

Inapproximability for Antiferromagnetic Spin Systems in the Tree Non-Uniqueness Region*

Andreas Galanis[†]

Daniel Štefankovič[‡]

Eric Vigoda[§]

November 6, 2014

Abstract

A remarkable connection has been established for antiferromagnetic 2-spin systems, including the Ising and hard-core models, showing that the computational complexity of approximating the partition function for graphs with maximum degree Δ undergoes a phase transition that coincides with the statistical physics uniqueness/non-uniqueness phase transition on the infinite Δ -regular tree. Despite this clear picture for 2-spin systems, there is little known for multi-spin systems. We present the first analog of the above inapproximability results for multi-spin systems.

The main difficulty in previous inapproximability results was analyzing the behavior of the model on random Δ -regular bipartite graphs, which served as the gadget in the reduction. To this end one needs to understand the moments of the partition function. Our key contribution is connecting: (i) induced matrix norms, (ii) maxima of the expectation of the partition function, and (iii) attractive fixed points of the associated tree recursions (belief propagation). The view through matrix norms allows a simple and generic analysis of the second moment for any spin system on random Δ -regular bipartite graphs. This yields concentration results for any spin system in which one can analyze the maxima of the first moment. The connection to fixed points of the tree recursions enables an analysis of the maxima of the first moment for specific models of interest.

For k -colorings we prove that for even k , in a tree non-uniqueness region (which corresponds to $k < \Delta$) there is no FPRAS, unless $\text{NP}=\text{RP}$, to approximate the number of colorings for triangle-free Δ -regular graphs. Our proof extends to the antiferromagnetic Potts model, and, in fact, to every antiferromagnetic model under a mild condition.

*A preliminary version of this paper appeared in *Proceedings of the 46th Annual ACM Symposium on Theory of Computing* (STOC), 823-831, 2014.

[†]University of Oxford, Wolfson Building, Parks Road, Oxford, OX1 3QD, UK. andreas.galanis@cs.ox.ac.uk. The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) ERC grant agreement no. 334828. The paper reflects only the authors' views and not the views of the ERC or the European Commission. The European Union is not liable for any use that may be made of the information contained therein.

[‡]Department of Computer Science, University of Rochester, Rochester, NY 14627. stefanko@cs.rochester.edu. Research supported in part by NSF grant CCF-1318374.

[§]School of Computer Science, Georgia Institute of Technology, Atlanta GA 30332. vigoda@cc.gatech.edu. Research supported in part by NSF grant CCF-1217458.

1 Introduction

1.1 Background

Spin systems are a general framework from statistical physics that captures classical physics models, including the Ising and Potts models, and models of particular combinatorial interest, including k -colorings and the hard-core lattice gas model defined on independent sets. We define these combinatorial models more precisely before presenting the context of our results.

The hard-core lattice gas model is an example of a 2-spin system. For a graph $G = (V, E)$, configurations of the model are the set Ω of independent sets of G . The model is parameterized by an activity $\lambda > 0$, and a configuration $\sigma \in \Omega$ is assigned weight $w(\sigma) = \lambda^{|\sigma|}$. The Gibbs distribution is $\mu(\sigma) = w(\sigma)/Z$ where the normalizing factor is known as the partition function and is defined as $Z = \sum_{\sigma \in \Omega} w(\sigma)$. In the hard-core model the spins correspond to occupied/unoccupied. Multi-spin systems are models with more than 2 spins, an example being the k -colorings problem. In the colorings problem, for a graph $G = (V, E)$, configurations are the set Ω of assignments of a set of k colors to vertices so that neighboring vertices receive different colors. The Gibbs distribution is the uniform distribution over Ω , and in this case the partition function $Z = |\Omega|$ is the number of k -colorings in G .

The hard-core model and colorings are examples of antiferromagnetic systems – neighboring vertices “prefer” to have different spins. In contrast, in ferromagnetic systems neighboring spins tend to align. We defer the formal definition of antiferromagnetic spin systems to Section 1.2.3 (see Definition 1.3), where we also discuss how our results extend to general spin systems.

The focus of this paper is the computational complexity of computing the partition function. Exact computation of the partition function is typically #P-complete, even for very restricted classes of graphs [Gre00]. Hence our focus is on the existence of a fully-polynomial approximation scheme – either a deterministic FPTAS or randomized FPRAS – for estimating the partition function. For any spin system, (approximate) sampling from the Gibbs distribution implies an FPRAS for estimating the partition function, and hence our hardness results also apply to the associated sampling problem.

The computational complexity of approximating the partition function is now well-understood for 2-spin systems, such as the Ising and hard-core models. For all ferromagnetic 2-spin systems, there is an FPRAS for estimating the partition function [GJP03]. The picture is more intricate (and fascinating) for antiferromagnetic 2-spin systems. We will detail the picture after introducing the statistical physics notion of a phase transition.

Let $\mathbb{T}_{\Delta, \ell}$ denote the complete Δ -regular tree of depth ℓ with root r . The question of interest is whether or not we can fix a configuration on the leaves of $\mathbb{T}_{\Delta, \ell}$ so that the root is influenced by this boundary configuration in the limit $\ell \rightarrow \infty$. For the example of colorings, fix a coloring σ_ℓ of the leaves (such that there is at least one coloring of the rest of the tree that is consistent with σ_ℓ). Look at a random coloring of the tree $\mathbb{T}_{\Delta, \ell}$ conditioned on the leaves having coloring σ_ℓ . For all sequences (σ_ℓ) of fixed leaf colorings, if in the limit $\ell \rightarrow \infty$, the marginal at the root is uniform over the k colors, then we say uniqueness holds, and otherwise we say non-uniqueness holds. (The terminology comes from statistical physics where the focus is on the set of infinite-volume Gibbs measures, see [Geo11].)

For the hard-core model the critical activity is $\lambda_c(\Delta) = (\Delta - 1)^{\Delta-1}/(\Delta - 2)^\Delta$ [Kel91]. [Wei06] presented an FPTAS for estimating the partition function in the tree uniqueness region (i.e., when $\lambda < \lambda_c(\Delta)$). On the other side, [Sly10] (extended in [SS12, GGS⁺14, GSV12]) proved that, unless NP=RP, it is NP-hard to obtain an FPRAS for Δ -regular graphs in the tree non-uniqueness region (i.e., when $\lambda > \lambda_c(\Delta)$). These results were extended to all 2-spin antiferromagnetic models by

[LLY13] (see also [SST12]) and [SS12]. For 2-spin antiferromagnetic models, this establishes a beautiful picture connecting the computational complexity of approximating the partition function to statistical physics phase transitions in the infinite tree.

1.2 Main Results

The picture for multi-spin systems (systems with $q > 2$ possible spins for vertices) is much less clear; the above approaches for 2-spin systems do not extend to multi-spin models in a straightforward manner. We aim to establish the analog of the above inapproximability results for the colorings problem, namely, NP-hardness in the tree non-uniqueness region. Our techniques and results generalize to a broad class of antiferromagnetic spin systems.

1.2.1 Results for Colorings

For the colorings problem, even understanding the uniqueness threshold is challenging. [Jon02] established uniqueness when $k \geq \Delta + 1$, and it is easy to show non-uniqueness when $k \leq \Delta$ since a fixed coloring on the leaves can “freeze” the internal coloring. For 2-spin systems uniqueness can be characterized by the existence of multiple solutions of a certain system of equations (22), called tree recursions, see Section 4 for additional explanation. In statistical physics terminology the solutions to these equations correspond to semi-translation invariant measures on the infinite tree \mathbb{T}_Δ . For colorings the uniqueness threshold and the semi-translation invariant uniqueness threshold no longer coincide. In particular, [BW02] established, for semi-translation invariant measures, uniqueness when $k \geq \Delta$ and non-uniqueness when $k < \Delta$.

We prove, for even k , that it is NP-hard to approximate the number of colorings (in other words, NP-hard to approximate the partition function) when there is non-uniqueness of semi-translation invariant Gibbs measures on \mathbb{T}_Δ , i.e., when $k < \Delta$. Moreover, our result proves hardness for the class of triangle-free Δ -regular graphs. Hence, our result is particularly interesting in the region $k = \Omega(\Delta/\log \Delta)$ since a seminal result of [Joh96] (see also [MR02]) shows that all triangle-free graphs are colorable with $O(\Delta/\log \Delta)$ colors. His proof, which uses the nibble method and the Lovász Local Lemma, can be made algorithmic using the constructive proof of [MT10]. For general graphs with maximum degree Δ , the interesting region is $k = \Delta - O(\sqrt{\Delta})$, since [MR01] showed, for sufficiently large constant Δ , a polynomial-time algorithm to determine if a graph with maximum degree Δ is k -colorable when $k \geq \Delta - \sqrt{\Delta} + 3$. We note that most parts of the proof extend to the odd k case as well, modulo the technical condition described in the end of Section 1.2.3.

Here is the formal statement of our inapproximability result for colorings.

Theorem 1.1. *For all even $k \geq 4$, all $\Delta \geq 3$, for the k -colorings problem, when $k < \Delta$, unless $\text{NP}=\text{RP}$, there is no FPRAS that approximates the partition function for triangle-free Δ -regular graphs. Moreover, there exists $\varepsilon = \varepsilon(k, \Delta)$ such that, unless $\text{NP}=\text{RP}$, one cannot approximate the partition function within a factor $2^{\varepsilon n}$ for triangle-free Δ -regular graphs (where n is the number of vertices).*

1.2.2 Results for Antiferromagnetic Potts

Our result also extends to the antiferromagnetic Potts model. In the q -state Potts model there is a parameter $B > 0$ which corresponds to the “temperature” and controls the strength of the interactions along an edge. For a graph $G = (V, E)$, the set Ω of configurations are assignments σ where $\sigma : V \rightarrow [q]$. Each configuration has a weight $w(\sigma) = B^{m(\sigma)}$ where $m(\sigma)$ is the number of monochromatic edges in σ . The Gibbs distribution is $\mu(\sigma) = w(\sigma)/Z$ where $Z = \sum_{\tau \in \Omega} w(\tau)$

is the partition function. The case $B > 1$ is the ferromagnetic Potts model, and $B < 1$ is the antiferromagnetic Potts model. Colorings corresponds to the $B = 0$ case, and the Ising model is the $q = 2$ case.

The uniqueness/non-uniqueness threshold for the infinite tree \mathbb{T}_Δ is not known for the antiferromagnetic Potts model. We prove that the uniqueness/non-uniqueness threshold for semi-translation invariant Gibbs measures on \mathbb{T}_Δ occurs at $B_c(\Delta) = \frac{\Delta-q}{\Delta}$. We believe this threshold coincides with the uniqueness/non-uniqueness threshold, unlike in the case of colorings. We prove, for even q , that approximating the partition function is NP-hard in the non-uniqueness region for semi-translation invariant measures.

Theorem 1.2. *For all even $q \geq 4$, all $\Delta \geq 3$, for the antiferromagnetic q -state Potts model, for all $B < \frac{\Delta-q}{\Delta}$, unless NP=RP, there is no FPRAS that approximates the partition function for triangle-free Δ -regular graphs. Moreover, there exists $\varepsilon = \varepsilon(q, \Delta)$ such that, unless NP=RP, one cannot approximate the partition function within a factor $2^{\varepsilon n}$ for triangle-free Δ -regular graphs (where n is the number of vertices).*

1.2.3 Results for General Antiferromagnetic Models

Our approach applies in much more generality and yields inapproximability of the partition function for any antiferromagnetic model when there is non-uniqueness of semi-translation invariant measures on \mathbb{T}_Δ and mild additional conditions.

We first need to define general antiferromagnetic models. A general q -spin system is specified by a symmetric $q \times q$ interaction matrix $\mathbf{B} = (B_{ij})_{i,j \in [q]}$ with non-negative entries, which specify the strength of the interaction between the spins. For example, the interaction matrix for the Potts model has off-diagonal entries equal to 1 and its diagonal entries equal to B . For a finite undirected graph $G = (V, E)$, a q -spin system is a probability distribution μ_G over the space Ω_G of all configurations, i.e., spin assignments $\sigma : V \rightarrow [q]$. The weight of a configuration $\sigma \in \Omega_G$ is the product of neighboring spin interactions, that is,

$$w_G(\sigma) = \prod_{(u,v) \in E} B_{\sigma(u)\sigma(v)}.$$

The Gibbs distribution μ_G is defined as $\mu_G(\sigma) = w_G(\sigma)/Z_G$ where the partition function Z_G is $Z_G = \sum_{\sigma \in \Omega_G} w_G(\sigma)$. We drop the subscript G when the graph under consideration is clear.

We use the following definition of antiferromagnetic models in terms of the signature of the interaction matrix \mathbf{B} , i.e., the signs of its eigenvalues. The interaction matrix \mathbf{B} is assumed to be symmetric and have non-negative entries. These are standard assumptions since we are interested in undirected graphs and the Gibbs distribution should be a probability distribution. W.l.o.g., we will also assume that \mathbf{B} is irreducible. Otherwise, by a suitable permutation of the spins, \mathbf{B} can be put into block diagonal form (which coincides with the normal form of the reducible \mathbf{B}) where each of the blocks is either irreducible or zero. Effectively, this says that the original spin model can be studied by considering the induced sub-models of each block which correspond to irreducible symmetric matrices (where our results apply). For connected graphs G , the partition function for the original model is simply the sum of the partition functions of each sub-model.

We are now ready to give the definition of antiferromagnetism we use.

Definition 1.3. *Let \mathbf{B} be the interaction matrix of a q -state spin system. Since \mathbf{B} is symmetric all of its eigenvalues are real. Also note that it has non-negative entries and by irreducibility, the Perron-Frobenius theorem implies that one of the eigenvalues of \mathbf{B} with the largest magnitude*

is positive and simple, i.e., the associated eigenspace is one-dimensional. The model is called antiferromagnetic if all the other eigenvalues are negative. Note that no eigenvalue is allowed to be zero and hence \mathbf{B} is regular.

The above definition generalizes antiferromagnetism for 2-spin systems (see [GJP03, LLY13, SS12]), and captures colorings as well as the antiferromagnetic region for the Potts models. Moreover, the above definition seems natural in that it implies that neighboring vertices prefer to have different spin assignments (see Corollary 6.4 in Section 6.2). Another nice feature of Definition 1.3 is that it does not depend on the presence of external fields. Specifically, for Δ -regular graphs, any external field can be pushed into the interaction matrix \mathbf{B} with a congruence transformation of the matrix \mathbf{B} . The resulting interaction matrix, by Sylvester's law of inertia, has the same number of positive, zero and negative eigenvalues and in particular remains antiferromagnetic.

We conclude this discussion by pointing out that some of our results for general models are more easily stated when \mathbf{B} is further assumed to be aperiodic. We shall refer to such matrices \mathbf{B} (irreducible and aperiodic) as *ergodic*. Note that if \mathbf{B} is periodic, its period must be two, since \mathbf{B} is symmetric. Such a model is only interesting on bipartite graphs (otherwise the partition function is zero). Definition 1.3 implies that the interaction matrix \mathbf{B} of an antiferromagnetic model is ergodic whenever $q \geq 3$ (note that it is trivial to compute the partition function on periodic models with $q = 2$).

We need several additional definitions concerning the moments of the partition function. For antiferromagnetic models on a random Δ -regular bipartite graph $G = (V, E)$ with bipartition $V = V_1 \cup V_2$, the goal is to understand the Gibbs distribution μ_G by looking at the distribution of spin values in V_1 and V_2 . Let $n = |V_1| = |V_2|$. For a configuration $\sigma : V \rightarrow [q]$, we shall denote the set of vertices assigned spin i by $\sigma^{-1}(i)$. Denote by Δ_q the simplex $\Delta_q = \{(x_1, x_2, \dots, x_q) \in \mathbb{R}^q \mid \sum_{i=1}^q x_i = 1 \text{ and } x_i \geq 0 \text{ for } i = 1, \dots, q\}$. For $\alpha, \beta \in \Delta_q$, let

$$\Sigma^{\alpha, \beta} = \left\{ \sigma : V \rightarrow \{1, \dots, q\} \mid |\sigma^{-1}(i) \cap V_1| = \alpha_i n, |\sigma^{-1}(i) \cap V_2| = \beta_i n \text{ for } i = 1, \dots, q \right\},$$

that is, configurations in $\Sigma^{\alpha, \beta}$ assign $\alpha_i n$ and $\beta_i n$ vertices in V_1 and V_2 the spin value i , respectively¹. We will be interested in the total weight $Z_G^{\alpha, \beta}$ of configurations in $\Sigma^{\alpha, \beta}$, namely

$$Z_G^{\alpha, \beta} = \sum_{\sigma \in \Sigma^{\alpha, \beta}} w(\sigma).$$

We study $Z_G^{\alpha, \beta}$ by looking at the moments $\mathbf{E}_G[Z_G^{\alpha, \beta}]$ and $\mathbf{E}_G[(Z_G^{\alpha, \beta})^2]$, where the expectation is over the distribution of the random Δ -regular bipartite graph, from hereon denoted by \mathcal{G} .

For $\alpha, \beta \in \Delta_q$, denote the leading term of the first and second moments as:

$$\Psi_1(\alpha, \beta) = \Psi_1^{\mathbf{B}}(\alpha, \beta) := \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbf{E}_G[Z_G^{\alpha, \beta}]. \quad (1)$$

$$\Psi_2(\alpha, \beta) = \Psi_2^{\mathbf{B}}(\alpha, \beta) := \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbf{E}_G \left[\left(Z_G^{\alpha, \beta} \right)^2 \right]. \quad (2)$$

We will refer to α, β that maximize Ψ_1 as *dominant phases*. Moreover, we say that a dominant phase (α, β) is *Hessian dominant* if the Hessian of Ψ_1 at (α, β) is negative definite. (Note this is a sufficient condition for α, β to be a local maximum.) In the uniqueness region there is a unique dominant phase and it has $\alpha = \beta$. In contrast, for 2-spin antiferromagnetic models and

¹Technically we need to define $\Sigma^{\alpha, \beta} = \left\{ \sigma : V \rightarrow [q] \mid |\sigma^{-1}(i) \cap V_1| = \hat{\alpha}_i, |\sigma^{-1}(i) \cap V_2| = \hat{\beta}_i \text{ for } i \in [q] \right\}$, where $\{\hat{\alpha}_i\}$ are $\{\alpha_i n\}$ rounded in a canonical fashion so that their sum is preserved (for example using ‘‘cascade rounding’’) and in the same way $\{\hat{\beta}_i\}$ are $\{\beta_i n\}$ rounded.

for colorings in the semi-translation non-uniqueness region, the dominant phases have $\alpha \neq \beta$, and one expects this would hold for all antiferromagnetic models. In our reduction we will need this additional condition that the dominant phases are not symmetric (i.e., $\alpha \neq \beta$).

Our main technical result relates the second moment to the first moment, for any model on random bipartite regular graphs.

Theorem 1.4. *For any spin system, for all $\Delta \geq 3$,*

$$\max_{\alpha, \beta} \Psi_2(\alpha, \beta) = 2 \max_{\alpha, \beta} \Psi_1(\alpha, \beta).$$

Crucially, Theorem 1.4 implies that $\Psi_2(\alpha, \beta) = 2\Psi_1(\alpha, \beta)$ for dominant phases, which is key for our arguments, since it will eventually allow us to find the asymptotic distribution of the random variables $Z_G^{\alpha, \beta}$ (as $n \rightarrow \infty$). We do this by applying the so-called small subgraph conditioning method. The asymptotic convergence is utilized to prove the properties of the gadget we use in the reduction. The gadget is a slight modification of a random Δ -regular bipartite graph and its properties are described in Section 6.4. The precise formulation of these properties does not matter at this stage, but rather that we can prove them when the dominant phases (α, β) satisfy the following conditions: (i) each dominant phase is Hessian dominant, (ii) the dominant phases are permutation symmetric, i.e., obtainable from one another by a suitable permutation of the set of spins (we clarify here that the permutations must be automorphisms of the interaction matrix \mathbf{B})², (iii) each dominant phase (α, β) has $\alpha \neq \beta$. Condition (iii) implies that the model is in the non-uniqueness region of \mathbb{T}_Δ and, further, that a typical configuration in the Gibbs distribution of the random graph is “unbalanced” between the two sides, which allows to encode a CSP (in our case MAX-CUT). Condition (i) ensures the asymptotic convergence of $Z_G^{\alpha, \beta}$. Condition (ii) ensures that the asymptotic distribution of $Z_G^{\alpha, \beta}$ is identical for all the dominant phases.

We want to remark why the permutation symmetry condition arises naturally. A generic multi-spin system in the semi-translational non-uniqueness region will have exactly two maxima of Ψ_1 and hardness (assuming $\text{NP} = \text{RP}$) follows easily. Models coming from statistical physics (for example, Potts model or Widom-Rowlinson model) are not generic since they usually come with permutation symmetries of the same type as condition (ii) in the previous paragraph. (The symmetries make the hardness result more difficult to state and prove.)

We now state our general inapproximability result.

Theorem 1.5. *Let $q \geq 2, \Delta \geq 3$. For an antiferromagnetic q -spin system whose interaction matrix \mathbf{B} is ergodic, if the dominant phases (α, β) of Ψ_1 are permutation symmetric and all of them are Hessian dominant and satisfy $\alpha \neq \beta$, then, unless $\text{NP} = \text{RP}$, there is no FPRAS for approximating the partition function for triangle free Δ -regular graphs. Moreover, there exists $\varepsilon = \varepsilon(q, \Delta)$ such that, unless $\text{NP} = \text{RP}$, one cannot approximate the partition function within a factor $2^{\varepsilon n}$ for triangle-free Δ -regular graphs (where n is the number of vertices).*

We remark here that, whenever the hypotheses of Theorem 1.5 are satisfied, the spin system with interaction matrix \mathbf{B} is in the tree non-uniqueness region of \mathbb{T}_Δ , see Section 4 for more details. However, the reverse direction is not necessarily true, that is, an antiferromagnetic spin system in the tree non-uniqueness region of \mathbb{T}_Δ does not necessarily have multiple dominant phases, an example is the k -colorings model when $k = \Delta$ (see Theorem 1.6 below).

²More precisely, the permutation symmetric property can be stated as follows: for any two dominant phases, say (α_1, β_1) and (α_2, β_2) , there exists a $q \times q$ permutation matrix \mathbf{P} such that $\mathbf{B} = \mathbf{P}\mathbf{B}\mathbf{P}^\top$ and $(\alpha_1, \beta_1) = (\mathbf{P}\alpha_2, \mathbf{P}\beta_2)$ or $(\alpha_1, \beta_1) = (\mathbf{P}\beta_2, \mathbf{P}\alpha_2)$. In other words, the dominant phases can be obtained from each other by interchanging α and β , by permuting the spins in a way that \mathbf{B} is left invariant, or a combination of the previous two operations.

For illustrative purposes, we first note that the inapproximability results for antiferromagnetic 2-spin systems in the tree non-uniqueness region [Sly10, SS12, GSV12] follow as corollaries of Theorem 1.5. In particular, for antiferromagnetic 2-spin systems it is well known that for any $\Delta \geq 3$, in the non-uniqueness region of \mathbb{T}_Δ , the maximizers of Ψ_1 are exactly two pairs (α, β) and (β, α) with $\alpha \neq \beta$. Note that these two dominant phases satisfy trivially the permutation symmetric property. Moreover, it can also be verified that they are Hessian dominant and hence the hypotheses of Theorem 1.5 are satisfied.

As a more indicative application of Theorem 1.5, let us deduce Theorems 1.1 and 1.2. To do this, we need the following theorem (proved in Section 7) which describes the dominant phases for the colorings and antiferromagnetic Potts models.

Theorem 1.6. *Let $q \geq 3$, $0 \leq B < 1$ and $\Delta \geq 3$. For the antiferromagnetic q -state Potts model with parameter B on a random Δ -regular bipartite graph (note that the k -colorings model corresponds to $B = 0$ and $q = k$ in the following), it holds that*

1. *When $B \geq \frac{\Delta-q}{\Delta}$, there is a unique dominant phase (α, β) which satisfies $\alpha = \beta$.*
2. *For all even $q \geq 4$, for all $\Delta \geq 3$, when $0 \leq B < \frac{\Delta-q}{\Delta}$, the dominant phases (α, β) are in one-to-one correspondence with subsets $T \subseteq [q]$ with $|T| = q/2$. Moreover, there exist $a(q, \Delta, B), b(q, \Delta, B)$ with $a \neq b$ such that for $T \subseteq [q]$ with $|T| = q/2$, the dominant phase (α, β) corresponding to T satisfies*

$$\begin{aligned} \alpha_i &= a \text{ if } i \in T, & \alpha_i &= b \text{ if } i \notin T, \\ \beta_i &= b \text{ if } i \in T, & \beta_i &= a \text{ if } i \notin T. \end{aligned} \tag{3}$$

Moreover, the dominant phases are Hessian.

Proof of Theorems 1.1 and 1.2. We verify the hypotheses of Theorem 1.5. Equation (3) of Theorem 1.6 establishes that the dominant phases (α, β) are permutation symmetric and each of them satisfies $\alpha \neq \beta$. Thus, the hypotheses of Theorem 1.5 hold in the regime $q < \Delta$ and $0 \leq B < \frac{\Delta-q}{\Delta}$. \square

Note that the restriction of even k, q in Theorems 1.1 and 1.2, respectively, is a technical one and comes from the second part of Theorem 1.6. For odd q , we are unable to establish whether the dominant phases are supported on vectors with two or three different entries, see Section 7 for more details. Classifying the dominant phases for odd q would also extend the inapproximability results of Theorems 1.1 and 1.2.

1.3 Proof Approach

The key gadget in the inapproximability results for 2-spin models is a random Δ -regular bipartite graph. The rough idea for the hard-core model is that in the tree non-uniqueness region, on a random Δ -regular bipartite graph, an independent set from the Gibbs distribution is “unbalanced” with high probability (the fraction of occupied vertices in the two parts of the bipartition differ by a constant). To analyze random regular bipartite graphs, the original inapproximability result of [Sly10] relied on a second moment analysis of [MWW09], which Sly called a technical tour-de-force. The optimization at the heart of that analysis was difficult enough that his result only held for λ close to the uniqueness threshold.

We present a new approach for the associated optimization problem which is at the heart of the second moment analysis. Our approach yields a simple, short analysis that holds for *any model* on

random Δ -regular bipartite graphs. The key idea is to define a new function Φ , which is represented as an induced matrix norm, and has the same critical points as the first moment. We can then use the fact that induced matrix norms are multiplicative over tensor product to analyze the second moment.

1.4 Paper Outline

In Section 2 we derive some basic expressions for the first and second moments. Then in Section 3 we analyze the second moment using matrix norms and thereby prove Theorem 1.4. In Section 4, we analyze the maxima of the function Ψ_1 . There, we further prove a connection between local maxima of Ψ_1 and stable fixpoints of the so-called tree recursions which is used in later sections.

The reduction for the inapproximability results uses an intermediate problem, which we call the “phase labeling problem”. Our inapproximability results hinge on showing that the phase labeling problem is hard to approximate. In Section 5, we give the main elements of this reduction for the colorings model to introduce the relevant concepts. The hardness of approximating the phase labeling problem for general antiferromagnetic models is proved in Section 6, where we also fill in the details which were omitted in the simplified exposition for the colorings model.

We show how the phase labeling problem reduces to the approximation of the partition function in Section 6.4, based on arguments in [SS12]. The reduction uses gadgets whose existence and construction are based on a slight variation of the random Δ -regular bipartite graph distribution. At this point, to establish the properties of the gadgets, we use the small subgraph conditioning method. The application of the method is fairly standard though technically intensive due to its use of precise asymptotics for the first and second moments. The technical details of applying the method in our case are given in Appendix A, while the asymptotics for the first and second moments are derived in Appendix B.

The proof of our general inapproximability result (Theorem 1.5) is given in Section 6.1. We saw in Section 1.2.3 how to deduce the inapproximability results for the colorings and Potts models (Theorems 1.1 and 1.2) from Theorem 1.5 using the classification of the dominant phases in Item 2 of Theorem 1.6. The proof of Item 2 in Theorem 1.6 is given in Section 7.

Finally, in Appendix C, we extend the argument of [BW02] to prove Item 1 of Theorem 1.6, that is, show uniqueness for semi-translation invariant Gibbs measures for the antiferromagnetic Potts model when $B \geq (\Delta - q)/\Delta$.

2 Expressions for the first and second moments

In this section we derive the expressions for the first and second moments of $Z_G^{\alpha,\beta}$ and, in particular, the expressions for Ψ_1 and Ψ_2 .

Let $\mathcal{G}_n(\Delta)$ be the probability distribution over bipartite graphs with $n + n$ vertices formed by taking the union of Δ random perfect matchings. We will use the simplified notation $\mathcal{G}_n := \mathcal{G}_n(\Delta)$ or even $\mathcal{G} := \mathcal{G}_n(\Delta)$ when n is clear from context. Strictly speaking, this distribution is over bipartite multi-graphs. However, since our results hold asymptotically almost surely (a.a.s.) over \mathcal{G}_n , as noted in [MWW09], by contiguity arguments they also hold a.a.s. for the uniform distribution over bipartite Δ -regular graphs. For a complete account of contiguity, we refer the reader to [JLR00, Chapter 9].

Let $G \sim \mathcal{G}$. We will denote the two sides of the bipartition of G as V_1, V_2 . We first compute the first moment $\mathbf{E}_{\mathcal{G}}[Z_G^{\alpha,\beta}]$. For $\sigma \in \Sigma^{\alpha,\beta}$ and a uniform matching between V_1 and V_2 , let x_{ij} denote the number of edges matching vertices in $\sigma^{-1}(i) \cap V_1$ and $\sigma^{-1}(j) \cap V_2$. Under the convention that

$0^0 \equiv 1$, we then have

$$\mathbf{E}_G[Z_G^{\alpha,\beta}] = \binom{n}{\alpha_1 n, \dots, \alpha_q n} \binom{n}{\beta_1 n, \dots, \beta_q n} \times \left(\sum_{\mathbf{x}} \frac{\prod_i \binom{\alpha_i n}{x_{i1}n, \dots, x_{iq}n} \prod_j \binom{\beta_j n}{x_{1j}n, \dots, x_{qj}n} \prod_{i,j} B_{ij}^{n x_{ij}}}{\binom{n}{x_{11}n, \dots, x_{qq}n}} \right)^\Delta, \quad (4)$$

where the sum ranges over $\mathbf{x} = (x_{11}, \dots, x_{qq})$ with $n\mathbf{x} \in \mathbb{Z}^{q^2}$ satisfying the following constraints:

$$\begin{aligned} \sum_j x_{ij} &= \alpha_i \quad (\forall i \in [q]), \quad \sum_i x_{ij} = \beta_j \quad (\forall j \in [q]), \\ x_{ij} &\geq 0 \quad (\forall (i, j) \in [q]^2). \end{aligned} \quad (5)$$

The first line in (4) accounts for the cardinality of $\Sigma^{\alpha,\beta}$, while the second line is $\mathbf{E}_G[w_G(\sigma)]$ for an arbitrary $\sigma \in \Sigma^{\alpha,\beta}$. Since the weight of a configuration is multiplicative over the edges and the matchings are independent, $\mathbf{E}_G[w_G(\sigma)]$ is the Δ -power of the expected contribution of a single matching. The latter is completely determined by \mathbf{x} and is equal to $\prod_{i,j} B_{ij}^{x_{ij}}$, scaled by the probability that the matching induces the prescribed \mathbf{x} .

We next calculate the second moment of $Z_G^{\alpha,\beta}$. To do this, for $(\sigma_1, \sigma_2) \in \Sigma^{\alpha,\beta} \times \Sigma^{\alpha,\beta}$, we need to compute $\mathbf{E}_G[w_G(\sigma_1)w_G(\sigma_2)]$. Let $\gamma_{ik} = |\sigma_1^{-1}(i) \cap \sigma_2^{-1}(k) \cap V_1|/n$, $\delta_{jl} = |\sigma_1^{-1}(j) \cap \sigma_2^{-1}(l) \cap V_2|/n$. The vectors γ and δ capture the overlap of configurations in V_1 and V_2 , respectively. For a uniform matching between V_1 and V_2 , let y_{ikjl} denote the number of edges matching vertices in $\sigma_1^{-1}(i) \cap \sigma_2^{-1}(k) \cap V_1$ and $\sigma_1^{-1}(j) \cap \sigma_2^{-1}(l) \cap V_2$ (scaled by n). Under the convention $0^0 \equiv 1$, we then have

$$\begin{aligned} \mathbf{E}_G[(Z_G^{\alpha,\beta})^2] &= \sum_{\gamma, \delta} \binom{n}{\gamma_{11}n, \dots, \gamma_{qq}n} \binom{n}{\delta_{11}n, \dots, \delta_{qq}n} \\ &\times \left(\sum_{\mathbf{y}} \frac{\prod_{i,k} \binom{\gamma_{ik}n}{y_{ik11}n, \dots, y_{ikqq}n} \prod_{j,l} \binom{\delta_{jl}n}{y_{11jl}n, \dots, y_{qqjl}n} \prod_{i,k,j,l} (B_{ij}B_{kl})^{n y_{ikjl}}}{\binom{n}{y_{1111}n, \dots, y_{qqqq}n}} \right)^\Delta, \quad (6) \end{aligned}$$

where the sums range over $\gamma = (\gamma_{11}, \dots, \gamma_{qq})$, $\delta = (\delta_{11}, \dots, \delta_{qq})$, $\mathbf{y} = (y_{1111}, \dots, y_{qqqq})$ with $n\gamma, n\delta \in \mathbb{Z}^{q^2}$ and $n\mathbf{y} \in \mathbb{Z}^{q^4}$ satisfying

$$\begin{aligned} \sum_k \gamma_{ik} &= \alpha_i \quad (\forall i \in [q]), \quad \sum_l \delta_{jl} = \beta_j \quad (\forall j \in [q]), \quad \sum_{j,l} y_{ikjl} = \gamma_{ik} \quad (\forall (i, k) \in [q]^2), \\ \sum_i \gamma_{ik} &= \alpha_k \quad (\forall k \in [q]), \quad \sum_j \delta_{jl} = \beta_l \quad (\forall l \in [q]), \quad \sum_{i,k} y_{ikjl} = \delta_{jl} \quad (\forall (j, l) \in [q]^2), \\ \gamma_{ik} &\geq 0 \quad (\forall (i, k) \in [q]^2), \quad \delta_{jl} \geq 0 \quad (\forall (j, l) \in [q]^2), \quad y_{ikjl} \geq 0 \quad (\forall (i, k, j, l) \in [q]^4). \end{aligned} \quad (7)$$

The first line in (6) accounts for the cardinality of $\Sigma^{\alpha,\beta} \times \Sigma^{\alpha,\beta}$, while the second line is $\mathbf{E}_G[w_G(\sigma_1)w_G(\sigma_2)]$ for $(\sigma_1, \sigma_2) \in \Sigma^{\alpha,\beta} \times \Sigma^{\alpha,\beta}$ with the prescribed γ, δ . Since the weight of a configuration is multiplicative over the edges and the matchings are independent, $\mathbf{E}_G[w_G(\sigma_1)w_G(\sigma_2)]$ is the Δ -power of the expected weight of a single matching. The latter is completely determined by \mathbf{y} and is equal to $\prod_{i,k,j,l} (B_{ij}B_{kl})^{y_{ikjl}}$, scaled by the probability that the matching induces the prescribed \mathbf{y} .

Remark 2.1. Note that (6) shows that the second moment can be interpreted as the first moment of a paired-spin model with interaction matrix $\mathbf{B} \otimes \mathbf{B}$. Indeed, we can interpret $B_{ij}B_{kl}$ as the activity between the paired spins (i, k) and (j, l) , thus giving the desired alignment.

The sums in (4) and (6) are typically exponential in n . The most critical component of our arguments is to find the quantitative structure of configurations which determine the exponential

order of the moments. Formally, we study the limits of $\frac{1}{n} \log \mathbf{E}_{\mathcal{G}}[Z_G^{\alpha, \beta}]$ and $\frac{1}{n} \log \mathbf{E}_{\mathcal{G}}[(Z_G^{\alpha, \beta})^2]$ as $n \rightarrow \infty$. Under the usual conventions that $\ln 0 \equiv -\infty$ and $0 \ln 0 \equiv 0$, standard application of Stirling's approximation yields the following:

$$\begin{aligned} \Psi_1(\alpha, \beta) &:= \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbf{E}_{\mathcal{G}}[Z_G^{\alpha, \beta}] = \max_{\mathbf{x}} \Upsilon_1(\alpha, \beta, \mathbf{x}), \\ \text{where } \Upsilon_1(\alpha, \beta, \mathbf{x}) &:= (\Delta - 1)f_1(\alpha, \beta) + \Delta g_1(\mathbf{x}) \\ f_1(\alpha, \beta) &:= \sum_i \alpha_i \ln \alpha_i + \sum_j \beta_j \ln \beta_j \\ g_1(\mathbf{x}) &:= \sum_{i,j} x_{ij} \ln B_{ij} - \sum_{i,j} x_{ij} \ln x_{ij}. \end{aligned} \tag{8}$$

And for the second moment:

$$\begin{aligned} \Psi_2(\alpha, \beta) &:= \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbf{E}_{\mathcal{G}}[(Z_G^{\alpha, \beta})^2] = \max_{\gamma, \delta} \max_{\mathbf{y}} \Upsilon_2(\gamma, \delta, \mathbf{y}), \\ \text{where } \Upsilon_2(\gamma, \delta, \mathbf{y}) &:= (\Delta - 1)f_2(\gamma, \delta) + \Delta g_2(\mathbf{y}) \\ f_2(\gamma, \delta) &:= \sum_{i,k} \gamma_{ik} \ln \gamma_{ik} + \sum_{j,l} \delta_{jl} \ln \delta_{jl} \\ g_2(\mathbf{y}) &:= \sum_{i,k,j,l} y_{ikjl} \ln(B_{ij} B_{kl}) - \sum_{i,k,j,l} y_{ikjl} \ln y_{ikjl} \end{aligned} \tag{9}$$

The functions Υ_1 and Υ_2 are defined on the regions (5) and (7), respectively. We also relax the integrality constraints of the vectors $\alpha, \beta, \mathbf{x}$ and $\gamma, \delta, \mathbf{y}$ which were imposed by the expressions (4) and (6). This does not affect our considerations in the limit $n \rightarrow \infty$. Moreover, note that the function Υ_2 depends on α, β due to the linear constraints (7). This dependence is omitted above since we are going to study the second moment for α, β fixed to some well chosen vectors.

The limits (8) and (9) can be justified using standard Laplace arguments (see for example [dB81, Chapter 4]).

Remark 2.2. *The maximization in the first moment depends only on the function $g_1(\mathbf{x})$ which is strictly concave in the convex region where it is defined. Hence, for any fixed α, β , the global maximum of $\Upsilon_1(\alpha, \beta, \mathbf{x})$ with respect to \mathbf{x} is achieved at a unique point. Similarly, for any fixed γ, δ , the maximum of $\Upsilon_2(\gamma, \delta, \mathbf{y})$ with respect to \mathbf{y} is achieved at a unique point. Crucially for the calculation of the asymptotics of the second moment in Appendix B, if α, β are global maximizers of Ψ_1 , the global maximum of $\Upsilon_2(\gamma, \delta, \mathbf{y})$ with respect to $\gamma, \delta, \mathbf{y}$ is also achieved at a unique point, see Lemma 3.2 in Section 3.4.*

A notational convention that we have adopted silently so far is perhaps useful to allude: the indices i, k “point” to the set V_1 , while indices j, l “point” to the set V_2 .

3 Second Moment Analysis

In this section we prove Theorem 1.4. We first present some basic definitions concerning matrix norms. We then show that the maximum of the first moment function Ψ_1 can be reformulated in terms of matrix norms. This then enables a short proof of Theorem 1.4.

3.1 Basic Definitions: Matrix Norms

We will reformulate the maxima of the first and second moments in terms of matrix norms. We first recall the basic definitions regarding matrix norms. The usual vector norms are denoted as:

$$\|\mathbf{x}\|_p = \left(\sum_{i=1}^n x_i^p \right)^{1/p}.$$

We will use the subordinate matrix norm (also known as the induced matrix norm) which will be denoted as $\|\cdot\|_{p \rightarrow q}$ and is defined as:

$$\|\mathbf{A}\|_{p \rightarrow q} = \max_{\|\mathbf{x}\|_p=1} \|\mathbf{A}\mathbf{x}\|_q.$$

Note that if \mathbf{A} has non-negative entries then one can restrict the maximization to \mathbf{x} with non-negative entries. A well-known example of an induced norm is the spectral norm $\|\cdot\|_{2 \rightarrow 2}$.

3.2 Reformulating the First Moment in Terms of Matrix Norms

A key component in the analysis of the second moment is the following function Φ . Let $p = \Delta/(\Delta - 1)$. For non-negative \mathbf{r}, \mathbf{c} , define $\Phi(\mathbf{r}, \mathbf{c})$ by:

$$\exp(\Phi(\mathbf{r}, \mathbf{c})/\Delta) = \frac{\mathbf{r}^\top \mathbf{B} \mathbf{c}}{\|\mathbf{r}\|_p \|\mathbf{c}\|_p}.$$

We will show that the critical points of Φ and Ψ_1 match in the sense that there is a one-to-one correspondence between them and their values are equal at the corresponding critical points. The full statement is contained in Theorem 4.1 in Section 4.1, but the important element for the current discussion is captured in the following lemma:

Lemma 3.1.

$$\max_{\alpha, \beta \in \Delta_q} \Psi_1(\alpha, \beta) = \max_{\mathbf{r}, \mathbf{c}} \Phi(\mathbf{r}, \mathbf{c}).$$

Therefore, to determine the dominant phases of Ψ_1 it suffices to study Φ . The maximum of Φ can be compactly expressed in terms of matrix norms as follows:

$$\max_{\mathbf{r}, \mathbf{c}} \exp(\Phi(\mathbf{r}, \mathbf{c})/\Delta) = \max_{\mathbf{c}} \max_{\mathbf{r}} \frac{\mathbf{r}^\top \mathbf{B} \mathbf{c}}{\|\mathbf{r}\|_p \|\mathbf{c}\|_p} = \max_{\mathbf{c}} \frac{\|\mathbf{B} \mathbf{c}\|_\Delta}{\|\mathbf{c}\|_p} = \|\mathbf{B}\|_{p \rightarrow \Delta}, \quad (10)$$

where the second equality follows from matrix norm duality.

Hence, the dominant phases of Ψ_1 can be expressed in terms of matrix norms:

$$\max_{\alpha, \beta \in \Delta_q} \exp(\Psi_1(\alpha, \beta)/\Delta) = \|\mathbf{B}\|_{\frac{\Delta}{\Delta-1} \rightarrow \Delta}. \quad (11)$$

3.3 Analyzing the Second Moment: Proof of Theorem 1.4

To analyze the second moment function Ψ_2 we will reduce it to the first moment optimization in the following manner. The key observation is that the associated optimization for the second moment is equivalent to a first moment optimization of a “paired-spin” model which is specified by the tensor product of the original interaction matrix with itself. This property enables us to relate the maximum for the second moment calculations with the maximum of the first moment calculations.

Proof of Theorem 1.4. The second moment considers a pair of configurations, say σ and σ' , which are constrained to have a given phase α for V_1 and β for V_2 , where $V = V_1 \cup V_2$. We capture this constraint using a pair of vectors γ, δ corresponding to the overlap between σ and σ' , in particular, γ_{ij} (and δ_{ij}) is the number of vertices in V_1 (and V_2 , respectively) with spin i in σ and spin j in σ' .

