Journal of Combinatorial Theory, Series B 135 (2019) 129–178

Contents lists available at ScienceDirect

Journal of Combinatorial Theory, Series B

www.elsevier.com/locate/jctb

Strong forms of stability from flag algebra calculations

Oleg Pikhurko [a,1], Jakub Sliačan [b,2,3], Konstantinos Tyros [c,2]

[a] *Mathematics Institute and DIMAP, University of Warwick, Coventry CV4 7AL, UK*
[b] *Department of Mathematics and Mathematical Statistics, Umeå University, Umeå, 901 87, Sweden*
[c] *Department of Mathematics, University of Athens, Panepistimiopolis, Athens, 157 84, Greece*

A R T I C L E I N F O

Article history:
Received 14 June 2017
Available online 13 August 2018

Keywords:
Edit distance
Flag algebras
Stability
Subgraph count

A B S T R A C T

Given a hereditary family $\mathcal{G}$ of *admissible* graphs and a function $\lambda(G)$ that linearly depends on the statistics of order-κ subgraphs in a graph G, we consider the extremal problem of determining $\lambda(n,\mathcal{G})$, the maximum of $\lambda(G)$ over all admissible graphs G of order n. We call the problem *perfectly B-stable* for a graph B if there is a constant C such that **every** admissible graph G of order $n \geqslant C$ can be made into a blow-up of B by changing at most $C(\lambda(n,\mathcal{G})-\lambda(G))\binom{n}{2}$ adjacencies. As special cases, this property describes all almost extremal graphs of order n within $o(n^2)$ edges and shows that every extremal graph of order $n \geqslant C$ is a blow-up of B.

We develop general methods for establishing stability-type results from flag algebra computations and apply them to concrete examples. In fact, one of our sufficient conditions for perfect stability is stated in a way that allows automatic verification by a computer. This gives a unifying way to obtain computer-assisted proofs of many new results.

E-mail address: jakub.sliacan@gmail.com (J. Sliačan).

[1] Supported by ERC grant 306493 and EPSRC grant EP/K012045/1.
[2] Part of this research was done when this author was supported by ERC grant 306493.
[3] Part of this research was done when this author was working in the School of Mathematics and Statistics, The Open University.

https://doi.org/10.1016/j.jctb.2018.08.001
0095-8956/

1. Introduction

By the term graph, we mean finite *simple* graph, that is, without loops and multiple edges. For a graph G, we refer to the cardinality of its vertex set as the *order* of G and we denote it by $v(G)$. Moreover, for a subset X of the vertex set $V(G)$ of G, we denote by $G[X]$ the subgraph *induced* by X in G, that is, the graph having X as the vertex set and two nodes $x, y \in X$ are connected in $G[X]$ if and only if they are connected in G.

Let F and G be graphs of orders $v(F) \leqslant v(G)$. Call F a *subgraph* of G if there is a subset X of $V(G)$ such that $G[X]$ is isomorphic to F. (Thus a subgraph means an induced subgraph.) Let $P(F,G)$ be the number of $v(F)$-subsets of $V(G)$ that induce a subgraph isomorphic to F. Also, let

$$p(F,G) := P(F,G) \Big/ \binom{v(G)}{v(F)}$$

be the *(induced) density* of F in G; equivalently, $p(F,G)$ is the probability that a random $v(F)$-subset X of $V(G)$ induces a subgraph isomorphic to F.

Suppose that we have a (possibly infinite) family $\mathcal{F}$ of *forbidden* graphs. Call a graph G *admissible* or *$\mathcal{F}$-free* (and denote this by $G \in \mathrm{Forb}(\mathcal{F})$) if no $F \in \mathcal{F}$ is a subgraph of G. Let $\mathcal{G} := \mathrm{Forb}(\mathcal{F})$ be the family of all admissible graphs; clearly, $\mathcal{G}$ is a hereditary graph property, that is, every subgraph of some member of $\mathcal{G}$ belongs to $\mathcal{G}$, too.

Let κ be a positive integer. We denote by $\mathcal{G}^0_\kappa$ the set obtained by taking one representative from each isomorphism class of graphs in $\mathcal{G}$ of order κ. Clearly $\mathcal{G}^0_\kappa$ is a finite set. Let γ be a function from $\mathcal{G}^0_\kappa$ to the reals. It gives rise to two other functions defined on graphs G with $v(G) \geqslant \kappa$:

$$\begin{aligned} \Lambda(G) &:= \sum_{H \in \mathcal{G}^0_\kappa} \gamma(H) P(H,G), \\ \lambda(G) &:= \sum_{H \in \mathcal{G}^0_\kappa} \gamma(H) p(H,G) = \binom{v(G)}{\kappa}^{-1} \cdot \Lambda(G). \end{aligned} \tag{1}$$

One can view $\lambda(G)$ as the expected value of $\gamma(G[X])$, where X is a random κ-subset of $V(G)$.

Under the above conventions, consider the problem of maximising $\Lambda(G)$ over admissible graphs G of given order n. Namely, we define the extremal function

$$\Lambda(n, \mathcal{G}) := \max\{\Lambda(G) : G \in \mathcal{G},\ v(G) = n\} \tag{2}$$

and its density version $\lambda(n,\mathcal{G}) := \Lambda(n,\mathcal{G})/\binom{n}{\kappa}$. It is not hard to show (see Lemma 2.2) that the sequence $(\lambda(n,\mathcal{G}))_{n=\kappa}^{\infty}$ is non-increasing and therefore the following limit exists:

$$\lambda(\mathcal{G}) := \lim_{n\to\infty} \lambda(n,\mathcal{G}). \tag{3}$$

This is a rather general setting. As an illustration, here is one example (and the reader is encouraged to consult other concrete examples that can be found in Section 1.1).

Example 1.1 *(Turán function).* Let $\kappa = 2$, $\gamma(\overline{K}_2) = 0$ and $\gamma(K_2) = 1$, where by K_m we denote the complete graph of order m and by $\overline{G}$ the complement of a graph G. (Thus $\Lambda(G) = e(G)$ is the number of edges in G.) If $\mathcal{H}$ is any family of graphs and $\mathcal{H}^{\uparrow} \supseteq \mathcal{H}$ consists of all graphs that can be obtained from $H \in \mathcal{H}$ by adding some missing edges, then $\Lambda(n, \mathrm{Forb}(\mathcal{H}^{\uparrow}))$ is the well-known Turán function $\mathrm{ex}(n, \mathcal{H})$.

Fix a graph B with vertex set $[m] := \{1, \dots, m\}$. For pairwise disjoint sets $V_1, \dots, V_m$ (some of which may be empty), let the *blow-up* $B(V_1, \dots, V_m)$ be obtained from the empty graph on $V_1 \cup \dots \cup V_m$ by adding for every edge $\{i, j\} \in E(B)$ the complete bipartite graph with parts V_i and V_j. Note that no part V_i spans an edge. Let $B()$ be the family of all possible blow-ups of B. It consists of graphs that can be obtained from B by a sequence of vertex duplications and vertex deletions.

Suppose that $B() \subseteq \mathcal{G}$. Trivially, we get the lower bound $\Lambda(n, \mathcal{G}) \geqslant \Lambda(n, B())$ valid for every integer $n \geqslant \kappa$, where, in accordance with our general notation, $\Lambda(n, B())$ is the maximum value of Λ over all blow-ups of B of order n. For a vector $\boldsymbol{a} = (a_1, \dots, a_m)$ in the *m-simplex*

$$\mathbb{S}_m := \{\boldsymbol{x} \in \mathbb{R}^m : \boldsymbol{x} \geqslant 0,\ x_1 + \dots + x_m = 1\},$$

define

$$\lambda(B(\boldsymbol{a})) := \lim_{n\to\infty} \lambda(B(V_{1,n}, \dots, V_{m,n})), \tag{4}$$

where $|V_{i,n}| = a_i n + O(1)$ for each $i \in [m]$. In other words, we look at the limiting value of the function λ evaluated on a blow-up G of B of order $n \to \infty$ where the i-th part occupies $(a_i + O(1/n))$-fraction of all vertices. It is easy to see that the limit in (4) exists (and does not depend on the choice of the sizes $|V_{i,n}|$). In fact, $\lambda(B(\boldsymbol{a}))$ is a polynomial in $\boldsymbol{a}$ of degree at most κ, so the rate of convergence in (4) is $O(1/n)$. An easy argument based on the compactness of $\mathbb{S}_m$ and the continuity of $\lambda(B(\boldsymbol{a}))$ as a function of $\boldsymbol{a} \in \mathbb{S}_m$ shows that

$$\lambda(B()) = \max\{\lambda(B(\boldsymbol{x})) : \boldsymbol{x} \in \mathbb{S}_m\}. \tag{5}$$

By our assumption, $B() \subseteq \mathcal{G}$ so $\lambda(B())$ is a lower bound on $\lambda(\mathcal{G})$.

Here is an illustration. In Example 1.1, if $\mathcal{H} = \mathcal{H}^{\uparrow} = \{K_t\}$ consists of the t-clique, then a good choice is to take $B = K_{t-1}$. Then $B()$ is a subset of $\mathcal{G} = \mathrm{Forb}(\mathcal{F})$ and $\lambda(B(\boldsymbol{a})) = 2\prod_{1\leqslant i<j\leqslant t-1} a_i a_j$. This is clearly maximised if all entries a_i are equal to each other, giving $\lambda(B()) = (t-2)/(t-1)$, which is a lower bound on the Turán density of K_t. The classical results of Mantel [28] (for $t = 3$) and Turán [43] (for any t) imply that this is an equality. More strongly, they showed that $\Lambda(n, \mathcal{G}) = \Lambda(n, B())$ for all n,

while an easy optimisation shows that $\Lambda(n, B())$ is attained by the (unique) blow-up of B of order n with parts as equal as possible.

In general, there is no hope for a theory that allows to determine $\lambda(\mathcal{G})$ for every λ and $\mathcal{G}$. Namely, as it was shown by Hatami and Norin [22], the question if $\lambda(\mathcal{G}) \leqslant c$ on input $(\mathcal{G}, \Lambda, c)$ is undecidable in general (even if $\mathcal{G}$ consists of all graphs). However, one can determine $\lambda(\mathcal{G})$ for various concrete examples of interest. Many of these proofs utilise the powerful flag algebra approach of Razborov [38,39], where a computer can be used to generate a *certificate* $\mathcal{C}$ proving $\lambda(\mathcal{G}) \leqslant c$.

We will discuss certificates in more detail in Section 3. For the time being, let us just remark that the desired inequality $\lambda(\mathcal{G}) \leqslant c$ follows by symbolically representing $c - \lambda(G)$, for $G \in \mathcal{G}_n^0$ as a sum of squares within error term of order $O(1/n)$ as $n \to \infty$. One illustrative example of such a sum is $s(G) := \binom{n}{3}^{-1} \sum_{x \in V(G)} (d_G(x) - d_{\overline{G}}(x))^2$, where $d_H(x)$ is the degree of x in a graph H. Clearly, $s(G)$ is non-negative while it is routine to see that $s(G) = 6p(K_3, G) + 6p(\overline{K_3}, G) - 2p(P_3, G) - 2p(\overline{P_3}, G) + O(1/n)$, where P_i is the path with i vertices. This gives an asymptotic inequality that always holds between 3-vertex subgraph densities. One advantage of the flag algebra approach is that it allows us to generate and manipulate such equalities automatically; here finding optimal coefficients amounts to solving a certain semi-definite programme (which is independent of n). We refer the reader to Section 3 for details and formal definitions.

Thus, if a flag algebra calculation proves $\lambda(\mathcal{G}) \leqslant c$ while we can find an order-m graph B with $B() \subseteq \mathcal{G}$ and $\boldsymbol{a} \in \mathbb{S}_m$ with $\lambda(B(\boldsymbol{a})) = c$, then we know $\lambda(n, \mathcal{G})$ within an error term of $O(1/n)$:

$$c - O(1/n) \leqslant \lambda(n, B(\boldsymbol{a})) \leqslant \lambda(n, \mathcal{G}) \leqslant c + O(1/n).$$

In addition to determining $\lambda(\mathcal{G})$, it is often desirable to obtain information on the structure of all large admissible graphs G for which the value $\lambda(G)$ is close to the maximal possible. In particular, we look for sufficient conditions establishing that every such G is necessarily close to a blow-up of B, in which case we regard the problem as stable. This paper will consider a few non-equivalent versions of stability, with the corresponding definitions following shortly. The stability is a very useful property in extremal graph theory as it is often indispensable in determining the exact value of $\lambda(n, \mathcal{G})$ as well as the set of all extremal graphs of large order n. Besides being an important property on its own, stability also helps in solving the randomised or counting versions of extremal results.

We will use only one notion of distance on graphs. Namely, the *(edit) distance* between two graphs G and H of the same order n is

$$\Delta_{\mathrm{edit}}(G, H) := \min_{\sigma} \Big| E(G) \bigtriangleup \big\{ \{\sigma(u), \sigma(v)\} : uv \in E(H) \big\} \Big|,$$

where the minimum is taken over all bijections $\sigma : V(H) \to V(G)$. In other words, $\Delta_{\mathrm{edit}}(G, H)$ is the minimum number of adjacencies that one needs to change in G in order to obtain a graph isomorphic to H. We define the *normalised (edit) distance* to be $\delta_{\rm edit}(G,H) := \Delta_{\rm edit}(G,H)/\binom{n}{2}$. For a family $\mathcal{H}$ of graphs we define $\Delta_{\rm edit}(G,\mathcal{H}) := \min\{\Delta_{\rm edit}(G,H) : H \in \mathcal{H}_n^0\}$ and $\delta_{\rm edit}(G,\mathcal{H}) := \min\{\delta_{\rm edit}(G,H) : H \in \mathcal{H}_n^0\}$.

We say that our problem (2) is *robustly B-stable* (resp. *perfectly B-stable*) if there is $C > 0$ such that for every graph $G \in \mathcal{G}$ of order $n \geqslant C$ we have

$$\delta_{\rm edit}(G, B()) \leqslant C \max\left(1/n, \lambda(n,\mathcal{G}) - \lambda(G)\right),$$

(resp. $\delta_{\rm edit}(G,B()) \leqslant C(\lambda(n,\mathcal{G})-\lambda(G))$). For comparison, the *classical B-stability* states that for every $\varepsilon > 0$ there is $\delta > 0$ such that $\delta_{\rm edit}(G,B()) \leqslant \varepsilon$ for every $G \in \mathcal{G}$ with $n \geqslant 1/\delta$ vertices and $\lambda(G) \geqslant \lambda(\mathcal{G}) - \delta$. Clearly, the perfect stability implies the robust stability which in turn implies the classical stability. (Our Theorem 1.11 will in particular show that these notions of stability are not equivalent, already for such a natural problem as the Turán function.) Also, if the problem is perfectly stable, then for all $n \geqslant C$ we have $\Lambda(n,\mathcal{G}) = \Lambda(n,B())$ and every extremal graph is a blow-up of B (which one may call an *exact result*).

For example, for the Turán function $\mathrm{ex}(n,\mathcal{F})$, the classical stability was established independently by Erdős [10] and Simonovits [41]. The perfect (and thus also robust) stability in the case when $\mathcal{F}$ consists of a clique K_t follows from results of Füredi [15] (who considered the distance to being $(t-1)$-partite instead of complete $(t-1)$-partite as we do in this paper). Very recently, Roberts and Scott [40] improved on Füredi's result by extending it to all colour critical graphs and giving a sharper bound on the distance.

As far as we know, the above results in [15,40] and some recent work of Norin and Yepremyan [33,34] (who considered the Turán problem for hypergraphs) are the only known examples where perfect stability was established for a non-trivial problem. Furthermore, almost all proofs where the classical stability and the exact result were derived from a flag algebra computation were rather ad-hoc and tailored to a particular problem.

The purpose of this paper is to present some general sufficient conditions that imply some version of stability. This allows us to give a unified proof of many previous stability and exactness results. Also, we can establish perfect stability (a new result) for a number of problems.

More specifically, Theorem 4.1 gives a sufficient condition for robust stability and Theorems 5.8, 5.13, and 7.1 give various sufficient conditions for perfect stability. Furthermore, all assumptions of Theorem 7.1 are stated in a way that allows automatic verification by a computer. We also present an openly available computer code (written in `sage` by adopting the `flagmatic` package of Emil Vaughan [44]) that allows us to both generate and verify certificates for general problems based on Theorem 7.1. In all the cases where we could verify assumptions of Theorem 7.1 and derive perfect stability, the procedure was essentially mechanical and the final high-level code is very short (having 6–10 lines, each invoking some function).

Even if one knows that $\lambda(\mathcal{G}) = \lambda(B())$ for a concrete B, the determination of asymptotically optimal part ratios (namely, finding all $\boldsymbol{a} \in \mathbb{S}_m$ with $\lambda(B()) = \lambda(B(\boldsymbol{a}))$) may still be a non-trivial task that amounts to polynomial maximisation. While the combination of Lagrange multipliers and Gröbner bases provides a general computational framework, in an extremal problem one often has a candidate $\boldsymbol{a} \in \mathbb{S}_m$ and wishes to prove that this is the only vector (up to a symmetry of B) that achieves $\lambda(B())$. Clearly, if a flag algebra certificate $\mathcal{C}$ proves that $\lambda(\mathcal{G}) \leqslant \lambda(B(\boldsymbol{a}))$ then this automatically implies that $\boldsymbol{a}$ is a maximiser and it is possible that the information in $\mathcal{C}$ is enough to imply the uniqueness of $\boldsymbol{a}$. We present such a sufficient condition on $\mathcal{C}$ in Lemma 6.2 which can be automatically verified by a computer and seems to work very well in practice.

The exact statements of the above sufficient conditions rely on understanding flag algebra certificates, so we postpone them until later. Here, let us list the extremal problems for which our method gives perfect stability. In almost all cases, Theorem 7.1 and Lemma 6.2 apply directly, immediately giving perfect stability and implying the uniqueness of asymptotically optimal part ratios.

However, there are a few natural problems where the assumptions of Theorem 7.1 are not satisfied. As the test case that our method can still give perfect stability, we chose the inducibility function for the paw graph, see Theorem 1.10. The asymptotic value of this function was determined by Hirst [24] but the classical stability and the exact result were not known. By utilising our other results (Theorem 5.8) we derive perfect stability for the inducibility problem for the paw graph. Since the proof is rather long and was meant mainly as an illustration of the flexibility of our method, we decided to include only one such example in this paper.

1.1. Examples of extremal problems for which we can prove perfect stability

1.1.1. Minimising the number of independent sets in triangle-free graphs

Erdős [9] asked for the value of $f(n, k, l)$, the minimum number of independent sets of size k in a K_l-free graph of order n.

Consider first the case $l = 3$, when we forbid a triangle. Goodman [17] determined $f(2n, 3, 3)$; his bounds also give the asymptotic value of $f(2n + 1, 3, 3)$. Lorden [26] showed that, for $n \geqslant 12$, the value of $f(n, 3, 3)$ is attained by taking $K_{\lfloor n/2\rfloor,\lceil n/2\rceil}$ and removing any (possibly empty) matching from it. Some partial results were obtained by Nikoforov [32,31]. The problem for $k \geqslant 4$ remained open until recently when the classical stability and the exact result were established independently by Das, Huang, Ma, Naves and Sudakov [7] (for $k \in \{4, 5\}$) and Pikhurko and Vaughan [36] (for $k \in \{4, 5, 6, 7\}$) when n is sufficiently large: if $k \in \{4, 5\}$, then all extremal graphs are blow-ups of C_5 and if $k \in \{6, 7\}$, then all extremal graphs are blow-ups of the *Clebsch graph* L. The *Clebsch graph* L has, as vertices, binary 5-sequences of even *weight* (that is, the number of ones), with two vertices being adjacent if the point-wise sum modulo 2 of the corresponding sequences has weight 4. It easily follows from this description that the graph L has 16 vertices and is triangle-free and 5-regular.

This question of Erdős is a special case of our general problem. It turns out that our computer codes can prove the perfect stability in the following cases. (Note that the $f(n, 3, 3)$-problem is not perfectly stable by the above mentioned result of Lorden [26].)

Theorem 1.2. *Let*

- $k \in \{4, 5\}$ *and* $B = C_5$, *or*
- $k \in \{6, 7\}$ *and* $B = L$.

Let $\mathcal{F} = \{K_3\}$ *(thus* $\mathcal{G}$ *consists of all triangle-free graphs) and let* $\gamma(H)$ *be 0 except* $\gamma(\overline{K}_k) = -1$. *Then the corresponding problem* $\Lambda(n, \mathcal{G})$ *(that is, Erdős' problem of determining* $f(n, k, 3)$*) is perfectly* B*-stable. Furthermore, for each* $k \in \{4, \dots, 7\}$ *the unique probability vector* $\boldsymbol{a} \in \mathbb{S}_{v(B)}$ *that maximises* $\lambda(B(\boldsymbol{a}))$ *is the uniform vector.*

If $l \geqslant 4$, then the asymptotic value of $f(n, k, l)$ is known only for $k = 3$ and $l \leqslant 7$, see [7,36]. The papers [7,36] also showed that in each of these cases the problem is classically B-stable, where $B = K_{l-1}$, and the value of $f(n, k, l)$ is attained by a blow-up of B for all large n. However, the problem is not perfectly B-stable since it is possible to remove a few edges from the blow-up of B (e.g. a matching between two parts) so that the number of copies of $\overline{K}_3$ does not change.

The above results and our Theorem 5.13 imply that, in fact, the $f(n, 3, l)$-problem is not robustly stable for $l \in \{3, 4, 5, 6, 7\}$. Alternatively, the same conclusion can be derived directly by taking the optimal blow-up $K_{l-1}(V_1, \dots, V_{l-1})$, fixing some sets $X \subseteq V_1$ and $Y \subseteq V_2$ each of size εn, where ε is a small constant, and then flipping all pairs between X and Y. (Such transformation is done in the proof of Theorem 5.13 and is carefully analysed there.)

1.1.2. Maximising the number of pentagons in triangle-free graphs

Erdős [11] asked if $c(5m) \leqslant m^5$ for every natural number m, where $c(n)$ is the maximum number of 5-cycles that a triangle-free graph of order n can have. Note that this bound is sharp for every $m \in \mathbb{N}$ which is witnessed by the balanced blow-up of C_5.

