Is Proposition 13.6 tight? It's not too hard to see that the answer is no. Indeed, already for $k = 3$, Proposition 13.6 implies that $r(3) > 4$, and it is not hard to show that in fact $r(3) > 5$, as witnessed by the following coloring.

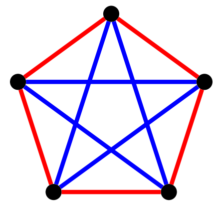

Nonetheless, it is not clear how to do much better than Proposition 13.6 in general. Indeed, in the 1940s, Turán believed that the Erdős–Szekeres bound is way off, and that the truth is $r(k) = \Theta(k^2)$ (i.e. that Proposition 13.6 is best possible up to a constant factor). As it turns out, this belief was *way* off.

Theorem 13.7 (Erdős)**.** *For any $k \geqslant 2$, we have $r(k) \geqslant 2^{k/2}$.*

Together with Theorem 13.4, this proves that $r(k)$ really does grow as an exponential function of k, although these theorems do not tell us the precise growth rate. Theorem 13.7 was a major breakthrough not only—or even primarily—because of the result itself. In proving Theorem 13.7, Erdős introduced the so-called *probabilistic method* to combinatorics. This method would quickly become one of the most important tools in combinatorics, and will recur frequently throughout this course.

Proof of Theorem 13.7. Fix k, and let[6] $N = 2^{k/2}$. The claimed bound is trivial for $k = 2$, so let's assume $k \geqslant 3$. Consider a *random* 2-coloring of $E(K_N)$. Namely, for each edge of K_N, we assign it color red or blue with probability $\frac{1}{2}$, making these choices independently over all edges. We begin by estimating the probability that this coloring contains a monochromatic K_k.

For any fixed set of k vertices, the probability that it forms a monochromatic K_k is precisely $2^{1-\binom{k}{2}}$. This is because we have $\binom{k}{2}$ coin tosses, which we need to all agree, and we have two options for the shared outcome (hence the extra $+1$ in the exponent). Moreover, there are exactly $\binom{N}{k}$ possible k-sets we need to consider. Therefore,

$$\Pr(\text{there is a monochromatic } K_k) \leqslant \binom{N}{k} 2^{1-\binom{k}{2}},$$

where we have applied the *union bound* $\binom{N}{k}$ times; this is the bound that says that the probability that A *or* B happens is at most the sum of the probability that A happens and the probability that B happens.

[6] The astute reader will notice that $2^{k/2}$ is not an integer unless k is even. Thus, we should really write here $N = \lceil 2^{k/2} \rceil$. However, once the computations we do become more complicated, keeping track of such floor and ceiling signs becomes not just annoying, but actively confusing. Therefore, for the rest of the course, we'll omit floor and ceiling signs unless they are actually crucial, and it will be understood that any quantity that should be an integer but doesn't look like one should be rounded up or down to an integer.

Note that $\binom{N}{k} < N^k/k!$ and that $k! > 2^{1+k/2}$ for all $k \geqslant 3$. Therefore, we have

$$\binom{N}{k} 2^{1-\binom{k}{2}} < \frac{N^k}{k!} \cdot 2^{1-\frac{k^2-k}{2}} < \frac{N^k}{2^{1+\frac{k}{2}}} \cdot 2^{1+\frac{k}{2}-\frac{k^2}{2}} = \left(N \cdot 2^{-\frac{k}{2}}\right)^k = 1, \tag{5}$$

where the final equality is our choice of N.

Putting this all together, we find that in this random coloring, the probability that there is a monochromatic K_k is *strictly less than one*. Therefore, there must exist *some* coloring of $E(K_N)$ with no monochromatic K_k, as if such a coloring did not exist, the probability above would be exactly one. This completes the proof. □

It's worth stressing the miraculous magic trick that takes place in the proof of Theorem 13.7. Unlike in Proposition 13.6, Erdős does not give any sort of explicit description of a coloring on $2^{k/2}$ vertices with no monochromatic K_k. Instead, he argues that such a coloring must exist for probabilistic reasons, but this argument gives absolutely no indication of what such a coloring looks like. In fact, the following remains a major open problem.

