Forum of Mathematics, Sigma (2023), **Vol. 11**:e84 1–15
doi:10.1017/fms.2023.89

CAMBRIDGE UNIVERSITY PRESS

RESEARCH ARTICLE

Tower Gaps in Multicolour Ramsey Numbers

Quentin Dubroff[1], António Girão[2], Eoin Hurley[3] and Corrine Yap[4]

[1]Department of Mathematics, Rutgers University, Piscataway, NJ, 08854, USA; E-mail: qcd2@math.rutgers.edu.
[2]Mathematical Institute, University of Oxford, Oxford, OX2 6GG, United Kingdom; E-mail: antonio.girao@maths.ox.ac.uk.
[3]Korteweg-de Vries Institute for Mathematics, Universiteit van Amsterdam, Amsterdam, Netherlands; E-mail: e.p.hurley@uva.nl.
[4]School of Mathematics, Georgia Institute of Technology, Atlanta, GA, 30332, USA; E-mail: math@corrineyap.com.

Received: 23 March 2022; **Revised:** 29 January 2023; **Accepted:** 25 August 2023

2020 Mathematics Subject Classification: *Primary* – 05C55; *Secondary* – 05D10, 05C65

Abstract

Resolvinga problem of Conlon, Fox and Rödl, we construct a family of hypergraphs with arbitrarily large tower height separation between their 2-colour and q-colour Ramsey numbers. The main lemma underlying this construction is a new variant of the Erdős–Hajnal stepping-up lemma for a generalized Ramsey number $r_k(t;q,p)$, which we define as the smallest integer n such that every q-colouring of the k-sets on n vertices contains a set of t vertices spanning fewer than p colours. Our results provide the first tower-type lower bounds on these numbers.

Contents

1. Introduction

Let $K_n^{(k)}$ denote the complete k-uniform hypergraph on n vertices. We define $r_k(G;q)$ for $k,q \in \mathbb{N}$ as the smallest integer n such that in every q-colouring of $K_n^{(k)}$, there is a monochromatic copy of the hypergraph G. For simplicity, when G is $K_t^{(k)}$, we write $r_k(G;q) = r_k(t;q)$. Observe that when $q = 2$, $r_k(G;2)$ and $r_k(t;2)$ coincide with the classical Ramsey numbers $r_k(G)$ and $r_k(t)$, and we will denote them as such. One of the most central open problems in Ramsey theory is determining the growth rate of the 3-uniform Ramsey number $r_3(t)$.

A famous result of Erdős, Hajnal and Rado [9] from the 60s shows that there exist constants c and c' such that

$$2^{ct^2} \leq r_3(t) \leq 2^{2^{c't}}.$$

Note that the upper bound is essentially exponential in the lower bound. Despite much attention, this remains the state of the art. Perhaps surprisingly, if we allow four colours instead of two, Erdős and Hajnal (see, for example, [12]) showed that the double-exponential upper bound is essentially correct (i.e., there is a $c > 0$ such that $r_3(t;4) \geq 2^{2^{ct}}$). More recently, Conlon, Fox and Sudakov [6] proved a super-exponential bound with three colours – that is, that there exists $c > 0$ such that $r_3(t;3) \geq 2^{t^{c \log t}}$. Erdős conjectured that the double-exponential bound should hold without using extra colours, offering \$500 dollars for a proof that $r_3(t) \geq 2^{2^{ct}}$ for some constant $c > 0$. Raising the stakes for this conjecture is the ingenious *stepping-up* construction of Erdős and Hajnal (see, for example, [12]), which shows that for all q and $k \geq 3$,

$$r_{k+1}(2t + k - 4; q) > 2^{r_k(t;q)-1}. \tag{1.1}$$

For the past 60 years, we have used (1.1) to stack our lower bounds for $r_k(t;q)$ upon that of $r_3(t;q)$, yielding that $r_k(t) \geq T_{k-1}(ct^2)$, where $T_k(x)$, *the tower of height k in x*, is defined by $T_1(x) = x$, $T_{i+1}(x) = 2^{T_i(x)}$. The corresponding upper bounds of $r_k(t) \leq T_k(O(t))$ (see [7; 8; 9]) are once again exponential in the lower bounds, and thus a positive resolution of Erdős's conjecture would be the decisive step in showing that $r_k(t) = T_k(\Theta(t))$ for all $k \geq 3$.

Due to the lack of progress on this central conjecture, it is natural to try to understand just how significant a role the number of colours can play in hypergraph Ramsey numbers and whether or not there could really be such a large difference between $r_3(t)$ and $r_3(t;4)$. One argument in favour of the conjecture is that the reliance on extra colours to prove a double-exponential lower bound may be a technical limitation of the stepping-up construction. This is challenged by a stunning discovery of Conlon, Fox and Rödl [5] who exhibited an infinite family of 3-uniform hypergraphs called *hedgehogs*, whose Ramsey numbers display a strong dependence on the number of colours. Namely, they showed that the 2-colour Ramsey number of hedgehogs is polynomial in their order, whereas the 4-colour Ramsey number is at least exponential. To understand just how significant a role the number of colours could play they asked the following:

Question 1.1. For any integer $h \geq 3$, do there exist integers k and q and a family of k-uniform hypergraphs for which the 2-colour Ramsey number grows as a polynomial in the number of vertices, while the q-colour Ramsey number grows as a tower of height h?

Our main contribution is to answer this in the affirmative. Define the *k-uniform balanced hedgehog* $\hat{H}_t^{(k)}$ with body of order t to be the graph constructed as follows: take a set S of t vertices, called the *body*, and for each subset $X \subset S$ of order $\lceil \frac{k}{2} \rceil$, add a k-edge e with $e \cap S = X$ such that for all $e, f \in E(\hat{H}_t^{(k)})$, we have $e \cap f \subset S$. The hedgehog $H_t^{(k)}$ as defined by Conlon, Fox and Rödl differs only in that they consider every $X \subset S$ of order $k-1$ rather than $\lceil \frac{k}{2} \rceil$. We observe that for $k = 3$, the two definitions coincide. When the uniformity is clear from the context, we shall drop the superscript.

Theorem 1.2. *There exist $c > 0$ and $q : \mathbb{N} \to \mathbb{N}$ such that for all $k \in \mathbb{N}$ and sufficiently large t, we have*

(a) $r_{2k+1}(\hat{H}_t) \leq t^{k+3}$, *and*
(b) $r_{2k+1}(\hat{H}_t; q(k)) \geq T_{\lfloor c \log_2 \log_2 k \rfloor}(t)$.

To prove this, we provide new stepping-up lemmas for a more general type of hypergraph Ramsey numbers. Let $r_k(G;q,p)$ for $q \geq p$ be the smallest integer n such that in every q-colouring of $K_n^{(k)}$, there is a copy of the hypergraph G whose edges span fewer than p colours. As before, we use $r_k(t;q,p)$ when $G = K_t^{(k)}$ and suppress p when $p = 2$.

A standard application of the first moment method (see, for example, [1]) shows that for any $k, q \in \mathbb{N}$, there exists $c > 0$ such that $r_k(t;q,q) \geq 2^{ct^{k-1}}$ for all $t \in \mathbb{N}$. We note that in the graph case ($k = 2$), the special case of $q = p$ was already investigated by Erdős and Szemerédi [10] in the 70s; in fact, the more general case when $p < q$ is also indirectly discussed. They showed the following rather precise bounds: for all $q \ll t$, $2^{\Omega(t/q)} \leq r_2(t;q,q) \leq q^{O(t/q)}$.

https://doi.org/10.1017/fms.2023.89 Published online by Cambridge University Press

These generalized hypergraph Ramsey numbers were also considered in a special case by Conlon, Fox and Rödl [5] who asked if there exist an integer q and number $c > 0$ such that $r_3(t; q, 3) \geq 2^{2^{ct}}$. To date, the only nontrivial improvement on the first moment bound has been made by Mubayi and Suk [16], who proved there exists $c > 0$ such that for $q \geq 9$, we have $r_3(t; q, 3) \geq 2^{t^{2+cq}}$ for $t \in \mathbb{N}$ sufficiently large; for all other values of $k, q, p \geq 3$, the random construction is essentially the state of the art. Our knowledge (or lack thereof) is thus summarised by the following bounds for $k, q, p \geq 3$ and sufficiently large $t \in \mathbb{N}$:

$$2^{t^c} \leq r_k(t; q, p) \leq T_k(O(t)),$$

where $c \geq 1$ is allowed to depend on k, q and p. Note that in this case, our upper bounds are a staggering tower of height $k - 2$ in the lower bounds.