Some partial progress on this problem was obtained by Győri [19] who proved $c(n) \leqslant 1.03\,(n/5)^5$ and Füredi (unpublished) who proved $c(n) \leqslant 1.01\,(n/5)^5$. Recently, Grzesik [18] and independently Hatami et al. [21] proved that, as $n \to \infty$, there can be at most $(1/5^5 + o(1))n^5$ copies of C_5. Furthermore, Hatami et al. [21] proved the exact result for all large n (and the classical stability can also be derived from their method). In fact, the validity of Erdős' conjecture follows from the asymptotic result by a simple blow-up trick (see [21, Corollary 3.3]). Interestingly, if $n = 8$, there is another extremal example which is not a blow-up of C_5 that was discovered by Michael [29]. Very recently, the value of $c(n)$ and the description of all extremal graphs for every n was obtained by Lidický and Pfender [25].

This problem fits into our general framework and we can prove (again in a completely automated way) that it is perfectly stable.

Theorem 1.3. *Let $\mathcal{F} = \{K_3\}$, $\kappa = 5$, and $\gamma(H)$ be zero, except $\gamma(C_5) = 1$. Let $B = C_5$. Then the corresponding problem is perfectly B-stable (with the unique maximiser for $\lambda(B(\boldsymbol{a}))$ being the vector $\boldsymbol{a} \in \mathbb{S}_5$ with each entry equal to $1/5$).*

1.1.3. Inducibility

Given a graph F, the *inducibility problem* for F asks for the maximal possible (induced) density of the graph F among all graphs of order n. In our general notation, it can be expressed as follows. Let $\kappa = v(F)$, $\mathcal{F} = \emptyset$ and let γ take the value 0 on every graph with κ vertices except F, where it takes value 1. Thus we are interested in $i(n, F) := \Lambda(n, \mathcal{G})$. We call $i(F) := \lambda(\mathcal{G})$ the *inducibility* of F. Equivalently,

$$i(F) := \lim_{n\to\infty} \max\{p(F, G) : v(G) = n\}.$$

Observe that the inducibility of each graph is equal to the inducibility of its complement. The inducibility problem has drawn a considerable amount of interest, see for example [2–4,6,13,14,16,20,23,24,42].

Before we look at concrete examples, let us mention the following general result of Brown and Sidorenko [6, Proposition 1]: if F is *complete partite* (i.e. a blow-up of some clique), then for every $n \in \mathbb{N}$ at least one $i(n, F)$-extremal graph is complete partite. The proof in [6] uses the symmetrisation method and it is not clear how to extract a stability-type result from it. Also, Even-Zohar and Linial [13, Table 2] systematically looked at the inducibility of 5-vertex graphs but without trying to convert the numerical bounds coming from flag algebra calculations into computer-verifiable mathematical proofs.

1.1.4. Inducibility of the cycle on four vertices

The inducibility of the 4-cycle, denoted by C_4, follows from the above mentioned result of Brown and Sidorenko [6]. Previously, Pippenger and Golumbic [16] determined $i(n, K_{k,l})$ for all k, l with $|k-l| \leqslant 1$, observing that the complete balanced bipartite graph is an extremal graph. Here we prove perfect stability for this problem (by invoking our computer code).

Theorem 1.4. *Let $\mathcal{F} = \emptyset$, $\kappa = 4$, and $\gamma(H)$ be zero, except $\gamma(C_4) = 1$. Let $B = K_2$ be a single edge. Then the corresponding problem is perfectly B-stable (with $\lambda(\mathcal{G}) = i(C_4) = 3/8$ and the unique maximiser for $\lambda(B(\boldsymbol{a}))$ being the vector $\boldsymbol{a} = (1/2, 1/2)$).*

1.1.5. Inducibility of K_4 minus an edge

Let K_4^- be the graph obtained by removing an edge from the complete graph on four vertices. The inducibility problem for K_4^- was considered by Hirst [24], who determined $i(K_4^-)$ using the flag algebra method. Our new result is that this problem is perfectly stable (and, in particular, that $i(n, K_4^-)$ is attained by a blow-up of the complete graph on five vertices K_5, for all large n).

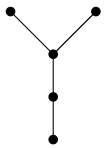

Fig. 1. The graph Y.

Theorem 1.5. *Let $\mathcal{F} = \emptyset$, $\kappa = 4$, and $\gamma(H)$ be zero, except $\gamma(K_4^-) = 1$. Also let $B = K_5$. Then the corresponding problem is perfectly B-stable (with $\lambda(\mathcal{G}) = i(K_4^-) = 72/125$ and the unique maximiser for $\lambda(B(\boldsymbol{a}))$ being the vector $\boldsymbol{a} \in \mathbb{S}_5$ with each entry equal to $1/5$).*

1.1.6. Inducibility of $K_{3,2}$

The function $i(n, K_{3,2})$ was calculated by Golumbic and Pippenger in [16], where the complete balanced bipartite graph is an extremal graph. We show that the problem is perfectly stable.

Theorem 1.6. *Let $\mathcal{F} = \emptyset$, $\kappa = 5$, and $\gamma(H)$ be zero, except $\gamma(K_{3,2}) = 1$. Also let $B = K_2$. Then the corresponding problem is perfectly B-stable (with $\lambda(\mathcal{G}) = i(K_{3,2}) = 5/8$ and the unique maximiser for $\lambda(B(\boldsymbol{a}))$ being the vector $\boldsymbol{a} = (1/2, 1/2)$).*

1.1.7. Inducibility of $K_{2,2,1}$

The function $i(n, K_{2,2,1})$ can be derived from the results of Brown and Sidorenko [6]. Here we prove perfect stability for the corresponding problem.

Theorem 1.7. *Let $\mathcal{F} = \emptyset$, $\kappa = 5$, and $\gamma(H)$ be zero, except $\gamma(K_{2,2,1}) = 1$. Also let $B = K_3$. Then the corresponding problem is perfectly B-stable (with $\lambda(\mathcal{G}) = i(K_{2,2,1}) = 10/27$ and the unique maximiser for $\lambda(B(\boldsymbol{a}))$ being the vector $\boldsymbol{a} = (1/3, 1/3, 1/3)$).*

1.1.8. Inducibility of $P_3 \cup K_2$

We consider the inducibility problem for the disjoint union of a path on 3 vertices and an edge, which we denote by $P_3 \cup K_2$. We prove the following.

Theorem 1.8. *Let $\mathcal{F} = \emptyset$, $\kappa = 5$, and $\gamma(H)$ be zero, except $\gamma(P_3 \cup K_2) = 1$. Also let $B = K_3 \cup K_3$, that is, the disjoint union of two triangles. Then the corresponding problem is perfectly B-stable (with $\lambda(\mathcal{G}) = i(P_3 \cup K_2) = 5/18$ and the unique maximiser for $\lambda(B(\boldsymbol{a}))$ being the vector $\boldsymbol{a} \in \mathbb{S}_6$ with each entry equal to $1/6$).*

1.1.9. Inducibility of the "Y" graph

We consider the inducibility problem for the graph depicted in Fig. 1, which we denote by Y. We prove the following.

Theorem 1.9. *Let $\mathcal{F} = \emptyset$, $\kappa = 5$, and $\gamma(H)$ be zero, except $\gamma(\mathrm{Y}) = 1$. Also let $B = C_5$. Then the corresponding problem is perfectly B-stable (with $\lambda(\mathcal{G}) = i(Y) = 24/125$ and the unique maximiser for $\lambda(B(\boldsymbol{a}))$ being the vector $\boldsymbol{a} \in \mathbb{S}_5$ with each entry equal to $1/5$).*

1.1.10. Inducibility for the paw graph

Let us denote by F_{paw} the graph obtained by adding to a triangle a pendant edge. Using flag algebras, Hirst [24] determined the value of $i(F_{\mathrm{paw}})$. Here, it is more convenient to work with the (equivalent) complementary problem. Thus we consider the inducibility problem of the disjoint union of a path on three vertices and a single node, that we denote by $P_3 \cup K_1$. Unfortunately, the perfect stability does not follow directly from Theorem 7.1. However, it can be proved using our methods combined with additional work, see Section 8.1 for the proof.

Theorem 1.10. *Let $\mathcal{F} = \emptyset$, $\kappa = 4$, and $\gamma(H)$ be zero, except $\gamma(P_3 \cup K_1) = 1$. Also let $B = K_2 \cup K_2$, that is, the disjoint union of two edges. Then the corresponding problem is perfectly B-stable (with $\lambda(\mathcal{G}) = i(P_3 \cup K_1) = i(F_{\mathrm{paw}}) = 3/8$ and the unique maximiser for $\lambda(B(\boldsymbol{a}))$ being the vector $\boldsymbol{a} \in \mathbb{S}_4$ with each entry equal to $1/4$).*

1.1.11. Turán problem

Recall that the Turán problem was introduced in Example 1.1. Given a family of graphs $\mathcal{H}$ we consider $\mathcal{H}^{\uparrow}$, the collection of graphs obtained by adding missing edges to the graphs in $\mathcal{H}$. While our computer code can automatically prove the perfect stability when $\mathcal{H} = \{K_t\}$ with $t \leqslant 7$, this is superceded by the following result whose proof does not require a computer. For integer $q \geqslant 1$ denote by K_m^q the balanced blow up of K_m on qm vertices (i.e. $K_m(V_1, \dots, V_m)$ with $|V_1| = \dots = |V_m| = q$).

Theorem 1.11. *Let $\mathcal{H}$ be a family of graphs and let*

$$m := \min\{\chi(H) : H \in \mathcal{H}\} - 1 \geqslant 2,$$

where by $\chi(H)$ we denote the chromatic number of the graph H, that is, the minimum number of colours needed in a colouring of the vertex set with no adjacent vertices of the same colour. Then the following hold.

1. *The Turán problem $\mathrm{ex}(n, \mathcal{H})$ is perfectly K_m-stable if and only if there is an integer q such that K_m^q plus one edge is not $\mathcal{H}^{\uparrow}$-free.*
2. *Assuming in addition that $\mathcal{H}$ is finite, we have the following. The Turán problem $\mathrm{ex}(n, \mathcal{H})$ is robustly K_m-stable if and only if there is an integer q such that K_m^q plus some forest in one of the parts of K_m^q is not $\mathcal{H}^{\uparrow}$-free.*

As we learned later, the non-trivial implication in Part 1 of the above theorem is apparently a folklore result. Since it follows from [40, Lemma 2.3], we omit its proof. The second part of Theorem 1.11 is proved in Section 9 of this paper.

2. Notation and preliminaries

Some of the definitions and proofs of this paper will be more natural when stated in a more analytic way. For example, the definition of $\lambda(B(\boldsymbol{a}))$ in (4) would not require a limit if instead we were working with vertex-weighted graphs. Such objects are quite common in extremal graph theory nowadays (appearing, for example, in the definition of the *Lagrangian* of a graph that goes back to Motzkin and Straus [30]) and, of course, they are generalised in a powerful and far-reaching way by graphons (see, for example, the excellent book by Lovász [27]). However, we believe that by staying within the universe of simple unweighted graphs, we make the paper and its ideas better accessible to a wider audience.

As usual, for each positive integer n, we denote by $[n]$ the set $\{1, ..., n\}$. Let $\mathbf{E}(X)$ denote the expected value of a random variable X. We will often abbreviate a pair $\{i, j\}$ as ij. For a finite set A and a positive integer k we denote by $\binom{A}{k}$ the set of all k-subsets of A.

Recall that K_m denotes the complete graph of order m and $\overline{G} := (V(G), \binom{V(G)}{2} \setminus E(G))$ denotes the complement of a graph G. Let $K_{m,n}$ be the complete bipartite graph with part sizes m and n. For a vertex $x \in V(G)$, let $\Gamma_G(x) = \{y \in V(G) : \{x, y\} \in E(G)\}$ denote the *neighbourhood* of x in G. We write $G \cong H$ if G and H are isomorphic graphs.

We call a graph B λ*-minimal* if $\lambda(B())$ strictly decreases when we remove any vertex of B. By (5), B is λ-minimal if and only if no point on the boundary of $\mathbb{S}_m$ achieves the maximum.

A graph B is *twin-free* if it contains no two vertices x and y with identical neighbourhoods (i.e., for all distinct $x, y \in V(B)$ we have $\Gamma_B(x) \neq \Gamma_B(y)$).

Recall that the *(edit) distance* $\Delta_{\mathrm{edit}}(G, H)$ between two graphs G and H of the same order n is the minimum of adjacencies one has to edit in G to make it isomorphic to H. Also, the *(edit) distance* $\Delta_{\mathrm{edit}}(G, \mathcal{H})$ from a graph G to a family $\mathcal{H}$ of graphs is the minimum of $\Delta_{\mathrm{edit}}(G, H)$ over all $H \in \mathcal{H}$ that have the same order as G; this is the minimum number of adjacency edits needed to transform G into a graph in $\mathcal{H}$. The respective *normalised distances* are $\delta_{\mathrm{edit}}(G, H) := \Delta_{\mathrm{edit}}(G, H)/\binom{n}{2}$ and $\delta_{\mathrm{edit}}(G, \mathcal{H}) := \Delta_{\mathrm{edit}}(G, \mathcal{H})/\binom{n}{2}$.

Throughout this paper we will work under the following assumptions which are collected together for future reference.

Assumption 2.1. Let κ, m be positive integers and $\mathcal{F}$ a family of graphs.

1. Set $\mathcal{G} = \mathrm{Forb}(\mathcal{F})$.
2. Let $\gamma : \mathcal{G}_\kappa^0 \to \mathbb{R}$ be a function and define Λ and λ as in (1).
3. Let B be a graph on $[m]$ such that $B() \subseteq \mathcal{G}$.

The next lemma provides some basic information on the behaviour of the sequence $(\lambda(n, \mathcal{G}))_{n=\kappa}^{\infty}$.

Lemma 2.2. *Let $\mathcal{G}$ be a graph property closed under taking induced subgraphs. Then for $\kappa \leqslant q \leqslant n$ with $q \to \infty$ we have*

$$0 \leqslant \lambda(q, \mathcal{G}) - \lambda(n, \mathcal{G}) \leqslant o_q(1).$$

Furthermore, if $\mathcal{G}$ is closed under taking blow-ups, then the error term is $O(1/q)$.

Proof. Take an optimal graph G for $\lambda(n, \mathcal{G})$. Let X be a random q-subset of $V(G)$ and $G' := G[X]$. Then $G' \in \mathcal{G}$. Thus $\lambda(q, \mathcal{G}) \geqslant \mathbf{E}(\lambda(G'))$. Clearly, if we take a uniform $X \in \binom{V(G)}{q}$ and then a uniform $Y \in \binom{X}{\kappa}$, then Y is uniformly distributed among all κ-subsets of $V(G)$. Thus $\mathbf{E}(\lambda(G'))$ equals $\lambda(G) = \lambda(n, \mathcal{G})$, giving $\lambda(q, \mathcal{G}) \geqslant \lambda(n, \mathcal{G})$. Thus $\lambda(q, \mathcal{G})$ is non-increasing in q and tends to a limit, implying the other desired inequality $\lambda(n, \mathcal{G}) \geqslant \lambda(q, \mathcal{G}) + o_q(1)$.

Finally, assume that $\mathcal{G}$ is also closed under taking blow-ups. To show $\lambda(q, \mathcal{G}) \leqslant \lambda(n, \mathcal{G}) + O(1/n)$, take an optimal graph G for $\lambda(q, \mathcal{G})$ on $[q]$. Consider a random map $\phi : [n] \to [q]$ with all q^n choices being equally likely and let G' be the graph on $[n]$ with $E(G') = \phi^{-1}(E(G))$. Take any κ-subset $X \subseteq [n]$. With probability $1 - O(1/q)$, the map ϕ is injective on X. If we condition on this, then $\phi(X) \in \binom{[q]}{\kappa}$ is uniform and the average of $\gamma(G'[X]) = \gamma(G[\phi(X)])$ is $\lambda(G)$. Thus

$$\lambda(n, \mathcal{G}) \geqslant \mathbf{E}_\phi(\lambda(G')) = (1 - O(1/q))\lambda(G) - O(1/q),$$

giving the desired. □

3. Flag algebra method

As we have already mentioned in the introduction of this paper, the flag algebra method is a powerful technique pioneered by Razborov [38,39]. In this section, we define what a certificate is and how it implies an upper bound on $\lambda(\mathcal{G})$. Recall that we always work under Assumption 2.1.

3.1. Types and flags

A *type* is a pair of the form (H, ϕ), where H is an admissible graph and $\phi : [v] \to V(H)$ is a bijection with $v = v(H)$. Given a type $\tau = (H, \phi)$ as above, a *τ-flag* is a pair of the form (G, ψ), where G is an admissible graph and $\psi : [v] \to V(G)$ is an injection such that $\psi \circ \phi^{-1} : V(H) \to V(G)$ is an *embedding* (that is, an injection that preserves both edges and non-edges). Informally, a τ-flag (G, ψ) is a partially labelled graph such that the labelled vertices induce τ. The *order* $v((G, \psi))$ of the flag is $v(G)$, the number of vertices in it.

For two τ-flags (G_1, ψ_1) and (G_2, ψ_2) with respectively $n_1 \leqslant n_2$ vertices, let the *sub-flag count* $P((G_1, \psi_1), (G_2, \psi_2))$ be the number of n_1-subsets X of $V(G_2)$ such that

$X \supseteq \psi_2([v])$ (i.e. X contains all labelled vertices) and the τ-flags (G_1, ψ_1) and $(G_2[X], \psi_2)$ are *isomorphic*, meaning that there is a graph isomorphism that preserves the labels. Also, define the *(flag) density* as

$$p((G_1, \psi_1), (G_2, \psi_2)) := \frac{P((G_1, \psi_1), (G_2, \psi_2))}{\binom{n_2 - v}{n_1 - v}}. \tag{6}$$

The latter quantity can be viewed as the probability that a random n_1-subset X of $V(G_2)$ with $X \supseteq \phi_2([v])$ induces a copy of the flag (G_1, ψ_1) in (G_2, ψ_2).

We will also need a variation of the above notion. Let $F_1 = (G_1, \psi_1)$, $F_2 = (G_2, \psi_2)$ and (G, ψ) be three τ-flags with respectively n_1, n_2 and n vertices. We define the *joint sub-flag count* $P(F_1, F_2, (G, \phi))$ to be the number of pairs (X, Y) such that X, Y are subsets of $V(G)$ with n_1 and n_2 elements respectively, $X \cap Y = \psi([v])$ and the τ-flags $(G[X], \psi)$ and $(G[Y], \psi)$ are isomorphic to F_1 and F_2, respectively.

The type with no vertices will be denoted by 0. Thus 0-flags are just unlabelled graphs. In this case, the 0-flag density as defined by (6) coincides with the notion of subgraph density from the Introduction.

3.2. Certificates

Definition 3.1. A *(flag algebra) certificate* is a triple

$$\mathcal{C} = (N, \mathcal{T}, (Q^{\tau})_{\tau \in \mathcal{T}}), \tag{7}$$

where

- $N \geqslant \kappa$ is an integer;
- $\mathcal{T} = (\tau_1, \ldots, \tau_t)$ is an ordered list of some types such that $N - v(\tau_i)$ is a positive even integer for each $i \in [t]$;
- Q^{τ_i} is an arbitrary positive semi-definite $g_i \times g_i$-matrix for $i \in [t]$, where we fix some enumeration $(F_1^{\tau_i}, \ldots, F_{g_i}^{\tau_i})$ of all τ_i-flags with exactly $(N + v(\tau_i))/2$ vertices, up to isomorphism of flags (and thus g_i is the number of these flags).

Note that the third component of the certificate $\mathcal{C}$ consists of exactly t matrices, one per each of the types $\tau_1, \ldots, \tau_t$; one can view the rows/columns of Q^{τ_i} as indexed by the τ_i-flags of order $(N + v(\tau_i))/2$.

To describe the upper bound of $\lambda(\mathcal{G})$ that a certificate $\mathcal{C}$ witnesses, we need to introduce several quantities.

Let $G_1, \ldots, G_g$ be the enumeration in some fixed order of all (up to an isomorphism) admissible N-vertex graphs. Thus $\mathcal{G}_N^0 = \{G_1, \ldots, G_g\}$ with no two listed graphs being isomorphic. For each $q \in [g]$ (that is, for each G_q), we define real numbers

$$a_q := \sum_{i=1}^{t}\sum_{h=1}^{g_i}\sum_{j=1}^{g_i} c^{\tau_i}_{h,j,q} Q^{\tau_i}_{h,j}, \quad \text{where} \quad c^{\tau_i}_{j,h,q} := \sum_{\phi} P(F^{\tau_i}_j, F^{\tau_i}_h, (G_q,\phi)) \tag{8}$$

and the sum in the definition of $c^{\tau_i}_{h,j,q}$ is taken over all injective maps $\phi : [v(\tau_i)] \to V(G_q)$ that induce a copy of the flag τ_i in G. Also, let

$$b_q := \lambda(G_q) = \sum_{H\in\mathcal{G}^0_\kappa} \gamma(H)p(H,G_q),$$

and

$$u_\lambda(\mathcal{C}) := \max\{a_q + b_q : q \in [g]\}.$$

A graph $G_q \in \mathcal{G}^0_N$ is called $(\mathcal{C},\lambda)$-*sharp* (or $\mathcal{C}$-*sharp*, or just *sharp*) if $a_q + b_q = u_\lambda(\mathcal{C})$. The following lemma motivates the above definitions.

Lemma 3.2. *Under the above notation, for every admissible graph G of order $n \geq N$ we have that*

$$u_\lambda(\mathcal{C}) - \lambda(G) + O(1/n) = \sum_{q=1}^{g} (u_\lambda(\mathcal{C}) - b_q)p(G_q,G) + O(1/n) \tag{9}$$

$$\geqslant \sum_{q=1}^{g} a_q p(G_q,G) + O(1/n) \ \geqslant\ 0. \tag{10}$$

In particular, we have that $\lambda(n,\mathcal{G}) \leqslant u_\lambda(\mathcal{C}) + O(1/n)$ and $\lambda(\mathcal{G}) \leqslant u_\lambda(\mathcal{C})$.

Proof. Let G be an arbitrary admissible graph of order $n \geq N$. We have

$$\lambda(G) = \sum_{H\in\mathcal{G}^0_\kappa} \gamma(H)p(H,G) = \sum_{H\in\mathcal{G}^0_\kappa} \gamma(H) \sum_{q=1}^{g} p(H,G_q)p(G_q,G) = \sum_{q=1}^{g} b_q p(G_q,G), \tag{11}$$

proving (9).

Next, we define a non-negative quantity a in the following way. Initially, we set $a = 0$. For each non-negative integer v such that $N - v$ is a positive even integer we work as follows. We enumerate all $n(n-1)\dots(n-v+1)$ injections $\psi : [v] \to V(G)$. If the induced type $(G[\psi([v])],\psi)$ is equal to some $\tau_i \in \mathcal{T}$, then we add the quantity $\boldsymbol{x}Q^{\tau_i}\boldsymbol{x}^T$ to a, where

$$\boldsymbol{x} := \big(P(F^{\tau_i}_1,(G,\psi)), \dots, P(F^{\tau_i}_{g_i},(G,\psi))\big). \tag{12}$$

Since each matrix Q^{τ_i} is positive semi-definite, we have that each $\boldsymbol{x}Q^{\tau_i}\boldsymbol{x} \geqslant 0$ and that the final a is non-negative.

