

Some Exact Inducibility-type Results for Graphs via Flag Algebras

Levente Bodnár

Oleg Pikhurko

Submitted: Sep 11, 2025; Accepted: Mar 29, 2026; Published: May 22, 2026

© The authors. Released under the CC BY license (International 4.0).

Abstract

The (κ, ℓ) -edge-inducibility problem asks for the maximum number of κ -subsets inducing exactly ℓ edges that a graph of given order n can have. Using flag algebras and stability approach, we resolve this problem for all sufficiently large n (including a description of all extremal and almost extremal graphs) in eleven new non-trivial cases when $\kappa \leq 7$.

We also compute the F -inducibility constant (the asymptotically maximum density of induced copies of F in a graph of given order n) and obtain some corresponding structure results for three new graphs F with 5 vertices: the 3-edge star plus an isolated vertex, the 4-cycle plus an isolated vertex, and the 4-cycle with a pendant edge.

Mathematics Subject Classifications: 05C35

1 Introduction

Let κ and ℓ be given non-negative integers with $\ell \leq \binom{\kappa}{2}$. A (κ, ℓ) -graph is a graph having exactly κ vertices and ℓ edges. For a graph G , let $A_{\kappa, \ell}(G)$ denote the number of (κ, ℓ) -subgraphs, meaning induced subgraphs of G with κ vertices and ℓ edges. In other words, $A_{\kappa, \ell}(G)$ counts κ -subsets of $V(G)$ that span exactly ℓ edges in G . The (κ, ℓ) -edge-inducibility problem (or (κ, ℓ) -problem for short) asks for

$$A_{\kappa, \ell}(n) := \max\{A_{\kappa, \ell}(G) : n\text{-vertex graph } G\},$$

the maximum number of (κ, ℓ) -subgraphs that a graph with n vertices can have. It is natural to consider the normalised function

$$\lambda_{\kappa, \ell}(n) := \frac{A_{\kappa, \ell}(n)}{\binom{n}{\kappa}}, \quad \text{for } n \geq \kappa.$$

Mathematics Institute and DIMAP, University of Warwick, Coventry CV4 7AL, U.K.
(bodnalev@gmail.com, O.Pikhurko@warwick.ac.uk).

The standard averaging argument shows that the limit

$$\lambda_{\kappa,\ell} := \lim_{n \rightarrow \infty} \lambda_{\kappa,\ell}(n) \tag{1}$$

exists, see e.g. [46, Lemma 2.2]. We refer to the value of the limit as the *edge-inducibility constant* of (κ, ℓ) .

Observe that, by replacing all graphs with their complements, the value of $\lambda_{\kappa,\ell}(n)$ will not change if we replace ℓ by $\binom{\kappa}{2} - \ell$. Trivially, it holds that $\lambda_{\kappa,\ell}(n) = 1$ if and only if $\ell = 0$ or $\binom{\kappa}{2}$. Also, the special case $\ell = 1$, which is equivalent to the inducibility problem for the unique up to isomorphism $(\kappa, 1)$ -graph, is resolved through the results in [14, 31, 39–41], where in particular Liu, Mubayi and Reiher [41, Theorem 1.13] obtained an explicit formula for $\lambda_{\kappa,1}$ valid for every $\kappa \geq 4$. Thus we restrict ourselves to $2 \leq \ell \leq \binom{\kappa}{2}/2$ only.

This problem was recently introduced by Alon, Hefetz, Krivelevich and Tyomkyn [3] and has received considerable attention. In particular, the Edge-Statistics Conjecture [3, Conjecture 1.1] that $\lambda_{\kappa,\ell} \leq 1/e + o_\kappa(1)$ was fully resolved by a sequence of papers by Kwan, Sudakov and Tran [36], Martinsson, Mousset, Noever and Trujić [43] and Fox and Sauerman [26]. Two other conjectures of Alon et al [3, Conjectures 6.1 and 6.2] on stronger bounds in the case when ℓ is well separated from 0 were resolved by Kwan, Sudakov and Tran [36] and Kwan and Sauerman [37]. Hypergraph versions of these results were very recently obtained by Jain, Kwan, Mubayi and Tran [32]. Motivated by these results, versions of this problem for hypercubes were studied by Alon, Axenovich and Goldwasser [1] and the authors [10].

Here, we systematically investigate the cases of the edge-inducibility problem when $\kappa \leq 7$ using the flag algebra method. Our aim is not only to find the value of the edge-inducibility constant $\lambda_{\kappa,\ell}$ (that is, to determine $\Lambda_{\kappa,\ell}(n)$ within an additive $o(n^\kappa)$ error term as $n \rightarrow \infty$) but also to prove perfect stability. The formal definition of this property will appear in Section 2.2; informally speaking, here it means that there is a blowup pattern B and a constant C such that every graph G with $n \geq C$ vertices can be made into a blowup of B by changing at most $C(\lambda_{\kappa,\ell}(n) - \lambda_{\kappa,\ell}(G))n^2$ adjacencies. In particular, every n -vertex graph G which maximizes the number of (κ, ℓ) -subgraphs is a blowup of B ; thus the determination of $\lambda_{\kappa,\ell}(n)$ (and of the set of extremal graphs) amounts to finding optimal part sizes of a B -blowup (that is, to maximising some polynomial over non-negative integers summing up to n). If this analytic problem is resolved with a description of all optimal part ratios, then perfect stability implies Erdős–Simonovits stability [22, 52] that aims to describe the structure of every graph G of order $n \rightarrow \infty$ with $\lambda_{\kappa,\ell}(G) = \lambda_{\kappa,\ell} + o(1)$ up to $o(n^2)$ adjacencies (see Section 2.2).

We were able to determine the value of $\lambda_{\kappa,\ell}$ for eleven new pairs (κ, ℓ) , also showing that perfect stability holds in each solved case except for $(\kappa, \ell) = (4, 3)$. Table 1 summarises our new findings, where we use the following notation for constructions: $T_{n_0, \dots, n_{m-1}}$ is the complete m -partite graph with parts of sizes n_0, \dots, n_{m-1} ; K_n is the clique with n vertices, $F + H$ is the vertex-disjoint union of graphs F and H , $mF := F + \dots + F$ is the

union of m copies of F , and the constants are

$$\alpha_1 := \frac{1}{16} \left(9 - \sqrt{17} \right), \tag{2}$$

$$\alpha_2 := \frac{1}{2} \left(1 - \sqrt{\frac{1}{3} \left(2\sqrt{10} - 5 \right)} \right), \tag{3}$$

$$\alpha_3 := \frac{1}{2} \left(1 - \sqrt{\frac{1}{3} \left(\frac{4}{5} \sqrt{10} - 1 \right)} \right). \tag{4}$$

We refer the reader to Section 3 for the formal statements and further details.

(κ, ℓ)	Construction	$\lambda_{\kappa, \ell}$	Stability	Reference
(4, 2)	$2T_{n/6, n/6, n/6}$	$\frac{1}{2}$	Perfect stability	Theorem 3
(4, 3)	$2K_{n/2}$ or $T_{n/2, n/2}$	$\frac{1}{2}$	See Section 4	Theorem 4
(5, 2)	$3T_{n/9, n/9, n/9}$	$\frac{280}{3^6}$	Perfect stability	Theorem 6
(5, 3)	$T_{\alpha_1 n, \alpha_1 n} + K_{(1-2\alpha_1)n}$	$\frac{255\sqrt{17}-535}{2^{10}}$	Perfect stability	Theorem 7
(5, 4)	$2K_{n/2}$	$\frac{5}{8}$	Perfect stability	Theorem 8
(6, 4)	$3K_{n/3}$	$\frac{40}{3^4}$	Perfect stability	Theorem 9
(6, 5)	$T_{\alpha_2 n, (1-\alpha_2)n}$	$\frac{10\sqrt{10}-28}{9}$	Perfect stability	Theorem 10
(6, 7)	$2K_{n/2}$	$\frac{15}{2^5}$	Perfect stability	Theorem 11
(7, 6)	$T_{\alpha_3 n, (1-\alpha_3)n}$	$\frac{28\sqrt{10}-35}{135}$	Perfect stability	Theorem 12
(7, 9)	$2K_{n/2}$	$\frac{35}{2^6}$	Perfect stability	Theorem 13
(7, 10)	$T_{n/3, 2n/3}$	$\frac{28}{3^4}$	Perfect stability	Theorem 14

Table 1: New values of edge-inducibility constants.

The case $(\kappa, \ell) = (4, 3)$ turned out to be special in many aspects (see Section 4 for details). This problem is self-complementary and there are two types of extremal graphs, namely, complete bipartite graphs or two disjoint cliques (with part sizes $n/2 + o(n)$). Interestingly, when we change the adjacency of a pair xy between the parts then the number of $(4, 3)$ -subgraphs changes by only $O(n)$, instead of a positive fraction of $\binom{n-2}{2} = \Theta(n^2)$ 4-sets containing xy . (In fact, if we start with parts of sizes exactly $\lfloor n/2 \rfloor$ and $\lceil n/2 \rceil$ then changing one adjacency across strictly increases the number of $(4, 3)$ -subgraphs.) It follows that perfect stability does not hold for this problem, even if we consider the weaker version where we allow finitely many possible patterns. However, we can prove a version of Erdős–Simonovits stability in Lemma 18: if an n -vertex graph G satisfies

$\lambda_{4,3}(G) = \lambda_{4,3} + o(1)$ then G is $o(n^2)$ -close to one of the above two constructions. The optimal part sizes can be computed exactly and they deviate from $n/2$ by $\sqrt{3n}/2 + O(1)$. This extra imbalance adds a “drift” that penalises wrong pairs across and suffices for us to prove in Theorem 19 that every extremal graph of large order n has no wrong pairs at all.

We also obtained new results on the following graph inducibility problem. (See also Section 6 for a short discussion on the related semi-inducibility problem.) For graphs F and G with $\kappa \leq n$ vertices respectively, let $P(F, G)$ denote the number of κ -subsets of $V(G)$ that span a graph isomorphic to F and let $p(F, G) := P(F, G)/\binom{n}{\kappa}$ be the *density* of F in G . The *inducibility problem* for a graph F asks for $\lambda_F(n)$, the maximum of $\lambda_F(G) := p(F, G)$ over all graphs G with n vertices. As before, it is easy to show that the limit $\lambda_F := \lim_{n \rightarrow \infty} \lambda_F(n)$ exists; we call it the *inducibility constant* of F .

The inducibility problem has drawn a great amount of interest since it was introduced by Pippenger and Golumbic [48] in 1975. For some sample of results, see e.g. [5, 12–14, 23, 29–31, 35, 38, 46, 53, 56, 57].

If F is complete partite then the result by Brown and Sidorenko [14] implies that, in order to determine the value of $\lambda_F(n)$, it is enough to consider complete partite graphs on $[n]$ and the problem in the limit reduces to some analytic-type optimisation on the space of part ratios. If the latter is fully solved, with the description of all extremal ratios, then the method of Liu, the second author, Sharifzadeh and Staden [39] can often be applied to decide if perfect stability holds or not. So we exclude complete partite F from our consideration. Since the inducibility constant λ_F does not change if we replace F by its complement, it is enough to consider only one graph from each complementary pair.

Each 3-vertex graph F or its complement is complete partite, so we exclude these (as they are covered by the above result of Brown and Sidorenko [14]).

All 4-vertex graphs F were resolved by the results in [13, 14, 24, 31] except when $F = P_4$ is the 4-vertex path. The best known lower bound $\lambda_{P_4} \geq 1173/5824 = 0.2014\dots$ is due to Even-Zohar and Linial [23] while the best known upper bound $\lambda_{P_4} \leq 0.204513\dots$ comes from flag algebras.

Here we look at 5-vertex graphs F . For notational convenience, we assume by default that the vertex set of F is $\{0, 1, 2, 3, 4\}$ while xy means a pair $\{x, y\}$. Even-Zohar and Linial [23, Table 2] produced a summary of known and new results for 5-vertex graphs F , in particular providing numerical upper bounds on λ_F coming from flag algebra calculations in the cases when the exact value was not known.

After the appearance of [23], two new 5-vertex cases (when F is not complete partite), namely when $E(F) = \{01, 12\}$ (the 2-edge path plus 2 isolated vertices) and $E(F) = \{01, 12, 23, 24\}$ (the “Y-graph”) were fully resolved for all large n (including perfect stability) by the second author, Sliacan and Tyros [46]. The authors of [46] also tried to solve some other open 5-vertex cases using Emil Vaughan’s package `flagmatic` but were not able to.

Using the first author’s new package `FlagAlgebraToolbox`, see [8], we are able to determine the value of λ_F for three new 5-vertex graphs F : the 3-star $T_{3,1}$ plus an isolated vertex, the 4-cycle $T_{2,2}$ plus an isolated vertex, and the 4-cycle $T_{2,2}$ with a pendant edge

attached. These results are summarised in Table 2, where we let $\beta := (3 + \sqrt{3})/12$ while $R(G, p)$ means a “typical” spanning subgraph of G when each edge of G is kept with probability p independently of the other edges. (See Section 3 for the formal statements and further details.)

Edges of F	Construction	λ_F	Stability	Reference
01, 02, 03	$T_{\beta n, \beta n} + T_{(1/2-\beta)n, (1/2-\beta)n}$	$\frac{5}{24}$	Perfect stability	Theorem 15
01, 12, 23, 30	$T_{\beta n, \beta n} + T_{(1/2-\beta)n, (1/2-\beta)n}$	$\frac{5}{32}$	Perfect stability	Theorem 16
01, 12, 23, 30, 04	$R(T_{n/2, n/2}, \frac{5}{6})$	$\frac{5^6}{2^8 \cdot 3^5}$	See Section 5	Theorem 17

Table 2: New values of the inducibility constant λ_F for a graph F with $V(F) = \{0, 1, 2, 3, 4\}$.

Note that, in the first two cases, the structure of large extremal graphs (which are disjoint unions of 2 complete bipartite graphs) and the limiting part ratios happen to be the same. Also, we can prove perfect stability in both cases.

In the last case when F is the 4-cycle with a pendant edge, we prove in Theorem 21 via extra arguments that every almost extremal graph G of order $n \rightarrow \infty$ admits a balanced vertex partition $V(F) = V_0 \cup V_1$ such that each part spans $o(n^2)$ edges while the induced bipartite graph $G[V_0, V_1]$ is $5/6$ -quasirandom, thus obtaining a good characterisation of almost extremal graphs. This is in an interesting contrast with the result of Jain, Michelen and Wei [33] that (non-bipartite) Erdős–Rényi random graphs $R(K_n, p)$ for constant $p \in (0, 1)$ cannot be almost extremal for any graph inducibility problem. Regarding extremal graphs, that is, n -vertex graphs G with $\lambda_F(G) = \lambda_F(n)$, we additionally prove in Theorem 24 that, for all large n , each such graph G admits a vertex partition $V(G) = V_0 \cup V_1$ into two independent sets and every vertex of G has $(5/12 + o(1))n$ neighbours in the other part. This reduces the F -inducibility problem for large n to its bipartite version (modulo the issue of finding optimal part sizes). Resolving this bipartite problem exactly seems challenging and we limit ourselves to the above partial description of extremal graphs.

2 Preliminaries

In this section, we present here some definitions and auxiliary results.

Let \mathbb{R} denote the set of reals. Let \mathbb{N} denote the set of non-negative integers and, for $n \in \mathbb{N}$, we define $[n] := \{0, \dots, n - 1\}$. Note that we start indexing from 0, merely to be consistent with the same convention as in our code. If the meaning is clear, we may abbreviate an unordered pair $\{u, w\}$ as uw , including the case when u and w are single-digit numbers. For a set X and an integer $\kappa \geq 0$, the set of all κ -subsets of X is denoted by $\binom{X}{\kappa}$. Also, $a = b \pm \varepsilon$ for $a, b \in \mathbb{R}$ means $b - \varepsilon \leq a \leq b + \varepsilon$. We may omit ceiling/floor signs when they are not essential.

A *pattern* $B = (V(B), E(B))$ is a graph where we additionally allow loops on some vertices (but we do not allow multiple edges). Its *order* is $v(B) := |V(B)|$. We write $\{u, u\} \in E(B)$ (or $uu \in E(B)$) to indicate that there is a loop on a vertex u . A *pattern automorphism* is a bijection $f : V(B) \rightarrow V(B)$ such that, for every $u, w \in V(B)$, we have $\{u, w\} \in E(B)$ if and only if $\{f(u), f(w)\} \in E(B)$; thus it is an automorphism of the underlying graph that also preserves loops and non-loops. The B -*neighbourhood* $\Gamma_B(u)$ of a vertex $u \in V(B)$ is the set $\{w \in V(B) : \{u, w\} \in E(B)\}$. Note that u itself is included into $\Gamma_B(u)$ if and only if u is a loop in B . The *degree* of u is $\deg_B(u) := |\Gamma_B(u)|$.

Of course, the above definitions also apply to graphs (which are patterns without loops). Also, for a graph F , its *complement* is $\bar{F} := (V(F), (V(F) \choose 2) \setminus E(F))$. For $X \subseteq V(F)$, the subgraph *induced* by X is $F[X] := (X, \{uw \in E(F) : u, w \in X\})$ and, for disjoint $X, Y \subseteq V(F)$, we denote $F[X, Y] := \{(u, w) \in X \times Y : uw \in E(F)\}$.

