

TIGHT BOUNDS ON THE COEFFICIENTS OF PARTITION FUNCTIONS VIA STABILITY

EWAN DAVIES, MATTHEW JENSSEN, WILL PERKINS, AND BARNABY ROBERTS

ABSTRACT. Partition functions arise in statistical physics and probability theory as the normalizing constant of Gibbs measures and in combinatorics and graph theory as graph polynomials. For instance the partition functions of the hard-core model and monomer-dimer model are the independence and matching polynomials respectively.

We show how stability results follow naturally from the recently developed occupancy method for maximizing and minimizing physical observables over classes of regular graphs, and then show these stability results can be used to obtain tight extremal bounds on the individual coefficients of the corresponding partition functions.

As applications, we prove new bounds on the number of independent sets and matchings of a given size in regular graphs. For large enough graphs and almost all sizes, the bounds are tight and confirm the Upper Matching Conjecture of Friedland, Krop, and Markström, and a conjecture of Kahn on independent sets, for a wide range of parameters. Additionally we prove tight bounds on the number of q -colorings of cubic graphs with a given number of monochromatic edges, and tight bounds on the number of independent sets of a given size in cubic graphs of girth at least 5.

1. INTRODUCTION

The matching polynomial (or matching generating function) of a graph G is the function

$$Z_G^{\text{match}}(\lambda) = \sum_{M \in \mathcal{M}(G)} \lambda^{|M|},$$

where the sum is over $\mathcal{M}(G)$, the set of all matchings in the graph G . Analogously, the independence polynomial of a graph G is

$$Z_G^{\text{ind}}(\lambda) = \sum_{I \in \mathcal{I}(G)} \lambda^{|I|},$$

where $\mathcal{I}(G)$ is the set of all independent sets of G . In statistical physics, $Z_G^{\text{match}}(\lambda)$ and $Z_G^{\text{ind}}(\lambda)$ are the partition functions of the monomer-dimer and hard-core models respectively.

The following theorems give a tight upper bound on $Z_G^{\text{match}}(\lambda)$ and $Z_G^{\text{ind}}(\lambda)$ over the family of all d -regular graphs.

Theorem 1 (Davies, Jenssen, Perkins, Roberts [9]). *For any d -regular graph G and any $\lambda > 0$,*

$$\frac{1}{|V(G)|} \log Z_G^{\text{match}}(\lambda) \leq \frac{1}{2d} \log Z_{K_{d,d}}^{\text{match}}(\lambda).$$

Date: May 14, 2018.

WP supported in part by EPSRC grant EP/P009913/1.

Theorem 2 (Kahn [21, 22], Galvin–Tetali [17], Zhao [32]). *For any d -regular graph G and any $\lambda > 0$,*

$$\frac{1}{|V(G)|} \log Z_G^{\text{ind}}(\lambda) \leq \frac{1}{2d} \log Z_{K_{d,d}}^{\text{ind}}(\lambda).$$

In particular, if we set $\lambda = 1$, then the two theorems say that if $2d$ divides n , then the total number of matchings and the number of independent sets in any d -regular graph on n vertices is at most that of the graph $H_{d,n}$ consisting of $n/2d$ copies of $K_{d,d}$. For much more on such extremal problems for regular graphs, see the notes of Galvin [16], survey of Zhao [33], and paper of Csikvari [5].

In this paper we will address two strengthenings of results of the form above and how they are related. The first possible strengthening of Theorems 1 and 2 is that $H_{d,n}$ might maximize the polynomials Z^{match} and Z^{ind} on the level of each individual coefficient. Let $m_k(G)$ and $i_k(G)$ denote the number of matchings and independent sets of size k respectively in a graph G . Then we can write $Z_G^{\text{match}}(\lambda) = \sum_{k \geq 0} m_k(G) \lambda^k$. Kahn and Friedland, Krop, and Markstrom conjectured the following.

Conjecture 3 (Kahn [21]). *Let $2d$ divide n . Then for any d -regular G on n vertices, and any $1 \leq k \leq n/2$,*

$$i_k(G) \leq i_k(H_{d,n}).$$

Conjecture 4 (Friedland, Krop, Markstrom [15]). *Let $2d$ divide n . Then for any d -regular G on n vertices, and any $1 \leq k \leq n/2$,*

$$m_k(G) \leq m_k(H_{d,n}).$$

These conjectures were known to be true in a very small number of cases ($k \leq 4$ and $k = n/2$). Several approximate versions of the two conjectures have been proved by Carroll, Galvin, and Tetali [3], Ilinca and Kahn [20], Perkins [28], and Davies, Jenssen, Perkins and Roberts [9]. The first three bounds were in general off by a factor exponential in n , while the fourth bound was off by a factor $2\sqrt{n}$. The case $k = n/2$ for matchings follows from Bregman’s theorem [2] on the permanents of 0/1 matrices with given row sums. (The case of independent sets in non-regular graphs with given minimum degree was proved by Gan, Loh, and Sudakov [18]).

Our first main result is that the conjectures of Kahn and Friedland, Krop, and Markstrom hold for a wide range of parameters.

Theorem 5. *For all $d \geq 2$ and $\varepsilon > 0$ there exists $N = N(d, \varepsilon)$ such that the following holds. Suppose that $n \geq N$ and n is divisible by $2d$. Let G be any d -regular graph on n vertices. Then for all $k > \varepsilon n$,*

$$\begin{aligned} i_k(G) &\leq i_k(H_{d,n}), \\ m_k(G) &\leq m_k(H_{d,n}). \end{aligned}$$

In other words, for sufficiently large graphs, and almost all values of k , Conjectures 3 and 4 hold.

The second strengthening of Theorems 1 and 2 we consider is the question of uniqueness and stability. The proofs of those theorems imply that equality is attained only by $K_{d,d}$ and unions of copies of $K_{d,d}$. The combinatorial notion of *stability* goes beyond this and states that any graph with a near extremal matching or independence polynomial must be

‘close’ to one of the extremal graphs under some natural distance on graphs. In extremal combinatorics, the theory of stability for Turán-type problems was developed by Erdős and Simonovits [13, 30]. Stability has also proved very useful in the extremal theory of dense graphs and graph limits [29], and in other areas of combinatorics (e.g. [14, 23, 26]).

Let $\delta_o(G, H)$ denote the sampling distance between two bounded degree graphs (see Section 3 below and Lovász’s book on graph limits [25]). This is a distance function that metrizes Benjamini–Schramm convergence [1]. Our next result is that Theorems 1 and 2 are stable under the $\delta_o(\cdot, \cdot)$ distance.

Theorem 6. *For any $d \geq 2$ there exist continuous functions $s^{\text{match}}(d, \lambda)$, $s^{\text{ind}}(d, \lambda)$ which satisfy $s^{\text{match}}(d, 0) = s^{\text{ind}}(d, 0) = 0$, and which are strictly increasing in λ such that the following holds. For any d -regular graph G ,*

$$\begin{aligned} \frac{1}{|V(G)|} \log Z_G^{\text{match}}(\lambda) &\leq \frac{1}{2d} \log Z_{K_{d,d}}^{\text{match}}(\lambda) - s^{\text{match}}(d, \lambda) \cdot \delta_o(G, K_{d,d}), \\ \frac{1}{|V(G)|} \log Z_G^{\text{ind}}(\lambda) &\leq \frac{1}{2d} \log Z_{K_{d,d}}^{\text{ind}}(\lambda) - s^{\text{ind}}(d, \lambda) \cdot \delta_o(G, K_{d,d}). \end{aligned}$$

Up to constant factors $\delta_o(G, K_{d,d})$ is simply the fraction of vertices of G that are not in a $K_{d,d}$ component, and so the reader may think of $\delta_o(G, K_{d,d})$ in this way, but in the general setting of extremal problems for bounded degree graphs, using a distance compatible with Benjamini–Schramm convergence is very natural (see e.g. [6, 24]) and so we stick with this notation.

A stability version of Theorem 2 for independent sets at $\lambda = 1$ was recently proved by Dmitriev and Dainiak [12] by examining the details of the entropy method proof. Our approach is more probabilistic and proceeds via the occupancy method of [9] and properties of linear programming. It also yields monotonicity of the functions $s^{\text{match}}(d, \lambda)$, $s^{\text{ind}}(d, \lambda)$ in λ which is helpful in proving the results below and does not obviously follow from entropy methods.

The method we use to prove extremality and stability via probabilistic observables and linear programming, and the method we use to deduce exact bounds on individual coefficients from stability versions of extremal results are both generally applicable. We provide two further examples to illustrate this.

Let HW be the Heawood graph, the unique $(3, 6)$ -cage graph (see Figure 1). For n divisible by 14, let HW_n be the graph consisting of a union of $n/14$ copies of HW . Perarnau and Perkins [27] proved that $\frac{1}{|V(G)|} \log Z_G^{\text{ind}}(\lambda)$ is maximized over all cubic graphs of girth at least 5 by HW . Here we use this and a corresponding stability result to prove tight upper bounds on the number of independent sets of a given size in such graphs.

Theorem 7. *For all $\varepsilon > 0$ there exists $N(\varepsilon)$ so that the following is true. Suppose that $n \geq N$ and n is divisible by 14. Let G be any 3-regular graph of girth at least 5 on n vertices. Then for all $k > \varepsilon n$,*

$$i_k(G) \leq i_k(HW_n).$$

Next we consider colorings of regular graphs. Galvin and Tetali [17] conjectured (and proved for bipartite graphs) that $H_{d,n}$ maximizes the number of proper q -colorings of any d -regular graph on n vertices. The general conjecture is still open, but the case $d = 3$ and arbitrary q was proved in [11]. To view this conjecture and these results in the setting




FIGURE 1. The Heawood graph, HW , the smallest 3-regular graph with girth 6.

of partition functions and graph polynomials, we consider the q -color Potts model from statistical physics. Let $c_k^q(G)$ be the number of q -colorings of the vertices of G with exactly k monochromatic edges. Then the function

$$Z_G^q(\lambda) = \sum_{k \geq 0} c_k^q(G) \lambda^k$$

is the Potts model partition function (the standard formulation has λ replaced by $e^{-\beta}$). The conjecture on proper colorings is that $H_{d,n}$ maximizes the coefficient $c_0^q(G)$ over n -vertex, d -regular graphs. Here we show that for cubic graphs, $H_{3,n}$ maximizes $c_k^q(G)$ for a wide range of k .

Theorem 8. *For all $q \geq 3$, $\varepsilon > 0$, there exists $N(q, \varepsilon)$ so that the following is true. Suppose that $n \geq N$ and n is divisible by 6. Let G be any 3-regular graph on n vertices. Then for $0 \leq k \leq (1 - \varepsilon)3n/2q$ we have*

$$c_k^q(G) \leq c_k^q(H_{3,n}).$$

The range of k is essentially optimal by the following observation. Note that the expected number of monochromatic edges in a uniformly random assignment of q colors to the vertices of a d -regular graph is $nd/2q$. For an n -vertex graph G the sum of all the coefficients c_k is equal to q^n , and so no single graph can maximize all of the coefficients. We expect that a union of cliques maximizes most coefficients c_k where $k \geq nd/2q$.

