

The Erdős-Rothschild problem on edge-colourings with forbidden monochromatic cliques

BY OLEG PIKHURKO[†] AND KATHERINE STADEN[‡]

Mathematics Institute and DIMAP, University of Warwick, Coventry, CV4 7AL

e-mail: {o.pikhurko,k.l.staden}@warwick.ac.uk

ZELEALEM B. YILMA

Carnegie Mellon University Qatar, Doha, Qatar

e-mail: zyilma@qatar.cmu.edu

(Received 18 May 2016)

Abstract

Let $\mathbf{k} := (k_1, \dots, k_s)$ be a sequence of natural numbers. For a graph G , let $F(G; \mathbf{k})$ denote the number of colourings of the edges of G with colours $1, \dots, s$ such that, for every $c \in \{1, \dots, s\}$, the edges of colour c contain no clique of order k_c . Write $F(n; \mathbf{k})$ to denote the maximum of $F(G; \mathbf{k})$ over all graphs G on n vertices. This problem was first considered by Erdős and Rothschild in 1974, but it has been solved only for a very small number of non-trivial cases.

We prove that, for every \mathbf{k} and n , there is a complete multipartite graph G on n vertices with $F(G; \mathbf{k}) = F(n; \mathbf{k})$. Also, for every \mathbf{k} we construct a finite optimisation problem whose maximum is equal to the limit of $\log_2 F(n; \mathbf{k}) / \binom{n}{2}$ as n tends to infinity. Our final result is a stability theorem for complete multipartite graphs G , describing the asymptotic structure of such G with $F(G; \mathbf{k}) = F(n; \mathbf{k}) \cdot 2^{o(n^2)}$ in terms of solutions to the optimisation problem.

1. Introduction and results

Let a sequence $\mathbf{k} = (k_1, \dots, k_s) \in \mathbb{N}^s$ of natural numbers be given. By an *s-edge-colouring* (or *colouring* for brevity) of a graph $G = (V, E)$ we mean a function $\sigma : E \rightarrow [s]$, where we denote $[s] := \{1, \dots, s\}$. Note that we do not require colourings to be proper, that is, adjacent edges can have the same colour. A colouring σ of G is called *\mathbf{k} -valid* if, for every $c \in [s]$, the colour- c subgraph $\sigma^{-1}(c)$ contains no copy of K_{k_c} , the complete graph of order k_c . Write $F(G; \mathbf{k})$ for the number of \mathbf{k} -valid colourings of G .

In this paper, we investigate $F(n; \mathbf{k})$, the maximum of $F(G; \mathbf{k})$ over all graphs G on n vertices, and the *\mathbf{k} -extremal* graphs, i.e. order- n graphs which attain this maximum. We assume throughout the paper that $s \geq 2$ and that $k_c \geq 3$ for all $c \in [s]$ (since $k_c = 2$ just forbids colour c and the problem reduces to one with $s - 1$ colours).

[†] O.P. was supported by ERC grant 306493 and EPSRC grant EP/K012045/1.

[‡] K.S. was supported by ERC grant 306493.

1.1. *Previous work*

The problem, namely the case when $k_1 = \dots = k_s =: k$, was first considered by Erdős and Rothschild in 1974 (see [5, 6]). Clearly, any colouring of a K_k -free graph is \mathbf{k} -valid. By Turán's theorem [19], the maximum such graph on n vertices is $T_{k-1}(n)$, the complete $(k-1)$ -partite graph with parts as equal as possible. This implies the trivial lower bound

$$F(n; (k, \dots, k)) \geq s^{t_{k-1}(n)}, \quad (1.1)$$

where $t_{k-1}(n)$ is the number of edges in $T_{k-1}(n)$. In particular, Erdős and Rothschild conjectured that, when $\mathbf{k} = (3, 3)$ and n is sufficiently large, the trivial lower bound (1.1) is in fact tight and, furthermore, $T_2(n)$ is the unique \mathbf{k} -extremal graph. The conjecture was verified for all $n \geq 6$ by Yuster [21] (who also computed $F(n; (3, 3))$ for smaller n). Yuster generalised the conjecture to $\mathbf{k} = (k, k)$ and proved an asymptotic version. The full conjecture for all $k \geq 3$ was proved by Alon, Balogh, Keevash and Sudakov [1] who further showed that an analogous result holds for three colours:

THEOREM 1 (Alon, Balogh, Keevash and Sudakov [1]). *Let $k, n \in \mathbb{N}$ where $k \geq 3$ and $n \geq n_0(k)$. Then*

$$F(n; (k, k)) = 2^{t_{k-1}(n)} \quad \text{and} \quad F(n; (k, k, k)) = 3^{t_{k-1}(n)}.$$

Moreover, $T_{k-1}(n)$ is the unique extremal graph in both cases. \square

The proof of Theorem 1 uses Szemerédi's Regularity Lemma. Unfortunately, this also means that the graphs to which it applies are very large indeed. In fact, the assertions are not true for all numbers n of vertices. As was remarked in [1], the conclusion of Theorem 1 fails when $k \leq n < s^{(k-2)/2}$, as in this case a random colouring of the edges of K_n with s colours contains no monochromatic K_k with probability more than $1/2$. Thus, for this range of n , we have $F(n; (k, \dots, k)) > s^{\binom{n}{2}/2} \geq s^{t_{k-1}(n)}$.

The authors of [1] noted that when more than three colours are used, the behaviour of $F(n; (k, \dots, k))$ changes, making its determination both harder and more interesting. Namely, it was shown in [1, page 287] that if $s \geq 4$ (and $k \geq 3$) then $F(n; (k, \dots, k))$ is exponentially larger than $s^{t_{k-1}(n)}$. In particular, any extremal graph has to contain many copies of K_k . In the case when $\mathbf{k} = (3, 3, 3, 3)$, they determined $\log F(n; \mathbf{k})$ asymptotically by showing that $F(n; (3, 3, 3, 3)) = (2^{1/8} 3^{1/2})^{n^2 + o(n^2)}$, where $T_4(n)$ achieves the right exponent. Similarly, they proved that $F(n; (4, 4, 4, 4)) = (3^{8/9})^{n^2 + o(n^2)}$, where $T_9(n)$ achieves the right exponent. Determining the exact answer in these two cases, the first and third author of this paper proved that, when $n \geq n_0$, $T_4(n)$ is the unique $(3, 3, 3, 3)$ -extremal graph on n vertices, and $T_9(n)$ is the unique $(4, 4, 4, 4)$ -extremal graph on n vertices.

It was also proved in [1, Proposition 5.1] that the limit

$$F(\mathbf{k}) := \lim_{n \rightarrow \infty} \frac{\log_2 F(n; \mathbf{k})}{n^2/2} \quad (1.2)$$

exists (and is positive) when $\mathbf{k} = (k, \dots, k)$. As it is easy to see, the proof from [1] extends to an arbitrary fixed sequence \mathbf{k} .

Erdős and Rothschild also considered the generalisation of the problem, where one forbids a monochromatic graph H (the same for each colour). In [1] the authors showed that the analogue of Theorem 1 holds when H is *colour-critical*, that is, the removal of any edge from H reduces its chromatic number. (Note that every clique is colour-critical.)

In a further generalisation, Balogh [3] considered edge-colourings which themselves do not contain a specific colouring of a fixed graph H . Other authors have addressed this question in the cases of forbidden monochromatic matchings, stars, paths, trees and some other graphs in [9, 10], matchings with a prescribed colour pattern in [11], and rainbow stars in [13]. Extending work in [11], Benevides, Hoppen and Sampaio considered forbidden cliques with a prescribed colour pattern, and using techniques similar to our own, obtained several results in this direction, including a version of Theorem 2 below. In [8], a colouring version of the Erdős-Ko-Rado theorem for families of ℓ -intersecting r -element subsets of an n -element set was considered; that is, one counts the number of colourings of families of r -sets such that every colour class is ℓ -intersecting. A so-called ‘ q -analogue’ was addressed in [12], which considers a colouring version of the Erdős-Ko-Rado theorem in the context of vector spaces over a finite field $GF(q)$.

Alon and Yuster [2] studied a directed version of the problem, to determine the maximum number of T -free orientations of an n -vertex graph, where T is a given k -vertex tournament. They showed that the answer is $2^{t_{k-1}(n)}$ for $n \geq n_0(k)$. This in fact answers the original question of Erdős [5], which he modified to ask about edge-colourings.

The problem of counting H -free edge-colourings in hypergraphs was studied in [8, 15, 16]. In an asymptotic hypergraph version of Theorem 1, Lefmann, Person and Schacht [16] proved that, for every k -uniform hypergraph H and $s \in \{2, 3\}$, the maximum number of H -free s -edge-colourings over all k -uniform hypergraphs with n vertices is $s^{\text{ex}(n, H) + o(n^k)}$, where the Turán function $\text{ex}(n, H)$ is the maximum number of edges in an H -free k -uniform hypergraph on n vertices. This is despite the fact that $\text{ex}(n, H)$ is known only for few H .

