Turandom: Structures in Ramsey theory

Yuval Wigderson

HUJI combinatorics seminar, November 7, 2022

> Notte senza lumicino, gola nera d'un camino,
> son più chiare degli enigmi di Turandot!
>
> Giacomo Puccini, *Turandot*
> (libretto: Giuseppe Adami and Renato Simoni)

Many of the questions I think about can be described, broadly speaking, as questions about structures in Ramsey theory. There are many such structures—and many significant advances in Ramsey theory require coming up with new structures—but two simple, basic structures show up again and again: the Turán coloring and the random coloring. In this talk, I'll describe three major topics[1] in Ramsey theory, and discuss how these two fundamental structures shed light on each of them.

1 First riddle: Classical Ramsey numbers

For a positive integer t, let $r(t)$ denote the least integer N such that, no matter how the edges of K_N are colored in red and blue, there is a monochromatic complete graph on t vertices. The fact that $r(t)$ exists, that is, that some finite N guarantees this property, is the content of Ramsey's theorem from 1930, and thus $r(t)$ is called the *Ramsey number* of t. A central question in Ramsey theory is to understand how $r(t)$ grows as a function of t.

Ramsey's original proof yielded a bound of $r(t) \leq t!$, and he wrote, "I have little doubt that the values for [Ramsey numbers] obtained below are far larger than is necessary." Indeed, a few years later, Ramsey's theorem was reproved by Erdős and Szekeres, who improved the upper bound to $r(t) \leq 4^t$, thus improving a super-exponential bound to an exponential one.

What about lower bounds? In order to prove $r(t) > N$, it suffices to exhibit a two-coloring of $E(K_N)$ with no monochromatic copy of K_t. The easiest way of doing so is to use the *Turán coloring:* We partition the vertices of K_N into $t-1$ equally-sized blocks, color all edges between the blocks blue, and color all edges inside a block red. This guarantees

[1] Corresponding to Turandot's three riddles.

that there is no blue K_t, by the pigeonhole principle: any t-tuple of vertices must have two vertices from the same block, and thus they will span a red edge. In order to not have a red K_t, however, we must ensure that each block has at most $t-1$ vertices, and thus we should take $N = (t-1)^2$; this proves $r(t) > (t-1)^2$.

This coloring is called the Turán coloring because of Turán's theorem, which says that it is the unique coloring on N vertices that maximizes the number of blue edges while not containing a blue K_t. This fact, as well as the difficulty of constructing other colorings which yield a stronger lower bound on $r(t)$, apparently led Turán to believe that this bound is basically tight, and he believed that $r(t) = \Theta(t^2)$.

However, as it turns out, the Turán coloring is extraordinarily far from optimal.

Theorem 1.1 (Erdős 1947). $r(t) > \sqrt{2}^{\,t}$.

Erdős proved this theorem by introducing the second key player in today's talk, the *random coloring.* Indeed, Erdős showed that if $N = \sqrt{2}^{\,t}$ and if one colors the edges of K_N uniformly at random, then with high probability, there will be no monochromatic K_t. Indeed, by the union bound, the probability that some t-set is monochromatic in the random coloring is at most

$$\binom{N}{t} 2^{1-\binom{t}{2}} < \frac{N^t}{t!} \cdot 2^{1+\frac{t}{2}} 2^{-\frac{t^2}{2}} = \frac{2^{1+\frac{t}{2}}}{t!} \cdot \sqrt{2}^{t^2} 2^{-\frac{t^2}{2}} = \frac{2^{1+\frac{t}{2}}}{t!} < 1.$$

Thus, we now know that $r(t)$ grows as an exponential function of t, although the base of the exponent remains very mysterious: since 1947, there has been no improvement to either of the exponential constants $\sqrt{2}$ and 4. It would be a major breakthrough to prove $r(t) > (\sqrt{2}+\varepsilon)^t$ or $r(t) < (4-\varepsilon)^t$ for any positive ε.

However, this is a talk on structures in Ramsey theory, so I want to return some more to the random coloring, which shows that $r(t) > \sqrt{2}^{\,t}$. We see that for this problem, a fully random coloring does *much* better than the deterministic Turán coloring. In fact, while we now know of explicit constructions that do significantly better than the Turán coloring, the following remains a fundamental open problem.

Open problem 1.2 (Erdős). *Can one explicitly construct a coloring on $(1+\delta)^t$ vertices with no monochromatic K_t, for some $\delta > 0$?*

The best known construction, due to Cohen, yields an explicit coloring on $2^{t^{1/(\log\log t)^c}}$ vertices with no monochromatic K_t, for some $c > 0$, and is thus just barely subexponential. The fact that this problem seems so difficult has led many people to wonder if, in some sense, a random coloring is truly "the best" way of coloring the edges of K_N to avoid a monochromatic K_t. One precise version of this question is the following.

Conjecture 1.3 (Sós). *If $N = r(t) - 1$, then any two-coloring of $E(K_N)$ with no monochromatic K_t is quasirandom.*

Here, a coloring of K_N is called *quasirandom* if, for every $S \subseteq V(K_N)$, the number of red edges in S differs from the number of blue edges in S by $o(N^2)$. This definition is due to Chung, Graham, and Wilson, who also proved that it is equivalent to many other natural notions of what it should mean for a coloring to be "random-like".

