



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

#BIS-hardness for 2-spin systems on bipartite bounded degree graphs in the tree non-uniqueness region

Citation for published version:

Cai, J-Y, Galanis, A, Goldberg, LA, Guo, H, Jerrum, M, Stefankovic, D & Vigoda, E 2016, '#BIS-hardness for 2-spin systems on bipartite bounded degree graphs in the tree non-uniqueness region', *Journal of Computer and System Sciences*, vol. 82, no. 5, pp. 690-711. <https://doi.org/10.1016/j.jcss.2015.11.009>

Digital Object Identifier (DOI):

[10.1016/j.jcss.2015.11.009](https://doi.org/10.1016/j.jcss.2015.11.009)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Journal of Computer and System Sciences

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



#BIS-Hardness for 2-Spin Systems on Bipartite Bounded Degree Graphs in the Tree Non-Uniqueness Region

Jin-Yi Cai^{a,1}, Andreas Galanis^{b,2,*}, Leslie Ann Goldberg^{b,2}, Heng Guo^{a,1}, Mark Jerrum^c, Daniel Štefankovič^{d,3}, Eric Vigoda^{e,4}

^a*Department of Computer Sciences, University of Wisconsin-Madison, Madison, WI 53706, USA.*

^b*Department of Computer Science, University of Oxford, Wolfson Building, Parks Road, Oxford OX1 3QD, UK.*

^c*School of Mathematical Sciences, Queen Mary, University of London, Mile End Road, London E1 4NS, UK.*

^d*Department of Computer Science, University of Rochester, Rochester, NY 14627, USA.*

^e*School of Computer Science, Georgia Institute of Technology, Atlanta, GA 30332, USA.*

Abstract

Counting independent sets on bipartite graphs (#BIS) is considered a canonical counting problem of intermediate approximation complexity. It is conjectured that #BIS neither has an FPRAS nor is as hard as #SAT to approximate. We study #BIS in the general framework of two-state spin systems on bipartite graphs. We define two notions, nearly-independent phase-correlated spins and unary symmetry breaking. We prove that it is #BIS-hard to approximate the partition function of any 2-spin system on bipartite graphs supporting these two notions. As a consequence, we classify the complexity of approximating the partition function of antiferromagnetic 2-spin systems on bounded-degree bipartite graphs.

Keywords: Spin systems, approximate counting, complexity, #BIS-hardness, phase transition.
2010 MSC: 68Q25, 68Q87

1. Introduction

There has been great progress in classifying the complexity of counting problems recently. One important success is for counting constraint satisfaction problems (#CSP), where a sweeping complexity dichotomy is proved [1, 10, 3]. While the landscape of exact counting becomes clearer, the complexity of approximate counting remains mysterious. Two typical classes of problems have been identified: 1) those that have a fully polynomial-time randomized approximation scheme (FPRAS), and 2) those that are #SAT-hard with respect to approximation preserving reductions (AP-reductions) [8]. If $\text{NP} \neq \text{RP}$

*Corresponding author.

Email addresses: `jyc@cs.wisc.edu` (Jin-Yi Cai), `andreas.galanis@cs.ox.ac.uk` (Andreas Galanis), `leslie.goldberg@cs.ox.ac.uk` (Leslie Ann Goldberg), `hguo@cs.wisc.edu` (Heng Guo), `m.jerrum@qmul.ac.uk` (Mark Jerrum), `stefanko@cs.rochester.edu` (Daniel Štefankovič), `vigoda@cc.gatech.edu` (Eric Vigoda)

¹Research supported by NSF grant CCF-1217549. Heng Guo is also supported by a 2013 Simons award for graduate students in theoretical computer science.

²The 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.

³Research supported by NSF grant CCF-1318374.

⁴Research supported by NSF grant CCF-1217458.

then #SAT admits no FPRAS⁵ [31], and therefore neither does any problem in the second class. These two classes are analogous to P-time tractable vs. NP-hard decision or optimization problems.

Interestingly, in approximate counting, there has emerged a third distinct class of natural problems, which seems to be of intermediate complexity. It is conjectured [8] that the problems in this class do not have an FPRAS but that they are not as hard as #SAT to approximate. A canonical problem in this class has been identified, which is to count the number of independent sets in a bipartite graph (#BIS). Despite many attempts, nobody has found an FPRAS for #BIS or an AP-reduction from #SAT to #BIS. The conjecture is that neither exists. Mossel et al. [26] showed that the Gibbs sampler for sampling independent sets in bipartite graphs mixes slowly even on bipartite graphs of degree at most 6. Another interesting attempted Markov Chain for #BIS by Ge and Štefankovič [14] was also shown later to be slowly mixing by Goldberg and Jerrum [17].

#BIS plays an important role in classifying counting problems with respect to approximation. A trichotomy theorem is shown for the complexity of approximately solving unweighted Boolean counting CSPs, where in addition to problems that are solvable by FPRASes and those that are AP-reducible from #SAT, there is the intermediate class of problems which are equivalent to #BIS [9]. Many counting problems are shown to be #BIS-hard and hence are conjectured to have no FPRAS [2, 7], including estimating the partition function of the ferromagnetic Potts model [18]. Moreover, under AP-reductions #BIS is complete in a logically defined class of problems, called #RHP₁, to which an increasing variety of problems have been shown to belong. Other typical complete problems in #RHP₁ include counting the number of downsets in a partially ordered set [8] and computing the partition function of the ferromagnetic Ising model with local external fields [16].

The problem of counting independent sets (#IS) can be viewed as a special case in the general framework of spin systems, which originated from statistical physics to model interactions between neighbors on graphs. In this paper, we focus on two-state spin systems. In general such a system is parameterized by edge weights $\beta, \gamma \geq 0$ and a vertex weight $\lambda > 0$. An instance is a graph $G = (V, E)$. A configuration σ is a mapping $\sigma : V \rightarrow \{0, 1\}$ from vertices to (two) spins. The weight $w(\sigma)$ of a configuration σ is given by

$$w(\sigma) = \beta^{m_0(\sigma)} \gamma^{m_1(\sigma)} \lambda^{n_1(\sigma)} \quad (1)$$

where $m_0(\sigma)$ is the number of (0,0) edges given by the configuration σ , $m_1(\sigma)$ is the number of (1,1) edges, and $n_1(\sigma)$ is the number of vertices assigned 1. We are interested in computing the partition function, which is defined by

$$Z_G(\beta, \gamma, \lambda) = \sum_{\sigma: V \rightarrow \{0,1\}} w(\sigma). \quad (2)$$

The partition function is the normalizing factor of the Gibbs distribution, which is the distribution in which a configuration σ is drawn with probability $\Pr_{G;\beta,\gamma,\lambda}(\sigma) = \frac{w(\sigma)}{Z_G(\beta,\gamma,\lambda)}$. The spin system is called *ferromagnetic* if $\beta\gamma > 1$ and *antiferromagnetic* if $\beta\gamma < 1$. In particular, when $\beta = \gamma$, such a system is the famous *Ising model*, and when $\beta = 1$ and $\gamma = 0$, it is the *hard-core gas model*, the partition function of which counts independent sets when $\lambda = 1$. The external field λ is typically referred to as the activity or fugacity of the hard-core model.

Approximating the partition function of the hard-core model is especially well studied. We now know that the complexity transition from easy to hard corresponds exactly to the uniqueness of the Gibbs measure in infinite $(\Delta - 1)$ -ary trees \mathbb{T}_Δ (for details of these notions, see [15]). Notice that $(\Delta - 1)$ -ary

⁵In fact, Zuckerman proves a stronger result — there is no FPRAS for the logarithm of the number of satisfying assignments unless NP=RP.

trees are graphs of maximum degree Δ , hence our use of the notation \mathbb{T}_Δ . On the algorithmic side, Weitz presented a fully polynomial-time approximation scheme (FPTAS) for the hard-core gas model on graphs of maximum degree Δ when uniqueness holds [30]. On the other hand, Sly showed that the approximation problem is $\#\text{SAT}$ -hard for a small interval beyond the uniqueness threshold [28]. Building on their work, it is now established that for all antiferromagnetic 2-spin systems there is an FPTAS for graphs of maximum degree Δ up to the uniqueness threshold [24] (see also [23, 27]), whereas non-uniqueness implies $\#\text{SAT}$ -hardness under AP-reductions on Δ -regular graphs [29] (see also [4, 12]). The only place that remains unclear is exactly at the uniqueness threshold.

A key feature of spin systems in the antiferromagnetic non-uniqueness region is the ability to support a gadget with many vertices whose spins are highly correlated with the phase of the gadget (which is either $+$ or $-$), but are nearly independent among themselves conditioned on that phase. Such a gadget was used by Sly [28] to show inapproximability of the partition function of the hard-core model when λ is just above the uniqueness threshold. A different gadget with similar properties was used by Goldberg et al. [19] to show inapproximability on a planar graph when λ is much larger. We abstract this notion of nearly-independent phase-correlated spins. It is this feature that enables us to reduce from $\#\text{SAT}$ to approximating the partition function of antiferromagnetic two-spin systems in the non-uniqueness region.

Restricted to bipartite graphs, it appears that supporting nearly-independent phase-correlated spins alone is not enough to imply $\#\text{BIS}$ -hardness. It was shown that Sly’s gadget is applicable to the antiferromagnetic Ising model without an external field by Galanis et al. [12]. However, such a system has an FPRAS on bipartite graphs. The reason is that this system is perfectly symmetric on bipartite graphs and therefore can be translated into a ferromagnetic Ising system, whose partition function can be approximated using the FPRAS of Jerrum and Sinclair [21] (see Corollary 13 in Section 6 for details). To get around this perfectly symmetric case, we introduce the second key concept called unary symmetry breaking. Unary symmetry breaking is about the availability of non-trivial “unary weights” which can be built in the model. The availability of unary weights is relevant to counting complexity, for example in the context of counting CSPs. Unary symmetry breaking simply means that some (any) non-trivial unary weight can be built. Formal definitions of the two notions – nearly-independent phase-correlated spins and unary symmetry breaking – can be found in Section 3. Our main technical theorem is the following.

Theorem 1. *Suppose a tuple of parameters $(\beta, \gamma, \lambda, \Delta)$ with $\beta\gamma \neq 1$ and $\Delta \geq 3$ supports nearly-independent phase-correlated spins and unary symmetry breaking. Then the partition function (2) of two-spin systems (β, γ, λ) is $\#\text{BIS}$ -hard to approximate on bipartite graphs with maximum degree Δ .*

Previous hardness proofs for the problem $\#\text{IS}$ and for the problem of estimating the partition function of antiferromagnetic 2-spin systems typically reduce from MAX-CUT or from the problem of counting certain types of cuts [20, 28, 29]. However such a technique sheds little light in the bipartite setting as cut problems are trivial on bipartite graphs. Reductions between $\#\text{BIS}$ -equivalent problems typically involve transformations that “blow up” vertices and edges into sets of vertices that are completely connected, so they do not apply to bounded-degree graphs either.

A key property of Sly’s gadget is that either phase occurs with probability bounded below by an inverse polynomial. This bound is sufficient in Sly’s setting to reduce from MAX-CUT , but it is not enough to construct AP-reductions for our use. We resolve this issue by introducing a balancing construction. The construction takes two copies of a gadget with nearly-independent phase-correlated spins, and produces a new gadget with similarly-correlated spins, but in the new gadget the two phases occur with probability close to $1/2$.

The proof of Theorem 1 utilizes an intermediate problem, that is, computing the partition function of antiferromagnetic Ising systems with non-uniform external fields on bipartite graphs. A non-uniform external field means that the instance specifies a subset of vertices on which the external field acts. A 2-spin system with a non-uniform external field is very similar to a Boolean #CSP with one binary symmetric non-negative valued function (corresponding to edge weights) and one unary non-negative valued function (corresponding to vertex weights) (see, for example [5]).

Our reduction implements an external field, and this is where unary symmetry breaking comes into play. As discussed earlier, the partition function of the Ising model without an external field has an FPRAS, so the unary symmetry breaking gadget seems necessary. In fact, we show (Lemma 9) that unary symmetry breaking holds for all 2-spin systems except for the Ising model without an external field or degenerate systems (i.e., systems satisfying $\beta\gamma = 1$). We also prove (Lemma 7) that all antiferromagnetic 2-spin systems support nearly-independent phase-correlated spins in the non-uniqueness region. Taking Lemmas 7 and 9 together with Theorem 1 yields our main result:

Theorem 2. *For all tuples of parameters $(\beta, \gamma, \lambda, \Delta)$ with $\Delta \geq 3$ and $\beta\gamma < 1$, except for the case $(\beta = \gamma, \lambda = 1)$, if the infinite Δ -regular tree \mathbb{T}_Δ is in the non-uniqueness region then approximating the partition function (2) on bipartite graphs with maximum degree Δ is #BIS-equivalent under AP-reductions.*

There is an FPRAS for the exceptional case $(\beta = \gamma, \lambda = 1)$ (Corollary 13). Let us now survey the approximability picture that this theorem helps establish. For general antiferromagnetic 2-spin models with soft constraints (i.e., $\beta\gamma > 0$), non-uniqueness holds if and only if $\sqrt{\beta\gamma} < \frac{\Delta-2}{\Delta}$ and $\lambda \in (\lambda_1, \lambda_2)$ for some critical values λ_1 and λ_2 depending on β , γ , and Δ (see [24, Lemma 5]). Hence, for all $\beta, \gamma > 0$ where $\beta\gamma < 1$, and all $\Delta \geq 3$ the following holds:

1. If $\sqrt{\beta\gamma} > \frac{\Delta-2}{\Delta}$, for all λ , there is an FPTAS to approximate the partition function for Δ -regular graphs [27, 24] (this extends to graphs of maximum degree Δ in an appropriate sense, see [24] for details).
2. If $\sqrt{\beta\gamma} < \frac{\Delta-2}{\Delta}$, then there exists $0 < \lambda_1 < \lambda_2$ so that:
 - (a) For all $\lambda \notin [\lambda_1, \lambda_2]$, there is an FPTAS to approximate the partition function for Δ -regular graphs [27, 24] (this again extends in an appropriate sense to graphs of maximum degree Δ [24]).
 - (b) For all $\lambda \in (\lambda_1, \lambda_2)$, it is #SAT-hard to approximate the partition function on Δ -regular graphs [29].
 - (c) For all $\lambda \in (\lambda_1, \lambda_2)$, it is #BIS-hard to approximate the partition function on bipartite graphs of maximum degree Δ (Theorem 2 in this paper).

For the particular case of the hard-core model the critical value (i.e., critical activity $\lambda_c(\Delta)$) is more easily stated. For the hard-core model (i.e., $\beta = 0$ and $\gamma = 1$) Kelly [22] showed that non-uniqueness holds on \mathbb{T}_Δ if and only if $\lambda > \lambda_c(\Delta) := \frac{(\Delta-1)^{\Delta-1}}{(\Delta-2)^\Delta}$. As a consequence we get the following corollary for the hard-core model.

Corollary 1. *For all $\Delta \geq 3$, all $\lambda > \lambda_c(\Delta) = \frac{(\Delta-1)^{\Delta-1}}{(\Delta-2)^\Delta}$, it is #BIS-hard to approximate the partition function of the hard-core model on bipartite graphs of maximum degree Δ .*

We also get a corollary concerning the more general partition function as long as β and γ are less than 1 and the degree bound Δ is sufficiently large. For β and γ satisfying $0 < \beta, \gamma < 1$ and any $\lambda > 0$, there exists a Δ such that (β, γ, λ) is in the non-uniqueness region of \mathbb{T}_Δ [24, Lemma 21.2]. This implies the following corollary.

Corollary 2. *For every $0 < \beta, \gamma < 1$ and $\lambda > 0$, there exists a Δ such that the 2-spin system with parameters β, γ and with uniform or non-uniform external field λ on bipartite graphs with degree bound Δ is #BIS-equivalent under AP-reductions, except when $\beta = \gamma$ and $\lambda = 1$, in which case it has an FPRAS.*

More generally, for antiferromagnetic 2-spin systems we get the following picture for the complexity of approximating the partition function on general graphs. As usual there is a difficulty classifying the complexity of approximating the partition function at the boundary between uniqueness and non-uniqueness. To address this issue, for parameters $(\beta, \gamma, \lambda, \Delta)$, [24] define a notion of *up-to- Δ unique* which is equivalent to the parameters lying in the interior of the uniqueness region for the infinite $(d - 1)$ -ary tree \mathbb{T}_d for all $3 \leq d \leq \Delta$ (see Definition 7 in [24]). Moreover, the parameters (β, γ, λ) satisfy *∞ -strict-uniqueness* if it is up-to- ∞ unique.⁶ On the other side, we say the parameters (β, γ, λ) satisfy *∞ -non-uniqueness* if for some $\Delta \geq 3$ the tree \mathbb{T}_Δ has non-uniqueness. The only gap between the notions of ∞ -strict-uniqueness and ∞ -non-uniqueness is the case when the parameters (β, γ, λ) are at the uniqueness/non-uniqueness threshold of \mathbb{T}_Δ for some Δ .

