



RESEARCH ARTICLE

WILEY

Ramsey games near the critical threshold

David Conlon¹ | Shagnik Das² | Joonkyung Lee³ | Tamás Mészáros²

¹Department of Mathematics, California Institute of Technology, Pasadena, CA

²Institut für Mathematik, Freie Universität Berlin, Berlin, Germany

³Department of Mathematics, University College London, London, UK

Correspondence

Shagnik Das, Institut für Mathematik, Freie Universität Berlin, 14195 Berlin, Germany.
Email: shagnik@mi.fu-berlin.de

Funding information

This research was supported by the Deutsche Forschungsgemeinschaft (DFG) project 415310276; the German-Israeli Foundation for Scientific Research and Development, GIF grant G-1347-304.6/2016 [S.D.]. The Dahlem Research School, DRS Fellowship Program and the Berlin Mathematics Research Center MATH+, project “learning hypergraphs” [T.M.]. The European Research Council (ERC), starting grant RanDM 676632 [D.C.] and ERC consolidator grant PEPCo 724903 [J.L.]. Open access funding enabled and organized by Projekt DEAL.

Abstract

A well-known result of Rödl and Ruciński states that for any graph H there exists a constant C such that if $p \geq Cn^{-1/m_2(H)}$, then the random graph $G_{n,p}$ is a.a.s. H -Ramsey, that is, any 2-coloring of its edges contains a monochromatic copy of H . Aside from a few simple exceptions, the corresponding 0-statement also holds, that is, there exists $c > 0$ such that whenever $p \leq cn^{-1/m_2(H)}$ the random graph $G_{n,p}$ is a.a.s. not H -Ramsey. We show that near this threshold, even when $G_{n,p}$ is not H -Ramsey, it is often extremely close to being H -Ramsey. More precisely, we prove that for any constant $c > 0$ and any strictly 2-balanced graph H , if $p \geq cn^{-1/m_2(H)}$, then the random graph $G_{n,p}$ a.a.s. has the property that every 2-edge-coloring without monochromatic copies of H cannot be extended to an H -free coloring after $\omega(1)$ extra random edges are added. This generalizes a result by Friedgut, Kohayakawa, Rödl, Ruciński, and Tetali, who in 2002 proved the same statement for triangles, and addresses a question raised by those authors. We also extend a result of theirs on the three-color case and show that these theorems need not hold when H is not strictly 2-balanced.

KEYWORDS

Ramsey theory, random graphs, positional games

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2020 The Authors. *Random Structures & Algorithms* published by Wiley Periodicals LLC.

1 | INTRODUCTION

The study of sparse generalizations of combinatorial theorems has attracted considerable interest in recent years and there are now several general mechanisms [2, 3, 16, 17] that allow one to prove that analogues of classical results such as Ramsey’s theorem, Turán’s theorem and Szemerédi’s theorem hold relative to sparse random graphs and sets of integers. Much of this work is based, in one way or another, on the beautiful random Ramsey theorem of Rödl and Ruciński [14, 15] from 1995. This seminal result gives a complete answer to the question of when the binomial random graph $G_{n,p}$ is (H, r) -Ramsey, that is, has the property that any r -coloring of its edges contains a monochromatic copy of the graph H .

To state the Rödl-Ruciński theorem precisely, we need some notation. For a graph H , we write $d_2(H) = 0$ if H has no edges, $d_2(H) = 1/2$ when $H = K_2$ and $d_2(H) = (e(H) - 1)/(v(H) - 2)$ in the general case. We then write $m_2(H) = \max_{H' \subseteq H} d_2(H')$ and call this quantity the 2-density of H . Though we will not use these definitions immediately, we also say that H is 2-balanced if $m_2(H') \leq m_2(H)$ and strictly 2-balanced if $m_2(H') < m_2(H)$ for all proper subgraphs H' of H .

Theorem (Rödl-Ruciński, 1995). *Let $r \geq 2$ be a positive integer and let H be a graph that is not a forest consisting of stars and paths of length 3. Then there are positive constants c and C such that*

$$\lim_{n \rightarrow \infty} \mathbb{P}[G_{n,p} \text{ is } (H, r)\text{-Ramsey}] = \begin{cases} 0 & \text{if } p < cn^{-1/m_2(H)}, \\ 1 & \text{if } p > Cn^{-1/m_2(H)}. \end{cases}$$

There has been much work extending this result. We will not attempt an exhaustive survey, but refer the interested reader instead to some of the latest progress on hypergraphs [8], the asymmetric case [11], establishing sharp thresholds [18] and the equivalent problem in settings other than the binomial random graph [5, 13]. Our particular concern here will be with the following surprising result of Friedgut, Kohayakawa, Rödl, Ruciński, and Tetali [6] regarding two-round Ramsey games against a random Builder.

Theorem (Friedgut-Kohayakawa-Rödl-Ruciński-Tetali, 2003). *Let $c > 0$ be fixed and, for $p = cn^{-1/2}$, let $G = G_{n,p}$. Then, with high probability, the following statements hold:*

- (a) *Let φ_2 be an arbitrary monochromatic- K_3 -free 2-edge-coloring of G . If $q_2 = \omega(n^{-2})$, then, with high probability, φ_2 cannot be extended to a monochromatic- K_3 -free 2-edge-coloring of $G \cup G_{n,q_2}$.*
- (b) *Let φ_3 be an arbitrary monochromatic- K_3 -free 3-edge-coloring of G . If $q_3 = \omega(n^{-1})$, then, with high probability, φ_3 cannot be extended to a monochromatic- K_3 -free 3-edge-coloring of $G \cup G_{n,q_3}$.*

When $H = K_3$, the Rödl-Ruciński theorem implies that if $p = Cn^{-1/2}$ for some sufficiently large C , then every 2-edge-coloring contains a monochromatic triangle. Part (a) of the theorem above says that for any $c > 0$, no matter how small, if $p = cn^{-1/2}$, then, even though there are 2-edge-colorings of $G_{n,p}$ containing no monochromatic K_3 , no such coloring can be extended to a monochromatic- K_3 -free 2-edge-coloring after $\omega(1)$ extra random edges are added. One interpretation of this result is that for any $c > 0$ the random graph $G_{n,p}$ with $p = cn^{-1/2}$ is, with high probability, already extremely close to being $(K_3, 2)$ -Ramsey. Part (b) gives a similar result for 3-edge-colorings, though in this case $\omega(n)$ extra edges may be needed in the second round of coloring to guarantee a monochromatic triangle.

Addressing a problem raised by Friedgut, Kohayakawa, Rödl, Ruciński, and Tetali [6], our main result says that a similar statement holds for all graphs H containing an edge h for which $m_2(H \setminus h) < m_2(H)$. In particular, the result applies when H is strictly 2-balanced, since any edge h works in this case.

Theorem 1.1. *Let H be a graph and suppose that there is some edge $h \in E(H)$ whose removal decreases the 2-density, that is, $m_2(H \setminus h) < m_2(H)$. Let $c > 0$ be fixed and, for $p = cn^{-1/m_2(H)}$, let $G = G_{n,p}$. Then, with high probability, the following statements hold:*

- (a) *Let φ_2 be an arbitrary monochromatic- H -free 2-edge-coloring of G . If $q_2 = \omega(n^{-2})$, then, with high probability, φ_2 cannot be extended to a monochromatic- H -free 2-edge-coloring of $G \cup G_{n,q_2}$.*
- (b) *Let φ_3 be an arbitrary monochromatic- H -free 3-edge-coloring of G . If $q_3 = \omega(n^{-1/m(H)})$, then, with high probability, φ_3 cannot be extended to a monochromatic- H -free 3-edge-coloring of $G \cup G_{n,q_3}$.*

Observe that the densities q_2 and q_3 of the random graphs that must be added to create monochromatic copies of H are best possible. Indeed, if $q = O(n^{-2})$, then with positive probability $G_{n,q}$ has no edges, so φ_2 trivially extends to $G \cup G_{n,q}$. Only slightly less trivially, if φ_3 only uses two of the three colors on the edges of G , then we can color all the edges of $G_{n,q}$ with the third color. If $q = O(n^{-1/m(H)})$, then with positive probability $G_{n,q}$ is H -free, thus giving a valid extension of φ_3 . Finally, note that these results cannot be extended to $r \geq 4$ colors, since the two random graphs $G_{n,p}$ and $G_{n,q}$ can be colored independently with disjoint pairs of colors, so we can avoid creating a monochromatic copy of H until the density of one of the two random graphs exceeds the random Ramsey threshold $Cn^{-1/m_2(H)}$ from Theorem 1.

2 | THE NECESSITY OF A CONDITION

In Theorem 1.1, we impose the condition that there is some edge $h \in E(H)$ such that $H \setminus h$ has a strictly lower 2-density than H . While this condition covers, for example, strictly 2-balanced graphs (where the edge h can be chosen arbitrarily), it is natural to ask whether it is necessary. In this section we show that Theorem 1.1 does not apply to all graphs H , so some condition is indeed required.