A related notion called the set-colouring Ramsey number was introduced by Erdős, Hajnal and Rado in [9] and subsequently studied in [19] and much more recently in [4] and [2]. Borrowing notation from [4], let $R_k(t; q, s)$ denote the minimum number of vertices such that every (q, s)-*set colouring* of $K_n^{(k)}$ – that is, a colouring in which each k-set is assigned an element of $\binom{[q]}{s}$ – contains a monochromatic $K_t^{(k)}$. Here, monochromatic means the intersection of the colour sets assigned to the edges is nonempty. Observe that certain cases of R_k and r_k coincide. For example, $R_k(t; q, q-1) = r_k(t; q, q)$ and in general, we have the bound

$$r_k(t; q, p) \leq R_k\left(t; \binom{q}{p-1}, \binom{q-1}{p-2}\right).$$

We prove lower bounds on $r_k(t; q, p)$, thus giving lower bounds on certain set-colouring Ramsey numbers. However, we are not able to definitively resolve any questions from [4] due to central gaps in our understanding of hypergraph Ramsey numbers. See Section 5 for more on this.

Our main tool in the proof of Theorem 1.2 is the development of two new stepping-up constructions which yield the first tower-type results of their kind. We show the following three stepping-up statements, listed in order of decreasing strength, with $C_k = \frac{1}{k+1}\binom{2k}{k}$ denoting the k-th Catalan number.

Theorem 1.3. *Let $k, q, p \geq 3$. There exist $c \geq 1$ and t_0 such that for all $t > t_0$,*

(a) *if $p \leq C_k - 2$, then $r_{k+1}(t^c; q, p) > 2^{r_k(t;q,p)-1}$,*
(b) *if $p \leq C_k$, then $r_{k+1}(t^c; 2q + p, p) > 2^{r_k(t;q,p)-1}$ and*
(c) *if $p \leq k!$, then $r_{2k}(t^c; qp, p) > 2^{r_k(t;q,p)-1}$.*

Note that the growth rate in k which is implied by Part (c) (approximately a tower of height $\log_2 k$) of Theorem 1.3 is much smaller than that of Parts (a) and (b) because we can only step up at the cost of doubling the uniformity size. Unfortunately, this does not allow us to answer the question of Conlon, Fox and Rödl on $r_3(t; q, 3)$ since $C_2 = 2$, but already for $k \geq 4$, we have the following two corollaries:

Corollary 1.4. *For all $k \geq 4$, there is $q \in \mathbb{N}$ and $c > 0$ such that $r_k(t; q, 5) \geq T_{k-1}(t^c)$.*

Corollary 1.5. *For all $k \geq 4$, there is $c > 0$ such that $r_k(t; 3, 3) \geq T_{k-1}(t^c)$.*

Observe that by the second corollary, the growth rate of $r_k(t; 3, 3)$ matches the current best lower bounds for $r_k(t)$ up to a polynomial in t. The reason we have an absolute constant c in the exponent is due to the use of an Erdős-Hajnal type result on sequences (see Section 2).

The second main element of our proof connects the problem of avoiding monochromatic balanced hedgehogs to that of avoiding cliques that span few colours. It is a straightforward adaptation of ideas from Conlon, Fox and Rödl [5].

Lemma 1.6. *Given $k, q, t \in \mathbb{N}$, let $p = \binom{2k+1}{k+1}$ and $q' = \binom{q}{p}$. Then*

$$r_{2k+1}(\hat{H}_t; q', 2) > r_{k+1}\left(t; q, p+1\right) - 1.$$

https://doi.org/10.1017/fms.2023.89 Published online by Cambridge University Press

Using this result along with Part (c) of Theorem 1.3 yields the lower bound in Theorem 1.2(b). It is natural to ask whether one can combine the growth rate in k given by Part (a) of Theorem 1.3 with the ability to impose as many colours as in Part (c). Unfortunately, the condition $p \leq C_k$ prevents us from using Part (a) as the right-hand side because $C_k = \frac{1}{k+1}\binom{2k}{k} < \binom{2k+1}{k+1}$. This is tantalisingly close, if not a little curious, as the dependence on C_k comes from our exact solution to a subsequence avoidance problem (Corollary 2.4). We show that C_k presents a natural barrier in this endeavour. This barrier is made concrete by some new and tight results on the Ramsey theory of sequences, including an Erdős-Hajnal-type result in Section 2.

1.1. Outline

The outline of the paper is as follows. In Section 2, we introduce and prove results on the Ramsey theory of sequences. In Section 3, we use these to prove our stepping up constructions – namely, Theorem 1.3, and in Section 4, we prove our hedgehog-related results, including Theorem 1.2 and Lemma 1.6, and provide a construction of a degenerate hypergraph that relates to the Burr-Erdős Conjecture. To finish, we pose some questions and problems highlighted by our results.

2. Ramsey theory of sequences

Let $V = \{0,1\}^m$ and given vectors $v, w \in V$, define

$$\delta(v,w) = \max\{i : v_i \neq w_i\}.$$

We say $v < w$ if $v_\delta < w_\delta$. For every set of vertices $v_1 < v_2 < \cdots < v_{k+1}$ in V, there is a corresponding sequence $(\delta_1, \delta_2, \ldots, \delta_k)$ given by $\delta_i = \delta(v_i, v_{i+1})$. In this section, we introduce some definitions that will be useful in order to analyse the structure of these δ-sequences.

We say a sequence $S = (a_1, a_2, \ldots, a_m)$ is *monotonic* if $a_1 \leq a_2 \leq \cdots \leq a_m$ or $a_1 \geq a_2 \geq \cdots \geq a_m$. Two sequences $(a_1, a_2, \ldots, a_t)$ and $(b_1, b_2, \ldots, b_t)$ have the same pattern if the relative ordering of every pair of elements is the same; that is, $a_i > a_j$, $a_i = a_j$, or $a_i < a_j$ if and only if $b_i > b_j$, $b_i = b_j$, or $b_i < b_j$ respectively for all $1 \leq i < j \leq t$. A *pattern* is then the equivalence class of sequences with respect to this relative ordering. Given sequences S and P, we say S *contains the pattern* P if we can find a subsequence of S which has the same pattern as P; otherwise, we say S *avoids* P. A pattern is a *permutation pattern* if there is a permutation in the equivalence class, and we often use this permutation as a representative for the pattern. In what follows, we will mainly be concerned with permutations that avoid the patterns 132 or 231.

2.1. Max-induced subsequences

With the aforementioned applications to hypergraph Ramsey theory in mind, we introduce 'max-induced' subsequences. We say $(a_{i_1}, a_{i_2}, \ldots, a_{i_t})$ with $i_1 < i_2 < \cdots < i_t$ is a *max-induced subsequence* of $(a_1, a_2, \ldots, a_m)$ if the maximum of $(a_{i_j}, a_{i_j+1}, \ldots, a_{i_{j+1}})$ is attained at a_{i_j} or $a_{i_{j+1}}$ (i.e., at the left or right extreme) for all $1 \leq j < t$. We say a sequence S contains a *max-induced pattern* P if there is a max-induced subsequence of S which has pattern P, and a family of sequences $\mathcal{F}$ has the *max-induced Erdős-Hajnal* property with exponent $c(\mathcal{F})$ if any sequence S that avoids every member of $\mathcal{F}$ as a max-induced subsequence has a monotonic max-induced subsequence of order $|S|^{c(\mathcal{F})}$, where $|S|$ denotes the length of the sequence. We are able to characterize the families with this property.

The following fact, whose short proof can be found in Section 3, relates the max-induced property to our original motivation:

Claim 2.1. *Suppose $v_1 < v_2 < \cdots < v_{\ell+1}$ are vectors in $\{0,1\}^m$ and let $\delta_i = \delta(v_i, v_{i+1})$. If $(\delta_{i_1}, \ldots, \delta_{i_k})$ is a max-induced subsequence of $(\delta_1, \ldots, \delta_\ell)$, then there are $v_{j_1}, \ldots, v_{j_{k+1}}$ such that $\delta(v_{j_s}, v_{j_{s+1}}) = \delta_{i_s}$ for each $s \in [k]$.*

https://doi.org/10.1017/fms.2023.89 Published online by Cambridge University Press

Define an *interval* in a sequence to be a subsequence of consecutive elements. The following theorem is our core result on max-induced patterns:

Theorem 2.2. *A finite family of patterns $\mathcal{F}$ has the max-induced Erdős-Hajnal property if and only if $\mathcal{F}$ contains a* 132*-avoiding permutation pattern and a* 231*-avoiding permutation pattern.*

Observe that if a single permutation is $\{132, 231\}$-avoiding, it must be decreasing up to some element and increasing for the rest of the permutation. We say such a permutation has a *unique local minimum*, and we immediately get the following corollary of Theorem 2.2:

Corollary 2.3. *A permutation P has the max-induced Erdős-Hajnal property (with exponent $c(P) = 16^{-|P|}$) if and only if it has a unique local minimum (in other words, it is $\{132, 231\}$-avoiding).*

An old result of Shelah [17] states that any graph which does not have large cliques or large independent sets must contain exponentially many non-isomorphic induced subgraphs. Taking inspiration from this, we may ask: given $k \in \mathbb{N}$, what is the largest integer $f(k)$ for which there is $\varepsilon > 0$ such that every sequence of length n either contains at least $f(k)$ distinct max-induced patterns on k elements or a max-induced monotonic subsequence of length n^{ε}? It is well-known that the number of 132-avoiding permutations on $[k]$ is exactly the Catalan number, C_k (see, for example, [18, Section 1.5]). Thus, Theorem 2.2 allows us to answer the above question exactly and is the origin of the Catalan number that appears in Theorem 1.3.