The following result detailing the complexity for general graphs is now established.

Corollary 3. *For all tuples of parameters (β, γ, λ) with $\beta\gamma < 1$, the following holds:*

1. *If the parameters satisfy ∞ -strict-uniqueness then there is a FPTAS for the partition function for all graphs [24, Theorem 2].*
2. *If the parameters satisfy ∞ -non-uniqueness then:*
 - (a) *it is #SAT-hard to approximate the partition function on graphs [29].*
 - (b) *except when $\beta = \gamma$ and $\lambda = 1$, it is #BIS-hard to approximate the partition function on bipartite graphs (Theorem 2 in this paper).*

A recent paper of Liu et al. [25] shows that our Theorem 1 can also be used to analyse the complexity of *ferromagnetic* partition functions (where $\beta\gamma > 1$). In particular, it uses Theorem 1 to show #BIS-hardness for approximating the partition function for ferromagnetic 2-spin systems when $\beta \neq \gamma$ for sufficiently large external field λ . An interesting problem that remains open is to prove #BIS-hardness for the entire non-uniqueness region for ferromagnetic 2-spin systems with $\beta \neq \gamma$.

The remainder of the paper is organized as follows. We begin by formally defining AP-reductions and the problems that we use in our reductions in Section 2. In Section 3 we formally define nearly-independent phase-correlated spins and unary symmetry breaking, and present the main lemmas concerning these concepts. We prove Theorem 1 in Section 4. Finally, in Sections 5, 6, and 7 we prove the various lemmas about the two key concepts that are presented in Section 3.

2. Approximation-Preserving Reductions and #BIS

We are interested in the complexity of approximate counting. Let Σ be a finite alphabet. We want to approximate the value of a function $f : \Sigma^* \rightarrow \mathbb{R}$. A *randomized approximation scheme* is an algorithm that takes an instance $x \in \Sigma^*$ and a rational error tolerance $\varepsilon > 0$ as inputs, and outputs a rational number z such that, for every x and ε ,

$$\Pr[e^{-\varepsilon} f(x) \leq z \leq e^\varepsilon f(x)] \geq \frac{3}{4}.$$

⁶To be precise, the notion of ∞ -strict-uniqueness is called *universally unique* in [24, Definition 7]).

A *fully polynomial randomized approximation scheme* (FPRAS) is a randomized approximation scheme which runs in time bounded by a polynomial in $|x|$ and ε^{-1} . Note that the quantity $\frac{3}{4}$ can be changed to any value in the interval $(\frac{1}{2}, 1)$ or even $1 - 2^{-n^c}$ for a problem of size n without changing the set of problems that have fully polynomial randomized approximation schemes since the higher accuracy can be achieved with only polynomial delay by taking a majority vote of multiple samples.

Dyer *et al.* [8] introduced the notion of approximation-preserving reductions. Suppose f and g are two functions from Σ^* to \mathbb{R} . An *approximation-preserving reduction* (AP-reduction) from f to g is a randomized algorithm \mathcal{A} to approximate f using an oracle for g . The algorithm \mathcal{A} takes an input $(x, \varepsilon) \in \Sigma^* \times (0, 1)$, and satisfies the following three conditions: (i) every oracle call made by \mathcal{A} is of the form (y, δ) , where $y \in \Sigma^*$ is an instance of g , and $0 < \delta < 1$ is an error bound satisfying $\delta^{-1} \leq \text{poly}(|x|, \varepsilon^{-1})$; (ii) the algorithm \mathcal{A} meets the specification for being a randomized approximation scheme for f whenever the oracle meets the specification for being a randomized approximation scheme for g ; (iii) the run-time of \mathcal{A} is polynomial in $|x|$ and ε^{-1} .

If an AP-reduction from f to g exists, we write $f \leq_{\text{AP}} g$, and say that f is *AP-reducible* to g . If $f \leq_{\text{AP}} g$ and $g \leq_{\text{AP}} f$, then we say that f and g are *AP-interreducible* or *AP-equivalent*, and write $f \equiv_{\text{AP}} g$. The problem #BIS is defined as follows.

Name. #BIS.

Instance. A bipartite graph B .

Output. The number of independent sets in B .

In this paper, we are interested in 2-spin systems over bounded degree bipartite graphs parametrized by a tuple (β, γ, λ) . We say a real number z is *efficiently approximable* if there is an FPRAS for the problem of computing z . Throughout the paper we only deal with non-negative real parameters that are efficiently approximable. For efficiently approximable non-negative real numbers β, γ, λ and a positive integer Δ , we define the problem of computing the partition function of the 2-spin system (β, γ) with external field λ on bipartite graphs of bounded degree Δ , as follows.

Name. BI-(M-)2-SPIN($\beta, \gamma, \lambda, \Delta$).

Instance. A bipartite (multi)graph $B = (V, E)$ with degree bound Δ .

Output. The quantity

$$Z_B(\beta, \gamma, \lambda) = \sum_{\sigma: V \rightarrow \{0,1\}} \lambda^{\sum_{v \in V} \sigma(v)} \prod_{(v,u) \in E} \beta^{(1-\sigma(v))(1-\sigma(u))} \gamma^{\sigma(v)\sigma(u)}.$$

Notice that we also introduced a multigraph version of the same problem. It will be useful later. We drop the parameter Δ when there is no degree bound, that is, BI-2-SPIN(β, γ, λ) is the same as BI-2-SPIN($\beta, \gamma, \lambda, \infty$).

We also found the notion of non-uniform external field useful in the reductions. The following problems are introduced as intermediate problems. We also introduce a multigraph version, but as intermediate problems we do not need the bounded degree variant.

Name. BI-(M-)NONUNIFORM-2-SPIN(β, γ, λ).

Instance. A bipartite (multi)graph $B = (V, E)$ and a subset $U \subseteq V$.

Output. The quantity

$$Z_{B,U}(\beta, \gamma, \lambda) = \sum_{\sigma: V \rightarrow \{0,1\}} \lambda^{\sum_{v \in U} \sigma(v)} \prod_{(v,u) \in E} \beta^{(1-\sigma(v))(1-\sigma(u))} \gamma^{\sigma(v)\sigma(u)}.$$

3. Key Properties of the Gadget

In this section we define two key concepts: nearly-independent phase-correlated spins and unary symmetry breaking.

We first describe the basic setup of a certain gadget. For positive integers Δ , t and n where n is even and is at least $2t$, let T^- and T^+ be disjoint vertex sets of size t and let V^- be a size- $n/2$ superset of T^- and V^+ be a size- $n/2$ superset of T^+ which is disjoint from V^- . Let $T = T^- \cup T^+$ and $V(t, n) = V^- \cup V^+$. Let $\mathcal{G}(t, n, \Delta)$ be the set of bipartite graphs with vertex partition (V^-, V^+) in which every vertex has degree at most Δ and every vertex in T has degree at most $\Delta - 1$. We refer to the vertices in T as “terminals”. Vertices in T^+ are “positive terminals” and vertices in T^- are “negative terminals”.

When the gadget G is drawn from $\mathcal{G}(t, n, \Delta)$, we use the notation $T(G)$ to refer to the set of terminals. Each configuration $\sigma: V(t, n) \rightarrow \{0, 1\}$ is assigned a unique phase $Y(\sigma) \in \{-, +\}$. Roughly in our applications of the definitions below the phase of a configuration σ is π if V^π contains more vertices with spin 1 than does $V^{-\pi}$.

We define measures Q^+ and Q^- . Fix some $0 < q^- < q^+ < 1$. For any positive integer t ,

- Q^+ is the distribution on configurations $\tau: T \rightarrow \{0, 1\}$ such that, for every $v \in T^+$, $\tau(v) = 1$ independently with probability q^+ and, for every $v \in V^-$, $\tau(v) = 1$ independently with probability q^- , and
- Q^- is the distribution on configurations $\tau: T \rightarrow \{0, 1\}$ such that, for every $v \in T^-$, $\tau(v) = 1$ independently with probability q^+ and, for every $v \in T^+$, $\tau(v) = 1$ independently with probability q^- .

To give a rough sense for the values q^- and q^+ they will correspond to the marginal probabilities of the root of an infinite tree obtained by taking limits of finite trees with appropriate boundary conditions, see Section 7 for more details.

To prove the #BIS-hardness we need a gadget where the spins of the terminals are drawn from distributions close to Q^+ or Q^- conditioned on the phase $+$ or $-$. The following definition formalises this notion. The definition will be crucial for obtaining our #BIS-hardness results for antiferromagnetic models in the non-uniqueness region. While the definition is stated for general 2-spin systems, it is most useful for antiferromagnetic systems; a ferromagnetic 2-spin system is not going to satisfy the conditions of the definition since one expects that for such systems $q^- = q^+$ or, equivalently, there is no notion of a binary “phase”.

Definition 4. *A tuple of parameters $(\beta, \gamma, \lambda, \Delta)$ supports nearly-independent phase-correlated spins if there are efficiently-approximable values $0 < q^- < q^+ < 1$ such that the following is true. There are functions $n(t, \varepsilon)$, $m(t, \varepsilon)$, and $f(t, \varepsilon)$, each of which is bounded from above by a polynomial in t and ε^{-1} , and for every t and ε there is a distribution on graphs in $\mathcal{G}(t, n(t, \varepsilon), \Delta)$ such that a gadget $G = (V, E)$ with terminals T can be drawn from the distribution within $m(t, \varepsilon)$ time, and the probability that the following inequalities hold is at least $3/4$:*

1. *The phases are roughly balanced, i.e.,*

$$\Pr_{G;\beta,\gamma,\lambda}(Y(\sigma) = +) \geq \frac{1}{f(t, \varepsilon)} \text{ and } \Pr_{G;\beta,\gamma,\lambda}(Y(\sigma) = -) \geq \frac{1}{f(t, \varepsilon)}. \quad (3)$$

2. *For a configuration $\sigma: V \rightarrow \{0, 1\}$ and any $\tau: T \rightarrow \{0, 1\}$,*

$$\left| \frac{\Pr_{G;\beta,\gamma,\lambda}(\sigma|_T = \tau \mid Y(\sigma) = +)}{Q^+(\tau)} - 1 \right| \leq \varepsilon \text{ and } \left| \frac{\Pr_{G;\beta,\gamma,\lambda}(\sigma|_T = \tau \mid Y(\sigma) = -)}{Q^-(\tau)} - 1 \right| \leq \varepsilon. \quad (4)$$

In fact, given a gadget with the above property, one can construct a gadget where the phases are (nearly) uniformly distributed as detailed in the following definition.

Definition 5. We say that the tuple of parameters $(\beta, \gamma, \lambda, \Delta)$ supports **balanced** nearly-independent phase-correlated spins if Definition 4 holds with (3) replaced by:

$$\Pr_{G;\beta,\gamma,\lambda}(Y(\sigma) = +) \geq \frac{1-\varepsilon}{2} \text{ and } \Pr_{G;\beta,\gamma,\lambda}(Y(\sigma) = -) \geq \frac{1-\varepsilon}{2}, \quad (5)$$

where ε is quantified as in Definition 4.

The following lemma shows that for essentially all 2-spin systems, Definition 4 implies Definition 5. Hence, in the proof of Theorem 1 we assume the existence of a gadget with balanced phases.

Lemma 6. If the parameter tuple $(\beta, \gamma, \lambda, \Delta)$ with $\beta\gamma \neq 1$ supports nearly-independent phase-correlated spins, then it supports balanced nearly-independent phase-correlated spins.

The main technical result for proving #BIS-hardness for 2-spin antiferromagnetic systems in the tree non-uniqueness region is the following lemma, which is proved in Section 7.

Lemma 7. For all $\Delta \geq 3$, all $\beta, \gamma, \lambda > 0$ where $\beta\gamma < 1$, if the infinite Δ -regular tree \mathbb{T}_Δ is in the non-uniqueness region then the tuple of parameters $(\beta, \gamma, \lambda, \Delta)$ supports balanced nearly-independent phase-correlated spins.

The second property of the gadget is the notion of unary symmetry breaking which is relatively simple.

Definition 8. We say that a tuple of parameters $(\beta, \gamma, \lambda, \Delta)$ **supports unary symmetry breaking** if there is a bipartite graph H whose vertices have degree at most Δ which has a distinguished degree-1 vertex v_H such that $\Pr_{H;\beta,\gamma,\lambda}(\sigma_{v_H} = 1) \notin \{0, \lambda/(1+\lambda), 1\}$.

We will prove in Section 6 that unary symmetry breaking holds for all 2-spin models except for two cases.

Lemma 9. Assume $\Delta \geq 3$. The parameters $(\beta, \gamma, \lambda, \Delta)$ support unary symmetry breaking unless (i) $\beta\gamma = 1$ or (ii) $\beta = \gamma$ and $\lambda = 1$.

4. General Reduction

In this section we prove Theorem 1. We first show how the two notions of “nearly-independent phase-correlated spins” and “unary symmetry breaking” lead to #BIS-hardness.

4.1. An Intermediate Problem

The goal of this section is to show that it is #BIS-hard to approximate the partition function of antiferromagnetic Ising models with non-uniform non-trivial external fields on bipartite graphs.

Lemma 10. For any $0 < \alpha < 1$, $\lambda > 0$ and $\lambda \neq 1$,

$$\#\text{BIS} \leq_{AP} \text{BI-M-NONUNIFORM-2-SPIN}(\alpha, \alpha, \lambda).$$

Proof. By flipping 0 to 1 and 1 to 0 for each configuration σ , we see that BI-M-NONUNIFORM-2-SPIN(α, α, λ) is in fact the same as BI-M-NONUNIFORM-2-SPIN($\alpha, \alpha, 1/\lambda$). Hence we may assume $\lambda < 1$.

Let $M = \begin{pmatrix} \alpha & 1 \\ 1 & \alpha \end{pmatrix}$, and $\begin{pmatrix} \rho_0 \\ \rho_1 \end{pmatrix} = M \begin{pmatrix} 1 \\ \lambda \end{pmatrix} = \begin{pmatrix} \alpha + \lambda \\ 1 + \alpha\lambda \end{pmatrix}$. Note that $\alpha < 1$ and $\lambda < 1$, so $\rho_1 > \rho_0$.

Let $B = (V, E)$ be an input to #BIS with $n = |V|$ and $m = |E|$. Let I_B be the number of independent sets of B . Let ε be the desired accuracy of the reduction. We will construct an instance $B' = (V', E')$ with a specified vertex subset $U \subset V'$ for BI-M-NONUNIFORM-2-SPIN(α, α, λ) such that

$$\exp\left(-\frac{\varepsilon}{2}\right) I_B \leq \frac{Z_{B',U}(\alpha, \alpha, \lambda)}{C} \leq \exp\left(\frac{\varepsilon}{2}\right) I_B,$$

where C is a quantity that is easy to approximate. Therefore it suffices to call oracle BI-M-NONUNIFORM-2-SPIN(α, α, λ) on B' with the specified subset U with accuracy $\frac{\varepsilon}{4}$ and approximate C within $\frac{\varepsilon}{4}$.

The construction of B' involves two positive integers t_1 and t_2 . Let t_1 be the least positive integer such that

$$\alpha^{2t_1} \leq \frac{\varepsilon}{6 \cdot 2^n}. \quad (6)$$

Note that t_1 depends on n and ε and there is a polynomial p in n and ε^{-1} such that $t_1 \leq p(n, \varepsilon^{-1})$. Let t_2 be the least positive integer depending on n, ε and t_1 such that

$$\left(\frac{\rho_0}{\rho_1}\right)^{t_2} \leq \frac{\alpha^{t_1 m} \cdot \varepsilon}{6 \cdot 2^{2t_1 m + n}}. \quad (7)$$

Once again, t_2 is bounded from above by a polynomial in n and ε^{-1} .

Given the integers t_1 and t_2 , the graph B' is constructed as follows (see Figure 1 for an illustrative diagram). Let $W_v = \{w_{v,j} \mid 1 \leq j \leq t_1 \deg(v)\}$ for each $v \in V$ where $\deg(v)$ is the degree of v in B . Let $U_{v,j} = \{u_{v,j,k} \mid 1 \leq k \leq t_2\}$ for any $v \in V$ and $1 \leq j \leq t_1 \deg(v)$. Let

$$W = \bigcup_{v \in V} W_v \quad \text{and} \quad U = \bigcup_{v \in V} \bigcup_{1 \leq j \leq t_1 \deg(v)} U_{v,j}.$$

The vertex set of B' is $V' = V \cup U \cup W$. Note that $|W| = 2t_1 m$ and $|U| = 2t_1 t_2 m$.