1.2. New results

Our first result states that it suffices to consider very special graphs G in order to determine the value of $F(n; \mathbf{k})$:

THEOREM 2. *For every $n, s \in \mathbb{N}$ and $\mathbf{k} \in \mathbb{N}^s$, at least one of the \mathbf{k} -extremal graphs of order n is complete multipartite.*

Our second result (Theorem 4 below) writes the limit in (1.2) as the value of a certain optimisation problem.

Problem Q_t : *Given a sequence $\mathbf{k} := (k_1, \dots, k_s) \in \mathbb{N}^s$ of natural numbers and $t \in \{0, 1, 2\}$, determine*

$$Q_t(\mathbf{k}) := \max_{(r, \phi, \alpha) \in \text{FEAS}_t(\mathbf{k})} q(r, \phi, \alpha), \quad (1.3)$$

the maximum value of

$$q(r, \phi, \alpha) := 2 \sum_{\substack{1 \leq i < j \leq r \\ \phi(ij) \neq \emptyset}} \alpha_i \alpha_j \log_2 |\phi(ij)| \quad (1.4)$$

over the set $\text{FEAS}_t(\mathbf{k})$ of feasible solutions, that is, triples (r, ϕ, α) such that

- $r \in \mathbb{N}$ and $r < R(\mathbf{k})$, where $R(\mathbf{k})$ is the Ramsey number of \mathbf{k} (i.e. the minimum R such that K_R admits no \mathbf{k} -valid s -edge-colouring);
- $\phi \in \Phi_t(r; \mathbf{k})$, where $\Phi_t(r; \mathbf{k})$ is the set of all functions $\phi : \binom{[r]}{2} \rightarrow 2^{[s]}$ such that

$$\phi^{-1}(c) := \left\{ ij \in \binom{[r]}{2} : c \in \phi(ij) \right\}$$

is K_{k_c} -free for every colour $c \in [s]$ and $|\phi(ij)| \geq t$ for all $ij \in \binom{[r]}{2}$;

- $\alpha = (\alpha_1, \dots, \alpha_r) \in \Delta^r$, where Δ^r is the set of all $\alpha \in \mathbb{R}^r$ with $\alpha_i \geq 0$ for all $i \in [r]$, and $\alpha_1 + \dots + \alpha_r = 1$.

Note that the maximum in (1.3) is attained. Indeed, for each of the finitely many allowed pairs (r, ϕ) , the function $q(r, \phi, \cdot)$ is continuous and hence attains its maximum over the non-empty compact set Δ^r . A triple (r, ϕ, α) is called Q_t -optimal if it attains the maximum, that is, $(r, \phi, \alpha) \in \text{FEAS}_t(\mathbf{k})$ and $q(r, \phi, \alpha) = Q_t(\mathbf{k})$.

As we will show later in Lemma 6, $Q_0(\mathbf{k}) = Q_1(\mathbf{k}) = Q_2(\mathbf{k})$ so we will denote this common value by $Q(\mathbf{k})$. Of course, if one wishes to determine the value of $Q(\mathbf{k})$, then one should work with Problem Q_2 as it has the smallest feasible set. Since one of our results is stated in terms of Q_1 -optimal triples (which may be a strict superset of Q_2 -optimal triples), we stated different versions of the optimisation problem. In Section 5, we explore how one might hope to solve this optimisation problem, and show that all previously obtained (asymptotic) results can be recovered.

First we show that $Q(\mathbf{k})$ gives rise to an asymptotic lower bound on $F(n; \mathbf{k})$.

LEMMA 3. *For every $s \in \mathbb{N}$ and $\mathbf{k} \in \mathbb{N}^s$, there exists C such that for all $n \in \mathbb{N}$ there is a graph G on n vertices with $F(G; \mathbf{k}) \geq 2^{Q(\mathbf{k})\binom{n}{2} - Cn}$.*

Proof. Let (r, ϕ, α) be Q_0 -optimal. For $n \in \mathbb{N}$, let $G_{\phi, \alpha}(n)$ be the graph of order n with vertex partition X_1, \dots, X_r , where $||X_i| - \alpha_i n| \leq 1$; and in which for all $i, j \in [r]$ and $x_i \in X_i$ and $y_j \in X_j$, we have that $x_i y_j$ is an edge of $G_{\phi, \alpha}(n)$ if and only if $i \neq j$ and $\phi(ij) \neq \emptyset$. Consider those colourings of $G_{\phi, \alpha}(n)$ in which $x_i y_j$ is coloured with some colour in $\phi(ij)$, for every $x_i \in X_i$, $y_j \in X_j$, where $1 \leq i < j \leq r$. Every such colouring is \mathbf{k} -valid because $\phi^{-1}(c)$ is K_{k_c} -free for all $c \in [s]$. The number of such colourings gives the desired lower bound for $F(n; \mathbf{k})$:

$$F(n; \mathbf{k}) \geq F(G_{\phi, \alpha}(n); \mathbf{k}) \geq \prod_{\substack{1 \leq i < j \leq r \\ \phi(ij) \neq \emptyset}} |\phi(ij)|^{|X_i||X_j|} \geq 2^{Q(\mathbf{k})\binom{n}{2} - Cn}, \quad (1.5)$$

where $C = C(\mathbf{k})$ is a constant due to rounding. \square

THEOREM 4. *For every $s \in \mathbb{N}$ and $\mathbf{k} \in \mathbb{N}^s$, we have $F(n; \mathbf{k}) = 2^{Q(\mathbf{k})\binom{n}{2} + o(n^2)}$, that is, $F(\mathbf{k}) = Q(\mathbf{k})$, where $F(\mathbf{k})$ is the limit in (1.2).*

So, as in the result of Lefmann, Person and Schacht [16] mentioned above, this theorem can be proved without knowledge of $Q(\mathbf{k})$. Our proof of Theorem 4 builds upon the techniques of [1, 17] and also uses the Regularity Lemma.

The structure of an arbitrary order- n graph G with $F(G; \mathbf{k}) = 2^{(Q(\mathbf{k}) + o(1))n^2/2}$ can be rather complicated (see a short discussion in Section 5 of the case $\mathbf{k} = (4, 3)$). However, the next result states that if G is assumed to be complete multipartite, then the part ratios have to be close to being Q_1 -optimal.

THEOREM 5. *For every $\delta > 0$ there are $\eta > 0$ and n_0 such that if $G = (V, E)$ is a complete multipartite graph of order $n \geq n_0$ with (non-empty) parts V_1, \dots, V_r and $F(G; \mathbf{k}) \geq 2^{(Q(\mathbf{k}) - \eta)n^2/2}$ then there is a Q_1 -optimal triple (r, ϕ, α') such that the ℓ^1 -distance between $\alpha' \in \Delta^r$ and $\alpha = (|V_1|/n, \dots, |V_r|/n)$ is at most δ : $\|\alpha - \alpha'\|_1 := \sum_{i=1}^r |\alpha_i - \alpha'_i| \leq \delta$.*

In a sense, a converse to Theorem 5 holds. Indeed, for every Q_1 -optimal triple (r, ϕ, α') , for all $n \in \mathbb{N}$, the proof of Lemma 3 gives a complete r -partite graph $G_{\phi, \alpha'}(n)$ on n vertices with parts X_1^n, \dots, X_r^n such that, setting $\alpha_n = (|X_1^n|/n, \dots, |X_r^n|/n)$, we have, as $n \rightarrow \infty$, that

$$\frac{\log_2 F(G_{\phi, \alpha'}(n); \mathbf{k})}{n^2/2} \rightarrow Q(\mathbf{k}) \text{ and } \|\alpha_n - \alpha'\| \rightarrow 0.$$

The rest of the paper is organised as follows. Theorem 2 is proved in Section 2. Section 3 contains a general lemma which is then used in Section 4 to prove Theorems 4 and 5. Section 5 contains some concluding remarks. We will use the following notation. For a set X and an integer $k \leq |X|$, let $\binom{X}{k}$ denote the set of all k -subsets of X . Also, let 2^X be the set of all subsets of X . If it is clear from the context, we may write ij to denote the set $\{i, j\}$ or the ordered pair (i, j) .

2. Symmetrisation and \mathbf{k} -extremal graphs

In this section we prove Theorem 2, which states that, for any instance of the problem (i.e. any choice of the parameters n, s, \mathbf{k}), there is a complete multipartite graph which is \mathbf{k} -extremal. The proof uses the well-known symmetrisation method that was introduced by Zykov [22].

