

The complexity of approximating averages on bounded-degree graphs

Andreas Galanis*

Daniel Štefankovič†

Eric Vigoda‡

July 19, 2020

Abstract

We prove that, unless $P=NP$, there is no polynomial-time algorithm to approximate within some multiplicative constant the average size of an independent set in graphs of maximum degree 6. This is a special case of a more general result for the hard-core model defined on independent sets weighted by a parameter $\lambda > 0$. In the general setting, we prove that, unless $P=NP$, for all $\Delta \geq 3$, all $\lambda > \lambda_c(\Delta)$, there is no FPTAS which applies to all graphs of maximum degree Δ for computing the average size of the independent set in the Gibbs distribution, where $\lambda_c(\Delta)$ is the critical point for the uniqueness/non-uniqueness phase transition on the Δ -regular tree. Moreover, we prove that for λ in a dense set of this non-uniqueness region the problem is NP-hard to approximate within some constant factor. Our work extends to the antiferromagnetic Ising model and generalizes to all 2-spin antiferromagnetic models, establishing hardness of computing the average magnetization in the tree non-uniqueness region.

Previously, Schulman, Sinclair and Srivastava (2015) showed that it is $\#P$ -hard to compute the average magnetization exactly, but no hardness of approximation results were known. Hardness results of Sly (2010) and Sly and Sun (2014) for approximating the partition function do not imply hardness of computing averages. The new ingredient in our reduction is an intricate construction of pairs of rooted trees whose marginal distributions at the root agree but their derivatives disagree. The main technical contribution is controlling what marginal distributions and derivatives are achievable and using Cauchy's functional equation to argue existence of the gadgets.

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

†Department of Computer Science, University of Rochester, USA. Research supported in part by NSF grant CCF-1563757.

‡Department of Computer Science, University of California, Santa Barbara, CA 93106, USA. Research supported in part by NSF grant CCF-1563838.

1 Introduction

This paper addresses the computational problem of computing averages of simple functions over combinatorial structures of a graph. Can we estimate elementary statistics of combinatorial structures in polynomial-time? This genre of problems is nicely illustrated for the example of independent sets.

Given a graph $G = (V, E)$ can we efficiently estimate the average size of an independent set in G ? For graph $G = (V, E)$, let $\mathcal{I}_G$ denote the set of independent sets of G , and let μ_G denote the uniform distribution over $\mathcal{I}_G$. Denote the average independent set size by $\mathcal{M}(G) = \mathbf{E}_{\sigma \sim \mu} [|\sigma|]$.

Schulman *et al.* [12] (see also [13]) proved that exactly computing the average independent set size is #P-hard for bounded-degree graphs. We investigate approximation algorithms for the problem. For a constant $C > 1$ we say there is a C -approximation algorithm for the average independent set size if the algorithm outputs an estimate EST for which $\frac{1}{C} \times \mathcal{M}(G) \leq \text{EST} \leq C \times \mathcal{M}(G)$. An FPTAS is an algorithm which, for any input $G = (V, E)$ and $\epsilon > 0$, achieves a $(1 + \epsilon)$ -approximation factor in time $\text{poly}(|V|, 1/\epsilon)$.

Weitz [17] presented an FPTAS for estimating $|\mathcal{I}_G|$, the number of independent sets, in graphs of maximum degree ≤ 5 , and this immediately yields an FPTAS for the average independent set size in graphs of maximum degree ≤ 5 , see also [2, 9, 10] for new algorithmic approaches. We prove that this result is optimal. In fact, we prove that approximating the average independent set size in graphs of maximum degree 6 is hard within a constant factor.

Theorem 1. *There is a constant $C > 1$ such that, for all integers $\Delta \geq 6$, for graphs G of maximum degree Δ there is no polynomial-time C -approximation algorithm for computing the average independent-set size of G , unless $\text{P} = \text{NP}$.*

This theorem is a special case of a more general result for the hard-core model, which is a statistical physics model of particular combinatorial interest. The hard-core model is defined on independent sets weighted by a parameter $\lambda > 0$, known as the fugacity. An independent set $\sigma \in \mathcal{I}_G$ has weight $w(\sigma) = \lambda^{|\sigma|}$. For a graph G and fugacity $\lambda > 0$, the Gibbs distribution is defined as $\mu_{G;\lambda}(\sigma) = w(\sigma)/Z_{G;\lambda}$ where the normalizing factor $Z_{G;\lambda} = \sum_{\tau \in \mathcal{I}_G} w(\tau)$ is known as the partition function.

The earlier case of unweighted independent sets corresponds to the hard-core model with $\lambda = 1$. Hence we use the same notation $\mathcal{M}(G) = \mathbf{E}_{\sigma \sim \mu} [|\sigma|]$ to denote the average independent set size in the Gibbs distribution.

On the Δ -regular tree, the hard-core model undergoes a phase transition at the critical point $\lambda_c(\Delta) = \frac{(\Delta-1)^{\Delta-1}}{(\Delta-2)^\Delta}$. When $\lambda \leq \lambda_c(\Delta)$ there is a unique infinite-volume Gibbs measure on the Δ -regular tree (roughly, this corresponds to the decay of the “influence” of the leaves on the root), whereas when $\lambda > \lambda_c(\Delta)$ there is non-uniqueness, i.e., there are multiple infinite-volume Gibbs measures.

There is an interesting computational phase transition for graphs of maximum degree Δ that occurs at this same tree threshold. For all constant Δ , all $\lambda < \lambda_c(\Delta)$, all graphs of maximum degree Δ , there is an FPTAS for the partition function [17]. On the other side, for all $\Delta \geq 3$, all $\lambda > \lambda_c(\Delta)$, there is no FPRAS for approximating the partition function on graphs of maximum degree Δ , unless $\text{NP} = \text{RP}$ [15, 16, 5]. However, hardness of computing partition functions does not imply hardness of computing averages in the Gibbs distribution; see the case of the antiferromagnetic Ising model with no external field discussed in Section 1.1.

We prove that computing the average independent set size also undergoes a computational phase transition at the tree critical point $\lambda_c(\Delta)$. As before, Weitz’s algorithmic result [17] yields, for all

constant Δ , all $\lambda < \lambda_c(\Delta)$, all graphs G of maximum degree Δ , an FPTAS for the average independent set size $\mathcal{M}(G)$ (more generally, an approximate sampling algorithm implies an algorithm for approximating averages). We prove that this result is optimal: when $\lambda > \lambda_c(\Delta)$ there is no FPTAS for the average independent set size.

Theorem 2. *Let $\Delta \geq 3$ be an integer and $\lambda > \lambda_c(\Delta)$. Then, for graphs G of maximum degree Δ , there is no FPTAS for computing the average independent-set size in the hard-core distribution $\mu_{G;\lambda}$, unless $P=NP$.*

In fact, our inapproximability result for general $\lambda > \lambda_c(\Delta)$ is stronger, it actually precludes approximation algorithms with factors of $1 \pm \frac{\delta}{\log n}$, where n is the number of the vertices of the input graph and δ is an appropriate constant, see Theorem 7 below for the precise statement. For a dense set of λ , we actually obtain constant-factor inapproximability (analogously to Theorem 1).

Theorem 3. *Let $\Delta \geq 3$ be an integer. Then, for every real $\lambda > \lambda_c(\Delta)$ and $\epsilon > 0$, there is an algebraic number $\hat{\lambda}$ with $|\hat{\lambda} - \lambda| \leq \epsilon$ and a constant $C = C(\hat{\lambda}, \Delta) > 1$ such that, for graphs G of maximum degree Δ , there is no poly-time C -approximation algorithm for computing the average independent-set size of G in the hard-core distribution $\mu_{G;\hat{\lambda}}$, unless $P=NP$.*

1.1 Results for the antiferromagnetic Ising model

Our results extend to the antiferromagnetic Ising model. Let $G = (V, E)$ be a graph. For $\beta, \lambda > 0$, let $\mu_{G;\beta,\lambda}$ denote the Ising distribution on G with edge activity β and external field λ , i.e., for $\sigma : V \rightarrow \{0, 1\}$ we have

$$\mu_{G;\beta,\lambda}(\sigma) = \frac{\lambda^{|\sigma|} \beta^{m(\sigma)}}{Z_{G;\beta,\lambda}},$$

where $m(\sigma)$ denotes the number of monochromatic edges in G under σ , i.e., edges whose endpoints have the same spin. The model is called antiferromagnetic if $\beta \in (0, 1)$ and ferromagnetic, otherwise. We define the average magnetization of G to be the average number of vertices with spin 1, i.e.,

$$\mathcal{M}_{\beta,\lambda}(G) = \frac{1}{Z_{G;\beta,\lambda}} \sum_{\sigma: V \rightarrow \{0,1\}} |\sigma| \lambda^{|\sigma|} \beta^{m(\sigma)} = \mathbf{E}_{\sigma \sim \mu} [|\sigma|],$$

where for a configuration $\sigma : V \rightarrow \{0, 1\}$, we use $|\sigma|$ to denote $\sum_{v \in V(G)} \sigma(v)$, i.e., the total number of vertices with spin 1.

In the ferromagnetic case, there is an FPRAS for approximating the magnetization for all $\beta > 1$ and $\lambda > 0$, due to the algorithm of Jerrum and Sinclair [6]. For $\Delta \geq 3$, let $\beta_c(\Delta) = \frac{\Delta-2}{\Delta}$. It is known that the antiferromagnetic Ising model with edge activity β and external field λ has non-uniqueness on the infinite Δ -regular tree iff $\beta \in (0, \beta_c)$ and $\lambda \in (1/\lambda_c^{\text{Ising}}, \lambda_c^{\text{Ising}})$ for some explicit $\lambda_c^{\text{Ising}} = \lambda^{\text{Ising}}(\beta, \Delta) > 1$. For all constant Δ , in the tree uniqueness region there is an FPTAS for the partition function for graphs of maximum degree Δ [14, 7], and, once again, this implies an FPTAS for the magnetization. In the tree non-uniqueness region, for all $\Delta \geq 3$, unless $NP=RP$ there is no FPRAS for the partition function for graphs of maximum degree Δ [16]. We prove that approximating the magnetization is also intractable in the tree non-uniqueness region, apart from the case $\lambda = 1$ (where the magnetization can be computed trivially for all graphs G since it equals $\frac{1}{2}|V(G)|$).

Theorem 4. *Let $\Delta \geq 3$ be an integer, $\beta \in (0, \beta_c(\Delta))$ and $\lambda \in (\frac{1}{\lambda_c}, \lambda_c)$ with $\lambda \neq 1$, where $\lambda_c = \lambda_c^{\text{Ising}}(\Delta, \beta)$. Then, for graphs G of maximum degree Δ , there is no FPTAS for computing the average magnetization in the Ising distribution $\mu_{G;\beta,\lambda}$, unless $P=NP$.*

As in Theorem 2, the inapproximability factor in Theorem 4 is in fact stronger, see Theorem 7 below for the precise statement. For a dense set of λ , we again obtain constant-factor inapproximability.

Theorem 5. *Let $\Delta \geq 3$ be an integer, rational $\beta \in (0, \beta_c(\Delta))$ and $\lambda_c = \lambda_c^{\text{Ising}}(\Delta, \beta)$. Then, for every $\lambda \in (\frac{1}{\lambda_c}, \lambda_c)$ and $\epsilon > 0$, there is an algebraic number $\hat{\lambda}$ with $|\hat{\lambda} - \lambda| \leq \epsilon$ and a constant $C = C(\beta, \hat{\lambda}, \Delta) > 1$ such that, for graphs G of maximum degree Δ , there is no poly-time C -approximation algorithm for computing the average magnetization $\mathcal{M}_{\beta, \hat{\lambda}}(G)$, unless $P=NP$.*

1.2 Results for general antiferromagnetic 2-spin systems

While the hard-core model and the Ising model are the most canonical 2-spin models, the results of the previous two sections will be obtained as special cases of the following results for general antiferromagnetic 2-spin systems. This more general perspective will also allow us to give a unified proof of Theorems 2, 3, 4, 5.

