of F_0^- in C.*

Before proving these claims, let us see how we can use them to contradict our assumption that C is a minimal counterexample. Firstly, consider the case that our positive-weight block X is a triangle. Let e be an edge of X from the assertion of Claim 3.16. As shown at the beginning of the proof, $C \setminus \{e\}$ has a very good colouring. We may extend this colouring to a very good colouring of C as follows. We colour e blue unless the other two edges of X are coloured blue, in which case we colour e red. As X is the only triangle containing e, the colouring remains triangle-free and every edge of C that is not in a triangle is still coloured blue. Thus, we just need to show that there are no copies of C_{rrbb} in C, see property 3 of Definition 3.1. Suppose that there was such a copy. As the colouring of $C \setminus \{e\}$ is very good, this copy of C_{rrbb} would contain e. Corollary 3.4 would then imply that e lies in the 4-cycle in a copy of K_4^-, contradicting the assumption that X is a block, a copy of F_0^-, contrary to our choice of e, or a copy of F_1^-, contradicting the property (iii) of being very well-behaved.

The case when X is a copy of K_4^- is resolved similarly. Without loss of generality, we can assume that the edges of X are labelled as in Claim 3.17 and neither e_1 nor f_1 belong to a 4-cycle in a copy of F_0^-. As above, $C \setminus \{e_1, f_1\}$ has a very good colouring, which we may extend fo a very good colouring of C as follows. We colour e_1 red and f_1 unless that creates a copy of C_{rrbb} with the edges e_2 and f_2, in which case we colour e_1 blue and f_1 red. We claim that this gives a very good colouring of C. Since the only triangle in C containing e_1 or f_1 is the triangle containing both of them, the colouring remains triangle-free; every edge not in a triangle is still blue. We just need to verify that there are no copies of C_{rrbb}. As in the previous case, we can apply Corollary 3.4 and rule out that e_1 and f_1 are in copies of F_1^- and F_0^- using condition (iii) of being very well-behaved and the key property of e_1 and f_1 coming from Claim 3.17. The only case left to consider then, is that e_1 or f_1 lie in some copy of C_{rrbb} that lies in a copy of K_4^- in C. Since C is $\{K_4, F_2, F_3\}$-free, the only copy of K_4^- containing our copy of C_{rrbb} is X itself, which is impossible, as we coloured X to avoid having a copy of C_{rrbb}.

Now that we have proved that Claims 3.16 and 3.17 contradict the assumption that C is a minimal counterexample, it remains only to prove these two claims.

Proof of Claim 3.16. Denote $V(X) = \{x, y, z\}$ and suppose towards a contradiction that each of xy, xz and yz belongs to a 4-cycle in some copy of F_0^-. To get a contradiction, due to Claim 3.15, it suffices to show that there is some $C' \subseteq C$ such that $\mathcal{B}(C) = \mathcal{B}(C')$ and $w_{C'}(X) \leq 0$.

We begin by considering C' to be the union of all the triangles in C (so that $\mathcal{B}(C) = \mathcal{B}(C')$). By our assumption, each edge of X must share a vertex with a triangle other than X; indeed, otherwise it cannot lie on a 4-cycle in a copy of F_0^-. Consequently, at least two of X's vertices, say y and z, are also in other triangles and hence blocks. Therefore, their contribution to X's weight (in stage 3 of the weight redistribution process) is at most $5/2$ each, and so $w_{C'}(X) \leq 1$. We may further assume that x does not lie in an additional triangle, since otherwise $w_{C'}(X) \leq -3/2$, see Figure 3.

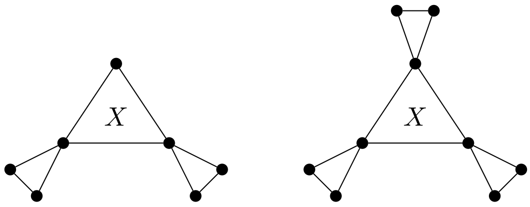

(i) $w_{C'}(X) \leq 1$ *(ii)* $w_{C'}(X) \leq -3/2$

Figure 3: Triangle configurations on the vertices of $X = K_3$.

Now, since both xy and xz are in some copies of F_0^- and X is the only triangle that x belongs to, C must have an edge xv that does not lie in a block; we now add xv to C'. If v supported a triangle (and hence a block) in C, then xv would contribute $-3/2$ to $w_{C'}(X)$ in stage 6 of the weight redistribution process and therefore $w_{C'}(X) \leq -1/2$ would be negative (see Figure 4 for illustration). We may thus further assume that v does not belong to a triangle in C. Now, add to C' all edges in copies of F_0^- that that contain one of xy, xz. If xv lied in a 4-cycle containing xy and in a 4-cycle containing xz, we would have that $d_{C'}(v) \geq 3$ and so xv would have weight at most $-3 + 5/3 = -4/3$ after stage 5, all of which would go to X in stage 6, and once again $w_{C'}(X) \leq -1/3$ would be negative. Thus, the 4-cycles containing xy and xz do not share edges there must be a vertex $u \neq v$ such that $xu \in C'$. The contribution of each of xu and xv to X's weight is at most $-1/2$ and therefore $w_{C'}(X) \leq 0$, a contradiction. This concludes the proof of Claim 3.16. ◀

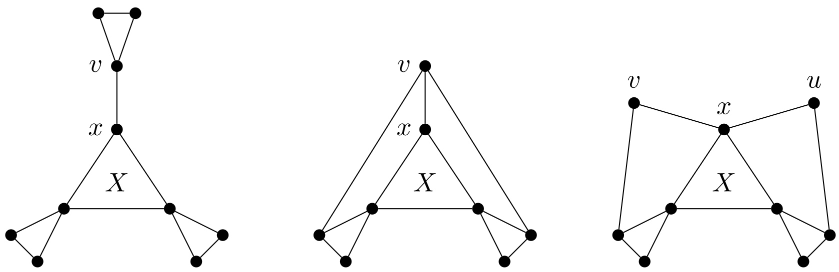

(i) $w_{C'}(X) \leq -1/2$ *(ii)* $w_{C'}(X) \leq -1/3$ *(iii)* $w_{C'}(X) \leq 0$

Figure 4: Edge configurations on the third vertex of $X = K_3$.

Proof of Claim 3.17. Denote $V(X) = \{x_1, x_2, y, z\}$ so that $g = yz$ and $e_i = x_i y, f_i = x_i z$ for $i \in \{1, 2\}$. Assume towards contradiction that one of e_1, f_1 as well as one of e_2, f_2 belongs to a 4-cycle in a copy of F_0^-. We let C' be the union of all triangles in C, so that $\mathcal{B}(C') = \mathcal{B}(C')$. By Claim 3.15, it is enough to show that $w_{C'}(X) \leq 0$.

Now, note that stage 4 of the redistribution process on C' moves weight from the edges of X to X. If two or more vertices of X belonged to blocks other than X, then the contribution to X's weight coming from its vertices, in stages 2 and 3, would be at most $2 \cdot 5 + 2 \cdot (5/2) \leq 15$ and we would have $w_{C'}(X) \leq 0$. Hence, it must be the case that at most one vertex in X belongs to a block that is not X. Moreover, if none of y, z, x_i belonged to a block other than X, then none of e_{3-i}, f_{3-i} would be contained in a 4-cycle of a copy of F_0^-. Consequently, one of y, z must belong to a block other than X; without loss of generality, assume that y is the only vertex of X contained in a block other than X. Since neither f_1 nor f_2 can lie in a 4-cycle of a copy of F_0^-, it must be that both e_1 and e_2 do.

For each $i \in \{1, 2\}$, denote by C_i the 4-cycle in a copy of F_0^- that passes through e_i. We claim that e_i is the only edge of X in C_i. Indeed, the union of a copy of K_4^- and a copy of F_0^- whose 4-cycle intersects this K_4^- in more than one edge contains a copy of F_2, F_3 or F_4 depicted in Figure 2. However, C cannot contain any of

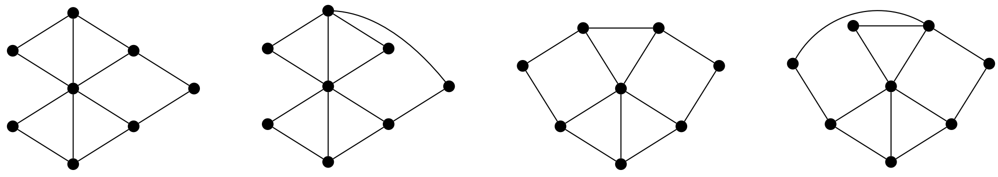

Figure 5: The four graphs appearing in $X \cup X_1' \cup X_2' \cup C_1 \cup C_2$, each of which has 8 vertices and 12 edges.

these graphs, see condition (iii) in Definition 3.11. Further, the second edge of C_i that is incident with y belongs to some block $X_i' \neq X$. We claim that X_i' is not a copy of K_4^-. Indeed, if it were, then $X \cup X_i' \cup C_i$ would be a copy of either F_5 or F_5' from Figure 5 (recall that C_i is not allowed to intersect a copy of K_4^- in more than one edge) and this graph is too dense to be contained in C, see condition (iii) in Definition 3.11. Therefore, both X_1' and X_2' are triangle blocks. Finally, if $X_1' \neq X_2'$, then, since the ordering σ prioritises triangles over copies of K_4^-, the vertex y would give none of its weight to X in stage 3 of the redistribution process, yielding $w_{C'}(X) \leq 0$, as desired. Thus $X' \coloneqq X_1' = X_2'$ is a triangle block. Moreover, each C_i shares only one edge with X', as otherwise $X' \cup C_i$ would be a copy of K_4^-, which is impossible due to the assumption that X' is a triangle block. Consequently, each C_i contains a unique vertex $w_i \notin V(X) \cup V(X')$. If $w_1 = w_2$, then $X \cup C_1 \cup C_2$ contains a copy of F_4, which is too dense to be contained in C, so we may assume that $w_1 \neq w_2$. But then, $X \cup X' \cup C_1 \cup C_2$ would be a copy of one of F_6 or F_6' from Figure 5, a contradiction. ◀

This concludes the proof of Lemma 3.13. □

4 Proof of the 1-statements

In this section, we prove our 1-statements, establishing the upper bounds on $q(n; K_3, p)$ in Theorem 1.4. Our aim is to prove that, if $q \gg q(n; K_3, p)$, then a.a.s. *no* K_3-free colouring of $G_{n,p}$ can be extended to the edges of an independent copy $G_{n,q}$ without creating monochromatic triangles. We will achieve this by showing that every K_3-free colouring of the edges of a typical $G_{n,p}$ results in many local obstructions – individual edges or copies of $K_{1,2}$ what one cannot colour without introducing a monochromatic triangle, see Figure 6. More precisely, we will show that there are either $\omega(q^{-1})$ such *dangerous edges* or $\omega(q^{-2})$ such *dangerous copies of* $K_{1,2}$. Standard probabilistic arguments will then show that a.a.s. at least one such local obstruction will appear in $G_{n,q}$, precluding the existence of a K_3-free extension of our colouring of $G_{n,p}$. In fact, with just a little more work, it will be enough for us to find either $\omega(q^{-1})$ copies of C_{rrbb} or $\omega(q^{-2})$ copies of C_{rbbbb}.

Proposition 4.1. *Suppose that* $n^{-2/3} \ll p \ll n^{-1/2}$, $t \geq n^7 p^{10}$ *and* $t^{-1} \ll q < 1$ *and let* $G_1 \sim G_{n,p}$ *and* $G_2 \sim G_{n,q}$ *be independent. Then a.a.s. any* G_1*-measurable colouring* $\varphi \colon E(G_1) \to \{\text{red}, \text{blue}\}$ *that contains at least* t *copies of* C_{rrbb} *cannot be extended to a* K_3*-free colouring of* $G_1 \cup G_2$.

Proposition 4.2. *Suppose that* $n^{-2/3} \ll p \ll n^{-1/2}$, $t \geq n^7 p^9$ *and* $t^{-1/2} \ll q \leq 1$ *and let* $G_1 \sim G_{n,p}$ *and* $G_2 \sim G_{n,q}$ *be independent. Then a.a.s. any* G_1*-measurable colouring* $\varphi \colon E(G_1) \to \{\text{red}, \text{blue}\}$ *that contains at least* t *copies of* C_{rbbbb} *cannot be extended to a* K_3*-free colouring of* $G_1 \cup G_2$.

In order to find the required number of copies of C_{rrbb} or C_{rbbbb}, we will use three different arguments. We first split our analysis depending on the structure of the colouring. If the colouring is *balanced*, in that a positive proportion of the edges of $G_{n,p}$ are coloured each colour, then a condition of $q \gg n^{-4}p^{-4}$ already guarantees the existence of the required number of C_{rrbb}s. We prove this in Section 4.1 using the method of hypergraph containers (Theorem 2.4). Noting that $n^{-4}p^{-4} \ll n^{-6}p^{-8} \ll n^{-3}p^{-7/2}$ in our full range $n^{-2/3} \ll p \ll n^{-1/2}$, this settles the desired result for balanced colourings in both the lower and upper ranges.

It thus remains to consider colourings that are *unbalanced*, that is, when there is one colour, say red, that appears only $o(n^2p)$ times. Here, we need more delicate arguments based on Janson's inequality and careful union bounds over all unbalanced colourings. In Section 4.2, we show that every unbalanced colouring contains $\Omega(n^6p^7)$ copies of C_{rbbbb}, which implies, by Proposition 4.2, that no such colouring can be extended to a typical copy of $G_{n,q}$ as soon as $q \gg n^{-3}p^{-7/2}$. Combining this result with the case of balanced colourings covers all possible colourings and shows that $q(n; K_3, p) \le n^{-3}p^{-7/2}$ in the full range of interest $n^{-2/3} \ll p \ll n^{-1/2}$. Finally, in Section 4.3, we improve on this in the upper range, when $p \gg n^{-3/5}$, showing that every unbalanced colouring contains $\Omega(n^6p^8)$ copies of C_{rrbb}, which renders any unbalanced colouring nonextendable to $G_{n,q}$ as soon as $q \gg n^{-6}p^{-8}$, see Proposition 4.1. Here, in order to perform a union bound over all unbalanced, K_3-free colourings of $G_{n,p}$, we face some serious technicalities when p approaches $n^{-3/5}$.

We complete this lengthy introduction with proofs of Propositions 4.1 and 4.2 that supply sufficient conditions on nonextendability of colourings in terms of the number of copies of C_{rrbb} and C_{rbbbb}.

Proof of Proposition 4.1. Fix some G_1 satisfying property **(P3)** of Lemma 2.5 (which occurs a.a.s.) and some colouring $\varphi\colon E(G_1) \to \{\text{red}, \text{blue}\}$ that contains at least t copies of C_{rrbb}. Recall that a pair of vertices $\{x, y\} \in E(K_n)$ is *dangerous* if it is the 'colour-splitting' diagonal of at least one such C_{rrbb}, that is, if there exist $u, v \in V(G_1) \setminus \{x, y\}$ with $xu, yu, xv, yv \in E(G_1)$, $\varphi(xu) = \varphi(yu) = \text{blue}$ and $\varphi(xv) = \varphi(yv) = \text{red}$, see Figure 6. We will show that there are at least $t/50$ dangerous pairs. Note that this immediately implies the assertion of the lemma. Indeed, the number of dangerous pairs that appear in $G_{n,q}$ is bounded from below by a $\text{Bin}(t/50, q)$, which is positive with probability $1 - o(1)$, by our assumption that $tq \gg 1$.