2.1 | Edge-rooted products of graphs

We first define the edge-rooted product of graphs.

Definition 2.1. Let G be a graph, let H be a graph rooted at an edge $h = \{u, v\} \in E(H)$ and let $k \in \mathbb{N}$. To build the k -fold edge-rooted product $G \Theta_k(H, h)$, we start with a central copy of G and then attach k copies of H to each edge $g = \{x, y\} \in E(G)$ such that $\{x, y\}$ is the root-edge h in each copy of H and all other vertices in each copy are new and distinct.

In other words, $V(G \Theta_k(H, h)) = V(G) \cup (E(G) \times [k] \times (V(H) \setminus \{u, v\}))$, $V(G)$ induces a copy of G and, for each $g = \{x, y\} \in E(G)$ and $i \in [k]$, $(\{g\} \times \{i\} \times (V(H) \setminus \{u, v\})) \cup \{x, y\}$ induces a copy of H with $\{x, y\}$ playing the role of $\{u, v\}$. (Note that there is some slack in this definition, since we have not prescribed an orientation for each attached copy of H . In practice, the particular choice of orientation makes no difference, so we will simply assume that some fixed choice has been made).

The reduced k -fold edge-rooted product, denoted $G \odot_k(H, h)$, is the subgraph obtained by removing all the edges from the central copy of G .

We have already defined $d_2(H)$, $m_2(H)$ and stated what it means for a graph to be 2-balanced or strictly 2-balanced. In a similar fashion, we write $d_1(H) = 0$ if H has no edges and $d_1(H) = e(H)/(v(H) - 1)$ otherwise. We then write $m_1(H) = \max_{H' \subseteq H} d_1(H')$ and call this quantity the 1-density of H . We say that H is 1-balanced if $m_1(H') \leq m_1(H)$ and strictly 1-balanced if $m_1(H') < m_1(H)$ for

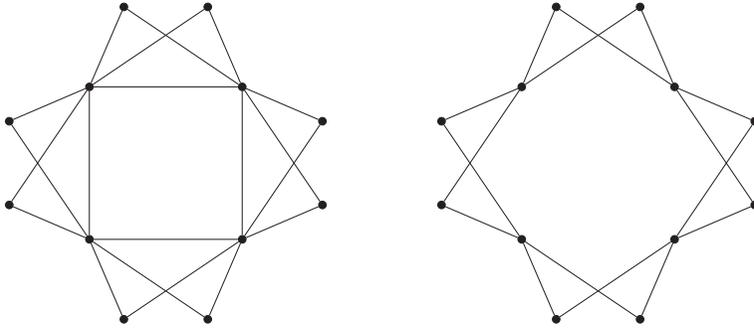


FIGURE 1 $C_4 \ominus_2 K_3$ (on the left) and $C_4 \odot_2 K_3$ (on the right)

all proper subgraphs H' of H . Finally, write $d(H) = e(H)/v(H)$ and $m(H) = \max_{H' \subsetneq H} d(H')$, which we call the *density* of H . We then say that H is *balanced* if $m(H') \leq m(H)$ and *strictly balanced* if $m(H') < m(H)$ for all proper subgraphs H' of H . We will make repeated use of the following simple lemma in what follows.

Lemma 2.2. *If H is 2-balanced with $d_2(H) > 1$, then H is strictly 1-balanced and strictly balanced.*

Proof. Suppose that H is not strictly 1-balanced and let $F \subset H$ be a subgraph with $d_1(F) \geq d_1(H)$. That is, $e(F)/(v(F) - 1) \geq e(H)/(v(H) - 1)$ or, by multiplying the expression out,

$$e(F)v(H) - e(F) \geq e(H)v(F) - e(H). \tag{1}$$

Since H is 2-balanced, we have $d_2(H) \geq d_2(F)$, which implies that $(e(H) - 1)/(v(H) - 2) \geq (e(F) - 1)/(v(F) - 2)$. Rearranging gives $e(H)v(F) - 2e(H) - v(F) + 2 \geq e(F)v(H) - 2e(F) - v(H) + 2$. Substituting (1), we get

$$e(H)v(F) - 2e(H) - v(F) + 2 \geq e(F)v(H) - 2e(F) - v(H) + 2 \geq e(H)v(F) - e(H) - e(F) - v(H) + 2.$$

Canceling the like terms gives $-e(H) - v(F) \geq -e(F) - v(H)$, which in turn implies that $(e(F) - 1) - (v(F) - 2) \geq (e(H) - 1) - (v(H) - 2)$, which can be rewritten as $(d_2(F) - 1)(v(F) - 2) \geq (d_2(H) - 1)(v(H) - 2)$. However, this is a contradiction, since by assumption $d_2(H) - 1 \geq d_2(F) - 1$ and $v(H) - 2 > v(F) - 2$. The argument in the strictly balanced case follows along almost exactly the same lines. ■

The key observation for our purposes is that the edge-rooted product behaves well with respect to the various graph densities.

Lemma 2.3. *For any graphs G and H of density at least 1, any edge $h \in E(H)$ and any $k \in \mathbb{N}$:*

- (a) *if G is strictly balanced, H is 2-balanced and $d(G) < d_2(H)$, then $G \ominus_k(H, h)$ is strictly balanced,*
- (b) *$m_2(G \ominus_k(H, h)) = \max\{m_2(G), m_2(H)\}$ and*
- (c) *if $m_2(H \setminus h) < m_2(H)$, then $m_2(G \odot_k(H, h)) < m_2(H)$.*

Proof. (a) Let $F \subseteq G \ominus_k(H, h)$ be a smallest (induced) subgraph maximizing $d(F) = e(F)/v(F)$. We wish to show that $F = G \ominus_k(H, h)$. We start with a lower bound on the density of

$G \Theta_k(H, h)$:

$$\begin{aligned}
 m(G \Theta_k(H, h)) &\geq d(G \Theta_k(H, h)) = \frac{e(G) + ke(G)(e(H) - 1)}{v(G) + ke(G)(v(H) - 2)} \\
 &= \frac{e(H) - (1 - \frac{1}{k})}{v(H) - 1 - (1 - \frac{1}{kd(G)})} \geq \frac{e(H)}{v(H) - 1} = d_1(H) = m_1(H), \quad (2)
 \end{aligned}$$

where the inequality on the second line follows since either $kd(G) = 1$, in which case we have equality, or $(1 - \frac{1}{k}) / (1 - \frac{1}{kd(G)}) \leq 1 \leq d(H) < d_1(H)$. Note that the final equality, $d_1(H) = m_1(H)$, is an application of Lemma 2.2.

We also observe that $d(G \Theta_k(H, h))$ is a convex combination of $d(G)$ and $d_2(H)$:

$$\frac{e(G) + ke(G)(e(H) - 1)}{v(G) + ke(G)(v(H) - 2)} = d(G) \frac{v(G)}{v(G \Theta_k(H, h))} + d_2(H) \left(1 - \frac{v(G)}{v(G \Theta_k(H, h))} \right). \quad (3)$$

Now, for each $g \in E(G)$ and $i \in [k]$, let $F_{g,i} \subseteq H$ be the subgraph induced by the vertices of F in the i th copy of H attached to the edge g in the central copy of G . Let $F_0 \subseteq G$ be the subgraph induced by the vertices of F in the central copy of G .

By the minimality of the size of F , we may assume that F is connected, as otherwise its densest component would be a smaller subgraph attaining the maximum density. We cannot have $F \subseteq H$ since, by (2), the density of $G \Theta_k(H, h)$ is at least $m_1(H)$, which is strictly larger than $m(H)$. Thus, $F_{g,i}$ must be nonempty for at least two pairs $(g, i) \in E(G) \times [k]$ and, hence, to be connected, each nonempty $F_{g,i}$ must contain at least one vertex of g .