Moreover, while a result similar to Theorem 8 is undoubtedly true for the case $q = 2$, the exact statement is not true ($H_{3,n}$ has no 2-colorings with exactly 1 or 2 monochromatic edges) and so instead of dealing with unenlightening technicalities we omit this case.

Outline of the remainder of the paper.

- In Section 2 we define the related probabilistic models from statistical physics and present some of their probabilistic properties.
- In Section 3 we define our notion of stability for partition functions of bounded degree graphs and prove Theorem 6 for matchings.
- In Section 4 we prove our first transference result, showing how stability theorems imply exact results for coefficients. Here we prove Theorem 5 for matchings.
- In Section 5 we describe a hierarchy of different types of extremal results for partition functions and graph polynomials and present some open problems.
- In Section 6 we give the modifications of the proofs in Sections 3 and 4 necessary to prove Theorems 5, 6, and 7 for independent sets and Theorem 8 for colorings.

2. PROBABILISTIC PROPERTIES OF PARTITION FUNCTIONS

2.1. Occupancy fractions and observables. Each of the partition functions defined above comes from a model in statistical physics. The hard-core model on a graph G at fugacity λ is a random independent set \mathbf{I} chosen from G with probability

$$\mathbb{P}_{G,\lambda}(\mathbf{I} = I) = \frac{\lambda^{|I|}}{Z_G^{\text{ind}}(\lambda)}.$$

(We indicate the graph and value of the fugacity in the subscript of probabilities and expectations but drop it from the notation when they are clear from context). The partition function $Z_G^{\text{ind}}(\lambda)$ is the normalizing constant of the probability distribution – it ensures that the sum of $\mathbb{P}_{G,\lambda}(\mathbf{I} = I)$ over all independent sets is 1. In statistical physics terminology, the scaling $\frac{1}{|V(G)|} \log Z_G^{\text{ind}}(\lambda)$ is the free energy (or free energy density) and is a convenient scaling to allow us to compare partition functions on graphs with different numbers of vertices. In particular for any graph G the free energy density of G is the same as the free energy density of multiple copies of G .

Likewise the monomer-dimer model at fugacity λ is a random matching \mathbf{M} chosen with probability

$$\mathbb{P}_{G,\lambda}(\mathbf{M} = M) = \frac{\lambda^{|M|}}{Z_G^{\text{match}}(\lambda)},$$

and the Potts model is a random vertex-coloring χ (not necessarily proper) chosen with probability

$$\mathbb{P}_{G,\lambda}(\chi = \chi) = \frac{\lambda^{m(G,\chi)}}{Z_G^q(\lambda)},$$

where $m(G, \chi)$ is the number of monochromatic edges in G with coloring χ . If $\lambda > 1$, we say the model is ferromagnetic (monochromatic edges preferred); if $\lambda \in [0, 1)$ we say the model is anti-ferromagnetic (bichromatic edges preferred); and in the specific case $\lambda = 0$ (and G has at least one proper q -coloring), then χ is a uniformly chosen proper q -coloring of G .

Associated to these models are physical *observables*: expectations of real-valued functions of the random configurations. For instance, the occupancy fraction of the hard-core model is the expected fraction of vertices of the graph that appear in the random independent set:

$$\alpha_G^{\text{ind}}(\lambda) = \frac{1}{|V(G)|} \mathbb{E}_{G,\lambda} |\mathbf{I}|.$$

Likewise, the occupancy fraction of the monomer-dimer model is the expected fraction of edges that appear in the random matching:

$$\alpha_G^{\text{match}}(\lambda) = \frac{1}{|E(G)|} \mathbb{E}_{G,\lambda} |\mathbf{M}|,$$

and the internal energy of the Potts model is the expected fraction of monochromatic edges:

$$\beta_G^q(\lambda) = \frac{1}{|E(G)|} \mathbb{E}_{G,\lambda} |m(G, \chi)|.$$

For our purposes, the importance of these particular observables is the fact that they can be expressed in terms of derivatives of the corresponding log partition functions:

$$\begin{aligned}\alpha_G^{\text{ind}}(\lambda) &= \frac{1}{|V(G)|} \mathbb{E}_{G,\lambda} |\mathbf{I}| = \frac{1}{|V(G)|} \frac{\sum_{I \in \mathcal{I}(G)} |I| \lambda^{|I|}}{Z_G^{\text{ind}}(\lambda)} \\ &= \frac{1}{|V(G)|} \frac{\lambda}{Z_G^{\text{ind}}(\lambda)} \frac{\partial}{\partial \lambda} Z_G^{\text{ind}}(\lambda) \\ &= \frac{\lambda}{|V(G)|} \frac{\partial}{\partial \lambda} \log Z_G^{\text{ind}}(\lambda),\end{aligned}$$

and similarly,

$$\begin{aligned}\alpha_G^{\text{match}}(\lambda) &= \frac{\lambda}{|E(G)|} \frac{\partial}{\partial \lambda} \log Z_G^{\text{match}}(\lambda), \\ \beta_G^q(\lambda) &= \frac{\lambda}{|E(G)|} \frac{\partial}{\partial \lambda} \log Z_G^q(\lambda).\end{aligned}$$

The method of [9] gives the following, from which we derive Theorem 1 and an alternative proof of Theorem 2.

Theorem 9 (Davies, Jenssen, Perkins, Roberts [9]). *For any d -regular graph G , and any $\lambda > 0$,*

$$\alpha_G^{\text{ind}}(\lambda) \leq \alpha_{K_{d,d}}^{\text{ind}}(\lambda),$$

and

$$\alpha_G^{\text{match}}(\lambda) \leq \alpha_{K_{d,d}}^{\text{match}}(\lambda).$$

The above result is a strengthening of Theorems 1 and 2 for the partition function with the additional benefit of a natural probabilistic interpretation. The analogous result for independent sets in cubic graphs of girth at least 5 was proved in [27]. We also know that in the anti-ferromagnetic Potts model on cubic graphs, that $K_{3,3}$ minimizes the internal energy, which in turn implies that $K_{3,3}$ maximizes the free energy.

Theorem 10 (Davies, Jenssen, Perkins, Roberts [11]). *For any 3-regular graph G , any $q \geq 2$, and any $\lambda \in [0, 1]$,*

$$\beta_G^q(\lambda) \geq \beta_{K_{3,3}}^q(\lambda),$$

and so for every $\lambda \in [0, 1]$,

$$\frac{1}{|V(G)|} \log Z_G^q(\lambda) \leq \frac{1}{6} \log Z_{K_{3,3}}^q(\lambda).$$

2.2. The canonical ensemble. The random configurations \mathbf{I} , \mathbf{M} , and χ defined above come from the *grand canonical ensemble* in the language of statistical physics, and these models have several appealing mathematical properties, most notably the spatial Markov property. The individual coefficients of $Z_G^{\text{ind}}(\lambda)$, $Z_G^{\text{match}}(\lambda)$, $Z_G^q(\lambda)$, however, are associated directly to a different probabilistic model, the *canonical ensemble*. In this setting \mathbf{I}_k is a uniformly random independent set of size k drawn from G ; \mathbf{M}_k is a uniformly random matching of size k ; and χ_k is a uniformly random coloring of $V(G)$ with q colors that has exactly k monochromatic edges. The partition function of the canonical ensemble is simply $i_k(G)$, the number of independent sets of size k , or likewise $m_k(G)$ or $c_k^q(G)$.

The two ensembles are closely related; if we condition on the event that $|\mathbf{I}| = k$ in the grand canonical ensemble then we obtain the distribution \mathbf{I}_k (assuming of course that G has at least one independent set of size k). Moreover if we choose λ so that $\mathbb{E}_\lambda|\mathbf{I}| = k$, then we would expect \mathbf{I} and \mathbf{I}_k to behave quite similarly. In physics the two ensembles are considered to be essentially equivalent for most purposes; one way to understand this is by imagining running the canonical ensemble on a large lattice or otherwise periodic graph; if we then zoom in to a small block of the large graph, the induced distribution looks almost identical to the distribution of the grand canonical ensemble on the small block. We use a precise instantiation of this idea in our proofs in Section 4.

Theorems 1 and 2 can be viewed as tight upper bounds on partition functions from the grand canonical ensemble, whereas Conjectures 3 and 4 (and Theorem 5) discuss tight upper bounds on the ‘partition function’ of the canonical ensemble. The method we use to prove Theorem 5 builds upon the occupancy method of [9], which relies heavily on the spatial Markov property of the grand canonical ensemble. The canonical ensemble does not have this property, and so these techniques do not apply directly. As the authors showed in [9] with a proof of an asymptotic version of Conjecture 4, the similarities between these ensembles mean that bounds relating to the grand canonical ensemble can give approximate versions of the results of this paper. Here we go further, exploiting stability to obtain tight results. In Section 5 we consider general relations between extremal bounds for the two ensembles.

3. STABILITY

3.1. Sampling distance. There are many notions of distances between graphs. We use the notion of *sampling distance* which is closely related to the notion of Benjamini–Schramm convergence for sequences of bounded-degree graphs, see [1, 25]. Let G be a finite graph of maximum degree at most d . Let $\rho_r(G)$ be the probability distribution on rooted graphs of depth at most r induced by choosing a vertex v uniformly at random from $V(G)$ and taking the depth r ball around v . Then we can define the *sampling distance* of two graphs G, G' to be

$$\delta_o(G, G') = \sum_{r=1}^{\infty} 2^{-r} \|\rho_r(G) - \rho_r(G')\|_{TV},$$

where $\|\cdot\|_{TV}$ is the total variation distance of two probability distributions on the same ground set. A sequence of graphs G_n is Benjamini–Schramm convergent if it is a Cauchy sequence in the distance function $\delta_o(\cdot, \cdot)$.

A simple fact about $\delta_o(\cdot, \cdot)$ is that if H is any graph and H' consists of $k \geq 1$ copies of H , then $\delta_o(H, H') = 0$ (in particular, $\delta_o(\cdot, \cdot)$ is not a metric).

3.2. Stability results. Theorem 6 is a direct consequence of the methods of [9] which show that $K_{d,d}$ is the unique connected graph which maximizes the occupancy fraction. These methods have been used to give tight bounds on quantities analogous to the occupancy fraction in a variety of models and classes of graphs [4, 8, 11, 10, 27], reducing extremal graph theory problems to solving linear programs over the space of probability distributions on local configurations (for more detail see the aforementioned references and Section 3.3 below). In cases in which there is a unique extremal distribution arising from a unique connected graph, a stability statement analogous to Theorem 6 follows from complementary slackness and basic

calculus. If desired, an explicit form for the constants (e.g. $c^{\text{ind}}(d, \lambda)$, $c^{\text{match}}(d, \lambda)$ below) can be computed via this method, but we omit such calculations.

The statement for the occupancy fraction from which we begin exists implicitly in [9], but here we give more detail.