Recall that K_n denotes the complete graph with n vertices, P_n is the n -vertex path, and $T_{n_0, \dots, n_{m-1}}$ denotes the complete m -partite graph with parts of sizes n_0, \dots, n_{m-1} . We may refer to $T_{2,1} \cong P_3$ as the *cherry*. Also, $F + H$ denotes the union of vertex-disjoint copies of graphs F and H . When we define a small graph/pattern, we may write it as (m, E) , meaning that the vertex set is $[m]$. For example, we denote the 4-vertex path as $(4, \{01, 12, 23\})$ and 2 isolated loops as $(2, \{00, 11\})$.

Let B be a pattern with vertex set $[m]$. For pairwise disjoint sets V_0, \dots, V_{m-1} (with some possibly empty), the *blowup* $B(V_0, \dots, V_{m-1})$ of B is the graph (without loops) on $V = \cup_{i=0}^{m-1} V_i$ where distinct $x \in V_i$ and $y \in V_j$ are adjacent if and only if $\{i, j\} \in E(B)$. In particular, a part V_i spans a clique (resp. an independent set) if i is (resp. is not) a loop of B . Let $B()$ denote the family of all blowups of B . A *homomorphism* from a graph F to a pattern B is a (not necessarily injective) function $f : V(F) \rightarrow V(B)$ such that for every distinct $x, y \in V(F)$ it holds that $\{x, y\} \in E(F)$ if and only if $\{f(x), f(y)\} \in E(B)$. Thus homomorphisms from F to B are exactly possible assignments of vertices of F to the parts of (sufficiently large) blowups of B that give induced copies of F . For graphs F and G , a function $f : V(F) \rightarrow V(G)$ is an *embedding* of F into G (written as $f : F \hookrightarrow G$) if f is injective and preserves both edges and non-edges; that is, f gives an isomorphism of F on its image.

To avoid confusion, let us repeat that our definition of homomorphism requires that both edges and non-edges are preserved. (In the rare cases when we have to consider maps that are required to preserve edges only, we will use the term *non-induced homomorphism*.) Thus, for example, an embedding can be defined as an injective homomorphism.

The *edit distance* $\delta_{\text{edit}}(G, H)$ between two graphs G and H of the same order is the minimum value of $|E(G) \Delta f(E(H))|$ over all bijections $f : V(H) \rightarrow V(G)$; in other words, it is the smallest number of *edits* (changes in adjacency) we have to do in one graph to make it isomorphic to the other. The distance from a graph G to a graph family \mathcal{G} is

$$\delta_{\text{edit}}(G, \mathcal{G}) := \min\{\delta_{\text{edit}}(G, H) : H \in \mathcal{G}, v(H) = v(G)\}.$$

We will be mostly interested in the case when $\mathcal{G} = B()$ is the family of all blowups of B ; thus $\delta_{\text{edit}}(G, B())$ is the smallest number of edits in G needed to make it a blowup of B .

For $m \in \mathbb{N}$, let \mathcal{F}_m^0 be the family of graphs (without loops) of order m consisting of one representative from each isomorphism class. For graphs F and G with κ and n vertices respectively, let $P(F, G)$ be the number of κ -subsets $X \subseteq V(G)$ that induce a subgraph isomorphic to F in G .

Suppose that $\kappa \leq n$. Then we let $p(F, G) := P(F, G) / \binom{n}{\kappa}$, calling it the (*induced*) *density* of F in G . Occasionally, it will be more convenient to work with the *embedding density* $t(F, G)$ which is defined as the probability that a random injective map $V(F) \rightarrow V(G)$ is an embedding, that is, preserves both edges and non-edges. Thus, informally speaking, we look at vertex labelled copies of F in G . One can easily transfer between these two densities using that

$$t(F, G) = \frac{|\text{aut}(F)|}{\kappa!} p(F, G), \quad (5)$$

where $\text{aut}(F)$ denotes the automorphism group of F .

We call a sequence of growing bipartite graphs $(G_n)_{n \in \mathbb{N}}$ with almost equal parts *c-quasirandom* if for every bipartite graph F the *bipartite non-induced homomorphism density* $t_{\text{bip}}(F, G)$ of F in G_n (which is the probability that a random part-preserving map $V(F) \rightarrow V(G_n)$ sends edges of F to edges of G) is $c^{|E(F)|} + o(1)$ as $n \rightarrow \infty$. Note that this is the value we observe in a typical c -random subgraph of $T_{n,n}$. As it can be shown by an easy adaptation of the classical proof of Chung, Graham and Wilson [17] (with details spelled in e.g. [18, Lemma 14]), it is enough to check this property only when F is the edge and the 4-cycle.

Next, let us introduce some notation that will allow to treat the problems studied in this paper in a uniform way. Given an integer $\kappa \geq 2$ and a function $\gamma : \mathcal{F}_\kappa^0 \rightarrow \mathbb{R}$, we consider the following function on graphs:

$$\Lambda_\gamma(G) := \sum_{F \in \mathcal{F}_\kappa^0} \gamma(F) P(F, G), \quad \text{for a graph } G,$$

and its density version

$$\lambda_\gamma(G) := \frac{\Lambda_\gamma(G)}{\binom{v(G)}{\kappa}} = \sum_{F \in \mathcal{F}_\kappa^0} \gamma(F) p(F, G), \quad \text{for a graph } G \text{ with } v(G) \geq \kappa.$$

We consider the corresponding extremal problem where we maximise these objective functions on order- n graphs, namely we are interested in

$$\Lambda_\gamma(n) := \max\{\Lambda_\gamma(G) : v(G) = n\}, \quad \text{for } n \in \mathbb{N},$$

and its density version $\lambda_\gamma(n) := \Lambda_\gamma(n) / \binom{n}{\kappa}$ for $n \geq \kappa$. Of course, these two functions are equivalent and interchangeable. We tend to use $\lambda_\gamma(\cdot)$ but switch to $\Lambda_\gamma(\cdot)$ when the involved quantity seems to be better to write or understand as some counting. Define $\lambda_\gamma := \lim_{n \rightarrow \infty} \lambda_\gamma(n)$, where the existence of the limit follows from e.g. [46, Lemma 2.2].

Let us observe that the problems studied here can be represented this way. For the (κ, ℓ) -edge-inducibility problem, we use the same κ and let γ be 1 on every graph with

exactly ℓ edges and be 0 otherwise. For the F -inducibility problem, we let $\kappa := v(F)$ and let γ take the value 1 on F (or, more precisely, the unique graph in \mathcal{F}_κ^0 isomorphic to F) and 0 on any other graph in \mathcal{F}_κ^0 .

For a graph G and a vertex u of G , let $\Lambda_\gamma(G, u)$ be the sum over all κ -subsets $X \subseteq V(G)$ that contain u of the value of γ evaluated at the induced subgraph $G[X]$ (more precisely, we take $\gamma(F)$ for the unique $F \in \mathcal{F}_\kappa^0$ isomorphic to $G[X]$). For example, it holds for any graph G that $\Lambda_\gamma(G) = (1/\kappa) \sum_{u \in V(G)} \Lambda_\gamma(G, u)$ and thus $(1/\kappa)\Lambda_\gamma(G, u)$ can be considered as the contribution of a vertex u to the global value $\Lambda_\gamma(G)$. For $v(G) \geq \kappa$, we also define $\lambda_\gamma(G, u) := \Lambda_\gamma(G, u) / \binom{n-1}{\kappa-1}$ to be its normalised version. For this normalisation, it holds that $\lambda_\gamma(G)$ is the average of $\lambda_\gamma(G, u)$ over $u \in V(G)$:

$$\lambda_\gamma(G) = \frac{1}{v(G)} \sum_{u \in V(G)} \lambda_\gamma(G, u), \quad \text{any graph } G \text{ with } v(G) \geq \kappa. \quad (6)$$

The $(m-1)$ -dimensional simplex is

$$\mathbb{S}_m := \{(x_0, \dots, x_{m-1}) \in \mathbb{R}^m : x_0 + \dots + x_{m-1} = 1 \text{ and } \forall i \in [m] \ x_i \geq 0\}. \quad (7)$$

For $\mathbf{x} = (x_0, \dots, x_{m-1}) \in \mathbb{S}_m$ and a pattern B on $[m]$, let $\lambda_\gamma(B(\mathbf{x}))$ be the limit as $n \rightarrow \infty$ of $\lambda_\gamma(B(V_0, \dots, V_{m-1}))$, where $|V_i| = (x_i + o(1))n$ for $i \in [m]$. This is a continuous function on \mathbb{S}_m (in fact, a polynomial). Let

$$\lambda_\gamma(B()) = \sup\{\lambda_\gamma(B(\mathbf{x})) : \mathbf{x} \in \mathbb{S}_m\}.$$

By the compactness of \mathbb{S}_m the supremum is attained by at least one $\mathbf{x} \in \mathbb{S}_m$; such vectors will be called (γ, B) -maximisers. We call the pattern B γ -minimal if, for every pattern B' (of order $m-1$) obtained from B by removing one vertex, it holds that $\lambda_\gamma(B'()) < \lambda_\gamma(B())$. By compactness and continuity, this holds if and only if the (closed) set of (γ, B) -optimal vectors is disjoint from the boundary of the simplex \mathbb{S}_m . The pattern B is called γ -optimal if $\lambda_\gamma = \lambda_\gamma(B())$, that is, we can attain the asymptotically optimal constant λ_γ by some blowups of B . If γ and/or B are understood, we may omit them from the above notation.

Asymptotic notation, such as $o(1)$, is taken with respect to $n \rightarrow \infty$ (where n is usually the order of the unknown graph G); the constants hidden in it may depend on κ and γ but not on any other parameters. We call a sequence of graphs $(G_n)_{n \in \mathbb{N}}$ with strictly increasing orders *almost γ -extremal* (resp. *almost c -regular*) if $\lambda_\gamma(G_n) = \lambda_\gamma + o(1)$ as $n \rightarrow \infty$ (resp. for every $\varepsilon > 0$ there is n_0 such that, for every $n \geq n_0$, at least $(1 - \varepsilon)v(G)$ vertices u of G satisfy $\deg_G(u) = (c \pm \varepsilon)v(G_n)$).

2.1 Flag algebras

Since the flag algebra approach of Razborov [49] is well established by now (and is described in detail in e.g. [4, 20, 28, 34, 50]), here we just give a bare minimum of definitions needed to define what a flag algebra certificate contains.

Let τ be a *type*, that is, is a graph with vertex set $[q]$ for some $q \in \mathbb{N}$; we view τ as having all its q vertices labelled. A τ -flag is a pair (F, f) where F is a graph and $f : [q] \rightarrow$

$V(F)$ is an embedding of τ into F (that is, an injection that preserves edges and non-edges). We view a τ -flag as a partially labelled graph where the labelled vertices, called *roots*, induce a copy of τ . For τ -flags (F, f) and (H, h) , let $P((F, f), (H, h))$ be the number of $(v(F) - q)$ -subsets X of $V(H) \setminus h([q])$ such that the τ -flag $(H[X \cup h([q])], h)$ is isomorphic to (F, f) , where isomorphisms between τ -flags have to preserve each labelled root (in addition to preserving edges and non-edges). If $v(F) \leq v(H)$, then the corresponding (*flag*) *density* is

$$p((F, f), (H, h)) := \frac{P((F, f), (H, h))}{\binom{v(H) - q}{v(F) - q}}.$$

For $s \geq q$, let \mathcal{F}_s^τ be the set of all τ -flags with s vertices up to isomorphism. (This is consistent with our previous notation \mathcal{F}_s^0 since we let 0 denote the empty type.) We fix, once and for all, an ordering of \mathcal{F}_s^τ to be used when we have some vectors or matrices indexed by \mathcal{F}_s^τ . For a τ -flag (H, h) with $v(H) \geq s$, let

$$\mathbf{V}_{(H, h)}^{\tau, s} := (P((F, f), (H, h)))_{(F, f) \in \mathcal{F}_s^\tau} \quad \text{and} \quad \mathbf{v}_{(H, h)}^{\tau, s} := (p((F, f), (H, h)))_{(F, f) \in \mathcal{F}_s^\tau} \quad (8)$$

be the (column) vectors listing the counts and densities of all s -vertex τ -flags in (H, h) (in the fixed ordering of \mathcal{F}_s^τ). Clearly, the entries of $\mathbf{v}_{(H, h)}^{\tau, s}$ are non-negative and sum up to 1.

Now, we can present the information that is contained in a flag algebra certificate of an upper bound $\lambda_\gamma \leq u$ for given $\gamma : \mathcal{F}_\kappa^0 \rightarrow \mathbb{R}$ and $u \in \mathbb{R}$. The certificate lists an integer $N \geq \kappa$ and, for each type τ on $[q]$ with $1 \leq q \leq N - 2$ and $q + N$ even, a positive semi-definite matrix X^τ whose rows and columns are indexed by \mathcal{F}_s^τ , where $s := (N + q)/2$. More precisely, it will be the case that, from each equivalence class \mathcal{C} of types under isomorphism as unlabelled graphs, we list only one representative $\tau \in \mathcal{C}$ and its matrix X^τ , effectively using the all-zero matrix for every other type in \mathcal{C} . Also, for every $F \in \mathcal{F}_N^0$, we list a real coefficient $c_F \geq 0$ (called the *slack* at F) such that the following identity holds for every graph G of order $n \rightarrow \infty$:

$$\begin{aligned} & \sum_{F \in \mathcal{F}_N^0} (u - \lambda_\gamma(F)) P(F, G) \\ &= \sum_{\substack{1 \leq q \leq N - 2 \\ q \equiv N \pmod{2}}} \sum_{f: [q] \hookrightarrow V(G)} \left(\mathbf{V}_{(G, f)}^{\tau, s} \right)^T X^\tau \mathbf{V}_{(G, f)}^{\tau, s} + \sum_{F \in \mathcal{F}_N^0} c_F P(F, G) + O(n^{N-1}), \end{aligned} \quad (9)$$

where, in the inner sum, $s := (N + q)/2$ and $\tau := ([q], f^{-1}(E(G)))$ is the graph on $[q]$ such that the injection $f : [q] \hookrightarrow V(G)$ is an embedding of τ into G . Note that the right-hand side of (9), apart from the error term, is non-negative (since all matrices X^τ are positive semi-definite and all slacks are non-negative by our assumptions). On the other hand, the left-hand side is exactly $(u - \lambda_\gamma(G)) \binom{n}{N}$. So the identity in (9) indeed proves that $\lambda_\gamma \leq u$.

Let us remark that, if we take any square matrices X^τ of the appropriate dimensions, then (for given $u \in \mathbb{R}$) there is a unique choice of the slack coefficients c_F (possibly negative) that makes (9) hold for every G . In brief, if we fix any type τ with $q \leq N - 2$

vertices, $q \equiv N \pmod{2}$, and two τ -flags $F_1, F_2 \in \mathcal{F}_s^\tau$, where $s := (N + q)/2$, then the sum $\sum_f P(F_1, (G, f))P(F_2, (G, f))$ over all embeddings f of τ into an order- n graph G can be expressed, up to an additive $O(n^{N-1})$ error term, as a linear combination (whose coefficients are independent of n) of counts of N -vertex graphs in G . Indeed, this sum counts the number of pairs of copies of F_1 and F_2 in G sharing the same q roots. Each such pair uses at most $2s - q = N$ vertices while the contribution of each N -set $X \subseteq V(G)$ depends only on the isomorphism class of the subgraph it induces in G . It follows that, for given N , the best possible upper bound u that can be proved via (9) is the value of an explicit (although usually very large) semi-definite program.

We will need the following lemma which states, roughly speaking, that the typical vectors of τ -rooted densities in every almost extremal graph have to be close to the zero eigenspace of X^τ .

Lemma 1. *Suppose that, for some N , we have a flag algebra certificate proving that $\lambda_\gamma \leq u$ as in (9). Let τ be any type present in the certificate, say with $V(\tau) = [q]$, and let $s := (N + q)/2$. (Thus s is an integer and $s \geq q + 1$). Then for every $\varepsilon > 0$ there are $\delta > 0$ and n_0 such that, for every graph G with $n \geq n_0$ vertices and $\lambda_\gamma(G) \geq u - \delta$, there are at most εn^q embeddings $f : \tau \hookrightarrow G$ such that*

$$\left\| X^\tau \mathbf{v}_{(G,f)}^{\tau,s} \right\|_\infty \geq \varepsilon, \tag{10}$$

where the vector $\mathbf{v}_{(G,f)}^{\tau,s}$, as defined in (8), lists the densities of s -vertex τ -flags in (G, f) .

Proof. Since the matrix X^τ is positive semi-definite, we have that $\mathbf{x}^T X^\tau \mathbf{x} = 0$ if and only if $X^\tau \mathbf{x}$ is the zero vector. By the compactness of \mathbb{S}_m , $m := |\mathcal{F}_s^\tau|$, and the continuity of the function that maps \mathbf{x} to $\mathbf{x}^T X^\tau \mathbf{x}$, there is $c > 0$ such that every vector $\mathbf{x} \in \mathbb{S}_m$ with $\|X^\tau \mathbf{x}\|_\infty \geq \varepsilon$ satisfies $\mathbf{x}^T X^\tau \mathbf{x} > c$.