Recall, $\Psi_1^{\mathbf{B}}$ indicates the dependence of the function Ψ_1 on the interaction matrix \mathbf{B} ; to simplify the notation we will drop the exponent if it is \mathbf{B} . We have (see Remark 2.1 in Section 2 for more details on this connection)

$$\Psi_2(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \max_{\boldsymbol{\gamma}, \boldsymbol{\delta}} \Psi_1^{\mathbf{B} \otimes \mathbf{B}}(\boldsymbol{\gamma}, \boldsymbol{\delta}), \quad (12)$$

where the optimization in (12) is constrained to $\boldsymbol{\gamma}$ and $\boldsymbol{\delta}$ such that

$$\sum_i \gamma_{ik} = \alpha_k, \quad \sum_k \gamma_{ik} = \alpha_i, \quad \sum_j \delta_{jl} = \beta_l \quad \text{and} \quad \sum_\ell \delta_{j\ell} = \beta_j. \quad (13)$$

Ignoring the four constraints in (13) can only increase the value of (12) and hence

$$\max_{\boldsymbol{\alpha}, \boldsymbol{\beta}} \exp(\Psi_2(\boldsymbol{\alpha}, \boldsymbol{\beta})/\Delta) \leq \max_{\boldsymbol{\gamma}, \boldsymbol{\delta}} \exp(\Psi_1^{\mathbf{B} \otimes \mathbf{B}}(\boldsymbol{\gamma}, \boldsymbol{\delta})/\Delta) = \|\mathbf{B} \otimes \mathbf{B}\|_{\frac{\Delta}{\Delta-1} \rightarrow \Delta}. \quad (14)$$

The key fact we now use is that for induced norms $\|\cdot\|_{p \rightarrow q}$ with $p \leq q$ it holds (c.f., [Ben77, Proposition 10.1]) that:

$$\|\mathbf{B} \otimes \mathbf{B}\|_{p \rightarrow q} = \|\mathbf{B}\|_{p \rightarrow q} \|\mathbf{B}\|_{p \rightarrow q}. \quad (15)$$

Therefore,

$$\max_{\boldsymbol{\alpha}, \boldsymbol{\beta}} \Psi_2(\boldsymbol{\alpha}, \boldsymbol{\beta}) \leq 2\Delta \log \|\mathbf{B}\|_{\frac{\Delta}{\Delta-1} \rightarrow \Delta} = 2 \max_{\boldsymbol{\alpha}, \boldsymbol{\beta}} \Psi_1(\boldsymbol{\alpha}, \boldsymbol{\beta}). \quad (16)$$

To complete the proof of Theorem 1.4 it just remains to prove the reverse inequality, which follows from the fact that $\mathbf{E}[X^2] \geq \mathbf{E}[X]^2$. \square

3.4 Optimal second moment configuration

We will need more detailed information about the $\boldsymbol{\gamma}, \boldsymbol{\delta}$ which achieve equality in Theorem 1.4 and equation (12). The following lemma is true whenever \mathbf{B} is regular (and hence for antiferromagnetic models as well, cf. Definition 1.3). Roughly, the lemma captures that the major contribution to the second moment comes from pairs of configurations which are uncorrelated. This is crucial to calculate the asymptotics of the second moment in Appendix B.

Lemma 3.2. *Assume that \mathbf{B} is regular. The $\boldsymbol{\gamma}, \boldsymbol{\delta}$ for which the equality in*

$$\max_{\boldsymbol{\alpha}, \boldsymbol{\beta}} \max_{\boldsymbol{\gamma}, \boldsymbol{\delta} \text{ satisfying (13)}} \Psi_1^{\mathbf{B} \otimes \mathbf{B}}(\boldsymbol{\gamma}, \boldsymbol{\delta}) = \max_{\boldsymbol{\alpha}, \boldsymbol{\beta}} \Psi_1^{\mathbf{B}}(\boldsymbol{\alpha}, \boldsymbol{\beta}), \quad (17)$$

is achieved satisfy (for all $i, j, k, l \in [q]$)

$$\gamma_{ik} = \alpha_i \alpha_k \quad \text{and} \quad \delta_{jl} = \beta_j \beta_l. \quad (18)$$

Proof. We will have to dig in to the proof of (15) and use (13). Bennett's proof of (15) is the following (our particular values are $q' = \Delta$ and $p = \Delta/(\Delta - 1)$):

$$\begin{aligned} \|(\mathbf{B} \otimes \mathbf{B})\mathbf{r}\|_{q'} &= \left(\sum_k \sum_i \left| \sum_j B_{ij} \sum_l B_{kl} R_{jl} \right|^{q'} \right)^{1/q'} \\ &\leq \|\mathbf{B}\|_{p \rightarrow q'} \left(\sum_k \left(\sum_j \left| \sum_l B_{kl} R_{jl} \right|^p \right)^{q'/p} \right)^{1/q'} \\ &\leq \|\mathbf{B}\|_{p \rightarrow q'} \left(\sum_j \left(\sum_k \left| \sum_l B_{kl} R_{jl} \right|^{q'} \right)^{p/q'} \right)^{1/p} \\ &\leq \|\mathbf{B}\|_{p \rightarrow q'}^2 \left(\sum_{j,l} R_{jl}^p \right)^{1/p}. \end{aligned}$$

Note that in the last inequality one uses $\|\mathbf{B} \mathbf{r}\|_{q'} \leq \|\mathbf{B}\|_{p \rightarrow q'} \|\mathbf{r}\|_p$, applied to the vectors $\mathbf{r}'_j := (R_{j1}, R_{j2}, \dots, R_{jq})$, for $j = 1, \dots, q$. Thus if \mathbf{r} is a maximizer of

$$\max_{\mathbf{r}} \frac{\|(\mathbf{B} \otimes \mathbf{B})\mathbf{r}\|_{q'}}{\|\mathbf{r}\|_p}, \quad (19)$$

then the vectors \mathbf{r}'_j are maximizers of

$$\max_{\mathbf{r}'} \frac{\|\mathbf{B} \mathbf{r}'\|_{q'}}{\|\mathbf{r}'\|_p}. \quad (20)$$

The same, by symmetry, applies to $\mathbf{r}''_l := (R_{1l}, R_{2l}, \dots, R_{ql})$, for $l = 1, \dots, q$.

The second inequality in Bennett's proof is Minkowski's inequality applied to vectors $\mathbf{B}\mathbf{r}'_1, \dots, \mathbf{B}\mathbf{r}'_q$. The equality is achieved only if $\mathbf{B}\mathbf{r}'_1, \dots, \mathbf{B}\mathbf{r}'_q$ generate space of dimension one, and since \mathbf{B} is regular we have also that $\mathbf{r}'_1, \dots, \mathbf{r}'_q$ generate space of dimension one. Hence, for a maximizer \mathbf{r} of (19) we have $\mathbf{r} = \mathbf{r}' \otimes \mathbf{r}''$, where \mathbf{r}' and \mathbf{r}'' are maximizers of (20). By Theorem 4.1 (equation (23)) we then have

$$\gamma_{ik} = \alpha'_i \alpha''_k \quad (21)$$

for the corresponding maximizers of $\Psi_1^{\mathbf{B} \otimes \mathbf{B}}(\gamma, \delta)$ and $\Psi_1^{\mathbf{B}}(\alpha, \beta)$. Equation (21) together with constraints

$$\sum_i \gamma_{ik} = \alpha_k \quad \text{and} \quad \sum_k \gamma_{ik} = \alpha_i,$$

from (13) imply $\gamma_{ik} = \alpha_i \alpha_k$ (since $\alpha_k = \sum_i \gamma_{ik} = \sum_i \alpha'_i \alpha''_k = \alpha''_k$ and similarly $\alpha_i = \alpha'_i$). The proof of $\delta_{jl} = \beta_j \beta_l$ is analogous. \square

4 Tree recursions, first moment, and matrix norms

The second moment results of the previous section will be used to establish that, with probability $1 - o(1)$ over the choice of a random Δ -regular bipartite graph, the Gibbs distribution has most of its mass on configurations whose spin frequencies on the two sides of the graph are (close to) dominant phases. To do this, it will be important to examine dominant phases, i.e., the maxima of $\Psi_1(\alpha, \beta)$ and, further, to characterize the local maxima. We will use this information to connect the functions Φ and Ψ_1 and thus prove Lemma 3.1 which was the critical component in the second moment analysis; in fact, Lemma 3.1 is an immediate corollary of the upcoming Theorem 4.1 which details further the connection between Φ and Ψ_1 .

For a spin system with interaction matrix \mathbf{B} , the following recursions are relevant for the analysis of the critical points of Ψ_1 .

$$\hat{R}_i \propto \left(\sum_{j=1}^q B_{ij} C_j \right)^{\Delta-1} \quad \text{and} \quad \hat{C}_j \propto \left(\sum_{i=1}^q B_{ij} R_i \right)^{\Delta-1}. \quad (22)$$

We refer to (22) as *tree recursions* since they emerge naturally in the analysis of spin systems on the infinite Δ -regular tree \mathbb{T}_Δ . More precisely, the fixpoints of the tree recursions correspond to semi-translation invariant Gibbs measures on \mathbb{T}_Δ (fixpoints of (22) are those R_i 's and C_j 's such that $\hat{R}_i \propto R_i$ and $\hat{C}_j \propto C_j$, for all $i, j \in [q]$). The fixpoints of the tree recursions correspond to critical points of Ψ_1 , as was first observed in [MWW09], see Section 4.1.2 for a derivation in our setting.

We prove the following result which connects tree recursions, the function Φ and the function Ψ_1 . Lemma 3.1 is a corollary of the following more general theorem.

Theorem 4.1. *There is a one-to-one correspondence between the fixpoints of the tree recursions and the critical points of Φ (both considered for $R_i \geq 0, C_j \geq 0$ in the projective space, that is, up to scaling by a constant).*

The following transformation $(\mathbf{r}, \mathbf{c}) \mapsto (\boldsymbol{\alpha}, \boldsymbol{\beta})$ given by:

$$\alpha_i = \frac{R_i^{\Delta/(\Delta-1)}}{\sum_i R_i^{\Delta/(\Delta-1)}} \quad \text{and} \quad \beta_j = \frac{C_j^{\Delta/(\Delta-1)}}{\sum_j C_j^{\Delta/(\Delta-1)}} \quad (23)$$

yields a one-to-one-to-one correspondence between the critical points of Φ and the critical points of Ψ_1 (in the region defined by $\alpha_i \geq 0, \beta_j \geq 0$ and $\sum_i \alpha_i = 1, \sum_j \beta_j = 1$).

Moreover, for the corresponding critical points (\mathbf{r}, \mathbf{c}) and $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ one has

$$\Phi(\mathbf{r}, \mathbf{c}) = \Psi_1(\boldsymbol{\alpha}, \boldsymbol{\beta}). \quad (24)$$

Finally, for spin systems whose interaction matrix \mathbf{B} is ergodic, the local maxima of Φ and Ψ_1 happen at the critical points (that is, there are no local maxima on the boundary).

To argue that the Gibbs distribution places most of its mass on configurations whose spin frequencies are given by dominant phases, we need a more explicit handle on *local maxima* of Ψ_1 . The latter will also be crucial to analyze the global maxima of Ψ_1 for specific models of interest.

We connect local maxima of Ψ_1 to attractive fixpoints of the associated tree recursions. Specifically, we call a fixpoint x of a function f a *Jacobian attractive fixpoint* if the Jacobian of f at x has spectral radius less than 1. We say that a critical point $\boldsymbol{\alpha}, \boldsymbol{\beta}$ is a *Hessian local maximum* if the Hessian of Ψ_1 at $\boldsymbol{\alpha}, \boldsymbol{\beta}$ is negative definite (note this is a sufficient condition for $\boldsymbol{\alpha}, \boldsymbol{\beta}$ to be a local maximum).

We prove the following theorem in Section 4.2.

Theorem 4.2. *Jacobian attractive fixpoints of the tree recursions (22) (considered as a function $(R_1, \dots, R_q, C_1, \dots, C_q) \mapsto (\hat{R}_1, \dots, \hat{R}_q, \hat{C}_1, \dots, \hat{C}_q)$) correspond to Hessian local maxima of Ψ_1 .*

Theorem 4.2 is important for analyzing the global maxima of Ψ_1 for colorings and antiferromagnetic Potts model (see Section 7). Moreover, it will be used to apply the small subgraph conditioning method (see Section A.2).

4.1 Connection between Φ and Ψ_1

In this section, we prove Theorem 4.1.

4.1.1 Preliminaries on maximum-entropy distributions

Let $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ be non-negative vectors in \mathbb{R}^q such that

$$\sum_i \alpha_i = 1 \quad \text{and} \quad \sum_j \beta_j = 1. \quad (25)$$

For $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ that satisfy (25) let

$$g(\alpha_1, \dots, \alpha_q, \beta_1, \dots, \beta_q) = \max \sum_{i=1}^q \sum_{j=1}^q x_{ij} (\ln(B_{ij}) - \ln x_{ij}), \quad (26)$$

where the maximum is taken over non-negative x_{ij} 's such that

$$\alpha_i = \sum_j x_{ij} \quad \text{and} \quad \beta_j = \sum_i x_{ij}. \quad (27)$$

Lemma 4.3. *The maximum of the right-hand-side of (26) is achieved at unique x_{ij} . The x_{ij} are given by*

$$x_{ij} = B_{ij}R_iC_j, \quad (28)$$

where \mathbf{r} and \mathbf{c} satisfy

$$R_i \sum_{j=1}^q B_{ij}C_j = \alpha_i \quad \text{and} \quad C_j \sum_{i=1}^q B_{ij}R_i = \beta_j, \quad (29)$$

and

$$\begin{aligned} \sum_{j=1}^q B_{ij}C_j = 0 &\implies R_i = 0; \\ \sum_{i=1}^q B_{ij}R_i = 0 &\implies C_j = 0. \end{aligned} \quad (30)$$

The value of g , in terms of R_i 's and C_j 's, is given by

$$g(\alpha_1, \dots, \alpha_q, \beta_1, \dots, \beta_q) = - \sum_{i=1}^q \sum_{j=1}^q B_{ij}R_iC_j \ln(R_iC_j). \quad (31)$$

Proof. From strict concavity of $-x \ln x$ it follows that the right-hand side of (26) has a unique critical point (if there were two critical points then the segment between the points lies in the linear space defined by (27); the function has a zero derivative on both ends of the segment; and the second derivative of the function is negative on the segment; a contradiction).

Using the method of Lagrange multipliers we obtain that the critical points of the right-hand side of (26) are x_{ij} given by (28) where R_i 's and C_j 's are solutions of (29). We can make any solution of (29) satisfy (30): if $\sum_{j=1}^q B_{ij}C_j = 0$ then set $R_i = 0$ (and symmetrically, if $\sum_{i=1}^q B_{ij}R_i = 0$ then set $C_j = 0$). We now argue that this change does not violate (29). Suppose that after the change for some $k \in [q]$ we have

$$R_k \sum_{j=1}^q B_{kj}C_j \neq \alpha_k. \quad (32)$$

Then $i = k$ (since only R_i changed) and since $\sum_{j=1}^q B_{ij}C_j = 0$ we also have $\alpha_i = 0$, a contradiction (with (32)). Now suppose that after the change for some $j \in [q]$ we have

$$C_j \sum_{k=1}^q B_{kj}R_k \neq \beta_j. \quad (33)$$

Then $B_{ij} > 0$ and $C_j > 0$ (otherwise changing R_i would not violate (33)). This then implies $\sum_{j=1}^q B_{ij}C_j > B_{ij}C_j > 0$, a contradiction. Thus the change does not violate (29).

Equation (31) is obtained by substituting (29) into (26). \square

Remark 4.4. *Scaling all the R_i 's up by the same factor while scaling all the C_j 's down by the same factor preserves (28) and (29). Modulo such scaling the R_i 's and C_j 's are unique, since the x_{ij} 's are unique and (28) determines the R_i 's and C_j 's once one value (say R_1) is fixed (here we use the fact that the matrix of the model is ergodic).*

Remark 4.5. Note that the condition (25) translates (using (29)) into the following condition on R_i 's and C_j 's

$$\sum_{i=1}^q \sum_{j=1}^q B_{ij} R_i C_j = 1. \quad (34)$$

Our goal now is to see how the value of (26) changes when we perturb α_i 's and β_j 's. We are going to view them as functions of a new variable z . All differentiation in this section will be with respect to z . Note that to stay in the subspace defined by (25) we should have, in particular,

$$\sum_i \alpha'_i = \sum_i \alpha''_i = 0 \quad \text{and} \quad \sum_j \beta'_j = \sum_j \beta''_j = 0. \quad (35)$$

Differentiating (29) we obtain

$$\sum_{j=1}^q B_{ij} (R_i C_j)' = \alpha'_i \quad \text{and} \quad \sum_{i=1}^q B_{ij} (R_i C_j)' = \beta'_j. \quad (36)$$

The following ratio of (29) and (36) will be useful later

$$\frac{\alpha'_i}{\alpha_i} = \frac{R'_i}{R_i} + \frac{\sum_{j=1}^q B_{ij} C'_j}{\sum_{j=1}^q B_{ij} C_j} \quad \text{and} \quad \frac{\beta'_j}{\beta_j} = \frac{C'_j}{C_j} + \frac{\sum_{i=1}^q B_{ij} R'_i}{\sum_{i=1}^q B_{ij} R_i}. \quad (37)$$

The scaling freedom for R_i 's and C_j 's (discussed in Remark 4.4) is equivalent to increasing all R_i/R_i 's by the same (additive) amount and decreasing all C_j/C_j 's by the same (additive) amount. We are going to remove this freedom by requiring

$$\sum_{i=1}^q \alpha_i \frac{R'_i}{R_i} = \sum_{j=1}^q \beta_j \frac{C'_j}{C_j}. \quad (38)$$

(Recall that we study the effect of perturbing g when we change α_i 's and β_j 's; equation (38) just fixes the corresponding change in R_i 's and C_j 's.)

Now we compute the derivatives of g .

Lemma 4.6. *We have*

$$g' = - \sum_{i=1}^q (\ln R_i) \alpha'_i - \sum_{j=1}^q (\ln C_j) \beta'_j, \quad (39)$$

$$g'' = - \sum_{i=1}^q \frac{R'_i}{R_i} \alpha'_i - \sum_{j=1}^q \frac{C'_j}{C_j} \beta'_j - \sum_{i=1}^q (\ln R_i) \alpha''_i - \sum_{j=1}^q (\ln C_j) \beta''_j. \quad (40)$$

Proof. Using $(f \ln f)' = (1 + \ln f)f'$ and equations (36) and (35) we obtain

$$g' = - \sum_{i=1}^q \sum_{j=1}^q B_{ij} (1 + \ln(R_i C_j)) (R_i C_j)' = - \sum_{i=1}^q (\ln R_i) \alpha'_i - \sum_j (\ln C_j) \beta'_j.$$

Differentiating (39) we obtain (40). □

Note the expressions (39) and (40) are independent of the choice of scaling of R_i 's and C_j 's (this follows from (35)). The particular tying of R'_i/R_i 's and C'_j/C_j 's to α'_i and β'_j (given by (38)) will be useful later.

4.1.2 Critical points of Ψ_1 and the tree recursions

In this section we establish the connection between the critical points of Ψ_1 and the fixpoints of the tree recursions.

Lemma 4.7. *Let α, β be a critical point of $\Psi_1(\alpha, \beta)$ in the subspace defined by (25). Let \mathbf{r}, \mathbf{c} be given by (29). Then*

$$\alpha_i \propto R_i^{\Delta/(\Delta-1)} \quad \text{and} \quad \beta_j \propto C_j^{\Delta/(\Delta-1)}. \quad (41)$$

Consequently, \mathbf{r}, \mathbf{c} satisfy the tree recursions stated in Section 4:

$$R_i \propto \left(\sum_{j=1}^q B_{ij} C_j \right)^{\Delta-1} \quad \text{and} \quad C_j \propto \left(\sum_{i=1}^q B_{ij} R_i \right)^{\Delta-1}. \quad (22)$$

Proof. At the critical points of Ψ the first derivative of Ψ has to vanish for all α'_i 's and β'_j 's from the subspace defined by (35), that is,

$$\begin{aligned} \Psi' &= (\Delta - 1) \left(\sum_{i=1}^q (1 + \ln \alpha_i) \alpha'_i + \sum_{j=1}^q (1 + \ln \beta_j) \beta'_j \right) - \Delta \left(\sum_{i=1}^q (\ln R_i) \alpha'_i + \sum_{j=1}^q (\ln C_j) \beta'_j \right) \\ &= \sum_{i=1}^q ((\Delta - 1)(1 + \ln \alpha_i) - \Delta \ln R_i) \alpha'_i + \sum_{j=1}^q ((\Delta - 1)(1 + \ln \beta_j) - \Delta \ln C_j) \beta'_j = 0, \end{aligned} \quad (42)$$

where the R_i 's and C_j 's are given by (29). Inspecting (42) we see that $(\Delta - 1)(1 + \ln \alpha_i) - \Delta \ln R_i$ have the same value. Indeed, if two of them, say with indices i_1, i_2 , had different values then we could increase α_{i_1} and decrease α_{i_2} by the same infinitesimal amount and violate (42). Similarly, $(\Delta - 1)(1 + \ln \beta_j) - \Delta \ln C_j$ have the same value and hence we have (41). Plugging (41) into (29) one obtains (22). \square

Lemma 4.8. *Let (\mathbf{r}, \mathbf{c}) be a solution of the tree recursions (22). Let (α, β) be given by (23). Then (α, β) is a critical point of $\Psi_1(\alpha, \beta)$ in the subspace defined by (25).*

Proof. Let

$$Z_R := (\Delta - 1)(1 + \ln \alpha_i) - \Delta \ln R_i = (\Delta - 1) \left(1 - \ln \sum_{i=1}^q R_i^{(\Delta+1)/\Delta} \right),$$

where the second equality follows from (23). Note that Z_R is independent of the choice of i . Similarly let

$$Z_C := (\Delta - 1)(1 + \ln \beta_j) - \Delta \ln C_j = (\Delta - 1) \left(1 - \ln \sum_{j=1}^q C_j^{(\Delta+1)/\Delta} \right).$$

For perturbations of α, β in the subspace given by (25) we have

$$\begin{aligned} \Psi'_1(\alpha, \beta) &= \sum_{i=1}^q ((\Delta - 1)(1 + \ln \alpha_i) - \Delta \ln R_i) \alpha'_i + \sum_{j=1}^q ((\Delta - 1)(1 + \ln \beta_j) - \Delta \ln C_j) \beta'_j \\ &= Z_R \sum_{i=1}^q \alpha'_i + Z_C \sum_{j=1}^q \beta'_j = 0, \end{aligned}$$

and hence (α, β) is a critical point. \square

4.1.3 Value of Ψ_1 at the critical points

Lemma 4.9. *Let (α, β) be critical point of $\Psi_1(\alpha, \beta)$. Let (\mathbf{r}, \mathbf{c}) be given by (29). Then*

$$\Phi(\mathbf{r}, \mathbf{c}) = \Psi_1(\alpha, \beta). \quad (24)$$

Moreover, (\mathbf{r}, \mathbf{c}) is a critical point of $\Phi(\mathbf{r}, \mathbf{c})$.

Proof. We have (see equation (31))

$$\Psi_1(\alpha, \beta) = (\Delta - 1) \left(\sum_{i=1}^q \alpha_i \ln \alpha_i + \sum_{j=1}^q \beta_j \ln \beta_j \right) - \Delta \sum_{i=1}^q \sum_{j=1}^q B_{ij} R_i C_j \ln(R_i C_j). \quad (43)$$

At the critical points we have (see equation (41))

$$\alpha_i = \frac{R_i^{\Delta/(\Delta-1)}}{\sum_{i=1}^q R_i^{\Delta/(\Delta-1)}} \quad \text{and} \quad \beta_j = \frac{C_j^{\Delta/(\Delta-1)}}{\sum_{j=1}^q C_j^{\Delta/(\Delta-1)}}. \quad (44)$$

Plugging (29) into (43) we obtain

$$\begin{aligned} \Psi_1(\alpha, \beta) &= (\Delta - 1) \left(\sum_{i=1}^q \alpha_i \ln \alpha_i + \sum_{j=1}^q \beta_j \ln \beta_j \right) - \Delta \left(\sum_{i=1}^q \alpha_i \ln R_i + \sum_{j=1}^q \beta_j \ln C_j \right) = \\ &= \sum_{i=1}^q \alpha_i \ln \frac{\alpha_i^{\Delta-1}}{R_i^\Delta} + \sum_{j=1}^q \beta_j \ln \frac{\beta_j^{\Delta-1}}{C_j^\Delta} = -(\Delta - 1) \left[\ln \left(\sum_{i=1}^q R_i^{\Delta/(\Delta-1)} \right) + \ln \left(\sum_{j=1}^q C_j^{\Delta/(\Delta-1)} \right) \right], \end{aligned} \quad (45)$$

where in the last equality we used (44) and the fact that α_i 's and β_j 's sum to 1. Recall that

$$\sum_{i=1}^q \sum_{j=1}^q B_{ij} R_i C_j = \sum_{i=1}^q \alpha_i = 1, \quad (46)$$

and hence the following is obtained by adding zero to the right-hand side of (45)

$$\begin{aligned} \Psi_1(\alpha, \beta) &= \Delta \ln \left(\sum_{i=1}^q \sum_{j=1}^q B_{ij} R_i C_j \right) - (\Delta - 1) \left[\ln \left(\sum_{i=1}^q R_i^{\Delta/(\Delta-1)} \right) + \ln \left(\sum_{j=1}^q C_j^{\Delta/(\Delta-1)} \right) \right] \\ &= \Phi(\mathbf{r}, \mathbf{c}). \end{aligned}$$

Now we argue that (\mathbf{r}, \mathbf{c}) is a critical point of $\Phi(\mathbf{r}, \mathbf{c})$. We have

$$\frac{\partial}{\partial R_i} \Phi(\mathbf{r}, \mathbf{c}) = \Delta \frac{\sum_{j=1}^q B_{ij} C_j}{\sum_{i=1}^q \sum_{j=1}^q B_{ij} R_i C_j} - (\Delta - 1) \frac{\frac{\Delta}{\Delta-1} R_i^{1/(\Delta-1)}}{\sum_{i=1}^q R_i^{\Delta/(\Delta-1)}} = \Delta \frac{\alpha_i}{R_i} - \Delta \frac{\alpha_i}{R_i} = 0. \quad (47)$$

where we used (44), (29), and (25). The same argument yields

$$\frac{\partial}{\partial C_j} \Phi(\mathbf{r}, \mathbf{c}) = \Delta \frac{\sum_{i=1}^q B_{ij} R_i}{\sum_{i=1}^q \sum_{j=1}^q B_{ij} R_i C_j} - (\Delta - 1) \frac{\frac{\Delta}{\Delta-1} C_j^{1/(\Delta-1)}}{\sum_{j=1}^q C_j^{\Delta/(\Delta-1)}} = 0. \quad (48)$$

and hence \mathbf{r}, \mathbf{c} is a critical point of Φ . \square

Lemma 4.10. *Let (\mathbf{r}, \mathbf{c}) be a critical point of $\Phi(\mathbf{r}, \mathbf{c})$. Let α, β be given by (23). Then α, β is a critical point of $\Psi_1(\alpha, \beta)$ in the subspace defined by (25).*

Proof. At a critical point of Φ we have that (47) is zero for $i \in [q]$. Note that the denominators do not depend on i and hence we have

$$R_i^{1/(\Delta-1)} \propto \sum_{j=1}^q B_{ij} C_j.$$

Similarly, from (48) we obtain

$$C_j^{1/(\Delta-1)} \propto \sum_{i=1}^q B_{ij} R_i.$$

Hence (\mathbf{r}, \mathbf{c}) satisfy the tree recursions. Now we use Lemma 4.8 to conclude that (α, β) is a critical point of $\Psi_1(\alpha, \beta)$ in the subspace defined by (25). \square

4.1.4 Local maxima of Ψ_1 are in the interior

In this section we show that for models with ergodic (irreducible and aperiodic) interaction matrix \mathbf{B} the maximum of $\Phi(\mathbf{r}, \mathbf{c})$ is achieved in the interior. A symmetric matrix is irreducible if the graph whose edges correspond to non-zero edges of \mathbf{B} is connected. A symmetric matrix is aperiodic if the graph whose edges correspond to non-zero edges of \mathbf{B} has an odd cycle.

Lemma 4.11. *Assume that \mathbf{B} is ergodic. Let $(\mathbf{r}, \mathbf{c}) \neq 0$ be a local maximum of Φ in the region $\mathbf{r}, \mathbf{c} \geq 0$. Then $R_i > 0$ for all $i \in [q]$ and $C_j > 0$ for all $j \in [q]$.*

Proof. Suppose not, that is, we have a maximum that has a zero on some coordinate of \mathbf{r} or \mathbf{c} . From the ergodicity of \mathbf{B} we have that there exist $i, j \in [q]$ such that i) $R_i = 0, C_j > 0$, and $B_{ij} > 0$ or ii) $R_i > 0, C_j = 0$, and $B_{ij} > 0$. (Suppose not. Let $Z_R \subseteq [q]$ be the set of i such that $R_i = 0$. Similarly let $Z_C \subseteq [q]$ be the set of j such that $C_j = 0$. If neither i) nor ii) happens then non-zero B_{ij} are possibly between $i \in Z_R$ and $j \in Z_C$ and $i \in [q] \setminus Z_R$ and $j \in [q] \setminus Z_C$. Thus in \mathbf{B}^2 the non-zero $(B^2)_{ij}$ are possibly between $i, j \in Z_R$ and $i, j \in [q] \setminus Z_R$. Thus \mathbf{B} is not ergodic.) W.l.o.g. assume that it is the case i) (the case ii) is handled analogously).

The derivative of Φ w.r.t. R_i is (we are using $R_i = 0$)

$$\frac{\partial}{\partial R_i} \Phi(\mathbf{r}, \mathbf{c}) = \Delta \frac{\sum_{j=1}^q B_{ij} C_j}{\sum_{i=1}^q \sum_{j=1}^q B_{ij} R_i C_j} > \Delta \frac{B_{ij} C_j}{\sum_{i=1}^q \sum_{j=1}^q B_{ij} R_i C_j} > 0,$$

and hence we are not at a maximum, a contradiction. \square

Lemma 4.12. *Assume that \mathbf{B} is ergodic. Let $\alpha, \beta \geq 0$ be a local maximum of $\Psi_1(\alpha, \beta)$ in the subspace defined by (25). Then $\alpha_i > 0$ for all $i \in [q]$ and $\beta_j > 0$ for all $j \in [q]$.*

Proof. It will be useful to view Ψ_1 as a function of (\mathbf{r}, \mathbf{c}) . Because of Lemma 4.3 we have (\mathbf{r}, \mathbf{c}) satisfying (46) and (30) (and any such (\mathbf{r}, \mathbf{c}) yields (α, β) satisfying (25)). We have (from (45))

$$\begin{aligned} \Psi_1(\alpha, \beta) &= \sum_{i=1}^q \sum_{j=1}^q B_{ij} R_i C_j \left((\Delta - 1) \left(\ln \left(\sum_{j=1}^q B_{ij} C_j \right) + \ln \left(\sum_{i=1}^q B_{ij} R_i \right) \right) - \ln R_i - \ln C_j \right) \\ &=: \hat{\Psi}_1(\mathbf{r}, \mathbf{c}). \end{aligned}$$

If \mathbf{r} has a zero coordinate then, by ergodicity of \mathbf{B} there exists $k, \ell \in [q]$ such that i) $R_k = 0, C_\ell > 0$, and $B_{k\ell} > 0$ or ii) $R_k > 0, C_\ell = 0$, and $B_{k\ell} > 0$ (see the argument in the proof of Lemma 4.11). W.l.o.g. it is the case i).

Note that we have

$$\frac{\partial}{\partial R_k} \sum_{i=1}^q \sum_{j=1}^q B_{ij} R_i C_j = \sum_{j=1}^q B_{kj} C_j \geq B_{k\ell} C_\ell > 0. \quad (49)$$

We have

$$\begin{aligned} \frac{\partial}{\partial R_k} \hat{\Psi}_1 &= \sum_{j=1}^q B_{kj} C_j \left((\Delta - 1) \ln \left(\sum_{i=1}^q B_{ij} R_i \right) - \ln C_j \right) \\ &\quad + \left((\Delta - 1) \ln \left(\sum_{j=1}^q B_{kj} C_j \right) - \ln R_k \right) \left(\sum_{j=1}^q B_{kj} C_j \right) + (\Delta - 2) \sum_{j=1}^q B_{kj} C_j. \end{aligned} \quad (50)$$

The first sum in (50) is finite since if $C_j > 0$ then $\sum_{i=1}^q B_{ij} R_i > 0$ (using (30)); if $C_j = 0$ then the contribution of the term to the sum is zero (we are using the usual convention $0 \ln 0 = 0$). The second term in (50) has value $+\infty$ since $\ln R_k = -\infty$ and (49). Finally, the last term in (50) is finite and hence we have $\frac{\partial}{\partial R_k} \hat{\Psi}_1 = +\infty$.

Recall that $C_\ell > 0$ and hence (using (30)):

$$\frac{\partial}{\partial C_\ell} \sum_{i=1}^q \sum_{j=1}^q B_{ij} R_i C_j = \sum_{i=1}^q B_{i\ell} R_i > 0. \quad (51)$$

Finally, we argue that $\frac{\partial}{\partial C_\ell} \hat{\Psi}_1$ is finite. We have (analogously to (50))

$$\begin{aligned} \frac{\partial}{\partial C_\ell} \hat{\Psi}_1 &= \sum_{i=1}^q B_{i\ell} C_i \left((\Delta - 1) \ln \left(\sum_{j=1}^q B_{ij} C_j \right) - \ln R_i \right) \\ &\quad + \left((\Delta - 1) \ln \left(\sum_{i=1}^q B_{i\ell} R_i \right) - \ln C_\ell \right) \left(\sum_{i=1}^q B_{i\ell} R_i \right) + (\Delta - 2) \sum_{i=1}^q B_{i\ell} R_i. \end{aligned} \quad (52)$$

The first and third term in (52) are finite by the same argument as for (50). In the second term we use (51) and $C_\ell > 0$.

Now we increase R_k by an infinitesimal amount and change C_ℓ to maintain (34) (and hence (25)). (This is possible because both C_ℓ and R_k change the value of (34), see equations (49) and (51).) This change will increase $\hat{\Psi}_1$ and hence Ψ_1 contradicting the local maximality of α, β . \square

4.1.5 Proof of Theorem 4.1

Proof of Theorem 4.1. Lemmas 4.7 and 4.8 give the connection between the critical points of Ψ_1 and the fixpoints of the tree recursions. Lemmas 4.9 and 4.10 give connection between the critical points of Ψ_1 and Φ and show that the values agree on the corresponding critical points. Finally, Lemmas 4.11 and 4.12 show that the maxima happen in the interior (that is, for $R_i > 0, C_j > 0$ in the case of Φ and for $\alpha_i > 0, \beta_j > 0$ in the case of Ψ_1). \square

4.2 Connecting Local Maxima and Stability of Tree Recursions

In this section we prove Theorem 4.2.

4.2.1 Maximum entropy configurations on random Δ -regular bipartite graphs

We analyze the critical points by looking at the second derivative. Using $(f \ln f)'' = (f')^2/f + (1 + \ln f)f''$ we have

$$\begin{aligned}
\Psi_1''(\alpha, \beta) &= (\Delta - 1) \sum_{i=1}^q \left((\alpha'_i)^2 / \alpha_i + (1 + \ln \alpha_i) \alpha''_i \right) - \Delta \sum_{i=1}^q \left(\alpha'_i \frac{R'_i}{R_i} + (\ln R_i) \alpha''_i \right) \\
&\quad + (\Delta - 1) \sum_{j=1}^q \left((\beta'_j)^2 / \beta_j + (1 + \ln \beta_j) \beta''_j \right) - \Delta \sum_{j=1}^q \left(\beta'_j \frac{C'_j}{C_j} + (\ln C_j) \beta''_j \right) \\
&= (\Delta - 1) \sum_{i=1}^q (\alpha'_i)^2 / \alpha_i - \Delta \sum_{i=1}^q \alpha'_i \frac{R'_i}{R_i} + \sum_{i=1}^q \alpha''_i \left((\Delta - 1)(1 + \ln \alpha_i) - \Delta \ln R_i \right) \\
&\quad + (\Delta - 1) \sum_{j=1}^q (\beta'_j)^2 / \beta_j - \Delta \sum_{j=1}^q \beta'_j \frac{C'_j}{C_j} + \sum_{j=1}^q \beta''_j \left((\Delta - 1)(1 + \ln \beta_j) - \Delta \ln C_j \right) \\
&= (\Delta - 1) \sum_{i=1}^q (\alpha'_i)^2 / \alpha_i - \Delta \sum_{i=1}^q \alpha'_i \frac{R'_i}{R_i} + (\Delta - 1) \sum_{j=1}^q (\beta'_j)^2 / \beta_j - \Delta \sum_{j=1}^q \beta'_j \frac{C'_j}{C_j},
\end{aligned} \tag{53}$$

where the last equality follows from (42) (replacing α'_i by α''_i and β'_j by β''_j ; note that they are both from the same subspace (35)).

Plugging (37) into (53) we obtain

$$\Psi_1''(\alpha, \beta) = \sum_{i=1}^q \alpha'_i \left((\Delta - 1) \frac{\sum_{j=1}^q B_{ij} C'_j}{\sum_{j=1}^q B_{ij} C_j} - \frac{R'_i}{R_i} \right) + \sum_{j=1}^q \beta'_j \left((\Delta - 1) \frac{\sum_{i=1}^q B_{ij} R'_i}{\sum_{i=1}^q B_{ij} R_i} - \frac{C'_j}{C_j} \right). \tag{54}$$

We are going to use the second partial derivative test (which gives a sufficient condition) to establish maxima of Ψ_1 . We will use the following terminology for local maxima established using this method.

Definition 4.13. A critical point x of a function $f : \mathcal{M} \rightarrow \mathbb{R}$ is called **Hessian local maximum** if the Hessian of f at x is negative definite.

Let \mathbf{L} be the (matrix of) linear map $(r_1, \dots, r_q, c_1, \dots, c_q) \mapsto (\hat{r}_1, \dots, \hat{r}_q, \hat{c}_1, \dots, \hat{c}_q)$ given by

$$\hat{r}_i = \sum_j \frac{B_{ij} R_i C_j}{\sqrt{\alpha_i \beta_j}} c_j \quad \text{and} \quad \hat{c}_j = \sum_i \frac{B_{ij} R_i C_j}{\sqrt{\alpha_i \beta_j}} r_i. \tag{55}$$

In the following, we denote by \mathbf{I} the identity matrix of dimension $2q \times 2q$.

Lemma 4.14. A critical point (α, β) is a Hessian local maximum of $\Psi_1(\alpha, \beta)$ in the subspace defined by (35) if and only if $\mathbf{w}^\top (\mathbf{I} + \mathbf{L}) ((\Delta - 1)\mathbf{L} - \mathbf{I}) \mathbf{w} < 0$ for all $\mathbf{w} = (r_1, \dots, r_q, c_1, \dots, c_q)^\top$ such that

$$\sum_{i=1}^q \sqrt{\alpha_i} r_i = 0 \quad \text{and} \quad \sum_{j=1}^q \sqrt{\beta_j} c_j = 0. \tag{56}$$

Proof. To check whether we are at a Hessian local maximum of $\Psi(\alpha, \beta)$ we have to have (54) negative for non-zero α'_i 's and β'_j 's from the subspace defined by (35) and (38).

Let $r_i = \sqrt{\alpha_i} R'_i / R_i$ and $c_j = \sqrt{\beta_j} C'_j / C_j$. Using (37) we have

$$\begin{aligned} \Psi'' &= \sum_i \alpha_i \left(\frac{R'_i}{R_i} + \frac{\sum_j B_{ij} C'_j}{\sum_j B_{ij} C_j} \right) \left((\Delta - 1) \frac{\sum_j B_{ij} C'_j}{\sum_j B_{ij} C_j} - \frac{R'_i}{R_i} \right) \\ &\quad + \sum_j \beta_j \left(\frac{C'_j}{C_j} + \frac{\sum_i B_{ij} R'_i}{\sum_i B_{ij} R_i} \right) \left((\Delta - 1) \frac{\sum_i B_{ij} R'_i}{\sum_i B_{ij} R_i} - \frac{C'_j}{C_j} \right) \\ &= \sum_i \left(r_i + \sum_j \frac{B_{ij} R_i C_j}{\sqrt{\alpha_i \beta_j}} c_j \right) \left(\sum_j (\Delta - 1) \frac{B_{ij} R_i C_j}{\sqrt{\alpha_i \beta_j}} c_j - r_i \right) \\ &\quad + \sum_j \left(c_j + \sum_i \frac{B_{ij} R_i C_j}{\sqrt{\alpha_i \beta_j}} r_i \right) \left(\sum_i (\Delta - 1) \frac{B_{ij} R_i C_j}{\sqrt{\alpha_i \beta_j}} r_i - c_j \right). \end{aligned}$$

Let $\mathbf{w} = (r_1, \dots, r_q, c_1, \dots, c_q)^\top$. In terms of \mathbf{L} and \mathbf{w} we have

$$\Phi'' = \mathbf{w}^\top (\mathbf{I} + \mathbf{L}) ((\Delta - 1) \mathbf{L} - \mathbf{I}) \mathbf{w}. \quad (57)$$

We have to examine when (57) is in the subspace defined by (35) and (38), which in terms of r_i 's and c_j 's become

$$\sum_i \alpha'_i = \sum_j \beta'_j = \sum_i \sqrt{\alpha_i} r_i + \sum_j \sqrt{\beta_j} c_j = 0, \quad (58)$$

$$\sum_i \alpha_i \frac{R'_i}{R_i} - \sum_j \beta_j \frac{C'_j}{C_j} = \sum_i \sqrt{\alpha_i} r_i - \sum_j \sqrt{\beta_j} c_j = 0. \quad (59)$$

We give more detail on the derivation of (58) below. We have

$$\begin{aligned} \sum_i \alpha'_i &= \sum_i \alpha_i \frac{\alpha'_i}{\alpha_i} = \sum_i \alpha_i \left(\frac{R'_i}{R_i} + \frac{\sum_j B_{ij} C'_j}{\sum_j B_{ij} C_j} \right) = \sum_i r_i \sqrt{\alpha_i} + \sum_i \sum_j B_{ij} R_i C'_j \\ &= \sum_i r_i \sqrt{\alpha_i} + \sum_j \frac{c_j}{\sqrt{\beta_j}} \sum_i B_{ij} R_i C_j = \sum_i r_i \sqrt{\alpha_i} + \sum_j c_j \sqrt{\beta_j}, \end{aligned}$$

the derivation for $\sum_j \beta'_j$ is analogous. □

4.2.2 Attractive fixpoints of tree recursions

The variables R_i , C_j , α_i , β_j in this section refer to a priori different quantities as the variables in Section 4.2.1. We feel that this conflict is justified since we will establish that they coincide.

For convenience we repeat the tree recursions as stated in the introduction:

$$\hat{R}_i \propto \left(\sum_{j=1}^q B_{ij} C_j \right)^{\Delta-1} \quad \text{and} \quad \hat{C}_j \propto \left(\sum_{i=1}^q B_{ij} R_i \right)^{\Delta-1}. \quad (22)$$

We are interested in the **fixpoints** of the tree recursions, that is, R_i 's and C_j 's such that

$$\hat{R}_i \propto R_i \quad \text{and} \quad \hat{C}_j \propto C_j$$

for all $i, j \in [q]$. Note that the fixpoints correspond to the critical points of Ψ_1 (using Theorem 4.1)).