Proposition 11. *There exists continuous functions $c^{\text{ind}}(d, \lambda)$, $c^{\text{match}}(d, \lambda) \geq 0$ which are strictly positive when $\lambda > 0$ such that for any d -regular graph G ,*

$$\begin{aligned}\alpha_G^{\text{match}}(\lambda) &\leq \alpha_{K_{d,d}}^{\text{match}}(\lambda) - c^{\text{match}}(d, \lambda) \cdot \delta_o(G, K_{d,d}), \\ \alpha_G^{\text{ind}}(\lambda) &\leq \alpha_{K_{d,d}}^{\text{ind}}(\lambda) - c^{\text{ind}}(d, \lambda) \cdot \delta_o(G, K_{d,d}).\end{aligned}$$

We prove Proposition 11 for matchings in Section 3.3 and give the required adaptation to prove the statement for independent sets in Section 6. We give the short derivation of Theorem 6 from Proposition 11 here.

Proof of Theorem 6. Using Proposition 11 and simple calculus, we have

$$\begin{aligned}\frac{1}{|V(G)|} \log Z_G^{\text{match}}(\lambda) &= \frac{1}{|V(G)|} \log Z_G^{\text{match}}(0) + \frac{1}{|V(G)|} \int_0^\lambda (\log Z_G^{\text{match}}(t))' dt \\ &= \frac{d}{2} \cdot \int_0^\lambda \frac{\alpha_G^{\text{match}}(t)}{t} dt \\ &\leq \frac{d}{2} \cdot \int_0^\lambda \left(\frac{\alpha_{K_{d,d}}^{\text{match}}(t)}{t} - \frac{c^{\text{match}}(d, t) \cdot \delta_o(G, K_{d,d})}{t} \right) dt \\ &= \frac{1}{2d} \log Z_{K_{d,d}}^{\text{match}}(\lambda) - \delta_o(G, K_{d,d}) \cdot \frac{d}{2} \int_0^\lambda \frac{c^{\text{match}}(d, t)}{t} dt.\end{aligned}$$

Hence we may take $s^{\text{match}}(d, \lambda) = \frac{d}{2} \int_0^\lambda \frac{c^{\text{match}}(d, t)}{t} dt$, which proves the theorem since $c^{\text{match}}(d, t) > 0$ for all $t > 0$. The proof for independent sets is essentially the same. \square

As a corollary of Theorem 6, we obtain the following stability version of Bregman's theorem on perfect matchings applied to regular graphs.

Corollary 12. *There exists a constant $c(d)$ so that for any d -regular G ,*

$$\frac{1}{|V(G)|} \log m_{\text{perf}}(G) \leq \frac{1}{2d} \log m_{\text{perf}}(K_{d,d}) - c(d) \cdot \delta_o(G, K_{d,d}).$$

(An analogue holds for independent sets with an elementary proof).

Proof. We can write $m_{\text{perf}}(G)$ in terms of the monomer-dimer partition function:

$$\frac{1}{|V(G)|} \log m_{\text{perf}}(G) = \lim_{\lambda \rightarrow \infty} \frac{1}{|V(G)|} \log Z_G^{\text{match}}(\lambda) - \frac{\log \lambda}{2}.$$

(If G has no perfect matching then the right-hand side tends to $-\infty$). Now applying Theorem 6,

$$\begin{aligned} \frac{1}{|V(G)|} \log m_{\text{perf}}(G) &= \lim_{\lambda \rightarrow \infty} \frac{1}{|V(G)|} \log Z_G^{\text{match}}(\lambda) - \frac{\log \lambda}{2} \\ &\leq \lim_{\lambda \rightarrow \infty} \frac{1}{2d} \log Z_{K_{d,d}}^{\text{match}}(\lambda) - \frac{\log \lambda}{2} - s^{\text{match}}(d, \lambda) \cdot \delta_o(G, K_{d,d}) \\ &= \frac{1}{2d} \log m_{\text{perf}}(K_{d,d}) - c(d) \cdot \delta_o(G, K_{d,d}), \end{aligned}$$

where we take $c(d) = \lim_{\lambda \rightarrow \infty} s^{\text{match}}(d, \lambda)$. The limit exists because $s^{\text{match}}(d, \lambda)$ is strictly increasing in λ and is bounded (for if were not bounded, the above inequality for λ large would contradict the existence of a perfect matching in a graph G with $\delta_o(G, K_{d,d}) > 0$). \square

3.3. Proof of Proposition 11. We give the argument in detail for matchings; the case of independent sets and an analogue for colorings is discussed in Section 6. We begin with a brief summary of the proof from [9] of Theorem 9.

To prove Theorem 9 for matchings, let G be any n -vertex d -regular graph and consider the following experiment. We write $N(e)$ for the set of edges incident to an edge e in G . Sample a matching \mathbf{M} from the monomer-dimer model at fugacity λ on G , and sample an edge e uniformly at random from $E(G)$. Then record the *local view* L which consists of the subgraph of G given by the edge e and any incident edges which are *externally uncovered* in the sense that they are not incident to any edge in $\mathbf{M} \setminus (\{e\} \cup N(e))$. We write \mathcal{L}_d for the set of possible local views in d -regular graphs. A graph G and fugacity parameter λ induce a distribution p on \mathcal{L}_d such that each $p(L)$ is the probability that the two-part experiment above yields the local view L .

The occupancy fraction of G can be computed from its corresponding distribution p . In [9] we derived a series of linear constraints on the values $p(L)$, $L \in \mathcal{L}_d$, that hold for any d -regular G . We then relaxed the maximization problem from distributions arising from graphs to all probability distributions on \mathcal{L}_d satisfying these constraints and showed that the linear programming relaxation is tight: that is, the optimum of the linear program, α^* , equals $\alpha_{K_{d,d}}^{\text{match}}(\lambda)$.

This linear program can be expressed in standard form as

$$(1) \quad \alpha^* = \max a^T p \quad \text{subject to } Ap \leq b \text{ and } p \geq 0,$$

where A is a matrix with rows indexed by constraints and columns indexed by local views, a is a vector indexed by local views, and b is a vector indexed by constraints. We will use the fact that A , a , and b are independent of n .

The distribution p^* corresponding to $K_{d,d}$ is a feasible solution for this linear program, and so $\alpha^* \geq a^T p^* = \alpha_{K_{d,d}}^{\text{match}}(\lambda)$. To show that $\alpha^* = \alpha_{K_{d,d}}^{\text{match}}(\lambda)$ we use linear programming duality.

The dual program (with dual variables given by a vector q) is

$$\alpha^* = \min b^T q \quad \text{subject to } A^T q \geq a \text{ and } q \geq 0,$$

where each dual constraint (given by a row of A^T) corresponds to a local view L . If we can find a feasible primal solution q^* so that $b^T q^* = \alpha_{K_{d,d}}^{\text{match}}(\lambda)$ then we have shown that $\alpha^* \leq b^T q^* = \alpha_{K_{d,d}}^{\text{match}}(\lambda)$, from which Theorem 9 follows. In [9] we computed q^* by solving the

dual constraints corresponding to local views in the support of p^* to hold with equality. We then define the *slack* as a vector $s \in \mathbb{R}^{\mathcal{L}_d}$ such that

$$s_L = A^T q^* - a_L.$$

We define \mathcal{L}_d^* as the set of local views L so that $s_L = 0$. Complementary slackness tells us that any primal solution p' that achieves the optimal value α^* must have support contained in \mathcal{L}_d^* . To prove Proposition 11 we show quantitatively that the only graphs with distributions supported in \mathcal{L}_d^* are $K_{d,d}$ (and $H_{d,n}$).

Lemma 13. *Let G be a d -regular graph, $\lambda > 0$, and suppose that L is a random local view obtained from the monomer-dimer model G at fugacity λ as described above. Then there exists a function $f(d, \lambda) > 0$ such that*

$$\mathbb{P}(L \notin \mathcal{L}_d^*) \geq \delta_o(G, K_{d,d}) \cdot f(d, \lambda).$$

This lemma gives an additional constraint on the distribution of local views for graphs G with $\delta_o(G, K_{d,d}) \geq \delta$, namely that

$$(2) \quad x^T p \geq \delta f(d, \lambda),$$

where x is the vector in $\{0, 1\}^{\mathcal{L}_d}$ indicating membership of $\mathcal{L}_d \setminus \mathcal{L}_d^*$.

To obtain Proposition 11 we augment the primal program (1) with the constraint (2) and consider the new dual, which has an additional variable we denote θ ,

$$(3) \quad \alpha_{\text{new}}^* = \min b^T q - \delta f \theta \quad \text{subject to } A^T q - \theta x \geq a, q \geq 0, \text{ and } \theta \geq 0.$$

If G is a d -regular graph with $\delta_o(G, K_{d,d}) \geq \delta$, then $\alpha_G^{\text{match}}(\lambda)$ is feasible in the augmented primal, hence any q and θ which are feasible in the new dual yield an upper bound on $\alpha_G^{\text{match}}(\lambda)$. We claim that $q = q^*$ and $\theta = \theta^* = \min\{s_L : L \in \mathcal{L}_d \setminus \mathcal{L}_d^*\}$ are dual-feasible, which is easy to verify from the constraints in (3). This follows from the facts that x indicates membership of $L \in \mathcal{L}_d \setminus \mathcal{L}_d^*$ and that θ^* is chosen to minimize the slack s_L over this set. Then directly from (3) we obtain Proposition 11 for matchings with $c^{\text{match}}(d, \lambda) = f(d, \lambda)\theta^* > 0$. It remains to prove Lemma 13, for which we need a simple fact that is used throughout the rest of the paper. The fact follows easily from the spatial Markov property, but for completeness we give a direct proof that avoids the technical description of this property.

Fact. *Let G be a graph and let M be a random matching drawn from the monomer-dimer model on G at fugacity λ . Then for any edge e of G*

$$\mathbb{P}_{G, \lambda}(e \in M) \leq \frac{\lambda}{1 + \lambda}.$$

In particular, $\alpha_G^{\text{match}}(\lambda) \leq \frac{\lambda}{1 + \lambda}$.

Proof. Let \mathcal{M}_1 denote the set of matchings in G that contain e and let \mathcal{M}_0 denote the set of matchings that do not. Note that the function $f : \mathcal{M}_1 \rightarrow \mathcal{M}_0$ given by $M \mapsto M - e$ is injective and so

$$\mathbb{P}(e \in M) = \frac{\sum_{M \in \mathcal{M}_1} \lambda^{|M|}}{\sum_{M \in \mathcal{M}_0} \lambda^{|M|} + \sum_{M \in \mathcal{M}_1} \lambda^{|M|}} \leq \frac{\sum_{M \in \mathcal{M}_1} \lambda^{|M|}}{\sum_{M \in f(\mathcal{M}_1)} \lambda^{|M|} + \sum_{M \in \mathcal{M}_1} \lambda^{|M|}} = \frac{\lambda}{1 + \lambda}. \quad \square$$

We remark that the analogous statement for independent sets, $\mathbb{P}_{G, \lambda}(v \in I) \leq \frac{\lambda}{1 + \lambda}$ for any $v \in V(G)$, holds by the same proof.