Corollary 2.4. *The value of $f(k)$ is the number of* 132*-avoiding permutations (equivalently* 231*-avoiding), which is equal to the Catalan number C_k.*

To prove Theorem 2.2, we require the following lemma:

Lemma 2.5. *For every $k \geq 1$, there is a* 231*-avoiding permutation S_k of length $2^{k+1} - 1$ that does not contain a max-induced monotonic subsequence of length greater than $k + 1$.*

Proof. We proceed by induction on k. Note that for $k = 1$, the permutation $(1, 3, 2)$ works. Let $S_{k-1} = (a_1, a_2, \ldots, a_{m_{k-1}})$ be a 231-avoiding permutation of length $2^k - 1$ with no max-induced monotonic subsequence of length greater than k. Define the sequence $S_k = (b_1, \ldots, b_{2^{k+1}-1})$ by

$$b_i = \begin{cases} a_i & \text{if } i < 2^k; \\ 2^{k+1} - 1 & \text{if } i = 2^k; \\ a_{i-2^k} + 2^k - 1 & \text{if } i > 2^k. \end{cases}$$

If $(b_{i_1}, \ldots, b_{i_\ell})$ is a max-induced monotonic subsequence of S_k, then either $i_\ell \leq 2^k$ or $i_1 \geq 2^k$. Thus, by our choice of S_{k-1}, there is no max-induced monotonic component of length greater than $k + 1$ in S_k. Similarly, any possible 231 pattern $(b_{i_1}, b_{i_2}, b_{i_3})$ has either $i_3 < 2^k$ or $i_1 > 2^k$ by construction, so the inductive hypothesis shows S_k is 231-avoiding. □

Proof of Theorem 2.2. For the forward direction, suppose (without loss of generality) that $\mathcal{F}$ contains no 231-avoiding permutation pattern. Then each sequence in $\mathcal{F}$ contains either the pattern 231 or two elements which are equal. The sequence S_k defined in Lemma 2.5 is an $\mathcal{F}$-avoiding sequence of length at least 2^k containing no max-induced monotonic subsequence of length greater than $k + 1$.

For the backwards direction, let L and R be two permutations such that L is 132-avoiding and R is 231-avoiding, and suppose S is a sequence of length n. We prove by induction on $t := |L| + |R|$ that S contains at least one of L, R or a max-induced monotonic sequence of length at least $n^{\varepsilon}/2$, where $\varepsilon = 4^{-t}$. If $t \leq 5$, then one of L or R has at most two elements, so the statement is clear. We therefore assume $t > 5$, in which case we may assume also that n is large enough for the inequalities below.

We first describe an algorithm that shows either S contains a max-induced monotonic sequence of length at least $n^{\varepsilon}/2$ or there is an index i such that a_i is the maximum in a substantial interval (i.e., $a_j \leq a_i$ for $i - n^{1-\varepsilon} \leq j \leq i + n^{1-\varepsilon}$).

https://doi.org/10.1017/fms.2023.89 Published online by Cambridge University Press

We will attempt to build an increasing sequence from the leftmost element and a decreasing sequence starting from the rightmost element. Initialize index sets $I = [n]$ and $F_{\text{left}} = F_{\text{right}} = \varnothing$. In the k^{th} iteration of the algorithm, suppose the index $j_k \in I$ is such that $a_{j_k} = \max_{i \in I} a_i$. Let $\ell = \max F_{\text{left}}$ and $r = \min F_{\text{right}}$. For $k = 1$, we set $\ell = 0$ and $r = n + 1$. Then

- if $j_k - \ell < n^{1-\varepsilon}$, add j_k to F_{left}, delete all $i \in I$ such that $i \le j_k$, and proceed to the next iteration;
- if $r - j_k < n^{1-\varepsilon}$, add j_k to F_{right}, delete all $i \in I$ such that $i \ge j_k$, and proceed to the next iteration;
- otherwise, the algorithm terminates.

Observe that $(a_i)_{i \in F_{\text{left}}}$ and $(a_i)_{i \in F_{\text{right}}}$ are max-induced monotonic sequences in S, so if either $|F_{\text{left}}| \ge n^{\varepsilon}/2$ or $|F_{\text{right}}| \ge n^{\varepsilon}/2$, then we have constructed a max-induced monotonic subsequence of length $n^{\varepsilon}/2$. This must be the case if the algorithm terminates due to deleting all of I.

Otherwise, suppose the algorithm terminates on the k^{th} iteration for some fixed k. Then j_k satisfies $a_i \le a_{j_k}$ for $j_k - n^{1-\varepsilon} \le i \le j_k + n^{1-\varepsilon}$. Furthermore, for $F_{\max} := \{i \in I : a_i = a_{j_k}\}$, the constant sequence $(a_i)_{i \in F_{\max}}$ is a max-induced monotonic sequence in S, so we may assume $|F_{\max}| < n^{\varepsilon}/2$.

Our goal now is to show that S must contain a max-induced copy of either L, R or a long monotonic sequence. Very roughly, we seek a large 'blow-up' of either 132 or 231, where a blow-up of 231 is a subsequence on $I \cup J \cup K$ such that (a_i, a_j, a_k) is a max-induced 231 pattern for any $i \in I$, $j \in J$, $k \in K$. Supposing we have such a blow-up of 231 with I and K large, we may apply the inductive hypothesis in I with the permutation patterns L_{left} and R and on K with the permutation patterns L_{right} and R, where L_{left} and L_{right} are the permutation patterns to the left and right of the maximum element in any representative of the pattern L. If we find a max-induced R or long monotonic sequence, we are done, so we may assume I contains the pattern L_{left} and K the pattern L_{right}, which can be joined together using the 231 blow-up structure to find the pattern L in $I \cup J \cup K$. We are not able to find a true blow-up of either 132 or 231, so first some care is needed to find a suitable replacement.

Let $I_{\text{left}} = \{i \in I : i < j_k\}$ and $I_{\text{right}} = \{i \in I : i > j_k\}$. Let M ($\supset F_{\max}$) be the indices of the largest $n^{(1-\varepsilon)/2}$ elements in $(a_i)_{i \in I}$, breaking ties arbitrarily. Either $M_{\text{left}} := M \cap I_{\text{left}}$ or $M_{\text{right}} := M \setminus M_{\text{left}}$ has at least $\frac{1}{2} n^{(1-\varepsilon)/2}$ elements. Suppose the former (the argument for the latter is similar).

- Since $|I_{\text{right}}| \ge n^{1-\varepsilon}$, there must be an interval $I' \subset I_{\text{right}}$ of length at least $\frac{1}{2} n^{(1-\varepsilon)/2}$ such that $I' \cap M_{\text{right}} = \varnothing$.
- Since $|M_{\text{left}}| \ge \frac{1}{2} n^{(1-\varepsilon)/2}$ and $|F_{\max}| < n^{\varepsilon}/2$, there is an interval $J \subset I_{\text{left}}$ such that $|J \cap M_{\text{left}}| \ge n^{(1-3\varepsilon)/2}$ and $J \cap F_{\max} = \varnothing$.

We are not quite done: observe that there may be an element indexed by I' equal to an element indexed by J. To break ties, we restrict I' further. As before, we define $I'_{\max} := \{i \in I' : a_i = \max_{j \in I'} a_j\}$. The constant sequence $(a_i)_{i \in I'_{\max}}$ is a max-induced monotonic sequence in S, so we may assume $|I'_{\max}| < n^{\varepsilon}/2$. Then,

- There is an interval $I'' \subseteq I'$ of length at least $\frac{1}{2} n^{(1-3\varepsilon)/2}$ such that $|I'' \cap I'_{\max}| = \varnothing$.

Now, let $A = J \cap M_{\text{left}}$ and $B = I''$. By construction, any max-induced subsequence in A or B is a max-induced subsequence of S. Also if $i \in A$, $j \in B$ then $a_j < a_i < a_{j_k}$. Recall that the permutation L is 132-avoiding. Define L_{left} to be the subpermutation preceding $\max L$ and L_{right} to be the subpermutation following $\max L$. Observe that L_{left} and L_{right} are both 132-avoiding.

We finally apply the inductive hypothesis to $S_A := (a_i)_{i \in A}$ with the permutations L_{left} and R, and to $S_B := (a_i)_{i \in B}$ with the permutations L_{right} and R. If we find a max-induced copy of R in S_A or S_B, we are done. If S_A contains a max-induced monotonic subsequence of length at least

$$\frac{1}{2}|A|^{4^{-(|L_{\text{left}}|+|R|)}} \ge \frac{1}{2}(n^{(1-3\varepsilon)/2})^{4^{-(t-1)}} = \frac{1}{2}(n^{(1-3\varepsilon)/2})^{4\varepsilon} \ge \frac{n^{\varepsilon}}{2}$$

or S_B contains a max-induced monotonic subsequence of length at least

$$\frac{1}{2}|B|^{4^{-(|L_r|+|R|)}} \ge \frac{1}{2}(n^{(1-3\varepsilon)/2})^{4^{-(t-1)}} = \frac{1}{2}(n^{(1-3\varepsilon)/2})^{4\varepsilon} \ge \frac{n^{\varepsilon}}{2},$$

https://doi.org/10.1017/fms.2023.89 Published online by Cambridge University Press

we are finished. Otherwise, there is a max-induced copy of L_{left} in S_A and L_r in S_B, which along with a_{j_k} forms a max-induced copy of L in S. □

Given $n, k \in \mathbb{N}$, recall $f(k)$ is the largest integer for which there exists some $\varepsilon > 0$ such that every sequence of length n contains either at least $f(k)$ distinct max-induced patterns on k elements or a monotonic subsequence of length n^{ε}.