Next we examine the stability of fixpoints. For a continuously differentiable map a sufficient condition for a fixpoint to be attractive is if the spectral radius of the derivative is less than one at the fixpoint. We will use the following terminology for fixpoints whose attractiveness is established using this method.

Definition 4.15. A fixpoint x of a function $f : \mathcal{M} \rightarrow \mathcal{M}$ is called **Jacobian attractive fixpoint** if the Jacobian of f at x has spectral radius less than 1.

Lemma 4.16. Let (\mathbf{r}, \mathbf{c}) be a fixpoint of the tree recursions. Let $\alpha_i = \sum_{j=1}^q B_{ij} R_i C_j$ and $\beta_j = \sum_{i=1}^q B_{ij} R_i C_j$ and let \mathbf{L} be the (matrix of the) map defined by (55). We have that (\mathbf{r}, \mathbf{c}) is Jacobian attractive if and only if $(\Delta - 1)\mathbf{L}$ has spectral radius less than 1 in the subspace of $\mathbf{w} = (r_1, \dots, r_q, c_1, \dots, c_q)$ that satisfy

$$\sum_{i=1}^q \sqrt{\alpha_i} r_i = 0 \quad \text{and} \quad \sum_{j=1}^q \sqrt{\beta_j} c_j = 0. \quad (56)$$

Proof. W.l.o.g. we can assume that (\mathbf{r}, \mathbf{c}) is scaled so that

$$\sum_{i=1}^q \sum_{j=1}^q B_{ij} R_i C_j = 1. \quad (60)$$

Note that the scaling does not affect the value of \mathbf{L} nor does it affect the constraint (56).

When we perturb the R_i 's and C_j 's and apply one step of the tree recursion we obtain

$$\frac{\hat{R}'_i}{\hat{R}_i} = (\Delta - 1) \frac{\sum_{j=1}^q B_{ij} C_j \frac{C'_j}{C_j}}{\sum_{j=1}^q B_{ij} C_j} \quad \text{and} \quad \frac{\hat{C}'_j}{\hat{C}_j} = (\Delta - 1) \frac{\sum_{i=1}^q B_{ij} R_i \frac{R'_i}{R_i}}{\sum_{i=1}^q B_{ij} R_i}. \quad (61)$$

We can rewrite (61) as follows

$$\frac{\hat{R}'_i}{\hat{R}_i} = (\Delta - 1) \frac{\sum_{j=1}^q B_{ij} R_i C_j \frac{C'_j}{C_j}}{\alpha_i} \quad \text{and} \quad \frac{\hat{C}'_j}{\hat{C}_j} = (\Delta - 1) \frac{\sum_{i=1}^q B_{ij} R_i C_j \frac{R'_i}{R_i}}{\beta_j}. \quad (62)$$

The perturbation that scales all R_i 's by the same factor does not change the messages (since they are in the projective space) and hence we need to exclude it when studying local stability of (61). Similarly scaling all C_j 's by the same factor does not change the messages. We need to locate an invariant subspace of (62) whose complement corresponds to the scaling. We obtain the following subspace (it corresponds to preserving (60)):

$$\sum_{i=1}^q \alpha_i \frac{R'_i}{R_i} = 0 \quad \text{and} \quad \sum_{j=1}^q \beta_j \frac{C'_j}{C_j} = 0. \quad (63)$$

Now we check that (63) is invariant under the map (62), indeed,

$$\sum_{i=1}^q \alpha_i \frac{\hat{R}'_i}{\hat{R}_i} = (\Delta - 1) \sum_{i=1}^q \sum_{j=1}^q B_{ij} R_i C_j \frac{C'_j}{C_j} = (\Delta - 1) \sum_{j=1}^q \beta_j \frac{C'_j}{C_j} = 0; \quad (64)$$

the argument for $\sum_{j=1}^q \beta_j \frac{\hat{C}'_j}{\hat{C}_j} = 0$ is analogous.

A fixpoint $(R_1, \dots, R_q, C_1, \dots, C_q)$ is Jacobian attractive if the linear transformation

$$\left(\frac{R'_1}{R_1}, \dots, \frac{R'_q}{R_q}, \frac{C'_1}{C_1}, \dots, \frac{C'_q}{C_q} \right) \mapsto \left(\frac{\hat{R}'_1}{\hat{R}_1}, \dots, \frac{\hat{R}'_q}{\hat{R}_q}, \frac{\hat{C}'_1}{\hat{C}_1}, \dots, \frac{\hat{C}'_q}{\hat{C}_q} \right)$$

given by (61) has spectral radius less than 1 in the subspace defined by (63).

Let $r_i = \sqrt{\alpha_i} R'_i / R_i$, $c_j = \sqrt{\beta_j} C'_j / C_j$, $\hat{r}_i = \sqrt{\alpha_i} \hat{R}'_i / \hat{R}_i$, and $\hat{c}_j = \sqrt{\beta_j} \hat{C}'_j / \hat{C}_j$. This linear transformation of variables turns (62) into

$$\hat{r}_i = (\Delta - 1) \sum_{j=1}^q \frac{B_{ij} R_i C_j}{\sqrt{\alpha_i \beta_j}} c_j \quad \text{and} \quad \hat{c}_j = (\Delta - 1) \sum_{i=1}^q \frac{B_{ij} R_i C_j}{\sqrt{\alpha_i \beta_j}} r_i. \quad (65)$$

Note that (65) is $(\Delta - 1)L$ where L is the map defined by (55). The constraint (63) becomes (56). \square

4.2.3 Connecting attractive fixpoints to maximum entropy configurations

Now we are ready to prove Theorem 4.2.

Proof of Theorem 4.2. Let S be the linear subspace defined by (56) (note that (58) together with (59) define the same subspace). The constraint for the fixpoint to be Jacobian attractive is that $(\Delta - 1)\mathbf{L}$ on S has spectral radius less than 1. The constraint for the critical point to be Hessian maximum is that the eigenvalues of $(\mathbf{I} + \mathbf{L})((\Delta - 1)\mathbf{L} - \mathbf{I})$ on S are negative (see equation (57)).

Note that \mathbf{L} is symmetric and it is a result of tensor product with the matrix $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$. Hence \mathbf{L} has symmetric real spectrum (symmetry means that if a is an eigenvalue then so is $-a$). Note that S is invariant under \mathbf{L} and hence the spectrum of \mathbf{L} on S is a subset of the spectrum of \mathbf{L} (it is still symmetric real; the restriction wiped out a pair of eigenvalues -1 and 1).

The constraint for the fixpoint to be Jacobian attractive, in terms of eigenvalues, is: for each eigenvalue x of \mathbf{L} on S

$$-1 < (\Delta - 1)x < 1. \quad (66)$$

The constraint for the critical point to be Hessian maximum, in terms of eigenvalues, is: for each eigenvalue x of \mathbf{L} on S

$$(1 + x)((\Delta - 1)x - 1) < 0 \quad \text{and} \quad (1 - x)(-(\Delta - 1)x - 1) < 0, \quad (67)$$

where the second constraint comes from the symmetry of the spectrum (thus $-x$ is an eigenvalue). Note that conditions (66) and (67) are equivalent (since $(1 + x)((\Delta - 1)x - 1)$ is negative for $-1 < x < 1/(\Delta - 1)$). \square

5 Reduction for Colorings

In this section we outline our proof of Theorem 1.1. We start by reviewing the main components of the reduction for 2-spin systems (as carried out in [Sly10, SS12]) and in particular the hardcore model. This will allow us to isolate the parts of the argument which do not extend to the multi-spin case and motivate our reduction scheme. The first step is a reduction from max-cut to a so-called phase labeling problem that we introduce. To present the main ideas of this particular key reduction we first present it in this section in the simplified setting of the colorings problem (see Lemma 5.1).

The basic gadget in the reduction is a bipartite random graph, which we denote by G . The sides of the bipartition have an equal number of vertices, and the sides are labelled with $+$ and $-$. Most vertices in G have degree Δ but there is also a small number of degree $\Delta - 1$ vertices (to allow to make connections between gadgets without creating degree $\Delta + 1$ vertices). For $s = \{+, -\}$, let

the vertices in the s -side be $U^s \cup W^s$ where the vertices in $U = U^+ \cup U^-$ have degree Δ and the vertices in $W = W^+ \cup W^-$ have degree $\Delta - 1$. The phase of an independent set I is $+$ (resp. $-$) if I has more vertices in U^+ (resp. U^-). Note that the phase depends only on the spins of the “large” portion of the graph, i.e., the spins of vertices in U .

In non-uniqueness regimes, the gadget G has two important properties, both of which can be obtained by building on the second moment analysis of Section 3. First, the phase of a random independent set I is equal to $+$ or $-$ with probability roughly equal to $1/2$. Second, conditioned on the phase of a random independent set I , the spins of the vertices in W are approximately independent, i.e., the marginal distribution on W is close to a product distribution. In this product distribution if the phase is $+$ (resp. $-$), a vertex in W^+ is in I with probability p^+ (resp. p^-), while a vertex in W^- is in I with probability p^- (resp. p^+). The values p^\pm correspond to maxima of the function Ψ_1 and, crucially (as we shall demonstrate shortly), they satisfy $p^+ \neq p^-$.

Using the second moment analysis of Section 3 and in particular Theorem 1.4, we can prove that an analogous phenomenon takes place for the k -colorings model in the semi-translation non-uniqueness regime (the precise statement of the gadget’s properties are given in Lemma 6.9). The main difference is that, instead of two phases, the number of phases is equal to the number of maximizers of the function Ψ_1 (as described in Theorem 1.6). In particular, for k even, the phase of a coloring is determined by the dominant set of $k/2$ colors on U^+ , i.e., the $k/2$ colors with largest frequencies among vertices of U^+ . Each of the $\binom{k}{k/2}$ phases appears with roughly equal probability and given the phase, the marginal distribution on W is close to a product distribution, which we now describe. We can compute explicit values $a' = a'(k, \Delta), b' = b'(k, \Delta)$ such that for a phase $T \in \binom{[k]}{k/2}$ the probability mass function \mathbf{x} of a vertex in W^+ has its i -th entry equal to a' if $i \in T$ and equal to b' if $i \notin T$. Similarly, the probability mass function \mathbf{y} of a vertex in W^- has its i -th entry equal to b' if $i \in T$ and equal to a' if $i \notin T$. (The values a', b' correspond to the values a, b described in Item 2 of Theorem 1.6, the correspondence is obtained using (23) in Theorem 4.1.³)

Let \mathcal{Q} be the union of the pairs (\mathbf{x}, \mathbf{y}) over all dominant phases. Hereafter, we will identify the phases with elements of \mathcal{Q} . Note that if $(\mathbf{x}, \mathbf{y}) \in \mathcal{Q}$, then $(\mathbf{y}, \mathbf{x}) \in \mathcal{Q}$ as well. We also denote by \mathcal{Q}' the union of unordered elements of \mathcal{Q} . Elements of \mathcal{Q}' are called unordered phases (we use \mathbf{p} to denote unordered phases). Given a phase $\mathbf{p} = \{\mathbf{x}, \mathbf{y}\}$ an ordering of the pair will be called “assigning spin to the phase”. The two ordered phases corresponding to the unordered phase \mathbf{p} will be denoted by \mathbf{p}^+ and \mathbf{p}^- .

The conditional independence property is crucial, it allows us to quantify the effect of using vertices of W as terminals to make connections between copies of the gadget G . For example, consider the following type of connection, which we refer to as parallel. Let $v^+ \in W^+, v^- \in W^-$ and consider two copies of the gadget G , say G_1, G_2 . For $i = 1, 2$ denote by v_i^+, v_i^- the images of v^+, v^- in G_i . Now add the edges (v_1^+, v_2^+) and (v_1^-, v_2^-) and denote the final graph by G_{12} . Thus, a parallel connection corresponds to joining the $+, +$ and $-, -$ sides of two copies of the gadget.

Clearly, random colorings of G_{12} can be generated by first generating random colorings of G_1, G_2 and keeping the resulting coloring if v_1^\pm, v_2^\pm have different colors. We thus have that the partition function of G_{12} is equal to $(Z_G)^2$ times the probability that v_1^\pm, v_2^\pm have different colors in random colorings of G_1, G_2 . The latter quantity can easily be computed if we condition on the phases $(\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2)$ of the colorings in G_1, G_2 , and this is equal to $(1 - \mathbf{x}_1^\top \mathbf{x}_2)(1 - \mathbf{y}_1^\top \mathbf{y}_2)$.

By taking logarithms, we can assume a parallel connection between gadgets with phases $(\mathbf{x}_1, \mathbf{y}_1)$ and $(\mathbf{x}_2, \mathbf{y}_2)$ incurs an (additive) weight

$$w_p((\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2)) = \ln(1 - \mathbf{x}_1^\top \mathbf{x}_2) + \ln(1 - \mathbf{y}_1^\top \mathbf{y}_2).$$

³In particular, a', b' can be readily obtained from a, b using the relations $a = a'^{\Delta/\Delta-1}/S$, $b = a'^{\Delta/\Delta-1}/S$, $\frac{a}{2}(a' + b') = 1$, where $S := \frac{a}{2}(a'^{\Delta/\Delta-1} + b'^{\Delta/\Delta-1})$.

In the hard-core model, parallel connections are sufficient to give hardness. In this case, we have that $\mathcal{Q}' = \{\mathbf{p}\}$ and $\mathcal{Q} = \{\mathbf{p}^+, \mathbf{p}^-\}$ and the respective function $w_p(\cdot, \cdot)$ satisfies

$$w_p(\mathbf{p}^+, \mathbf{p}^+) = w_p(\mathbf{p}^-, \mathbf{p}^-) < w_p(\mathbf{p}^+, \mathbf{p}^-). \quad (68)$$

Thus, in this case, $w_p(\cdot, \cdot)$ takes only two values and neighboring gadgets prefer to have different phases. Now assume that H is an instance of MAX-CUT and replace each vertex in H by a copy of the gadget G , while for each edge of H , connect the respective gadgets in parallel. The partition function of the final graph is dominated from phase assignments which correspond to large cuts in H . This intuition is the basis of the reduction in [Sly10, SS12].

For the colorings model, reducing from MAX-CUT poses an extra challenge. While for every unordered phase \mathbf{p} equation (68) continues to hold, a short calculation shows that the optimal configuration for a triangle of gadgets connected in parallel is to give all three gadgets different phases. To bypass this entanglement, we need to introduce some sort of ferromagnetism in the reduction to enforce gadgets corresponding to vertices of H to use a single (unordered) phase. To achieve this, we use *symmetric* connections, which correspond to having not only $(+, +), (-, -)$ connections of the gadgets, but also $(+, -)$ and $(-, +)$. Thus, a symmetric connection whose endpoints have phases $(\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2)$ incurs (additive) weight

$$w_s((\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2)) = w_p((\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2)) + w_p((\mathbf{x}_1, \mathbf{y}_1), (\mathbf{y}_2, \mathbf{x}_2)).$$

Symmetric connections will allow us to enforce a single unordered phase to all gadgets, while parallel connections will allow us to recover a maximum-cut partition. To have some modularity in our construction, rather than reducing from MAX-CUT directly, we use the following “phase labeling problem”.

COLORINGS PHASE LABELING PROBLEM(\mathbf{B}, \mathcal{Q}):

INPUT: undirected edge-weighted multigraph $H = (V, E)$ and a partition of the edges $\{E_p, E_s\}$.

OUTPUT: $\text{MAXLWT}(H) := \max_{\mathcal{Y}} \text{LWT}_H(\mathcal{Y})$, where the maximization is over all possible phase labelings $\mathcal{Y} : V \rightarrow \mathcal{Q}$ and

$$\text{LWT}_H(\mathcal{Y}) := \sum_{\{u,v\} \in E_s} w_s(\mathcal{Y}(u), \mathcal{Y}(v)) + \sum_{\{u,v\} \in E_p} w_p(\mathcal{Y}(u), \mathcal{Y}(v)).$$

Edges in E_p (resp. E_s) correspond to parallel (resp. symmetric) connections and we shall refer to them as parallel (resp. symmetric) edges. The arguments in [SS12], which we sketched earlier, can easily be adapted to show that an algorithm for approximating the partition function to an arbitrarily small exponential factor yields a PTAS for the phase labeling problem, see Lemma 6.1 and its proof in Section 6.4. It then remains to prove that a PTAS for the phase labeling problem yields a PTAS for MAX-CUT on 3-regular graphs. This is the scope of the next lemma, which we focus on proving in the remainder of this section.

Lemma 5.1. *A (randomized) algorithm that approximates the solution to the COLORINGS PHASE LABELING PROBLEM(\mathbf{B}, \mathcal{Q}) on bounded degree graphs within a factor of $1 - o(1)$ yields a (randomized) algorithm that approximates MAXCUT on 3-regular graphs within a factor of $1 - o(1)$.*

Our reduction relies on the following gadget which “prefers” the unordered phase of two distinguished vertices u and v to agree. For a phase assignment \mathcal{Y} with ordered phases, we denote by \mathcal{Y}' the respective phase assignment with unordered phases.

Lemma 5.2. *A constant sized gadget J_1 with two distinguished vertices u, v can be constructed with the following property: all edges of J_1 are symmetric and the following is true,*

$$\max_{\mathcal{Y}; \mathcal{Y}'(u)=\mathcal{Y}'(v)} \text{LWT}_{J_1}(\mathcal{Y}) > \varepsilon_1 + \max_{\mathcal{Y}; \mathcal{Y}'(u) \neq \mathcal{Y}'(v)} \text{LWT}_{J_1}(\mathcal{Y}), \quad (69)$$

where $\varepsilon_1 > 0$ is a constant depending only on k and Δ .

We give the proof of the critical Lemma 5.2 after the (simpler) proof of Lemma 5.1.

Proof of Lemma 5.1. Let ε_1 be as in Lemma 5.2 and

$$t := 2 \lceil (\max_{\mathbf{p}_1, \mathbf{p}_2} w_p(\mathbf{p}_1, \mathbf{p}_2) - \min_{\mathbf{p}_1, \mathbf{p}_2} w_p(\mathbf{p}_1, \mathbf{p}_2)) / \varepsilon_1 \rceil.$$

Given a 3-regular instance $H = (V, E)$ of MAX-CUT, we first declare all edges of H to be parallel. Moreover, for every edge (u', v') of H , take t copies of gadget J_1 from Lemma 5.2, identify (merge) their u vertices with u' , and identify (merge) their v vertices with v' . Let H' be the final graph.

To find the optimal phase labeling of H' , we may focus on the phase assignment restricted to vertices in H , since each gadget J_1 can be independently set to its optimal value conditioned on the phases for its distinguished vertices u and v . We claim that

$$\text{MAXLWT}(H') = C_1 \text{MAXCUT}(H) + (C_2 + C_3 t) |E|, \quad (70)$$

for constants C_1, C_2, C_3 to be specified later (depending only on k, Δ). Using the trivial bound $\text{MAXCUT}(H) \geq |E|/2 = 3|V|/4$, the lemma follows easily from (70). We thus focus on proving (70).

The key idea is that for any phase labeling $\mathcal{Y} : V \rightarrow \mathcal{Q}$, changing the unordered phases of vertices in H to the same unordered phase $\mathbf{p} \in \mathcal{Q}'$, while keeping the spins, can only increase the weight of the labeling. Indeed, for $(u, v) \in E$ such that $\mathcal{Y}'(u) = \mathcal{Y}'(v)$, no change in the weight of the labeling occurs, using (69). For $(u, v) \in E$ such that $\mathcal{Y}'(u) \neq \mathcal{Y}'(v)$, the potential (weight) loss from the parallel edge (u, v) is compensated by the gain on the t copies of J_1 by (69) and the choice of t .

For phase labelings which assign vertices of H the same unordered phase \mathbf{p} , to attain the maximum weight for a phase labeling, we only need to choose the spins, in order to maximize the contribution from parallel edges (the edges of H). The same argument we discussed for the hard-core model, (68) yields that the optimal choice of spins to the phases induces a maximum-cut partition of H . For such a spin assignment, the contribution from parallel edges is $C_1 \text{MAXCUT}(H) + C_2 |E|$, where

$$C_1 := w_p(\mathbf{p}^+, \mathbf{p}^-) - w_p(\mathbf{p}^-, \mathbf{p}^-) \text{ and } C_2 := w_p(\mathbf{p}^-, \mathbf{p}^-).$$

The contribution from symmetric edges is $C_3 t |E|$, where

$$C_3 := \max_{\mathcal{Y}; \mathcal{Y}'(u)=\mathcal{Y}'(v)=\mathbf{p}} \text{LWT}_{J_1}(\mathcal{Y}).$$

This proves (70). □

We conclude this section by giving the proof of Lemma 5.2.

Proof of Lemma 5.2. Let $\mathcal{Q}' := \{\mathbf{p}_1, \dots, \mathbf{p}_{Q'}\}$ and $\mathbf{p}_i := \{\mathbf{x}_i, \mathbf{y}_i\}$ for $i \in [Q']$. Denote by K the multigraph on Q' vertices $b_1, b_2, \dots, b_{Q'}$ with the following symmetric edges: self-loop on b_i for $i \in [Q']$ and two edges between b_i and b_j for every $i, j \in [Q']$ with $i \neq j$. We first prove that the optimal phase assignments \mathcal{Y} of K are those which assign each vertex b_i a distinct phase from \mathcal{Q}' (note that the spin of the phase does not matter since all edges of K are symmetric). The desired gadget J_1 will be constructed afterwards.

Let \mathcal{Y} be a phase labeling of K and s_i be the number of vertices assigned phase \mathbf{p}_i . Denote by \mathbf{s} the vector $(s_1, \dots, s_{Q'})^\top$. Note that $\mathbf{1}^\top \mathbf{s} = Q'$, where $\mathbf{1}$ is the all one vector with dimension Q' . Then

$$\text{LWT}_K(\mathcal{Y}) = \sum_{i,j \in [Q']} s_i s_j w_s(\mathbf{p}_i, \mathbf{p}_j) = \mathbf{s}^\top \mathbf{A} \mathbf{s},$$

where \mathbf{A} is the $Q' \times Q'$ matrix whose (i, j) entry equals $w_s(\mathbf{p}_i, \mathbf{p}_j)$. Note that \mathbf{A} is symmetric and $\mathbf{1}$ is an eigenvector of \mathbf{A} (because of the transitive symmetry of phases). Moreover, if we let $\mathbf{s}' = \mathbf{s} - \mathbf{1}$, then $\mathbf{1}^\top \mathbf{s}' = 0$. It follows that

$$\mathbf{s}^\top \mathbf{A} \mathbf{s} = \mathbf{1}^\top \mathbf{A} \mathbf{1} + (\mathbf{s}')^\top \mathbf{A} \mathbf{s}'. \quad (71)$$

If \mathbf{A} is negative definite, equation (71) shows that the all ones labeling is better than any other labeling. Hence the result will follow if we prove that \mathbf{A} is negative definite.

Let $\mathbf{z}_1, \dots, \mathbf{z}_{Q'} := \mathbf{x}_1, \dots, \mathbf{x}_{Q'}, \mathbf{y}_1, \dots, \mathbf{y}_{Q'}$ and let $\hat{\mathbf{A}}$ be the $Q \times Q$ matrix whose ij -entry is $\ln(1 - \mathbf{z}_i^\top \mathbf{z}_j)$. Using the definition of the weights $w_s(\cdot, \cdot)$, it is easy to check that for any vector \mathbf{s} it holds that

$$\mathbf{s}^\top \mathbf{A} \mathbf{s} = (\mathbf{s}, \mathbf{s})^\top \hat{\mathbf{A}} (\mathbf{s}, \mathbf{s}),$$

so it suffices to prove that $\hat{\mathbf{A}}$ is negative definite. We will show here that $\hat{\mathbf{A}}$ is negative semi-definite; the proof that $\hat{\mathbf{A}}$ is regular (and hence negative definite) is trickier and is given in the proof of the more general Lemma 6.5. Note that the entries of $\hat{\mathbf{A}}$ are obtained by applying $z \mapsto \ln(1 - z)$ to each entry of the Gram matrix of the vectors $\mathbf{z}_1, \dots, \mathbf{z}_Q$. Since for $|z| < 1$ we have $\ln(1 - z) = -z - z^2/2 - z^3/3 - \dots$, by Schur's product theorem (see Corollary 7.5.9 in [HJ13]) we obtain that $\hat{\mathbf{A}}$ is negative semi-definite, as desired.

To construct the gadget J_1 , we overlay two copies of K as follows. Let K_u (resp. K_v) be a copy of K , where the image of $b_{Q'}$ is renamed to u (resp. v). Overlay K_u, K_v by identifying the images of $b_1, \dots, b_{Q'-1}$ in the two copies. Thus, the resulting graph J_1 has two self loops on b_i for $i \in [Q' - 1]$, four edges between b_i and b_j for every $i, j \in [Q' - 1]$ with $i \neq j$, two edges between u and b_i for $i \in [Q' - 1]$, two edges between v and b_i for $i \in [Q' - 1]$ and a self loop on u, v .

Note that for every phase labeling \mathcal{Y} of J_1 , we have $\text{LWT}_{J_1}(\mathcal{Y}) = \text{LWT}_{K_u}(\mathcal{Y}) + \text{LWT}_{K_v}(\mathcal{Y})$ and hence $\text{MAXLWT}(J_1) \leq 2\text{MAXLWT}(K)$. Using that the optimal phase labelings for K are those which assign each vertex a distinct phase from \mathcal{Q}' , we obtain that the inequality holds at equality for those (and only those) phase labelings which assign u, v a common phase $\mathbf{p} \in \mathcal{Q}'$ and vertices $b_1, \dots, b_{Q'-1}$ a distinct phase from $\mathcal{Q}' - \{\mathbf{p}\}$. This yields the ε_1 in the statement of the lemma. Note that ε_1 depends only on Q' , which in turn is completely determined by k, Δ . \square

6 General Reduction

6.1 Phase labeling Problem

We first introduce the phase labeling problem for a general antiferromagnetic spin system (which satisfies the hypotheses of Theorem 1.5). As in the case for the colorings model (see Section 5), we

let \mathcal{Q}' be the union of $\{\mathbf{x}, \mathbf{y}\}$ over all phases, i.e.,

$$\mathcal{Q}' = \{\{\mathbf{x}_1, \mathbf{y}_1\}, \dots, \{\mathbf{x}_{Q'}, \mathbf{y}_{Q'}\}\}.$$

Henceforth, we will refer to elements of \mathcal{Q}' as phases. Note that for fixed q, Δ, \mathbf{B} the global maxima of Ψ_1 correspond to fixpoints of (22) and hence can be approximated to any desired polynomial accuracy of their values. The values of \mathbf{x}, \mathbf{y} may then be recovered using (23) (see Footnote 3 for an explicit description of the correspondence in the case of colorings). The assumption of Theorem 1.5 translates into $\mathbf{x}_i \neq \mathbf{y}_i$ for all $i \in [Q']$.

Given an unordered phase $\{\mathbf{x}, \mathbf{y}\}$ an ordering of the pair will be called “assigning spin to the phase”. Let

$$\mathcal{Q} = \{(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_Q, \mathbf{y}_Q)\}$$

be the collection of ordered phases. Note that $Q = 2Q'$. We will denote unordered phases using \mathbf{p} ; the two ordered phases corresponding to the unordered phase \mathbf{p} will be denoted by \mathbf{p}^+ and \mathbf{p}^- . Given a graph H with vertex set V we will assign ordered phases to its vertices—the labeling (called phase assignment) will be denoted by $\mathcal{Y} : V \rightarrow \mathcal{Q}$. The corresponding labeling by unordered phases (where the ordering is removed) will be denoted by \mathcal{Y}' .

Now we define the weight of a phase assignment. We will have two types of edges in H : parallel or symmetric; the type of an edge will only impact the weight of a phase assignment. In particular, a parallel edge whose endpoints have labels $(\mathbf{x}_1, \mathbf{y}_1)$ and $(\mathbf{x}_2, \mathbf{y}_2)$ incurs weight

$$w_p((\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2)) = \ln(\mathbf{x}_1^\top \mathbf{B} \mathbf{x}_2) + \ln(\mathbf{y}_1^\top \mathbf{B} \mathbf{y}_2),$$

while a symmetric edge incurs weight

$$w_s((\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2)) = w_p((\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2)) + w_p((\mathbf{x}_1, \mathbf{y}_1), (\mathbf{y}_2, \mathbf{x}_2)).$$

Note that if we flip $(\mathbf{x}_1, \mathbf{y}_1)$, that is, replace it by $(\mathbf{y}_1, \mathbf{x}_1)$, the weight of the symmetric edge does not change.

We will use the following problem in our reduction.

PHASE LABELING PROBLEM(\mathbf{B}, \mathcal{Q}):

INPUT: undirected edge-weighted multigraph $H = (V, E)$ and a partition of the edges $\{E_p, E_s\}$.

OUTPUT: $\text{MAXLWT}(H) := \max_{\mathcal{Y}} \text{LWT}_H(\mathcal{Y})$, where the maximization is over all possible phase labelings $\mathcal{Y} : V \rightarrow \mathcal{Q}$ and

$$\text{LWT}_H(\mathcal{Y}) = \sum_{\{u,v\} \in E_s} w_s(\mathcal{Y}(u), \mathcal{Y}(v)) + \sum_{\{u,v\} \in E_p} w_p(\mathcal{Y}(u), \mathcal{Y}(v)).$$

The motivation for the PHASE LABELING PROBLEM is the following lemma. The proof roughly follows the lines of [SS12] and is given in Section 6.4.

Lemma 6.1. *In the setting of Theorem 1.5, the following holds. A (randomized) algorithm that approximates the partition function on triangle free Δ -regular graphs within an arbitrarily small exponential factor yields a (randomized) algorithm that approximates the solution to the phase labeling problem with parameters \mathbf{B}, \mathcal{Q} on bounded degree graphs within a factor of $1 - o(1)$.*

The following lemma requires more work in our setting and is proved in Section 6.3.

Lemma 6.2. *A (randomized) algorithm that approximates the solution to the phase labeling problem with parameters \mathbf{B}, \mathcal{Q} on bounded degree graphs within a factor of $1 - o(1)$ yields a (randomized) algorithm that approximates MAXCUT on 3-regular graphs within a factor of $1 - o(1)$.*

Using Lemmas 6.1 and 6.2, we obtain Theorem 1.5.

Proof of Theorem 1.5. Suppose that there exists a (randomized) algorithm to approximate the partition function on Δ -regular graphs with interaction matrix \mathbf{B} up to an arbitrarily small exponential factor. Then, combining Lemmas 6.1 and 6.2, we obtain a (randomized) algorithm to approximate MAXCUT on 3-regular graphs within a factor of $1 - o(1)$. This contradicts the result of [AK97]. \square

6.2 Properties of Antiferromagnetic Spin Systems

In this section we prove two basic properties of antiferromagnetic systems that will be used in our general reductions.

As a consequence of the Perron-Frobenius theorem and the antiferromagnetism definition (cf. Definition 1.3), we may decompose the interaction matrix \mathbf{B} of an antiferromagnetic model as

$$\mathbf{B} = \mathbf{u}\mathbf{u}^\top - \mathbf{P}^\top\mathbf{P}, \quad (72)$$

where the vector \mathbf{u} has positive entries and \mathbf{P} is a square matrix. Using the decomposition (72), we prove the following two lemmas which are used in the reduction.

Lemma 6.3. *For antiferromagnetic \mathbf{B} , and vectors $\mathbf{z}_1, \mathbf{z}_2 \in \mathbb{R}_{\geq 0}^q$ with $\|\mathbf{z}_1\|_1 = \|\mathbf{z}_2\|_1 = 1$, we have*

$$(\mathbf{z}_1^\top \mathbf{B} \mathbf{z}_1)(\mathbf{z}_2^\top \mathbf{B} \mathbf{z}_2) \leq (\mathbf{z}_1^\top \mathbf{B} \mathbf{z}_2)^2.$$

Equality holds iff $\mathbf{z}_1 = \mathbf{z}_2$.

Proof. Set $\mathbf{w}_1 = \mathbf{P}\mathbf{z}_1$, $\mathbf{w}_2 = \mathbf{P}\mathbf{z}_2$, $a_1 = \mathbf{u}^\top \mathbf{z}_1$, $a_2 = \mathbf{u}^\top \mathbf{z}_2$. Then

$$\mathbf{z}_1^\top \mathbf{B} \mathbf{z}_1 = a_1^2 - \mathbf{w}_1^\top \mathbf{w}_1, \quad \mathbf{z}_2^\top \mathbf{B} \mathbf{z}_2 = a_2^2 - \mathbf{w}_2^\top \mathbf{w}_2, \quad \mathbf{z}_1^\top \mathbf{B} \mathbf{z}_2 = a_1 a_2 - \mathbf{w}_1^\top \mathbf{w}_2.$$

Since $\mathbf{B}, \mathbf{z}_1, \mathbf{z}_2$ have nonnegative entries, the above equalities imply $a_1^2 - \mathbf{w}_1^\top \mathbf{w}_1, a_2^2 - \mathbf{w}_2^\top \mathbf{w}_2, a_1 a_2 - \mathbf{w}_1^\top \mathbf{w}_2 \geq 0$. The inequality reduces to

$$(a_1^2 - \mathbf{w}_1^\top \mathbf{w}_1)(a_2^2 - \mathbf{w}_2^\top \mathbf{w}_2) \leq (a_1 a_2 - \mathbf{w}_1^\top \mathbf{w}_2)^2.$$

This is known as Aczél's inequality. The fastest proof goes as follows: set $b_1^2 = a_1^2 - \mathbf{w}_1^\top \mathbf{w}_1$ and $b_2^2 = a_2^2 - \mathbf{w}_2^\top \mathbf{w}_2$, so that by Cauchy-Schwarz $a_1 a_2 \geq b_1 b_2 + \mathbf{w}_1^\top \mathbf{w}_2$, implying the inequality.

Equality can only hold if $a_1 = \lambda a_2$ and $\mathbf{w}_1 = \lambda \mathbf{w}_2$, yielding $\mathbf{u}^\top (\mathbf{z}_1 - \lambda \mathbf{z}_2) = 0$ and $\mathbf{P}(\mathbf{z}_1 - \lambda \mathbf{z}_2) = 0$. We easily obtain $\mathbf{B}(\mathbf{z}_1 - \lambda \mathbf{z}_2) = 0$ and since \mathbf{B} is invertible, $\mathbf{z}_1 = \lambda \mathbf{z}_2$. The assumption $\|\mathbf{z}_1\|_1 = \|\mathbf{z}_2\|_1 = 1$ implies $\lambda = 1$, as wanted. \square

Corollary 6.4. *By plugging in the inequality of Lemma 6.3 the vectors with a single 1 in the positions i and j respectively, we obtain that any two spins i, j induce an antiferromagnetic two-spin system.*

Lemma 6.5. *Let $\mathbf{z}_1, \dots, \mathbf{z}_n \in \mathbb{R}^d$ be a collection of distinct non-negative vectors such that $\|\mathbf{z}_i\|_1 = 1$ for $i \in [n]$. Let $a_i = \mathbf{z}_i^\top \mathbf{u}$, where \mathbf{u} is as in (72). Let \mathbf{A}' be the $n \times n$ matrix whose ij -th entry is $\ln(\mathbf{z}_i^\top \mathbf{B} \mathbf{z}_j) - \ln(a_i) - \ln(a_j)$. Then \mathbf{A}' is negative definite.*

Proof. Let $\mathbf{w}_i = \frac{1}{a_i} \mathbf{P} \mathbf{z}_i$ and let \mathbf{W} be the $q \times n$ matrix whose columns are $\mathbf{w}_1, \dots, \mathbf{w}_n$. We first argue $\mathbf{w}_i \neq \mathbf{w}_j$ for $i \neq j$. Suppose $\mathbf{w}_i = \mathbf{w}_j$. Let $\mathbf{z} = \frac{1}{a_i} \mathbf{z}_i - \frac{1}{a_j} \mathbf{z}_j$. We have $\mathbf{P} \mathbf{z} = \mathbf{w}_i - \mathbf{w}_j = 0$ and $\mathbf{u}^\top \mathbf{z} = 1 - 1 = 0$ and hence $\mathbf{B} \mathbf{z} = 0$. Since \mathbf{B} is regular we have $\mathbf{z} = 0$. Thus $0 = \mathbf{z}^\top \mathbf{1} = \frac{1}{a_i} - \frac{1}{a_j}$ which implies $a_i = a_j$ which in turn implies $\mathbf{z}_i = \mathbf{z}_j$, a contradiction. Thus $\mathbf{w}_i \neq \mathbf{w}_j$ for $i \neq j$.

Note that we have

$$\ln(1 - \mathbf{w}_i^\top \mathbf{w}_j) = \ln(a_i a_j - \mathbf{z}_i^\top \mathbf{P}^\top \mathbf{P} \mathbf{z}_j) - \ln(a_i a_j) = A'_{ij}.$$

Thus the ij -th entry in \mathbf{A}' is obtained by applying $z \mapsto \ln(1 - z)$ to each entry of the Gramm matrix $\mathbf{W}^\top \mathbf{W}$. Note that for $|z| < 1$ we have $\ln(1 - z) = -z - z^2/2 - z^3/3 - \dots$ and hence by Schur product theorem \mathbf{A}' is negative semi-definite (see Corollary 7.5.9 in [HJ13]).

Now we argue that \mathbf{A}' is regular (and hence negative definite). We have

$$-\mathbf{A}' = \sum_{k=1}^{\infty} \frac{1}{k} \mathbf{W}_k^\top \mathbf{W}_k, \quad (73)$$

where \mathbf{W}_k is the $q^k \times n$ matrix whose columns are $w_1^{\otimes k}, \dots, w_n^{\otimes k}$. Note that if \mathbf{A}' is singular then there exists a non-zero vector \mathbf{v} such that $\mathbf{v}^\top \mathbf{A}' \mathbf{v} = 0$ and for this to happen we would have to have

$$\mathbf{W}_k \mathbf{v} = 0 \quad (74)$$

for all $k \geq 1$ (the terms on the right-hand side of (73) are non-negative and if even one of them is positive then $\mathbf{v}^\top \mathbf{A}' \mathbf{v} < 0$).

There exists a vector $\mathbf{r} \in \mathbb{R}^q$ such that $\alpha_i = \mathbf{r}^\top \mathbf{w}_i, i = 1, \dots, n$ are distinct real numbers (the \mathbf{w}_i 's are distinct and hence for any $i \neq j$ the measure of $r \in [0, 1]^q$ such that $\mathbf{r}^\top \mathbf{w}_i = \mathbf{r}^\top \mathbf{w}_j$ is zero). Note that $(\mathbf{r}^{\otimes k})^\top \mathbf{W}_k$ is $(\alpha_1^k, \dots, \alpha_n^k)$. From (74) we obtain that for every integer $k \geq 1$ we have $(\alpha_1^k, \dots, \alpha_n^k) \mathbf{v} = 0$ and hence $\mathbf{v} = 0$ (by considering the Vandermonde matrix $\{\alpha_i^k\}$), a contradiction. Hence \mathbf{A}' is regular and negative definite. \square

6.3 Reducing MaxCut to Phase Labeling

In this section, we prove Lemma 6.2.

6.3.1 An intermediate gadget

We will use the following gadget which “prefers” the unordered phase of two vertices to agree.

Lemma 6.6. *A constant sized gadget J_1 with two distinguished vertices u, v can be constructed with the following property: all edges of J_1 are symmetric and the following is true,*

$$\max_{\mathcal{Y}; \mathcal{Y}'(u)=\mathcal{Y}'(v)} \text{LWT}_{J_1}(\mathcal{Y}) > \varepsilon_1 + \max_{\mathcal{Y}; \mathcal{Y}'(u) \neq \mathcal{Y}'(v)} \text{LWT}_{J_1}(\mathcal{Y}), \quad (75)$$

where $\varepsilon_1 > 0$ is a constant depending only on the spin model and Δ .

Note that Lemma 5.2 which was proved in Section 5 is a special case of Lemma 6.6 in the case of the colorings model. The proof of Lemma 6.6 follows roughly the same lines with slightly more intricate technical details.

Proof of Lemma 6.6. Let $\mathbf{z}_1, \dots, \mathbf{z}_Q := \mathbf{x}_1, \dots, \mathbf{x}_{Q'}, \mathbf{y}_1, \dots, \mathbf{y}_{Q'}$. Let \mathbf{u} be defined as in Equation (72). In Section 6.2, Lemma 6.5 it is proved that the $Q \times Q$ matrix $\hat{\mathbf{A}}$ whose ij -th entry is $\ln(\mathbf{z}_i^\top \mathbf{B} \mathbf{z}_j) - \ln(\mathbf{z}_i^\top \mathbf{u}) - \ln(\mathbf{z}_j^\top \mathbf{u})$ is negative definite. Let \mathbf{A}' be the $Q' \times Q'$ matrix obtained by the following “folding” of $\hat{\mathbf{A}}$:

$$\mathbf{A}'_{ij} = \hat{\mathbf{A}}_{i,j} + \hat{\mathbf{A}}_{i+Q',j} + \hat{\mathbf{A}}_{i,j+Q'} + \hat{\mathbf{A}}_{i+Q',j+Q'}.$$

We have that \mathbf{A}' is also negative definite (since $\mathbf{x}^\top \mathbf{A}' \mathbf{x} = \mathbf{y}^\top \hat{\mathbf{A}} \mathbf{y}$, where $\mathbf{y}^\top = (\mathbf{x}^\top, \mathbf{x}^\top)$). Note that

$$\mathbf{A}'_{ij} = w_s((\mathbf{x}_i, \mathbf{y}_i), (\mathbf{x}_j, \mathbf{y}_j)) - a'_i - a'_j,$$

where $a'_i := 2 \ln(\mathbf{x}_i^\top \mathbf{u}) + 2 \ln(\mathbf{y}_i^\top \mathbf{u})$.

Let λ_1 be largest eigenvalue of $-\mathbf{A}'$ and let λ_2 be the smallest eigenvalue of $-\mathbf{A}'$. Note that $0 < \lambda_2 \leq \lambda_1$. Define \mathbf{A} to be the $Q' \times Q'$ matrix with $A_{ij} = A'_{ij} + a'_i + a'_j$ and consider the following maximization problem

$$\max_{\mathbf{x}; \mathbf{x}^\top \mathbf{1} = 1, \mathbf{x} \geq 0} \mathbf{x}^\top \mathbf{A} \mathbf{x}. \quad (76)$$

Note that for \mathbf{x} with $\mathbf{x}^\top \mathbf{1} = 1$ we have

$$\mathbf{x}^\top \mathbf{A} \mathbf{x} = 2\mathbf{a}'^\top \mathbf{x} + \mathbf{x}^\top \mathbf{A}' \mathbf{x}, \quad (77)$$

where \mathbf{A}' is negative definite. Note that if \mathbf{x} and \mathbf{y} are distinct optimal solutions of (76) then $(\mathbf{x} + \mathbf{y})/2$ satisfies all the constraints, and from (77) and negative definiteness of \mathbf{A}' we have

$$((\mathbf{x} + \mathbf{y})/2)^\top \mathbf{A} ((\mathbf{x} + \mathbf{y})/2) > (\mathbf{x}^\top \mathbf{A} \mathbf{x} + \mathbf{y}^\top \mathbf{A} \mathbf{y})/2,$$

a contradiction (with optimality of both \mathbf{x} and \mathbf{y}). Thus (76) has a unique maximum; let \mathbf{x}^* be the value of \mathbf{x} achieving it. Let $O^* = (\mathbf{x}^*)^\top \mathbf{A} \mathbf{x}^*$. Let S be the set of non-zero coordinates in \mathbf{x}^* .