Open problem 13.8 (Erdős). *For some $\varepsilon > 0$ and all sufficiently large k, explicitly construct a 2-coloring on $(1+\varepsilon)^k$ vertices with no monochromatic K_k.*

There was a great deal of partial progress over the years, much of it exploiting a deep and surprising connection to the topic of *randomness extraction* in theoretical computer science. Very recently, there was a major breakthrough on this problem.

Theorem 13.9 (Li). *For some absolute constant $\varepsilon > 0$ and all sufficiently large k, there is an explicit 2-coloring on 2^{k^ε} vertices with no monochromatic K_k.*

The central open problem in Ramsey theory is to narrow the gap between the lower and upper bounds $2^{k/2} \leqslant r(k) \leqslant 4^k$. For over 75 years, there was a great deal of interest in this question, and while there were several important developments, none of them were able to improve either of the constants in the bases of the exponents. But very recently, there was a huge breakthrough on this problem.

Theorem 13.10 (Campos–Griffiths–Morris–Sahasrabudhe). *There is an absolute constant $\delta > 0$ such that $r(k) \leqslant (4-\delta)^k$ for all k.*

Their original proof showed roughly that $r(k) \leqslant 3.993^k$. A later result by Gupta, Ndiaye, Norin, and Wei, which both optimized the original technique and introduced beautiful new ideas, shows that $r(k) \leqslant 3.8^k$, which remains the current record. The proof of Theorem 13.10 is far too complex to cover in this course, but I have written an exposition of it that you can find on my website.

Additionally, just *yesterday*, there was another breakthrough, due to Ma, Shen, and Xie, this time on the lower bound. Erdős's random argument naturally extends to the off-diagonal setting, and proves that

$$r(k, Ck) \geqslant (f(C) + o(1))^k,$$

for some explicit function f, where we think of $C \geqslant 1$ as an absolute constant and let $k \to \infty$. Ma, Shen, and Xie improved this bound for any $C > 1$, proving that for any $C > 1$, there exists some $\varepsilon > 0$ such that

$$r(k, Ck) \geqslant (f(C) + \varepsilon + o(1))^k.$$

They also use the probabilistic method, but consider a different probability distribution coming from high-dimensional geometry. Unfortunately, their technique does not, at the moment, give any improvement on the *diagonal* Ramsey number $r(k)$.

14 Hypergraph Ramsey numbers

14.1 The hypergraph Ramsey theorem

We saw in the proof of the Erdős–Stone–Simonovits theorem that it can be very useful to study hypergraph analogues of graph-theoretic resutls. In the present context, there is a natural analogue of Ramsey's theorem for hypergraphs, which was also proved by Ramsey. Recall that $K_N^{(t)}$ denotes the complete t-uniform hypergraph on N vertices.

Theorem 14.1 (Ramsey). *For all integers $k \geqslant t \geqslant 2, q \geqslant 2$, there exists some N such that the following holds. In any q-coloring $\chi : E(K_N^{(t)}) \to [\![q]\!]$, there is a monochromatic copy of $K_k^{(t)}$. In other words, there exist k vertices such that each of the $\binom{k}{t}$ t-tuples among them receive the same color under χ.*

Continuing our earlier practice, we define the *t-uniform Ramsey number* $r_t(k; q)$ to be the least N for which Theorem 14.1 is true, and we use the shorthand $r_t(k)$ when $q = 2$. We also define the *off-diagonal t-uniform Ramsey number* $r_t(k_1, \ldots, k_q)$ to be the least N so that in any q-coloring of $E(K_N^{(t)})$, there is a monochromatic copy of $K_{k_i}^{(t)}$ in color i, for some $i \in [\![q]\!]$. Similarly, for any t-uniform hypergraphs $\mathcal{H}_1, \ldots, \mathcal{H}_q$, we denote by $r_t(\mathcal{H}_1, \ldots, \mathcal{H}_q)$ the least N such that any q-coloring of $K_N^{(t)}$ contains a monochromatic copy of $\mathcal{H}_i$ in color i, for some $i \in [\![q]\!]$, and write $r_t(\mathcal{H}; q)$ for shorthand if $\mathcal{H}_1 = \cdots = \mathcal{H}_q = \mathcal{H}$.