Sós's conjecture remains open, and it is plausibly very hard; in particular, one may need to really understand the asymptotic behavior of $r(t)$ in order to attack it. This is basically the state of our knowledge about classical Ramsey numbers: the random coloring is the best tool we have—far outstripping known explicit constructions, such as the Turán coloring—but we don't know whether it is literally optimal (i.e. whether $r(t) \approx \sqrt{2}^t$), nor whether it is optimal in some structural sense, namely whether any optimal construction "looks random".

2 Second riddle: Ramsey goodness and books

In the previous section, we quickly stopped talking about the Turán coloring, because Erdős showed that it performs so much worse than the random coloring. But we haven't actually used the Turán coloring to its full potential. To explain what I mean, we need some notation: for graphs H_1, H_2, let $r(H_1, H_2)$ be the minimum N such that every two-coloring of $E(K_N)$ contains a blue copy of H_1 or a red copy of H_2. Thus, $r(t) = r(K_t, K_t)$ in the new notation. As before, the Turán coloring immediately gives us a general lower bound for $r(H_1, H_2)$.

Proposition 2.1 (Chvátal–Harary, Burr). *Let H_1 be a graph with chromatic number $k+1$, and let H_2 be a connected graph on n vertices. Then*

$$r(H_1, H_2) \geq k(n-1) + 1.$$

Proof. Let $N = k(n-1)$, and consider the Turán coloring of K_N, where we partition the vertex set into k blocks, each comprising $n-1$ vertices, and color all internal edges red and all cross-edges blue. Since H_2 is connected and has n vertices, there is no red copy of H_1: it can't "fit" in the blocks, which have only $n-1$ vertices. On the other hand, the blue graph is $(\chi(H_1)-1)$-partite, and so cannot contain any copy of H_1. □

Of course, plugging in $H_1 = H_2 = K_t$, we recover our earlier bound of $r(t) \geq (t-1)^2 + 1$, which is very far from the truth. Amazingly, however, the lower bound in Proposition 2.1 turns out to be *exactly* tight in certain cases. The earliest result of this type is due to Chvátal, who proved that if $H_1 = K_{k+1}$ is a clique and if $H_2 = T_n$ is a tree on n vertices, then $r(K_{k+1}, T_n) = k(n-1)+1$. Following Burr and Erdős, we say that an n-vertex connected graph H is $(k+1)$-*good* if $r(K_{k+1}, H) = k(n-1)+1$. In this language, Chvátal's theorem says that all trees are $(k+1)$-good for all k.

Burr and Erdős began systematically investigating Ramsey goodness, and observed that it seemed to be a very general phenomenon. Namely, they conjectured, based on some partial results, that if k is fixed and if H is a sufficiently large sparse graph, then H should be $(k+1)$-good. In other words, while the Turán coloring yields a very weak bound on $r(H_1, H_2)$ in case both H_1 and H_2 are large dense graphs (e.g. both equal to K_t, as $t \to \infty$), Burr and Erdős conjectured that the Turán coloring should yield a *tight* bound when H_1 is a fixed clique and H_2 is a large, sparse graph.

Burr and Erdős formulated many conjectures along these lines. Their most general conjecture was that if k and Δ are fixed, then any sufficiently large connected graph H with maximum degree Δ is $(k+1)$-good. As it turns out, this conjecture is too optimistic: Brandt used a simple, elegant argument to show that if H is a good expander, then H is not $(k+1)$-good for any $k \geq 2$.

Proposition 2.2 (Brandt 1996)**.** *For any $k \geq 2$ and all sufficiently large Δ and n, there exists a Δ-regular n-vertex graph which is not $(k+1)$-good.*

Proof sketch for $k = 2$. For a fixed $\varepsilon > 0$, let $N = 5(\frac{1}{2} - \varepsilon)n$. We color K_N by letting the blue graph be a balanced blowup of C_5: we split the vertex set into five equally-sized parts $A_1, \ldots, A_5$, color all edges between A_i and $A_{i+1 \pmod 5}$ blue, and all other edges red. There is no blue K_3 in this coloring.

The key claim is that for an appropriately chosen Δ-regular n-vertex graph H, this coloring also contains no red copy of H. Indeed, we choose H to be a good expander, in the sense that there is an edge between any two disjoint vertex subsets of size at least εn. By the pigeonhole principle, if we choose n vertices from K_N, then there will be at least εn vertices in two consecutive blocks $A_i, A_{i+1 \pmod 5}$. Therefore, no matter how we try to embed the vertices of H in K_N, there will be at least one edge colored blue which is supposed to be red. For sufficiently small ε, this shows that

$$r(K_3, H) > N = \frac{5}{2}n - 5\varepsilon n > 2(n-1) + 1,$$

and thus H is not 3-good. □

Since almost every Δ-regular graph is a good expander, Brandt's result implies that there exist many counterexamples to this strong conjecture.