Let $G = (V, E)$ be a graph. For $\beta, \gamma, \lambda > 0$, let $\mu_{G; \beta, \gamma, \lambda}$ denote the Gibbs distribution on G with edge activities β, γ and external field λ , i.e., for $\sigma : V \rightarrow \{0, 1\}$ we have

$$\mu_{G; \beta, \lambda}(\sigma) = \frac{\lambda^{|\sigma|} \beta^{m_0(\sigma)} \gamma^{m_1(\sigma)}}{Z_{G; \beta, \gamma, \lambda}},$$

where $m_0(\sigma), m_1(\sigma)$ denotes the number of edges in G whose endpoints are assigned under σ the pair of spins $(0, 0)$ and $(1, 1)$, respectively.

The parameter pair (β, γ) is called *antiferromagnetic* if $\beta\gamma \in [0, 1)$ and at least one of β, γ is non-zero, and it is called *ferromagnetic*, otherwise. Note that the hard-core model is the case $\beta = 1, \gamma = 0$ (under the convention that $0^0 \equiv 1$) whereas the antiferromagnetic Ising model is the case $0 < \beta = \gamma < 1$.

We next define the range of parameters (β, γ, λ) where our inapproximability results for the magnetization apply; these are precisely the parameters where the spin system exhibits non-uniqueness on the infinite Δ -regular tree, apart from the case of the antiferromagnetic Ising model with $\lambda = 1$, where as noted in Section 1.1 the magnetization can be computed trivially for all graphs G .

Definition 6. *Let $\Delta \geq 3$ be an integer. We let $\mathcal{U}_\Delta$ be the set of (β, γ, λ) such that (β, γ) is antiferromagnetic, and the (unique) fixpoint $x^* > 0$ of the function $f(x) = \frac{1}{\lambda} \left(\frac{\beta x + 1}{x + \gamma} \right)^{\Delta-1}$ satisfies $|f'(x^*)| > 1$.*

We let $\mathcal{U}_\Delta^ = \mathcal{U}_\Delta \setminus \bigcup_{\beta \in (0, 1)} \{(\beta, \beta, 1)\}$ be the set of parameters in $\mathcal{U}_\Delta$ other than those where computing the magnetization is trivial.*

We note that Li, Lu, Yin [7] define a notion of “up-to- Δ ” uniqueness which requires (β, γ, λ) to be in uniqueness for every $d \leq \Delta$; they obtain an FPTAS for the partition function in that region. The complement of their region corresponds to non-uniqueness in the sense of Definition 6 for some $d \leq \Delta$. Our inapproximability results extend to this bigger region by applying our theorems for smaller values of Δ .

Our first inapproximability result for general antiferromagnetic 2-spin models, which is a generalization/strengthening of Theorems 2 and 4, is the following. The proof is given in Section 5.5.

Theorem 7. *Let $\Delta \geq 3$ be an integer and $(\beta, \gamma, \lambda) \in \mathcal{U}_\Delta^*$. Then, for graphs $G \in \mathbf{G}_\Delta$, there is no FPTAS for computing the average magnetization in the Gibbs distribution $\mu_{G; \beta, \gamma, \lambda}$, unless $P=NP$. In fact, there is a constant $\kappa = \kappa(\Delta, \beta, \gamma, \lambda) > 0$ such that there is no poly-time $(1 + \frac{\kappa}{\log n})$ -approximation algorithm for computing the average magnetization $\mathcal{M}_{\beta, \gamma, \lambda}(G)$, where $n = |V(G)|$.*

We remark that a “no-FPTAS” result can be strengthened to an inapproximability factor of $(1 \pm \frac{1}{n^\epsilon})$ for any constant $\epsilon > 0$ via standard powering techniques; the tighter hardness factor of $(1 \pm \frac{\kappa}{\log n})$ in Theorem 7 requires a substantially more delicate argument.

Theorem 8. *Let $\Delta \geq 3$ be an integer and $(\beta, \gamma, \lambda) \in \mathcal{U}_\Delta$ with β, γ rational numbers. Then, for every $\epsilon > 0$, there is an algebraic number $\hat{\lambda}$ with $|\hat{\lambda} - \lambda| \leq \epsilon$ and a constant $C = C(\beta, \gamma, \hat{\lambda}, \Delta) > 1$ such that, for graphs G of maximum degree Δ , there is no poly-time C -approximation algorithm for computing the average magnetization $\mathcal{M}_{\beta, \gamma, \hat{\lambda}}(G)$, unless $P=NP$.*

2 Proof Outline

In this section, we give some of the key elements of the techniques needed to prove our inapproximability results and conclude with the proof of Theorem 1. We start by describing “field gadgets” and state the main lemmas that we will use in our reduction.

2.1 Gadgets with approximately equal effective fields and different averages

For a graph G and $\sigma : V \rightarrow \{0, 1\}$, we say that a vertex $v \in V$ is occupied if $\sigma(v) = 1$, and unoccupied otherwise.¹ For a subset $S \subseteq V$, we use σ_S to denote the configuration on S which is restriction of σ on S .

Definition 9. *A field gadget is a rooted tree $\mathcal{T}$ whose root ρ has degree one. For antiferromagnetic (β, γ) and $\lambda > 0$, let $\mu = \mu_{\mathcal{T}; \beta, \gamma, \lambda}$.*

1. *The effective field of the gadget, denoted by $R_{\mathcal{T}} = R_{\mathcal{T}}(\beta, \gamma, \lambda)$, is $1/\lambda$ times the ratio of the weight of configurations where the root is occupied to the weight of configurations where the root is unoccupied, i.e., $R_{\mathcal{T}} = \frac{1}{\lambda} \frac{\mu(\sigma_\rho=1)}{\mu(\sigma_\rho=0)}$. (The division by λ is to avoid double-counting the contribution of the root later on.)*
2. *The magnetization gap of the gadget, denoted by $M_{\mathcal{T}} = M_{\mathcal{T}}(\beta, \gamma, \lambda)$, is the expected number of occupied vertices conditioned on the root being occupied minus the expected number of occupied vertices conditioned on the root being unoccupied (the root is included in the count), i.e.,*

$$M_{\mathcal{T}} = \mathbf{E}_{\sigma \sim \mu}[|\sigma| \mid \sigma(\rho) = 1] - \mathbf{E}_{\sigma \sim \mu}[|\sigma| \mid \sigma(\rho) = 0].$$

In the special case that $\lambda = \frac{1-\beta}{1-\gamma}$ with $\beta \neq \gamma$, a field gadget consists of a rooted bipartite graph obtained from a rooted tree where some of the leaves have been replaced by a distinct cycle of length four (by identifying the leaf with a vertex of the cycle).

Note that, in the case that $\lambda = \frac{1-\beta}{1-\gamma}$, the effective field of any tree gadget can be shown to be equal to 1 and the magnetization gap to 0, so that is why we need to consider the “mildly non-tree” construction in Definition 9.


It will be useful to illustrate Definition 9 with a few examples.

Example 10. *For example the rooted tree with one edge has effective field $R = \frac{1+\gamma\lambda}{\beta+\lambda}$. The degenerate example where the root has degree zero has effective field $R = 1$.*

The following more interesting example in the case of independent sets will be used to prove Theorem 1; it gives a pair of field gadgets with the same effective field but different magnetization gaps.

¹This terminology is standard for the hard-core model, but it will be convenient to use it for general 2-spin systems.

Example 11. Consider the independent set model (with $\lambda = 1$). Consider the trees $\mathcal{T}_1, \mathcal{T}_2$ with roots ρ_1, ρ_2 below.



Then $R_{\mathcal{T}_1} = R_{\mathcal{T}_2}$ since $R_{\mathcal{T}_1} = \frac{2}{3}$ and $R_{\mathcal{T}_2} = \frac{24}{36}$, but $M_{\mathcal{T}_1} \neq M_{\mathcal{T}_2}$ since $M_{\mathcal{T}_1} = \frac{3}{2} - \frac{2}{3} = \frac{5}{6}$ and $M_{\mathcal{T}_2} = \frac{35}{12} - \frac{13}{6} = \frac{3}{4}$. These values can be either verified by enumeration and linearity of expectation, or else using the recursions of the upcoming Lemma 17.

We will be interested in finding pairs of field gadgets analogous to Example 11. Our interest in such pairs of field gadgets is justified by the following theorem.

Theorem 12. Let $\Delta \geq 3$ be an integer and $(\beta, \gamma, \lambda) \in \mathcal{U}_{\Delta}^*$. Suppose that there exists a pair of field gadgets $\mathcal{T}_1, \mathcal{T}_2$ with maximum degree Δ such that $R_{\mathcal{T}_1} = R_{\mathcal{T}_2}$ but $M_{\mathcal{T}_1} \neq M_{\mathcal{T}_2}$.

Then, there is constant $c = c(\Delta, \beta, \gamma, \lambda) > 1$ such that, for graphs G of maximum degree Δ , there is no poly-time c -approximation algorithm for computing the average magnetization $\mathcal{M}_{\beta, \gamma, \lambda}(G)$, unless $P=NP$.

When can we find pairs of trees as in Example 11 with the same effective fields but different magnetization gaps? Even in the case of the hard-core model, a pair of trees either have the same effective field for only finitely many values of λ or else have the same magnetization gap for all λ . Thus we cannot hope for a universal pair of gadgets. Actually, the set of λ where the effective fields can be equal, over all pairs of trees, must be algebraic, hence measure zero.

Our main theorem for the construction of field gadgets is the following. The theorem roughly asserts that we can construct field gadgets with arbitrarily close effective fields but substantially different magnetization gaps. For a rational $r = p/q$ with integers p, q satisfying $\gcd(p, q) = 1$, we let $\text{bits}(r)$ denote the total number of bits needed to represent p, q .

Theorem 13. Let (β, γ, λ) be antiferromagnetic with $(\beta, \gamma, \lambda) \neq (\beta, \beta, 1)$. There exist $\hat{R}, \hat{M}, \Xi > 0$ and an algorithm which on input a rational $r \in (0, 1/2)$ outputs in time $\text{poly}(\text{bits}(r))$ a pair of field gadgets $\mathcal{T}_1, \mathcal{T}_2$, each of maximum degree 3 and size $O(|\log r|)$, such that

$$|R_{\mathcal{T}_1} - \hat{R}|, |R_{\mathcal{T}_2} - \hat{R}| \leq r, \text{ but } |M_{\mathcal{T}_1} - M_{\mathcal{T}_2}| > \hat{M}.$$

Moreover, the magnetization gaps $M_{\mathcal{T}_1}, M_{\mathcal{T}_2}$ are bounded in absolute value by the constant Ξ .

The proof of Theorem 13 is given in Section 3. In fact, we can bootstrap Theorem 13 to obtain pairs of trees with the same effective fields but different magnetization gaps for a dense set of (algebraic numbers) λ , and this gives a constant factor inapproximability result for the magnetization by applying Theorem 12.

Theorem 14. Let (β, γ) be antiferromagnetic. There exists a set S of algebraic numbers λ , dense in the interval $(0, \infty)$, such that for each $\lambda \in S$ the following holds. There is a pair of field gadgets $\mathcal{T}_1, \mathcal{T}_2$, each of maximum degree 3, such that $R_{\mathcal{T}_1} = R_{\mathcal{T}_2}$ but $M_1 \neq M_2$.

In the following, we sketch the proof of Theorem 12 and in Section 3 we will give a detailed outline of the proof of Theorems 13 and 14.

2.2 The reduction: using field gadgets to obtain inapproximability

Fix $\Delta \geq 3$ and $\beta, \gamma, \lambda \in \mathcal{U}_\Delta$. Our results will be based on showing that approximating MAX-CUT on 3-regular graphs H reduces to approximating the magnetization will be via a reduction from MAX-CUT. The reduction uses two types of gadgets.

The first type of gadget is a bipartite graph G which is an almost Δ -regular graph with n vertices on each side and ℓ “ports” on each side of degree $\Delta - 1$, which will be used to connect distinct copies of gadgets. These gadgets were used in previous inapproximability results for the partition function and analysing them was one of the key difficulties in those results. The main idea in these results is to replace each vertex of H by a distinct copy of G and make appropriate connections between ports to encode the edges of H ; then due to the antiferromagnetic interaction the resulting graph settles in configurations which correspond to Max-Cut configurations of H (the Max-Cut assignment can roughly be read off by looking at the “phases” of the gadgets G , i.e., which side has the most occupied vertices, see Section 2.3 for details).

