

Hypergraph Ramsey numbers of cliques versus stars

David Conlon* Jacob Fox† Xiaoyu He‡ Dhruv Mubayi§ Andrew Suk¶
Jacques Verstraëte||

Abstract

Let $K_m^{(3)}$ denote the complete 3-uniform hypergraph on m vertices and $S_n^{(3)}$ the 3-uniform hypergraph on $n+1$ vertices consisting of all $\binom{n}{2}$ edges incident to a given vertex. Whereas many hypergraph Ramsey numbers grow either at most polynomially or at least exponentially, we show that the off-diagonal Ramsey number $r(K_4^{(3)}, S_n^{(3)})$ exhibits an unusual intermediate growth rate, namely,

$$2^{c \log^2 n} \leq r(K_4^{(3)}, S_n^{(3)}) \leq 2^{c' n^{2/3} \log n}$$

for some positive constants c and c' . The proof of these bounds brings in a novel Ramsey problem on grid graphs which may be of independent interest: what is the minimum N such that any 2-edge-coloring of the Cartesian product $K_N \square K_N$ contains either a red rectangle or a blue K_n ?

1 Introduction

A k -uniform hypergraph (henceforth, k -graph) $G = (V, E)$ consists of a vertex set V and an edge set $E \subseteq \binom{V}{k}$. In particular, we write $K_n^{(k)}$ for the complete k -graph with $V = [n]$ and $E = \binom{[n]}{k}$. Ramsey's theorem states that for any k -graphs H_1 and H_2 there is a positive integer N such that any k -graph G of order N either contains H_1 (as a subgraph) or its complement \overline{G} contains H_2 . The Ramsey number $r(H_1, H_2)$ is the smallest such N and the main objective of hypergraph Ramsey theory is to determine the size of these Ramsey numbers. In particular, for fixed $k \geq 3$, the two central problems in the area are to determine the growth rate of the “diagonal” Ramsey number $r(K_n^{(k)}, K_n^{(k)})$ as $n \rightarrow \infty$ and the “off-diagonal” Ramsey number $r(K_m^{(k)}, K_n^{(k)})$ where $m > k$ is fixed and $n \rightarrow \infty$.

Seminal work of Erdős–Rado [11] and Erdős–Hajnal (see, e.g., [8]) reduces the estimation of diagonal Ramsey numbers for $k \geq 3$ to the $k = 3$ case. For off-diagonal Ramsey numbers, the only case for which the tower height of the growth rate is not known is $r(K_{k+1}^{(k)}, K_n^{(k)})$, though it was noted in [17] that this

*Department of Mathematics, California Institute of Technology, Pasadena, CA 91125. Email: dconlon@caltech.edu.
Research supported by NSF Award DMS-2054452.

†Department of Mathematics, Stanford University, Stanford, CA 94305. Email: jacobfox@stanford.edu. Research supported by a Packard Fellowship and by NSF Awards DMS-1800053 and DMS-2154169.

‡Department of Mathematics, Princeton University, Princeton, NJ 08544. Email: xiaoyuh@princeton.edu. Research supported by NSF Award DMS-2103154.

§Department of Mathematics, Statistics and Computer Science, University of Illinois, Chicago, IL 60607. Email: mubayi@uic.edu. Research partially supported by NSF Awards DMS-1763317, DMS-1952767 and DMS-2153576 and by a Humboldt Research Award.

¶Department of Mathematics, University of California at San Diego, La Jolla, CA 92093. Email: asuk@ucsd.edu. Research supported by an NSF CAREER Award and by NSF Award DMS-1952786.

||Department of Mathematics, University of California at San Diego, La Jolla, CA 92093. Email: jacques@ucsd.edu. Research supported by NSF Award DMS-1800332.

tower height could be determined by proving that the 4-uniform Ramsey number $r(K_5^{(4)}, K_n^{(4)})$ is double exponential in a power of n . Moreover, it was shown in [16] that if $r(K_n^{(3)}, K_n^{(3)})$ grows double exponentially in a power of n , then the same is also true for $r(K_5^{(4)}, K_n^{(4)})$. Hence, the growth rate for all diagonal and off-diagonal hypergraph Ramsey numbers with $k \geq 4$ would follow from knowing the growth rate of the diagonal Ramsey number when $k = 3$. Because of this pivotal role, we will restrict our attention in the discussion below to the case $k = 3$.

Despite considerable progress in this area in recent years, our state of knowledge about the two central problems mentioned above remains rather dismal. The best known bounds in the diagonal case (see, e.g., the survey [9]) are of the form

$$2^{cn^2} \leq r(K_n^{(3)}, K_n^{(3)}) \leq 2^{2^{c'n^2}} \quad (1.1)$$

for some positive constants c and c' , differing by an entire exponential order. For the off-diagonal case, when $s \geq 4$ is fixed, the best known bounds [7] are of the form

$$2^{cn \log n} \leq r(K_s^{(3)}, K_n^{(3)}) \leq 2^{c'n^{s-2} \log n} \quad (1.2)$$

for some positive constants c and c' , differing by a power of n in the exponent.

As a possible approach to improving the lower bounds in (1.1) and (1.2), Fox and He [13] gave new lower-bound constructions for the Ramsey numbers $r(K_n^{(3)}, S_t^{(3)})$, where $S_t^{(3)}$ is the “star 3-graph” on $t+1$ vertices whose edges are all $\binom{t}{2}$ triples containing a given vertex. For $n \rightarrow \infty$, they showed that $r(K_n^{(3)}, S_n^{(3)}) \geq 2^{cn^2}$ for some positive constant c , giving a new proof of the lower bound in (1.1), while, for $t \geq 3$ fixed, they showed that $r(K_n^{(3)}, S_t^{(3)}) = 2^{\Theta(n \log n)}$, giving another proof of the lower bound in (1.2). In particular, when $t = 3$, this implies that $r(K_n^{(3)}, K_4^{(3)} - e) = 2^{\Theta(n \log n)}$. One surprising feature of these results is that they give the same bounds that were previously known for clique Ramsey numbers, but with one of the cliques replaced by the sparser corresponding star.

Fox and He essentially settled the growth rate of $r(K_n^{(3)}, S_t^{(3)})$ in two cases: when the star is fixed in size and the clique grows; and when the star and clique grow together. The present paper studies the remaining regime: when the clique is fixed in size and the star grows. First, a standard application of the Lovász Local Lemma yields the following proposition.

Proposition 1.1. *There are positive constants c and c' such that*

$$c \frac{n^2}{\log^2 n} \leq r(K_4^{(3)} - e, S_n^{(3)}) \leq c' \frac{n^2}{\log n}.$$

Our main result is that if we replace $K_4^{(3)} - e$ by $K_4^{(3)}$, then the Ramsey number has an exotic growth rate intermediate between polynomial and exponential. On the other hand, once we move up to $K_5^{(3)}$ the growth rate stabilizes to exponential.