Nevertheless, it turns out that many of the remaining conjectures of Burr and Erdős are true. Moreover, it turns out that in a certain sense, good expanders are the "only non-good graphs". Recall that a family $\mathcal{F}$ of graphs is called *hereditary* if it is closed under taking induced subgraphs. Let's say that $\mathcal{F}$ *has small separators* if there exists some $\varepsilon > 0$ such that for every n-vertex $H \in \mathcal{F}$, we can delete $n^{1-\varepsilon}$ vertices from H so that H breaks into connected components, each of size at most $\frac{2}{3}n$. In some sense, having small separators is the opposite of being a good expander: graphs with small separators can be easily disconnected, whereas it is very difficult to disconnect a good expander.

Theorem 2.3 (Nikiforov–Rousseau 2009)**.** *Let $k \geq 2$ and let $\mathcal{F}$ be a hereditary family with small separators. If n is sufficiently large, then every n-vertex connected $H \in \mathcal{F}$ is $(k+1)$-good, i.e.*

$$r(K_{k+1}, H) = k(n-1) + 1.$$

Using this theorem, Nikiforov and Rousseau were able to resolve a huge number of questions (some raised by Burr and Erdős and some not) about Ramsey goodness. Here are just a few examples.

- All sufficiently large connected planar graphs are $(k+1)$-good for all $k \geq 2$.
- Fix some graph M and some $k \geq 2$. Then every sufficiently large connected graph without M as a minor is $(k+1)$-good.
- If $k \geq 2$ and t is sufficiently large, then the subdivision of K_t is $(k+1)$-good.
- Let Γ_n^d denote the $n \times n \times \cdots \times n$ grid graph in $\mathbb{R}^d$. If d and k are fixed and n is sufficiently large, then Γ_n^d is $(k+1)$-good.

We remark that one final question raised by Burr and Erdős was whether the hypercube graph Q_d is $(k+1)$-good for sufficiently large d. The techniques of Nikiforov and Rousseau were not sufficient to answer this question, but it was eventually proved by Fiz Pontiveros, Griffiths, Morris, Saxton, and Skokan, building on earlier work of Conlon, Fox, Lee, and Sudakov.

In fact, Nikiforov and Rousseau proved a more general result than Theorem 2.3, where they could replace the fixed graph K_{k+1} by a more general graph of chromatic number $k+1$, so long as the graph was "fairly small". I again won't state their result in full generality, but I will state the following corollary. For positive integers $m > k$, the *book graph*[2] $B_m^{(k)}$ consists of a clique K_k, together with $m-k$ common neighbors of the clique, with no other edges between them. Equivalently, one can view $B_m^{(k)}$ as $m-k$ copies of K_{k+1}, glued along a common K_k. The terminology comes from the case $k=2$, where we can visualize $B_m^{(2)}$ as a book with $m-2$ triangular pages. Because of this, the clique K_k is called the *spine* of the book, and the $m-k$ additional vertices are called the *pages*.

Note that $\chi(B_m^{(k)}) = k+1$, and that $B_{k+1}^{(k)} = K_{k+1}$. Because, of this, the following result of Nikiforov and Rousseau generalizes Theorem 2.3, which corresponds to the case $m=1$.

Theorem 2.4 (Nikiforov–Rousseau 2009)**.** *Let $k \geq 2$ and let $\mathcal{F}$ be a hereditary family with small separators. There exists some $c_0 = c_0(k, \mathcal{F}) > 0$ such that if n is sufficiently large, then*

$$r(B_m^{(k)}, H) = k(n-1)+1$$

for all n-vertex connected $H \in \mathcal{F}$, and for all $m \leq c_0 n$.

In other words, the Turán coloring is still optimal even if we are not searching for a blue book $B_m^{(k)}$, rather than a blue clique K_{k+1}, so long as m is not too large relative to n.

The family of books $\{B_m^{(k)}\}$ has small separators for any fixed k, since deleting the spine turns $B_m^{(k)}$ into $m-k$ isolated vertices. Plugging this fact into Theorem 2.4, we get the following corollary.

Corollary 2.5. *Fix $k, \ell \geq 2$. There exists some $c_0 > 0$ such that for all sufficiently large n and all $m \leq c_0 n$, we have*

$$r(B_m^{(k)}, B_n^{(\ell)}) = k(n-1)+1.$$

[2] My notation here is slightly non-standard; usually, what I'm calling $B_m^{(k)}$ would be called $B_{m-k}^{(k)}$.

These remarkable theorems of Nikiforov and Rousseau have one major drawback. In all of them, the proofs use Szemerédi's regularity lemma, and therefore they obtain extremely poor control on the value of c_0 appearing, as well as on the "sufficiently large" condition for n. For example, in Corollary 2.5, their proof yields a constant $c_0 > 0$ such that $1/c_0$ is bounded by a tower-type function of k and ℓ, and similarly only applies once n is at least a tower-type function of k and ℓ.

In recent work with Jacob Fox and Xiaoyu He, we were able to eliminate the use of the regularity lemma from Corollary 2.5, and consequently obtain the following result with much stronger quantitative information.