However, for our inapproximability results we would need to analyze the magnetization of such gadgets by taking into account the effect of conditioning the spins of the ports; this task seems even more challenging than in the previous settings since it seems to require very refined estimates.

The field gadgets, which is the second type of gadgets, give a new reduction technique to bypass the need to perform this delicate, and likely difficult, task. In particular, by using a pair of field gadgets with the same effective fields and different magnetization gaps we will be able to observe the value of MAX-CUT(H) by taking the difference of the magnetizations when we append our field gadgets appropriately; crucially, the fact that the effective fields are the same implies that the underlying distribution does not change but that the difference of the magnetization gap of the two gadgets manifests itself in the magnetization.

2.3 The bipartite “phase-gadget” of Sly and Sun

Fix an integer $\Delta \geq 3$. Let $\mathcal{G}_n^\ell$ be the set of $(2n)$ -vertex bipartite graphs whose two sides are labelled with $+$, $-$ and are obtained from a Δ -regular $(2n)$ -vertex bipartite graph by deleting a matching of size ℓ . For a graph $G \in \mathcal{G}_n^\ell$, we denote its bipartition by (U^+, U^-) where U^+, U^- are vertex sets with $|U^+| = |U^-| = n$, and we denote by W^+, W^- the sets of vertices with degree $\Delta - 1$ on each side of the bipartition, i.e., the vertices incident to the edges of the matching, so that $|W^+| = |W^-| = \ell$. We will refer to set $W = W^+ \cup W^-$ as the *ports* of G .

For $\sigma : U^+ \cup U^- \rightarrow \{0, 1\}$, we define the *phase* $\mathcal{Y}(\sigma)$ of the configuration σ as the side which has the most occupied vertices under σ , i.e.,

$$\mathcal{Y}(\sigma) = \begin{cases} + & \text{if } |\sigma_{U^+}| \geq |\sigma_{U^-}|, \\ - & \text{otherwise.} \end{cases}$$

Sly and Sun [15] (see also [4]) establish that, whenever $\beta, \gamma, \lambda \in \mathcal{U}_\Delta$, for arbitrary $\epsilon > 0$ and integer $\ell \geq 1$, there exist n and $G \in \mathcal{G}_n^\ell$ such that, in the Gibbs distribution of the gadget graph G , each phase appears with probability close to $1/2 \pm \epsilon$, and that the spins of the ports of G , conditional on the phase, are roughly independent, with occupation probabilities that are asymmetric between the two sides.

The detailed properties of the gadget can be found in Lemma 33 of Section 5.1, though the exact details will not be important at this stage.

2.4 The reduction

Let H be a 3-regular instance of the MAX-CUT problem. Let k be an arbitrary positive integer, and let n and $G \in \mathcal{G}_n^{3k}$ be a bipartite gadget satisfying Lemma 33. Let also $\mathcal{T}$ be a field gadget with effective field R and magnetization gap M .

Let $\hat{H}_G^k$ be the graph with m disconnected components obtained by replacing each vertex of H with a distinct copy of G ; for $v \in V(H)$ we denote by G_v the copy of G in $\hat{H}_G^k$ corresponding to the vertex v , and by $U_v^\pm, W_v^\pm$ the vertex sets and ports of G_v in each side of the bipartition. Finally, we let $U_v = U_v^+ \cup U_v^-$ and $W_v = W_v^+ \cup W_v^-$.

Let $H_{G,\mathcal{T}}^k$ be the graph obtained from $\hat{H}_G^k$ as follows. For each edge $e = (u, v) \in E(H)$, pick k distinct ports from $W_u^\pm, W_v^\pm$ and connect them using a path with three edges, for a total of $2k$ edge-disjoint paths between the gadgets G_u and G_v : k paths between $+/+$ sides and k paths between $-/-$ sides. Then, for each path P append two distinct copies of the field gadget $\mathcal{T}$ by identifying the roots of the copies of $\mathcal{T}$, say ρ_1, ρ_2 , with the internal vertices of P , say t_1, t_2 . It will be helpful to consider also the graph H_G^k which is the graph $H_{G,\mathcal{T}}^k$ when $\mathcal{T}$ is the single-vertex graph, i.e., H_G^k is the graph after attaching the edge-disjoint paths but without the field gadgets.

For a configuration $\sigma : V(H_{G,\mathcal{T}}^k) \rightarrow \{0, 1\}$, we let $\hat{\mathcal{Y}}(\sigma) : V(H) \rightarrow \{+, -\}$ be the phases over the gadgets G_v with $v \in V(H)$, i.e., $\hat{\mathcal{Y}}(\sigma) = \{\mathcal{Y}(\sigma_{U_v})\}_{v \in V(H)}$. We define $\hat{\mathcal{Y}}(\cdot)$ similarly for configurations on $\hat{H}_G^k, H_G^k$. For a phase assignment $Y : V(H) \rightarrow \{+, -\}$, we let $\text{Cut}_H(Y) = \{(u, v) \in E(H) \mid Y_u \neq Y_v\}$ be the set of edges that are cut in H by viewing the phase assignment Y as a bipartition of $V(H)$ in the natural way. For $\mu = \mu_{H_{G,\mathcal{T}}^k; \beta, \gamma, \lambda}$, let

$$\text{AVG-CUT}_\mu(H) := \sum_{Y: V(H) \rightarrow \{+, -\}} \mu(\hat{\mathcal{Y}}(\sigma) = Y) |\text{Cut}_H(Y)|$$

be the size of the average cut in H when phase assignments are weighted by μ .

The following lemma associates the magnetization on the graph $H_{G,\mathcal{T}}^k$ with the quantity $\text{AVG-CUT}_\mu(H)$ and, in turn, with $\text{MAX-CUT}(H)$. The main idea is that the paths of length 3 between the ports cause antiferromagnetic interaction and the spin system on $H_{G,\mathcal{T}}^k$ prefers phase configurations Y that have large $|\text{Cut}_H(Y)|$ value. By choosing k large enough we can further ensure that $\text{AVG-CUT}_\mu(H)$ arbitrarily close to $\text{MAX-CUT}(H)$.

Lemma 15. *Let $\Delta \geq 3$ and $(\beta, \gamma, \lambda) \in \mathcal{U}_\Delta$. There are rational functions $A(R), B(R), C(R)$ defined for all $R \geq 0$, satisfying*

$$(i) A(0) = 1, (ii) A(R) > 1 \text{ for } R > 0, \text{ and } (iii) B(R) - C(R) = \lambda R \cdot (\log A(R))' \text{ for } R > 0 \quad (1)$$

so that the following holds for any field gadget $\mathcal{T}$ with effective field $R > 0$ and magnetization gap M . Let A, B, C be the values of the functions at λR and define $\mathcal{A}' := \mathbf{E}_{\tau \sim \mu_{\mathcal{T}; \beta, \gamma, \lambda}}[|\tau| \mid \tau(\rho) = 0]$, where ρ is the root of $\mathcal{T}$.

Let $\epsilon \in (0, 1/10)$, $k \geq 10/\log A$ be an integer, and n and $G \in \mathcal{G}_n^{3k}$ satisfy Lemma 33 with $\ell = 3k$ and the given ϵ . Then, for every 3-regular graph H with sufficiently large $|V(H)|$, for $\mu = \mu_{H_{G,\mathcal{T}}^k; \beta, \gamma, \lambda}$, we have that the average magnetization of the graph $H_{G,\mathcal{T}}^k$ satisfies

$$\mathcal{M}_{\beta, \gamma, \lambda}(H_{G,\mathcal{T}}^k) = 4k\mathcal{A}'|E(H)| + \mathbf{E}_{\sigma \sim \mu}[\|\sigma_{V(\hat{H}_G^k)}\|] + (1 \pm 8\epsilon)kMQ$$

with $Q := (B - C)\text{AVG-CUT}_\mu(H) + C|E(H)|$, while $\text{AVG-CUT}_\mu(H)$ satisfies

$$1/K \leq \frac{\text{AVG-CUT}_\mu(H)}{\text{MAX-CUT}(H)} \leq 1 \quad \text{for } K := 1 + \frac{6}{k \log A}.$$

We defer the detailed proof of Lemma 15 to Section 5.3.

Note that the expression for the magnetization of $H_{G,\mathcal{T}}^k$ accounts for the contributions of the phase gadgets G only implicitly. This is by design; when we append our pair of field gadgets $\mathcal{T}_1, \mathcal{T}_2$ which have the same effective fields, vertices in $V(\hat{H}_G^k)$ will have the same marginal distribution in both $\mu_{H_{G,\mathcal{T}_1}^k}$ and $\mu_{H_{G,\mathcal{T}_2}^k}$; on the other hand the magnetization gaps of $\mathcal{T}_1, \mathcal{T}_2$ are different. Therefore, when we take the difference of the magnetizations in $H_{G,\mathcal{T}_1}^k$ and $H_{G,\mathcal{T}_2}^k$, the contributions of the phase gadgets G cancel and, in the left-over quantity, the major contribution comes from the value of $\text{AVG-CUT}_\mu(H)$.

We are almost ready to prove Theorem 12. Recall, the assumption therein is that we have a pair of field gadgets with the same effective fields and different magnetization gaps. We will need to bootstrap this slightly to obtain additional pairs of field gadgets, as stated in the following lemma.

Lemma 16. *Let $\Delta \geq 3$ be an integer and (β, γ, λ) be antiferromagnetic with $(\beta, \gamma, \lambda) \neq (\beta, \beta, 1)$. Suppose that there exists a pair of field gadgets $\mathcal{T}_1, \mathcal{T}_2$ with maximum degree Δ such that $R_{\mathcal{T}_1} = R_{\mathcal{T}_2}$ but $M_{\mathcal{T}_1} \neq M_{\mathcal{T}_2}$. Then, we can construct for $j = 0, 1, 2, \dots$ an infinite sequence of pairs of field gadgets $\mathcal{T}_{1,j}, \mathcal{T}_{2,j}$ of maximum degree Δ with effective fields $R_{1,j}, R_{2,j}$ and magnetization gaps $M_{1,j}, M_{2,j}$ such that $R_{1,j} = R_{2,j}$ but $M_{1,j} \neq M_{2,j}$; moreover the values of $R_{1,j}$, and hence of $R_{2,j}$ as well, are pairwise distinct.*

The proof of Lemma 16 is given in Section 4.5. With Lemmas 16 and 33 at hand, we can now give the proof of Theorem 12.

Proof of Theorem 12. For $i \in \{1, 2\}$, let $R_i := R_{\mathcal{T}_i}$ be the effective field of $\mathcal{T}_i$, $M_i := M_{\mathcal{T}_i}$ be the magnetization gap of $\mathcal{T}_i$ and $\mathcal{A}'_i := \mathbf{E}_{\tau \sim \mu_{\mathcal{T}_i; \beta, \gamma, \lambda}}[|\tau| \mid \tau(\rho_i) = 0]$. Note that $R_1 = R_2$ but $M_1 \neq M_2$, and $\mathcal{A}'_1, \mathcal{A}'_2$ are absolute constants.

Let $A(R), B(R), C(R)$ be the functions in Lemma 15 and set $D(R) = B(R) - C(R)$. For convenience, we let A, B, C, D denote the common values of $A(R_i), B(R_i), C(R_i), D(R_i)$, respectively, for $i \in \{1, 2\}$. We will need to ensure in our argument that $D \neq 0$. The first observation is that (1) implies that for all but finitely many values of R we have that $D(R) \neq 0$.² The second observation is that, from Lemma 16, using $\mathcal{T}_1, \mathcal{T}_2$ we can construct for $j = 0, 1, 2, \dots$ an infinite sequence of pairs of field gadgets $\mathcal{T}_{1,j}, \mathcal{T}_{2,j}$ of maximum degree Δ with effective fields $R_{1,j}, R_{2,j}$ and magnetization gaps $M_{1,j}, M_{2,j}$ such that $R_{1,j} = R_{2,j}$ but $M_{1,j} \neq M_{2,j}$; moreover the values of $R_{1,j}$ are pairwise distinct. From these two observations, it follows that we can assume without loss of generality that $D \neq 0$.