Let $\mathbf{y} \in \mathbb{R}^{Q'}$ be such that $\mathbf{y}^\top \mathbf{1} = 0$ and \mathbf{y} is zero on coordinates outside S . Then from (local) optimality of \mathbf{x}^* we have

$$(\mathbf{x}^* + \mathbf{y})^\top \mathbf{A} (\mathbf{x}^* + \mathbf{y}) = O^* + 2(\mathbf{a}'^\top + (\mathbf{x}^*)^\top \mathbf{A}') \mathbf{y} + \mathbf{y}^\top \mathbf{A}' \mathbf{y} = O^* + \mathbf{y}^\top \mathbf{A}' \mathbf{y} \geq O^* - \lambda_1 \|\mathbf{y}\|_2^2. \quad (78)$$

Equation (78) tells us that moving slightly from the optimum the objective decreases at most quadratically in the length of \mathbf{y} .

Let $\mathbf{y} \in \mathbb{R}^{Q'}$ be such that $\mathbf{y}^\top \mathbf{1} = 0$ and \mathbf{y} is non-negative on coordinates outside S . Then from (local) optimality of \mathbf{x}^* we have

$$(\mathbf{x}^* + \mathbf{y})^\top \mathbf{A} (\mathbf{x}^* + \mathbf{y}) = O^* + 2(\mathbf{a}'^\top + (\mathbf{x}^*)^\top \mathbf{A}') \mathbf{y} + \mathbf{y}^\top \mathbf{A}' \mathbf{y} = O^* + \mathbf{y}^\top \mathbf{A}' \mathbf{y} \geq O^* - \lambda_2 \|\mathbf{y}\|_2^2. \quad (79)$$

Equation (79) tells us that moving slightly from the optimum the objective decreases at least quadratically in the length of \mathbf{y} .

Let $Z \geq (4Q'\lambda_1/\lambda_2)^{Q'}$. Note that Z is a constant depending only on the model and Δ . Let $z_1/z, \dots, z_{Q'}/z$ be the optimal simultaneous Diophantine approximation of $x_1^*, \dots, x_{Q'}^*$ with $z_1, \dots, z_{Q'}, z \in \mathbb{Z}$ and $1 \leq z \leq Z$. By Dirichlet's theorem we have

$$|zx_i^* - z_i| \leq Z^{-1/Q'} < 1. \quad (80)$$

Note that (80) implies

$$\text{if } x_i^* = 0 \text{ then } z_i = 0. \quad (81)$$

Also note that

$$\left| \sum_{i=1}^{Q'} zx_i^* - \sum_{i=1}^{Q'} z_i \right| \leq \sum_{i=1}^{Q'} |zx_i^* - z_i| \leq Q' Z^{-1/Q'} < 1,$$

and since z and z_i 's are integers and $(\mathbf{x}^*)^\top \mathbf{1} = 1$ we have

$$\sum_{i=1}^{Q'} \frac{z_i}{z} = 1. \quad (82)$$

From (81) and (82) we have that for $\mathbf{y} := (z_1/z, \dots, z_{Q'}/z) - \mathbf{x}^*$ we can apply (78) and hence

$$(z_1/z, \dots, z_{Q'}/z) \mathbf{A} (z_1/z, \dots, z_{Q'}/z)^\top \geq O^* - \lambda_1 Q' Z^{-2/Q'} z^{-2}. \quad (83)$$

Now we are ready to construct the gadget J_1 . First, let K be the multigraph on z vertices b_1, b_2, \dots, b_z with the following symmetric edges: self-loop on b_i for $i \in [z]$ and two edges between b_i and b_j for every $i, j \in [z]$ with $i \neq j$. To obtain J_1 , we overlay two copies of K as follows. Let K_u (resp. K_v) be a copy of K , where the image of b_z is renamed to u (resp. v). Overlay K_u, K_v by identifying the images of b_1, \dots, b_{z-1} in the two copies. Thus, the resulting graph J_1 has $z+1$ vertices and the following edges: two self loops on b_i for $i \in [z-1]$, four edges between b_i and b_j for every $i, j \in [Q'-1]$ with $i \neq j$, two edges between u and b_i for $i \in [z-1]$, two edges between v and b_i for $i \in [z-1]$ and a self loop on u, v .

Note that the weight of a phase assignment on J_1 is the sum of the induced phase assignments on K_u and K_v . Consider an assignment of phases \mathcal{Y}_o such that in each complete graph z_i vertices get phase i (note that this forces the phases of u and v to be the same). The weight of the phase assignment \mathcal{Y}_o is

$$\text{LWT}_{J_1}(\mathcal{Y}_o) = S_1 := 2(z_1, \dots, z_{Q'}) \mathbf{A} (z_1, \dots, z_{Q'})^\top \geq 2z^2 O^* - 2\lambda_1 Q' Z^{-2/Q'}. \quad (84)$$

Now suppose that we have a phase assignment \mathcal{Y} for J_1 where the phases of u and v are different. Let $\hat{\mathbf{u}}$ be the vector with \hat{u}_i counting the number of vertices with phase i in K_u and define similarly $\hat{\mathbf{v}}$.

Note that $\|\hat{\mathbf{u}} - \hat{\mathbf{v}}\|_2^2 = 2$ (since $\hat{\mathbf{u}}$ and $\hat{\mathbf{v}}$ differ in two coordinates—the phases of u and v in the assignment). By triangle inequality we have $\|\hat{\mathbf{u}}/z - \mathbf{x}^*\|_2 \geq 1/(z\sqrt{2})$ or $\|\hat{\mathbf{v}}/z - \mathbf{x}^*\|_2 \geq 1/(z\sqrt{2})$ (otherwise we would have $\|\hat{\mathbf{u}}/z - \hat{\mathbf{v}}/z\|_2 < \sqrt{2}/z$). W.l.o.g. assume that $\hat{\mathbf{u}}/z$ has the greater distance from \mathbf{x}^* . We have

$$\text{LWT}_{J_1}(\mathcal{Y}_o) = S_2 := \hat{\mathbf{u}}^\top \mathbf{A} \hat{\mathbf{u}} + \hat{\mathbf{v}}^\top \mathbf{A} \hat{\mathbf{v}} \leq z^2(2O^* - \lambda_2/(2z^2)) = 2z^2 O^* - \lambda_2/2. \quad (85)$$

By our choice of Z we have $S_1 > S_2$ and hence in an optimal phase assignment for J_1 we have that u and v get the same phase. Note that we did not show which phase assignment is optimal; we only found a phase assignment in which u, v have the same phase that is better than any assignment in which u, v have different phases. \square

6.3.2 The reduction

In Section 6.2, Lemma 6.3 we proved that for a parallel edge and any phase \mathbf{p} we have $w(\mathbf{p}^+, \mathbf{p}^+) = w_p(\mathbf{p}^-, \mathbf{p}^-) < w_p(\mathbf{p}^+, \mathbf{p}^-)$ and hence there exists a constant $\varepsilon_2 > 0$ depending only on the model and Δ such that for every phase $\mathbf{p} \in \mathcal{Q}$ we have

$$w_p(\mathbf{p}^+, \mathbf{p}^+) = w_p(\mathbf{p}^-, \mathbf{p}^-) < w_p(\mathbf{p}^+, \mathbf{p}^-) - \varepsilon_2. \quad (86)$$

Combining Lemma 6.6 with equation (86) we can construct a gadget that “prefers” the unordered phase of two vertices to agree and also “prefers” the spin assignment to disagree.

Lemma 6.7. *A constant sized gadget J_2 can be constructed with two distinguished vertices u, v and the following property: there exists a phase $\mathbf{p} \in \mathcal{Q}'$ satisfying simultaneously all of the following:*

1. $A_1(\mathbf{p}) = \text{MAXLWT}(J_2)$, where

$$A_1(\mathbf{p}) := \max_{\mathcal{Y}; \mathcal{Y}(u)=\mathbf{p}^+, \mathcal{Y}(v)=\mathbf{p}^-} \text{LWT}_{J_2}(\mathcal{Y}) = \max_{\mathcal{Y}; \mathcal{Y}(u)=\mathbf{p}^-, \mathcal{Y}(v)=\mathbf{p}^+} \text{LWT}_{J_2}(\mathcal{Y}). \quad (87)$$

2. Among \mathbf{p} that satisfy Item 1, \mathbf{p} maximizes

$$A_2(\mathbf{p}) := \max_{\mathcal{Y}; \mathcal{Y}(u)=\mathbf{p}^+, \mathcal{Y}(v)=\mathbf{p}^+} \text{LWT}_{J_2}(\mathcal{Y}) = \max_{\mathcal{Y}; \mathcal{Y}(u)=\mathbf{p}^-, \mathcal{Y}(v)=\mathbf{p}^-} \text{LWT}_{J_2}(\mathcal{Y}). \quad (88)$$

3. The following inequalities hold

$$A_1(\mathbf{p}) > A_2(\mathbf{p}) + \varepsilon_3 \quad \text{and} \quad A_2(\mathbf{p}) > \varepsilon_3 + \max_{\mathcal{Y}; \mathcal{Y}'(u) \neq \mathcal{Y}'(v)} \text{LWT}_{J_2}(\mathcal{Y}), \quad (89)$$

where $\varepsilon_3 > 0$ is a constant (depending only on the model and Δ).

Proof. To construct J_2 we take $t := 3\lceil (\max_{\mathbf{p}_1, \mathbf{p}_2} w_p(\mathbf{p}_1, \mathbf{p}_2) - \min_{\mathbf{p}_1, \mathbf{p}_2} w_p(\mathbf{p}_1, \mathbf{p}_2)) / \varepsilon_1 \rceil$ copies of gadget J_1 from Lemma 6.6, identify (merge) their u vertices, and identify (merge) their v vertices. Finally we add a parallel edge between u and v .

Let \mathbf{p} be the unordered phase that is the common value of $\mathcal{Y}'(u)$ and $\mathcal{Y}'(v)$ for which the maximum on the left-hand side of (75) is achieved (note that \mathbf{p} is not unique; we just take one such \mathbf{p}). Let

$$A_4 := \max_{\mathcal{Y}; \mathcal{Y}'(u)=\mathbf{p}, \mathcal{Y}'(v)=\mathbf{p}} \text{LWT}_{J_2}(\mathcal{Y}) \quad \text{and} \quad A_5 := \max_{\mathcal{Y}; \mathcal{Y}'(u) \neq \mathcal{Y}'(v)} \text{LWT}_{J_2}(\mathcal{Y}).$$

Then applying (75) on each copy of J_1 in J_2 we obtain

$$A_4 > A_5 + 2(\max_{\mathbf{p}_1, \mathbf{p}_2} w_p(\mathbf{p}_1, \mathbf{p}_2) - \min_{\mathbf{p}_1, \mathbf{p}_2} w_p(\mathbf{p}_1, \mathbf{p}_2)). \quad (90)$$

Thus the maximizer of $\max_{\mathcal{Y}} \text{LWT}_{J_2}(\mathcal{Y})$ happens for \mathcal{Y} with $\mathcal{Y}'(u) = \mathcal{Y}'(v)$. Only the parallel edge is influenced by the spin and hence, by (86), we have

$$\max_{\mathcal{Y}} \text{LWT}_{J_2}(\mathcal{Y}) = \max_{\mathbf{p}} \max_{\mathcal{Y}; \mathcal{Y}(u)=\mathbf{p}^+, \mathcal{Y}(v)=\mathbf{p}^-} \text{LWT}_{J_2}(\mathcal{Y}). \quad (91)$$

Let \mathbf{p} be the maximizer on the right-hand side of (91) that (secondarily) maximizes the second expression in (88). Note that \mathbf{p} satisfies the first and second condition of the lemma. The first part of the third condition is satisfied for any $\varepsilon_3 \leq \varepsilon_2$ (using (86)). Recall that $\varepsilon_2 > 0$. The second part of the third condition is satisfied for $\varepsilon_3 \leq \max_{\mathbf{p}_1, \mathbf{p}_2} w_p(\mathbf{p}_1, \mathbf{p}_2) - \min_{\mathbf{p}_1, \mathbf{p}_2} w_p(\mathbf{p}_1, \mathbf{p}_2)$. Recall that $\max_{\mathbf{p}_1, \mathbf{p}_2} w_p(\mathbf{p}_1, \mathbf{p}_2) - \min_{\mathbf{p}_1, \mathbf{p}_2} w_p(\mathbf{p}_1, \mathbf{p}_2) > 0$. Thus we can take $\varepsilon_3 > 0$ to be the smaller of the two upper bounds (each of which is a constant depending on the model and Δ only). \square

Lemma 6.8. *Let \mathbf{B} be the interaction matrix of an antiferromagnetic spin model. Let A_1, A_2 be the constants defined in Lemma 6.7. There exists constants D_1, D_2, D_3 depending only on the model and Δ such that the following is true. Given a cubic graph H we can, in polynomial-time, construct a max-degree- D_1 graph G with $|V(G)| \leq D_2|V(H)|$ such that*

$$\text{MAXLWT}(G) = (A_1 - A_2)\text{MAXCUT}(H) + A_2|E(H)| + A_1D_3|V(H)|.$$

We can now go back and prove the inapproximability result for the phase labeling problem.

Proof of Lemma 6.2. Since A_1, A_2, D_3 are constants depending only on the model and Δ , the trivial algorithm gives the bound $\text{MAXCUT}(H) \geq 1/2|E(H)| = 3/4|V(H)|$. Together with Lemma 6.8 we obtain the result. \square

Proof of Lemma 6.8. Replace each edge of H by gadget J_2 and for each vertex $w \in V(H)$ add D_3 new vertices w_1, \dots, w_{D_3} and add a gadgets J_2 between w and w_i (for $i \in [D_3]$), where D_3 will be determined shortly.

The purpose of the D_3 copies of J_2 is to force phase \mathbf{p} (from Lemma 6.7) to be used on the distinguished vertices in a labeling of G with maximum weight. A phase $\mathbf{r} \neq \mathbf{p}$ can have

$$\ell_1(\mathbf{r}) := \max_{\mathcal{Y}; \mathcal{Y}(u)=\mathbf{r}^+, \mathcal{Y}(v)=\mathbf{r}^+} \text{LWT}_{J_2}(\mathcal{Y}) - \max_{\mathcal{Y}; \mathcal{Y}(u)=\mathbf{p}^+, \mathcal{Y}(v)=\mathbf{p}^+} \text{LWT}_{J_2}(\mathcal{Y}) > 0, \quad (92)$$

but then by the choice of \mathbf{p}

$$\ell_2(\mathbf{r}) := \max_{\mathcal{Y}; \mathcal{Y}(u)=\mathbf{p}^+, \mathcal{Y}(v)=\mathbf{p}^-} \text{LWT}_{J_2}(\mathcal{Y}) - \max_{\mathcal{Y}; \mathcal{Y}(u)=\mathbf{r}^+, \mathcal{Y}(v)=\mathbf{r}^-} \text{LWT}_{J_2}(\mathcal{Y}) > 0. \quad (93)$$

Let

$$D_3 = 4 + 3 \left\lceil \max_{\mathbf{r}} \frac{\ell_1(\mathbf{r})}{\ell_2(\mathbf{r})} \right\rceil,$$

where the maximum is taken over \mathbf{r} such that (92) is satisfied (if no such \mathbf{r} exists we can take $D_3 = 0$). Note that D_3 is a constant depending on the model and Δ only.

Now we want to find the maximum weight labeling of G . We are only going to focus on labeling of the distinguished vertices (u 's and v 's in the J_2 gadgets), since once those are fixed one just finds the optimal labeling in each gadget (conditioned on the labels of distinguished vertices). Let W be a labeling of the distinguished vertices that leads to the maximum weight labeling of G . Let \hat{W} be the labeling obtained from W by changing the phase of each distinguished vertex to \mathbf{p} while (1) keeping the original spin on the vertices of H , and (2) making the spin of w_1, \dots, w_{D_3} the opposite of the spin of w (for each $w \in V(H)$). Now we compare W and \hat{W} for each J_2 gadget corresponding to edge of H :

- if in W the phase of u and v were different then \hat{W} has higher weight than W on the gadget, using (89);
- if in W the phase of u and v is the same but the spin is different then \hat{W} has greater or equal weight than W on the gadget, using (87);
- if in W the phase of u and v is the same and the spin is the same that the loss of \hat{W} on the gadget (compared to W) is $\ell_1(\mathbf{r})$ (where \mathbf{r} is the phase of u, v in W).

For the J_2 gadgets connecting w to w_1, \dots, w_{D_3} we have

- if the phase of w in W was \mathbf{r} such that $\ell_1(\mathbf{r}) > 0$ then the gain of \hat{W} on each gadget (compared to W) is at least $\ell_2(\mathbf{r})$;
- otherwise, by (87) then \hat{W} has greater or equal weight than W on the gadget.

For each vertex whose phase in W was \mathbf{r} such that $\ell_1(\mathbf{r}) > 0$ there are 3 edges where \hat{W} can lose $\ell_1(\mathbf{r})$ (compared to W) but there are D_3 edges where \hat{W} gains $\ell_2(\mathbf{r})$ (compared to W). Since $D_3 \ell_2(\mathbf{r}) > 3 \ell_1(\mathbf{r})$ we have that \hat{W} has at least as large weight as W (and hence is also optimal).

Now we just argue how the spins should be assigned. The largest number of J_2 gadgets with opposite spins on the distinguished vertices arises when we take the max-cut of H and assign the spin according to the cut. \square

6.4 Connection between approximating the partition function and the phase labeling problem

In this section we prove Lemma 6.1.

Let $H = (V, E)$ be an instance of the phase labeling problem, where $\{E_p, E_s\}$ is a partition of the edges of H . Let $|V| = m$. The degree of a vertex $v \in V$ will be defined as $2d_s + d_p + 4l_s + 2l_p$, where d_s, d_p are the numbers of symmetric and parallel edges joining v to a distinct vertex u and l_s, l_p are the numbers of symmetric and parallel loops from v to itself. The bounded degree assumption means there is an absolute constant D (not depending on m) which bounds the degree of any $v \in V$.

To approximate the phase labeling problem on H with parameters \mathbf{B}, \mathcal{Q} , we will replace each vertex in the graph H by a suitable graph in a family of gadgets \mathcal{F} . The construction has a parameter k which roughly controls the accuracy of the approximation we want to achieve. The family \mathcal{F} will be of the form $\{G^d\}_{d \in [D]}$ and the gadget for a vertex v will be G^d where d is the degree of v . Note that the cardinality of \mathcal{F} is bounded by the absolute constant D . The gadgets G^d are selected from a graph distribution \mathcal{G}_n^{kd} for some appropriate n to be specified later. For integer r, n satisfying $n > r \geq 0$, we next describe the graph distribution $\mathcal{G}_n^r := \mathcal{G}_n^r(\Delta)$.

1. \mathcal{G}_n^r is supported on bipartite graphs. The two parts of the bipartite graph are labeled by $+, -$ and each is partitioned as $U^s \cup W^s$ where $|U^s| = n$, $|W^s| = r$ for $s = \{+, -\}$. U denotes the set $U^+ \cup U^-$ and similarly W denotes the set $W^+ \cup W^-$.
2. To sample $G \sim \mathcal{G}_n^r$, sample uniformly and independently Δ matchings: (i) $(\Delta - 1)$ perfect matchings between $U^+ \cup W^+$ and $U^- \cup W^-$, (ii) a n -matching between U^+ and U^- . The edge set of G is the union of the Δ matchings. Thus, vertices in U have degree Δ , while vertices in W have degree $\Delta - 1$.

Note that in the special case $r = 0$, the distribution \mathcal{G}_n^r is identical to the graph distribution \mathcal{G}_n defined in Section 2.

Before further specifying the family \mathcal{F} , we first describe the properties that a gadget in \mathcal{F} should have. We assume throughout that r is an arbitrarily large constant (independent of n). Let $G \sim \mathcal{G}_n^r$ and denote by μ_G the Gibbs distribution on G with interaction matrix \mathbf{B} . Note that G is a random graph on $2(n + r)$ vertices.

For $\sigma : U \cup W \rightarrow [q]$, the footprint of σ is a pair of q -dimensional vectors $(\alpha_\sigma, \beta_\sigma)$. The i -th entry of α_σ (resp. β_σ) is equal to $|\sigma^{-1}(i) \cap U^+|/n$ (resp. $|\sigma^{-1}(i) \cap U^-|/n$). Let $\mathbf{p} \in \mathcal{Q}$ and recall that \mathbf{p} corresponds to a dominant phase (α, β) of Ψ_1 . The phase of a configuration $\sigma : U \cup W \rightarrow [q]$ will be denoted by $Y(\sigma)$ and equals \mathbf{p} if the closest⁴ dominant phase to the footprint $(\alpha_\sigma, \beta_\sigma)$ of σ is (α, β) . Note that the phase of σ depends only on the spins of vertices in U .

We shall display shortly that, conditioned on $Y(\sigma) = \mathbf{p}$, the marginal distribution of μ_G on the vertices in W can be well approximated by an appropriate product measure $\nu_{\mathbf{p}}^\otimes(\cdot)$. To do this, recall that every phase $\mathbf{p} \in \mathcal{Q}$ corresponds to a fixpoint of the tree recursions (22). Let $(\hat{R}_1, \dots, \hat{R}_q)$ be a scaled version of (R_1, \dots, R_q) so that $\sum_i \hat{R}_i = 1$ (and define similarly $\hat{C}_1, \dots, \hat{C}_q$). We now define a product measure $\nu_{\mathbf{p}}^\otimes(\cdot)$ on the space of spin assignments to vertices in W . For $\eta : W \rightarrow [q]$ and $\mathbf{p} \in \mathcal{Q}$, let

$$\nu_{\mathbf{p}}^\otimes(\eta) = \prod_{i \in [q]} (\hat{R}_i)^{|\eta^{-1}(i) \cap W^+|} \prod_{j \in [q]} (\hat{C}_j)^{|\eta^{-1}(j) \cap W^-|}. \quad (94)$$

For $\sigma : U \cup W \rightarrow [q]$, denote by σ_W the restriction of σ to vertices in W .

⁴See Appendix B, equation (153) for the precise definition.

Lemma 6.9. *Let r be an arbitrarily large constant. In the setting of Theorem 1.5, for every $\varepsilon > 0$, there exists $N(\varepsilon)$ such that for $n \geq N$, a random graph $G \sim \mathcal{G}_n^r$ satisfies with positive probability the following:*

1. *The graph G is simple.*
2. *For each $\mathbf{p} \in \mathcal{Q}$, $(1 - \varepsilon)/|\mathcal{Q}| \leq \mu_G(Y(\sigma) = \mathbf{p}) \leq (1 + \varepsilon)/|\mathcal{Q}|$. That is, the phases in \mathcal{Q} appear with roughly equal probability.*
3. *Let $\sigma \sim \mu_G$. Then, $\mu_G(\sigma_W = \eta | Y(\sigma) = \mathbf{p}) / \nu_{\mathbf{p}}^{\otimes}(\eta) \in [1 - \varepsilon, 1 + \varepsilon]$ for all $\eta : W \rightarrow [q]$. That is, conditioned on the phase \mathbf{p} of the configuration, the spins of the vertices in W are roughly independent and the marginal distribution on them can be approximated by the distribution $\nu_{\mathbf{p}}^{\otimes}(\cdot)$.*
4. *There is no edge between W^+ and W^- . Moreover, there is no vertex in G which has two neighbors in $W^+ \cup W^-$.*

Lemma 6.9 is proved in Section 6.4.1. An immediate consequence of Lemma 6.9 is the following.

Corollary 6.10. *Let k be an arbitrarily large constant. For $d \in [D]$, let $G^d \sim \mathcal{G}_n^{kd}$ and set $\mathcal{F} = \{G^d\}_{d \in [D]}$. Then, for all sufficiently large n , G^d satisfies Items 1, 2, 3 and 4 of Lemma 6.9 with positive probability for every $d \in [D]$.*

Corollary 6.10 also yields a trivial randomized algorithm to construct the family \mathcal{F} for an arbitrary constant k . In fact, since all the parameters are constants, one can construct the family \mathcal{F} by brute force search. With the family \mathcal{F} in our hands, we can now give the details of the construction.

The first step consists of replacing each vertex $v \in H$ with degree d with a distinct copy of the gadget $G^d \in \mathcal{F}$. We will denote the gadget corresponding to vertex v by G_v and the images of the sets W, W^\pm, U^\pm in G_v by W_v, W_v^\pm, U_v^\pm . Further, denote by \hat{H} the graph obtained by the disconnected copies of the gadgets.

The second step consists of encoding the edges of H in \hat{H} , that is, making connections between the gadgets. The final graph will be denoted by $H_{\mathcal{F}}$. The edges we are going to place will form a perfect matching on $\cup_{v \in H} W_v$ and as a result $H_{\mathcal{F}}$ will be Δ -regular. Every parallel edge of H corresponds to $2k$ edges in $H_{\mathcal{F}}$, while every symmetric to $4k$. Roughly, parallel and symmetric indicate which parts of two gadgets get connected (recall that the gadgets are bipartite). Loops are treated as if they were connecting distinct vertices.

In detail, let (u, v) be an edge e of H . Suppose first that $u \neq v$. If e is parallel, place k edges between W_u^s and W_v^s for $s \in \{+, -\}$. If e is symmetric, place k edges between W_u^s and W_v^s and k edges between W_u^s and W_v^{-s} for $s \in \{+, -\}$. Suppose now that $u = v$. If e is parallel, place k edges between distinct vertices in W_v^+ and k edges between distinct vertices in W_v^- . If e is symmetric, place $2k$ edges between W_v^+ and W_v^- , k edges between distinct vertices in W_v^+ and k edges between distinct vertices in W_v^- .

The first step of the construction guarantees that the second step can be done in a (deterministic) way so that $H_{\mathcal{F}}$ is Δ -regular. Moreover, by Corollary 6.10 and item 4 of Lemma 6.9, $H_{\mathcal{F}}$ is a simple, triangle-free graph.

Proof of Lemma 6.1. Using Corollary 6.10 and specifically items 2 and 3 of Lemma 6.9, the argument in [SS12, Lemma 4.3] almost verbatim gives

$$\frac{(1 - \varepsilon)^{2m}}{|\mathcal{Q}|^m} \leq \frac{Z_{H_{\mathcal{F}}} / Z_{\hat{H}}}{\exp(k \cdot \text{MAXLWT}(H))} \leq (1 + \varepsilon)^m.$$

This can be rearranged into

$$\frac{1}{k} \log \left(\frac{Z_{H_{\mathcal{F}}}}{Z_{\hat{H}}} \right) - \frac{m}{k} \log(1 + \varepsilon) \leq \text{MAXLWT}(H) \leq \frac{1}{k} \log \left(\frac{Z_{H_{\mathcal{F}}}}{Z_{\hat{H}}} \right) - \frac{m}{k} [2 \log(1 - \varepsilon) - \log |\mathcal{Q}|]. \quad (95)$$

The argument in [SS12, Proof of Theorems 1 and 2] gives the desired result. We give the short details. The graph \hat{H} consists of m disconnected subgraphs, each of constant size. Hence, we can compute $Z_{\hat{H}}$ exactly in polynomial time. Assume now that $Z_{H_{\mathcal{F}}}$ can be approximated within a factor of $\exp(c|\hat{H}|)$ in polynomial time for any $c > 0$. Since $\log(Z_{H_{\mathcal{F}}})$ is bounded above by $O(|\hat{H}|)$, the ratio $\log(Z_{H_{\mathcal{F}}}/Z_{\hat{H}})$ can be approximated within an additive $O(c|\hat{H}|) = O[cm(n + kD)] = O(cnm)$ since $n > kD$. Thus, by (95), we obtain upper and lower bounds for $\text{MAXLWT}(H)$ which differ by $O[(cn + 1)m/k]$. A random phase labeling yields the lower bound $\text{MAXLWT}(H) \geq \Omega(m)$. Thus, the final approximation is within a multiplicative factor $O[(cn + 1)/k]$ of $\text{MAXLWT}(H)$. To make the multiplicative factor arbitrarily small, we need to take k large. This might increase n , but we can compensate by taking c small. This concludes the proof. \square

6.4.1 Proof of Lemma 6.9

Let $G \sim \mathcal{G}_n^r$. To get a handle on Items 2 and 3 of Lemma 6.9, we first define the partition functions conditioned on a phase $\mathbf{p} \in \mathcal{Q}$. Similar definitions appear in [Sly10]. Let $\Omega^{\mathbf{p}}$ be the configurations $\sigma \in \Omega$ whose phase $Y(\sigma)$ equals \mathbf{p} , i.e.,

$$\Omega^{\mathbf{p}} = \{\sigma \in \Omega \mid Y(\sigma) = \mathbf{p}\}. \quad (96)$$

Similarly, for a configuration $\eta : W \rightarrow [q]$, let

$$\Omega^{\mathbf{p}}(\eta) = \{\sigma \in \Omega \mid Y(\sigma) = \mathbf{p}, \sigma_W = \eta\}. \quad (97)$$

Note that $\Omega^{\mathbf{p}} = \cup_{\eta} \Omega^{\mathbf{p}}(\eta)$ and $\Omega = \cup_{\mathbf{p} \in \mathcal{Q}} \Omega^{\mathbf{p}}$. The conditioned partition functions $Z_G^{\mathbf{p}}$ and $Z_G^{\mathbf{p}}(\eta)$ are defined as

$$Z_G^{\mathbf{p}}(\eta) := \sum_{\sigma \in \Omega^{\mathbf{p}}(\eta)} w_G(\sigma), \quad Z_G^{\mathbf{p}} := \sum_{\sigma \in \Omega^{\mathbf{p}}} w_G(\sigma) = \sum_{\eta : W \rightarrow [q]} Z_G^{\mathbf{p}}(\eta). \quad (98)$$

The following equalities display the relevance of these quantities to Lemma 6.9.

$$\mu_G(Y(\sigma) = \mathbf{p}) = \frac{Z_G^{\mathbf{p}}}{\sum_{\mathbf{p} \in \mathcal{Q}} Z_G^{\mathbf{p}}}, \quad \mu_G(\sigma_W = \eta \mid Y(\sigma) = \mathbf{p}) = \frac{Z_G^{\mathbf{p}}(\eta)}{Z_G^{\mathbf{p}}}. \quad (99)$$

Note that the definition of $Z_G^{\mathbf{p}}$ also makes sense in the case $r = 0$. Note that for $r = 0$ there are no vertices of degree $\Delta - 1$ (and hence no set W), so the graph distribution \mathcal{G}_n^0 is identical to the graph distribution \mathcal{G}_n defined in Section 2.

To start, we are going to show that Items 2 and 3 of Lemma 6.9 hold in expectation. This is the scope of the following lemma which expresses $\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}], \mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta)]$ in terms of $\mathbf{E}_{\mathcal{G}_n}[Z_G^{\mathbf{p}}]$. Note that $o(1)$ refers to quantities that tend to 0 as $n \rightarrow \infty$.

Lemma 6.11. *Let r be a fixed constant and let \mathbf{p} be a phase, i.e., $\mathbf{p} \in \mathcal{Q}$. There exists a constant $C(\mathbf{p})$ such that for every $\eta : W \rightarrow [q]$, it holds that*

$$\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta)] = (1 + o(1)) C^r \nu_{\mathbf{p}}^{\otimes}(\eta) \mathbf{E}_{\mathcal{G}_n}[Z_G^{\mathbf{p}}], \text{ and thus } \max_{\eta : W \rightarrow [q]} \left| \frac{\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta)]}{\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}]} - \nu_{\mathbf{p}}^{\otimes}(\eta) \right| = o(1). \quad (100)$$

Moreover, when the phases \mathcal{Q} are permutation symmetric, $\mathbf{E}_{\mathcal{G}_n}[Z_G^{\mathbf{p}}] = (1 + o(1))\mathbf{E}_{\mathcal{G}_n}[Z_G^{\mathbf{p}'}]$ for any two phases $\mathbf{p}, \mathbf{p}' \in \mathcal{Q}$ and the constant C in (100) does not depend on the particular phase \mathbf{p} . Consequently, for $\mathbf{p}, \mathbf{p}' \in \mathcal{Q}$

$$\mathbf{E}_{\mathcal{G}_n}[Z_G^{\mathbf{p}}] = (1 + o(1))\mathbf{E}_{\mathcal{G}_n}[Z_G^{\mathbf{p}'}], \text{ and thus } \frac{\mathbf{E}_{\mathcal{G}_n}[Z_G^{\mathbf{p}}]}{\sum_{\mathbf{p} \in \mathcal{Q}} \mathbf{E}_{\mathcal{G}_n}[Z_G^{\mathbf{p}}]} = (1 + o(1)) \frac{1}{|\mathcal{Q}|}. \quad (101)$$

Proof. The second equalities in each of (100) and (101) follow immediately from the first. The latter may be proved by explicit calculations following the same arguments as in [Sly10, Lemma 3.3]. It is worthy to note that (100) holds even if the phases are not permutation symmetric, which is not in general true for (101). \square

In light of Equations (99), (100) and (101), the path to obtain Items 2 and 3 of Lemma 6.9 is now paved: it suffices to show that the conditioned partition functions $Z_G^{\mathbf{p}}(\eta)$ are (with positive probability) arbitrarily close to their expectations for large n . Note that we want this to be simultaneously true for all \mathbf{p} and η , that is, for the same graph G . This in turn requires using in full strength a theorem by [Jan95], which is an extension of the small subgraph conditioning method introduced by [RW94].

We do an exposition of these theorems and their application in Appendix A. For satisfying the reader who is more interested in the proof of Lemma 6.9, the following lemma is a distilled version of the results in Appendix A, yet at the same point containing some important bits which will allow us to motivate it.

Lemma 6.12. *Let $G \sim \mathcal{G}_n^r$ and denote by X_{in} , $i = 1, 2, \dots$, the number of cycles of length $2i$ in G . There exist random variables $W_{mn}^{\mathbf{p}}$, a deterministic function of $X_{1n}, X_{2n}, \dots, X_{mn}$, such that for every $\varepsilon > 0$*

$$\lim_{m \rightarrow \infty} \limsup_{n \rightarrow \infty} \Pr_{\mathcal{G}_n^r} \left(\bigcup_{\mathbf{p}} \bigcup_{\eta} \left[\left| \frac{Z_G^{\mathbf{p}}(\eta)}{\mathbf{E}_{\mathcal{G}_n}[Z_G^{\mathbf{p}}(\eta)]} - W_{mn}^{\mathbf{p}} \right| > \varepsilon \right] \right) = 0. \quad (102)$$

There also exists a positive constant $c > 0$ such that $W_{mn}^{\mathbf{p}} > c$ uniformly in m, n . Moreover, when the phases \mathcal{Q} are permutation symmetric, the random variables $W_{mn}^{\mathbf{p}}$ do not depend on the phase \mathbf{p} .

Lemma 6.12 provides a straightforward proof of Lemma 6.9, so we shall elucidate its most important aspects in an attempt to demystify its rather unintuitive statement. Equation (102) says that for all sufficiently large m, n the random variables $Z_G^{\mathbf{p}}(\eta)/\mathbf{E}_{\mathcal{G}_n}[Z_G^{\mathbf{p}}(\eta)]$ are well-approximated by the variables $W_{mn}^{\mathbf{p}}$, with large probability. To get a feeling about this statement, it is well known fact that a random Δ -regular graph is locally tree-like and its girth diverges as $n \rightarrow \infty$. That is, as n grows large, for any positive integer t , for all but $o(n)$ vertices, the t -depth neighborhood of a vertex is isomorphic to the first t levels of the infinite Δ -regular tree. This is in alignment with the fact that $\mathbf{E}_{\mathcal{G}_n}[Z_G^{\mathbf{p}}(\eta)]$ is determined by the Gibbs measure on the infinite Δ -regular tree associated to the phase \mathbf{p} . On the other hand, a graph $G \sim \mathcal{G}_n^r$ does have $o(n)$ vertices which are contained in constant sized cycles. Thus, it is reasonable to expect that $Z_G^{\mathbf{p}}(\eta)$ fluctuates from its expectation. It is equally reasonable to expect the fluctuations to depend on the presence of small cycles which occur with small but non-zero probability. Equation (102) thus provides an explicit handle on these fluctuations, given by the variables $W_{mn}^{\mathbf{p}}$, which are a deterministic function of the small cycle counts in G . Crucially for our proof of Lemma 6.9, when the phases are permutation symmetric, the fluctuations from the expectation are captured by a single random variable, which allows us to control them uniformly over all the phases \mathbf{p} and configurations η .

We should point out that the notation $W_{mn}^{\mathbf{p}}$ should not be confused by any means to the labeling of the degree $\Delta - 1$ vertices in G , i.e., the set of vertices W .

Proof of Lemma 6.9. We assume that the ε in the statement of the lemma is fixed. Let $\varepsilon' > 0$ be sufficiently small, to be picked later.

By Lemma 6.12, for all m, n sufficiently large the random variables $Z_G^{\mathbf{p}}(\eta)/\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta)]$ are well approximated by $W_{mn}^{\mathbf{p}}$ with large probability. That is, there exist $M(\varepsilon')$, $N(\varepsilon')$ such that for $m \geq M$ and $n \geq N$, it holds with probability $1 - \varepsilon'$ over the choice of the graph G that, for every phase \mathbf{p} and every configuration $\eta : W \rightarrow [q]$,

$$Z_G^{\mathbf{p}}(\eta) = (W_{mn}^{\mathbf{p}} \pm \varepsilon')\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta)]. \quad (103)$$

We will show that whenever this is the case, Items 2 and 3 hold. To do this, sum (103) over η to obtain that for each phase \mathbf{p} , it holds

$$Z_G^{\mathbf{p}} = (W_{mn}^{\mathbf{p}} \pm \varepsilon')\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}], \quad (104)$$

Using the positive constant c in Lemma 6.12, we obtain that for ε' sufficiently smaller than c , the ratio $Z_G^{\mathbf{p}}(\eta)/Z_G^{\mathbf{p}}$ is within a multiplicative $(1 \pm \varepsilon)$ from $\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta)]/\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}]$. This gives Item 3 of the lemma, when used in conjunction with (99) and (101). Note that this part of the argument did not use that the phases \mathbf{p} are permutation symmetric.

To obtain Item 2, we have to use that the phases \mathbf{p} are permutation symmetric. Then $W_{mn}^{\mathbf{p}} =: W_{mn}$ by the last assertion in Lemma 6.12. Thus, a summation of (104) over $\mathbf{p} \in \mathcal{Q}$ gives $Z_G^{\mathbf{p}} = (W_{mn} \pm \varepsilon')\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}]$. Exactly the same reasoning as before yields the thesis.

It is a standard union bound to show that Item 4 holds with probability $1 - O(1/n)$ over the choice of the graph G , essentially because G is an expander. Perhaps the second assertion there requires a brief proof sketch. Let $v \in V_1$, $w_1, w_2 \in W^-$ and let E_i be the event that (v, w_i) is an edge of G . The events E_1, E_2 are negatively correlated since v has a fixed number of edges incident to it, either Δ or $\Delta - 1$. It is also easy to see that $\Pr_{\mathcal{G}_n^r}(E_i) \leq 1 - (1 - 1/n)^\Delta = O(1/n)$, so that $\Pr_{\mathcal{G}_n^r}(E_1 \cap E_2) = O(1/n^2)$. A union bound over the roughly $nr^2 = O(n)$ possibilities of the vertices v, w_1, w_2 gives the desired bound.

Thus, a graph $G \sim \mathcal{G}_n^r$ satisfies Items 2, 3 and 4 with large probability for all sufficiently large n . The first assertion in Item 1 of Lemma 6.12 can hence be guaranteed by contiguity, see [Jan95, Section 2]. \square

7 Dominant phases for Potts Model and Colorings

7.1 Proof outline

In this section we prove Theorem 1.6 which establishes the hypotheses of Theorem 1.5 for the dominant phases of the antiferromagnetic Potts and colorings models on random Δ -regular bipartite graphs (and, as we showed in Section 1.2.3, Theorems 1.1 and 1.2 follow as corollaries).

Recall, the interaction matrix \mathbf{B} for the Potts model is completely determined by a parameter B , which is equal to $\exp(-\beta)$ where β is the inverse temperature in the standard notation for the Potts model. The antiferromagnetic regime corresponds to $0 < B < 1$. The coloring model is the zero temperature limit of the Potts model and corresponds to the particular case $B = 0$ in what follows. We should note that in Statistical Physics terms, the arguments of this section are closely related to the phase diagrams of the models.

By Theorem 4.1 specified to the antiferromagnetic Potts and colorings models, studying the global maxima of Ψ_1 is equivalent to studying the global maxima of Φ . Moreover, the global maxima

of Φ and Ψ_1 occur at their critical points. Since there is a one-to-one correspondence between the critical points of Φ and the critical points of Ψ_1 (given by (23)), we will freely interchange our focus between critical points of Φ and Ψ_1 .

The critical points of Φ , by the first part of Theorem 4.1, are given by fixpoints of the tree recursions (22), which for the Potts model read as:

$$R_i \propto \left(BC_i + \sum_{j \neq i} C_j \right)^d, \quad C_j \propto \left(BR_j + \sum_{i \neq j} R_i \right)^d, \quad (105)$$

where $i, j = 1, \dots, q$ and d is the notational convenient substitution $d := \Delta - 1 \geq 2$. Given a fixpoint of the tree recursions (105), we will classify whether it is a Hessian local maximum of Ψ_1 using Theorem 4.2.

Once we find the global maxima of Ψ_1 , it will be simple to prove that they are Hessian and permutation symmetric. Finding however the global maxima of Ψ_1 is going to be more intricate, mainly because the number of local maxima varies according to the value of B . We will thus have to compare the values of Ψ_1 at the critical points. Rather than doing this directly (which seems as a difficult task), we solve a relaxed optimisation problem, which for q even can be tied to the maximization of Ψ_1 . We next give the details.

We begin our considerations by examining when a fixpoint (105) is *translation invariant*, i.e., satisfies $R_i \propto C_i$ for every $i \in [q]$.

Lemma 7.1. *Let $0 \leq B < 1$ and $\Delta \geq 3$. If a solution of (105) satisfies $R_i \propto C_i$ for $i \in [q]$, then it holds that $R_1 = \dots = R_q$ and $C_1 = \dots = C_q$.*

Proof of Lemma 7.1. By the symmetries of the model, we may assume an arbitrary ordering of the R_i 's. Since $0 \leq B < 1$, (105) easily implies the reverse ordering of the C_i 's. Thus, $R_i \propto C_i$ for every $i \in [q]$ yields that the ordering must be trivial, i.e, $R_1 = \dots = R_q$ and $C_1 = \dots = C_q$. \square

Corollary 7.2. *Translation invariant fixpoints of (105) always exist and are unique up to scaling.*

We next explore in which regimes of B , the critical points of Φ consist solely of translation invariant fixpoints. In this regime, we immediately obtain by Theorem 4.1 that the global maximum of Ψ_1 (and hence the global maximum of Φ as well) is achieved at a translation invariant fixpoint.

Lemma 7.3. *Let $0 \leq B < 1$ and $q, \Delta \geq 3$. When $B \geq \frac{\Delta-q}{\Delta}$, the solution of the system of equations (105) satisfies $R_1 = \dots = R_q$ and $C_1 = \dots = C_q$.*

The proof of Lemma 7.3 is an extension of an argument in [BW02] for colorings and is given in Appendix C. The next lemma states that in the complementary regime of Lemma 7.3, the translation invariant fixpoint does not correspond to a local maximum of Ψ_1 and hence, by Theorem 4.1, the global maximum of Ψ_1 occurs at a fixpoint of (105) which is not translation invariant. In particular, in this regime we have semi-translational non-uniqueness.