Theorem 2.6 (Fox–He–W. 2021). *Fix $k, \ell \geq 2$. There exist $c_0 \geq 2^{-2^{\mathrm{poly}(k,\ell)}}$ and $n_0 \leq 2^{2^{\mathrm{poly}(k,\ell)}}$ such that for all $n \geq n_0$ and all $m \leq c_0 n$, we have*

$$r(B_m^{(k)}, B_n^{(\ell)}) = k(n-1)+1.$$

The proof of Theorem 2.6 is fairly complicated. A somewhat simpler result, which captures many of the key ideas, is to prove the result for $m = k+1$, i.e. to prove that $r(K_{k+1}, B_n^{(\ell)}) = k(n-1)+1$. In other words, we wish to prove that $B_n^{(\ell)}$ is $(k+1)$-good.

Proof sketch for $m = k+1$. The key idea in the proof is that to find a red book or a blue clique in a coloring of $E(K_N)$, it suffices to find a small (constant-sized) Turán subcoloring. This is the content of the following lemma.

Lemma 2.7. *Let $N = k(n-1)+1$, and suppose that a coloring of $E(K_N)$ contains disjoint sets $A_1, \ldots, A_k$ of vertices with $|A_i| = \ell$, such that each A_i is monochromatic red and all edges between A_i and A_j for $i \neq j$ are blue. Then the coloring contains a red $B_n^{(\ell)}$ or a blue K_{k+1}.*

Proof. If some vertex $v \notin (A_1 \cup \cdots \cup A_k)$ has a blue neighbor in each A_i, then we immediately find a blue K_{k+1}. So we can assume that for each such $v \notin (A_1 \cup \cdots \cup A_k)$, there exists some i so that v is monochromatic red to A_i. Let B_i be the set of vertices that are monochromatic red to A_i, and note that we have $A_i \subseteq B_i$ since each A_i is a red clique. By the argument above, every vertex of K_N is in at least one B_i. Therefore, there exists some i for which

$$|B_i| \geq \left\lceil \frac{N}{k} \right\rceil = \left\lceil \frac{k(n-1)+1}{k} \right\rceil = n.$$

The vertices of this largest B_i form a red copy of $B_n^{(\ell)}$. □

The main part of the proof involves finding a Turán subcoloring in the sense above, so that we can apply Lemma 2.7. We build up the Turán subcoloring one part at a time, at each step substantially shrinking the sets we've already found. The crucial point is that since we are searching for a *constant-sized* object, losing substantially at each step doesn't really bother us.

We proceed as follows. Fix a coloring of $E(K_N)$, where $N = k(n-1)+1$. If there is no blue K_{k+1}, then there is a large red clique X, by applying Ramsey's theorem. This is a

Turán subcoloring with one part. To convert this to a Turán subcoloring with two parts, consider all vertices outside of X. If many vertices outside of X have at least $(1-\varepsilon)|X|$ red neighbors in X (for some small ε), then it is easy to find a red $B_n^{(\ell)}$ whose spine is in X. So we may assume that we have a large set Z of vertices outside of X, each with at least $\varepsilon|X|$ blue neighbors in X. There are only $\binom{|X|}{\varepsilon|X|}$ possible blue neighborhoods, so by the pigeonhole principle, there is some $Y \subseteq Z$ with $|Y| \geq |Z|/\binom{|X|}{\varepsilon|X|}$ such that all vertices in Y have the same blue neighborhood in X; call this blue neighborhood $X' \subseteq X$. Note that X' is a red clique, and all edges between X' and Y are blue. To convert this to a Turán subcoloring with two parts, we again apply Ramsey's theorem inside Y (using the fact that the coloring has no blue K_{k+1}) to find a large red clique $Y' \subseteq Y$. Then (X', Y') is the desired two-part Turán subcoloring.

We now repeat this same argument to shrink X' and Y' even more, and add a third set to the Turán subcoloring. We can continue in this way until we construct the Turán subcoloring with k parts in the assumption of Lemma 2.7, at which point we are done. $\square$

As indicated, our proof of Theorem 2.6 really uses the structure of book graphs. We actually prove a slightly more general result than Theorem 2.6, where $B_m^{(k)}$ can be replaced by a somewhat more general $(k+1)$-partite graph with $m \leq c_0 n$ vertices, but our proof really seems to require the graph we're searching for in red to be a book $B_n^{(\ell)}$. This leaves open the following natural problem.

Open problem 2.8. *Can one prove the full Ramsey goodness result of Nikiforov and Rousseau, or its consequences Theorems 2.3 and 2.4, without invoking Szemerédi's regularity lemma (and thus obtaining stronger quantitative control)?*

For simplicity, let's set $k = \ell$ in Corollary 2.5 and Theorem 2.6. Then these results tell us that there is some small $c_0 > 0$ such that $r(B_{cn}^{(k)}, B_n^{(k)}) = k(n-1)+1$ for all $c < c_0$ and all sufficiently large n. Imagine we fix k and let n be very large, and then start increasing c from 0 to 1, and ask how $r(B_{cn}^{(k)}, B_n^{(k)})$ changes. These results tell us that for a while, it actually doesn't change at all: it's stuck on the fixed value $k(n-1)+1$, matching the Turán coloring lower bound.