Proof of Lemma 13. We refer to [25] for properties of the sampling distance $\delta_o(\cdot, \cdot)$. In particular, note that if $\delta_o(G, K_{d,d}) \geq \delta$ then there exists a radius $r \geq 2$ such that the variation distance between the distributions on depth- r neighborhoods obtained by uniform random sampling in G and $K_{d,d}$ is at least $\delta/2$. Since $r \geq 2$, the only such neighborhood which occurs in $K_{d,d}$ is $K_{d,d}$ itself, and we obtain that the probability of sampling a different neighborhood in G is at least δ . That is, the fraction of vertices of G which are not in a copy of $K_{d,d}$ is at least $\delta_o(G, K_{d,d})$. Since G is d -regular the same fact holds for edges.

Fix an arbitrary $e = u_1u_2 \in E(G')$ which is not in a copy of $K_{d,d}$ and consider the local view generated by choosing a random matching \mathbf{M} from the monomer-dimer model on G and selecting the edge e as the central edge. We have two cases.

If e is contained in a triangle $T = u_1u_2u_3$, then this triangle is present in the local view at e if and only if the $d - 2$ edges outside T which are incident to u_3 are not in \mathbf{M} . Let A_e be this event. In the monomer-dimer model on any graph, an edge is in the random matching with probability at most $\lambda/(1 + \lambda)$, so by successive conditioning $\mathbb{P}(A_e) \geq (1 + \lambda)^{-(d-2)}$.

If e is not contained in a triangle then, because e is not in a $K_{d,d}$, there must be a pair of distinct vertices u_3, u_4 in G such that

- (1) $e = u_1u_2$, $f = u_2u_3$, and $g = u_3u_4$ are edges of G ,
- (2) u_1u_4 is not an edge of G .

In this case, let A_e be the event that g is the only edge of G at distance 2 from e in the matching \mathbf{M} . When no edge incident to g and no edge incident to e are in \mathbf{M} , we have $g \in \mathbf{M}$ with probability $\lambda/(1 + \lambda)$. Since the number of edges incident to g or at distance ≤ 2 from e is $< 2d^2$ we have $\mathbb{P}(A_e) \geq \lambda(1 + \lambda)^{-2d^2}$.

The only local views which may arise in $K_{d,d}$ are triangle-free with equal numbers of edges on each side of the ‘central’ edge. Then in each of the cases above, the event A_e implies the local view centered on e is not in \mathcal{L}_d^* ; for the second case note that A_e implies f is not present but all other edges incident to e are.

Since e was arbitrary, and at least a fraction $\delta_o(G, K_{d,d})$ of the edges of G are not in a copy of $K_{d,d}$, with probability at least $\delta_o(G, K_{d,d}) \min \{(1 + \lambda)^{-(d-2)}, \lambda(1 + \lambda)^{-2d^2}\}$ the random local view L from G is not in \mathcal{L}_d^* . \square

4. FROM STABILITY TO COEFFICIENTS

To transfer stability results of the form of Theorem 6 to sharp bounds on individual coefficients (Theorems 5, 7, and 8) we use the probabilistic properties of the relevant statistical physics models. We use these properties to analyze the coefficients of the partition functions of the optimizing graphs, $H_{d,n}$ for Theorems 6 and 8 and HW_n for Theorem 7. In these cases the optimizing graph is the union of small isomorphic components. When we run, say, the monomer-dimer model on e.g. $H_{d,n}$, because of the spatial Markov property of the model, the sizes of the intersections of the random matching with each component are i.i.d. random variables bounded between 0 and d , and this allows us to use probability theory related to the sums of independent random variables. Note that in the canonical model, the sizes of intersections of the random matching with each component are not independent.

The following is a Local Central Limit Theorem due to Gnedenko [19] (see the particular statement below and exposition by Tao [31]).

Theorem 14 (Gnedenko's Local Limit Theorem). *Let X_1, \dots, X_n be i.i.d. integer valued random variables with mean μ , variance σ^2 , and suppose that the support of X_1 includes two consecutive integers. Let $S_n = X_1 + \dots + X_n$. Then*

$$\mathbb{P}(S_n = k) = \frac{1}{\sqrt{2\pi n\sigma}} \exp \left[-(k - n\mu)^2 / 2n\sigma^2 \right] + o(n^{-1/2}),$$

with the error term $o(n^{-1/2})$ uniform in k .

The condition on the support of X_1 rules out cases in which, say X_1 only takes even values and so the probability S_n is odd is 0.

We apply Theorem 14 to derive bounds on the relative size of nearby coefficients in the optimizing graphs.

Lemma 15. *For any $\delta, \varepsilon > 0$, $d, n_1 > 0$, there exists N large enough so that the following is true. Suppose $n > N$ and $\varepsilon n \leq k \leq (1 - \varepsilon)n/2$. Then there exists $\lambda \geq \frac{2\varepsilon}{d}$ so that for all $0 \leq r \leq n_1$,*

$$(1 - \delta)m_{k-r}(H_{d,n}) \leq \lambda^r \cdot m_k(H_{d,n}) \leq (1 + \delta)m_{k-r}(H_{d,n}).$$

Proof. Choose λ so that the expected size of the matching drawn from the monomer-dimer model on $H_{d,n}$ is k . By our choice of λ , $\frac{\lambda}{1+\lambda} \geq \alpha_{H_{d,n}}^{\text{match}}(\lambda) = \frac{2k}{nd} \geq \frac{2\varepsilon}{d}$ and so $\lambda \geq \frac{2\varepsilon}{d}$. Note that $\alpha_{H_{d,n}}^{\text{match}}(t) = \alpha_{K_{d,d}}^{\text{match}}(t)$ is strictly increasing in t with $\lim_{t \rightarrow \infty} \alpha_{K_{d,d}}^{\text{match}}(t) = \frac{1}{d}$. Since $k \leq (1 - \varepsilon)n/2$ (so that $\alpha_{K_{d,d}}^{\text{match}}(\lambda) \leq \frac{1-\varepsilon}{d}$) it follows that λ must be bounded from above by a function of d and ε (but not n).

Let $K = n/2d$ and let X be the size of the random matching drawn from $H_{d,n}$ at fugacity λ . Note that X is the sum of K i.i.d. random variables X_1, \dots, X_K , where X_1 is the size of a random matching chosen from $K_{d,d}$ according to the monomer-dimer model at fugacity λ and so takes the values $\{0, 1, \dots, d\}$ with positive probability. Let σ^2 denote the variance of X_1 and note that σ is independent of n . Now Theorem 14 gives:

$$\mathbb{P}(X = k - t) = \frac{1}{\sqrt{2\pi K\sigma}} e^{-t^2/2K\sigma^2} + o(K^{-1/2}),$$

where hidden constants in the error term depend only on d and ε since λ is bounded above and below in terms of d and ε .

For $0 \leq r \leq n_1$ and K large enough as a function of n_1, δ , we have

$$(1 - \delta)\mathbb{P}[X = k - r] \leq \mathbb{P}[X = k] \leq (1 + \delta)\mathbb{P}[X = k - r].$$

Since $\mathbb{P}(X = k - r) = \frac{m_{k-r}(H_{d,n})\lambda^{k-r}}{Z_{H_{d,n}}^{\text{match}}(\lambda)}$, we get

$$(1 - \delta)m_{k-r}(H_{d,n})\lambda^{-r} \leq m_k(H_{d,n}) \leq (1 + \delta)m_{k-r}(H_{d,n})\lambda^{-r},$$

and the result follows. \square

The same result holds for independent sets in both $H_{d,n}$ and HW_n . For colorings the result holds as well, and we state it here for completeness.

Lemma 16. *For any $\delta, \varepsilon > 0$, $d, q, n_1 > 0$, there exists n_0 large enough so that the following is true. Suppose $n > n_0$ and $\varepsilon n < k < (1 - \varepsilon)dn/2q$. Then there exists $\lambda \leq 1 - g(q, d, \varepsilon)$, so*

that for all $0 \leq r \leq n_1$,

$$(1 - \delta)c_{k-r}^q(H_{d,n}) \leq \lambda^r \cdot c_k^q(H_{d,n}) \leq (1 + \delta)c_{k-r}^q(H_{d,n}),$$

where $g(q, d, \varepsilon) > 0$.

The proof is identical to that of Lemma 15, we just need the fact that for $k < (1 - \varepsilon)dn/2q$, $\lambda \leq 1 - g(q, d, \varepsilon)$ for some $g(q, d, \varepsilon) > 0$. This holds since $\beta_G^q(1) = 1/q$ for all G and $\beta_{K_{d,d}}^q(\lambda)$ is strictly increasing in λ .

4.1. Matchings of a given size. Here we prove Theorem 5 for matchings.

Divide G into two parts, G' and H , where H is a union of $K_{d,d}$'s and G' contains no copy of $K_{d,d}$. Let $n_1 = |G'|$ and $n_2 = n - n_1$. Let $H' = H_{d,n_1}$.

Given $\varepsilon > 0$ and d , we choose auxiliary parameters with the following dependences:

- $\delta = \delta(d)$.
- $N_1 = N_1(d, \varepsilon)$.
- $\varepsilon' = \varepsilon'(N_1, \delta, d)$.
- $\delta' = \delta'(\varepsilon, d)$.
- $N = N(N_1, \varepsilon', \delta, \delta', d)$.

We break our argument into cases depending on whether G' is large or small. If G' is small the intuition is that when uniformly choosing a matching of size k from the canonical model on G , the distribution of the matching induced on G' looks as though it has come from the grand canonical model with a particular λ . This will allow us to use Theorem 6. If G' is large then we can use Theorem 6 directly in combination with the local limit theorem. We break the case where G' is small further depending on the size of k . If k is very large Lemma 15 does not apply and so we use a slightly different approach there.

Small-1: If n_1 is small ($n_1 \leq N_1$) and k is large (at least $(1 - \varepsilon')n/2$) we show that only the top coefficient of the matching polynomial of G' matters, and that there is a constant factor gap between the top coefficient of G' and H' (regardless of how small n_1 is).

Small-2: If n_1 is small and k is moderate, $\varepsilon n < k < (1 - \varepsilon')n/2$, we show that a random matching of size k in G approximately induces the monomer-dimer model on H with an appropriate λ , and thus we can simply compare the partition functions of G' and H' . This uses the local limit theorem.

Large: If n_1 is large ($n_1 > N_1$) then we use stability (Theorem 6) and the local limit theorem (Theorem 14).

Small-1: $n_1 \leq N_1$ and $k \geq (1 - \varepsilon')n/2$.