Probably the most natural way to prove Theorem 14.1 is via the following argument, directly mimicking the proof of Theorem 13.4.

Proof of Theorem 14.1. Let us only deal with the case $q = 2$. We prove by induction on t the statement that $r_t(k, \ell)$ exists for all $k, \ell \geqslant t$, and for any fixed t we prove this statement by induction on $k + \ell$. Note that the base case $t = 2$ is already done by Theorem 13.1, so we fix some $t \geqslant 3$ and assume the result has been proved for $t - 1$. For this fixed t, the base case $k = t$ or $\ell = t$ is trivial, so we may assume the result has been proved for the pairs $(k - 1, \ell)$ and $(k, \ell - 1)$.

The key claim is that the following recursive bound holds, analogously to (4):

$$r_t(k, \ell) \leqslant r_{t-1}(r_t(k - 1, \ell), r_t(k, \ell - 1)) + 1. \tag{6}$$

Note that we are done if we prove (6), since by induction, we know that the numbers $a \coloneqq r_t(k-1,\ell)$ and $b \coloneqq r_t(k,\ell-1)$ are finite, as is the number $r_{t-1}(a,b)$. Thus, (6) implies Theorem 14.1, at least in the case $q=2$.

To prove (6), let $N = r_{t-1}(r_t(k-1,\ell), r_t(k,\ell-1)) + 1$, and consider any 2-coloring $\chi : E(K_N^{(t)}) \to \{\text{red}, \text{blue}\}$. Fix a vertex $v \in V(K_N^{(t)})$. There is a bijection between hyperedges containing v and $(t-1)$-tuples of vertices in $V(K_N^{(t)}) \setminus \{v\}$. That is, we can use χ to define a coloring $\psi : E(K_{N-1}^{(t-1)}) \to \{\text{red}, \text{blue}\}$, by setting

$$\psi(\{w_1, \ldots, w_{t-1}\}) \coloneqq \chi(\{w_1, \ldots, w_{t-1}, v\}).$$

By the definition of N, we know that ψ contains a monochromatic red clique of order $r_t(k-1,\ell)$, or a monochromatic blue clique of order $r_t(k,\ell-1)$. The two cases are symmetric, so let us assume we are in the first. Looking at χ on these $r_t(k-1,\ell)$ vertices, we can either find a monochromatic blue $K_\ell^{(t)}$, or a monochromatic red $K_{k-1}^{(t)}$. In the first case we are done. In the second case, we have $k-1$ vertices, such that each of the t-tuples among them are colored red. Moreover, by the definition of ψ, if we combine any $(t-1)$-tuple from this set with v, we obtain another t-tuple that is colored red by χ. That is, we have found a monochromatic red $K_k^{(t)}$, showing that we are done in this case as well. □

Remark. While this proof is clearly reminiscient of the proof of Theorem 13.4, you might think that some things are different. For example, (6) is a bit different from (4), in that the former has this strange r_{t-1} term, whereas the latter simply has a sum. It is worth pondering what a 1-*uniform hypergraph* should be, and what the 1-uniform version of Theorem 14.1 should say. If you think about this enough, you'll come to realize that the proof above really is nothing more than a generalization of the proof of Theorem 13.4.

14.2 A geometric application

The paper of Erdős and Szekeres in which they proved Theorem 13.4—one of the most influential and foundational papers in the field—was titled "A combinatorial problem in geometry". We will now study this geometric problem, and see how it relates to Ramsey theory.

Definition 14.2. Let $p_1, \ldots, p_k$ be points in $\mathbb{R}^d$. A point $p \in \mathbb{R}^d$ is *in their convex hull* if there exist numbers $\lambda_1, \ldots, \lambda_k \geqslant 0$ with $\sum_{i=1}^k \lambda_i = 1$ such that

$$p = \sum_{i=1}^k \lambda_i p_i.$$

That is, p is in the convex hull of $p_1, \ldots, p_k$ if p is a weighted average of them.