Let us show that, for example, $\delta := \varepsilon c / N!$ satisfies the lemma if n_0 is sufficiently large. Take any graph G of sufficiently large order n with $\lambda_\gamma(G) \geq \lambda_\gamma - \delta$. Thus the left-hand side of (9) is at most $\delta \binom{n}{N}$. Every embedding $f : \tau \hookrightarrow G$ for which (10) holds contributes at least $c \binom{n-q}{s-q}^2$ to the right-hand side of (9). By the non-negativity of all other terms in the right-hand side of (9), this identity implies that the number of such embeddings f is at most $(\delta \binom{n}{N} + O(n^{N-1})) / (c \binom{n-q}{s-q}^2) < \varepsilon n^q$, as desired. \square

2.2 Perfect stability

Let $\kappa \geq 2$ and $\gamma : \mathcal{F}_\kappa^0 \rightarrow \mathbb{R}$ be given. We define two notions of stability (namely, Erdős–Simonovits stability and perfect stability) and present the sufficient condition for perfect stability from [46] that can be automatically verified by computer. There is one problem solved by us, namely the (4, 3)-edge-inducibility problem with two different optimal patterns, for which only the former type of stability could hold. So we allow the pattern in our definition of Erdős–Simonovits stability to depend on G (so that our definition applies to the (4, 3)-problem). Namely, we call the λ_γ -problem *Erdős–Simonovits stable* if for every $\varepsilon > 0$ there are $\delta > 0$ and n_0 such that if G is a graph with $n \geq n_0$

vertices and $\lambda_\gamma(G) \geq \lambda_\gamma - \delta$ then there is a γ -optimal and γ -minimal pattern B with $\delta_{\text{edit}}(G, B()) \leq \varepsilon \binom{n}{2}$. Recall that the last inequality means that there is a partition $V(G) = V_0 \cup \dots \cup V_{m-1}$ with $m := v(B)$ such that

$$|E(G) \Delta E(B(V_0, \dots, V_{m-1}))| \leq \varepsilon \binom{n}{2}. \quad (11)$$

Of course, if there is a unique γ -optimal and γ -minimal pattern B up to isomorphism and the (γ, B) -optimal vector in \mathbb{S}_m is unique (up to an automorphism of the pattern B) then our definition implies the more common formulation of Erdős–Simonovits stability that any two graphs G and G' of the same order $n \rightarrow \infty$ with both $\lambda_\gamma(G)$ and $\lambda_\gamma(G')$ being $\lambda_\gamma + o(1)$ are $o(n^2)$ -close to each other in the edit distance. This property is very useful as the first step towards characterizing graphs of sufficiently large order n with $\lambda_\gamma(G) = \lambda_\gamma(n)$, see for example the proof of Theorem 2 here. This approach was pioneered by Erdős [22] and Simonovits [52].

Perfect stability is a stronger property which, roughly speaking, states there is a constant C so that (11) holds for any function $\varepsilon(n) \geq 0$ with $\delta := \varepsilon/C$. Following [46], we call the λ_γ -problem *perfectly B -stable* for a pattern B if there is $C > 0$ such that for every graph G of order $n \geq C$ we have

$$\delta_{\text{edit}}(G, B()) \leq C (\lambda_\gamma(n) - \lambda_\gamma(G)) n^2. \quad (12)$$

In particular, it follows that, for all $n \geq C$, every order- n graph G with $\lambda_\gamma(G) = \lambda_\gamma(n)$ is a blowup of B and then the problem reduces to just maximising an explicit polynomial of degree at most κ over (integer) part sizes summing up to n .

A version of perfect stability for the Turán problem for K_t , that is, for maximising the number of edges in a K_t -free graph of given order n (with $B = K_{t-1}$), was proved by Füredi [27] while Roberts and Scott [51] extended this result to forbidding any colour critical graph. The perfect stability of some hypergraph Turán problems was established by Norin and Yepremyan [44, 45]. The second author, Sliacan and Tyros [46, Theorem 7.1] presented a sufficient condition for a flag algebra proof to give perfect stability and successfully applied it to a number of problems (including some instances of the graph inducibility problem).

Let us present a version of [46, Theorem 7.1] that we will need here. This result is stated in [46] for a more general kind of optimisation where we have a hereditary graph family \mathcal{G} (that is, we forbid some induced subgraphs) and we maximise our objective function over n -vertex graphs in \mathcal{G} only. Since every graph is allowed for the inducibility-type questions considered here (and thus \mathcal{G} consists of all graphs), we do not list \mathcal{G} in our notation and omit those assumptions of [46, Theorem 7.1] that vacuously hold. On the other hand, we need a generalisation where the pattern B can have loops (which indicate those parts in blowups of B where we put cliques). We have to adapt some definitions from [46] accordingly. Since the proof from [46] rather straightforwardly extends to patterns with loops, we just sketch it. In fact, if we are allowed to replace a problem by the complementary one (where all involved graphs are replaced by their complements) then

there is only one new solved case (namely the (5, 3)-edge-inducibility problem) where we need the pattern to have loops.

Recall that a homomorphism of a graph F to a pattern B is a (not necessarily injective) map that preserves both edges and non-edges; in other words it is an assignment of the vertices of F to the parts of a blowup of B that gives an induced copy of F . Also, recall that B is γ -minimal if, for every pattern B' obtained from B by removing a vertex, it holds that $\lambda_\gamma(B'()) < \lambda_\gamma(B())$.

Theorem 2 ([46, Theorem 7.1]). *Given an integer $\kappa \geq 2$ and a function $\gamma : \mathcal{F}_\kappa^0 \rightarrow \mathbb{R}$, let $\lambda_\gamma(G) := \sum_{F \in \mathcal{F}_\kappa^0} \gamma(F)p(F, G)$ be the corresponding objective function on graphs. Let B be a pattern (possibly with loops) on $[m]$ and let $\mathbf{a} \in \mathbb{S}_m$ be a vector with all entries positive. Suppose that all of the following statements hold.*

1. *We have a flag algebra certificate \mathcal{C} on N vertices proving that $\lambda_\gamma \leq \lambda_\gamma(B(\mathbf{a}))$ as in (9).*

2. *We have a graph τ (without loops) with at most $N - 2$ vertices such that*

(a) *if we restrict the maximisation of λ_γ to graphs without an induced copy of τ then the limit strictly decreases, that is,*

$$\lim_{n \rightarrow \infty} \max\{\lambda_\gamma(G) : v(G) = n, G \text{ is } \tau\text{-free}\} < \lambda_\gamma; \quad (13)$$

(b) *up to an automorphism of B (which by definition has to preserve loops and non-loops), there is a unique homomorphism from τ to B ;*

(c) *every two distinct vertices $x, y \in V(B)$ have distinct neighborhoods in $f(V(\tau))$, that is, $\Gamma_B(x) \cap f(V(\tau)) \neq \Gamma_B(y) \cap f(V(\tau))$, for some (or, equivalently by Item (b), for every) homomorphism f of τ to B .*

3. *Every $F \in \mathcal{F}_N^0$ with $c_F = 0$ admits a homomorphism to B .*

Suppose further that at least one of the following three statements holds:

(i) *the certificate \mathcal{C} contains the graph τ as a type and the corresponding matrix X^τ in \mathcal{C} is of co-rank 1 (that is, its kernel has dimension 1);*

(ii) *the pattern B has no loops and if we restrict maximisation of λ_γ to graphs without any induced copy of B then the limit strictly decreases;*

(iii) *the pattern B is γ -minimal.*

Then the problem of maximizing $\lambda_\gamma(G)$ over n -vertex graphs G is perfectly B -stable. Moreover, if Item (i) holds (in addition to Items 1–3) then \mathbf{a} is the unique vector $\mathbf{x} \in \mathbb{S}_m$ that maximises $\lambda_\gamma(B(\mathbf{x}))$.

Let us remark that the limit in the left-hand side of (13) exists by an easy double-counting argument, see e.g. [46, Lemma 2.2]. Also, note that we allow non-injective homomorphisms in Item (b) and such maps may be potentially required (e.g. we may need to force that a specific vertex x of τ is mapped to a loop, which can be done by adding a clone x' of x and making them adjacent).

Sketch of proof of Theorem 2. Suppose on the contrary that perfect stability does not hold. Thus for every C there is a graph G that violates (12). We let $C \rightarrow \infty$; then the order $n := v(G)$ of G goes to infinity and $\lambda_\gamma(G) = \lambda_\gamma + o(1)$.

Let us say that the problem is *robustly B -stable* if there is a constant $C_1 > 0$ such that for every graph H of order $n \geq C_1$ it holds that

$$\delta_{\text{edit}}(H, B()) \leq C_1 \max \{n, (\lambda_\gamma(n) - \lambda_\gamma(H))n^2\}.$$

Informally speaking, this is the same as the definition of perfect B -stability except we do not stipulate anything about order- n graphs H with $\lambda_\gamma(H) = \lambda_\gamma(n) + O(1/n)$.

Let us show that the λ_γ -problem is robustly B -stable, following [46, Theorem 4.1]. Let $u := \lambda_\gamma$. By [46, Lemma 2.2] (which is based on the standard blowup trick) we have that $\lambda_\gamma(n) = u + O(1/n)$. It follows from (9) that our hypothetical counterexample G satisfies that

$$p(F, G) = O(\max\{1/n, \lambda_\gamma(n) - \lambda_\gamma(G)\}), \quad \text{for every } F \in \mathcal{F}_N^0 \text{ with } c_F > 0. \quad (14)$$

Thus, by Assumption 3, every N -vertex graph F not admitting a homomorphism to B satisfies (14).

Furthermore, by Assumption (a) and the supersaturation argument of Erdős and Simonovits [21], we have that $p(\tau, G) = \Omega(1)$. Assume that the vertex set of τ is $[q]$. Fix a homomorphism f from τ to B . Every copy τ' of τ in G , say given by an injection $g : [q] \rightarrow V(G)$, defines m pairwise disjoint sets V_0, \dots, V_{m-1} , where V_i consists of those $x \in V(G) \setminus g([q])$ such that $g^{-1}(\Gamma_G(x)) = f^{-1}(\Gamma_B(i))$. Informally speaking, we put x into V_i if the G -adjacencies of x to τ' are the same as the B -adjacencies of i to τ . Distribute the vertices of the remainder $V(G) \setminus (V_0 \cup \dots, V_{m-1})$ arbitrarily into the parts, e.g. let all be assigned to V_0 . Assumptions (b), (c) and 3 of the theorem imply by $q \leq N - 2$ that, for every *wrong* pair xy , that is, $xy \in E(G) \triangle E(B(V_0, \dots, V_{m-1}))$, every N -set $Y \supseteq \{x, y\} \cup g([q])$ spans a subgraph in G not admitting a homomorphism to B . It follows from (14) that the expected number of wrong pairs over a uniformly random induced copy of τ in G is at most $O(\max\{n, (\lambda_\gamma(n) - \lambda_\gamma(G))n^2\})$. If we change the adjacency of every wrong pair, we obtain a blowup of B ; thus robust stability holds.

Let us show that the pattern B from the statement of Theorem 2 is γ -minimal. This is assumed in Item (iii) so we have to consider the other two cases.

First, suppose that Item (ii) holds. Suppose on the contrary that some vector $\mathbf{x} \in \mathbb{S}_{m-1}$ and a pattern B' obtained by removing a vertex from B satisfy $\lambda_\gamma(B'(\mathbf{x})) = \lambda_\gamma(B(\mathbf{a}))$, that is, we can asymptotically attain the maximum value λ_γ using blowups of B' . Then, by Item (ii), some sufficiently large blowups of B' contain a copy of B . This means that B admits a homomorphism h to B' . But then some two vertices of B that are sent by h to the same vertex of B' contradict Assumption (c).

So consider the case that Item (i) holds. Let us show that \mathbf{a} is the unique vector in \mathbb{S}_m maximising $\lambda_\gamma(B(\mathbf{a}))$, following [46, Lemma 6.2]. Let H be any n -vertex blowup of B with $\Lambda_\gamma(H) = \Lambda_\gamma(n) + o(n^\kappa)$. When we evaluate the identity (9) on H then the left-hand side is $o(n^N)$. So, each term in the right-hand side is also $o(n^N)$ by the non-negativity of the terms. In particular, this applies to the term obtained by summing over all embeddings f of τ into H of $(\mathbf{V}_{(H,f)}^{\tau,s})^T X^\tau \mathbf{V}_{(H,f)}^{\tau,s}$, where $s := (N + q)/2$ and $\mathbf{V}_{(H,f)}^{\tau,s}$ is the vector of the counts of s -vertex τ -flags in (H, f) , as defined in (8). Recall that we have $\Omega(n^q)$ embeddings f by Item (a) and supersaturation. Each such embedding comes from some homomorphism g of τ to B . There are only a bounded number of possible functions g (at most m^q) so fix the most frequent g . All maps f corresponding to this g give the same (normalised) vector $\mathbf{v} := \mathbf{v}_{(H,f)}^{\tau,s}$ and this vector \mathbf{v} necessarily satisfies $\mathbf{v}^T X^\tau \mathbf{v} = o(1)$. Thus \mathbf{v} is close to the 1-dimensional zero eigenspace of $X^\tau \succcurlyeq 0$. Also, the part sizes of H are determined by \mathbf{v} , since each part ratio can be written by Items (b) and (c) as the rooted density in (H, g) of a certain τ -flag with $q + 1 \leq s$ vertices. Passing to the limit as $n \rightarrow \infty$, we obtain that the vector \mathbf{a} is the unique maximiser in \mathbb{S}_m . Now, since the unique maximiser \mathbf{a} has all entries positive, the pattern is γ -minimal, as desired.

Thus B is γ -minimal in all cases (and we also proved the uniqueness of \mathbf{a} when Item (i) holds).

Next, we need two further definitions from [46], slightly simplified as we do not have any forbidden subgraphs. We call the pattern B *flip-averse* if there are $\delta > 0$ and n_0 such that for every blowup $H = B(V_0, \dots, V_{m-1})$ of order $n \geq n_0$ with $\lambda_\gamma(H) \geq \lambda_\gamma - \delta$ and any pair $xy \in \binom{V(H)}{2}$ it holds that $\Lambda_\gamma(H) - \Lambda_\gamma(H \oplus xy) \geq \delta n^{\kappa-2}$, where $H \oplus xy$ is obtained from H by changing the adjacency between x and y . Informally speaking, we require that every single adjacency change in an almost optimal blowup decreases the objective function by the maximum possible amount in the order of magnitude. Of course, for every $i, j \in [m]$, it is enough to check this property only for one pair $(x, y) \in V_i \times V_j$. Next, we call B *strict* if for every $\varepsilon > 0$ there are $\delta > 0$ and n_0 such that, for every $\mathbf{x} \in \mathbb{S}_m$ that maximises $\lambda_\gamma(B(\mathbf{x}))$, it holds that if $\lambda_\gamma(H', y) \geq \lambda_\gamma - \delta$ for a graph H' obtained by attaching a new vertex y to a blowup $H := B(V_0, \dots, V_{m-1})$ with $n \geq n_0$ vertices, where $|V_i| = x_i n \pm 1$ for $i \in [m]$, then there is $i \in [m]$ such that y can be made into a clone of a vertex in V_i after doing at most εn adjacency edits at y . Using compactness, these two definitions can be equivalently restated in the language of step graphons; for example, the latter would state that the only way to optimally attach to an optimal blowup of B is to be a clone (up to a null set) of an existing vertex.

We can show that the pattern B is both flip-averse and strict, following the argument from [46, Theorem 5.13]. If the first property is violated by an n -vertex blowup $H = B(V_0, \dots, V_{m-1})$ and x, y with $x \in V_i$ and $y \in V_j$ then the graph H' obtained from H by changing all adjacencies between X and Y for some disjoint εn -subsets $X \subseteq V_i$ and $Y \subseteq V_j$ for small fixed $\varepsilon > 0$ would violate the robust stability that we already proved. Indeed, H' can be shown to be $\Omega(\varepsilon^2 n^2)$ -far in the edit distance from a blowup of B while $\lambda_\gamma(H) - \lambda_\gamma(H') = o(\varepsilon^2)$ so no $\delta > 0$ satisfying robust stability can exist. (Also, note that such sets X and Y always exists, even if $i = j$: since B is γ -minimal, there is some $\varepsilon > 0$ such that $x_i > 3\varepsilon$ for every optimal vector $\mathbf{x} \in \mathbb{S}_m$.) Similarly, if the strictness property

fails then we can violate the robust stability by adding εn new copies of the same vertex v .

Now, we are ready to get the final contradiction, as in [46, Theorem 5.8]. Following [46], pick sufficiently small positive constants $c_6 \gg \dots \gg c_1$. Assume by taking n sufficiently large that $\lambda_\gamma(G) \geq \lambda_\gamma - c_1/2$. By Erdős-Simonovits stability (which directly follows from the established robust stability), take a blowup $H := B(V_0, \dots, V_{m-1})$ with $|W| \leq c_2 \binom{n}{2}$, where $W := E(G) \triangle E(H)$ is the set of wrong pairs. It must hold that the vector $\mathbf{b} := (|V_0|/n, \dots, |V_{m-1}|/n)$ is c_3 -close in the ℓ_1 -distance to an optimal vector $\mathbf{a}' = (a'_0, \dots, a'_{m-1})$. By the minimality of B , we have that each a'_i is at least c_6 (and so $b_i \geq c_6/2$). Also, for every $x \in V(G)$, we have $\lambda_\gamma(G, x) \leq \lambda_\gamma + c_2$ (as otherwise by adding e.g. $c_2 n$ clones of the vertex x we can bring $\lambda_\gamma(G)$ well over λ_γ , a contradiction). This allows us to conclude that $|S| \leq (2c_2/c_4)n$, where we let

$$S := \{x \in V(G) : \lambda_\gamma(G, x) \leq \lambda_\gamma - c_4\}.$$

Since the set S is small, it follows from the strictness of B that each vertex outside of S has the W -degree less than $c_5 n$ ([46, Equation (29)]). This in turn can be used to prove that the effect on the objective function with respect to G of a single edge flip for a pair inside $V(G) \setminus S$ is within additive $c_6 n^{\kappa-2}$ the same as the effect of the flip with respect to H ([46, Claim 5.12]). Again by strictness, we can “fix” all vertices in S to have correct adjacencies to the parts V_i so that, in terms of $\Lambda_\gamma(G)$, the objective function $\Lambda_\gamma(G)$ increases by at least $c_4 n^{\kappa-1}$ per each fixed vertex.