We add t_1 parallel edges in B' between u and v for each $(u, v) \in E$ and add edges between v and every vertex in W_v , and between $w_{v,j}$ and every vertex in $U_{v,j}$ for each $v \in V$ and $1 \leq j \leq t_1 \deg(v)$. Formally the edge set of B' is

$$E' = \left(\bigsqcup_{1 \leq i \leq t_1} E \right) \cup \bigcup_{v \in V} E_v \cup \bigcup_{\substack{v \in V \\ 1 \leq j \leq t_1 \deg(v)}} E_{v,j},$$

where \bigsqcup denotes a disjoint union as a multiset of t_1 copies of E , $E_v = \{(v, w) \mid w \in W_v\}$ and $E_{v,j} = \{(w_{v,j}, u) \mid u \in U_{v,j}\}$ for each v and j .

Let $C = \rho_1^{2t_1 t_2 m} \alpha^{t_1 m}$ and $N = \begin{pmatrix} 1 & 1 \\ 1 & \alpha^{2t_1} \end{pmatrix}$.

For each $\sigma : V \cup W \rightarrow \{0, 1\}$, let $w(\sigma)$ be the contribution to $Z_{B',U}(\alpha, \alpha, \lambda)$ of configurations that are consistent with σ . First consider configurations σ such that $\sigma(w) = 1$ for all $w \in W$. Denote by Σ




Figure 1: An illustrative diagram of the reduction in the proof of Lemma 10. As an example, we consider the (bipartite) graph B in the left part of the picture to be the input to $\#\text{BIS}$. The respective input B' to $\text{BI-M-NONUNIFORM-2-SPIN}(\alpha, \alpha, \lambda)$ is shown in the right part of the picture. For this graph B , where $\deg(u) = \deg(v) = 1$, t_1 is the number of parallel edges between u and v in B' , as well as the cardinalities of the sets W_u and W_v (where t_1 is an appropriate integer, cf. (6)). Also, t_2 (cf. (7)) equals the cardinalities of the sets $U_{u,j}$ and $U_{v,j}$ for $j = 1, \dots, t_1$. For a general graph B and a vertex u in B , the cardinality of the set W_u depends on the degree of the vertex u and equals $t_1 \deg(u)$. Other than that, the construction of the multigraph B' is analogous by iterating the construction in the figure for each edge in B .

the set of all such configurations on $V \cup W$. Then for $\sigma \in \Sigma$, we have

$$\begin{aligned} w(\sigma) &= \rho_1^{t_2|W|} \prod_{(u,v) \in E} (M_{1,\sigma(u)} M_{\sigma(u),\sigma(v)} M_{\sigma(v),1})^{t_1} \\ &= C \prod_{(u,v) \in E} N_{\sigma(u),\sigma(v)}. \end{aligned} \quad (8)$$

To see (8), note that W separates the vertices in U from the rest of the vertices of B' . Since σ specifies the spins of all vertices in W , the weight $w(\sigma)$ is the product of: (i) the weights induced on edges that do not involve vertices in U , and (ii) the weights induced on edges with endpoints in U . Item (i) is just $\prod_{(u,v) \in E} (M_{1,\sigma(u)} M_{\sigma(u),\sigma(v)} M_{\sigma(v),1})^{t_1}$. For Item (ii), we need to sum over all possible assignments on U ; the factor $\rho_1^{t_2|W|}$ gives the value of this sum as can be seen by expanding the expression

$$\prod_{(w,u) \in E' \cap (W \times U)} (1 + \alpha\lambda) = (1 + \alpha\lambda)^{t_2|W|} = \rho_1^{t_2|W|}.$$

Let $\Sigma^{ind} \subset \Sigma$ be the subset of configurations which induce an independent set on the vertices V and Z^{ind} be its contribution to $Z_{B',U}(\alpha, \alpha, \lambda)$. Let $\Sigma^{bad} = \Sigma \setminus \Sigma^{ind}$ and Z^{bad} be its contribution. If $\sigma \in \Sigma^{ind}$ then $w(\sigma) = C$. Otherwise, $w(\sigma) \leq \alpha^{2t_1} C$. It implies

$$Z^{ind} = I_B \cdot C \quad \text{and} \quad Z^{bad} \leq 2^n \alpha^{2t_1} C \leq \frac{\varepsilon}{6} \cdot C, \quad (9)$$

since t_1 satisfies Eq. (6).

Next consider configurations σ on $V \cup W$ such that $\sigma(w) = 0$ for at least one $w \in W$. Denote this set by Σ' and its contribution by Z^{small} . Then for $\sigma \in \Sigma'$,

$$w(\sigma) \leq \left(\rho_0 \rho_1^{|W|-1} \right)^{t_2} \leq \left(\frac{\rho_0}{\rho_1} \right)^{t_2} \rho_1^{t_2|W|} = \left(\frac{\rho_0}{\rho_1} \right)^{t_2} \frac{C}{\alpha^{t_1 m}}.$$

It implies

$$Z^{small} \leq 2^{2t_1 m+n} \left(\frac{\rho_0}{\rho_1}\right)^{t_2} \frac{C}{\alpha^{t_1 m}} \leq \frac{\varepsilon}{6} \cdot C, \quad (10)$$

since $|\Sigma'| \leq 2^{2t_1 m+n}$ and t_2 satisfies Eq. (7).

By Eq. (9) and Eq. (10) we have

$$\begin{aligned} Z_{B',U}(\alpha, \alpha, \lambda) &= Z^{ind} + Z^{bad} + Z^{small} \\ &\leq I_B \cdot C + \frac{\varepsilon}{6} \cdot C + \frac{\varepsilon}{6} \cdot C \\ &\leq \exp\left(\frac{\varepsilon}{3}\right) I_B \cdot C, \end{aligned}$$

and clearly $Z_{B',U}(\alpha, \alpha, \lambda) \geq I_B \cdot C$. It is also clear that C can be approximated accurately enough given FPRASes for λ and α . This finishes our proof. \square

4.2. Simulating the Antiferromagnetic Ising Model

In this section we prove the following lemma.

Lemma 11. *Suppose β , γ and λ are efficiently approximable reals satisfying $\beta, \gamma \geq 0$, $\lambda > 0$ and $\beta\gamma \neq 1$. Suppose that Δ is either an integer that is at least 3 or $\Delta = \infty$ (indicating that we do not have a degree bound). If $(\beta, \gamma, \lambda, \Delta)$ supports nearly-independent phase-correlated spins and unary symmetry breaking, then there exist efficiently approximable $0 < \alpha < 1$ and $\lambda' > 0$ such that $\lambda' \neq 1$ and*

$$\text{BI-M-NONUNIFORM-2-SPIN}(\alpha, \alpha, \lambda') \leq_{AP} \text{BI-2-SPIN}(\beta, \gamma, \lambda, \Delta).$$

Proof. We prove the antiferromagnetic case first, that is, $\beta\gamma < 1$. α and λ' are chosen as follows. Recall that $M = \begin{pmatrix} \beta & 1 \\ 1 & \gamma \end{pmatrix}$ and $M^+ = \begin{pmatrix} 1 - q^- & q^- \\ 1 - q^+ & q^+ \end{pmatrix}$. Let $N = M^+ M (M^+)^T = \begin{pmatrix} N_{--} & N_{-+} \\ N_{+-} & N_{++} \end{pmatrix}$. Then $\det(N) = (\beta\gamma - 1)(q^+ - q^-)^2 < 0$. Therefore $N_{--}N_{++} < N_{-+}N_{+-}$ and let $\alpha = \frac{N_{--}N_{++}}{N_{-+}N_{+-}} < 1$. Moreover, suppose H is the unary symmetry breaking gadget with distinguished vertex v_H . Let $\rho = \begin{pmatrix} \rho_0 \\ \rho_1 \end{pmatrix}$ where ρ_i denote $\Pr_{H;\beta,\gamma,\lambda}(\sigma_{v_H} = i)$ for spin $i \in \{0, 1\}$ and $\rho_0 + \rho_1 = 1$. Let $\rho' = \begin{pmatrix} \rho'_0 \\ \rho'_1 \end{pmatrix} = M^+ \begin{pmatrix} \rho_0 \\ \rho_1/\lambda \end{pmatrix}$, and $\lambda' = \frac{\rho'_1}{\rho'_0}$. It is easy to verify that $\lambda' \neq 1$ as $\rho_1 \neq \lambda/(1 + \lambda)$ by the unary symmetry breaking assumption.

Given $0 < \varepsilon < 1$ and a bipartite multigraph $B = (V, E)$ with a subset $U \subseteq V$ where $|V| = n$, $|E| = m$, and $|U| = n'$, our reduction first constructs a bipartite graph B' with degree at most Δ . The construction of B' involves a gadget G . Since $(\beta, \gamma, \lambda, \Delta)$ supports nearly-independent phase-correlated spins, by Lemma 6 $(\beta, \gamma, \lambda, \Delta)$ also supports balanced nearly-independent phase-correlated spins. Therefore we draw $G \sim \mathcal{G}(t, n(t, \varepsilon'), \Delta)$ such that Eq. (5) and Eq. (4) hold with probability at least $3/4$, where $t = m + 1$ and $\varepsilon' = \frac{\varepsilon}{8n}$. Assume G satisfies them and otherwise the reduction fails. We will construct B' such that

$$\exp\left(-\frac{\varepsilon}{2}\right) Z_{B,U}(\alpha, \alpha, \lambda') \leq \frac{Z_{B'}}{(N_{+-}N_{-+})^m (\rho'_0 Z_H)^{n'} \left(\frac{Z_G}{2}\right)^n} \leq \exp\left(\frac{\varepsilon}{2}\right) Z_{B,U}(\alpha, \alpha, \lambda'),$$

where we use the abbreviated expressions $Z_{B'} = Z_{B'}(\beta, \gamma, \lambda)$, $Z_H = Z_H(\beta, \gamma, \lambda)$, and $Z_G = Z_G(\beta, \gamma, \lambda)$. The lemma follows by one oracle call for $Z_{B'}$ with accuracy $\frac{\varepsilon}{6}$, one oracle call for Z_G with accuracy $\frac{\varepsilon}{6n}$,

and an approximation of other terms in the denominator with accuracy $\frac{\varepsilon}{6}$ using FPRASEs for q^- , q^+ , β , γ and λ .

The graph B' is constructed as follows. For each vertex $v \in V$ we introduce a copy of G , denoted by G_v with vertex set $V(G_v)$. Moreover, for each vertex $u \in U$ we introduce a copy of H , denoted by H_u . Whenever a terminal vertex is used in the construction once, we say it is occupied. For each $(u, v) \in E$, we connect one currently unoccupied positive (and respectively negative) terminal of G_u to one currently unoccupied positive (and respectively negative) terminal of G_v . Denote by E' all these edges between terminals. For each $u \in U$, we identify an unoccupied positive terminal of G_u with the distinguished vertex v_{H_u} of H_u . We denote this terminal by t_u . The resulting graph is B' . It is clear that B' is bipartite and has bounded degree Δ .

Let $\tilde{\sigma}: V \rightarrow \{-, +\}$ be a configuration of the phases of the G_v 's. Let $Z_{B'}(\tilde{\sigma})$ be the contribution to $Z_{B'}$ from the configurations σ that are consistent with $\tilde{\sigma}$ in the sense that, for each $v \in V$, $Y(\sigma_{V(G_v)}) = \tilde{\sigma}(v)$. Then $Z_{B'} = \sum_{\tilde{\sigma}} Z_{B'}(\tilde{\sigma})$. Let T be the set of all terminals $T = \cup_{v \in V} T(G_v)$ and $\tau: T \rightarrow \{0, 1\}$ be a configuration on T . Let $\tau_{T(G_v)}$ be the configuration τ restricted to $T(G_v)$. Recall that for $\pi \in \{-, +\}$, $Z_{G_v}^\pi(\tau_{T(G_v)})$ is the contribution to Z_{G_v} from configurations that have phase π and are consistent with $\tau_{T(G_v)}$. Also,

$$\Pr_{G_v; \beta, \gamma, \lambda}(\tau_{T(G_v)} \mid Y(\sigma_{V(G_v)}) = \pi) = \frac{Z_{G_v}^\pi(\tau_{T(G_v)})}{Z_{G_v}^\pi}.$$

Moreover, for each $u \in U$ and each spin $i \in \{0, 1\}$, let $Z_{H_u}(i)$ be the contribution to Z_{H_u} from configurations σ with $\sigma(t_u) = i$. Hence,

$$\rho_i = \Pr_{H_u; \beta, \gamma, \lambda}(\sigma(t_u) = i) = \frac{Z_{H_u}(i)}{Z_{H_u}}.$$

We express $Z_{B'}(\tilde{\sigma})$ as

$$Z_{B'}(\tilde{\sigma}) = \sum_{\tau: T \rightarrow \{0, 1\}} w_{E'}(\tau) \prod_{v \in V} Z_{G_v}^{\tilde{\sigma}(v)}(\tau_{T(G_v)}) \prod_{u \in U} \frac{Z_{H_u}(\tau(t_u))}{\lambda^{\tau(t_u)}},$$

where $w_{E'}(\tau)$ is the contribution of edges in E' given configuration τ . Notice that we divide the last factor by λ when $\tau(t_u) = 1$ because we counted the vertex weight twice in that case. Define $\widetilde{Z}_{B'}(\tilde{\sigma})$ to be an approximation version of the partition function where on each $T(G_v)$ the spins are chosen exactly according to $Q^{\tilde{\sigma}(v)}$. That is,

$$\begin{aligned} \widetilde{Z}_{B'}(\tilde{\sigma}) &= \sum_{\tau: T \rightarrow \{0, 1\}} w_{E'}(\tau) \prod_{v \in V} Z_{G_v}^{\tilde{\sigma}(v)} Q^{\tilde{\sigma}(v)}(\tau_{T(G_v)}) \prod_{u \in U} \frac{Z_{H_u}(\tau(t_u))}{\lambda^{\tau(t_u)}} \\ &= \left(\prod_{v \in V} Z_{G_v}^{\tilde{\sigma}(v)} \right) \cdot \left(\sum_{\tau: T \rightarrow \{0, 1\}} w_{E'}(\tau) \prod_{v \in V} Q^{\tilde{\sigma}(v)}(\tau_{T(G_v)}) \prod_{u \in U} \frac{Z_{H_u}(\tau(t_u))}{\lambda^{\tau(t_u)}} \right). \end{aligned} \quad (11)$$

Let $\widetilde{Z}_{B'} = \sum_{\tilde{\sigma}} \widetilde{Z}_{B'}(\tilde{\sigma})$. Eq. (4) implies that $Z_{B'}(\tilde{\sigma})$ and $\widetilde{Z}_{B'}(\tilde{\sigma})$ are close, that is,

$$(1 - \varepsilon')^n \leq \frac{Z_{B'}(\tilde{\sigma})}{\widetilde{Z}_{B'}(\tilde{\sigma})} \leq (1 + \varepsilon')^n. \quad (12)$$

Moreover Eq. (5) implies that

$$\left(\frac{1-\varepsilon'}{2}\right)^n \leq \frac{\prod_{v \in V} Z_{G_v}^{\tilde{\sigma}(v)}}{(Z_G)^n} \leq \left(\frac{1+\varepsilon'}{2}\right)^n. \quad (13)$$

Notice that here Z_{G_v} is the same for any $v \in V$ as the G_v 's are identical copies of G .

Then we calculate the following quantity given $\tilde{\sigma}$

$$\sum_{\tau: T \rightarrow \{0,1\}} w_{E'}(\tau) \prod_{v \in V} Q^{\tilde{\sigma}(v)}(\tau_{T(G_v)}) \prod_{u \in U} \frac{Z_{H_u}(\tau(t_u))}{\lambda^{\tau(t_u)}}.$$

As the measure $Q^{\tilde{\sigma}(v)}$ is i.i.d., we may count the weight of each edge in E' independently. Notice that $N_{\pi_1 \pi_2}$ is the edge contribution when one end point is chosen with probability q^{π_1} and the other q^{π_2} . For an edge $(u, v) \in V$, if u and v are assigned the same phase $+$, then an edge in E' connecting one $+$ terminal of G_u and one $+$ terminal of G_v gives a weight of N_{++} and an edge connecting two $-$ terminals gives N_{--} . The total weight is $\mu_1 = N_{++}N_{--}$. Similarly if u and v are assigned the same phase $-$, the total weight is μ_1 as well. On the other hand if u and v are assigned distinct phases $+$ and $-$, the total weight is $\mu_2 = N_{+-}N_{-+}$. Recall that $\alpha = \frac{\mu_1}{\mu_2}$. Moreover, for each $u \in U$, if $\tilde{\sigma}(u) = +$, then the contribution of H_u is $\rho'_1 Z_{H_u}$ and otherwise $\rho'_0 Z_{H_u}$. Notice that here Z_{H_u} is the same for any $u \in U$ as the H_u 's are identical copies of H . Recall that $\lambda' = \frac{\rho'_1}{\rho'_0}$.