Let $i \in [t]$ and set $v = v(\tau_i)$. Take any $j, h \in [g_i]$. Notice that the sum of the products $P(F_j^{\tau_i}, (G, \psi))\, P(F_h^{\tau_i}, (G, \psi))$ over all injections $\psi : [v] \to V(G)$ such that $(G[\psi([v])], \psi)$ is isomorphic to τ_i, is equal to

$$\sum_{q=1}^{g} c_{j,h,q}^{\tau_i} P(G_q, G) + O(n^{N-1}),$$

where $c_{j,h,q}^{\tau_i}$ is defined in (8), see e.g. [38, Lemma 2.3]. (Informally speaking, we just count in two different ways the number of embeddings of $F_j^{\tau_i}$ and $F_h^{\tau_i}$ into G so that the corresponding labelled vertices coincide; the error term $O(n^{N-1})$ comes from embeddings where some unlabelled vertices happen to collide.) Thus, summing over $i \in [t]$, as well as, over injections ψ and expanding each quadratic form $\boldsymbol{x} Q^{\tau} \boldsymbol{x}^T$, we get the representation

$$0 \leqslant \frac{a}{\binom{n}{N}} = O(1/n) + \sum_{q=1}^{g} a_q p(G_q, G), \tag{13}$$

where a_q's are as in (8). Adding this to (11), we obtain that

$$\lambda(G) + O(1/n) \leqslant \sum_{q=1}^{g} (a_q + b_q) p(G_q, G) \leqslant u_\lambda(\mathcal{C}) \sum_{q=1}^{g} p(G_q, G) = u_\lambda(\mathcal{C}). \tag{14}$$

The inequalities in (10) follow readily by (13) and (14). □

Lemmas 2.2 and 3.2 have the following immediate consequence.

Corollary 3.3. *Suppose that Assumption 2.1 holds. Let $\boldsymbol{a}$ be a vector in $\mathbb{S}_m$ and $\mathcal{C}$ a certificate such that $\lambda(B(\boldsymbol{a})) \geqslant u_\lambda(\mathcal{C})$. Then $\lambda(\mathcal{G}) = \lambda(B(\boldsymbol{a})) = u_\lambda(\mathcal{C})$. Moreover, if $\mathcal{G}$ is closed under taking blow-ups, then $\lambda(n, \mathcal{G}) = u_\lambda(\mathcal{C}) + O(1/n)$.* □

Finally, we close this section with the following lemma.

Lemma 3.4. *Under Assumption 2.1, suppose that $\mathcal{G}$ is closed under taking blow-ups and that a certificate $\mathcal{C}$ and a vector $\boldsymbol{a} = (a_1, \ldots, a_m) \in \mathbb{S}_m$ with no zero entry satisfy $u_\lambda(\mathcal{C}) = \lambda(B(\boldsymbol{a}))$. Fix $i \in [t]$ and set $v := v(\tau_i)$. Also, let n be a large positive integer, $V_1, \ldots, V_m$ a partition of $[n]$ with $|V_i| = a_i n + O(1)$ and $G := B(V_1, \ldots, V_m)$. Finally, let $\psi : [v] \to V(G)$ be an injection such that (G, ψ) is a τ_i-flag. Define $\boldsymbol{x}$ as in (12). Then $\boldsymbol{x} Q^{\tau_i} \boldsymbol{x}^T = O(n^{N-v-1})$.*

Proof. Observe that any modification of the injection ψ such that its values stay in the same parts V_i is an embedding. These new injections give the same vector $\boldsymbol{x}$. There are $\Omega(n^v)$ such injections since each part V_i has size $a_i n + O(1) = \Omega(n)$. Let a be the quantity defined in the proof of Lemma 3.2 when applied to G. Note that a is the sum of non-negative quantities, some of which correspond to the above $\Omega(n^v)$ embeddings of τ_i into G. Thus

$$0 \leqslant \Omega(n^v) \cdot \boldsymbol{x} Q^{\tau_i} \boldsymbol{x}^T \leqslant a. \tag{15}$$

Also, observe that $\lambda(G) = \lambda(B(\boldsymbol{a})) + O(1/n)$. Since $u_\lambda(\mathcal{C}) = \lambda(B(\boldsymbol{a}))$, by (10) and (13) we have that $a = O(n^{N-1})$. Invoking (15), the result follows. □

4. Robust stability from flag algebra proofs

The main result of this section is Theorem 4.1 below that provides a sufficient condition for a problem to be robustly stable. Let H, G be two graphs. We say that a map $f : V(H) \to V(G)$ is a *strong homomorphism* if it preserves both adjacency and non-adjacency. Observe that a strong homomorphism, in contrast to an embedding, does not need to be injective, allowing pairwise non-adjacent vertices to be mapped to the same image. Moreover, let us note that a graph H admits a strong homomorphism in a graph B if and only if H is a blow-up of B.

Theorem 4.1 *(Robust stability). Suppose that in addition to Assumption 2.1 the following holds.*

1. *We have a vector $\boldsymbol{a} \in \mathbb{S}_m$ and a certificate $\mathcal{C} = (N, \mathcal{T}, (Q^\tau)_{\tau\in\mathcal{T}})$ with $u_\lambda(\mathcal{C}) \leqslant \lambda(B(\boldsymbol{a}))$.*
2. *There is a graph τ of order at most $N-2$ satisfying the following.*
   (a) *$\lambda(\mathrm{Forb}(\mathcal{F})) > \lambda(\mathrm{Forb}(\mathcal{F} \cup \{\tau\}))$.*
   (b) *There exists a unique (up to automorphisms of τ and B) strong homomorphism f from τ into B.*
   (c) *For every distinct x_1 and x_2 in $V(B)$ we have $\Gamma_B(x_1) \cap f(V(\tau)) \neq \Gamma_B(x_2) \cap f(V(\tau))$.*
3. *Every $\mathcal{C}$-sharp graph of order N admits a strong homomorphism into B.*

Then the problem is robustly B-stable.

Proof. By Corollary 3.3, we know that $\lambda(\mathcal{G}) = \lambda(B(\boldsymbol{a})) = u_\lambda(\mathcal{C})$. For notational convenience, assume that $V(\tau) = [q]$. Choose large constants in the order $C_1 \ll C$. In particular, we assume that $C > 2/(\lambda(\mathcal{G}) - \lambda(\mathrm{Forb}(\mathcal{F} \cup \{\tau\})))$. Take any $\mathcal{F}$-free graph G of order $n > C$. Note that Condition 2(c) of Theorem 4.1 implies that B is twin-free.

We can assume that $\lambda(G) \geqslant (\lambda(\mathcal{G}) + \lambda(\mathrm{Forb}(\mathcal{F} \cup \{\tau\})))/2$ for otherwise

$$C(\lambda(\mathcal{G}) - \lambda(G)) \geqslant C(\lambda(\mathcal{G}) - \lambda(\mathrm{Forb}(\mathcal{F} \cup \{\tau\})))/2 > 1 \geqslant \delta_{\mathrm{edit}}(G, B()),$$

and there is nothing to do.

Since G is $\mathcal{F}$-free but $\lambda(G)$ is strictly larger than $\lambda(\mathrm{Forb}(\mathcal{F}\cup\{\tau\}))$, the supersaturation argument of Erdős and Simonovits [12] or an application of the Removal Lemma shows that

$$p(\tau, G) \geqslant 1/C_1, \tag{16}$$

that is, G has at least $\binom{n}{q}/C_1$ copies of τ.

For every embedding $\psi : \tau \to G$, we define the following. For each binary string $\boldsymbol{b} = (b_1, ..., b_q)$ of length q, let $V_{\psi,\boldsymbol{b}}$ consist of those vertices $x \in V(G)$ such that the neighbourhood of x in $\psi([q])$ is given by $\boldsymbol{b}$, that is, $\{i \in [q] : \{x, \psi(i)\} \in E(G)\} = \{i \in [q] : b_i = 1\}$. Thus, the sets $V_{\psi,\boldsymbol{b}}$, $\boldsymbol{b} \in \{0,1\}^q$, form a partition of $V(G)$. Observe that, if we apply the above definition to the (fixed) map $f : \tau \to B$ (instead of $\psi : \tau \to G$), then each part in the obtained partition of $V(B) = [m]$ has at most one vertex by Condition 2. Let $\boldsymbol{b}^j \in \{0,1\}^q$ be the binary sequence corresponding to the part $\{j\}$ for each $j \in [m]$; thus $\boldsymbol{b}^j$ encodes the adjacencies of $j \in V(B)$ to the fixed copy of τ in B. We call all other length-q binary sequences *singular*. Also, we call a part $V_{\psi,\boldsymbol{b}}$ *singular* if $\boldsymbol{b}$ is singular, that is, not one of $\boldsymbol{b}^1, \dots, \boldsymbol{b}^m$. Finally, we call a pair of distinct vertices $x_1, x_2 \in V(G)$ *singular* if at least one of them is in a singular part or both of them are in non-singular parts but the adjacency relations between x_1, x_2 in G and between j_1, j_2 in B mismatch (that is, one is an edge and the other is a non-edge), where j_l is the unique element of $[m]$ satisfying $x_l \in V_{\psi,\boldsymbol{b}^{j_l}}$ for $l = 1, 2$. Note if we have $j_1 = j_2$ above, then $\{x_1, x_2\}$ is singular if and only if x_1 and x_2 are connected in G.

Observe that due to Condition 2, we have that the union of $\psi([q])$ with every singular pair $\{x_1, x_2\}$ induces a graph that does not embed into a blow-up of B. For example, if x_1 is in a singular part then already $\psi([q]) \cup \{x_1\}$ spans a subgraph in G that does not belong to $B()$. If we add an arbitrary disjoint $(N - |X|)$-set Y of vertices to $X := \psi([q]) \cup \{x_1, x_2\}$, we get a subgraph of G of order N that does not belong to $B()$. Condition 3 and inequality (10) give that the total number of such subgraphs in G is at most $C_1\binom{n}{N}\max(1/n, \lambda(\mathcal{G}) - \lambda(G))$, where we assume that $1/C_1$ is smaller than $\min\{u_\lambda(\mathcal{C}) - a_q - b_q : G_q \text{ is non-sharp}\}$. Also, each such subgraph H of G can arise for at most $N!$ triples $(\psi, \{x_1, x_2\}, Y)$, a rough bound on the number of ways to embed τ into H, then choose two more vertices in H and let Y be the rest of $V(H)$. Thus, the number of triples $(\psi, \{x_1, x_2\}, Y)$ as above is at most $C_1\binom{n}{N}\max(1/n, \lambda(\mathcal{G}) - \lambda(G)) \times N!$. Clearly, if we fix the first two entries, namely $(\psi, \{x_1, x_2\})$, then any choice of Y will do and there are at least $\binom{n}{N-q-2}$ choices of Y (as $|X|$ is always at most $q + 2$). Thus the total number of possible choices of $(\psi, \{x_1, x_2\})$ as above is at most

$$C_1\binom{n}{N}\max(1/n, \lambda(\mathcal{G}) - \lambda(G)) \times N!/\binom{n}{N-q-2}.$$

Choose ψ for which the number of singular pairs is at most the average. By (16) and Corollary 3.3, it is at most

$$\frac{C_1\binom{n}{N}\max(1/n,\lambda(\mathcal{G})-\lambda(G))\times N!/\binom{n}{N-q-2}}{\binom{n}{q}/C_1}<C\binom{n}{2}\max(1/n,\lambda(n,\mathcal{G})-\lambda(G)).$$

Observe that one can convert G into a blow-up of B by flipping all singular pairs between non-singular parts of G and merging the singular parts into non-singular ones in an arbitrary way. Thus, for every (and in particular this) ψ, the number of singular pairs is at least $\Delta_{\mathrm{edit}}(G,B())$, which is by definition the minimum number of pairs that one needs to change in G to make is a blow-up of B. This finishes the proof of the theorem. □

5. Sufficient conditions for perfect stability

The aim of this section is to present sufficient conditions for perfect stability. To state our results, we need the notions of strictness and flip-aversion. Their definitions require several other concepts that we introduce in the next section.

5.1. Notation and some preliminary results

Throughout this section we work under the following set of assumptions.

Assumption 5.1. In addition to Assumption 2.1, we assume the following.

1. Each graph in $\mathcal{F}$ is twin-free and
2. $\lambda(\mathcal{G})=\lambda(B())$.

Observe that a trivial consequence of twin-freeness of each $F\in\mathcal{F}$ is the following.

Lemma 5.2. *The set of admissible graphs $\mathcal{G}$ is closed under taking blow-ups.* □

We will also need the following pieces of notation. If G is a graph and x,y is a pair of distinct nodes of G, then by $G\oplus xy$ we denote the graph obtained by flipping the adjacency of x and y, while by $G-x$ we denote the graph obtained by deleting the node x in G. Moreover, if κ is a positive integer, for a graph G of order $n\geqslant\kappa$ and a vertex x of G, we define

$$\begin{aligned}\Lambda(G,x) &:= \Lambda(G)-\Lambda(G-x)\text{ and}\\ \lambda(G,x) &:= \binom{n-1}{\kappa-1}^{-1}\cdot\Lambda(G,x).\end{aligned}\tag{17}$$

The value of $\Lambda(G,x)$ can be determined by summing $\gamma(G[X])$ over all κ-subsets X of $V(G)$ containing x. Also, $\lambda(G,x)$ is the conditional expectation of $\gamma(G[X])$ where X is a random κ-subset of $V(G)$ conditioned on $X\ni x$.

Let $\boldsymbol{a} = (a_1, ..., a_m)$ in $\mathbb{S}_m$ be arbitrary. Consider a blow-up $B' := B(V_1, \ldots, V_m)$ of order n, where $|V_i| = a_i n + O(1)$. Let B'' be obtained from it by adding a new vertex w. Then $\Lambda(B'', w)$ is determined within additive error $O(n^{\kappa-2})$ by the vector of ratios

$$\boldsymbol{y} := \left(\frac{|\Gamma_{B''}(w) \cap V_1|}{|V_1|}, \ldots, \frac{|\Gamma_{B''}(w) \cap V_m|}{|V_m|} \right) \in [0,1]^m. \tag{18}$$

In fact, we have

$$\lambda(B'', w) = R_{\boldsymbol{a}}(\boldsymbol{y}) + O(1/n), \tag{19}$$

where $R_{\boldsymbol{a}} = R_{B,\lambda,\boldsymbol{a}}$ is some real polynomial in $\boldsymbol{y}$. One can write $R_{\boldsymbol{a}}$ explicitly as follows.

First, for a (not necessarily injective) map $\phi : [t] \to [m]$ and a (binary) vector $\boldsymbol{b} = (b_1, \ldots, b_t)$ in $\{0,1\}^t$, let $B(\phi, \boldsymbol{b})$ be the graph on $[t+1]$ such that two elements i and j of $[t]$ are adjacent if and only if $\phi(i)$ and $\phi(j)$ are adjacent in B, and $\{i, t+1\}$ is an edge if and only if $b_i = 1$. Informally speaking, $B(\phi, \boldsymbol{b})$ is a graph that we can form from a blow-up of B on $[t]$ by adding a new vertex whose neighbourhood in $[t]$ is given by the binary vector $\mathbf{b}$. Then the value of the polynomial $R_{\boldsymbol{a}}$ at $\boldsymbol{y} = (y_1, ..., y_m)$ is

$$\sum_{\phi:[\kappa-1]\to[m]} \sum_{\boldsymbol{b}\in\{0,1\}^{\kappa-1}} (\kappa-1)!\gamma(B(\phi,\boldsymbol{b})) \prod_{p=1}^{m} \prod_{q=0}^{1} \frac{(a_p(qy_p + (1-q)(1-y_p)))^{|\{i:\phi(i)=p,\, b_i=q\}|}}{|\{i : \phi(i) = p,\, b_i = q\}|\,!}.$$

Let us call a vector $\boldsymbol{y} \in [0,1]^m$ *admissible* if for every $t \in \mathbb{N}$, every map $\phi : [t] \to [m]$, and every binary vector $\mathbf{b} = (b_1, \ldots, b_t) \in \{0,1\}^t$ such that $y_{\phi(i)} = 0$ implies $b_i = 0$ and $y_{\phi(i)} = 1$ implies $b_i = 1$ (while b_i can be arbitrary if $0 < y_{\phi(i)} < 1$), the graph $B(\phi, \boldsymbol{b})$ is $\mathcal{F}$-free. In other words, this condition says that if we take a blow-up $B(V_1, \ldots, V_m)$ with each $|V_i|$ large and add a vertex w with $y_i|V_i|$ neighbours in V_i for each $i \in [m]$, then the obtained graph is still $\mathcal{F}$-free. Clearly, whether $\boldsymbol{y} = (y_1, \ldots, y_m)$ is admissible or not, depends only on the sets $\{i \in [m] : y_i = 0\}$ and $\{i \in [m] : y_i = 1\}$ and therefore the next claim follows easily.

Claim 5.3. *The set of the admissible vectors forms a closed subset of* $[0,1]^m$. □

Let us point out that, since $\mathcal{F}$ is twin-free, it suffices to check the condition in the definition of an admissible $\boldsymbol{y}$ only for those choices of $t, \phi, \boldsymbol{b}$ for which $B(\phi, \boldsymbol{b})$ is twin-free. In particular, it suffices to consider t to be at most $2m$.

The following vectors will play a special role. For every $i \in [m]$, we define $\boldsymbol{v}_i = (v_{i,1}, ..., v_{i,m})$ in $\{0,1\}^m$ by setting $v_{i,j} = 1$ if ij belongs to $E(B)$ and $v_{i,j} = 0$ otherwise for all $j \in [m]$. Informally speaking, the assignment $\boldsymbol{y} = \boldsymbol{v}_i$ corresponds to adding one extra vertex in part V_i. Thus each vector $\boldsymbol{v}_i \in [0,1]^m$ is admissible. Under this terminology, we have, in particular, for each $i_0 \in [m]$ and $x \in V_{i_0}$ that

$$R_{\boldsymbol{a}}(\boldsymbol{v}_{i_0}) = \lambda(B', x) + O(1/n). \tag{20}$$

Indeed, both sides of (20) measure the (normalised) change in the objective function λ when we remove one vertex from the i_0-th part of an $(\boldsymbol{a} + O(1/n))$-blow-up of B of order n.

There is the following connection between $\lambda(B', x)$ and $\frac{\partial}{\partial a_{i_0}}\lambda(B(\boldsymbol{a}))$.

Claim 5.4.

$$R_{\boldsymbol{a}}(\boldsymbol{v}_{i_0}) = \lambda(B', x) + O(1/n) = \frac{1}{\kappa}\frac{\partial}{\partial a_{i_0}}\lambda(B(\boldsymbol{a})).$$

Proof. First, for every positive integer t and map $\phi : [t] \to [m]$ we define $B(\phi)$ to be the graph having $[t]$ as the vertex set with i and j being adjacent if and only if $\phi(i)$ and $\phi(j)$ are adjacent. For each $H \in \mathcal{G}^0_\kappa$ we define

$$\Phi_H := \{\phi : [\kappa] \to [m] : B(\phi) \cong H\}$$

and

$$\Phi_H^{i_0} := \{\phi \in \Phi_H : i_0 \in \phi([\kappa])\}.$$

Then we have that

$$\lambda(B(\boldsymbol{a})) = \sum_{H\in\mathcal{G}^0_\kappa} \gamma(H)\kappa! \sum_{\phi\in\Phi_H} \prod_{i=1}^{m} \frac{a_i^{|\phi^{-1}(i)|}}{|\phi^{-1}(i)|!}. \tag{21}$$

On the other hand, we have that

$$\begin{aligned}\lambda(B', x) &= \sum_{H\in\mathcal{G}^0_\kappa} \gamma(H)(\kappa-1)! \sum_{\phi\in\Phi_H^{i_0}} \frac{a_{i_0}^{|\phi^{-1}(i_0)|-1}}{(|\phi^{-1}(i_0)|-1)!} \prod_{\substack{i=1\\ i\neq i_0}}^{m} \frac{a_i^{|\phi^{-1}(i)|}}{|\phi^{-1}(i)|!} + O(1/n)\\ &\overset{(21)}{=} \frac{1}{\kappa}\frac{\partial}{\partial a_{i_0}}\lambda(B(\boldsymbol{a})). \quad \square\end{aligned}$$

Let us illustrate some of the above concepts in the special case of Example 1.1 with $\mathcal{H} = \{K_t\}$ (namely, the Turán function $\mathrm{ex}(n, K_t)$). Here $m = t-1$ and $B = K_m$. Ignoring rounding errors, if we create B' from the complete m-partite graph $K_{a_1 n,\dots,a_m n} = B(V_1, \dots, V_m)$ by adding a new vertex w having $y_1 a_1 n, \dots, y_m a_m n$ neighbours in $V_1, \dots, V_m$ respectively, then $\Lambda(B', w)$ is just $\sum_{i=1}^m y_i a_i n$, the number of edges at w. Thus $R_{\boldsymbol{a}}(\boldsymbol{y}) = \lim_{n\to\infty} \Lambda(B', w)/n = \sum_{i=1}^m y_i a_i$. Since we forbid K_{m+1}, a vector $\boldsymbol{y}$ is admissible if and only if at least one y_i is 0. Here $\boldsymbol{v}_i$ is the $\boldsymbol{y}$-vector corresponding to w being a twin of the vertices in V_i, that is, $\boldsymbol{v}_i$ consists of 1s except one 0 at position i. Note that $\boldsymbol{a} = (1/m, \dots, 1/m)$ is the (unique) maximiser of $\lambda(B(\boldsymbol{x}))$ for $\boldsymbol{x} \in \mathbb{S}_m$. If we fix this $\boldsymbol{a}$ and maximise $R_{\boldsymbol{a}}(\boldsymbol{y})$ over admissible $\boldsymbol{y} \in [0, 1]^m$, then trivially the set of maximisers is $\{\boldsymbol{v}_1, \dots, \boldsymbol{v}_m\}$. The following lemma states that one part of this inclusion (namely, that each $\boldsymbol{v}_i$ is a maximiser) holds whenever $\boldsymbol{a}$ has no zero entries. This makes a perfect combinatorial sense: in every extremal configuration $B(V_1, \dots, V_m)$ all vertices must make asymptotically the same contribution to Λ.