Let D be the collection of dangerous pairs. For each pair $\rho \in D$, let $r_\rho \ge 1$ be the number of red copies of $K_{1,2}$ in G_1 that form a triangle with ρ and likewise let b_ρ be the number of blue copies of $K_{1,2}$ in G_1 that form a triangle with ρ, so that ρ is the colour-splitting diagonal for $r_\rho b_\rho$ copies of C_{rrbb} in G_1 and $\sum_{\rho \in D} r_\rho b_\rho \ge t$. We further say that $\rho \in D$ is *heavy* if $r_\rho b_\rho \ge 25$ and let $D_H \subseteq D$ be the collection of heavy dangerous pairs. Now, for each heavy pair ρ, by the AM-GM inequality, we have that $r_\rho + b_\rho \ge 2\sqrt{r_\rho b_\rho} \ge 10$ and ρ forms the part of size 2 in $\binom{r_\rho + b_\rho}{10}$ copies of $K_{2,10}$. Therefore, by the assumed conclusion of Lemma 2.5 **(P3)**,

$$\sum_{\rho \in D_H} \frac{r_\rho b_\rho}{25} \le \sum_{\rho \in D_H} \left(\frac{r_\rho + b_\rho}{10}\right)^2 \le \sum_{\rho \in D_H} \left(\frac{r_\rho + b_\rho}{10}\right)^{10} \le \sum_{\rho \in D_H} \binom{r_\rho + b_\rho}{10} \le N_{K_{2,10}}(G_1) \le n^{11}p^{18} \le \frac{t}{50},$$

where in the last inequality we used that $n^{11}p^{18} \ll n^7p^{10} \le t$ due to the fact that $p \ll n^{-1/2}$. Hence $\sum_{\rho \in D \setminus D_H} r_\rho b_\rho \ge t/2$ and, as each $\rho \in D \setminus D_H$ has $r_\rho b_\rho \le 25$, we indeed obtain $|D| \ge |D \setminus D_H| \ge t/50$. □

Proof of Proposition 4.2. Fix some G_1 satisfying properties **(P1)** and **(P3)** of Lemma 2.5 (which occur a.a.s.) and some colouring $\varphi\colon E(G_1) \to \{\text{red}, \text{blue}\}$ that contains at least t copies of C_{rbbbb}. A copy K of $K_{1,2}$ in G_1 with vertices w, u_1, u_2 (so that K is formed from edges wu_i for $i = 1, 2$) is *dangerous* if there are distinct vertices $w_1, w_2 \in V(G) \setminus \{w, u_1, u_2\}$ such that $u_1u_2, u_1w_1, w_1w, ww_2, w_2u_2 \in E(G_1)$, $\varphi(u_1u_2) = \text{red}$ and $\varphi(u_1w_1) = \varphi(w_1w) = \varphi(ww_2) = \varphi(w_2u_2) = \text{blue}$. We say that K *hosts* this copy of C_{rbbbb} on vertices u_1, u_2, w_2, w, w_1. See Figure 6 for a depiction.

Moreover, for such a dangerous copy K of $K_{1,2}$, we define x_1^K to be the number of choices of w_1 such that $\varphi(u_1w_1) = \varphi(w_1w) = \text{blue}$ and x_2^K to be the number of choices of w_2 such that $\varphi(u_2w_2) = \varphi(w_2w) = \text{blue}$. Therefore, each K hosts at most $x_1^K x_2^K$ copies of C_{rbbbb} (this is not equality as some choices could have $w_1 = w_2$). Taking $\mathcal{K}$ to be the collection of dangerous copies of $K_{1,2}$ on $V(G)$, we then have that $\sum_{K \in \mathcal{K}} x_1^K x_2^K \ge t$. As in the proof of Proposition 4.1, our aim is to prove that $\mathcal{K}$ is large. Define a copy K to be *heavy* if

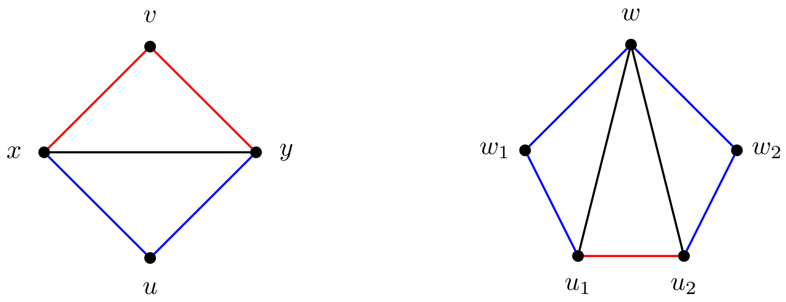

Figure 6: A dangerous edge and a dangerous copy of $K_{1,2}$ in G.

$m^K := \max\{x_1^K, x_2^K\} \geq 10$ and let $\mathcal{K}_H$ be the collection of heavy dangerous copies of $K_{1,2}$. We then have that

$$\sum_{K\in\mathcal{K}_H} x_1^K x_2^K \leq \sum_{K\in\mathcal{K}_H} (m^K)^2 \leq \sum_{K\in\mathcal{K}_H} 100\left(\frac{m^K}{10}\right)^{10} \leq 100\sum_{K\in\mathcal{K}_H}\binom{m^K}{10}.$$

We claim that the sum in the right-hand side of the above inequality is at most $N_{K^+_{2,10}}(G_1)$, where $K^+_{2,10}$ is the graph obtained from $K_{2,10}$ by adding a pendant edge to one of its vertices of degree 10. Indeed, fixing some K in the summand, label the vertices u_1, u_2 and w as we did above and suppose that m^K is achieved by x_i^K for $i \in [2]$ (if $x_1^K = x_2^K$ then choose i arbitrarily). Then, for every set W_i of 10 vertices that can play the role of w_i in the sense that they are all connected to both u_i and w by blue edges in G_1, we get a copy of $K^+_{2,10}$ on the vertices w, u_1, u_2 and W_i. This gives $\binom{m^K}{10}$ copies of $K^+_{2,10}$ in the summand corresponding to K; these copies are distinct, as each copy of $K^+_{2,10}$ determines K completely. So we have that

$$\sum_{K\in\mathcal{K}_H} x_1^K x_2^K \leq 100 N_{K^+_{2,10}}(G_1) \leq 100\cdot n^{11}p^{18}\cdot 2np \leq 200n^{12}p^{19} \leq t/2,$$

where we used properties **(P1)** and **(P3)** of Lemma 2.5 to bound $N_{K^+_{2,10}}(G_1)$ and we used that $n^{12}p^{19} \ll n^7p^9 \leq t$ in the last inequality. This implies that $\sum_{K\in\mathcal{K}\setminus\mathcal{K}_H} x_1^K x_2^K \geq t/2$ and as every $K \in \mathcal{K}\setminus\mathcal{K}_H$ has $x_1^K x_2^K \leq 9^2 \leq 100$, we have that $|\mathcal{K}| \geq |\mathcal{K}\setminus\mathcal{K}_H| \geq t/200$.

Now, taking $G_2 \sim G_{n,q}$, for each dangerous copy $K \in \mathcal{K}$, let $I_K := \mathbb{1}[K \subseteq G_2]$ be the indicator random variable for the event that both edges of K appear in G_2 and let $X := \sum_{K\in\mathcal{K}} I_K$. As each dangerous copy of $K_{1,2}$ appears with probability q^2, we have that $\mu := \mathbb{E}[X] \geq tq^2/200 \gg 1$. Moreover, writing $K \sim K'$ when a pair K, K' of dangerous copies of $K_{1,2}$ share at least one edge, we have

$$\Delta := \sum_{K\sim K'} \mathbb{E}[I_K I_{K'}] \leq \mu\cdot\big(1 + 4\Delta(G_1)q\big) \leq 2\max\{\mu, 8npq\mu\},$$

using again the assumed conclusion of Lemma 2.5 **(P1)** in G_1. Indeed, given some copy K of $K_{1,2}$, we can obtain an upper bound on the number of dangerous K' that intersect K (but are not equal to K) by the number of choices of an edge e of K and a G_1-neighbour of one of the endpoints of e. Using that

$$\frac{\mu^2}{npq\mu} \geq \frac{tq}{200np} \geq \frac{t^{1/2}}{200np} \geq \frac{n^{5/2}p^{7/2}}{200} \gg 1,$$

it follows from Janson's inequality (Lemma 2.2) that

$$\Pr[X = 0] \leq \exp\left(-\frac{\mu^2}{2\Delta}\right) \ll 1.$$

Therefore, a.a.s. there is a dangerous copy of $K_{1,2}$, on vertices w, u_1, u_2 say, such that $E(K) = \{wu_1, wu_2\} \subseteq E(G_2)$. This precludes the possibility of extending φ to G_2. Indeed, as K is dangerous, it hosts some copy of C_{rbbbb} in G (under φ). If either wu_1 or wu_2 are coloured blue, then they will form a blue triangle with edges in the copy of C_{rbbbb} whilst if they are both red then there is a red triangle formed with the edge u_1u_2. This completes the proof. □

4.1 Balanced colourings

In this section, we prove the following proposition which deals with balanced colourings of $G_{n,p}$ in our full range of interest.

Proposition 4.3. *For every $\beta > 0$, there exists a $\lambda > 0$ such that the following holds. Suppose that $n^{-2/3} \ll p \ll n^{-1/2}$ and let $G \sim G_{n,p}$. Then, a.a.s. every colouring $\varphi\colon E(G) \to \{\text{red}, \text{blue}\}$ such that $|\varphi^{-1}(c)| \geq \beta n^2 p$ for both $c \in \{\text{red}, \text{blue}\}$, contains at least $\lambda n^4 p^4$ copies of C_{rrbb}.*

As $n^4p^4 \gg n^6p^8$ for $p \ll n^{-1/2}$, Propositions 4.1 and 4.3 give that, for $q \gg n^{-4}p^{-4}$, a.a.s. no balanced K_3-free colouring of $G_{n,p}$ with $n^{-2/3} \ll p \ll n^{-1/2}$ can be extended to the edges of $G_{n,q}$ without creating monochromatic triangles. Before embarking on the proof of Proposition 4.3, we need a deterministic lemma that deals with the complete graph, that is, the case $p = 1$ in the proposition.

Lemma 4.4. *For every $\beta > 0$, there exists a λ such that the following holds. For all sufficiently large n, every $\psi\colon E(K_n) \to \{\text{red}, \text{blue}\}$ that satisfies $|\psi^{-1}(c)| \geq \beta n^2$ for both $c \in \{\text{red}, \text{blue}\}$ contains at least λn^4 copies of C_{rrbb}.*

Proof. For every ordered pair x, y of distinct vertices, let $V_{x,y}$ denote the number of z such that $\psi(xz) = \text{red}$ and $\psi(yz) = \text{blue}$. Letting C be the number of copies of C_{rrbb}, we have, by convexity,

$$C = \sum_{x,y} \binom{V_{x,y}}{2} \geq n(n-1) \cdot \binom{\bar{V}}{2},$$

where

$$\bar{V} = \frac{1}{n(n-1)} \cdot \sum_{x,y} V_{x,y}.$$

Observe that, if a, b, c, d are four distinct vertices such that $\psi(ab) = \text{red}$ and $\psi(cd) = \text{blue}$, then either ab and bc or bc and cd are counted by $V_{a,c}$ or $V_{b,d}$, respectively. This implies that, for all large n,

$$\bar{V} \geq \frac{1}{n(n-1)} \cdot \frac{\beta n^2 \cdot (\beta n^2 - 2n)}{2n} \geq \frac{\beta^2 n}{4},$$

which gives the claimed lower bound on C. □

We now turn to proving Proposition 4.3

Proof of Proposition 4.3. We can assume that $0 < \beta < 1/4$. Let $\mathcal{H}$ be the 4-uniform hypergraph with vertex set $E(K_n) \times \{\text{red}, \text{blue}\}$ whose edges are all copies of C_{rrbb} in K_n, that is, sets of the form

$$\big\{(uv, \text{red}), (uw, \text{red}), (vx, \text{blue}), (wx, \text{blue})\big\},$$

where u, v, w, and x are any four distinct vertices of K_n. Observe that

$$v(\mathcal{H}) = 2\binom{n}{2}, \quad e(\mathcal{H}) = 12 \cdot \binom{n}{4}, \quad \Delta_2(\mathcal{H}) \leq n, \quad \Delta_3(\mathcal{H}) = \Delta_4(\mathcal{H}) = 1.$$

Let $\varepsilon := \lambda_{4.4}(\beta/2)/2$ and let $\delta > 0$ and $t \in \mathbb{N}$ be the constants provided by Theorem 2.4 invoked with $k = 4$ and $K = 2$. Set

$$\gamma := \min\left\{\varepsilon, \frac{\beta}{32}\right\}, \qquad \sigma := \min\left\{\frac{\beta}{4t}, \frac{\gamma^2}{2^t}\right\}, \qquad \text{and} \qquad \lambda := \frac{\delta\sigma^4}{2}.$$

Let $\mathcal{I}(\mathcal{H})$ be the family of all sets $I \subseteq V(\mathcal{H})$ that induce fewer than $\lambda n^4 p^4$ edges of $\mathcal{H}$. Since $p \gg n^{-2/3}$, we may apply Theorem 2.4 to $\mathcal{H}$ with $\tau := \sigma p$ to obtain a function $f\colon \mathcal{P}(V(\mathcal{H}))^t \to \mathcal{P}(V(\mathcal{H}))$ such that:

(i) For every $I \in \mathcal{I}(\mathcal{H})$, there are $S_1, \ldots, S_t \subseteq V(\mathcal{H})$ each of size at most $\tau v(\mathcal{H})$ such that $S_1 \cup \cdots \cup S_t \subseteq I \subseteq f(S_1, \ldots, S_t)$.

(ii) For each $S_1, \dots, S_t \subseteq V(\mathcal{H})$, the set $f(S_1, \dots, S_t)$ induces fewer than εn^4 edges in $\mathcal{H}$.

Suppose that φ is a *bad* colouring of G, that is, a colouring with $|\varphi^{-1}(c)| \ge \beta n^2 p$ for both $c \in \{\text{red}, \text{blue}\}$ but fewer than $\lambda n^4 p^4$ copies of C_{rrbb}. Then $\varphi \in \mathcal{I}(\mathcal{H})$ and thus $\varphi \subseteq f(S_1, \dots, S_t)$ for some $S_1, \dots, S_t \subseteq \varphi$. We will call such $\mathbf{S} := (S_1, \dots, S_t)$ the *signature* of φ and denote it by $\operatorname{sig}(\varphi)$. Denote by $\pi\colon V(\mathcal{H}) \to E(K_n)$ the projection to the first coordinate and, with slight abuse of notation, let $\pi(\mathbf{S}) := \pi(S_1 \cup \dots \cup S_t)$; note that $\pi(\operatorname{sig}(\varphi)) \subseteq G$.

Claim 4.5. *For every $S_1, \dots, S_t \subseteq V(\mathcal{H})$, letting $\mathbf{S} := (S_1, \dots, S_t)$, we have*

$$\Pr\big[G \textit{ has a bad colouring } \varphi \textit{ with } \operatorname{sig}(\varphi) = \mathbf{S}\big] \le \Pr\big[\pi(\mathbf{S}) \subseteq G\big] \cdot \exp(-\gamma n^2 p).$$

Proof. Suppose that G has a bad colouring φ with $\operatorname{sig}(\varphi) = \mathbf{S}$. This means, in particular, that $\pi(\mathbf{S}) \subseteq G$, so it is enough to show that

$$\Pr\big[G \text{ has a bad colouring } \varphi \text{ with } \operatorname{sig}(\varphi) = \mathbf{S} : \pi(\mathbf{S}) \subseteq G\big] \le \exp(-\gamma n^2 p).$$

Define

$$\begin{aligned} R(\mathbf{S}) &:= \big\{e \in E(K_n) : (e, \text{red}) \in f(\mathbf{S})\big\}, \\ B(\mathbf{S}) &:= \big\{e \in E(K_n) : (e, \text{blue}) \in f(\mathbf{S})\big\}, \\ X(\mathbf{S}) &:= E(K_n) \setminus \big(R(\mathbf{S}) \cup B(\mathbf{S})\big) \end{aligned}$$

and observe that $\varphi \subseteq f(\mathbf{S})$ means that G is disjoint from $X(\mathbf{S})$ and that

$$\varphi^{-1}(\text{red}) \subseteq R(\mathbf{S}) \cap G \qquad \text{and} \qquad \varphi^{-1}(\text{blue}) \subseteq B(\mathbf{S}) \cap G.$$

We claim that at least one of the following must be true:

(a) The set $X(\mathbf{S})$ has at least εn^2 edges.