Repeating iteratively this for every vertex of S and flipping the remaining wrong pairs (which all lie inside $V(G) \setminus S$), we increase the objective function $\Lambda_\gamma(G)$ by at least $c_3(n|S| + |W|)n^{\kappa-2}$ ([46, Equation (33)]) while doing at most $n|S| + |W|$ changes. Thus our assumption that G fails perfect stability with $C \rightarrow \infty$ leads to a contradiction. \square

3 Computer-generated results

In this section we present the results whose proof was generated by computer, postponing those stability and exact structure results that do not directly follow from Theorem 2 but require some extra arguments to Sections 4–5.

The upper bound on the (edge/graph) inducibility constant λ_γ for the appropriate γ in each theorem of this section is proved via a flag algebra identity (9), and we just give the value of N in the statement of the corresponding theorem. Also, if perfect stability is claimed then it was derived by applying Theorem 2; we just indicate which of its alternatives (i) or (ii) was applied and the used graph τ . (In this paper, we do not use Item (iii) of Theorem 2 but we included it for future reference.) In all cases that are solved in this paper, the provided maximiser \mathbf{a} of $\lambda_\gamma(B(\mathbf{a}))$ happens to be unique up to an automorphism of B (and if perfect stability is proved via Theorem 2 then the same value of N is used for this). Recall that the uniqueness of the maximiser \mathbf{a} follow automatically in the cases where Theorem 2(i) applies; in the remaining cases we include a proof that a maximising vector is unique up to an automorphism of B . Also, our scripts verify the stated lower bounds on λ_γ by computing the value of $\lambda_\gamma(B(\mathbf{a}))$; these calculations can also be found in the arxiv version [11] of this paper.

We formed the SDPs coming from (9) and then analysed the solutions returned by computer, using the SageMath package `FlagAlgebraToolbox` (commit 22b8765) of the first author. This package is still under development; a short guide on how to install it and an overview of its current functionality can be found in [8]. The scripts that we used to generate the certificates, verify them and check that Theorem 2 applies in each stated case can be found in a separate GitHub repository https://github.com/bodnalev/supplementary_files/tree/main/graph_inducibility while the certificates themselves can be found in its sub-folder `certificates`. Alternatively, all these files can be found in the ancillary folder of the arXiv version of this paper.

The certificate for e.g. the $(4, 3)$ -edge-inducibility problem is named `stats43.pickle`, while a certificate for the graph inducibility starts with `ind` and so should be easy to identify. Our scripts also verify (using exact arithmetical calculations) that the matrices and slacks listed in each exact-value certificate indeed prove the claimed upper bound via an identity as in (9). Alternatively, a reader can use their own verifier; a description of how the data are arranged in each `pickle` file can be found in the readme file of the same [supplementary_files](https://github.com/bodnalev/supplementary_files) repository.

Let us now give details of our results on edge-inducibility (that were briefly summarised in Table 1).

Theorem 3. *It holds that $\lambda_{4,2} = 1/2$, the problem is perfectly B -stable for $B := 2K_3$ by Theorem 2(i) with $N = 7$ and $\tau = K_3 + K_2$ and the unique maximiser of $\lambda_{4,2}(B(\mathbf{a}))$ is $\mathbf{a} = (1/6, \dots, 1/6)$.*

Theorem 4. *It holds that $\lambda_{4,3} = 1/2$. The lower bound comes from blowups of $B = (2, \{00, 11\})$ (two loops) and of $B = (2, \{10\})$ (a single edge), with $\mathbf{a} = (1/2, 1/2)$ being the unique maximiser of $\lambda_{4,3}(B(\mathbf{a}))$ in both cases. The upper bound can be proved by flag algebras as in (9) using $N = 6$.*

Proof. Given the output of our scripts, it remains to show that the given maximiser \mathbf{a} is unique. Observe that a 4-subset of a blowup of B spans a $(4, 3)$ -subgraph if and only if it has exactly 3 vertices in one part (and 1 vertex in the other part). If x denotes the fraction of vertices in the first part then the asymptotic density of such 4-sets is

$$p(x) := 4(x^3(1-x) + x(1-x)^3), \quad (15)$$

where the factor 4 is the number of ways to choose the “lonely” vertex. The derivative $p'(x) = -4(2x-1)^3$ has the unique (triple) root at $x = 1/2$, which gives the global maximum of p and is the unique argument on which p assumes the maximum value $1/2$. \square

Let us observe that perfect stability does not hold for the $(4, 3)$ -problem even if we adapt its definition to allow multiple possible optimal patterns B . Indeed, fix small $\varepsilon > 0$, let $n \rightarrow \infty$ and let G be obtained from $K_{n/2} + K_{n/2}$ by adding all edges between two fixed εn -subsets in different parts. It is routine to see that G is $\varepsilon^2 n^2$ -far from a union of two cliques (and from $T_{n/2, n/2}$) while $\lambda_{4,3}(G) = 1/2 + O(\varepsilon^3)$. Since $\varepsilon > 0$ is arbitrary, the edit distance cannot be upper bounded by $C(\lambda_{4,3}(n) - \lambda_{4,3}(G))n^2$ for an absolute constant C .

Even though perfect stability does not hold for the $(4, 3)$ -problem, we establish Erdős-Simonovits stability (in Lemma 18) and prove using it that, for all large n , every extremal graph is a complete bipartite graph or a union of two cliques (in Theorem 19). For this we will need the following lemma that can be derived by examining our flag algebra certificate proving the upper bound $\lambda_{4,3} \leq 1/2$.

Lemma 5. *For every $\varepsilon > 0$ there are $\delta > 0$ and n_0 such that every graph G with $n \geq n_0$ vertices and $\lambda_{4,3}(G) \geq 1/2 - \delta$ satisfies $\lambda_{4,1}(G) \leq \varepsilon$ and $\lambda_{4,5}(G) \leq \varepsilon$.*

Proof. Our script checks that, for every graph F with $N = 6$ vertices that contains a $(4, 1)$ - or $(4, 5)$ -subgraph, its slack constant c_F from the certificate is positive and thus lower bounded by some absolute constant $c > 0$. Thus each such F has density at most $\delta/c + o(1)$ in G . By

$$\lambda_{4,i}(G) = \sum_{F \in \mathcal{F}_N^0} \lambda_{4,i}(F) p(F, G) + o(1),$$

the lemma can be satisfied by taking e.g. $\delta := 0.9 c\varepsilon/2^{\binom{N}{2}}$ and then sufficiently large n_0 . \square

Theorem 6. *It holds that $\lambda_{5,2} = 280/3^6$, the problem is perfectly B -stable for $B := 3K_3$ by Theorem 2(i) with $N = 8$ and $\tau = 3K_2$ and the unique maximizer of $\lambda_{5,2}(B(\mathbf{a}))$ is $\mathbf{a} = (1/9, \dots, 1/9)$.*

Theorem 7. *It holds that $\lambda_{5,3} = (255\sqrt{17} - 535)/2^{10}$, the problem is perfectly B -stable for $B := (3, \{01, 22\})$ (that is, B consists of an edge plus an isolated loop) by Theorem 2(i) with $N = 7$ and $\tau = T_{2,1}$ (the cherry), and the unique maximiser of $\lambda_{5,3}(B(\mathbf{a}))$ is $\mathbf{a} = (\alpha_1, \alpha_1, 1 - 2\alpha_1)$, where α_1 was defined in (2).*

Theorem 8. *It holds that $\lambda_{5,4} = 5/8$, the problem is perfectly B -stable with $B := (2, \{00, 11\})$ (that is, B consists of two isolated loops) by Theorem 2(i) with $N = 5$ and $\tau = (3, \{01\})$ and the unique maximiser of $\lambda_{5,4}(B(\mathbf{a}))$ is $\mathbf{a} = (1/2, 1/2)$.*

Theorem 9. *It holds that $\lambda_{6,4} = 40/3^4$, the problem is perfectly B -stable for $B := (3, \{00, 11, 22\})$ by Theorem 2(i) with $N = 6$ and $\tau = (2, \{\})$ and the unique maximizer of $\lambda_{6,4}(B(\mathbf{a}))$ is $\mathbf{a} = (1/3, 1/3, 1/3)$.*

Theorem 10. *It holds that $\lambda_{6,5} = (10\sqrt{10} - 28)/9$, the problem is perfectly B -stable with $B := K_2$ by Theorem 2(ii) with $N = 7$ and $\tau = (3, \{01, 02\})$, and every maximiser of $\lambda_{6,5}(B(\mathbf{a}))$ is $\mathbf{a} = (\alpha_2, 1 - \alpha_2)$ up to an automorphism of B , where α_2 was defined in (3).*

Proof. A 6-set spans 5 edges in a blowup $B((x, 1 - x))$ if and only if it has 5 vertices in one part (and 1 vertex in the other). This happens with limiting probability

$$6 \cdot (x^5(1 - x) + x(1 - x)^5),$$

whose derivative is

$$-6(6x^4 - 12x^3 + 14x^2 - 8x + 1)(2x - 1),$$

with roots (including complex ones)

$$\frac{1}{2} \pm \frac{1}{2} \sqrt{\pm \frac{2}{3} \sqrt{10} - \frac{5}{3}}$$

and $1/2$. The highest density is attained at and only at α_2 or $1 - \alpha_2$, both giving the final bound $(10\sqrt{10} - 28)/9$. The two maximising vectors are the same, up to the automorphism of B that swaps the parts. \square

Theorem 11. *It holds that $\lambda_{6,7} = 15/2^5$, the problem is perfectly B -stable for $B := (2, \{00, 11\})$ by Theorem 2(i) with $N = 6$ and $\tau = (2, \{\})$ and the unique maximizer of $\lambda_{6,7}(B(\mathbf{a}))$ is $\mathbf{a} = (1/2, 1/2)$.*

Theorem 12. *It holds that $\lambda_{7,6} = (28\sqrt{10} - 35)/135$, the problem is perfectly B -stable with $B := K_2$ by Theorem 2(ii) with $N = 7$ and $\tau = (3, \{\})$, and every maximiser of $\lambda_{7,6}(B(\mathbf{a}))$ is $\mathbf{a} = (\alpha_3, 1 - \alpha_3)$ up to an automorphism of B , where α_3 was defined in (4).*

Proof. A 7-set spans 6 edges if and only if it has 6 vertices in one part (and 1 in the other). Thus the limiting density in $B((x, 1 - x))$ is

$$7 \cdot (x^6(1 - x) + x(1 - x)^6),$$

whose derivative is

$$-7(15x^4 - 30x^3 + 25x^2 - 10x + 1)(2x - 1)$$

with roots (including complex ones)

$$\frac{1}{2} \pm \frac{1}{2} \sqrt{\pm \frac{4}{15} \sqrt{10} - \frac{1}{3}}$$

and $1/2$. The highest density is attained at and only at α_3 or $1 - \alpha_3$, both giving the final bound $(28\sqrt{10} - 35)/135$. The two resulting maximising vectors are the same up to the automorphism of B that swaps the two parts. \square

Theorem 13. *It holds that $\lambda_{7,9} = 35/2^6$, the problem is perfectly B -stable for $B := (2, \{00, 11\})$ by Theorem 2(i) with $N = 7$ and $\tau = (3, \{01\})$ and the unique maximizer of $\lambda_{7,9}(B(\mathbf{a}))$ is $\mathbf{a} = (1/2, 1/2)$.*

Theorem 14. *It holds that $\lambda_{7,10} = 28/3^4$, the problem is perfectly B -stable for $B := K_2$ by Theorem 2(ii) with $N = 7$ and $\tau = (3, \{\})$ and every maximizer of $\lambda_{7,10}(B(\mathbf{a}))$ is $\mathbf{a} = (1/3, 2/3)$ up to an automorphism of B .*

Proof. A 7-set spans 10 edges if and only if it has 5 vertices in one part (and 2 in the other). Thus the limiting density in $B((x, 1 - x))$ is

$$\binom{7}{2} (x^5(1 - x)^2 + x^2(1 - x)^5),$$

whose maximum is attained at and only at $1/3$ or $2/3$, the optimal roots of the derivative

$$-21(3x-1)(3x-2)(2x-1)(x-1)x,$$

both giving the highest density $28/3^4$, as claimed. The two resulting maximising vectors are the same under the automorphism of B that swaps the parts. \square

Next, we turn to new inducibility results. Interestingly, even though the optimal part ratios are irrational in the next two results, the final density happens to be a rational number. In fact, the limiting density of every graph is a rational number in this construction.

Theorem 15. *Let $F := (5, \{01, 02, 03\})$, be the 3-edge star $T_{3,1}$ plus an isolated vertex. Let the pattern $B := (4, \{01, 23\})$ consist of two disjoint edges (i.e. $B := 2K_2$) and let $\beta := (3 + \sqrt{3})/12$. Then $\lambda_F = 5/24$ and the F -inducibility problem is perfectly B -stable by Theorem 2(ii) with $N = 7$ and $\tau = (5, \{02, 03, 14\})$. Moreover, every maximiser of $\lambda_F(B(\mathbf{a}))$ is $\mathbf{a} = (\beta, \beta, 1/2 - \beta, 1/2 - \beta)$ up to an automorphism of B .*

Proof. Let G be an $(x\mu_0, x(1-\mu_0), (1-x)\mu_1, (1-x)(1-\mu_1))$ -blowup of B with connectivity components V_0 and V_1 . Thus we have a pair of complete bipartite graphs on V_0 and on V_1 with part ratios μ_0, μ_1 respectively, and the size ratio between V_0 and V_1 is x . A 5-set spans a copy of F if and only if some 4 vertices are inside one part V_i and span the 3-star $T_{3,1}$, while the remaining fifth vertex is in the other part V_{1-i} . Note that $p(T_{3,1}, K_{\mu_n, (1-\mu)n}) = h(\mu) + o(1)$, where we denote $h(\mu) := 4(\mu^3(1-\mu) + (1-\mu)^3\mu)$. This gives that the total probability of seeing a copy of F is

$$p(F, G) = 5(x^4(1-x)h(\mu_0) + x(1-x)^4h(\mu_1)) + o(1).$$

When $x \in \{0, 1\}$, this probability is 0. Otherwise, the probability depends strictly monotonically on $h(\mu_0)$ and $h(\mu_1)$, hence we can optimize them separately. The polynomial h takes its unique maximum at $1/2$. Thus assume from now on that $\mu_0 = \mu_1 = 1/2$. We have

$$p(F, G) = \frac{5}{2}(x^4(1-x) + x(1-x)^4) + o(1).$$

The maximum of this polynomial is attained at 2β and $1-2\beta$, which are two roots of the derivative

$$-\frac{5}{2}(6x^2 - 6x + 1)(2x - 1),$$

giving the maximal density $5/24$, as claimed. Also, the cases $x = 2\beta$ and $x = 1 - 2\beta$ are the same up to any automorphism of B that swaps the two edges of B , implying the uniqueness of the maximiser. \square

Theorem 16. *Let $F := (4, \{01, 12, 23, 30\})$ be the 4-cycle plus an isolated vertex. Let the pattern $B := (4, \{01, 23\})$ consist of two disjoint edges and let $\beta := (3 + \sqrt{3})/12$. Then $\lambda_F = 5/32$ and the F -inducibility problem is perfectly B -stable by Theorem 2(ii) with $N = 7$ and $\tau = (5, \{02, 03, 14\})$. Moreover, every maximiser of $\lambda_F(B(\mathbf{a}))$ is $\mathbf{a} = (\beta, \beta, 1/2 - \beta, 1/2 - \beta)$ up to an automorphism of B .*

Proof. As before, let G be an $(x\mu_0, x(1-\mu_0), (1-x)\mu_1, (1-x)(1-\mu_1))$ -blowup of B with connectivity components V_0 and V_1 . A 5-set spans a copy of F if and only if some 4 vertices are inside one part V_i and span the 4-cycle $T_{2,2}$, while the remaining fifth vertex is in the other part V_{1-i} . Note that $p(T_{2,2}, K_{\mu, (1-\mu)n}) = h(\mu) + o(1)$, where we denote $h(\mu) := 6\mu^2(1-\mu)^2$. Thus

$$p(F, G) = 5(x^4(1-x)h(\mu_0) + x(1-x)^4h(\mu_1)) + o(1).$$

When $x \in \{0, 1\}$, this probability is 0. Otherwise, the probability depends strictly monotonically on $h(\mu_0)$ and $h(\mu_1)$ whose unique maximum is at $1/2$. Thus we get the same polynomial as above, except with a different coefficient in front:

$$p(F, G) = \frac{15}{8}(x^4(1-x) + x(1-x)^4) + o(1).$$

As we know, the maximum is attained at and only at 2β and $1-2\beta$, giving the maximal density $5/32$ and the uniqueness of the maximiser up to an automorphism of B . \square

Our last graph inducibility result deals with the 4-cycle plus a pendant edge. Here the inducibility constant is attained by a sequence of bipartite $5/6$ -quasirandom graphs. In Theorem 21, we will prove that all almost extremal order- n graphs are, within $o(n^2)$ edits, of this form. For this structural result, we need some additional information from the flag algebra proof of the upper bound, which we collect here.

Theorem 17. *Let $F := (5, \{01, 12, 23, 30, 04\})$ be the 4-cycle plus a pendant edge. Then*

$$\lambda_F = \frac{5^6}{2^8 \cdot 3^5} = \frac{15625}{62208}.$$

Moreover, there is a flag algebra proof of the upper bound with $N = 5$ that satisfies the following properties.