Let $\delta = \delta(d) > 0$ be such that for any G' (with $|G'| = n_1$ divisible by $2d$) containing no copies of $K_{d,d}$ we have

$$(4) \quad m_{n_1/2}(G') \leq (1 - \delta)m_{n_1/2}(H_{d,n_1}).$$

Such a δ exists by Corollary 12. We choose $\varepsilon' = \varepsilon'(\delta, d, N_1)$ small enough and $N = N(\varepsilon', N_1)$ large enough so that the following holds. For $n > N - N_1$, $k > (1 - \varepsilon')n/2$,

$$(5) \quad \frac{m_{k+1}(H_{d,n})}{m_k(H_{d,n})} < \delta \cdot 2^{-dN_1/2}.$$

To see this is possible, note that given a matching M of size $k + 1$ in a graph, we can obtain $k + 1$ matchings of size k by removing one of the edges of M . On the other hand given a matching M' of size k , there are at most $d(n - 2k)/2$ edges which can be added to M' to obtain a matching of size $k + 1$. It follows that $\frac{m_{k+1}(H_{d,n})}{m_k(H_{d,n})} \leq \frac{d(n-2k)}{2(k+1)} < 2d\varepsilon'$ and so choosing $\varepsilon' = \frac{\delta \cdot 2^{-dN_1/2}}{2d}$ will suffice. We then see that

$$\begin{aligned}
m_k(G) &= \sum_{r=0}^{n_1/2} m_r(G') m_{k-r}(H) \\
&\leq m_{n_1/2}(G') m_{k-n_1/2}(H) + 2^{dn_1/2} \max_{k' > k-n_1/2} m_{k'}(H) \\
&\leq (1 - \delta) m_{n_1/2}(H') m_{k-n_1/2}(H) + 2^{dn_1/2} \max_{k' > k-n_1/2} m_{k'}(H) \\
&\leq (1 - \delta) m_{n_1/2}(H') m_{k-n_1/2}(H) + \delta \cdot m_{k-n_1/2}(H) \\
&\leq m_{n_1/2}(H') m_{k-n_1/2}(H) \\
&\leq m_k(H_{d,n}),
\end{aligned}$$

where the second inequality follows from (4) and the third from (5). \square

Small-2: $n_1 \leq N_1$ and $k \leq (1 - \varepsilon')n/2$.

Choose $0 < \delta' = \delta'(\varepsilon, d) < 1/2$ small enough that for $\lambda \geq \frac{2\varepsilon}{d}$, and any d -regular G' not containing a copy of $K_{d,d}$, we have $Z_{G'}^{\text{match}}(\lambda) \leq \frac{1-\delta'}{1+2\delta'} Z_{H_{d,|G'|}}^{\text{match}}(\lambda)$. Such a δ' exists by Theorem 6. Choose $N = N(N_1, \varepsilon, \varepsilon')$ large enough that by Lemma 15, for $n \geq N$, $\varepsilon n < k < (1 - \varepsilon')n/2$, $n_1 \leq N_1$ and $0 \leq r \leq n_1$ we have

$$(6) \quad (1 - \delta') m_{k-r}(H_{d,n_2}) \leq \lambda^r \cdot m_k(H_{d,n_2}) \leq (1 + \delta') m_{k-r}(H_{d,n_2})$$

for some $\lambda \geq \frac{2\varepsilon}{d}$. Now we bound

$$\begin{aligned}
m_k(G) &= \sum_{r=0}^{n_1/2} m_r(G') m_{k-r}(H) \\
&\leq (1 + 2\delta') \sum_{r=0}^{n_1/2} m_r(G') \lambda^r m_k(H) \quad \text{by (6)} \\
&= (1 + 2\delta') m_k(H) Z_{G'}^{\text{match}}(\lambda) \\
&\leq (1 - \delta') m_k(H) Z_{H'}^{\text{match}}(\lambda) \\
&= (1 - \delta') \sum_{r=0}^{n_1/2} m_r(H') \lambda^r m_k(H) \\
&\leq \sum_{r=0}^{n_1/2} m_r(H') m_{k-r}(H) \quad \text{by (6)} \\
&= m_k(H_{d,n}).
\end{aligned}$$

\square

Large: $n_1 > N_1$.

Choose λ so that the expected size of the random matching in $H_{d,n}$ at fugacity λ is k , and choose s so that

$$m_s(G')m_{k-s}(H) = \max_{0 \leq r \leq n_1/2} m_r(G')m_{k-r}(H).$$

In other words, s is the most likely size of matching induced on G' when choosing a random matching from G of size exactly k (note that the choice of s may not be unique). We write $\mathbb{P}_{G,\lambda}[|\mathbf{M}| = t]$ to mean the probability that a random matching from the monomer-dimer model on G at fugacity λ has t edges. Then we have

$$\begin{aligned} m_k(G) &\leq \left(\frac{n_1}{2} + 1\right) \cdot m_s(G')m_{k-s}(H) \\ &= \left(\frac{n_1}{2} + 1\right) \cdot \frac{Z_{G'}^{\text{match}}(\lambda)}{\lambda^s} \mathbb{P}_{G',\lambda}[|\mathbf{M}| = s] \cdot \frac{Z_{H_{d,n_2}}^{\text{match}}(\lambda)}{\lambda^{k-s}} \mathbb{P}_{H_{d,n_2},\lambda}[|\mathbf{M}| = k-s], \end{aligned}$$

and

$$\begin{aligned} m_k(H_{d,n}) &\geq m_{\lfloor kn_1/n \rfloor}(H')m_{\lceil kn_2/n \rceil}(H) \\ &= \frac{Z_{H_{d,n_1}}^{\text{match}}(\lambda)}{\lambda^{\lfloor kn_1/n \rfloor}} \mathbb{P}_{H_{d,n_1},\lambda}[|\mathbf{M}| = \lfloor kn_1/n \rfloor] \cdot \frac{Z_{H_{d,n_2}}^{\text{match}}(\lambda)}{\lambda^{\lceil kn_2/n \rceil}} \mathbb{P}_{H_{d,n_2},\lambda}[|\mathbf{M}| = \lceil kn_2/n \rceil]. \end{aligned}$$

Cancelling terms, it is enough to show that

$$(7) \quad \frac{Z_{H_{d,n_1}}^{\text{match}}(\lambda)}{Z_{G'}^{\text{match}}(\lambda)} \cdot \frac{\mathbb{P}_{H_{d,n_1},\lambda}[|\mathbf{M}| = \lfloor kn_1/n \rfloor]}{\mathbb{P}_{G',\lambda}[|\mathbf{M}| = s]} \cdot \frac{\mathbb{P}_{H_{d,n_2},\lambda}[|\mathbf{M}| = \lceil kn_2/n \rceil]}{\mathbb{P}_{H_{d,n_2},\lambda}[|\mathbf{M}| = k-s]} \geq \left(\frac{n_1}{2} + 1\right).$$

Using Theorem 14 (Gnedenko's Local Limit Theorem), we see that since $\lceil kn_2/n \rceil$ is within 1 of the mean size of matching in H_{d,n_2} at fugacity λ , we have

$$\mathbb{P}_{H_{d,n_2},\lambda}[|\mathbf{M}| = \lceil kn_2/n \rceil] \geq \Omega(1/\sqrt{n_2}),$$

and

$$\mathbb{P}_{H_{d,n_2},\lambda}[|\mathbf{M}| = k-s] \leq O(1/\sqrt{n_2}).$$

Therefore the factor $\frac{\mathbb{P}_{H_{d,n_2},\lambda}[|\mathbf{M}| = \lceil kn_2/n \rceil]}{\mathbb{P}_{H_{d,n_2},\lambda}[|\mathbf{M}| = k-s]}$ is lower bounded by a constant depending only on ε and d , independent of n and n_1 , call this $\eta(\varepsilon, d)$. Likewise by Theorem 14 again since $\lfloor kn_1/n \rfloor$ is within 1 of the mean size of the random matching in H' at fugacity λ , the factor $\mathbb{P}_{H_{d,n_1},\lambda}[|\mathbf{M}| = \lfloor kn_1/n \rfloor]$ is lower bounded by $1/\sqrt{n_1}$ times another constant $\eta'(\varepsilon, d)$. Finally, $\mathbb{P}_{G',\lambda}[|\mathbf{M}| = s] \leq 1$. Using Theorem 6, we can choose $N_1 = N_1(\varepsilon, d)$ large enough that for all $n_1 > N_1$,

$$\frac{Z_{H_{d,n_1}}^{\text{match}}(\lambda)}{Z_{G'}^{\text{match}}(\lambda)} \geq \left(\frac{n_1}{2} + 1\right) \cdot \frac{\sqrt{n_1}}{\eta(\varepsilon, d) \cdot \eta'(\varepsilon, d)},$$

which gives (7). □

5. A HIERARCHY OF EXTREMAL RESULTS FOR PARTITION FUNCTIONS AND THEIR COEFFICIENTS

In this section we investigate several possible ways in which one partition function can dominate another. We discuss notions of dominance that correspond to extremal results in the literature, and prove general implications between them.

In Section 1 we discussed dominance in terms of counting (evaluating the partition function at $\lambda = 1$); evaluating the partition function at any λ (Theorems 1 and 2); in terms of individual coefficients (Conjectures 3 and 4); and in terms of the highest degree coefficient (Bregman's Theorem [2]). In Section 2 we discussed dominance in terms of the logarithmic derivative or occupancy fraction (Theorems 9 and 10).

Another notion of dominance relates to the canonical ensemble. The occupancy fractions and internal energy are not interesting observables in the canonical ensemble; they are fixed at $k/|V(G)|$ and $k/|E(G)|$ respectively by definition. Instead we consider the *free volume*. For an independent set I of a graph G we let $F_G^{\text{ind}}(I)$ denote the set of vertices of G that are neither in I nor adjacent to I . That is, they are precisely the vertices v which can be added to I to make a larger independent set. This is the free volume of I in G . For a matching M we let $F_G^{\text{match}}(M)$ denote the set of edges of G that are neither in M nor incident to an edge of M .

Now consider the expected free volume when choosing a random independent set or matching from the corresponding canonical ensemble; we denote these by

$$f_k^{\text{ind}}(G) = \mathbb{E}|F_G^{\text{ind}}(\mathbf{I}_k)|, \\ f_k^{\text{match}}(G) = \mathbb{E}|F_G^{\text{match}}(\mathbf{M}_k)|.$$

Observe that we can write $f_k^{\text{ind}}(G)$ and $f_k^{\text{match}}(G)$ in terms of the coefficients of the respective partition functions:

$$f_k^{\text{ind}}(G) = (k+1) \frac{i_{k+1}(G)}{i_k(G)}, \\ f_k^{\text{match}}(G) = (k+1) \frac{m_{k+1}(G)}{m_k(G)}.$$

We can ask for dominance of the expected free volume in the canonical ensemble, that is, showing that $f_k^{\text{ind}}(G) = (k+1) \frac{i_{k+1}(G)}{i_k(G)}$ is maximized for all k by some particular extremal graph. In [9] we conjectured $H_{d,n}$ has this property for matchings and independent sets in a regular graphs.