Theorem 1.2. *There are positive constants c and c' such that*

$$2^{c \log^2 n} \leq r(K_4^{(3)}, S_n^{(3)}) \leq 2^{c' n^{2/3} \log n}.$$

Moreover, for each $s \geq 5$, there are positive constants c_s and c'_s such that $2^{c_s n} \leq r(K_s^{(3)}, S_n^{(3)}) \leq 2^{c'_s n}$.

The intermediate growth rate of $r(K_4^{(3)}, S_n^{(3)})$ is in striking contrast with the “polynomial-to-exponential transition” conjecture of Erdős and Hajnal [10], whose exact statement is technical but roughly states that all off-diagonal hypergraph Ramsey numbers against cliques are either at most polynomial or at least exponential. This conjecture was proved to be true infinitely often when $k = 3$ by Conlon, Fox and

Sudakov [7] and was then settled in the affirmative for all $k \geq 4$ by Mubayi and Razborov [15]. As a corollary of Theorem 1.2, we see that no such transition can occur for hypergraph Ramsey numbers against stars.

Since our results for $r(K_4^{(3)} - e, S_n^{(3)})$ and $r(K_s^{(3)}, S_n^{(3)})$ when $s \geq 5$ are straightforward applications of known methods, we will focus our attention in the remainder of the introduction on how we estimate $r(K_4^{(3)}, S_n^{(3)})$. The key idea is to reduce to a novel Ramsey problem involving grids, somewhat reminiscent of the grid case of the famous cube lemma (see, e.g., [6]) developed by Shelah [18] in his proof of primitive-recursive bounds for Hales–Jewett and van der Waerden numbers. To say more, we need some further definitions.

Let the $m \times n$ grid graph $G_{m \times n}$ be the Cartesian product $K_m \square K_n$, i.e., the graph whose vertex set is the rectangular grid $[m] \times [n]$ and whose edges are all pairs of distinct vertices sharing exactly one coordinate. If $m \geq a \geq 1$ and $n \geq b \geq 1$, then an $a \times b$ subgrid in $G_{m \times n}$ is an induced copy of $G_{a \times b}$. In particular, we call a 2×2 subgrid, i.e., a set of four vertices $(x, y), (x, y'), (x', y'), (x', y)$ with the four axis-parallel edges between them, a *rectangle*. We will be interested in the Ramsey number $\text{gr}(G_{2 \times 2}, K_n)$, which is the smallest N such that in any 2-edge coloring of $G_{N \times N}$ there is either a monochromatic red rectangle or a monochromatic blue K_n . For such Ramsey numbers, we show the following.

Theorem 1.3. *There are positive constants c and c' such that*

$$2^{c \log^2 n} \leq \text{gr}(G_{2 \times 2}, K_n) \leq 2^{c' n^{2/3} \log n}.$$

In passing, we remark that the best known lower bound for the grid case of Shelah’s cube lemma, due to Conlon, Fox, Lee and Sudakov [6], is of the form $2^{c(\log r)^{2.5}/\sqrt{\log \log r}}$ for some positive constant c , curiously similar to the lower bound in Theorem 1.3, though there r refers to the number of colors. However, despite the similarities, we were not able to find any nontrivial connection between the two problems.

The connection between $r(K_4^{(3)}, S_n^{(3)})$ and $\text{gr}(G_{2 \times 2}, K_n)$ is that the latter is equivalent to a natural bipartite variant of the former. Let $B^{(3)}(a, b)$ denote the complete bipartite 3-graph on $[a + b]$ whose edges are all triples intersecting both $[a]$ and $[a + 1, a + b]$. Observe that the 3-graph $B^{(3)}(a, b)$ and the grid graph $G_{a \times b}$ have the same number $a \binom{b}{2} + b \binom{a}{2}$ of edges and we can give an explicit correspondence between their edge sets by sending horizontal edges $(x, y) \sim (x', y)$ in the grid to triples $\{x, x', a + y\} \in E(B^{(3)}(a, b))$ and vertical edges $(x, y) \sim (x, y')$ to triples $\{x, a + y, a + y'\} \in E(B^{(3)}(a, b))$. It is easy to check that rectangles in the square grid $G_{N \times N}$ correspond to copies of $K_4^{(3)}$ in the corresponding bipartite 3-graph $B^{(3)}(N, N)$, while n -cliques in $G_{N \times N}$ correspond to copies of $S_n^{(3)}$ in $B^{(3)}(N, N)$. Thus, $\text{gr}(G_{2 \times 2}, K_n)$ is exactly equal to the smallest N such that any 2-edge-coloring of $B^{(3)}(N, N)$ contains either a red $K_4^{(3)}$ or a blue $S_n^{(3)}$.

As $B^{(3)}(N, N) \subseteq K_{2N}^{(3)}$, it follows immediately that $r(K_4^{(3)}, S_n^{(3)}) \leq 2\text{gr}(G_{2 \times 2}, K_n)$, so the upper bound in Theorem 1.3 implies that in Theorem 1.2. We do not know of a direct implication between the lower bounds, but we will be able to glue together copies of our lower-bound construction for $\text{gr}(G_{2 \times 2}, K_n)$ to manufacture one for $r(K_4^{(3)}, S_n^{(3)})$ of comparable size. This suggests, and we strongly believe, that $r(K_4^{(3)}, S_n^{(3)})$ and $\text{gr}(G_{2 \times 2}, K_n)$ are of roughly the same order.

The final result that we mention here is a generalization of Theorem 1.3 to larger grids.

Theorem 1.4. *There is a positive constant c and, for all fixed $a \geq b \geq 2$, a positive constant $c' = c'_a$ such that*

$$2^{c \log^2 n} \leq \text{gr}(G_{a \times b}, K_n) \leq 2^{c' n^{1-(2^b-1)^{-1}} \log n}.$$

The lower bound is an obvious corollary of the lower bound in Theorem 1.3, while the upper bound involves some extra effort, in particular drawing on recent work of the authors on set-coloring Ramsey

numbers [5]. One interesting corollary of this result is that there are positive constants c and c' such that

$$2^{c \log^2 n} \leq r(K_5^{(3)} - e, S_n^{(3)}) \leq 2^{c' n^{2/3} \log n}.$$

The lower bound here simply follows from the fact that $r(K_5^{(3)} - e, S_n^{(3)}) \geq r(K_4^{(3)}, S_n^{(3)})$. For the upper bound, note that, just as $\text{gr}(G_{2 \times 2}, K_n)$ is equivalent to a bipartite variant of $r(K_4^{(3)}, S_n^{(3)})$, we also have that $\text{gr}(G_{3 \times 2}, K_n)$ is equivalent to a bipartite variant of $r(K_5^{(3)} - e, S_n^{(3)})$, so that $r(K_5^{(3)} - e, S_n^{(3)}) \leq 2\text{gr}(G_{3 \times 2}, K_n)$, yielding the required upper bound. Together with the fact that $r(K_5^{(3)}, S_n^{(3)}) = 2^{\Theta(n)}$, this gives a more complete picture of the transition window.

Throughout the paper, for the sake of clarity of presentation, we systematically omit floor and ceiling signs whenever they are not essential. Moreover, unless otherwise specified, all logarithms are base 2.