Now suppose there was some (g, i) such that $F_{g,i}$ contained only one of the two endpoints of g . Then, by removing $F_{g,i}$ from F , we lose $e(F_{g,i})$ edges and $v(F_{g,i}) - 1$ vertices. Since $F_{g,i} \subset H$, the ratio $e(F_{g,i}) / (v(F_{g,i}) - 1)$ is at most $m_1(H)$, which by (2) is at most $m(G \Theta_k(H, h))$. Removing $F_{g,i}$ would therefore not decrease the density of F , contradicting the minimality of its size. Thus, if $F_{g,i}$ is nonempty, we must have $g \in E(F_{g,i})$. Hence,

$$v(F) = v(F_0) + \sum_{g \in E(F_0)} \sum_{i=1}^k (v(F_{g,i}) - 2) \quad \text{and} \quad e(F) = e(F_0) + \sum_{g \in E(F_0)} \sum_{i=1}^k (e(F_{g,i}) - 1). \quad (4)$$

Thus,

$$d(F) = \frac{e(F_0) + \sum_{g \in E(F_0)} \sum_{i=1}^k (e(F_{g,i}) - 1)}{v(F_0) + \sum_{g \in E(F_0)} \sum_{i=1}^k (v(F_{g,i}) - 2)} = d(F_0) \frac{v(F_0)}{v(F)} + \sum_{g \in E(F_0)} \sum_{i=1}^k d_2(F_{g,i}) \frac{v(F_{g,i}) - 2}{v(F)}.$$

Since $F_0 \subseteq G$ and G is balanced, $d(F_0) \leq d(G)$. Similarly, for each g and i , $d_2(F_{g,i}) \leq d_2(H)$. We therefore have

$$d(F) \leq d(G) \frac{v(F_0)}{v(F)} + d_2(H) \left(1 - \frac{v(F_0)}{v(F)} \right).$$

Comparing this to (3), since $d(G) < d_2(H)$, for $d(F) \geq d(G \Theta_k(H, h))$ to hold we require

$$\frac{v(F_0)}{v(F)} \leq \frac{v(G)}{v(G \Theta_k(H, h))} = \frac{1}{1 + kd(G)(v(H) - 2)}. \quad (5)$$

Now $v(F) = v(F_0) + \sum_{g \in E(F_0)} \sum_{i=1}^k (v(F_{g,i}) - 2) \leq v(F_0) + ke(F_0)(v(H) - 2)$, with equality if and only if $F_{g,i} = H$ for all $g \in E(F_0)$ and $i \in [k]$. Therefore,

$$\frac{v(F_0)}{v(F)} \geq \frac{v(F_0)}{v(F_0) + ke(F_0)(v(H) - 2)} = \frac{1}{1 + kd(F_0)(v(H) - 2)}.$$

Thus, in order to satisfy the inequality of (5), $d(F_0) \geq d(G)$. As G is strictly balanced, it follows that $F_0 = G$ and then, since $F_{g,i} = H$ for all g and i , we have $F = G \odot_k (H, h)$, as required.

- (b) Since $G, H \subseteq G \odot_k (H, h)$, we immediately have $m_2(G \odot_k (H, h)) \geq \max\{m_2(G), m_2(H)\}$. The proof of the upper bound follows the same lines as in part (a). Let $F \subseteq G \odot_k (H, h)$ be a smallest subgraph realizing the 2-density, that is, $m_2(G \odot_k (H, h)) = d_2(F) = (e(F) - 1)/(v(F) - 2)$. Let F_0 and, for each $(g, i) \in E(G) \times [k]$, $F_{g,i}$ be defined as in part (a). We may assume that $F_{g,i} \neq \emptyset$ for at least two pairs (g, i) , since otherwise $F \subseteq H$ and thus $d_2(F) \leq m_2(H)$. By the minimality of the size of F , we may further assume that F is 2-connected, as otherwise one of the blocks B of F will satisfy $m_2(B) \geq m_2(F)$ (see, for instance, Lemma 8 of [12]). In particular, this implies that $g \in E(F_{g,i})$ whenever $F_{g,i} \neq \emptyset$.

The vertices and edges of F can then be enumerated as in (4), so

$$d_2(F) = \frac{e(F) - 1}{v(F) - 2} = \frac{e(F_0) - 1 + \sum_{g \in E(F_0)} \sum_{i=1}^k (e(F_{g,i}) - 1)}{v(F_0) - 2 + \sum_{g \in E(F_0)} \sum_{i=1}^k (v(F_{g,i}) - 2)}. \tag{6}$$

Since $e(F_0) - 1 \leq m_2(G)(v(F_0) - 2)$ and $e(F_{g,i}) - 1 \leq m_2(H)(v(F_{g,i}) - 2)$ for each (g, i) , it follows that $d_2(F) = m_2(G \odot_k (H, h)) \leq \max\{m_2(G), m_2(H)\}$.

- (c) The product $G \odot_k (H, h)$ is obtained by deleting the edges of the central copy of G from the product $G \odot_k (H, h)$. We show $m_2(G \odot_k (H, h)) < m_2(H)$ by following the argument of part (b). To start, let $F \subseteq G \odot_k (H, h)$ be a smallest subgraph attaining the 2-density.

As before, let F_0 be the subgraph of G induced by the vertices of F from the central copy of G and, for $g \in E(G)$ and $i \in [k]$, let $F_{g,i}$ be the subgraph of H induced by the vertices of F in the i th copy of H attached to the edge g . Note that, since the edges from the central copy of G are deleted in $G \odot_k (H, h)$, neither the edges of F_0 nor the edge g in $F_{g,i}$ (if present) appear in F . However, it will be convenient for us to include them in F_0 and $F_{g,i}$ for our calculations.

If $F_{g,i}$ is only nonempty for one pair of (g, i) , then $F = F_{g,i} \setminus g \subseteq H \setminus h$, so $d_2(F) \leq m_2(H \setminus h) < m_2(H)$. Otherwise, since F must be 2-connected, g must be in $F_{g,i}$ whenever $F_{g,i}$ is nonempty. We can then compute the 2-density of F as in part (b), arriving at an expression similar to (6), except the edges in F_0 do not appear in F . Thus,

$$d_2(F) = \frac{-1 + \sum_{g \in E(F_0)} \sum_{i=1}^k (e(F_{g,i}) - 1)}{v(F_0) - 2 + \sum_{g \in E(F_0)} \sum_{i=1}^k (v(F_{g,i}) - 2)} < \frac{\sum_{g \in E(F_0)} \sum_{i=1}^k (e(F_{g,i}) - 1)}{\sum_{g \in E(F_0)} \sum_{i=1}^k (v(F_{g,i}) - 2)} \leq m_2(H),$$

since $F_{g,i} \subseteq H$ implies $e(F_{g,i}) - 1 \leq m_2(H)(v(F_{g,i}) - 2)$. ■

2.2 | Graphs requiring unusually many extra random edges

Part (c) of Lemma 2.3 shows the role played by the assumption of the existence of the edge h in Theorem 1.1. We will show how to use this to prove Theorem 1.1 in the next section, but first we use

the other parts of this lemma to construct graphs for which the conclusion of Theorem 1.1 does not hold.

Theorem 2.4. *Let F be a 2-balanced graph containing a cycle, let $f \in E(F)$ be an arbitrary edge of F and let $H = F \ominus_1 (F, f)$. Let $G = G_{n,p}$ for $p = cn^{-1/m_2(H)}$, where $c > 0$ is a sufficiently small constant. Then, with high probability, the following statements hold:*

- (a) *There is a monochromatic- H -free 2-edge-coloring of G such that if $q = o(n^{-v(F)/e(F)})$ the coloring can with high probability be extended to a coloring of $G \cup G_{n,q}$ without monochromatic copies of H .*
- (b) *There is a monochromatic- H -free 3-edge-coloring of G and $\delta = \delta(H) > 0$ such that if $q = o(n^{-1/m(H)+\delta})$ the coloring can with high probability be extended to a coloring of $G \cup G_{n,q}$ without monochromatic copies of H .*

Proof. Since F is 2-balanced and contains a cycle, we have $d(F) \geq 1$ and $d_2(F) > 1$. Using Lemma 2.3(b), $m_2(H) = m_2(F)$. Hence, if the constant c is sufficiently small, the Rödl-Ruciński theorem implies that with high probability we can find a 2-coloring φ of $E(G)$ without any monochromatic copy of F . This is the edge-coloring we extend in both cases.

- (a) In this case, extend φ to the edges of $G_{n,q}$ arbitrarily. Observe that $H = F \ominus_1 (F, f)$ consists of $e(F)$ edge-disjoint copies of F . Since there are no monochromatic copies of F in G , any monochromatic copy of H in $G \cup G_{n,q}$ must contain at least $e(F)$ edges from $G_{n,q}$.

There are at most $n^{v(H)}$ potential copies of H and $2^{e(H)}$ ways to distribute its edges between G and $G_{n,q}$. Since $q \leq p$, the probability that a copy with at least $e(F)$ edges from $G_{n,q}$ appears in $G \cup G_{n,q}$ is at most $p^{e(H)-e(F)} q^{e(F)}$. Thus, by the union bound, the probability that there is a copy of H in $G \cup G_{n,q}$ with at least $e(F)$ edges from $G_{n,q}$ is at most

$$n^{v(H)} 2^{e(H)} p^{e(H)-e(F)} q^{e(F)} = 2^{e(H)} c^{e(H)-e(F)} n^{v(F)+e(F)(v(F)-2)-e(F)(e(F)-1)/m_2(H)} q^{e(F)},$$

where we used that $v(H) = v(F) + e(F)(v(F) - 2)$ and $e(H) = e(F)^2$. As $m_2(H) = m_2(F) = \frac{e(F)-1}{v(F)-2}$, this simplifies to $2^{e(H)} c^{e(H)-e(F)} n^{v(F)} q^{e(F)}$, which is $o(1)$ by our choice of q . Hence, with high probability our arbitrary extension of φ to $G \cup G_{n,q}$ does not create a monochromatic copy of H .

- (b) The coloring φ uses two colors, say red and blue. This leaves us with one unused color, say green, that we can use when extending φ to the edges of $G_{n,q}$.