Does this behavior last forever? In other words, is the Turán bound just optimal for all $c \in (0,1]$? It turns out that the answer is no, thanks to our old friend the random coloring. For example, it is easy to check that a random coloring yields

$$r(B_n^{(k)}, B_n^{(k)}) \geq (2^k - o(1))n,$$

since if we randomly color N vertices, then any k-set of vertices will have roughly $2^{-k}N$ common neighbors in either of the two colors. This is much larger than the bound of $(k - o(1))n$ coming from the Turán coloring.

More generally, we can use random colorings to lower bound $r(B_{cn}^{(k)}, B_n^{(k)})$ for any $c \in (0,1]$. A fairly straightforward computation shows that the optimal thing is to color every edge red with probability $p = 1/(c^{1/k}+1)$ and blue with probability $1-p$, which yields the lower bound

$$r(B_{cn}^{(k)}, B_n^{(k)}) \geq \left(c^{1/k}+1\right)^k n - o(n),$$

which beats the random bound once c is sufficiently far from 0.

Thus, the Turán bound is tight for $c < c_0$, and must eventually stop being tight, since it is eventually outstripped by the random bound. Is the random bound tight, or is there an even better coloring? It turns out that there is not.

Theorem 2.9 (Conlon 2019 for $c = 1$, Conlon–Fox–W. 2021+ in full generality)**.**
For every $k \geq 2$, there exists some $c_1 = c_1(k) \in (0, 1]$ such that for all $c \in [c_1, 1]$ and all sufficiently large n, we have

$$r(B_{cn}^{(k)}, B_n^{(k)}) = \left(c^{1/k} + 1\right)^k n + o(n),$$

In other words, the random construction is asymptotically best possible.

Our proof of Theorem 2.9 is, in some sense, similar to the proof of Theorem 2.6: in order to find the monochromatic books, we first find an "approximate Turán sub-coloring". This consists of k vertex subsets, such that a substantial portion of the internal edges are red, a substantial portion of the cross edges are blue, and everything "looks random", where this last point can be made precise using the notion of ε-regularity. Such a structure can be used to build large books, by finding the spine of a red book inside one of the parts, and finding the spine of a blue book between the parts. Crucially, as in the proof of Theorem 2.6, it's OK if this approximate sub-coloring is very small, since we only use it to find the spine.

Here is an idealized version of this argument, which essentially suffices to prove Theorem 2.9 in the case $c = 1$. Suppose that in a coloring of $E(K_N)$ we find disjoint vertex sets $A_1, \ldots, A_k$, and suppose that "everything looks random" in the following sense: each set A_i contains many red K_k, which are "evenly distributed", and there are many blue K_k with one vertex in each A_i, and these are again "evenly distributed". Fix a vertex $v \notin A_1 \cup \cdots \cup A_k$. Let $r_i(v)$ denote the number of red neighbors of v in A_i, and let $x_i(v) = r_i(v)/|A_i|$.

Suppose we perform the following random experiment. First, we pick a color from {red, blue} uniformly at random. If the color is red, we pick an index $i \in [k]$ uniformly at random, and then pick a uniformly random red K_k in A_i. If the color is blue, we pick a uniformly random blue K_k with one vertex in each of $A_1, \ldots, A_k$. We now ask for the probability that v extends this random K_k into a monochromatic K_{k+1}. By our "even distribution" assumption, this probability is roughly

$$\frac{1}{2}\left(\frac{1}{k}\sum_{i=1}^{k} x_i(v)^k + \prod_{i=1}^{k}(1 - x_i(v))\right).$$

Indeed, the probability that all vertices of a red K_k in some A_i are adjacent to v should be roughly $x_i(v)^k$, and the probability that all vertices of a blue K_k between the parts are adjacent to v should be roughly $\prod_{i=1}^{k}(1 - x_i(v))$. We now use a simple analytic fact about real numbers, namely that for all $x_1, \ldots, x_k \in [0, 1]$, we have that

$$\frac{1}{2}\left(\frac{1}{k}\sum_{i=1}^{k} x_i^k + \prod_{i=1}^{k}(1 - x_i)\right) \geq 2^{-k}.$$

Adding up the contributions over all vertices v, we find that in expectation, our random monochromatic K_k has at least $2^{-k}N$ extensions to a monochromatic K_{k+1}. In other words, we can find a monochromatic $B^{(k)}_{2^{-k}N}$. This entire heuristic argument can be made precise with the techniques of ε-regularity, though the technical details rapidly become complicated.

Beyond determining the Ramsey numbers $r(B^{(k)}_{cn}, B^{(k)}_n)$, we are able to prove structural results as well: if $c > c_1$, then all nearly-extremal colorings must be quasirandom, whereas if $c < c_0$, then all nearly-extremal colorings are close to the Turán coloring.