Proof of Theorem 2. Let $G = (V, E)$ be a \mathbf{k} -extremal graph on n vertices. Consider distinct vertices $u, v \in V$ with $uv \notin E$. Let $G' = G - \{u, v\}$, where $G - X = G[V \setminus X]$ is the graph obtained from G by removing every vertex of a set $X \subseteq V$ and every edge adjacent to a vertex of X . For a graph H , let $\mathcal{F}(H)$ denote the set of \mathbf{k} -valid colourings of H . (Thus $F(H; \mathbf{k}) = |\mathcal{F}(H)|$.) Let σ_u and σ_v denote the number of \mathbf{k} -valid extensions of $\sigma \in \mathcal{F}(G')$ to $G - \{v\}$ and $G - \{u\}$ respectively. Since $uv \notin E$ and each forbidden graph is a clique, we have that the number of \mathbf{k} -valid extensions of σ to G is $\sigma_u \sigma_v$. Thus

$$F(G; \mathbf{k}) = \sum_{\sigma \in \mathcal{F}(G')} \sigma_u \sigma_v. \quad (2.1)$$

Let G_u be the graph obtained from G by deleting v and adding a new vertex u' which is a clone of u in G . Define G_v analogously. From (2.1), it follows that

$$F(G_u; \mathbf{k}) = \sum_{\sigma \in \mathcal{F}(G')} \sigma_u^2 \text{ and } F(G_v; \mathbf{k}) = \sum_{\sigma \in \mathcal{F}(G')} \sigma_v^2. \quad (2.2)$$

Since G is \mathbf{k} -extremal, we have that

$$0 \leq 2F(G; \mathbf{k}) - F(G_u; \mathbf{k}) - F(G_v; \mathbf{k}) \stackrel{(2.1), (2.2)}{=} - \sum_{\sigma \in \mathcal{F}(G')} (\sigma_u - \sigma_v)^2 \leq 0, \quad (2.3)$$

and hence we have equality everywhere. Therefore G_u and G_v are both \mathbf{k} -extremal. In order to finish the proof, it is enough to show that we can reach a complete multipartite graph by starting with G and iteratively performing the above operation.

We say that two vertices x and y are *twins* (and write $x \sim y$) if they have the same sets of neighbours. Note that twins are necessarily non-adjacent. It is easy to see that \sim is an equivalence relation. Let $[x]_{\sim}$ denote the equivalence class of x .

Let $G^1 := G$. Repeat the following for as long as possible. Suppose that we have defined graphs G^1, \dots, G^i for some $i \geq 1$, which are all \mathbf{k} -extremal. Suppose that G^i contains a

pair u, v of non-adjacent vertices which are not twins. Choose such a pair so that $||[u]_{\sim}|$ is maximal. Let $G^{i+1} = (G^i)_u$ be the graph obtained from G^i by deleting v and adding a new vertex u' which is a clone of u . As was argued above, G^{i+1} is necessarily \mathbf{k} -extremal.

For each $i \geq 1$, call an equivalence class $[x]_{\sim}$ in the graph G^i *frozen* if G^i is complete between $[x]_{\sim}$ and its complement, and *unfrozen* otherwise. Let $f(G^i)$ be the sum of sizes of all frozen classes plus the largest size of an unfrozen one. It is easy to see that $f(G^i)$ is strictly increasing with i . Since $f(G^i)$ is bounded above by n , the process terminates in at most $n - 1$ steps with some \mathbf{k} -extremal graph H . Since every pair of non-adjacent vertices in H are twins, H is complete multipartite, as desired. \square

Also, the symmetrisation can be applied to Q_t -optimal solutions. In particular, one can prove the following.

LEMMA 6. *For every \mathbf{k} , we have $Q_0(\mathbf{k}) = Q_1(\mathbf{k}) = Q_2(\mathbf{k})$.*

Proof. Since trivially $\text{FEAS}_0(\mathbf{k}) \supseteq \text{FEAS}_1(\mathbf{k}) \supseteq \text{FEAS}_2(\mathbf{k})$, we have $Q_0(\mathbf{k}) \geq Q_1(\mathbf{k}) \geq Q_2(\mathbf{k})$.

On the other hand, among all Q_0 -optimal solutions (r, ϕ, α) , fix one with r as small as possible. Then, in particular, we have that each α_i is non-zero. We claim that necessarily $(r, \phi, \alpha) \in \text{FEAS}_2(\mathbf{k})$ (which will give the required inequality $Q_2(\mathbf{k}) \geq Q_0(\mathbf{k})$). If this is not true, then $|\phi(ij)| \leq 1$ for some $ij \in \binom{[r]}{2}$, say for $\{i, j\} = \{r-1, r\}$. For a real c , consider α' defined by $\alpha'_{r-1} = \alpha_{r-1} + c$, $\alpha'_r = \alpha_r - c$ and $\alpha'_h := \alpha_h$ for all $h \in [r-2]$. In other words, we shift weight c from α_r to α_{r-1} . Since $q(r, \phi, \alpha')$ is a linear function $f(c)$ of c and $(r, \phi, \alpha') \in \text{FEAS}_0(\mathbf{k})$ when $|c|$ is at most $\min\{\alpha_{r-1}, \alpha_r\} > 0$, it must be the case that $f(c)$ is a constant function. Thus $f(c) = f(0) = Q_0(\mathbf{k})$ regardless of c . In particular, by taking $c = \alpha_r$, that is, by shifting all weight from α_r to α_{r-1} , we obtain a Q_0 -optimal solution (r, ϕ, α') with $\alpha'_r = 0$, whose restriction to $[r-1]$ gives another Q_0 -optimal solution, contradicting the minimality of r . \square

3. A unifying lemma

The proofs of Theorems 4 and 5 will both follow from the next lemma, which states that the number of \mathbf{k} -valid colourings of any complete r -partite graph H can be bounded above by evaluating q for a triple $(r, \phi, \beta) \in \text{FEAS}_1(\mathbf{k})$, where β is given by the ratios of the parts of H .

LEMMA 7. *For all $s \in \mathbb{N}$, $\mathbf{k} \in \mathbb{N}^s$ and $\eta > 0$, there exists $n_0 \in \mathbb{N}$ such that for every complete multipartite graph H of order $N \geq n_0$ with (non-empty) parts Y_1, \dots, Y_r with at least one \mathbf{k} -valid colouring, there is some $\phi \in \Phi_1(r; \mathbf{k})$ such that*

$$\frac{\log_2 F(H; \mathbf{k})}{N^2/2} \leq q(r, \phi, \beta) + \eta,$$

where $\beta := (|Y_1|/N, \dots, |Y_r|/N)$.

In outline, the argument to prove Lemma 7 is as follows. The main idea of the proof is to use Szemerédi's Regularity Lemma to pass from a \mathbf{k} -valid colouring σ of H to a set of feasible solutions that come from r -tuples of clusters which are transversal with respect to the r -partition of H . For each obtained solution $(r, \phi, \beta) \in \text{FEAS}_0(\mathbf{k})$, an upper bound on $q(r, \phi, \beta)$ can be translated via regularity into an upper bound on the number of restrictions of possible colourings σ to the involved clusters (an idea already used in [1]).

Then we estimate $F(H; \mathbf{k})$ by taking an appropriately weighted sum of logarithms of these bounds. It turns out that the dominant contribution is from those triples (r, ϕ, β) that belong to $\text{FEAS}_1(\mathbf{k})$, and so the bound obtained for $F(H; \mathbf{k})$ is in terms of the largest $q(r, \phi, \beta)$ among such triples.

3.1. Regularity tools

We will need the following definitions related to Szemerédi's Regularity Lemma.

DEFINITION 8 (Edge density, ε -regular, (ε, γ) -regular, equitable partition). *Given a graph G and disjoint non-empty sets $A, B \subseteq V(G)$, we define the edge density between A and B to be*

$$d(A, B) := \frac{|E(G[A, B])|}{|A||B|}.$$

Given $\varepsilon, \gamma > 0$, the pair (A, B) is called

- ε -regular if for every $X \subseteq A$ and $Y \subseteq B$ with $|X| \geq \varepsilon|A|$ and $|Y| \geq \varepsilon|B|$, we have that $|d(X, Y) - d(A, B)| \leq \varepsilon$;
- (ε, γ) -regular if it is ε -regular and has edge density at least γ .

We call a partition $V(G) = V_1 \cup \dots \cup V_m$

- equitable if $||V_i| - |V_j|| \leq 1$ for all $i, j \in [m]$;
- ε -regular if it is equitable, $m \geq 1/\varepsilon$, and all but at most $\varepsilon \binom{m}{2}$ of the pairs (V_i, V_j) with $1 \leq i < j \leq m$ are ε -regular.

Our first tool states that an induced subgraph of a regular pair is still regular, provided both parts are not too small.