As F is 2-balanced, Lemma 2.2 implies that it is also strictly balanced. By Lemma 2.3(a), it follows that H is strictly balanced. As a consequence, any union of two copies of H that share at least an edge must be strictly denser than H itself. Indeed, the subgraph common to both copies of H is a proper subgraph and therefore strictly sparser than H . Hence, the vertices and edges added in the second copy in the union must increase the overall density.

There is thus some $\delta_1 = \delta_1(H) > 0$ such that, whenever $q = o(n^{-1/m(H)+\delta_1})$, intersecting copies of H do not appear in $G_{n,q}$. That is, the copies of H appearing in $G_{n,q}$ are with high probability pairwise edge-disjoint. We shall choose our $\delta(H)$ to be less than this $\delta_1(H)$.

We now order the edges of $G_{n,q}$ arbitrarily and process them one-by-one. We color each edge green, unless that would create a green copy of H , in which case we color the edge red. When coloring in this fashion, if we create a monochromatic copy of H , it clearly must be red.

Consider a red copy H_0 of H in our coloring of $G \cup G_{n,q}$. Since H_0 is an edge-disjoint union of $e(F)$ copies of F and the coloring φ of G has no monochromatic copy of F , each copy of F in H_0 must contain at least one red edge from $G_{n,q}$. An edge e from $G_{n,q}$ is only red if it is the last edge of an otherwise green copy H_e of H , which must be wholly contained in $G_{n,q}$. Moreover, for $e \neq e'$, the copies H_e and $H_{e'}$ of H are edge-disjoint.

This gives us a subgraph of $G \cup G_{n,q}$ with at most $v(F) + e(F)(v(F) - 2 + v(H) - 2)$ vertices and $e(F)(e(F) + e(H) - 1)$ edges, of which at least $e(F)e(H)$ edges come from $G_{n,q}$. Since $q \leq p$ and there are at most some constant K ways of building such a subgraph and dividing its edges between G and $G_{n,q}$, the probability of finding such a structure is at most

$$Kn^{v(F)+e(F)(v(F)-2+v(H)-2)} p^{e(F)(e(F)-1)} q^{e(F)e(H)}.$$

Since $p \leq n^{-1/m_2(H)} = n^{-(v(F)-2)/(e(F)-1)}$, this is at most $Kn^{v(F)+e(F)(v(H)-2)} q^{e(F)e(H)}$. Now $q = o(n^{-1/m(H)+\delta}) = o(n^{-v(H)/e(H)+\delta})$, since H is strictly balanced. Thus the upper bound on the probability of the appearance of a red copy of H is $o(n^{v(F)-2e(F)+\delta e(F)e(H)}) = o(n^{e(F)(\delta e(H)-1)})$, since $e(F) \geq v(F)$. Hence, if we choose $\delta \leq \min\{1/e(H), \delta_1(H)\}$, this probability is $o(1)$, so with high probability we can extend the coloring to the edges of $G_{n,q}$ without creating a monochromatic copy of H . ■

3 | THE PROOF OF THEOREM 1.1

Having shown in the previous section that some condition on the graph H is necessary in Theorem 1.1, we now show that our condition is sufficient. We begin with a sketch of the proof and then recall several useful results before providing the details of the argument.

3.1 | An overview of the proof

We shall assume the colors used are red, blue and, in the case of three-colorings, green. Our goal is to find structures in the first random graph, G , that force the creation of a monochromatic copy of H no matter how the edges of the second random graph, $G_{n,q}$, are colored. To that end, we make the following definitions.

Definition 3.1 (Color-forced edges). A copy of $H \setminus h$ in G is *supported* on the pair $\{x, y\}$ if $\{x, y\}$ maps to the missing edge h . We then call $\{x, y\}$ the *base* of the copy. Given an edge-coloring φ , we say $\{x, y\}$ is a *red, blue or green base* if it is the base of a monochromatic copy of $H \setminus h$ of the corresponding color. Finally, we say a pair $\{x, y\}$ is *green-forced* if it is both a red and a blue base simultaneously, with *blue-forced* and *red-forced* defined similarly.

In the two-color case, observe that it is impossible to extend φ_2 to a green-forced pair, since coloring it either red or blue would create a monochromatic copy of H . For the first assertion of Theorem 1.1, we shall show that with high probability G is such that every two-coloring φ_2 admits quadratically many green-forced pairs. Then, again with high probability when $q_2 = \omega(n^{-2})$, one of these pairs will be an edge of the second random graph G_{n,q_2} , so any extension of φ_2 to $G \cup G_{n,q_2}$ will create a monochromatic copy of H .

When dealing with three colors, our goal will instead be to show that there is some color, say green, such that the green-forced pairs in G are sufficiently dense that, when $q_3 = \omega(n^{-1/m(H)})$, we will find

a copy of H in $G_{n,q}$ consisting solely of green-forced pairs. If any one of its edges is colored red or blue, it will complete a monochromatic copy of H with edges from G . On the other hand, if all of its edges are colored green, we obtain a green copy of H instead.

To find these color-forced structures, we consider the reduced graph of a regular partition of G (with respect to the coloring φ_2 or φ_3). In this reduced graph we will find two colors, say red and blue, and a copy of $H \odot_2 (H, h)$ such that for each (removed) edge from the central copy of H , one of the attached copies of (H, h) is monochromatic red and the other is monochromatic blue. By applying the sparse counting lemma, we will deduce the existence of many potential copies of H consisting of green-forced edges, from which we will be able to draw the desired conclusion.

Although the proof can be simplified in the two-colored setting, for the sake of brevity we shall present a single unified argument allowing for three colors throughout, and only differentiate between the two cases at the end of the proof.

3.2 | Some preliminaries

Here we collect several results about random graphs and sparse regularity that we shall use in our proof.

3.2.1 | Random graphs

The Hoeffding inequality shows that $G_{n,p}$ does not have any subgraphs that are far sparser or denser than expected with high probability.

Proposition 3.2. *Let $\eta > 0$ be fixed and suppose $p = \omega(n^{-1})$. Then, with high probability, $G_{n,p}$ is such that the following holds for any disjoint sets X, Y of vertices with $|X|, |Y| \geq \eta n$:*

- (i) $\frac{1}{2} \binom{|X|}{2} p \leq e(G_{n,p}[X]) \leq 2 \binom{|X|}{2} p$ and
- (ii) $\frac{1}{2} |X| |Y| p \leq e(G_{n,p}[X, Y]) \leq 2 |X| |Y| p$.

A simple application of Markov’s inequality also shows that $G_{n,p}$ is unlikely to contain many more copies of any subgraph than expected.

Proposition 3.3. *Given any graph F with v vertices and e edges and any $K > 1$, the probability that there are more than $Kn^v p^e$ copies of F in $G_{n,p}$ is at most $1/K$.*

In the other direction, we can use Chebyshev’s inequality to establish the existence of subgraphs in $G_{n,p}$ when p is suitably large. More precisely, it follows from Theorem 4.4.5 in The Probabilistic Method by Alon and Spencer [1] that if $p = \omega(n^{-1/m(F)})$, then the number of copies of F in $G_{n,p}$ is concentrated around its expectation.

Proposition 3.4. *Given a graph F on v vertices and a constant $\zeta > 0$, let \mathcal{F} be a collection of ζn^v potential copies of F . If $p = \omega_n n^{-1/m(F)}$ with some $\omega_n = \omega(1)$, then the probability that $G_{n,p}$ does not contain a copy of F from \mathcal{F} is at most $\frac{v! 2^v}{\zeta \omega_n}$.*

If the edge probability p is even larger, then the following result, a consequence of Theorem 3.29 from the book Random Graphs by Janson, Łuczak, and Ruciński [9], shows that there will be many pairwise edge-disjoint copies of H in $G_{n,p}$.

Proposition 3.5. *For every graph H with $m_2(H) > 1$, there is a constant $\kappa = \kappa(H)$ such that, given constants $\rho, c > 0$ and setting $p = cn^{-1/m_2(H)}$, with high probability every induced subgraph of $G_{n,p}$ on at least $\frac{1}{2}\rho n$ vertices contains at least $\kappa c^{e(H)-1} \rho^{v(H)} n^2 p$ edge-disjoint copies of H .*

3.2.2 | Sparse regularity and counting

Given an n -vertex graph G , two disjoint sets of vertices X and Y form an (ϵ, p) -regular pair of density d if $d(X, Y) = d$ and, for all $X' \subseteq X$ with $|X'| \geq \epsilon |X|$ and $Y' \subseteq Y$ with $|Y'| \geq \epsilon |Y|$, we have $|d(X', Y') - d(X, Y)| < \epsilon p$, where $d(U, V)$ denotes $\frac{e(U, V)}{|U||V|}$. This notion of regularity is inherited by induced and random subgraphs (see Lemma 4.3 in [7]).