- A) *For the single-vertex type 1, the zero eigenspace of the matrix X^1 has dimension 1.*
- B) *For the cherry $\nu := (3, \{01, 02\})$, the zero eigenspace of the matrix X^ν has dimension 1.*
- C) *The zero eigenspace of the matrix X^σ for the type $\sigma := (3, \{\})$ (that has 3 vertices and no edges) does not contain any non-zero linear combination of 4-vertex σ -flags where the unlabelled vertex sends at most one edge to the roots.*
- D) *The zero eigenspace of the matrix X^μ for the type $\mu := (3, \{01\})$ (that has 3 vertices and one edge) does not contain any non-zero linear combination of 4-vertex μ -flags where the free vertex sends at most one edge to the roots.*

Proof. The claimed upper bound on λ_F is proved by a standard flag algebra application with $N = 5$. The provided script also verifies the additional Properties A)–D) of the obtained certificate. \square

4 Stability and exact result for the (4, 3)-problem

First, we show that Erdős-Simonovits stability (which was defined here to allow multiple constructions) holds for the (4, 3)-edge-inducibility problem. Recall that $\lambda_{4,3} = 1/2$, as proved in Theorem 4.

Lemma 18. *For every $\varepsilon > 0$ there are $\delta > 0$ and n_0 such that every graph G with $n \geq n_0$ vertices and $\lambda_{4,3}(G) \geq \frac{1}{2} - \delta$ is within edit distance at most $\varepsilon \binom{n}{2}$ from $T_{n/2, n/2}$ or $2K_{n/2}$.*

Proof. Since we do not compute the dependence of δ on ε , we present a proof which is short but not efficient. (For example, the application of the Induced Removal Lemma can be avoided by fixing a “typical” 4-set spanning a clique or an independent set, and defining the two parts depending on the adjacencies of a vertex to the set.)

Suppose for contradiction that the lemma fails for some $\varepsilon > 0$. Thus for every large integer s there is a graph G_s that satisfies $v(G_s) \geq s$ and $\lambda_{4,3}(G_s) \geq \frac{1}{2} - \frac{1}{s}$ but is $\varepsilon \binom{v(G_s)}{2}$ -far from the stated blowups.

Let $s \rightarrow \infty$. We know by Lemma 5 that $\lambda_{4,1}(G_s)$ and $\lambda_{4,5}(G_s)$ are both $o(1)$. By the Induced Removal Lemma of Alon, Fischer, Krivelevich and Szegedy [2], we can change $o(1)$ -fraction of adjacencies in G_s and destroy all copies of (4, 1)- and (4, 5)-subgraphs. Of course, this changes any subgraph density by $o(1)$ so the new graph, call it H , satisfies $\lambda_{4,3}(H) = \frac{1}{2} + o(1)$ as $s \rightarrow \infty$. Let $n := v(H)$.

By Ramsey’s theorem, H contains a clique or an independent set with at least 4 vertices. By passing to graph complements if needed, we can assume that the former holds. Let $V_0 \subseteq V(H)$ be a maximum subset spanning a clique. If some u in $V_1 := V(H) \setminus V_0$ has two distinct neighbours $w_0, w_1 \in V_0$ then for any w in the set $V_0 \setminus \Gamma(u)$ (which is non-empty by the maximality of V_0) the set $\{u, w_0, w_1, w\}$ spans a (4, 5)-subgraph in H . This contradiction shows that every vertex outside of V_0 has at most one neighbour in V_0 .

Suppose next that we have two non-adjacent vertices $u, w \in V_1$. By above and since $|V_0| \geq 4$, there are distinct $w_0, w_1 \in V_0 \setminus (\Gamma(u) \cup \Gamma(w))$. Thus the set $\{u, w, w_0, w_1\}$ spans a (4, 1)-subgraph in H , a contradiction again.

We see that, apart at most $|V_1| \leq n$ crossing edges, H is the union of the cliques on V_0 and V_1 . Thus, for $x := |V_0|/n$, we have that

$$\Lambda_{4,3}(H) \geq \Lambda_{4,3}(K_{xn} + K_{(1-x)n}) + O(n^3) = p(x) \binom{n}{4} + O(n^3), \quad (16)$$

where $p(x) := 4(x^3(1-x) + x(1-x)^3)$. As we observed in the proof of Theorem 4, $x = 1/2$ is the unique maximiser of this polynomial (with $p(1/2) = 1/2$). Thus, by $\lambda_{4,3}(H) = 1/2 + o(1)$, each part V_i has $(\frac{1}{2} + o(1))n$ vertices. Thus the original graph G_s is $o(n^2)$ -close in the edit distance to $2K_{n/2}$, a contradiction to our assumption. \square

Our next result shows that every $\lambda_{4,3}$ -extremal graph G of sufficiently large order n is the disjoint union of two cliques $K_m + K_{n-m}$ or a complete bipartite graph $T_{m, n-m}$ for some integer $m \in [0, n]$, without a single wrong adjacency. Clearly, the number of

(4, 3)-subgraphs is exactly $f_m := m \binom{n-m}{3} + \binom{m}{3}(n-m)$. It is possible to describe all integers m that maximise it. For this, routine calculations show that

$$f_{m+1} - f_m = \frac{1}{6}(n - 2m - 1)(4m^2 - 4mn + 4m + n^2 - 5n + 6).$$

This cubic in m polynomial has three real roots, namely $(n-1)/2$ and $m_{\pm} := (n - 1 \pm \sqrt{3n-5})/2$. Thus (for all large n) there are 4 optimal choices of m if $3n-5$ is a square of different parity than n and 2 optimal choices otherwise. For example, in the latter case, the (unique) optimal m at least $n/2$ is $\lceil m_+ \rceil$, the maximal integer with $f_m - f_{m-1} > 0 > f_{m+1} - f_m$. In any case, every optimal m is $(n \pm \sqrt{3n})/2 + O(1)$, which gives, after routine calculations that, for such m ,

$$\Lambda_{4,3}(n) \geq \Lambda_{4,3}(K_m + K_{n-m}) = \frac{1}{2} \binom{n}{4} + \frac{1}{8} n^2 + O(n), \quad \text{as } n \rightarrow \infty. \quad (17)$$

For comparison, to see the effect of this imbalance between the part sizes, note that

$$\Lambda_{4,3}(K_{\lfloor n/2 \rfloor} + K_{\lceil n/2 \rceil}) = \frac{1}{2} \binom{n}{4} - \frac{1}{16} n^2 + O(n), \quad \text{as } n \rightarrow \infty.$$

Theorem 19. *There is n_0 such that every graph G with $n \geq n_0$ vertices and $\Lambda_{4,3}(G) = \Lambda_{4,3}(n)$ is a union of two cliques or a complete bipartite graph.*

Proof. Choose positive constants in this order, each being sufficiently small depending on the previous ones:

$$\varepsilon_2 \gg \varepsilon_1 \gg \varepsilon_0. \quad (18)$$

Let G be any graph with $n \rightarrow \infty$ vertices such that $\Lambda_{4,3}(G) = \Lambda_{4,3}(n)$.

By Lemma 18, we can assume that G is within $\varepsilon_0 \binom{n}{2}$ in the edit distance from two disjoint cliques or a complete bipartite graph. Since the problem is self-complementary, we can assume that G is close to the union of two cliques. Pick a partition $V_0 \cup V_1 = V(G)$ such that the symmetric difference

$$W := E(G) \Delta \left(\binom{V_0}{2} \cup \binom{V_1}{2} \right)$$

has the smallest possible size. We call pairs in W *wrong*. Of course, $|W| \leq \varepsilon_0 \binom{n}{2}$. Let $x := |V_0|/n$. As we already established in the proof of Theorem 4, the polynomial $p(x)$ in the right-hand side of (16) is at most $1/2$ with equality if and only if $x = 1/2$. By continuity and compactness, we can additionally assume that $|x - 1/2| \leq \varepsilon_1$. As before, we do not write explicit dependencies between the constants even though this can in principle be done for all steps; for example, for this step it suffices to take, say, $\varepsilon_0 \leq \varepsilon_1^4/100$.

Recall that, for $u \in V(G)$, $A_{\kappa,\ell}(G, u)$ denotes the number of (κ, ℓ) -subgraphs in G that contain the vertex u . By (17) the expected value of $\Lambda_{4,3}(G, u)$ for a uniformly random vertex u is at least $\frac{1}{2} \binom{n}{4} \cdot \frac{4}{n} = \frac{1}{2} \binom{n-1}{3}$. Fix a vertex u with $\Lambda_{4,3}(G, u)$ being at least average.

Since we cannot increase the total number of (κ, ℓ) -subgraphs when replacing some vertex w by a clone of u , it holds that

$$\Lambda_{4,3}(G, w) \geq \frac{1}{2} \binom{n-1}{3} - \binom{n-2}{2}, \quad \text{for every } w \in V(G). \quad (19)$$

Next, we would like to show that each vertex is incident to small number of wrong pairs. For this we first define a polynomial $Q(x, y, z)$ that gives the limiting density of newly created $(4, 3)$ -subgraphs when we add to $H := K_{xn} + K_{(1-x)n}$ a new vertex w with yn and zn neighbours in the first and the second cliques respectively. It is routine to see that a triple X in H makes a $(4, 3)$ -subgraph with w if and only if X spans a triangle and sends no edge to w , or X spans exactly 1 edge and sends 2 edges to w . Thus we define

$$Q(x, y, z) := (x - y)^3 + (1 - x - z)^3 + 3y^2(1 - x - z) + 3(x - y)z^2 + 6yz(1 - y - z).$$

The following claim implies that every almost optimal way for a vertex to attach to H is to almost follow the pattern structure (which is basically the strictness property from the proof of Theorem 2, when adapted to having more than one optimal pattern B).

Claim 20. *The maximum value of $q(y, z) := Q(1/2, y, z)$ on $[0, 1/2]^2$ is $1/2$, and is attained only at $(1/2, 0)$ and $(0, 1/2)$.*

Also, for every $x \in [1/2 - \varepsilon_1, 1/2 + \varepsilon_1]$ and for every $(y, z) \in [0, x] \times [0, 1 - x]$ at ℓ_1 -distance at least ε_2 from $\{(0, 1 - x), (x, 0)\}$, it holds that $Q(x, y, z) \leq 1/2 - \varepsilon_1$.

Proof. Let $(y, z) \in [0, 1/2]^2$ maximise the polynomial q . Of course, $q(y, z) \geq q(1/2, 0) = 1/2$. Suppose first that (y, z) is an interior point of $[0, 1/2]^2$. Then it is a critical point of q . Routine calculations show that the difference of the partial derivatives at (y, z) is

$$\frac{\partial q}{\partial y}(y, z) - \frac{\partial q}{\partial z}(y, z) = 6(y - z)(y + z),$$

and thus $y = z$. The derivative of the cubic polynomial $q(y, y) = -20y^3 + 12y^2 - 3y/2 + 1/4$ has roots $(4 \pm \sqrt{6})/20$ and the values of $q(y, y)$ on these points, which are $(27 \pm 3\sqrt{6})/100$, are both strictly less than $1/2$. Thus no interior point can maximise q .

So (y, z) lies on the boundary. Suppose first that $z = 0$. Then the derivative of $q(y, 0) = -y^3 + 3y^2 - 3y/4 + 1/4$ has two roots $1 \pm \sqrt{3}/2$, one belonging in $[0, 1/2]$ and the other being larger than $1/2$. Thus the only possible values for optimal y are 0 and $1/2$. We have to rule out the former as $q(0, 0) = 1/4 < 1/2$ (while the latter gives the maximum). Finally, by the symmetry between y and z , it remains to consider only the case $z = 1/2$. Here $q(y, 1/2) = -y^3 - 3y^2/2 + 1/2$, its derivative has roots -1 and 0, and $y = 0$ is the only point in $[0, 1/2]$ giving the maximum value $1/2$, proving the first part of the claim.

The second part can be derived from the first one via a compactness argument. Suppose on the contrary that, for some given $\varepsilon_2 > 0$, it cannot be satisfied for any sufficiently small $\varepsilon_1 > 0$. We let $\varepsilon_1 := 1/s$ with integer $s \rightarrow \infty$ and for each s pick a counterexample

(x_s, y_s, z_s) . By passing to a subsequence of s , we can assume that these triples converge to some (x, y, z) . We have that $x = 1/2$, $(y, z) \in [0, 1/2]^2$ and, by the continuity of the polynomial Q , that $q(y, z) = 1/2$. By the first part, (y, z) is $(1/2, 0)$ or $(0, 1/2)$ contradicting our assumption that each (y_s, z_s) is ε_2 -far from both $(0, 1 - x_s) = (0, 1/2 + o(1))$ and $(x_s, 0) = (1/2 + o(1), 0)$. \square

Let us show that every vertex w satisfies $\deg_W(w) \leq \varepsilon_2 n$, that is, is incident to at most $\varepsilon_2 n$ wrong pairs. Recall that we defined $x = |V_0|/n$. Let $y := |\Gamma(w) \cap V_0|/n$ and $z := |\Gamma(w) \cap V_1|/n$. The difference between $\Lambda_{4,3}(G, w)$ and $Q(x, y, z) \binom{n-1}{3}$ is at most $\deg_W(w)n^2 + |W|n + O(n^2)$, where the first two terms upper bound the number of 4-sets that contain w and at least one wrong pair while the last term is due to approximations (such as of $\binom{xn}{2}$ by $x^2 n^2 / 2 + O(n)$). Now, (19) and the second part of Claim 20 give that (y, z) is ε_2 -close in the ℓ_1 -distance to $(x, 0)$ or $(0, 1 - x)$. Since moving w to the other part cannot decrease $|W|$ by the choice of the partition $V_0 \cup V_1$, the former alternative holds, giving the required.

Next, let us show that each part V_i spans a clique, without a single missing edge. Suppose on the contrary that some distinct $u, w \in V_i$ are non-adjacent. Let G' be obtained from G by making uw an edge. Of course, every 4-subset of $V(G)$ not containing the pair uw induces the same subgraph in G and G' , and so contributes the same amount to each of $\Lambda_{4,3}(G')$ and $\Lambda_{4,3}(G)$. So consider a 4-set $X \subseteq V(G)$ that contain both u and w but no other wrong pair (which excludes at most $(\deg_W(u) + \deg_W(w))n + |W| \leq 3\varepsilon_2 n^2$ quadruples). The set X cannot span a $(4, 3)$ -subgraph in G (spanning always 1, 2 or 5 edges). However, if the remaining two vertices of X lie in different parts (that is, $|X \cap V_i| = 3$) then X spans a $(4, 3)$ -subgraph in G' . Thus

$$\Lambda_{4,3}(G') - \Lambda_{4,3}(G) \geq (|V_i| - 2) |V_{1-i}| - 3\varepsilon_2 n^2 > 0,$$

which contradicts the maximality of G .

Thus all wrong pairs go between the two parts. This gives us quite strong control: for example, the only way to get a $(4, 3)$ -subgraph containing a wrong pair uw is that uw is the only wrong pair and there are exactly two vertices in each part V_i .

We would like to further bound possible wrong degrees. Let u be a vertex of the maximum wrong degree. Let $d := \deg_W(u)$. Assume that $d \geq 1$ as otherwise we are done. Among the d wrong neighbours of u (which are all in the other part), let w be one of the maximum wrong degree and let $f := \deg_W(w)$. Of course, $1 \leq f \leq d$. Without loss of generality, assume that $u \in V_0$ and $w \in V_1$. Recall that we denote $x = |V_0|/n$. Let y be $1 - x = |V_1|/n$. Let G' be the graph obtained from G by removing the edge uw . Consider the difference $\Lambda_{4,3}(G) - \Lambda_{4,3}(G')$ which is non-negative by the extremality of G . As before, we have to analyse only those 4-sets X that contain the pair uw . Let u' and w' be the other two vertices of X . If such X spans a $(4, 3)$ -subgraph in G then, up to re-ordering of u' and w' , we have that $u' \in V_0 \setminus \Gamma(w)$, $w' \in V_1 \setminus \Gamma(u)$ and $u'w' \notin E(G)$. Since the number of edges in G connecting $V_0 \setminus \Gamma(w)$ to $V_1 \setminus \Gamma(u)$ is at least $|W| - 2d^2$, we get

$$\Lambda_{4,3}(G, uw) \leq (xn - f)(yn - d) - |W| + 2d^2, \tag{20}$$

where $\Lambda_{4,3}(G, uw)$ denotes the number of 4-sets $X \subseteq V(G)$ spanning exactly 3 edges in G and containing both u and w . On the other hand, we have

$$\Lambda_{4,3}(G', uw) \geq \binom{xn-f}{2} + \binom{yn-d}{2} + (d-1)(xn-f) + (f-1)(yn-d) - 2d^2. \quad (21)$$

Indeed, the first two terms count those quadruples $X \ni u, w$ that have 3 vertices in some part (and 1 vertex in the other part) such that uw is the only wrong edge within X ; note that each such X spans exactly 3 edges in G' . The next two terms count the number of those 4-sets $X = \{u, w, u', w'\}$ such that $u' \in V_0, w' \in V_1$ and exactly one of uw' and $u'w$ is in W . Each such X spans exactly 3 edges in G' , unless $u'w' \in W$; the number of such wrong pairs $u'w'$ can be upper bounded by $(d-1)^2 + (f-1)(d-1) \leq 2d^2$, using the definitions of d and f .