PROPOSITION 9. *Let ε, δ be such that $0 < 2\delta \leq \varepsilon < 1$. Suppose that (X, Y) is a δ -regular pair, and let $X' \subseteq X$ and $Y' \subseteq Y$. If*

$$\min \left\{ \frac{|X'|}{|X|}, \frac{|Y'|}{|Y|} \right\} \geq \frac{\delta}{\varepsilon},$$

then the pair (X', Y') is ε -regular.

Proof. Let $X'' \subseteq X'$ and $Y'' \subseteq Y'$ be such that $|X''| \geq \varepsilon|X'|$ and $|Y''| \geq \varepsilon|Y'|$. Then $|X''|/|X|, |Y''|/|Y| \geq \delta$. Since (X, Y) is δ -regular, we have that $|d(X'', Y'') - d(X, Y)| \leq \delta$. Note further that $|X'|/|X|, |Y'|/|Y| \geq \delta/\varepsilon > \delta$, so $|d(X', Y') - d(X, Y)| \leq \delta$. By the Triangle Inequality, $|d(X'', Y'') - d(X', Y')| \leq 2\delta \leq \varepsilon$. This implies that (X', Y') is ε -regular. \square

We use the following multicolour version of Szemerédi's Regularity Lemma [18] (see e.g. Theorem 1.18 in Komlós and Simonovits [14]).

LEMMA 10 (Multicolour Regularity Lemma). *For every $\varepsilon > 0$ and $s \in \mathbb{N}$, there exists $M \in \mathbb{N}$ such that for any graph G on $n \geq M$ vertices and any s -edge-colouring $\sigma : E(G) \rightarrow [s]$, there is an (equitable) partition $V(G) = V_1 \cup \dots \cup V_m$ with $m \leq M$, which is ε -regular simultaneously with respect to all graphs $(V(G), \sigma^{-1}(c))$, with $c \in [s]$. \square*

Finally, we need the following bound.

PROPOSITION 11. *Let $s, r \in \mathbb{N}$ and $\mathbf{k} \in \mathbb{N}^s$. Let $\phi \in \Phi_0(r; \mathbf{k})$ and $\alpha, \beta \in \Delta^r$. Then*

$$|q(r, \phi, \alpha) - q(r, \phi, \beta)| \leq 2\|\alpha - \beta\|_1 \log_2 s.$$

Proof. We have that

$$\begin{aligned}
& |q(r, \phi, \boldsymbol{\alpha}) - q(r, \phi, \boldsymbol{\beta})| \\
&= \left| \sum_{i \in [r]} \alpha_i \sum_{j \in [r] \setminus \{i\}} \alpha_j \log |\phi(ij)| - \sum_{i \in [r]} \beta_i \sum_{j \in [r] \setminus \{i\}} \beta_j \log |\phi(ij)| \right| \\
&\leq \left| \sum_{i \in [r]} (\alpha_i - \beta_i) \sum_{j \in [r] \setminus \{i\}} \alpha_j \log |\phi(ij)| \right| + \left| \sum_{j \in [r]} (\alpha_j - \beta_j) \sum_{i \in [r] \setminus \{j\}} \beta_i \log |\phi(ij)| \right| \\
&\leq 2 \log_2(s) \cdot \|\boldsymbol{\alpha} - \boldsymbol{\beta}\|_1.
\end{aligned}$$

□

3.2. Proof of Lemma 7

Let $\eta > 0$ (assumed without loss of generality to be sufficiently small) and choose an additional constant γ so that $0 < \gamma \ll \eta \ll 1/R(\mathbf{k})$. By the (standard) Embedding Lemma (see, for example, [14, Theorem 2.1]), there exist $\varepsilon > 0$ and $m_0 \in \mathbb{N}$ such that the following holds for all $c \in [s]$: if G is a graph with a partition $V(G) = W_1 \cup \dots \cup W_{k_c}$ such that $|W_i| \geq m_0$ for all $i \in [k_c]$ and every pair (W_i, W_j) for $1 \leq i < j \leq k_c$ is (ε, γ) -regular, then $K_{k_c} \subseteq G$.

We may assume that $0 < 1/m_0 \ll \varepsilon \ll \gamma$ since whenever $\varepsilon' \leq \varepsilon$, we have that an ε' -regular pair is also an ε -regular pair. Let M be the integer returned by Lemma 10 when applied with parameters ε^2 and s . Choose $n_0 \in \mathbb{N}$ and assume, without loss of generality, that $1/n_0 \ll 1/M \ll 1/m_0$. We have the hierarchy

$$0 < 1/n_0 \ll 1/M \ll 1/m_0 \ll \varepsilon \ll \gamma \ll \eta \ll 1/R(\mathbf{k}). \quad (3.1)$$

Let $N \geq n_0$ be arbitrary. Let H be a complete multipartite graph on N vertices with parts Y_1, \dots, Y_r . We may assume that $r < R(\mathbf{k})$ otherwise $F(H; \mathbf{k}) = 0$. Let $G = (V, E)$ be a graph obtained from H by removing all but one vertex from every part Y_i of size at most $\eta^2 N$ (and all edges incident with the removed vertices). Write $n := |V|$ and $X_i := Y_i \cap V$ for all $i \in [r]$. Then $N - n \leq R(\mathbf{k}) \cdot \eta^2 N$. So

$$F(G, \mathbf{k}) \geq F(H, \mathbf{k}) \cdot s^{-R(\mathbf{k})\eta^2 N^2}$$

and so

$$\begin{aligned}
\frac{\log_2 F(G; \mathbf{k})}{n^2/2} &\geq \frac{\log_2 F(H; \mathbf{k})}{N^2/2} \geq \frac{\log_2 F(H; \mathbf{k})}{N^2/2} - 3R(\mathbf{k})\eta^2 \log_2 s \\
&\geq \frac{\log_2 F(H; \mathbf{k})}{N^2/2} - \frac{\eta}{3}.
\end{aligned} \quad (3.2)$$

Define $\boldsymbol{\alpha} := (|X_1|/n, \dots, |X_r|/n)$ and $\boldsymbol{\beta} := (|Y_1|/N, \dots, |Y_r|/N)$. Then

$$\|\boldsymbol{\alpha} - \boldsymbol{\beta}\|_1 \leq \frac{R(\mathbf{k})\eta^2 N}{n} \leq 2R(\mathbf{k})\eta^2. \quad (3.3)$$

Without loss of generality, there is some $w \in [r]$ such that $X_i = \{x_i\}$ is a singleton for all $i \in [w]$, and $|X_j| > \eta^2 n$ for all $w < j \leq r$.

For the rest of the proof, we will work with G rather than H . Informally, the reason for passing to G is the following. After applying the Regularity Lemma to H with a valid colouring σ , we do not *a priori* have control on the distribution of coloured edges incident to small parts of H . If the statement of Lemma 7 asked for a $\phi \in \Phi_0(r; \mathbf{k})$, we

could simply neglect these parts; but since we require $\phi \in \Phi_1(r; \mathbf{k})$ we cannot do this. Therefore we introduce G in which each small part X_i is replaced by a token vertex x_i , which merely asserts the existence of its part. But for each $x \in V(G)$, there are only constantly many possible values for $\{\sigma(xx_i) : i \in [w]\}$ for all s -edge-colourings σ . Thus we can refine our regularity partition into parts according to these values. Now we have good control between *all* pairs of parts: if both are large then regularity provides good control; and if one of them is small it is necessarily a single vertex and σ is constant on all edges between the parts.

Let $\sigma : E \rightarrow [s]$ be a \mathbf{k} -valid colouring of G . By the choice of M (that is, by Lemma 10 applied to G and σ with parameters ε^2 and s), there is an (equitable) partition $V = V_1 \cup \dots \cup V_m$, with $m \leq M$, which is ε^2 -regular simultaneously with respect to all graphs $(V, \sigma^{-1}(c))$, $c \in [s]$.