2.2. Separated subsequences

Up to this point, our sequence results do not use a useful fact about the δ-sequences which features often in stepping-up. A sequence $(a_1, a_2, \dots, a_k)$ has the *unique maximum property* if for any interval I, $|\{i \in I : a_i = \max(a_i)_{i \in I}\}| = 1$. We capitalize on this useful property in Lemma 2.6, which studies a second notion of subsequence that is used in our proof of Theorem 1.3(c). Call $(a_{i_1}, a_{i_2}, \dots, a_{i_t})$ with $0 < i_1 < \cdots < i_t \le m$ a *separated subsequence* of $(a_1, a_2, \dots, a_m)$ if $i_{j+1} > i_j + 1$ for all $j \in [t]$. Similar to max-inducedness, the following simple fact captures the usefulness of this definition for stepping up: suppose $v_1 < v_2 < \cdots < v_{\ell+1}$ are vectors in $\{0,1\}^m$ and let $\delta_i = \delta(v_i, v_{i+1})$. If $(\delta_{i_1}, \dots, \delta_{i_k})$ is a separated subsequence of $(\delta_1, \dots, \delta_\ell)$, then there are $v_{j_1}, \dots, v_{j_{2k}}$ such that $\delta(v_{j_{2s-1}}, v_{j_{2s}}) = \delta_{i_s}$ for each $s \in [k]$. Given a sequence $S = (s_1, \dots, s_n)$, define $\|S\| := |\{s_1, \dots, s_n\}|$ – that is, the number of distinct values in the sequence. We can now state our key result on separated subsequences.

Lemma 2.6. *Let $A = (a_i)_{i=1}^n$ be a sequence with the unique maximum property and $k \in \mathbb{N}$. If $\|A\| < n^{1/(k+1)}$ and n is large enough, A contains every permutation pattern on $[k]$ as a separated subsequence.*

The proof relies on a simple density argument to find a very rich substructure that contains within it all possible small structures.

Proof of Lemma 2.6. We first find constant subsequences $A_1, \dots, A_k$ such that

(i) $|A_i| \ge n^{1-i/(k+1)}$ for each $i \in [k]$,
(ii) if $a \in A_i$ (meaning a is equal to some element of A_i) and $a' \in A_j$ with $i < j$, then $a < a'$, and
(iii) (interlacing property) if $a_j, a_\ell \in A_i$ such that $j < \ell$, then there is some $a_m \in A_{i-1}$ such that $j < m < \ell$.

Given such subsequences, embedding an arbitrary permutation $(\sigma(1), \dots, \sigma(k))$ is straightforward: index the elements of A_k by $(a_{\ell_1}, \dots, a_{\ell_s})$. For each $j \in [k]$, use the interlacing property (and the fact $s \ge n^{1/(k+1)}$) to find a subsequence $(a_{j_1}, \dots, a_{j_k})$ of A_j with $\ell_{m2^{k+1}} < j_m < \ell_{(m+1)2^{k+1}}$. The (separated) subsequence $(a_{\sigma(1)_1}, \dots, a_{\sigma(k)_k})$ has the same pattern as $(\sigma(1), \dots, \sigma(k))$ by property (ii).

We construct the desired subsequences by induction. Since $\|A\| < n^{\frac{1}{k+1}}$, there is a constant subsequence A_1 such that $|A_1| \ge n^{1-1/(k+1)}$. Note that by the unique maximum property, A_1 must be a separated subsequence. For $i \ge 1$, given a constant separated subsequence $A_i = (a_{i_j})_j$ satisfying the above properties, we construct A_{i+1} as follows:

Let $A_i' = (a_{i_j}')_j$ be the subsequence of A consisting of the maximum values between the elements of A_i, meaning if $a_{i_j}, a_{i_{j+1}} \in A_i$, then $a_{i_j}' = \max\{a_\ell : i_j \le \ell \le i_{j+1}\}$. By the unique maximum property of A, a_{i_j}' is well-defined and is strictly greater than a_{i_j} and $a_{i_{j+1}}$. Then $|A_i'| = |A_i| - 1 \ge n^{1-i/(k+1)} - 1$, and $\|A_i'\| \le n^{1/(k+1)} - 1$ from the assumption that $\|A\| \le n^{1/(k+1)}$. Thus, there is a constant subsequence of A_i' with length at least

$$\frac{|A_i'|}{\|A_i'\|} \ge n^{1-(i+1)/(k+1)},$$

and we define A_{i+1} as this subsequence. It is clear from the construction that properties (i), (ii) and (iii) are satisfied. □

We are now ready to employ these results in our stepping-up constructions.

https://doi.org/10.1017/fms.2023.89 Published online by Cambridge University Press

3. Stepping up with many colours

In this section, we explain how the results on max-induced and separated patterns in sequences allow us to step up colourings in which all large sets of vertices span many colours. We prove Theorem 1.3 as well as Corollary 1.4. Throughout, all colourings discussed will be (hyper)edge-colourings. We say a colouring of a complete hypergraph with q colours is $(t;q,p)$-*rainbow* if every set of t vertices spans at least p colours. In this language, $r_k(t;q,p)-1$ is the largest integer n for which there exists a $(t;q,p)$-rainbow colouring of $K_n^{(k)}$. Let $V(K_n^{(r+1)})=\{0,1\}^n$ and given $v,w\in V$, recall that

$$\delta(v,w):=\max\{i : v_i\neq w_i\},$$

and that $v<w$ if and only if $v_\delta<w_\delta$. To each set of vertices $v_1<v_2<\cdots<v_{k+1}$ in $K_n^{(r+1)}$ corresponds a sequence $(\delta_1,\delta_2,\dots,\delta_k)$ given by $\delta_i=\delta(v_i,v_{i+1})$, which we refer to as the *corresponding δ-sequence*. The crucial properties for our application of the results of Section 2 to hypergraph Ramsey theory are the following:

$$\text{For an interval } I, \text{ let } \delta_I=\max(\delta_i)_{i\in I}. \text{ Then } |\{i\in I:\delta_i=\delta_I\}|=1. \tag{3.1}$$

$$\text{For any } v_1<v_2<\cdots<v_{k+1},\ \delta(v_1,v_{k+1})=\max\{\delta_1,\delta_2,\dots,\delta_k\}. \tag{3.2}$$

For examples, see [12] for proofs. A useful perspective is gained by viewing $[2^n]$ as the leaves of a binary tree of depth n, ordered from left to right as you would draw them. Given two leaves $v,w\in[2^n]$, the number $2\delta(v,w)$ is simply the length of the unique path between them.

Recall that Corollary 2.4 shows that a sequence $(\delta_1,\delta_2,\dots,\delta_\ell)$ either contains many distinct max-induced patterns or a long monotonic max-induced subsequence. In order to make use of this fact, we recall Claim 2.1, which simply says the following: suppose $v_1<v_2<\cdots<v_{\ell+1}$ are vectors in $\{0,1\}^m$ and let $\delta_i=\delta(v_i,v_{i+1})$. If $(\delta_{i_1},\dots,\delta_{i_k})$ is a max-induced subsequence of $(\delta_1,\dots,\delta_\ell)$, then there are $v_{j_1},\dots,v_{j_{k+1}}$ such that $\delta(v_{j_s},v_{j_{s+1}})=\delta_{i_s}$.

Proof of Claim 2.1. Let $v_{j_1}=v_{i_1}$ and $v_{j_{k+1}}=v_{i_k+1}$. Then for $s>1$, define $v_{j_s}=v_{i_{s-1}+1}$ if $\delta_{i_{s-1}}<\delta_{i_s}$ and $v_{j_s}=v_{i_s}$ otherwise. It is straightforward to check, using property (3.2), that $\delta(v_{j_s},v_{j_{s+1}})=\delta_{i_s}$ for each s. □

We are ready to use Theorem 2.2 to prove the first two parts of Theorem 1.3.

Proof of Theorem 1.3(a) and (b). Let $k,p,q\geq 3$ be such that $p\leq C_k$, the k^{th} Catalan number, and let $c:=16^k+1$. We first prove Part (b) – the inequality $r_{k+1}(t^c;2q+p,p)>2^{r_k(t;q,p)-1}$ – by showing such that we can construct a $(t^c;2q+p,p)$-rainbow colouring χ' of $K_{2^n}^{(k+1)}$ from any $(t;q,p)$-rainbow colouring χ of $K_n^{(k)}$. Let χ be such a rainbow colouring.