Suppose for the sake of contradiction that, for arbitrarily small $\kappa > 0$, there is a polynomial-time algorithm that, on input a graph G of maximum degree Δ , produces a $(1 + \kappa)$ -approximation of the magnetization $\mathcal{M}_{G; \beta, \gamma, \lambda}$. We will show that we can approximate MAX-CUT on 3-regular graphs within a constant factor arbitrarily close to 1, contradicting the inapproximability result of [1].

For $i \in \{1, 2\}$, consider the graph $H_{G, \mathcal{T}_i}^k$ for some large integer $k > 1 + 10/\log A$ and small $\epsilon > 0$ to be specified later. Let μ_i denote the Gibbs distribution on $H_{G, \mathcal{T}_i}^k$ with parameters β, γ, λ . By Lemma 15, the average magnetization of the graph $H_{G, \mathcal{T}_i}^k$ satisfies

$$\mathcal{M}_{\beta, \gamma, \lambda}(H_{G, \mathcal{T}_i}^k) = 4k\mathcal{A}'_i|E(H)| + \mathbf{E}_{\sigma \sim \mu_i} \left[|\sigma_{V(\hat{H}_G^k)}| \right] + (1 \pm 8\epsilon)kM_iQ_i, \quad (2)$$

²If $B(R)$ and $C(R)$ agree on infinitely many values of R , then since both of them are rational functions of R , it must be the case that $B(R) = C(R)$ for all $R \geq 0$. Since for all $R \geq 0$, we have that $\lambda R \cdot (\log A(R))' = B(R) - C(R)$, this would give that $A(R)$ is a constant function; contradiction, since (1) asserts that $A(0) = 1$ and $A(R) > 1$ for all $R > 0$.

where $Q_i = D \text{AVG-CUT}_{\mu_i}(H) + C|E(H)|$ and $\text{AVG-CUT}_{\mu_i}(H)$ satisfies

$$1/K \leq \frac{\text{AVG-CUT}_{\mu_i}(H)}{\text{MAX-CUT}(H)} \leq 1 \text{ for } K := 1 + \frac{6}{k \log A}. \quad (3)$$

Note that since $R_1 = R_2$ we have that (see the upcoming (77) for a more detailed explanation)

$$\mathbf{E}_{\sigma \sim \mu_1} \left[|\sigma_{V(\hat{H}_G^k)}| \right] = \mathbf{E}_{\sigma \sim \mu_2} \left[|\sigma_{V(\hat{H}_G^k)}| \right] \quad \text{and} \quad \text{AVG-CUT}_{\mu_1}(H) = \text{AVG-CUT}_{\mu_2}(H). \quad (4)$$

Let $\text{AVG-CUT}_{\mu}(H)$ denote the common value of $\text{AVG-CUT}_{\mu_1}(H)$, $\text{AVG-CUT}_{\mu_2}(H)$ and Q be the common value of Q_1, Q_2 .

Let $\mathcal{D} := \mathcal{M}_{\beta, \gamma, \lambda}(H_{G, \mathcal{T}_1}^k) - \mathcal{M}_{\beta, \gamma, \lambda}(H_{G, \mathcal{T}_2}^k)$. From (2) and (4)

$$\mathcal{D} = 4k(\mathcal{A}'_1 - \mathcal{A}'_2)|E(H)| + (1 \pm 8\epsilon)k(M_1 - M_2)Q.$$

Let $L := |V(\mathcal{T}_1)| + |V(\mathcal{T}_2)|$. Using the purported algorithm for the magnetization on $H_{G, \mathcal{T}_i}^k$, and since the latter graph has at most $|V(H)||V(G)| + 4k|E(H)||V(\mathcal{T}_i)| \leq k(|V(G)| + 4L)|E(H)| =: X$ vertices, we can compute an estimate of $\mathcal{M}_{\beta, \gamma, \lambda}(H_{G, \mathcal{T}_i}^k)$ that is off by at most (an additive) κX . By subtracting these estimates for $i \in \{1, 2\}$, we obtain an estimate $\hat{\mathcal{D}}$ of $\mathcal{D}$ satisfying

$$|\hat{\mathcal{D}} - \mathcal{D}| \leq 2k(|V(G)| + 4L)|E(H)|\kappa$$

for some small κ . Then, we have that $\widehat{\text{MC}} = \frac{\hat{\mathcal{D}} - 4k(\mathcal{A}'_1 - \mathcal{A}'_2)|E(H)|}{k(M_1 - M_2)D} - \frac{C}{D}|E(H)|$ satisfies

$$|\widehat{\text{MC}} - \text{AVG-CUT}_{\mu}(H)| \leq \left(\frac{2(|V(G)| + 4L)}{|M_1 - M_2|} \kappa + 8\left(1 + \frac{|C|}{|D|}\right)\epsilon \right) |E(H)|, \quad (5)$$

where $\text{AVG-CUT}_{\mu}(H)$ denotes the common value of $\text{AVG-CUT}_{\mu_1}(H)$, $\text{AVG-CUT}_{\mu_2}(H)$ (see (4)). By choosing k sufficiently large, we have from (3) that $\text{AVG-CUT}_{\mu}(H)$ is within a factor arbitrarily close to 1 from $\text{MAX-CUT}(H)$. We will also choose $\epsilon > 0$ to be arbitrarily close to 0. This has the potential effect of increasing the size of G , but by choosing $\kappa > 0$ to be sufficiently small we can nevertheless ensure from (5) that our approximation $\widehat{\text{MC}}$ is within a factor arbitrarily close to 1 from $\text{AVG-CUT}_{\mu}(H)$, and hence from $\text{MAX-CUT}(H)$ as well. This finishes the contradiction argument and completes the proof of Theorem 12. $\square$

With Theorem 12, we can easily conclude the inapproximability result for the average-size of an independent set. Theorems 3, 5, 8 will follow analogously, once we give the field gadget constructions of Section 2.1. The proofs of these Theorems are given in Section 5.4.

Proof of Theorem 1. By Example 11, we have two trees $\mathcal{T}_1$ and $\mathcal{T}_2$ with the same effective fields and different magnetization gaps. Also, independent sets correspond to $\beta = 1, \gamma = 0, \lambda = 1$ and it is well-known (see, e.g., [15]) that $(\beta, \gamma, \lambda) \in \mathcal{U}_{\Delta}^*$ iff $\Delta \geq 6$. Therefore, the desired inapproximability result follows by applying Theorem 12. $\square$

The proof of Theorems 2, 4, and 7 build on similar reduction ideas, but the details are slightly more technical because of the different type of gadgets that we use (cf. Theorem 13). The details can be found in Section 5.5.

3 Field Gadget construction

In this section, we give the proof of Theorem 13 by first outlining the key ideas of our field-gadget constructions. We start with the following lemma which describes the effective field and magnetization gap of field gadgets built out of smaller field gadgets. The proof is given in Section 4.1.

Lemma 17. *Suppose we have field gadgets $\mathcal{T}_1, \dots, \mathcal{T}_k$ with roots $\rho_1, \dots, \rho_k$, effective fields $R_1, \dots, R_k$ and magnetization gaps $M_1, \dots, M_k$, respectively. Let $\mathcal{T}$ be the tree with root ρ , edge $\{\rho, u\}$ and the roots of $\mathcal{T}_1, \dots, \mathcal{T}_k$ identified with u . Let R and M be the effective field and magnetization gap for $\mathcal{T}$. Then*

$$R = \frac{1 + \gamma \lambda \prod_{i=1}^k R_i}{\beta + \lambda \prod_{i=1}^k R_i}, \quad M = 1 - \omega(R) \left(1 + \sum_{i=1}^k (M_i - 1) \right),$$

where

$$\omega(R) := \frac{1 + \beta\gamma - \beta R - \gamma/R}{1 - \beta\gamma}.$$

For antiferromagnetic (β, γ) and any $\lambda > 0$, it holds that $R \in (\gamma, 1/\beta)$ and $0 < \omega(R) < 1$.

The first step in the construction is to create a family of field gadgets with sufficiently dense effective fields in an interval and magnetization gaps that are in a small interval. To define the interval, consider x^* and ω^* defined as follows

$$x^* = \frac{1 + \gamma \lambda (x^*)^2}{\beta + \lambda (x^*)^2}, \quad \omega^* = \omega(x^*) = \frac{1 + \beta\gamma - \beta x^* - \gamma/x^*}{1 - \beta\gamma}, \quad (6)$$

and note that $\omega^* \in (0, 1)$ from Lemma 17.

Lemma 18. *Fix $\lambda > 0$. For any $\delta > 0$ and any sufficiently small $\tau_1 > 0$ we can find $\tau \in (0, \tau_1)$ and construct a family of field gadgets $\mathcal{L} = \{\mathcal{T}_1, \dots, \mathcal{T}_k\}$ whose effective fields are in the interval $[x^* - \tau, x^* + \tau]$ and for any $x \in [x^* - \tau, x^* + \tau]$ there exist a field gadget in the family whose effective field is in the interval $[x - \tau\delta, x + \tau\delta]$.*

The proof of Lemma 18 builds on techniques from [3] and is given in Section 4.2.

The second step of the construction is to use the field gadgets in $\mathcal{L}$ constructed in Lemma 18 in an iterative way as follows. At time $t = 0$, we will start with some field gadget $\mathcal{T}_0$. Suppose that at some stage we have constructed a rooted field gadget $\mathcal{T}$ with effective field R and magnetization gap M . To construct a new rooted field gadget, we merge $\mathcal{T}$ with one field gadget from $\mathcal{L}$ using the operation from Lemma 17. We are going to analyze what pairs (effective field, magnetization gap) we can achieve with this procedure. Let R_i and M_i be the effective field and magnetization gap for the field gadget $\mathcal{T}_i$ in our collection. By Lemma 17, merging the field gadget $\mathcal{T}$ with $\mathcal{T}_i$ yields a rooted field gadget with effective field $\phi_i(R)$ and magnetization gap $\psi_i(R, M)$ where the pair of maps (ϕ_i, ψ_i) are given by

$$\phi_i(R) = \frac{1 + \gamma \lambda R R_i}{\beta + \lambda R R_i} \quad \text{and} \quad \psi_i(R, M) = 1 - \omega(\phi_i(R))(M + M_i - 1). \quad (7)$$

For a small constant $\tau > 0$ to be specified later, let

$$I' = \left[x^* - 2\tau \frac{|\omega^*|}{1 - |\omega^*|}, x^* + 2\tau \frac{|\omega^*|}{1 - |\omega^*|} \right] \quad (8)$$

The choice of the interval I' is such that the maps ϕ_i are uniformly contracting and map the interval I' to itself. Namely, we show the following in Section 4.3.

Lemma 19. *There exists $0 < C_{\min} < C_{\max} < 1$ and $\tau_0 > 0$ such that for all $\tau \in (0, \tau_0)$ for all $R \in I'$ and all $R_i \in [x^* - \tau, x^* + \tau]$ it holds that*

$$C_{\min} \leq |\phi'_i(R)| \leq C_{\max}, \quad \omega(R) \leq C_{\max}, \quad \phi_i(R) \in I'$$

The second step of our construction will actually take place in the following smaller sub-interval of I' :

$$I = [x^* - |\omega|\tau/2, x^* + |\omega|\tau/2] \subseteq I', \quad (9)$$

The choice of I is to ensure the following “well-covered” property for the maps ϕ_i on I which is obtained using the density of the effective fields from Lemma 18, the proof is given in Section 4.3.

Lemma 20. *Suppose $\delta < |\omega|/100$. For every two points $x_1, x_2 \in I$ such that $|x_1 - x_2| \leq |\omega|\tau\delta/2$ there exists i such that $x_1, x_2 \in \phi_i(I)$.*

Lemmas 19 and 20 ensure that for any $x \in I$ we can construct a sequence of field gadgets whose effective field approaches x . Consider the following process **Build-gadget**(x, t) where $x \in I$ and $t \geq 0$ is an integer. If $t = 0$ we return the degenerate tree. If $t \geq 1$ then we let ϕ_i be any map such that $x \in \phi_i(I)$. Let $y = \phi_i^{-1}(x)$ and $\mathcal{T}' = \text{Build-gadget}(y, t-1)$. Return the tree $\mathcal{T}$ obtained by merging $\mathcal{T}'$ and $\mathcal{T}_i$ using the operation of Lemma 17.

The point behind the process **Build-gadget**(x, t) is that it allows us to construct, for arbitrary $x \in I$, a field gadget whose effective field is arbitrary close to x , with error that decays exponentially fast with t (using the contraction properties of the ϕ_i 's). This is detailed in the following lemma.