2 Basic bounds

In this short section, we prove Proposition 1.1 and the second part of Theorem 1.2. We will use the Lovász Local Lemma in the following standard form (see, e.g., [2, Lemma 5.1.1]).

Lemma 2.1 (Lovász Local Lemma). *Let A_1, A_2, \dots, A_n be events in an arbitrary probability space. A directed graph $D = (V, E)$ on the set of vertices $V = [n]$ is called a dependency digraph for the events A_1, \dots, A_n if for each i , $1 \leq i \leq n$, the event A_i is mutually independent of all the events $\{A_j : (i, j) \notin E\}$. Suppose that $D = (V, E)$ is a dependency digraph for the above events and suppose there are real numbers x_1, \dots, x_n such that $0 \leq x_i < 1$ and $\Pr[A_i] \leq x_i \prod_{(i, j) \in E} (1 - x_j)$ for all $1 \leq i \leq n$. Then $\Pr[\bigwedge_{i=1}^n \overline{A_i}] \geq \prod_{i=1}^n (1 - x_i)$. In particular, with positive probability no event A_i holds.*

The proof of Proposition 1.1 is now a direct application of Lemma 2.1 to a suitable random 3-graph.

Proof of Proposition 1.1. The upper bound follows from the fact that $r(K_4^{(3)} - e, S_n^{(3)}) \leq r(K_3, K_n) + 1$, by specializing to the link of a single vertex, and then applying the bound $r(K_3, K_n) = O(\frac{n^2}{\log n})$ due to Ajtai, Komlós and Szemerédi [1].

For the lower bound, we will use Lemma 2.1, assuming throughout that n is sufficiently large. Consider a random hypergraph $\Gamma = G^{(3)}(N, p)$ on $N = 10^{-3}n^2(\log n)^{-2}$ vertices (which we identify with $[N]$ for convenience) with $p = 4(\log n)/n$. We would like to show that, with positive probability, Γ contains no $K_4^{(3)} - e$ and its complement $\overline{\Gamma}$ contains no $S_n^{(3)}$. Let S be the collection of all $N \binom{N-1}{3}$ copies of $K_4^{(3)} - e$ on $[N]$ and, for $s \in S$, let A_s be the event that $s \subseteq \Gamma$. Let T be the collection of all $N \binom{N-1}{n}$ copies of $S_n^{(3)}$ on $[N]$ and, for $t \in T$, let B_t be the event that $t \subseteq \overline{\Gamma}$. The probabilities of these events are $\Pr[A_s] = p^3$ and $\Pr[B_t] = (1 - p)^{\binom{n}{2}}$.

Let D be the digraph whose vertex set is $S \cup T$ and whose edges are pairs $(i, j) \in (S \cup T)^2$ which intersect in at least one edge. Thus, D is a dependency digraph for the events $\{A_s\}_{s \in S} \cup \{B_t\}_{t \in T}$. Each vertex in S is adjacent to at most $9N$ other vertices in S and at most $|T|$ vertices in T , while each vertex in T is adjacent to at most $\binom{n}{2}3N < 2n^2N$ vertices in S and at most $|T|$ vertices in T . Let $x = 3p^3$ be the local lemma weight for all the A_s events and $y = 1/|T|$ be the local lemma weight for all the B_t events. With this choice of parameters and using that n is sufficiently large and $|T| = N \binom{N-1}{n} < N(eN/n)^n$, we have

$$\begin{aligned} x(1 - x)^{9N}(1 - y)^{|T|} &\geq p^3, \\ y(1 - x)^{2n^2N}(1 - y)^{|T|} &\geq (1 - p)^{\binom{n}{2}}, \end{aligned}$$

so the conditions of Lemma 2.1 are satisfied. Thus, with positive probability none of the events A_s or B_t hold and we obtain $r(K_4^{(3)} - e, S_n^{(3)}) \geq N$, as desired. \square

The close connection between $r(K_4^{(3)} - e, S_n^{(3)})$ and $r(K_3, K_n)$ suggests that $r(K_4^{(3)} - e, S_n^{(3)}) = \Theta(\frac{n^2}{\log n})$. It seems likely that a proof of this may be possible through a careful analysis of the $(K_4^{(3)} - e)$ -free process (see, e.g., [4] for results of this type in a similar context). However, we have chosen not to pursue this here.

We now prove the bounds on $r(K_s^{(3)}, S_n^{(3)})$ for $s \geq 5$ stated in Theorem 1.2.

Proof of Theorem 1.2 for $s \geq 5$. The upper bound follows from [13, Theorem 1.4], which says that

$$r(K_s^{(3)}, S_n^{(3)}) < (2s)^{sn}$$

for all $s, n \geq 3$. The lower bound construction is as follows. Let K_N be a complete graph on N vertices, labelled v_1, \dots, v_N , and let ϕ be a 3-coloring of the edges of K_N which independently colors each edge by a uniform random element of $\mathbb{Z}/3\mathbb{Z}$. If $K_N^{(3)}$ is a complete 3-graph on the same vertex set, define a 2-coloring χ of the edges of $K_N^{(3)}$ where $\chi(v_i, v_j, v_k)$ is red if $\phi(v_i, v_j) + \phi(v_i, v_k) + \phi(v_j, v_k) \equiv 1 \pmod{3}$ and blue otherwise.

Suppose, for the sake of contradiction, that a red $K_5^{(3)}$ appears in the coloring χ at vertices u_1, \dots, u_5 . If we sum up $\phi(u_i, u_j) + \phi(u_i, u_k) + \phi(u_j, u_k)$ across all $\binom{5}{3} = 10$ triples of these vertices, the sum is 1 $\pmod{3}$, since each triple sums to 1 $\pmod{3}$. On the other hand, each summand $\phi(u_i, u_j)$ appears three times in this sum, so the total sum of all these triples must be 0 $\pmod{3}$. This is a contradiction, so no such coloring can have a red $K_5^{(3)}$.

Next, we show that for $N = (3/2)^{n/2}$ and $n \geq 5$, the probability of a blue $S_n^{(3)}$ appearing in χ is less than 1. Indeed, consider any given copy of $S_n^{(3)}$ in $K_N^{(3)}$, with central vertex u and clique w_1, \dots, w_n in the link of u . In order for every edge in this copy of $S_n^{(3)}$ to be blue, every color $\phi(w_i, w_j)$ must satisfy $\phi(w_i, w_j) + \phi(u, w_i) + \phi(u, w_j) \not\equiv 1 \pmod{3}$. Each such event is independent with probability $2/3$, so we obtain that the probability this particular copy of $S_n^{(3)}$ is blue in χ is exactly $(2/3)^{\binom{n}{2}}$. As there are $N \binom{N-1}{n}$ copies of $S_n^{(3)}$ in $K_N^{(3)}$, we see that the expected number of blue $S_n^{(3)}$ is

$$N \binom{N-1}{n} \cdot (2/3)^{\binom{n}{2}} < N(eN/n)^n (2/3)^{n(n-1)/2} = (3e/2n)^n < 1.$$

Hence, there is such a coloring with no blue $S_n^{(3)}$. □

3 The lower bound

The key ingredient for our lower bound on $r(K_4^{(3)}, S_n^{(3)})$ is the following lemma, which states that it is possible to construct a random subgraph of the grid where each row and column looks like a sparse Erdős–Rényi random graph, but they are coupled in such a way that there are no rectangles and their edge unions are sparse.