We identify $V(K_{2^n}^{(k+1)})$ and $V(K_n^{(k)})$ with $\{0,1\}^n$ and $[n]$, respectively. Partition the set of patterns on $[k]$, excluding the two monotonic permutations, into $p-2$ classes $\mathcal{C}_1,\dots,\mathcal{C}_{p-2}$ such that each class contains at least one 132-avoiding permutation pattern and one 231-avoiding permutation pattern. Let $\mathcal{C}_{p-1}$ be the family of patterns containing only the strictly increasing permutation and $\mathcal{C}_p$ the family containing the strictly decreasing permutation. Observe that such a partition is guaranteed to exist by Corollary 2.4 and the assumption that $C_k\geq p$ (note that a permutation may be both 132-avoiding and 231-avoiding). For a sequence S, we say $S\in\mathcal{C}_i$ if the pattern of S is in $\mathcal{C}_i$. Assign to each class $\mathcal{C}_i$ a distinct colour c_i. Then, for an edge $e=\{v_1,\dots,v_{k+1}\}$ of $K_{2^n}^{(k+1)}$, with corresponding δ-sequence $(\delta_1,\dots,\delta_k)$, let

$$\chi'(e)=\begin{cases}\chi(\{\delta_1,\dots,\delta_k\})\times 1 & \text{if } \delta_1<\delta_2<\cdots<\delta_k;\\ \chi(\{\delta_1,\dots,\delta_k\})\times 2 & \text{if } \delta_1>\delta_2>\cdots>\delta_k;\\ c_i & \text{if } (\delta_1,\dots,\delta_k)\in\mathcal{C}_i,\ i\leq p-2.\end{cases}$$

https://doi.org/10.1017/fms.2023.89 Published online by Cambridge University Press

We claim χ' is the desired colouring. Consider a set of t^c vertices $\{v_1, \dots, v_{t^c}\} \subset V(K_{2^n}^{(k+1)})$ and its corresponding δ-sequence $(\delta_1, \dots, \delta_{t^c-1})$. Since each class $\mathcal{C}_i$ contains a 132-avoiding permutation pattern and a 231-avoiding permutation pattern, Theorem 2.2 tells us that $(\delta_1, \dots, \delta_{t^c-1})$ either contains a max-induced pattern from each of $\mathcal{C}_1, \dots, \mathcal{C}_p$ or a max-induced monotonic sequence of length t (the value of c was chosen for this purpose). Claim 2.1 shows that for every max-induced sequence $S \subset (\delta_1, \dots, \delta_{t^c-1})$ of length k, we can find $\{v_{i_1}, \dots, v_{i_{k+1}}\} \subset \{v_1, \dots, v_{t^c}\}$ whose corresponding δ-sequence is S. In other words, we have that $\chi'(\{v_{i_1}, \dots, v_{i_{k+1}}\}) = c_i$ if the pattern of S is in $\mathcal{C}_i$. Therefore, $(\delta_1, \dots, \delta_{t^c-1})$ containing a max-induced pattern from each of $\mathcal{C}_1, \dots, \mathcal{C}_p$ implies that $\{v_1, \dots, v_{t^c}\}$ spans at least p colours.

If instead, $(\delta_1, \dots, \delta_{t^c-1})$ contains a max-induced monotonic subsequence, say $(\delta_1, \dots, \delta_t)$, then by Claim 2.1, we can find $\{v_{i_1}, \dots, v_{i_{t+1}}\} \subset \{v_1, \dots, v_{t^c}\}$ whose corresponding pattern is $(\delta_1, \dots, \delta_t)$. We may assume without loss of generality that $(\delta_1, \dots, \delta_t)$ is increasing, and by property (3.1), $\delta_1 < \cdots < \delta_t$. Every subsequence of a strictly increasing max-induced sequence is a strictly increasing max-induced subsequence. Thus, again by Claim 2.1, for every subsequence $S \subset (\delta_1, \dots, \delta_t)$ of length k, we can find a subsequence $S' \subset (v_{i_1}, \dots, v_{i_{t+1}})$ of length $k+1$ whose corresponding δ-sequence is S. Observe that $\chi'(S') = (\chi(S), 1)$, so the set of $(k+1)$-edges on $\{v_{i_1}, \dots, v_{i_{t+1}}\}$ spans at least as many colours as the set of k-edges on $\{\delta_1, \dots, \delta_t\}$. By assumption, this is at least p.

All that remains is to count the number of colours used by χ'. This is at most 2 colours for each of the q colours used by χ plus an extra $p-2$ for the c_i's. Combined with the above arguments, this shows that χ' is a $(t^c; 2q+p-2, p)$-rainbow colouring and, in particular, a $(t^c; 2q+p, p)$-rainbow colouring.

We now describe the construction of the $(t^c; q, p-2)$-rainbow colouring, say χ'', which proves Part (a) (i.e., that $r_{k+1}(t^c; q, p-2) > 2^{r_k(t;q,p-2)-1}$). Let a $(t; q, p-2)$-rainbow colouring χ of $K_n^{(k)}$ be given. Partition the set of patterns on $[k]$ as before and assign to each class $\mathcal{C}_i$ a distinct colour c_i for $i \in [p-2]$; only this time identify these colours with the first $[p-2]$ colours used by χ (trivially we have $q \geq p-2$). For an edge $e = \{v_1, \dots, v_{k+1}\}$ with corresponding δ-sequence $(\delta_1, \dots, \delta_k)$, let

$$\chi''(e) = \begin{cases} \chi(\{\delta_1, \dots, \delta_k\}) & \text{if } \delta_1 < \delta_2 < \cdots < \delta_k \text{ or } \delta_1 > \delta_2 > \cdots > \delta_k; \\ c_i & \text{if } (\delta_1, \dots, \delta_k) \in \mathcal{C}_i. \end{cases}$$

We omit the proof that χ'' is in fact the desired colouring, as it is almost identical to the proof for χ'. We simply note that now we can only force $p-2$ colours in any set of t^c vertices, as an edge with a monotonic pattern may receive any one of the colours c_i for $i \in [p-2]$. □

In applying our stepping-up results, we will need the following result on rainbow colorings which is proven by a standard use of the first moment method (see, for example, [1]).

Proposition 3.1. *For every k and q, there exists $\varepsilon > 0$ and t_0 so that for $t > t_0$, there is a $(t; q, q)$-rainbow colouring of $K_n^{(k)}$ with $n = 2^{\varepsilon t^{k-1}}$.*

Proof. Set $\varepsilon = 1/(qk!)$ and $t_0 = eq$. Consider a uniformly random q-colouring of $K_n^{(k)}$ with $n = 2^{\varepsilon t^{k-1}}$. The probability a given t-set contains fewer than q colours is at most $q(1-1/q)^{\binom{t}{k}}$. Therefore, the expected number of t sets with fewer than q colours is at most

$$\binom{n}{t} q(1-1/q)^{\binom{t}{k}} < \left(\frac{en}{t}\right)^t q e^{-\binom{t}{k}/q} < \left(\frac{q^{1/t} en}{t e^{-\binom{t}{k}/(tq)}}\right)^t < 1.$$

As the expectation is strictly less than one, there must exist some $(t; q, q)$-rainbow colouring. □

Proof of Corollary 1.4. Proposition 3.1 shows that there is an $\varepsilon > 0$ such that $r_3(t; q, q) > 2^{\varepsilon t}$. Since the Catalan number $C_i \geq 5$ for $i \geq 3$, we may apply the inequality $r_{i+1}(t^c; 2q+5, 5) > 2^{r_i(t;q,5)-1}$ from Theorem 1.3, once for each $i \in \{3, \dots, k-1\}$ and with a different $c \geq 1$ each round. This ultimately yields a lower bound of $r_k(t^{c'}; q', 5) \geq T_{k-1}(t)$, where $q' = 3^k q$ and $c' \geq 1$. □

https://doi.org/10.1017/fms.2023.89 Published online by Cambridge University Press

The proof of Corollary 1.5 follows the exact same procedure as the proof of Corollary 1.4, simply using the bound $r_{k+1}(t^c;3,3) > 2^{r_k(t;3,3)-1}$ in place of the bound $r_{k+1}(t^c;2q+5,3) > 2^{r_k(t;q,5)-1}$. We thus omit it. We now prove the final part of Theorem 1.3, which combines Lemma 2.6 with a slightly different stepping-up technique.

Proof of Theorem 1.3(c). Let $k, p, q \geq 3$ be such that $p \leq k!$ and let $c := k+2$. We prove the inequality $r_{2k}(t^c; pq, p) > 2^{r_k(t;q,p)-1}$ by showing such that we can construct a $(t^c; pq, p)$-rainbow colouring χ' of $K_{2^n}^{(2k)}$ from any $(t;q,p)$-rainbow colouring χ of $K_n^{(k)}$. Let such a χ be given.