Lemma 21. *There exists $C > 0$ such that for any $x \in I$ and any $t \geq 0$ the effective field R of the field gadget returned by **Build-gadget**(x, t) (for any choice of the ϕ_i 's) satisfies $|R - x| \leq CC_{\max}^t$.*

We will prove Lemma 21 in Section 4.4. For the field gadgets we construct we will always maintain the effective field in the interval I' , cf. Lemma 19. Now we show that the magnetization gaps of the field gadgets constructed using our process also stay restricted to an interval J . Let

$$T = \frac{2 + \max |M_i|}{1 - C_{\max}} \text{ and let } J \text{ be the interval } [-T, T]. \quad (10)$$

Lemma 22. *Suppose $R \in I'$ and $M \in J$ then $\psi_i(R, M) \in J$.*

Proof. We have $\psi_i(R, M) = 1 - \omega(R)(M_i + M - 1)$. Hence $|\psi_i(R, M)| \leq 1 + C_{\max}(T + \max |M_i| + 1) \leq T$, where the last inequality follows from (10). $\square$

For any $x \in I$ we are going to construct a sequence of families of pairs of maps $\mathcal{F}_{x,0}, \mathcal{F}_{x,1}, \dots$ as follows. In each pair, the first map is from $\mathbb{R}$ to $\mathbb{R}$, and the second map is from $\mathbb{R}^2$ to $\mathbb{R}$ (similarly to (7)). Let $\mathcal{F}_{x,0}$ contain the pair of maps $(x \mapsto x, (x, y) \mapsto y)$. To construct $\mathcal{F}_{x,t+1}$ we take every ϕ_i such that $x \in \phi_i(I)$ and every $(f, g) \in \mathcal{F}_{\phi_i^{-1}(x),t}$ and place into $\mathcal{F}_{x,t+1}$ the map

$$(R, M) \mapsto (\phi_i(f(R)), \psi_i(f(R), g(R, M))). \quad (11)$$

Every $(f, g) \in \mathcal{F}_{x,t}$ corresponds to a sequence of ϕ_i 's with length t , which in turn corresponds to a rooted field gadget built using the procedure described just above (7) for t steps. In this procedure, we will now view the initial field gadget $\mathcal{T}_0$ as an “input” (to the procedure). If the input has effective field R and magnetization gap M then the final field gadget will have effective field $f(R)$ and magnetization gap $g(R, M)$. In particular, with the right choice of input (i.e., $f^{-1}(x)$), we will obtain a field gadget with effective field x . We will usually view $f(R)$ and $g(R, M)$ as the effective

field and the magnetization gap induced by the **Build-gadget**(x, t) procedure when at the base step we use $\mathcal{T}_0$ instead of the degenerate tree (provided the choice of the ϕ_i 's in **Build-gadget** matches up with the sequence for the pair (f, g)). However, we will not be able to use the exact input needed to obtain the effective field x but rather some approximation of it, and the pair (f, g) will allow us to track the influence of the last t steps when building a rooted field gadget $\mathcal{T}'$, which intuitively have larger influence both on the effective field and magnetization gap of $\mathcal{T}'$ (since the maps ϕ_i are contracting). In particular, once we have the value of the magnetization gap and effective field of the “input” field gadget, by applying (f, g) we know precisely where the effective field and magnetization gap will end up after applying the sequence of ϕ_i 's corresponding to (f, g) .

We are going to distinguish two possible cases for the families $\mathcal{F}_{x,t}$. In the first case we will obtain the gadgets we need immediately using the **Build-gadget** procedure. In the second case we will construct a continuous function from the families. Then we will argue that the function cannot satisfy a functional equation and this will yield the gadgets we require.

Lemma 23 (Case I). *Suppose there exists $x \in I$ such that for some t_1, t_2 there exist $(f_1, g_1) \in \mathcal{F}_{x,t_1}$ and $(f_2, g_2) \in \mathcal{F}_{x,t_2}$ such that $g_1(I' \times J)$ and $g_2(I' \times J)$ are disjoint. Let $\hat{M}$ be the distance of $g_1(I' \times J)$ and $g_2(I' \times J)$.*

Then, there is an algorithm which, on input a rational $r \in (0, 1/2)$, outputs in time $\text{poly}(\text{bits}(r))$ a pair of field gadgets $\mathcal{T}_1$ and $\mathcal{T}_2$, each of maximum degree 3 and size $O(|\log r|)$, such that

$$|R_{\mathcal{T}_1} - x|, |R_{\mathcal{T}_2} - x| \leq r, \text{ but } |M_{\mathcal{T}_1} - M_{\mathcal{T}_2}| \geq \hat{M}.$$

Proof. Note that for a fixed t the union $\bigcup_x \mathcal{F}_{x,t}$ contains finitely many functions (because there are only finitely many choices for ϕ_i in each step of the construction). If t_1, t_2 and x assumed by the lemma exist then we can find them by examining finitely many functions. For each (f_1, g_1) and (f_2, g_2) we check whether $f_1(I) \cap f_2(I) \neq \emptyset$ and $g_1(I' \times J) \cap g_2(I' \times J) = \emptyset$; if we find we such a pair we take $x \in f_1(I) \cap f_2(I)$. Note the running time for this process is a constant depending on β, γ, λ . (Note that $f(I)$ and $g(I' \times J)$ is always an interval that we can find inductively.)

To construct $\mathcal{T}_i$, $i \in \{1, 2\}$ we use the **Build-gadget** procedure following the choices of (f_i, g_i) in the first t_i steps, cf. the discussion below (11). Using Lemma 21, we run the procedure for $O(|\log r|)$ steps achieving $|R_i - x| \leq r$, $i \in \{1, 2\}$. We have $M_i \in g_i(I' \times J)$, $i \in \{1, 2\}$ and hence $|M_1 - M_2| \geq \hat{M}$. $\square$

Lemma 24 (Case II). *Suppose that for every x , every t_1, t_2 and every two functions $(f_1, g_1) \in \mathcal{F}_{x,t_1}$ and $(f_2, g_2) \in \mathcal{F}_{x,t_2}$ we have that $f_2(I' \times J)$ and $g_2(I' \times J)$ intersect. Then there exists a continuous function $F : I \rightarrow J$ such that for every $x \in I$, every $\epsilon > 0$ there exists t_0 such that for every $t \geq t_0$ and every $(f, g) \in \mathcal{F}_{x,t}$ and every $R \in I'$ and every $M \in J$ we have*

$$|g(R, M) - F(x)| \leq \epsilon. \quad (12)$$

The proof of Lemma 24 is given in Section 4.4.

Lemma 25. *Suppose CASE II happens, that is, the assumption of Lemma 24 is satisfied; let F be the continuous function guaranteed by Lemma 24. Suppose that there exist $x_1, x_2 \in I$ such that the following equation is violated.*

$$F\left(\frac{1 + \gamma\lambda x_1 x_2}{\beta + \lambda x_1 x_2}\right) = 1 - \frac{(1 - \beta\gamma)\lambda x_1 x_2}{(1 + \gamma\lambda x_1 x_2)(\beta + \lambda x_1 x_2)}(F(x_1) + F(x_2) - 1) \quad (13)$$

Then, for any integer $k \geq 1$, we can find in constant time (where the constant depends on $\beta, \gamma, \lambda, k$) rational numbers $x_1, x_2 \in I$ such that (13) is violated and, moreover rationals $\hat{R}_1, \dots, \hat{R}_k \in I$ such that the following holds.

There is an algorithm which on input a rational $r \in (0, 1/2)$ and any $i \in [k]$ outputs in time $\text{poly}(\text{bits}(r))$ a pair of field gadgets $\mathcal{T}_1$ and $\mathcal{T}_2$, each of maximum degree 3 and size $O(|\log r|)$, such that $|R_{\mathcal{T}_1} - \hat{R}_i| \leq r$, $|R_{\mathcal{T}_2} - \hat{R}_i| \leq r$ and $|M_{\mathcal{T}_1} - M_{\mathcal{T}_2}| \geq \hat{M}$.

Proof of Lemma 25. Let $k \geq 1$ be an arbitrary integer and fix $x_1, x_2 \in I$ that violate (13). Let $\delta > 0$ be the absolute value of the difference between the two sides. From the continuity of F , more precisely equation (53), there exists $\epsilon > 0$ such that for any $y_1, y_2 \in I$ with $|y_1 - x_1| \leq \epsilon$ and $|y_2 - x_2| \leq \epsilon$ we have that y_1, y_2 violate (13) with difference at least $\delta/2$ between the two sides.

Let $\mathcal{C}$ be a finite set of pairs (y_1, y_2) which form an $\frac{\epsilon}{20k}$ -net for $I \times I$; by perturbing slightly the set of points in $\mathcal{C}$ we can obtain an $\frac{\epsilon}{10k}$ -net for $I \times I$, say $\mathcal{C}'$, such that for any two pairs (y_1, y_2) and (y'_1, y'_2) it holds that $y_1 y_2 \neq y'_1 y'_2$. Now, to check whether a pair (y_1, y_2) in $\mathcal{C}'$ violates (12), we run **Build-gadget** for y_1, y_2 and $y_3 := \frac{1+\gamma\lambda y_1 y_2}{\beta+\lambda y_1 y_2}$ with the value of t given by Lemma 24 to achieve bound $\delta/100$ on the right-hand side of (12). Then, we will find at least k different pairs $(y_{1,j}, y_{2,j})$, $j = 1, \dots, k$ that violate (13). For $j \in [k]$, let $\hat{R}_j := y_{3,j} = \frac{1+\gamma\lambda y_{1,j} y_{2,j}}{\beta+\lambda y_{1,j} y_{2,j}}$ and note that by the construction of $\mathcal{C}'$, the $\hat{R}_j$'s are pairwise distinct.

Now, on input $j \in [k]$, we use the **Build-gadget** procedure to construct $\hat{\mathcal{T}}_i$ for $y_{i,j}$, $i \in \{1, 2, 3\}$ for $t = O(|\log r|)$ steps, using Lemma 21. The tree $\mathcal{T}_1$ is obtained by merging $\hat{\mathcal{T}}_1$ and $\hat{\mathcal{T}}_2$ and the tree $\mathcal{T}_2$ is $\hat{\mathcal{T}}_3$. $\square$

Now assume that equation (13) is satisfied for all $x_1, x_2 \in I$. We are going to derive a contradiction, thereby showing that (13) must be violated. We will do this in two steps. First we use a special case of equation (13) to obtain a functional equation that constrains the possible solutions of F . Second we show that none of these solutions satisfies (13).

Suppose $x_1, x_2, x_3 \in I$ are such that $x_1 x_2 = x_3 x^*$. Plugging into (13) we obtain that F has to satisfy the following equation

$$F(x_1) + F(x_2) = F(x_3) + F(x^*). \quad (14)$$

Lemma 26. *Suppose F is a continuous function on I . Suppose that for $x_1, x_2, x_3 \in I$ such that $x_1 x_2 = x_3 x^*$ we have (14). Then there exists c such that for all $x \in I$ we have*

$$F(x) = c \log(x/x^*) + F(x^*). \quad (15)$$

Proof. We will use the following parametrization to turn (14) into Cauchy's functional equation. Let $x_1 = x^* \exp(y_1)$, $x_2 = 2x^* \exp(y_2)$, and $x_3 = x^* \exp(y_3)$. The condition $x_1 x_2 = x^* x_3$ is equivalent to $y_1 + y_2 = y_3$. Let

$$G(y) = F(x^* \exp y) - F(x^*). \quad (16)$$

Note that G is defined on the interval $[\log(I_L/x^*), \log(I_R/x^*)]$ that contains 0 (since $I_L < x^* < I_R$). Equation (14) becomes

$$G(y_1) + G(y_2) = G(y_1 + y_2). \quad (17)$$

From continuity of F we have continuity of G . Since (17) is Cauchy's functional equation on an interval containing zero the only continuous solutions are $G(y) = cy$ for some constant c . Plugging in (16) we obtain (15). $\square$

Finally we show that (13) has to be violated.