Lemma 5.5. *Fix any $\boldsymbol{a} \in \mathbb{S}_m$ that maximises $\lambda(B(\cdot))$. Suppose that $\boldsymbol{a}$ has no zero entries. Then the maximum of $R_{\boldsymbol{a}}(\boldsymbol{y})$ over admissible $\boldsymbol{y} \in [0,1]^m$ is $\lambda(\mathcal{G})$ and, furthermore, $R_{\boldsymbol{a}}(\boldsymbol{v}_i) = \lambda(\mathcal{G})$ for each $i \in [m]$ (that is, each of the vectors $\boldsymbol{v}_i$ is a maximiser).*

Proof. Since $\boldsymbol{a}$ achieves a maximum and lies in the interior of $\mathbb{S}_m$, we have that $\frac{\partial}{\partial_i}\lambda(B(\boldsymbol{a})) = \frac{\partial}{\partial_j}\lambda(B(\boldsymbol{a}))$ for all $i, j \in [m]$. Denote this common value by R. By Claim 5.4, we have $R = \frac{1}{k} R_{\boldsymbol{a}}(\boldsymbol{v}_i)$ for all $i \in [m]$. Since $\lambda(B(\boldsymbol{a}))$ is a homogeneous polynomial of degree κ, we have that

$$\lambda(B(\boldsymbol{a})) = \frac{1}{\kappa}\sum_{i=1}^{m} a_i \frac{\partial}{\partial a_i}\lambda(B(\boldsymbol{a})).$$

This, the fact that $\boldsymbol{a}$ maximises $\lambda(B(\boldsymbol{a}))$, Claim 5.4, and equality $\sum_{i=1}^m a_i = 1$ imply that

$$\lambda(\mathcal{G}) = \lambda(B(\boldsymbol{a})) = \frac{1}{\kappa}\sum_{i=1}^{m} a_i \frac{\partial}{\partial a_i}\lambda(B(\boldsymbol{a})) = \sum_{i=1}^{m} a_i R_{\boldsymbol{a}}(\boldsymbol{v}_i) = R.$$

Thus $R_{\boldsymbol{a}}(\boldsymbol{v}_i) = \lambda(\mathcal{G})$ for all $i \in [m]$.

To prove the first part of the lemma, we derive a contradiction by assuming that some admissible $\boldsymbol{y} \in [0,1]^m$ achieves a strictly greater value. Let $c := R_{\boldsymbol{a}}(\boldsymbol{y}) - \lambda(\mathcal{G}) > 0$ and pick some real ε with $0 < \varepsilon \ll c$.

Here we can start with $B' = B(V_1, \dots, V_m)$ of order $n \to \infty$ with $|V_i|/n \to a_i$ and form B'' by adding a set Y of εn new vertices that span an independent set with the identical adjacencies to $V_1, \dots, V_m$ governed by $\boldsymbol{y}$. Since the vector $\boldsymbol{y}$ is admissible, the obtained graph is $\mathcal{F}$-free. Indeed, $B(V_1, \dots, V_m)$ plus one vertex $v \in Y$ is $\mathcal{F}$-free by the admissibility of $\boldsymbol{y}$; by blowing up the vertex v we cannot violate $\mathcal{F}$-freeness because each member of $\mathcal{F}$ is twin-free.

The contribution of the new vertices to λ is $c\varepsilon - O(\varepsilon^2)$. Indeed, if we take a random κ-subset X of $V(B'')$, then with probabilities respectively $1 - \kappa\varepsilon + O(\varepsilon^2)$, $\kappa\varepsilon + O(\varepsilon^2)$, and $O(\varepsilon^2)$, the set X intersects Y in zero, one and at least two vertices; thus

$$\lambda(B'') = (1-\kappa\varepsilon)\lambda(B') + \kappa\varepsilon(\lambda(B') + c) + O(\varepsilon^2) = \lambda(B') + c\kappa\varepsilon + O(\varepsilon^2).$$

So we see that $\lambda(B'') - \lambda(B')$ can be made strictly positive by choosing small constant $\varepsilon \ll c$. Thus the $(\boldsymbol{a} + o(1))$-blow-up B' of B is not asymptotically optimal, contradicting the optimality of $\boldsymbol{a}$. $\square$

Let us say that B is $(\lambda, \boldsymbol{a})$-*strict* if the set of maximisers of $R_{\boldsymbol{a}}(\boldsymbol{y})$ over the admissible $\boldsymbol{y}$'s in $[0,1]^m$ is exactly $\{\boldsymbol{v}_1, \ldots, \boldsymbol{v}_m\}$. Call the graph B λ-*strict* if B is λ-minimal and B is $(\lambda, \boldsymbol{a})$-strict for every $\boldsymbol{a} \in \mathbb{S}_m$ that maximises $\lambda(B(\boldsymbol{a}))$.

Recall that B is λ-*minimal* if $\lambda(B'())$ is strictly smaller than $\lambda(B())$ for any proper subgraph B' of B. It trivially follows that such B is necessarily twin-free and every maximiser $\boldsymbol{a} \in \mathbb{S}_m$ has all coordinates non-zero (and by compactness at least one maximiser $\boldsymbol{a}$ exists). Thus, if B is λ-strict, then, for each optimal $\boldsymbol{a}$, each of $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_m \in [0,1]^m$ is a maximiser of $R_{\boldsymbol{a}}$ by Lemma 5.5 while the strictness property requires that there are no other maximisers.

Lemma 5.6. *If B is λ-strict, then for every $\varepsilon > 0$ there is $\delta > 0$ such that if $\boldsymbol{a} \in \mathbb{S}_m$ and admissible $\boldsymbol{y} \in [0,1]^m$ satisfy $\lambda(B(\boldsymbol{a})) \geqslant \lambda(\mathcal{G}) - \delta$ and $R_{\boldsymbol{a}}(\boldsymbol{y}) \geqslant \lambda(\mathcal{G}) - \delta$, then $\boldsymbol{y}$ is ε-close to some $\boldsymbol{v}_i$.*

Proof. Suppose there is $\varepsilon > 0$ that violates the lemma, that is, for every $j \in \mathbb{N}$ there are $\boldsymbol{a}^j$ and $\boldsymbol{y}^j$ that violate the conclusion for $\delta = 1/j$. By passing to a subsequence, we may assume that these vectors converge to $\boldsymbol{a}$ and $\boldsymbol{y}$ respectively. By the continuity of $\lambda(B(\boldsymbol{x}))$, $\boldsymbol{a}$ is a maximiser. By λ-minimality of B we have that each $a_i > 0$. By Claim 5.3, we have that $\boldsymbol{y}$ is admissible, while by Lemma 5.5 and Assumption 5.1 we have that $R_{\boldsymbol{a}}(\boldsymbol{y}) = \lambda(\mathcal{G})$. Since B is λ-strict, we have that $\boldsymbol{y}$ is equal to $\boldsymbol{v}_i$ for some $i \in [m]$. But then $\boldsymbol{y}^j$ has to get ε-close to $\boldsymbol{v}_i$ leading to a contradiction. □

Here is another easy consequence of the compactness of $\mathbb{S}_m$.

Lemma 5.7. *If B is λ-minimal, then there is $\delta > 0$ such that for every $\boldsymbol{a} \in \mathbb{S}_m$ satisfying $\lambda(B(\boldsymbol{a})) \geqslant \lambda(\mathcal{G}) - \delta$ we have that each a_i is at least δ.* □

Finally, we call a graph B λ-*flip-averse* if there is $\delta > 0$ such that the following holds. If we take a blow-up $B' = B(V_1, \ldots, V_m)$ with $n \geqslant 1/\delta$ vertices such that $\lambda(B') \geqslant \lambda(B()) - \delta$ and obtain $B' \oplus xy$ by changing the adjacency between a pair of distinct nodes $x, y \in V(B')$ (possibly from the same part), then either $B' \oplus xy$ contains some $H \in \mathcal{F}$ with $v(H) \leqslant m + 2$ as a subgraph or we have that

$$\Lambda(B') - \Lambda(B' \oplus xy) \geqslant \delta n^{\kappa-2}. \tag{22}$$

By compactness, the property of being flip-averse can be equivalently re-formulated in terms of the polynomial $\lambda(B(\boldsymbol{x}))$, where n disappears from the definition but then its combinatorial meaning will be less clear.

5.2. Main results for perfect stability

This section consists of two results, each providing a sufficient condition for perfect stability. The first one is the following.

Theorem 5.8 *(Perfect stability I). Suppose that, in addition to Assumptions 2.1 and 5.1, the following assumptions hold.*

1. *The problem is classically B-stable.*
2. *The graph B is λ-strict.*
3. *The graph B is λ-flip-averse.*

Then the problem is perfectly B-stable.

Proof. Given $\lambda, B, \mathcal{F}$, we fix sufficiently small positive constants $c_6 \gg c_5 \gg c_4 \gg c_3 \gg c_2 \gg c_1$. To prove the perfect stability we pick some large enough real number C (depending on the previous constants). In this proof, let asymptotic notation such as $O(1)$ or $\Omega(1)$ hide constants that depend on $\mathcal{F}$, κ, $\gamma(\cdot)$, and B only (but not on the constants c_i).

Let n be an integer with $n > C$. Choosing C large enough, we may assume that $\lambda(\mathcal{G}) + c_1/2 \geqslant \lambda(n, \mathcal{G})$. Let G be an arbitrary admissible graph on $[n]$. Assume that $\lambda(G) \geqslant \lambda(n, \mathcal{G}) - c_1/2$ for otherwise the result follows trivially, since $Cc_1/2 > 1$ and the normalised distance δ_{edit} is always bounded by 1. By Condition 1, that is, the classical B-stability, there is a partition $[n] = V_1 \cup \ldots \cup V_m$ such that

$$|W| \leqslant c_2 \binom{n}{2}, \tag{23}$$

where $W := E(G) \triangle E(B')$ and $B' := B(V_1, \ldots, V_m)$. We call pairs in W *wrong*. Assume that the parts $V_1, \ldots, V_m$ were chosen so that $|W|$ is minimum. Clearly, this choice of parts V_i implies that (23) still holds. Since the number of κ-subsets X of $[n]$ such that $G[X] \not\cong B'[X]$ is at most $|W|\binom{n}{\kappa-2}$, we conclude that

$$\Lambda(B') \geqslant \Lambda(G) - |W|\binom{n}{\kappa - 2} \cdot 2\|\gamma\|_\infty \geqslant (\lambda(\mathcal{G}) - O(c_2))\binom{n}{\kappa}, \tag{24}$$

where $\|\gamma\|_\infty := \max\{|\gamma(H)| : H \in \mathcal{G}^0_\kappa\}$.

Let us call a vertex x *special* if $\lambda(G, x) < \lambda(\mathcal{G}) - c_4$. We set S to be the set of special vertices and $\sigma := |S|/n$.

For each $i \in [m]$, we set $b_i := |V_i|/n$. By (24), the continuity of $\lambda(B(\cdot))$, and the compactness of $\mathbb{S}_m$, we can assume that the vector $\boldsymbol{b} = (b_1, \ldots, b_m)$ is c_3-close to a maximiser $\boldsymbol{a}$ of $\lambda(B(\cdot))$, that is,

$$\|\boldsymbol{a} - \boldsymbol{b}\|_1 \leqslant c_3. \tag{25}$$

By Lemma 5.7, we can assume that each $a_i \geqslant c_6$; thus we conclude that $b_i \geqslant c_6 - c_3 \geqslant c_6/2$ for each $i \in [m]$.

At this point, we can give an informal overview of the rest of the proof. First, Claim 5.9 shows that, for every vertex x of G, the normalised contribution $\lambda(G, x)$ of a vertex x to

$\lambda(G)$ is less than $\lambda(\mathcal{G})+c_2$ for otherwise the addition of an appropriate number of clones of x to G will bring $\lambda(G)$ well over $\lambda(\mathcal{G})$, which is impossible. It follows that, in order to avoid $\lambda(G)$ being too small, we have that $\sigma = O(c_2/c_4)$. Furthermore, the adjacency of each vertex $x \in [n] \setminus S$ essentially follows the ideal adjacency of part-i vertices, for some $i \in [m]$, as this is the only possibility to have $\lambda(G, x)$ close to $\lambda(\mathcal{G})$ by the assumed λ-strictness. Since our choice of the parts V_i minimises the number of wrong adjacencies, this vertex x has to belong to V_i and thus its wrong degree $|\Gamma_W(x)|$ is necessarily small, see (29). (Also, somewhat conversely, each vertex $x \in S$ has high wrong degree, just to account for the drop $\lambda(x, G) < \lambda(\mathcal{G}) - c_4$.) This, the near-optimality of $\boldsymbol{b}$, the fact that $|S| = \sigma n$ is small and the λ-flip-aversion give that every edge-flip inside $[n]\setminus S$ has negative effect on λ (Claim 5.11), not only with respect to G but also with respect an arbitrary graph $\widetilde{G}$ obtained from G by changing some adjacencies inside W (Claim 5.12). Thus if we flip W', all wrong pairs outside S, and "fix" each vertex of S, then Λ increases by at least $\Omega(c_4 n^{\kappa-2})$ per one changed edge. (Note that, since all vertices of high W-degree are inside the small set S, the "pairwise" effects can be shown to be negligible.) On the other hand, $|W'| + n|S|$ is clearly an upper bound on the edit distance from G to the family $B()$. These two estimates give the perfect stability.

Let us provide all the remaining details now.

Claim 5.9. *For every vertex x, we have that $\lambda(G, x) < \lambda(\mathcal{G}) + c_2$.*

Proof of Claim. We assume on the contrary that there exists a node x_0 satisfying $\lambda(G, x_0) \geqslant \lambda(\mathcal{G}) + c_2$. Set $\varepsilon = {c_2}^2$. Consider G' obtained from G by adding εn clones of x_0. We view $\lambda(G')$ as the expectation of $\gamma(G'[X])$ for a random κ-set X. With probability at least $1 - \kappa\varepsilon$, the set X is disjoint from the added clones and its conditional expectation is exactly $\lambda(G)$. With probability $\kappa\varepsilon + O(\varepsilon^2)$, the set X has exactly one element from the added clones and avoids x_0. Conditioned on the latter event, $G'[X]$ is the same as $G[Y]$ where we take a random κ-subset Y of $V(G)$ conditioned on $Y \ni x_0$; thus the conditional expectation of $\gamma(G'[X])$ is exactly $\lambda(G, x)$ (which we assumed to be at least $\lambda(\mathcal{G}) + c_2$). Finally, the contribution from the remaining sets is in the absolute value at most $2\|\gamma\|_\infty$ times their probability $O(\varepsilon^2)$. Also, note our choice of G such that $\lambda(G) \geqslant \lambda(n, \mathcal{G}) - c_1/2 \geqslant \lambda(\mathcal{G}) - c_1$. Thus

$$\begin{aligned}\lambda(G') &\geqslant (1-\kappa\varepsilon)\lambda(G) + \kappa\varepsilon\lambda(G, x) + O(\varepsilon^2)\\ &\geqslant (1-\kappa\varepsilon)(\lambda(\mathcal{G}) - c_1/2) + \kappa\varepsilon(\lambda(\mathcal{G}) + c_2) - O(\varepsilon^2)\\ &\geqslant \lambda(\mathcal{G}) + \kappa {c_2}^3 - c_1 - O({c_2}^4).\end{aligned}$$

This is strictly larger than $\lambda(\mathcal{G})$. On the other hand, by Lemma 5.2, we have that G' is admissible and therefore, invoking Lemma 2.2, we get that $\lambda(G') \leqslant \lambda(\mathcal{G}) + O(1/n) \leqslant \lambda(\mathcal{G}) + O(1/C)$, a contradiction. □

If we pick a uniform random $x \in [n]$, then the difference $\lambda(\mathcal{G}) - \lambda(G, x)$ is never below $-c_2$ by Claim 5.9, while with probability σ it is at least c_4. On the other hand, the average of $\lambda(\mathcal{G}) - \lambda(G, x)$ over $x \in V(G)$ is $\lambda(\mathcal{G}) - \lambda(G) \leqslant c_1$. Thus $-(1-\sigma)c_2 + \sigma c_4 \leqslant c_1$ and, roughly, $\sigma \leqslant 2c_2/c_4$.

Take any $x \in [n]$. Let B'_x be obtained from B' by changing adjacencies at x so that $\Gamma_{B'_x}(x) = \Gamma_G(x)$. We have that $\lambda(B'_x, x) = R_{\boldsymbol{b}}(\boldsymbol{y}_x) + O(1/n)$, where $\boldsymbol{y}_x = (y_{x,1}, \ldots, y_{x,m})$ is an element of $[0,1]^m$ defined by $y_{x,i} := |\Gamma_G(x) \cap V_i|/|V_i|$ for all $i \in [m]$. We also define another element $\boldsymbol{y}'_x = (y'_{x,1}, \ldots, y'_{x,m})$ of $[0,1]^m$ by setting $y'_{x,i} := y_{x,i}$ unless if $y_{x,i} \leqslant c_3/m$ (resp. $y_{x,i} \geqslant 1 - c_3/m$), then we set $y'_{x,i} := 0$ (resp. $y'_{x,i} := 1$). Clearly,

$$\|\boldsymbol{y}_x - \boldsymbol{y}'_x\|_1 \leqslant c_3. \tag{26}$$

Claim 5.10. *The vector $\boldsymbol{y}'_x$ is admissible.*

Proof of Claim. Suppose that the claim does not hold. Let this be witnessed by a vector $\boldsymbol{b} \in \{0,1\}^v$ and a map $\phi : [v] \to [m]$. Then $y'_{x,\phi(i)} \in \{0,1\}$ implies $b_i = y'_{x,\phi(i)}$, while the graph $B(\phi, \boldsymbol{b})$ is of order $v+1$ and not $\mathcal{F}$-free. As we observed after the definition of an admissible vector, one can assume that $v \leqslant 2m$. If $y'_{x,i}$ does not belong to $\{0,1\}$, then $y_{x,i}$ is c_3/m-far from 0 and 1. Also, we know that each V_i has at least $c_6 n/2$ vertices.

Let us show that B'_x has at least $\Omega((c_3 c_6 n)^v)$ copies of $B(\phi, \boldsymbol{b})$ via x. In fact, it is enough to consider only the copies where the vertex $v+1$ of $B(\phi, \boldsymbol{b})$ is mapped into x. For $i \in [v]$, let T_i be $V_i \setminus \Gamma_G(x)$ if $b_i = 0$ and $V_i \cap \Gamma_G(x)$ if $b_i = 1$; note that T_i always has at least $|V_i| \times c_3/m - O(1)$ vertices. Now, if we map each $i \in [v]$ arbitrarily into T_i, then these vertices together with x form a copy of $B(\phi, \boldsymbol{b})$ in B'_x, giving at least the stated number of copies.

Each of the above copies contains a wrong pair which is not adjacent to x. (Recall that G is $\mathcal{F}$-free but $B(\phi, \boldsymbol{b})$ is not and that the vertex x has the same neighbourhoods in G and B'_x.) On the other hand, each wrong pair disjoint from x can be counted at most n^{v-2} times. This gives at least $\Omega((c_3 c_6 n)^v)/n^{v-2}$ wrong pairs, contradicting (23) since c_2 is sufficiently small with respect to c_3 and c_6 (and $v \leqslant 2m$). □

By (25) we also have that $|R_{\boldsymbol{b}}(\boldsymbol{y}_x) - R_{\boldsymbol{a}}(\boldsymbol{y}_x)| \leqslant O(c_3)$. On the other hand, there are at most $|W|\binom{n-3}{\kappa-3}$ κ-subsets X of $[n]$ satisfying $x \in X$ and $G[X] \not\cong B'_x[X]$, because each such set must contain a wrong pair disjoint from x. Thus by (23), we have that $|\lambda(G, x) - \lambda(B'_x, x)| \leqslant O(c_2)$. Also, observe that $\lambda(B'_x, x) = R_{\boldsymbol{b}}(\boldsymbol{y}_x) + O(1/n)$. By (26), we get that $|R_{\boldsymbol{a}}(\boldsymbol{y}_x) - R_{\boldsymbol{a}}(\boldsymbol{y}'_x)| \leqslant O(c_3)$. By the Triangle Inequality, we derive that

$$|\lambda(G, x) - R_{\boldsymbol{a}}(\boldsymbol{y}'_x)| \leqslant O(c_3).$$

Suppose furthermore that $x \in [n] \setminus S$. By the definition of S, we have that $\lambda(G, x) \geqslant \lambda(\mathcal{G}) - c_4$ and therefore

$$R_{\boldsymbol{a}}(\boldsymbol{y}'_x) \geqslant \lambda(G, x) - O(c_3) \geqslant \lambda(\mathcal{G}) - O(c_4). \tag{27}$$

By Assumption 2, Lemma 5.6 and Inequality (27), we conclude that $\boldsymbol{y}'_x$ is $(c_5/2)$-close (in the L_1-norm) to the "adjacency vector" $\boldsymbol{v}_i$ of some $i \in [m]$. By (26),

$$\|\boldsymbol{y}_x - \boldsymbol{v}_i\|_1 \leqslant c_5. \tag{28}$$

Next, let us show that x belongs to V_i. Suppose on the contrary that x belongs to V_j for some $j \neq i$. By the twin-freeness of B (which trivially follows from the λ-minimality of B), there is some $h \in [m]$ which is adjacent to exactly one of i and j, say $ih \in E(B)$ but $jh \notin E(B)$. The vertex x is adjacent in G to $y_{x,h}|V_h|$ vertices of V_h. But we know that $|V_h| \geqslant c_6 n/2$ and, by (28), $y_{x,h} \geqslant v_{i,h} - c_5 = 1 - c_5$. On the other hand, B' has no edges between $V_i \ni x$ and V_h. Thus x belongs to at least $(1-c_5)c_6 n/2$ wrong pairs having an endpoint in V_h. Let as denote this set of edges by A. Consider changing the partition $V_1 \cup \ldots \cup V_m$ by moving x to V_i. Observe that the new set of wrong pairs will differ from the old one only on edges containing x. By (28), at most $c_5 n$ edges can be introduced into the set of wrong pairs, while every edge in A will not be contained, anymore, in the new set of wrong pairs. Thus the number of wrong pairs will strictly decrease. This contradicts the choice of the partition $V_1 \cup \ldots \cup V_m$ and, in particular, the minimality of $|W|$. Thus indeed $x \in V_i$, as claimed.