Lemma 3.1. *There exists a positive constant c such that, for all n sufficiently large and $N = 2^{c \log^2 n}$, there is a random subgraph $H \subseteq G_{N \times N}$ with the following properties:*

- (1) *For every row r_y , $H[r_y] \sim G(N, n^{-3/4})$, i.e., the marginal distribution of the induced subgraph $H[r_y]$ is $G(N, n^{-3/4})$. Similarly, for every column c_x , $H[c_x] \sim G(N, n^{-3/4})$.*
- (2) *There are no rectangles, i.e., no x, x', y, y' with $(x, y) \sim (x, y') \sim (x', y') \sim (x', y) \sim (x, y)$, in H .*
- (3) *The edge union of all the row graphs $H[r_y]$ lies in a $G(N, n^{-1/8})$. Similarly, the edge union of all the column graphs $H[c_x]$ lies in a $G(N, n^{-1/8})$.*

Note that properties (1) and (2) of the random graph H are already enough to prove the lower bound in Theorem 1.3, since if we let H determine the set of red edges in $G_{N \times N}$ and its complement the blue edges, then (2) shows that there are no red rectangles and (1) implies that w.h.p. there are no blue K_n . We will use the additional property (3) to prove the lower bound on $r(K_4^{(3)}, S_n^{(3)})$ in Theorem 1.2.

Proof of Lemma 3.1. We give a construction for H which starts by choosing the column graphs in such a way that every pair of columns is edge-disjoint on some large vertex subset. This then allows us to place many edges between these columns without creating a rectangle.

Setting up. Let $N = 2^{c \log^2 n}$. We start by picking a family of subsets $\{U_i\}_{i \leq T}$ of $[N]$ with $T = n^{1/2}(\log n)^{10}$ such that each element y lies in $d(y)$ sets, where $d(y) \in [\frac{1}{2}(\log n)^{10}, \frac{3}{2}(\log n)^{10}]$, and every pair of elements lies in at most $\frac{1}{4} \log n$ subsets. Such a family exists by the probabilistic method. To see this, pick the sets U_i independently such that each $j \in [N]$ lies in U_i independently with probability $n^{-1/2}$. The expected value of $d(y)$ is $(\log n)^{10}$. The multiplicative Chernoff bound for a binomial random variable X implies that $\Pr(|X - \mathbb{E}X| \geq \frac{1}{2}\mathbb{E}X) < 2 \exp(-\frac{1}{12}\mathbb{E}X)$. Together with a union bound over the N elements y , this yields that the probability there exists y with $d(y) \notin [\frac{1}{2}(\log n)^{10}, \frac{3}{2}(\log n)^{10}]$ is less than $N \cdot 2 \exp(-\frac{1}{12}(\log n)^{10}) < \frac{1}{2}$ for n large. The probability that distinct elements y, y' are both in a given U_i is n^{-1} , so the probability they are both in at least $\frac{1}{4} \log n$ of the U_i is at most

$$T^{\frac{1}{4} \log n} \cdot (n^{-1})^{\frac{1}{4} \log n} \leq 2^{-\frac{1}{16} \log^2 n}$$

for n large enough. Hence, for c sufficiently small, the probability that there exists a pair of vertices y, y' that lies in at least $\frac{1}{4} \log n$ sets is at most $\binom{N}{2} 2^{-\frac{1}{16} \log^2 n} < \frac{1}{2}$. Therefore, the required family $\{U_i\}_{i \leq T}$ exists.

Next we show that there is a collection of bipartitions $P_i \sqcup Q_i = [N]$, one for each $1 \leq i \leq T$, of the set of columns satisfying the following two properties:

- (a) Every pair of columns x, x' lies on opposite sides of $(\frac{1}{2} + o(1))T$ bipartitions.
- (b) For every horizontal edge $(x, y) \sim (x', y)$, the number of i for which $y \in U_i$ and x, x' lie on opposite sides of the bipartition $P_i \sqcup Q_i$ is $\Theta((\log n)^{10})$.

To see that properties (a) and (b) can be satisfied simultaneously, we show that for a random choice of the bipartitions $P_i \sqcup Q_i = [N]$, both properties hold with high probability. Indeed, if $D_{x, x'}$ is the set of all i for which x, x' lie on opposite sides of the bipartition $P_i \sqcup Q_i$, then $|D_{x, x'}| \sim \text{Bin}(T, 1/2)$ for all choices of x and x' . By the Chernoff bound,

$$\Pr[|\text{Bin}(T, 1/2) - T/2| > \varepsilon T] < e^{-\Omega_\varepsilon(T)},$$

so, even after a union bound over all $\binom{N}{2} = e^{O(\log^2 n)}$ choices of x and x' , we have that w.h.p. $|D_{x, x'}| = (1/2 + o(1))T$ for all x, x' . That is, property (a) holds w.h.p. To check that property (b) also holds w.h.p., note that another application of the Chernoff bound shows that if $D_{x, x'}(y)$ is the set of all $i \in D_{x, x'}$ satisfying the additional condition that $y \in U_i$, then $|D_{x, x'}(y)| \sim \text{Bin}(d(y), 1/2)$ must be tightly concentrated around $d(y)/2 = \Theta((\log n)^{10})$, even after taking a union bound over all choices of x, x' and y . We may therefore fix a partition $P_i \sqcup Q_i$ for each $i \in [T]$ such that the collection of such partitions satisfies (a) and (b).

To force property (3), we sample two random graphs $R \sim G(N, n^{-1/8})$ and $C \sim G(N, n^{-1/8})$ in advance; we will make sure that the rows of H only take edges from R and the columns of H only take edges from C . Finally, for each $i \in [T]$, let $A_i = G(|U_i|, 1/2)$ be a random graph on vertex set U_i and let $B_i = \overline{A_i}$ be the edge-complement of A_i .

We emphasize here that for each $i \in [T]$, the sets U_i and the pairs (P_i, Q_i) are now fixed. Our goal is to define a probability space (a random subgraph $H \subset G_{N \times N}$) and, thus, all probabilistic statements that follow are with respect to the product space $(\prod A_i) \times R \times C$.

The columns. We first decide the columns of H . Let c_x be the column indexed x in $G_{N \times N}$. We define H_x to be the (random) graph with vertex set $c_x \cong [N]$ such that (y, y') is an edge of H_x if and only if, for every U_i containing both y and y' , either $x \in P_i$ and $(y, y') \in E(A_i)$ or $x \in Q_i$ and $(y, y') \in E(B_i)$. In words, on each column we stipulate that within each of the subsets U_i , the induced subgraph $H_x[U_i]$ is a subgraph of one of the two complementary random graphs A_i or B_i , according to which part of the partition $P_i \sqcup Q_i$ the x -coordinate falls into.