We will now take a common refinement of X_1, \dots, X_r and V_1, \dots, V_m which also takes into account attachments to $W := \{x_1, \dots, x_w\}$. Namely, for all $j \in [m]$, subdivide V_j into at most $r(s^w + w)$ parts as follows. Put each vertex in $W \cap V_j$ into a separate part. Now, for any vertices y, y' remaining in V_j , put y and y' in the same part if and only if there is some $\ell \in [r]$ such that $\{y, y'\} \subseteq X_\ell$, and $\sigma(x_h y) = \sigma(x_h y')$ for all $h \in [w]$. Thus we obtain a (not necessarily equitable) partition $U_{i,1} \cup \dots \cup U_{i,m_i}$ of X_i for each $i \in [r]$, where $m_i \leq M(s^w + w)$. Let \mathcal{U} be the collection of sets $U_{i,j}$. It is indexed by

$$I := \{ij : i \in [r] \text{ and } j \in [m_i]\}.$$

For a colour $c \in [s]$, let P^c consist of all pairs of indices $\{ig, jh\} \in \binom{I}{2}$ such that $\sigma^{-1}(c)[U_{i,g}, U_{j,h}]$ is (ε, γ) -regular, and at least one of the following holds: $U_{i,g}$ is a vertex of W ; $U_{j,h}$ is a vertex of W ; or $\min\{|U_{i,g}|, |U_{j,h}|\} \geq m_0$. (So if, say, $U_{i,g}$ is a vertex of W , then $\{ig, jh\} \in P^c$ for some $c \in [s]$ since $G[U_{i,g}, U_{j,h}]$ is a monochromatic star under σ .) We define $E^c \subseteq E$ to be the union of $\sigma^{-1}(c)[U_{i,g}, U_{j,h}]$ over all pairs $\{ig, jh\} \in P^c$. Let $E_0 := E \setminus (E^1 \cup \dots \cup E^s)$. Thus E_0 consists of edges without endpoints in W which are incident with a part of size less than m_0 ; and edges which come from coloured pairs that are not ε -regular or have edge density less than γ . The following claim, whose proof is fairly standard, shows that E_0 cannot contain many edges.

CLAIM 12. $|E_0| \leq s\gamma n^2$.

Proof. Call a part $U_{i,g} \subseteq V_\ell$ *small* if $|U_{i,g}| < \varepsilon|V_\ell|$. Let $E_{\text{small}} \subseteq E$ be the set of edges that have at least one vertex in a small part. Since each V_ℓ is subdivided into at most $r(s^w + w) < 2R(\mathbf{k})s^{R(\mathbf{k})}$ new parts, the number of vertices in small parts is at most $2\varepsilon R(\mathbf{k})s^{R(\mathbf{k})}n$ and, trivially,

$$|E_{\text{small}}| \leq 2\varepsilon R(\mathbf{k})s^{R(\mathbf{k})}n^2.$$

Let $E_{\text{irr}} \subseteq E$ consist of those edges of G that lie inside some V_ℓ or belong to some colour- c bipartite subgraph $\sigma^{-1}(c)[V_\ell, V_{\ell'}]$ which is not ε^2 -regular. Since $V_1 \cup \dots \cup V_m$ is an ε^2 -regular (equitable) partition, we have

$$|E_{\text{irr}}| \leq m \binom{\lceil n/m \rceil}{2} + s\varepsilon^2 \binom{m}{2} \left\lceil \frac{n}{m} \right\rceil^2$$

which is by $m \geq 1/\varepsilon^2$ at most, say, εn^2 .

Next, we bound the size of $E_0 \setminus (E_{\text{small}} \cup E_{\text{irr}})$. Let e be any edge from this set. Since each $U_{i,g}$ is an independent set in G , we have $e \in E(G[U_{i,g}, U_{j,h}])$ for some distinct

$ig, jh \in I$. Let $\ell, \ell' \in [m]$ satisfy $V_\ell \supseteq U_{i,g}$ and $V_{\ell'} \supseteq U_{j,h}$. Since $e \notin E_{\text{small}}$, we have

$$\min\{|U_{i,g}|, |U_{j,h}|\} \geq \min\{\varepsilon|V_\ell|, \varepsilon|V_{\ell'}|\} \geq \varepsilon \lfloor n/m \rfloor,$$

which is at least m_0 by our choice of constants. Let $c = \sigma(e)$ be the colour of e . Since $e \notin E_{\text{irr}}$, we have that $\ell \neq \ell'$ and $\sigma^{-1}(c)[V_\ell, V_{\ell'}]$ is ε^2 -regular. Thus Proposition 9 implies that $\sigma^{-1}(c)[U_{i,g}, U_{j,h}]$ is ε -regular. Since $e \notin E^c$, it must be the case that $\sigma^{-1}(c)[U_{i,g}, U_{j,h}] \ni e$ has edge density less than γ . We conclude that $E_0 \setminus (E_{\text{small}} \cup E_{\text{irr}})$ has edge density at most $s\gamma$ between any pair $(U_{i,g}, U_{j,h})$. Thus

$$|E_0| \leq |E_{\text{small}}| + |E_{\text{irr}}| + \sum_{\{ig, jh\} \in \binom{I}{2}} s\gamma |U_{i,g}| |U_{j,h}| \leq 2\varepsilon R(\mathbf{k}) s^{R(\mathbf{k})} n^2 + \varepsilon n^2 + s\gamma \binom{n}{2} < s\gamma n^2,$$

proving the claim. \blacksquare

Define $\phi : \binom{I}{2} \rightarrow 2^{[s]}$ by setting, for all $\{ig, jh\} \in \binom{I}{2}$,

$$\phi(ig, jh) := \{c \in [s] : \{ig, jh\} \in P^c\}.$$

If neither $U_{i,g}$ nor $U_{j,h}$ is a vertex of W but $\min\{|U_{i,g}|, |U_{j,h}|\} < m_0$, then $\phi(ig, jh)$ is empty. Otherwise, $\phi(ig, jh)$ consists of those c for which $\sigma^{-1}(c)[U_{i,g}, U_{j,h}]$ is (ε, γ) -regular. Also, let $\sigma_0 = \sigma|_{E_0}$ be the restriction of σ to E_0 .

For each \mathbf{k} -valid colouring σ of G , fix one partition $V = V_1 \cup \dots \cup V_m$ as above and then define the tuple $(\mathcal{U}, I, \phi, E_0, \sigma_0)$ accordingly.

CLAIM 13. *The number of possible tuples $(\mathcal{U}, I, \phi, E_0, \sigma_0)$ is at most $2^{\eta n^2/4}$.*

Proof. Clearly, there are at most $(M(s^w + w))^n \leq (M(s^{R(\mathbf{k})} + R(\mathbf{k})))^n < 2^{\eta n^2/12}$ possible partitions of V in which, for all $i \in [r]$, every $x \in X_i$ lies in one of at most $M(s^w + w)$ parts. Each such partition determines \mathcal{U} and I uniquely (since the partition $V = X_1 \cup \dots \cup X_r$ is fixed throughout the whole proof).

Given \mathcal{U} and I , the number of possible ϕ is at most $(2^s)^{\binom{r(Ms^w + w)}{2}} < 2^{\eta n^2/12}$. By Claim 12, the number of ways to choose E_0 and colour these edges (i.e. choose σ_0) is, very roughly, at most

$$\binom{\binom{n}{2}}{s\gamma n^2} (s+1)^{s\gamma n^2} < 2^{\eta n^2/12}.$$

The claim is proved by multiplying these three bounds. \blacksquare

Fix a tuple $(\mathcal{U}, I, \phi, E_0, \sigma_0)$ such that $\mathcal{C} \neq \emptyset$, where \mathcal{C} is the set of colourings σ which generate it. Our next step is to provide an upper bound for $|\mathcal{C}|$. For every $\sigma \in \mathcal{C}$, we have $\sigma|_{E_0} = \sigma_0$. Also, by the definition of E_0 , every $e \in E \setminus E_0$ lies in some (ε, γ) -regular bipartite graph $\sigma^{-1}(c)[U_{i,g}, U_{j,h}]$ with $c \in [s]$ and $\{ig, jh\} \in \binom{I}{2}$ such that $\min\{|U_{i,g}|, |U_{j,h}|\} \geq m_0$ or at least one of $U_{i,g}, U_{j,h}$ is a vertex of W . Thus $\{ig, jh\} \in P^c$, that is, $\sigma(e) \in \phi(ig, jh)$. Therefore

$$|\mathcal{C}| \leq \prod_{ij \in \binom{[r]}{2}} \prod_{\substack{gh \in [m_i] \times [m_j] \\ \phi(ig, jh) \neq \emptyset}} |\phi(ig, jh)|^{|U_{i,g}| |U_{j,h}|}.$$

Let us agree that $\log_2 0 := 0$. Then

$$\log_2 |\mathcal{C}| \leq \sum_{ij \in \binom{[r]}{2}} \sum_{gh \in [m_i] \times [m_j]} |U_{i,g}| |U_{j,h}| \log_2 |\phi(ig, jh)|. \quad (3.4)$$

Let $T := [m_1] \times \dots \times [m_r]$. We use T to index all ‘transversal’ r -tuples of parts from \mathcal{U} , where we take one part from each of X_1, \dots, X_r . For each $\mathbf{t} = (t_1, \dots, t_r)$ in T , define $\phi_{\mathbf{t}} : \binom{[r]}{2} \rightarrow 2^{[s]}$ by setting, for $ij \in \binom{[r]}{2}$,

$$\phi_{\mathbf{t}}(ij) := \phi(it_i, jt_j).$$

Recall the definition of α after (3.2).