We identify $V(K_{2^n}^{(2k)})$ and $V(K_n^{(k)})$ with $\{0,1\}^n$ and $[n]$, respectively. Order the $k!$ permutations of $[k]$ and let P_i be the set of sequences whose pattern is the i^{th} permutation for $i \in [p]$. Then, for an edge $e = \{v_1, \dots, v_{2k}\}$ with corresponding δ-sequence $(\delta_1, \dots, \delta_{2k-1})$, we set

$$\chi'(e) = \chi(\{\delta_1, \delta_3, \dots, \delta_{2k-1}\}) \times i$$

if $(\delta_1, \delta_3, \dots, \delta_{2k-1}) \in P_i$ for some $i \in [p]$ and we let $\chi'(e)$ be arbitrary otherwise.

We now check that this is the desired colouring. It is clear that it uses at most pq colours. Suppose $v_1 < \cdots < v_{t^c}$ are vertices of $K_{2^n}^{(2k)}$ and let $(\delta_1, \dots, \delta_{t^c-1})$ be the corresponding δ-sequence. If $\|(\delta_i)_{i=1}^{t^c-1}\| \geq 2t$, then we can find an index set $I \subset [t^c - 1]$ such that $(\delta_i)_{i\in I}$ is a separated subsequence with $\|(\delta_i)_{i\in I}\| \geq t$. Let $f = \{\delta_{i_1}, \dots, \delta_{i_k}\} \subset (\delta_i)_{i\in I}$, where $i_1 < i_2 < \cdots < i_k$. Then the edge $e = \{v_{i_1}, v_{i_1+1}, v_{i_2}, v_{i_2+1}, \dots, v_{i_k}, v_{i_k+1}\}$ has colour $\chi'(e) = (\chi(f), \cdot)$. As f was arbitrary, we have that $\{v_1, \dots, v_{t^c}\}$ spans at least as many colours as $\{\delta_i\}_{i\in I}$. But $\|(\delta_i)_{i\in I}\| \geq t$ and so $(\delta_i)_{i\in I}$ spans at least p colours by assumption.

Now suppose that $\|(\delta_i)_{i=1}^{t^c-1}\| < 2t$. By (3.1), $(\delta_i)_{i=1}^{t^c-1}$ satisfies the unique maximum property. Therefore, by our choice of c and assumption that t is large, we can apply Lemma 2.6 with $A = (\delta_i)_{i=1}^{t^c-1}$ to conclude that A contains every permutation of $[k]$ as a separated subsequence. Suppose we have a separated subsequence $S \subset A$ whose pattern is the i^{th} permutation for some $i \in [p]$. As before, we can find a $(2k)$-edge $e \subseteq \{v_1, \dots, v_{t^c}\}$ for whom the corresponding δ-pattern is S. Thus, $\chi'(e) = (\cdot, i)$. Repeating this for the first p permutations implies that $\{v_1, \dots, v_{t^c}\}$ spans at least p colours. □

Both of our stepping-up constructions (like the original) rely on breaking the stepping-up into two parts: a *lifted* colouring from a lower uniformity graph and a *new* colouring. The constructions have the property that every large set of vertices satisfies one of the following: either it spans a colouring that is *lifted* from a large clique on the lower graph or it sees many distinct *new* colours. The number of new colours one can guarantee using this approach depends strongly on the uniformity of the hypergraphs in question. We ask a question in the concluding remarks concerning this.

4. Hedgehogs

In this section, we relate our results on many-coloured Ramsey numbers of complete hypergraphs to a Ramsey problem on a certain class of hypergraphs which we call *generalized hedgehogs*. We note our idea is very much inspired by the Conlon-Fox-Rödl construction described in the introduction. The main result is Theorem 1.2(a), and we also prove Lemma 1.6 for a more general family of hedgehogs.

We define the *generalized hedgehog* $\mathrm{H}_t^{(k)}(s)$ for $s, k, t \in \mathbb{N}$ with $k > s$ and $t \geq s$ to be the following k-uniform hypergraph: fix a set of t vertices called the *body*. The edge set of $\mathrm{H}_t^{(k)}(s)$ consists of the following edges: for every s-subset S of the body, we add a k-edge e containing S along with $k-s$ additional vertices that are contained in no other edge of $\mathrm{H}_t^{(k)}(s)$.

The k-uniform balanced hedgehog $\hat{H}_t^{(k)}$ is $\mathrm{H}_t^{(k)}(\lceil \frac{k}{2} \rceil)$, and the hedgehog introduced by Conlon, Fox and Rödl is $\mathrm{H}_t^{(k)}(k-1)$. We note that $\mathrm{H}_t^{(k)}(s)$ has $\binom{t}{s}$ edges and $t + (k-s)\binom{t}{s}$ vertices.

Lemma 1.6 is a corollary of the following result, which we state in its full generality.

https://doi.org/10.1017/fms.2023.89 Published online by Cambridge University Press

Lemma 4.1. *Given $k, q, p', t \in \mathbb{N}$, let $p = \binom{k}{s}$ and $q' = \binom{q}{p}$. Then*

$$r_k(H_t^{(k)}(s); q', p'+1) > r_s(t; q, p'p+1) - 1.$$

Proof of Lemma 1.6. Apply Lemma 4.1 with $2k+1, k+1$ and 1 playing the roles of k, s and p', respectively. □

Proof of Lemma 4.1. We prove $r_k(\mathrm{H}_t^{(k)}(s); q', p'+1) > r_s(t; q, p'p+1) - 1$ by showing that given a $(t; q, p'p+1)$-rainbow colouring of $K_n^{(s)}$, we can construct a q'-colouring χ' of $K_n^{(k)}$ in which every copy of $\mathrm{H}_t^{(k)}(s)$ spans at least $p'+1$ colours. Let such a χ be given. Identify the vertex sets of $K_n^{(k)}$ and $K_n^{(s)}$. We colour $e' \in E(K_n^{(k)})$ by $\chi'(e') = \{\chi(e) : e \in E(K_n^{(s)}[e'])\}$ – that is, the set of all colours of the s-edges contained in e'. The number of s-edges contained in e' is exactly $p = \binom{k}{s}$, so $\chi'(e')$ will be a set of at most that many colours. If $|\chi'(e')| < p$, then we add arbitrary colours to the set until $|\chi'(e')| = p$. Thus, the new colouring uses at most $q' = \binom{q}{p}$ colours.

Now consider a copy $H \subset K_n^{(k)}$ of $\mathrm{H}_t^{(k)}(s)$. As the body contains t vertices, by assumption, it spans at least $p'p+1$ colours under χ. Each of these colours appears as an element of $\chi'(e')$ for some $e' \in E(H)$ (i.e., $|\cup_{e' \in E(H)} \chi'(e')| \geq p'p+1$). However, p colours appear in $\chi'(e')$ for a given e', so by pigeonhole there must be more than p' edges $e' \in E(H)$ with distinct colours in $\chi'(H)$. □

We now derive Theorem 1.2(b).

Proof of Theorem 1.2(b). Throughout the proof, we assume $t \in \mathbb{N}$ is sufficiently large. Let $m = \lceil \frac{e(k+1)}{\log_2(k+1)} \rceil$, which is chosen so that $m! > (\frac{m}{\mathrm{e}})^m > 2^{2k} > \binom{2k+1}{k+1}$ for k sufficiently large. Recall that by Proposition 3.1, $r_m(t; m!, m!) \geq T_2(\varepsilon t)$ for some $\varepsilon > 0$. We now use the inequality $r_{2k}(t^c; pq, p) \geq 2^{r_k(t;q,p)}$ from Theorem 1.3 exactly $\lfloor \log_2 \frac{k+1}{m} \rfloor > \log_2 \log_2 k - 3$ times, to obtain that there exists q' and $c' \geq 1$ such that $r_{k+1}(t^{c'}; q', m!) \geq T_{\log_2 \log_2 k - 1}(t)$ (if $k+1 > 2^j m$ with $j = \lfloor \log_2 \frac{k+1}{m} \rfloor$ we use the simple observation $r_{i+1}(t; q, p) > r_i(t; q', p)$ for some $q' > q$). This is possible because the uniformity increases at each step while the number of colours imposed remains at $m!$. Since $m! > \binom{2k+1}{k+1}$, we can apply Lemma 1.6 to get that for some q'', $r_{2k+1}(\hat{H}_t; q'', 2) \geq T_{\log_2 \log_2 k - 2}(t)$, where we removed the exponent c' at the cost of a tower height. We choose $c > 0$ such that if $m! \leq \binom{2k+1}{k+1}$, then $c \log_2 \log_2 k \leq 1$.