Definition 14.3. A collection $p_1, \ldots, p_k$ of points in $\mathbb{R}^d$ is *in convex position* if no p_i is in the convex hull of $p_1, \ldots, p_{i-1}, p_{i+1}, \ldots, p_k$.

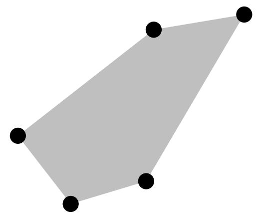

Five points in convex position
(the gray region is their convex hull)

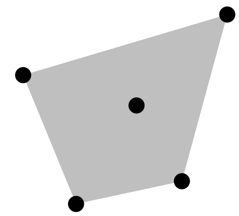

Five points *not* in convex position
(the gray region is their convex hull)

The question studied by Erdős and Szekeres begins with a simple observation of Klein.

Proposition 14.4 (Klein)**.** *Among any five points in* $\mathbb{R}^2$*, no three of them collinear, there are four points in convex position.*

Proof. Consider the convex hull of the five points. It is a polygon with at most five vertices. If it has four or five vertices, then four of these vertices yield our four desired points in convex position. So we may assume that the convex hull is a triangle, meaning that the final two points lie inside the triangle, as shown in the following picture.

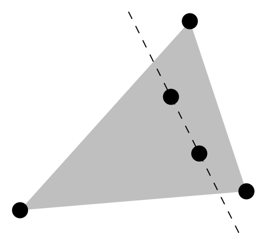

Consider the line through the two interior points. Since no three points are collinear, two of the vertices of the triangle must lie on one side of this line. But then these two vertices, plus the two interior points, yield four points in convex position. □

Although this was before Ramsey theory really existed, Klein realized that there was a Ramsey-theoretic flavor to this result. She asked Erdős and Szekeres whether Proposition 14.4 could be generalized to finding arbitrarily large collections of points in convex position. Erdős and Szekeres proved that the answer is yes.

Theorem 14.5 (Erdős–Szekeres)**.** *For every* $k \geqslant 4$*, there exists some* N *such that the following holds. Among any* N *points in* $\mathbb{R}^2$*, no three of them collinear, there are* k *points in convex position.*

We will see two proofs of this theorem (and another proof is in the homework); the first is the original proof of Erdős and Szekeres.

Erdős and Szekeres's proof of Theorem 14.5. We will show that the theorem holds with $N = r_4(5, k)$. Fix N points $p_1, \ldots, p_N$ in $\mathbb{R}^2$, no three of them collinear. We identify $V(K_N^{(4)})$ with $\{p_1, \ldots, p_N\}$, and define a two-coloring of $E(K_N^{(4)})$ as follows. Given a 4-tuple $\{p_a, p_b, p_c, p_d\}$, we color it blue if these four points are in convex position, and red otherwise.

The first observation is that we cannot have a monochromatic red $K_5^{(4)}$. Indeed, this would correspond to five points in the plane, no three collinear, such that *every* 4-tuple among them is not in convex position. Proposition 14.4 says that such a configuration cannot exist.

Therefore, by the choice of N, there must exist k points, say $p_1, \ldots, p_k$, such that each hyperedge among them is colored blue. That is, every 4-tuple among them *is* in convex position. To complete the proof, we require the following simple lemma.

Lemma 14.6 (Carathéodory's theorem)**.** *Let $p_1, \ldots, p_k$ be a collection of points in $\mathbb{R}^2$, such that each 4-tuple among them is in convex position. Then $p_1, \ldots, p_k$ are in convex position.*

In a moment, we will give a formal proof of Lemma 14.6, but the intuitive proof is the following. Suppose for contradiction that $p_1, \ldots, p_k$ are not in convex position, and say without loss of generality that p_k is in the convex hull of $p_1, \ldots, p_{k-1}$, and call this convex hull P. Then P is a convex polygon, whose vertices are (some subset of) $p_1, \ldots, p_{k-1}$. Pick an arbitrary triangulation of P, that is, a partition of P into triangles whose vertices are vertices of P itself. Since $p_k \in P$, we must have that p_k is contained in one of the triangles of the triangulation. But that means that p_k is in the convex hull of three vertices of P; this yields four points out of $p_1, \ldots, p_k$ which are not in convex position.