Conjecture 17 ([9]). *Let $2d$ divide n . Then for all k ; over d -regular, n -vertex graphs G , the quantities $f_k^{\text{ind}}(G)$ and $f_k^{\text{match}}(G)$ are maximized by $H_{d,n}$.*

As we show below in Proposition 19, dominance of the expected free volume in the canonical model implies all of the other notions of dominance of partition functions and their coefficients mentioned above. The notion was already investigated implicitly for lower bounds in the hard-core model by Cutler and Radcliffe [7].

Theorem 18 (Cutler, Radcliffe [7]). *Let $d+1$ divide n . Then for any d -regular G on n vertices, and any $1 \leq k \leq n/(d+1)$,*

$$f_k^{\text{ind}}(G) \geq f_k^{\text{ind}}(CL_{d,n}),$$

where $CL_{d,n}$ is the union of $n/(d+1)$ disjoint cliques on $d+1$ vertices. As a consequence, for all k we have

$$i_k(G) \geq i_k(CL_{d,n}).$$

In fact, the free volume interpretation gives a very short proof of the theorem: in a d -regular graph each vertex in an independent set covers d other vertices; in a union of cliques the sets covered by each vertex in an independent are disjoint, and so $CL_{d,n}$ maximizes the

expected number of covered vertices in the canonical model, or equivalently, it minimizes the expected free volume.

In the next section we describe the rich picture of implications between the different notions of dominance described above.

5.1. A hierarchy of counting theorems. Let $Z_G(\lambda) = \sum_{k=0}^n c_k(G)\lambda^k$ and $Z_H(\lambda) = \sum_{k=0}^n c_k(H)\lambda^k$, and suppose $c_0(G) = c_0(H)$ and that $c_k(G), c_k(H) \geq 0$ for all k .

Consider the following statements:

- COUNT:** $\sum_{k=0}^n c_k(G) \geq \sum_{k=0}^n c_k(H)$.
- PART:** $Z_G(\lambda) \geq Z_H(\lambda)$ for all $\lambda \geq 0$.
- COEF:** $c_k(G) \geq c_k(H)$ for all $1 \leq k \leq n$.
- OCC:** $\frac{\lambda Z'_G(\lambda)}{Z_G(\lambda)} \geq \frac{\lambda Z'_H(\lambda)}{Z_H(\lambda)}$ for all $\lambda > 0$.
- MAX:** $c_n(G) \geq c_n(H)$.
- FV:** $\frac{c_{k+1}(G)}{c_k(G)} \geq \frac{c_{k+1}(H)}{c_k(H)}$ for all $0 \leq k \leq n-1$.

For the independence polynomial of a regular graph, the theorems of Kahn, Galvin and Tetali, Zhao, and Theorem 9 state that **COUNT**, **PART**, and **OCC** hold with $G = H_{d,n}$ and H any d -regular graph on n vertices; the statements **COEF** and **FV** are Conjectures 3 and 17.

For the matching polynomial of a regular graph, Bregman's theorem is **MAX**, while Theorem 9 is **OCC**. By the following proposition this implies **COUNT** and **PART** and provides an alternative proof of **MAX**. The statements **COEF** and **FV** are Conjectures 4 and 17.

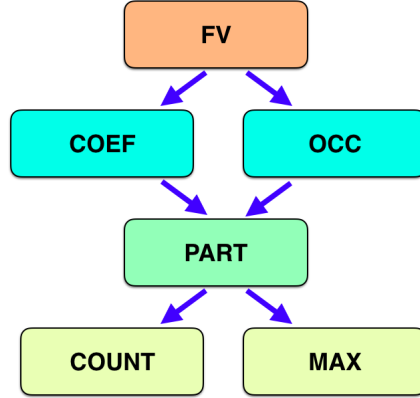
Proposition 19. *Let $Z_G(\lambda)$, $Z_H(\lambda)$ be defined as above. Then*

- (1) **PART** implies **COUNT**.
- (2) **PART** implies **MAX**.
- (3) **COEF** implies **PART**.
- (4) **OCC** implies **PART**.
- (5) **FV** implies **COEF** and **OCC**.
- (6) **COUNT** and **MAX** are incomparable in general.
- (7) **COEF** and **OCC** are incomparable in general.

Proof.

- (1) Immediate by taking $\lambda = 1$.
- (2) Take $\lambda \rightarrow \infty$.
- (3) Immediate since the coefficients are all non-negative.

FIGURE 2. The implications of Proposition 19



(4) We calculate

$$\begin{aligned}
 \log Z_G(\lambda) &= \log Z_G(0) + \int_0^\lambda (\log Z_G(t))' dt \\
 &= \log c_0(G) + \int_0^\lambda \frac{Z'_G(t)}{Z_G(t)} dt \\
 &\geq \log c_0(H) + \int_0^\lambda \frac{Z'_H(t)}{Z_H(t)} dt \\
 &= \log Z_H(0) + \int_0^\lambda (\log Z_H(t))' dt \\
 &= \log Z_H(\lambda).
 \end{aligned}$$

(5) **FV** implies **COEF** by induction on k .

To show **FV** implies **OCC**, we need to show

$$\left(\sum_{j=1}^n j c_j(G) \lambda^{j-1} \right) \left(\sum_{j=0}^n c_j(H) \lambda^j \right) - \left(\sum_{j=1}^n j c_j(H) \lambda^{j-1} \right) \left(\sum_{j=0}^n c_j(G) \lambda^j \right) \geq 0.$$

Writing the left hand side as $\sum_{j=0}^n r_j \lambda^j$, we will in fact show that $r_j \geq 0$ for all j . Let $a_i = c_i(G)$ and $b_i = c_i(H)$. We have $r_0 = a_1 b_0 - b_1 a_0$ and so $r_0 \geq 0$ by **FV** and $a_0 = b_0$. Similarly we have $r_1 = a_2 b_0 - b_2 a_0$ and so we see that $r_1 \geq 0$ by multiplying two cases of **FV**. Continuing, we have

$$r_2 = 3a_3 b_0 + a_2 b_1 - 3b_3 a_0 - a_1 b_2$$

and so $r_2 \geq 0$ since $a_0 = b_0$, $a_3 \geq b_3$, and $a_2/a_1 \geq b_2/b_1$.

In general we have

$$\begin{aligned}
r_k &= \sum_{j=0}^k (k-j+1)b_j a_{k-j+1} - \sum_{j=0}^k (k-j+1)a_j b_{k-j+1} \\
&= \sum_{j=0}^k (k-j+1)b_j a_{k-j+1} - \sum_{j=1}^{k+1} j a_{k-j+1} b_j \\
&= \sum_{j=0}^{k+1} (k-j+1)b_j a_{k-j+1} - \sum_{j=0}^{k+1} j a_{k-j+1} b_j \\
&= \sum_{j=0}^{k+1} (k-2j+1)b_j a_{k-j+1} \\
&= \sum_{j=0}^{\lfloor (k+1)/2 \rfloor} (k-2j+1)b_j a_{k-j+1} - \sum_{j=0}^{\lfloor (k+1)/2 \rfloor} (k-2j+1)a_j b_{k-j+1}.
\end{aligned}$$

Now considering the last line term by term, we see the inequality $r_k \geq 0$ reduces to showing that

$$\frac{a_{k-j+1}}{a_j} \geq \frac{b_{k-j+1}}{b_j}$$

for $j \leq \lfloor (k+1)/2 \rfloor$. This follows by multiplying successive cases of **FV**.

- (6) **COUNT** and **MAX** are incomparable: take, for example, $Z_G(\lambda) = 1 + 5\lambda + 2\lambda^2$ and $Z_H(\lambda) = 1 + 2\lambda + 3\lambda^2$.
- (7) **COEF** and **OCC** are incomparable: take first $Z_G(\lambda) = 1 + 2\lambda + \lambda^2$ and $Z_H(\lambda) = 1 + 3\lambda + \lambda^2$. Then Z_H dominates Z_G coefficient by coefficient, but for large λ its logarithmic derivative is smaller. Next take $Z_G(\lambda) = 1 + 5\lambda + 5\lambda^2 + 5\lambda^3$ and $Z_H(\lambda) = 1 + 4\lambda + 6\lambda^2 + \lambda^3$. Then $Z_G(\lambda)$ has a larger logarithmic derivative at all λ but does not dominate $Z_H(\lambda)$ coefficient by coefficient. \square

In general while **COEF** implies **PART**, the converse does not hold and so Conjectures 3 and 4 remain open. Nevertheless, the method of Section 4 gives a partial converse: that with uniqueness and stability, **PART** implies **COEF** for large enough graphs and almost all values of k .

6. PROOFS FOR OTHER MODELS

In this section we complete the proofs of Theorems 5 and 6, and prove Theorems 7 and 8.

6.1. Independent sets. We first discuss the modifications of the arguments of Sections 3 and 4 necessary to obtain Theorems 5 and 6 for independent sets in d -regular graphs, and the analogous Theorem 7 for cubic graphs of girth at least 5.

To adapt Section 3 it suffices to prove Proposition 11 for independent sets, and the first part of the following result for independent sets in cubic graphs of girth at least 5. The derivation of stability for the partition function is identical to the case of matchings, see the proof of Theorem 6 in Section 3.2.

Proposition 20. *There exists a continuous function $c^{\text{girth}}(\lambda)$ which is strictly positive when $\lambda > 0$ so that for any 3-regular graph G of girth at least 5, and any $\lambda \geq 0$,*

$$\alpha_G^{\text{ind}}(\lambda) \leq \alpha_{HW}^{\text{ind}}(\lambda) - c^{\text{girth}}(\lambda) \cdot \delta_{\circ}(G, HW),$$

As a consequence, there is a function $s^{\text{girth}}(\lambda)$, strictly increasing in λ with $s^{\text{girth}}(0) = 0$, so that for any 3-regular graph G and any $\lambda \geq 0$,

$$\frac{1}{|V(G)|} \log Z_G^{\text{ind}}(\lambda) \leq \frac{1}{14} \log Z_{HW}^{\text{ind}}(\lambda) - s^{\text{girth}}(\lambda) \cdot \delta_{\circ}(G, HW).$$

The manipulation of the linear programs from [9, 27] which show that distributions arising in $K_{d,d}$ and HW (or unions thereof) are the unique maximizers of $\alpha_G^{\text{ind}}(\lambda)$ over the relevant graph classes is the same as in the proof of Proposition 11, hence it suffices to give an analogue of Lemma 13 for each case. For the case of 3-regular graphs of girth at least 5 there is a minor additional consideration; the sets consisting of local views which occur in HW and of local views with slack equal to zero do not coincide. In the proof we work with the latter set, so the methods are identical.

For a vertex v in a graph G , write $N_t(v)$ for the set of vertices at distance exactly t from v in G , and $\overline{N}_t(v)$ for the vertices at distance at most t from v in G . We will consider the subgraphs of G induced by these sets, denoted e.g. $G[N_t(v)]$.