Lemma 7.4. *For $0 \leq B < \frac{\Delta-q}{\Delta}$, the global maximum of Ψ_1 is not achieved at the translation invariant fixpoint.*

Proof of Lemma 7.4. We apply Theorem 4.2 by showing that the translation invariant fixpoint is Jacobian unstable and hence not a local maximum of Ψ_1 . By Lemma 4.16, for a general interaction matrix \mathbf{B} , the condition for Jacobian stability of a fixpoint of the tree recursions is related to the spectrum of $\mathbf{L} = \begin{bmatrix} \mathbf{0} & \mathbf{A} \\ \mathbf{A}^\top & \mathbf{0} \end{bmatrix}$, where \mathbf{A} is the $q \times q$ matrix whose ij -entry is given by $A_{ij} = B_{ij}R_iC_j/\sqrt{\alpha_i\beta_j}$ and α_i, β_j are given by (29). Recall that ± 1 are eigenvalues of \mathbf{L} and the condition

for Jacobian stability is that all the other eigenvalues have absolute value less than $1/(\Delta - 1)$ (see for details the proof of Theorem 4.2 in Section 4.2.3).

In the setting of the lemma, the matrix \mathbf{A} for the translation invariant fixpoint has off-diagonal entries equal to $1/(B + q - 1)$ and diagonal entries equal to $B/(B + q - 1)$. It follows that the eigenvalues of \mathbf{L} are ± 1 by multiplicity 1 and $\pm(1 - B)/(B + q - 1)$ by multiplicity $q - 1$. The absolute value of the latter is greater than $\frac{1}{\Delta - 1}$ for $0 \leq B < \frac{\Delta - q}{\Delta}$, as claimed. \square

We summarize the above results into the following corollary.

Corollary 7.5. *Let $0 \leq B < 1$ and $q, \Delta \geq 3$. When $B \geq \frac{\Delta - q}{\Delta}$, Ψ_1 has a unique global maximum for $\alpha_1 = \dots = \alpha_q = \beta_1 = \dots = \beta_q = 1/q$ or, in other words, the global maximum of Ψ_1 is achieved by the fixpoint which corresponds to the (unique) translation invariant Gibbs measure. In the complementary regime $0 \leq B < \frac{\Delta - q}{\Delta}$, the maximum of Ψ_1 is not achieved at the translation invariant fixpoint, and hence it is achieved at a semi-translation invariant fixpoint which is not translation invariant.*

Corollary 7.5 is not sufficient to obtain Theorems 1.1 and 1.2, since we need to verify that the global maxima of Ψ_1 in semi-translational non-uniqueness are Hessian and permutation symmetric. We do this by identifying the critical points which are maxima of Ψ_1 .

To state the result, we first need the following structural statement for the solutions of equations (105), namely that solutions of (105) are supported on at most 3 values for the R_i 's and similarly for the C_i 's.

Lemma 7.6. *Let $(R_1, \dots, R_q, C_1, \dots, C_q)$ be a positive solution of the system (105). Let t_R be the number of values on which the R_i 's are supported and define similarly t_C . Then $t_R, t_C \leq 3$ and $t_R = t_C =: t$.*

The proof of Lemma 7.6 is given in Section 7.5. Lemma 7.6 motivates the following definition.

Definition 7.7. *From Lemma 7.6, the R_i 's and C_j 's of a fixpoint of (105) attain at most $t \leq 3$ different values. Let $\tilde{R}_1, \dots, \tilde{R}_t$ and $\tilde{C}_1, \dots, \tilde{C}_t$ be their values and let $q_1, \dots, q_t \geq 1$ be their multiplicities. When $t = 1$, define $q_2 = q_3 = 0$; when $t = 2$, define $q_3 = 0$; when $q_i = 0$, define the values of \tilde{R}_i, \tilde{C}_i to be zero. The corresponding solution of (105) or equivalently the fixpoint of the tree recursions is then defined to be of type (q_1, q_2, q_3) . Note that $q_1 + q_2 + q_3 = q$ and the q_i 's are non-negative integers. Call a (q_1, q_2, q_3) -type fixpoint to be t -supported if the number of q_i 's which are non-zero equals t .*

Finding the types of fixpoints which correspond (via (23)) to global maxima of Ψ_1 is not a trivial task. While 2-supported fixpoints are simple to handle, this is not the case for 3-supported fixpoints. The main lemma we prove is the following, which identifies the type of fixpoints which maximize Ψ_1 .

Lemma 7.8. *For $0 \leq B < \frac{\Delta - q}{\Delta}$ and even $q \geq 3$, the maximum of Ψ_1 over (q_1, q_2, q_3) -type solutions of (105) is attained at fixpoints of type $(q/2, q/2, 0)$.*

The final piece is to show that fixpoints of type $(q/2, q/2, 0)$ are Hessian maxima of Ψ_1 and permutation symmetric. This is the scope of the next lemma, whose proof is given in Section 7.5.

Lemma 7.9. *For $0 \leq B < \frac{\Delta - q}{\Delta}$ and even $q \geq 3$, fixpoints of type $(q/2, q/2, 0)$ are Jacobian stable and hence correspond to Hessian maxima of Ψ_1 . The values of R_i 's and C_j 's for fixpoints of type $(q/2, q/2, 0)$ are unique up to scaling and permutations of the colours.*

We are now ready to prove Theorem 1.6.

Proof of Theorem 1.6. Item 1 follows from Corollary 7.5 (see also Lemmas 7.3 and 7.4). Item 2 follows from Lemmas 7.8 and 7.9, after using the correspondence between fixpoints of the tree recursions (105) and dominant phases of Theorem 4.1 (equation (23)). \square

7.2 Proof of Lemma 7.8

In this section, we outline the proof of Lemma 7.8. We need to find the type(s) of the fixpoints which maximize Ψ_1 . Let $\mathbf{q} = (q_1, q_2, q_3)$ specify the type of a fixpoint of (105) and let $\mathbf{r} = (R_1, R_2, R_3)$, $\mathbf{c} = (C_1, C_2, C_3)$ be the respective values of the R_i 's and C_j 's, see Definition 7.7. Note that the q_i 's are non-negative integers satisfying $q_1 + q_2 + q_3 = q$.

Using Theorem 4.1, we obtain that the value of $\Psi_1(\boldsymbol{\alpha}, \boldsymbol{\beta})$ corresponding to this fixpoint of (105) is given by the value of the function $\overline{\Phi}^S$, where

$$\begin{aligned} \overline{\Phi}^S(\mathbf{q}, \mathbf{r}, \mathbf{c}) := & (d+1) \ln \left(\sum_{i=1}^3 q_i R_i \sum_{j=1}^3 q_j C_j + (B-1) \sum_i q_i R_i C_i \right) \\ & - d \ln \left(\sum_{i=1}^3 q_i R_i^{(d+1)/d} \right) - d \ln \left(\sum_{j=1}^3 q_j C_j^{(d+1)/d} \right), \end{aligned} \quad (106)$$

and $d = \Delta - 1$. It is a non-trivial task to directly compare the values of $\overline{\Phi}^S$ over fixpoints of (105). Instead, we will solve a relaxed version of the problem, seeking to maximize $\overline{\Phi}^S$ over non-negative q_i 's which satisfy $q_1 + q_2 + q_3 = q$. If this maximum happens to occur for integer \mathbf{q} and the respective values of R_i 's and C_j 's are solutions of (105), then we have also found the solution to the original maximization problem. It turns out that all of the above are satisfied iff q is even.

To formalize the argument, for non-negative q_i 's such that $q_1 + q_2 + q_3 = q$, define

$$\overline{\Phi}(\mathbf{q}) := \max_{\mathbf{r}, \mathbf{c}} \overline{\Phi}^S(\mathbf{q}, \mathbf{r}, \mathbf{c}) \quad (107)$$

where the maximum is over $\mathbf{r} = (R_1, R_2, R_3)^\top$, $\mathbf{c} = (C_1, C_2, C_3)^\top$ which satisfy

$$\begin{aligned} \sum_{i=1}^3 q_i R_i \sum_{j=1}^3 q_j C_j + (B-1) \sum_{i=1}^3 q_i R_i C_i &> 0, \\ R_1, R_2, R_3, C_1, C_2, C_3 &\geq 0. \end{aligned} \quad (108)$$

It is simple to see that in the region (108), $\overline{\Phi}^S$ is well defined. It is not completely immediate that the maximum in (107) is well defined since the region (108) is not compact. This is a consequence of the following scale-free property of $\overline{\Phi}^S$ with respect to \mathbf{r} and \mathbf{c} :

$$\text{for every } c_1, c_2 > 0 \text{ it holds that } \overline{\Phi}^S(\mathbf{q}, c_1 \mathbf{r}, c_2 \mathbf{c}) = \overline{\Phi}^S(\mathbf{q}, \mathbf{r}, \mathbf{c}). \quad (109)$$

Using (109), it is simple to obtain the following.

Lemma 7.10. *Let $B \geq 0$ and $q \geq 2$. For all $q_1, q_2, q_3 \geq 0$ which satisfy $q_1 + q_2 + q_3 = q$, the maximum in (107) is well defined. Moreover, the maximum of $\overline{\Phi}(q_1, q_2, q_3)$ over all such q_1, q_2, q_3 is attained.*

We next seek to connect the maximizers of (107) with solutions of (105). To do this, we first need to consider whether the maximum in (107) happens on the boundary of the region (108); it turns out that the maximum can happen at the boundary $R_i = 0$ or $C_i = 0$ if q_i is close to zero. While the boundary cases are an artifact of allowing q_i 's to be non-integer, we will need to treat them explicitly to find the maximum of $\overline{\Phi}$.

Definition 7.11. A triple $\mathbf{q} = (q_1, q_2, q_3)$ is good if the \mathbf{r}, \mathbf{c} which achieve the maximum in (107) satisfy: for $i = 1, 2, 3$, $q_i > 0$ implies $R_i, C_i > 0$. A triple $\mathbf{q} = (q_1, q_2, q_3)$ is bad if it is not good.

To complete the connection, we need to further restrict the set of triples \mathbf{q} . To motivate this restriction, note that if we consider the region (108) in the subspace $R_1 = R_2$ and $C_1 = C_2$, we obtain $\bar{\Phi}(q_1 + q_2, q_3, 0) \leq \bar{\Phi}(q_1, q_2, q_3)$. To avoid degenerate cases, we consider only triples \mathbf{q} where such simple inequalities do not hold at equality.

Definition 7.12. Let $t = 2$ or 3 . A triple $\mathbf{q} = (q_1, q_2, q_3)$ is called t -maximal if exactly t of the q_i 's are non-zero and for all distinct $i, j, k \in \{1, 2, 3\}$ (with $q_i q_j > 0$) it holds that $\bar{\Phi}(q_i + q_j, q_k, 0) < \bar{\Phi}(\mathbf{q})$.

Our interest is in maximal good triples $\mathbf{q} = (q_1, q_2, q_3)$. This is justified by the following lemma, whose proof is given in Section 7.5.

Lemma 7.13. Suppose that q_1, q_2, q_3 are non-negative integers and the triple $\mathbf{q} = (q_1, q_2, q_3)$ is t -maximal and good. Then, the \mathbf{r}, \mathbf{c} which achieve the maximum in (105) specify a t -supported fixpoint of (105) of type (q_1, q_2, q_3) .

Thus to prove Lemma 7.8, it suffices to prove that the triple $(q/2, q/2, 0)$ is 2-maximal and good and that the maximum of $\bar{\Phi}(\mathbf{q})$ is achieved at $(q/2, q/2, 0)$. The next lemma examines which maximal good triples can be a maximum of $\bar{\Phi}$.

Lemma 7.14. Let $q \geq 3$ and $0 \leq B < 1$. There do not exist 3-maximal good triples \mathbf{q} which maximize $\bar{\Phi}(\mathbf{q})$. The only 2-maximal good triples \mathbf{q} where a maximum of $\bar{\Phi}(\mathbf{q})$ can occur are $(q/2, q/2, 0)$ or its permutations.

Lemma 7.14 is not sufficient to yield Lemma 7.8 because the maximum of $\bar{\Phi}(\mathbf{q})$ can occur at a bad triple \mathbf{q} . This possibility is excluded by the following lemma.

Lemma 7.15. Let $q \geq 3$ and $0 \leq B < \frac{\Delta - q}{\Delta}$. There do not exist bad triples \mathbf{q} which maximize $\bar{\Phi}(\mathbf{q})$.

Using Lemmas 7.14 and 7.15, we can now give the proof of Lemma 7.8.

Proof of Lemma 7.8. The maximum of $\bar{\Phi}(\mathbf{q})$ over triples \mathbf{q} is attained by Lemma 7.10. This maximum can happen either at a bad or a good triple \mathbf{q} . Maxima at bad triples \mathbf{q} are excluded by Lemma 7.15. Maxima at 3-maximal good triples are excluded by the first part of Lemma 7.14. Thus, the maximum must happen at a (good) triple of the form $\mathbf{q} = (q_1, q_2, 0)$. The latter can be either 2-maximal or not. If it is not 2-maximal, the maximum must equal $\bar{\Phi}(\mathbf{q})$, which in the regime $0 \leq B < \frac{\Delta - q}{\Delta}$ is excluded by Lemma 7.4. Thus, the maximum must happen at a 2-maximal good triple, which Lemma 7.14 asserts that it must be the triple $(q/2, q/2, 0)$. Finally, for q even, by Lemma 7.13 the \mathbf{r}, \mathbf{c} which achieve the maximum in (107) correspond to a 2-supported fixpoint of (105) of type $(q/2, q/2, 0)$, as wanted. \square

For the proofs of Lemmas 7.14 and 7.15, we will often perturb the values of q_i 's. The following lemma, which is proved in Section 7.5 will be very helpful.

Lemma 7.16. Let $\mathbf{q} = (q_1, q_2, q_3)$ and $I = \{i \mid q_i > 0\}$. Suppose that \mathbf{r}, \mathbf{c} achieve the maximum in (107). Then, for $i \in I$ it holds that

$$\frac{\partial \bar{\Phi}^S}{\partial q_i}(\mathbf{q}, \mathbf{r}, \mathbf{c}) = \frac{R_i \sum_j q_j C_j + C_i \sum_j q_j R_j + (d-1)(1-B)R_i C_i}{\sum_j q_j R_j \sum_j q_j C_j + (B-1) \sum_j q_j R_j C_j}. \quad (110)$$

Moreover, if there exist $i, j \in I$ such that $\frac{\partial \bar{\Phi}^S}{\partial q_i} - \frac{\partial \bar{\Phi}^S}{\partial q_j} \neq 0$, the maximum of $\bar{\Phi}$ is not achieved at the triple \mathbf{q} .

7.3 Good triples: proof of Lemma 7.14

We first prove the statement of the lemma for 3-maximal good triples $\mathbf{q} = (q_1, q_2, q_3)$, the proof for 2-maximal good triples will easily be inferred by appropriately modifying the arguments in the special case $q_2 = 0$.

Let $\mathbf{q} = (q_1, q_2, q_3)$ be a 3-maximal good triple. Since \mathbf{q} is 3-maximal all of the q_i 's are positive. Moreover, \mathbf{q} is good, and hence the maximum in (107) for \mathbf{q} is attained at positive R_i 's and C_j 's. Thus, the R_i 's and C_j 's satisfy $\partial\Phi^S/\partial R_i = \partial\Phi^S/\partial C_j = 0$ which give

$$R_i^{1/d} \propto q_1 C_1 + q_2 C_2 + q_3 C_3 + (B-1)C_i, \quad C_j^{1/d} \propto q_1 R_1 + q_2 R_2 + q_3 R_3 + (B-1)R_j. \quad (111)$$

Since \mathbf{q} is 3-maximal, we may assume that \mathbf{r} is such that $R_i \neq R_j$ for all $i \neq j$. Otherwise, if for example $R_1 = R_2$, by (111), we have $C_1 = C_2$ as well, so that $\bar{\Phi}(q_1, q_2, q_3) = \bar{\Phi}(q_1 + q_2, q_3, 0)$, contradicting the 3-maximality of \mathbf{q} . Thus, we may assume a strict ordering of the R_i 's, which by (111) implies the reverse ordering of the C_j 's. W.l.o.g., we will use the following ordering:

$$R_1 > R_2 > R_3 > 0 \quad \text{and} \quad 0 < C_1 < C_2 < C_3. \quad (112)$$

The following lemma, together with the second part of Lemma 7.16, establishes that the maximum of $\bar{\Phi}$ cannot occur at a 3-maximal triple.

Lemma 7.17. *Suppose that R_i 's and C_j 's satisfy (111) and (112). If $R_1/R_3 \neq C_3/C_1$ then $\frac{\partial\Phi^S}{\partial q_1} - \frac{\partial\Phi^S}{\partial q_3} \neq 0$. If $R_1/R_3 = C_3/C_1$ then $\frac{\partial\Phi^S}{\partial q_1} - \frac{\partial\Phi^S}{\partial q_2} \neq 0$.*

We next give the proof of Lemma 7.17. We will utilize Lemma 7.16 by specifying a particular scaling of the R_i 's and C_j 's which will be beneficial. To do this, set

$$r_1^d = R_1/R_3, r_2^d = R_2/R_3, c_2^d = C_2/C_1, c_3^d = C_3/C_1. \quad (113)$$

The R_i 's and C_j 's may be recovered from r_i 's, c_j 's using

$$R_1 \propto r_1^d, R_2 \propto r_2^d, R_3 \propto 1, \quad \text{and} \quad C_1 \propto 1, C_2 \propto c_2^d, C_3 \propto c_3^d. \quad (114)$$

Translating (112) into r_1, r_2, c_2, c_3 gives

$$r_1 > r_2 > 1 \quad \text{and} \quad c_3 > c_2 > 1. \quad (115)$$

Moreover, dividing appropriate pairs of (111), we also obtain

$$\begin{aligned} r_1 &= \frac{B + q_1 - 1 + q_2 c_2^d + q_3 c_3^d}{q_1 + q_2 c_2^d + (B + q_3 - 1) c_3^d}, \quad c_3 = \frac{B + q_3 - 1 + q_2 r_2^d + q_1 r_1^d}{q_3 + q_2 r_2^d + (B + q_1 - 1) r_1^d}, \\ r_2 &= \frac{q_1 + (B + q_2 - 1) c_2^d + q_3 c_3^d}{q_1 + q_2 c_2^d + (B + q_3 - 1) c_3^d}, \quad c_2 = \frac{q_3 + (B + q_2 - 1) r_2^d + q_1 r_1^d}{q_3 + q_2 r_2^d + (B + q_1 - 1) r_1^d}. \end{aligned} \quad (116)$$

It can easily be verified that this system of equations gives

$$q_1 = \frac{(1-B)f(r_1, c_3) + q_2 P(c_2^d - c_3^d r_2^d)}{P(r_1^d c_3^d - 1)}, \quad q_3 = \frac{(1-B)f(c_3, r_1) + q_2 P(r_2^d - r_1^d c_2^d)}{P(r_1^d c_3^d - 1)} \quad (117)$$

$$r_2 = \frac{r_1 c_3^d - 1 - c_2^d(r_1 - 1)}{c_3^d - 1}, \quad r_2^d = \frac{r_1^d c_3 - 1 - c_2(r_1^d - 1)}{c_3 - 1}, \quad (118)$$

$$f(x, y) := x^{d+1} y^{d+1} - x^d y^{d+1} - x y^{d+1} + y^d + y - 1, \quad P := (r_1 - 1)(c_3 - 1) > 0.$$

We will need the following lemma.

Lemma 7.18. Assume that $q_1, q_2, q_3, r_1, r_2, c_2, c_3$ satisfy (115), (117), (118). If $r_1 = c_3$ then $r_2 = c_2$ and $q_1 = q_3$.

Proof of Lemma 7.18. We prove that $r_1 = c_3$ implies $r_2 = c_2$. Once this is done, (117) easily gives that $r_1 = c_3$ implies $q_1 = q_3$ as well, thus proving the lemma.

So, suppose that $z = r_1 = c_3$ and for the sake of contradiction assume $r_2 \neq c_2$. By (115) we obtain that $r_2, c_2 \in (1, z)$. Eliminating r_2 from (118) we obtain that c_2 (and by a symmetric argument r_2) satisfies

$$g(s) := \left(\frac{z^{d+1} - 1 - s^d(z-1)}{z^d - 1} \right)^d + \frac{s(z^d - 1) - (z^{d+1} - 1)}{z - 1} = 0.$$

In fact, $g(1) = g(z) = 0$ as well, so that g has at least four distinct roots in $[1, z]$. It follows that $g'(s) = 0$ has at least three distinct solutions in $[1, z]$, say s_i for $i = 1, 2, 3$. As a consequence of $g'(s_i) = 0$, we easily obtain that the s_i 's satisfy $h(s_i) = c$ where $h(s) := (z^{d+1} - 1)s - s^{d+1}(z - 1)$ and c is a constant which depends only on z, d . Thus, $h'(s) = 0$ has at least two distinct solutions in $[1, z]$ which is clearly absurd. \square

Proof of Lemma 7.17. Set

$$DIF_{13} := \frac{\partial \overline{\Phi^S}}{\partial q_1} - \frac{\partial \overline{\Phi^S}}{\partial q_3}, \quad DIF_{12} := \frac{\partial \overline{\Phi^S}}{\partial q_1} - \frac{\partial \overline{\Phi^S}}{\partial q_2}, \quad S := \sum_i q_i R_i \sum_j q_j C_j + (B-1) \sum_i q_i R_i C_i.$$

We use the expressions (110) for the derivatives. The denominators in the expressions are the same, so we may ignore them. Moreover, the expressions therein are scale-free, consequently in order to write the derivatives with respect to r_i 's and c_j 's we just need to make the substitutions (114).

To prove the first part of the lemma, we eliminate q_1, q_3 from the resulting expression for DIF_{13} using (117). This substitution has the beneficial effect of eliminating q_2, r_2 from the final expression. After straightforward calculations, we obtain the following:

$$DIF_{13} = -\frac{(1-B)g(r_1, c_3)}{S(r_1-1)(c_3-1)}, \quad \text{where} \quad (119)$$

$$g(r_1, c_3) := (r_1 - c_3)(r_1^d - 1)(c_3^d - 1) - d(r_1 - 1)(c_3 - 1)(r_1^d - c_3^d).$$

It can easily be seen that for $r_1, c_3 > 1$, it holds that $g(r_1, c_3) = 0$ iff $r_1 = c_3$ iff $R_1/R_3 = C_3/C_1$ as desired.

We next prove the second part of the lemma. Since $R_1/R_3 = C_3/C_1$, we have $r_1 = c_3$ and by Lemma 7.18, $r_2 = c_2$ and $q_1 = q_3$. Using these, (117) and (118) simplify to

$$q_1 = \frac{(1-B)(r_1^{d+1} - 1) - q_2 r_2^d (r_1 - 1)}{(r_1 - 1)(r_1^d + 1)}, \quad r_2 = \frac{r_1^{d+1} - 1 - r_2^d (r_1 - 1)}{r_1^d - 1}. \quad (120)$$

Moreover, using the substitutions (114) and $q_1 = q_3$, we obtain

$$DIF_{12} = \frac{(q_1 r_1^d + q_2 r_2^d + q_1)(r_1^d - 2r_2^d + 1) + (d-1)(1-B)(r_1^d - r_2^{2d})}{S}$$

$$= -\frac{(1-B)[(d-1)(r_1 - 1)r_2^{2d} + 2r_2^d(r_1^{d+1} - 1) - (r_1^{2d+1} + dr_1^{d+1} - dr_1^d - 1)]}{(r_1 - 1)S},$$

where in the second equality we substituted the value of q_1 from (120). Observe that the numerator is a quadratic polynomial in r_2^d and, by inspection, for $r_1 > 1$, its roots are of opposite sign. Thus, $DIF_{12} = 0$ iff $r_2^d = \rho_1$, where

$$\rho_1(r_1) := \frac{\sqrt{D} - (r_1^{d+1} - 1)}{(d-1)(r_1 - 1)} \text{ and } D := (dr_1^{d+1} - (d-1)r_1^d + 1)(r_1^{d+1} + (d-1)r_1 - d).$$

For the sake of contradiction, suppose that $r_2^d = \rho_1$. Then (120) gives that $r_2 = \rho_2$, where

$$\rho_2(r_1) := \frac{d(r_1^{d+1} - 1) - \sqrt{D}}{(d-1)(r_1^d - 1)}.$$

Thus $\rho_1 = \rho_2^d$. We obtain a contradiction by showing that for every $r_1 > 1$, it holds that $\rho_2^d < \rho_1$ or equivalently $d \ln \rho_2 < \ln \rho_1$. It is easy to see that in the limit $r_1 \downarrow 1$ the inequality is satisfied at equality, thus it suffices to prove that the derivative of the rhs w.r.t r_1 is greater than the respective derivative of the l.h.s. for $r_1 > 1$.

This differentiation is cumbersome but otherwise straightforward. The final result is

$$\begin{aligned} \frac{1}{\rho_1} \frac{\partial \rho_1}{\partial r_1} - \frac{d}{\rho_2} \frac{\partial \rho_2}{\partial r_1} &= \frac{(d+1)g(r_1)h(r_1)}{2(r_1 - 1)(r_1^d - 1) \left(\sqrt{D} - (r_1^{d+1} - 1) \right) \left(d(r_1^{d+1} - 1) - \sqrt{D} \right)}, \\ g(r_1) &:= r_1^{2d} - d^2 r_1^{d+1} + 2(d^2 - 1)r_1^d - d^2 r_1^{d-1} + 1, \\ h(r_1) &:= (d+1)(r_1^{d+1} - 1) - (d-1)(r_1^d - 1) - 2\sqrt{D}. \end{aligned} \quad (121)$$

Note that the denominator in the r.h.s. of (121) is positive for $r_1 > 1$: the terms involving \sqrt{D} are positive since they are the numerators of ρ_1, ρ_2 . The final part of the proof consists of proving that $g(r_1) > 0$ and $h(r_1) > 0$ for $r_1 > 1$.

The polynomial g has 4 sign changes and hence, by the Descartes' rule of signs has at most 4 positive roots. In fact, a tedious calculation shows that $r_1 = 1$ is a root by multiplicity 4, thus proving that $g(r_1) > 0$ for $r_1 > 1$. To prove that $h(r_1) > 0$ for $r_1 > 1$, note the identity

$$[(d+1)(r_1^{d+1} - 1) - (d-1)(r_1^d - 1)]^2 - 4D = (d-1)^2(r_1 - 1)^2(r_1^d - 1)^2.$$

This completes the proof. \square

To prove the second part of Lemma 7.14, assume that $\mathbf{q} = (q_1, q_2, q_3)$ is a 2-maximal good triple. Since \mathbf{q} is 2-maximal, w.l.o.g. we may assume that $q_2 = 0$. Note that the values of R_2, C_2 do not affect the value of the derivatives $\partial \bar{\Phi}^S / \partial q_1, \partial \bar{\Phi}^S / \partial q_3$ when $q_2 = 0$. Similarly, (117) continues to hold even when $q_2 = 0$. Thus, the proof of the first part of Lemma 7.17 carries through verbatim. In particular, if $R_1/R_3 \neq C_3/C_1$, then $\partial \bar{\Phi}^S / \partial q_1 - \partial \bar{\Phi}^S / \partial q_3 \neq 0$. By the second part of Lemma 7.16, it follows that $\mathbf{q} = (q_1, 0, q_3)$ cannot be a maximum unless $R_1/R_3 = C_3/C_1$. In this case, (117) gives $q_1 = q_3$. Since $q_1 + q_3 = q$, we obtain that the only 2-maximal good triples where the maximum of $\bar{\Phi}$ may occur are $(q/2, q/2, 0)$ or its permutations, as desired.

This concludes the proof of Lemma 7.14.

7.4 Bad triples: proof of Lemma 7.15

To get a handle on bad triples, we first give necessary conditions so that the maximum in (107) happens at the boundary. The proof of the following lemma is given in Section 7.5.

Lemma 7.19. *Let $0 \leq B < 1$. For a triple $\mathbf{q} = (q_1, q_2, q_3)$, let \mathbf{r}, \mathbf{c} achieve the maximum in (107). Then, if $q_i > 0$, the following implications hold:*

$$R_i = 0 \Rightarrow \sum_j q_j C_j \leq (1 - B)C_i, \quad C_i = 0 \Rightarrow \sum_j q_j R_j \leq (1 - B)R_i.$$

In particular, if $q_i > 1 - B$ it holds that $R_i, C_i > 0$. Hence, for every $q \geq 3$ there exists $i \in \{1, 2, 3\}$ such that $R_i, C_i > 0$.

We next examine bad triples. Note that a bad triple $\mathbf{q} = (q_1, q_2, q_3)$, by the second part of Lemma 7.16, must have at least two positive entries. We consider cases whether the triple \mathbf{q} has two or three positive entries. We start with the case where exactly two of the q_i 's are positive. We assume throughout the rest of the section that \mathbf{r}, \mathbf{c} achieve the maximum in (107).

Let $\mathbf{q} = (q_1, q_2, 0)$ be a bad triple where $q_1, q_2 > 0$. Since \mathbf{q} is bad, at least one of R_1, R_2, C_1, C_2 is zero. Wlog, we may assume $C_2 = 0$. By the second part of Lemma 7.19, it follows that $R_1, C_1 > 0$. There are two cases to consider.

$$(I) \ R_2 = 0, \ (II) \ R_2 > 0. \tag{122}$$

Case (I) is straightforward: by the first part of Lemma 7.16, we trivially have $\frac{\partial \bar{\Phi}^S}{\partial q_1} > 0$ and $\frac{\partial \bar{\Phi}^S}{\partial q_2} = 0$, so that the second part of Lemma 7.16 yields that \mathbf{q} does not maximize $\bar{\Phi}$.

We next examine case (II). Since $\bar{\Phi}^S$ is scale-free (see (109)), we may assume that $C_1 = 1$. Since R_1, R_2 are positive, it holds that $\partial \bar{\Phi}^S / \partial R_1 = \partial \bar{\Phi}^S / \partial R_2 = 0$, yielding

$$R_1 \propto y^d, \ R_2 \propto 1, \text{ where } y = (q_1 + B - 1)/q_1.$$

Expressing q_1, q_2 in terms of y and substituting in $\bar{\Phi}^S$, we obtain the value of $\bar{\Phi}(\mathbf{q})$:

$$\bar{\Phi}(\mathbf{q}) = \log h(y), \text{ where } h(y) := \frac{(1 - B)(q(1 - y) - (1 - B)(1 - y^{d+1}))}{(1 - y)^2}.$$

Let I be the interval $[0, (q + B - 1)/q]$. Note that for any $y \in I$, there exists a positive $q_1 \in [0, q]$ such that $y = (q_1 + B - 1)/q_1$. Obviously, if \mathbf{q} maximizes $\bar{\Phi}$, it must be the case that y maximizes $h(y)$ in the interval I . We compute $h'(y)$.

$$h'(y) = \frac{(1 - B)r(y)}{(1 - y)^3}, \text{ where } r(y) := q(1 - y) - (1 - B)((d - 1)y^{d+1} - (d + 1)y^d + 2).$$

It is immediate to see that $r(y)$ is convex for $y \in [0, 1]$. Since $r(0) = q - 2(1 - B) > 0$ and $r(1) = 0$, we obtain that either

- (i) $r(y) > 0$ for all $y \in I$, or
- (ii) $\exists y_o \in I: r(y_o) = 0, r(y) > 0$ iff $y < y_o$.

In case (i), $h(y)$ is increasing and hence $h(y)$ is maximized at $y = (q + B - 1)/q$. This value of y corresponds to $q_1 = q$ and thus $\bar{\Phi}(\mathbf{q}) = \bar{\Phi}(q, 0, 0)$.

In case (ii), we have $h(y) \leq h(y_o)$. The value of q_1 corresponding to y_o is $q_o := (1 - B)/(1 - y_o)$. We will show that the maximum in (107) does not happen at the boundary $C_2 = 0$ when $\mathbf{q} = (q_o, q - q_o, 0)$, implying that $h(y_o)$ does not equal $\bar{\Phi}(\mathbf{q})$ and hence the maximum of $\bar{\Phi}$ as well. To prove the former, we utilize the first part of Lemma 7.19. In particular, we prove that

$$q_o y_o^d + (q - q_o) > (1 - B). \tag{123}$$

Note that $r(y_o) = 0$ yields $q = (1-B)((d-1)y_o^{d+1} - (d+1)y_o^d + 2)/(1-y_o)$. Plugging this expression into (123), we only need to show that

$$\frac{(d-1)y_o^{d+1} - dy_o^d + 1}{1-y_o} > 1 \text{ or } (d-1)y_o^d + 1 > dy_o^{d-1}, \quad (124)$$

which holds by the AM-GM inequality for any positive $y_o \neq 1$.

Let $\mathbf{q} = (q_1, q_2, q_3)$ be a bad triple where all of the q_i 's are positive. Since \mathbf{q} is bad, at least one of the R_i 's and C_j 's is zero. W.l.o.g. we may assume $C_2 = 0$. Moreover, by the second part of Lemma 7.19, we may also assume that $R_1, C_1 > 0$. There are four cases to consider.

$$(I) R_2 = 0, \quad (II) R_2, R_3 > 0, C_3 = 0, \quad (III) R_2, R_3, C_3 > 0, \quad (IV) R_2, C_3 > 0, R_3 = 0.$$

We omitted the case $R_2 > 0$ and $R_3 = C_3 = 0$, which is identical to case (I) after renaming the q_i 's.

Case (I) is straightforward: since $R_2 = C_2 = 0$, (110) gives $\partial \bar{\Phi}^S / \partial q_2 = 0$. Since at least one of $\partial \bar{\Phi}^S / \partial q_1$, $\partial \bar{\Phi}^S / \partial q_3$ is positive, the second part of Lemma 7.16 yields that \mathbf{q} does not maximize $\bar{\Phi}$.

We next examine case (II). Since $\bar{\Phi}^S$ is scale-free (see (109)), we may substitute $C_1 = 1$. Setting the derivatives of $\partial \bar{\Phi}^S / \partial R_1, \partial \bar{\Phi}^S / \partial R_2, \partial \bar{\Phi}^S / \partial R_3$ equal to zero, we obtain

$$R_1 \propto (q_1 + B - 1)^d / q_1^d, \quad R_2 \propto 1, \quad R_3 \propto 1.$$

It follows that $\bar{\Phi}(\mathbf{q}) = \bar{\Phi}(q_1, q_2 + q_3, 0)$ and hence the maximum of $\bar{\Phi}$ does not occur at \mathbf{q} by the argument for case (II) in (122).

We next examine case (III). The partial derivatives of $\bar{\Phi}^S$ with respect to R_1, R_2, R_3, C_1, C_3 must vanish so we obtain

$$\begin{aligned} R_1^{1/d} &\propto q_1 C_1 + q_2 C_3 - (1-B)C_1, \quad R_2^{1/d} \propto q_1 C_1 + q_3 C_3, \quad R_3^{1/d} \propto q_1 C_1 + q_3 C_3 - (1-B)C_3, \\ C_1^{1/d} &\propto q_1 R_1 + q_2 R_2 + q_3 R_3 - (1-B)R_1, \quad C_3^{1/d} \propto q_1 R_1 + q_2 R_2 + q_3 R_3 - (1-B)R_3. \end{aligned} \quad (125)$$

If $C_1 = C_3$, then $R_1 = R_3$ and thus we obtain $\bar{\Phi}(q_1, q_2, q_3) = \bar{\Phi}(q_1 + q_3, q_2, 0)$, contradicting the maximality of \mathbf{q} by the argument for case (II) in (122). Thus, wlog we may assume $C_1 < C_3$. By (125), this yields

$$R_2 > R_1 > R_3, \quad C_1 < C_3. \quad (126)$$

We have the following analogue of Lemma 7.17, which proves that the maximum cannot occur at \mathbf{q} by the second part in Lemma 7.16.

Lemma 7.20. *Suppose that R_i 's and C_j 's satisfy (125) and (126). If $R_1/R_3 \neq C_3/C_1$ then $\frac{\partial \bar{\Phi}^S}{\partial q_1} - \frac{\partial \bar{\Phi}^S}{\partial q_3} \neq 0$. If $R_1/R_3 = C_3/C_1$ then $\frac{\partial \bar{\Phi}^S}{\partial q_1} - \frac{\partial \bar{\Phi}^S}{\partial q_2} \neq 0$.*

Proof of Lemma 7.20. The proof is analogous to the proof of Lemma 7.17, we highlight the main differences. Let $r_1^d = R_1/R_3, r_2^d = R_2/R_3, c_3^d = C_3/C_1$. The R_i 's and C_j 's may be recovered by the r_i 's and c_j 's by

$$R_1 \propto r_1^d, R_2 \propto r_2^d, R_3 \propto 1, \text{ and } C_1 \propto 1, C_3 \propto c_3^d. \quad (127)$$

By (126), we have

$$r_2 > r_1 > 1 \text{ and } c_3 > 1.$$

The expressions for r_1, r_2, c_3 in (116) are exactly the same after substituting $c_2 = 0$. The same is true for (117), (118). It follows that the proof for the first part of Lemma 7.17 holds verbatim in this case as well (note that the ordering of r_1, r_2 is different here but that part of the argument does not use the ordering).

While the proof for the second part of Lemma 7.17 does not carry through as simply, the changes are minor. We assume that $r_1 = c_3$ and set $DIF_{12} := \frac{\partial \Phi^S}{\partial q_1} - \frac{\partial \Phi^S}{\partial q_2}$. Plugging $r_1 = c_3$ and $c_2 = 0$ in (117), (118) and then substituting the resulting expressions in DIF_{12} we obtain

$$DIF_{12} = \frac{(1-B)h(r_1)}{r_1 - 1}, \text{ where } h(r_1) := (r_1^{2d+1} + dr_1^{d+1} - dr_1^d - 1) - \frac{(r_1^{d+1} - 1)^{d+1}}{(r_1^d - 1)^d}.$$

By a first derivative argument, the function

$$g(r_1) := \log \left(\frac{(r_1^{d+1} - 1)^{d+1}}{(r_1^d - 1)^d (r_1^{2d+1} + dr_1^{d+1} - dr_1^d - 1)} \right),$$

is strictly increasing for $r_1 > 1$. Thus, $g(r_1) \geq g(+\infty) = 0$, which gives $h(r_1) > 0$ for all $r_1 > 1$. This proves that $DIF_{12} \neq 0$, as desired. \square

Finally, we examine case (IV). The partial derivatives of Φ^S with respect to R_1, R_2, C_1, C_3 must vanish so we obtain

$$\begin{aligned} R_1^{1/d} &\propto q_1 C_1 + q_3 C_3 - (1-B)C_1, & R_2^{1/d} &\propto q_1 C_1 + q_3 C_3, \\ C_1^{1/d} &\propto q_1 R_1 + q_2 R_2 - (1-B)R_1, & C_3^{1/d} &\propto q_1 R_1 + q_2 R_2. \end{aligned} \quad (128)$$

Note that we have $R_1 < R_2$ and $C_1 < C_3$.

Lemma 7.21. *If $R_2/R_1 \neq C_3/C_1$ then either $\frac{\partial \Phi^S}{\partial q_2} - \frac{\partial \Phi^S}{\partial q_3} \neq 0$ or $\frac{\partial \Phi^S}{\partial q_1} - \frac{\partial \Phi^S}{\partial q_2} \neq 0$. If $R_2/R_1 = C_3/C_1$ and $\frac{\partial \Phi^S}{\partial q_1} - \frac{\partial \Phi^S}{\partial q_2} = 0$, then the maximum in (107) does not happen at the boundary $C_2 = 0$.*

Proof of Lemma 7.21. The approach for the first part is similar the proof of Lemma 7.17. Set $r_2^d = R_2/R_1$ and $c_3^d = C_3/C_1$, so that $r_1, c_3 > 1$. Dividing appropriate pairs in (128), we obtain

$$r_2 = \frac{q_1 + q_3 c_3^d}{(q_1 + B - 1) + q_3 c_3^d}, \quad c_3 = \frac{q_1 + q_2 r_2^d}{(q_1 + B - 1) + q_2 r_2^d}. \quad (129)$$

It follows that

$$q_2 = \frac{q_1 - (q_1 + B - 1)c_3}{r_2^d(c_3 - 1)}, \quad q_3 = \frac{q_1 - (q_1 + B - 1)r_2}{c_3^d(r_2 - 1)}.$$

Using these, we obtain

$$\frac{\partial \Phi^S}{\partial q_2} - \frac{\partial \Phi^S}{\partial q_3} = 0 \Rightarrow f(r_2) = f(c_3), \text{ where } f(x) := \frac{x^{d+1}}{x-1}, \quad (130)$$

$$\frac{\partial \Phi^S}{\partial q_1} - \frac{\partial \Phi^S}{\partial q_3} = 0 \Rightarrow r_2^{d+1}(c_3 - 1) - (d+1)r_2 c_3 + d(r_2 + c_3) - (d-1) = 0. \quad (131)$$

From (131), we obtain

$$c_3 = g(r_2), \text{ where } g(r_2) := \frac{r_2^{d+1} - dr_2 + (d-1)}{r_2^{d+1} - (d+1)r_2 + d}. \quad (132)$$

It follows that $r_2 = c_3$ is equivalent to

$$r_2^{d+1} = (d+1)r_2 - (d-1). \quad (133)$$

It is straightforward to check that (133) has exactly one solution for $r_2 > 1$, say $r_2 = x$. Using the expression for c_3 from (132), (130) gives

$$h(r_2) = 0, \text{ where } h(r_2) := r_2^{d+1} - \frac{\left(r_2^{d+1} - dr_2 + (d-1)\right)^{d+1}}{\left(r_2^{d+1} - (d+1)r_2 + d\right)^d}.$$

A standard calculation (albeit lengthy) shows that h is strictly increasing for $r_2 > 1$ (and every $d \geq 2$). Moreover, it holds that $h(x) = 0$, so that (130) and (131) can only hold simultaneously when $r_2 = c_3$, which yields the first part of the lemma.