Let $E_x(y, y')$ be the event that a given edge (y, y') appears in the random graph H_x . By definition, $E_x(y, y')$ occurs if and only if for every U_i containing both y and y' , either $x \in P_i$ and $(y, y') \in E(A_i)$ or $x \in Q_i$ and $(y, y') \in E(B_i)$. We have

$$\Pr[(x \in P_i \wedge (y, y') \in E(A_i)) \vee (x \in Q_i \wedge (y, y') \in E(B_i))] = 1/2.$$

There are at most $\frac{1}{4} \log n$ choices of i for which U_i contains both y and y' and these events are independent over i . Thus, $\Pr[E_x(y, y')] \geq 2^{-\frac{1}{4} \log n} = n^{-1/4}$. We observe further that $E_x(y, y')$ depends only on the randomness of the single edge (y, y') in the A_i and B_i , so, for a fixed x , these events are mutually independent as (y, y') varies through the possible edges of H_x . Thus, we may choose a random subgraph $H'_x \subseteq H_x$ with distribution exactly $G(N, n^{-5/8})$. Finally, we take $H[c_x] = H'_x \cap R$, which is a random graph with distribution exactly $G(N, n^{-3/4})$, proving properties (1) and (3) for the columns.

The rows. Next, we define the horizontal edges of H by picking the edges between each pair of columns independently. For each pair of columns $c_x, c_{x'}$, recall that $D_{x,x'}$ is the set of all $i \in [N]$ for which x, x' fall on opposite sides of the partition $P_i \sqcup Q_i$. By our choice of the bipartitions, we know that $|D_{x,x'}| = (1/2 + o(1))T$. For each pair x, x' , pick a uniform random $i_{x,x'} \in D_{x,x'}$. Now, in each row y , let H_y be the random graph whose edges are exactly those pairs (x, x') for which $U_{i_{x,x'}} \ni y$. A given edge (x, x') appears in H_y if and only if the random index $i_{x,x'}$ is chosen to be one of the $d(y) \in [\frac{1}{2}(\log n)^{10}, \frac{3}{2}(\log n)^{10}]$ indices $i \in D_{x,x'}$ for which $y \in U_i$, while $i_{x,x'}$ is uniform out of $|D_{x,x'}| = (\frac{1}{2} + o(1))T = (\frac{1}{2} + o(1))n^{1/2}(\log n)^{10}$ choices. Thus, each edge appears in H_y with probability $\Theta(n^{-1/2})$. Furthermore, for fixed y these events are mutually independent over all choices of possible (x, x') , since the random indices $i_{x,x'}$ are chosen independently. We may therefore find a random subgraph $H'_y \subseteq H_y$ with distribution exactly $G(N, n^{-5/8})$. Finally, we take $H[r_y] = H'_y \cap C$, which is again a random graph with distribution exactly $G(N, n^{-3/4})$, proving properties (1) and (3) for the rows.

Property (2). Suppose that there is a rectangle in H , say $(x, y), (x, y'), (x', y), (x', y')$. By the way we picked the horizontal edges, this means that, for $i = i_{x,x'}$, we have $y, y' \in U_i$ and x, x' fall on opposite sides of the bipartition $P_i \sqcup Q_i$. If, say, $x \in P_i$ and $x' \in Q_i$, then we see that $H[c_x][U_i] \subseteq A_i$ and $H[c_{x'}][U_i] \subseteq B_i$ are disjoint graphs on the set U_i , so at most one of the two vertical edges $(x, y) \sim (x, y')$ and $(x', y) \sim (x', y')$ can lie in H . Hence, there are no rectangles in H , as desired. \square

Now that Lemma 3.1 is proved, we layer $\log N$ copies of this bipartite construction on top of each other to obtain the lower bound on $r(K_4^{(3)}, S_n^{(3)})$ in Theorem 1.2, namely, $r(K_4^{(3)}, S_n^{(3)}) \geq 2^{c \log^2 n}$ for some $c > 0$.

Proof of the lower bound on $r(K_4^{(3)}, S_n^{(3)})$. For N as in Lemma 3.1, draw $t = \log N$ independent samples H_1, \dots, H_t from the distribution H . Identify the vertices of $K_N^{(3)}$ with $[N] := \{0, \dots, N-1\}$ (we use this convention so that each vertex has at most t bits when written in binary).

Each H_ℓ , $1 \leq \ell \leq t$, gives rise to a two-edge-coloring χ_ℓ of a bipartite subgraph of $K_N^{(3)}$ as follows. Let $\ell(i_1, i_2, i_3)$ denote the maximum binary bit on which three distinct $i_1, i_2, i_3 \in [N]$ do not agree. Let Γ_ℓ denote the spanning subgraph of $K_N^{(3)}$ consisting of all edges with $\ell(i_1, i_2, i_3) = \ell$. For clarity, we write the vertices of an edge in Γ_ℓ as $\{x, y, y'\}$ if x is the vertex which is 0 on bit ℓ and y, y' are the vertices which are 1, calling these “vertical edges”, and as $\{x, x', y\}$ if x, x' are 0 on bit ℓ and y is 1, calling these “horizontal edges”. Let χ_ℓ denote the coloring of Γ_ℓ for which vertical edges $\{x, y, y'\} \in \Gamma_\ell$ are colored red if and only if $(x, y) \sim (x, y')$ is an edge of H_ℓ and horizontal edges $\{x, x', y\} \in \Gamma_\ell$ are colored red if and only if $(x, y) \sim (x', y)$ is an edge of H_ℓ .

We first claim that the colorings χ_ℓ contain no red copies of $K_4^{(3)}$. Indeed, since Γ_ℓ is bipartite, a copy of $K_4^{(3)}$ in Γ_ℓ must lie on four vertices $\{x, x', y, y'\}$ where x, x' are 0 on bit ℓ and y, y' are 1 on bit ℓ . This induced subhypergraph is red if and only if the four edges $\{x, y, y'\}$, $\{x', y, y'\}$, $\{x, x', y\}$ and $\{x, x', y'\}$ are all red in χ_ℓ . This in turn means that the four edges $(x, y) \sim (x, y')$, $(x', y) \sim (x', y')$, $(x, y) \sim (x', y)$ and $(x, y') \sim (x', y')$ are edges of H_ℓ , forming a rectangle in H and contradicting property (2) of Lemma 3.1. Thus, no red $K_4^{(3)}$ appears in any of the colorings χ_ℓ .

Write $N_\ell(v)$ for the set of all $u \in [N]$ which disagree with v on bit ℓ but agree on all higher bits and write $L_\ell(v)$ for the link of v in Γ_ℓ restricted to $N_\ell(v)$. By the definition of Γ_ℓ , $L_\ell(v)$ is a complete graph. Moreover, the coloring χ_ℓ induces a coloring on $L_\ell(v)$ for each v , which, by property (1) of Lemma 3.1, has red edges distributed as in $G(N, n^{-3/4})$.