Lemma 27. *A solution of the form (15) cannot satisfy (13).*

Proof. We will plug-in (15) into (13) with $x_1 = x^*$ and $x_2 = y$. We obtain

$$c \log \left(\frac{1 + \gamma \lambda x^* y}{\beta + \lambda x^* y} \right) - c \log x^* + F(x^*) = 1 - \frac{(1 - \beta \gamma) \lambda x^* y}{(1 + \gamma \lambda x^* y)(\beta + \lambda x^* y)} (c \log(y) - c \log(x^*) - 1 + 2F(x^*)). \quad (18)$$

Differentiating (w.r.t. y) we obtain

$$\begin{aligned} -c \frac{\lambda(1 - \beta \gamma) x^*}{(\beta + \lambda x^* y)(1 + \gamma \lambda x^* y)} &= \left(-\frac{(1 - \beta \gamma) \lambda x^* (-\beta + \gamma \lambda^2 (x^* y)^2)}{(1 + \gamma \lambda x^* y)^2 (\beta + \lambda x^* y)^2} \right) (c \log(y) - c \log(x^*) - 1 + 2F(x^*)) \\ &\quad - \frac{(1 - \beta \gamma) \lambda x^* y}{(1 + \gamma \lambda x^* y)(\beta + \lambda x^* y)} \left(c \frac{1}{y} \right), \end{aligned}$$

which simplifies to

$$0 = \frac{(1 - \beta \gamma) \lambda x^* (-\beta + \gamma \lambda^2 (x^* y)^2)}{(1 + \gamma \lambda x^* y)^2 (\beta + \lambda x^* y)^2} (c \log(y) - c \log(x^*) - 1 + 2F(x^*)) \quad (19)$$

The only solution for (19) has to have $c = 0$ and $F(x^*) = 1/2$. Plugging $c = 0$ and $F(x^*) = 1/2$ into (18) we obtain $1/2 = 1 + 0$, a contradiction. $\square$

Finally, we combine the various pieces from the previous subsection to prove Theorem 13. In fact, we will prove a slight strengthening of Theorem 13, given below, which will be used in our reduction, analogously to Lemma 16.

Theorem 28. *Let (β, γ, λ) be antiferromagnetic with $(\beta, \gamma, \lambda) \neq (\beta, \beta, 1)$. For every integer $k \geq 1$, there exist constants $\hat{M}, \Xi > 0$ and k distinct numbers $\hat{R}_1, \dots, \hat{R}_k > 0$ such that the following holds. There is an algorithm, which, on input $i \in [k]$ and a rational $r \in (0, 1/2)$, outputs in time $\text{poly}(\text{bits}(r))$ a pair of field gadgets $\mathcal{T}_1, \mathcal{T}_2$, each of maximum degree 3 and size $O(|\log r|)$, such that*

$$|R_{\mathcal{T}_1} - \hat{R}_i|, |R_{\mathcal{T}_2} - \hat{R}_i| \leq r, \text{ but } |M_{\mathcal{T}_1} - M_{\mathcal{T}_2}| \geq \hat{M}.$$

Moreover, the magnetization gaps $M_{\mathcal{T}_1}, M_{\mathcal{T}_2}$ are bounded in absolute value by the constant Ξ .

Proof. There are two cases to consider: either Lemma 23 or Lemma 24. In the latter case, by Lemmas 26 and 27, we obtain that (13) is violated, and hence Lemma 25 yields the desired algorithm. In the former case, we are also done, modulo that Lemma 23 only guarantees the existence of a single $\hat{R}$, namely the value x . Let t_1, t_2 be such that $(f_1, g_1) \in \mathcal{F}_{x, t_1}$ and $(f_2, g_2) \in \mathcal{F}_{x, t_2}$ satisfy $g_1(I' \times J) \cap g_2(I' \times J) = \emptyset$. The functions in $\mathcal{F}_{x, t}$ are continuous and defined on a closed interval, so we have that for sufficiently small $\epsilon > 0$, for all y such that $|y - x| \leq \epsilon$, we can find $(\tilde{f}_1, \tilde{g}_1) \in \mathcal{F}_{y, t_1}$ and $(\tilde{f}_2, \tilde{g}_2) \in \mathcal{F}_{y, t_2}$ which satisfy $\tilde{g}_1(I' \times J) \cap \tilde{g}_2(I' \times J) = \emptyset$. We pick k such y 's for the values of the $\hat{R}_i$'s for $i = 1, \dots, k$ and produce the desired trees by running the **Build-gadget** procedure for $O(\log |r|)$ steps. Finally note that the magnetizations all lie in the interval J , see (10) and Lemma 22, which is bounded by absolute constants, finishing the proof. $\square$

Note that Theorem 13 corresponds to the case $k = 1$ in Theorem 28.

3.1 Proof of Theorem 14

To prove Theorem 14, we will just need an absolute bound on the second derivative of the effective-field constructions of the previous subsection, when viewed as functions of the parameter λ . The following lemma gives this bound.

Lemma 29. Let $R_1(\lambda), \dots, R_k(\lambda)$ be a family of functions from a closed interval I to $(\gamma, 1/\beta)$. Consider the following family $\mathcal{F}$ of functions from I to $(\gamma, 1/\beta)$. Let $Q_0(\lambda)$ be a function from I to $(\gamma, 1/\beta)$. For every sequence $a_0, \dots, a_{m-1}$ of numbers in $[k]$ for $i \in [m-1]$ let

$$Q_{i+1}(\lambda) = \frac{1 + \gamma \lambda R_{a_i}(\lambda) Q_i(\lambda)}{\beta + \lambda R_{a_i}(\lambda) Q_i(\lambda)}.$$

and place function Q_m into $\mathcal{F}$.

Suppose each $R_\ell(\lambda)$ and $Q_0(\lambda)$ is twice continuously differentiable on I . There exists bounded intervals I_1, I_2 such that for any $\lambda \in I$ and any $Q(\lambda) \in \mathcal{F}$ we have $Q'_i(\lambda) \in I_1$ and $Q''_i(\lambda) \in I_2$.

Proof. Let $\ell = a_i$. We have

$$Q'_{i+1}(\lambda) = \frac{(\gamma \lambda R'_\ell(\lambda) Q_i(\lambda) + \gamma \lambda R_\ell(\lambda) Q'_i(\lambda))}{\beta + \lambda R_i(\lambda) Q_i(\lambda)} - \frac{(1 + \gamma \lambda R_\ell(\lambda) Q_i(\lambda))(\lambda R'_\ell(\lambda) R(\lambda) + \lambda R_\ell(\lambda) Q'_i(\lambda))}{(\beta + \lambda R_\ell(\lambda) Q_i(\lambda))^2}. \quad (20)$$

We can re-write (20) as follows

$$\frac{Q'_{i+1}(\lambda)}{Q_{i+1}(\lambda)} = \frac{\lambda R_\ell(\lambda) Q_i(\lambda) (\beta \gamma - 1)}{(1 + \gamma \lambda R_\ell(\lambda) Q_i(\lambda))(\beta + \lambda R_\ell(\lambda) Q_i(\lambda))} + \frac{p_1(R_\ell(\lambda), Q_i(\lambda), R'_\ell(\lambda))}{(\beta + \lambda R_\ell(\lambda) Q_i(\lambda))}, \quad (21)$$

where p_1 is a constant degree polynomial. Note that the coefficient of $Q'_i(\lambda)/Q_i(\lambda)$ in (20) satisfies

$$\left| \frac{\lambda R_\ell(\lambda) Q_i(\lambda) (\beta \gamma - 1)}{(1 + \gamma \lambda R_\ell(\lambda) Q_i(\lambda))(\beta + \lambda R_\ell(\lambda) Q_i(\lambda))} \right| \leq C < 1, \quad (22)$$

where $C = \frac{1 - \sqrt{\beta \gamma}}{\sqrt{\beta \gamma} + 1}$ if $\beta \gamma \neq 0$ (the case $\beta \gamma = 0$ is discussed below). To see the bound in (22) let $W = \lambda R_\ell(\lambda) Q_i(\lambda)$. The coefficient becomes $\frac{W(\beta \gamma - 1)}{(1 + \gamma W)(\beta + W)}$, which achieves maximum for $W = \sqrt{\beta/\gamma}$ with value $\frac{\sqrt{\beta \gamma} - 1}{\sqrt{\beta \gamma} + 1}$. In the case $\beta = 0$ the coefficient becomes $-1/(W\gamma + 1)$, which since W is uniformly bounded from below has absolute value uniformly bounded from 1. In the case $\gamma = 0$ the coefficient becomes $-W/(\beta + W)$ and since W is uniformly bounded from above has absolute value uniformly bounded from 1.

Each $R'_\ell(\lambda)$ is bounded (since it is a continuous function on a closed interval). Hence there exists constant M_1 such that for any $x \in (\gamma, 1/\beta)$ and any $\lambda \in I$

$$\left| \frac{p_1(R_\ell(\lambda), x, R'_\ell(\lambda))}{(\beta + \lambda R_\ell(\lambda) x)^2} \right| \leq M_1.$$

From continuous differentiability of Q_0 we also have that there exists a constant M_2 such that for any $\lambda \in I$ we have

$$\left| \frac{Q'_0(\lambda)}{Q_0(\lambda)} \right| \leq M_2.$$

Now using (21) by induction $\left| \frac{Q'_i(\lambda)}{Q_i(\lambda)} \right| \leq (M_1 + M_2)/(1 - C)$.

The argument for the second derivative is almost the same. We can write an expression for $Q''_{i+1}(\lambda)/Q_{i+1}(\lambda)$ as follows

$$\frac{Q''_{i+1}(\lambda)}{Q_{i+1}(\lambda)} = \frac{\lambda R_\ell(\lambda) Q_i(\lambda) (\beta \gamma - 1)}{(1 + \gamma \lambda R_\ell(\lambda) Q_i(\lambda))(\beta + \lambda R_\ell(\lambda) Q_i(\lambda))} \frac{Q''_i(\lambda)}{Q_i(\lambda)} + \frac{p_2(R_\ell(\lambda), Q'_i(\lambda), Q_i(\lambda), R'_\ell(\lambda), R''_\ell(\lambda))}{(\beta + \lambda R_\ell(\lambda) Q_i(\lambda))^3}, \quad (23)$$

where p_2 is a constant degree polynomial. Again all the parameters for p_2 are bounded (for $Q'_i(\lambda)$ we use the $(M_1 + M_2)/(1 - C)(1/\beta)$ bound) and hence there exists a constant M_3 such that for any $x \in (\gamma, 1/\beta)$ and any $|y| \leq (M_1 + M_2)/(1 - C)$ we have

$$\left| \frac{p_2(R_\ell(\lambda), y, x, R'_\ell(\lambda), R''_\ell(\lambda))}{(\beta + \lambda R_\ell(\lambda)x)^3} \right| \leq M_3.$$

We also have that there exists a constant M_4 such that for any $\lambda \in I$ we have

$$\left| \frac{Q''_0(\lambda)}{Q_0(\lambda)} \right| \leq M_4.$$

Note that the coefficient of $Q''_i(\lambda)/Q_i(\lambda)$ in (23) satisfies (22). Now using (23) by induction $|Q''_i(\lambda)/Q_i(\lambda)| \leq (M_3 + M_4)/(1 - C)$. $\square$

We now give the proof of Theorem 14.

Proof of Theorem 14. By Theorem 13, for antiferromagnetic (β, γ) and $\lambda > 0$ with $(\beta, \gamma, \lambda) \neq (\beta, \beta, 1)$, there exist $\hat{R}, \hat{M} > 0$ such that for any $r > 0$, there exist a pair of field gadgets $\mathcal{T}_1, \mathcal{T}_2$, each of maximum degree 3 and size $O(|\log r|)$, such that

$$|R_{\mathcal{T}_1} - \hat{R}|, |R_{\mathcal{T}_2} - \hat{R}| \leq r, \text{ but } |M_{\mathcal{T}_1} - M_{\mathcal{T}_2}| > \hat{M}.$$

The effective fugacities of the gadgets $\mathcal{T}_1, \mathcal{T}_2$, viewed as functions of λ , are of the form given in Lemma 29 and hence their second derivatives are uniformly bounded. Note that the magnetization gap is the first derivative of the logarithm of the effective fugacity. So, the lower bound on the difference of $M_{\mathcal{T}_1}, M_{\mathcal{T}_2}$ implies an absolute lower bound on the difference of the derivatives of $R_{\mathcal{T}_1}, R_{\mathcal{T}_2}$. Since the second derivative is uniformly bounded (by Lemma 29), we can find $\hat{\lambda}$ arbitrarily close to λ such that $R_{\mathcal{T}_1}(\hat{\lambda}), R_{\mathcal{T}_2}(\hat{\lambda})$ are equal. Since $R_{\mathcal{T}_1}, R_{\mathcal{T}_2}(\lambda)$ are rational functions of λ (using that β, γ are rational), we obtain that any such $\hat{\lambda}$ must in fact be algebraic, finishing the proof. $\square$

4 Field Gadgets: Remaining Proofs

In this section, we give the remaining proofs of Section 3.