(b) One of the sets $R(\mathbf{S})$ or $B(\mathbf{S})$ has at most $\beta n^2/2$ edges.

Suppose that (b) does not hold and let $\psi\colon E(K_n) \to \{\text{red}, \text{blue}\}$ be an arbitrary colouring of K_n satisfying $\psi^{-1}(\text{red}) \subseteq R(\mathbf{S}) \cup X(\mathbf{S})$, $\psi^{-1}(\text{blue}) \subseteq B(\mathbf{S}) \cup X(\mathbf{S})$ and $|\psi^{-1}(c)| \ge \beta n^2/2$ for each $c \in \{\text{red}, \text{blue}\}$; such a colouring exists as $\lceil \beta n^2/2 \rceil \le \lfloor \binom{n}{2}/2 \rfloor$ due to our upper bound on β. It follows from Lemma 4.4 and our definition of ε that ψ has at least $2\varepsilon n^4$ copies of C_{rrbb}. Any such copy corresponds to an edge of $\mathcal{H}[f(\mathbf{S})]$ unless it contains an edge of $X(\mathbf{S})$. However, the number of 4-cycles with an edge of $X(\mathbf{S})$ is at most $X(\mathbf{S}) \cdot n^2$. This implies that $e(\mathcal{H}[f(\mathbf{S})]) \ge 2\varepsilon n^4 - |X(\mathbf{S})| \cdot n^2$, which gives $|X(\mathbf{S})| > \varepsilon n^2$ due to condition (ii) on $f(\mathbf{S})$ from the outcome of Theorem 2.4.

If (a) holds, then

$$\Pr\big[G \cap X(\mathbf{S}) = \emptyset : \pi(\mathbf{S}) \subseteq G\big] \le (1-p)^{|X(\mathbf{S})|} \le \exp(-\varepsilon n^2 p),$$

so we may assume that (b) holds; without loss of generality, $|R(\mathbf{S})| \le \beta n^2/2$. Conditioned on the event that $\pi(\mathbf{S}) \subseteq G$, the distribution of $e\big(R(\mathbf{S}) \cap G\big)$ is stochastically dominated by the random variable $|\pi(\mathbf{S})| + \operatorname{Bin}\big(|R(\mathbf{S})|, p\big)$. Since $|\pi(\mathbf{S})| \le t\tau n^2 \le t\sigma n^2 p \le \beta n^2 p/4$, we have

$$\Pr\big[e\big(R(\mathbf{S}) \cap G\big) \ge \beta n^2 p : \pi(\mathbf{S}) \subseteq G\big] \le \Pr\big[\operatorname{Bin}(\beta n^2/2, p) \ge (3/4)\beta n^2 p\big] \le \exp(-\beta n^2 p/32),$$

by Lemma 2.1. This proves the assertion of the claim. ◀

Using Claim 4.5, we may conclude that

$$\Pr\big[G \text{ has a bad colouring}\big] \le \sum_{\mathbf{S}} \Pr\big[G \text{ has a bad colouring } \varphi \text{ with } \operatorname{sig}(\varphi) = \mathbf{S}\big]$$
$$\le \exp(-\gamma n^2 p) \cdot \sum_{\mathbf{S}} \Pr\big[\pi(\mathbf{S}) \subseteq G\big].$$

Finally, since there are at most $2^{t|U|}$ sequences $\mathbf{S} := (S_1, \ldots, S_t)$ satisfying $\pi(\mathbf{S}) = U$, we have

$$\sum_{\mathbf{S}} \Pr\big[\pi(\mathbf{S}) \subseteq G\big] \le \sum_{u \le t\tau v(\mathcal{H})} \binom{\binom{n}{2}}{u} \cdot 2^{tu} \cdot p^u \le \sum_{u \le \sigma n^2 p} \left(\frac{2^t e n^2 p}{u}\right)^u \le n^2 \left(\frac{2^t e}{\sigma}\right)^{\sigma n^2 p} \le e^{\gamma n^2 p/2},$$

where the penultimate inequality follows from the fact that, for every $a > 0$, the function $u \mapsto (ea/u)^u$ is increasing when $u \in (0, a]$. Hence we have that a.a.s. there are no bad colourings of G, concluding the proof. □

4.2 Unbalanced colourings in the lower range

In this section, we establish the following theorem, proving that $q(n; K_3, p) \le n^{-3}p^{-7/2}$ when $n^{-2/3} \ll p \ll n^{-1/2}$ and hence giving the 1-statement for the lower range in Theorem 1.4.

Theorem 4.6. *Suppose that $n^{-2/3} \ll p \ll n^{-1/2}$ and $q \gg n^{-3}p^{-7/2}$ and let $G_1 \sim G_{n,p}$ and $G_2 \sim G_{n,q}$ be independent. Then, a.a.s. no G_1-measurable K_3-free colouring $\varphi\colon E(G_1) \to \{\text{red}, \text{blue}\}$ can be extended to a K_3-free colouring of $G_1 \cup G_2$.*

This theorem will follow from the following proposition which deals with unbalanced colourings.

Proposition 4.7. *There exist $\beta, \zeta > 0$ such that the following holds. Suppose that $n^{-2/3} \ll p \ll n^{-1/2}$ and let $G \sim G_{n,p}$. Then, a.a.s. every K_3-free $\varphi\colon E(G) \to \{\text{red}, \text{blue}\}$ such that $\left|\varphi^{-1}(\text{red})\right| < \beta n^2 p$ results in at least $\zeta n^6 p^7$ copies of C_{rbbbb}.*

Indeed, with Proposition 4.7 and our previous results, Theorem 4.6 follows readily.

Proof of Theorem 4.6. Let $\beta, \zeta > 0$ be the constants from the statement of Proposition 4.7. Further, let $\lambda > 0$ be the constant output by Proposition 4.3 with input β and let $t_1 := \lambda n^4 p^4 \ge n^7 p^{10}$ and $t_2 := \zeta n^6 p^7 \ge n^7 p^9$. Now fixing $G_1 \sim G_{n,p}$ and $G_2 \sim G_{n,q}$, we have that a.a.s. the conclusions of Propositions 4.1 with $t_{4.1} = t_1$, 4.2 with $t_{4.2} = t_2$, 4.3 and 4.7 all hold. We claim that this implies the theorem. Indeed, consider some G_1-measurable K_3-free colouring $\varphi\colon E(G_1) \to \{\text{red}, \text{blue}\}$. Suppose first that $\left|\varphi^{-1}(c)\right| \ge \beta n^2 p$ for both $c \in \{\text{red}, \text{blue}\}$. By the assumed conclusion of Proposition 4.3, there are at least t_1 copies of C_{rrbb} induced by φ. Since $q \gg n^{-3}p^{-7/2} \gg t_1^{-1}$, the assumed conclusion of Proposition 4.1 gives that φ cannot be extended to G_2 whilst avoiding monochromatic triangles. Likewise, if $|\varphi^{-1}(\text{red})| < \beta n^2 p$, then Proposition 4.7 gives that there are at least t_2 copies of C_{rbbbb} induced by φ and Proposition 4.2 then gives that we cannot extend φ to G_2 without getting monochromatic triangles, using that $q \gg t_2^{-1/2}$. Since both colours play symmetric roles, the same conclusion holds under the assumption $|\varphi^{-1}(\text{blue})| < \beta n^2 p$. This covers all colourings and completes the proof. □

It remains to prove Proposition 4.7. Our proof works by taking a union bound over all possibilities T for the red subgraph. For each T, we use Janson's inequality (Lemma 2.2) to prove that it is very unlikely that we avoid creating many C_{rbbbb} when we colour T red and $G \setminus T$ blue. This simple approach almost works – it turns out that in order to get strong enough error probabilities in the Janson argument, we need to consider only red subgraphs T that are well behaved, in that they satisfy a maximum degree condition. Before embarking on the proof of Proposition 4.7, we prove that any red subgraph T that we are interested in contains a large induced subgraph that is well behaved.

Lemma 4.8. *For any $c > 0$, there exists a $\beta > 0$ such that the following holds for all sufficiently large $n \in \mathbb{N}$ and $p = p(n) \in [0, 1]$. Let T be a graph on n vertices such that $e(T) < \beta n^2 p$ and $e(T[U]) \geq 1$ for every $U \subseteq V(G)$ with $|U| \geq \frac{n}{2}$. Then, there exists a vertex subset $W \subseteq V(T)$ such that $|W| \geq \frac{n}{2}$ and $S := T[W]$ satisfies $\Delta(S) \leq \frac{cnp}{\log(n^2p)-\log(e(S))}$.*

Proof. Suppose that T satisfies the assumptions of the lemma. We can assume that $0 < c < 1/10$ and we fix some $\varepsilon \in (0, c^2)$ and $\beta \in (0, \varepsilon c/8)$. Consider the following iterative process of peeling off vertices:

Let $T_0 := T$ and $t_0 := e(T_0)$.
for $i = 1, 2, \ldots$ **do**
  if $\Delta(T_{i-1}) \leq \frac{cnp}{\log(n^2p)-\log(t_{i-1})}$ **then**
    Terminate with $S := T_{i-1}$.
  else
    Let v_i be an arbitrary vertex of T_{i-1} with degree exceeding $\frac{cnp}{\log(n^2p)-\log(t_{i-1})}$, let $T_i := T_{i-1} - v_i$ and $t_i := e(T_i)$.

We claim that the process terminates after fewer than $n/2$ steps and thus outputs an appropriate S. Suppose for a contradiction that this is not the case and that the process is still running after $n/2$ steps. We will show that this contradicts our upper bound on $e(T)$. Firstly, note that

$$t_{n/4} \geq \frac{n}{4} \cdot \frac{cnp}{\log(n^2p) - \log(t_{n/2})} \geq \frac{n}{4} \cdot \frac{cnp}{\log(n^2p)} \geq \frac{cn^2p}{2^3 \log n},$$

using here that $t_{n/2} \geq 1$ due to our assumption on T. Now, define $\tau_0 := 0$ and

$$\tau_i := \min \left\{ \tau : t_{n/4-\tau} \geq 2^{i-3} \frac{cn^2p}{\log n} \right\}$$

for $i = 1, 2, \ldots, k := \log_2(\varepsilon \log n)$.

Claim 4.9. *$\tau_k \leq n/4$ (and hence all τ_i are well defined).*

Note that the claim implies that

$$e(T) = t_0 \geq t_{\tau_k} \geq 2^{k-3} \frac{cn^2p}{\log n} = \frac{\varepsilon cn^2p}{2^3} > \beta n^2 p,$$

contradicting our upper bound on $e(T)$. It thus remains to prove the claim.

For this, note that, for each $0 \leq i \leq k-1$ and all $\tau \in \mathbb{N}$ with $\tau \leq n/4 - \tau_i$, we have that

$$t_{n/4-\tau_i-\tau} - t_{n/4-\tau_i} \geq \tau \cdot \frac{cnp}{\log(n^2p) - \log(t_{n/4-\tau_i})} \geq \tau \cdot \frac{cnp}{\log(2^{3-i}c^{-1}\log n)}.$$

Consequently, for $i \in \{0, \ldots, k-1\}$,

$$\tau_{i+1} - \tau_i \leq 2^{i-2} \cdot \frac{cn^2p}{\log n} \cdot \frac{\log\left(2^{3-i}c^{-1}\log n\right)}{cnp} = \frac{2^i n \left(\log\left(2^3 c^{-1}\log n\right) - i\log 2\right)}{4\log n}$$

and so

$$\begin{aligned}
\tau_k &= \tau_0 + \sum_{i=0}^{k-1}(\tau_{i+1} - \tau_i) \\
&\leq \frac{n}{4\log n}\left(\sum_{i=0}^{k-1} 2^i \log\left(2^3c^{-1}\log n\right) - \log 2 \cdot \sum_{i=0}^{k-1} i2^i\right) \\
&\leq \frac{n}{4\log n}\left(2^k \log\left(2^3c^{-1}\log n\right) - \log 2 \cdot (k-2)2^k\right) \\
&= \frac{2^k n}{4\log n}\log(2^{5-k}c^{-1}\log n) \\
&\leq \frac{\varepsilon n}{4}\log\left(2^5c^{-1}\varepsilon^{-1}\right) \leq \frac{n}{4},
\end{aligned}$$

as required, where we used our upper bounds on ε and c in the final inequality. ∎

We now use Lemma 4.8 to establish Proposition 4.7.

Proof of Proposition 4.7. Let θ be the constant from the statement of Lemma 2.5 **(P4)**, let $\zeta = \theta/2^{10}$, let $c = 2^{-17}$ and let $\beta := \beta_{4.8}(c)$ be the constant from the statement of Lemma 4.8. Define $\mathcal{S}$ to be the set of graphs S such that $\theta n^3p^3/8 \leq e(S) < \beta n^2 p$ and $\Delta(S) \leq \frac{cnp}{\log(n^2p)-\log(e(S))}$. Given an $S \in \mathcal{S}$, let $B(S)$ be the event that $G \sim G_{n,p}$ contains fewer than $\zeta n^6 p^7$ copies of C_5 whose vertices all lie in $V(S)$ and that have one edge in S and four edges of $G \setminus S$. The following key claim which bounds the probability of $B(S)$ for all $S \in \mathcal{S}$.

Claim 4.10. *For any $S \in \mathcal{S}$, letting $s = e(S)$, we have that*

$$\Pr[B(S)] \leq \left(\frac{s}{n^2p}\right)^{2s}.$$

Before proving this claim, let us see how it implies the proposition. Firstly, let $\mathcal{F}$ be the family of all graphs T on $V(G)$ that satisfy the following:

(i) $e(T) < \beta n^2 p$;

(ii) for every $U \subseteq V(G)$ with $|U| \geq \frac{n}{2}$, we have that $e(T[U]) \geq \theta |U|^3 p^3$.

Now, if G satisfies property **(P4)** of Lemma 2.5, then every K_3-free colouring $\varphi\colon E(G) \to \{\text{red}, \text{blue}\}$ that colours fewer than $\beta n^2 p$ edges red satisfies $\varphi^{-1}(\text{red}) \in \mathcal{F}$. Indeed, **(P4)** gives a collection of at least $\theta|U|^3p^3$ edge-disjoint triangles in each $U \subseteq V(G)$ with $|U| \geq \frac{n}{2}$ and at least one edge in each triangle must be coloured red. Further, Lemma 4.8 implies that, for every $T \in \mathcal{F}$, there is some $S = S(T) \in \mathcal{S}$ such that $S = T[W]$ for some $W \subseteq V(G)$ with $|W| \geq n/2$; indeed, the fact that $T \in \mathcal{F}$ gives that $e(S) \geq \theta n^3 p^3/8$. This implies that, if there is a K_3-free $\varphi\colon E(G) \to \{\text{red}, \text{blue}\}$ with $|\varphi^{-1}(\text{red})| < \beta n^2 p$ and fewer than $\zeta n^6 p^7$ copies of C_{rbbbb} (and G satisfies property **(P4)** of Lemma 2.5), then there is some $S \in \mathcal{S}$ such that $B(S)$ occurs and $S \subseteq G$. Indeed, $T := \varphi^{-1}(\text{red}) \in \mathcal{F}$ and $B(S(T))$ occurs as otherwise we get at least $\zeta n^6 p^7$ copies of C_5 on $V(S)$ each of which has exactly one edge in S and the other 4 edges in $G \setminus S$ and hence gives a copy of C_{rbbbb} in G. Finally, note that, for each $S \in \mathcal{S}$, the events $B(S)$ and $S \subseteq G$ are independent. Therefore, the probability that G has a colouring $\varphi\colon E(G) \to \{\text{red}, \text{blue}\}$ with $|\varphi^{-1}(\text{red})| < \beta n^2 p$ and fewer than $\zeta n^6 p^7$ copies of C_{rbbbb} is less than

$$\sum_{S \in \mathcal{S}} \Pr[B(S) \wedge S \subseteq G] + \Pr[G \notin \textbf{(P4)}] \leq \sum_{S \in \mathcal{S}} \Pr[B(S)] \cdot \Pr[S \subseteq G] + \Pr[G \notin \textbf{(P4)}].$$

By Lemma 2.5, we have that $\Pr[G \notin \textbf{(P4)}] \ll 1$. We split the sum over $S \in \mathcal{S}$ depending on $e(S) = s$. As there are at most $2^n\binom{\binom{n}{2}}{s}$ graphs $S \in \mathcal{S}$ with s edges (the factor of 2^n bounds the number of choices for $V(S)$), appealing to Claim 4.10, we therefore have that

$$\begin{aligned}
\sum_{S \in \mathcal{S}} \Pr[B(S)] \cdot \Pr[S \subseteq G] &\leq \sum_s 2^n \binom{\binom{n}{2}}{s} \cdot \left(\frac{s}{n^2p}\right)^{2s} \cdot p^s \\
&\leq \sum_s 2^n \cdot \left(\frac{en^2}{2s} \cdot \left(\frac{s}{n^2p}\right)^2 \cdot p\right)^s \\
&\leq \sum_s 2^n \cdot \left(\frac{es}{n^2p}\right)^s \ll 1,
\end{aligned}$$

where the sum goes over all $s \in (\theta n^3 p^3/8, \beta n^2 p)$ and, in the last inequality, we used

$$s \log\left(\frac{n^2p}{s}\right) \geq s \geq \theta n^3 p^3/8 \gg n.$$

Therefore, it remains only to establish Claim 4.10.