We now build a coloring χ of $K_N^{(3)}$ out of the colorings χ_ℓ as follows. Note that $\Gamma_1, \dots, \Gamma_t$ form an edge partition of $K_N^{(3)}$. As a starting point, we let $\chi'(e) = \chi_\ell(e)$ for that ℓ such that $e \in \Gamma_\ell$. However, this coloring χ' may now contain some red copies of $K_4^{(3)}$, so we modify it as follows. For each red $K_4^{(3)}$ in χ' , say with vertices $\{i_1, i_2, i_3, i_4\}$, mark the triple of vertices $\{i_1, i_2, i_3\}$ that has the smallest value ℓ of $\ell(i_1, i_2, i_3)$. To see that this triple is unique, suppose it were not and $\ell(i_1, i_2, i_3) = \ell(i_1, i_2, i_4) = \ell$. But then all four vertices agree on all higher bits than ℓ , so the ℓ -values of all four 3-tuples are at most ℓ by definition. Thus, all four 3-tuples among $\{i_1, i_2, i_3, i_4\}$ lie in Γ_ℓ , giving a red $K_4^{(3)}$ in the coloring χ_ℓ of Γ_ℓ , which is a contradiction. We also observe that if $\{i_1, i_2, i_3\}$ is marked by a red clique on $\{i, i_1, i_2, i_3\}$, then $\ell(i, i_1, i_2) = \ell(i, i_1, i_3) = \ell(i, i_2, i_3) = \ell'$ for some $\ell' > \ell$. That is, the other three edges in this red clique all belong to the same χ_ℓ , so we may say that the edge $\{i_1, i_2, i_3\}$ is *marked by level ℓ'* (note that a single edge can be marked by multiple levels). The coloring χ is now defined as follows: the red edges of χ are exactly the unmarked red edges of χ' .

We claim that χ is a coloring of $K_N^{(3)}$ that contains no red $K_4^{(3)}$ and, with positive probability, no $S_{nt}^{(3)}$. Since every $K_4^{(3)}$ which is red under χ' contains a marked triple, χ indeed has no red $K_4^{(3)}$. It remains to bound the probability of finding a blue $S_{nt}^{(3)}$.

Fix a vertex $u \in V(K_N^{(3)})$ and suppose a blue $S_{nt}^{(3)}$ appears in χ with u as the central vertex. The sets $N_1(u), \dots, N_t(u)$ form a partition of $V(K_N^{(3)})$, so at least one of these contains at least n of the vertices of our K_{nt} . Thus, for some ℓ , there must be $v_1, \dots, v_n \in N_\ell(u)$ forming a blue $S_n^{(3)}$ with u as the central vertex. By the union bound, it will suffice to show that the probability of such an occurrence is smaller than N^{-n-1} .

Let ϕ be the coloring of the copy of K_n formed by the vertices of this copy of $S_n^{(3)}$ other than u , with colors given by $\phi(v_i, v_j) = \chi(u, v_i, v_j)$. By construction, the red edges in ϕ correspond to red edges in $L_\ell(u)$ in χ_ℓ that are unmarked. We first bound the number of marked edges. For each $\ell' > \ell$, let $M_{\ell'}$ be the graph on $\{v_1, \dots, v_n\}$ whose edges are pairs $\{v_i, v_j\}$ for which $\{u, v_i, v_j\}$, if it were red, would be marked by level ℓ' , as defined previously. In other words, $v_i \sim v_j$ in $M_{\ell'}$ if and only if there exists a fourth vertex w for which $\{w, u, v_i\}$, $\{w, u, v_j\}$ and $\{w, v_i, v_j\}$ are all red in $\chi_{\ell'}$.

We claim that for each level $\ell' > \ell$, the graph $M_{\ell'}$ is contained inside a copy of $G(n, n^{-1/8})$. Indeed, because of property (3) of Lemma 3.1 and the definition of $\chi_{\ell'}$, there exists a random graph $M'_{\ell'} \sim G(n, n^{-1/8})$

such that if $\{w, v_i, v_j\}$ is red for any w , then $\{v_i, v_j\}$ is an edge of $M'_{\ell'}$. In particular, if $\{v_i, v_j\}$ is an edge of $M_{\ell'}$, then $\{w, v_i, v_j\}$ is red in $\chi'_{\ell'}$ for some w and thus $\{v_i, v_j\}$ is an edge of $M'_{\ell'}$ as well. That is, $M_{\ell'}$ is a spanning subgraph of $M'_{\ell'}$, which has distribution $G(n, n^{-1/8})$.

Whether an edge is marked by a particular level is independent for each level, so the edge union of the graphs $M_{\ell'}$ with $\ell' > \ell$ is contained inside the edge union of the independent random graphs $M'_{\ell'}$, which is in turn contained inside a single random graph M with distribution $G(n, tn^{-1/8})$. Let E be the event that M has at least $\frac{1}{2}\binom{n}{2}$ edges. That is, E is the event that the edge count of M , distributed like $\text{Bin}(\binom{n}{2}, tn^{-1/8})$, is at least $\frac{1}{2}\binom{n}{2}$. By the Chernoff bound, $\Pr[E] \leq 2^{-\Omega(n^2)}$.

An edge (v_i, v_j) is red in ϕ if it does not appear in M and the edge $\{u, v_i, v_j\}$ is red in $\chi'_{\ell'}$. The latter occurs with probability $n^{-3/4}$, by property (1) of Lemma 3.1. Thus,

$$\Pr[\phi(v_i, v_j) \text{ red} | (v_i, v_j) \notin E(M)] \geq n^{-3/4}.$$

Given any particular choice of M , all such events are mutually independent, so the probability that ϕ is monochromatic blue is at most

$$\Pr[E] + \Pr[\phi \text{ is monochromatic blue} | \overline{E}] \leq 2^{-\Omega(n^2)} + \left(1 - n^{-3/4}\right)^{\frac{1}{2}\binom{n}{2}} \leq 2^{-\Omega(n^{5/4})},$$

which suffices to union bound over all N^{n+1} choices of u, v_1, \dots, v_n , as desired. This completes the proof. \square

4 The upper bound

In this section, we first prove the upper bound in Theorem 1.3, which states that

$$\text{gr}(G_{2 \times 2}, K_n) \leq 2^{c'n^{2/3} \log n} \tag{4.1}$$

for some positive constant c' . As observed in the introduction, $r(K_4^{(3)}, S_n^{(3)}) \leq 2\text{gr}(G_{2 \times 2}, K_n)$, so this immediately implies the upper bound in Theorem 1.2 as well.

The main technical tool used is the following Ramsey-type result of Erdős and Szemerédi.

Lemma 4.1 (Erdős–Szemerédi [12]). *There exists a positive constant c_0 such that if the edges of the complete graph K_N are colored in r colors, then there exists a clique of order $n = c_0 \frac{r}{\log r} \log N$ and a color i that does not appear on any edge in that clique.*

We are now ready to prove (4.1).