Plug these calculation into Eq. (11), we have

$$\begin{aligned} \widetilde{Z_{B'}}(\tilde{\sigma}) &= \left(\prod_{v \in V} Z_{G_v}^{\tilde{\sigma}(v)} \right) \cdot \left(\mu_1^{m_+(\tilde{\sigma})} \mu_2^{m_-(\tilde{\sigma})} (\rho'_1 Z_H)^{n_+(\tilde{\sigma})} (\rho'_0 Z_H)^{n_-(\tilde{\sigma})} \right) \\ &= \mu_2^m (\rho'_0 Z_H)^{n'} \left(\prod_{v \in V} Z_{G_v}^{\tilde{\sigma}(v)} \right) \cdot \left(\alpha^{m_+(\tilde{\sigma})} (\lambda')^{n_+(\tilde{\sigma})} \right), \end{aligned} \quad (14)$$

where $m_+(\tilde{\sigma})$ denotes the number of edges of which the two endpoints are of the same phase given $\tilde{\sigma}$, and $n_+(\tilde{\sigma})$ denotes the number of vertices in U that are assigned $+$ given $\tilde{\sigma}$. Apply Eq.(13) to Eq.(14),

$$(1-\varepsilon')^n \left(\alpha^{m_+(\tilde{\sigma})} (\lambda')^{n_+(\tilde{\sigma})} \right) \leq \frac{\widetilde{Z_{B'}}(\tilde{\sigma})}{\mu_2^m (\rho'_0 Z_H)^{n'} \left(\frac{Z_G}{2} \right)^n} \leq (1+\varepsilon')^n \left(\alpha^{m_+(\tilde{\sigma})} (\lambda')^{n_+(\tilde{\sigma})} \right). \quad (15)$$

Then we sum over $\tilde{\sigma}$ in Eq. (15),

$$(1-\varepsilon')^n \left(\sum_{\tilde{\sigma}} \alpha^{m_+(\tilde{\sigma})} (\lambda')^{n_+(\tilde{\sigma})} \right) \leq \frac{\widetilde{Z_{B'}}}{\mu_2^m (\rho'_0 Z_H)^{n'} \left(\frac{Z_G}{2} \right)^n} \leq (1+\varepsilon')^n \left(\sum_{\tilde{\sigma}} \alpha^{m_+(\tilde{\sigma})} (\lambda')^{n_+(\tilde{\sigma})} \right) \quad (16)$$

However notice that $Z_{B,U}(\alpha, \alpha, \lambda') = \sum_{\tilde{\sigma}} \alpha^{m_+(\tilde{\sigma})} (\lambda')^{n_+(\tilde{\sigma})}$ by just mapping $-$ to 0 and $+$ to 1 in each configuration $\tilde{\sigma}$. Combine Eq.(12), and Eq.(16),

$$(1-\varepsilon')^{2n} Z_{B,U}(\alpha, \alpha, \lambda') \leq \frac{Z_{B'}}{\mu_2^m (\rho'_0 Z_H)^{n'} \left(\frac{Z_G}{2} \right)^n} \leq (1+\varepsilon')^{2n} Z_{B,U}(\alpha, \alpha, \lambda').$$

Recall that $\varepsilon' = \frac{\varepsilon}{8n}$ and we get the desired bounds.

The other case is ferromagnetic, that is, $\beta\gamma > 1$. Notice that in this case $\det(N) = (\beta\gamma - 1)(q^+ - q^-)^2 > 0$, So we choose $\alpha = \frac{N_{+-}N_{-+}}{N_{++}N_{--}} < 1$ and λ' to be the same as the antiferromagnetic case. The construction of B' is similar to the previous case, with the following change. For each $(u, v) \in E$, we connect one unoccupied positive terminal of G_u to one unoccupied negative terminal of G_v , and vice versa. The rest of the construction is the same. With this change, given a configuration $\tilde{\sigma}: V \rightarrow \{-, +\}$, if two endpoints are assigned the same spin, the contribution is $N_{+-}N_{-+}$ and otherwise $N_{++}N_{--}$. Therefore the effective edge weight is $\alpha < 1$ when the spins are the same, after normalizing the weight to 1 when the spins are distinct. The rest of the proof is the same. \square

4.3. Completing the Proof of Theorem 1

We can now conclude the proof of Theorem 1.

Proof of Theorem 1. #BIS-hardness in Theorem 1 follows directly from Lemma 10 and Lemma 11. The other direction, #BIS-easiness, follows fairly directly from Theorem 47 of [6] (the full version of [7]). An edge in the instance graph can be viewed as a constraint of arity 2. If $\beta\gamma > 1$, then the constraint on the edge is “weakly log-supermodular” and the vertex weight can be viewed as a unary constraint, which is taken as given in a “conservative” CSP. If $\beta\gamma \leq 1$, then reverse the interpretation of 0 and 1 on one side of the bipartition of the instance graph, so that the effective interaction along an edge is given by the matrix $\begin{pmatrix} 1 & \beta \\ \gamma & 1 \end{pmatrix}$. This constraint is also “weakly log-supermodular” since $1 \cdot 1 \geq \beta\gamma$. After the reversing there are two vertex weights λ and λ^{-1} , which are also allowed for “conservative” CSP instances. \square

5. Balanced Nearly-Independent Phase-Correlated Spins

In this section we prove Lemma 6 that a gadget with nearly-independent phases can be used to construct a gadget with balanced phases.

Before proving the lemma, we introduce some notation. Let

$$M = \begin{pmatrix} \beta & 1 \\ 1 & \gamma \end{pmatrix}.$$

Let q^+ and q^- be the quantities from Definition 4 and let

$$M^+ = \begin{pmatrix} 1 - q^- & q^- \\ 1 - q^+ & q^+ \end{pmatrix}$$

The two columns of M^+ correspond to spin 0 and spin 1. The first row corresponds to the distribution induced on a positive terminal from Q_t^- and the second to the distribution induced from Q_t^+ . Similarly the first row also corresponds to the distribution induced on a negative terminal from Q_t^+ and the second to the distribution induced from Q_t^- . Notice that $\det(M^+) = q^+ - q^- > 0$.

When the parameters β, γ and λ are clear, we sometimes make the notation more concise, referring to the partition function as Z_G rather than as $Z_G(\beta, \gamma, \lambda)$. Also, given a configuration $\sigma: V(G) \rightarrow \{0, 1\}$ and a subset S of $V(G)$, we often use the notation σ_S to denote the restriction $\sigma|_S$. For a gadget G drawn from $\mathcal{G}(t, n(t, \varepsilon), \Delta)$, let Z_G^π be the contribution of phase $\pi \in \{-, +\}$ to the partition function Z_G . Moreover, for a subset $S \subseteq T(G)$, suppose $\tau_S: T(G) \rightarrow \{0, 1\}$ is a configuration on terminals in S . Let $Z_G^\pi(\tau_S)$ be the contribution of configurations that are consistent with τ_S and belong to phase π , that is,

$$Z_G^\pi(\tau_S) = \sum_{\substack{\sigma: Y(\sigma)=\pi \\ \sigma_S=\tau_S}} w(\sigma),$$

where $w(\sigma)$ is the weight of configuration σ defined in (1). It is easy to see that for $\pi \in \{-, +\}$,

$$\Pr_{G;\beta,\gamma,\lambda}(Y(\sigma) = \pi) = \frac{Z_G^\pi}{Z_G},$$

and

$$\Pr_{G;\beta,\gamma,\lambda}(\sigma_{T(G)} = \tau_{T(G)} \mid Y(\sigma) = \pi) = \frac{Z_G^\pi(\tau_{T(G)})}{Z_G^\pi}.$$

We are now prepared to prove the lemma which is the focus of this section.

Proof of Lemma 6. Let ε satisfy $0 < \varepsilon < 1$. By assumption we may draw a gadget G from $\mathcal{G}(t+t', n(t+t', \varepsilon'), \Delta)$ such that it satisfies Eq. (3) and Eq. (4) with probability at least $3/4$, where t' and ε' will be specified later. Assume G does. Otherwise the construction fails.

We consider first the antiferromagnetic case $\beta\gamma < 1$. We construct a gadget K such that K satisfies Eq. (5) and Eq. (4). We make two copies of G , say G_1 and G_2 . Let the terminals of G_i be $T(G_i) = T^+(G_i) \cup T^-(G_i)$ for $i = 1, 2$. For each $\pi \in \{-, +\}$, we add a set of edges that form a perfect matching between t' terminals in $T^\pi(G_1)$ and t' terminals in $T^\pi(G_2)$. Denote by P the edges of the two perfect matchings. K is the resulting graph. Denote by C_i the vertices of G_i that are endpoints of P . The terminals of K are those $2t$ terminal nodes in $T(G_1)$ that are still unmatched, that is $T^\pi(K) = T^\pi(G_1) \setminus C_1$ for $\pi \in \{-, +\}$, and $T(K) = T^+(K) \cup T^-(K)$. Denote by I the terminals of G_2 that are unmatched.

We define the phase of K to be the phase of G_1 , that is, K is said to have phase $+$ or $-$ if and only if G_1 has the same $+$ or $-$ phase regardless of the phase of G_2 . Let (π_1, π_2) be a vector denoting the phases of G_1 and G_2 where $\pi_1, \pi_2 \in \{-, +\}$. Then the $+$ phase of K corresponds to the vector $\{(+, +), (+, -)\}$ and the $-$ phase corresponds to $\{(-, +), (-, -)\}$. For two configurations τ_V and τ_U with $V \cap U = \emptyset$, let (τ_V, τ_U) be the joint configuration on $V \cup U$. Then we have the following:

$$\begin{aligned} Z_K^\pi(\tau_{T(K)}) &= \sum_{\tau_{C_1}, \tau_{C_2}, \tau_I} Z_{G_1}^\pi(\tau_{T(K)}, \tau_{C_1}) \left(Z_{G_2}^+(\tau_{C_2}, \tau_I) w(\tau_{C_1}, \tau_{C_2}) + Z_{G_2}^-(\tau_{C_2}, \tau_I) w(\tau_{C_1}, \tau_{C_2}) \right) \\ &= \sum_{\tau_{C_1}, \tau_{C_2}} Z_{G_1}^\pi(\tau_{T(K)}, \tau_{C_1}) w(\tau_{C_1}, \tau_{C_2}) \left(Z_{G_2}^+(\tau_{C_2}) + Z_{G_2}^-(\tau_{C_2}) \right), \end{aligned} \quad (17)$$

where $w(\tau_{C_1}, \tau_{C_2})$ denote the contribution from edges of P given configurations τ_{C_1} and τ_{C_2} , and $Z_K^\pi = \sum_{\tau_{T(K)}} Z_K^\pi(\tau_{T(K)})$. Moreover by Eq. (4), for $i = 1, 2$ and any subset $S \subseteq T(G_i)$, we have

$$(1 - \varepsilon') Q^\pi(\tau_S) Z_{G_i}^\pi \leq Z_{G_i}^\pi(\tau_S) \leq (1 + \varepsilon') Q^\pi(\tau_S) Z_{G_i}^\pi,$$

where we have used $Q^\pi(\tau_S)$ to denote the probability that the configuration on terminals in S is τ_S in the distribution Q^π . Therefore by Eq. (17),

$$Z_K^+(\tau_{T(K)}) \leq (1 + \varepsilon')^2 Q^+(\tau_{T(K)}) Z_{G_1}^+ Z_{G_2}^- \sum_{\tau_{C_1}, \tau_{C_2}} Q^+(\tau_{C_1}) w(\tau_{C_1}, \tau_{C_2}) \left(\frac{Z_{G_2}^+}{Z_{G_2}^-} Q^+(\tau_{C_2}) + Q^-(\tau_{C_2}) \right). \quad (18)$$

We need to calculate the quantity

$$\mu(\pi_1, \pi_2) := \sum_{\tau_{C_1}, \tau_{C_2}} Q^{\pi_1}(\tau_{C_1}) Q^{\pi_2}(\tau_{C_2}) w(\tau_{C_1}, \tau_{C_2}). \quad (19)$$

Recall our definitions of M and M^+ . Let $N = M^+M(M^+)^T$ where T means transposition. Then $\det(N) = (q^+ - q^-)^2(\beta\gamma - 1) < 0$. We write

$$N = \begin{pmatrix} N_{--} & N_{-+} \\ N_{+-} & N_{++} \end{pmatrix}$$

and let $c = \frac{N_{++}N_{--}}{N_{+-}N_{-+}} < 1$. Here $N_{\pi_1\pi_2}$ is the edge contribution when one end point is chosen with probability q^{π_1} and the other q^{π_2} . Also notice that each edge is independent under Q^\pm so we can count them separately. Then the quantity in Eq.(19) is

$$\begin{aligned} \mu(+, +) &= \mu(-, -) = (N_{++}N_{--})^{t'}; \\ \mu(+, -) &= \mu(-, +) = (N_{+-}N_{-+})^{t'}. \end{aligned} \quad (20)$$

Plug Eq. (20) in Eq. (18),

$$\begin{aligned} Z_K^+(\tau_{T(K)}) &\leq (1 + \varepsilon')^2 Z_{G_1}^+ Z_{G_2}^- \left(\frac{Z_{G_2}^+}{Z_{G_2}^-} (N_{++}N_{--})^{t'} + (N_{+-}N_{-+})^{t'} \right) \cdot Q^+(\tau_{T(K)}) \\ &= (1 + \varepsilon')^2 Z_{G_1}^+ Z_{G_2}^- (N_{+-}N_{-+})^{t'} \left(\frac{Z_{G_2}^+}{Z_{G_2}^-} c^{t'} + 1 \right) \cdot Q^+(\tau_{T(K)}). \end{aligned} \quad (21)$$

Summing over $\tau_{T(K)}$ in Eq. (21) we get

$$Z_K^+ \leq (1 + \varepsilon')^2 Z_{G_1}^+ Z_{G_2}^- (N_{+-}N_{-+})^{t'} \left(\frac{Z_{G_2}^+}{Z_{G_2}^-} c^{t'} + 1 \right). \quad (22)$$

Similarly we get an estimate for Z_K^- :

$$Z_K^- \geq (1 - \varepsilon')^2 Z_{G_1}^- Z_{G_2}^+ (N_{+-}N_{-+})^{t'} \left(\frac{Z_{G_2}^-}{Z_{G_2}^+} c^{t'} + 1 \right). \quad (23)$$

Let $r = \frac{Z_{G_2}^-}{Z_{G_2}^+}$. Notice that $Z_{G_1}^\pi = Z_{G_2}^\pi$ as G_1 and G_2 are identical copies. Combine Eq. (22) and Eq. (23),

$$\frac{Z_K^-}{Z_K^+} \geq \left(\frac{1 - \varepsilon'}{1 + \varepsilon'} \right)^2 \cdot \frac{1 + c^{t'} r}{1 + c^{t'}/r}. \quad (24)$$

By Eq. (3) there is $f(t + t', \varepsilon')$ such that

$$\frac{1}{f(t + t', \varepsilon')} \leq r \leq f(t + t', \varepsilon'),$$

and $f(t + t', \varepsilon')$ is bounded above by a polynomial in $t + t'$ and $1/\varepsilon'$. To show Eq. (5) it suffices to show $\frac{Z_K^-}{Z_K^+} \geq \frac{1-\varepsilon}{1+\varepsilon}$ and $\frac{Z_K^+}{Z_K^-} \geq \frac{1-\varepsilon}{1+\varepsilon}$. Clearly $\frac{Z_K^-}{Z_K^+} \geq \frac{1-\varepsilon}{1+\varepsilon}$ can be achieved by picking $\varepsilon' = \frac{\varepsilon}{3}$ and $t' = O(\log(t + \varepsilon^{-1}))$ in Eq. (24). To show $\frac{Z_K^+}{Z_K^-} \geq \frac{1-\varepsilon}{1+\varepsilon}$ is similar and therefore omitted.