Thus, again by (28), we have that

$$|\Gamma_W(x)| \leqslant c_5 n, \quad \forall\, x \in [n] \setminus S. \tag{29}$$

Claim 5.11. *For every pair xy in $W \cap \binom{[n]\setminus S}{2}$, the graph $B' \oplus xy$ (which is obtained from B' by changing the adjacency of xy) is $\mathcal{F}$-free and satisfies $\Lambda(B') - \Lambda(B' \oplus xy) \geqslant c_6 n^{\kappa-2}$.*

Proof of Claim. Suppose on the contrary that $B' \oplus xy$ contains a forbidden subgraph $H \in \mathcal{F}$. Since $\mathcal{F}$ consists of twin-free graphs and $B' \oplus xy$ has at most $m+2$ pairwise non-twin vertices, we can assume that H has $v \leqslant m+2$ vertices. In fact, we must have at least $\binom{(c_6/2-\sigma)n}{v-2}$ copies of H on $[n] \setminus S$ via xy in $B' \oplus xy$, since $|V_i| \geqslant c_6 n/2$ for each $i \in [m]$. Notice that the vertex set of each such copy contains a pair from W different from xy. By (29), we have at most $2c_5 n$ wrong pairs adjacent to xy, each in at most n^{v-3} copies of H; while every other wrong pair appears in at most n^{v-4} copies of H. This gives that the total number of H-subgraphs on $[n] \setminus S$ via xy is at most

$$2c_5 n \cdot n^{v-3} + |W| \cdot n^{v-4} \leqslant 4c_5 n^{v-2},$$

where we used (23). This is strictly less than $\binom{c_6 n/2}{v-2}$, a contradiction. This contradiction shows that no such H exists, proving the first part of the claim.

The second part follows from Assumption 3 of the theorem. □

Claim 5.12. *Let $\widetilde{G}$ be an arbitrary (not necessarily $\mathcal{F}$-free) graph having $[n]$ as a vertex set and such that $\widetilde{W} \subseteq W$, where $\widetilde{W} := E(\widetilde{G}) \bigtriangleup E(B')$. Then for every pair xy in $\widetilde{W} \cap \binom{[n]\setminus S}{2}$ we have*

$$\Sigma' := \sum_{X\in\binom{[n]\setminus S}{\kappa}} \Big(\gamma((\widetilde{G}\oplus xy)[X]) - \gamma(\widetilde{G}[X])\Big) > c_6 n^{\kappa-2}/2. \tag{30}$$

Proof of Claim. Let us estimate $\Sigma' - \Sigma''$, where we define

$$\Sigma'' := \sum_{X\in\binom{[n]\setminus S}{\kappa}} \left(\gamma(B'[X]) - \gamma((B'\oplus xy)[X])\right).$$

Let X be a κ-subset of $[n]\setminus S$ that contributes different amounts to Σ' and Σ''. Clearly, both x and y belong to X; also X has to contain at least one further pair $ab \in \widetilde{W}$. The number of the κ-subsets X containing a pair $ab \in \widetilde{W}$ satisfying $\{a,b\}\cap\{x,y\} = \emptyset$ is at most $|\widetilde{W}| \leqslant |W|$ (the number of choices of ab) times $\binom{n-4}{\kappa-4}$ (the number of choices of $X\setminus\{a,b,x,y\}$). Likewise, the number of the κ-subsets X containing a pair $ab\in\widetilde{W}$ satisfying $\{a,b\}\cap\{x,y\}\neq\emptyset$ is at most the number of wrong pairs adjacent to x or y, which by (29) satisfies

$$|\Gamma_{\widetilde{W}}(x)| + |\Gamma_{\widetilde{W}}(y)| \leqslant |\Gamma_W(x)| + |\Gamma_W(y)| \leqslant 2c_5 n,$$

times $\binom{n-3}{\kappa-3}$. Thus, (23) gives that $|\Sigma' - \Sigma''| \leqslant O(c_5 n^{\kappa-2})$. On the other hand, the sum

$$\Sigma''' := \sum_{X\in\binom{[n]}{\kappa}\setminus\binom{[n]\setminus S}{\kappa}} \left(\gamma(B'[X]) - \gamma((B'\oplus xy)[X])\right)$$

has at most $|S|n^{\kappa-3}$ non-zero terms (all such X have to contain the pair xy as well as intersect S). Observe that $\Sigma''+\Sigma''' = \Lambda(B')-\Lambda(B'\oplus xy)$ is at least $c_6 n^{\kappa-2}$ by Claim 5.11. Thus $\Sigma' \geqslant c_6 n^{\kappa-2}/2$, as desired. □

Enumerate $W' := W\cap\binom{[n]\setminus S}{2}$ as $\{e_1,\dots,e_w\}$. Let $G_0 := G$ and for $i=1,\dots,w$, let $G_i = G_{i-1}\oplus e_i$; that is, we flip the wrong pairs on $[n]\setminus S$ in some order. The final graph G_w coincides with B' on $[n]\setminus S$. By using Claim 5.12 to estimate the effect of each of the w flips, we conclude that

$$\sum_{X\in\binom{[n]\setminus S}{\kappa}} \left(\gamma(B'[X]) - \gamma(G[X])\right) \geqslant w c_6 n^{\kappa-2}/2. \tag{31}$$

On the other hand, we have that

$$\sum_{\substack{X\in\binom{[n]}{\kappa}\\ X\cap S\neq\emptyset}} \left(\gamma(B'[X]) - \gamma(G[X])\right) \geqslant \sum_{x\in S} \left(\Lambda(B',x) - \Lambda(G,x)\right) - O(|S|^2 n^{\kappa-2}). \tag{32}$$

For each vertex $x\in S$, the value $\lambda(G,x)$ is at most $\lambda(\mathcal{G}) - c_4$ by the definition of S. By Claim 5.4, the value $\lambda(B',x)$ is equal to $\frac{1}{\kappa}\frac{\partial}{\partial_i}\lambda(B(\boldsymbol{b}))+O(1/n)$ where $i\in[m]$ is the index of the part V_i that contains x. Since $\boldsymbol{b}$ is c_3-close to an optimal vector (namely, a vector $\boldsymbol{a} \in \mathbb{S}_m$ that satisfies $\lambda(B(\boldsymbol{a})) = \lambda(\mathcal{G})$), we have that

$$\frac{1}{\kappa}\frac{\partial}{\partial_i}\lambda(B(\boldsymbol{b})) \geqslant \frac{1}{\kappa}\frac{\partial}{\partial_i}\lambda(B(\boldsymbol{a})) - O(c_3) = \lambda(B()) - O(c_3).$$

Thus $\lambda(B', x) - \lambda(G, x) \geqslant c_4 - O(c_3) \geqslant c_4/2$ for each $x \in S$ and invoking (31) and (32), we get

$$\Lambda(B') - \Lambda(G) \geqslant |W'|c_6 n^{\kappa-2}/2 + |S|\frac{c_4}{2}\binom{n-1}{\kappa-1} - O(|W'|\sigma n^{\kappa-2} + |S|^2 n^{\kappa-2}). \tag{33}$$

By (33) (and our bounds on $|W'| \leqslant |W| \leqslant c_2\binom{n}{2}$ and $|S|/n = \sigma \leqslant 2c_2/c_4 \ll \min(c_6, c_4)$), we have that, for example,

$$\Lambda(n, \mathcal{G}) - \Lambda(G) \geqslant \Lambda(B') - \Lambda(G) \geqslant |W'|c_6 n^{\kappa-2}/4 + |S|\frac{c_4}{4}\binom{n-1}{\kappa-1} \geqslant c_3(|W'| + |S|n)\frac{\binom{n}{k}}{\binom{n}{2}}.$$

Observing that $|W'| + |S|n \geqslant |W|$ and $\delta_{\mathrm{edit}}(G, B()) = |W|/\binom{n}{2}$, we derive the perfect stability. □

Theorem 5.13 *(Perfect stability II). Suppose that Assumptions 2.1 and 5.1 are satisfied, the problem is robustly B-stable and B is λ-minimal. Then the problem is perfectly B-stable.*

Proof. Clearly, the perfect stability will follow by Theorem 5.8 if we show that its Assumptions 1, 2 and 3 are satisfied. Assuming that the problem is robustly B-stable, we trivially have that the problem is classically B-stable, that is, Assumptions 1 of Theorem 5.8 is satisfied. Thus it is enough to verify Assumptions 2 and 3 of Theorem 5.8.

Roughly speaking, our proof is based on the following idea. For example, suppose that Assumption 2 (the strictness of λ) fails. Let this be witnessed by a vector $\boldsymbol{y} \in [0, 1]^m$. Then we take a blow-up $G = B(V_1, \ldots, V_m)$ of order n with optimal part ratios and add a set Z of εn twin vertices, each attached to G according to $\boldsymbol{y}$. Since $\boldsymbol{y}$ is $\Omega(1)$-far from each canonical attachment $\boldsymbol{v}_i$, the new graph G' has normalised edit distance $\Omega(\varepsilon)$ to the family $B()$. On the other hand, if we take a random κ-subset $X \subseteq V(G')$ then it either is disjoint from Z (and the conditional expectation of $\gamma(G'[X])$ is exactly $\lambda(G)$), or contains exactly one vertex of Z (and the conditional expectation of $\gamma(G'[X])$ is $\lambda(\mathcal{G}) + o(1)$ by the choice of $\boldsymbol{y}$), or contains at least two vertices of Z (which has probability $O(\varepsilon^2)$). We conclude that $|\lambda(G') - \lambda(G)| = O(\varepsilon^2)$, a contradiction to the robust stability. Likewise, if some edge flip violates Assumption 3 (the flip aversion of λ), then one "magnifies" this by flipping all pairs between two appropriately placed sets of size εn.

Let us continue with the formal proof. Let the robust stability of the problem be satisfied with constant C. Given λ, B and C, we choose a small enough quantity $c > 0$.

In order to prove that the problem is strict, we assume on the contrary that there exist a maximiser $\boldsymbol{a}$ in $\mathbb{S}_m$ of $\lambda(B(\cdot))$ and an admissible $\boldsymbol{y}$ in $[0,1]^m$ violating λ-strictness. Since B is λ-minimal, we have that each $a_i \geqslant c$. We set

$$\delta := \min_{i\in[m]} \|\boldsymbol{y} - \boldsymbol{v}_i\|_1 > 0,$$

and we pick some positive real ε satisfying $\varepsilon \ll \min(c, \delta)$.

Let G be a blow-up $B(V_1, \dots, V_m)$ on $n \to \infty$ vertices with $|V_i|/n \to a_i$. Let G' be obtained from G by adding a set Z of εn twins whose attachment to $V(G)$ is given by the vector $\boldsymbol{y} + o(1)$, where we insist that if $y_i = 0$ (resp. $y_i = 1$), then each $z \in Z$ is adjacent to no vertex in V_i (resp. every vertex in V_i). Since $\boldsymbol{y}$ is admissible and the graphs in $\mathcal{F}$ are twin-free, G' is $\mathcal{F}$-free. Since $R_{\boldsymbol{a}}(\boldsymbol{y}) = \lambda(\mathcal{G})$, we have that the average of $\gamma(G'[X])$ over the κ-subsets X of $V(G')$ with $|X \cap Z| = 1$ is $\lambda(\mathcal{G}) + o(1)$. Thus it follows that $\lambda(\mathcal{G}) - \lambda(G')$ is at most $O(\varepsilon^2)$. By robust stability, the normalised distance from G' to some blow-up $B' = B(U_1, \dots, U_m)$ of B is $O(\varepsilon^2)$. Clearly, $\lambda(B') \geqslant \lambda(G') - O(\varepsilon^2) \geqslant \lambda(\mathcal{G}) - O(\varepsilon^2)$.

Recall that we have partitions $V_1 \cup \dots \cup V_m \cup Z = U_1 \cup \dots \cup U_m$. We have that each $|U_i| \geqslant cn$ for otherwise we obtain the contradiction that

$$\lambda(B') \leqslant \lambda' + O(c) < \lambda(B()) - O(\varepsilon^2),$$

where $\lambda' < \lambda(B())$ is the maximum of λ over all blow-ups of proper subgraphs of B. Similarly, each V_i has at least cn elements.

Claim 5.14. *There is an automorphism $\sigma : [m] \to [m]$ of B such that for each i*

$$|U_{\sigma(i)} \triangle V_i| \leqslant 2\varepsilon n/c. \tag{34}$$

Proof of Claim. We show first that for each $i \in [m]$ there exists $\sigma(i) \in [m]$ satisfying (34) (and then observe that the map $\sigma : [m] \to [m]$ is an automorphism of B). Take any $i \in [m]$. Suppose that there is no choice of $\sigma(i)$ satisfying (34). We pick $x \in [m]$ such that $|U_x \cap V_i| \geqslant |V_i|/m \geqslant cn/m > \varepsilon n/cm$. We distinguish the following two cases.

Case I: There exists $y \in [m]$ such that $y \neq x$ and $|U_y \cap V_i| > \varepsilon n/cm$.

Since B is twin-free (which follows by the λ-minimality of B), pick $h \in [m]$ such that exactly one of x, y is a B-neighbour of h. Then, every $v \in U_h \setminus Z$ is incident to at least $\varepsilon n/cm$ pairs on which the graphs G' and B' differ, because each $v \in U_h \setminus Z$ has different B'-adjacencies to $V_i \cap U_x$ and $V_i \cap U_y$ but the same G'-adjacency to all vertices of $V_i \supseteq (V_i \cap U_x) \cup (V_i \cap U_y)$. Thus $\Delta_{\text{edit}}(G', B') \geqslant (1/2) \cdot |U_h \setminus Z| \cdot (\varepsilon n/cm)$ which is not $O(\varepsilon^2 n^2)$, a contradiction.

Case II: For every $y \in [m]$ such that $y \neq x$ we have that $|U_y \cap V_i| \leqslant \varepsilon n/cm$.

It holds that $|V_i \setminus U_x| \leqslant \varepsilon n/c$. Since we work under the assumption that there is no appropriate choice of $\sigma(i)$, we have, in particular, that $|U_x \triangle V_i| > 2\varepsilon n/c$ and therefore

$|U_x \setminus V_i| > \varepsilon n/c$. We pick $j \in [m]$ with $j \neq i$ such that $U_x \cap V_j > \varepsilon n/cm$. Arguments similar to the ones used in Case I lead to a contradiction.

To complete the proof we show that σ is an automorphism of B. Let us observe that σ is an injection. Indeed, suppose on the contrary that there exist i, j and x in $[m]$ such that $i \neq j$ and $\sigma(i) = \sigma(j) = x$. Then we have that

$$|U_x \triangle V_j| \geqslant |U_x \setminus V_j| \geqslant |U_x \cap V_i| \overset{(34)}{\geqslant} |V_i| - 2\varepsilon n/c \geqslant cn - 2\varepsilon n/c$$

contradicting (34). To prove that σ is edge and non-edge preserving, we assume on the contrary that there exists a pair of nodes ij such that σ does not preserve adjacency. Then the graphs G' and B' differ on every pair uv with $u \in V_i \cap U_{\sigma(i)}$ and $v \in V_j \cap U_{\sigma(j)}$ generating at least $((c - 2\varepsilon/c)n)^2 \gg \varepsilon^2 n^2$ such pairs. The latter is a contradiction to $\Delta_{\text{edit}}(G', B') = O(\varepsilon^2 n^2)$. The claim is proved. □

By relabelling $U_1, \dots, U_m$, we can assume that the bijection σ of Claim 5.14 is the identity map. We expand $(V_i)_{i=1}^m$ to a partition $(V_i')_{i=1}^m$ of the vertex set of G' setting $V_i' = V_i \cup (U_i \cap Z)$ for each $i \in [m]$. Clearly

$$|V_i' \triangle U_i| \leqslant 2\varepsilon n/c \tag{35}$$

for all $i \in [m]$. Finally, we set

$$\Delta_1 := E(G') \triangle E(B(V_1', \dots, V_m')) \text{ and } \Delta_2 := E(B') \triangle E(B(V_1', \dots, V_m')).$$

Each vertex $v \in Z$ is adjacent to at least $\delta n/2$ pairs in Δ_1, because $\boldsymbol{y}$ is δ-far from $\boldsymbol{v}_1, \dots, \boldsymbol{v}_m$, and at most $2\varepsilon mn/c$ pairs in Δ_2. Thus the symmetric difference between G' and B' is at least $\varepsilon n \times (\delta/2 - 2\varepsilon m/c)n \gg \varepsilon^2 v(G')^2$, a contradiction which shows that the graph B is λ-strict.

Next, let us prove the λ-flip-aversion of B. We pick some positive real $\varepsilon \ll c$ and towards a contradiction we assume that there exists some integer n with $n > 1/\varepsilon^3$, an almost optimal blow-up $B' = B(V_1, \dots, V_m)$ on $[n]$ and some pair x, y of distinct nodes such that the graph $B' \oplus xy$ contains no forbidden graph of order at most $m + 2$ and

$$\Lambda(B') - \Lambda(B' \oplus xy) < \varepsilon^3 n^{\kappa - 2}. \tag{36}$$

Let $i, j \in [m]$ be such that $x \in V_i$ and $y \in V_j$. We pick subsets X and Y of V_i and V_j respectively with cardinality εn each. If $i = j$ then we choose X and Y to be disjoint. Let $\mathcal{B}$ be the set of all pairs of nodes with one node in X and one in Y. Also let G be the graph obtained by flipping the adjacency between each pair in $\mathcal{B}$. Since each of X and Y consists of twins, G does not contain any forbidden subgraph.

Let us show that

$$\lambda(B') - \lambda(G) \leqslant O(\varepsilon^3). \tag{37}$$

Indeed, let $\mathcal{A}$ be the set of all κ-element subsets of $V = V_1 \cup \ldots \cup V_m$. We partition $\mathcal{A}$ into $\mathcal{A}_0$, $\mathcal{A}_1$ and $\mathcal{A}_{\geqslant 2}$, the set of all $Z \in \mathcal{A}$ containing respectively zero, one and at least two pairs of $\mathcal{B}$. Finally, for each $e \in \mathcal{B}$, we set $\mathcal{A}_1^e$ and $\mathcal{A}_{\geqslant 2}^e$ to be the set of all $Z \in \mathcal{A}_1$ and $\mathcal{A}_{\geqslant 2}$ respectively, containing e. Note that if $Z \in \mathcal{A}_{\geqslant 2}$, then $|Z \cap (X \cup Y)| \geqslant 3$ and thus $|\mathcal{A}_{\geqslant 2}| = O(\varepsilon^3 n^\kappa)$. We are going to use this fact a couple of times in the following chain of equalities.

$$\begin{aligned}
\Lambda(B') - \Lambda(G) &= \sum_{Z \in \mathcal{A}} (\lambda(B'[Z]) - \lambda(G[Z])) \\
&= \sum_{e \in \mathcal{B}} \sum_{Z \in \mathcal{A}_1^e} (\lambda(B'[Z]) - \lambda(G[Z])) + O(\varepsilon^3 n^\kappa) \\
&= \sum_{e \in \mathcal{B}} \sum_{Z \in \mathcal{A}_1^e} (\lambda(B'[Z]) - \lambda((B' \oplus e)[Z])) + O(\varepsilon^3 n^\kappa) \\
&= \sum_{e \in \mathcal{B}} (\Lambda(B') - \Lambda(B' \oplus e)) + O(\varepsilon^3 n^\kappa) \\
&= \sum_{e \in \mathcal{B}} (\Lambda(B') - \Lambda(B' \oplus xy)) + O(\varepsilon^3 n^\kappa) \overset{(36)}{\leqslant} O(\varepsilon^3 n^\kappa).
\end{aligned}$$

Therefore, by the almost-optimality of B' we have that $|\lambda(G) - \lambda(\mathcal{G})| \leqslant O(\varepsilon^3)$. By the assumed robust stability, there exists some blow-up $B'' = B(U_1, \ldots, U_m)$ of B such that $\delta_{\text{edit}}(B'', G) = O(\varepsilon^3)$. Following arguments as in the proof of Claim 5.14, we can assume that $|V_h \bigtriangleup U_h| \leqslant \varepsilon n/c$ for every $h \in [m]$.

Then we distinguish the following three (non-exclusive) cases.

(i) $|X \setminus U_i| \geqslant \varepsilon n/2$.
(ii) $|Y \setminus U_j| \geqslant \varepsilon n/2$.
(iii) $|X \cap U_i| > \varepsilon n/2$ and $|Y \cap U_j| > \varepsilon n/2$.

We complete the proof by showing that each case leads to a contradiction and, in particular, we show that each case yields that $\delta_{\text{edit}}(G, B'') = \Omega(\varepsilon^2)$. Indeed, let us assume (i). Then there is $i' \neq i$ such that $|X \cap U_{i'}| \geqslant \varepsilon n/2m$. Pick $h \in [m]$ such that the B-adjacencies of $\{h, i\}$ and $\{h, i'\}$ differ. We have at least $(cn - \varepsilon/c)n$ vertices in $U_h \cap V_h$. Thus the symmetric difference between G and B'' is at least $(cn - \varepsilon/c)n \times \varepsilon n/2m \gg \varepsilon^2 n$. Likewise, case (ii) leads to a contradiction. Finally assuming case (iii) we have that G and B'' differ on every pair with one node in $X \cap U_i$ and one in $Y \cap U_j$. Thus the symmetric difference between G and B'' is at least $(\varepsilon n/2)^2$. □

6. Finding optimal asymptotic part ratios

In this section, we provide some analysis related to the values of $\boldsymbol{a}$ in $\mathbb{S}_m$ that maximise the function $\lambda(B(\cdot))$.

While in all examples from Section 1.1 the optimal vector $\boldsymbol{a}$ was uniform, this is not always the case. For example, it was conjectured in [36] (based on the numerical evidence from Flagmatic) that the asymptotically extremal value for Erdős' $f(n,4,4)$-problem is attained by a blow-up of a specific 8-part graph B. If the conjecture is true, then the optimal blow-up of B that minimises the number of $\overline{K}_4$-subgraphs is not uniform (in fact, the optimal part ratios are some irrational numbers). Alternatively, here is a simple although rather artificial example that illustrates the point.

Example 6.1 *(Simple problem with a non-uniform optimal vector).* Let $\mathcal{F}$ consist of all odd cycles plus the graph with 3 vertices and one edge. Then $\mathcal{F}$-free graphs on $[n]$ are exactly complete bipartite graphs, that is, blow-ups of $B = K_2$. Let $\kappa = 6$. Let $\gamma(H) = 0$ except one defines $\gamma(H)$ for $H \in \{K_{0,6}, K_{1,5}, K_{2,4}, K_{3,3}\}$ so that $\lambda(K_{xn,(1-x)n}) = p(x) + o(1)$, where, e.g.

$$p(x) = 12(x - 1/2)^6 - 217(x - 1/2)^4 + 24(x - 1/2)^2.$$

This polynomial p is symmetric around $1/2$ and its maximum on $[0, 1]$ is attained at $x_0 = (3 - \sqrt{2})/6 = 0.264...$ and $1 - x_0$. Finding the maximum of $\lambda(\mathrm{Forb}(\mathcal{F})) = \lambda(B())$ over $\mathbb{S}_2 = \{(x, 1 - x) : x \in [0, 1]\}$ amounts to optimising $p(x)$ over $x \in [0, 1]$ which is not attained for $(1/2, 1/2)$. □

Let us prove a sufficient condition that implies the uniqueness of the maximiser and happens to apply to many concrete problems.