Proof of (4.1). Let $c' = \max(2, 1/c_0)$, where c_0 is the constant in Lemma 4.1. Let $N = 2^{c'n^{2/3} \log n}$. We would like to show that in any 2-edge-coloring of $G_{N \times N}$, there is either a red rectangle or a blue K_n . Letting $r = n^{1/3}$, we will in fact prove the stronger statement that the same result holds for the rectangular grid $G_{N \times M}$, where the height is chosen to be the Ramsey number $M = r(K_r, K_n) \leq n^r \leq N$.

Fix a 2-edge-coloring of $G_{N \times M}$. Each column is a 2-edge-colored K_M , so, by the definition of M , each column contains either a red K_r or a blue K_n . In the latter case we are already done, so we may assume that every column of the grid contains a red K_r . Associate with each column x the y -coordinates $\vec{y}(x) = (y_1, \dots, y_r)$ of the vertices of some red K_r in that column. Since $y(x)$ can take at most M^r possible values, there exist $N' = N/M^r$ columns with the same value $\vec{y} = \vec{y}(x)$. Since $M \leq n^r$, we have $M^r \leq n^{r^2} = 2^{n^{2/3} \log n}$. As $c' \geq 2$, we have $N' \geq \sqrt{N}$.

Restrict to the $N' \times r$ subgrid where the rows are the r rows indexed by the coordinates of \vec{y} and the columns are the N' columns x with $\vec{y}(x) = \vec{y}$. In this subgrid, every column is a monochromatic red K_r . If

any pair of columns has at least two red edges between them, then we have a red rectangle and we are done. Thus, we may assume that there is an edge-coloring $h : E(K_{N'}) \rightarrow [r]$ of the complete graph on $[N']$ such that the horizontal edge $(x, y) \sim (x', y)$ is blue whenever $y \neq h(x, x')$.

By Lemma 4.1, there is a clique of order

$$c_0 \frac{r}{\log r} \log N' \geq c_0 \frac{3n^{1/3}}{\log n} \cdot \frac{1}{2} \log N \geq n$$

and a color y such that $y \neq h(x, x')$ for every pair of vertices x, x' in the clique. If the vertices of the clique are $\{x_1, \dots, x_n\}$, then all edges between the vertices $\{(x_1, y), \dots, (x_n, y)\}$ in the original grid are blue, forming the desired blue K_n . This completes the proof. \square

Next, we generalize this upper bound to arbitrary grids. Recall that $G_{a \times b}$ is the $a \times b$ grid graph and $\text{gr}(G_{a \times b}, K_n)$ is the smallest N such that in any 2-edge-coloring of $G_{N \times N}$ there is either a monochromatic red copy of $G_{a \times b}$ or a monochromatic blue copy of K_n . We next prove the upper bound in Theorem 1.4, that, for all $a \geq b \geq 2$, there is a positive constant $c' = c'_a$ such that

$$\text{gr}(G_{a \times b}, K_n) \leq 2^{c' n^{1-(2^b-1)^{-1}} \log n}. \quad (4.2)$$

The proof is essentially the result of iterating the argument for Theorem 1.3. However, we will also require a generalization of Lemma 4.1 from a recent paper of the authors [5]. Define the *set-coloring Ramsey number* $R(n; r, s)$ to be the smallest positive integer N such that if every edge of K_N receives a set of s colors from a palette of r colors, then there must exist a copy of K_n where a single common color appears on every edge.

Lemma 4.2 (Corollary of [5, Theorem 1.1]). *There is a constant C_0 such that the following holds. For any integers $n \geq 3$ and $r > s \geq r/2 \geq 1$,*

$$R(n; r, s) \leq 2^{C_0 n(r-s)^2 r^{-1} \log(r/(r-s))}.$$

Combining the above lemma with Turán's theorem, we obtain the following result, needed for the iterative step in our proof of (4.2).

Lemma 4.3. *There is a constant C such that the following holds. Suppose r, a, n, N, N' are positive integers satisfying $r \geq 2a$, $n \geq (r/a)^2$ and $N \geq 2^{C n a^2 r^{-1} \log(r/a)} N'$. If the vertical edges of $G := G_{N \times r}$ are colored red and the horizontal edges of G are colored red or blue, then G contains either a blue K_n or a copy of $G_{N' \times a}$ where all vertices in some column are only incident to red edges.*

Proof. Suppose that G contains no blue K_n . Let $T = 2^{C_0 n a^2 r^{-1} \log(r/a)}$, so that, by Lemma 4.2, $T \geq R(n; r, r-a)$. Let $C = \max\{1, 3C_0\}$. We claim that any $T \times r$ subgrid of G contains a red $G_{2 \times a}$. Indeed, let G' be a $T \times r$ subgrid of G and define an edge-coloring χ on the complete graph K_T whose vertices are the columns of G' which colors each edge (x, x') by the set of all y for which the edge $(x, y) \sim (x', y)$ is blue. If there is no copy of a red $G_{2 \times a}$ in G' , the edge-coloring χ assigns at least $r-a$ colors to every edge, so, by the definition of T , we obtain a monochromatic K_n in some color in χ . But this implies that we have a monochromatic blue K_n in the original graph G , a contradiction.

Define another auxiliary graph H whose vertices are the columns of G and edges are pairs of columns containing a red $G_{2 \times a}$. Color each edge (x, x') of H by a set of a y -coordinates y_1, \dots, y_a such that the induced subgraph $G[\{x, x'\} \times \{y_1, \dots, y_a\}]$ forms a monochromatic red $G_{2 \times a}$. We know that among every T vertices of H , there is at least one edge. Hence, by Turán's theorem, there are at least $N^2/2(T-1) - N/2 \geq N^2/T^2$

edges in H , at least $N^2/T^2 \binom{r}{a}$ of which receive the same color. Thus, there must be a vertex with at least $N/T^2 \binom{r}{a} \geq N/T^3 \geq N'$ neighbors in a single color. But this corresponds exactly to a copy of $G_{N' \times a}$ with a column incident to only red edges, completing the proof. \square

It remains to iterate the above lemma to obtain (4.2).

Proof of (4.2). To begin, let $a = r_1 < r_2 < \dots < r_b$ and $1 = N_1 < N_2 < \dots < N_b$ be positive integers satisfying $2 \leq r_{i+1}/r_i \leq \sqrt{n}$ and

$$N_{i+1} = 2^{Cnr_i^2 r_{i+1}^{-1} \log(r_{i+1}/r_i)} (N_i + 1)$$

for every $1 \leq i \leq b-1$, where C is the constant from Lemma 4.3. Let $N := n^{r_b^2} \cdot N_b$. We claim that $\text{gr}(G_{b \times a}, K_n) \leq N$.