Proposition 3.6. *Suppose that $c \in (0, \frac{1}{2}]$, G is a graph and U, W are disjoint vertex sets, both of size N , with (U, W) an (ϵ, p) -regular pair of density $d = \omega(N^{-1})$. Then the following is true:*

- (i) *for $X \subset U$ and $Y \subset W$ with $|X|, |Y| \geq cN$, the pair (X, Y) is $(\epsilon/c, p)$ -regular with density at least $d - \epsilon p$ and*
- (ii) *for $m \geq cdN^2$, the subgraph G' of G obtained by choosing m edges from $G[U, W]$ uniformly at random forms a $(2\epsilon, p)$ -regular pair with high probability.*

An (ϵ, p) -regular partition \mathcal{P} of G is a partition $V(G) = V_0 \cup V_1 \cup \dots \cup V_k$ such that $|V_0| \leq \epsilon n$, $|V_1| = |V_2| = \dots = |V_k|$ and all but at most ϵk^2 pairs (V_i, V_j) , $1 \leq i < j \leq k$, are (ϵ, p) -regular. When the graph G is edge-colored, we say a partition is (ϵ, p) -regular if for all but at most ϵk^2 pairs of parts the edges of each color between the two parts form an (ϵ, p) -regular subgraph. If G has density d , we say it is (η, D) -upper-uniform if, for all disjoint sets X and Y of size at least ηn , we have $d(X, Y) \leq Dd$. With these definitions in place, we may state a version of the sparse regularity lemma, originally due to Kohayakawa and Rödl [10].

Theorem 3.7. *For all $\epsilon, D > 0$ and $r, t \in \mathbb{N}$, there are $\eta > 0$ and $T \in \mathbb{N}$ such that every r -coloring of the edges of an (η, D) -upper-uniform graph G of density d on at least T vertices has an (ϵ, d) -regular partition \mathcal{P} with k parts for some $k \in [t, T]$.*

The final ingredient we will need is a sparse counting lemma due to Conlon, Gowers, Samotij and Schacht [4]. Given a graph H , integers N and m , and $\epsilon, p > 0$, we define the family $\mathcal{G}(H, N, m, p, \epsilon)$ to be all graphs obtained by replacing each vertex of H by an independent set of size N and replacing each edge of H by an (ϵ, p) -regular bipartite graph with exactly m edges. Given such a graph G , let $G(H)$ denote the number of canonical copies of H in G (by which we mean that each vertex of H in the copy belongs to the corresponding independent set in G).

Theorem 3.8. *For every graph H and every $d > 0$, there exist $\epsilon, \xi > 0$ with the following property. For every $\eta > 0$, there is $C > 0$ such that if $p \geq Cn^{-1/m_2(H)}$, then, with high probability, for every $N \geq \eta n$, $m \geq dpN^2$ and every subgraph G of $G_{n,p}$ in $\mathcal{G}(H, N, m, p, \epsilon)$, $G(H) \geq \xi N^{v(H)} \left(\frac{m}{N^2}\right)^{e(H)}$.*

3.3 | The reduced graph

With these preliminaries in hand, we can proceed with the proof of Theorem 1.1. We begin by describing the (standard) construction of the reduced graph and proving it has some useful properties.

Let t, ϵ, α be defined such that $1/t \leq \epsilon \ll \alpha \ll \kappa, c$, where $\kappa = \kappa(H)$ is the constant from Proposition 3.5, $p = cn^{-1/m_2(H)}$ and ‘ \ll ’ means these parameters are sufficiently small for the subsequent calculations to hold.

Now consider a monochromatic- H -free 3-edge-coloring φ of the edges of $G \sim G_{n,p}$ (where, as in the case of φ_2 , we may only be using two of the three colors) and let G_{red}, G_{blue} and G_{green} represent the red, blue and green subgraphs of G , respectively. Given our choice of ϵ and t and setting $r = 3$ and $D = 4$, let η and T be as in Theorem 3.7. Proposition 3.2 shows that G is with high probability $(\eta, 4)$ -upper-uniform. Hence, there is an (ϵ, p) -regular partition $V(G) = V_0 \cup V_1 \cup \dots \cup V_k$, where $t \leq k \leq T$.

We next define three graphs, $\Gamma_{red}, \Gamma_{blue}$ and Γ_{green} , on the same vertex set $[k]$. Γ_{red} has an edge between i and j if and only if the bipartite induced subgraph $G_{red}[V_i, V_j]$ forms an (ϵ, p) -regular pair of density at least αp , with Γ_{blue} and Γ_{green} defined similarly with respect to G_{blue} and G_{green} , respectively. The reduced (multi)graph Γ is the colored union of Γ_{red} in red, Γ_{blue} in blue and Γ_{green} in green. Given a vertex $i \in [k]$, we write $N_{red}(i), N_{blue}(i)$ and $N_{green}(i)$ for its neighborhoods in $\Gamma_{red}, \Gamma_{blue}$ and Γ_{green} , respectively, and write $d_{red}(i), d_{blue}(i)$ and $d_{green}(i)$ for the sizes of these sets.

We first show that any induced subgraph of Γ with linearly many vertices has a vertex with large degree in at least two of the colors.

Lemma 3.9. Define $f(\rho) = \frac{1}{24} \kappa c^{e(H)-1} \rho^{v(H)-1}$. Suppose ρ satisfies

$$6\rho f(\rho) \geq 3\epsilon + \frac{1}{2}\alpha. \tag{7}$$

Then, with high probability, for any subset $U \subseteq [k]$ of ρk vertices of the reduced graph Γ , we can find a vertex $u \in U$, two disjoint sets $X_1, X_2 \subset U$ of size at least $f(\rho)k$ and two distinct colors χ_1, χ_2 such that, for each $i \in [2]$, u is adjacent to all vertices in X_i with edges of color χ_i .

Proof. Let $W = \cup_{u \in U} V_u$ be the vertices in the parts of G corresponding to the vertices of U and note that $|W| \geq (1 - \epsilon)\rho n \geq \frac{1}{2}\rho n$. Hence, by Proposition 3.5, we may assume $G[W]$ contains at least $24\rho f(\rho)n^2p$ edge-disjoint copies of H . Since there are no monochromatic copies of H in the 3-edge-coloring of G , each such copy must contain two edges of distinct colors. It easily follows that there are two colors, say red and blue, that each appear on at least $12\rho f(\rho)n^2p$ edges of $G[W]$.

Using Proposition 3.2, we observe that all but at most $(3\epsilon + \frac{1}{2}\alpha)n^2p$ red edges of $G[W]$ are contained within dense (ϵ, p) -regular pairs. Indeed, at most $k \cdot 2\binom{n/k}{2}p \leq n^2p/k \leq \epsilon n^2p$ edges can be contained within the parts V_i , at most $\epsilon k^2 \cdot 2(n/k)^2p \leq 2\epsilon n^2p$ edges can be within irregular pairs (V_i, V_j) and at most $\binom{k}{2} \cdot (n/k)^2 \alpha p \leq \frac{1}{2}\alpha n^2p$ edges are within (ϵ, p) -regular pairs (V_i, V_j) of density less than αp . From (7), it follows that there are at least $6\rho f(\rho)n^2p$ red edges in $G[W]$ that are contained in (ϵ, p) -regular pairs of density at least αp . Again by Proposition 3.2, each such pair can account for at most $2(n/k)^2p$ edges in $G[W]$, so there must be at least $3\rho f(\rho)k^2$ such pairs, each of which corresponds to an edge of $\Gamma_{red}[U]$. By symmetry, we also find at least $3\rho f(\rho)k^2$ edges in $\Gamma_{blue}[U]$.

Now let $A = \{a \in U : d_{red}(a, U) \geq 2f(\rho)k\}$. By summing the red degrees of vertices in U , distinguishing between those in A and those not, we have

$$6\rho f(\rho)k^2 \leq \rho k \cdot |A| + 2f(\rho)k \cdot \rho k,$$

from which we deduce that $|A| \geq 4f(\rho)k$. Defining $B = \{b \in U : d_{blue}(b, U) \geq 2f(\rho)k\}$, we similarly have $|B| \geq 4f(\rho)k$. If $A \cap B \neq \emptyset$, let $u \in A \cap B$. Since $d_{red}(u, U), d_{blue}(u, U) \geq 2f(\rho)k$, we can find the required disjoint sets X_1 and X_2 of size $f(\rho)k$ of red and blue neighbors, respectively.

Otherwise, for every $a \in A$ and $b \in B$, by Proposition 3.2, there are at least $\frac{1}{2}(n/k)^2p$ edges in G between V_a and V_b , so one of the three colors appears on at least $\frac{1}{6}(n/k)^2p > \alpha(n/k)^2p$ edges. Let χ_1 be the color that appears most commonly as the majority color in these $|A| |B|$ pairs. Ignoring the pairs that give rise to irregular pairs in Γ_{χ_1} , it follows that there are at least $\frac{1}{3} |A| |B| - \epsilon k^2$ edges in Γ_{χ_1} between A and B . Provided α is sufficiently large with respect to ϵ , (7) and our lower bound on $|A|, |B|$ imply this is at least $\frac{1}{4} |A| |B|$ edges.