Let $c := (x-1/2)\sqrt{n}$; thus $|V_0| = n/2 + c\sqrt{n}$. Also, let $a := |W|/n$. By using the bounds in (20) and (21) (and then that $d \geq f$) we get after routine calculations that

$$\begin{aligned} 0 &\leq \Lambda_{4,3}(G) - \Lambda_{4,3}(G') \leq \text{RHS of (20)} - \text{RHS of (21)} \\ &= \frac{3-2a-4c^2-d-f}{2}n + \frac{7d^2+6df-f^2}{2} + 3(f-d)c\sqrt{n} - \frac{3d+3f}{2} \\ &\leq \frac{3-2a-4c^2-d-f}{2}n + \frac{13}{2}d^2. \end{aligned}$$

Thus, by $1 \leq f \leq d \leq \varepsilon_2 n$, it holds that $d \leq 2$ and $f = 1$. Furthermore, since $d+f \geq 2$, we have that

$$2a + 4c^2 \leq 1 + o(1). \quad (22)$$

Next, we write an upper bound on the global function $\Lambda_{4,3}(G)$ via some version of inclusion-exclusion and then argue that it is incompatible with the lower bound in (17) and our assumption that $d \geq 1$. In order to state it, we need to give some definitions first. A set $X \subseteq V_0 \cup V_1$ is an $\{i, j\}$ -set if $\{|X \cap V_0|, |X \cap V_1|\} = \{i, j\}$. Let

- $t_0 := \binom{xn}{3}yn + xn\binom{yn}{3}$ be the number of $\{3, 1\}$ -sets;
- $t_1 := |W|(xn-1)(yn-1)$ be the number of pairs $(X, \{u, w\})$ where X is a $\{2, 2\}$ -set, $\{u, w\} \in W$ and $u, w \in X$;
- $t_2 := |W|(\binom{xn-1}{2} + \binom{yn-1}{2})$ be the number of pairs $(X, \{u, w\})$ where X is a $\{3, 1\}$ -set, $\{u, w\} \in W$ and $u, w \in X$;
- t_3 be the number of triples $(X, \{u, w\}, \{u', w'\})$, where $X = \{u, w, u', w'\}$ is a $\{2, 2\}$ -set and $\{u, w\}, \{u', w'\} \in W$;
- t_4 be the number of $\{3, 1\}$ -sets X that span two edges in W .

Let us show that

$$\Lambda_{4,3}(G) \leq t_0 + t_1 - t_2 - t_3 + t_4. \quad (23)$$

We prove this by listing in Table 3 the contribution of every 4-set $X \subseteq V(G)$ to both sides of (23). Note that there are no 4-sets spanning 3 or more wrong edges (because W has maximum degree $d \leq 2$ and each neighbour of a degree-2 vertex must have degree 1, as we proved earlier). The inequality in (23) follows from the fact that, in each row of Table 3, the sum of the entries in the last five columns is at least the entry in the $\Lambda_{4,3}(G)$ -column.

$\{ X \cap V_0 , X \cap V_1 \}$	$\binom{X}{2} \cap W$	$\Lambda_{4,3}(G)$	t_0	t_1	$-t_2$	$-t_3$	t_4
$\{4,0\}$	necessarily empty	0	0	0	0	0	0
$\{3,1\}$	empty	1	1	0	0	0	0
$\{3,1\}$	single edge	0	1	0	-1	0	0
$\{3,1\}$	2-edge path	0	1	0	-2	0	1
$\{2,2\}$	empty	0	0	0	0	0	0
$\{2,2\}$	single edge	1	0	1	0	0	0
$\{2,2\}$	2-edge matching	0	0	2	0	-2	0
$\{2,2\}$	2-edge path	0	0	2	0	0	0

Table 3: Proof of the inequality in (23).

By $d \leq 2$, we have that $t_3 \geq |W|^2 - 2n$. Also, since $uw, uw' \in W$ with $w \neq w'$ implies that each of the vertices w, w' has wrong degree 1, we have that $t_4 \leq |W|/2 \cdot n/2 + o(n^2) = (a/4 + o(1))n^2$. It follows via routine calculations that if we substitute these bounds into (23) and then subtract the lower bound $\Lambda_{4,3}(G) \geq \binom{n}{4}/2 + (1/8 + o(1))n^2$ coming from (17) (and use that $|W| = an$ while x and y are respectively $1/2 \pm c/\sqrt{n}$) then the terms in front of n^4 and n^3 cancel each other, while the coefficient at n^2 is $p(a, c) + o(1)$, where

$$p(a, c) := -a^2 - 2ac^2 + \frac{3a}{4} - \frac{c^4}{3} + \frac{c^2}{2} - \frac{3}{16}.$$

Thus it holds that $0 \leq p(a, c) + o(1)$. Since p is an even function of c , we can assume for the calculations in this paragraph that $c \geq 0$. The roots of the partial derivative $\partial p/\partial c$ are 0 and $\pm c_0$, where $c_0 := \sqrt{3 - 12a}/2$. If $a \geq 1/4$ then p as a function of $c \geq 0$ is decreasing and

$$p(a, c) \leq p(a, 0) = -a^2 + \frac{3a}{4} - \frac{3}{16} \leq \left(\frac{3}{8}\right)^2 + \frac{3 \cdot \frac{3}{8}}{4} - \frac{3}{16} = -\frac{3}{64} < 0,$$

a contradiction. So suppose that $a \leq 1/4$. Here it holds that $p(a, c) \leq p(a, c_0) = 2a^2 - 3a/4$. The maximum of this function on $[0, 1/4]$ is 0 which is attained if and only if $a = 0$. Furthermore, $p(0, c) = -(4c^2 - 3)^2/48$ is 0 only if $c = c_0 = \sqrt{3}/2$.

Thus, by the standard compactness argument, $p(a, c) \geq o(1)$ implies that $a = o(1)$ and $c = \pm\sqrt{3}/2 + o(1)$. However then (22) is violated. This contradiction shows that $d = 0$ and finishes the proof of Theorem 19. \square

5 Structural results for the 4-cycle plus a pendant edge

Here we prove an Erdős–Simonovits type stability result as well as some partial structural information about large extremal graphs when F is the 4-cycle plus a pendant edge.

Theorem 21. *Let $F := (5, \{01, 12, 23, 30, 04\})$. Then every sequence of graphs $(G_n)_{n \in \mathbb{N}}$ of growing orders with $\lambda_F(G_n) = \lambda_F + o(1)$ as $n \rightarrow \infty$ consists, up to $o(v(G_n)^2)$ edits in each G_n , of balanced bipartite $5/6$ -quasirandom graphs.*

Proof. Choose positive constants, each being sufficiently small depending on the previous ones:

$$\varepsilon_5 \gg \varepsilon_4 \gg \varepsilon_3 \gg \varepsilon_2 \gg \varepsilon_1 \gg \varepsilon_0.$$

Let G be an arbitrary graph with $n \geq 1/\varepsilon_0$ vertices and $\lambda_F(G) \geq \lambda_F - \varepsilon_0$. Recall that, for a type τ on $[q]$, an integer $s > q$, and an embedding $f : \tau \hookrightarrow G$, the vector $\mathbf{v}_{(G,f)}^{\tau,s}$ lists the densities of s -vertex τ -flags in (G, f) as defined in (8). We will need to compare these with the analogous densities in a bipartite random graph R_n which consists of two independent $(n/2)$ -sets V_0 and V_1 where each pair across is an edge with probability $5/6$ and these events are mutually independent. If two maps f and f' assign each element of $[q]$ to the same part V_i then, by standard concentration results (e.g. the Azuma-Hoeffding Inequality), the corresponding density vectors are close to each other. Thus we write R_∞ to refer to the limit of R_n as $n \rightarrow \infty$ and, instead of f , we list just the set F_0 , where $F_i := f^{-1}(V_i)$ for $i \in [2]$ (that is, F_i consists of the vertices of τ that are mapped by f into V_i).

An explicit formula for $\mathbf{v}_{(R_\infty,A)}^{\tau,s}$ can be written as follows. Namely, if at least one of $F_0 := A$ or $F_1 := [q] \setminus A$ spans at least one edge in τ then output 0. Otherwise, starting with the graph τ on $[q]$, add $S := [s] \setminus [q]$ to the vertex set, take a uniform random partition $S_0 \cup S_1$ of the $(s - q)$ -set S and make each pair in $(F_0 \times S_1) \cup (F_1 \times S_0) \cup (S_0 \times S_1)$ an edge with probability $5/6$ with all choices being mutually independent. Now, for a τ -flag $F \in \mathcal{F}_s^\tau$, the F -th entry of $\mathbf{v}_{(R_\infty,A)}^{\tau,s}$ is the probability that the obtained random τ -flag on $[s]$ with roots $0, \dots, q - 1$ is isomorphic to F . Of course, by the symmetry between the parts, this value will not change if we replace A by $[q] \setminus A$.

Since ε_0 is sufficiently small depending on ε_1 , we can assume that the conclusion of Lemma 1 holds with respect to ε_1 for τ being each of the four types $1, \nu, \sigma, \mu$ appearing in Properties A)–D) of Theorem 17. Since the matrix X^τ is positive semi-definite, any vector \mathbf{x} with non-negative entries summing to 1 and satisfying $\|X^\tau \mathbf{x}\| \leq o(1)$ is $o(1)$ -close, say in the supremum norm, to the zero eigenspace of X^τ . As the bipartite quasirandom graph R_n for large n also satisfies Lemma 1, the limiting vector $\mathbf{v}_{(R_\infty,A)}^{\tau,s}$, where $s := (N + q)/2$, lies in the zero eigenspace of $X^\tau \not\approx 0$ for every $A \subseteq [q]$. If the zero eigenspace of X^τ

happens to be 1-dimensional, as it is the case for $\tau = 1$ and $\tau = \nu$ by [A](#)) and [B](#)), then the vectors $\mathbf{v}_{(G,f)}^{\tau,s}$ and $\mathbf{v}_{(R_\infty, f^{-1}(V_0))}^{\tau,s}$ must be close to each other, except for a small number of embeddings f . We conclude that the rooted densities of 3-vertex 1-flags (resp. 4-vertex ν -flags) in G are close to what we observe in R_n . Since there is only one way to assign the vertices of 1 or ν to the parts of R_n (up to swapping V_0 and V_1) we omit $f^{-1}(V_0)$ in the formulas below.

Claim 22.

(i) For all but at most $\varepsilon_2 n$ maps $f : \{0\} \rightarrow V(G)$ it holds that $\left\| \mathbf{v}_{(G,f)}^{1,3} - \mathbf{v}_{R_\infty}^{1,3} \right\|_\infty \leq \varepsilon_2$.

(ii) For all but at most $\varepsilon_2 n^3$ embeddings f of ν into G , it holds that $\left\| \mathbf{v}_{(G,f)}^{\nu,4} - \mathbf{v}_{R_\infty}^{\nu,4} \right\|_\infty \leq \varepsilon_2$. □

The first part of Claim 22 implies by simple averaging over $f(0) \in V(G)$ that, for every 3-vertex (unlabelled) graph F , its density in G is within, say, $3\varepsilon_2$ of its limit density $p(F, R_\infty) := \lim_{n \rightarrow \infty} p(F, R_n)$, that is,

$$|p(F, G) - p(F, R_\infty)| \leq 3\varepsilon_2, \quad \text{for every } F \in \mathcal{F}_3^0. \tag{24}$$

Note that $p(\nu, R_\infty) = (3/4) \cdot (5/6)^2 = 25/48$, as it equals to the probability that 3 random vertices are not all in one part times the probability that the two corresponding crossing pairs are edges. Thus G has many cherries by (24) so, in particular, the second part of Claim 22 is not vacuous. Also, we can express the edge density in G via 3-vertex densities, namely, $p(K_2, G) = (1/3) \sum_{F \in \mathcal{F}_3^0} |E(F)| p(F, G)$. Thus it follows from (24) that $p(K_2, G)$ is within $(1/3 + 2/3 + 1) \cdot 3\varepsilon_2 = 6\varepsilon_2$ of the value in R_∞ , which is $(1/2) \cdot (5/6) = 5/12$.

If we have a cherry with edges ab and ac then we say that bc is its *base* or that the cherry is *based* on bc . Note that all pairs of non-adjacent vertices $u_0 u_1$ in R_n are of two types: if the vertices are in two different parts then there are no cherries at all based on them and otherwise, the density of cherries based on $u_0 u_1$ is approximately $(1/2) \cdot (5/6)^2 = 25/72$. Let us show that a similar classification of non-adjacent pairs is possible in G . Define B (resp. C) to consist of those $u_0 u_1 \in E(\overline{G})$ such that the number of cherries based on $u_0 u_1$ is $(25/72 \pm \varepsilon_3)n$ (resp. at most $\varepsilon_3 n$).

Claim 23. $|E(\overline{G}) \setminus (B \cup C)| \leq \varepsilon_3 \binom{n}{2}$.

Proof of Claim. Call a cherry in G *bad* if the corresponding embedding $f : [3] \rightarrow V(G)$ fails the conclusion of Claim 22(ii). By Claim 22, we know that there are at most $\varepsilon_2 n^3$ (vertex labelled) bad cherries. Call a pair *bad* if it is the base of at least $\varepsilon_3 n/2$ bad cherries. Clearly, the number of (ordered) bad pairs is at most $\varepsilon_2 n^3 / (\varepsilon_3 n/2) < \varepsilon_3 n(n-1)$.

So it is enough to show that any pair $u_0 u_1 \in E(\overline{G})$ which is not bad is in $B \cup C$. Assume that at least $\varepsilon_3 n$ cherries are based on it as otherwise $u_1 u_0 \in C$, as desired. Of these cherries, less than $\varepsilon_3 n/2$ are bad so there is a vertex u that makes a good cherry with the base $u_0 u_1$. If we take a uniform random vertex w in $V(G) \setminus \{u_0, u_1, u\}$ then the probabilities of the possible adjacencies of w to the good cherry on $\{u_0, u_1, u\}$ are each

within ε_2 from the corresponding values for a cherry in R_∞ . Therefore, the probability that w is attached to both u_0 and u_1 (which can be written as the sum of two densities depending on whether uw is an edge or not) is approximated within $2\varepsilon_2 < \varepsilon_3$ by the analogous probability in R_∞ . The latter is exactly $25/72$, giving the required. \square

By (24), $p(\nu, G)$ is within $3\varepsilon_2$ from $p(\nu, R_\infty) = 25/48$. On the other hand, the total number of cherries $P(\nu, G)$ can be computed as the sum of the number of cherries based on u_0u_1 for all $u_0u_1 \in E(\overline{G})$, which is at least $|B| \cdot (25/72 - \varepsilon_3)n$. Thus

$$|B| \leq \frac{\left(\frac{25}{48} + 3\varepsilon_2\right) \binom{n}{3}}{\left(\frac{25}{72} - \varepsilon_3\right)n} \leq \left(\frac{1}{2} + 2\varepsilon_3\right) \binom{n}{2}. \quad (25)$$

Let $H := (V(G), E(G) \cup C)$. Let us show that there are at most $4\varepsilon_4 \binom{n}{3}$ triangles in H . There are four different types of a triangle, depending on how many of its edges are in G . So it is enough to bound the number of triangles in H of each type by $\varepsilon_4 \binom{n}{3}$.

The number of triangles that take all three edges from G is obviously $p(K_3, G) \binom{n}{3} < \varepsilon_4 \binom{n}{3}$.

Next, note that every pair in $C \subseteq E(\overline{G})$ is in at most $\varepsilon_3 n$ cherries by definition; thus the number of triangles in H with exactly two edges from G is at most $\binom{n}{2} \cdot \varepsilon_3 n \leq \varepsilon_4 \binom{n}{3}$.

Let us turn to triangles with exactly one edge from G . Call an edge $u_0u_1 \in E(G)$ *bad* if it is in at least $\varepsilon_3 n$ triangles of G . Since the total number of triangles in G is at most $3\varepsilon_2 \binom{n}{3}$ by (24), we have at most $3 \cdot 3\varepsilon_2 \binom{n}{3} / (\varepsilon_3 n) \leq (\varepsilon_4/6) \binom{n}{2}$ bad edges and thus at most $(\varepsilon_4/2) \binom{n}{3}$ triangles in H whose unique G -edge is bad. Next, take any triangle $u_0u_1u_2$ in H whose unique edge from G , say $u_0u_1 \in E(G)$, is not bad. If we take a uniform random vertex $w \in V(G) \setminus \{u_0, u_1, u_2\}$ then the probability that it sends at least two edges to $u_0u_1u_2$ is at most $3\varepsilon_3 n / (n - 3) < 4\varepsilon_3$, since each pair u_iu_j (which is a non-bad edge or a pair in C) belongs to at most $\varepsilon_3 n$ triangles. Thus the vector $\mathbf{v}_{(G,f)}^{\mu,4}$ of the densities of 4-vertex μ -flags in (G, f) for the embedding f of μ into G that sends i to u_i is $O(\varepsilon_3)$ -close to being supported on μ -flags where the free vertex sends at most one edge to the roots. By Property D), every such vector of ℓ_1 -norm 1 must be $\Omega(1)$ -far from the zero eigenspace of the matrix X^μ . Since $\varepsilon_3 \ll \varepsilon_4$ is sufficiently small, there are at most $\varepsilon_4 \binom{n}{3}$ such maps f and thus at most $(\varepsilon_4/2) \binom{n}{3}$ corresponding unordered triples $\{u_0, u_1, u_2\}$ (since each such triple gives 2 embeddings of μ). We conclude that the total number of triangles in H with exactly one edge from G is at most $\varepsilon_4 \binom{n}{3}$, as desired.

Finally, the argument for upper bounding the number of triangles with all three pairs coming from C is similar to the one from the previous paragraph, except each such triangle spans a copy of the edgeless type σ in G and we use Property C).