Theorem 2.10 (Conlon–Fox–W. 2021+)**.** *For every $k \geq 2$, there exist $c_0, c_1 \in (0,1]$ such that the following hold for all sufficiently large n.*

- *If $c \leq c_0$, then $r(B^{(k)}_{cn}, B^{(k)}_n) = k(n-1)+1$. Moreover, any two-coloring on $N = kn - o(n)$ vertices with no blue $B^{(k)}_{cn}$ and no red $B^{(k)}_n$ can be turned into the Turán coloring by recoloring $o(N^2)$ edges.*

- *If $c \geq c_1$, then $r(B^{(k)}_{cn}, B^{(k)}_n) = (c^{1/k}+1)^k n + o(n)$. Moreover, any two-coloring on $N = (c^{1/k}+1)^k n + o(n)$ with no blue $B^{(k)}_{cn}$ and no red $B^{(k)}_n$ is quasirandom with red edge density $p = 1/(c^{1/k}+1)$.*

Thus, there are two subintervals of $(0,1]$, one near 0 and one near 1, where the Turán coloring and the random coloring are each asymptotically optimal, and moreover they are the unique optimal structure in these intervals. In particular, one can view the second part of Theorem 2.10 as a book version of Sós's Conjecture 1.3, which said that all the extremal colorings for $r(t)$ are quasirandom.

Additionally, we were able to obtain both upper and lower bounds for $c_1(k)$ of the form $((1+o(1))\log k/k)^k$. In particular, we see that $c_1(k) \to 0$ as $k \to \infty$, meaning that for large k, the random bound is tight for "most" $c \in (0,1]$. Our understanding of $c_0(k)$ is very poor, since we use Szemerédi's regularity lemma to deduce the structural result, meaning that we get an upper bound on $1/c_0(k)$ which is of tower type in k. Plausibly, one could modify the techniques used in Theorem 2.6 to improve this tower-type behavior.

Finally, our results say nothing about the interval (c_0, c_1). For c in this interval, neither the Turán coloring nor the random coloring yield the optimal bound on $r(B^{(k)}_{cn}, B^{(k)}_n)$. In this range, some other structure takes over, which can do better than both the Turán coloring and the random coloring; at the moment, we do not have any sort of conjecture for what such a structure might look like.

3 Third riddle: Ramsey multiplicity

Let H be a fixed t-vertex graph. We know from Ramsey's theorem that if N is sufficiently large, then every two-coloring of $E(K_N)$ contains at least one monochromatic copy of H. However, we can say much more: a simple averaging argument, due to Erdős, shows that any two-coloring of $E(K_N)$ contains at least $(c(H)-o(1))N^t$ labeled monochromatic copies of H, for some constant $c(H) > 0$. This constant is called the *Ramsey multiplicity constant* of H.

The earliest result on Ramsey multiplicity, predating even the definition, is due to Goodman, who showed that $c(K_3) \geq \frac{1}{4}$. In other words, every two-coloring of $E(K_N)$ contains at least $(\frac{1}{4} - o(1))N^3$ labeled monochromatic triangles. This is tight, as shown by a random coloring: each labeled triangle is monochromatic red with probability $2^{-3} = \frac{1}{8}$, and similarly it's monochromatic blue with probability $\frac{1}{8}$, and so the random coloring on N vertices has $(\frac{1}{4} + o(1))N^3$ monochromatic triangles with high probability. Thus, Goodman's result shows that the random coloring asymptotically minimizes the number of monochromatic triangles among all N-vertex colorings.

Similarly, the random coloring has $(2^{1-\binom{t}{2}} + o(1))N^t$ monochromatic copies of K_t, for any t. This led Erdős to conjecture that for every $t \geq 4$, we have $c(K_t) = 2^{1-\binom{t}{2}}$, i.e. that the random coloring also asymptotically minimizes the number of monochromatic cliques of any size. This conjecture was extended by Burr and Rosta to apply to all graphs: they conjectured that the random coloring asymptotically minimizes the number of monochromatic copies of H, for any graph H.

Conjecture 3.1 (Burr–Rosta 1980)**.** *If H has m edges, then $c(H) = 2^{1-m}$.*

Graphs for which the Burr–Rosta conjecture is true are called *common*, and many natural families of graphs are known to be common, including all trees and all cycles. Sidorenko's conjecture, a major open problem, implies that all bipartite graphs are common. In general, there is a rich theory of common graphs, which I won't say much more about.

Given the topic of this talk, you can probably guess what comes next: this conjecture is false, as shown by the Turán coloring, which has fewer copies of H than the random coloring. That is what we'll get to soon, but that's actually not where the original counterexamples came from. The first counterexample to the Burr–Rosta conjecture was due to Sidorenko, who showed that $c(H) < 2^{1-4}$ for $H =$ 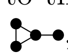, a triangle with a pendant edge. At roughly the same time, Thomason disproved Erdős's original conjecture about cliques, for all $t \geq 4$.

Theorem 3.2 (Thomason 1989)**.** *For every $t \geq 4$, we have $c(K_t) < 0.976 \cdot 2^{1-\binom{t}{2}}$.*

As it turns out, neither Sidorenko nor Thomason used the Turán coloring, and neither did most of the researchers who found new counterexamples to the Burr–Rosta conjecture over the years. Thomason's proof, in particular, used a fairly intricate coloring, coming from certain discrete geometries over $\mathbb{F}_2$, in which he could count monochromatic K_t.