Lemma 21. *Let G be a d -regular graph and $\lambda > 0$. Let L be the random local view obtained by drawing an independent set \mathbf{I} from the hard-core model on G at fugacity λ , drawing a vertex v in G uniformly at random, and setting L to be the induced subgraph of G on the neighbors of v which are not adjacent to any vertex in $\mathbf{I} \setminus N(v)$. Write $\mathcal{L}_d^* = \{\emptyset, \overline{K_d}\}$ for the set of local views which may arise in $K_{d,d}$ (and which are exactly the local views which have zero slack in the dual linear program for independent sets from [9]). Then*

$$\mathbb{P}[L \notin \mathcal{L}_d^*] > \delta_{\circ}(G, K_{d,d}) \cdot \frac{\lambda}{(1+\lambda)^{2d+1}}.$$

Proof. Fix an arbitrary $v \in V(G)$ which is not in a copy of $K_{d,d}$, and let u_1, u_2 both be neighbors of v so that $N(u_1) \neq N(u_2)$. Such u_1, u_2 must exist otherwise v is in a copy of $K_{d,d}$. Let w be a vertex neighboring u_1 but not u_2 and write A_v for the event that $\{\mathbf{I} \cap N(u_2) = \emptyset\} \cap \{w \in \mathbf{I}\}$, where \mathbf{I} is a random independent set from the hard-core model on G . Note that A_v implies that the local view rooted at v cannot arise in $K_{d,d}$ because u_2 is uncovered while u_1 is not.

By successive conditioning, the probability that $\mathbf{I} \cap \{N(u_2) \cup N(w)\} = \emptyset$ is at least $(1+\lambda)^{-2d}$. Now conditioned on $\mathbf{I} \cap N(w) = \emptyset$, $\mathbb{P}[w \in \mathbf{I}] = \frac{\lambda}{1+\lambda}$, and so $\mathbb{P}[A_v] \geq \frac{\lambda}{(1+\lambda)^{2d+1}}$. Since v was arbitrary, and at least a fraction $\delta_{\circ}(G, K_{d,d})$ of the vertices in $V(G)$ are not in a copy of $K_{d,d}$, the required bound follows. \square

Lemma 22. *Let G be a cubic graph of girth at least 5 and $\lambda > 0$. Let L be the random local view obtained by drawing an independent set \mathbf{I} from the hard-core model on G at fugacity λ , drawing a vertex v in G uniformly at random, and setting L to be the induced subgraph of G on $\overline{N}_2(v)$ which are not adjacent to any vertex in $\mathbf{I} \setminus \overline{N}_2(v)$. Write \mathcal{L}^* for the set of such local views which have zero slack in the dual linear program from [27]. Then*

$$\mathbb{P}[L \notin \mathcal{L}^*] > \delta_{\circ}(G, HW) \cdot \min \left\{ \frac{\lambda}{(1+\lambda)^{15}}, \frac{1}{(1+\lambda)^2} \right\}.$$

Proof. The proof relies on the fact that HW is the only connected cubic graph of girth at least 5 containing a vertex v with which is not in a 5-cycle and which has $|N_3(v)| = 4$. To see this note that in such a graph G , the subgraph $G[\bar{N}_2(v)]$ must be a tree with 6 leaves. If $|N_3(v)| = 4$ then it is easy to check that there is a unique way to connect $N_3(v)$ to $N_2(v)$ without creating cycles of length at most 4, such that each vertex has degree 3, and that this construction gives a copy of HW . Now suppose that v is a vertex of G that not contained in a copy of HW . There are two cases.

If v is not contained in a 5-cycle, then by the above fact there must be at least 5 vertices in $N_3(v)$. Hence there exists a vertex $w \in N_3(v)$ which is connected to either 1 or 2 vertices in $N_2(v)$. Let A_v be the event that $\{w\} = \mathbf{I} \cap N_3(v)$. If w is connected to 2 vertices in $N_2(v)$ they cannot have second common neighbor, else G would contain a 4-cycle. Hence the event A_v implies that, writing $N_1(v) = \{u_1, u_2, u_3\}$, the sizes of $N_2(v) \cap N_1(u_i)$ in L are $(1, 2, 2)$ or $(1, 1, 2)$. Using the facts that any vertex is in \mathbf{I} with probability at most $\lambda/(1 + \lambda)$, and that this holds with equality for any w conditioned on $\mathbf{I} \cap N(w) = \emptyset$, and that $|N_3(v)| \leq 12$, we have $\mathbb{P}[A_v] \geq \lambda(1 + \lambda)^{-15}$.

If v is contained in a 5-cycle then let A_v be the event that the 5-cycle is present in L . Then we have $\mathbb{P}[A_v] \geq (1 + \lambda)^{-2}$.

In each case, the event A_v implies that the local view at v is one that has positive slack. The lemma follows. \square

It is now straightforward to adapt the methods of Section 4.1 and prove Theorem 5 and 7 for independent sets: essentially identical variants of all the required results are in place.

6.2. Colorings. We now turn to our remaining example, colorings of cubic graphs. Here the proofs are slightly different but follow the same idea. In proving Theorem 8, the main difference is that Theorem 10 holds only for $\lambda < 1$, and $H_{3,n}$ only maximizes $Z_G^q(\lambda)$ in this range. In fact, for $\lambda = 1$ the free energy is the same for all graphs, so we cannot hope to show that $H_{3,n}$ maximizes all of the individual coefficients of $Z_G^q(\lambda)$. Instead we can show that it maximizes the ‘anti-ferromagnetic’ coefficients; i.e. those for which k is less than the mean number of monochromatic edges in the non-interacting case $\lambda = 1$: $\mathbb{E}_{G,1}[m(G, \chi)] = \frac{3n}{2q}$ for any 3-regular graph G . Thus Theorem 8 holds for all $0 \leq k \leq (1 - \varepsilon)3n/2q$.

The results for colorings which correspond to those of Section 3 are as follows.

Proposition 23. *There exists a continuous function $c^q(\lambda)$ which is strictly positive when $\lambda \in (0, 1)$ such that for any 3-regular graph G , any $q \geq 2$, and $\lambda \in [0, 1]$,*

$$\beta_G^q(\lambda) \geq \beta_{K_{3,3}}^q(\lambda) + c^q(\lambda) \cdot \delta_o(G, K_{3,3}).$$

As a consequence, there is a function $s^q(\lambda)$, strictly decreasing in λ with $s^q(1) = 0$, so that for any 3-regular graph G , any $q \geq 2$, and $\lambda \in [0, 1]$,

$$\frac{1}{|V(G)|} \log Z_G^q(\lambda) \leq \frac{1}{6} \log Z_{K_{3,3}}^q(\lambda) - s^q(\lambda) \cdot \delta_o(G, K_{3,3}).$$

To prove Proposition 23 we require an analogue of Lemma 13.

Lemma 24. *Let G be a d -regular graph, $q \geq 2$, $0 \leq \lambda < 1$, and let L be the random local view obtained by drawing a random q -coloring χ from the q -color Potts model on G , sampling a uniformly random vertex v in G , and setting L to be the $G[\bar{N}_2(v)]$ together with the vertices*

of $N_2(v)$ colored by χ . Write \mathcal{L}_d^* for the set of such local views which have zero slack in the dual linear program from [11]. Then there exists $f(q, d, \lambda) > 0$ such that

$$\mathbb{P}[L \notin \mathcal{L}_d^*] > \delta_\circ(G, K_{d,d}) \cdot f(q, d, \lambda).$$

Proof. We show that there is a constant $f'(q, d, \lambda) > 0$ such that every vertex not in a copy of $K_{d,d}$ has a local view not in \mathcal{L}_d^* with probability at least $f'(q, d, \lambda)$.

For any vertex v in $K_{d,d}$ (or $H_{d,n}$) any two neighbors of v have the same neighborhood as each other. Unions of complete bipartite graphs are the only graphs with this property. If v is a vertex not in a copy of $K_{d,d}$ it must have two neighbors (u and w) with distinct neighborhoods. We argue that with probability bounded below by a function of q, d and λ these two neighbors of v have neighborhoods with distinct colorings. It is clear that the probability that the neighborhoods of u and w have different colorings is positive since every coloring is possible. We must show that there is a lower bound independent of n . For any set of t vertices every choice of coloring of these vertices happens with probability at least $(\lambda^d/q)^t$. This is because there are at most $d \cdot t$ edges incident vertices in the set, so the contribution to the partition function is at least $\lambda^{d \cdot t}$ and the partition function is at most q^t whenever $\lambda \leq 1$. From this we see that there is at least a $(\lambda^d/q)^{2d}$ probability that the neighborhoods of u and w are colored differently. \square

Proof of Proposition 23. As with independent sets, the manipulation of the linear program from [11] with the additional constraint offered by Lemma 24 is the same as in Section 3.3. The required statement for the internal energy follows.

We now show how the partition function result follows from the stability of the internal energy. For $\lambda \in [0, 1]$, we have:

$$\begin{aligned} \frac{1}{|V(G)|} \log Z_G^q(\lambda) &= \frac{1}{|V(G)|} \log Z_G^q(1) - \frac{d}{2} \cdot \int_\lambda^1 \frac{\beta_G^q(t)}{t} dt \\ &= \log q - \frac{d}{2} \cdot \int_\lambda^1 \frac{\beta_G^q(t)}{t} dt \\ &\leq \log q - \frac{d}{2} \cdot \int_\lambda^1 \left(\frac{\beta_{K_{3,3}}^q(t)}{t} + \frac{c^q(t) \cdot \delta_\circ(G, K_{3,3})}{t} \right) dt \\ &= \frac{1}{6} \log Z_{K_{3,3}}^q(\lambda) - s^q(\lambda) \cdot \delta_\circ(G, K_{3,3}), \end{aligned}$$

where we can set $s^q(\lambda) = \frac{d}{2} \int_\lambda^1 \frac{c^q(t)}{t} dt$. \square

As a corollary of Proposition 23 we obtain the following stability result on the number of proper colorings.

Corollary 25. *For all $q \geq 2$, there exists $c(q) > 0$ such that for any n divisible by 6 and any 3-regular graph G on n vertices, we have*

$$\frac{1}{n} \log c_0^q(G) \leq \frac{1}{n} \log c_0^q(H_{3,n}) - c(q) \cdot \delta_\circ(G, K_{3,3})$$

The proof is identical to that of Corollary 12.

6.2.1. *Proof of Theorem 8.* Here we keep d as a variable instead of setting $d = 3$, to show that if a version of Theorem 10 can be proved for $d \geq 4$, then the corresponding result on the individual coefficients will follow. We again split G into two parts, G' of size n_1 with no $K_{d,d}$ components, and $H = H_{d,n_2}$, where $n_2 = n - n_1$, consisting of all $K_{d,d}$ components.

Given $\varepsilon > 0$, d , and q , we choose auxiliary parameters with the following dependence:

- $\delta = \delta(d, q)$.
- $N_1 = N_1(d, q, \varepsilon)$.
- $\varepsilon' = \varepsilon'(N_1, \delta, d, q)$.
- $\delta' = \delta'(\varepsilon, d, q)$.
- $N = N(N_1, \varepsilon', \delta, \delta', d, q)$.