Proof of Claim 4.10. Fix some $S \in \mathcal{S}$ and let $s := e(S)$ and $W := V(S)$. We will appeal to Janson's inequality (Lemma 2.2) to obtain the required upper bound on $\Pr[B(S)]$. Let $\Gamma := E(K_n[W]) \setminus S$ and let $\mathcal{C}$ be set of all 5-cycles in $K_n[W]$ comprising of one edge of S and four edges of Γ. For each such $C \in \mathcal{S}$, let I_C be the indicator random variable for the event that $C \cap \Gamma \subseteq \Gamma_p$ and note that $\mathbb{E}[I_C] = p^4$. For two cycles $C, C' \in \mathcal{C}$, write $C \sim C'$ if $C \cap C' \cap \Gamma \neq \emptyset$. Then, following the notation of Lemma 2.2, we define

$$X := \sum_{C \in \mathcal{C}} I_C, \qquad \mu := \mathbb{E}[X], \qquad \text{and} \qquad \Delta := \sum_{C \sim C'} \mathbb{E}[I_C I_{C'}],$$

where the sum in the definition of Δ ranges over all pairs $(C, C') \in \mathcal{C} \times \mathcal{C}$ such that $C \sim C'$. Now $B(S)$ is precisely the event that $X \leq \zeta n^6 p^7$ and we can use Lemma 2.2 to upper bound the probability of this event occurring.

We begin by estimating $\mu = \mathbb{E}[X]$. We first observe that, for each $u_0 u_4 \in E(S)$, there are at least $(n/4)^3 = n^3/2^6$ choices of $u_1, u_2, u_3 \in W$ such that $u_i u_{i+1} \in \Gamma$ for $i = 0, 1, 2, 3$. Indeed, this follows from the fact that $|W| \geq n/2$ and $\Delta(S) \leq cnp < n/8$ and so u_1, u_2, u_3 can be chosen greedily, avoiding edges of S, with at least $n/4$ choices at each step. Consequently,

$$\mu = \mathbb{E}[X] \geq \frac{sn^3p^4}{2^6} \geq \frac{\theta n^6 p^7}{2^9} \geq 2\zeta n^6 p^7, \tag{1}$$

using that $s \geq \theta n^3 p^3/8$, due to the fact that $S \in \mathcal{S}$.

In order to estimate Δ, we fix some arbitrary $C' \in \mathcal{C}$ and estimate the number of $C \in \mathcal{C}$ (whose vertices we will label $u_0, \ldots, u_4$ as above) that intersect C'. We split the analysis into cases.

1. Firstly assume that $|C \cap C' \cap \Gamma| = 1$. There are at most

$$4 \cdot (4 \cdot \Delta(S) \cdot n^2 + 4 \cdot s \cdot n) \leq 32\Delta(S) \cdot n^2 \tag{2}$$

choices of C that intersect C' in one edge (outside of S), using that $s \leq \Delta(S) \cdot n$ in the inequality. The factor 4 comes from choosing an edge of $C' \cap \Gamma$, say e. The first summand then comes from considering the case where $e = u_0 u_1$ (or analogously $e = u_3 u_4$, resulting in a factor of 2). Given that $e = u_0 u_1$ and choice of labelling of the vertices (another factor of 2), there are at most $\Delta(S)$ choices for u_4 and at most n further choices for each of u_2 and u_3. The second summand stems from the case where $e = u_1 u_2$ (or analogously $e = u_2 u_3$), where after labelling e there are at most s choices for $u_0 u_4$ and at most n further choices for u_3.

2. Next assume $|C \cap C' \cap \Gamma| = 2$. There are at most

$$6 \cdot (2 \cdot \Delta(S) \cdot n + 4 \cdot n + 2s + 8 \cdot \Delta(S)) \leq 96\Delta(S) \cdot n \tag{3}$$

choices of C that intersect C' in two edges (outside of S). Indeed, the factor 6 bounds the number of choices of two edges of $C' \cap \Gamma$, say e_1 and e_2. The first summand then treats the case where $\{e_1, e_2\} = \{u_0 u_1, u_1 u_2\}$ (equivalently, the case where $\{e_1, e_2\} = \{u_2 u_3, u_3 u_4\}$). We then have two options for choosing how to label the endpoints of the path $e_1 e_2$ as u_0 and u_2, then at most $\Delta(S)$ choices for u_4, and n choices for u_3. In the second summand, we consider the case where $\{e_1, e_2\} = \{u_0 u_1, u_3 u_4\}$, which means that there are at most four choices for the edge of $C \cap S$ and at most n further choices for u_2. The third summand treats the case where $\{e_1, e_2\} = \{u_1 u_2, u_2 u_3\}$ and a choice of the edge in S and a labelling of its vertices determines C. Finally, in the fourth summand, we consider the case where $\{e_1, e_2\} = \{u_0 u_1, u_2 u_3\}$ (or $\{e_1, e_2\} = \{u_1 u_2, u_3 u_4\}$) and a choice of the edge $u_0 u_4 \in S$ adjacent to u_0 determines C.

3. Next, consider the case where $|C \cap C' \cap \Gamma| = 3$. There are at most

$$4 \cdot (4 \cdot \Delta(S) + 8) \leq 48\Delta(S) \tag{4}$$

choices of C that intersect C' in three edges (outside of S). Indeed, there are at most 4 choices for the edge $f \in C' \cap \Gamma$ which is *not* on C. If $f = u_0u_1$ (or $f = u_3u_4$), all vertices of C apart from u_0 are fixed and so a choice of neighbour of u_4 in S defines C. If $f = u_1u_2$ (or $f = u_2u_3$), then after labelling, C is already completely determined, leading to the upper bound in the second summand.

4. Finally, if $|C \cap C' \cap \Gamma| = 4$, then clearly there is just one choice for C.

We can now put together the bounds from above to conclude that

$$\Delta \leq \mu \cdot \big(32\Delta(S)n^2p^3 + 96\Delta(S)np^2 + 48\Delta(S)p + 1\big) \leq \mu \cdot 40\Delta(S)n^2p^3.$$

Therefore, we have, using (1) and Lemma 2.2, that

$$\begin{aligned}\Pr[B(S)] = \Pr[X \leq \zeta n^6p^7] \leq \Pr[X \leq \mu/2] &\leq \exp\left(-\frac{\mu^2}{8\Delta}\right) \\ &\leq \exp\left(-\frac{\mu}{2^{10}\Delta(S)n^2p^3}\right) \leq \exp\left(-\frac{snp}{2^{16}\Delta(S)}\right).\end{aligned}$$

Finally, since $\Delta(S) \leq \frac{cnp}{\log(n^2p)-\log(s)}$, which follows from the fact that $S \in \mathcal{S}$, and $c = 2^{-17}$, we have

$$\Pr[B(S)] \leq \exp\big(-2s \cdot \big(\log(n^2p) - \log s\big)\big) = \left(\frac{s}{n^2p}\right)^{2s}.$$

as claimed. ◀

The proof of Proposition 4.7 is now complete. □

4.3 Unbalanced colourings in the upper range

In this section, we improve on Theorem 4.6 when $n^{-3/5} \ll p \ll n^{-1/2}$ and show that in this range we have that $q(n; K_3, p) \leq n^{-6}p^{-8}$. This gives the 1-statement for the upper range in Theorem 1.4.

Theorem 4.11. *Suppose that $n^{-3/5} \ll p \ll n^{-1/2}$ and $q \gg n^{-6}p^{-8}$ and let $G_1 \sim G_{n,p}$ and $G_2 \sim G_{n,q}$ be independent. Then, a.a.s. no G_1-measurable K_3-free colouring $\varphi\colon E(G_1) \to \{\text{red}, \text{blue}\}$ can be extended to a K_3-free colouring of $G_1 \cup G_2$.*

As in the previous section, we first reduce Theorem 4.11 to the following proposition.

Proposition 4.12. *There exist $\beta, \zeta > 0$ such that the following holds. Suppose that $n^{-3/5} \ll p \ll n^{-1/2}$ and let $G \sim G_{n,p}$. Then, a.a.s. every K_3-free $\varphi\colon E(G) \to \{\text{red}, \text{blue}\}$ such that $\left|\varphi^{-1}(\text{red})\right| < \beta n^2p$ results in at least ζn^6p^8 copies of C_{rrbb}.*

With Proposition 4.12 and Proposition 4.3, the proof of Theorem 4.11 follows almost immediately.

Proof of Theorem 4.11. Let $\beta, \zeta > 0$ be the constants from the statement of Proposition 4.12. Further, let $\lambda > 0$ be the constant output by Proposition 4.3 with input β, let $t := \zeta n^6p^8$ and note that $t \leq \lambda n^4p^4$. Now, with $G_1 \sim G_{n,p}$ and $G_2 \sim G_{n,q}$, we have that a.a.s. the conclusions of Propositions 4.1, 4.3 and 4.7 all hold. In particular, any G_1-measurable K_3-free colouring $\varphi\colon E(G_1) \to \{\text{red}, \text{blue}\}$ gives rise to at least t copies of C_{rrbb}. Indeed, this follows from Proposition 4.3 if $\left|\varphi^{-1}(c)\right| \geq \beta n^2p$ for both $c \in \{\text{red}, \text{blue}\}$, or from Proposition 4.12 if $|\varphi^{-1}(c)| < \beta n^2p$ for some $c \in \{\text{red}, \text{blue}\}$. The conclusion of the theorem then follows from Proposition 4.1 as $q \gg t^{-1}$. □

It remains to prove Proposition 4.12. Before embarking on this, we make some definitions and prove several auxiliary lemmas. As in the proof of Proposition 4.7, presented in the previous section, we will condition on the *red* subgraph T of $G_{n,p}$. The following definition captures important properties of T that hold a.a.s. in $G_{n,p}$ and that we will thus be able to assume hold in our proof. Throughout this section, we write θ for the constant from Lemma 2.5.

Definition 4.13. *For $\beta > 0$ and $n^{-3/5} \ll p \ll n^{-1/2}$, let $\mathcal{F} = \mathcal{F}(\beta; p)$ be the set of subgraphs $T \subseteq K_n$ such that*

(i) $\theta n^3 p^3 \leq e(T) < \beta n^2 p$;

*(ii) T satisfies conditions **(P1)** (upper bounding the maximum degree), **(P2)** (upper bounding the number of small subgraphs F), **(P3)** (upper bounding the number of $K_{2,10}$) and **(P5)** (upper bounding the number of edges between vertex sets) of Lemma 2.5.*

In proving Proposition 4.12, we will show that a.a.s. any K_3-free colouring $\varphi\colon E(G) \to \{\text{red}, \text{blue}\}$ such that $\left|\varphi^{-1}(\text{red})\right| < \beta n^2 p$ will have $\varphi^{-1}(\text{red}) \in \mathcal{F}(\beta; p)$. Again, similarly to Proposition 4.7, we will not be able to take a union bound over all possible red subgraphs $T \in \mathcal{F}$ and will instead consider only carefully chosen subgraphs of such T that we can enumerate more efficiently. Given that we aim to find many C_{rrbb}, the following definitions will be useful.

Definition 4.14. *Let $S \subseteq K_n$ be a graph on n vertices. We define the following parameters:*

- *$X_2(S)$ denotes the number of copies of $K_{1,2}$ in S;*
- *$\Pi(S)$ denotes the edges in K_n that complete a triangle with a copy of $K_{1,2}$ that lies in S;*
- *$\mathcal{X}_S$ denotes the family of all copies of $K_{1,2}$ in K_n that form a 4-cycle with some copy of $K_{1,2}$ in S.*

Our next simple lemma gives a lower bound on $X_2(S)$ in terms of the number of edges of a subgraph $S \subseteq K_n$, given that S is not too small.

Lemma 4.15. *If S is a graph on n vertices with at least $2n$ edges, then $X_2(S) \geq \frac{3e(S)^2}{2n}$.*

Proof. By convexity, we have that

$$X_2(S) = \sum_{v \in V} \binom{d_S(v)}{2} \geq n \cdot \binom{\sum_{v \in V} d_S(v)/n}{2} = n \cdot \binom{2e(S)/n}{2} \geq \frac{3e(S)^2}{2n},$$

where the last inequality holds due to our assumption that $e(S) \geq 2n$. □

Next, for certain subgraphs $S \subseteq K_n$, we show that $|\Pi(S)|$ can be lower bounded by $X_2(S)$.

Lemma 4.16. *Suppose that $n^{-3/5} \ll p \ll n^{-1/2}$ and $S \subseteq K_n$ is an n-vertex graph such that $s := e(S) \geq \theta n^3 p^3/2$ and S satisfies **(P3)** (upper bounding the number of $K_{2,10}$) of Lemma 2.5. Then*

$$|\Pi(S)| \geq \frac{X_2(S)}{12} \geq \frac{s^2}{8n}.$$

Proof. For each pair of vertices $\rho \in \binom{[n]}{2} = E(K_n)$, let d_ρ be the number of copies of $K_{1,2}$ in S that form a triangle with ρ and call ρ is *heavy* if $d_\rho \geq 10$. Then $\Pi = \Pi(S) \subseteq E(K_n)$ are the pairs $\rho \in E(K_n)$ such that $d_\rho \geq 1$ and let $\Pi_H \subseteq \Pi$ be the heavy pairs. Then we have that

$$\sum_{\rho \in \Pi_H} \frac{d_\rho}{10} \leq \sum_{\rho \in \Pi_H} \left(\frac{d_\rho}{10}\right)^{10} \leq \sum_{\rho \in \Pi_H} \binom{d_\rho}{10} = N_{K_{2,10}}(S) \leq n^{11} p^{18} \ll n^5 p^6, \tag{5}$$

using that property **(P3)** of Lemma 2.5 holds in S in the penultimate inequality and the fact that $p \ll n^{-1/2}$ in the final inequality. On the other hand,

$$\sum_{\rho \in \Pi} d_\rho = X_2(S) \geq \frac{3s^2}{2n} \geq \frac{\theta^2 n^5 p^6}{4}, \tag{6}$$

by appealing to Lemma 4.15 and our lower bound on $s = e(S)$. Combining (5) and (6) then gives that $\sum_{\rho \in \Pi \setminus \Pi_H} d_\rho \geq 5X_2(S)/6$ and so

$$|\Pi(S)| \geq |\Pi \setminus \Pi_H| \geq \frac{1}{10} \sum_{\rho \in \Pi \setminus \Pi_H} d_\rho \geq \frac{X_2(S)}{12} \geq \frac{s^2}{8n},$$

using Lemma 4.15, which completes the proof. □

Our next lemma identifies, for each $T \in \mathcal{F}$, some subgraph $S = S(T) \subseteq T$ for which the collection $\mathcal{X}_S$ is large and well-spread in K_n. Following the notation of Lemma 2.2, for a subgraph $S \subseteq K_n$ and $p = p(n)$, we let

$$\mu(\mathcal{X}_S) := |\mathcal{X}_S| p^2 \in |\Pi(S)| p^2 \cdot \{n-2, n-3\} \tag{7}$$

be the expected number of copies of $K_{1,2}$ in $\mathcal{X}_S$ that appear in $G_{n,p}$ and let

$$\Delta(\mathcal{X}_S) := \sum_{K,K'} p^{e(K \cup K')},$$

where the sum goes over all pairs of copies $K, K' \in \mathcal{X}_S$ such that $K \cap K' \neq \emptyset$.