Establishing that Eq. (4) still holds is easy to show. We will show

$$\frac{Z_K^+(\tau_{T(K)})}{Z_K^+} \geq (1 - \varepsilon)Q^+(\tau_{T(K)}),$$

and the other bounds are similar. By an argument similar to Eq. (18), we have

$$Z_K^+(\tau_{T(K)}) \geq (1 - \varepsilon')^2 Q^+(\tau_{T(K)}) Z_{G_1}^+ Z_{G_2}^- \sum_{\tau_{C_1}, \tau_{C_2}} Q^+(\tau_{C_1}) w(\tau_{C_1}, \tau_{C_2}) \left(\frac{Z_{G_2}^+}{Z_{G_2}^-} Q^+(\tau_{C_2}) + Q^-(\tau_{C_2}) \right).$$

Moreover, summing over $\tau_{T(K)}$ in Eq. (18) we get:

$$Z_K^+ \leq (1 + \varepsilon')^2 Z_{G_1}^+ Z_{G_2}^- \sum_{\tau_{C_1}, \tau_{C_2}} Q^+(\tau_{C_1}) w(\tau_{C_1}, \tau_{C_2}) \left(\frac{Z_{G_2}^+}{Z_{G_2}^-} Q^+(\tau_{C_2}) + Q^-(\tau_{C_2}) \right).$$

The desired bound follows as $\varepsilon' = \frac{\varepsilon}{3}$.

The other case is ferromagnetic, that is, $\beta\gamma > 1$. We construct K in the same way as in the previous case, but with the following change. To form the perfect matching P , we match $+/-$ terminals of G_1 to $-/+$ terminals of G_2 . The proof goes similarly but $\det(N) = (q^+ - q^-)^2(\beta\gamma - 1) > 0$. However since we made a twist in connecting G_1 and G_2 , it follows that

$$\begin{aligned} \mu(+, +) &= \mu(-, -) = (N_{+-} N_{-+})^{t'}; \\ \mu(+, -) &= \mu(-, +) = (N_{--} N_{++})^{t'}. \end{aligned}$$

Therefore we continue with the new constant $c' = \frac{N_{+-} N_{-+}}{N_{++} N_{--}} < 1$ and the rest of the proof is the same. \square

6. Unary Symmetry Breaking

In this section we prove Lemma 9 that almost all 2-spin models support unary symmetry breaking.

Proof of Lemma 9. Consider the sequence of gadgets $(H_k : k \in \mathbb{N})$, defined as follows. The vertex set of H_k is $V(H_k) = \{u, u', u'', v_1, v_2, \dots, v_k\}$, and u is considered the attachment vertex. The edge set of H_k is

$$E(H_k) = \{\{u', v_i\} : 1 \leq i \leq k\} \cup \{\{v_i, u''\} : 1 \leq i \leq k\} \cup \{\{u'', u\}\}.$$

We shall argue that if the first three graphs in the sequence, namely H_0 , H_1 and H_2 , all fail to be symmetry breakers then one of conditions (i) or (ii) holds. The graph H_0 has an isolated vertex that could clearly be removed; we leave it in to make the calculations uniform. Note that the maximum degree of any vertex in these graphs is 3.

Let $a = \beta^2 + \lambda$, $b = \beta + \lambda\gamma$ and $c = 1 + \lambda\gamma^2$. Then the effective weight of vertex u is given by the column vector,

$$\rho^k = \begin{pmatrix} 1 & 0 \\ 0 & \lambda \end{pmatrix} \begin{pmatrix} \beta & 1 \\ 1 & \gamma \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & \lambda \end{pmatrix} \begin{pmatrix} a^k & b^k \\ b^k & c^k \end{pmatrix} \begin{pmatrix} 1 \\ \lambda \end{pmatrix} = T \begin{pmatrix} a^k + \lambda b^k \\ b^k + \lambda c^k \end{pmatrix}$$

where

$$T = \begin{pmatrix} \beta & \lambda \\ \lambda & \lambda^2\gamma \end{pmatrix}.$$

For h_k to be a symmetry breaker we require the vector ρ^k not to be a multiple of $(1, \lambda)^T$.

Suppose ρ^0 , ρ^1 and ρ^2 all fail to be symmetry breakers. Then they must all lie in a one-dimensional subspace of \mathbb{R}^2 . One way this can happen is if the matrix T is rank 1, i.e., if $\beta\gamma = 1$. This is case (i).

Otherwise, $(1 + \lambda, 1 + \lambda)^T$ and $(a + \lambda b, b + \lambda c)^T$ and $(a^2 + \lambda b^2, b^2 + \lambda c^2)^T$ lie in a one-dimensional subspace, namely the one generated by $(1, 1)^T$. This implies $a + \lambda b = b + \lambda c$ and $a^2 + \lambda b^2 = b^2 + \lambda c^2$, or recasting,

$$\begin{aligned} a - b &= \lambda(c - b) \\ (a - b)(a + b) &= \lambda(c - b)(c + b). \end{aligned}$$

So either $a = b = c$, or (dividing the second equation by the first) $a = c$ and $\lambda = 1$. In either case, substituting for a , b and c in terms of β , γ , λ , we obtain either $\beta = \gamma = 1$ (which belongs to case (i)) or $\beta = \gamma$ and $\lambda = 1$ (which is case (ii)). \square

There are two exceptional cases. The first is $\beta\gamma = 1$. It is well-known that in this case the 2-spin system can be decomposed and hence tractable. The other case of $\beta = \gamma$ and $\lambda = 1$ is perfectly symmetric and this symmetry cannot be broken. This system is the Ising model without external fields. For this system, the marginal probability of any vertex in any graph is exactly $1/2$.

Regarding the ferromagnetic case, Jerrum and Sinclair [21] presented an FPRAS for the ferromagnetic Ising model with consistent external fields. On the other hand, anti-ferromagnetic Ising models without external field on bipartite graphs are actually equivalent to ferromagnetic Ising models. The trick is to flip the assignment on only one part of the bipartition, which has been used before by Goldberg and Jerrum [16].

Lemma 12. *For $0 < \alpha < 1$, $\text{BI-2-SPIN}(\alpha, \alpha, 1) \equiv_T \text{BI-2-SPIN}(1/\alpha, 1/\alpha, 1)$.*

Proof. Let $B = (V_1, V_2, E)$ be a bipartite graph where V_1 and V_2 are two partitions of vertices. Let $|E| = m$. Then we have

$$\begin{aligned} Z_B(\alpha, \alpha, 1) &= \sum_{\sigma_{V_1}: V_1 \rightarrow \{0,1\}} \sum_{\sigma_{V_2}: V_2 \rightarrow \{0,1\}} \prod_{(v,w) \in E} \alpha^{(1-\sigma_{V_1}(v))(1-\sigma_{V_2}(w))} \alpha^{\sigma_{V_1}(v)\sigma_{V_2}(w)} \\ &= \sum_{\sigma_{V_1}: V_1 \rightarrow \{0,1\}} \sum_{\sigma_{V_2}: V_2 \rightarrow \{0,1\}} \prod_{(v,w) \in E} \alpha^{(1-\sigma_{V_1}(v))\sigma_{V_2}(w)} \alpha^{\sigma_{V_1}(v)(1-\sigma_{V_2}(w))} \\ &= \sum_{\sigma_{V_1}: V_1 \rightarrow \{0,1\}} \sum_{\sigma_{V_2}: V_2 \rightarrow \{0,1\}} \prod_{(v,w) \in E} \alpha^{\sigma_{V_2}(w) - \sigma_{V_1}(v)\sigma_{V_2}(w) + \sigma_{V_1}(v) - \sigma_{V_1}(v)\sigma_{V_2}(w)} \\ &= \alpha^m \sum_{\sigma_{V_1}: V_1 \rightarrow \{0,1\}} \sum_{\sigma_{V_2}: V_2 \rightarrow \{0,1\}} \prod_{(v,w) \in E} \alpha^{-(1-\sigma_{V_1}(v))(1-\sigma_{V_2}(w)) - \sigma_{V_1}(v)\sigma_{V_2}(w)} \\ &= \alpha^m Z_B(\alpha^{-1}, \alpha^{-1}, 1), \end{aligned}$$

where in the second line we flip the assignment of σ_{V_2} . \square

Combining Lemma 12 with the FPRAS by Jerrum and Sinclair [21] for the ferromagnetic Ising model yields the following corollary.

Corollary 13. *For any $\alpha > 0$, $\text{BI-2-SPIN}(\alpha, \alpha, 1)$ has an FPRAS.*

As explained in the introduction, this corollary helps explain why the notion of unary symmetry breaking is necessary to achieve $\#\text{BIS}$ -hardness.

7. Nearly-independent Phases for Antiferromagnetic Systems

In this section we prove Lemma 7 that for 2-spin antiferromagnetic systems in the tree non-uniqueness region there is a gadget with nearly-independent phase-correlated spins.

We first give some necessary background which details how the values q^- and q^+ are derived.

7.1. Gibbs Measures on Infinite Trees

Recall that the Gibbs distribution is the distribution in which a configuration σ is drawn with probability

$$\Pr_{G;\beta,\gamma,\lambda}(\sigma) = \frac{w(\sigma)}{Z_G(\beta, \gamma, \lambda)}. \quad (25)$$

Let \mathbb{T}_Δ be the infinite $(\Delta - 1)$ -ary tree. A *Gibbs measure* on \mathbb{T}_Δ is a measure such that for any finite subtree $T \subset \mathbb{T}_\Delta$, the induced distribution on T conditioned on the outer boundary is the same as that given by (25). There may be one or more Gibbs measures (see, e.g., [15] for more details). A Gibbs measure is called *translation-invariant* if it is invariant under all automorphisms of \mathbb{T}_Δ , and is *semi-translation-invariant* if it is invariant under all parity-preserving automorphisms of \mathbb{T}_Δ . In our context, the Gibbs measures that will be of interest are the two extremal semi-translation-invariant Gibbs measures corresponding to the all 1's and all 0's boundary conditions. These two measures are different in the non-uniqueness region of \mathbb{T}_Δ . Let $0 < q^- < q^+ < 1$ be the two marginal probabilities of the root having spin 1 in the two extremal semi-translation-invariant Gibbs measures. More precisely, one can define q^+, q^- as follows. For $s \in \{0, 1\}$, let $q_{\ell,s}$ denote the probability that the root is assigned spin 1 in the $(\Delta - 1)$ -ary tree of depth ℓ in the Gibbs distribution where the leaves are fixed to spin s . In standard terminology, fixing the configuration on the leaves to all 1's or all 0's is most commonly referred to as the $+, -$ boundary conditions, respectively (and hence the indices $+, -$ in our notation of q^+, q^-). It is not hard to show that $q_{2\ell,1}$ is decreasing in ℓ , while $q_{2\ell,0}$ is increasing. Let $q^+ = \lim_{\ell \rightarrow \infty} q_{2\ell,1}$ and let $q^- = \lim_{\ell \rightarrow \infty} q_{2\ell,0}$. The two quantities q^+ and q^- satisfy the standard tree recursion in the following sense. Let $r^\pi = \frac{q^\pi}{1-q^\pi}$ for $\pi \in \{-, +\}$, and $f(x) = \lambda \left(\frac{1+\beta x}{\gamma+x} \right)^{\Delta-1}$. Then

$$r^+ = f(r^-), \text{ and } r^- = f(r^+).$$

7.2. Sly's Gadget

Sly ([28], Theorem 2.1) showed that for every $\Delta \geq 3$, there exists ε_Δ such that the hard-core model $(1, 0, \lambda, \Delta)$ supports nearly-independent phase-correlated spins for any λ satisfying $\lambda_c(\mathbb{T}_\Delta) < \lambda < \lambda_c(\mathbb{T}_\Delta) + \varepsilon_\Delta$. This region is a small interval just above the uniqueness threshold. In the same paper Sly also showed that $(1, 0, 1, 6)$ supports nearly-independent phase-correlated spins. The quantities defining measures Q^\pm in Definition 4 are exactly the marginal probabilities q^\pm on \mathbb{T}_Δ .

Galanis et al. [11] extended the applicable region of Sly's gadget for the hard-core model to all $\lambda > \lambda_c(\mathbb{T}_\Delta)$ for $\Delta = 3$ and $\Delta \geq 6$. The gap of $\Delta = 4$ and $\Delta = 5$ is later closed [12]. In the later paper Galanis et al. [12] also verified Sly's gadget for parameters $(\beta, \gamma, \lambda, \Delta)$ such that $\beta\gamma < 1$, $\sqrt{\beta\gamma} \geq \frac{\sqrt{\Delta-1}-1}{\sqrt{\Delta-1}+1}$, and (β, γ, λ) is in the non-uniqueness region of infinite tree \mathbb{T}_Δ . Using [13] we extend the applicable region of Sly's gadget to the entire non-uniqueness region for all 2-spin antiferromagnetic models.

Sly's construction is based on a technical analysis of the configurations with the largest contribution to the partition function of a random bipartite Δ -regular graph. This required a second moment argument involving a optimization which posed technical difficulties. Following recent developments in [13], we describe how to obtain the properties of Sly's gadget for all 2-spin antiferromagnetic systems (including the case with an external field) in the non-uniqueness region. We will give an overview of the second moment approach and how to tackle it using the technique in [13]. We begin by presenting the relevant definitions and describing the construction of the gadget.

For integers r, n , let \mathcal{G}_n^r be the following graph distribution:

1. \mathcal{G}_n^r is supported on bipartite graphs. The two parts of the bipartite graph are labelled by $+, -$ and each is partitioned as $V^\pi := U^\pi \cup W^\pi$ where $|U^\pi| = n$, $|W^\pi| = r$ for $\pi = \{+, -\}$. 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) $(\Delta - 1)$ random perfect matchings between $U^+ \cup W^+$ and $U^- \cup W^-$, (ii) a random perfect 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 $\Delta - 1$.

The case $r = 0$ will also be important, in which case we denote the distribution as \mathcal{G}_n . Note that \mathcal{G}_n is supported on bipartite Δ -regular graphs. Strictly speaking, $\mathcal{G}_n^r, \mathcal{G}_n$ are supported on multi-graphs, but it is well known that every statement that holds asymptotically almost surely on these spaces continues to hold asymptotically almost surely conditioned on the event that the graph is simple.

We are now ready to give the construction of the gadget. For constants $0 < \theta, \psi < 1/8$, let $m' = (\Delta - 1)^{\lfloor \theta \log_{\Delta-1} n \rfloor + 2 \lfloor \frac{\psi}{2} \log_{\Delta-1} n \rfloor}$. Note that $m' = o(n^{1/4})$. First, sample G from the distribution $\mathcal{G}_n^{m'}$ conditioning on G being simple. Next, for $\pi \in \{+, -\}$, attach t disjoint $(\Delta - 1)$ -ary trees of even depth ℓ (with $t = (\Delta - 1)^{\lfloor \theta \log_{\Delta-1} n \rfloor}$ and $\ell = 2 \lfloor \frac{\psi}{2} \log_{\Delta-1} n \rfloor$) to W^π , so that every vertex in W^π is a leaf of exactly one tree (this is possible since $m' = |W| = t(\Delta - 1)^\ell$). Denote by T^π the roots of the trees, so that $|T^\pi| = t$. The trees do not share common vertices with the graph G , apart from the vertices in W . The final graph \tilde{G} is the desired gadget, where the terminals T are the roots of the trees. We denote the family of graphs that can be constructed this way by $\tilde{\mathcal{G}}(t, n(t), \Delta)$. Note that the size of the construction is $(2 + o(1))n$ which is bounded above by a polynomial in t when Δ is a fixed constant. Moreover, any \tilde{G} drawn from $\tilde{\mathcal{G}}(t, n(t), \Delta)$ is bipartite. The terminals of \tilde{G} are $T(\tilde{G}) = T^+ \cup T^-$, and T^+ and T^- are on distinct partitions of the bipartite graph.

Next we show that if β, γ, λ lie in the non-uniqueness regime of \mathbb{T}_Δ , then the gadget \tilde{G} satisfies Definition 4. Lemma 19 in [12] shows that this is true, assuming that a certain technical condition is true [12, Condition 1]. We will introduce this condition and then show that it always holds when β, γ, λ lie in the non-uniqueness regime of \mathbb{T}_Δ .