CLAIM 14. $\log_2 |\mathcal{C}| \leq (q^* + \sqrt{\gamma})n^2/2$, where

$$q^* := \max\{q(r, \phi_{\mathbf{t}}, \alpha) : (r, \phi_{\mathbf{t}}, \alpha) \in \text{FEAS}_1(\mathbf{k}), \mathbf{t} \in T\}.$$

Proof. We will first show that, for every $c \in [s]$ and $\mathbf{t} \in T$, the graph $\phi_{\mathbf{t}}^{-1}(c)$ is K_{k_c} -free. Indeed, suppose that i_1, \dots, i_{k_c} span a copy of K_{k_c} in $\phi_{\mathbf{t}}^{-1}(c)$. First consider the case when $U_{i_1, t_{i_1}}$ is not a vertex of W but $|U_{i_1, t_{i_1}}| < m_0$. Then, by the definition of ϕ , we have that $U_{i_q, t_{i_q}}$ is a vertex of W for all $2 \leq q \leq k_c$. Moreover, for every $pq \in \binom{[k_c]}{2}$, every edge in $G[U_{i_p, t_{i_p}}, U_{i_q, t_{i_q}}]$ is coloured with c by σ , a contradiction.

So, without loss of generality, we may assume that there is some $0 \leq \ell \leq \min\{k_c, w\}$ such that each of $U_{i_1, t_{i_1}}, \dots, U_{i_\ell, t_{i_\ell}}$ consists of a vertex of W and $|U_{i_q, t_{i_q}}| > m_0$ for all $\ell + 1 \leq q \leq k_c$. Then, by the definition of \mathcal{U} , we have that $\sigma(e) = c$ for all $e \in G[U_{i_p, t_{i_p}}, U_{i_q, t_{i_q}}]$ with $p \in [\ell]$ and $q \in [k_c] \setminus \{p\}$. By the definition of $P^c \supseteq \phi^{-1}(c)$ and the Embedding Lemma (that is, our choice of parameters at the beginning of the proof), for all $\ell + 1 \leq q \leq k_c$, there is $z_q \in U_{i_q, t_{i_q}}$ such that together these vertices z_q span a copy of $K_{k_c - \ell}$ in $\sigma^{-1}(c)$. Then $\sigma^{-1}(c)$ spans a copy of K_{k_c} , contradicting the \mathbf{k} -validity of σ . This and the trivial bound $r < R(\mathbf{k})$ imply that $\phi_{\mathbf{t}} \in \Phi(r; \mathbf{k})$. Therefore, for each $\mathbf{t} \in T$, we have that $(r, \phi_{\mathbf{t}}, \alpha) \in \text{FEAS}_0(\mathbf{k})$, and so

$$\sum_{ij \in \binom{[r]}{2}} \alpha_i \alpha_j \log_2 |\phi(it_i, jt_j)| \leq b(\mathbf{t}), \quad (3.5)$$

where we define

$$b(\mathbf{t}) = \begin{cases} q^*/2 & \text{if } (r, \phi_{\mathbf{t}}, \alpha) \in \text{FEAS}_1(\mathbf{k}) \\ r^2 \log_2(s)/2 & \text{otherwise} \end{cases}$$

(i.e. if $(r, \phi_{\mathbf{t}}, \alpha) \notin \text{FEAS}_1(\mathbf{k})$ we take a (somewhat arbitrary) trivial bound for $b(\mathbf{t})$). The claim will follow from taking a weighted average of (3.5) by multiplying by $\prod_{\ell \in [r]} |U_{\ell, t_\ell}|$ and summing over all $\mathbf{t} \in T$. First consider the right hand side of (3.5). Let T_0 be the set of $\mathbf{t} \in T$ such that $\phi_{\mathbf{t}}(ij) = \emptyset$ for some $ij \in \binom{[r]}{2}$. We will show that the sum of $\prod_{\ell \in [r]} |U_{\ell, t_\ell}|$ over all $\mathbf{t} \in T \setminus T_0$ is not much less than the sum taken over the whole of T .

To this end, fix a pair $\{ig, jh\} \in \binom{[r]}{2}$ such that $\phi(ig, jh) = \emptyset$. If at least one edge e in $G[U_{i,g}, U_{j,h}]$ is not in E_0 , then there is some $c \in [s]$ such that $e \in E^c$. Then $\{ig, jh\} \in P^c$ and so $\phi(ig, jh) \ni c$ is non-empty, a contradiction. Therefore $E(G[U_{i,g}, U_{j,h}]) \subseteq E_0$. Furthermore, by our definition of ϕ , we have that $|X_i|, |X_j| \geq \eta^2 n$. Observe that, if one sums only over those $\mathbf{t} \in T$ that contain $\{ig, jh\}$, then one gets

$$\sum_{\substack{\mathbf{t} \in T: \\ t_i=g, t_j=h}} \prod_{\ell \in [r]} |U_{\ell, t_\ell}| = |U_{i,g}| |U_{j,h}| \prod_{\ell \in [r] \setminus \{i,j\}} |X_\ell| \leq \frac{|U_{i,g}| |U_{j,h}|}{\eta^4 n^2} \prod_{\ell \in [r]} |X_\ell|.$$

Then, using the upper bound on $|E_0|$ from Claim 12, we have that

$$\begin{aligned} \sum_{\mathbf{t} \in T_0} \prod_{\ell \in [r]} |U_{\ell, \mathbf{t}_\ell}| &\leq \sum_{\substack{\{ig, jh\} \in \binom{I}{2}: \\ E(G[U_{i,g}, U_{j,h}]) \subseteq E_0}} \sum_{\substack{\mathbf{t} \in T: \\ t_i = g, t_j = h}} \prod_{\ell \in [r]} |U_{\ell, \mathbf{t}_\ell}| \leq \frac{|E_0|}{\eta^4 n^2} \prod_{\ell \in [r]} |X_\ell| \\ &\leq \frac{s\gamma}{\eta^4} \prod_{\ell \in [r]} |X_\ell|. \end{aligned} \quad (3.6)$$

We can now give an upper bound for the weighted average of the right hand of (3.5) as follows:

$$\begin{aligned} \sum_{\mathbf{t} \in T} \prod_{\ell \in [r]} |U_{\ell, \mathbf{t}_\ell}| b(\mathbf{t}) &\leq \frac{q^*}{2} \sum_{\mathbf{t} \in T} \prod_{\ell \in [r]} |U_{\ell, \mathbf{t}_\ell}| + \frac{r^2 \log_2 s}{2} \sum_{\mathbf{t} \in T_0} \prod_{\ell \in [r]} |U_{\ell, \mathbf{t}_\ell}| \\ &\stackrel{(3.6)}{\leq} \prod_{\ell \in [r]} |X_\ell| \left(\frac{q^*}{2} + \frac{r^2 s \gamma \log_2 s}{2\eta^4} \right) \leq \prod_{\ell \in [r]} |X_\ell| \frac{q^* + \sqrt{\gamma}}{2}. \end{aligned} \quad (3.7)$$

Using this bound together with a weighted average of the left hand side of (3.5), we have that

$$\begin{aligned} &\frac{q^* + \sqrt{\gamma}}{2} \prod_{\ell \in [r]} |X_\ell| \\ &\geq \sum_{\mathbf{t} \in T} \sum_{ij \in \binom{[r]}{2}} \alpha_i \alpha_j \log_2 |\phi(it_i, jt_j)| \prod_{\ell \in [r]} |U_{\ell, \mathbf{t}_\ell}| \\ &= \sum_{ij \in \binom{[r]}{2}} \alpha_i \alpha_j \sum_{gh \in [m_i] \times [m_j]} |U_{i,g}| |U_{j,h}| \log_2 |\phi(ig, jh)| \sum_{\substack{\mathbf{t} \in T: \\ t_i = g, t_j = h}} \prod_{\ell \in [r] \setminus \{i, j\}} |U_{\ell, \mathbf{t}_\ell}| \\ &= \sum_{ij \in \binom{[r]}{2}} \frac{|X_i|}{n} \cdot \frac{|X_j|}{n} \sum_{gh \in [m_i] \times [m_j]} |U_{i,g}| |U_{j,h}| \log_2 |\phi(ig, jh)| \prod_{\ell \in [r] \setminus \{i, j\}} |X_\ell| \\ &\stackrel{(3.4)}{\geq} \frac{1}{n^2} \log_2 |\mathcal{C}| \prod_{\ell \in [r]} |X_\ell|, \end{aligned}$$

proving Claim 14. ■

Let $\mathbf{t}^* \in T$ be such that $q^* = q(r, \phi_{\mathbf{t}^*}, \boldsymbol{\alpha})$. Recall that $\boldsymbol{\beta} = (|Y_1|/N, \dots, |Y_r|/N)$. Then $(r, \phi_{\mathbf{t}^*}, \boldsymbol{\alpha})$ and hence $(r, \phi_{\mathbf{t}^*}, \boldsymbol{\beta})$ lies in $\text{FEAS}_1(\mathbf{k})$. Now Claims 13 and 14 and Proposition 11 imply that