Lemma 4.17. *Suppose $0 < \beta < 2^{-100}$ and $n^{-3/5} \ll p \ll n^{-1/2}$ and let $T \in \mathcal{F} = \mathcal{F}(\beta; p)$ with $t := e(T)$. Then there exists a subgraph $S = S(T) \subseteq T$ with $e(S) \geq t/2$ and such that either*

(a) $\frac{\mu(\mathcal{X}_S)^2}{\Delta(\mathcal{X}_S)} \geq 10t \log\left(\frac{2n^2 p}{t}\right)$*; or,*

(b) $10t \log\left(\frac{2n^2 p}{t}\right) > \frac{\mu(\mathcal{X}_S)^2}{\Delta(\mathcal{X}_S)} \geq \frac{\mu(\mathcal{X}_S)}{3}$.

Proof. Fix some $T \in \mathcal{F}$ and denote $t := e(T)$. Now for any subgraph $S \subseteq T$, we define

$$\Delta_1(\mathcal{X}_S) := |\{K, K' \in \mathcal{X}_S : K \cup K' \text{ is a path with 3 edges }\}| \cdot p^3 \text{ and}$$
$$\Delta_2(\mathcal{X}_S) := |\{K, K' \in \mathcal{X}_S : K \cup K' \text{ is a copy of } K_{1,3}\}| \cdot p^3,$$

and we note that, for every S, we have $\Delta(\mathcal{X}_S) = \Delta_1(\mathcal{X}_S) + \Delta_2(\mathcal{X}_S) + \mu(\mathcal{X}_S)$. Indeed, the sum in the definition of $\Delta(\mathcal{X}_S)$ ranges over all pairs $K, K' \in \mathcal{X}_S$ that intersect in at least one edge. If they intersect in two edges, then $K = K'$ and the contribution to $\Delta(\mathcal{X}_S)$ is counted by $\mu(\mathcal{X}_S)$ and if they intersect in precisely one edge, then their union is either a path, in which case they are counted by $\Delta_1(\mathcal{X}_S)$, or a star in which case they are counted by $\Delta_2(\mathcal{X}_S)$. The following claim is the key step in proving the lemma.

Claim 4.18. *There is an $S \subseteq T$ with at least $t/2$ edges such that*

$$\frac{\mu(\mathcal{X}_S)^2}{\Delta_2(\mathcal{X}_S)} \geq 30t \log\left(\frac{2n^2 p}{t}\right).$$

With Claim 4.18, the lemma follows quickly. Indeed, fix $S \subseteq T$ as output by the claim and note that

$$\frac{\mu(\mathcal{X}_S)^2}{\Delta(\mathcal{X}_S)} = \frac{\mu(\mathcal{X}_S)^2}{\Delta_1(\mathcal{X}_S) + \Delta_2(\mathcal{X}_S) + \mu(\mathcal{X}_S)} \geq \frac{1}{3} \min\left\{\frac{\mu(\mathcal{X}_S)^2}{\Delta_1(\mathcal{X}_S)}, \frac{\mu(\mathcal{X}_S)^2}{\Delta_2(\mathcal{X}_S)}, \mu(\mathcal{X}_S)\right\}. \tag{8}$$

Firstly, suppose that the minimum is achieved by the last term. In this case we have that $\frac{\mu(\mathcal{X}_S)^2}{\Delta(\mathcal{X}_S)} \geq \frac{\mu(\mathcal{X}_S)}{3}$ and the conclusion of the lemma is satisfied, with S satisfying (a) if $\frac{\mu(\mathcal{X}_S)}{3} \geq 10t \log\left(\frac{2n^2 p}{t}\right)$ and (b) otherwise.

Likewise, if the minimum is achieved by the middle term, then by Claim 4.18 we have that $\frac{\mu(\mathcal{X}_S)^2}{\Delta(\mathcal{X}_S)}$ satisfies (a). It remains to consider the case where the minimum is achieved by the first term. For this, note that we have $\Delta_1(\mathcal{X}_S) \leq 4|\Pi(S)|^2 p^3$. Indeed, it $v_1v_2v_3v_4$ is a path of length three in S labeled so that K is the path $v_1v_2v_3$ and K' is the path $v_2v_3v_4$, see Figure 7, then $v_1v_3, v_2v_4 \in \Pi(S)$ and thus we can count the number of pairs K, K' whose union is a path with three edges by the number of pairs in $\Pi(S)$ with a labelling of the endpoints of each pair. Using (7), we therefore have that

$$\frac{\mu(\mathcal{X}_S)^2}{\Delta_1(\mathcal{X}_S)} \geq \frac{|\Pi(S)|^2 n^2 p^4}{5|\Pi(S)|^2 p^3} \geq \frac{n^2 p}{5} = t \log\left(\frac{2n^2 p}{t}\right) \cdot \frac{n^2 p}{5t} \cdot \log^{-1}\left(\frac{2n^2 p}{t}\right) \geq 30 t \log\left(\frac{2n^2 p}{t}\right),$$

using here that $t < \beta n^2 p$ (property (i) of Definition 4.13) and the fact that $\beta < 2^{-100}$. Therefore, we also satisfy part (a) of the lemma when the minimum in (8) is achieved by the first term.

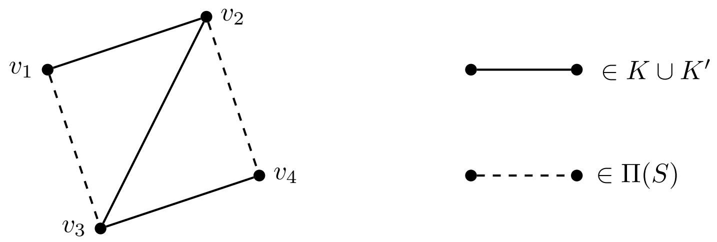

Figure 7: Two copies of $K_{1,2}$ in $\mathcal{X}_S$ whose union is a path.

It remains to prove Claim 4.18, which we do now.

Proof of Claim 4.18. In order to derive the lower bound in the Claim 4.18, we need to upper bound $\Delta_2(\mathcal{X}_S)$ and hence we need an upper bound on the count of pairs $K, K' \in \mathcal{X}_S$ such that $K \cup K'$ forms a copy of $K_{1,3}$. Given such a pair K and K', denote the vertex of degree three in $K \cup K'$ by u and the remaining vertices by w_1, w_2, w_3 so that K lies on vertices u, w_1, w_2, K' lies on vertices u, w_2, w_3 and thus $w_1w_2, w_2w_3 \in \Pi(S)$, see Figure 8. Hence, we have that

$$\Delta_2(\mathcal{X}_S) \leq X_2(\Pi(S)) n p^3, \tag{9}$$

where $X_2(\Pi(S))$ is the number of copies of $K_{1,2}$ in $\Pi(S)$ when considered as a graph on n vertices. Indeed, the number of possible pairs K, K' with $K \cup K'$ a copy of $K_{1,3}$ can be bounded by choosing a copy of $K_{1,2}$ in $\Pi(S)$ and a choice of vertex u (at most n choices). We proceed by splitting our analysis into cases, depending on whether or not T contains a large subgraph with maximum degree at most $d := \sqrt{\frac{tp}{2000 \log(2n^2p/t)}}$.

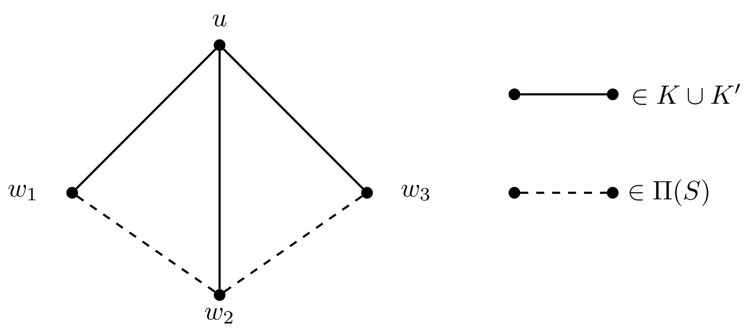

Figure 8: Two copies of $K_{1,2}$ in $\mathcal{X}_S$ whose union is a copy of $K_{1,3}$.

Case 1. T contains a subgraph with at least $t/2$ edges and maximum degree at most d.

We let S be one such subgraph. Note that $X_2(\Pi(S)) \leq |\Pi(S)|\Delta(\Pi(S))$ where $\Delta(\Pi(S))$ is the maximum degree of a vertex in $\Pi(S)$ when considered as a graph on n vertices. Since $\Delta(\Pi(S)) \leq \Delta(S)^2 \leq d^2$, appealing to (7) and (9), we have that

$$\frac{\mu(\mathcal{X}_S)^2}{\Delta_2(\mathcal{X}_S)} \geq \frac{|\Pi(S)|^2 n^2 p^4}{2|\Pi(S)|d^2 np^3} \geq \frac{1000|\Pi(S)|n\log(2n^2p/t)}{t} \geq 30t\log\left(\frac{2n^2p}{t}\right),$$

as claimed, where we used Lemma 4.16, noting that the conditions are satisfied as $S \subseteq T \in \mathcal{F}$ and $e(S) \geq t/2$.

Case 2. Every subgraph of T with maximum degree at most d has fewer than $t/2$ edges. In this case, we define $S = T$ as the subgraph with the desired properties. First, we claim that

$$|\Pi(T)| \geq \frac{td}{24}. \tag{10}$$

Indeed, let $H \subseteq T$ be a maximal subgraph with respect to inclusion such that $\Delta(H) \leq d$. By our assumption, $e(H) < t/2$. By the definition of H, for any $e \in E(T) \setminus E(H)$, one of its endpoints has degree at least $d+1$ when added to T and hence e is contained in at least d copies of $K_{1,2}$ with edges in H. Summing over all edges in $E(T) \setminus E(H)$ gives that $X_2(T) \geq td/2$ and (10) follows from Lemma 4.16. We will also show that

$$X_2(\Pi(T)) \leq 20|\Pi(T)|tp. \tag{11}$$

We claim that this suffices to prove Claim 4.18. Indeed, appealing to (7), (9) and (10), we get that

$$\frac{\mu(\mathcal{X}_T)^2}{\Delta_2(\mathcal{X}_T)} \geq \frac{|\Pi(T)|^2 n^2 p^4}{40|\Pi(T)|tp \cdot np^3} \geq \frac{|\Pi(T)|n}{40t} \geq \frac{dn}{1000} = \frac{1}{1000} \cdot \sqrt{\frac{tp}{2000\log(2n^2p/t)}} \cdot n$$
$$\geq 30t\log\left(\frac{2n^2p}{t}\right) \cdot \sqrt{\frac{n^2p}{2^{50}t} \cdot \log^{-3}\left(\frac{2n^2p}{t}\right)} \geq 30t\log\left(\frac{2n^2p}{t}\right),$$

as desired, using that $t < \beta n^2 p$ and the fact that $\beta \leq 2^{-100}$ in the last inequality here.

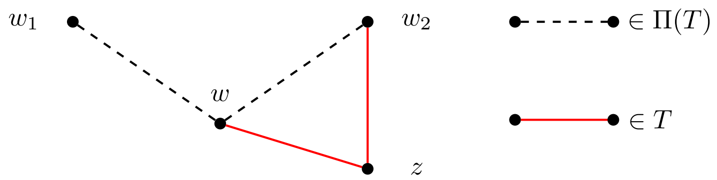

Figure 9: A copy of $K_{1,2}$ in $\Pi(T)$.

To show that (11) holds, note first that $X_2(\Pi(T))$ is at most the number of paths w_1wzw_2 in K_n such that $w_1w \in \Pi(T)$ and $wz, zw_2 \in T$, see Figure 9. Consequently, denoting by $d_\Pi(v)$ the number of neighbours of v in $\Pi(T)$, we have

$$X_2(\Pi(T)) \leq \sum_{e=wz\in T} d_\Pi(w)d_T(z) = \sum_{w\in[n]}\sum_{z\in[n]} d_\Pi(w)d_T(z)\mathbb{1}[wz \in T]. \tag{12}$$

We will split this sum further by grouping together vertices depending on their degrees. To this end, for $0 \leq \alpha \leq \log_2(np)$ and $0 \leq \beta \leq 2\log_2(np)+1$ let

$$A_\alpha := \left\{z \in V(T) : \frac{d_T(z)}{2np} \in \left[2^{-\alpha-1}, 2^{-\alpha}\right]\right\}, \qquad B_\beta := \left\{w \in V(T) : \frac{d_\Pi(w)}{4n^2p^2} \in \left[2^{-\beta-1}, 2^{-\beta}\right]\right\},$$

and note that the sets A_α partition the vertices and so do the sets B_β, using here that $d_\Pi(w) \leq \Delta(\Pi(S)) \leq \Delta(T)^2 \leq 4n^2p^2$ for all $w \in V(T)$ due to the fact that $T \in \mathcal{F}(\beta, p)$ and so $\Delta(T) \leq 2np$ (see Definition 4.13). Hence, returning to the upper bound (12), we have that

$$X_2(\Pi(T)) \leq \sum_{\alpha=0}^{\log_2(np)} \sum_{\beta=0}^{2\log_2(np)+1} e_T(A_\alpha, B_\beta) \cdot 8n^3p^3 2^{-\alpha-\beta}. \tag{13}$$

We now turn to bounding $e_T(A_\alpha, B_\beta)$ for each $0 \le \alpha \le \log_2(np)$ and $0 \le \beta \le 2\log_2(np) + 1$. For each such α and β, set

$$a_\alpha := \frac{2^{\alpha+1}t}{np}, \qquad b_\beta := \frac{2^\beta|\Pi(T)|}{n^2p^2},$$

and note that $|A_\alpha| \le a_\alpha$ and $|B_\beta| \le b_\beta$. Indeed, $2^{-\alpha}np|A_\alpha| \le \sum_{z\in V} d_T(z) \le 2t$, and similarly for $|B_\beta|$. Moreover

$$a_\alpha = \frac{2^{\alpha+1}t}{np} \ge \frac{t}{np} \ge \theta n^2p^2 \gg \frac{(\log n)^7}{p},$$

using that $t \ge \theta n^3p^3$ as $T \in \mathcal{F}$ (see Definition 4.13) and the fact that $p \ge n^{-3/5}$. Similarly, we have that

$$b_\beta = \frac{2^\beta|\Pi(T)|}{n^2p^2} \ge \frac{td}{24n^2p^2} \ge \frac{t^{3/2}p^{1/2}}{1200n^2p^2\log(2n^2p/t)} \ge \frac{\theta^{3/2}n^{5/2}p^3}{1200\log(2n^2p/t)} \gg \frac{(\log n)^7}{p},$$

where we used (10) to lower bound $|\Pi(T)|$ as well as our lower bounds on t and p. Now as $T \in \mathcal{F}(\beta; p)$ has property **(P5)** of Lemma 2.5, see Definition 4.13,

$$e_T(A_\alpha, B_\beta) \le |A_\alpha|\cdot|B_\beta|\cdot p + \frac{a_\alpha \cdot b_\beta \cdot p}{\log^3 n} \le |A_\alpha|\cdot|B_\beta|\cdot p + \frac{2^{\alpha+\beta+1}t|\Pi(T)|}{n^3p^2\log^3 n},$$

for all α, β in our ranges of interest. Therefore, plugging these upper bounds into (13), we get

$$\begin{aligned}
X_2(\Pi(T)) &\le \sum_{\alpha=0}^{\log_2(np)} \sum_{\beta=0}^{2\log_2(np)+1} \left(|A_\alpha|\cdot|B_\beta|\cdot p + \frac{2^{\alpha+\beta+1}t|\Pi(T)|}{n^3p^2\log^3 n}\right)\cdot 8n^3p^32^{-\alpha-\beta} \\
&\le \frac{16|\Pi(T)|\cdot tp}{\log n} + 8p\cdot\left(\sum_{\alpha=0}^{\log_2(np)} |A_\alpha|\cdot 2^{-\alpha}np\right)\cdot\left(\sum_{\beta=0}^{2\log_2(np)+1} |B_\beta|\cdot 2^{-\beta}n^2p^2\right) \\
&\le |\Pi(T)|\cdot tp + 8p\cdot\left(\sum_{z\in[n]} d_T(z)\right)\cdot\left(\sum_{w\in[n]} d_\Pi(w)/2\right) \\
&\le |\Pi(T)|\cdot tp + 8p\cdot 2t\cdot|\Pi(T)| \le 20|\Pi(T)|tp,
\end{aligned}$$

establishing (11) and completing the proof of this case, the claim and the lemma. ◀

□

We will split our analysis in the proof of Proposition 4.12 depending on whether $T \in \mathcal{F}(\beta; p)$ outputs an $S(T)$ that satisfies (a) or (b) when Lemma 4.17 is applied to T. If (a) holds, then we say that T is of *type* (a) and, likewise, if (b) holds for $S(T)$, we say that T is of *type* (b). Our next lemma states that type (b) subgraphs can only occur when both p and $e(T)$ are very close ot their minimal values.