Nonetheless, the Turán coloring is actually quite useful for this problem, as first observed by Fox. Indeed, suppose H is connected, has chromatic number $k+1$, and has t vertices, and consider the Turán coloring of K_N, where we partition the vertex set into k equally-sized blocks, color all edges inside a block red, and all edges between blocks blue. Then there is no blue copy of H since the blue graph is k-partite. Since H is connected, each red copy of H must lie in one of the k parts, which each have size N/k. Thus, the total number of labeled monochromatic copies of H is roughly $k \cdot (N/k)^t = k^{1-t}N^t$, which shows that $c(H) \leq k^{1-t}$.

Using this, Fox showed that the Burr–Rosta conjecture is "very false": there exist m-edge graphs H whose Ramsey multiplicity constant is super-exponentially small in m, rather than the exponential behavior predicted by Burr and Rosta.

Theorem 3.3 (Fox 2007)**.** *There exists a graph H with m edges and*

$$c(H) = 2^{-\Omega(m \log m)}$$

as $m \to \infty$.

Proof. For positive integers $t > k$, form a graph K_{k+1}^{+t} by starting with a clique K_{k+1} and adding $t-k-1$ vertices of degree 1, each adjacent to exactly one vertex of the K_{k+1}. Then K_{k+1}^{+t} has t vertices and chromatic number $k+1$. Moreover, it has $\binom{k+1}{2}+(t-k-1) = \Theta(k^2+t)$ edges. By the discussion above, the Turán coloring of K_N shows that

$$c(K_{k+1}^{+t}) \leq k^{1-t} = 2^{-\Omega(t \log k)}.$$

By setting $k = \Theta(\sqrt{m})$ and $t = \Theta(m)$, we see that K_{k+1}^{+t} has m edges and Ramsey multiplicity constant at most $2^{-\Omega(m \log m)}$. □

In the above proof, the optimal choice of k and t is, as above, $k = \Theta(\sqrt{m})$ and $t = \Theta(m)$. However, for any fixed k and any $t = \Omega(k^2)$, this construction yields a graph with $m = \Theta(t)$ edges and Ramsey multiplicity constant $2^{-\Omega(m \log k)}$. Thus, no matter how fast t grows as a function of k, this construction yields a family of graphs with super-exponential Ramsey multiplicity constant, and in particular yields a family of counterexamples to the Burr–Rosta conjecture. Once t is sufficiently large in terms of k, Fox and I were able to show that the Turán coloring is exactly optimal for this problem.

Theorem 3.4 (Fox–W. 2022+)**.** *For all $k \geq 3$ and $t \geq k^{100k}$, we have $c(K_{k+1}^{+t}) = k^{1-t}$. Moreover, for all sufficiently large N, the Turán coloring is the unique coloring on N vertices with the minimum number of monochromatic copies of K_{k+1}^{+t}.*

Thus, there are common graphs, for which the random coloring asymptotically minimizes the number of monochromatic copies, and also other graphs, like K_{k+1}^{+t}, where the Turán coloring minimizes the number of monochromatic copies; we call such graphs *bonbons*. We are actually able to prove a more general theorem than Theorem 3.4, where K_{k+1} is replaced by any graph H_0 with chromatic number $k+1$ and with the property that there is an edge whose deletion decreases this chromatic number. Given such a graph H_0, one can form H_0^{+t} by adding vertices of degree 1 adjacent to a single vertex of H_0, until there are t vertices total. We can prove that if $t \geq (1000kh)^{10}h^{10k}$, where $h = |V(H_0)|$, then H_0^{+t} is a bonbon. Additionally, one can show that all the assumptions in the theorem are in some sense necessary—H_0^{+t} is *not* a bonbon if $\chi(H_0) \leq 3$, if H_0 does not have an edge whose deletion decreases the chromatic number, or if t is not sufficiently large.

We also conjecture that many other natural graphs are bonbons. For example, we believe that if H is obtained by gluing arbitrary trees to the vertices of H_0, then H is a bonbon so long as it has $t \gg h$ vertices; this generalizes the case of H_0^{+t}, where we glued stars to a the vertices of H_0.

Theorem 3.4 has a fairly long and intricate proof, but the basic idea is to slowly "discover" the Turán coloring. Given a coloring of $E(K_N)$ with the minimum number of copies of K_{k+1}^{+t}, we iteratively prove structural lemmas, each of which shows that the coloring is approximately like a Turán coloring. For example, we first prove that (potentially after inverting the colors) almost all vertices have red degree not much larger than N/k. We then show that there are very few blue copies of K_{k+1}. Using these, we are able to find an approximate Turán partition, namely a partition of the vertex set into k parts of roughly equal sizes, so that almost all edges inside a part are red and almost all edges between parts are blue. Finally, we are able to remove both of these "almosts" and prove that the coloring really is the Turán coloring.