For the second part, we have $r_2 = c_3$ so r_2 satisfies (133). To prove that the maximum does not happen at the boundary $C_2 = 0$, we use the first part of Lemma 7.19. It suffices to prove that

$$q_1 + q_2 r_2^d > (1-B)r_2^d. \quad (134)$$

We have that $q_2 = q_3 = (q - q_1)/2$, so that (129) gives

$$q_1 = \frac{q(r_2^{d+1} - r_2^d) - 2r_2(1-B)}{(r_2 - 1)(r_2^d - 2)}, \quad q_2 = q_3 = \frac{q - r_2(q + B - 1)}{(r_2 - 1)(r_2^d - 2)}. \quad (135)$$

Plugging (135) into (134) gives the equivalent inequality

$$\frac{(1-B)(r_2^d + r_2 - r_2^{d+1})}{r_2 - 1} > 0$$

To see the latter, use (133) to obtain

$$r_2^d + r_2 - r_2^{d+1} = r_2^d + (d-1) - dr_2 > 0, \text{ for all } r_2 > 1 \text{ by the AM-GM inequality.}$$

This completes the proof. \square

7.5 Remaining proofs

Proof of Lemma 7.6. W.l.o.g. we may assume that the scaling factors in (105) are equal to 1. Let $R_i = r_i^d$, $C_i = c_i^d$, $r = \sum_{i=1}^q r_i^d$, and $c = \sum_{i=1}^q c_i^d$. We have

$$r_i = c - (1-B)c_i^d \quad \text{and} \quad c_i = r - (1-B)r_i^d,$$

It is clear from this equation that $R_i = R_j$ iff $C_i = C_j$ and hence also $t_R = t_C$. We also obtain that for $i = 1, \dots, q$,

$$r_i = c - (1-B)(r - (1-B)r_i^d)^d. \quad (136)$$

Since r is the sum of r_i^d and the r_i are positive, we have $(1-B)r_i^d < r$. Fix the values of r, c and let I be the interval where $(1-B)x^d < r$. Using (136), we shall prove that $t_R \leq 3$ by arguing that $f(x) = c - (1-B)(r - (1-B)x^d)^d - x$ has at most 3 positive roots in the interval I , counted by multiplicities. We have

$$f'(x) = (1-B)^2 d^2 (r - (1-B)x^d)^{d-1} x^{d-1} - 1 = \left(\sum_{i=0}^{d-2} g(x)^i \right) (g(x) - 1),$$

where

$$g(x) = ((1-B)d)^{2/(d-1)} (r - (1-B)x^d)x.$$

Note that $g(x) > 0$ in the interval I and hence all roots of $f'(x)$ in this interval come from $g(x) - 1$. The polynomial $g(x) - 1$ has at most two positive roots by Descartes' rule of signs, hence $f'(x)$ has at most two positive roots in I . Thus, $f(x)$ has at most three positive roots in I , all roots counted by their multiplicities. This concludes the proof. \square

Proof of Lemma 7.9. Let $q' = q/2$. To better align with the results of Section 7.3, let us assume that the fixpoint $(q', 0, q')$ maximizes Ψ_1 . In Section 7.3, we proved that this can be the case only if $R_1/R_3 = C_3/C_1$ or (in the parameterization of Section 7.3) $r_1 = c_3 =: x$ where $x > 1$. Equation (116) for $q_2 = 0$, $q_1 = q_3 = q'$ gives that x satisfies

$$x = \frac{B + q' - 1 + q'x^d}{q' + (B + q' - 1)x^d}. \quad (137)$$

It is straightforward to check that (137) has exactly one solution $x > 1$ for all $0 \leq B < \frac{\Delta - q}{\Delta}$. The values of R_1, C_1, R_3, C_3 may be recovered by (114), which in the case $q_2 = 0$ give

$$R_1 \propto x^d, R_3 \propto 1 \text{ and } C_1 \propto 1, C_3 \propto x^d.$$

This proves the second part of the lemma. For the first part, to check Jacobian stability, we proceed as in the proof of Lemma 7.4. The eigenvalues of the matrix \mathbf{L} in this case can be computed easily as well. They are given by ± 1 by multiplicity 1, $\pm \lambda_1$ by multiplicity $q - 2$ and $\pm(B + q - 1)\lambda_1^2$ by multiplicity 1, where

$$\lambda_1 := \frac{(1 - B)x^{d/2}}{\sqrt{(q' + (B + q' - 1)x^d)(B + q' - 1 + q'x^d)}}.$$

To prove that the absolute value of the eigenvalues different from 1 is less than $1/d$, it suffices to prove that $\lambda_1 < 1/d$. Use (137) to solve for q' and plug the value into the expression for λ_1 . This yields that λ_1 is equal to $x^{(d-1)/2}(x-1)/(x^d-1)$, which by the AM-GM inequality is less than $1/d$ for $x > 1$. \square

Proof of Lemma 7.10. For non-negative $\mathbf{q} = (q_1, q_2, q_3)$ with $q_1 + q_2 + q_3 = q$, consider the function

$$\bar{F}(\mathbf{q}) = \max_{\mathbf{r}, \mathbf{c}} F(\mathbf{q}, \mathbf{r}, \mathbf{c}), \text{ where } F(\mathbf{q}, \mathbf{r}, \mathbf{c}) := \sum_{i=1}^3 q_i R_i \sum_{j=1}^3 q_j C_j + (B-1) \sum_{i=1}^3 q_i R_i C_i, \quad (138)$$

and the maximum is over the compact region (by restricting to $R_i = C_i = 0$ whenever $q_i = 0$)

$$\begin{aligned} \sum_{i=1}^3 q_i R_i^{(d+1)/d} &\leq 1, \quad \sum_{j=1}^3 q_j C_j^{(d+1)/d} \leq 1, \\ R_1, R_2, R_3, C_1, C_2, C_3 &\geq 0. \end{aligned} \quad (139)$$

Note that $\bar{F}(\mathbf{q}) > 0$, since we can set all of the R_i 's and C_j 's equal to x , where $qx^{(d+1)/d} = 1$. Clearly, $\bar{\Phi}(\mathbf{q}) \geq \ln \bar{F}(\mathbf{q})$. Since $\bar{\Phi}^S(\mathbf{q}, \mathbf{r}, \mathbf{c})$ is scale-free with respect to \mathbf{r} and \mathbf{c} (see (109)), we may scale \mathbf{r}, \mathbf{c} to satisfy (139) and hence $\bar{\Phi}(\mathbf{q}) = \sup_{\mathbf{r}, \mathbf{c}} \bar{\Phi}^S(\mathbf{q}, \mathbf{r}, \mathbf{c}) \leq \ln \bar{F}(\mathbf{q})$, proving that $\bar{\Phi}(\mathbf{q}) = \ln \bar{F}(\mathbf{q})$ and consequently the supremum is attained.

To prove that $\sup_{\mathbf{q}} \bar{\Phi}(\mathbf{q})$ is attained, it clearly suffices to prove that $L := \sup_{\mathbf{q}} \bar{F}(\mathbf{q})$ is attained. This can be accomplished by using variants of Berge's Maximum Theorem and showing that the function $\bar{F}(\mathbf{q})$ is upper semi-continuous. We give a more direct argument, which is similar to the proof of Berge's Maximum Theorem and can also easily be adapted to show that $\bar{F}(\mathbf{q})$ is upper semi-continuous.

Note first that $L < \infty$ by a simple application of Hölder's inequality. Let \mathbf{q}_n , $n = 1, 2, \dots$ be a sequence such that $\bar{F}(\mathbf{q}_n) \uparrow L$. Since the \mathbf{q}_n lie in a compact region, by restricting to a subsequence we may assume that $\mathbf{q}_n \rightarrow \mathbf{q}$. Let $\mathbf{r}_n, \mathbf{c}_n$ be maximizers for $\bar{F}(\mathbf{q}_n)$ in (138).

Suppose first that \mathbf{q} has positive entries. Then, for sufficiently large n , the maximizers $\mathbf{r}_n, \mathbf{c}_n$ lie in a compact set and hence a standard diagonalisation argument yields a convergent subsequence $(\mathbf{q}_{n_k}, \mathbf{r}_{n_k}, \mathbf{c}_{n_k}) \rightarrow (\mathbf{q}, \mathbf{r}, \mathbf{c})$. By continuity, \mathbf{r}, \mathbf{c} must lie in the region (139) defined by \mathbf{q} and moreover $\bar{F}(\mathbf{q}_{n_k}) = F(\mathbf{q}_{n_k}, \mathbf{r}_{n_k}, \mathbf{c}_{n_k}) \rightarrow F(\mathbf{q}, \mathbf{r}, \mathbf{c})$. Thus $L = F(\mathbf{q}, \mathbf{r}, \mathbf{c})$ and the supremum is attained.

Suppose now that \mathbf{q} has an entry equal to zero, say q_1 , so that $q_{1n} \rightarrow 0$ (with the natural notation for entries of the subsequences). In this setting, R_{1n}, C_{1n} might escape to infinity, so assume that $R_{1n}, C_{1n} \uparrow \infty$, by restricting to a subsequence if necessary. (139) implies $q_{1n} R_{1n}^{(d+1)/d}, q_{1n} C_{1n}^{(d+1)/d} \leq 1$ and hence $q_{1n} R_{1n}, q_{1n} C_{1n} \rightarrow 0$. Note that $q_{1n} R_{1n} C_{1n} \rightarrow 0$ as well; otherwise there exists a subsequence with $q_{1n_k} R_{1n_k} C_{1n_k} \geq \varepsilon > 0$. This contradicts that $\mathbf{r}_{n_k}, \mathbf{c}_{n_k}$ maximize $F(\mathbf{q}_{n_k}, \cdot, \cdot)$, since setting $R_{1, n_k} = C_{1, n_k} = 0$ would maintain feasibility in (139) and achieve a bigger value of F for all sufficiently large k (recall that $B < 1$). Thus $q_{1n} R_{1n}, q_{1n} C_{1n}, q_{1n} R_{1n} C_{1n} \rightarrow 0$, yielding once again $L = F(\mathbf{q}, \mathbf{r}, \mathbf{c})$. \square

Proof of Lemmas 7.13 and 7.16. We first prove Lemma 7.16. Let $I_R = \{i \in I \mid R_i > 0\}$. For $i \in I_R$, it must hold that $\partial \bar{\Phi}^S / \partial R_i = 0$. Since $q_i > 0$ for $i \in I$, it follows that

$$R_i^{1/d} \propto \sum_j q_j C_j - (1 - B) C_i \quad \text{for all } i \in I_R, \quad (140)$$

and hence

$$R_i^{(d+1)/d} \propto R_i (\sum_j q_j C_j - (1 - B) C_i) \quad \text{for all } i \in I.$$

Thus, for $i \in I$ it holds that

$$\frac{R_i^{(d+1)/d}}{\sum_j q_j R_j^{(d+1)/d}} = \frac{R_i (\sum_j q_j C_j - (1 - B) C_i)}{\sum_j q_j R_j \sum_j q_j C_j + (B - 1) \sum_j q_j R_j C_j}, \quad (141)$$

and an analogous argument for the C_i 's gives

$$\frac{C_i^{(d+1)/d}}{\sum_j q_j C_j^{(d+1)/d}} = \frac{C_i (\sum_j q_j R_j - (1 - B) R_i)}{\sum_j q_j R_j \sum_j q_j C_j + (B - 1) \sum_j q_j R_j C_j}. \quad (142)$$

Moreover, by a direct calculation we have

$$\frac{\partial \bar{\Phi}^S}{\partial q_i} = \frac{(d+1)(R_i \sum_j q_j C_j + C_i \sum_j q_j R_j + (B-1) R_i C_i)}{\sum_j q_j R_j \sum_j q_j C_j + (B-1) \sum_j q_j R_j C_j} - \frac{d R_i^{(d+1)/d}}{\sum_j q_j R_j^{(d+1)/d}} - \frac{d C_i^{(d+1)/d}}{\sum_j q_j C_j^{(d+1)/d}}. \quad (143)$$

Plugging (141), (142) in (143) proves the first part of Lemma 7.16.

For the second part of Lemma 7.16, assume w.l.o.g. that $q_1, q_2 > 0$ and $\frac{\partial \bar{\Phi}^S}{\partial q_1} - \frac{\partial \bar{\Phi}^S}{\partial q_2} > 0$. For $\varepsilon > 0$, consider $\mathbf{q}' = (q_1 + \varepsilon, q_2 - \varepsilon, q_3)$. Since q_1, q_2 are positive, for small enough ε , \mathbf{q}' has positive entries which sum to q . Moreover, for small enough ε the value of $\bar{\Phi}^S$ increases, while still maintaining feasibility in the region (108). Hence, \mathbf{q} does not maximize $\bar{\Phi}$, as desired.

Lemma 7.13 follows easily: just use (140) and the fact that q_1, q_2, q_3 are integers to get the alignment with (105). \square

Proof of Lemma 7.19. Suppose that $q_i > 0$ and $\sum_j q_j C_j > (1 - B)C_i$. We look at the derivative $\partial \overline{\Phi^S} / \partial R_i$ evaluated at $R_i = 0$:

$$\frac{\partial \overline{\Phi^S}}{\partial R_i} = \frac{q_i(q_1 C_1 + q_2 C_2 + q_3 C_3 - (1 - B)C_i)}{\sum_j q_j R_j \sum_j q_j C_j + (B - 1) \sum_j q_j R_j C_j} > 0$$

Thus, increasing the value of R_i by a sufficiently small amount, increases the value of $\overline{\Phi^S}$. Hence, the maximum cannot be obtained at the boundary $R_i = 0$. The second part of the lemma follows immediately from the first part. \square

References

- [AK97] Paola Alimonti and Viggo Kann. Hardness of approximating problems on cubic graphs. In Gian Carlo Bongiovanni, Daniel P. Bovet, and Giuseppe Di Battista, editors, *Algorithms and Complexity, Third Italian Conference, CIAC '97, Rome, Italy, March 12-14, 1997, Proceedings*, volume 1203 of *Lecture Notes in Computer Science*, pages 288–298. Springer, 1997.
- [Ben77] G. Bennett. Schur multipliers. *Duke Math. J.*, 44(3):603–639, 1977.
- [BW02] G. R. Brightwell and P. Winkler. Random colorings of a Cayley tree. In *Contemporary combinatorics*, volume 10 of *Bolyai Soc. Math. Stud.*, pages 247–276. János Bolyai Math. Soc., Budapest, 2002.
- [dB81] N. G. de Bruijn. *Asymptotic methods in analysis*. Dover Publications, Inc., New York, third edition, 1981.
- [GGS⁺14] Andreas Galanis, Qi Ge, Daniel Štefankovič, Eric Vigoda, and Linji Yang. Improved inapproximability results for counting independent sets in the hard-core model. *Random Structures Algorithms*, 45(1):78–110, 2014.
- [GSV12] Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. Inapproximability of the partition function for the antiferromagnetic Ising and hard-core models. *CoRR*, abs/1203.2226, 2012.
- [Geo11] Hans-Otto Georgii. *Gibbs measures and phase transitions*, volume 9 of *de Gruyter Studies in Mathematics*. Walter de Gruyter & Co., Berlin, second edition, 2011.
- [GJP03] Leslie Ann Goldberg, Mark Jerrum, and Mike Paterson. The computational complexity of two-state spin systems. *Random Structures Algorithms*, 23(2):133–154, 2003.
- [Gre00] Catherine Greenhill. The complexity of counting colourings and independent sets in sparse graphs and hypergraphs. *Comput. Complexity*, 9(1):52–72, 2000.
- [HJ13] Roger A. Horn and Charles R. Johnson. *Matrix analysis*. Cambridge University Press, Cambridge, second edition, 2013.
- [Jan95] Svante Janson. Random regular graphs: asymptotic distributions and contiguity. *Combin. Probab. Comput.*, 4(4):369–405, 1995.

- [JLR00] Svante Janson, Tomasz Łuczak, and Andrzej Ruciński. *Random graphs*. Wiley-Interscience Series in Discrete Mathematics and Optimization. Wiley-Interscience, New York, 2000.
- [Joh96] A. Johansson. Asymptotic choice number for triangle free graphs. Technical Report 91-5, Center for Discrete Mathematics and Theoretical Computer Science (DIMACS), Rutgers University, NJ, USA, 1996.
- [Jon02] Johan Jonasson. Uniqueness of uniform random colorings of regular trees. *Statist. Probab. Lett.*, 57(3):243–248, 2002.
- [Kel91] F. P. Kelly. Loss networks. *Ann. Appl. Probab.*, 1(3):319–378, 1991.
- [LLY13] Liang Li, Pinyan Lu, and Yitong Yin. Correlation decay up to uniqueness in spin systems. In Sanjeev Khanna, editor, *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pages 67–84. SIAM, 2013.
- [LY08] David G. Luenberger and Yinyu Ye. *Linear and nonlinear programming*. International Series in Operations Research & Management Science, 116. Springer, New York, third edition, 2008.
- [MR02] Michael Molloy and Bruce Reed. *Graph colouring and the probabilistic method*, volume 23 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 2002.
- [MR01] Michael Molloy and Bruce A. Reed. Colouring graphs when the number of colours is nearly the maximum degree. In Jeffrey Scott Vitter, Paul G. Spirakis, and Mihalis Yannakakis, editors, *Proceedings on 33rd Annual ACM Symposium on Theory of Computing, July 6-8, 2001, Heraklion, Crete, Greece*, pages 462–470. ACM, 2001.
- [MT10] Robin A. Moser and Gábor Tardos. A constructive proof of the general Lovász local lemma. *J. ACM*, 57(2):Art. 11, 15, 2010.
- [MWW09] Elchanan Mossel, Dror Weitz, and Nicholas Wormald. On the hardness of sampling independent sets beyond the tree threshold. *Probab. Theory Related Fields*, 143(3-4):401–439, 2009.
- [RW94] R. W. Robinson and N. C. Wormald. Almost all regular graphs are Hamiltonian. *Random Structures Algorithms*, 5(2):363–374, 1994.
- [SST12] Alistair Sinclair, Piyush Srivastava, and Marc Thurley. Approximation algorithms for two-state anti-ferromagnetic spin systems on bounded degree graphs. In Yuval Rabani, editor, *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17-19, 2012*, pages 941–953. SIAM, 2012.
- [Sly10] Allan Sly. Computational transition at the uniqueness threshold. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 287–296. IEEE Computer Society, 2010.
- [SS12] Allan Sly and Nike Sun. The computational hardness of counting in two-spin models on d-regular graphs. In *53rd Annual IEEE Symposium on Foundations of Computer*

Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012, pages 361–369. IEEE Computer Society, 2012.

- [Wei06] Dror Weitz. Counting independent sets up to the tree threshold. In Jon M. Kleinberg, editor, *Proceedings of the 38th Annual ACM Symposium on Theory of Computing, Seattle, WA, USA, May 21-23, 2006*, pages 140–149. ACM, 2006.

A The Small Subgraph Conditioning Method

In this section, we prove Lemma 6.12 by applying the small subgraph conditioning method.

A.1 Overview

The small subgraph conditioning method was introduced by [RW94] to prove that a random Δ -regular contains asymptotically almost surely (a.a.s.) a Hamilton cycle. Roughly speaking, the method provides a way to get a.a.s results when the second moment method fails, in the particular case (though common in the random regular setting) where the ratio of the second moment of a variable to the first moment squared converges to a constant strictly greater than 1.

The method was first used to analyze spin models on random regular graphs in [MWW09] and was subsequently used in [Sly10, GSV12]. In our setting, applying the small subgraph conditioning method of [RW94] as in the previous works [MWW09, Sly10, GSV12] would not be sufficient, since it guarantees a polynomial multiplicative deviation from the expectation, which is weak in the setting of Lemma 6.9. We instead use an extension of the method given by [Jan95].

More generally, the method of [RW94] is sufficient when the interest is in proving concentration of a variable within a polynomial factor from its expectation. Janson’s refinement of the method gives the distributional limit of the variable and explicitly attributes the fluctuations from the expectation to the presence of specific subgraph structures. For the convenience of the reader, we include both versions of the method in Theorem A.1, which is a concatenated version of the respective Theorems in [RW94, Jan95]. The theorem can be extrapolated from [Jan95], after combining [Jan95, Lemma 1, Remark 4, Remark 9]. The notation $[X]_m$ refers to the m -th order falling factorial of the variable X . We shall discuss the theorem statement afterwards.

Theorem A.1. *Let S be a set of finite cardinality. For $s \in S$ and $i = 1, 2, \dots$, let $\mu_i > 0$ and $\delta_i^{(s)} > -1$ be constants and assume that for each n there are random variables X_{in} , $i = 1, 2, \dots$, and $Y_n^{(s)}$, $s \in S$, all defined on the same probability space $\mathcal{G} = \mathcal{G}_n$ such that X_{in} is non-negative integer valued, $Y_n^{(s)} \geq 0$ and $\mathbf{E}[Y_n^{(s)}] > 0$ (for n sufficiently large). Furthermore, for every $s \in S$, the following hold:*

(A1) $X_{in} \xrightarrow{d} Z_i$ as $n \rightarrow \infty$, jointly for all i , where $Z_i \sim \text{Po}(\mu_i)$ are independent Poisson random variables;

(A2) for every finite sequence j_1, \dots, j_m of non-negative integers,

$$\frac{\mathbf{E}_{\mathcal{G}}[Y_n^{(s)}[X_{1n}]_{j_1} \cdots [X_{mn}]_{j_m}]}{\mathbf{E}_{\mathcal{G}}[Y_n^{(s)}]} \rightarrow \prod_{i=1}^m \left(\mu_i (1 + \delta_i^{(s)}) \right)^{j_i} \quad \text{as } n \rightarrow \infty; \quad (144)$$

(A3) $\sum_i \mu_i (\delta_i^{(s)})^2 < \infty$;

$$(A4) \quad \mathbf{E}_{\mathcal{G}}[(Y_n^{(s)})^2]/(\mathbf{E}_{\mathcal{G}}[Y_n^{(s)}])^2 \leq \exp\left(\sum_i \mu_i (\delta_i^{(s)})^2\right) + o(1) \text{ as } n \rightarrow \infty;$$

Then, the following conclusions hold:

(C1) Let $r(n)$ be a function such that $r(n) \rightarrow 0$ as $n \rightarrow \infty$. For each $s \in S$, it holds that $Y_n^{(s)} > r(n)\mathbf{E}_{\mathcal{G}}[Y_n^{(s)}]$ asymptotically almost surely.

(C2) For $s \in S$,

$$\frac{Y_n^{(s)}}{\mathbf{E}_{\mathcal{G}}[Y_n^{(s)}]} \xrightarrow{d} W^{(s)} = \prod_i \left[1 + \delta_i^{(s)}\right]^{Z_i} \exp(-\mu_i \delta_i^{(s)}). \quad (145)$$

This and the convergence in (A1) hold jointly. The infinite product defining $W^{(s)}$ converges a.s. and in L^2 , with

$$\mathbf{E}[W^{(s)}] = 1 \text{ and } \mathbf{E}[(W^{(s)})^2] = \lim_{n \rightarrow \infty} \mathbf{E}_{\mathcal{G}}[(Y_n^{(s)})^2]/(\mathbf{E}_{\mathcal{G}}[Y_n^{(s)}])^2.$$

Moreover, $W^{(s)} > 0$ a.s. iff $\delta_i^{(s)} > -1$ for all i ;

The random variables $Y_n^{(s)}$ in Theorem A.1 are the ones we are interested in obtaining “concentration” type results, where s is simply an index allowing us to treat simultaneously more than one variables. In our setting, for $G \sim \mathcal{G}_n$, $Y_n^{(s)}$ are going to be the variables $Z_G^{\mathbf{p}}(\eta)$ for phases $\mathbf{p} \in \mathcal{Q}$ and configurations η on W . The random variables X_{in} , for graphs with no small multicyclic components, correspond to cycles of length i . For example, in our setting and because the graph G is bipartite, X_{in} is the number of cycles of length i in G where i is even.

The conclusion (C1) of Theorem A.1 is essentially due to [RW94], while the conclusion (C2) is an extension of conclusion (C1) due to [Jan95]. At this point, to obtain Lemma 6.12 (which was the important part to prove Lemma 6.9) we will not explicitly use either of (C1) or (C2) but rather the following variant. The variant was observed in [Jan95, p.5], who discusses it without proof in a specific setting, and is also implicit in [RW94]. As such, we write and prove a formal statement in the setup of Theorem A.1. The proof follows Janson’s proof of Theorem A.1 but uses a different finish.

Lemma A.2. Assume that the conditions in Theorem A.1 hold. For an integer $m > 0$ and $s \in S$, let

$$W_{mn}^{(s)} = \prod_{i=1}^m (1 + \delta_i^{(s)})^{X_{in}} \exp(-\mu_i \delta_i^{(s)}).$$

Then, for every $\varepsilon > 0$, it holds that

$$\lim_{m \rightarrow \infty} \limsup_{n \rightarrow \infty} \Pr_{\mathcal{G}_n} \left(\bigcup_{s \in S} \left[\left| \frac{Y_n^{(s)}}{\mathbf{E}_{\mathcal{G}_n}[Y_n^{(s)}]} - W_{mn}^{(s)} \right| > \varepsilon \right] \right) = 0. \quad (146)$$

Proof of Lemma A.2. We prove the statement for a fixed $s \in S$, the extension of the argument to prove (146) is straightforward (e.g. by a union bound) and is omitted. To lighten notation we will drop s from the notation and w.l.o.g. we also assume $\mathbf{E}_{\mathcal{G}_n}[Y_n] = 1$. We will prove that

$$\limsup_{n \rightarrow \infty} \Pr_{\mathcal{G}_n} (|Y_n - W_{mn}| > \varepsilon) \leq \frac{1}{4} \varepsilon^{-2} \left[\exp\left(\sum_{i=1}^{\infty} \mu_i \delta_i^2\right) - \exp\left(\sum_{i=1}^m \mu_i \delta_i^2\right) \right]. \quad (147)$$

This clearly gives the statement of the lemma, since by assumption (A3) of Theorem A.1, the lhs is finite and goes to 0 as $m \rightarrow \infty$. To prove (147), we follow [Jan95, Proof of Theorem 1] up to a certain point but avoid the use of Skorokhod's theorem in the argument. Janson's proof goes as follows. For a positive integer m define the functions

$$\begin{aligned} f_n(x_1, \dots, x_m) &= \mathbf{E}_{\mathcal{G}_n}[Y_n \mid X_{1n} = x_1, \dots, X_{mn} = x_m], \\ f_\infty(x_1, \dots, x_m) &= \lim_{n \rightarrow \infty} f_n(x_1, \dots, x_m) = \prod_{i=1}^m (1 + \delta_i)^{x_i} e^{-\mu_i \delta_i}. \end{aligned} \quad (148)$$

The second equality follows by assumption (A2) of Theorem A.1 and [Jan95, Lemma 1]. Define also the random variable

$$Y_n^{(m)} = \mathbf{E}_{\mathcal{G}_n}[Y_n \mid X_{1n}, \dots, X_{mn}].$$

Using assumptions (A1) and (A2), Fatou's Lemma and that $Y_n^{(m)}$ is a conditional expectation of Y_n , one obtains

$$\limsup_{n \rightarrow \infty} \mathbf{E}_{\mathcal{G}_n}[|Y_n - Y_n^{(m)}|^2] \leq \exp\left(\sum_{i=1}^{\infty} \mu_i \delta_i^2\right) - \exp\left(\sum_{i=1}^m \mu_i \delta_i^2\right),$$

see [Jan95, Equation (5.2)] for details. We now give the main deviation point from Janson's proof, which amounts to proving that for fixed m , we have

$$\lim_{n \rightarrow \infty} \Pr_{\mathcal{G}_n}(|Y_n^{(m)} - W_{mn}| > \varepsilon) = 0. \quad (149)$$

as $n \rightarrow \infty$. Fix $M > 0$. By (148), there is N such that for $n \geq N$ it holds that

$$|f_n(x_1, \dots, x_m) - f_\infty(x_1, \dots, x_m)| < \varepsilon \text{ for all integer } x_1, \dots, x_m \in [0, M].$$

It follows that for $n \geq N$, we have

$$\Pr_{\mathcal{G}_n}(|Y_n^{(m)} - W_{mn}| > \varepsilon) \leq \Pr_{\mathcal{G}_n}\left(\bigcup_{i=1}^m [X_{in} > M]\right)$$

Note that as $n \rightarrow \infty$, the rhs by assumption (A1) converges to $\Pr(\bigcup_{i=1}^m [Z_i > M])$. The latter can be made arbitrarily small by letting $M \rightarrow \infty$. This proves (149).

The final step is to bound

$$\begin{aligned} \limsup_{n \rightarrow \infty} \Pr_{\mathcal{G}_n}(|Y_n - W_{mn}| > \varepsilon) &\leq \limsup_{n \rightarrow \infty} \Pr_{\mathcal{G}_n}(|Y_n - Y_n^{(m)}| > \varepsilon/2) + \limsup_{n \rightarrow \infty} \Pr_{\mathcal{G}_n}(|Y_n^{(m)} - W_{mn}| > \varepsilon/2) \\ &\leq \frac{1}{4} \varepsilon^{-2} \left[\exp\left(\sum_{i=1}^{\infty} \mu_i \delta_i^2\right) - \exp\left(\sum_{i=1}^m \mu_i \delta_i^2\right) \right] + 0, \end{aligned}$$

which finishes the proof of (147). □

A.2 Application of the Small Subgraph Conditioning Method

The application of Theorem A.1, and similarly Lemma A.2, requires a verification of its assumptions. This check is routine for the most part, but it is nevertheless technically arduous, mainly because of assumption (A3), which requires precise calculation of the moments' asymptotics. We suppress the verification in the following lemma whose proof is given later in this section. The lemma includes some details on a few quantities which will be relevant for the proof of Lemma 6.12.

Lemma A.3. *Let $G \sim \mathcal{G}_n^r$ and X_{in} be the number of cycles of even length i appearing in G , $i = 2, 4, \dots$. Let $S = \{(\mathbf{p}, \eta) \mid \mathbf{p} \in \mathcal{Q}, \eta : W \rightarrow [q]\}$ and for $s \in S$ with $s = (\mathbf{p}, \eta)$, set $Y_n^{(s)} = Z_G^{\mathbf{p}}(\eta)$. In the setting of Theorem 1.5, the assumptions of Theorem A.1 hold.*

Further, for $s \in S$ with $s = (\mathbf{p}, \eta)$ and all even $i \geq 2$, $\delta_i^{(s)}$ satisfies (i) $\delta_i^{(s)} > 0$, (ii) $\delta_i^{(s)}$ depends on \mathbf{p} but not on η , (iii) $\sum_i \mu_i \delta_i^{(s)} < \infty$, (iv) if the phases are permutation symmetric, $\delta_i^{(s)}$ depends on the spin model but not on the particular phase \mathbf{p} .

Using Lemmas A.2 and A.3, we are ready to prove Lemma 6.12.

Proof of Lemma 6.12. To see (102), note that the $W_{mn}^{(s)}$ of Lemma A.2 depend on the particular s only through the $\delta_i^{(s)}$'s. By Item (ii) of Lemma A.3, these depend only on \mathbf{p} in general and specifically for the permutation symmetric case, only on the spin model by Item (iv).

It remains to prove that $W_{mn}^{\mathbf{p}}$ are lower bounded uniformly in \mathbf{p} by a positive constant. Since the number of phases \mathbf{p} is bounded by a constant depending only on the spin model, it suffices to show that this is the case for a fixed phase \mathbf{p} . Using Item (i) of Lemma A.3 and that the random variables X_{in} are non-negative integer valued, we have everywhere the bound

$$W_{mn}^{\mathbf{p}} = \prod_{i=1}^m (1 + \delta_i^{\mathbf{p}})^{X_{in}} \exp(-\mu_i \delta_i^{\mathbf{p}}) \geq \prod_{i=1}^m \exp(-\mu_i \delta_i^{\mathbf{p}}) > \prod_{i=1}^{\infty} \exp(-\mu_i \delta_i^{\mathbf{p}}).$$

Note that we have identified the $\delta_i^{(s)}$'s with the respective $\delta_i^{\mathbf{p}}$'s, this is justified by Item (ii) of Lemma A.3. The last quantity is finite and positive by Item (iii) in Lemma A.3. \square

We next prove Lemma A.3 which amounts to checking the validity of the assumptions (A1)-(A4) of Theorem A.1 for $Z_G^{\mathbf{p}}(\eta)$ for $\mathbf{p} \in \mathcal{Q}$ and $\eta : W \rightarrow [q]$.

Let us fix first some notation. Recall that a phase $\mathbf{p} \in \mathcal{Q}$ corresponds to a global maximum (α, β) of Ψ_1 . Let $\mathbf{x} = (x_{ij})_{i,j \in [q]}$ be as in Lemma 4.3, i.e., the unique vector which maximizes $\Upsilon_1(\alpha, \beta, \mathbf{x})$ when α, β are fixed. In the setting of Theorem 1.5, we may assume that (α, β) is a Hessian local maximum of Ψ_1 . The following lemma puts together some relevant quantities and information which we derived in Section 4.2 in the course of proving Theorem 4.2.

Lemma A.4. *For a random Δ -regular graph, suppose that (α, β) is a Hessian local maximum of Ψ_1 . Define the vector $\mathbf{x} = (x_{ij})_{i,j \in [q]}$ as in Lemma 4.3.*

Let \mathbf{J} be the matrix $\begin{bmatrix} \mathbf{0} & \mathbf{L} \\ \mathbf{L}^T & \mathbf{0} \end{bmatrix}$, where \mathbf{L} is the $q \times q$ matrix whose ij -entry is given by $x_{ij}/\sqrt{\alpha_i}\sqrt{\beta_j}$. Then, the spectrum of \mathbf{J} is

$$\pm 1, \pm \lambda_1, \dots, \pm \lambda_{q-1},$$

for some positive λ_i which satisfy $\max_i \lambda_i < \frac{1}{\Delta-1}$. Relevant to Lemma A.3, observe that if the phases \mathbf{p} are permutation symmetric, then the λ_i 's are common for all phases.

Let $G \sim \mathcal{G}_n^r$ and $X_i := X_{in}$ be the number of cycles in G of even length i . Let $\mathbf{p} = (\alpha, \beta) \in \mathcal{Q}$. We next verify the assumptions of Theorem A.1 for the random variables $Z_G^{\mathbf{p}}(\eta)$, $\eta : W \rightarrow [q]$. We have the following lemmas.

Lemma A.5 (Lemma 7.3 in [MWW09]). *Assumption (A1) of Theorem A.1 holds for even i with*

$$\mu_i = \frac{r(\Delta, i)}{i} = \frac{(\Delta - 1)^i + (-1)^i(\Delta - 1)}{i},$$

where $r(\Delta, i)$ is the number of ways to properly edge color a cycle of length i with Δ colors.

The proof of Lemma A.5 is given in [MWW09] and is omitted.

Lemma A.6. *Let λ_j , $j \in [q - 1]$ be as in Lemma A.4. Then, for all even $i \geq 2$ it holds that*

$$\frac{\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{P}}(\eta)X_i]}{\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{P}}(\eta)]} \rightarrow \mu_i(1 + \delta_i) \text{ as } n \rightarrow \infty, \text{ where } \delta_i := \sum_{j=1}^{q-1} \lambda_j^i.$$

In particular, δ_i is positive for every even $i \geq 2$.

The proof of Lemma A.6 is given in Section A.3.

Lemma A.7. *Let δ_i , $i = 2, 4, \dots$ be as in Lemma A.6. For every finite sequence m_1, \dots, m_k of nonnegative integers, it holds that*

$$\frac{\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{P}}(\eta)[X_2]_{m_1} \cdots [X_{2k}]_{m_k}]}{\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{P}}(\eta)]} \rightarrow \prod_{i=1}^k (\mu_i(1 + \delta_i))^{m_i} \text{ as } n \rightarrow \infty.$$

Once we give the proof of Lemma A.6, the proof of Lemma A.7 is identical to [MWW09, Proof of Lemma 7.5] and is omitted.

Lemma A.8. *In the notation and setting of Lemma A.4, it holds that*

$$\exp\left(\sum_{\text{even } i \geq 2} \mu_i \delta_i^2\right) = \prod_{i=1}^{q-1} \prod_{j=1}^{q-1} (1 - (\Delta - 1)^2 \lambda_i^2 \lambda_j^2)^{-1/2} \prod_{i=1}^{q-1} \prod_{j=1}^{q-1} (1 - \lambda_i^2 \lambda_j^2)^{-(\Delta-1)/2}.$$

Moreover, $\sum_{\text{even } i \geq 2} \mu_i \delta_i < \infty$.

The proof of Lemma A.8 is given in Section A.3.

Finally, we find the asymptotics of the second moment over the first moment squared.

Lemma A.9. *In the notation and setting of Lemma A.4, it holds that*

$$\lim_{n \rightarrow \infty} \frac{\mathbf{E}_{\mathcal{G}_n^r}[(Z_G^{\mathbf{P}}(\eta))^2]}{(\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{P}}(\eta)])^2} = C, \text{ where } C := \prod_{i=1}^{q-1} \prod_{j=1}^{q-1} (1 - (\Delta - 1)^2 \lambda_i^2 \lambda_j^2)^{-1/2} \prod_{i=1}^{q-1} \prod_{j=1}^{q-1} (1 - \lambda_i^2 \lambda_j^2)^{-(\Delta-1)/2}.$$

The proof of Lemma A.9 is quite extensive. In Appendix B, we first reduce the lemma to the case $r = 0$. This part of the proof is standard and is analogous to the proof of Lemma 6.11. Then, we compute the asymptotics in terms of determinants of relevant Hessian matrices. These determinants are computed in Appendix B.1.2, where also the proof of Lemma A.9 is completed.

With Lemmas A.5–A.9 at hand, the proof of Lemma A.3 is immediate.

Proof of Lemma A.3. We prove the first part of the lemma by verifying the assumptions of Theorem A.1. Lemma A.5 verifies assumption (A1), Lemma A.7 verifies assumption (A2) and Lemmas A.8 and A.9 verify assumptions (A3) and (A4). This proves the first part of the lemma.

For the second part, just use the second parts in Lemmas A.4, A.6, A.8 to establish Items (i)–(iv). \square

A.3 Proofs of Lemmas A.6 and A.8

In this section, we give the proofs of Lemmas A.6 and A.8.

Proof of Lemma A.6. The proof is close to [MWW09, Proof of Lemma 7.4], the approach only needs a few modifications to account for the q -spin setting. We make the minor notation change from X_i to X_ℓ .

We will do the computations for the case of $G \sim \mathcal{G}_n$ and the random variables $Z_G^{\alpha, \beta}$, the extension to the case $G \sim \mathcal{G}_n^r$ and the random variables $Z_G^{\mathbf{p}}(\eta)$ has slightly more complicated expressions, but otherwise the derivation is completely analogous (see for example [Sly10, Proof of Lemma 3.8]). We will show that

$$\frac{\mathbf{E}_{\mathcal{G}_n}[Z_G^{\alpha, \beta} X_\ell]}{\mathbf{E}_{\mathcal{G}_n}[Z_G^{\alpha, \beta}]} \rightarrow \mu_\ell(1 + \delta_\ell) \text{ as } n \rightarrow \infty,$$

where μ_ℓ, δ_ℓ are as in the statement of the lemma. For simplicity, let $\mathcal{G} := \mathcal{G}_n$.

Let $\mathcal{S} = \{S_1, \dots, S_q\}$ and $\mathcal{T} = \{T_1, \dots, T_q\}$ be partitions of V_1 and V_2 respectively such that $|S_i| = \alpha_i n$ and $|T_j| = \beta_j n$ for all $i, j \in [q]$. Denote by $Y_{\mathcal{S}, \mathcal{T}}$ the weight of the configuration σ that \mathcal{S}, \mathcal{T} induce, i.e. for a vertex $v \in V_1$, $\sigma(v) = i$ iff $v \in S_i$ and similarly for vertices in V_2 .

Fix a specific pair of \mathcal{S}, \mathcal{T} . By symmetry,

$$\frac{\mathbf{E}_{\mathcal{G}}[Z_G^{\alpha, \beta} X_\ell]}{\mathbf{E}_{\mathcal{G}}[Z_G^{\alpha, \beta}]} = \frac{\mathbf{E}_{\mathcal{G}}[Y_{\mathcal{S}, \mathcal{T}} X_\ell]}{\mathbf{E}_{\mathcal{G}}[Y_{\mathcal{S}, \mathcal{T}}]}. \quad (150)$$

We now decompose X_ℓ as follows:

- ξ will denote a proper Δ -edge colored, rooted and oriented ℓ -cycle ($r(\Delta, \ell)$ possibilities), in which the vertices are colored with $\{Y_1, \dots, Y_q, G_1, \dots, G_q\}$ and edges are colored with $\{1, \dots, \Delta\}$.

A vertex colored with Y_i (resp. G_i) for some $i \in [q]$ will be loosely called yellow (resp. green) and signifies that the vertex belongs to S_i (resp. T_i). Since a yellow vertex belongs to V_1 , and a green vertex belongs to V_2 , a vertex coloring is consistent with the bipartiteness of the random graph if adjacent vertices of the cycle are not both yellow or green, that is, the vertex assignments which are prohibited for neighboring vertices in the cycle are (Y_i, Y_j) and (G_i, G_j) , $\forall (i, j) \in [q]^2$. Note here that we do not explicitly prohibit assignments (Y_i, G_j) in the presence of a hard constraint $B_{ij} = 0$; this will be accounted otherwise. The color of the edges will prescribe which of the Δ perfect matchings an edge of a (potential) cycle will belong to.

- Given ξ , ζ denotes a position that an ℓ -cycle can be, i.e., the exact vertices it traverses in order, such that the prescription of the vertex colors of ξ is satisfied.
- $\mathbf{1}_{\xi, \zeta}$ is the indicator function whether a cycle specified by ξ, ζ is present in the graph G .

Note that each possible cycle corresponds to exactly 2ℓ different configurations ξ (the number of ways to root and orient the cycle). For each of those ξ , the respective sets of configurations ζ are the same. Hence, we may write

$$X_\ell = \frac{1}{2\ell} \sum_{\xi} \sum_{\zeta} \mathbf{1}_{\xi, \zeta}.$$

Let $p_1 := \Pr_{\mathcal{G}}[\mathbf{1}_{\xi, \zeta} = 1]$. We have

$$\mathbf{E}_{\mathcal{G}}[Y_{\mathcal{S}, \mathcal{T}} X_{\ell}] = \frac{1}{2\ell} \sum_{\xi} \sum_{\zeta} p_1 \cdot \mathbf{E}[Y_{\mathcal{S}, \mathcal{T}} | \mathbf{1}_{\xi, \zeta} = 1].$$

In light of (150), we need to study the ratio $\mathbf{E}_{\mathcal{G}}[Y_{\mathcal{S}, \mathcal{T}} | \mathbf{1}_{\xi, \zeta} = 1] / \mathbf{E}_{\mathcal{G}}[Y_{\mathcal{S}, \mathcal{T}}]$. At this point, to simplify notation, we may assume that ξ, ζ are fixed.

We have shown in Section 2 that

$$\mathbf{E}_{\mathcal{G}}[Y_{\mathcal{S}, \mathcal{T}}] = \left(\sum_{\mathbf{x}} \binom{n}{x_{11}n, \dots, x_{qq}n}^{-1} \prod_i \binom{\alpha_i n}{x_{i1}n, \dots, x_{iq}n} \prod_j \binom{\beta_j n}{x_{1j}n, \dots, x_{qj}n} \prod_{i,j} B_{ij}^{nx_{ij}} \right)^{\Delta}, \quad (151)$$

where the variables $\mathbf{x} = (x_{11}, \dots, x_{qq})$ denote the number of edges between \mathcal{S}, \mathcal{T} in one matching. In particular nx_{ij} is the number of edges between the sets S_i and T_j .