Thus, for all k, $r_{2k+1}(\hat{H}_t; q'', 2) \geq T_{c \log_2 \log_2 k}(t)$. Furthermore, it is clear that q'' is only a function of k and of m, which is itself a function of k, as required. □

Conlon, Fox and Rödl showed the following using a similar argument to Theorem 1.2(a):

Theorem 4.2 [5]. *For all $k \geq 4$, there exists $c > 0$ such that*

$$r_k(H_t^{(k)}; 2, 2) \leq T_{k-2}(ct).$$

For the case $k = 4$, a construction of Kostochka and Rödl [13] shows this is approximately sharp (i.e., that $r_4(\mathrm{H}_t^{(4)}(3); 2, 2) = T_2(\Omega(t))$). We cannot prove a matching lower bound for $k = 5$, as when we attempt to apply Lemma 1.6, we need to impose 6 colours on 4-uniform graphs. When stepping up to uniformity 4, Theorem 1.3 allows us to impose at most $C_3 = 5$ colours, so we cannot beat the random argument here. We do, however, obtain the following:

Lemma 4.3. *For $5 \leq k \leq 13$, there exist $c > 0$ and $q \in \mathbb{N}$ such that for all t,*

$$r_k(H_t^{(k)}; q, 2) \geq T_{k-3}(t^c).$$

Proof. The proof mimics that of Theorem 1.2(b), but now we start from the fact that $r_4(t; 14, 14) \geq T_2(ct)$ (using Proposition 3.1) and leverage that $C_4 = 14$. Starting from $\ell = 4$, we apply the relation $r_{\ell+1}(t^C; 2q+p, p) \geq 2^{r_\ell(t;q,p)}$ from Theorem 1.3 $k-5$ times. This gives us $r_{k-1}(t; q', 14) \geq T_{k-3}(t^{c'})$

https://doi.org/10.1017/fms.2023.89 Published online by Cambridge University Press

for some $q' \in \mathbb{N}$ and $c' \geq 1$. Applying Lemma 1.6 then gives the result using q'' colours and with c' as c. □

We now prepare to prove our upper bounds on the Ramsey numbers of balanced hedgehogs, Theorem 1.2(a). Let H be an r-uniform hypergraph and $A \subset V(H)$ with $|A| < r$. The *piercing number* of A (denoted $\tau_H(A)$) is the size of the smallest set of vertices from $V(H)\backslash A$ that intersects every edge of H containing A. Equivalently, it is the minimum number of vertices that must be deleted from $V(H)\backslash A$ to delete all edges containing A.

Proposition 4.4. *Suppose H is an r-uniform hypergraph and $v \in V(H)$ has $\tau(v) \geq (r-1)m$. Then there exist m edges whose pairwise intersections are all precisely v.*

Proof of Proposition 4.4. Let $X \subset V(H)$ be a witness to $\tau(v)$. It is clear the set of edges incident to v has order at least $(r-1)m$. As each such edge e contains at most $r-1$ elements of X, and $|X|$ is minimal with respect to intersecting all edges incident to v, we can greedily find the desired m edges. □

Proof of Theorem 1.2(a). Let $G = K_n^{(2k+1)}$, where $n \geq t^{k+3}$, and let χ be an arbitrary red/blue colouring of the edges of G. We will show that there exists a monochromatic copy of $H_t^{(2k+1)}(k+1)$ in G, and thus that $r_{2k+1}(H_t^{(2k+1)}(k+1); 2, 2) \leq n$. We will do so by using χ to define a partial edge 2-colouring χ' of $G' := K_n^{(k+1)}$ and, in turn, using χ' to define a 2-colouring χ'' of $V(G)$. We will then find a large monochromatic set of vertices in χ'', use this to find a large (red or blue) independent set in G' and finally use this to find our monochromatic $H_t^{(2k+1)}(k+1)$ in G.

Throughout the proof, let $H_t := H_t^{(2k+1)}(k+1)$. We say a set of vertices is in *red danger* if its red piercing number (its piercing number in the subgraph of red edges) is less than t^{k+1} and similarly define *blue danger*. Then for $e \in G' := K_n^{(k+1)}$, let $\chi'(e)$ be red if e is in red danger and blue if e is in blue danger. As $\tau_G(e) > 2t^{k+1}$, we have that this partial colouring χ' is well-defined.

Now say a vertex v is in *red peril* if its red piercing number in G' (under χ') is at most $2kt^{k+1}$ and similarly define *blue peril*. Let $\chi''(v)$ be red if v is in red peril and blue if v is in blue peril, with ties broken arbitrarily. We claim that χ'' assigns a colour to every vertex. Indeed, suppose that some vertex v has red piercing number and blue piercing number at least $2kt^{k+1}$ under χ'.

By Proposition 4.4, there exist $(k+1)$-edges $e_1, \ldots, e_s, f_1, \ldots, f_s \in G'$, where $s = 2t^{k+1}$ such that

- $e_i \cap e_j = f_i \cap f_j = \{v\}$ for all $i \neq j$,
- $\chi'(e_i)$ is red for all $i \in [s]$ and
- $\chi'(f_j)$ is blue for all $j \in [s]$

Let $A_1, A_2, \ldots, A_s$ be disjoint subsets each of size $k-1$ in $V(G) \setminus (\bigcup_{i=1}^{s}(e_i \cup f_i))$. For each $i, j \in [s]$, define a $(2k+1)$-edge g_{ij} to be $e_i \cup f_j$ along with an arbitrary choice of $(2k+1) - |e_i \cup f_j|$ vertices from $A_{(i+j-1) \bmod s}$. Observe that $e_i \cup f_j \neq e_{i'} \cup f_{j'}$ for $(i, j) \neq (i', j')$, so in particular, these edges are all distinct.

At least half of these edges must have been, without loss of generality, red under the colouring χ of G. Therefore, there is some $i \in [s]$ such that e_i is contained in at least $\frac{s^2}{2s} = \frac{s}{2} > t^{k+1}$ red $(2k+1)$-edges under χ. But since $g_{ij} \cap g_{ij'} = e_i$ for $j \neq j'$, this contradicts the fact that e_i was a red-danger edge of G'. Thus, the colouring χ'' of $V(G)$ colours every vertex.

We now choose a set X of red vertices (without loss of generality) which has order $\frac{n}{2}$. By the definition of χ'', we have that the red piercing number under χ' of each $v \in X$ is at most $2kt^{k+1}$. Therefore, we can greedily find a set Y of order $\frac{n}{2 \cdot 2kt^{k+1}} \geq t$ which contains no red edges of G' (here we use that t is large relative to k).

All edges of G' in Y are not in red danger and thus have red piercing number at least t^{k+1} under χ. By the definition of piercing number and as $t^{k+1} > t + k\binom{t}{k+1} = |V(H_t^{(2k+1)}(k+1))|$, we can build the hedgehog greedily using Y as the body. □

https://doi.org/10.1017/fms.2023.89 Published online by Cambridge University Press

4.1. The Burr-Erdős Conjecture in hypergraphs

The *degeneracy* of a hypergraph is the minimum d such that every induced subgraph contains a vertex incident to at most d edges; such a hypergraph is called *d-degenerate*. Burr and Erdős conjectured that for every d, there exists a constant $c_d \geq 1$ such that every d-degenerate graph G on n vertices satisfies $r_2(G) < c_d n$ [3]. This was finally proven by Lee in [15], building on the previous work of several authors [(11),(14)]. In the case of hypergraphs, however, the conjecture fails: Kostochka and Rödl [13] showed $r_4(H_t^{(4)}) \geq 2^{ct}$ and were able to construct for every large enough d, d-degenerate 3-uniform hypergraphs on n vertices whose ramsey number is at least $n^{d^{1/4}}$. If one allows 3 or more colours, Conlon, Fox and Rödl [5] proved $r_3(H_t^{(3)}; 3) \geq \Omega(t^3/\log^6 t)$. Note that the degeneracy of hedgehogs is 1.

This shows that the Burr-Erdős Conjecture fails for k-uniform hypergraphs where $k \geq 4$ and for 3-uniform hypergraphs provided the number of colours is at least 3 or the degeneracy is large enough. For the sake of completeness, we give a simple construction of 8-degenerate 3-uniform hypergraphs whose Ramsey number is not linear in the number of vertices.

Proposition 4.5. *There exists a 3-uniform hypergraph on Cn^2 vertices which is 8-degenerate and for which the 2-colour Ramsey number is at least Cn^3.*

Proof. Let V be a set of n vertices, and with $m = \binom{n}{2}$, let B be a set of $m+1$ vertices disjoint from V with a total ordering $x_1, x_2, \ldots x_m, x_{m+1}$. Consider an ordering on the edges E of the complete graph on V, say $e_1, e_2, ..., e_m$. Let H be a 3-uniform hypergraph consisting of the following: for each edge $e_i = (x, y) \in E$, add the two triples $\{x, y, x_i\}$ and $\{x, y, x_{i+1}\}$, and for every $i \in [m-3]$, add all triples within $\{x_i, x_{i+1}, x_{i+2}, x_{i+3}, x_{i+4}\}$. By considering the smallest element of B in each subset of vertices, we see that H is 8-degenerate.

We claim $r_3(H) \geq n^3/8$. Indeed, given $K^{(3)}_{n^3/8}$, let $W_1, \ldots W_{n/4}$ be an arbitrary partition of the vertices where $|W_i| = n^2/2$ for all $i \in [n/4]$. We colour blue all edges inside any W_i and all edges with vertices in distinct sets $W_{i_1}, W_{i_2}, W_{i_3}$. We colour red all edges with exactly two vertices in a set W_i, for some $i \in [n/4]$.