Given Lemma 14.6, the proof is complete: we have found k points from our original collection that are in convex position. □

While the geometric proof sketch presented above can be made rigorous, there is also a fairly simple linear-algebraic proof of Lemma 14.6, which we now present.

Proof of Lemma 14.6. We may assume that $k \geqslant 5$, for otherwise there is nothing to prove. Suppose for contradiction that one of the points, say p_k, is in the convex hull of of the remaining points. This means that there exist numbers $\lambda_1, \ldots, \lambda_{k-1} \geqslant 0$ with $\sum \lambda_i = 1$ and

$$p_k = \sum_{i=1}^{k-1} \lambda_i p_i.$$

Let us fix such a collection $\lambda_1, \ldots, \lambda_{k-1}$ with the fewest number of non-zero elements. That is, we may assume by renaming the points that $\lambda_1, \ldots, \lambda_t > 0$, that $\lambda_{t+1}, \ldots, \lambda_{k-1} = 0$, and that no such representation is possible with fewer than t non-zero coefficients.

If $t \leqslant 3$, then we have shown that the points p_1, p_2, p_3, p_k are not in convex position (since p_k is in the convex hull of p_1, p_2, p_3), contradicting our assumption that all 4-tuples are in convex position. Therefore we may assume that $t \geqslant 4$. Consider the vectors

$$v_1 \coloneqq p_1 - p_t, \qquad v_2 \coloneqq p_2 - p_t, \qquad \ldots, \qquad v_{t-1} \coloneqq p_{t-1} - p_t.$$

These are $t - 1 \geqslant 3$ vectors in $\mathbb{R}^2$, so they must be linearly dependent. That is, there exist $\alpha_1, \ldots, \alpha_{t-1} \in \mathbb{R}$, at least one of which is non-zero, such that $\sum_{i=1}^{t-1} \alpha_i v_i = 0$. Now note that for any $\varepsilon \geqslant 0$, we have

$$\begin{aligned} p_k &= \sum_{i=1}^{t} \lambda_i p_i \\ &= \lambda_t p_t + \sum_{i=1}^{t-1} \lambda_i p_i + \sum_{i=1}^{t-1} \varepsilon \alpha_i v_i \\ &= \lambda_t p_t + \sum_{i=1}^{t-1} [(\lambda_i + \varepsilon \alpha_i) p_i - \varepsilon \alpha_i p_t] \\ &= \sum_{i=1}^{t-1} (\lambda_i + \varepsilon \alpha_i) p_i + \left(\lambda_t - \varepsilon \sum_{i=1}^{t-1} \alpha_i \right) p_t \\ &=: \sum_{i=1}^{t-1} \mu_i(\varepsilon) p_i + \mu_t(\varepsilon) p_t. \end{aligned}$$

Notice that each $\mu_i(\varepsilon)$ is a continuous (in fact, linear) function of ε. Also, by assumption, we have that $\mu_i(0) > 0$ for all $i \in [\![t]\!]$. Also, by construction, we have that $\sum_i \mu_i(\varepsilon) = 1$ for all ε. However, since one of the α_i is non-zero, we see that in the limit $\varepsilon \to \infty$, at least one of the $\mu_i(\varepsilon)$ must become negative. Therefore, there is some smallest value ε^* such that $\mu_i(\varepsilon^*) = 0$ for at least one i, and $\mu_j(\varepsilon^*) \geqslant 0$ for all $j \neq i$. However, this gives us a new representation of p_k as a convex combination of $p_1, \ldots, p_{k-1}$ with fewer non-zero coefficients, contradicting our choice of $\lambda_1, \ldots, \lambda_{k-1}$. □

An alternative proof of Theorem 14.5 was found by Tarsi, who showed how to obtain the same result by using a diagonal 3-uniform Ramsey theorem, rather than the off-diagonal 4-uniform Ramsey theorem used by Erdős and Szekeres.