To calculate $\mathbf{E}_{\mathcal{G}}[Y_{\mathcal{S}, \mathcal{T}} | \mathbf{1}_{\xi, \zeta} = 1]$, we need some notation. For colors $c_1, c_2 \in \{Y_1, \dots, Y_q, G_1, \dots, G_q\}$, we say that an edge is of type $\{c_1, c_2\}$ if its endpoints have colors c_1, c_2 . Let y_i, g_j denote the number of vertices colored with Y_i, G_j respectively. For $k = 1, \dots, \Delta$, let $a_{ij}(k)$ denote the number of edges of color k and type $\{Y_i, G_j\}$. Finally, for $i, j \in [q]$ let $a_{ij} = \sum_k a_{ij}(k)$. By considering the sum of the degrees of vertices colored Y_i , the sum of the degrees of vertices colored G_j and the total number of edges of the cycle, we obtain the following equalities.

$$\sum_j a_{ij} = 2y_i, \quad \sum_i a_{ij} = 2g_j, \quad \sum_{i,j} a_{ij} = 2\ell. \quad (152)$$

We are almost set to compute $\mathbf{E}[Y_{\mathcal{S}, \mathcal{T}} | \mathbf{1}_{\xi, \zeta} = 1]$. We denote by \mathbf{x}_k the same set of variables as in (151) but for the k -th matching. Namely, $nx_{ij,k}$ is the number of (undetermined) edges between sets S_i and T_j in the k -th matching. This number includes the $a_{ij}(k)$ edges prescribed by ξ, ζ . To simplify the following fomulas, let $nx'_{ij,k} = nx_{ij,k} - a_{ij}(k)$ and set $E = \mathbf{E}_{\mathcal{G}}[Y_{\mathcal{S}, \mathcal{T}} | \mathbf{1}_{\xi, \zeta} = 1]$. We have

$$E = \prod_{k=1}^{\Delta} \left[\sum_{\mathbf{x}_k} \binom{n - \sum_{i,j} a_{ij}(k)}{x'_{11,k}n, \dots, x'_{qq,k}n}^{-1} \prod_i \binom{\alpha_i n - \sum_j a_{ij}(k)}{nx'_{i1,k}, \dots, nx'_{iq,k}} \prod_j \binom{\beta_j n - \sum_i a_{ij}(k)}{nx'_{1j,k}, \dots, nx'_{qj,k}} \prod_{i,j} B_{ij}^{nx_{ij,k}} \right].$$

In the above sums, for any $\varepsilon > 0$ and all sufficiently large n , terms whose \mathbf{x}_k 's are ε -far from the optimal value of \mathbf{x} given in Lemma 4.3 have exponentially small contribution and may be ignored. Standard approximations of binomial coefficients, see for example [GSV12, Lemma 27], give

$$\frac{\binom{\alpha_i n - \sum_j a_{ij}(k)}{x'_{i1,k}n, \dots, x'_{iq,k}n}}{\binom{\alpha_i n}{x_{i1,k}n, \dots, x_{iq,k}n}} \sim \frac{\prod_j (x_{ij,k})^{a_{ij}(k)}}{\alpha_i^{\sum_j a_{ij}(k)}}, \quad \frac{\binom{\beta_j n - \sum_i a_{ij}(k)}{x'_{1j,k}n, \dots, x'_{qj,k}n}}{\binom{\beta_j n}{x_{1j,k}n, \dots, x_{qj,k}n}} \sim \frac{\prod_i (x_{ij,k})^{a_{ij}(k)}}{\beta_j^{\sum_i a_{ij}(k)}}, \quad \frac{\binom{n - \sum_{i,j} a_{ij}(k)}{x'_{11,k}n, \dots, x'_{qq,k}n}}{\binom{n}{x_{11,k}n, \dots, x_{qq,k}n}} \sim \prod_{i,j} (x_{ij,k})^{a_{ij}(k)}.$$

Thus, we obtain

$$\frac{\mathbf{E}_{\mathcal{G}}[Y_{\mathcal{S}, \mathcal{T}} | \mathbf{1}_{\xi, \zeta} = 1]}{\mathbf{E}_{\mathcal{G}}[Y_{\mathcal{S}, \mathcal{T}}]} \sim \frac{\prod_{i,j} (x_{ij})^{a_{ij}}}{\prod_i \alpha_i^{\sum_j a_{ij}} \prod_j \beta_j^{\sum_i a_{ij}}}.$$

We have $p_1 \sim n^{-\ell}$ and for given ξ , the number of possible ζ is asymptotic to $n^{\ell} \prod_i \alpha_i^{y_i} \prod_j \beta_j^{g_j}$. Thus, for the given ξ , we have

$$\frac{\sum_{\zeta} p_1 \mathbf{E}_{\mathcal{G}}[Y_{\mathcal{S}, \mathcal{T}} | \mathbf{1}_{\xi, \zeta} = 1]}{\mathbf{E}_{\mathcal{G}}[Y_{\mathcal{S}, \mathcal{T}}]} \sim \frac{\prod_i \alpha_i^{y_i} \prod_j \beta_j^{g_j} \prod_{i,j} (x_{ij})^{a_{ij}}}{\prod_i \alpha_i^{\sum_j a_{ij}} \prod_j \beta_j^{\sum_i a_{ij}}} = \prod_{i,j} \left(\frac{x_{ij}}{\sqrt{\alpha_i \beta_j}} \right)^{a_{ij}}$$

Note that the rhs evaluates to 0 whenever there exist i, j such that $B_{ij} = 0$ but $a_{ij} \neq 0$, since then we have $x_{ij} = 0$. This is in complete accordance with the fact that the configuration induced by the partition $\{\mathcal{S}, \mathcal{T}\}$ has zero weight. Thus, by (150), we have

$$\frac{\mathbf{E}_{\mathcal{G}}[Z_G^{\alpha, \beta} X_{\ell}]}{\mathbf{E}_{\mathcal{G}}[Z_G^{\alpha, \beta}]} \sim \frac{r(\Delta, \ell)}{2\ell} \cdot \sum_{\xi} N_{\mathbf{a}} \left(\frac{x_{ij}}{\sqrt{\alpha_i \beta_j}} \right)^{a_{ij}},$$

where $\mathbf{a} = \{a_{11}, \dots, a_{qq}\}$ and $N_{\mathbf{a}}$ is the number of possible ξ with a_{ij} edges having assignment (Y_i, G_j) . To analyze this sum, we employ a technique given in [Jan95]. The idea is to define a weighted transition matrix and view it as the (weighted) adjacency matrix of a weighted graph. The powers of the matrix count the (multiplicative) weight of walks in the graph and a closed walk in this graph will correspond to a specification ξ . By defining the weights appropriately, one can also ensure that each closed walk will correctly capture the weight of the specification ξ .

In our setting, the transition matrix is simply the matrix \mathbf{J} of Lemma A.4. The first q rows and q columns correspond to the colors Y_i and the remaining rows and columns to colors G_j . The total weight of closed walks of length ℓ is given by $\text{Tr}(\mathbf{J}^{\ell})$. Using the description of the eigenvalues given in Lemma A.4, we obtain that for even ℓ , $\text{Tr}(\mathbf{J}^{\ell}) = 2 \left(1 + \sum_{i=1}^{q-1} \lambda_i^{\ell} \right)$. This concludes the proof. \square

Proof of Lemma A.8. Using Lemma A.5, we have

$$\sum_{\text{even } i \geq 2} \mu_i \delta_i^2 = \sum_{\text{even } i \geq 2} \frac{r(\Delta, i)}{i} \cdot \left(\sum_{j=1}^{q-1} \lambda_j^i \right)^2 = \sum_{\text{even } i \geq 2} \frac{(\Delta-1)^i + (\Delta-1)}{i} \cdot \left(\sum_{j=1}^{q-1} \sum_{j'=1}^{q-1} \lambda_j^i \lambda_{j'}^i \right).$$

Observe that $\sum_{j \geq 1} \frac{x_j^2}{2j} = -\frac{1}{2} \ln(1-x^2)$ for all $|x| < 1$. By Lemma A.4, the λ_j 's satisfy $(\Delta-1)\lambda_j < 1$ for all j , so that $(\Delta-1)\lambda_j \lambda_{j'} < 1$ for all j, j' . It follows that

$$\sum_{\text{even } i \geq 2} \mu_i \delta_i^2 = -\frac{1}{2} \left(\sum_{i,j} \ln(1 - (\Delta-1)^2 \lambda_i^2 \lambda_j^2) + (\Delta-1) \sum_{i,j} \ln(1 - \lambda_i^2 \lambda_j^2) \right),$$

thus proving the first part of the lemma. The proof of $\sum_i \mu_i \delta_i < \infty$ is completely analogous. \square

B Moment Asymptotics

In this section, we prove Lemma A.9. For the purposes of this section, we will identify \mathcal{Q} with the dominant phases of a random Δ -regular bipartite graph. Thus, we will use $\mathbf{p} \in \mathcal{Q}$ to denote a dominant phase (α, β) .

We first recall some relevant definitions from Section 6.4. For $r \geq 0$, let $G \sim \mathcal{G}_n^r$ and $\sigma : U \cup W \rightarrow [q]$ be a configuration on G . The *footprint* of σ is a pair of q -dimensional vectors $\alpha_{\sigma}, \beta_{\sigma}$ whose i -th entries are equal to $|\sigma^{-1} \cap U^+|/n, |\sigma^{-1} \cap U^-|/n$, respectively. The *phase* $Y(\sigma)$ of σ is the dominant phase (α, β) which is closest to $(\alpha_{\sigma}, \beta_{\sigma})$, precisely:

$$Y(\sigma) = \arg \max_{\mathbf{p}=(\alpha, \beta) \in \mathcal{Q}} \left(\|\alpha - \alpha_{\sigma}\|_2^2 + \|\beta - \beta_{\sigma}\|_2^2 \right)^{1/2} \quad (153)$$

Finally, recall that for $\mathbf{p} \in \mathcal{Q}$, $Z_G^{\mathbf{p}}$ is the partition function “conditioned on the phase \mathbf{p} ”, i.e., the contribution to the partition function of G from configurations σ with $Y(\sigma) = \mathbf{p}$, and for $\eta : W \rightarrow [q]$, $Z_G^{\mathbf{p}}(\eta)$ is the contribution to the partition function of G from configurations σ with $Y(\sigma) = \mathbf{p}$ and $\sigma_W = \eta$, see (98) for more details.

In the setting of Lemma A.9, we need to compute the asymptotics of $\mathbf{E}_{\mathcal{G}_n^r}[(Z_G^{\mathbf{p}}(\eta))^2]/(\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta)])^2$ for $\mathbf{p} \in \mathcal{Q}$ and configurations $\eta : W \rightarrow [q]$. The following lemma reduces the computation to the case $r = 0$. Note that for $r = 0$, the set of vertices W is empty and the distribution \mathcal{G}_n^r coincides with the distribution $\mathcal{G} := \mathcal{G}_n$ on random Δ -regular bipartite graphs from Section 2.

Lemma B.1. *Let $\mathbf{p} = (\alpha, \beta) \in \mathcal{Q}$ be a Hessian dominant phase. Then, for every fixed $r > 0$, for every $\eta : W \rightarrow [q]$ it holds that*

$$\lim_{n \rightarrow \infty} \frac{\mathbf{E}_{\mathcal{G}_n^r}[(Z_G^{\mathbf{p}}(\eta))^2]}{(\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta)])^2} = \lim_{n \rightarrow \infty} \frac{\mathbf{E}_{\mathcal{G}_n}[(Z_G^{\mathbf{p}})^2]}{(\mathbf{E}_{\mathcal{G}_n}[Z_G^{\mathbf{p}}])^2}.$$

Proof. By Lemma 6.11, we have

$$\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta)] = (1 + o(1))C^r \nu_{\mathbf{p}}^{\otimes}(\eta) \mathbf{E}_{\mathcal{G}_n}[Z_G^{\mathbf{p}}], \quad (100)$$

where $C(\mathbf{p})$ is the constant in Lemma 6.11 and $\nu_{\mathbf{p}}^{\otimes}(\eta)$ is defined in (94). As in (the proof of) Lemma 6.11, we also obtain

$$\mathbf{E}_{\mathcal{G}_n^r}[(Z_G^{\mathbf{p}}(\eta))^2] = (1 + o(1))C^{2r}(\nu_{\mathbf{p}}^{\otimes}(\eta))^2 \mathbf{E}_{\mathcal{G}_n}[(Z_G^{\mathbf{p}})^2], \quad (154)$$

where $C(\mathbf{p})$ is again the constant in Lemma 6.11. Combining (100) and (154) proves the lemma. \square

In light of Lemma B.1, we need to compute the limiting ratio of $\mathbf{E}_{\mathcal{G}_n}[(Z_G^{\mathbf{p}})^2]/(\mathbf{E}_{\mathcal{G}_n}[Z_G^{\mathbf{p}}])^2$ for $\mathbf{p} \in \mathcal{Q}$. We do this by computing separately the asymptotics of $\mathbf{E}_{\mathcal{G}_n}[Z_G^{\mathbf{p}}]$ and $\mathbf{E}_{\mathcal{G}_n}[(Z_G^{\mathbf{p}})^2]$. We begin with an observation that will allow us to deduce the asymptotics of $\mathbf{E}_{\mathcal{G}}[(Z_G^{\mathbf{p}})^2]$ from the asymptotics of $\mathbf{E}_{\mathcal{G}}[Z_G^{\mathbf{p}}]$ applied to the spin system with interaction matrix $\mathbf{B} \otimes \mathbf{B}$.

Lemma B.2. *Let $\mathbf{p} = (\alpha, \beta) \in \mathcal{Q}$ be a dominant phase for the spin system with interaction matrix \mathbf{B} . Then $\mathbf{p}' = (\alpha \otimes \alpha, \beta \otimes \beta)$ is a dominant phase for the spin system with interaction matrix $\mathbf{B} \otimes \mathbf{B}$.*

Let $Z_G^{\mathbf{p}'}$ equal the partition function for the spin system with interaction matrix $\mathbf{B} \otimes \mathbf{B}$ conditioned on the phase \mathbf{p}' . Then, $\lim_{n \rightarrow \infty} \frac{\mathbf{E}_{\mathcal{G}}[(Z_G^{\mathbf{p}})^2]}{\mathbf{E}_{\mathcal{G}}[Z_G^{\mathbf{p}}]} = 1$.

Proof of Lemma B.2. The first part of the lemma is an immediate consequence of Lemma 3.2. In fact, the proof of Lemma 3.2 shows the stronger fact that the set of dominant phases for the spin system with interaction matrix $\mathbf{B} \otimes \mathbf{B}$ is given by $\mathcal{Q}^{\otimes 2} := \{(\alpha \otimes \alpha', \beta \otimes \beta') \mid (\alpha, \beta), (\alpha', \beta') \in \mathcal{Q}\}$.

For the second part, let

$$\Sigma_1 = \left\{ (\alpha', \beta') \mid \alpha', \beta' \in \Delta_q, \mathbf{p} = \arg \min_{\mathbf{p}^* = (\alpha^*, \beta^*) \in \mathcal{Q}} (\|\alpha' - \alpha^*\|^2 + \|\beta' - \beta^*\|^2)^{1/2} \right\}, \quad (155)$$

so that

$$Z_G^{\mathbf{p}} = \sum_{(\alpha', \beta') \in \Sigma_1} Z_G^{\alpha', \beta'} \text{ and } (Z_G^{\mathbf{p}})^2 = \sum_{(\alpha', \beta'), (\alpha'', \beta'') \in \Sigma_1} Z_G^{\alpha', \beta'} Z_G^{\alpha'', \beta''}.$$

It follows that $\mathbf{E}_{\mathcal{G}}[(Z_G^{\mathbf{p}})^2]$ is given by the sum in the r.h.s. in (6), but now the sum is over γ, δ which satisfy

$$\begin{aligned} \sum_k \gamma_{ik} &= \alpha'_i \quad (\forall i \in [q]), & \sum_l \delta_{jl} &= \beta'_j \quad (\forall j \in [q]), \\ \sum_i \gamma_{ik} &= \alpha''_k \quad (\forall k \in [q]), & \sum_j \delta_{jl} &= \beta''_l \quad (\forall l \in [q]), \end{aligned}$$

and $(\alpha', \beta'), (\alpha'', \beta'')$ range over Σ_1 . Note that by the definition of Σ_1 , $\mathbf{p} = (\alpha, \beta)$ is the unique dominant phase (of the spin system with interaction matrix \mathbf{B}) contained in Σ_1 . By Lemma 3.2, for

any $\varepsilon > 0$ and all sufficiently large n , terms in the sum with $(\|\gamma - \alpha \otimes \alpha\|_2^2 + \|\delta - \beta \otimes \beta\|_2^2)^{1/2} \geq \varepsilon$ have exponentially small contribution and hence may be ignored. Similarly, for the spin system with interaction matrix $\mathbf{B} \otimes \mathbf{B}$, $\mathbf{E}_{\mathcal{G}}[Z_G^{\mathbf{p}'}]$ is given by the sum in the r.h.s. in (6), where now the sum is over γ, δ with $(\gamma, \delta) \in \Sigma_2$, where

$$\Sigma_2 := \left\{ (\gamma', \delta') \mid \gamma', \delta' \in \Delta_{q^2}, \mathbf{p}' = \arg \min_{\mathbf{p}^*=(\gamma^*, \delta^*) \in \mathcal{Q}^{\otimes 2}} (\|\gamma' - \gamma^*\|^2 + \|\delta' - \delta^*\|^2)^{1/2} \right\}.$$

Once again, by Lemma 3.2, for any $\varepsilon > 0$ and all sufficiently large n , terms in the sum with $(\|\gamma - \alpha \otimes \alpha\|_2^2 + \|\delta - \beta \otimes \beta\|_2^2)^{1/2} \geq \varepsilon$ have exponentially small contribution and hence may be ignored. It follows that for all sufficiently small $\varepsilon > 0$, the remaining terms in the two sums are identical which completes the proof. \square

As a consequence of Lemma B.2, we may focus on the asymptotics of the first moment $\mathbf{E}_{\mathcal{G}}[Z_G^{\mathbf{p}}]$. Let

$$P_1 = \{(i, j) \in [q]^2 \mid B_{ij} > 0\}. \quad (156)$$

In the presence of a hard constraint $B_{ij} = 0$, edge assignments (i, j) yield a zero-weight configuration. In the maximization of Υ_1 , the hard constraint $B_{ij} = 0$ was not directly relevant, since for $x_{ij} > 0$ the function Υ_1 evaluates to $-\infty$. Indeed, we found that the optimal x_{ij} is of the form $B_{ij}R_iC_j$ and hence zero. However, the asymptotics of $\mathbf{E}_{\mathcal{G}}[Z_G^{\mathbf{p}}]$ include products of the optimal values of the x_{ij} and to correctly capture them, we need to explicitly rule out the zero values.

To do so, in the formulation (5), we hard-code $x_{ij} = 0$ for a pair $(i, j) \notin P_1$ and hence the variables $\alpha, \beta, \mathbf{x}$ are restricted to the space

$$\begin{aligned} \sum_i \alpha_i &= 1, & \sum_j \beta_j &= 1, \\ \sum_j x_{ij} &= \alpha_i \quad (\forall i \in [q]), & \sum_i x_{ij} &= \beta_j \quad (\forall j \in [q]), \\ x_{ij} &= 0 \quad (\forall (i, j) \in [q]^2 \setminus P_1), & x_{ij} &\geq 0 \quad (\forall (i, j) \in P_1). \end{aligned} \quad (157)$$

We will also need to have a set of affinely independent variables which describe the polytope (157). Note that the dimension of the polytope (157) is $(q^2 + 2q) - (2q + 1) - (q^2 - |P_1|) = |P_1| - 1$. To get affinely independent variables $\alpha, \beta, \mathbf{x}$, we use the equalities in (157) and substitute an appropriate set of $(q + 1)^2 - |P_1|$ variables. We will not need to understand these substitutions till Appendix B.1.1, yet in the integrations which follow it is preferable to have integration variables rather than integrate over subspaces.

After this process, we are going to have $|P_1| - 1$ variables lying in a full dimensional space. We refer to this set of variables as the full dimensional representation of (157). For simplicity, we will still use $\alpha, \beta, \mathbf{x}$ for these variables and refer, e.g., to x_{ij} even if x_{ij} is not in the full dimensional representation of (157), under the understanding that this is just a shorthand for the substituted expression. Using these conventions, we may view $\Upsilon_1(\alpha, \beta, \mathbf{x})$ as a function of the full dimensional representation of (157), and we will refer to this setup as the full dimensional representation of Υ_1 .

The following lemma expresses the asymptotics of $\mathbf{E}_{\mathcal{G}}[Z_G^{\mathbf{p}}]$ in terms of suitable determinants. The computation of these determinants is given in Section B.1.2, where also the proof of Lemma A.9 is completed.

Lemma B.3. *Let $\mathbf{p} = (\alpha^*, \beta^*)$ be a dominant phase, i.e., (α^*, β^*) maximizes $\Psi_1(\alpha, \beta)$. Let \mathbf{x}^* be the (unique) maximizer of $\Upsilon_1(\alpha^*, \beta^*, \mathbf{x})$ (given in Lemma 4.3). Denote by \mathbf{H}_1^f be the Hessian of the full dimensional representation of $\Upsilon_1(\alpha, \beta, \mathbf{x})$ scaled by $1/\Delta$ (evaluated at $\alpha^*, \beta^*, \mathbf{x}^*$) and by*

$\mathbf{H}_{1,\mathbf{x}}^f$ the square submatrix of \mathbf{H}_1^f corresponding to rows and columns indexed by \mathbf{x} . Then

$$\lim_{n \rightarrow \infty} \frac{\mathbf{E}_{\mathcal{G}}[Z_G^{\mathbf{p}}]}{e^{n\Upsilon_1(\boldsymbol{\alpha}^*, \boldsymbol{\beta}^*, \mathbf{x}^*)}} = \frac{\left(\prod_i \alpha_i^* \prod_j \beta_j^*\right)^{(\Delta-1)/2} \left(\prod_{(i,j) \in P_1} x_{ij}^*\right)^{-\Delta/2}}{\Delta^{q-1} (\text{Det}(-\mathbf{H}_1^f))^{1/2} (\text{Det}(-\mathbf{H}_{1,\mathbf{x}}^f))^{(\Delta-1)/2}}.$$

Proof of Lemma B.3. We assume that $\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}$ is a full dimensional representation of (157). We denote by $(\boldsymbol{\alpha}^*, \boldsymbol{\beta}^*, \mathbf{x}^*)$ the optimal vector which maximizes the full dimensional representation of $\Upsilon_1(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x})$. We have that $\alpha_i^*, \beta_j^* > 0$ for all i, j and $x_{ij}^* > 0$ for $(i, j) \in P_1$. Pick δ sufficiently small such that:

$$\|(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}) - (\boldsymbol{\alpha}^*, \boldsymbol{\beta}^*, \mathbf{x}^*)\|_2 \leq \delta \text{ implies } \alpha_i, \beta_j > 0 \text{ for all } i, j \text{ and } x_{ij} > 0 \text{ for } (i, j) \in P_1.$$

Since Υ_1 has the unique global maximum $(\boldsymbol{\alpha}^*, \boldsymbol{\beta}^*, \mathbf{x}^*)$ at the intersection of the spaces (155) and (157), standard compactness arguments imply that there exists $\varepsilon(\delta) > 0$ such that $\|(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}) - (\boldsymbol{\alpha}^*, \boldsymbol{\beta}^*, \mathbf{x}^*)\|_2 \geq \delta$ implies $\Upsilon_1(\boldsymbol{\alpha}^*, \boldsymbol{\beta}^*, \mathbf{x}^*) - \Upsilon_1(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}) \geq \varepsilon$. It follows that the contribution of terms with $\|(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}) - (\boldsymbol{\alpha}^*, \boldsymbol{\beta}^*, \mathbf{x}^*)\|_2 \geq \delta$ to $\mathbf{E}_{\mathcal{G}}[(Z_G^{\mathbf{p}})^2]$ is exponentially small and may be ignored. Hence we may restrict our attention to $\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}$ satisfying $\|(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}) - (\boldsymbol{\alpha}^*, \boldsymbol{\beta}^*, \mathbf{x}^*)\|_2 < \delta$. Moreover, using Taylor's expansion, we may choose δ small enough such that Υ_1 decays quadratically in a δ -ball around $(\boldsymbol{\alpha}^*, \boldsymbol{\beta}^*, \mathbf{x}^*)$.

Utilizing the choice of δ and Stirling's approximation for factorials, we thus obtain

$$\begin{aligned} \frac{\mathbf{E}_{\mathcal{G}}[Z_G^{\mathbf{p}}]}{e^{n\Upsilon_1(\boldsymbol{\alpha}^*, \boldsymbol{\beta}^*, \mathbf{x}^*)}} &= \left(1 + O(n^{-1})\right) \sum_{\boldsymbol{\alpha}, \boldsymbol{\beta}} \left(\frac{1}{\sqrt{2\pi n}}\right)^{2(q-1)} \left(\prod_i \alpha_i \prod_j \beta_j\right)^{(\Delta-1)/2} \\ &\quad \left[\sum_{\mathbf{x}} \left(\frac{1}{\sqrt{2\pi n}}\right)^{|P_1|-(2q-1)} \left(\prod_{(i,j) \in P_1} \frac{1}{\sqrt{x_{ij}}}\right) e^{n(\Upsilon_1(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}) - \Upsilon_1(\boldsymbol{\alpha}^*, \boldsymbol{\beta}^*, \mathbf{x}^*))/\Delta} \right]^{\Delta}. \end{aligned}$$

We now compute

$$L := \lim_{n \rightarrow \infty} \frac{\mathbf{E}_{\mathcal{G}}[Z_G^{\mathbf{p}}]}{e^{n\Upsilon_1(\boldsymbol{\alpha}^*, \boldsymbol{\beta}^*, \mathbf{x}^*)}}.$$

Standard techniques of rewriting sums as integrals and an application of the dominated convergence theorem (see for example [JLR00, Section 9.4]) ultimately give

$$\begin{aligned} L &= \left(\prod_i \alpha_i^* \prod_j \beta_j^*\right)^{(\Delta-1)/2} \left(\prod_{(i,j) \in P_1} x_{ij}^*\right)^{-\Delta/2} \\ &\quad \left(\frac{1}{\sqrt{2\pi}}\right)^{2(q-1)} \int_{-\infty}^{\infty} \dots \int_{-\infty}^{\infty} \left[\left(\frac{1}{\sqrt{2\pi}}\right)^{|P_1|-(2q-1)} \int_{-\infty}^{\infty} \dots \int_{-\infty}^{\infty} e^{\frac{1}{2}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}) \cdot \mathbf{H} \cdot (\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x})^{\top}} d\mathbf{x}\right]^{\Delta} d\boldsymbol{\alpha} d\boldsymbol{\beta}, \end{aligned} \tag{158}$$

where \mathbf{H} denotes the Hessian matrix of Υ_1 evaluated at $(\boldsymbol{\alpha}^*, \boldsymbol{\beta}^*, \mathbf{x}^*)$ scaled by $1/\Delta$ and the operator \cdot stands for matrix multiplication.

We thus focus on computing the integral in (158). We begin with the inner integration. Let

$$I_1 = \left[\left(\frac{1}{\sqrt{2\pi}}\right)^{|P_1|-(2q-1)} \int_{-\infty}^{\infty} \dots \int_{-\infty}^{\infty} e^{\frac{1}{2}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}) \cdot \mathbf{H} \cdot (\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x})^{\top}} d\mathbf{x}\right]^{\Delta}.$$

To calculate I_1 , we first decompose the exponent to isolate the terms involving \mathbf{x} . We obtain

$$\frac{1}{2}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}) \cdot \mathbf{H} \cdot (\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x})^{\top} = \frac{1}{2}(\boldsymbol{\alpha}, \boldsymbol{\beta}) \cdot \mathbf{H}_{\boldsymbol{\alpha}, \boldsymbol{\beta}} \cdot (\boldsymbol{\alpha}, \boldsymbol{\beta})^{\top} - \frac{1}{2}\mathbf{x} \cdot (-\mathbf{H}_{\mathbf{x}}) \cdot \mathbf{x}^{\top} + \mathbf{T} \cdot \mathbf{x}^{\top},$$

where $\mathbf{H} = \begin{bmatrix} \mathbf{H}_{\boldsymbol{\alpha}, \boldsymbol{\beta}} & \mathbf{H}_{\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}} \\ \mathbf{H}_{\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}}^{\top} & \mathbf{H}_{\mathbf{x}} \end{bmatrix}$ and $\mathbf{T} = (\boldsymbol{\alpha}, \boldsymbol{\beta}) \cdot \mathbf{H}_{\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}}$. Specifically:

- $\mathbf{H}_{\alpha,\beta}$ is the square submatrix of \mathbf{H} corresponding to the rows indexed by α, β and the columns indexed by α, β ,
- $\mathbf{H}_{\mathbf{x}}$ is the square submatrix of \mathbf{H} corresponding to the rows indexed by \mathbf{x} and the columns indexed by \mathbf{x} ,
- $\mathbf{T} = (\alpha, \beta) \cdot \mathbf{H}_{\alpha\beta,\mathbf{x}}$, where $\mathbf{H}_{\alpha\beta,\mathbf{x}}$ is the submatrix of \mathbf{H} corresponding to the rows indexed by α, β and the columns indexed by \mathbf{x} .

Note that $\mathbf{H}_{\mathbf{x}}$ is the Hessian of $g_1(\mathbf{x})$ evaluated at \mathbf{x}^* . Since $g_1(\mathbf{x})$ is concave, we have that $\mathbf{H}_{\mathbf{x}}$ is negative definite. Utilizing this decomposition, we obtain

$$\begin{aligned} I_1 &= e^{\frac{\Delta}{2}(\alpha,\beta) \cdot \mathbf{H}_{\alpha,\beta} \cdot (\alpha,\beta)^\top} \left[\left(\frac{1}{\sqrt{2\pi}} \right)^{|P_1|-(2q-1)} \int_{-\infty}^{\infty} \dots \int_{-\infty}^{\infty} e^{-\frac{1}{2}\mathbf{x} \cdot (-\mathbf{H}_{\mathbf{x}}) \cdot \mathbf{x}^\top + \mathbf{T} \cdot \mathbf{x}^\top} d\mathbf{x} \right]^\Delta \\ &= \frac{1}{(\text{Det}(-\mathbf{H}_{\mathbf{x}}))^{\Delta/2}} e^{\frac{\Delta}{2}(\mathbf{T} \cdot (-\mathbf{H}_{\mathbf{x}})^{-1} \cdot \mathbf{T}^\top + (\alpha,\beta) \cdot \mathbf{H}_{\alpha,\beta} \cdot (\alpha,\beta)^\top)}. \end{aligned}$$

We are left with the task of computing the integral

$$I_2 = \left(\frac{1}{\sqrt{2\pi}} \right)^{2(q-1)} \int_{-\infty}^{\infty} \dots \int_{-\infty}^{\infty} e^{\frac{\Delta}{2}(\mathbf{T} \cdot (-\mathbf{H}_{\mathbf{x}})^{-1} \cdot \mathbf{T}^\top + (\alpha,\beta) \cdot \mathbf{H}_{\alpha,\beta} \cdot (\alpha,\beta)^\top)} d\alpha d\beta. \quad (159)$$

Using the definition of \mathbf{T} , we have

$$\mathbf{T} \cdot (-\mathbf{H}_{\mathbf{x}})^{-1} \cdot \mathbf{T}^\top + (\alpha,\beta) \cdot \mathbf{H}_{\alpha,\beta} \cdot (\alpha,\beta)^\top = (\alpha,\beta) \cdot (\mathbf{H}_{\alpha,\beta} - \mathbf{H}_{\alpha\beta,\mathbf{x}} \cdot \mathbf{H}_{\mathbf{x}}^{-1} \cdot \mathbf{H}_{\alpha\beta,\mathbf{x}}^\top) \cdot (\alpha,\beta)^\top.$$

The matrix $\mathbf{M} = \mathbf{H}_{\alpha,\beta} - \mathbf{H}_{\alpha\beta,\mathbf{x}} \cdot \mathbf{H}_{\mathbf{x}}^{-1} \cdot \mathbf{H}_{\alpha\beta,\mathbf{x}}^\top$ is the Schur complement of the block $\mathbf{H}_{\mathbf{x}}$ of \mathbf{H} . In fact, we have the identity $\text{Det}(\mathbf{H}) = \text{Det}(\mathbf{H}_{\mathbf{x}})\text{Det}(\mathbf{M})$ and in particular \mathbf{M} is negative definite. A Gaussian integration then yields

$$I_2 = \left(\frac{1}{\Delta^{2(q-1)} \text{Det}(-\mathbf{M})} \right)^{1/2} = \left(\frac{\text{Det}(-\mathbf{H}_{\mathbf{x}})}{\Delta^{2(q-1)} \text{Det}(-\mathbf{H})} \right)^{1/2}. \quad (160)$$

Combining equations (158), (159), (160), we obtain the statement of the Lemma. \square

B.1 The Determinants

This section addresses the computation of the determinants of the Hessians in Lemma B.3. The calculations are quite complex since one has to make a choice of free variables, do the substitutions, differentiate, and then hope that the structure of the problem will prevail in the determinants. Pushing this procedure in our setting leads to complications since the choice of free variables takes away much of the combinatorial structure of the problem. We follow a different path, which amongst other things, reveals that the determinants, via the matrix-tree theorem, correspond to counting weighted trees in appropriate graphs.

The proof has two parts. The first part connects different formulations of the Hessian of a constrained maximization in an abstract setting. Essentially, this puts together well known concepts from optimization in a way that will allow to stay as close as possible to the combinatorial structure of the determinants. The second part specialises the work of the first part to compute the required determinants and is unavoidably more computational.

B.1.1 Hessian formulations for Constrained problems

The setting of this section is the following: we are given Υ , a function of $\mathbf{z} \in \mathbb{R}^n$, subject to the linear constraints $\mathbf{A}\mathbf{z} = \mathbf{b}$, where $\mathbf{A} \in \mathbb{R}^{m \times n}$. The assumption of linear constraints stems from the setting of Lemma B.3, yet the arguments extend to other constraints as well by considering gradients of these constraints at the point \mathbf{z}_0 and implicit functions. W.l.o.g., we will also assume that $\mathbf{b} = \mathbf{0}$.

We are interested in the Hessian \mathbf{H}^f of a full dimensional representation of Υ . A full dimensional representation of Υ consists essentially of substituting an appropriate subset of the variables \mathbf{z} using the constraints $\mathbf{A}\mathbf{z} = \mathbf{0}$. Note that the representation is not as much tied to Υ as it is tied to the space $\mathbf{A}\mathbf{z} = \mathbf{0}$. Specifically, assume that the row rank of \mathbf{A} is r . In all the relevant constrained functions we consider, the constraints are not linearly independent so such an assumption is necessary. A full dimensional representation of Υ is specified by two submatrices of \mathbf{A} denoted by $(\mathbf{A}_f, \mathbf{A}_{fs})$. The matrix \mathbf{A}_f is a submatrix of \mathbf{A} consisting of r linearly independent rows of \mathbf{A} , so that $\mathbf{A}\mathbf{z} = \mathbf{0}$ iff $\mathbf{A}_f \mathbf{z} = \mathbf{0}$. Then, \mathbf{A}_{fs} is an $r \times r$ submatrix of \mathbf{A}_f which is invertible. The variables corresponding to columns of \mathbf{A}_{fs} are denoted by \mathbf{z}_s . The remaining variables \mathbf{z}_f are called free and \mathbf{A}_{ff} is the submatrix of \mathbf{A}_f induced by the columns indexed by \mathbf{z}_f . Renaming if needed, the equation $\mathbf{A}_f \mathbf{z} = \mathbf{0}$ may be naturally decomposed as

$$\begin{bmatrix} \mathbf{A}_{ff} & \mathbf{A}_{fs} \end{bmatrix} \begin{bmatrix} \mathbf{z}_f \\ \mathbf{z}_s \end{bmatrix} = \mathbf{0}, \text{ so that } \mathbf{z} = \begin{bmatrix} \mathbf{z}_f \\ \mathbf{z}_s \end{bmatrix} = \begin{bmatrix} \mathbf{I} \\ -(\mathbf{A}_{fs})^{-1} \mathbf{A}_{ff} \end{bmatrix} \mathbf{z}_f.$$

Thus, we can now think of Υ as a function which is completely determined by the variables \mathbf{z}_f which, in contrast with the variables \mathbf{z} , span a full dimensional space.

Denote by \mathbf{H} the unconstrained Hessian of Υ with respect to the variables \mathbf{z} and by \mathbf{H}^f the Hessian of the full dimensional representation of Υ with respect to the variables \mathbf{z}_f . The Hessians \mathbf{H} , \mathbf{H}^f are connected by the following equation, which follows by straightforward matrix calculus and its proof is omitted.

$$\mathbf{H}^f = \mathbf{S}^\top \mathbf{H} \mathbf{S}, \text{ where } \mathbf{S} = \begin{bmatrix} \mathbf{I} \\ -(\mathbf{A}_{fs})^{-1} \mathbf{A}_{ff} \end{bmatrix}. \quad (161)$$

Note that \mathbf{H}^f is different, though closely related, from the constrained Hessian \mathbf{H}^c of Υ in the subspace $\mathbf{A}\mathbf{z} = \mathbf{0}$, see for example [LY08, Chapter 10]. The constrained Hessian \mathbf{H}^c has infinitely many matrix representations, all of which correspond to similar matrices, that is, matrices with the same set of eigenvalues. A matrix representation may be obtained by first picking an orthonormal basis of the $(n - r)$ -dimensional space $\{\mathbf{z} \mid \mathbf{A}\mathbf{z} = \mathbf{0}\}$. Let \mathbf{E} denote the $n \times (n - r)$ matrix whose columns are the vectors in the basis. Then a matrix representation of \mathbf{H}^c is given by

$$\mathbf{H}^c = \mathbf{E}^\top \mathbf{H} \mathbf{E}, \quad (162)$$

where \mathbf{H} is as before the unconstrained Hessian of Υ with respect to the variables \mathbf{z} . We are ready to prove the following. It is useful to recall here that congruent matrices have the same number of negative, zero and positive eigenvalues.

Lemma B.4. *\mathbf{H}^f is congruent to any matrix representation of \mathbf{H}^c . Moreover, it holds that*

$$\text{Det}(\mathbf{H}^f) = \text{Det}(\mathbf{H}^c) \text{Det}(\mathbf{A}_f \mathbf{A}_f^\top) / \text{Det}(\mathbf{A}_{fs})^2.$$

Proof of Lemma B.4. The columns of the matrix \mathbf{S} defined in equation (161) form a basis of the space $\{\mathbf{z} \mid \mathbf{A}\mathbf{z} = \mathbf{0}\}$. Indeed, \mathbf{S} has clearly full column rank and also $\mathbf{A}_f \mathbf{S} = \mathbf{0}$ implying $\mathbf{A} \mathbf{S} = \mathbf{0}$ as well. For future use, by a direct evaluation

$$\mathbf{S}^\top \mathbf{S} = \mathbf{I} + \mathbf{A}_{ff}^\top (\mathbf{A}_{fs} \mathbf{A}_{fs}^\top)^{-1} \mathbf{A}_{ff}, \text{ so } \text{Det}(\mathbf{S}^\top \mathbf{S}) = \text{Det}(\mathbf{I} + \mathbf{A}_{ff} \mathbf{A}_{ff}^\top (\mathbf{A}_{fs} \mathbf{A}_{fs}^\top)^{-1}),$$

where the latter equality uses Sylvester's determinant theorem. This clearly yields

$$\text{Det}(\mathbf{S}^\top \mathbf{S}) = \text{Det}(\mathbf{A}_f \mathbf{A}_f^\top) / \text{Det}(\mathbf{A}_{fs})^2. \quad (163)$$

Comparing (161) and (162), the only difference is that \mathbf{S} does not necessarily encode an orthonormal basis. Nevertheless, there clearly exists an invertible matrix \mathbf{P} such that $\mathbf{S} \mathbf{P}$ consists of orthonormal columns, for example by the Gram-Schmidt process on the columns of \mathbf{S} . It follows that $\mathbf{P}^\top \mathbf{H}^f \mathbf{P}$ is a matrix representation of \mathbf{H}^c . This proves the first part of the lemma and also gives $\text{Det}(\mathbf{H}^c) = \text{Det}(\mathbf{H}^f) \text{Det}(\mathbf{P})^2$.

For the second part, the selection of \mathbf{P} implies that $(\mathbf{S} \mathbf{P})^\top \mathbf{S} \mathbf{P}$ is the identity matrix and hence $\text{Det}(\mathbf{S}^\top \mathbf{S}) \text{Det}(\mathbf{P})^2 = 1$. The desired equality follows. \square

Lemma B.4 allows us to focus on the determinant of \mathbf{H}^c or equivalently the product of its eigenvalues. The latter may be handled using bordered Hessians. Specifically, let \mathbf{A}_f be any submatrix of \mathbf{A} induced by r linearly independent rows. Then, λ is an eigenvalue of \mathbf{H}^c iff it is a root of the polynomial

$$p(\lambda) = \text{Det} \left(\begin{bmatrix} \mathbf{0} & \mathbf{A}_f \\ -\mathbf{A}_f^\top & \mathbf{H} - \lambda \mathbf{I}_n \end{bmatrix} \right). \quad (164)$$

In our case, deleting rows of \mathbf{A} to obtain \mathbf{A}_f would cause undesirable complications. In the following, we circumvent such deletions by adding suitable ‘‘perturbations’’. We will also allow for certain degrees of freedom to select the perturbations which will be exploited in the computations. We first prove the following.

For a polynomial $p(s)$, $[s^t]p(s)$ denotes the coefficient of s^t in $p(s)$.

Lemma B.5. *Let $\mathbf{M} \in \mathbb{R}^{m \times m}$ be a symmetric matrix with rank r and let $\mu_i, i = 1, \dots, m$ be the eigenvalues of \mathbf{M} with corresponding unit eigenvectors \mathbf{v}_i , where $\{\mathbf{v}_1, \dots, \mathbf{v}_m\}$ is an orthonormal basis of \mathbb{R}^m . Then, for any symmetric matrix $\mathbf{T} \in \mathbb{R}^{m \times m}$, it holds that*

$$[\varepsilon^{m-r}] \text{Det}(\varepsilon \mathbf{T} + \mathbf{M}) = \prod_{i: \mu_i \neq 0} \mu_i \prod_{i: \mu_i = 0} \mathbf{v}_i^\top \mathbf{T} \mathbf{v}_i, \quad (165)$$

In particular, if \mathbf{T} is positive semidefinite and $[\mathbf{T} \ \mathbf{M}]$ has full row rank, the rhs of (165) is non-zero.

Proof of Lemma B.5. Let $\mathbf{M}(\varepsilon) = \varepsilon \mathbf{T} + \mathbf{M}$ and denote by $\mu_i(\varepsilon), \mathbf{v}_i(\varepsilon)$ the eigenvalues and unit eigenvectors of $\mathbf{M}(\varepsilon)$. Rellich's theorem asserts that $\mu_i(\varepsilon)$ and $\mathbf{v}_i(\varepsilon)$ are analytic functions of ε around $\varepsilon = 0$. By Hadamard's first variation formula, we have $\frac{\partial \mu_i}{\partial \varepsilon} = \mathbf{v}_i^\top \frac{\partial \mathbf{M}}{\partial \varepsilon} \mathbf{v}_i$. At $\varepsilon = 0$, \mathbf{M} has rank r and hence exactly $m - r$ eigenvalues are zero. Thus, for small enough ε ,

$$\text{Det}(\mathbf{M}(\varepsilon)) = \varepsilon^{m-r} \prod_{i: \mu_i \neq 0} \mu_i \prod_{i: \mu_i = 0} \mathbf{v}_i^\top \mathbf{T} \mathbf{v}_i + O(\varepsilon^{m-r+1}).$$

Hence, $[\varepsilon^{m-r}] \text{Det}(\mathbf{M}(\varepsilon)) \neq 0$ if for every $\mathbf{v}_i \neq \mathbf{0}$ such that $\mathbf{M} \mathbf{v}_i = \mathbf{0}$, we have $\mathbf{v}_i^\top \mathbf{T} \mathbf{v}_i \neq 0$. The latter is true. Otherwise, using the positive semidefiniteness of \mathbf{T} , we obtain $\mathbf{v}_i^\top [\mathbf{T} \ \mathbf{M}] = \mathbf{0}$, contradicting that $[\mathbf{T} \ \mathbf{M}]$ has full row rank. \square

The following lemma gives the promised extension of (164).

Lemma B.6. *Suppose that \mathbf{T} is a diagonal positive semidefinite $m \times m$ matrix such that $[\mathbf{T} \ \mathbf{A}]$ has full row rank. Let \mathbf{H} (resp. \mathbf{H}^c) be the unconstrained (resp. constrained) Hessian of Υ evaluated at a point \mathbf{z}_0 . Then, λ is an eigenvalue of \mathbf{H}^c iff it is a root of the polynomial*

$$p(\lambda) = [\varepsilon^{m-r}] \text{Det}(\mathbf{H}_\lambda) \text{ where } \mathbf{H}_\lambda = \begin{bmatrix} \varepsilon \mathbf{T} & \mathbf{A} \\ -\mathbf{A}^\top & \mathbf{H} - \lambda \mathbf{I}_n \end{bmatrix}. \quad (166)$$

Further, if \mathbf{H} is invertible, then $\text{Det}(\mathbf{H}^c) = (-1)^r \text{Det}(\mathbf{H}) \frac{[\varepsilon^{m-r}] \text{Det}(\varepsilon \mathbf{T} + \mathbf{A} \mathbf{H}^{-1} \mathbf{A}^\top)}{[\varepsilon^{m-r}] \text{Det}(\varepsilon \mathbf{T} - \mathbf{A} \mathbf{A}^\top)}$.