Suppose there is a red copy of H. In particular, the edges induced by $S = \{x_1, x_2, x_3, x_4, x_5\}$ are all red. If S intersects distinct parts W_i, W_j and W_k for some $i, j, k \in [n/4]$, then the vertices in the intersections form a blue edge. However, if $S \subseteq W_i \cup W_j$, then $|S \cap W_i| \geq 3$ or $|S \cap W_j| \geq 3$, which again forms a blue edge. So we have a contradiction.

Suppose instead there is a blue copy of H. By pigeonhole, there are two vertices $v, w \in V$ which both lie in the same W_i. Let $e_j = \{v, w\}$. The edges $\{v, w, x_j\}$ and $\{v, w, x_{j+1}\}$ are both blue, so $x_j, x_{j+1} \in W_i$. But $\{x_j, x_{j+1}, x_{j+2}\}$ and $\{x_{j-1}, x_j, x_{j+1}\}$ are also blue, so this implies $x_{j-1}, x_{j+2} \in W_i$. Proceeding inductively, we conclude that all of B must lie in W_i.

However, $|W_i| = n^2/2$ and $|B| = \binom{n}{2}$, so there are $n/2$ vertices of V outside of W_i. Two of these vertices z and u must lie in the same W_j with $j \neq i$. Letting $e_k = \{z, u\}$, we get that the edge $\{z, u, x_k\}$ is coloured red, a contradiction which finishes the proof. □

5. Concluding remarks

We have proved that for every positive integer h, there exist q, k and an infinite family of k-uniform hypergraphs whose 2-colour Ramsey numbers differ by a tower of height h from the q-colour Ramsey numbers. This reinforces the fact that the number of colours plays an important role in the behaviour of Ramsey numbers of hypergraphs and casts a shadow on Erdős's conjecture on the 2-colour Ramsey number of a 3-uniform clique.

Observe that both of our new stepping-up constructions rely on a dichotomy: either we can find many suitable substructures within the δ-sequences (which give rise to many colours) or we must have a long monotonic sequence (which allows us to use induction). Since for every k-edge there are at most $k!$ distinct permutations, our methods fail to give good lower bounds for $r_k(t; q, p)$ whenever $k \ll p$. Even

https://doi.org/10.1017/fms.2023.89 Published online by Cambridge University Press

in the simplest case $r_3(t;q,3)$, we were not able to prove a double exponential lower bound, leaving open the following question of Conlon, Fox and Rödl on $r_3(t;q,3)$.

Problem 5.1. [5, Problem 1] Is there an integer q, a positive constant c and a q-colouring of the 3-uniform hypergraph on $2^{2^{ct}}$ vertices such that every subset of order t receives at least 3 colours?

We propose here a much weaker problem than Problem 5.1, which we were not able to resolve. We note that a negative answer would uncover a radical new phenomenon in the Ramsey numbers of hypergraphs.

Problem 5.2. Does there exist $k \in \mathbb{N}$ such that the following holds? For all $p \in \mathbb{N}$, there exist $q \in \mathbb{N}$ and $c > 0$ such that $r_k(t;q,p) \geq 2^{2^{t^c}}$ for all t sufficiently large.

A similar but much more ambitious problem was posed in [4].

Problem 5.3. [4, Problem 6.3] Determine the tower height of $R_k(n;r,r-1) = r_k(n;r,r)$ for all $k \geq 3$ and $r \geq 2$.

The authors of [4] note the apparent difficulty of Problem 5.3 and ask the following weaker question. Is there a fixed integer c such that $R_k(n;r,r-1) \geq T_{k-c}(n)$ for every $k \geq 3$ and $r \geq 2$? We cannot answer this question, but using Theorem 1.3(a), we can prove $R_k(n;r,r-1)$ is at least a tower of height roughly $k - 0.5\log_2 r$. Any improvement beyond this bound would likely be very interesting.

We make the following conjecture regarding the Ramsey numbers of k-uniform hedgehogs. This would, in particular, demonstrate that the 2-colour and q-colour Ramsey numbers of these hedgehogs, unlike those of balanced hedgehogs, do not differ by arbitrarily large tower heights.

Conjecture 5.4. *There is $\ell \in \mathbb{N}$ such that for every positive integer k, for every sufficiently large t,*

$$r_k(H_t^{(k)}) \geq T_{k-\ell}(t).$$

Finally, recall Proposition 4.5 shows there is an infinite family of 3-uniform hypergraphs which are 8-degenerate and for which the the 2-colour Ramsey numbers grow faster than linear in the order of the hypergraphs. It would be interesting to improve the quantitative aspects of this result.

Problem 5.5. Give an infinite family of 1-degenerate 3-uniform hypergraphs whose 2-colour Ramsey number is not polynomial in the order of the hypergraph.

Kostochka and Rödl [13] indicate that this may not be possible. We share this suspicion, noting that the following could be true: for any d, there is $f(d)$ such that the Ramsey number of any d-degenerate 3-uniform hypergraph (on n vertices) is at most $n^{f(d)}$.

Acknowledgements. We thank Bhargav Narayanan for helpful conversations on this topic and the anonymous referees for their detailed and invaluable feedback, especially regarding the organization of the first two sections.

Competing interest. The authors have no competing interest to declare.

Funding statement. The first author was supported in part by Simons Foundation grant 332622. The first and fourth authors were supported in part by NSF Grant CCF-1814409 and NSF Grant DMS-1800521. The second author was supported by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy EXC-2181/1 - 390900948 (the Heidelberg STRUCTURES Cluster of Excellence) and by EPSRC grant EP/V007327/1. The third author was supported in part by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 428212407.

References

[1] N. Alon and J. Spencer, *The Probabilistic Method* (John Wiley & Sons, 2016).
[2] L. Aragão, M. Collares, J. P. Marciano, T. Martins and R. Morris, 'A lower bound for set-colouring Ramsey numbers', Preprint, 2022, arXiv:2212.06802.
[3] S. Burr and P. Erdős, 'On the magnitude of generalized Ramsey numbers for graphs,' in *Infinite and Finite Sets I* vol. 10 (Colloq. Math. Soc. János Bolyai, North-Holland, Amsterdam-London, 1975).

https://doi.org/10.1017/fms.2023.89 Published online by Cambridge University Press

[4] D. Conlon, J. Fox, X. He, D. Mubayi, A. Suk and J. Verstraëte, 'Set-coloring Ramsey numbers via codes', Preprint, 2022, arXiv:2206.11371.
[5] D. Conlon, J. Fox and V. Rödl, 'Hedgehogs are not color blind', *J. Comb.* **8**(3) (2017), 475–485.
[6] D. Conlon, J. Fox and B. Sudakov, 'Hypergraph Ramsey numbers', *J. Amer. Math. Soc.* **23**(1) (2010), 247–266.
[7] P. Erdős and A. Hajnal, 'On Ramsey like theorems, problems and results', in *Combinatorics* (Proc. Conf. Combinatorial Math., Math. Inst., Oxford, 1972), 123–140.
[8] P. Erdős and R. Rado, 'Combinatorial theorems on classifications of subsets of a given set', *Proc. London Math. Soc.* **3**(1) (1952), 417–439.
[9] P. Erdős, A. Hajnal and R. Rado, 'Partition relations for cardinal numbers', *Acta Math. Acad. Sci. Hung.* **16** (1965), 93–196.
[10] P. Erdős and A. Szemerédi, 'On a Ramsey type theorem', *Period. Math. Hungar.* **2**(1-4) (1972), 295–299.
[11] J. Fox and B. Sudakov, 'Two remarks on the Burr–Erdős conjecture', *European J. Combin.* **30** (2009), 1630–1645.
[12] R. Graham, B. Rothschild and J. Spencer, *Ramsey Theory* vol. 20 (John Wiley & Sons, 1990).
[13] A. Kostochka and V. Rödl, 'On Ramsey numbers of uniform hypergraphs with given maximum degree', *J. Combin. Theory Ser. A* **113**(7) (2006), 1555–1564.
[14] A. V. Kostochka and B. Sudakov, 'On Ramsey numbers of sparse graphs', *Combin. Probab. Comput.* **12** (2003), 627–641.
[15] C. Lee, 'Ramsey numbers of degenerate graphs', *Ann. of Math.* **185**(3) (2017), 791–829.
[16] D. Mubayi and A. Suk, 'Cliques with many colors in triple systems', Preprint, 2020, arXiv:2005.03078.
[17] S. Shelah, 'Erdős and Rényi conjecture', *J. Combin. Theory Ser. A* **82**(2) (1998), 179–185.
[18] R. P. Stanley, *Enumerative Combinatorics* (Cambridge Studies in Advanced Mathematics) vol. 1 (Cambridge University Press, New York, 2011).
[19] X. Xu, Z. Shao, W. Su and Z. Li, 'Set-coloring of edges and multigraph Ramsey numbers', *Graphs Combin.* **25**(6) (2009), 863–870.

https://doi.org/10.1017/fms.2023.89 Published online by Cambridge University Press