The technical condition involves asymptotics of the leading terms of the first and second moments. We will derive expressions for these shortly. Moreover, the condition is stated for the case when $r = 0$, as $r = o(n^{1/4})$ is relatively small and [12] shows that it does not affect the leading terms. Suppose that $G \sim \mathcal{G}_n$. We look at the contribution of configurations which assign a prescribed fraction of vertices in U^+, U^- the spin 1, with a view to identifying the phases of the gadget. Here $V^+ = U^+$ and $V^- = U^-$ and $n = |V^+| = |V^-|$. For $1 \geq \chi^+, \chi^- \geq 0$, let

$$\Sigma^{\chi^+, \chi^-} := \{\sigma : V^+ \cup V^- \rightarrow \{0, 1\} \mid |\sigma^{-1}(1) \cap U^+| = \lceil \chi^+ n \rceil, |\sigma^{-1}(1) \cap U^-| = \lceil \chi^- n \rceil, \}$$

where $\sigma^{-1}(1)$ denotes the set of vertices assigned the spin 1 in the configuration σ . Denote by $Z_G^{\chi^+, \chi^-}$ the total contribution to the partition function of G by the configurations in Σ^{χ^+, χ^-} , that is,

$$Z_G^{\chi^+, \chi^-} = \sum_{\sigma \in \Sigma^{\chi^+, \chi^-}} w_G(\sigma).$$

We will calculate the leading term in the exponent (as a function of n) of the first and second moment of $Z_G^{\chi^+, \chi^-}$.

By linearity we have $\mathbf{E}_{\mathcal{G}_n} [Z_G^{\chi^+, \chi^-}] = \sum_{\sigma \in \Sigma^{\chi^+, \chi^-}} \mathbf{E}_{\mathcal{G}_n} [w_G(\sigma)]$. Fix $\sigma \in \Sigma^{\chi^+, \chi^-}$. For convenience, set $\chi_n^+ = \lceil n\chi^+ \rceil / n$, $\chi_n^- = \lceil n\chi^- \rceil / n$. Let nx_n be the number of $(1, 1)$ edges in a random perfect matching

between V^+, V^- under the configuration σ . Note that $\chi_n \leq x_n \leq \chi'_n$, where $\chi_n = \max\{0, \chi_n^+ + \chi_n^- - 1\}$ and $\chi'_n = \min\{\chi_n^+, \chi_n^-\}$. It is not hard to see that

$$\mathbf{E}_{\mathcal{G}_n}[w_G(\sigma)] = \lambda^{n(\chi_n^+ + \chi_n^-)} \left(\sum_{\chi_n \leq x_n \leq \chi'_n} \kappa_{x_n}^{\chi_n^+, \chi_n^-} \right)^\Delta, \text{ where } \kappa_{x_n}^{\chi_n^+, \chi_n^-} = \frac{\binom{n\chi_n^+}{nx_n} \binom{n(1-\chi_n^+)}{n(\chi_n^- - x_n)}}{\binom{n}{n\chi_n^-}} \beta^{n(1-\chi_n^+ - \chi_n^- + x_n)} \gamma^{nx_n}. \quad (26)$$

By the linearity of expectation and symmetry, we then have

$$\mathbf{E}_{\mathcal{G}_n}[Z_G^{\chi^+, \chi^-}] = |\Sigma^{\chi^+, \chi^-}| \mathbf{E}_{\mathcal{G}_n}[w_G(\sigma)], \text{ where } |\Sigma^{\chi^+, \chi^-}| = \binom{n}{n\chi_n^+} \binom{n}{n\chi_n^-}. \quad (27)$$

Let $\chi := \max\{0, \chi^+ + \chi^- - 1\}$ and $\chi' := \min\{\chi^+, \chi^-\}$. Note that as $n \rightarrow \infty$, $(\chi_n^+, \chi_n^-, \chi_n, \chi'_n) \rightarrow (\chi^+, \chi^-, \chi, \chi')$. In Section 7.2.1, using Stirling's approximation, it is shown that

$$\forall \chi^+, \chi^-, \quad \frac{1}{n} \log \mathbf{E}_{\mathcal{G}_n}[Z_G^{\chi^+, \chi^-}] = \Psi_{1;\beta,\gamma,\lambda}(\chi^+, \chi^-) + o(1), \quad (28)$$

where, under the convention $0 \log 0 \equiv 0$,

$$\begin{aligned} \Psi_{1;\beta,\gamma,\lambda}(\chi^+, \chi^-) &:= \max_{\chi \leq x \leq \chi'} \Psi'_{1;\beta,\gamma,\lambda}(x, \chi^+, \chi^-), \\ \Psi'_{1;\beta,\gamma,\lambda}(x, \chi^+, \chi^-) &:= (\chi^+ + \chi^-) \log \lambda + (\Delta - 1) f_1(\chi^+, \chi^-) + \Delta g_{1;\beta,\gamma}(x, \chi^+, \chi^-), \\ f_1(\chi^+, \chi^-) &:= \chi^+ \log \chi^+ + (1 - \chi^+) \log(1 - \chi^+) + \chi^- \log \chi^- + (1 - \chi^-) \log(1 - \chi^-), \\ g_{1;\beta,\gamma}(x, \chi^+, \chi^-) &:= (1 - \chi^+ - \chi^- + x) \log \beta + x \log \gamma \\ &\quad - x \log x - (\chi^+ - x) \log(\chi^+ - x) - (\chi^- - x) \log(\chi^- - x) \\ &\quad - (1 - \chi^+ - \chi^- + x) \log(1 - \chi^+ - \chi^- + x). \end{aligned} \quad (29)$$

Note that all the quantities in (29) are independent of n .

To calculate the second moment, the approach is completely analogous though we need to introduce more variables. For $1 \geq v^+, v^- \geq 0$, let

$$\Sigma_{v^+, v^-}^{\chi^+, \chi^-} := \left\{ (\sigma_1, \sigma_2) \mid \sigma_1, \sigma_2 \in \Sigma^{\chi^+, \chi^-}, |\sigma_1^{-1}(1) \cap \sigma_2^{-1}(1) \cap U^+| = \lceil v^+ n \rceil, |\sigma_1^{-1}(1) \cap \sigma_2^{-1}(1) \cap U^-| = \lceil v^- n \rceil \right\}. \quad (30)$$

Let $Y_{v^+, v^-}^{\chi^+, \chi^-} = \sum_{(\sigma_1, \sigma_2) \in \Sigma_{v^+, v^-}^{\chi^+, \chi^-}} w_G(\sigma_1) w_G(\sigma_2)$. Notice that $Y_{v^+, v^-}^{\chi^+, \chi^-}$ is the contribution to $(Z_G^{\chi^+, \chi^-})^2$ from pairs of configurations in Σ^{χ^+, χ^-} with overlap $\lceil v^+ n \rceil$ and $\lceil v^- n \rceil$. Then, we have

$$\mathbf{E}_{\mathcal{G}_n} \left[(Z_G^{\chi^+, \chi^-})^2 \right] = \sum_{v^+, v^-} \mathbf{E}_{\mathcal{G}_n} \left[Y_{v^+, v^-}^{\chi^+, \chi^-} \right], \quad (31)$$

where the sum is over integer multiples v^+ and v^- of $1/n$. Later we will replace this sum with a maximum and consider general real numbers v^+ and v^- . To unify the notation in both cases, set $v_n^+ = \lceil nv^+ \rceil / n$ and $v_n^- = \lceil nv^- \rceil / n$. Note that the overlap v_n^+ is at most $\chi_n^+ n$ and at least $2\chi_n^+ n - n$ and similarly for the minus side.

To get an expression for $\mathbf{E}_{\mathcal{G}_n} \left[Y_{v^+, v^-}^{\chi^+, \chi^-} \right]$, fix arbitrary $(\sigma_1, \sigma_2) \in \Sigma_{v^+, v^-}^{\chi^+, \chi^-}$. The configurations σ_1 and σ_2 divide each of U^+ and U^- into 4 subsets. Let the subsets chosen for spin 1 by both σ_1 and σ_2 be numbered 1 on the + side and the - side. Similarly, let the subsets chosen for spin 1 in σ_1 and for

spin 0 in σ_2 be numbered 2 and the subsets chosen for spin 0 in σ_1 and spin 1 in σ_2 be numbered 3. Finally, let the subsets chosen for spin 0 in both σ_1 and σ_2 be numbered 4. Let ny_{ijn} be the number of edges between subset i on one side and subset j on the other, and $\mathbf{y}_n = \{y_{ijn}\}$. Define

$$\begin{aligned} L_{1n} &= v_n^+, & L_{2n} &= L_{3n} = \chi_n^+ - v_n^+, & L_{4n} &= 1 - 2\chi_n^+ + v_n^+, \\ R_{1n} &= v_n^-, & R_{2n} &= R_{3n} = \chi_n^- - v_n^-, & R_{4n} &= 1 - 2\chi_n^- + v_n^-. \end{aligned}$$

Observe that the \mathbf{y}_n satisfy

$$\begin{aligned} \sum_j y_{ijn} &= L_{in} \quad \text{for } i \in [4], & \sum_i y_{ijn} &= R_{jn}, \quad \text{for } j \in [4], \\ y_{ijn} &\geq 0 \quad \text{and } y_{ijn} \text{ is an integral multiple of } 1/n. \end{aligned} \tag{32}$$

Then we have

$$\mathbf{E}_{\mathcal{G}_n}[w(\sigma_1)w(\sigma_2)] = \lambda^{2n(\chi_n^+ + \chi_n^-)} \left(\sum_{\mathbf{y}_n} \xi_{v_n^+, v_n^-, \mathbf{y}_n}^{\chi_n^+, \chi_n^-} \right)^\Delta, \tag{33}$$

where

$$\begin{aligned} \xi_{v_n^+, v_n^-, \mathbf{y}_n}^{\chi_n^+, \chi_n^-} &:= \binom{n}{ny_{11n}, ny_{12n}, \dots, ny_{44n}}^{-1} \prod_{i=1}^4 \binom{nL_{in}}{ny_{i1n}, ny_{i2n}, ny_{i3n}, ny_{i4n}} \prod_{j=1}^4 \binom{nR_{jn}}{ny_{1jn}, ny_{2jn}, ny_{3jn}, ny_{4jn}} \\ &\times \beta^{n(y_{22n} + y_{24n} + y_{33n} + y_{34n} + y_{42n} + y_{43n} + 2y_{44n})} \gamma^{(2y_{11n} + y_{12n} + y_{13n} + y_{21n} + y_{31n} + y_{22n} + y_{33n})n} \end{aligned} \tag{34}$$

To see why (33) is true note that a random matching may be obtained by fixing the permutation within each of the 16 subsets y_{ij} of the “−” side and then choosing the 16 corresponding blocks of the “+” side giving the denominator $n! = \binom{n}{ny_{11n}, ny_{12n}, \dots, ny_{44n}} \prod_{i,j} (ny_{ijn})!$. Then the numerator counts all matchings consistent with σ_1 and σ_2 giving $\prod_{i=1}^4 \binom{nL_{in}}{ny_{i1n}, ny_{i2n}, ny_{i3n}, ny_{i4n}} \prod_{j=1}^4 \binom{nR_{jn}}{ny_{1jn}, ny_{2jn}, ny_{3jn}, ny_{4jn}} \prod_{i,j} (ny_{ijn})!$. The factors $\prod_{i,j} (ny_{ijn})!$ cancel, giving the first line in (34). The second line is just $w(\sigma_1)w(\sigma_2)$.

To account for the cardinality of $\Sigma_{v^+, v^-}^{\chi^+, \chi^-}$, by the linearity of expectation and symmetry, we obtain

$$\mathbf{E}_{\mathcal{G}_n} \left[Y_{v^+, v^-}^{\chi^+, \chi^-} \right] = \binom{n}{nL_{1n}, nL_{2n}, nL_{3n}, nL_{4n}} \binom{n}{nR_{1n}, nR_{2n}, nR_{3n}, nR_{4n}} \mathbf{E}_{\mathcal{G}_n}[w(\sigma_1)w(\sigma_2)] \tag{35}$$

Note that as $n \rightarrow \infty$, $L_{in} \rightarrow L_i$, $R_{jn} \rightarrow R_j$ ($\forall i, j \in [4]$), where

$$\begin{aligned} L_1 &= v^+, & L_2 &= L_3 = \chi^+ - v^+, & L_4 &= 1 - 2\chi^+ + v^+, \\ R_1 &= v^-, & R_2 &= R_3 = \chi^- - v^-, & R_4 &= 1 - 2\chi^- + v^-. \end{aligned} \tag{36}$$

Using again Stirling’s approximation, in Section 7.2.1 it is shown that

$$\frac{1}{n} \log \mathbf{E}_{\mathcal{G}_n} \left[Y_{v^+, v^-}^{\chi^+, \chi^-} \right] = \Psi'_{2; \beta, \gamma, \lambda}(\chi^+, \chi^-, v^+, v^-) + o(1), \tag{37}$$

where

$$\begin{aligned}
\Psi'_{2;\beta,\gamma,\lambda}(\chi^+, \chi^-, v^+, v^-) &:= \max_{\mathbf{y}} \Psi''_{2;\beta,\gamma,\lambda}(\chi^+, \chi^-, v^+, v^-, \mathbf{y}); \\
\Psi''_{2;\beta,\gamma,\lambda}(\chi^+, \chi^-, v^+, v^-, \mathbf{y}) &:= 2(\chi^+ + \chi^-) \log \lambda + (\Delta - 1)f_2(\chi^+, \chi^-, v^+, v^-) + \Delta g_{2;\beta,\gamma,\lambda}(\mathbf{y}); \\
f_2(\chi^+, \chi^-, v^+, v^-) &:= 2(\chi^+ - v^+) \log(\chi^+ - v^+) + v^+ \log v^+ + (1 - 2\chi^+ + v^+) \log(1 - 2\chi^+ + v^+) \\
&\quad + 2(\chi^- - v^-) \log(\chi^- - v^-) + v^- \log v^- + (1 - 2\chi^- + v^-) \log(1 - 2\chi^- + v^-); \\
g_{2;\beta,\gamma}(\mathbf{y}) &:= (y_{22} + y_{24} + y_{33} + y_{34} + y_{42} + y_{43} + 2y_{44}) \log \beta \\
&\quad + (2y_{11} + y_{12} + y_{13} + y_{21} + y_{22} + y_{31} + y_{33}) \log \gamma - \sum_{ij} y_{ij} \log y_{ij},
\end{aligned}$$

where the maximum is over non-negative (real vectors) \mathbf{y} satisfying

$$\sum_j y_{ij} = L_i, \quad \text{for } i \in [4], \quad \sum_i y_{ij} = R_j, \quad \text{for } j \in [4]. \quad (38)$$

To obtain the leading term of the second moment $\mathbf{E}_{\mathcal{G}_n} \left[(Z_G^{\chi^+, \chi^-})^2 \right]$, define

$$\mathcal{D} := \left\{ (v^+, v^-) \mid \max\{0, 2\chi^+ - 1\} \leq v^+ \leq \chi^+ \text{ and } \max\{0, 2\chi^- - 1\} \leq v^- \leq \chi^- \right\}.$$

In Section 7.2.1, we will show

$$\forall \chi^+, \chi^-, \quad \frac{1}{n} \log \mathbf{E}_{\mathcal{G}_n} \left[(Z_G^{\chi^+, \chi^-})^2 \right] = \Psi_{2;\beta,\gamma,\lambda}(\chi^+, \chi^-) + o(1), \quad (39)$$

where

$$\Psi_{2;\beta,\gamma,\lambda}(\chi^+, \chi^-) := \max_{(v^+, v^-) \in \mathcal{D}} \Psi'_{2;\beta,\gamma,\lambda}(\chi^+, \chi^-, v^+, v^-).$$

We want to focus on (χ^+, χ^-) which maximize $\Psi_{1;\beta,\gamma,\lambda}$. It is shown in [12, Lemma 1] that when non-uniqueness holds, there exist p^+, p^- with $p^+ > p^-$ such that $(\chi^+, \chi^-) = (p^+, p^-)$ is the unique maximizer (up to symmetry) of $\Psi_{1;\beta,\gamma,\lambda}$. We should note here that the values p^+, p^- are the analogues of q^+, q^- for the infinite Δ -regular tree (observe that the infinite Δ -regular tree differs from \mathbb{T}_Δ only at the degree of the root). The technical condition from [12] is as follows.

Condition 1 ([12, Condition 1]). $\Psi'_{2;\beta,\gamma,\lambda}(p^+, p^-, v^+, v^-)$ is maximized at $v^+ = (p^+)^2$, $v^- = (p^-)^2$.

To see that our version of Condition 1 is equivalent to the version in [12], use (37). Our interest in Condition 1 is justified by the following lemma which is proved in [12].

Lemma 14 ([12, Lemma 19]). *When the parameter set (β, γ, λ) lies in the non-uniqueness region of \mathbb{T}_Δ , and Condition 1 holds, the gadget \tilde{G} satisfies Definition 4.*

We will establish Condition 1 when (β, γ, λ) lies in the non-uniqueness region of \mathbb{T}_Δ . To do so, we apply results from [13]. The next lemma is a specialization of Theorem 3 from [13] to our setting for the case that there is no field.