Proof of Lemma B.6. Let $\mathbf{T} = (t_{i,j})_{i,j \in [m]}$ and $\mathbf{H}_\lambda = (h_{i,j})_{i,j \in [m+n]}$. Let $\mathcal{W} = \binom{[m]}{m-r}$ and for $W \in \mathcal{W}$ let $P_W = \{\sigma \in S_{m+n} \mid \{i \in [m] \mid \sigma(i) = i\} = W\}$. Since \mathbf{T} is diagonal, by Leibniz's formula,

$$p(\lambda) = [\varepsilon^{m-r}] \text{Det}(\mathbf{H}_\lambda) = \sum_{W \in \mathcal{W}} \prod_{i \in W} t_{i,i} \sum_{\sigma \in P_W} \text{sgn}(\sigma) \prod_{i \in [m+n] \setminus W} h_{i,\sigma(i)}. \quad (167)$$

Let $\mathbf{A}_{[m] \setminus W}$ be the $r \times n$ submatrix of \mathbf{A} which is obtain by excluding the rows indexed by W . Identifying permutations in P_W with permutations of $[n+r]$ in the natural way, we obtain

$$\sum_{\sigma \in P_W} \text{sgn}(\sigma) \prod_{i \in [m+n] \setminus W} h_{i,\sigma(i)} = \text{Det} \left(\begin{bmatrix} \mathbf{0} & \mathbf{A}_{[m] \setminus W} \\ -\mathbf{A}_{[m] \setminus W}^\top & \mathbf{H} - \lambda \mathbf{I}_n \end{bmatrix} \right) \equiv q_W(\lambda). \quad (168)$$

If $\mathbf{A}_{[m] \setminus W}$ has row rank $< r$, then $q_W(\lambda)$ is 0. Otherwise, the roots of $q_W(\lambda)$ are the eigenvalues of \mathbf{H}^c , c.f. (164). By (167), this is also the case for $p(\lambda)$, provided it is not identically zero.

To prove that $p(\lambda)$ is nonzero, we prove that the leading coefficient of $p(\lambda)$ is nonzero. Starting from (168) and plugging into (167), the leading coefficient of $p(\lambda)$ can easily be seen to equal

$$[\varepsilon^{m-r}] \text{Det} \left(\begin{bmatrix} \varepsilon \mathbf{T} & \mathbf{A} \\ -\mathbf{A}^\top & -\mathbf{I}_n \end{bmatrix} \right) = [\varepsilon^{m-r}] (-1)^n \text{Det}(\varepsilon \mathbf{T} - \mathbf{A} \mathbf{A}^\top),$$

where in the latter equality we used the Schur complement of the block $-\mathbf{I}_n$. The last expression is non-zero by Lemma B.5.

The determinant of \mathbf{H}^c is the product of its eigenvalues, which in turn equals $(-1)^{n-r} p(0)$ divided by the leading coefficient of $p(\lambda)$. The latter has already been computed. The former, using the Schur complement of the invertible \mathbf{H} , is equal to $[\varepsilon^{m-r}] \text{Det}(\mathbf{H}) \text{Det}(\varepsilon \mathbf{T} - \mathbf{A} \mathbf{H}^{-1} \mathbf{A}^\top)$. This concludes the proof. \square

Finally, we combine the above lemmas to obtain the following.

Lemma B.7. *Let Υ be a function of $\mathbf{z} \in \mathbb{R}^n$ subject to the linear constraints $\mathbf{A} \mathbf{z} = \mathbf{b}$, where $\mathbf{A} \in \mathbb{R}^{m \times n}$ and \mathbf{A} has rank r . Let $(\mathbf{A}_f, \mathbf{A}_{fs})$ specify a full dimensional representation of Υ and let \mathbf{H}^f be the corresponding Hessian of Υ evaluated at a point \mathbf{z}_0 .*

Suppose \mathbf{T} is a positive semidefinite diagonal matrix with dimensions $m \times m$ such that $[\mathbf{T} \ \mathbf{A}]$ has full row rank. Let \mathbf{H} be the unconstrained Hessian of Υ evaluated at \mathbf{z}_0 . If \mathbf{H} is invertible, then

$$\text{Det}(-\mathbf{H}^f) = \frac{L(\mathbf{A}_f, \mathbf{A}, \mathbf{T})}{\text{Det}(\mathbf{A}_{fs})^2} \text{Det}(-\mathbf{H}) [\varepsilon^{m-r}] \text{Det}(\varepsilon \mathbf{T} - \mathbf{A} \mathbf{H}^{-1} \mathbf{A}^\top), \quad (169)$$

where $L(\mathbf{A}_f, \mathbf{A}, \mathbf{T}) = (-1)^r \text{Det}(\mathbf{A}_f \mathbf{A}_f^\top) / [\varepsilon^{m-r}] \text{Det}(\varepsilon \mathbf{T} - \mathbf{A} \mathbf{A}^\top)$.

Proof of Lemma B.7. Just combine Lemmas B.4 and B.6. The minor sign change $-\mathbf{H}^f$ in the statement can easily be accounted by applying the lemmas to the function $-\Upsilon$. \square

The rhs of (169) has two qualitatively different factors: the factor $L(\mathbf{A}_f, \mathbf{A}, \mathbf{T})/\text{Det}(\mathbf{A}_{fs})^2$ depends on the specific full dimensional representation, while the remaining factor is tied to the Hessian of Υ . The technical convenience of Lemma B.7 is dual: first, it gives an explicit formula for $\text{Det}(-\mathbf{H}^f)$ without doing substitutions which would hinder the combinatorial view of the constraints \mathbf{A} ; second, it isolates the deletions of rows of \mathbf{A} in the factor $L(\mathbf{A}_f, \mathbf{A}, \mathbf{T})$ and leaves untouched the more complicated matrix $\mathbf{A}\mathbf{H}^{-1}\mathbf{A}^\top$.

B.1.2 The Computations

In this section, we utilize Lemma B.7 to compute the determinants in Lemma B.3.

Notation: For a vector $\mathbf{z} \in \mathbb{R}^n$ we denote by \mathbf{z}^D the $n \times n$ diagonal matrix $\text{diag}\{z_1, \dots, z_n\}$. For vectors $\mathbf{z}_i \in \mathbb{R}^{m_i}$, $i = 1, \dots, t$ we denote by $[\mathbf{z}_1, \dots, \mathbf{z}_t]^\top$ the $\mathbb{R}^{\sum_i m_i}$ vector which is the concatenation of the vectors $\mathbf{z}_1, \dots, \mathbf{z}_t$. For matrices \mathbf{A} and \mathbf{B} , $\mathbf{A} \otimes \mathbf{B}$ will denote the Kronecker product of \mathbf{A}, \mathbf{B} , while $\mathbf{A} \oplus \mathbf{B}$ is the direct sum of \mathbf{A}, \mathbf{B} , that is, the block diagonal matrix $\text{diag}\{\mathbf{A}, \mathbf{B}\}$. The expression $\oplus_2 \mathbf{A}$ is a shorthand for $\mathbf{A} \oplus \mathbf{A}$. Further, \mathbf{I}_n denotes the identity matrix of dimensions $n \times n$. Finally, $\mathbf{1}_n, \mathbf{0}_n$ denote the all-one and all-zero n -dimensional vector.

To start, the equality constraints in (157) may be written in the form

$$\mathbf{A}_1 [\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}]^\top = \mathbf{0}.$$

The matrix \mathbf{A}_1 has dimensions $(2q + 2) \times (|P_1| + 2q)$ (cf. (156) for the definition of P_1). Note that we exclude from consideration variables x_{ij} which are hard-coded to zero. This is done to ensure that the unconstrained Hessians are invertible, so that Lemma B.7 applies directly. It will be useful to decompose the matrix \mathbf{A}_1 as

$$\mathbf{A}_1 = \begin{bmatrix} \mathbf{A}_{1,\alpha\beta} & \mathbf{0} \\ -\mathbf{I}_{2q} & \mathbf{A}_{1,\mathbf{x}} \end{bmatrix}, \quad (170)$$

where $\mathbf{A}_{1,\alpha\beta}, \mathbf{A}_{1,\mathbf{x}}$ have dimensions $2 \times 2q$ and $2q \times |P_1|$, respectively.

The easiest way to handle the matrix $\mathbf{A}_{1,\mathbf{x}}$ is as the incidence matrix of a bipartite graph $G_{\mathbf{x}}$. First, we introduce some notation: for an undirected graph G , we denote by \mathbf{A}_G the 0,1 incidence matrix of G , by \mathbf{R}_G the adjacency matrix of G , by \mathbf{D}_G the diagonal matrix whose diagonal entries are equal to the degrees of the vertices in G and by $\boldsymbol{\Lambda}_G$ the matrix $\mathbf{D}_G + \mathbf{R}_G$. We will also be interested in the case where the graph G is weighted, in which case we assume that the weights on the edges are given by the diagonal entries of a square diagonal matrix \mathbf{W}_G . We denote by $\mathbf{R}_G^w, \mathbf{D}_G^w, \boldsymbol{\Lambda}_G^w$ the weighted versions of the matrices $\mathbf{R}_G, \mathbf{D}_G, \boldsymbol{\Lambda}_G$. It is well known that

$$\mathbf{A}_G \mathbf{A}_G^\top = \mathbf{D}_G, \quad \mathbf{A}_G \mathbf{W}_G \mathbf{A}_G^\top = \boldsymbol{\Lambda}_G^w. \quad (171)$$

The bipartite graph $G_{\mathbf{x}}$ has vertex bipartition $([q], [q])$ and an edge (i, j) is present iff $(i, j) \in P_1$, that is, $B_{ij} > 0$. Since \mathbf{B} is symmetric and irreducible, $G_{\mathbf{x}}$ is undirected and connected. An edge (i, j) in $G_{\mathbf{x}}$ has weight x_{ij} . In the language of (171), $\mathbf{W}_{G_{\mathbf{x}}} = \mathbf{x}^D$ (the choice of $\mathbf{W}_{G_{\mathbf{x}}}$ will become apparent when we consider the unconstrained Hessian). Applying (171) to the graph $G_{\mathbf{x}}$ is useful to do explicitly in order to decompose the resulting matrices. In particular, since these graphs are undirected and bipartite, we have

$$\boldsymbol{\Lambda}^w(G_{\mathbf{x}}) = \begin{bmatrix} \boldsymbol{\alpha}^D & \mathbf{S}_{\mathbf{x}} \\ \mathbf{S}_{\mathbf{x}}^\top & \boldsymbol{\beta}^D \end{bmatrix}, \quad (172)$$

where $\mathbf{S}_{\mathbf{x}}$ is the $q \times q$ matrix whose (i, j) entry is x_{ij} . Note that the total weight of the edges incident to vertices in $G_{\mathbf{x}}$ (in other words, the diagonal entries of the matrix \mathbf{D}^w) was substituted using (157).

We next state the unconstrained Hessians that will be of interest to us. From Lemma B.3, these are: (i) $\mathbf{H}_{1,\mathbf{x}}$, the Hessian of Υ_1/Δ with respect to \mathbf{x} when α, β are fixed, (ii) \mathbf{H}_1 , the Hessian of Υ_1/Δ with respect to $\alpha, \beta, \mathbf{x}$. These two matrices are all diagonal and by inspection one can check that

$$\begin{aligned} (\mathbf{H}_{1,\mathbf{x}})^{-1} &= -\mathbf{x}^D, \quad (\mathbf{H}_1)^{-1} = \frac{\Delta}{\Delta-1} \alpha^D \oplus \frac{\Delta}{\Delta-1} \beta^D \oplus (\mathbf{H}_{1,\mathbf{x}})^{-1}, \\ \text{Det}(-\mathbf{H}_{1,\mathbf{x}})^{-1} &= \prod_{(i,j) \in P_1} x_{ij}, \quad \text{Det}(-\mathbf{H}_1)^{-1} = \text{Det}(-\mathbf{H}_{1,\mathbf{x}})^{-1} \left(\frac{\Delta}{\Delta-1} \right)^{2q} \prod_{i \in [q]} \alpha_i \prod_{j \in [q]} \beta_j. \end{aligned} \quad (173)$$

We are now ready to evaluate these matrices at a global maximum $(\alpha^*, \beta^*, \mathbf{x}^*)$ of Υ_1 . Henceforth, we will not explicitly use asterisks in the notation with the understanding that the values of all the variables are fixed to their optimal values. We will apply Lemma B.7 to the matrices $\mathbf{H}_{1,\mathbf{x}}^f, \mathbf{H}_1^f$ using the matrices

$$\mathbf{T}_{1,\mathbf{x}} = \alpha^D \oplus \beta^D, \quad \mathbf{T}_1 = \mathbf{I}_2 \oplus \mathbf{0}_{2q \times 2q}, \quad (174)$$

respectively (in (174), $\mathbf{0}_{2q \times 2q}$ denotes the $2q \times 2q$ matrix with all zeros). We first compute the determinants of $\mathbf{M}_{1,\mathbf{x}} := \varepsilon \mathbf{T}_{1,\mathbf{x}} - \mathbf{A}_{1,\mathbf{x}} (\mathbf{H}_{1,\mathbf{x}})^{-1} \mathbf{A}_{1,\mathbf{x}}^\top$, $\mathbf{M}_1 := \varepsilon \mathbf{T}_1 - \mathbf{A}_1 (\mathbf{H}_1)^{-1} \mathbf{A}_1^\top$, which contribute the most interesting factors in Lemma B.7.

We begin with the simplest of these matrices, $\mathbf{M}_{1,\mathbf{x}}$. Note that $\mathbf{A}_{1,\mathbf{x}}$ has rank $2q-1$, so by Lemma B.7 we want to compute $[\varepsilon] \text{Det}(\mathbf{M}_{1,\mathbf{x}})$. Using (171), (173), (174), it is straightforward to check that $\mathbf{M}_{1,\mathbf{x}}$ has the following form

$$\mathbf{M}_{1,\mathbf{x}} = \begin{bmatrix} \alpha^D (\varepsilon \mathbf{I}_q + \mathbf{I}_q) & \mathbf{S}_{\mathbf{x}} \\ \mathbf{S}_{\mathbf{x}}^\top & \beta^D (\varepsilon \mathbf{I}_q + \mathbf{I}_q) \end{bmatrix}, \quad \text{so } \text{Det}(\mathbf{M}_{1,\mathbf{x}}) = \left(\prod_{i \in [q]} \alpha_i \prod_{j \in [q]} \beta_j \right) \text{Det}(\varepsilon \mathbf{I}_q + \mathbf{I}_q + \mathbf{J}), \quad (175)$$

where \mathbf{J} is the matrix in Lemma A.4. Note that in Equation (175), to get the second equality, we did the following operations on $\mathbf{M}_{1,\mathbf{x}}$: for $i = 1, \dots, q$, we divided the i -th row of by $\sqrt{\alpha_i}$, the i -th column by $\sqrt{\alpha_i}$, the $(i+q)$ -th row by $\sqrt{\beta_i}$, the $(i+q)$ -th column by $\sqrt{\beta_i}$. The eigenvalues of the matrix $\varepsilon \mathbf{I}_q + \mathbf{I}_q + \mathbf{J}$ are shifts of the eigenvalues of \mathbf{J} and are given by

$$\varepsilon, \varepsilon + 2, \varepsilon + 1 \pm \lambda_1, \dots, \varepsilon + 1 \pm \lambda_{q-1},$$

c.f., Lemma A.4 for the definition of the λ_i and their properties. We thus obtain

$$[\varepsilon] \text{Det}(\mathbf{M}_{1,\mathbf{x}}) = 2 \prod_{i \in [q]} \alpha_i \prod_{j \in [q]} \beta_j \prod_{i \in [q-1]} (1 - \lambda_i^2). \quad (176)$$

The determinant of the matrix \mathbf{M}_1 is more complicated to compute due to its more intricate block structure, which requires using Schur's complement formula to handle. As in the previous argument, we first write out its block structure and then appropriately normalize the resulting matrix. Here the normalization is slightly more intricate. The analog of (175) is

$$\text{Det}(\mathbf{M}_1) = \text{Det}(\mathbf{H}_1') \prod_{i \in [q]} \alpha_i \prod_{j \in [q]} \beta_j, \quad (177)$$

where

$$\mathbf{H}'_1 := \frac{\Delta}{\Delta-1} \begin{bmatrix} (\varepsilon \frac{\Delta-1}{\Delta} - 1) \mathbf{I}_2 & \mathbf{V} \\ -\frac{\Delta-1}{\Delta} \mathbf{V}^\top & \mathbf{W} \end{bmatrix}, \quad \mathbf{W} := \frac{1}{\Delta-1} \mathbf{I}_{2q} - \mathbf{J}, \quad \mathbf{V}^\top := \begin{bmatrix} \sqrt{\alpha} & \mathbf{0}_q \\ \mathbf{0}_q & \sqrt{\beta} \end{bmatrix}. \quad (178)$$

(\mathbf{J} is the matrix in Lemma A.4; note also that \mathbf{V} has dimension $2 \times 2q$ and its two rows are given by $\sqrt{\alpha_1}, \dots, \sqrt{\alpha_q}, 0, \dots, 0$ and $0, \dots, 0, \sqrt{\beta_1}, \dots, \sqrt{\beta_q}$ where each row has q zeros.) Equation (177) is obtained by performing the following operations on \mathbf{M}_1 : for $i, j = 1, \dots, q$, divide the $2+i$ row by $\sqrt{\alpha_i}$ and the $2+q+j$ row by $\sqrt{\beta_j}$; and the same operations on columns.

In light of (177), it suffices to compute $\text{Det}(\mathbf{H}'_1)$. To do this, we proceed by taking the Schur complement of the matrix \mathbf{W} . The spectrum of \mathbf{W} is

$$t \pm 1, t \pm \lambda_1, \dots, t \pm \lambda_{q-1},$$

where $t = 1/(\Delta-1)$. It follows that

$$\text{Det}(\mathbf{W}) = -\frac{\Delta(\Delta-2)}{(\Delta-1)^{2q}} \prod_{i \in [q-1]} (1 - (\Delta-1)^2 \lambda_i^2), \quad (179)$$

where $\lambda_i, i \in [q-1]$ are as in Lemma A.4. Note that \mathbf{W} is invertible, since the λ_i 's are non-negative and $\max \lambda_i < \frac{1}{\Delta-1}$. By taking the Schur complement of the matrix \mathbf{W} in \mathbf{H}'_1 , we obtain

$$\text{Det}(\mathbf{H}'_1) = \left(\frac{\Delta}{\Delta-1} \right)^2 \text{Det}(\mathbf{W}) \text{Det} \left(\varepsilon \frac{\Delta-1}{\Delta} \mathbf{I}_2 + \mathbf{Z} \right), \quad \text{where } \mathbf{Z} = -\mathbf{I}_2 + \frac{\Delta}{\Delta-1} \mathbf{V} \mathbf{W}^{-1} \mathbf{V}^\top. \quad (180)$$

We are left with the evaluation of $\text{Det}(\varepsilon \frac{\Delta-1}{\Delta} \mathbf{I}_2 + \mathbf{Z})$. The complication here is the nontrivial inverse of \mathbf{W} appearing in the formulation of \mathbf{Z} . The key idea to circumvent the computation of \mathbf{W}^{-1} is the following equality

$$\mathbf{V} \mathbf{W} = \left(\frac{1}{\Delta-1} \mathbf{I}_2 - \mathbf{J}' \right) \mathbf{V}, \quad \text{where } \mathbf{J}' = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}.$$

The equality can be checked using the relations $\sum_j x_{ij} = \alpha_i$ and $\sum_i x_{ij} = \beta_j$. Using that $\mathbf{V} \mathbf{V}^\top = \mathbf{I}_2$, we obtain

$$\mathbf{Z} = -\mathbf{I}_2 + \frac{\Delta}{\Delta-1} \left(\frac{1}{\Delta-1} \mathbf{I}_2 - \mathbf{J}' \right)^{-1} \mathbf{V} \mathbf{V}^\top = -\frac{\Delta-1}{\Delta-2} \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}, \quad (181)$$

We thus obtain

$$[\varepsilon] \text{Det} \left(\varepsilon \frac{\Delta-1}{\Delta} \mathbf{I}_2 + \mathbf{Z} \right) = -\frac{2(\Delta-1)^2}{\Delta(\Delta-2)}. \quad (182)$$

Plugging (179) and (182) in (180), we obtain

$$[\varepsilon] \text{Det}(\mathbf{H}'_1) = \frac{2\Delta^2}{(\Delta-1)^{2q}} \prod_{i \in [q-1]} (1 - (\Delta-1)^2 \lambda_i^2).$$

Using this and (177), we obtain

$$[\varepsilon] \text{Det}(\varepsilon \mathbf{T}_1 - \mathbf{A}_1 (\mathbf{H}_1)^{-1} \mathbf{A}_1^\top) = \frac{2\Delta^2}{(\Delta-1)^{2q}} \prod_{i \in [q]} \alpha_i \prod_{j \in [q]} \beta_j \prod_{i \in [q-1]} (1 - (\Delta-1)^2 \lambda_i^2). \quad (183)$$

Equations (173), (176), (183) deal with the factors in Lemma B.7 which are tied to the Hessians of the functions. While these contribute the most interesting factors, some care is needed to deal with the remaining factors. This is accomplished in the following lemma, which is given in the end of this section.

Lemma B.8. Let $((\mathbf{A}_{1,\mathbf{x}})_f, (\mathbf{A}_{1,\mathbf{x}})_{fs}), ((\mathbf{A}_1)_f, (\mathbf{A}_1)_{fs})$ specify arbitrary full dimensional representations of the spaces $\mathbf{A}_{1,\mathbf{x}}\mathbf{x} = \mathbf{0}$, $\mathbf{A}_1[\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}]^\top = \mathbf{0}$, respectively. Then:

$$\text{Det}((\mathbf{A}_{1,\mathbf{x}})_{fs})^2 = \text{Det}((\mathbf{A}_1)_{fs})^2 = 1, \quad (184)$$

$$L((\mathbf{A}_{1,\mathbf{x}})_f, \mathbf{A}_{1,\mathbf{x}}, \mathbf{T}_{1,\mathbf{x}}) = 1/2, \quad L((\mathbf{A}_1)_f, \mathbf{A}_1, \mathbf{T}_1) = 1/2, \quad (185)$$

where $\mathbf{T}_{1,\mathbf{x}}, \mathbf{T}_1$ are given by (174) and the quantities in (185) are defined in Lemma B.7.

We are now ready to finish the proof of Lemma A.9.

Proof of Lemma A.9. Apply Lemma B.7 two times to unravel the determinants appearing in Lemma B.3. Each of the resulting quantities has been computed and appears in one of (173), (176), (183) or Lemma B.8. Straightforward substitutions yield

$$\begin{aligned} \text{Det}(-\mathbf{H}_{1,\mathbf{x}}^f) &= \left(\prod_{(i,j) \in P_1} x_{ij} \right)^{-1} \prod_{i \in [q]} \alpha_i \prod_{j \in [q]} \beta_j \prod_{i \in [q-1]} (1 - \lambda_i)^2, \\ \text{Det}(-\mathbf{H}_1^f) &= \frac{1}{\Delta^{2(q-1)}} \left(\prod_{(i,j) \in P_1} x_{ij} \right)^{-1} \prod_{i \in [q]} \alpha_i \prod_{j \in [q]} \beta_j \prod_{i \in [q-1]} (1 - \lambda_i^2). \end{aligned}$$

Thus, Lemma B.3 gives

$$\lim_{n \rightarrow \infty} \frac{\mathbf{E}_{\mathcal{G}}[Z_G^{\mathbf{p}}]}{e^{n\Upsilon_1(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x})}} = \prod_{i=1}^{q-1} (1 - (\Delta - 1)^2 \lambda_i^2)^{-1/2} \prod_{i=1}^{q-1} (1 - \lambda_i^2)^{-(\Delta-1)/2}. \quad (186)$$

Note that in the last expression, only the eigenvalues of the matrix \mathbf{J} (different from 1) in Lemma A.4 appear. For the asymptotics of $\mathbf{E}_{\mathcal{G}}[(Z_G^{\mathbf{p}})^2]$, by Lemma B.2, it suffices to consider the spin system with interaction matrix $\mathbf{B} \otimes \mathbf{B}$ (and dominant phase $\boldsymbol{\gamma} = \boldsymbol{\alpha} \otimes \boldsymbol{\alpha}, \boldsymbol{\delta} = \boldsymbol{\beta} \otimes \boldsymbol{\beta}, \mathbf{y} = \mathbf{x} \otimes \mathbf{x}$). The eigenvalues (different from 1) of the matrix $\mathbf{J} \otimes \mathbf{J}$ are λ_i for $i \in [q-1]$ and $\lambda_i \lambda_j$ for $i, j \in [q-1]$. Thus, we obtain

$$\lim_{n \rightarrow \infty} \frac{\mathbf{E}_{\mathcal{G}}[(Z_G^{\mathbf{p}})^2]}{e^{n\Upsilon_2(\boldsymbol{\gamma}, \boldsymbol{\delta}, \mathbf{y})}} = C \cdot \prod_{i=1}^{q-1} (1 - (\Delta - 1)^2 \lambda_i^2)^{-1/2} \prod_{i=1}^{q-1} (1 - \lambda_i^2)^{-(\Delta-1)/2}, \quad (187)$$

where C is the constant in the statement of the lemma. Combining (186) and (187) with Lemma B.1 yields the result. \square

Finally, we give the proof of Lemma B.8.

Proof of Lemma B.8. We first prove (184). Since $\mathbf{A}_{1,\mathbf{x}}$ is the incidence matrix of the bipartite graph $G_{\mathbf{x}}$, it is a totally unimodular matrix. By the way full dimensional representations are chosen, the matrix $(\mathbf{A}_{1,\mathbf{x}})_{fs}$ is invertible and hence its determinant squared equals 1. For $(\mathbf{A}_1)_{fs}$, observe that $(\mathbf{A}_1)_{fs}$ has the block decomposition

$$(\mathbf{A}_1)_{fs} = \begin{bmatrix} (\mathbf{A}_{1,\boldsymbol{\alpha}\boldsymbol{\beta}})_{fs} & \mathbf{0} \\ -\mathbf{I} & (\mathbf{A}_{1,\mathbf{x}})_{fs} \end{bmatrix}, \text{ so that } \text{Det}((\mathbf{A}_1)_{fs}) = \text{Det}((\mathbf{A}_{1,\boldsymbol{\alpha}\boldsymbol{\beta}})_{fs}) \text{Det}((\mathbf{A}_{1,\mathbf{x}})_{fs}).$$

Since $\mathbf{A}_{1,\boldsymbol{\alpha}\boldsymbol{\beta}}, \mathbf{A}_{1,\mathbf{x}}$ are totally unimodular, any invertible submatrix of them has determinant ± 1 . This concludes the proof of (184).

We next turn to (185). We begin with $L((\mathbf{A}_{1,\mathbf{x}})_f, \mathbf{A}_{1,\mathbf{x}}, \mathbf{T}_{1,\mathbf{x}})$. The argument is closely related to the proof of Kirchoff's Matrix-Tree Theorem, but is written in a way that it easily extends to the more complicated $L((\mathbf{A}_1)_f, \mathbf{A}_1, \mathbf{T}_1)$.

Denote by μ_1, \dots, μ_{2q-1} the non-zero eigenvalues of $\mathbf{A}_{1,\mathbf{x}} \mathbf{A}_{1,\mathbf{x}}^\top$; there are exactly $2q-1$ of those since $G_{\mathbf{x}}$ is a connected bipartite graph. Moreover, $\mathbf{v}_0^\top = \frac{1}{\sqrt{2q}}[-\mathbf{1}_q \ \mathbf{1}_q]$ is the unit eigenvector of $\mathbf{A}_{1,\mathbf{x}} \mathbf{A}_{1,\mathbf{x}}^\top$ with eigenvalue 0. We claim that

$$[\varepsilon] \text{Det}(\varepsilon \mathbf{T}_{1,\mathbf{x}} - \mathbf{A}_{1,\mathbf{x}} \mathbf{A}_{1,\mathbf{x}}^\top) = -\frac{\prod_{i \in [2q-1]} \mu_i}{q}, \quad \text{Det}((\mathbf{A}_{1,\mathbf{x}})_f (\mathbf{A}_{1,\mathbf{x}})_f^\top) = \frac{\prod_{i \in [2q-1]} \mu_i}{2q}, \quad (188)$$

which yields that $L((\mathbf{A}_{1,\mathbf{x}})_f, \mathbf{A}_{1,\mathbf{x}}, \mathbf{T}_{1,\mathbf{x}}) = 1/2$, as wanted. The first equality is a direct application of Lemma B.5, after observing that $\mathbf{v}_0^\top \mathbf{T}_{1,\mathbf{x}} \mathbf{v}_0 = 1/q$. The second can be proved as follows. The matrix $(\mathbf{A}_{1,\mathbf{x}})_f (\mathbf{A}_{1,\mathbf{x}})_f^\top$ is a principal minor of $\mathbf{A}_{1,\mathbf{x}} \mathbf{A}_{1,\mathbf{x}}^\top$, the specific principal minor is clearly determined by which row of \mathbf{A}_1 we chose to delete to obtain $(\mathbf{A}_1)_f$. Since $\mathbf{A}_{1,\mathbf{x}} \mathbf{A}_{1,\mathbf{x}}^\top$ has exactly one zero eigenvalue, we have

$$\prod_{i \in [2q-1]} \mu_i = \sum_{W \in \binom{[2q]}{2q-1}} \text{Det}((\mathbf{A}_{1,\mathbf{x}})_W (\mathbf{A}_{1,\mathbf{x}})_W^\top), \quad (189)$$

where $(\mathbf{A}_{1,\mathbf{x}})_W$ is the submatrix of $\mathbf{A}_{1,\mathbf{x}}$ induced by the rows indexed with W . It is easily checked that for any $W, W' \in \binom{[2q]}{2q-1}$, there exists a unitary matrix \mathbf{P} such that $(\mathbf{A}_{1,\mathbf{x}})_W = \mathbf{P}(\mathbf{A}_{1,\mathbf{x}})_{W'}$, so that all summands in (189) are equal. Indeed, since $\mathbf{A}_{1,\mathbf{x}}$ corresponds to the incidence matrix of a bipartite graph, the sum of the first q rows (as vectors) equals the sum of the last q rows. It follows that any row of $\mathbf{A}_{1,\mathbf{x}}$ can be expressed as a $\{1, -1\}$ linear combination of the remaining rows, which easily yields the existence of \mathbf{P} with the desired properties. Hence, for any $(\mathbf{A}_{1,\mathbf{x}})_f$ as in the statement of the Lemma, the second equality in (188) holds as well.

We finally give a proof sketch for $L((\mathbf{A}_1)_f, \mathbf{A}_1, \mathbf{T}_1) = 1/2$. The matrix $\mathbf{A}_1 \mathbf{A}_1^\top$ has zero as an eigenvalue by multiplicity one. Denote by $\sigma_1, \dots, \sigma_{2q+1}$ the non-zero eigenvalues of $\mathbf{A}_1 \mathbf{A}_1^\top$. By looking at the space $\mathbf{z} \mathbf{A}_1 = \mathbf{0}$, it is easy to see that $\mathbf{v}_1 = \frac{1}{\sqrt{2(q+1)}}[-1, 1, -\mathbf{1}_q, \mathbf{1}_q]^\top$ is a unit length eigenvector for the eigenvalue 0. Moreover, the analog of (189) is

$$\prod_{i \in [2q+1]} \sigma_i = \sum_{W \in \binom{[2q+2]}{2q+1}} \text{Det}((\mathbf{A}_1)_W (\mathbf{A}_1)_W^\top). \quad (190)$$

Hence, the equality $L((\mathbf{A}_1)_f, \mathbf{A}_1, \mathbf{T}_1) = 1/2$ is obtained by the following analog of (188)

$$[\varepsilon] \text{Det}(\varepsilon \mathbf{T}_1 - \mathbf{A}_1 \mathbf{A}_1^\top) = -\frac{\prod_{i \in [2q+1]} \sigma_i}{q+1}, \quad \text{Det}((\mathbf{A}_1)_f (\mathbf{A}_1)_f^\top) = \frac{\prod_{i \in [2q+2]} \sigma_i}{2(q+1)}.$$

□

C Uniqueness of semi-translation invariant measures (Antiferromagnetic Potts)

In this section, we prove Lemma 7.3. As noted earlier, the proof extends the respective argument in [BW02] for colorings in the antiferromagnetic Potts model setting. The technical details, due to the presence of the extra parameter B , are relatively more intricate.

Proof of Lemma 7.3. W.l.o.g. we may assume that the scaling factors in (105) are equal to 1. We may also assume that $R_1 \geq \dots \geq R_q$. Then the equations easily imply $C_1 \leq \dots \leq C_q$. Define

$$\alpha = \frac{R_1}{R_q}, \quad \beta = \frac{R_1 + \dots + R_{q-1}}{(q-1)R_q}, \quad S = R_1 + \dots + R_{q-1}.$$

We clearly have $\alpha \geq \beta \geq 1$, and we may assume for the sake of contradiction that $\beta > 1$. Note that

$$\begin{aligned} \alpha^{1/d} &= \left(\frac{R_1}{R_q} \right)^{1/d} = 1 + \frac{(1-B)(C_q - C_1)}{C_1 + \dots + C_{q-1} + BC_q} \\ C_q &= (R_1 + \dots + R_{q-1} + BR_q)^d = [(q-1)\beta + B]^d R_q^d \\ C_1 &= (BR_1 + R_2 + \dots + R_q)^d = [(q-1)\beta + 1 - (1-B)\alpha]^d R_q^d \end{aligned}$$

Moreover, by Holder's inequality or otherwise, we have

$$\begin{aligned} C_1 + \dots + C_{q-1} + BC_q &= \sum_{i=1}^{q-1} [S + R_q - (1-B)R_i]^d + B(S + BR_q)^d \\ &\geq (q-1) \left[\frac{q-2+B}{q-1} S + (q-1)R_q \right]^d + B(S + BR_q)^d \\ &= (q-1) [(q-2+B)\beta + 1]^d R_q^d + B[(q-1)\beta + B]^d R_q^d. \end{aligned}$$

Thus, we obtain that every solution must satisfy

$$\begin{aligned} \alpha^{1/d} &\leq 1 + \frac{(1-B) \left\{ [(q-1)\beta + B]^d - [1 - (1-B)\alpha + (q-1)\beta]^d \right\}}{(q-1) [(q-2+B)\beta + 1]^d + B[(q-1)\beta + B]^d} \iff \\ 0 &\leq 1 - \alpha^{1/d} + \frac{(1-B) \left[1 - \left(1 - \frac{(1-B)(\alpha-1)}{(q-1)\beta+B} \right)^d \right]}{(q-1) \left[1 - \frac{(1-B)(\beta-1)}{(q-1)\beta+B} \right]^d + B} =: f(\alpha, \beta, B). \end{aligned}$$

To obtain a contradiction, our goal is to prove that for q and B as in the statement of the lemma, when $(q-1)\beta > \alpha \geq \beta > 1$, it holds that $f(\alpha, \beta, B) < 0$.

It is easy to see that f is decreasing in B . This immediately yields the lemma for $q \geq \Delta$: it holds that $f(\alpha, \beta, B) \leq f(\alpha, \beta, 0) < 0$, since the last inequality was proved by [BW02]. For $q \leq d$ and $B \geq \frac{d+1-q}{d+1} := B_c$, this yields

$$f(\alpha, \beta, B) \leq f(\alpha, \beta, B_c) =: g(\alpha, \beta).$$

We first prove that $g(\alpha, \beta) \leq g(\beta, \beta)$. For $q = 2$ there is nothing to prove. Hence we may assume that $d \geq q \geq 3$. Clearly it suffices to prove that g is decreasing in α . This requires a fair bit of work, so we state it as a Lemma to prove later.

Lemma C.1. *For $d \geq q \geq 3$ and $B_c = \frac{d+1-q}{d+1}$, the function $g(\alpha, \beta)$ is decreasing in α for $\alpha \geq \beta > 1$.*

We finish the proof by showing that for $\beta \geq 1$, it holds that $g(\beta, \beta) \leq 0$ with equality iff $\beta = 1$. After massaging the inequality, this reduces to

$$1 \leq \left[1 - \frac{(1-B_c)(\beta-1)}{(q-1)\beta+B_c} \right]^d \left[(q-1) (\beta^{1/d} - 1) + 1 - B_c \right] + B_c \beta^{1/d} =: h(\beta)$$

Note that the inequality holds at equality for $\beta = 1$, so it suffices to prove $h'(\beta) > 0$ for $\beta > 1$, which is the assertion of the next lemma.

Lemma C.2. For $d \geq q \geq 2$ and $B_c = \frac{d+1-q}{d+1}$, the function $h(\beta)$ is increasing for $\beta \geq 1$.

Modulo the proofs of Lemmas C.1 and C.2, which are given below, the proof is complete. \square

Proof of Lemma C.1. We compute

$$\frac{\partial g}{\partial \alpha} = -\frac{1}{d}\alpha^{-(d-1)/d} + \frac{(1-B_c)^2}{(q-1)\beta + B_c} \cdot \frac{d \left[1 - \frac{(1-B_c)(\alpha-1)}{(q-1)\beta + B_c}\right]^{d-1}}{(q-1) \left[1 - \frac{(1-B_c)(\beta-1)}{(q-1)\beta + B_c}\right]^d + B_c}$$

Let $F(x) = x \left[1 - \frac{(1-B_c)(x-1)}{(q-1)\beta + B_c}\right]^d$ for $x \in [\beta, (q-1)\beta]$. Straightforward manipulations show that $\frac{\partial g}{\partial \alpha} < 0$ is equivalent to

$$d^2(1-B_c)^2 F(\alpha)^{(d-1)/d} \leq [(q-1)\beta + B_c] \left[(q-1) \left(1 - \frac{(1-B_c)(\beta-1)}{(q-1)\beta + B_c}\right)^d + B_c \right]. \quad (191)$$

We prove that $F(x)$ is decreasing in $[\beta, (q-1)\beta]$. It is simple to check that

$$F'(x) = \left[1 - \frac{(1-B_c)(x-1)}{(q-1)\beta + B_c}\right]^{d-1} \frac{(q-1)\beta + 1 - (d+1)(1-B_c)x}{(q-1)\beta + B_c}.$$

For $x \in [\beta, (q-1)\beta]$, we have $(d+1)(1-B_c)x = qx = (q-1)x + x > (q-1)\beta + 1$, where in the last inequality we used that $\beta > 1$. It follows that $F(x)$ is indeed decreasing and thus $F(\alpha) \leq F(\beta)$.

To prove (191), it thus suffices to argue that for $\beta > 1$ it holds

$$d^2(1-B_c)^2 F(\beta)^{(d-1)/d} \leq [(q-1)\beta + B_c] \left[(q-1) \left(1 - \frac{(1-B_c)(\beta-1)}{(q-1)\beta + B_c}\right)^d + B_c \right]. \quad (192)$$

Note that $q-1+B_c = d(1-B_c)$ so that the inequality is tight for $\beta = 1$. By the weighted AM-GM inequality on A^d and 1 with weights $(q-1)$ and B_c respectively, we obtain

$$(q-1)A^d + B_c \geq (q-1+B_c)A^{d(q-1)/(q-1+B_c)} = d(1-B_c)A^{(q-1)(d+1)/q}.$$

We use this for $A = \left(1 - \frac{(1-B_c)(\beta-1)}{(q-1)\beta + B_c}\right)^d$ so that, after simplifications, it suffices to show that

$$d(1-B_c)\beta^{(d-1)/d} \leq [(q-1)\beta + B_c] \left[1 - \frac{(1-B_c)(\beta-1)}{(q-1)\beta + B_c}\right]^{-(d+1-2q)/q}.$$

This can further be massaged into

$$G(\beta) := \beta^{(d-1)/d} [(q-1)\beta + B_c]^{-(d+1-q)/q} [(q-2+B_c)\beta + 1]^{(d+1-2q)/q} \leq \frac{1}{d(1-B_c)}.$$

Once again, note that the inequality holds at equality for $\beta = 1$, so it suffices to prove that $G'(\beta) < 0$ for $\beta > 1$. This has nothing special, apart from tedious, but otherwise straightforward, calculations. We include the details briefly. Differentiating $\ln G(\beta)$, we obtain

$$\frac{G'(\beta)}{G(\beta)} = \frac{(d-1)}{d\beta} - \frac{(d+1-q)(q-1)}{q[(q-1)\beta + B_c]} + \frac{(d+1-2q)(q-2+B_c)}{q[(q-2+B_c)\beta + 1]}.$$

By clearing denominators, it suffices to check that the following second order polynomial $p(\beta)$ is negative whenever $\beta > 1$:

$$p(\beta) := (d-1)q[(q-1)\beta + B_c] [(q-2+B_c)\beta + 1] - (d+1-q)(q-1)d\beta[(q-1)\beta + B_c] \\ + (d+1-2q)(q-2+B_c)\beta[(q-1)\beta + B_c].$$

Using again that $q-1+B_c = d(1-B_c)$ it is easy to verify that $p(1) = 0$. The factorization of $p(\beta)$ (using the value of B_c) is given by

$$p(\beta) = -\frac{q(\beta-1)[\beta(d(q-1)^2 - (q-1)) + d(d-q) + q-1]}{d+1},$$

which is obviously negative for $\beta > 1$, whenever $d \geq q \geq 2$. □

Proof of Lemma C.2. We compute

$$h'(\beta) = \frac{1}{d}\beta^{-(d-1)/d} \left[(q-1) \left(1 - \frac{(1-B_c)(\beta-1)}{(q-1)\beta + B_c} \right)^d + B_c \right] \\ - d \left[1 - \frac{(1-B_c)(\beta-1)}{(q-1)\beta + B_c} \right]^{d-1} \frac{(1-B_c)(q-1+B_c)[(q-1)\beta^{1/d} - (q-2+B_c)]}{[(q-1)\beta + B_c]^2}.$$

Thus, to prove $h'(\beta) > 0$ it suffices to check (using $q-1+B_c = d(1-B_c)$ and the function F defined in Lemma C.1)

$$d^3(1-B_c)^2 F(\beta)^{(d-1)/d} \leq \\ \leq \frac{[(q-1)\beta + B_c]^2}{(q-1)\beta^{1/d} - (q-2+B_c)} \left[(q-1) \left(1 - \frac{(1-B_c)(\beta-1)}{(q-1)\beta + B_c} \right)^d + B_c \right],$$

This is similar to (192) and in fact follows from (192), once we prove that

$$\frac{(q-1)\beta^{1/d} - (q-2+B_c)}{(q-1)\beta + B_c} \leq \frac{1}{d}.$$

To see the last inequality, observe that $\beta + d - 1 \geq d\beta^{1/d}$ as a consequence of the weighted AM-GM inequality (or otherwise). Hence,

$$\frac{(q-1)\beta^{1/d} - (q-2+B_c)}{(q-1)\beta + B_c} \leq \frac{(q-1)\beta + (d-1)(q-1) - d(q-2+B_c)}{d[(q-1)\beta + B_c]} = \frac{1}{d},$$

completing the proof. □