As in the case of book Ramsey numbers, there is some "intermediate regime" where our understanding is very limited. Namely, for certain graphs, such as K_t for $t \geq 4$, we know that neither the Turán coloring nor the random coloring asymptotically minimizes the number of monochromatic copies. In such cases, no one really even knows what to conjecture; there is some mysterious structure, neither Turán nor random, which is asymptotically optimal, but we don't know much about it.

To conclude this section, let me just mention one connection between book Ramsey numbers and Ramsey multiplicity. Recall that Conlon proved in Theorem 2.9 that $r(B_n^{(k)}, B_n^{(k)}) = 2^k n + o(n)$. This implies that in any coloring of K_N, there is some monochromatic K_k which lies in at least $(2^{-k} - o(1))N$ monochromatic K_{k+1}, i.e. in at least as many monochromatic K_{k+1} as it would lie in in a random coloring. This can be viewed as a "local" version of Erdős's conjecture that $c(K_t) = 2^{1-\binom{t}{2}}$: the random coloring asymptotically minimizes the number of monochromatic K_{k+1} which a worst-case K_k lies in. Thus, while Erdős's original conjecture is false, its local version is true. Moreover, our structural result, Theorem 2.10, actually shows that something stronger is true: if a coloring has *fewer* monochromatic K_{k+1} than a random coloring, then some K_k lies in *more* monochromatic K_{k+1} than it would in a random coloring.

Theorem 3.5 (Conlon–Fox–W. 2020). *For every $k \geq 2$ and $\varepsilon > 0$, there exists some $\delta > 0$ such that the following holds for sufficiently large N. If a two-coloring of $E(K_N)$ contains at most $(2^{1-\binom{k+1}{2}} - \varepsilon)N^{k+1}$ monochromatic labeled copies of K_{k+1}, then some monochromatic copy of K_k lies in at least $(2^{-k} + \delta)N$ monochromatic copies of K_{k+1}.*

4 Bonus riddle: deleting a vertex

Let me end with one further riddle, which isn't directly related to the main story of the riddles above, but which does again demonstrate the utility of the Turán coloring for various Ramsey-theoretic problems. We are interested in the following conjecture.

Conjecture 4.1 (Conlon–Fox–Sudakov 2020)**.** *Let G be a graph, and let H be obtained from G by deleting a vertex. Then*

$$r(G) \le C \cdot r(H)$$

for some absolute constant $C > 0$.

The original motivation for this conjecture came from the study of Ramsey numbers of random graphs, as it would imply a concentration result by Azuma's inequality. But it's an interesting question in its own right, simply because it is interesting to study how natural quantities like the Ramsey number change under natural operations such as vertex deletion. Additionally, it's a natural conjecture because it's true in a certain "average" sense: if we imagine building G up one vertex at a time, as $G_1, G_2, \dots, G_n = G$, then we have that $r(G_1) = 1$ and $r(G_n) \le 4^n$. So for an average value of i, we have that $r(G_{i+1}) \le 4 \cdot r(G_i)$.

Conlon, Fox, and Sudakov proved a weakening of Conjecture 4.1, namely that it holds if one replaces the constant C by $2|V(G)|$. They also proved that Conjecture 4.1 holds in an important special case, namely that of dense graphs: if G has n vertices and at least δn^2 edges, then $r(G) \le C(\delta) \cdot r(H)$. This special case should, in some sense, be the hardest one, because an n-vertex graph with $o(n^2)$ edges has Ramsey number $2^{o(n)}$, so the average-case argument above suggests that an even stronger result than Conjecture 4.1 should be true for sparse graphs.

Despite all of this, the conjecture is false.

Theorem 4.2 (W. 2022+)**.** *For all n, there exists an $(n+1)$-vertex graph G with $r(G) = \Omega(n \log n)$, but by deleting a vertex from G, we obtain a graph H with $r(H) = n$.*

Proof. Let $H_{k,n}$ consist of a copy of K_k plus $n-k$ additional isolated vertices. Form $G_{k,n}$ from $H_{k,n}$ by adding a single vertex adjacent to all vertices of $H_{k,n}$, so that $G_{k,n}$ consists of a copy of K_{k+1} plus $n-k$ additional vertices adjacent to one vertex of the clique. The theorem now follows from the following two claims:

(1) $r(G_{k,n}) \ge kn + 1$.

(2) If $n \ge 4^k$, then $r(H_{k,n}) = n$.

These two claims imply the theorem, since we may set $k = \lfloor \frac{1}{2} \log_2 n \rfloor$ and set $G = G_{k,n}$, so that $r(G) \ge kn + 1 = \Omega(n \log n)$ and $r(H) = n$.

Claim (2) is straightforward. The lower bound $r(H_{k,n}) \ge n$ is immediate since $H_{k,n}$ has n vertices. For the upper bound, consider a coloring of $E(K_n)$, where $n \ge 4^k$. As $r(K_k) \le 4^k \le n$, there is a monochromatic copy of K_k; together with the remaining $n-k$ vertices, we obtain a monochromatic copy of H.

For claim (1), we use the Turán coloring. Namely, since $G_{k,n}$ is a connected graph with chromatic number $k+1$ and with $n+1$ vertices, Proposition 2.1 tells us that

$$r(G) \ge kn + 1,$$

as claimed. □