4.1 Proof of Lemma 17

Proof of Lemma 17. For convenience, we will write $\mu_{\mathcal{T}}$ instead of $\mu_{\mathcal{T};\beta,\gamma,\lambda}$. Let $Z_{\mathcal{T},\rho}^{\text{in}}$ be the weight of all configurations on $\mathcal{T}$ rooted at ρ where ρ is occupied (has state 1). Similarly let $Z_{\mathcal{T},r}^{\text{out}}$ be the weight of all configurations on $\mathcal{T}$ rooted at ρ where ρ is unoccupied (has state 1). Define analogously $Z_{\mathcal{T}_i,\rho_i}^{\text{in}}$ for $i = 1, \dots, k$. Then, we have

$$\begin{aligned} Z_{\mathcal{T},\rho}^{\text{in}} &= \lambda \left(\prod_{i=1}^k Z_{\mathcal{T}_i,\rho_i}^{\text{out}} + \gamma \lambda \prod_{i=1}^k \frac{Z_{\mathcal{T}_i,\rho_i}^{\text{in}}}{\lambda} \right) \\ Z_{\mathcal{T},\rho}^{\text{out}} &= \beta \prod_{i=1}^k Z_{\mathcal{T}_i,\rho_i}^{\text{out}} + \lambda \prod_{i=1}^k \frac{Z_{\mathcal{T}_i,\rho_i}^{\text{in}}}{\lambda} \end{aligned}$$

We therefore have that

$$R = \frac{1}{\lambda} \frac{Z_{\mathcal{T},\rho}^{\text{in}}}{Z_{\mathcal{T},\rho}^{\text{out}}} = \frac{\prod_{i=1}^k Z_{\mathcal{T}_i,\rho_i}^{\text{out}} + \gamma \lambda \prod_{i=1}^k \frac{Z_{\mathcal{T}_i,\rho_i}^{\text{in}}}{\lambda}}{\beta \prod_{i=1}^k Z_{\mathcal{T}_i,\rho_i}^{\text{out}} + \lambda \prod_{i=1}^k \frac{Z_{\mathcal{T}_i,\rho_i}^{\text{in}}}{\lambda}} = \frac{1 + \gamma \lambda \prod_{i=1}^k R_i}{\beta + \lambda \prod_{i=1}^k R_i}.$$

For the magnetization gap, we view $Z_{\mathcal{T},\rho}^{\text{in}}, Z_{\mathcal{T},\rho}^{\text{out}}$, and $Z_{\mathcal{T}_i,\rho_i}^{\text{in}}, Z_{\mathcal{T}_i,\rho_i}^{\text{out}}$ as polynomials in λ , and therefore the R_i 's and R as rational functions of λ . Then observe that

$$M_i = \lambda \frac{\partial \log Z_{\mathcal{T}_i,\rho_i}^{\text{in}}}{\partial \lambda} - \lambda \frac{\partial \log Z_{\mathcal{T}_i,\rho_i}^{\text{out}}}{\partial \lambda} = \lambda \frac{\partial \log(\lambda R_i)}{\partial \lambda} = 1 + \frac{\lambda}{R_i} \frac{\partial R_i}{\partial \lambda},$$

and similarly $M = 1 + \frac{\lambda}{R} \frac{\partial R}{\partial \lambda}$. It follows that

$$\begin{aligned} M &= 1 - \frac{(1 - \beta\gamma)\lambda}{R(\beta + \lambda \prod_{i=1}^k R_i)^2} \frac{\partial(\lambda \prod_{i=1}^k R_i)}{\partial \lambda} = 1 - \frac{(1 - \beta\gamma)\lambda}{(1 + \gamma\lambda \prod_{i=1}^k R_i)(\beta + \lambda \prod_{i=1}^k R_i)} \frac{\partial(\lambda \prod_{i=1}^k R_i)}{\partial \lambda} \\ &= 1 - \frac{(1 - \beta\gamma)\lambda \prod_{i=1}^k R_i}{(1 + \gamma\lambda \prod_{i=1}^k R_i)(\beta + \lambda \prod_{i=1}^k R_i)} \left(1 + \sum_{i=1}^k \frac{\lambda}{R_i} \frac{\partial R_i}{\partial \lambda}\right) = 1 - \omega(R) \left(1 + \sum_{i=1}^k (M_i - 1)\right). \quad \square \end{aligned}$$

4.2 Proof of Lemma 18

In this section, we give the proof of Lemma 18. We first create two field gadgets $\mathcal{T}_1$ and $\mathcal{T}_2$ with effective fields R_1, R_2 where R_1 is close to x^* and R_2 is even closer to x^* . We generalize a “bounding-tree” construction from [11], which was used in the case of the hard-core model.

To handle the special case $\lambda = \frac{1-\beta}{1-\gamma}$ for some $\beta, \gamma \in (0, 1)$, cf. Definition 9 and the remark below, we will use the following field gadget whose effective field is different than 1 and satisfies $R \in (\gamma, 1/\beta)$.

Lemma 30. *Let (β, γ) be antiferromagnetic with $\beta, \gamma \in (0, 1)$ and $\beta \neq \gamma$. Let $\lambda = \frac{1-\beta}{1-\gamma}$ and consider the field gadget $\mathcal{T}$ rooted at ρ , consisting of a 4-cycle together with the extra vertex ρ that is attached to one of the vertices of the cycle. Then, the effective field R of $\mathcal{T}$ is different than 1 and satisfies $R \in (\gamma, 1/\beta)$.*

Proof. A direct calculation gives that $R = \frac{1+\gamma\lambda r}{\beta+\lambda r}$ where $r = \frac{\beta^2+\lambda+2\lambda\beta\gamma+3\lambda^2\gamma^2+\lambda^3\gamma^4}{\beta^4+3\lambda\beta^2+\lambda^2+2\lambda^2\beta\gamma+\lambda^3\gamma^2}$. Since $\lambda = \frac{1-\beta}{1-\gamma}$, we have that $R \neq 1$ iff $r \neq 1$. Using the value of λ again we obtain that $r = 1$ is equivalent to $(1 - \beta)^3(\beta - \gamma) = 0$, which cannot be the case when $\beta \neq \gamma$ and $\beta, \gamma \in (0, 1)$. The fact that $R \in (\gamma, 1/\beta)$ follows from $r \in (\gamma, 1/\beta)$. $\square$

Lemma 31. *For every $\epsilon_1, \epsilon_2 > 0$ there exist field gadgets $\mathcal{T}_1$ and $\mathcal{T}_2$ such that*

$$|R_1 - x^*| < \epsilon_1 \quad \text{and} \quad |R_2 - x^*| < \epsilon_2 |R_1 - x^*|. \quad (24)$$

Proof. We will inductively construct two sequences of field gadgets $\mathcal{T}_1^L, \mathcal{T}_2^L, \dots$ and $\mathcal{T}_1^U, \mathcal{T}_2^U, \dots$ with effective fields $L_1, L_2, \dots$ and $U_1, U_2, \dots$. Throughout the construction we will maintain $1/\beta < L_i \leq x^* \leq U_i < \gamma$ and $L_i \neq U_i$. We will use the following inequalities

$$x < x^* \implies \frac{1 + \gamma\lambda x^2}{\beta + \lambda x^2} > x^* \quad \text{and} \quad x > x^* \implies \frac{1 + \gamma\lambda x^2}{\beta + \lambda x^2} < x^*; \quad (25)$$

these follow from the fact that $\frac{1+\gamma\lambda x^2}{\beta+\lambda x^2}$ is decreasing and that x^* is the fixpoint.

If $1 < x^*$ then we let $\mathcal{T}_1^L$ to be the degenerate tree (see Example 10) and $\mathcal{T}_1^U$ be the tree with one edge. Note that $\mathcal{T}_1^U$ is obtained from the degenerate tree using the operation of Lemma 17 and hence $L_1 = 1$ and $U_1 = (1 + \gamma\lambda)/(\beta + \lambda)$ which by (25) is greater than x^* . In the case $x^* < 1$ by the same argument we can take $\mathcal{T}_1^L$ the tree with one edge and $\mathcal{T}_1^U$ the degenerate tree and we again have $L_1 < x^* < U_1$. In the case where $x^* = 1$ we must have that $\lambda = \frac{1-\beta}{1-\gamma}$ for some $\beta, \gamma \in (0, 1)$

with $\beta \neq \gamma$; then, we can use the field gadget $\mathcal{T}$ of Lemma 30: if $R < 1$ then we let $\mathcal{T}_1^L$ be $\mathcal{T}$ and we take $\mathcal{T}_1^U$ to be the single edge tree, and if $R > 1$, vice versa.

Now assume $L_i \leq x^* \leq U_i$ and $L_i \neq U_i$. Using (25) we have that

$$\frac{1 + \gamma \lambda U_i^2}{\beta + \lambda U_i^2} \leq x^* \leq \frac{1 + \gamma \lambda L_i^2}{\beta + \lambda L_i^2}. \quad (26)$$

From the monotonicity of the map $x \mapsto (1 + \gamma \lambda x)/(\beta + \lambda x)$ and (26) we have that there is $k \in \{0, 1\}$ such that

$$L_{i+1} := \frac{1 + \gamma \lambda U_i^k L_i^{2-k}}{\beta + \lambda U_i^k L_i^{2-k}} \leq x^* \leq \frac{1 + \gamma \lambda U_i^{k+1} L_i^{1-k}}{\beta + \lambda U_i^{k+1} L_i^{1-k}} =: U_{i+1}. \quad (27)$$

Equation (27) corresponds to 1) creating $\mathcal{T}_{i+1}^L$ using k copies of $\mathcal{T}_i^U$ and $2 - k$ copies of $\mathcal{T}_i^L$ and 2) creating $\mathcal{T}_{i+1}^U$ using $k + 1$ copies of $\mathcal{T}_i^U$ and $1 - k$ copies of $\mathcal{T}_i^L$.

Note that $U_i \neq L_i$ implies $U_{i+1} \neq L_{i+1}$. We will now argue that $\lim_{n \rightarrow \infty} L_n = \lim_{n \rightarrow \infty} U_n = x^*$. From (27) we obtain

$$\frac{U_i}{L_i} = \frac{(L_{i+1}\beta - 1)(\gamma - U_{i+1})}{(U_{i+1}\beta - 1)(\gamma - L_{i+1})}. \quad (28)$$

We will now argue

$$\frac{(L_{i+1}\beta - 1)(\gamma - U_{i+1})}{(U_{i+1}\beta - 1)(\gamma - L_{i+1})} > \frac{U_{i+1}}{L_{i+1}}. \quad (29)$$

We can rewrite (29) as follows

$$(U_{i+1} - L_{i+1})(\gamma(1 - \beta L_{i+1}) + \beta L(U_{i+1} - \gamma)) > 0,$$

which follows from $R \in (\gamma, 1/\beta)$ (see Lemma 17). Note that U_i/L_i is a decreasing sequence, $U_i/L_i \geq 1$ and hence it has a limit. We are going to argue that the limit is 1. Suppose the limit is $\alpha \geq 1$. Note that the values of L_i are restricted to $(\gamma, 1/\beta)$ and hence the set $\{L_i \mid i \geq 1\}$ has an accumulation point x . At the accumulation point we have

$$\alpha = \frac{(x\beta - 1)(\gamma - \alpha x)}{(\alpha x\beta - 1)(\gamma - x)}, \quad (30)$$

since for every $\epsilon > 0$ there exist i such that $|L_{i+1} - x| \leq \epsilon$ and $|U_{i+1} - \alpha x| \leq \epsilon$ and $|U_i/L_i - \alpha| \leq \epsilon$ (which together with (28) implies (30)).