Thus we have shown that H has at most $4\varepsilon_4 \binom{n}{3}$ triangles and, by Claim 23 and (25), at least $\binom{n}{2} - |B| - \varepsilon_3 \binom{n}{2} \geq (1/2 - 21\varepsilon_3) \binom{n}{2}$ edges. The Erdős–Simonovits Stability Theorem [22, 52] implies that there is a partition $V_0 \cup V_1$ such that H (and thus G) has at most $\varepsilon_5 \binom{n}{2}$ edges inside V_0 or V_1 . Since we can choose ε_5 arbitrarily small, the graph G (or, more precisely, any sequence of almost extremal graphs) is almost bipartite and its $(5/6)$ -quasirandomness follows (by e.g. [18, Lemma 14]) since G has the correct density of edges and of all 4-vertex subgraphs containing a cherry (and thus the correct non-induced homomorphic density of 4-cycles) by Claim 22(ii). \square

Some further information can be derived about the structure of large extremal graphs.

Theorem 24. *Let F be the 4-cycle with a pendant edge. For any $\varepsilon > 0$ there is n_0 such that any graph G of order $n \geq n_0$ with $\lambda_F(G) = \lambda_F(n)$ admits a partition $V(G) = V_0 \cup V_1$ such that each V_i is an independent set and each vertex of G has $(5/12 \pm \varepsilon)n$ neighbours in the other part.*

Proof. Let G be any λ_F -extremal graph with $n \rightarrow \infty$ vertices. For brevity and since the meaning will be clear, we will hide all negligible constants under $o(1)$ terms.

By Theorem 21, there is a partition $U_0 \cup U_1 = V(G)$ such that each part has $(1/2 + o(1))n$ vertices and spans $o(n^2)$ edges while the bipartite graph $G[U_0, U_1]$ (or rather the corresponding sequence of graphs as $n \rightarrow \infty$) is p -quasirandom, where we denote $p := 5/6$. Let $V_0 \cup V_1$ be a partition of $V(G)$ that maximises the number of crossing edges. Note that this number is at least $|E(G[U_0, U_1])| = p(n/2)^2 + o(n^2)$ and at most $|E(G)| = p(n/2)^2 + o(n^2)$. Thus almost all edges of G must go between V_0 and V_1 . By the quasirandomness, we have up to swapping the parts that $|V_i \Delta U_i| = o(n)$ for each $i \in [2]$. Thus all conclusions of Theorem 21 also apply to the partition $V_0 \cup V_1$ (with error terms that are worse but are still $o(1)$).

Here it is more convenient to count embeddings of F into G (rather than 5-subsets that span a subgraph isomorphic to F). Since F has only two automorphisms, we have by (5) that $t(F, G)$ is $\gamma + o(1)$, where

$$\gamma := \frac{2}{5!} \lambda_F = \frac{1}{60} \frac{5^6}{2^8 \cdot 3^5} = \frac{5^5}{2^{10} \cdot 3^6} = \frac{p^5(1-p)}{2^4}.$$

Thus, by (6), the expected number of embeddings of F into G that use a uniformly random vertex is $v(F) \cdot t(F, G)n^5/n = (5\gamma + o(1))n^4$. Since every two distinct vertices of G are simultaneously in at most $5 \cdot 4 \cdot n^3 = o(n^4)$ embeddings of F into G and we cannot increase $t(F, G)$ by replacing a vertex by a clone of another vertex, we conclude that every vertex of G is in least $(5\gamma + o(1))n^4$ embeddings.

The rest of the proof relies on rather laborious calculations, which are also included in the ancillary notebook. Namely, the notebook verifies Claim 25 below by computing the polynomial q and checking local optimums both in the interior and on the boundary. Also, the notebook calculates the rate of decrease in the induced copies of F under the addition of an edge inside a part, a quantity that we have to estimate later when we argue that V_i spans no edges.

Take any vertex u of G . If u has $x_i n$ neighbours in V_i and $y_i n$ non-neighbours in V_i for the max-cut partition $V(G) = V_0 \cup V_1$ then the number of embeddings f of F containing u is by Theorem 17 (that is, by the 5/6-quasirandomness of $G[V_0, V_1]$) equal to

$$\sum_{i=0}^1 \left((x_i x_{1-i}^2 y_i p^2 (1-p)^2 + x_{1-i}^3 y_i p^2 (1-p)) \right. \\ \left. + 2 \cdot (x_i^2 y_{1-i}^2 p^3 (1-p)) + x_i^2 y_i y_{1-i} p^3 + x_i y_i y_{1-i}^2 p^4 \right) n^4 + o(n^4), \quad (26)$$

where the four terms under the sum count those f that map 0, one of 1 or 3, 2, and 4 respectively to the vertex u . (Recall that $E(F) = \{01, 12, 23, 30, 04\}$.) For example, let us show that the first term counts those embeddings f that send 0 to u . Since we have $o(n^2)$ edges inside each part, there are only $o(n^4)$ embeddings that map an edge of F inside some part of $G - u$; so let us exclude these. Thus, if $i \in [2]$ is the index with $f(2) \in V_i$ then $f(1), f(3) \in V_{1-i}$. Since 4 sends no edges to $\{1, 2, 3\}$ in the graph F , its image $f(4)$ can be in both V_i or V_{1-i} , which will give the two summands making the first term. For example, if $f(4) \in V_i$ then we have to count the number of choices of $f(2) \in V_i \setminus \Gamma(u)$, $f(1), f(3) \in V_{1-i} \cap \Gamma(u)$ and $f(4) \in V_i \cap \Gamma(u)$ such that $f(2)$ (resp. $f(4)$) is adjacent to both (resp. none) of $f(1)$ and $f(3)$. Since the induced bipartite graph $G[V_0, V_1]$ is p -quasirandom, the number of choices as above is $y_i n \cdot (x_{1-i} n)^2 \cdot x_i n \cdot p^2 (1-p)^2 + o(n^4)$, giving the first summand. The other cases are analogous.

Substituting $y_i = 1/2 - x_i + o(1)$, we get that the expression in (26) is $(q(x_0, x_1) + o(1))n^4$, where for $a, b \in \mathbb{R}$ we define $q(a, b)$ to be

$$\left(\frac{25ab}{54} + \frac{125}{1728}\right)(a^2 + b^2) - \frac{25}{108}(a^3 + b^3) - \frac{50a^2b^2}{81} + \left(\frac{325ab}{1296} + \frac{625}{10368}\right)(a + b) - \frac{625ab}{1296}.$$

A routine calculation confirms that $q(0, 5/12) = 5\gamma$, which is in accordance with our counting that led to the definition of q . Let us show that this is in fact the maximum value of q over $[0, 1/2]^2$ and the only maximisers of q are those pairs (a, b) that we expect to see in an extremal construction:

Claim 25. *If $(a, b) \in [0, 1/2]^2$ satisfies $q(a, b) \geq 5\gamma$ then $\{a, b\} = \{0, 5/12\}$.*

Proof of Claim. Suppose that (a, b) is a counterexample.

Suppose first that $c := q(a, b) - 5\gamma > 0$. Take n sufficiently large depending on c . Consider a graph G' obtained from G by removing a ‘‘typical’’ vertex w and adding a new vertex w' with $(a+o(1))n$ and $(b+o(1))n$ neighbours in the two parts. The new vertex is in at least $(5\gamma + c + o(1))n^4$ embeddings of F which is by $(c+o(1))n^4 > 20n^3$ larger than the number of embeddings destroyed by removing w . Thus $\Lambda_F(G') > \Lambda_F(G)$, a contradiction to the maximality of G . We conclude that $q(a, b) = 5\gamma$ and it is a maximiser of q on $[0, 1/2]^2$.

Next, suppose that (a, b) lies in the interior of $[0, 1/2]^2$. Since q as a polynomial is a differentiable function, (a, b) must be a critical point. We add a new variable $z \in \mathbb{R}$ and run the Buchberger Algorithm (which is a standard function in SageMath) to eliminate variables x and y from the following system of 3 polynomial equations: $z - q(x, y) = 0$, $\frac{\partial q}{\partial x}(x, y) = 0$ and $\frac{\partial q}{\partial y}(x, y) = 0$. The algorithm produces an explicit polynomial $Q(z)$ depending only on z which is in the ideal generated by these 3 polynomials, namely

$$\begin{aligned} Q(z) = & 98795242729650675885117146136576 z^6 \\ & + 18009465447674020572152337530880 z^5 - 400416701599337361136680960000 z^4 \\ & + 30661545256257839254732800000 z^3 - 264291334776391680000000000 z^2 \\ & - 674085801678710937500000 z - 370573902130126953125. \end{aligned}$$

Since the triple $(x, y, z) = (a, b, 5\gamma)$ is a root of each of the 3 equations, it must hold that $Q(5\gamma) = 0$. However, one can verify using exact arithmetic (see e.g. the calculations in our script) that 5γ is not a root of Q , a contradiction.

Thus the point (a, b) lies on the boundary of $[0, 1/2]^2$. Up to symmetry, there are only the following two cases to consider. First, let $b = 1/2$. We have that

$$q(a, 1/2) = \frac{25a^2}{576} - \frac{625a}{10368} + \frac{25}{1296}.$$

The derivative of this quadratic polynomial has zero at $a = 25/36 > 1/2$, so its maximum value on $[0, 1/2]$ is $q(0, 1/2) = 25/1296 < 5\gamma$, a contradiction. Finally, let $b = 0$. We have

$$q(a, 0) = -\frac{25a^3}{108} + \frac{125a^2}{1728} + \frac{625a}{10368}.$$

The derivative of this cubic polynomial has roots $-5/24$ and $5/12$ so its unique maximiser on $[0, 1/2]$ is $a = 5/12$ (when the value of q is 5γ), contradicting our initial assumption. \square

Similarly as we derived the second part of Claim 20 from the first, Claim 25 gives by a compactness argument that $q(a, b)$ is approximately 5γ only if $\{a, b\}$ is close to $\{0, 5/12\}$, which in particular applies to $\{x_0, x_1\}$. Since $V_0 \cup V_1$ is a max-cut partition, we conclude that every vertex of G has $o(n)$ neighbours in its part and $(5/12 + o(1))n$ neighbours in the other part.

We can now show that each part spans no edges. Suppose on the contrary that some edge $uw \in \binom{V_i}{2}$ violates this. Let

$$y := \frac{|V_{1-i} \cap \Gamma(u) \cap \Gamma(w)|}{n/2}$$

be the proportion of vertices of V_{1-i} that are adjacent to both u and w . Since the bipartite graph $G[V_0, V_1]$ is almost $5/6$ -regular, we have that $2p - 1 \leq y + o(1) \leq p$.

Let us compare G with the graph G' which is obtained from G by making the pair uw a non-edge. First, let us count embeddings of F into G that use the edge uw . Since the maximum degree inside each part is $o(n)$, there are only $o(n^3)$ embeddings of F into G that use an edge inside a part different from the pair uw . Every other embedding has to map $\{0, 4\}$ to $\{u, w\}$, having 2 choices here. Then the images of both $1, 3 \in V(F)$ have to be from the set of non-neighbours of $f(4)$ intersected with the neighbours of $f(0)$, whose size is $(p - y)n/2 + o(n)$. Finally, we have to choose $f(2)$ adjacent to both $f(1)$ and $f(3)$. By quasirandomness, the total number of embeddings is at most $2 \cdot ((p - y)n/2)^2 \cdot n/2 \cdot p^2 + o(n^3)$.

On the other hand, if we make uw a non-edge then the number of new F -embeddings that use both u and w can be lower bounded as follows. Again, up to an additive error term $o(n^3)$, we can pretend that each V_i is an independent set. We look at embeddings that map the non-edge $\{1, 3\} \in E(\overline{F})$ to $\{u, w\}$. There are 2 choices for the bijection between $\{1, 3\}$ and $\{u, w\}$. Then we send each of 0 and 2 to the common neighbourhood of u and w which has size $y(n/2) + o(1)$ and send 4 to any non-neighbour of u and w in V_i , having $n/2 + o(n)$ choices. We have to ensure additionally that $f(4)$ is adjacent

to $f(0)$ but not to $f(2)$, so we have to multiply the total number of choices by $p(1-p)$ by quasirandomness. Thus the number of new embeddings f with $f(13) = uw$ is at least $2 \cdot (y(n/2))^2 \cdot n/2 \cdot p(1-p) + o(n^3)$. Actually, this lower bound alone suffices for a contradiction (so we do not need to look at embeddings where some other non-edge of F is mapped to uw).

Thus if we take the difference between the number of embeddings of F into G' and G , and normalise this by n^3 then the main term is at least

$$\frac{5y^2}{36 \cdot 4} - \frac{25(5/6 - y)^2}{36 \cdot 4} =: r(y).$$

This quadratic polynomial is clearly increasing with y so the minimum value of r on $[2/3, 5/6]$ is $r(2/3) = 55/5184 > 0$. This contradicts the maximality of G and proves that each V_i is an independent set, finishing the proof of Theorem 24. \square

6 Concluding remarks

There are a number of pairs (κ, ℓ) for which we could not determine the edge-inducibility constant. Table 4 presents numerical upper bounds returned by computer (using the maximum computationally feasible value $N = 8$) and the lower bounds that we could find, where $P_{n/4, n/4, n/4, n/4}^{\bullet \circ \circ \bullet}$ denotes the uniform blowup of the pattern $(4, \{00, 01, 12, 23, 33\})$, which is the 4-vertex path with loops at its two endpoints, while $\alpha_5, \alpha_6, \alpha_7$ are the real roots of

$$\begin{aligned} 404\alpha_5^3 - 310\alpha_5^2 + 82\alpha_5 - 7 &= 0, \\ 14\alpha_6^2 - 14\alpha_6 + 3 &= 0, \\ 294\alpha_7^4 - 390\alpha_7^3 + 190\alpha_7^2 - 40\alpha_7 + 3 &= 0, \end{aligned}$$

satisfying $\alpha_5 = 0.17301747\dots$, $\alpha_6 = 0.31101776\dots$, and $\alpha_7 = 0.17645827\dots$.

In these cases, we were content with an example giving a lower bound just fairly close to the upper bound (without making a serious effort of finding a best possible construction). So it is possible that some of our lower bounds listed in Table 4 can be easily improved.

There are a few problems where the numerical bound coming from flag algebras is very close to a known lower bound. The most interesting such case is probably the $(5, 5)$ -edge-inducibility problem, where we conjecture that $\lambda_{5,5} = 45/128$, with the lower bound coming from the uniform blowups of the 4-vertex path with loops at end-points. This problem is self-complementary but, unlike the $(4, 3)$ -case when there are two complementary constructions, here we happen to have only one construction (which is self-complementary). We were not able to round the floating-point matrices returned by the solver into rational matrices that prove the sharp upper bound of $45/128$. One difficulty is that working with $N = 7$ is not enough so one has to use 8-vertex flags (obtaining an SDP program with $|\mathcal{F}_8^0| = 12346$ linear constraints). Also, the rounding step is complicated since there are so-called ‘‘phantom’’ pairs, namely the pairs witnessing that the problem is not flip-averse (such as the pairs across the two parts for the $(4, 3)$ -edge-inducibility problem).

(κ, ℓ)	Construction	Lower bound	Numerical SDP value
(5, 5)	$P_{n/4, n/4, n/4, n/4}^{\bullet \circ \circ \bullet}$	$\frac{45}{128} = 0.3515625$	0.3515625031...
(6, 2)	$6K_{n/6}$	$\frac{25}{72} = 0.3472222222\dots$	0.3513749476...
(6, 3)	$2T_{\alpha_5 n, \alpha_5 n} + K_{(1-4\alpha_5)n}$	0.3650891908...	0.3653600283...
(6, 6)	$2T_{n/22, n/22, \dots, n/22}$	$\frac{21675}{58564} = 0.3701079161\dots$	0.370111574...
(7, 2)	$8K_{n/8}$	$\frac{11025}{32768} = 0.3364562988\dots$	0.3367351897...
(7, 3)	$K_{3n/7} + \overline{K}_{4n/7}$	$\frac{34560}{117649} = 0.293755153\dots$	0.2990366379...
(7, 4)	$3T_{n/6, n/6}$	$\frac{35}{108} = 0.3240740741\dots$	0.3269092898...
(7, 5)	$3K_{n/3}$	$\frac{70}{243} = 0.2880658436\dots$	0.2965188293...
(7, 7)	$T_{\alpha_6 n, \alpha_6 n, (1-2\alpha_6)n}$	0.2880864973...	0.2925927027...
(7, 8)	$T_{\alpha_7 n, \alpha_7 n, \alpha_7 n} + K_{(1-3\alpha_7)n}$	0.3538476174...	0.3538476175...

Table 4: The remaining unknown edge-inducibility constants for $\kappa \leq 7$. Column 2 describes a (block) construction while Column 3 presents the best lower bound on the edge-inducibility constant coming from this construction (that is, by taking optimal part ratios). Column 4 lists the numerical value returned by SDP solver for $N = 8$.

By looking at the lower order terms, one can show that phantom pairs force some further coefficients c_F to be zero and further vectors to be zero eigenvectors of some matrices X^τ in (9), when compared to the standard restrictions coming from evaluating this identity on the blowup construction (see [47, Equation (3.10)] or [25, Lemma 16]). The presence of such pairs seems to make rounding harder.

The (7, 8)-edge-inducibility problem is another instance where the bounds are close to each other. Here $N = 7$ does not suffice and, of course, the fact that the optimal part ratios involve the irrational number α_7 (a root of an irreducible polynomial of degree 4) makes this task even more challenging.