Indeed, suppose we are given a red-blue edge-coloring of $G_{N \times N}$. We may further suppose that there is no blue K_n . Let H be the induced (complete) subgraph on some single column of $G_{N \times N}$, together with the induced edge-coloring. Since $r(K_{r_b}, K_n) \leq n^{r_b}$, any set of n^{r_b} vertices in H contains at least one red K_{r_b} . It follows that the number of pairs (c, U) where $U \in \binom{V(H)}{n^{r_b}}$ and c is a monochromatic red K_{r_b} in U is at least $\binom{N}{n^{r_b}}$. On the other hand, each such c lies in at most $\binom{N-r_b}{n^{r_b}-r_b}$ choices of U , so we find that there are at least

$$\binom{N}{n^{r_b}} / \binom{N-r_b}{n^{r_b}-r_b} = \binom{N}{r_b} / \binom{n^{r_b}}{r_b}$$

monochromatic red copies of K_{r_b} in H . Moreover, the same conclusion holds for any given column of the grid. Hence, by the pigeonhole principle, there are some r_b rows such that a $\binom{n^{r_b}}{r_b}^{-1}$ -fraction of the N columns are monochromatic red in between these rows. Since

$$N / \binom{n^{r_b}}{r_b} \geq N / n^{r_b^2} = N_b,$$

we obtain an induced subgrid of dimensions $N_b \times r_b$ with monochromatic red columns. Applying Lemma 4.3, we find inside this subgrid another subgrid of dimensions $(N_{b-1} + 1) \times r_{b-1}$ where one column is complete in red to the others. Iterating this process $b-1$ times and setting the distinguished column aside at each step, we find a monochromatic $G_{b \times a}$, as desired.

In order to obtain (4.2), for $1 \leq i \leq b$, we choose $r_i = a \cdot x^{2^{i-1}-1}$ with $x = n^{(2^b-1)^{-1}}$. One can check that these choices imply $2 \leq r_{i+1}/r_i \leq \sqrt{n}$ and $N_{i+1} = 2^{O_a(nx^{-1} \log n)} N_i$ for all $1 \leq i \leq b-1$. Consequently, $N_b = 2^{O_a(nx^{-1} \log n)}$ and, as $r_b^2 = a^2 n/x$, we also have $N = 2^{O_a(nx^{-1} \log n)}$. Since, by symmetry, $\text{gr}(G_{a \times b}, K_n) = \text{gr}(G_{b \times a}, K_n)$, this yields the required bound. \square

5 Concluding Remarks

The main problem left open by this paper is what the true bounds are for $r(K_4^{(3)}, S_n^{(3)})$ and the closely related function $\text{gr}(G_{2 \times 2}, K_n)$. In particular, we have the following question.

Problem 5.1. *Does there exist $c > 0$ such that $\text{gr}(G_{2 \times 2}, K_n) \geq 2^{n^c}$?*

We will not hazard a guess on which direction the truth should lie, though it would be much more interesting were the answer to turn out negative.

Theorem 1.4 gives a sub-exponential bound for grid Ramsey numbers of the form

$$\text{gr}(G_{a \times b}, K_n) \leq 2^{c' n^{1-(2^b-1)^{-1}} \log n} \tag{5.1}$$

when $a \geq b \geq 2$. The dependence of the exponent of n on b , which comes from iterating Lemma 4.2, is inverse exponential. The authors suggested in [5, Problem 6.1] that stronger bounds than Lemma 4.2 might be true for the set-coloring Ramsey number $R(n; r, s)$, especially in the regime $s \approx r - \sqrt{r}$ where we are applying it here. Such improved upper bounds on $R(n; r, s)$ would immediately improve the dependence on b in (5.1).

Note added. After the first version of this paper appeared on arXiv, Aragão, Collares, Marciano, Martins and Morris [3] showed that the bound in Lemma 4.2 is tight up to a logarithmic factor when $s \approx r - \sqrt{r}$, thereby ruling out any hope of substantially improving (5.1) by the route suggested above.

Acknowledgements. This research was initiated during a visit to the American Institute of Mathematics under their SQuaREs program. We would like to thank Noga Alon, Matija Bucić, Benjamin Gunby and Yuval Wigderson for helpful conversations.

References

- [1] M. Ajtai, J. Komlós and E. Szemerédi, A note on Ramsey numbers, *J. Combin. Theory Ser. A* **29** (1980), 354–360.
- [2] N. Alon and J. H. Spencer, **The Probabilistic Method**, Fourth edition, Wiley Series in Discrete Mathematics and Optimization, John Wiley & Sons, Inc., Hoboken, NJ, 2016.
- [3] L. Aragão, M. Collares, J. P. Marciano, T. Martins and R. Morris, A lower bound for set-colouring Ramsey numbers, preprint available at arXiv:2212.06802 [math.CO].
- [4] T. Bohman, D. Mubayi and M. Picollelli, The independent neighborhoods process, *Israel J. Math.* **214** (2016), 333–357.
- [5] D. Conlon, J. Fox, X. He, D. Mubayi, A. Suk and J. Verstraëte, Set-coloring Ramsey numbers via codes, preprint available at arXiv:2206.11371.
- [6] D. Conlon, J. Fox, C. Lee and B. Sudakov, On the grid Ramsey problem and related questions, *Int. Math. Res. Not.* **2015** (2015), 8052–8084.
- [7] D. Conlon, J. Fox and B. Sudakov, Hypergraph Ramsey numbers, *J. Amer. Math. Soc.* **23** (2010), 247–266.
- [8] D. Conlon, J. Fox and B. Sudakov, An improved bound for the stepping-up lemma, *Discrete Appl. Math.* **161** (2013), 1191–1196.
- [9] D. Conlon, J. Fox and B. Sudakov, Recent developments in graph Ramsey theory, in *Surveys in combinatorics 2015*, 49–118, London Math. Soc. Lecture Note Ser., 424, Cambridge Univ. Press, Cambridge, 2015.
- [10] P. Erdős and A. Hajnal, On Ramsey like theorems. Problems and results, in *Combinatorics (Proc. Conf. Combinatorial Math., Math. Inst., Oxford, 1972)*, 123–140, Inst. Math. Appl., Southend-on-Sea, 1972.
- [11] P. Erdős and R. Rado, Combinatorial theorems on classifications of subsets of a given set, *Proc. London Math. Soc.* **3** (1952), 417–439.

- [12] P. Erdős and A. Szemerédi, On a Ramsey type theorem, *Period. Math. Hungar.* **2** (1972), 295–299.
- [13] J. Fox and X. He, Independent sets in hypergraphs with a forbidden link, *Proc. London Math. Soc.* **123** (2021), 384–409.
- [14] R. L. Graham, B. L. Rothschild and J. H. Spencer, **Ramsey theory**, Second edition, Wiley-Interscience Series in Discrete Mathematics and Optimization, John Wiley & Sons, Inc., New York, 1990.
- [15] D. Mubayi and A. Razborov, Polynomial to exponential transition in Ramsey theory, *Proc. London Math. Soc.* **122** (2021), 69–92.
- [16] D. Mubayi and A. Suk, Off-diagonal hypergraph Ramsey numbers, *J. Combin. Theory Ser. B* **125** (2017), 168–177.
- [17] D. Mubayi and A. Suk, New lower bounds for hypergraph Ramsey numbers, *Bull. London Math. Soc.* **50** (2018), 189–201.
- [18] S. Shelah, Primitive recursive bounds for van der Waerden numbers, *J. Amer. Math. Soc.* **1** (1989), 683–697.