If χ_1 is not red, then take χ_2 to be red and, by averaging, find some $u \in A$ with a set X_1 of at least $\frac{1}{4} |B| \geq f(\rho)k$ neighbors in B in the color χ_1 . Since $u \in A$, we have $d_{red}(u, U) \geq 2f(\rho)k$, so we can find a disjoint set X_2 of $f(\rho)k$ red neighbors of u , as required. Otherwise, if χ_1 is red, we take χ_2 to be blue. By the same argument, we can find some $u \in B$ with a set X_1 of at least $f(\rho)k$ red neighbors in A and, since $u \in B$, it has large enough degree in $\Gamma_{blue}[U]$ to guarantee a disjoint set X_2 of blue neighbors. ■

Through repeated use of this lemma, we can build large multicolored structures in Γ .

Corollary 3.10. *Given $t \in \mathbb{N}$, let $\rho_0 = 1$ and, for $1 \leq i \leq 2t - 2$, let $\rho_i = f(\rho_{i-1})$. Provided $6\rho_{2t-3}f(\rho_{2t-3}) \geq 3\epsilon + \frac{1}{2}\alpha$, there is with high probability a vertex v_0 of Γ contained in two monochromatic t -cliques of distinct colors.*

Proof. Applying Lemma 3.9 with $\rho = \rho_0 = 1$ and $U = [k]$, we find a vertex u_0 with large degrees in two colors. Without loss of generality, let the colors be red and blue. In the first stage of this algorithm, we iterate within the red neighborhood of u_0 , finding either a vertex in blue and green t -cliques, in which case we are done, or a red t -clique containing u_0 .

To start, apply Lemma 3.9 again, this time taking U to be the set of red neighbors of u_0 . This gives a vertex with large degrees in two colors. While one of those colors is red, we repeat the process, giving us a sequence of vertices with large nested red neighborhoods. If this sequence (including u_0) has length $t - 1$, by choosing an arbitrary vertex in the final red neighborhood, we obtain a red t -clique containing u_0 and can proceed to the second stage.

Otherwise, after some $h \leq t - 2$ steps we obtain a vertex u_1 that has large blue and green neighborhoods. In this case, we first iterate within the blue neighborhood of u_1 . Each subsequent vertex has either a large red neighborhood or a large blue neighborhood within which we can proceed. Once we have obtained a sequence of $2t - 3$ vertices, there are either $t - 1$ of them (including u_0) for which we iterated within a red neighborhood or $t - 1$ of them (including u_1) for which we iterated within a blue neighborhood. In the first case, we choose an arbitrary vertex in the final neighborhood to create a red t -clique containing u_0 and can then proceed to the second stage.

In the second case, choosing an arbitrary vertex in the final neighborhood gives a blue t -clique containing u_1 . We can then return to the green neighborhood of u_1 and repeatedly iterate, at each point proceeding with a red or green neighborhood of the latest vertex. Once we reach a sequence of length $2t - 3$ (including the vertices between u_0 and u_1), we again either have $t - 1$ vertices with green neighborhoods or $t - 1$ vertices with red neighborhoods. In the first case, we can complete a green t -clique containing u_1 that, together with the earlier blue t -clique, completes the desired structure. In the second case, choosing a vertex in the final neighborhood again completes a red t -clique containing u_0 , with which we proceed to the second stage.

If we proceed to the second stage, we will have already found a red t -clique containing u_0 . The second stage consists of mirroring the above process in the blue neighborhood of u_0 . This results in a blue t -clique containing u_0 or a vertex u_2 in the blue neighborhood that is contained in both red and green t -cliques; in either case, we are done. ■

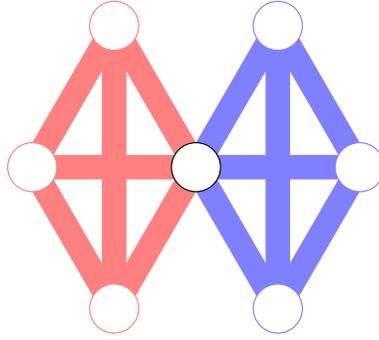


FIGURE 2 The parts of G corresponding to the two cliques from Corollary 3.10 [Colour figure can be viewed at wileyonlinelibrary.com]

3.4 | Building color-forced structures

Let $\mathcal{K}_n(H)$ denote the family of all copies of H in K_n . Using the cliques from Corollary 3.10, we will prove the following key proposition.

Proposition 3.11. *There are positive constants $\kappa = \kappa(c, H)$ and $\zeta = \zeta(c, H)$ such that, for any $K > 1$, with probability at least $1 - \kappa K^{-1} - o(1)$, for every monochromatic- H -free 3-edge-coloring φ of G , there is some color χ with at least $\zeta K^{-1} n^{v(H)}$ χ -forced copies of H in $\mathcal{K}_n(H)$.*

Proof. For convenience, we write $v = v(H)$ and $e = e(H)$. Setting $t = e(v - 2) + 1$ and applying Corollary 3.10, which holds with high probability, we find a vertex x in the reduced graph Γ that is in, say, both a red and a blue K_t . Let u_1, u_2, \dots, u_{t-1} be the other vertices from the red clique and w_1, w_2, \dots, w_{t-1} be the other vertices from the blue clique. Consider the corresponding parts in the graph G . We know that, for all $1 \leq i < j \leq t - 1$, the pairs $G_{red}[V_x, V_{u_i}], G_{red}[V_{u_i}, V_{u_j}], G_{blue}[V_x, V_{w_i}]$ and $G_{blue}[V_{w_i}, V_{w_j}]$ are all (ϵ, p) -regular pairs of density at least αp . This situation is illustrated below in the case $H = K_3$.

Partition the part V_x into v equal-sized subsets, X_1, X_2, \dots, X_v , letting N denote the size of these sets. Define η by $N = \eta n$, noting that $\eta \geq \frac{1-\epsilon}{kv}$, where we recall that $k \leq T$ is the number of parts in the (ϵ, p) -regular partition of G . For each i , let $R_i \subset V_{u_i}$ and $B_i \subset V_{w_i}$ be arbitrary subsets of size N . Let $\mathcal{X} = \{X_1, \dots, X_v\}$, $\mathcal{R} = \{R_1, \dots, R_{t-1}\}$ and $\mathcal{B} = \{B_1, \dots, B_{t-1}\}$. By Proposition 3.6(i), it follows that the pairs $G_{red}[X_i, R_j], G_{red}[R_i, R_j], G_{blue}[X_i, B_j]$ and $G_{blue}[B_i, B_j]$ are all $(\epsilon v, p)$ -regular of density at least $(\alpha - \epsilon)p$.

Next consider the graph $H \odot_2(H, h)$ and note that it has precisely $v+2(t-1)$ vertices, with one central copy H_0 of H , whose edges are deleted, and each deleted edge $g \in E(H_0)$ supporting two otherwise vertex-disjoint copies $H_{g,1}$ and $H_{g,2}$ of H . We can build a bijection $\psi : V(H \odot_2(H, h)) \rightarrow \mathcal{X} \cup \mathcal{R} \cup \mathcal{B}$ such that:

- $\psi(H_0) = \mathcal{X}$ and
- for all $g \in E(H_0)$, $\psi(V(H_{g,1}) \setminus g) \subset \mathcal{R}$ and $\psi(V(H_{g,2}) \setminus g) \subset \mathcal{B}$.

That is, for each edge $g \in E(H_0)$, we send one of the attached copies of H to the red parts \mathcal{R} and the other copy to the blue parts \mathcal{B} .

Let $m = \frac{1}{2} \alpha p N^2$ and consider an edge $f = \{y, z\} \in E(H \odot_2(H, h))$. If $f \in E(H_{g,1})$ for some $g \in E(H_0)$, then the pair $G_{red}[\psi(y), \psi(z)]$ is an $(\epsilon v, p)$ -regular pair of density at least $(\alpha - \epsilon)p$. Define $\psi(f) \subseteq G_{red}[\psi(y), \psi(z)]$ to be the subgraph obtained from this pair by choosing m edges uniformly at random.

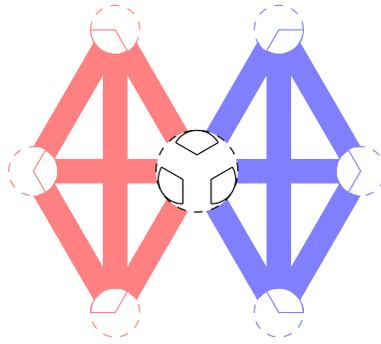


FIGURE 3 We divide the central part into $v(H)$ subsets and shrink the other parts accordingly [Colour figure can be viewed at wileyonlinelibrary.com]

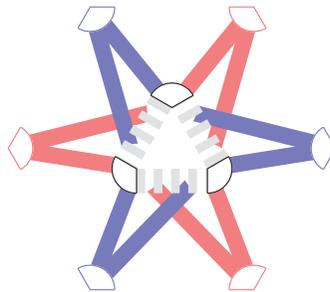


FIGURE 4 We imagine a copy of H between the subsets of the central part, with each edge supporting both a red and a blue copy of $H \setminus h$ using the parts from the red and blue cliques [Colour figure can be viewed at wileyonlinelibrary.com]

By Proposition 3.6, $\psi(f)$ is $(2\varepsilon v, p)$ -regular with high probability. Otherwise, $f \in E(H_{g,2})$ for some $g \in E(H_0)$, in which case we define $\psi(f)$ to be the subgraph obtained by selecting m edges uniformly at random from $G_{blue}[\psi(y), \psi(z)]$. We again have, with high probability, that $\psi(f)$ is $(2\varepsilon v, p)$ -regular.