Lemma 6.2. *Let all assumptions of Theorem 4.1 apply. View the graph τ from Assumption 2 also as a type and assume additionally that the flag algebra certificate $\mathcal{C}$ includes a matrix Q^τ of co-rank 1 associated to τ. Then the vector $\boldsymbol{a}$ is the unique maximiser of $\lambda(B(\cdot))$ in $\mathbb{S}_m$.*

Proof. Let $\boldsymbol{b} \in \mathbb{S}_m$ be a maximiser of $\lambda(B(\cdot))$. By Assumption 2(b), we have $\lambda(\mathrm{Forb}(\mathcal{F}) \cup \{\tau\}) < \lambda(\mathrm{Forb}(\mathcal{F})) = \lambda(B(\boldsymbol{b})))$. Thus there is a strong homomorphism f from τ into $B[\{i \in [m] : b_i > 0\}]$. Fix one such f.

For large n, let $G = B(V_1, \ldots, V_m)$ with $|V_i| = b_i n + O(1)$ and take an (injective) embedding $\psi : V(\tau) \to V(G)$ such that $\psi(x) \in V_{f(x)}$ for every $x \in V(\tau)$. Define $\boldsymbol{x_b}$ to be the limit as $n \to \infty$ of the vector $\boldsymbol{x}$ from (12) normalised so that the sum of entries is 1. Clearly, the limit does not depend on the choice of ψ. Arguing as in the proof of Lemma 3.4, we conclude that $\boldsymbol{x_b}$ is a zero eigenvector of Q^τ. Of course, the same applies to the vector $\boldsymbol{x_a}$. Since Q^τ is of co-rank 1, we have that $\boldsymbol{x_b} = \boldsymbol{x_a}$. However, $\boldsymbol{b}$ is uniquely determined from $\boldsymbol{x_b}$. Namely, by Assumption 2(c), the i-th entry b_i is the ℓ-th root, $\ell := (N - v(\tau))/2$, of the entry of $\boldsymbol{x_b}$ that corresponds to the τ-flag obtained by adding some ℓ new vertices from V_i to the ψ-image of τ; this follows by recalling that the vector $\boldsymbol{x_b}$ encodes the limiting distribution of the τ-subflag of G induced by a random κ-subset containing $\psi(V(\tau))$. Thus $\boldsymbol{b} = \boldsymbol{a}$ and $\boldsymbol{a}$ is indeed the unique maximiser of $\lambda(B(\cdot))$ in $\mathbb{S}_m$. □

7. Computer implementation

Combining Theorems 4.1 and 5.13 we obtain the following result, which provides sufficient conditions for perfect stability. The verification of these conditions can be carried out by a computer. In the next section we include such applications.

Theorem 7.1. *Let Assumption 2.1 and Part 1 of Assumption 5.1 apply. Also, we assume all of the following.*

1. *We have a vector $\boldsymbol{a} \in \mathbb{S}_m$ with no zero entries and a certificate $\mathcal{C}$ on N vertices that proves $\lambda(\mathcal{G}) \leqslant \lambda(B(\boldsymbol{a}))$. (Thus, by Assumption 2.1.3, we know that $\lambda(\mathcal{G}) = \lambda(B(\boldsymbol{a}))$.)*
2. *There is a graph τ of order at most $N - 2$ satisfying the following.*
   (a) $\lambda(\mathrm{Forb}(\mathcal{F})) > \lambda(\mathrm{Forb}(\mathcal{F} \cup \{\tau\}))$.
   (b) There exists the unique (up to automorphisms of τ and B) strong homomorphism f from τ into B.
   (c) For every distinct x_1 and x_2 in $V(B)$ we have $\Gamma_B(x_1) \cap f(V(\tau)) \neq \Gamma_B(x_2) \cap f(V(\tau))$.
3. *Every $\mathcal{C}$-sharp graph of order N admits a strong homomorphism into B.*

Additionally, suppose that at least one of the following two statements holds:

(i) the certificate $\mathcal{C}$ contains (as a type) the graph τ from Assumption 2 above and the corresponding matrix Q^τ in $\mathcal{C}$ is of co-rank 1, or
(ii) $\lambda(\mathrm{Forb}(\mathcal{F} \cup \{B\})) < \lambda(\mathrm{Forb}(\mathcal{F}))$.

Then the problem is perfectly B-stable.

Proof. Clearly, all assumptions of Theorem 4.1 are satisfied, so the problem is robustly B-stable. By Theorem 5.13, it is enough to check only that B is λ-minimal.

If Condition (i) holds, then the λ-minimality of B follows from Lemma 6.2 (and the assumption that $\boldsymbol{a}$ has no zero entries). So assume that Condition (ii) holds. Let B' be an arbitrary proper subgraph of B and let B'' be any blow-up of B' on $n \to \infty$ vertices. Since B is twin-free by Condition 2(c), we have that B'' is B-free and thus B'' belongs to $\mathrm{Forb}(\mathcal{F} \cup \{B\})$. Thus

$$\lambda(B'') \leqslant \lambda(\mathrm{Forb}(\mathcal{F} \cup \{B\})) + O(1/n) < \lambda(\mathrm{Forb}(\mathcal{F})) + O(1/n).$$

Again, we conclude that B is λ-minimal, as desired. □

Remark. If Assumptions 1, 2, 3 and (i) of Theorem 7.1 are satisfied, then we have that $\lambda(B(\cdot))$ admits $\boldsymbol{a}$ as a unique maximiser (see Lemma 6.2). This is not the case if Assumptions 1, 2, 3 and (ii) of Theorem 7.1 are satisfied, when the uniqueness of $\boldsymbol{a}$ as a maximiser of $\lambda(B(\cdot))$ is not guaranteed. In this case, one has to investigate the uniqueness of $\boldsymbol{a}$ by other means.

8. Applications of the general theorems

Below is a list of results that directly follow by Theorem 7.1 by running our computer code. The ancillary folder of the arxiv version of this paper contains, for each problem except $f(n, 6, 3)$ and $f(n, 7, 3)$ problems discussed in Section 1.1.1, the flagmatic script `*.sage` which was used to generate the certificate and the transcript of the session `*.txt` when the code is run. Due to arxiv's file size limitations, ancillary folder only contains some certificates `*.js`. All certificate files are in Flagmatic's Github directory at:

https://github.com/jsliacan/flagmatic/tree/master/certificates.

For example, for the $f(n, 4, 3)$-problem discussed in Section 1.1.1, these are the files `f43.sage`, `f43.txt` and `f43.js` respectively.

The reader who would like to verify these results has the following options.

Generate certificates from scratch using flagmatic: For this the reader would need to install our version of *Flagmatic* (which is built upon version 2.0 of Emil Vaughan), the *Sage* environment, and an SDP solver such as *CSDP* or *SDPA/SDPA-DD*. The required version of *Flagmatic* can be downloaded from this URL:

https://github.com/jsliacan/flagmatic

which in particular contains a `README.md` file with directions on how to install it and run our scripts.

Run our verifier script `inspect_certificate.py`: This stand-alone script (which is written in *Python/Sage* and uses exact arithmetic) can be used to verify the bound given by each certificate. It is available at the above URL. Its source code is relatively short and well-documented. (Also, the Appendix to the arXiv version of this paper [35, Appendix] contains some further notes on our implementation.) For example, the complete verification of the certificates `f43.js`, `f43_stab.js` can be invoked with the following shell command:

```
sage -python inspect_certificate.py f43.js -stability 3/25 "4:121324" "5:1223344551" f43_stab.js
```

The full details on how to use the `inspect_certificate.py` verification script can be found at the end of the `README.md` file at https://github.com/jsliacan/flagmatic.

Write an independent verifier: The information on how the data inside our certificate files are organised can be found in [35, Appendix].

In the following, we describe some of the input values (such as N and $\boldsymbol{a}$) that determine $\lambda(\mathcal{G})$ and prove perfect stability in Theorems 1.2–1.9.

Minimising the number of independent sets in triangle-free graphs (Theorem 1.2): Recall that $k \in \{4, ..., 7\}$, $\mathcal{F} = \{K_3\}$, $\kappa = k$ and γ is equal to zero except $\gamma(\overline{K}_k) = -1$. Theorem 1.2 for $k = 4, 5$ follows by Theorem 7.1 for $N = 5$, $B = C_5$, $\boldsymbol{a} = (1/5, \dots, 1/5) \in \mathbb{S}_5$ and $\tau = K_2 \cup K_1$, that is, the disjoint union of an edge and a single node (see scripts `f43.sage` and `f53.sage`).
Unfortunately, our code could not generate certificates when $k \in \{6, 7\}$. This computationally demanding task (with $N = 8$) seems to be very sensitive to the obtained numerical SDP solution and the version of *Sage*. However, the corresponding certificates have already been produced by Pikhurko and Vaughan [36] and we include them in the arxiv version of this paper. By running our script `inspect_certificate.py` on them, one can confirm that the problem is perfectly stable in these two cases, where we let B be the Clebsch graph, $\boldsymbol{a} = (1/16, \dots, 1/16) \in \mathbb{S}_{16}$, and τ be the 5-cycle C_5 with one isolated vertex added. (Interestingly, the correct asymptotic of $f(n, 6, 3)$ can be obtained already for $N = 7$ but we could not satisfy Condition 2 of Theorem 7.1 with this N.) In all the cases above, the uniqueness of $\boldsymbol{a}$ follows from Lemma 6.2, since the corank of Q_τ is 1 for each $k \in \{4, \dots, 7\}$.

Maximising the number of pentagons in triangle-free graphs (Theorem 1.3): Recall that the problem is defined by $\mathcal{F} = \{K_3\}$, $\kappa = 5$, and $\gamma(H)$ equals zero, except $\gamma(C_5) = 1$. Theorem 1.3 follows by Theorem 7.1 for $N = 5$, $B = C_5$, $\boldsymbol{a}$ the vector in $\mathbb{S}_5$ having each entry equal to $1/5$ and $\tau = K_2 \cup K_1$, that is, the disjoint union of an edge and a single node. The uniqueness of $\boldsymbol{a}$ follows from Lemma 6.2, since the co-rank of Q_τ is 1.

Inducibility of the cycle on four vertices (Theorem 1.4): Recall that the problem is defined by $\mathcal{F} = \emptyset$, $\kappa = 4$, and $\gamma(H)$ equals zero, except $\gamma(C_4) = 1$. Theorem 1.4 follows by Theorem 7.1 for $N = 5$, $B = K_2$, $\boldsymbol{a} = (1/2, 1/2)$ and $\tau = K_1$. The uniqueness of $\boldsymbol{a}$ follows from Lemma 6.2, since the co-rank of Q_τ is 1.

Inducibility of K_4 minus an edge (Theorem 1.5): Recall that the problem is defined by $\mathcal{F} = \emptyset$, $\kappa = 4$, and $\gamma(H)$ equals zero, except $\gamma(K_4^-) = 1$. Theorem 1.5 follows by Theorem 7.1 for $N = 7$, $B = K_5$, $\boldsymbol{a}$ the vector in $\mathbb{S}_5$ having each entry equal to $1/5$ and $\tau = K_5$. The uniqueness of $\boldsymbol{a}$ follows from Lemma 6.2, since the co-rank of Q_τ is 1.

Inducibility of $K_{3,2}$ (Theorem 1.6): Recall that the problem is defined by $\mathcal{F} = \emptyset$, $\kappa = 5$, and $\gamma(H)$ equals zero, except $\gamma(K_{3,2}) = 1$. Theorem 1.6 follows by Theorem 7.1 for $N = 6$, $B = K_2$, $\boldsymbol{a} = (1/2, 1/2)$ and $\tau = K_2$. The uniqueness of $\boldsymbol{a}$ follows from Lemma 6.2, since the co-rank of Q_τ is 1.

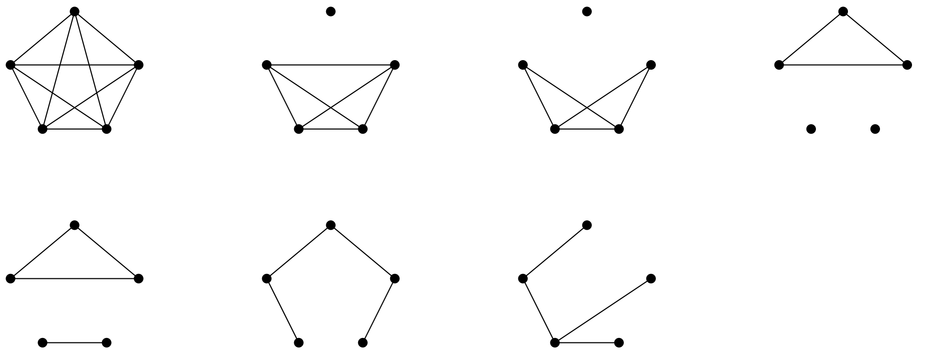

Fig. 2. Sharp graphs that are not a blow of $K_2 \cup K_2$.

Inducibility of $K_{2,2,1}$ ***(Theorem 1.7):*** Recall that the problem is defined by $\mathcal{F} = \emptyset$, $\kappa = 5$, and $\gamma(H)$ equals zero, except $\gamma(K_{2,2,1}) = 1$. Theorem 1.7 follows by Theorem 7.1 for $N = 6$, $B = K_3$, $\boldsymbol{a} = (1/3, 1/3, 1/3)$ and $\tau = K_2$. The uniqueness of $\boldsymbol{a}$ follows from Lemma 6.2, since the co-rank of Q_τ is 1.

Inducibility of $P_3 \cup K_2$ ***(Theorem 1.8):*** Recall that the problem is defined by $\mathcal{F} = \emptyset$, $\kappa = 5$, and $\gamma(H)$ equals zero, except $\gamma(P_3 \cup K_2) = 1$. Theorem 1.8 follows by Theorem 7.1 for $N = 6$, $B = K_3 \cup K_3$, that is, the disjoint union of two triangles, $\boldsymbol{a}$ the vector in $\mathbb{S}_6$ having each entry equal to $1/6$ and $\tau = K_2 \cup K_2$, that is, the disjoint union of two edges. The uniqueness of $\boldsymbol{a}$ follows from Lemma 6.2, since the co-rank of Q_τ is 1.

Inducibility of the "Y" graph (Theorem 1.9): Recall that the problem is defined by $\mathcal{F} = \emptyset$, $\kappa = 5$, and $\gamma(H)$ equals zero, except $\gamma(\mathrm{Y}) = 1$. Theorem 1.8 follows by Theorem 7.1 for $N = 6$, $B = C_5$, that is, the cycle on 5 vertices, $\boldsymbol{a}$ the vector in $\mathbb{S}_5$ having each entry equal to $1/5$ and $\tau = P_4$, that is, the path on 4 vertices. The uniqueness of $\boldsymbol{a}$ follows from Lemma 6.2, since the co-rank of Q_τ is 1.

8.1. Inducibility of the Paw graph

The value $i(F_{\text{paw}})$ has been calculated by Hirst [24], where F_{paw} is the paw graph, that is, the graph obtained by adding a pendant edge to a triangle. We work on the complementary problem. We set $\mathcal{F} = \emptyset$ and γ the map taking the value 0 on every graph of order 4 except the disjoint union of P_3 and a single vertex, that we denote by F, where it takes the value 1. From Hirst's work it follows that $i(F) = 3/8$ and an asymptotically extremal construction is a balanced blow-up of the graph consisting of two disjoint edges. In this section, we show that the problem is $K_2 \cup K_2$-perfectly stable.

Unfortunately, our result does not follow directly by Theorem 7.1, since Condition 3 does not hold for our flag algebra certificate. In particular, according to our certificate the sharp graphs consist of the blow-ups of $K_2 \cup K_2$ on 5 vertices and the graphs listed in Fig. 2. Let us denote by $\mathcal{S}$ the set of sharp graph on 5 vertices and by $\mathcal{NS}$ the set of the non-sharp ones. However, letting $B = K_2 \cup K_2$, $N = 5$ and τ be the disjoint union of an edge and a single vertex, we have that Assumptions 1 and 2 of Theorem 7.1 are satisfied. We refer to them as P1 and P2 respectively. The perfect stability of the problem follows by a sequence of lemmas.

Lemma 8.1. *The graph $K_2 \cup K_2$ is λ-minimal.*

Proof. Let $\mathbf{a} = (a_1, a_2, a_3, a_4)$ in $\mathbb{S}_4$ and set $B = K_2 \cup K_2$. It is easy to see that

$$\begin{aligned}\lambda(B(\mathbf{a})) &= 12(a_1^2 a_2(a_3 + a_4) + a_1 a_2^2(a_3 + a_4) + (a_1 + a_2)a_3^2 a_4 + (a_1 + a_2)a_3 a_4^2) \\ &= 12(a_1 a_2 + a_3 a_4)(a_1 + a_2)(a_3 + a_4).\end{aligned}$$

To prove the λ-minimality of B, it suffices by symmetry to show that the maximum value achieved by $\lambda(B(\mathbf{a}))$ for $\mathbf{a} = (a_1, a_2, a_3, a_4)$ in $\mathbb{S}$ with $a_4 = 0$ is strictly less than $\lambda(B((1/4, 1/4, 1/4, 1/4))) = 3/8$. Equivalently, it suffices to show that the maximum of the map

$$f(x_1, x_2) = 12x_1 x_2(x_1 + x_2)(1 - x_1 - x_2)$$

for (x_1, x_2) in $D = \{(x_1, x_2) : x_1, x_2 \geqslant 0 \text{ and } x_1 + x_2 \leqslant 1\}$, is strictly less than 3/8.

Indeed, observe that f vanishes on the boundary of D and therefore the maximum is achieved in the interior of D. The partial derivatives of f at the maximum satisfy the following:

$$\frac{\partial f}{\partial x_1} = 12(x_2(x_1 + x_2)(1 - x_1 - x_2) + x_1 x_2(1 - 2x_1 - 2x_2)) = 0 \tag{38}$$

and

$$\frac{\partial f}{\partial x_2} = 12(x_1(x_1 + x_2)(1 - x_1 - x_2) + x_1 x_2(1 - 2x_1 - 2x_2)) = 0. \tag{39}$$

Subtracting (38) and (39), we obtain that $x_1 = x_2$. Plugging it into (38), we get that f achieves a maximum at $(3/8, 3/8)$. Thus the maximum of f is 81/256, which is strictly smaller that 3/8 and the proof of the lemma is complete. □

Lemma 8.2. *The problem is classically $K_2 \cup K_2$-stable.*

Proof. Let ε be a positive real. By the Induced Removal Lemma of Alon et al. [1] there exists a positive real η such that for every graph G of order at least $1/\eta$ satisfying $p(H, G) \leqslant \eta$ for all $H \in \mathcal{NS}$, we have that there exists a graph G' of the same order as G such that $\delta_{\text{edit}}(G', G) \leqslant \varepsilon$ and each induced subgraph of G' belongs to $\mathcal{S}$. By (14), there exists a positive real δ such that for each graph G of order at least $1/\delta$ satisfying $\lambda(\mathcal{G}) - \lambda(G) \leqslant \delta$, we have that $p(H, G) \leqslant \eta$ for all $H \in \mathcal{S}$ and therefore there exists some graph G' of the same order as G such that $\delta_{\text{edit}}(G, G') \leqslant \varepsilon$ and each induced subgraph of G' belongs to $\mathcal{S}$.

Since G' is close to G and $\lambda(G)$ is close to $\lambda(\mathcal{G})$, we get that $\lambda(G')$ is close to $\lambda(G)$ and therefore, by P2(a), we have that τ embeds into G'. Recall that τ is the disjoint union of an edge and a single vertex. Without loss of generality, we may assume that $V(\tau) = [3]$ and $\{1, 2\}$ forms an edge in τ. Since G' admits an induced copy of τ, there exists an injective strong homomorphism $\psi : [3] \to V(G')$ between τ and G'. For every $s \in 2^{[3]}$, where we view $2^{[3]}$ as the set of maps from [3] to $\{0, 1\}$, we define

$$V'_s = \{x \in V(G') \setminus \mathrm{Im}(\psi) : \{x, \psi(j)\} \in E(G') \text{ iff } s(j) = 1, \text{ for all } j \in [3]\}. \tag{40}$$

Clearly, $(V'_s)_{s\in 2^{[3]}}$ forms a partition of $V(G') \setminus \mathrm{Im}(\psi)$. Let G'' be the graph obtained by deleting the nodes of G' that belong to some V'_s of cardinality at most 3. Finally, set $V_s = V'_s \cap V(G'')$ for all $s \in 2^{[3]}$. Thus we have that $(V_s)_{s\in 2^{[3]}}$ forms a partition of $V(G'') \setminus \mathrm{Im}(\psi)$, each V_s is either empty or contains at least four elements and every induced subgraph of G' on five vertices belongs to $\mathcal{S}$.

Claim 8.3. *The graph G'' is a blow-up of the disjoint union of two edges, or a disjoint union of a complete graph and a blow-up of an edge, or the disjoint union of a complete graph and an empty graph, or the disjoint union of two complete graphs.*

Before we give the proof of Claim 8.3, let us show how it implies Lemma 8.2. Observe that an isolated clique can contain at most one vertex of an F-subgraph. Thus if we remove all edges inside such cliques in G'', then we do not decrease the number of F-subgraphs. By Claim 8.3, the resulting graph G''' is a blow-up of $K_2 \cup K_2$. By Lemma 8.1, G''' cannot be a blow-up of $K_2 \cup K_1$. This easily implies that G'' itself is a blow-up of $K_2 \cup K_2$. Since G and G'' are close to each other, Lemma 8.2 follows.