$$\begin{aligned} \frac{\log_2 F(H; \mathbf{k})}{N^2/2} &\stackrel{(3.2)}{\leq} \frac{\log_2 F(G; \mathbf{k})}{n^2/2} + \frac{\eta}{3} \leq \frac{5\eta}{6} + q^* + \sqrt{\gamma} < q(r, \phi_{\mathbf{t}^*}, \boldsymbol{\alpha}) + \frac{6\eta}{7} \\ &\leq q(r, \phi_{\mathbf{t}^*}, \boldsymbol{\beta}) + 2 \log_2(s) \|\boldsymbol{\alpha} - \boldsymbol{\beta}\|_1 + \frac{6\eta}{7} \stackrel{(3.3)}{\leq} q(r, \phi_{\mathbf{t}^*}, \boldsymbol{\beta}) + \eta, \end{aligned}$$

completing the proof of the lemma. □

4. Proofs of Theorems 4 and 5

4.1. Proof of Theorem 4

By Lemma 3, it suffices to show that for every $\eta > 0$, there exists $n_0 \in \mathbb{N}$ such that $\log_2 F(n; \mathbf{k}) \leq (Q(\mathbf{k}) + \eta)n^2/2$ for all $n \geq n_0$. Fix $\eta > 0$ and obtain n_0 from Lemma 7.

Now let $n \geq n_0$. By Theorem 2, there exists a complete multipartite graph G on n vertices with $F(G; \mathbf{k}) = F(n; \mathbf{k})$. The required upper bound on $\log_2 F(G; \mathbf{k})$ follows immediately from Lemma 7. \square

4.2. Proof of Theorem 5

Suppose that there is $\delta > 0$ which contradicts the claim. We need the following claim, which uses a compactness argument to show that a triple in $\text{FEAS}_1(\mathbf{k})$ which is almost optimal is in fact ‘close’ to a Q_1 -optimal triple.

CLAIM 15. *There exists $\eta > 0$ such that for all $(r, \phi, \alpha) \in \text{FEAS}_1(\mathbf{k})$ with $q(r, \phi, \alpha) \geq Q(\mathbf{k}) - 2\eta$, there is a Q_1 -optimal triple (r, ϕ, α') such that $\|\alpha' - \alpha\|_1 \leq \delta$.*

Proof: Suppose this is not the case. Then for all $n \in \mathbb{N}$, there exists $(r, \phi, \alpha_n) \in \text{FEAS}_1(\mathbf{k})$ with

$$q(\phi, \alpha_n) \geq Q(\mathbf{k}) - \frac{1}{n}, \quad (4.1)$$

but for all $\alpha'_n \in \Delta^r$ with $\|\alpha_n - \alpha'_n\|_1 < \delta$, we have that (r, ϕ, α'_n) is not Q_1 -optimal.

Consider the sequence $(\alpha_1, \alpha_2, \dots)$. Since Δ^r is closed and bounded, the Heine-Borel theorem implies that it is compact. Therefore there is some subsequence $(\alpha_{n_1}, \alpha_{n_2}, \dots)$ of $(\alpha_1, \alpha_2, \dots)$ which converges (in any norm, since r is finite). Let $\lambda := \lim_{k \rightarrow \infty} \alpha_{n_k}$. Observe that $\lambda \in \Delta^r$, so $(r, \phi, \lambda) \in \text{FEAS}_1(\mathbf{k})$. Having fixed r, ϕ , observe that $q(r, \phi, \lambda) = 2 \sum_{ij \in \binom{[r]}{2}} \lambda_i \lambda_j \log |\phi(ij)|$ is a continuous function of λ . Therefore

$$\lim_{k \rightarrow \infty} q(r, \phi, \alpha_{n_k}) = q(r, \phi, \lambda).$$

Together with (4.1), this implies that $q(r, \phi, \lambda) = Q(\mathbf{k})$, and so (r, ϕ, λ) is Q_1 -optimal. Now, since $\alpha_{n_k} \rightarrow \lambda$, we can choose $N \in \mathbb{N}$ such that $\|\alpha_N - \lambda\|_1 < \delta$. This contradicts our assumption and hence proves the claim. \blacksquare

Choose η as in the claim. Obtain $n_0 \in \mathbb{N}$ by applying Lemma 7 with η . Since we supposed that $\delta > 0$ contradicts the statement of Theorem 5, there exists a complete multipartite graph G on $n \geq n_0$ vertices such that $F(G; \mathbf{k}) \geq 2^{(Q(\mathbf{k}) - \eta)n^2/2}$ and G is a counterexample to the statement. Let V_1, \dots, V_r be the parts of G and define $\alpha := (|V_1|/n, \dots, |V_r|/n)$. Then, for all Q_1 -optimal triples (r, ϕ, α') , we have that $\|\alpha - \alpha'\|_1 > \delta$. Lemma 7 and our assumption on G imply that there exists $\phi \in \Phi_1(r; \mathbf{k})$ such that

$$Q(\mathbf{k}) - \eta \leq \frac{\log_2 F(G; \mathbf{k})}{n^2/2} \leq q(r, \phi, \alpha) + \eta. \quad (4.2)$$

Claim 15 immediately gives a contradiction, completing the proof of the Theorem 5. \square

5. Concluding remarks

The referee of this paper asked if the cases where $F(\mathbf{k})$ was determined in [1] can be done using our optimisation problem. While the answer is in the affirmative, some claims from [1] are more conveniently derived by working with graphs rather than feasible solutions. For example, following [1] let us show that

$$\sum_{\ell \in [s]} \ell d_\ell \leq s \left(1 - \frac{1}{k-1} \right), \quad (5.1)$$

where $\mathbf{k} := (k, \dots, k)$ has length s , $(r, \phi, \alpha) \in \text{FEAS}_0(\mathbf{k})$ is an arbitrary feasible solution, and we define

$$d_\ell := 2 \sum_{\substack{ij \in \binom{[r]}{2} \\ |\phi(ij)| = \ell}} \alpha_i \alpha_j, \quad \text{for } \ell \in [s].$$

The shortest way is probably to consider the graph $G_{\phi, \alpha}(n)$ from the proof of Lemma 3. For $c \in [s]$, let H_c be the subgraph of $G_{\phi, \alpha}(n)$ spanned by pairs of parts (X_i, X_j) such that $c \in \phi(ij)$. Then H_c is K_k -free for all colours $c \in [s]$ and so Turán's theorem implies that $e(H_c) \leq (1 - \frac{1}{k-1})n^2/2$. Thus we have that, as $n \rightarrow \infty$,

$$\sum_{\ell \in [s]} \ell d_\ell = 2 \sum_{c \in [s]} \sum_{\substack{ij \in \binom{[r]}{2} \\ c \in \phi(ij)}} \alpha_i \alpha_j = 2 \sum_{c \in [s]} \frac{e(H_c) + O(n)}{n^2} \leq s \left(1 - \frac{1}{k-1}\right) + o(1),$$

which gives the claimed inequality (5.1). Interestingly, (5.1) and the trivial constraints $d_\ell \geq 0$ for $\ell \in [s]$ imply the sharp upper bound on $q(\phi, \alpha) = \sum_{\ell=1}^s d_\ell \log_2 \ell$ when $s \in \{2, 3\}$ and when $\mathbf{k} = (4, 4, 4, 4)$. (If $\mathbf{k} = (3, 3, 3, 3)$, then an additional constraint, analogous to (5.1), suffices to determine $Q(\mathbf{k})$, see [1].)

Unfortunately, the problem of (numerically) solving Problem Q_2 seems rather difficult even for moderately small \mathbf{k} . If we have a candidate pair (r, ϕ) , then the Lagrange Multiplier Method gives a linear program which either returns a best possible α for this (r, ϕ) in the interior of Δ^r , or implies that there is an optimal solution on the boundary so we can reduce r by one. This calculation can be efficiently implemented. However, the number of possible pairs (r, ϕ) becomes large very quickly. Here, the quest of replacing the crude bound $r < R(\mathbf{k})$ by a better one leads to the following Ramsey-type question. Namely, r can be bounded by $R_2(\mathbf{k}) - 1$, where we define $R_2(\mathbf{k})$ to be the smallest r such that for every choice of a *list-colouring* $\phi : \binom{[r]}{2} \rightarrow \binom{[s]}{2}$ there is $c \in [s]$ with $\phi^{-1}(c)$ containing a k_c -clique. Clearly, the definition would not change if we restrict ourselves to lists of size at least 2, so we can assume $r < R_2(\mathbf{k})$ in the statement of Problem Q_2 . The problem of estimating $R_2(\mathbf{k})$ runs into similar difficulties as those for the classical version $R(\mathbf{k})$. It is a special case of a parameter studied in [20], and seems to grow fast. For example, in [20] it was shown that $R_2(5, 5, 5) \geq 20$, which is already too large for a naïve enumeration of feasible ϕ by computer.