Further motivation for studying the edge-inducibility problem comes from the question of V.T.Sós [55] to describe graphons that are *size forcible*, that is, are determined up to weak isomorphism by the vector of the densities of (κ, ℓ) -configurations over all (κ, ℓ) . We refer the reader to [18, 19, 54] for the precise definitions and known results. Without going into the theory of graphons (for which we refer the reader to the book by Lovász [42]), let us just observe that if Erdős-Simonovits stability holds for the (κ, ℓ) -edge-inducibility problem with respect to a unique minimal blowup pattern B and $\mathbf{a} \in \mathbb{S}_m$ is the unique maximiser of $\lambda_{\kappa, \ell}(B(\mathbf{a}))$ then the limit of \mathbf{a} -blowups of B is size forcible. This is the case in all our new results listed in Table 1 except for $(\kappa, \ell) = (4, 3)$, so they provide many new examples of graphons that are size forcible (in fact, in the much stronger sense that

just one pair (κ, ℓ) suffices).

In terms of the inducibility constants of 5-vertex graphs, as far as we see, there are 5 remaining open cases, namely where $E(F)$ is $\{01, 12\}$ (2-edge path plus 2 isolated vertices), $\{01, 02, 03, 12\}$ (triangle with a pendant edge plus 1 isolated vertex), $\{01, 12, 23, 34\}$ (4-edge path), $\{01, 02, 03, 04, 12\}$ (triangle with two pendant edges adjacent to the same vertex) and $\{01, 02, 03, 12, 14\}$ (triangle with two pendant edges adjacent to two different vertices). In the first of these cases, when $F = P_3 + 2K_1$, the numerical bound given by flag algebra comes very close to the lower bound $\lambda_F \geq 174/625$ of Even-Zohar and Linial [23, Table 2] given by non-uniform blowups of $3K_2$, namely by $T_{n/10, n/10} + 2T_{n/5, n/5}$ as $n \rightarrow \infty$. In the remaining four cases, there seems to be a non-zero gap between the upper and lower bounds; we refer the reader to [23] for details.

In an ongoing project with Jun Gao, Jared León, Xizhi Liu and Shumin Sun, we started a systematic study of the F -inducibility problem for 6-vertex graphs F .

Recently, Basit, Granet, Horsley, Kündgen and Staden [7] introduced the following problem for a given blue-red edge-coloured graph H with κ vertices (and proved a number of initial results on it). For an (uncoloured) graph G , let $\Lambda_H(G)$ be the number of *embeddings* of H into G , that is injections $V(H) \rightarrow V(G)$ that map red (resp. blue) edges of H to edges (resp. non-edges) of G . Note that we do not stipulate anything about the images of the pairs that are not adjacent in H . The *H -semi-inducibility problem* asks for $\Lambda_H(n)$, the maximum of $\Lambda_H(G)$ over all n -vertex graphs G . One can consider the normalised version $\lambda_H(n) := \Lambda_H(n) / \prod_{i=0}^{\kappa-1} (n-i)$. Thus $\lambda_H(n)$ is the probability that a random injection $V(H) \rightarrow V(G)$ is an embedding of H into G . It is easy to show (see [7, Proposition 3.1]) that the limit $\lambda_H := \lim_{n \rightarrow \infty} \lambda_H(n)$ exists. We call λ_H the *semi-inducibility constant* of H . If H is a colouring of the complete graph on $V(H)$ then the H -semi-inducibility problem is, apart from a scaling factor, exactly the inducibility problem for the red subgraph of H , which was the main motivation in [7] for introducing this question.

Using flag algebras, we also determined in [9] the semi-inducibility constant λ_H for all H with 4 vertices except when H is the 3-edge path coloured blue-blue-red in this order (or is equivalent to this case). Also, the first author obtained a number of results for 5-vertex graphs H ; these findings will be presented in a follow-up paper. Here we present only the following two results, as they address two open problems highlighted by Basit et al [7, Problems 9.1 and 9.2].

Theorem 26. *Let H be the 6-edge cycle with its edges alternatively coloured blue and red as we go along the cycle. Then $\lambda_H = 1/2^6$.*

Proof. The upper bound can be obtained from flag algebras with $N = 6$ (with the certificate named `semiind_c6`) while the lower bound comes from $R(K_n, 1/2)$, that is, from $1/2$ -quasirandom graphs. \square

Theorem 27. *Let H be the 3-edge path with its edges coloured blue, red and blue in this order. Then $\lambda_H = 2^2/3^3$. Moreover, every sequence of almost extremal graphs is almost $1/3$ -regular.*

Proof. The upper bound can be proved via flag algebras with $N = 4$ (with the certificate named `semiind_p4`) while the lower bound comes from almost $1/3$ -regular graphs.

Furthermore, our scripts verify that if τ is any of the 2-vertex types $\tau_0 := (2, \{\})$ and $\tau_1 := (2, \{01\})$ (the non-edge and the edge), then every vector \mathbf{v} in the kernel of X^τ satisfies that

$$\sum_{F \in \mathcal{F}_3^\tau} (p(F_0, F) - p(F_1, F)) \mathbf{v}_F = 0, \quad (27)$$

where, for $i \in [2]$, F_i denotes the 3-vertex τ -flag with the (unique) unlabelled vertex adjacent to root i but not to root $1 - i$. (Thus each coefficient $p(F_i, F)$ in (27) is 0 or 1.)

Take any almost λ_H -extremal graph G of order $n \rightarrow \infty$. For any two distinct vertices $u_0, u_1 \in V(G)$, let τ be τ_0 if u_0 and u_1 are non-adjacent and let τ be τ_1 otherwise. If we evaluate the left-hand side of (27) when \mathbf{v} is the vector $\mathbf{v}_{(G, (u_0, u_1))}^{\tau, s}$ of the densities of 3-vertex τ -flags in the τ -flag G rooted at (u_0, u_1) , as defined in (8), then we obtain $(\deg_G(u_0) - \deg_G(u_1))/(n - 2)$. By Lemma 1, for most choices of $u_0, u_1 \in V(G)$, the vector $\mathbf{v}_{(G, (u_0, u_1))}^{\tau, s}$ is $o(1)$ -close to the null-space of X^τ ; then the normalised degrees of u_0 and u_1 must be $o(1)$ -close to each other. It follows that G is almost α -regular for some $\alpha = \alpha(n)$ (that potentially depends on n). Since G has $(\alpha + o(1))\binom{n}{2}$ edges and, for most of these edges, both endpoints have the complementary degree $(1 - \alpha + o(1))n$, it holds that $\lambda_H(G) = \alpha(1 - \alpha)^2 + o(1)$. This implies that $\alpha(n) = 1/3 + o(1)$, that is, extremal graphs are almost $1/3$ -regular, as claimed. \square

Note that every almost $1/3$ -regular graph G satisfies $\lambda_H(G) = 2^3/3^6 + o(1)$, so the above theorem gives a characterisation of almost extremal graphs. More can be said about an extremal graph G of order $n \rightarrow \infty$. For example, every vertex u of G must have degree $(1/3 + o(1))n$. Indeed, since we cannot increase $\lambda_H(G)$ by replacing u by a clone of any other vertex of G , it must be the case that number of embeddings of H that use u is at least $4\lambda_H(G)/n - 12n^2 = 16n^3/27 + O(n^2)$. On the other hand, by the almost $1/3$ -regularity of G this number is equal to

$$2 \cdot (n - \deg(u)) \cdot \frac{1}{3}n \cdot \frac{2}{3}n + 2 \cdot \deg(u) \cdot (n - \deg(u)) \cdot \frac{2}{3}n + o(n^3),$$

where the first (resp. second) term counts those embeddings when u corresponds to one of the two endpoints (resp. interior points) of the alternating 3-edge path. The obtained inequality implies that $d = (1/3 + o(1))n$. By bootstrapping and analysing more carefully the effect of various local changes on the objective function, further more precise results about the degrees and the structure of any extremal graph should be possible.

Since this paper is already rather long, we decided to restrict ourselves here to the above semi-inducibility results.

Independently of this work, the semi-inducibility problem for the alternating path of any odd length was resolved by Chen, Clemen and Noel [15] using entropy while the semi-inducibility constant for the alternating 6-cycle was determined by Chen and Noel [16]; see also Balogh, Lidicky, Mubayi, Pfender and Volec [6].

Acknowledgements

The authors thank the anonymous referee for many useful comments. Both authors were supported by ERC Advanced Grant 101020255.

References

- [1] N. Alon, M. Axenovich, and J. Goldwasser, *On hypercube statistics*, 2024. [arXiv:2410.20498](#).
- [2] N. Alon, E. Fischer, M. Krivelevich, and M. Szegedy, *Efficient testing of large graphs*, *Combinatorica* **20** (2000), 451–476.
- [3] N. Alon, D. Hefetz, M. Krivelevich, and M. Tyomkyn, *Edge-statistics on large graphs*, *Combin. Probab. Computing* **29** (2020), 163–189.
- [4] R. Baber and J. Talbot, *Hypergraphs do jump*, *Combin. Probab. Computing* **20** (2011), 161–171.
- [5] J. Balogh, P. Hu, B. Lidický, and F. Pfender, *Maximum density of induced 5-cycle is achieved by an iterated blow-up of 5-cycle*, *Europ. J. Combin.* **52** (2016), 47–58.
- [6] J. Balogh, B. Lidický, D. Mubayi, F. Pfender, and J. Volec, *Semi-inducibility of some small graphs*, 2026. [arXiv:2601.03433](#).
- [7] A. Basit, B. Granet, D. Horsley, A. Kündgen, and K. Staden, *The semi-inducibility problem*, 2025. [arXiv:2501.09842](#).
- [8] L. Bodnár, *FlagAlgebraToolbox: Flag algebra computations in SageMath*, 2026. [arXiv:2601.06590](#), the package is available at <https://github.com/bodnalev/sage>.
- [9] L. Bodnár and O. Pikhurko, *Semi-inducibility of 4-vertex graphs*, 2025. [arXiv:2510.24336](#).
- [10] L. Bodnár and O. Pikhurko, *Some exact values of the inducibility and statistics constants for hypercubes*, 2025. [arXiv:2503.03408](#).
- [11] L. Bodnár and O. Pikhurko, *Some exact inducibility-type results for graphs via flag algebras*, 2026. [arXiv:2507.01596](#).
- [12] B. Bollobás, Y. Egawa, A. Harris, and G. Jin, *The maximal number of induced r -partite subgraphs*, *Graphs Combin.* **11** (1995), 1–19.
- [13] B. Bollobás, C. Nara, and S. Tachibana, *The maximal number of induced complete bipartite graphs*, *Discrete Math.* **62** (1986), 271–275.
- [14] J. I. Brown and A. Sidorenko, *The inducibility of complete bipartite graphs*, *J. Graph Theory* **18** (1994), 629–645.
- [15] H. Chen, F. C. Clemen, and J. A. Noel, *Maximizing alternating paths via entropy*, 2025. [arXiv:2505.03903](#).

- [16] H. Chen and J. A. Noel, *On alternating 6-cycles in edge-coloured graphs*, 2025. [arXiv:2505.09809](https://arxiv.org/abs/2505.09809).
- [17] F. R. K. Chung, R. L. Graham, and R. M. Wilson, *Quasi-random graphs*, *Combinatorica* **9** (1989), 345–362.
- [18] O. Cooley, M. Kang, and O. Pikhurko, *On a question of Vera T. Sós about size forcing of graphons*, *Acta Math. Hungar.* **168** (2022), 1–26.
- [19] E. Csóka, *On the graph limit question of Vera T. Sós*, *J. Combin. Theory (B)* **116** (2016), 306–311.
- [20] M. K. de Carli Silva, F. M. de Oliveira Filho, and C. M. Sato, *Flag algebras: a first glance*, *Nieuw Arch. Wiskd. (5)* **17** (2016), 193–199.
- [21] P. Erdős and M. Simonovits, *Supersaturated graphs and hypergraphs*, *Combinatorica* **3** (1983), 181–192.
- [22] P. Erdős, *On bipartite subgraphs of graphs*, *Mat. Lapok* **18** (1967), 283–288. In Hungarian.
- [23] C. Even-Zohar and N. Linial, *A note on the inducibility of 4-vertex graphs*, *Graphs Combin.* **31** (2015), 1367–1380.
- [24] G. Exoo, *Dense packings of induced subgraphs*, *Ars Combinatoria* **22** (1986), 5–10.
- [25] V. Falgas-Ravry, E. Marchant, O. Pikhurko, and E. R. Vaughan, *The codegree threshold for 3-graphs with independent neighbourhoods*, *SIAM J. Discr. Math.* **23** (2015), 1504–1539.
- [26] J. Fox and L. Saueremann, *A completion of the proof of the edge-statistics conjecture*, *Adv. Comb.* (2020), Paper No. 4, 52.
- [27] Z. Füredi, *A proof of the stability of extremal graphs, Simonovits’ stability from Szemerédi’s regularity*, *J. Combin. Theory (B)* **115** (2015), 66–71.
- [28] S. Gilboa, R. Glebov, D. Hefetz, N. Linial, and A. Morgenstern, *On the local structure of oriented graphs—a case study in flag algebras*, *Electron. J. Combin.* **29** (2022), #P3.39.
- [29] H. Hatami, J. Hirst, and S. Norine, *The inducibility of blow-up graphs*, *J. Combin. Theory (B)* **109** (2014).
- [30] D. Hefetz and M. Tyomkyn, *On the inducibility of cycles*, *Journal of Combinatorial Theory, Series B* **133** (2018), 243–258.
- [31] J. Hirst, *The inducibility of graphs on four vertices*, *J. Graph Theory* **75** (2014), 231–243.
- [32] V. Jain, M. Kwan, D. Mubayi, and T. Tran, *The edge-statistics conjecture for hypergraphs*, 2025. [arXiv:2505.03954](https://arxiv.org/abs/2505.03954).
- [33] V. Jain, M. Michelen, and F. Wei, *The binomial random graph is a bad inducer*, 2023. [arXiv:2306.13014](https://arxiv.org/abs/2306.13014).

- [34] G. Jeong, S. Park, and H. Yang, *An introduction to Razborov’s flag algebra as a proof system for extremal graph theory*, 2026. [arXiv:2601.12741](#).
- [35] D. Král’, S. Norin, and J. Volec, *A bound on the inducibility of cycles*, J. Combin. Theory (A) **161** (2019), 359–363.
- [36] M. Kwan, B. Sudakov, and T. Tran, *Anticoncentration for subgraph statistics*, J. Lond. Math. Soc. (2) **99** (2019), 757–777.
- [37] M. Kwan and L. Saueremann, *Resolution of the quadratic Littlewood–Offord problem*, 2023. [arXiv:2312.13826](#).
- [38] B. Lidický, C. Mattes, and F. Pfender, *C_5 is almost a fractalizer*, J. Graph Theory **104** (2023), 220–244.
- [39] H. Liu, O. Pikhurko, M. Sharifzadeh, and K. Staden, *Stability from graph symmetrisation arguments with applications to inducibility*, J. Lond. Math. Soc. (2) **108** (2023), 1121–1162.
- [40] X. Liu, J. Ma, and T. Zhu, *The inducibility of Turán graphs*, 2026. [arXiv:2601.10548](#).
- [41] X. Liu, D. Mubayi, and C. Reiher, *The feasible region of induced graphs*, J. Combin. Theory (B) **158** (2023), 105–135.
- [42] L. Lovász, *Large networks and graph limits*, Colloquium Publications, Amer. Math. Soc., 2012.
- [43] A. Martinsson, F. Mousset, A. Noever, and M. Trujić, *The edge-statistics conjecture for $\ell \ll k^{6/5}$* , Israel J. Math. **234** (2019), 677–690.
- [44] S. Norin and L. Yepremyan, *Turán number of generalized triangles*, J. Combin. Theory (A) **146** (2017), 312–343.
- [45] S. Norin and L. Yepremyan, *Turán numbers of extensions*, J. Combin. Theory (A) **155** (2018), 476–492.
- [46] O. Pikhurko, J. Sliachan, and K. Tyros, *Strong forms of stability from flag algebra calculations*, J. Combin. Theory (B) **135** (2019), 129–178.
- [47] O. Pikhurko and E. R. Vaughan, *Minimum number of k -cliques in graphs with bounded independence number*, Combin. Probab. Computing **22** (2013), 910–934.
- [48] N. Pippenger and M. C. Golumbic, *The inducibility of graphs*, J. Combin. Theory (B) **19** (1975), 189–203.
- [49] A. Razborov, *Flag algebras*, J. Symb. Logic **72** (2007), 1239–1282.
- [50] A. Razborov, *On 3-hypergraphs with forbidden 4-vertex configurations*, SIAM J. Discr. Math. **24** (2010), 946–963.
- [51] A. Roberts and A. Scott, *Stability results for graphs with a critical edge*, Europ. J. Combin. **74** (2018), 27–38.

- [52] M. Simonovits, *A method for solving extremal problems in graph theory, stability problems*, Theory of Graphs (Proc. Colloq., Tihany, 1966), 1968, pp. 279–319.
- [53] J. Širáň, *A new lower bound for the inducibility of a graph*, Mathematica Slovaca **34** (1984), 365–370.
- [54] J. Sliáčan, *Flag algebras and flagmatic*, 2015. MPhil Thesis, University of Warwick.
- [55] V. T. Sós, *Some open problems on dense graphs*, 2012. Conference on Graphs and Analysis, IAS, Princeton.
- [56] R. Ueltzen, *Characterizing graphs with high inducibility*, 2024. [arXiv:2411.17362](https://arxiv.org/abs/2411.17362).
- [57] R. Yuster, *On the exact maximum induced density of almost all graphs and their inducibility*, J. Combin. Theory (B) **136** (2019), 81–109.