Lemma 15 ([13, Theorem 3]). *Suppose $\Delta \geq 3$. Then*

$$\max_{\chi^+, \chi^-} \Psi_{2;\beta,\gamma,1}(\chi^+, \chi^-) = 2 \max_{\chi^+, \chi^-} \Psi_{1;\beta,\gamma,1}(\chi^+, \chi^-)$$

The main difference between the setting in [13] and ours is that we allow external fields whereas they do not. We use the following lemma to apply Lemma 15 to our setting.

Lemma 16. *For all $\beta, \gamma, \lambda, \Delta, \chi^+, \chi^-, v^+$ and v^- ,*

$$\begin{aligned}\Psi_{1;\beta,\gamma,\lambda}(\chi^+, \chi^-) &= \Psi_{1;\beta/\lambda^{1/\Delta}, \gamma\lambda^{1/\Delta}, 1}(\chi^+, \chi^-) + \log \lambda, \\ \Psi'_{2;\beta,\gamma,\lambda}(\chi^+, \chi^-, v^+, v^-) &= \Psi'_{2;\beta/\lambda^{1/\Delta}, \gamma\lambda^{1/\Delta}, 1}(\chi^+, \chi^-, v^+, v^-) + 2 \log \lambda, \\ \Psi_{2;\beta,\gamma,\lambda}(\chi^+, \chi^-) &= \Psi_{2;\beta/\lambda^{1/\Delta}, \gamma\lambda^{1/\Delta}, 1}(\chi^+, \chi^-) + 2 \log \lambda\end{aligned}$$

Proof. For the first equation, note from the definition of g_1 that

$$g_{1;\beta,\gamma}(x, \chi^+, \chi^-) - g_{1;\beta/\lambda^{1/\Delta}, \gamma\lambda^{1/\Delta}}(x, \chi^+, \chi^-) = \log(\lambda)(1 - \chi^+ - \chi^-)/\Delta.$$

Thus, from the definition of Ψ_1 , we have

$$\Psi_{1;\beta,\gamma,\lambda}(\chi^+, \chi^-) - \Psi_{1;\beta/\lambda^{1/\Delta}, \gamma\lambda^{1/\Delta}, 1}(\chi^+, \chi^-) = \log \lambda.$$

Similarly, from the definition of Ψ''_2 , we have

$$\begin{aligned}\Psi''_{2;\beta,\gamma,\lambda}(\chi^+, \chi^-, v^+, v^-, \mathbf{y}) - \Psi''_{2;\beta/\lambda^{1/\Delta}, \gamma\lambda^{1/\Delta}, 1}(\chi^+, \chi^-, v^+, v^-, \mathbf{y}) &= 2(\chi^+ + \chi^-) \log \lambda \\ &\quad + \Delta(g_{2;\beta,\gamma}(\mathbf{y}) - g_{2;\beta/\lambda^{1/\Delta}, \gamma\lambda^{1/\Delta}}(\mathbf{y})),\end{aligned}$$

and from the definition of g_2 ,

$$\begin{aligned}\Delta(g_{2;\beta,\gamma}(\mathbf{y}) - g_{2;\beta/\lambda^{1/\Delta}, \gamma\lambda^{1/\Delta}}(\mathbf{y})) &= (y_{22} + y_{24} + y_{33} + y_{34} + y_{42} + y_{43} + 2y_{44} \\ &\quad - 2y_{11} - y_{12} - y_{13} - y_{21} - y_{22} - y_{31} - y_{33}) \log \lambda.\end{aligned}$$

However by (38)

$$\begin{aligned}(y_{22} + y_{24} + y_{33} + y_{34} + y_{42} + y_{43} + 2y_{44} - 2y_{11} - y_{12} - y_{13} - y_{21} - y_{22} - y_{31} - y_{33}) &= L_4 + R_4 - L_1 - R_1 \\ &= 2 - 2(\chi^+ + \chi^-).\end{aligned}$$

Therefore we have

$$\Psi''_{2;\beta,\gamma,\lambda}(\chi^+, \chi^-, v^+, v^-, \mathbf{y}) - \Psi''_{2;\beta/\lambda^{1/\Delta}, \gamma\lambda^{1/\Delta}, 1}(\chi^+, \chi^-, v^+, v^-, \mathbf{y}) = 2 \log \lambda.$$

As Ψ'_2 is the maximum of Ψ''_2 over \mathbf{y} and Ψ_2 is the maximum of Ψ'_2 over (v^+, v^-) , the second and the third equations hold. \square

Lemma 16 shows that modulo a constant term, the expressions Ψ_1 and Ψ_2 are preserved by the standard transformation on Δ -regular graphs whereby an external field λ is pushed into the edge interactions. Using this lemma, we can now draw conclusions about the maximisation of these quantities in our setting.

Lemma 17. *For any parameter set (β, γ, λ) and $\Delta \geq 3$, if (β, γ, λ) lies in the non-uniqueness region of \mathbb{T}_Δ , then $\Psi_{2;\beta,\gamma,\lambda}(p^+, p^-) = 2\Psi_{1;\beta,\gamma,\lambda}(p^+, p^-)$ and (p^+, p^-) maximizes both $\Psi_{1;\beta,\gamma,\lambda}$ and $\Psi_{2;\beta,\gamma,\lambda}$.*

Proof. First, using equalities (28) and (39), and the fact that $\mathbf{E}[X^2] \geq \mathbf{E}[X]^2$,

$$\begin{aligned}\Psi_{2;\beta,\gamma,\lambda}(p^+, p^-) &= \frac{1}{n} \log \mathbf{E}_{\mathcal{G}_n} \left[(Z_G^{p^+, p^-})^2 \right] + o(1) \\ &\geq \frac{1}{n} \log \left(\mathbf{E}_{\mathcal{G}_n} [Z_G^{p^+, p^-}] \right)^2 + o(1) \\ &= 2 \frac{1}{n} \log \mathbf{E}_{\mathcal{G}_n} [Z_G^{p^+, p^-}] + o(1) \\ &= 2\Psi_{1;\beta,\gamma,\lambda}(p^+, p^-) + o(1).\end{aligned}$$

But since $\Psi_{2;\beta,\gamma,\lambda}(p^+, p^-)$ and $\Psi_{1;\beta,\gamma,\lambda}(p^+, p^-)$ don't depend upon n , we have

$$\Psi_{2;\beta,\gamma,\lambda}(p^+, p^-) \geq 2\Psi_{1;\beta,\gamma,\lambda}(p^+, p^-).$$

For the other direction, we can apply Lemma 15 and Lemma 16 which together give

$$\begin{aligned}\max_{\chi^+, \chi^-} \Psi_{2;\beta,\gamma,\lambda}(\chi^+, \chi^-) &= \max_{\chi^+, \chi^-} \Psi_{2;\beta/\lambda^{1/\Delta}, \gamma\lambda^{1/\Delta}, 1}(\chi^+, \chi^-) + 2 \log \lambda \\ &= 2 \left(\max_{\chi^+, \chi^-} \Psi_{1;\beta/\lambda^{1/\Delta}, \gamma\lambda^{1/\Delta}, 1}(\chi^+, \chi^-) + \log \lambda \right) \\ &= 2 \max_{\chi^+, \chi^-} \Psi_{1;\beta,\gamma,\lambda}(\chi^+, \chi^-).\end{aligned}$$

Since (p^+, p^-) maximizes $\Psi_{1;\beta,\gamma,\lambda}$ when (β, γ, λ) lies in the non-uniqueness region of \mathbb{T}_Δ ([12, Lemma 1]), we have $\Psi_{2;\beta,\gamma,\lambda}(p^+, p^-) \leq 2\Psi_{1;\beta,\gamma,\lambda}(p^+, p^-)$. Therefore $\Psi_{2;\beta,\gamma,\lambda}(p^+, p^-) = 2\Psi_{1;\beta,\gamma,\lambda}(p^+, p^-)$, and (p^+, p^-) also maximizes $\Psi_{2;\beta,\gamma,\lambda}$. \square

By the same reasoning as in the proof above, we may apply other results from [13]. Recall that $\Psi_{2;\beta,\gamma,\lambda}(\chi^+, \chi^-) = \max_{v^+, v^-} \Psi'_{2;\beta,\gamma,\lambda}(\chi^+, \chi^-, v^+, v^-)$. Let

$$\mathcal{D}'(\chi^+, \chi^-) = \{(v^+, v^-) \in \mathcal{D}(\chi^+, \chi^-) \mid v^+ \neq (\chi^+)^2 \text{ or } v^- \neq (\chi^-)^2\}.$$

Lemma 18 ([13, Lemma 29]). *For all $0 \leq \chi^+ \leq 1$ and $0 \leq \chi^- \leq 1$ and all $(v^+, v^-) \in \mathcal{D}'(\chi^+, \chi^-)$, we have*

$$\Psi'_{2;\beta,\gamma,1}(\chi^+, \chi^-, v^+, v^-) < 2 \max_{0 \leq x \leq 1, 0 \leq y \leq 1} \Psi_{1;\beta,\gamma,1}(x, y).$$

To see that our version of Lemma 18 is the same as the one in [13], see [13, Equation (3)]. So using Lemma 16 we obtain

Corollary 19. *For all $0 \leq \chi^+ \leq 1$ and $0 \leq \chi^- \leq 1$ and all $(v^+, v^-) \in \mathcal{D}'(\chi^+, \chi^-)$, we have*

$$\Psi'_{2;\beta,\gamma,\lambda}(\chi^+, \chi^-, v^+, v^-) < 2 \max_{0 \leq x \leq 1, 0 \leq y \leq 1} \Psi_{1;\beta,\gamma,\lambda}(x, y).$$

Now we can show that Condition 1 holds when non-uniqueness holds.

Lemma 20. *For any parameter set (β, γ, λ) and $\Delta \geq 3$, if (β, γ, λ) lies in the non-uniqueness region of \mathbb{T}_Δ , then Condition 1 holds.*

Proof. From Corollary 19, we have that for all $0 \leq \chi^+ \leq 1$ and $0 \leq \chi^- \leq 1$ and all $(v^+, v^-) \in \mathcal{D}'(\chi^+, \chi^-)$,

$$\Psi'_{2;\beta,\gamma,\lambda}(\chi^+, \chi^-, v^+, v^-) < 2 \max_{0 \leq x \leq 1, 0 \leq y \leq 1} \Psi_{1;\beta,\gamma,\lambda}(x, y).$$

In particular, this holds for $\chi^+ = p^+$ and $\chi^- = p^-$. Together with Lemma 17, it gives us that for all $(v^+, v^-) \in \mathcal{D}'(p^+, p^-)$,

$$\begin{aligned} \Psi'_{2;\beta,\gamma,\lambda}(p^+, p^-, v^+, v^-) &< 2 \max_{0 \leq x \leq 1, 0 \leq y \leq 1} \Psi_{1;\beta,\gamma,\lambda}(x, y) \\ &= \Psi_{2;\beta,\gamma,\lambda}(p^+, p^-) \\ &= \max_{(v^+, v^-) \in \mathcal{D}'(\chi^+, \chi^-)} \Psi'_{2;\beta,\gamma,\lambda}(p^+, p^-, v^+, v^-). \end{aligned}$$

It implies that $\Psi'_{2;\beta,\gamma,\lambda}(p^+, p^-, v^+, v^-)$ is maximized at $v^+ = (p^+)^2, v^- = (p^-)^2$, which is Condition 1. \square

Lemma 7 follows directly.

Proof of Lemma 7. By Lemma 14, this lemma follows from Condition 1, and Condition 1 is established by Lemma 20. \square

7.2.1. Proof of Equations (28), (37) and (39)

The proof is a rather lengthy formalisation of the following main idea: as n grows large, the sums in the expressions for the first and second moments are dominated by their maximum terms and the integrality conditions for the variables may be dropped introducing only $o(1)$ (additive) error in the asymptotics of the logarithms.

We will use the following lemma. Recall that for a metric space (\mathcal{X}, d) , an ε -cover of a set $\mathcal{K} \subseteq \mathcal{X}$ is a set $\mathcal{C} \subseteq \mathcal{X}$ such that every point in \mathcal{K} is within distance ε from a point in \mathcal{C} . Note that we do *not* require $\mathcal{C} \subseteq \mathcal{K}$. Recall also that $x \in \mathcal{X}$ is a limit point of \mathcal{K} if x is in the closure of $\mathcal{K} \setminus \{x\}$.

Lemma 21. *Let $h : \mathcal{R}_h \rightarrow \mathbb{R}$ be a continuous function on $\mathcal{R}_h \subseteq \mathbb{R}^k$. Let \mathcal{R} be a compact subregion of \mathcal{R}_h . For $n = 1, 2, \dots$, let $\mathcal{C}_n \subseteq \mathcal{R}_h$ be an ε_n -cover of \mathcal{R} , where $\varepsilon_n \rightarrow 0$ as $n \rightarrow \infty$. Assume further that the \mathcal{C}_n are uniformly bounded and that the set of limit points of $\bigcup_n (\mathcal{C}_n \setminus \mathcal{R})$ is a subset of \mathcal{R} . Then, as $n \rightarrow \infty$, it holds that*

$$\sup_{\mathbf{x} \in \mathcal{C}_n} h(\mathbf{x}) \rightarrow \max_{\mathbf{x} \in \mathcal{R}} h(\mathbf{x}).$$

Proof. Let

$$L_n := \sup_{\mathbf{x} \in \mathcal{C}_n} h(\mathbf{x}), \quad L := \max_{\mathbf{x} \in \mathcal{R}} h(\mathbf{x}).$$

Note that the maximum in the definition of L is justified by compactness of the region \mathcal{R} and continuity of h . We will prove that

$$L \leq \liminf_{n \rightarrow \infty} L_n \leq \limsup_{n \rightarrow \infty} L_n \leq L. \quad (40)$$

We first prove the right-most inequality in (40). Assume for the sake of contradiction that $\limsup_{n \rightarrow \infty} L_n > L$. Then, there exist sequences n_m, \mathbf{x}_{n_m} with $n_m \rightarrow \infty$ and $\mathbf{x}_{n_m} \in \mathcal{C}_{n_m}$ such that $h(\mathbf{x}_{n_m}) > L + \varepsilon$ for some $\varepsilon > 0$. Since the \mathcal{C}_n are uniformly bounded, \mathbf{x}_{n_m} are bounded, and so they have a convergent subsequence, whose limit we denote by $\bar{\mathbf{x}}$. From the continuity of h , we obtain $h(\bar{\mathbf{x}}) \geq L + \varepsilon$. We claim that $\bar{\mathbf{x}} \in \mathcal{R}$, and thus obtain a contradiction to the choice of L . Indeed, if there exists m_0 such

that $\mathbf{x}_{n_m} \in \mathcal{R}$ for all $m \geq m_0$, we have $\bar{\mathbf{x}} \in \mathcal{R}$ by compactness of \mathcal{R} . Otherwise, by restricting to a subsequence of \mathbf{x}_{n_m} if necessary, $\bar{\mathbf{x}}$ is a limit point of $\bigcup_n (\mathcal{C}_n \setminus \mathcal{R})$, and hence $\bar{\mathbf{x}} \in \mathcal{R}$.