Lemma 4.19. *Suppose* $0 < \beta \le 2^{-100}$, $n^{-3/5} \ll p \ll n^{-1/2}$ *and* $T \in \mathcal{F} = \mathcal{F}(\beta; p)$ *is of type (b) with* $t := e(T)$. *Then*

(i) $p \le Cn^{-3/5}\log^{1/5} n$ *for some* $C = C(\theta)$;

(ii) $t \ll n^3p^3\log n$;

(iii) $\log\left(\frac{2n^2p}{t}\right) \ge \frac{\log n}{6}$.

Proof. Suppose that $T \in \mathcal{F}(\beta; p)$ is of type (b), let $S := S(T) \subseteq T$ be the graph output by Lemma 4.17 when applied to T and let $s := e(S) \ge t/2$. If (i) did not hold, then (7), Lemma 4.16 and the fact that $t \ge \theta n^3p^3$ (see Definition 4.13) would imply that, if $C = C(\theta)$ is sufficiently large,

$$\frac{\mu(X_S)}{3} \ge \frac{|\Pi(S)|np^2}{4} \ge \frac{s^2p^2}{32} \ge \frac{t^2p^2}{128} \ge \frac{\theta tn^3p^5}{128} \ge \frac{C^5\theta t\log n}{128} \ge 10t\log\left(\frac{2n^2p}{t}\right),$$

contradicting the fact that S satisfies part (b) of Lemma 4.17. Similarly, if (ii) did not hold, then, for some positive constant c, we would have, using that $p \gg n^{-3/5}$,

$$\frac{\mu(\mathcal{X}_S)}{3} \geq \frac{t^2p^2}{128} \geq \frac{ctn^3p^5\log n}{240} \gg t\log n,$$

again contradicting the fact that S satisfies part (b) of Lemma 4.17. This means that both (i) and (ii) must hold. Finally, the third assertion (iii) follows from the other two. Indeed, we have that $p \ll n^{-7/12}(\log n)^{-1/2}$ from part (i) and so using (ii), we get that

$$\frac{2n^2p}{t} \geq \frac{2}{np^2\log n} \gg n^{1/6}. \qquad \square$$

Finally, we prove Proposition 4.12, completing this section.

Proof of Proposition 4.12. Let $\alpha = 1/6400$, $\beta \coloneqq 2^{-100}$ and let $\zeta \coloneqq \min\{\alpha\theta^2/96, \alpha\theta/(6C^5)\}$, where $C = C(\theta)$ is the constant from the statement of Lemma 4.19. Let $G \sim G_{n,p}$, let $\mathcal{F} \coloneqq \mathcal{F}(\beta; p)$ be the collection of graphs from Definition 4.13 and let $\mathcal{H}$ be the collection of n-vertex graphs that satisfy properties **(P1)–(P5)** of Lemma 2.5. Further, let $\mathcal{F}_{(a)}, \mathcal{F}_{(b)} \subseteq \mathcal{F}$ be the graphs of types (a) and (b), respectively. Now, for a graph $R \subseteq K_n$, let $A(R)$ be the event that $R = \varphi^{-1}(\text{red})$ for some K_3-free colouring $\varphi\colon E(G) \to \{\text{red}, \text{blue}\}$ of G with fewer than ζn^6p^8 copies of C_{rrbb}. As $\Pr[G \notin \mathcal{H}] \ll 1$, due to Lemma 2.5, it suffices to show that a.a.s. the event $\bigcup\{A(R) : e(R) < \beta n^2p\} \cap \{G \in \mathcal{H}\}$ does not occur. To this end, we first claim that the event $A(R) \cap \{G \in \mathcal{H}\}$ is empty unless $R \in \mathcal{F}$. To see this, suppose that $G \in \mathcal{H}$ and the event $A(R)$ happens for some R with $e(R) < \beta n^2p$. As properties **(P1)–(P3)** and **(P5)** of Lemma 2.5 are all monotone decreasing, the graph $R \subseteq G \in \mathcal{H}$ must satisfy condition (ii) of Definition 4.13. Moreover, the upper bound on $e(R)$ in condition (i) of Definition 4.13 is satisfied by assumption. As for the lower bound, property **(P4)** of Lemma 2.5 supplies a collection of at least θn^3p^3 edge-disjoint triangles in G and each colours class of every K_3-free colouring of G must contain at least one edge from each triangle in this collection. Therefore, it remains to show that a.a.s. no event $A(T) \cap \{G \in \mathcal{H}\}$ occurs with $T \in \mathcal{F} = \mathcal{F}_{(a)} \cup \mathcal{F}_{(b)}$. We first deal with the type (a) graphs T.

Claim 4.20. *For every $T \in \mathcal{F}_{(a)}$ with $e(T) = t$, we have* $\Pr[A(T)] \leq p^t\left(\frac{t}{2n^2p}\right)^t$.

Proof. Applying Lemma 4.17, we get some subgraph $S \coloneqq S(T) \subseteq T$ with $s \coloneqq e(S) \geq t/2$ and $\frac{\mu(\mathcal{X}_S)^2}{\Delta(\mathcal{X}_S)} \geq 10t\log\left(\frac{2n^2p}{t}\right)$. Now, let $\mathcal{Y}_T \subseteq \mathcal{X}_S$ be the family of all copies of $K_{1,2}$ in $K_n \setminus T$ that form a 4-cycle with some copy of $K_{1,2}$ in S and avoid the edges of T. Using the notation of Lemma 2.2, let $\mu \coloneqq |\mathcal{Y}_T|p^2$ be the expected number of copies of $K_{1,2}$ in $\mathcal{Y}_T$ that appear in $G_{n,p}$ and let $\Delta \coloneqq \sum_{K,K'} p^{e(K\cup K')}$, where the sum goes over all pairs of copies $K, K' \in \mathcal{Y}_T$ such that $K \cap K' \neq \emptyset$. Note that $\Delta \leq \Delta(\mathcal{X}_S)$, as $\mathcal{Y}_T \subseteq \mathcal{X}_S$, and we also have that

$$\mu = |\mathcal{Y}_T|p^2 \geq |\Pi(S)|(n - 2\Delta(T))p^2 \geq \frac{2}{\sqrt{5}}|\Pi(S)|np^2 \geq \frac{2\mu(\mathcal{X}_S)}{\sqrt{5}},$$

appealing to (7) and the fact that $\Delta(T) \leq 2np \ll n$, as $T \in \mathcal{F}$, here.

Now, let $B(T)$ be the event that fewer than $\mu/2$ copies of $K_{1,2}$ from $\mathcal{Y}_T$ appear in G. By Lemma 2.2,

$$\Pr[B(T)] \leq \exp\left(-\frac{\mu^2}{8\Delta}\right) \leq \exp\left(-\frac{\mu(\mathcal{X}_S)^2}{10\Delta(\mathcal{X}_S)}\right) \leq \exp\left(-t\log\left(\frac{2n^2p}{t}\right)\right) = \left(\frac{t}{2n^2p}\right)^t.$$

We claim that $A(T) \subseteq B(T)$. Indeed, suppose that $A(T)$ occurs and fix some colouring $\varphi\colon E(G) \to \{\text{red}, \text{blue}\}$ with $\varphi^{-1}(\text{red}) = T$ and fewer than ζn^6p^8 copies of C_{rrbb}. Since every copy K of $K_{1,2}$ in $\mathcal{Y}_T$ that appears in G is coloured blue (as its edges are not in T), it gives rise to at least one copy of C_{rrbb}, as K forms a copy of C_4 with two edges of S (which φ colours red). This implies that $B(T)$ occurs as otherwise, by (7), Lemma 4.16 and the fact that $s \geq t/2 \geq \theta n^3p^3/2$, see Definition 4.13, we would get

$$\frac{\mu}{2} \geq \frac{\mu(\mathcal{X}_S)}{4} \geq \frac{|\Pi(S)|np^2}{8} \geq \frac{s^2p^2}{64} \geq \zeta n^6p^8$$

copies of C_{rrbb}, a contradiction. Finally, as the event $A(T)$ occurring implies that $T \subseteq G$ and the events $T \subseteq G$ and $B(T)$ are independent (the copies of $K_{1,2}$ in $\mathcal{Y}_T$ avoid the edges of T), we have

$$\Pr[A(T)] \le \Pr[\{T \subseteq G\} \cap B(T)] \le \Pr[T \subseteq G] \cdot \Pr[B(T)] \le p^t \left(\frac{t}{2n^2p}\right)^t,$$

as required. ◀

Using Claim 4.20 and appealing to a union bound, we thus have that

$$\Pr\left[A(T) \cap \{G \in \mathcal{H}\} \text{ for some } T \in \mathcal{F}_{(a)}\right] \le \sum_{T \in \mathcal{F}_{(a)}} \Pr\left[A(T)\right] \le \sum_{T \in \mathcal{F}_{(a)}} p^{e(T)} \cdot \left(\frac{e(T)}{2n^2p}\right)^{e(T)}$$
$$\le \sum_{t=\theta n^3p^3}^{\beta n^2 p} \binom{\binom{n}{2}}{t} \cdot p^t \cdot \left(\frac{t}{2n^2p}\right)^t \le \beta n^2 p \left(\frac{e}{4}\right)^n \ll 1.$$

It remains to consider the events $A(T) \cap \{G \in \mathcal{H}\}$ for type (b) graphs $T \in \mathcal{F}_{(b)}$. More precisely, we need to show that a.a.s. G does not belong to the family $\mathcal{H}'$ defined by

$$\mathcal{H}' := \left\{H \in \mathcal{H} : \exists \varphi\colon E(H) \to \{\text{red}, \text{blue}\} \text{ with } \varphi^{-1}(\text{red}) \in \mathcal{F}_{(b)} \text{ and fewer than } \zeta n^6 p^8 \text{ copies of } C_{rrbb}\right\}.$$

In order to do this, for each $H \in \mathcal{H}'$, we will identify some $\ell = \ell(H) \in \mathbb{N}$ and a pair of increasing sequences $\mathbf{J}(H) = (J_0, J_1, \ldots, J_\ell)$ and $\mathbf{R}(H) = (R_0, R_1, \ldots, R_\ell)$ of subgraphs of H; we will refer to this pair as the *stamp* of H and denote it by $\mathbf{S}(H) = (\mathbf{J}(H), \mathbf{R}(H))$. Our proof will then provide an upper bound on the probability that $\mathbf{S}(G) = \mathbf{S}$, for any given stamp $\mathbf{S}$, that is strong enough to survive a union bound over all possible stamps $\mathbf{S}$. Before proceeding, we remark that, by Lemma 4.19, we can assume that $p \le Cn^{-3/5} \log^{1/5} n$, as otherwise $\mathcal{F}_{(b)} = \emptyset$ and so $\mathcal{H}' = \emptyset$.

Now, for each $H \in \mathcal{H}'$, we define $\ell(H)$ and construct the sequences $\mathbf{J}(H)$ and $\mathbf{R} = \mathbf{R}(H)$ (and hence the stamp $\mathbf{S}(H)$) by considering the following process:

1. Fix some K_3-free colouring $\varphi\colon E(H) \to \{\text{red}, \text{blue}\}$ with fewer than $\zeta n^6 p^8$ copies of C_{rrbb} and $T := \varphi^{-1}(\text{red}) \in \mathcal{F}_{(b)}$.
2. Choose a collection $\mathcal{C}$ of $c_0 := \theta n^3 p^3$ edge-disjoint triangles in H (this is possible as $H \in \mathcal{H}$ and so it satisfies property **(P4)** of Lemma 2.5), fix J_0 to be the collection of edges featuring in $\mathcal{C}$ and $R_0 := J_0 \cap T$ to be the collection of edges in J_0 which are coloured red by φ.

At this point note that, for any n-vertex graph R with $R_0 \subseteq R \subseteq T$, we have that $R \in \mathcal{F}$. Indeed, $|R| \ge |R_0| \ge c_0 = \theta n^3 p^3$, as there is at least one red edge in each triangle in $\mathcal{C}$, and the other conditions of Definition 4.13 follow from the fact that $R \subseteq T \in \mathcal{F}_{(b)} \subseteq \mathcal{F}$. We now continue to form our sequences. We will maintain that $J_{i-1} \subseteq J_i$ for all $i \ge 1$ and define $R_i := J_i \cap T$ to be the collection of edges in J_i that are coloured red by φ (as is the case with $R_0 \subseteq J_0$).

3. Suppose that $i \ge 0$ and that J_i and R_i have already been defined. As $R_0 \subseteq R_i \subseteq T$, we have that $R_i \in \mathcal{F}$. In particular, Lemma 4.17 gives us $S_i := S(R_i)$ such that $e(S_i) \ge e(R_i)/2$. Now fix
$$M_i := \alpha \min\left\{|\Pi(S_i)| np^2, e(R_i) \log\left(\frac{2n^2p}{e(R_i)}\right)\right\}, \tag{14}$$
recalling the definition of $\Pi(S_i)$ from Definition 4.14. Let $\mathcal{X}(H, J_i, S_i) \subseteq \mathcal{X}_{S_i}$ be the set of copies of $K_{1,2}$ in $H \setminus J_i$ that avoid the edges of J_i and form a 4-cycle with some copy of $K_{1,2}$ in S_i. Moreover, let $\mathcal{Z}$ be a largest collection of edge-disjoint copies of $K_{1,2}$ in $\mathcal{X}(H, J_i, S_i)$. If $|\mathcal{Z}| < M_i$ then terminate the process and fix $\ell(H) := i$. Otherwise, if $|\mathcal{Z}| \ge M_i$, then choose some $\mathcal{Z}' \subseteq \mathcal{Z}$ with $|\mathcal{Z}'| = M_i$, let J_{i+1} be the graph obtained by adding the edges in copies of $K_{1,2}$ in $\mathcal{Z}'$ to J_i and let $R_{i+1} = J_{i+1} \cap T$. Repeat step 3 with $i+1$ replacing i.