Proof of Claim 8.3. To prove the claim, it suffices to show that

1. $V_{(1,1,1)}$ is empty,
2. both $V_{(1,0,1)}$ and $V_{(0,1,1)}$ are empty,
3. $G''[V_{(1,1,0)}]$ is complete,
4. both $G''[V_{(1,0,0)}]$ and $G''[V_{(0,1,0)}]$ are empty graphs,
5. $G''[V_{(0,0,1)}]$ is either complete or empty graph,
6. $G''[V_{(0,0,0)}]$ is an empty graph,
7. for every $(z_1, z_2) \in V_{(1,0,0)} \times V_{(0,1,0)}$, we have that z_1, z_2 form an edge in G'',
8. for every $(z_1, z_2) \in V_{(1,0,0)} \times V_{(0,0,0)}$, we have that z_1, z_2 do not form an edge in G'',
9. for every $(z_1, z_2) \in V_{(0,0,0)} \times V_{(0,0,1)}$, we have that z_1, z_2 form an edge in G'',
10. for every $(z_1, z_2) \in V_{(0,1,0)} \times V_{(0,0,0)}$, we have that z_1, z_2 do not form an edge in G'',
11. there is no edge between $V_{(1,0,0)}$ and $V_{(0,0,1)}$, as well as, between $V_{(0,1,0)}$ and $V_{(0,0,1)}$,
12. for every $(z_1, z_2) \in V_{(0,0,0)} \times V_{(1,1,0)}$, we have that z_1, z_2 do not form an edge in G'',
13. for every $(z_1, z_2) \in V_{(0,0,1)} \times V_{(1,1,0)}$, we have that z_1, z_2 do not form an edge in G'',
14. if $V_{(0,0,0)} \neq \emptyset$, then $G''[V_{(0,0,1)}]$ is an empty graph and
15. if $V_{(0,1,0)} \neq \emptyset$ or $V_{(1,0,0)} \neq \emptyset$ then $V_{(1,1,0)} = \emptyset$.

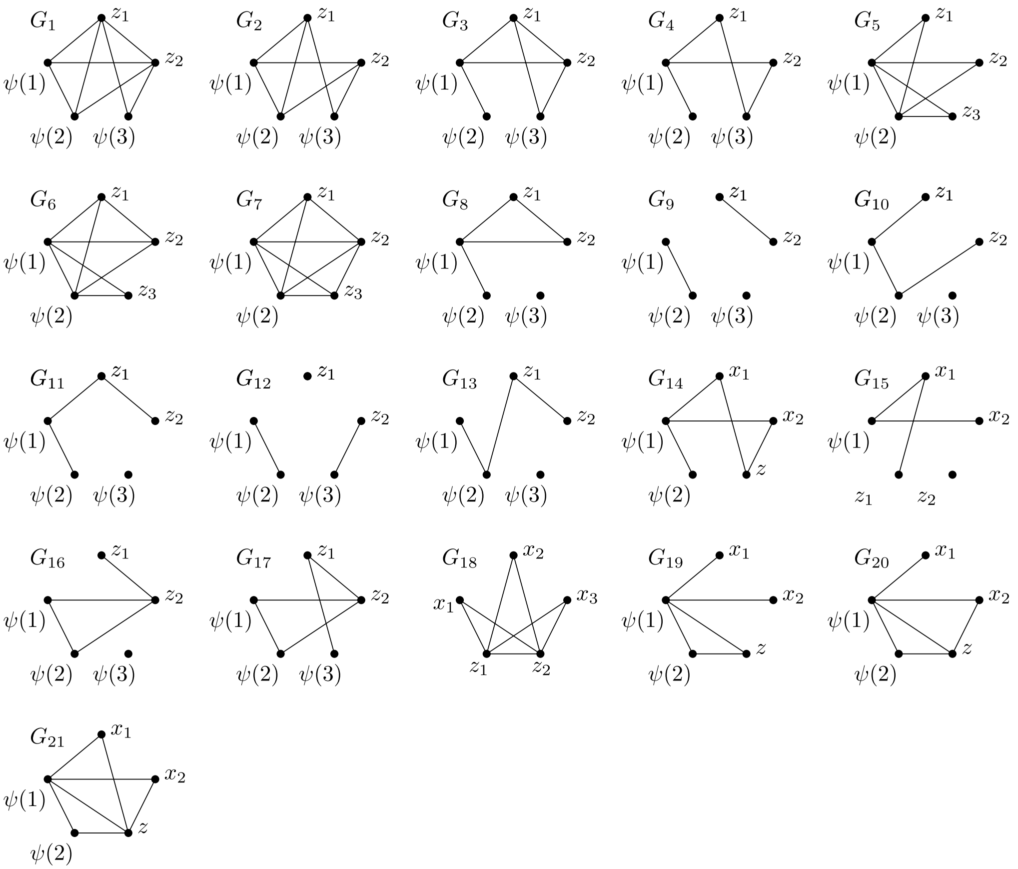

Fig. 3. Non-sharp graphs that are used in the proof of Claim 8.3.

To prove 1, we assume, on the contrary, that $V_{(1,1,1)}$ is non-empty. Thus $V_{(1,1,1)}$ contains at least four elements. We pick distinct z_1 and z_2 in $V_{(1,1,1)}$. There are two cases: either z_1 and z_2 form an edge in G'' or not. The induced subgraphs $G''[\mathrm{Im}(\psi) \cup \{z_1, z_2\}]$ of G'' are the graphs G_1 and G_2 in Fig. 3 respectively. Neither of them belongs to $\mathcal{S}$.

Concerning 2, the arguments justifying that $V_{(1,0,1)}$ and $V_{(0,1,1)}$ are empty are identical. So we only show that $V_{(1,0,1)}$ is empty. Again we assume the contrary and pick two distinct elements z_1 and z_2 in $V_{(1,0,1)}$. Then the induced subgraph of G'' on $\mathrm{Im}(\psi) \cup \{z_1, z_2\}$ is either G_3 or G_4 in Fig. 3, depending on whether z_1z_2 forms an edge or not in G''. Neither G_3 nor G_4 belongs to $\mathcal{S}$.

Towards 3, assuming the contrary, we pick distinct z_1, z_2 and z_3 in $V_{(1,1,0)}$ that do not form a triangle. Then the induced subgraph of G'' on $\{\psi(1), \psi(2), z_1, z_2, z_3\}$ is either the graph G_5 or G_6 or G_7 from Fig. 3, depending on whether z_1, z_2, z_3 span zero, one or two edges respectively. None of these graphs belongs to $\mathcal{S}$.

Concerning 4, the arguments justifying that $G''[V_{(1,0,0)}]$ and $G''[V_{(0,1,0)}]$ are empty graphs are identical. So we only show that $G''[V_{(1,0,0)}]$ is an empty graph. Assume on the contrary that there exist distinct z_1 and z_2 in $V_{(1,0,0)}$ that form an edge in G''. Then $G''[\mathrm{Im}(\psi) \cup \{z_1, z_2\}]$ is the graph G_8 in Fig. 3 and does not belong to $\mathcal{S}$.

To see 5, recall that $V_{(0,0,1)}$ is either empty or contains at least four elements. If $V_{(0,0,1)}$ is empty then our claim holds trivially. So let us assume that $V_{(0,0,1)}$ is of cardinality at least 4 and pick $z_1, z_2, z_3, z_4 \in V_{(0,0,1)}$. Let $H = G''[\{\psi(3), z_1, z_2, z_3, z_4\}]$. Observe that $\psi(3)$ is of degree 4 in H. The only graphs in $\mathcal{S}$ that contain a node of degree 4 are the star and the complete graph. Thus H has to be isomorphic to one of these two, yielding that $G''[\{z_1, z_2, z_3, z_4\}]$ is either an empty or a complete graph, respectively, on 4 vertices, and $G''[V_{(0,0,1)}]$ is either an empty or a complete graph, respectively.

To prove 6, we assume the contrary and pick z_1, z_2 in $V_{(0,0,0)}$ that form an edge in G''. Then the induced subgraph of G'' on $\mathrm{Im}(\psi) \cup \{z_1, z_2\}$ is the graph G_9 in Fig. 3, which does not belong to $\mathcal{S}$.

To prove 7, assuming the contrary, we pick z_1 in $V_{(1,0,0)}$ and z_2 in $V_{(0,1,0)}$ that do not form an edge. Then graph G_{10} in Fig. 3 is the induced subgraph of G'' on $\mathrm{Im}(\psi) \cup \{z_1, z_2\}$ and does not belong to $\mathcal{S}$.

To prove 8, assuming the contrary, we pick z_1 in $V_{(1,0,0)}$ and z_2 in $V_{(0,0,0)}$ that form an edge. Then graph G_{11} in Fig. 3 is the induced subgraph of G'' on $\mathrm{Im}(\psi) \cup \{z_1, z_2\}$ and does not belong to $\mathcal{S}$.

Similarly, to prove 9, assuming the contrary, we pick z_1 in $V_{(0,0,0)}$ and z_2 in $V_{(0,0,1)}$ that do not form an edge. Then graph G_{12} in Fig. 3 is the induced subgraph of G'' on $\mathrm{Im}(\psi) \cup \{z_1, z_2\}$ and does not belong to $\mathcal{S}$.

To prove 10, assuming the contrary, we pick z_1 in $V_{(0,1,0)}$ and z_2 in $V_{(0,0,0)}$ that form an edge. Then graph G_{13} in Fig. 3 is the induced subgraph of G'' on $\mathrm{Im}(\psi) \cup \{z_1, z_2\}$ and does not belong to $\mathcal{S}$.

Both assertions in 11 follow by identical arguments. So let us show that there is no edge between $V_{(1,0,0)}$ and $V_{(0,0,1)}$. First, we show that there is no vertex in one of these sets having more than one neighbour in the other. There are two cases and the arguments are similar. So we show that there is no vertex in $V_{(0,0,1)}$ having at least two neighbours in $V_{(1,0,0)}$. Indeed, assuming the contrary, we have that there exists $z \in V_{(0,0,1)}$ having at least two neighbours in $V_{(1,0,0)}$, say x_1, x_2. Then G'' induces on $\{\psi(1), \psi(2), x_1, x_2, z\}$ the graph G_{14}, which does not belong to $\mathcal{S}$. Finally, assuming that there is an edge between $V_{(1,0,0)}$ and $V_{(0,0,1)}$, we can find $x_1, x_2 \in V_{(1,0,0)}$ and $z_1, z_2 \in V_{(0,0,1)}$ such that x_1, z_1 form an edge, while x_2, z_2 do not. Then G'' induces on $\{\psi(1), x_1, x_2, z_1, z_2\}$ the graph G_{15}, which does not belong to $\mathcal{S}$.

To prove 12, assuming the contrary, we pick z_1 in $V_{(0,0,0)}$ and z_2 in $V_{(1,1,0)}$ that form an edge. Then the graph G_{16} in Fig. 3 is the induced subgraph of G'' on $\mathrm{Im}(\psi) \cup \{z_1, z_2\}$ and does not belong to $\mathcal{S}$.

Similarly, to prove 13, assuming the contrary, we pick z_1 in $V_{(0,0,1)}$ and z_2 in $V_{(1,1,0)}$ that form an edge. Then the graph G_{17} in Fig. 3 is the induced subgraph of G'' on $\mathrm{Im}(\psi) \cup \{z_1, z_2\}$ and does not belong to $\mathcal{S}$.

Towards 14, we assume on the contrary that $V_{(0,0,0)}$ is non-empty and $G''[V_{(0,0,1)}]$ is not an empty graph. Thus there exist $z_1, z_2 \in V_{(0,0,1)}$ that form an edge in G''. Pick distinct $x_1, x_2, x_3 \in V_{(0,0,0)}$. Invoking Items 6 and 9, we have that induced subgraph of G'' on $\{x_1, x_2, x_3, z_1, z_2\}$ is the graph G_{18} in Fig. 3, which does not belong to $\mathcal{S}$.

Concerning 15, the arguments yielding that $V_{(1,1,0)} = \emptyset$ assuming $V_{(1,0,0)} \neq \emptyset$ are identical to the ones yielding that $V_{(1,1,0)} = \emptyset$ assuming $V_{(0,1,0)} \neq \emptyset$. So let us show the first implication. Assume on the contrary that both $V_{(1,1,0)}$ and $V_{(1,0,0)}$ are non-empty. We pick $x_1, x_2 \in V_{(1,0,0)}$ and $z \in V_{(1,1,0)}$. By item 4, we have that x_1, x_2 do not form an edge in G''. Thus the induced subgraph of G'' on $\{\psi(1), \psi(2), x_1, x_2, z\}$ is either the graph G_{19} or G_{20} or G_{21} from Fig. 3, depending on whether z, x_1, x_2 span zero, one or two edges respectively. None of these graphs belongs to $\mathcal{S}$. This finishes the proof of Claim 8.3 (and thus of Lemma 8.2). □

We have the following strengthening of Lemma 8.1.

Lemma 8.4. *The only maximiser of* $\lambda(K_2 \cup K_2(\cdot))$ *is the vector* $(1/4, 1/4, 1/4, 1/4)$.

Proof. Let $\mathbf{a} = (a_1, a_2, a_3, a_4)$ in $\mathbb{S}_4$ and set $B = K_2 \cup K_2$. As we have already mentioned

$$\lambda(B(\mathbf{a})) = 12(a_1a_2 + a_3a_4)(a_1 + a_2)(a_3 + a_4).$$

To prove that the only maximiser of $\lambda(B(\cdot))$ is the vector $(1/4, 1/4, 1/4, 1/4)$, we show, equivalently, that the map

$$f(x_1, x_2, x_3) = 12(x_1x_2 + x_3(1 - x_1 - x_2 - x_3))(x_1 + x_2)(1 - x_1 - x_2)$$

with (x_1, x_2, x_3) in $D = \{(x_1, x_2, x_3) : x_1, x_2, x_3 \geqslant 0 \text{ and } x_1 + x_2 + x_3 \leqslant 1\}$ admits the vector $(1/4, 1/4, 1/4)$ as the unique maximiser. By Lemma 8.1, no maximiser of f is on the boundary of D. Thus, we are interested in the points belonging to the interior of D, where all the partial derivatives of f vanish. Hence, the following equations should be satisfied.

$$\frac{\partial f}{\partial x_1} = 12\big((x_2 - x_3)(x_1 + x_2)(1 - x_1 - x_2) + (x_1x_2 + x_3(1 - x_1 - x_2 - x_3))(1 - 2x_1 - 2x_2)\big) = 0, \tag{41}$$

$$\frac{\partial f}{\partial x_2} = 12\big((x_1 - x_3)(x_1 + x_2)(1 - x_1 - x_2) + (x_1x_2 + x_3(1 - x_1 - x_2 - x_3))(1 - 2x_1 - 2x_2)\big) = 0 \tag{42}$$

and

$$\frac{\partial f}{\partial x_3} = 12(1 - x_1 - x_2 - 2x_3)(x_1 + x_2)(1 - x_1 - x_2) = 0. \tag{43}$$

By (43) and recalling that we are only interested in points that belong to the interior of D, we get that $x_3 = 1 - x_1 - x_2 - x_3$, while by subtracting (41) from (42), we get that $x_1 = x_2$. Combining these two we get, in particular, that $x_1 + x_3 = 1/2$. Plugging the last three equalities into (41), we have that

$$0 = 12\big((2x_1 - 1/2)2x_1(1 - 2x_1) + (x_1^2 + (1/2 - x_1)^2)(1 - 4x_1)\big) = 3(1 - 4x_1)^3.$$

Thus $x_1 = 1/4$ and the result follows readily. □

Lemma 8.5. *The graph $K_2 \cup K_2$ is λ-flip-averse.*

Proof. Set $B = K_2 \cup K_2$ and let $B' = B(V_1, V_2, V_3, V_4)$ be a blow-up of B on n vertices, with $|V_i| = n/4 + O(1)$ for all $i = 1, 2, 3, 4$. Let $i, j \in [4]$, $x \in V_i$ and $y \in V_j$, with $x \neq y$. It suffices to distinguish the following three cases.

If ij forms an edge in B, then the number of F-subgraphs in B' (resp. $B' \oplus xy$) that use the pair xy is $n^2/4 + O(n)$ (resp. $n^2/16 + O(1)$) and we have that

$$\Lambda(B') - \Lambda(B' \oplus xy) = n^2/4 - n^2/16 + O(n) = 3n^2/16 + O(n). \tag{44}$$

If $i \neq j$ and ij do not form an edge, then

$$\Lambda(B') - \Lambda(B' \oplus xy) = 3n^2/16 - n^2/8 + O(n) = n^2/16 + O(n). \tag{45}$$

If $i = j$, then $B' \oplus xy$ has no copies via xy and

$$\Lambda(B') - \Lambda(B' \oplus xy) = n^2/8 + O(n). \tag{46}$$

By (46), (44) and (45), the result follows. □

Lemma 8.6. *The graph $K_2 \cup K_2$ is λ-strict.*

Proof. We set $\mathbf{a} = (1/4, 1/4, 1/4, 1/4)$ and $B = K_2 \cup K_2$. By Lemma 8.4, we have that $\mathbf{a}$ is the unique maximiser of $\lambda(B(\cdot))$. Thus it suffices to check that B is $(\lambda, \mathbf{a})$-strict. Indeed, let us fix some $\mathbf{y}$ in $[0, 1]^4$ that maximises $R_{\mathbf{a}}(\cdot)$. We will show that $\mathbf{y}$ has exactly one non-zero entry which is equal to 1.

Let B' be a balanced blow-up of B of order n. Let us denote by G the graph obtained by attaching to B' a new node w not belonging to $V(B')$ with adjacencies governed by $\mathbf{y}$. In particular, if $y_i = 0$ (resp. $y_i = 1$) for some $i \in [4]$, then w is attached to no vertex (resp. to all vertices) in V_i. We also define $\mathcal{H}$ to be the set of all graphs H on 5 vertices satisfying $H \sim G[X]$ for some $X \in \binom{V(G)}{5}$ with $w \in X$. We have the following claim.

Claim 8.7. *$\mathcal{H} \subseteq \mathcal{S}$.*

Proof of Claim 8.7. Let ε be a positive real. We denote by G' the graph obtained by adding εn twins of w in G. Set $V = V(B')$ and $V' = V(G') \setminus V$. Let $\mathcal{A}_0 = \binom{V}{5}$, $\mathcal{A}_1$

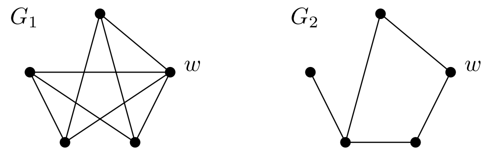

Fig. 4. The graphs G_1 and G_2 used in the proof of Lemma 8.6.

to be the set of all X in $\binom{V(G')}{5}$ having exactly one element in V' and $\mathcal{A}_2$ the set of all $X \in \binom{V(G')}{5}$ having at least two elements in V'. Since $\mathbf{y}$ maximises $R_{\mathbf{a}}(\cdot)$, by Lemma 5.5, we have that $R_{\mathbf{a}}(\mathbf{y}) = \lambda(\mathcal{G})$. Therefore,

$$\begin{aligned}\lambda(G') &= \binom{n+\varepsilon n}{k}^{-1} \sum_{X\in\binom{V(G')}{k}} \lambda(G'[X]) \\ &= \frac{|\mathcal{A}_0|}{\binom{n+\varepsilon n}{k}} \sum_{X\in\mathcal{A}_0} \lambda(G'[X]) + \frac{|\mathcal{A}_1|}{\binom{n+\varepsilon n}{k}} \sum_{X\in\mathcal{A}_1} \lambda(G'[X]) + \frac{|\mathcal{A}_2|}{\binom{n+\varepsilon n}{k}} \sum_{X\in\mathcal{A}_2} \lambda(G'[X]) \\ &\geqslant (1-k\varepsilon)\lambda(\mathcal{G}) + k\varepsilon(1-k\varepsilon)\lambda(\mathcal{G}) + O(\varepsilon^2) + O(1/n)\end{aligned} \tag{47}$$

and hence we get

$$\lambda(\mathcal{G}) - \lambda(G') \leqslant O(\varepsilon^2) + O(1/n). \tag{48}$$

By (14), there exists some positive real η independent from ε and n satisfying

$$\lambda(\mathcal{G}) - \lambda(G') \geqslant \eta p(H, G') + O(1/n) \tag{49}$$

for all $H \in \mathcal{NS}$. Finally, observe that for every $H \in \mathcal{H}$ we have that $p(H, G') = \Omega(\varepsilon)$. Thus by (48), (49) and a choice of a sufficiently small ε, $\mathcal{H}$ and $\mathcal{NS}$ are disjoint. □

Since $R_{\mathbf{a}}(0,0,0,0) = 3/16$, $\mathbf{y}$ has at least one non-zero coordinate. Next, let us observe that $\mathbf{y}$ cannot have two positive coordinates corresponding to adjacent nodes of B. Indeed, assuming the contrary, we pick four nodes in $V(B')$ adjacent to w and inducing a balanced complete bipartite graph. Together with w, they induce the graph G_1 in Fig. 4, which does not belong to $\mathcal{S}$, though it belongs to $\mathcal{H}$ by definition of $\mathcal{H}$, contradicting Claim 8.7. Hence, $\mathbf{y}$ has either one positive coordinate, or two positive coordinates that correspond to non-adjacent nodes i, j in B. Assuming that the second case occurs, picking two nodes adjacent to w from V_i, one node adjacent to w from V_j and one node non-adjacent to w from $V_{i'}$, where i' is the element of $V(B)$ adjacent to i in B, we have that these nodes together with w induce the graph G_2 from Fig. 4 that once again does not belong to $\mathcal{S}$ though it belongs to $\mathcal{H}$, contradicting Claim 8.7. Therefore, $\mathbf{y}$ has exactly one positive coordinate. Finally, we observe that the non-zero coordinate of $\mathbf{y}$ is equal to 1. Indeed, let us assume the contrary and let y_i be the non-zero coordinate of $\mathbf{y}$. Also let i' be the adjacent node to i in B. Picking two nodes in V_i adjacent to w, a non-adjacent node to w in V_i and a node in $V_{i'}$, together with w we induce the graph G_2 in Fig. 4 that belongs to $\mathcal{H}$ and not to $\mathcal{S}$ contradicting Claim 8.7. □

By Lemmas 8.2, 8.5 and 8.6, the assumptions of Theorem 5.8 are satisfied and therefore the problem is perfectly $K_2 \cup K_2$-stable.

9. Turán problem

This section is devoted to the proof of Theorem 1.11. As we have already mentioned Part 1 of Theorem 1.11 is known (see [40, Lemma 2.3]), we focus on the proof of the second part of Theorem 1.11. Let

$$v_0 := \max\{v(H') : H' \in \mathcal{H}\}, \tag{50}$$

which is finite as $\mathcal{H}$ was assumed to be finite. Recall that we defined $\mathcal{H}^{\uparrow}$ to be the collection of graphs obtained by adding missing edges to the graphs in $\mathcal{H}$. Before we start the proof, we provide an equivalent reformulation of the property in Part 2 of Theorem 1.11. As in the theorem, let $m = \min\{\chi(H) : H \in \mathcal{H}\} - 1$. Then the following two statements are equivalent.

(i) There is a constant D such that for every q if we add at least Dq edges into a part of K_m^q then the obtained graph is not $\mathcal{H}^{\uparrow}$-free.
(ii) There is a forest W such that the graph obtained from $K_m^{q_0}$, $q_0 := v(W)$, by adding W into one part is not $\mathcal{H}^{\uparrow}$-free.