To show the left-most inequality in (40), let $\mathbf{x}^* \in \mathcal{R}$ be such that $h(\mathbf{x}^*) = L$. Since the \mathcal{C}_n are ε_n -covers of \mathcal{R} and $\varepsilon_n \rightarrow 0$, there exists a sequence \mathbf{x}_n with $\mathbf{x}_n \in \mathcal{C}_n$ such that $\mathbf{x}_n \rightarrow \mathbf{x}^*$. Note that $L_n \geq h(\mathbf{x}_n)$ and thus, from the continuity of h , we obtain $\liminf_{n \rightarrow \infty} L_n \geq h(\mathbf{x}^*) = L$. This finishes the proof of (40). \square

We will now prove Equations (28), (37) and (39). We start with the first moment $\mathbf{E}_{\mathcal{G}_n} [Z_G^{\chi^+, \chi^-}]$. Under the convention $0 \log 0 = 0$, $f_1(z, w)$ is continuous for $0 \leq z \leq 1$ and $0 \leq w \leq 1$. We will write $g_1(x, z, w)$ for $g_{1; \beta, \gamma}(x, z, w)$. The function $g_1(x, z, w)$ is continuous over the region

$$\mathcal{R}_{g_1} = \{(x, z, w) \mid x, z, w \geq 0, x \leq z \leq 1, x \leq w \leq 1, 0 \leq 1 - z - w + x\}. \quad (41)$$

The following asymptotic expressions are obtained from Stirling's approximation for $|\Sigma^{\chi^+, \chi^-}|$ and $\kappa_{x_n}^{\chi_n^+, \chi_n^-}$ using equations (26) and (27).

$$\frac{1}{n} \log |\Sigma^{\chi^+, \chi^-}| = -f_1(\chi_n^+, \chi_n^-) + o(1), \quad \frac{1}{n} \log \kappa_{x_n}^{\chi_n^+, \chi_n^-} = g_1(x_n, \chi_n^+, \chi_n^-) + f_1(\chi_n^+, \chi_n^-) + o(1). \quad (42)$$

Note that the sum over x_n in the expression (26) which is used in (27) for $\mathbf{E}_{\mathcal{G}_n} [Z_G^{\chi^+, \chi^-}]$ has at most n terms, so obviously the number of terms is at most a polynomial in n . Thus, we may approximate the sum by its maximum term to obtain

$$\frac{1}{n} \log \mathbf{E}_{\mathcal{G}_n} [Z_G^{\chi^+, \chi^-}] = \log(\lambda)(\chi_n^+ + \chi_n^-) + (\Delta - 1)f_1(\chi_n^+, \chi_n^-) + \Delta \max_{x_n \in [\chi_n, \chi_n']} g_1(x_n, \chi_n^+, \chi_n^-) + o(1). \quad (43)$$

where, recall, $\chi_n = \max\{0, \chi_n^+ + \chi_n^- - 1\}$, $\chi_n' = \min\{\chi_n^+, \chi_n^-\}$. Moreover, as $n \rightarrow \infty$, we have

$$\log(\lambda)(\chi_n^+ + \chi_n^-) + (\Delta - 1)f_1(\chi_n^+, \chi_n^-) \rightarrow \log(\lambda)(\chi^+ + \chi^-) + (\Delta - 1)f_1(\chi^+, \chi^-), \quad (44)$$

$$\max_{x_n \in [\chi_n, \chi_n']} g_1(x_n, \chi_n^+, \chi_n^-) \rightarrow \max_{x \in [\chi, \chi']} g_1(x, \chi^+, \chi^-), \quad (45)$$

where recall that $\chi = \max\{0, \chi^+ + \chi^- - 1\}$, $\chi' = \min\{\chi^+, \chi^-\}$. The limit (44) follows from $(\chi_n^+, \chi_n^-) \rightarrow (\chi^+, \chi^-)$ and the continuity of f_1 . To establish the limit (45) we use Lemma 21 for the function g_1 , with \mathcal{R}_{g_1} as defined in (41), $\mathcal{R} = \{(x, \chi^+, \chi^-) \mid \chi \leq x \leq \chi'\}$ and $\mathcal{C}_n = \{(x_n, \chi_n^+, \chi_n^-) \mid \chi_n \leq x_n \leq \chi_n', n x_n \in \mathbb{Z}\}$. The \mathcal{C}_n are $\frac{10}{n}$ -covers of \mathcal{R} (with room to spare for the constant 10) under the Euclidean metric on \mathbb{R}^3 . It is also immediate to verify that the set of limit points of $\bigcup_n (\mathcal{C}_n \setminus \mathcal{R})$ is a subset of \mathcal{R} as a consequence of $(\chi_n^+, \chi_n^-, \chi_n, \chi_n') \rightarrow (\chi^+, \chi^-, \chi, \chi')$.

From the definition (29),

$$\Psi_{1; \beta, \gamma, \lambda}(\chi^+, \chi^-) = \log(\lambda)(\chi^+ + \chi^-) + (\Delta - 1)f_1(\chi^+, \chi^-) + \Delta \max_{\chi \leq x \leq \chi'} g_1(x, \chi^+, \chi^-), \quad (46)$$

Combining this with (43), (44) and (45) yields (28), as wanted.

We next turn to the second moment $\mathbf{E}_{\mathcal{G}_n} [(Z_G^{\chi^+, \chi^-})^2]$. The reasoning is almost identical, though carrying out the arguments is more cumbersome due to the larger number of variables. Let

$$\mathcal{R}_{f_2} = \{(z, w, u, v) \mid z \geq u \geq 0, 1 - 2z + u \geq 0, w \geq v \geq 0, 1 - 2w + v \geq 0\}.$$

The function $f_2(z, w, u, v)$ is continuous over \mathcal{R}_{f_2} . We will also write $g_2(\mathbf{y})$ instead of $g_{2;\beta,\gamma}(\mathbf{y})$. Note that $g_2(\mathbf{y})$ is continuous over $\mathcal{R}_{g_2} = \{\mathbf{y} \mid \forall i, j \in [4], y_{ij} \geq 0\}$. By analogy to (36), define

$$\begin{aligned} l_1 &= u, & l_2 &= l_3 = z - u, & l_4 &= 1 - 2z + u, \\ r_1 &= v, & r_2 &= r_3 = w - v, & r_4 &= 1 - 2w + v. \end{aligned}$$

For $(z, w, u, v) \in \mathcal{R}_{f_2}$, define

$$\mathcal{T}(z, w, u, v) = \{\mathbf{y} = \{y_{ij}\}_{i,j \in [4]} \mid \forall i, j \in [4], y_{ij} \geq 0, \sum_j y_{ij} = l_i, \sum_i y_{ij} = r_j\}.$$

Let $\mathcal{R} = \mathcal{T}(\chi^+, \chi^-, v^+, v^-)$ and let $\mathcal{C}_n = \mathcal{T}(\chi_n^+, \chi_n^-, v_n^+, v_n^-)$.

By Stirling's approximation, the functions f_2 and g_2 capture the asymptotics of $\Sigma_{v^+, v^-}^{\chi^+, \chi^-}$ and $\xi_{v_n^+, v_n^-}^{\chi_n^+, \chi_n^-}$ (see (30) and (34)):

$$\frac{1}{n} \log |\Sigma_{v^+, v^-}^{\chi^+, \chi^-}| = -f_2(\chi_n^+, \chi_n^-, v_n^+, v_n^-) + o(1), \quad \frac{1}{n} \log \xi_{v_n^+, v_n^-}^{\chi_n^+, \chi_n^-} = g_2(\mathbf{y}_n) + f_2(\chi_n^+, \chi_n^-, v_n^+, v_n^-) + o(1). \quad (47)$$

We are now set to compute the asymptotics of $\mathbf{E}_{\mathcal{G}_n} [Y_{v^+, v^-}^{\chi^+, \chi^-}]$, for which we have derived an expression in (33). Approximating the sums in the expression by their maximum terms only introduces $o(1)$ error, so we obtain

$$\frac{1}{n} \log \mathbf{E}_{\mathcal{G}_n} [Y_{v^+, v^-}^{\chi^+, \chi^-}] = 2 \log(\lambda)(\chi_n^+ + \chi_n^-) + (\Delta - 1)f_2(\chi_n^+, \chi_n^-, v_n^+, v_n^-) + \Delta \max_{\mathbf{y}_n \in \mathcal{C}_n} g_2(\mathbf{y}_n) + o(1). \quad (48)$$

Taking limits as $n \rightarrow \infty$, we will now derive analogues of (44) and (45):

$$2 \log(\lambda)(\chi_n^+ + \chi_n^-) + (\Delta - 1)f_2(\chi_n^+, \chi_n^-, v_n^+, v_n^-) \rightarrow 2 \log(\lambda)(\chi^+ + \chi^-) + (\Delta - 1)f_2(\chi^+, \chi^-, v^+, v^-), \quad (49)$$

$$\max_{\mathbf{y}_n \in \mathcal{C}_n} g_2(\mathbf{y}_n) \rightarrow \max_{\mathbf{y} \in \mathcal{R}} g_2(\mathbf{y}). \quad (50)$$

Equation (49) follows from the continuity of f_2 . Equation (50) follows by Lemma 21. Combining (48), (49) and (50) gives

$$\frac{1}{n} \log \mathbf{E}_{\mathcal{G}_n} [Y_{v^+, v^-}^{\chi^+, \chi^-}] = 2 \log(\lambda)(\chi^+ + \chi^-) + (\Delta - 1)f_2(\chi^+, \chi^-, v^+, v^-) + \Delta \max_{\mathbf{y} \in \mathcal{T}(\chi^+, \chi^-, v^+, v^-)} g_2(\mathbf{y}) + o(1).$$

The right-hand side is $\Psi'_{2;\beta,\gamma,\lambda}(\chi^+, \chi^-, v^+, v^-) + o(1)$, so we have established (37).

To obtain (39), the sum (31) yields

$$\begin{aligned} \frac{1}{n} \log \mathbf{E}_{\mathcal{G}_n} [(Z_G^{\chi^+, \chi^-})^2] &= 2 \log(\lambda)(\chi_n^+ + \chi_n^-) \\ &+ \max_{(v_n^+, v_n^-, \mathbf{y}_n) \in \mathcal{O}_n} \{(\Delta - 1)f_2(\chi_n^+, \chi_n^-, v_n^+, v_n^-) + \Delta g_2(\mathbf{y}_n)\} + o(1), \end{aligned} \quad (51)$$

where

$$\begin{aligned} \mathcal{D}_n &:= \{(v_n^+, v_n^-) \mid \max\{0, 2\chi_n^+ - 1\} \leq v_n^+ \leq \chi_n^+, \max\{0, 2\chi_n^- - 1\} \leq v_n^- \leq \chi_n^-\}, \\ \mathcal{O}_n &:= \{(v_n^+, v_n^-, \mathbf{y}_n) \mid (v_n^+, v_n^-) \in \mathcal{D}_n, \mathbf{y}_n \in \mathcal{C}_n\}. \end{aligned}$$

Yet another application of Lemma 21 yields

$$\begin{aligned} & \max_{(v_n^+, v_n^-, \mathbf{y}_n) \in \mathcal{O}_n} \{(\Delta - 1)f_2(\chi_n^+, \chi_n^-, v_n^+, v_n^-) + \Delta g_2(\mathbf{y}_n)\} \\ & \rightarrow \max_{(v^+, v^-, \mathbf{y}) \in \mathcal{O}(\chi^+, \chi^-)} \{(\Delta - 1)f_2(\chi^+, \chi^-, v^+, v^-) + \Delta g_2(\mathbf{y})\}, \end{aligned} \quad (52)$$

where $\mathcal{O}(\chi^+, \chi^-) := \{(v^+, v^-, \mathbf{y}) \mid (v^+, v^-) \in \mathcal{D}, \mathbf{y} \in \mathcal{T}(1, \chi^+, \chi^-, v^+, v^-)\}$. Combining (51) and (52) yields (39), as desired.

References

- [1] Andrei A. Bulatov. The complexity of the counting constraint satisfaction problem. In *ICALP*, pages 646–661. Springer-Verlag, 2008.
- [2] Andrei A. Bulatov, Martin Dyer, Leslie Ann Goldberg, Mark Jerrum, and Colin McQuillan. The expressibility of functions on the Boolean domain, with applications to counting CSPs. *J. ACM*, 60(5):32:1–32:36, October 2013.
- [3] Jin-Yi Cai and Xi Chen. Complexity of counting CSP with complex weights. In *STOC*, pages 909–920. ACM, 2012.
- [4] Jin-Yi Cai, Xi Chen, Heng Guo, and Pinyan Lu. Inapproximability after uniqueness phase transition in two-spin systems. In *COCO*, pages 336–347, 2012.
- [5] Jin-Yi Cai, Pinyan Lu, and Mingji Xia. Holant problems and counting CSP. In *STOC*, pages 715–724. ACM, 2009.
- [6] Xi Chen, Martin E. Dyer, Leslie Ann Goldberg, Mark Jerrum, Pinyan Lu, Colin McQuillan, and David Richerby. The complexity of approximating conservative counting CSPs. Preprint, 2012. Available from the arXiv at: <http://arxiv.org/abs/1208.1783>
- [7] Xi Chen, Martin E. Dyer, Leslie Ann Goldberg, Mark Jerrum, Pinyan Lu, Colin McQuillan, and David Richerby. The complexity of approximating conservative counting CSPs. In *STACS*, pages 148–159, 2013.
- [8] Martin E. Dyer, Leslie Ann Goldberg, Catherine S. Greenhill, and Mark Jerrum. The relative complexity of approximate counting problems. *Algorithmica*, 38(3):471–500, 2003.
- [9] Martin E. Dyer, Leslie Ann Goldberg, and Mark Jerrum. An approximation trichotomy for Boolean #CSP. *J. Comput. Syst. Sci.*, 76(3-4):267–277, 2010.
- [10] Martin E. Dyer and David Richerby. An effective dichotomy for the counting constraint satisfaction problem. *SIAM J. Comput.*, 42(3):1245–1274, 2013.
- [11] Andreas Galanis, Qi Ge, Daniel Štefankovič, Eric Vigoda, and Linji Yang. Improved inapproximability results for counting independent sets in the hard-core model. In *APPROX-RANDOM*, pages 567–578, 2011.
- [12] Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. Inapproximability of the partition function for the antiferromagnetic Ising and hard-core models. Preprint, 2012. Available from the arXiv at: <http://arxiv.org/abs/1203.2226v3>

- [13] Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. Inapproximability for antiferromagnetic spin systems in the tree non-uniqueness region. Preprint, 2013. To appear in *STOC*, 2014. Available from the arXiv at: <http://arxiv.org/abs/1305.2902v2>
- [14] Qi Ge and Daniel Štefankovič. A graph polynomial for independent sets of bipartite graphs. In *FSTTCS*, pages 240–250, 2010.
- [15] Hans-Otto Georgii. *Gibbs Measures and Phase Transitions, 2nd edition*. Walter de Gruyter, 2011.
- [16] Leslie Ann Goldberg and Mark Jerrum. The complexity of ferromagnetic Ising with local fields. *Combinatorics, Probability & Computing*, 16(1):43–61, 2007.
- [17] Leslie Goldberg and Mark Jerrum. A counterexample to rapid mixing of the Ge-Štefankovič process. *Electron. Commun. Probab.*, 17:no. 5, 1–6, 2012.
- [18] Leslie Ann Goldberg and Mark Jerrum. Approximating the partition function of the ferromagnetic Potts model. *J. ACM*, 59(5):25, 2012.
- [19] Leslie Ann Goldberg, Mark Jerrum, and Colin McQuillan. Approximating the partition function of planar two-state spin systems. Preprint, 2012. Available from the arXiv at: <http://arxiv.org/abs/1208.4987>
- [20] Leslie Ann Goldberg, Mark Jerrum, and Mike Paterson. The computational complexity of two-state spin systems. *Random Struct. Algorithms*, 23(2):133–154, 2003.
- [21] Mark Jerrum and Alistair Sinclair. Polynomial-time approximation algorithms for the Ising model. *SIAM J. Comput.*, 22(5):1087–1116, 1993.
- [22] Frank Kelly. Stochastic models of computer communication systems. *Journal of the Royal Statistical Society. Series B (Methodological)*, 47(3):pp. 379–395, 1985.
- [23] Liang Li, Pinyan Lu, and Yitong Yin. Approximate counting via correlation decay in spin systems. In *SODA*, pages 922–940, 2012.
- [24] Liang Li, Pinyan Lu, and Yitong Yin. Correlation decay up to uniqueness in spin systems. In *SODA*, pages 67–84, 2013. Available from the arXiv at: <http://arxiv.org/abs/1111.7064v2>
- [25] Jingcheng Liu, Pinyan Lu, and Chihao Zhang. The complexity of ferromagnetic two-spin systems with external fields. Preprint, 2014. Available from the arXiv at: <http://arxiv.org/abs/1402.4346>
- [26] Elchanan Mossel, Dror Weitz, and Nicholas Wormald. On the hardness of sampling independent sets beyond the tree threshold. *Prob. Theory Related. Fields*, 143:401–439, 2009.
- [27] Alistair Sinclair, Piyush Srivastava, and Marc Thurley. Approximation algorithms for two-state anti-ferromagnetic spin systems on bounded degree graphs. In *SODA*, pages 941–953, 2012.
- [28] Allan Sly. Computational transition at the uniqueness threshold. In *FOCS*, pages 287–296, 2010.
- [29] Allan Sly and Nike Sun. The computational hardness of counting in two-spin models on d-regular graphs. In *FOCS*, pages 361–369, 2012.
- [30] Dror Weitz. Counting independent sets up to the tree threshold. In *STOC*, pages 140–149, 2006.
- [31] David Zuckerman. On unapproximable versions of NP-complete problems. *SIAM J. Comput.*, 25(6):1293–1304, 1996.