Now define the subgraph $G' \subset G$ to be the union of all these subgraphs $\psi(f)$, that is,

$$V(G') = \bigcup_{y \in V(H \odot_2(H, h))} \psi(y) \quad \text{and} \quad E(G') = \bigcup_{f \in E(H \odot_2(H, h))} \psi(f).$$

From the above discussion, it is clear that $G' \in \mathcal{C}(H \odot_2(H, h), N, m, p, 2\varepsilon v)$, where this family of graphs is as defined before Theorem 3.8. Since $p = cn^{-1/m_2(H)}$ and, by Lemma 2.3(c), $m_2(H \odot_2(H, h)) < m_2(H)$, we can apply Theorem 3.8. This gives some constant $\xi > 0$ such that there are with high probability at least $\xi N^{v(H \odot_2(H, h))} \binom{m}{N^2}^{e(H \odot_2(H, h))}$ copies of $H \odot_2(H, h)$ in G' , where each vertex y comes

from the set $\psi(y)$. To simplify this expression, we define $c' = \xi \eta^{v(H \odot_2(H, h))} \left(\frac{1}{2}\alpha\right)^{e(H \odot_2(H, h))}$ and $\mu = n^{v-2} p^{e-1}$. Our lower bound on the number of copies of $H \odot_2(H, h)$ can then be written as $c' \mu^{2e} n^v$. Note that $c' > 0$ is a constant, while, since $p = cn^{-1/m_2(H)} \geq cn^{-(v-2)/(e-1)}$, $\mu = \Omega(1)$.

In each such copy of $H \odot_2(H, h)$, each missing edge $g \in E(H_0)$ in the central copy of H supports both a red copy $H_{g,1}$ of $H \setminus h$ and a blue copy $H_{g,2}$ of $H \setminus h$. In particular, this means g is green-forced and, as this holds for all edges g , this shows that the central copy H_0 forms a green-forced copy of H in $\mathcal{K}_n(H)$.

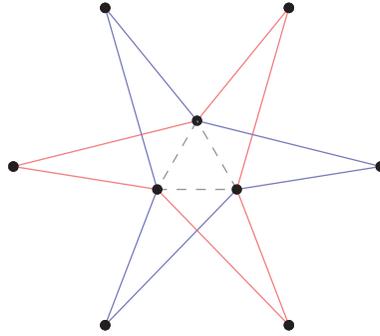


FIGURE 5 Applying Theorem 3.8 gives many copies of $H \odot_2 (H, h)$ in which the central copy of H is green-forced [Colour figure can be viewed at wileyonlinelibrary.com]

However, we are not quite done, as these green-forced copies of H may contribute to multiple copies of $H \odot_2 (H, h)$, in which case they will have been overcounted. To rectify this, and complete the proof, we now show that most of these copies of H are not counted too often.

To this end, suppose we have found r distinct green-forced central copies H_0 of H above and enumerate them as $H^{(1)}, H^{(2)}, \dots, H^{(r)}$. For each $1 \leq i \leq r$, let Z_i denote the number of copies of $H \odot_2 (H, h)$ found above in which $H^{(i)}$ is the central copy H_0 . We have thus far established that

$$\sum_{i=1}^r Z_i \geq c' \mu^{2e} n^v, \tag{8}$$

while we wish to show that $r \geq \frac{c}{K} n^v$.

Now consider the quantity Z_i^2 . This counts the number of ordered pairs (A, B) of canonical copies of $H \odot_2 (H, h)$ with $H^{(i)}$ as the central copy H_0 . Given such a pair, let $J = A \cup B$. In J , each edge $g \in H_0$ is contained in a red copy of $H \setminus h$ from A and one from B as well. The same holds true for the blue copies of $H \setminus h$. These attached copies of $H \setminus h$ in J are mostly disjoint outside the central H_0 , except that the two copies of the same color supported on the same edge g may share some vertices. We consider such a graph J as a degenerate copy of $H \odot_4 (H, h)$. There are several isomorphism classes J could belong to, depending on which vertices are shared by A and B .

For each edge $g \in H_0$, let $F_{g,1}$ be the subgraph of H induced by the vertices shared between the red copies of $H \setminus h$ in A and B supported on g and define $F_{g,2}$ analogously for the blue copies. We include the edge g in $F_{g,1}$ and $F_{g,2}$, even though it does not appear in J . Note that the union $\cup_{g \in E(H_0)} \cup_{j=1}^2 F_{g,j}$ determines the isomorphism class of J . Hence, there are at most $2^{v(H \odot_2 (H, h))}$ possible isomorphism types, as for each vertex in $H \odot_2 (H, h)$, we can decide whether or not it belongs to the corresponding $F_{g,j}$. Set $\kappa = 2^{v(H \odot_2 (H, h))}$.

We shall use Proposition 3.3 to show that, regardless of isomorphism type, there cannot be many copies of J in G . Indeed, we have

$$v(J) = v(H) + \sum_{g \in E(H_0)} \sum_{j=1}^2 (2(v(H) - 2) - (v(F_{g,j}) - 2)) = v + 4e(v - 2) - \sum_{g,j} (v(F_{g,j}) - 2)$$

and

$$e(J) = \sum_{g \in E(H_0)} \sum_{j=1}^2 (2(e(H) - 1) - (e(F_{g,j}) - 1)) = 4e(e - 1) - \sum_{g,j} (e(F_{g,j}) - 1).$$

This gives

$$n^{v(J)} p^{e(J)} = n^{v+4e(v-2) - \sum_{g,j} (v(F_{g,j})-2)} p^{4e(e-1) - \sum_{g,j} (e(F_{g,j})-1)} = \frac{\mu^{4e} n^v}{\prod_{g,j} (n^{v(F_{g,j})-2} p^{e(F_{g,j})-1})}.$$

Since $F_{g,j} \subseteq H$ and $p = cn^{-1/m_2(H)}$, we have $n^{v(F_{g,j})-2} p^{e(F_{g,j})-1} \geq c^{e(F_{g,j})-1}$ for all g, j . Thus, $n^{v(J)} p^{e(J)} \leq c^{-2e^2} \mu^{4e} n^v$. Hence, by Proposition 3.3, with probability at least $1 - K^{-1}$ there are at most $Kc^{-2e^2} \mu^{4e} n^v$ copies of J in G . Taking a union bound over all isomorphism classes, we find that with probability at least $1 - \kappa K^{-1}$, there are at most $\kappa Kc^{-2e^2} \mu^{4e} n^v$ of these degenerate copies of $H \odot_4 (H, h)$ in G .

We noted earlier that each pair (A, B) of copies of $H \odot_2 (H, h)$ counted by $\sum_i Z_i^2$ gives rise to a degenerate copy $J = A \cup B$ of $H \odot_4 (H, h)$. To reverse the correspondence, for each vertex in $H \odot_4 (H, h) \setminus (A \cap B)$, we must decide how to assign the corresponding vertices of J to A and B . Thus, there are at most $\kappa^2 \geq 2^{v(H \odot_4 (H, h))}$ pairs (A, B) giving rise to the same $J = A \cup B$.

Putting all this together, we have, with probability at least $1 - \kappa K^{-1}$,

$$\sum_{i=1}^r Z_i^2 \leq \kappa^3 Kc^{-2e^2} \mu^{4e} n^v.$$

Define $I = \{i : Z_i \geq \frac{2}{c'} \cdot \kappa^3 Kc^{-2e^2} \mu^{2e}\}$. It then follows from the above inequality that $\sum_{i \in I} Z_i \leq \frac{1}{2} c' \mu^{2e} n^v$. Plugging this into (8), we obtain $\sum_{i \notin I} Z_i \geq \frac{1}{2} c' \mu^{2e} n^v$. As there are at most r summands, each of which has size less than $\frac{2}{c'} \cdot \kappa^3 Kc^{-2e^2} \mu^{2e}$, we can conclude that

$$r \geq \left(\frac{(c')^2 c^{2e^2}}{4\kappa^3 K} \right) n^v.$$

Setting $\zeta = \frac{1}{4} (c')^2 c^{2e^2} \kappa^{-3}$ completes the proof. ■

3.5 | Finishing the proof

We begin with part (a). Suppose we have a monochromatic- H -free 2-edge-coloring φ_2 of G and $q_2 = \omega_n n^{-2}$ for some $\omega_n \rightarrow \infty$. Set $K = \omega_n^{1/2}$. By Proposition 3.11, with probability $1 - \kappa K^{-1} - o(1) = 1 - o(1)$, there is some color χ such that there are at least $\zeta K^{-1} n^{v(H)}$ χ -forced copies of H . As the coloring φ_2 only has red and blue edges, the color χ must be green.