Let us first assume (i) and prove (ii). Let Z be a graph with minimum degree at least $2D$ and girth strictly greater than v_0. Set $q := v(Z)$. Also, let $V_1, ..., V_m$ be disjoint sets, each of cardinality q, and let G to be the graph obtained from $K_m(V_1, ..., V_m)$ by adding a copy of Z in V_1. Since Z is of minimum degree $2D$, we have that Z contains at least Dq edges. By (i), G has a (not necessarily induced) subgraph $H \in \mathcal{H}$. Since $v(H) \leqslant v_0$, we conclude that $H[V_1]$ contains no cycle and thus H is as desired.

Assuming (ii), we claim that (i) holds with $D := q_0$. This is a consequence of the well-known fact that if G is a graph with $|E(G)| > (q_0 - 1)v(G)$, then G contains a copy of the forest W (not necessarily as an induced subgraph). Indeed, by e.g. [5, Theorem 2.5], G contains a non-empty subgraph G' of minimum degree at least q_0 where the required copy of W can be easily found.

Moreover, we will need the following result, which follows by the Ramsey Theorem [37] and elementary probabilistic estimates (see e.g. [8, Lemma 2.7] for a proof).

Lemma 9.1. *Let ε, θ be reals with $0 < \theta < \varepsilon$ and ℓ_1, ℓ_2 be positive integers with $\ell_1 < \ell_2$. Then there exists a positive integer ℓ_3 with the following property. For every probability space (Ω, Σ, μ) and every sequence $(A_j)_{j=1}^{\ell_3}$ such that $\mu(A_j) \geqslant \varepsilon$ for all $j \in [\ell_3]$, we have* *that there exists a subset L of $[\ell_3]$ of cardinality ℓ_2 such that for every subset K of L of cardinality ℓ_1 we have that*

$$\mu\Big(\bigcap_{j\in K} A_j\Big) \geqslant \theta^{\ell_1}. \quad \square$$

An iterated use of the above lemma yields the following, which we will use in the proof of the second part of Theorem 1.11.

Lemma 9.2. *Let ε, θ be reals with $0 < \theta < \varepsilon$ and q, ℓ be positive integers. Then there exists a positive integer $k = k(q, \ell, \theta, \varepsilon)$ with the following property. Let $(\Omega_1, \Sigma_1, \mu_1), ..., (\Omega_q, \Sigma_q, \mu_q)$ be probability spaces and for each $i \in [q]$ let $(A^i_j)_{j=1}^k$ be a sequence in Σ_i such that $\mu_i(A^i_j) \geqslant \varepsilon$ for all $j \in [k]$. Then there exists a subset L of $[k]$ of cardinality ℓ such that for every $i \in [q]$ we have that*

$$\mu_i\Big(\bigcap_{j\in L} A^i_j\Big) \geqslant \theta^{\ell}. \quad \square$$

Proof of Part 2 of Theorem 1.11. Recall that the theorem of Erdős [10] and Simonovits [41] states that the problem is classically stable with $B = K_m$. Thus the only twin-free graph B, which can have the property that the problem is robustly B-stable, is K_m. Let $t_m(n)$ be the maximum size of a K_m-blow-up of order n; it is easy to see that the maximum is attained if and only if any two part sizes differ at most by 1.

According to the discussion in the beginning of this section, it suffices to prove equivalence between robust K_m-stability and Condition (i) stated above. If Condition (i) fails, then for each D we can construct an $\mathcal{H}^{\uparrow}$-free graph G_D by adding Dq edges to K_m^q for some $q = q(D)$. This graph G_D of order $n := mq$ exceeds $t_m(n)$, the maximum size of a K_m-blow-up on n vertices, by Dn/m. Thus problem is not robustly K_m-stable.

Let us show the converse direction. Let D satisfy Condition (i) and define v_0 by (50). Given $\mathcal{H}$ and D, we choose positive constants in this order

$$c \gg c_3 \gg c_2 \gg c_1 \gg c_0,$$

each being sufficiently small depending on the previous ones. Assume on the contrary that the problem is not robustly K_m-stable. Hence, there exists an $\mathcal{H}^{\uparrow}$-free graph G with $n \geqslant 1/c_0$ vertices satisfying

$$t_m(n) - e(G) + n < c_1 \Delta_{\mathrm{edit}}(G, K_m()). \tag{51}$$

Let $V_1, ..., V_m$ be a max-cut partition of $V(G)$ and set $T := K_m(V_1, ..., V_m)$. Since $e(G) \geqslant t_m(n) - 2c_1\binom{n}{2}$, we have by the Erdős–Simonovits Stability Theorem [10,41] that

$$|E(G) \triangle E(T)| \leqslant c_2 \binom{n}{2}. \tag{52}$$

It routinely follows that

$$(1/m - c_3)n \leqslant |V_i| \leqslant (1/m + c_3)n, \quad \text{for all } i \in [m]. \tag{53}$$

Next we observe that in each $G[V_i]$ there are only a few vertices of high degree. More precisely, we have the following claim.

Claim 9.3. *For every $i \in [m]$, the induced subgraph $G[V_i]$ has at most $k(m, 2D, 2mc/3, 3mc/4)$ vertices of degree at least cn, where $k()$ satisfies Lemma 9.2.*

Proof. We set $k = k(m, 2D, 2mc/3, 3mc/4)$ and assume on the contrary that there exist $i_0 \in [m]$ and $x_1, ..., x_k \in V_{i_0}$ such that the degree of each x_j in $G[V_{i_0}]$ is at least cn. By the max-cut property of $V_1, ..., V_m$ we have for each $i \in [m]$ and $j \in [k]$ that the set of all neighbours of x_j in V_i, which we denote by A^i_j, is of cardinality at least cn and therefore, by (53), of uniform density at least $\frac{c}{1/m+c_3} \geqslant 3mc/4$. Applying Lemma 9.2, we obtain a subset L of $[k]$ of cardinality $2D$ such that for each $i \in [m]$, setting $B_i := \bigcap_{j\in L} A^i_j$ (which is the set of vertices in V_i that are G-adjacent to x_j for all $j \in L$), we have that

$$|B_i| \geqslant (2mc/3)^{2D}\,|V_i| \geqslant (c^{2D}/2^{2D})n.$$

We pick arbitrary subsets $Y_1, ..., Y_m$ of $B_1, ..., B_m$ respectively, of cardinality $(c^{2D}/2^{2D})n$ each. We set $Y := Y_1 \cup ... \cup Y_m$ and $Z := \{x_j : j \in L\}$. Observe that $G[Y]$ cannot contain a copy (not necessarily induced) of K^{4D}_m. Indeed, assume on the contrary that there exist pairwise disjoint $4D$-subsets $W_1, ..., W_m$ of $Y_1, \ldots, Y_m$, respectively, such that $E(G) \supseteq E(K_m(W_1, ..., W_m))$. Let W_1' be the set obtained by deleting $2D$ vertices from W_1 and adding the set Z. Then $E(G)$ is a superset of $E(G[W_1']) \cup E(K_m(W_1', W_2, ..., W_m))$. Observing that $G[W_1']$ contains at least $2D \cdot 2D = D \cdot |W_1'|$ edges, we get that G is not $\mathcal{H}^{\uparrow}$-free, a contradiction.

Thus, for every choice of a $4D$-subset W_j of Y_j, for $j = 1, ..., m$, there should be at least one missing edge (that is, an edge of T but not of G). Notice that there are $\binom{c^{2D}n/2^{2D}}{4D}^m$ choices of $(W_1, \ldots, W_m)$. On the other hand, a missing edge can be overcounted at most

$$\binom{c^{2D}n/2^{2D}-1}{4D-1}^2 \binom{c^{2D}n/2^{2D}}{4D}^{m-2} = \frac{(4D)^2}{(c^{2D}n/2^{2D})^2} \binom{c^{2D}n/2^{2D}}{4D}^m \tag{54}$$

times. Thus $E(T) \setminus E(G)$ is of cardinality at least $(c^{4D}/2^{4+4D}D^2)n^2$ contradicting (52). □

We set $K := m \cdot k(m, 2D, 2mc/3, 3mc/4)$. Let U' be the set all vertices having at least cn neighbours within their part. By Claim 9.3 we have that

$$|U'| \leqslant K \leqslant c_2 n. \tag{55}$$

We also set U'' to be the set of all vertices x in $V(G) \setminus U'$ such that $d_T(x) - d_G(x) \geqslant cn$. By (52), we get that

$$|U''| \leqslant \frac{c_2}{c}\, n \leqslant (c_3 - c_2)n. \tag{56}$$

Thus, setting $\mathcal{E}''$ to be the set of all pairs e of vertices in $V(G)$ satisfying $e \cap U'' \neq \emptyset$, we have that

$$|E(T) \cap \mathcal{E}''| - |E(G) \cap \mathcal{E}''| \geqslant |U''|\, cn - \binom{|U''|}{2} \overset{(56)}{\geqslant} |U''|(c - c_3)n \geqslant \frac{c}{2}\,|U''|\, n. \tag{57}$$

Moreover, setting $U := U' \cup U''$, by (55) and (56), we have that $|U| \leqslant c_3 n$. Also, set $V' := V \setminus U$ and $V_i' := V_i \setminus U$ for each $i \in [m]$. We have the following claim.

Claim 9.4. *Let $i \in [m]$ and X be a subset of $V' \setminus V_i$ with at most v_0 elements. Then V_i' has at least $(1 - 3cmv_0)|V_i'|$ vertices G-adjacent to every node in X.*

Proof. For every $x \in V' \setminus V_i$ we have the following. Let j be the unique element of $[m]$ satisfying $x \in V_j$. Since $x \notin U''$, we have that $d_G(x) \geqslant d_T(x) - cn$. Invoking the fact that x has at most cn G-neighbours in V_j, since $c \notin U'$, and x is T-adjacent to all vertices in V_i', it follows that x is G-adjacent to all but at most $2cn$ vertices in V_i'. By (53), (55) and (56), we get that $n/m \leqslant |V_i'|/(1 - 2c_3 m)$ and therefore x is G-adjacent to at least

$$|V_i'| - 2cm\frac{n}{m} \leqslant \left(1 - \frac{2cm}{1 - 2c_3 m}\right)|V_i'| \leqslant (1 - 3cm)|V_i'|$$

vertices in V_i'. Since X has at most v_0 elements, the claim follows. □

Next, we observe that in each V_j' we have a few edges. Namely, we have the following.

Claim 9.5. *For every $j \in [m]$, we have that $G[V_j']$ contains less than $(5D/2m)n$ edges.*

Proof. We assume on the contrary that there is some $j \in [m]$ such that $G[V_j']$ contains at least $5D(n/2m)$ edges. Without loss of generality, let $j = 1$. Let X_1 be a random subset of V_1' of size $n/2m$. Then the expected number of edges in $G[X_1]$ is at least

$$5D(n/2m)\binom{|V_1'| - 2}{n/2m - 2}\binom{|V_1'|}{n/2m}^{-1} \overset{(53)}{>} D \cdot n/2m.$$

We pick X_1 so that $G[X_1]$ has at least $D \cdot n/2m$ edges and, for each $i \in \{1, \dots, m\}$, we pick an arbitrary $(n/2m)$-subset X_i of V_i'. We set F to be the graph obtained by adding to $K_m(X_1, \dots, X_m)$ the edges of $G[X_1]$. By the choice of D there is an injective homomorphism f, that is, an injective map sending edges to edges, from some $H \in \mathcal{H}$ into F. We will arrive to a contradiction by constructing an injective homomorphism

f' from H into G. To this end, we set $Y_i := \{h \in H : f(h) \in X_i\}$ for each $i \in [m]$. We inductively define f' on each Y_i. For every $h \in Y_1$ we set $f'(h) = f(h)$. Then, for each $i = 2, ..., m$, assuming that f' has been defined on $\bigcup_{j=1}^{i-1} Y_j$, we extend f' on Y_i by using arbitrary elements of V_i' which are adjacent in G to every element of $f'(\cup_{j=1}^{i-1} Y_j)$. Claim 9.4 guarantees that such a selection is feasible. It follows easily that f' is indeed an injective homomorphism. Thus its image contains a (not necessarily induced) copy of H, which contradicts that G is $\mathcal{H}^{\uparrow}$-free. $\square$

We have

$$\begin{aligned}
t_m(n) - e(G) &\geqslant e(T) - e(G) \\
&\geqslant e(T[V']) - e(G[V']) + |E(T) \cap \mathcal{E}''| - |E(G) \cap \mathcal{E}''| - |U'|\,n \\
&\overset{(55),\,(57)}{\geqslant} e(T[V']) - e(G[V']) + \frac{c}{2}\,|U''|\,n - Kn \\
&\overset{\text{Claim 9.5}}{\geqslant} |E(T[V']) \triangle E(G[V'])| + \frac{c}{2}\,|U''|\,n - (5D + K)n \\
&\overset{(55)}{\geqslant} |E(T[V']) \triangle E(G[V'])| + |U'|\,n + \frac{c}{2}\,|U''|\,n - (5D + 2K)n \\
&\geqslant \frac{c}{2}\left(|E(T[V']) \triangle E(G[V'])| + |U|\,n\right) - (5D + 2K)n \\
&\geqslant \frac{c}{2}\,|E(T) \triangle E(G)| - (5D + 2K)n \\
&\geqslant \frac{c}{2}\,\Delta_{\text{edit}}(G, K_m()) - (5D + 2K)n.
\end{aligned}$$

By combining this with (51), we get

$$\frac{1}{c_1}\Big(t_m(n) - e(G) + n\Big) \leqslant \Delta_{\text{edit}}(G, K_m()) \leqslant \frac{2}{c}\big(t_m(n) - e(G) + (5D + 2K)n\big),$$

which is a contradiction since we have assumed that $c_1 \ll c$. $\square$

10. Concluding remarks

Theorem 1.11 implies that the three notions of stability introduced in Section 1 are non-equivalent. Indeed, let $K_{a,b,c}$ denote the complete 3-partite graph with part sizes a, b, c. Then the Turán problem $\mathrm{ex}(n, K_{2,2,2})$ is classically K_2-stable (by [10,41]) but not robustly K_2-stable by Part 2 of Theorem 1.11. (Namely, one can add a C_4-free bipartite graph of size $\Omega(q^{3/2})$ into one part of K_2^q, which will not violate the property of being $\{K_{2,2,2}\}^{\uparrow}$-free.) Also, $\mathrm{ex}(n, K_{2,2,1})$ is robustly but not perfectly K_2-stable by Theorem 1.11.

Theoretically, one should be able to write a computer code that takes as input only a family $\mathcal{F}$ of twin-free graphs and Λ and then tries to figure out everything else (namely

B, $\boldsymbol{a}$, N, and $\mathcal{C}$) automatically. For lower bounds, computer can enumerate all small B such that $B()$ is $\mathcal{F}$-free and then use Gröbner bases calculations to calculate $\lambda(B())$, thus identifying best possible B. For upper bounds, computer may start with largest feasible N (which is 8 for graphs unless $\mathcal{G}$ is rather structured), outputting some floating-point number c as an upper bound. Furthermore, if c seems to coincide with $\lambda(B())$, then the steps of finding smallest N that works and rounding the solution (using B as conjectured extremal configuration) could be also automated. However, the human intuition (based on various heuristics, symmetries, structure of admissible graphs, etc) is usually superior to the brute force search for plausible extremal configurations. Of course, the more powerful combination would be when computer search is restricted to a narrow set of plausible examples suggested by the user. It would be interesting to write a computer code that has this wider functionality and yet requires little coding from the user.

If the maximiser $\boldsymbol{a}$ of $\lambda(B())$ is unique (up to symmetry), one may be tempted to define another version of stability where one wishes to relate $\Lambda(G)-\Lambda(n,\mathcal{G})$ to the distance from G to $B(V_1,\dots,V_m)$ with $|V_i| = a_i/n + O(1)$. However, here the dependence is in general worse. For example, consider the Turán problem for triangle which is perfectly K_2-stable. Here the optimal $\boldsymbol{a}$ is unique: $(1/2, 1/2)$. However, for $G := K_{(1/2-\varepsilon)n,(1/2+\varepsilon)n}$ we have that $\lambda(n,\mathcal{G}) - \lambda(G) = O(\varepsilon^2)$, which is much smaller than $\delta_{\mathrm{edit}}(G, K_{n/2,n/2}) = \Omega(\varepsilon)$ when $\varepsilon \to 0$.

References

[1] N. Alon, E. Fischer, M. Krivelevich, M. Szegedy, Efficient testing of large graphs, Combinatorica 20 (2000) 451–476.
[2] J. Balogh, P. Hu, B. Lidický, F. Pfender, Maximum density of induced 5-cycle is achieved by an iterated blow-up of 5-cycle, European J. Combin. 52 (2016) 47–58.
[3] B. Bollobás, Y. Egawa, G. Jin, The maximal number of induced r-partite subgraphs, Graphs Combin. 11 (1995) 1–19.
[4] B. Bollobás, C. Nara, S. Tachibana, The maximal number of induced complete bipartite graphs, Discrete Math. 62 (1986) 271–275.
[5] J.A. Bondy, U.S.R. Murty, Graph Theory, Springer, Berlin, 2008.
[6] J.I. Brown, A. Sidorenko, The inducibility of complete bipartite graphs, J. Graph Theory 18 (1994) 629–645.
[7] S. Das, H. Huang, J. Ma, H. Naves, B. Sudakov, A problem of Erdős on the minimum number of k-cliques, J. Combin. Theory Ser. B 103 (2013) 344–373.
[8] P. Dodos, V. Kanellopoulos, K. Tyros, Measurable events indexed by products of trees, Combinatorica 34 (2014) 427–470.
[9] P. Erdős, On a theorem of Rademacher–Turán, Illinois J. Math. 6 (1962) 122–127.
[10] P. Erdős, On bipartite subgraphs of graphs (in Hungarian), Mat. Lapok 18 (1967) 283–288.
[11] P. Erdős, On some problems in graphs theory, combinatorial analysis, and combinatorial number theory, in: Graph Theory and Combinatorics, Cambridge, 1983, Academic Press, 1984, pp. 1–17.
[12] P. Erdős, M. Simonovits, Supersaturated graphs and hypergraphs, Combinatorica 3 (1983) 181–192.
[13] C. Even-Zohar, N. Linial, A note on the inducibility of 4-vertex graphs, J. Graph Theory 81 (2016) 109–119.
[14] G. Exoo, Dense packings of induced subgraphs, Ars Combin. 22 (1986) 5–10.
[15] Z. Füredi, A proof of the stability of extremal graphs, Simonovits' stability from Szemerédi's regularity, J. Combin. Theory Ser. B 115 (2015) 66–71.
[16] M.C. Golumbic, N. Pippenger, The inducibility of graphs, J. Combin. Theory Ser. B 19 (1975) 189–203.
[17] A.W. Goodman, On sets of acquaintances and strangers at any party, Amer. Math. Monthly 66 (1959) 778–783.

[18] A. Grzesik, On the maximum number of five-cycles in a triangle-free graph, J. Combin. Theory Ser. B 102 (2012) 1061–1066.
[19] E. Győri, On the number of C_5's in a triangle-free graph, Combinatorica 9 (1989) 101–102.
[20] H. Hatami, J. Hirst, S. Norine, The inducibility of blow-up graphs, J. Combin. Theory Ser. B 109 (2014) 196–212.
[21] H. Hatami, J. Hladký, D. Král', S. Norine, A. Razborov, On the number of pentagons in triangle-free graphs, J. Combin. Theory Ser. A 120 (2013) 722–732.
[22] H. Hatami, S. Norine, Undecidability of linear inequalities in graph homomorphism densities, J. Amer. Math. Soc. 24 (2011) 547–565.
[23] D. Hefetz, M. Tyomkyn, On the inducibility of cycles, available at arXiv:1702.07342, 2017.
[24] J. Hirst, The inducibility of graphs on four vertices, J. Graph Theory 75 (2014) 231–243.
[25] B. Lidický, F. Pfender, Pentagons in triangle-free graphs, available at arXiv:1712.08869, 2018.
[26] G. Lorden, Blue-empty chromatic graphs, Amer. Math. Monthly 69 (1962) 114–120.
[27] L. Lovász, Large Networks and Graph Limits, Colloquium Publications, Amer. Math. Soc., 2012.
[28] W. Mantel, Problem 28, Wiskd. Opgaven 10 (1907) 60–61.
[29] T.S. Michael, Cycles of length 5 in triangle-free graphs: a sporadic counterexample to a characterization of equality, preprint, 2011.
[30] T.S. Motzkin, E.G. Straus, Maxima for graphs and a new proof of a theorem of Turán, Canad. J. Math. 17 (1965) 533–540.
[31] V. Nikiforov, On the minimum number of k-cliques in graphs with restricted number of independence, Combin. Probab. Comput. 10 (2001) 361–366.
[32] V. Nikiforov, The minimum number of 4-cliques in a graph with triangle-free complement, available at arXiv:math/0501211, 2005.
[33] S. Norin, L. Yepremyan, Turán number of extensions, J. Combin. Theory Ser. A 155 (2018) 476–492.
[34] S. Norin, L. Yepremyan, Turán number of generalized triangles, J. Combin. Theory Ser. A 146 (2017) 312–343.
[35] O. Pikhurko, J. Sliačan, K. Tyros, Strong forms of stability from flag algebra calculations, available at arXiv:1706.02612, 2017.
[36] O. Pikhurko, E.R. Vaughan, Minimum number of k-cliques in graphs with bounded independence number, Combin. Probab. Comput. 22 (2013) 910–934.
[37] F.P. Ramsey, On a problem of formal logic, Proc. Lond. Math. Soc. 30 (1930) 264–286.
[38] A. Razborov, Flag algebras, J. Symbolic Logic 72 (2007) 1239–1282.
[39] A. Razborov, On 3-hypergraphs with forbidden 4-vertex configurations, SIAM J. Discrete Math. 24 (2010) 946–963.
[40] A. Roberts, A. Scott, Stability results for graphs with a critical edge, European J. Combin. 74 (2018) 27–38.
[41] M. Simonovits, A method for solving extremal problems in graph theory, stability problems, in: Theory of Graphs, Proc. Colloq., Tihany, 1966, Academic Press, 1968, pp. 279–319.
[42] J. Širáň, A new lower bound for the inducibility of a graph, Math. Slovaca 34 (1984) 365–370.
[43] P. Turán, On an extremal problem in graph theory (in Hungarian), Mat. Fiz. Lapok 48 (1941) 436–452.
[44] E.R. Vaughan, Flagmatic: a tool for researchers in extremal graph theory, Version 1.5, available at http://www.maths.qmul.ac.uk/~ev/flagmatic/, 2012.