We begin by collecting some observations about the process with the following claims.

Claim 4.21. *For every $i \in \{0, \dots, \ell - 1\}$, we have*

$$e(R_{i+1}) - e(R_i) \geq \frac{e(J_{i+1}) - e(J_i)}{3} = \frac{2M_i}{3} \geq 2\zeta n^6 p^8.$$

Consequently, $e(R_i) \geq e(J_i)/3 \geq 2\zeta i n^6 p^8/3$ for every $i \in \{0, \dots, \ell\}$.

Proof. Since we have already shown that $e(R_0) \geq e(J_0)/3 = c_0 = \theta n^3 p^3$, as each triangle in $\mathcal{C}$ must contain at least one red edge, the second assertion of the lemma easily follows from the first assertion. We begin by showing that $M_i \geq 3\zeta n^6 p^8$, which follows from the stronger inequality

$$\frac{M_i}{e(R_i)} \geq \frac{3\zeta n^3 p^5}{\theta}, \tag{15}$$

as $e(R_i) \geq e(R_0) \geq \theta n^3 p^3$. To see that (15) holds, consider two cases. If M_i is equal to the first term in (14), this follows from Lemma 4.16 as

$$\frac{|\Pi(S_i)| n p^2}{e(R_i)} \geq \frac{e(S_i)^2 p^2}{8e(R_i)} \geq \frac{e(R_i)p^2}{32} \geq \frac{e(R_0)p^2}{32} \geq \frac{\theta n^3 p^5}{32} \geq \frac{3\zeta n^3 p^5}{\alpha\theta}.$$

Similarly, if M_i is equal to the second term, then, recalling that we have assumed that $n^3p^5 \leq C^5 \log n$,

$$\log\left(\frac{2n^2p}{e(R_i)}\right) \geq \log\left(\frac{2n^2p}{e(T)}\right) \geq \frac{\log n}{6} \geq \frac{n^3p^5}{6C^5} \geq \frac{3\zeta n^3 p^5}{\alpha\theta},$$

where we also used Lemma 4.19 (iii) and the fact that $R_i \subseteq T \in \mathcal{F}_{(b)}$.

Observe now that at least two thirds among the collection $\mathcal{Z}'$ of M_i copies of $K_{1,2}$, which are added to J_i to get J_{i+1}, must contain an edge of T. Indeed, if this was not the case, then more than $M_i/3 \geq \zeta n^6 p^8$ such copies would be coloured completely blue. However, each of those forms a copy of C_{rrbb} with two edges of $S_i \subseteq R_i \subseteq T$, which are all coloured red. This contradicts the assumption that there are fewer than $\zeta n^6 p^8$ copies of C_{rrbb} in H. Therefore, it must be that $e(R_{i+1}) - e(R_i) \geq 2M_i/3 = (e(J_{i+1}) - e(J_i))/3$. ◀

Since $n^6p^8 \gg n^3p^3$, by our assumption that $p \gg n^{-3/5}$, Claim 4.21 implies that $e(R_\ell) \geq \ell n^3 p^3$. Consequently, since $R_\ell \subseteq T \in \mathcal{F}_{(b)}$, we must have that $\ell = \ell(H) \leq \log n$.

Claim 4.22. *For every $i \in \{0, \dots, \ell - 1\}$, we have $M_i = \alpha|\Pi(S_i)|np^2$.*

Proof. If this was not the case, then, for some $i \in \{0, \dots, \ell - 1\}$, we would have that

$$M_i = \alpha e(R_i) \log\left(\frac{2n^2p}{e(R_i)}\right) \geq \frac{\alpha e(J_i) \log n}{18} \geq \frac{\alpha e(J_0) \log n}{18} \geq \frac{\alpha\theta n^3 p^3 \log n}{18},$$

using Lemma 4.19 (iii) and Claim 4.21 here. But then, appealing again to Claim 4.21, we have that

$$e(R_\ell) \geq \frac{e(J_\ell)}{3} \geq \frac{e(J_{i+1})}{3} \geq \frac{2M_i}{3} \geq \frac{\alpha\theta n^3 p^3 \log n}{27},$$

which is a contradiction, as $R_\ell \subseteq T \in \mathcal{F}_{(b)}$ and hence $e(R_\ell) \ll n^3 p^3 \log n$ by Lemma 4.19. ◀

We are finally in a position to bound the probability that $\mathbf{S}(G) = \mathbf{S}$ for each possible stamp $\mathbf{S} = (\mathbf{J}, \mathbf{R})$. Recall the definition of $\mathcal{H}'$ and let

$$\boldsymbol{\mathcal{S}} = \{\mathbf{S}(H) = (\mathbf{J}(H), \mathbf{R}(H)) : H \in \mathcal{H}'\}$$

be the set of stamps obtained by running the above process on all possible graphs $H \in \mathcal{H}'$. Further, for $0 \leq k \leq \log n$, let $\boldsymbol{\mathcal{S}}^k := \{\mathbf{S}(H) : H \in \mathcal{H}', \ell(H) = k\} \subseteq \boldsymbol{\mathcal{S}}$ be the stamps of length k.

Claim 4.23. *For any* $0 \le k \le \log n$ *and all* $\mathbf{S} = ((J_0, \dots, J_k), (R_0, \dots, R_k)) \in \boldsymbol{\mathcal{S}}^k$*, we have that*

$$\Pr[\mathbf{S}(G) = \mathbf{S}] \le p^{e(J_k)} \exp(-\zeta n^3 p^5 e(J_k)).$$

Since $\mathbf{S}(G)$ is only defined for $G \in \mathcal{H}'$, the event that $\mathbf{S}(G) = \mathbf{S}$ implicitly implies that $G \in \mathcal{H}'$.

Proof of Claim 4.23. Fix some $0 \le k \le \log n$ and $\mathbf{S} = ((J_0, \dots, J_k), (R_0, \dots, R_k)) \in \boldsymbol{\mathcal{S}}^k$ as in the statement of the claim. Further, let $M := M_k$, as in (14), where $S := S_k = S(R_k)$ is the graph obtained from Lemma 4.17 with input R_k. Now, let let Z be the largest size of a collection of edge-disjoint copies of $K_{1,2}$ in $G \setminus J_k$ that form a 4-cycle with some copy of $K_{1,2}$ in S. We claim that $\mathbf{S}(G) = \mathbf{S}$ implies both $J_k \subseteq G$ and $Z < M$. Indeed, certainly any $H \in \mathcal{H}'$ with $\mathbf{S}(H) = \mathbf{S}$ must satisfy $J_k \subseteq H$. Further, if $Z \ge M$ and $G \in \mathcal{H}'$, then the process defining $\mathbf{S}(G)$ would not terminate at step k, and thus $\ell(G) > k$, precluding $\mathbf{S}(G) = \mathbf{S}$. Since the events $J_k \subseteq G$ and $Z < M$ are independent, as Z depends only on $G \setminus J_k$, we have

$$\Pr[\mathbf{S}(G) = \mathbf{S}] \le \Pr[\{J_k \subseteq G\} \cap \{Z < M\}] \le \Pr[J_k \subseteq G] \Pr[Z < M] = p^{e(J_k)} \Pr[Z < M].$$

It thus remains to bound $\Pr[Z < M]$.

To this end, we let $\mathcal{Y} \subseteq \mathcal{X}_S$ be the family of copies of $K_{1,2}$ in $K_n \setminus J_k$ that form a 4-cycle with two edges in S. As in the setting of Lemma 2.2, let $\mu := |\mathcal{Y}|p^2$ be the expected number of copies of $K_{1,2}$ in $\mathcal{Y}$ that appear in G and $\Delta := \sum_{K,K'} p^{e(K \cup K')}$, where the sum goes over all pairs of copies $K, K' \in \mathcal{Y}$ with $K \cap K' \neq \emptyset$. Note that $\Delta \le \Delta(\mathcal{X}_S)$, as $\mathcal{Y} \subseteq \mathcal{X}_S$, and we also have that

$$\mu = |\mathcal{Y}|p^2 \ge |\Pi(S)|(n - 2\Delta(J_k))p^2 \ge \frac{1}{\sqrt{2}}|\Pi(S)|np^2 \ge \frac{\mu(\mathcal{X}_S)}{\sqrt{2}},$$

using (7) and the inequality $\Delta(J_k) \le 2np \ll n$, which holds as $J_k \subseteq H$ for some $H \in \mathcal{H}' \subseteq \mathcal{H}$. Note also that $\mu(\mathcal{X}_S)/3 \ge |\Pi(S)|np^2/4$. Therefore, by Lemma 4.17 and our choice of α,

$$D := \frac{\mu^2}{800\Delta} \ge \frac{\mu(\mathcal{X}_S)^2}{1600\Delta(\mathcal{X}_S)} \ge \min\left\{\frac{1}{6400}|\Pi(S_i)|np^2, \frac{1}{160}e(R_k)\log\left(\frac{2n^2p}{e(R_k)}\right)\right\} \ge M.$$

In the notation of Corollary 2.3, letting $\mathcal{A}$ be the graph with vertex set $\Gamma = E(K_n)$ whose edges encode copies of $K_{1,2}$ in $\mathcal{Y}$, we have that $Z \sim \nu(\mathcal{A}[\Gamma_p])$, the size of the largest matching in $\mathcal{A}$. In particular, we may apply Corollary 2.3 to conclude that $\Pr[Z < M] \le \Pr[Z < D] \le \exp(-D) \le \exp(-M)$; thus, the claim will follow after showing that $M \ge \zeta n^3 p^5 e(J_k)$. This follows from Claim 4.21 and inequality (15):

$$\frac{M_k}{e(J_k)} \ge \frac{M_k}{3e(R_k)} \ge \frac{\zeta n^3 p^5}{\theta} \ge \zeta n^3 p^5. \qquad ◀$$

It remains to perform a union bound over all possible stamps $\mathbf{S} \in \boldsymbol{\mathcal{S}}$. We have that

$$\Pr[G \in \mathcal{H}'] = \sum_{\mathbf{S} \in \boldsymbol{\mathcal{S}}} \Pr[\mathbf{S}(G) = \mathbf{S}] = \sum_{k=0}^{\log n} \underbrace{\sum_{\mathbf{S} \in \boldsymbol{\mathcal{S}}^k} \Pr[\mathbf{S}(G) = \mathbf{S}]}_{\Sigma^k}.$$

In order to get a grasp on Σ^k, consider the following random process that constructs increasing random sequences $\mathbf{J}^* = (J_0^*, \dots, J_k^*)$ and $\mathbf{R}^* = (R_0^*, \dots, R_k^*)$ of subgraphs of K_n:

Let $\mathcal{C}^*$ be a uniformly chosen random family of $c_0 = \theta n^3 p^3$ triangles in K_n.

Let J_0^* be the union of $\mathcal{C}^*$ and let $R_0^* \subseteq J_0^*$ be a uniformly chosen random subset.

for $i = 0, \dots, k-1$ **do**

  if $R_i^* \in \mathcal{F}$ **then**

    Let $S_i^* := S(R_i^*)$, the graph output by Lemma 4.17.

    Let $\mathcal{X}^* \subseteq \mathcal{X}_{S_i^*}$ be a uniformly chosen random subset with $\alpha|\Pi(S_i^*)|np^2$ elements.

    Let L^* be the union of all copies of $K_{1,2}$ in $\mathcal{X}^*$.

    Let Q^* be a uniformly chosen random subset of L^*.

    Let $J_{i+1}^* := J_i^* \cup L^*$ and $R_{i+1}^* := R_i^* \cup Q^*$.

  else

    Let $J_{i+1}^* := J_i^*$ and $R_{i+1}^* := R_i^*$.

The key observation is that, for any $\mathbf{S} = ((J_0, \dots, J_k), (R_0, \dots, R_k)) \in \boldsymbol{\mathcal{S}}^k$,

$$\Pr\big((\mathbf{J}^*, \mathbf{R}^*) = \mathbf{S}\big) \ge q^*(\mathbf{S}) := \left(\binom{\binom{n}{3}}{c_0} 2^{3c_0} \prod_{i=0}^{k-1} \left(\binom{|\Pi(S_i)|n}{\alpha|\Pi(S_i)|np^2} 2^{2\alpha|\Pi(S_i)|np^2} \right) \right)^{-1}, \tag{16}$$

where, for each i, we denote by $S_i = S(R_i)$ the graph output by Lemma 4.17 with input R_i. Since clearly $\sum_{\mathbf{S} \in \boldsymbol{\mathcal{S}}^k} q^*(\mathbf{S}) \le 1$, we have

$$\Sigma^k \le \max\big\{ \Pr[\mathbf{S}(G) = \mathbf{S}] \cdot q^*(\mathbf{S})^{-1} : \mathbf{S} \in \boldsymbol{\mathcal{S}}^k \big\}.$$

Now, fix some $\mathbf{S} = ((J_0, \dots, J_k), (R_0, \dots, R_k)) \in \boldsymbol{\mathcal{S}}^k$ and let $S_i := S(R_i)$ and $M_i := \alpha|\Pi(S_i)|np^2$ for each $i \in \{0, \dots, k-1\}$. By Claims 4.22 and 4.23, we have that $e(J_k) = 3c_0 + \sum_{i=0}^{k-1} 2M_i$ and thus

$$\begin{aligned}
\Pr[\mathbf{S}(G) = \mathbf{S}] \cdot q^*(\mathbf{S})^{-1} &\le p^{e(J_k)} \exp(-\zeta n^3 p^5 e(J_k)) \cdot \left(\frac{8en^3/6}{\theta n^3 p^3} \right)^{c_0} \prod_{i=0}^{k-1} \left(\frac{4e|\Pi(S_i)|n}{\alpha|\Pi(S_i)|np^2} \right)^{M_i} \\
&\le \exp(-\zeta n^3 p^5 e(J_k)) \cdot \left(\frac{4}{\theta} \right)^{c_0} \prod_{i=0}^{k-1} \left(\frac{4e}{\alpha} \right)^{M_i} \\
&\le \exp\big(-e(J_k) \cdot \big(\zeta n^3 p^5 + \log\theta + \log\alpha\big)\big) \le \exp(-e(J_k)),
\end{aligned}$$

as $n^3p^5 \gg 1$. Since $e(J_k) \ge c_0 \gg n$, we conclude that $\Sigma^k \le e^{-n}$ and, consequently, $\Pr(G \in \mathcal{H}') \ll 1$. ☐

References

[1] N. Alon and J. H. Spencer. *The Probabilistic Method*. Wiley series in Discrete Mathematics and Optimization. Wiley, 4th edition, 2015.

[2] J. Balogh and J. Butterfield. Online Ramsey games for triangles in random graphs. *Discrete Math.*, 310(24):3653–3657, 2010.

[3] J. Balogh, R. Morris, and W. Samotij. Independent sets in hypergraphs. *Journal of the American Mathematical Society*, 28(3):669–709, 2015.

[4] J. Balogh, R. Morris, and W. Samotij. The method of hypergraph containers. In *Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018. Vol. IV. Invited lectures*, pages 3059–3092. World Sci. Publ., Hackensack, NJ, 2018.

[5] H. Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Annals of Mathematical Statistics*, 23:493–507, 1952.

[6] D. Conlon, S. Das, J. Lee, and T. Mészáros. Ramsey games near the critical threshold. *Random Structures Algorithms*, 57(4):940–957, 2020.

[7] P. Frankl and V. Rödl. Large triangle-free subgraphs in graphs without K_4. *Graphs and Combinatorics*, 2(1):135–144, 1986.

[8] E. Friedgut, Y. Kohayakawa, V. Rödl, A. Ruciński, and P. Tetali. Ramsey games against a one-armed bandit. *Combinatorics, Probability and Computing*, 12:515–545, 2003.

[9] E. Friedgut, E. Kuperwasser, W. Samotij, and M. Schacht. Sharp thresholds for ramsey properties. *arXiv preprint arXiv:2207.13982*, 2022.