Tarsi's proof of Theorem 14.5. Let $N = r_3(k)$, and fix points $p_1, \ldots, p_N$ in $\mathbb{R}^2$. By rotating the plane if necessary, we may assume that all the points $p_1, \ldots, p_N$ have distinct x-coordinates. Let us also relabel them so that they are sorted by x-coordinate, that is, so that p_1 is to the left of p_2, which is to the left of p_3, and so on. We identify $V(K_N^{(3)})$ with $\{p_1, \ldots, p_N\}$, and color $E(K_N^{(3)})$ as follows. For $1 \leqslant i < j < \ell \leqslant N$, we color the hyperedge $\{p_i, p_j, p_\ell\}$ red if p_j lies above the line $p_i p_\ell$, and blue if p_j lies below the line $p_i p_\ell$.

By the choice of N, there is a monochromatic $K_k^{(3)}$, say $p_{i_1}, \ldots, p_{i_k}$, where $i_1 < \cdots < i_k$. Let us suppose this $K_k^{(3)}$ is red. This means that every point in this set lies above the line between its neighbors on the left and right; intuitively, this means that the points need to look like this:

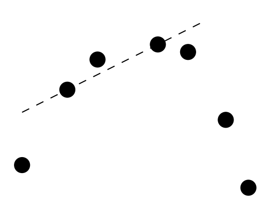

In particular, the points $p_{i_1}, \dots, p_{i_k}$ are in convex position, as is hopefully intuitive from the picture. This is in fact true, and is a discrete version of the well-known fact that a function with non-positive second derivative is concave.

To prove that $p_{i_1}, \dots, p_{i_k}$ are in convex position, it suffices by Lemma 14.6 to show that any four of them are in convex position. So let p_a, p_b, p_c, p_d be four points, ordered from left to right, with the property that each of the triples they define is red, that is, that each point lies above the line connecting its two neighbors. If they are not in convex position, then one of them must be in the convex hull of the other three, and it is not hard to see that the interior point must be either p_b or p_c (p_a and p_d are necessarily extreme points because they minimize and maximize, respectively, the x-coordinate among these four points). If, say, p_b is in the convex hull of p_a, p_c, p_d, then we see that p_b lies below the line between p_a and p_c, a contradiction.

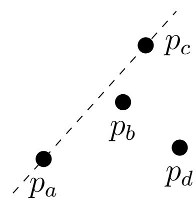

Similarly, if p_c is an interior point, it lies below the line joining p_b, p_d, another contradiction. This shows that all 4-tuples are indeed in convex position, and thus we have found our desired k-set in convex position by Lemma 14.6. In case $\{p_{i_1}, \dots, p_{i_k}\}$ form a blue clique, the same argument works after vertically reflecting the whole picture. □

14.3 Bounds on hypergraph Ramsey numbers

The proof we saw of Theorem 14.1 shows that $r_t(k, \ell)$ is finite for all t, k, ℓ. However, the bound it gives is absolutely enormous. For example, just trying to upper-bound $r_3(k, k)$, we find from (6) that

$$r_3(k) \leqslant r_2(r_3(k-1,k), r_3(k,k-1)) + 1.$$

Plugging in our bound $r_2(a) < 4^a$, this implies that

$$r_3(k) \leqslant 4^{r_3(k-1,k)}.$$

That is, a single step of the recursion has cost us an exponential! Continuing in this way, this proof yields a bound roughly of the form

$$r_3(k) \leqslant 4^{4^{\cdot^{\cdot^{\cdot^{4}}}}} \Big\} {\scriptstyle 2k \text{ times}}.$$

But then the bound in uniformity 4 is then much worse—a single step of the recursion (6) for $t = 4$ shows that $r_4(k)$ is bounded as a tower-type function of $r_4(k-1, k)$. That is, this proof yields a wowzer-type bound on $r_4(k)$, and in general, the bounds it gives for uniformity t are at the $(t-1)$th level of the Ackermann hierarchy.

Are such abysmal bounds necessary? At first glance, one might suspect that they are—exponential bounds really are the truth for $r_2(k)$, so the argument above is not particularly wasteful for uniformity 2. However, Erdős and Rado discovered an alternative proof of Theorem 14.1, which gives a much stronger bound.