Each edge can be in at most $n^{v(H)-2}$ green-forced copies of H , so there must be at least $\zeta K^{-1} n^2$ green-forced edges. If any of these edges were to appear in G_{n, q_2} , we would not be able to extend the coloring φ_2 , as coloring the edge red or blue creates a monochromatic copy of H . Hence, the probability that φ_2 extends to $G \cup G_{n, q_2}$ is at most

$$(1 - q_2)^{\zeta K^{-1} n^2} \leq \exp(-\zeta K^{-1} n^2 q_2) = \exp(-\zeta K) = o(1),$$

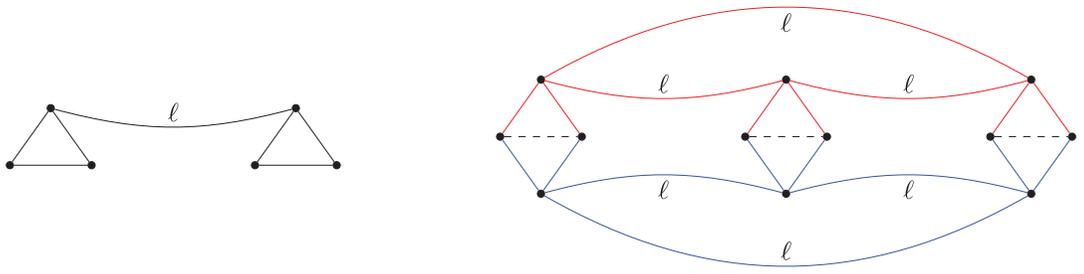


FIGURE 6 H , drawn on the left, consists of two triangles joined by a path of length ℓ . On the right, F_{red} is drawn in red, F_{blue} in blue and the matching M is drawn with dashed lines [Colour figure can be viewed at wileyonlinelibrary.com]

as required.

Part (b) follows the same lines. We begin as before: given the coloring φ_3 and some $q_3 = \omega_n n^{-1/m(H)}$, where $\omega_n \rightarrow \infty$, we set $K = \omega_n^{1/2}$. By Proposition 3.11, with probability $1 - o(1)$, there is some color χ with at least $\zeta K^{-1} n^{v(H)}$ χ -forced copies of H .

If any χ -forced copy of H appears in G_{n,q_3} , then φ_3 cannot be extended. Indeed, coloring all of its edges with the color χ clearly creates a monochromatic copy of H , but since all the edges are χ -forced, using any other color on an edge also completes a monochromatic copy. By Proposition 3.4, the probability that none of the χ -forced copies of H appear in G_{n,q_3} is at most

$$\frac{v(H)! 2^{v(H)} K}{\zeta \omega_n} = \frac{v(H)! 2^{v(H)}}{\zeta K} = o(1),$$

as desired. This completes the proof of Theorem 1.1.

4 | CONCLUDING REMARKS

Our investigations point to several open problems, perhaps the most interesting of which is to classify all graphs H for which Theorem 1.1 holds. We have shown that our condition, that there exists an edge h such that $m_2(H \setminus h) < m_2(H)$, cannot be entirely dispensed with. However, there are also examples of graphs which do not satisfy this condition, but still satisfy some of the conclusions of Theorem 1.1.

Indeed, our proof of Theorem 1.1(a) readily generalizes to the following statement.

Theorem 4.1. *Given a graph H , suppose there are graphs F_{red} , F_{blue} and a matching M such that*

- (i) $V(F_{red}) \cap V(F_{blue}) = V(M)$, with $V(M)$ forming an independent set in F_{red} and F_{blue} ,
- (ii) $m_2(H) > m_2(F_{red} \cup F_{blue})$,
- (iii) $m_2(H) \geq \frac{e(J)}{v(J) - v(M)}$ for all $J \subseteq F_{red} \cup F_{blue}$ with $V(M) \subset V(J)$ and $e(J) \geq 1$ and
- (iv) for any partition of the matching $M = M_{red} \cup M_{blue}$, H is a subgraph of $F_{red} \cup M_{red}$ or $F_{blue} \cup M_{blue}$.

Let $c > 0$ be fixed and, for $p = cn^{-1/m_2(H)}$, let $G = G_{n,p}$. Then, with high probability, the following holds. Let φ be an arbitrary monochromatic- H -free 2-edge-coloring of G . If $q = \omega(n^{-2})$, then, with high probability, φ cannot be extended to a monochromatic- H -free 2-edge-coloring of $G \cup G_{n,q}$.

One of the simplest examples satisfying the conditions of Theorem 4.1 is the graph H consisting of two triangles joined by a path of length $\ell \geq 2$. In this case we can take M to have size three with the corresponding graphs F_{red} and F_{blue} depicted above.

This example shows that the condition in Theorem 1.1 is not best possible, as Theorem 4.1 applies to a wider class of graphs. However, there is a subtle trade-off in finding appropriate forcing structures $F_{red} \cup F_{blue}$ for Theorem 4.1—we need them to be sparse enough to satisfy (ii) and (iii), but to have enough copies of H for (iv).

ACKNOWLEDGMENTS

Part of this work was carried out while the third author visited the second and fourth authors at FU Berlin and he is grateful for their hospitality. Open access funding enabled and organized by Projekt DEAL.

REFERENCES

1. N. Alon and J. Spencer. *The probabilistic method*, Wiley-Interscience Series in Discrete Mathematics and Optimization, Wiley-Interscience, New Jersey, 2008.
2. J. Balogh, R. Morris, and W. Samotij, *Independent sets in hypergraphs*, J. Amer. Math. Soc. **28** (2015), 669–709.
3. D. Conlon and W. T. Gowers, *Combinatorial theorems in sparse random sets*, Ann. Math. **184** (2016), 367–454.
4. D. Conlon, W. T. Gowers, W. Samotij, and M. Schacht, *On the KLR conjecture in random graphs*, Israel J. Math. **203** (2014), 535–580.
5. S. Das and A. Treglown, *Ramsey properties of randomly perturbed graphs: cliques and cycles*, Combin. Probab. Comput. (to appear).
6. E. Friedgut, Y. Kohayakawa, V. Rödl, A. Ruciński, and P. Tetali, *Ramsey games against a one-armed bandit*, Combin. Probab. Comput. **12** (2003), 515–545.
7. S. Gerke and A. Steger. *The sparse regularity lemma and its applications*, Surveys in combinatorics 2005. London Math. Soc. Lecture Note Ser., 327, Cambridge Univ. Press, Cambridge, 2005, pp. 227–258.
8. L. Gugelmann, R. Nenadov, Y. Person, N. Škorić, A. Steger, and H. Thomas, *Symmetric and asymmetric Ramsey properties in random hypergraphs*, Forum Math. Sigma. **5** (2017), e28, 47 pp.
9. S. Janson, T. Łuczak, and A. Ruciński. *Random graphs*, Wiley-Interscience Series in Discrete Mathematics and Optimization, Wiley-Interscience, New York, 2000.
10. Y. Kohayakawa. *Szemerédi's regularity lemma for sparse graphs*, Foundations of Computational Mathematics (Rio de Janeiro, 1997), 216–230, Springer, Berlin, 1997.
11. F. Mousset, R. Nenadov, and W. Samotij, *Towards the Kohayakawa-Kreuter conjecture on asymmetric Ramsey properties*, Combin. Probab. Comput. (to appear).
12. R. Nenadov and A. Steger, *A short proof of the random Ramsey theorem*, Combin. Probab. Comput. **25** (2016), 130–144.
13. E. Powierski, *Ramsey properties of randomly perturbed dense graphs*, preprint available at arXiv:1902.02197 [math.CO].
14. V. Rödl and A. Ruciński. *Lower bounds on probability thresholds for Ramsey properties*, Combinatorics, Paul Erdős is eighty, Bolyai Soc. Math. Stud., János Bolyai Math. Soc., Budapest, Vol. 1, 1993, pp. 317–346.
15. V. Rödl and A. Ruciński, *Threshold functions for Ramsey properties*, J. Amer. Math. Soc. **8** (1995), 917–942.
16. D. Saxton and A. Thomason, *Hypergraph containers*, Invent. Math. **201** (2015), 925–992.
17. M. Schacht, *Extremal results for random discrete structures*, Ann. Math. **184** (2016), 333–365.
18. M. Schacht and F. Schulenburg, *Sharp thresholds for Ramsey properties of strictly balanced nearly bipartite graphs*, Random Structures Algorithms. **52** (2018), 3–40.

How to cite this article: Conlon D, Das S, Lee J, Mészáros T. Ramsey games near the critical threshold. *Random Struct Alg.* 2020;57:940–957. <https://doi.org/10.1002/rsa.20959>