Small-1: $n_1 \leq N_1$ and $k < \varepsilon' n$.

Let $\delta = \delta(q, d)$ be such that for any G' (with $|G'| = n_1$ divisible by $2d$) containing no copies of $K_{d,d}$ we have

$$(8) \quad c_0^q(G') \leq (1 - \delta)c_0^q(H_{d,n_1}).$$

Such a δ exists by Corollary 25.

Claim 26. *For $q \geq 3$ there exists ε' small enough and N large enough (as functions of d, q, δ, N_1) such that the following holds. For all $n > N$, $k < \varepsilon' n$,*

$$\frac{c_{k-1}^q(H_{d,n})}{c_k^q(H_{d,n})} < \delta \cdot q^{-N_1}.$$

Proof. For each q -coloring of $H_{d,n}$ with exactly $k - 1$ monochromatic edges we can create a q -coloring with exactly k monochromatic edges in the following way. There are at least $\frac{n}{2d} - k \geq (\frac{1}{2d} - \varepsilon')n \geq \frac{n}{4d}$ copies of $K_{d,d}$ with no monochromatic edges so we choose one of these. We recolor this copy of $K_{d,d}$ in any way that gives a single monochromatic edge, which is possible whenever $q \geq 3$. We had at least $n/4d$ choices of the $K_{d,d}$ to recolor and each coloring of $H_{d,n}$ with k monochromatic edges can be created by this method at most $k \cdot q^{2d}$ times as there are k monochromatic edges and at most q^{2d} colorings of the $K_{d,d}$ that we altered. Thus

$$c_k^q(H_{d,n}) \geq \frac{n}{4d \cdot k \cdot q^{2d}} \cdot c_{k-1}^q(H_{d,n}),$$

and so the result follows if

$$\frac{1}{4d\varepsilon'q^{2d}} > \delta \cdot q^{-N_1}. \quad \square$$

We then see that

$$\begin{aligned} c_k^q(G) &= \sum_{r=0}^{dn_1/2} c_r^q(G') c_{k-r}^q(H) \\ &\leq c_0^q(G') c_k^q(H) + q^{n_1} \max_{k' < k} c_{k'}^q(H) \\ &\leq (1 - \delta) c_0^q(H') c_k^q(H) + q^{n_1} \max_{k' < k} c_{k'}^q(H) \quad \text{by (8)} \\ &\leq (1 - \delta) c_0^q(H') c_k^q(H) + \delta \cdot c_k^q(H) \quad \text{by Claim 26} \\ &\leq c_0^q(H') c_k^q(H) \\ &\leq c_k^q(H_{d,n}). \end{aligned} \quad \square$$

Small-2: $n_1 \leq N_1$ and $k \geq \varepsilon' n$.

Choose $0 < \delta' = \delta'(\varepsilon, q, d) < 1/2$ small enough that for $\lambda \leq 1 - g(q, d, \varepsilon)$, and any d -regular G' not containing a copy of $K_{d,d}$, we have $Z_{G'}^q(\lambda) \leq \frac{1-\delta'}{1+2\delta'} Z_{H_{d,|G'|}}^q(\lambda)$, where $g(q, d, \varepsilon) > 0$ is the function from Lemma 16. Such a δ' exists by Proposition 23. Choose $N = N(q, N_1, \varepsilon, \varepsilon')$ large enough that by Lemma 16, for $n \geq N$, $\varepsilon' n < k < (1 - \varepsilon)nd/2q$, $n_1 \leq N_1$ and $0 \leq r \leq dn_1/2$ we have

$$(9) \quad (1 - \delta') c_{k-r}^q(H_{d,n_2}) \leq \lambda^r \cdot c_k^q(H_{d,n_2}) \leq (1 + \delta') c_{k-r}^q(H_{d,n_2})$$

for some $\lambda \leq 1 - g(q, d, \varepsilon)$. Now we bound

$$\begin{aligned} c_k^q(G) &= \sum_{r=0}^{dn_1/2} c_r^q(G') c_{k-r}^q(H) \\ &\leq (1 + 2\delta') \sum_{r=0}^{dn_1/2} c_r^q(G') \lambda^r c_k^q(H) \quad \text{by (9)} \\ &= (1 + 2\delta') c_k^q(H) Z_{G'}^q(\lambda) \\ &\leq (1 - \delta') c_k^q(H) Z_{H'}^q(\lambda) \\ &= (1 - \delta') \sum_{r=0}^{dn_1/2} c_r^q(H') \lambda^r c_k^q(H) \\ &\leq \sum_{r=0}^{dn_1/2} c_r^q(H') c_{k-r}^q(H) \quad \text{by (9)} \\ &= c_k^q(H_{d,n}). \end{aligned} \quad \square$$

Large: $n_1 > N_1$.

This case is the same as that of **Large** for matchings: we use stability in the form of Proposition 23 to lower bound $\frac{Z_{H_{d,n_1}}^q(\lambda)}{Z_{G'}^q(\lambda)}$ and the local limit theorem, Theorem 14, to bound $\frac{\mathbb{P}_\lambda[m(H_{d,n_2}, \chi) = \lceil kn_2/n \rceil]}{\mathbb{P}_\lambda[m(H_{d,n_2}, \chi) = k-s]}$ and $\mathbb{P}_\lambda[m(H', \chi) = \lfloor kn_1/n \rfloor]$. \square

REFERENCES

- [1] I. Benjamini and O. Schramm. Recurrence of distributional limits of finite planar graphs. *Electronic Journal of Probability*, 6, 2001.
- [2] L. Bregman. Some properties of nonnegative matrices and their permanents. *Soviet Math. Dokl*, 14(4):945–949, 1973.
- [3] T. Carroll, D. Galvin, and P. Tetali. Matchings and independent sets of a fixed size in regular graphs. *Journal of Combinatorial Theory, Series A*, 116(7):1219–1227, 2009.
- [4] E. Cohen, W. Perkins, and P. Tetali. On the Widom–Rowlinson occupancy fraction in regular graphs. *Combinatorics, Probability and Computing*, 26(2):183–194, 2017.
- [5] P. Csikvári. Extremal regular graphs: the case of the infinite regular tree. *arXiv preprint arXiv:1612.01295*, 2016.
- [6] P. Csikvári. Lower matching conjecture, and a new proof of schrijver’s and gurvits’s theorems. *Journal of the European Mathematical Society*, 19(6):1811–1844, 2017.
- [7] J. Cutler and A. Radcliffe. The maximum number of complete subgraphs in a graph with given maximum degree. *Journal of Combinatorial Theory, Series B*, 104:60–71, 2014.
- [8] J. Cutler and A. Radcliffe. Minimizing the number of independent sets in triangle-free regular graphs. *Discrete Mathematics*, 341(3):793–800, 2018.

- [9] E. Davies, M. Jenssen, W. Perkins, and B. Roberts. Independent sets, matchings, and occupancy fractions. *Journal of the London Mathematical Society*, 96(1):47–66, 2017.
- [10] E. Davies, M. Jenssen, W. Perkins, and B. Roberts. On the average size of independent sets in triangle-free graphs. *Proceedings of the American Mathematical Society*, 146(1):111–124, 2018.
- [11] E. Davies, M. Jenssen, W. Perkins, and B. Roberts. Extremes of the internal energy of the Potts models on cubic graphs. *Random Structures & Algorithms*, to appear.
- [12] A. Dmitriev and A. Dainiak. Uniqueness of the extremal graph in the problem of maximizing the number of independent sets in regular graphs. *arXiv preprint arXiv:1602.08736*, 2016.
- [13] P. Erdős. On some new inequalities concerning extremal properties of graphs. In *Theory of Graphs (Proc. Colloq., Tihany, 1966)*, pages 77–81, 1966.
- [14] E. Friedgut. On the measure of intersecting families, uniqueness and stability. *Combinatorica*, 28(5):503–528, 2008.
- [15] S. Friedland, E. Krop, and K. Markström. On the number of matchings in regular graphs. *the Electronic Journal of Combinatorics*, 15(1):R110, 2008.
- [16] D. Galvin. Three tutorial lectures on entropy and counting. *arXiv preprint arXiv:1406.7872*, 2014.
- [17] D. Galvin and P. Tetali. On weighted graph homomorphisms. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 63:97–104, 2004.
- [18] W. Gan, P.-S. Loh, and B. Sudakov. Maximizing the number of independent sets of a fixed size. *Combinatorics, Probability and Computing*, 24(03):521–527, 2015.
- [19] B. V. Gnedenko. On a local limit theorem of the theory of probability. *Uspekhi Matematicheskikh Nauk*, 3(3):187–194, 1948.
- [20] L. Ilinca and J. Kahn. Asymptotics of the upper matching conjecture. *Journal of Combinatorial Theory, Series A*, 120(5):976–983, 2013.
- [21] J. Kahn. An entropy approach to the hard-core model on bipartite graphs. *Combinatorics, Probability and Computing*, 10(03):219–237, 2001.
- [22] J. Kahn. Entropy, independent sets and antichains: a new approach to Dedekind’s problem. *Proceedings of the American Mathematical Society*, 130(2):371–378, 2002.
- [23] P. Keevash and D. Mubayi. Stability theorems for cancellative hypergraphs. *Journal of Combinatorial Theory, Series B*, 92(1):163–175, 2004.
- [24] M. Lelarge. Counting matchings in irregular bipartite graphs and random lifts. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2230–2237. SIAM, 2017.
- [25] L. Lovász. *Large networks and graph limits*, volume 60. American Mathematical Society Providence, 2012.
- [26] D. Mubayi. Structure and stability of triangle-free set systems. *Transactions of the American Mathematical Society*, 359(1):275–291, 2007.
- [27] G. Perarnau and W. Perkins. Counting independent sets in cubic graphs of given girth. *Journal of Combinatorial Theory, Series B*, to appear.
- [28] W. Perkins. Birthday inequalities, repulsion, and hard spheres. *Proceedings of the American Mathematical Society*, 144(6):2635–2649, 2016.
- [29] O. Pikhurko. An analytic approach to stability. *Discrete Mathematics*, 310(21):2951–2964, 2010.
- [30] M. Simonovits. A method for solving extremal problems in graph theory, stability problems. In *Theory of Graphs (Proc. Colloq., Tihany, 1966)*, pages 279–319, 1968.
- [31] T. Tao. Variants of the central limit theorem. <https://terrytao.wordpress.com/2015/11/19/275a-notes-5-variants-of-the-central-limit-theorem/>, 2015. Accessed 2017-04-10.
- [32] Y. Zhao. The number of independent sets in a regular graph. *Combinatorics, Probability and Computing*, 19(02):315–320, 2010.
- [33] Y. Zhao. Extremal regular graphs: independent sets and graph homomorphisms. *American Mathematical Monthly*, 124(9):827–843, 2017.

LONDON SCHOOL OF ECONOMICS

Email address: {e.s.davies,m.o.jenssen,b.j.roberts}@lse.ac.uk

UNIVERSITY OF BIRMINGHAM

Email address: math@willperkins.org