Theorem 14.7 (Erdős–Rado)**.** *For all integers $t \geqslant 3, q \geqslant 2$, and $k_1, \ldots, k_q > t$, we have*

$$r_t(k_1, \ldots, k_q) \leqslant q^{1+\binom{r_{t-1}(k_1-1,\ldots,k_q-1)}{t-1}}.$$

In particular,

$$r_t(k; q) \leqslant q^{1+\binom{r_{t-1}(k-1)}{t-1}}.$$

Theorem 14.7 is sometimes called the *stepping-down* argument; it shows that we can bound a t-uniform Ramsey number by (an exponential function of) a $(t-1)$-uniform Ramsey number, that is, we step down one level in the uniformity. As an immediate consequence, we obtain much stronger bounds on hypergraph Ramsey numbers: for any fixed t, the bound is a fixed tower of 2s.

Corollary 14.8. *We have*

$$r_3(k; q) \leqslant 2^{2^{(Cq \log q)k}}$$

for some absolute constant $C > 0$. Similarly,

$$r_4(k; q) \leqslant 2^{2^{2^{(C'q \log q)k}}},$$

and in general,

$$r_t(k; q) \leqslant 2^{2^{\cdot^{\cdot^{2^{(C_t q \log q)k}}}}}\Bigg\} t-1 \text{ twos},$$

for some constant C_t depending only on t.

Additionally, there is a beautiful argument, called the *stepping-up lemma* of Erdős–Hajnal–Rado, which yields nearly matching lower bounds. At a high level, it allows us to convert a lower bound for $r_{t-1}(k/2; q)$ into a lower bound for $r_t(k; q)$ which is *exponentially* larger. In particular, it "should" allow us to close the gap above, by acting in concert with the stepping-down argument Theorem 14.7, as the two yield upper and lower bounds on $r_t(k; q)$ which are exponential in the $(t-1)$-uniform Ramsey number. However, there is an important catch: the stepping-up lemma only works if we start with a construction in uniformity 3 or above.

Theorem 14.9 (Erdős–Hajnal–Rado)**.** *For every $k \geqslant t \geqslant 3, q \geqslant 2$, we have*

$$r_{t+1}(2k + t - 4; q) > 2^{r_t(k;q)-1}.$$

As a corollary, we get a lower bound which "almost" matches Corollary 14.8, but there is a gap of 1 in the height of the tower.

Corollary 14.10. *We have*

$$r_4(k) \geqslant 2^{2^{ck^2}},$$

for some absolute constant $c > 0$. In general, for every $t \geqslant 4$, there is a constant $c_t > 0$ such that

$$r_t(k) \geqslant \left.2^{2^{\cdot^{\cdot^{2^{c_t k^2}}}}}\right\} t-2 \text{ twos}.$$

The most important open problem about hypergraph Ramsey numbers is to close this exponential gap. Note that if one closes this gap for any uniformity $t \geqslant 3$, then one automatically closes it for all higher uniformities, thanks to the stepping-down and stepping-up lemmas, Theorems 14.7 and 14.9. In particular, closing the gap for uniformity 3 would close it for all uniformities. It is generally believed that the upper bound is closer to the truth.

Conjecture 14.11 (Erdős–Hajnal–Rado)**.** *There exists an absolute constant $c > 0$ such that $r_3(k) \geqslant 2^{2^{ck}}$. As a consequence, for every $t \geqslant 3$, there exist constants $c_t, C_t > 0$ such that*

$$t-1 \text{ twos}\left\{2^{2^{\cdot^{\cdot^{2^{c_t k}}}}}\right. \leqslant r_t(k) \leqslant \left.2^{2^{\cdot^{\cdot^{2^{C_t k}}}}}\right\} t-1 \text{ twos}.$$

One important reason to believe this conjecture is that it is known to be true once the number of colors is at least four, via a variant of the stepping-up lemma due to Hajnal.

Theorem 14.12 (Hajnal)**.** *For every $k, q \geqslant 2$, we have*

$$r_3(k; 2q) > 2^{r_2(k-1;q)-1}.$$

In particular,

$$r_3(k; 4) > 2^{2^{ck}}$$

for some absolute constant $c > 0$.