As we mentioned, the existence of the limit in (1.2) can be shown by an easy modification of the proof for the case $k_1 = \dots = k_s$ in [1]. In fact, there are two different proofs. The one that appears in the published version of [1] was suggested by an anonymous referee and uses an entropy inequality of Shearer to show that $\log F(n; \mathbf{k})/n^2$ is a non-increasing function of n .

The other proof, which was the original argument by Alon et al [1], is similar to our proof of Theorem 4. In our language, it can be sketched as follows. Fix a large N such that $\log_2 F(N; \mathbf{k})/\binom{N}{2}$ is close to the limit superior of (1.2). Take an ε -regular partition $V(G) = V_1 \cup \dots \cup V_m$ of an arbitrary \mathbf{k} -extremal order- N graph G with a ‘typical’ colouring σ . Let $\phi(ij)$ be the set of those colours $c \in [s]$ for which $\sigma^{-1}(c)[V_i, V_j]$ is an (ε, γ) -regular pair. As in Lemma 3, use this function $\phi : \binom{[m]}{2} \rightarrow 2^{[s]}$ with the uniform vector $\alpha = (1/m, \dots, 1/m)$ to produce graphs of order $n \rightarrow \infty$ with at least $2q(m, \phi, \alpha)n^{2/2-O(n)}$ valid colourings. Since $q(m, \phi, \alpha)$ can be made arbitrarily close to the limit superior of (1.2) by choosing small $\gamma \gg \varepsilon \gg 1/N$, the limit in (1.2) exists.

The latter proof can be adopted to prove Theorem 4 (by applying symmetrisation to reduce the triple (m, ϕ, α) to one with fewer than $R(\mathbf{k})$ parts). However, our proof (where the Regularity Lemma is applied after the symmetrisation) has the advantages of giving some explicit (although rather bad) bound on the rate of convergence in (1.2) and implying Theorem 5 as well.

Despite Theorem 5, there may be order- n graphs G with $F(G; \mathbf{k}) = 2^{(Q(\mathbf{k})+o(1))n^2/2}$ which are very far in edit distance from being complete multipartite. For example, if $\mathbf{k} = (4, 3)$, then one can take for G an equitable complete bipartite graph with parts $A \cup B$ and add any triangle-free graph into A (e.g. a blow-up of a pentagon which is far from being complete partite). Here, we can colour edges between A and B arbitrarily provided all edges inside A have colour 1. Thus $F(G; (4, 3)) \geq 2^{|A||B|} = 2^{\frac{1}{2}\binom{n}{2}+O(n)}$, while $Q((4, 3))$ is easily seen to be equal to $1/2$.

Interestingly, our follow-up results (in preparation) show that all $(4, 3)$ -extremal graphs of sufficiently large order n happen to be in fact 3-partite. For example, if $n = 2m + 1$ is odd (and large), then the unique extremal graph is $K_{m, m-1, 2}$. In order to illustrate how a small part can increase the number of colourings, let us show that

$$F(K_{m, m-1}; (4, 3)) \geq 2 \cdot 2^{m(m+1)} - 2^{m^2}, \quad (5.2)$$

that is, the number of $(4, 3)$ -valid colourings of $H := K_{m, m-1}$ is by factor $2 - o(1)$ larger than that for the Turán graph $K_{m+1, m}$. If H has parts $V_1 \cup V_2 \cup V_3$ with $|V_3| = 1$, then H has $2^{m(m+1)}$ colourings where $G[V_1 \cup V_3, V_2]$ is coloured arbitrarily while all edges between V_1 and V_3 have colour 1. Similarly we have $2^{m(m+1)}$ colourings where V_3 is ‘bundled’ with V_2 (and all edges between V_2 and V_3 get colour 1). All colourings that appear twice are exactly those that assign colour 1 to all edges incident to V_3 , so there are $2^{|V_1||V_2|} = 2^{m^2}$ of them, giving (5.2).

The above example shows that one can have parts of size $o(n)$ in Theorem 5 even for \mathbf{k} -extremal graphs. (These parts will correspond to zero entries of α in the limit.) Nonetheless, we conjecture that Theorem 2 captures all extremal graphs:

CONJECTURE 16. *For every $n, s \in \mathbb{N}$ and $\mathbf{k} \in \mathbb{N}^s$, every n -vertex \mathbf{k} -extremal graph is complete multipartite.*

In a future paper, we hope to provide a sufficient condition for this to be true for all $n \geq n_0(\mathbf{k})$ and apply the developed theory to solving the problem for new values of \mathbf{k} . We note that, in the different setting of forbidden cliques with prescribed colour patterns explored in [4], the corresponding version of Conjecture 16 holds in some cases.

REFERENCES

- [1] N. Alon, J. Balogh, P. Keevash and B. Sudakov, The number of edge colorings with no monochromatic cliques, *J. London Math. Soc.* **70** (2004), 273–288.
- [2] N. Alon and R. Yuster, The number of orientations having no fixed tournament, *Combinatorica* **26** (2006), 1–16.
- [3] J. Balogh, A remark on the number of edge colorings of graphs, *Europ. J. Comb.* **27** (2006), 565–573.
- [4] F. S. Benevides, C. Hoppen and R. M. Sampaio, Edge-colorings of graphs avoiding complete graphs with a prescribed coloring, preprint (arXiv:1605.08013).
- [5] P. Erdős, Some new applications of probability methods to combinatorial analysis and graph theory, Proceedings of the Fifth Southeastern Conference on Combinatorics, Graph Theory and Computing, Congress Numerantium X (1974), 39–51.
- [6] P. Erdős, Some of my favorite problems in various branches of combinatorics, *Matematiche (Catania)* **47** (1992), 231–240.

- [7] P. Erdős and G. Szekeres, A combinatorial problem in geometry, *Comp. Math.* **2** (1935), 463–470.
- [8] C. Hoppen, Y. Kohayakawa and H. Lefmann, Kneser colorings of uniform hypergraphs, *Elec. Notes in Disc. Math.* **34** (2009), 219–223.
- [9] C. Hoppen, Y. Kohayakawa and H. Lefmann, Edge colourings of graphs avoiding monochromatic matchings of a given size, *Comb. Prob. Comp.* **21** (2012), 203–218.
- [10] C. Hoppen, Y. Kohayakawa and H. Lefmann, Edge-colorings of graphs avoiding fixed monochromatic subgraphs with linear Turán number, *Europ. J. Comb.* **35** (2014), 354–373.
- [11] C. Hoppen and H. Lefmann, Edge-colorings avoiding a fixed matching with a prescribed color pattern, *Europ. J. Comb.* **47** (2015), 75–94.
- [12] C. Hoppen, H. Lefmann and K. Odermann, A coloring problem for intersecting vector spaces, *Disc. Math.* **339** (12) (2016), 2941–2954.
- [13] C. Hoppen, H. Lefmann, K. Odermann and J. Sanches, Edge-colorings avoiding fixed rainbow stars, *Elec. Notes in Disc. Math.* **50** (2015), 275–280.
- [14] J. Komlós and M. Simonovits, ‘Szemerédi’s regularity lemma and its applications to graph theory’ in *Combinatorics, Paul Erdős is Eighty*, D. Miklós, V. T. Sós and T. Szőni, Eds., vol. 2, Bolyai Math. Soc., 1996, pp. 295–352.
- [15] H. Lefmann, Y. Person, V. Rödl and M. Schacht, On colorings of hypergraphs without monochromatic Fano planes, *Comb. Prob. Comp.* **18** (2009), 803–818.
- [16] H. Lefmann, Y. Person and M. Schacht, A structural result for hypergraphs with many restricted edge colorings, *J. Comb.* **1** (2010), 441–475.
- [17] O. Pikhurko and Z. Yilma, The maximum number of K_3 -free and K_4 -free edge 4-colorings, *J. London Math. Soc.* **85** (2012), 593–615.
- [18] E. Szemerédi, ‘Regular partitions of graphs’ in *Proc. Colloq. Int. CNRS*, Paris, 1976, pp. 309–401.
- [19] P. Turán, On an extremal problem in graph theory (in Hungarian), *Mat. Fiz. Lapok* **48** (1941), 436–452.
- [20] X. Xu, Z. Shao, W. Su and Z. Li, Set-coloring of edges and multigraph Ramsey numbers, *Graphs. Comb.* **25** (6) (2009), 863–870.
- [21] R. Yuster, The number of edge colorings with no monochromatic triangle, *J. Graph Theory* **21** (1996), 441–452.
- [22] A. A. Zykov, On some properties of linear complexes (in Russian), *Mat. Sbornik N.S.* **24** (1949), 163–188.