[10] E. Friedgut, V. Rödl, A. Ruciński, and P. Tetali. A sharp threshold for random graphs with a monochromatic triangle in every edge coloring. *Memoirs of the American Mathematical Society*, 179(845):vi+66, 2006.

[11] S. Janson. Poisson approximation for large deviations. *Random Structures & Algorithms*, 1(2):221–229, 1990.

[12] S. Janson, T. Łuczak, and A. Ruciński. *Random Graphs.* Wiley-Interscience series in Discrete Mathematics and Optimization. John Wiley and Sons, 2000.

[13] T. Łuczak, A. Ruciński, and B. Voigt. Ramsey properties of random graphs. *Journal of Combinatorial Theory, Series B*, 56(1):55–68, 1992.

[14] M. Marciniszyn, R. Spöhel, and A. Steger. Online Ramsey games in random graphs. *Combin. Probab. Comput.*, 18(1-2):271–300, 2009.

[15] M. Marciniszyn, R. Spöhel, and A. Steger. Upper bounds for online Ramsey games in random graphs. *Combin. Probab. Comput.*, 18(1-2):259–270, 2009.

[16] R. Nenadov and A. Steger. A short proof of the random Ramsey theorem. *Combin. Probab. Comput.*, 25(1):130–144, 2016.

[17] J. Nešetřil and V. Rödl. Sparse Ramsey graphs. *Combinatorica*, 4:71–78, 1984.

[18] A. Noever. Online Ramsey games for more than two colors. *Random Structures Algorithms*, 50(3):464–492, 2017.

[19] F. P. Ramsey. On a problem of formal logic. *Proceedings of the London Mathematical Society*, 2:264–286, 1930.

[20] V. Rödl and A. Ruciński. Lower bounds on probability thresholds for Ramsey properties. *Combinatorics, Paul Erdős is eighty*, 1:317–346, 1993.

[21] V. Rödl and A. Ruciński. Random graphs with monochromatic triangles in every edge coloring. *Random Structures & Algorithms*, 5(2):253–270, 1994.

[22] V. Rödl and A. Ruciński. Threshold functions for Ramsey properties. *Journal of the American Mathematical Society*, 8(4):917–942, 1995.

[23] D. Saxton and A. Thomason. Hypergraph containers. *Inventiones mathematicae*, 201(3):925–992, 2015.

A Proof of Corollary 2.3

Our proof follows the approach of [1, Lemma 8.4.2]. We denote

$$\tilde{\Delta} \coloneqq \Delta - \mu = \sum_{i \neq j} \mathbb{1}[A_i \cap A_j \neq \emptyset] \cdot p^{|A_i \cup A_j|}$$

and consider two cases, depending on the ratio $\tilde{\Delta}/\mu$.

Case 1. $\tilde{\Delta} \le \mu/3$. For an index $i \in [m]$, let $\delta_i \coloneqq \sum_{j \neq i} \mathbb{1}[A_i \cap A_j \neq \emptyset] \cdot p^{|A_j \setminus A_i|}$ and call $i \in [m]$ *good* if $\delta_i \le \frac{4\tilde{\Delta}}{\mu}$. Furthermore, let $\Lambda \subseteq [m]$ be the set of good indices and let $\mu_g \coloneqq \sum_{i \in \Lambda} p^{|A_i|}$. We will look for a matching of size $D^* \coloneqq \frac{\mu^2}{50\Delta}$ in $\mathcal{A}[\Gamma_p]$ only among the edges with good indices.

We will call a family $I \subseteq \Lambda$ *disjoint* if $A_i \cap A_j = \emptyset$ for all $i \neq j \in I$. The crucial observation is that, if the event $\nu(\mathcal{A}[\Gamma_p]) \le D^*$ occurs, then there must be some (possibly empty) disjoint family of indices $I \subseteq \Lambda$ of size at most D^* whose all corresponding edges appear in Γ_p that is maximal in the sense that no edge corresponding to the family $\Lambda_I \coloneqq \{j \in \Lambda : \forall i \in I : A_i \cap A_j = \emptyset\}$ appears in Γ_p. Denote the former event (that $A_i \subseteq \Gamma_p$ for all $i \in I$) by Q_I and the latter event (that $A_j \nsubseteq \Gamma_p$ for all $i \in \Lambda_I$) by M_I. Note that $\Pr[Q_I] = \prod_{i \in I} p^{|A_i|}$ and, crucially, that Q_I and M_I are independent.

In order to estimate the probability of M_I, observe first that ignoring bad indices does not have a big effect on the expected number of sets appearing in Γ_p and we have that $\mu_g \ge 3\mu/4$. Indeed, if this were not the case, then we would have that

$$\tilde{\Delta} = \sum_{i \in [m]} p^{|A_i|} \delta_i \ge \frac{4\tilde{\Delta}}{\mu} \sum_{i \in [m] \setminus \Lambda} p^{|A_i|} = \frac{4\tilde{\Delta}}{\mu} \cdot (\mu - \mu_g) > \tilde{\Delta},$$

a contradiction. We further note that

$$\mu_I \coloneqq \sum_{j \in \Lambda_I} p^{|A_j|} \ge \mu_g - \sum_{i \in I} \sum_{j \in \Lambda} \mathbb{1}[A_i \cap A_j \neq \emptyset] \cdot p^{|A_j|} \ge \mu_g - \sum_{i \in I} (1 + \delta_i) \ge \mu_g - |I| \left(1 + \frac{4\tilde{\Delta}}{\mu}\right).$$

In particular, if $|I| \le \frac{\mu^2}{16\Delta}$, then

$$\mu_I \ge \mu_g - \frac{\mu^2}{16\Delta} \cdot \left(1 + \frac{4\tilde{\Delta}}{\mu}\right) = \mu_g - \frac{\mu^2}{16\Delta} \cdot \left(\frac{4\Delta}{\mu} - 3\right) \ge \mu_g - \frac{\mu}{4} \ge \frac{\mu}{2}.$$

We also clearly have that

$$\Delta_I \coloneqq \sum_{i,j \in \Lambda_I} \mathbb{1}[A_i \cap A_j \neq \emptyset] \cdot \mathbb{E}[I_i I_j] \le \Delta.$$

Hence, by Lemma 2.2

$$\Pr[M_I] \le \exp\left(-\frac{\mu_I^2}{2\Delta_I}\right) \le \exp\left(-\frac{\mu^2}{8\Delta}\right).$$

This gives the following estimate:

$$\begin{aligned}
\Pr[\nu(\mathcal{A}[\Gamma_p]) \le D^*] &\le \sum_{|I| \le D^*} \Pr[Q_I \wedge M_I] = \sum_{k=0}^{D^*} \sum_{\substack{I \subseteq \Lambda \\ |I| = k}} \Pr[M_I] \cdot \Pr[Q_I] \\
&\le \exp\left(-\frac{\mu^2}{8\Delta}\right) \cdot \sum_{k=0}^{D^*} \sum_{\substack{I \subseteq \Lambda \\ |I| = k}} \prod_{i \in I} p^{|A_i|} \le \exp\left(-\frac{\mu^2}{8\Delta}\right) \cdot \sum_{k=0}^{D^*} \frac{1}{k!} \left(\sum_{i \in \Lambda} p^{|A_i|}\right)^k \\
&\le \exp\left(-\frac{\mu^2}{8\Delta}\right) \cdot \sum_{k=0}^{D^*} \frac{\mu^k}{k!} \le \exp\left(-\frac{\mu^2}{8\Delta}\right) \cdot \left(\frac{e\mu}{D^*}\right)^{D^*},
\end{aligned}$$

where the last inequality is the well-known concentration inequality for the Poisson distribution that states that, if $x \le \mu$, then $\Pr[\mathrm{Poisson}(\mu) \le x] \le \left(\frac{e\mu}{x}\right)^x e^{-\mu}$. Finally, since we assume that $\Delta = \mu + \tilde{\Delta} \le 3\mu/4$, we conclude that

$$\Pr[\nu(\mathcal{A}[\Gamma_p]) \le D^*] \le \exp\left(D^* \cdot \left(-\frac{25}{4} + \log\frac{50e\Delta}{\mu}\right)\right) \le \exp(-D^*).$$

Case 2. $\tilde{\Delta} > \mu/3$. In this case, we simply find a subset of indices that satisfies the assumption of Case 1. For a set $S \subseteq [m]$ of indices, denote

$$\mu_S \coloneqq \sum_{i \in S} p^{|A_i|} \qquad \text{and} \qquad \tilde{\Delta}_S \coloneqq \sum_{i \neq j \in S} \mathbb{1}[A_i \cap A_j \neq \emptyset] \cdot p^{|A_i \cup A_j|}.$$

It is clearly enough to show that there exists a set S with $\tilde{\Delta}_S \le \mu_S/3$ and $D_S^* \coloneqq \frac{\mu_S^2}{50(\mu_S + \tilde{\Delta}_S)} \ge D$. Let S be a random subset of $[m]$ obtained by independently retaining each element with probability $q \coloneqq \frac{\mu}{6\tilde{\Delta}}$. Then

$$\mathbb{E}[\mu_S - 3\tilde{\Delta}_S] = q\mu - 3q^2\tilde{\Delta} = \frac{\mu^2}{12\tilde{\Delta}},$$

so there is an S that satisfies both $\mu_S \ge \frac{\mu^2}{12\tilde{\Delta}}$ and $\tilde{\Delta}_S \le \mu_S/3$ and thus also

$$D_S^* = \frac{\mu_S^2}{50(\mu_S + \tilde{\Delta}_S)} \ge \frac{3\mu_S}{200} \ge \frac{\mu^2}{800\tilde{\Delta}} \ge \frac{\mu^2}{800\Delta} = D.$$

In particular, arguing as in Case 1, with $\mathcal{A}$ replaced by $\mathcal{A}_S \coloneqq \{A_i \in \mathcal{A} : i \in S\}$, we obtain,

$$\Pr[\nu(\mathcal{A}[\Gamma_p]) \le D] \le \Pr[\nu(\mathcal{A}_S[\Gamma_p]) \le D] \le \Pr[\nu(\mathcal{A}_S[\Gamma_p]) \le D_S^*] \le \exp(-D_S^*) \le \exp(-D).$$

B Derivation of Theorem 2.4

Here we derive our container lemma, Theorem 2.4, which we restate for convenience.

Theorem 2.4. *For every positive integer $2 \le k \in \mathbb{N}$ and all $\varepsilon \in (0,1)$ and $1 \le K \in \mathbb{N}$, there exist $t \in \mathbb{N}$ and $\delta > 0$ such that the following holds. Suppose that a nonempty k-uniform hypergraph $\mathcal{H}$ with vertex set V and $\tau \in (0, 1/t)$ satisfy*

$$\Delta_\ell(\mathcal{H}) \le K\tau^{\ell-1} \cdot \frac{e(\mathcal{H})}{v(\mathcal{H})}$$

for every $\ell \in \{2, \dots, k\}$. Then, there exists a function $f \colon \mathcal{P}(V)^t \to \mathcal{P}(V)$ with the following properties:

(i) For every set $I \subseteq V$ satisfying $e(\mathcal{H}[I]) \le \delta\tau^k e(\mathcal{H})$, there are $S_1, \dots, S_t \subseteq I$, each of size at most $\tau v(\mathcal{H})$ and such that $I \subseteq f(S_1, \dots, S_t)$.

(ii) For every $S_1, \dots, S_t \subseteq V$, the set $f(S_1, \dots, S_t)$ induces fewer than $\varepsilon e(\mathcal{H})$ edges in $\mathcal{H}$.

Theorem 2.4 is a slight reformulation of the following theorem of Saxton and Thomason.

Theorem B.1 ([23, Corollary 3.6])**.** *For every $2 \le k \in \mathbb{N}$ and $\varepsilon > 0$, there exists an integer $s \in \mathbb{N}$ such that the following holds. Suppose that a nonempty k-uniform hypergraph $\mathcal{H}$ on vertex set V and $\tau \in (0, 1/2)$ satisfy*

$$\delta(\mathcal{H}, \tau) \coloneqq 2^{\binom{k}{2}-1} \sum_{j=2}^{k} 2^{-\binom{j-1}{2}} \delta_j(\mathcal{H}, \tau) \le \frac{\varepsilon}{12k!},$$

where

$$\delta_j(\mathcal{H}, \tau) \coloneqq \frac{\tau^{1-j}}{ke(\mathcal{H})} \cdot \sum_{v \in V} \max\{d_{\mathcal{H}}(T) : v \in T \subseteq V \text{ and } |T| = j\}.$$

Then there exists a function $C \colon \mathcal{P}(V)^s \to \mathcal{P}(V)$ such that, letting

$$\mathcal{T} \coloneqq \big\{(T_1, \dots, T_s) \in \mathcal{P}(V)^s : |T_i| \le s\tau|V| \text{ for all } i \in [s]\big\},$$

we have:

(a) *For every set* $I \subseteq V$ *satisfying* $e(\mathcal{H}[I]) \leq 24\varepsilon k!k\tau^k e(\mathcal{H})$, *there exists* $T = (T_1, \ldots, T_s) \in \mathcal{T} \cap \mathcal{P}(I)^s$ *with* $I \subseteq C(T)$.

(b) *For every* $T \in \mathcal{T}$, *the set* $C(T)$ *induces at most* $\varepsilon e(\mathcal{H})$ *edges in* $\mathcal{H}$.

Derivation of Theorem 2.4 from Theorem B.1. Let $\mathcal{H}$ be a nonempty k-uniform hypergraph with vertex set V and let $\varepsilon > 0$ and $K \in \mathbb{N}$. We set $s := s_{B.1}(k, \varepsilon)$ to be the constant output by Theorem B.1 with input k and ε and let

$$L := \left\lceil \frac{12k!2^{\binom{k}{2}}K}{\varepsilon} \right\rceil, \qquad t := 2Ls^2, \qquad \text{and} \qquad \delta := 24\varepsilon k!kL^k.$$

Suppose that the maximum degrees of $\mathcal{H}$ satisfy the assumptions of the theorem for some $\tau \leq 1/t$. Note that, for every $j \in \{2, \ldots, k\}$,

$$\delta_j(\mathcal{H}, L\tau) = \frac{(L\tau)^{1-j}}{ke(\mathcal{H})} \cdot v(\mathcal{H})\Delta_j(\mathcal{H}) \leq \frac{KL^{1-j}}{k}$$

and thus, as $L \geq 2$,

$$\delta(\mathcal{H}, L\tau) \leq 2^{\binom{k}{2}-1} \cdot \sum_{j=2}^{k} \frac{KL^{1-j}}{k} \leq \frac{2^{\binom{k}{2}}K}{L} \leq \frac{\varepsilon}{12k!}.$$

Consequently, Theorem B.1, invoked with $\tau_{B.1} = L\tau$, implies that there exist a function $C\colon \mathcal{P}(V)^s \to \mathcal{P}(V)$ that satisfies conditions (a) and (b) of that theorem. Here, we used that $\tau \leq 1/t$ implies that $L\tau < 1/2$. Now define $f\colon \mathcal{P}(V)^t \to \mathcal{P}(V)$ by letting

$$f(S_1, \ldots, S_t) := C\big(S_1 \cup \cdots \cup S_{t/s}, \ldots, S_{(s-1)t/s+1} \cup \cdots \cup S_t\big).$$

If $I \subseteq V$ satisfies

$$e(\mathcal{H}[I]) \leq \delta\tau^k e(\mathcal{H}) \leq 24\varepsilon k!k(L\tau)^k e(\mathcal{H}),$$

then there are $T_1, \ldots, T_s \subseteq I$, with $|T_i| \leq sL\tau|V|$ for each i, such that $I \subseteq C(T_1, \ldots, T_s)$. This gives the assertion of part (i) of the theorem, as we may partition each T_i into t/s sets $S_{t(i-1)/s+1}, \ldots, S_{t(i+1)/s}$, each of size at most $\lceil s/t \cdot sL\tau|V|\rceil \leq \tau|V|$. Part (ii) of the theorem follows directly from our definition of f and part (b) of Theorem B.1. ∎