Equation (30) simplifies to

$$(1 - \alpha)(\alpha\beta\gamma x - \alpha\beta x^2 + \beta\gamma x - \gamma) = 0. \quad (31)$$

We will show that for $x \in (\gamma, 1/\beta)$ we have

$$\alpha\beta\gamma x - \alpha\beta x^2 + \beta\gamma x - \gamma < 0. \quad (32)$$

For $x = \gamma$ the left-hand side of (32) simplifies to $(\beta\gamma - 1)\gamma < 1$. The derivative of the left-hand side of (32) with respect to x is

$$\beta(\alpha\gamma - 2\alpha x + \gamma). \quad (33)$$

We will show that for $x \in (\gamma, 1/\beta)$ the expression (33) is negative; because of linearity we only need to argue at the endpoints of the interval. For $x = \gamma$ the value is $\gamma - \alpha\gamma < 0$. For $x = 1/\beta$ the value is $(1 + \alpha)\beta\gamma - 2\alpha < 0$. This proves (33) and hence (31) implies $\alpha = 1$, that is, the limit of U_i/L_i goes to 1 as $i \rightarrow \infty$.

To construct $\mathcal{T}_1$ we iterate the process until $U_i/L_i \leq 1 + \epsilon_1/x^*$; this implies $|U_i - x^*| \leq \epsilon_1$ and $|L_i - x^*| \leq \epsilon_1$. Since $U_i \neq L_i$ at least one of them is different from x^* ; this will be the effective field of our field gadget $\mathcal{T}_1$. To construct $\mathcal{T}_2$ we iterate the process until $U_j/L_j \leq 1 + \epsilon_2|R_1 - x^*|/x^*$ and take either the field gadget for L_j or U_j . $\square$

Proof of Lemma 18. We will use the Taylor approximation of the map of Lemma 17 for R_1, R_2 close to x^* . More precisely, there exists $C > 0$ and $\tau_0 > 0$ such that for any $\tau \in (0, \tau_0)$, $|y_1| \leq \tau$, and $|y_2| \leq \tau$ for $R_1 = x^* + y_1$ and $R_2 = x^* + y_2$ we have

$$\left| \frac{1 + \gamma \lambda R_1 R_2}{\beta + \lambda R_1 R_2} - R \right| \leq C\tau^2, \quad (34)$$

where $R = x^* - \omega^*(y_1 + y_2)$ (see equation (40)).

For $i \in \{0, \dots, Q\}$ (where $Q = \lceil 2 \log_2(1/\delta) + 2 \log_2(1/\omega^*) \rceil$) we will construct a sequence of field gadgets $\mathcal{T}_{i,0}, \dots, \mathcal{T}_{i,2^i}$ with effective fields $R_{i,0}, \dots, R_{i,2^i}$ that are close to an arithmetic sequence. More precisely for each $i \in \{0, \dots, Q\}$ there will be a value α_i and $\xi_i > 0$ such that for all $j \in \{0, \dots, 2^i\}$ we have

$$|R_{i,j} - (x^* + j\alpha_i)| \leq \xi_i |\alpha_i|. \quad (35)$$

We will aim to have ξ_i small; the exact bound will be established inductively. To simplify the notation we assume $C \geq 1$ (increasing C preserves (34)) and $\tau_0 < 1/100$ (decreasing τ_0 preserves (34)).

We start with $\mathcal{T}_{0,0} = \mathcal{T}_2$ and $\mathcal{T}_{0,1} = \mathcal{T}_1$ where $\mathcal{T}_1, \mathcal{T}_2$ are the field gadgets guaranteed by Lemma 31 with parameters $\epsilon_2 = \tau_0(\omega^*)^{2Q}/(C^{2Q}9^Q)$, $\epsilon_1 = \epsilon_2^2$. Note that we satisfy (35) with $\alpha_0 = R_{0,1}$ and $\xi_0 = \epsilon_2$.

Suppose we constructed $\mathcal{T}_{i,0}, \dots, \mathcal{T}_{i,2^i}$ and $\xi_i \leq 3^i \xi_0 < 1/100$. For $j \in \{0, \dots, 2^{i+1}\}$ we let $\mathcal{T}_{i+1,j}$ be obtained from $\mathcal{T}_{i,k}$ and $\mathcal{T}_{i,\ell}$ where $k + \ell = j$ (any choice of k and ℓ is fine). We let $\alpha_{i+1} = -\omega^* \alpha_i$. From (34) we have

$$|R_{i+1,j} - (x^* - \omega^*((R_{i,k} - x^*) + (R_{i,\ell} - x^*)))| \leq C(2^{i+1}|\alpha_i|)^2,$$

which combined with (35) for $R_{i,k}$ and $R_{i,\ell}$ implies

$$|R_{i+1,j} - (x^* + j\alpha_{i+1})| \leq C2^{2i+2}|\alpha_i|\xi_0\xi_i + 2\xi_i\omega^*|\alpha_i| \leq 3\xi_i|\alpha_{i+1}|, \quad (36)$$

where in the first inequality we used $|\alpha_i| \leq |\alpha_0| \leq \epsilon_1 = (\xi_0)^2 \leq \xi_0\xi_i$. Hence we have $\mathcal{T}_{i+1,0}, \dots, \mathcal{T}_{i+1,2^{i+1}}$ that satisfy (35) with $\xi_{i+1} \leq 3^{i+1}\xi_0 < 1/100$.

Our final family of field gadgets is $\mathcal{T}_{Q,i}, i \in \{0, \dots, 2^Q\}$ and $\mathcal{T}_{Q-1,i}, i \in \{0, \dots, 2^{Q-1}\}$ (the reason we take field gadgets for two values of the first index is to cover both sides of x^* ; note that α_i 's alternate sign). We take $\tau = 2^{Q-1}\omega^*|\alpha_Q|$. Let $x \in [x^* - \tau, x^* + \tau]$. If the sign of $x - x^*$ agrees with the sign of α_Q we let $k = \lfloor (x - x^*)/\alpha_Q \rfloor$ (note $0 \leq k \leq 2^{Q-1}$) and take the tree $\mathcal{T}_{Q,k}$. We have

$$|R_{Q,k} - (x^* + \alpha_Q k)| \leq (1/100)|\alpha_Q| \quad (37)$$

and hence

$$|R_{Q,k} - x| \leq 2|\alpha_Q| \leq 2^{-Q+2}\tau \leq \delta\tau. \quad (38)$$

If the sign of $x - x^*$ agrees with the sign of α_{Q-1} we let $k = \lfloor (x - x^*)/\alpha_{Q-1} \rfloor$ (note $0 \leq k \leq 2^{Q-1}$) and take the tree $\mathcal{T}_{Q-1,k}$. The argument is analogous to (37) and (38). $\square$

4.3 Proof of Lemmas 19 and 20

Proof of Lemma 19. We will consider the functions ϕ_i for R and R_i close to x^* . The gradient of $\phi_i(R)$ (viewed as a function of R and R_i) is

$$\left(\frac{\lambda R_i(\beta\gamma - 1)}{(\beta + \lambda R R_i)^2}, \frac{\lambda R(\beta\gamma - 1)}{(\beta + \lambda R R_i)^2} \right), \quad (39)$$

so we let $R_i = R = x^*$ in (39) then we obtain that the gradient of ϕ_i at (x^*, x^*) is (ω^*, ω^*) , since

$$\frac{\lambda x^*(\beta\gamma - 1)}{(\beta + \lambda(x^*)^2)^2} = \omega^*.$$

We therefore obtain the following linear approximation (using Taylor's theorem) of ϕ_i around x^* . There exists C such that for any sufficiently small $\tau > 0$ the following is true. Suppose $R_i = x^* + y_i$ and $R = x^* + y$ where $|y_i| \leq \tau$ and $|y| \leq \tau$. Then

$$|\phi_i(R) - (x^* + \omega^*(y_i + y))| \leq C\tau^2 \quad (40)$$

and

$$|\phi'_i(R) - \omega^*| \leq C|\tau|. \quad (41)$$

We will choose $\tau > 0$ small enough so that

$$\tau < \frac{|\omega^*|(1 - |\omega^*|)}{100C}. \quad (42)$$

Recall that $|\omega^*| < 1$ and hence (41) implies that for every $\tau > 0$ small enough we have that for $R_i \in [x^* - \tau, x^* + \tau]$ we have that the map ϕ_i on the interval

$$I' = \left[x^* - 2\tau \frac{|\omega^*|}{1 - |\omega^*|}, x^* + 2\tau \frac{|\omega^*|}{1 - |\omega^*|} \right] \quad (8)$$

is uniformly contracting, using also Lemma 17. The bound on $\omega(R)$ follows from $\omega(x^*) < 1$ (see Lemma 17). This proves the first two inequalities in the lemma.

For the third inequality, consider $R_i \in [x^* - \tau, x^* + \tau]$ and let $R_i = x^* + y_i$, so that $|y_i| \leq \tau$. Let also $R = x^* + y$; then, since $R \in I'$, we have $|y| \leq 2\tau \frac{|\omega^*|}{1 - |\omega^*|}$.

From (40) we have

$$|\phi_i(R) - x^*| \leq C\tau^2 + \left(\tau + 2\tau \frac{|\omega^*|}{1 - |\omega^*|} \right) |\omega^*|.$$

To prove the lemma it is enough to show

$$C\tau^2 + \left(\tau + 2\tau \frac{|\omega^*|}{1 - |\omega^*|} \right) |\omega^*| \leq 2\tau \frac{|\omega^*|}{1 - |\omega^*|}. \quad (43)$$

Equation (43) is equivalent to $C\tau^2 \leq 2\tau|\omega^*|$, which follows from our choice of τ in (42). $\square$

Proof of Lemma 20. Let $x = (x_1 + x_2)/(2\omega)$. Note that $x \in [x^* - \tau, x^* + \tau]$. Let R_i be the effective field of the field gadget guaranteed by Lemma 18, that is, $|R_i - x| \leq \tau\delta$. We will show

$$\phi_i(x^* + |\omega|\tau/2) < x_1 \quad \text{and} \quad x_2 < \phi_i(x^* - |\omega|\tau/2). \quad (44)$$

Using (40) we have

$$\phi_i(x^* + |\omega|\tau/2) - x_1 \leq x^* + \omega((R_i - x^*) + |\omega|\tau/2) - x_1 + C|\tau|^2 \leq \quad (45)$$

$$x^* + |\omega|\tau\delta + \omega((x - x^*) + |\omega|\tau/2) - x_1 + C|\tau|^2 \leq 2|\omega|\tau\delta + \omega|\omega|\tau/2 + C|\tau|^2 < 0, \quad (46)$$

the last inequality is true for τ satisfying (42) and $\delta < |\omega|/100$ (the negative middle term “overwhelms” the remaining positive terms). The second inequality in (44) follows the same argument yielding

$$\phi_i(x^* - |\omega|\tau/2) - x_2 \geq -2|\omega|\tau\delta - \omega|\omega|\tau/2 - C|\tau|^2 > 0. \quad \square$$

4.4 Proof of Lemma 24

In this section, we consider Case 2 and prove Lemma 24. We first argue that the maps $\mathcal{F}_{x,t}$ are getting “flat” as t goes to infinity.

Lemma 32. *There exist constants $C < 1, D < 1$ and $C' > 0, D' > 0$ such that for any $x \in I$ any $t \geq 0$ and any $(f, g) \in \mathcal{F}_{x,t}$ we have*

$$\max_{R, R' \in I'} |f(R) - f(R')| \leq C' C^t. \quad (47)$$

and

$$\max_{R, R' \in I', M, M' \in J} |g(R, M) - g(R', M')| \leq D' D^t. \quad (48)$$

Proof. We will prove the lemma by induction on t . Consider $(f, g) \in \mathcal{F}_{x,t}$. Let i be such that $x \in \phi_i^{-1}(I)$ and $(\hat{f}, \hat{g}) \in \mathcal{F}_{\phi_i^{-1}(x), t-1}$ be such that

$$f(R) = \phi_i(\hat{f}(R)) \quad \text{and} \quad g(R, M) = \psi_i(\hat{f}(R), \hat{g}(R, M)).$$

Note that $\hat{f}(R) \in I'$ and $\hat{f}(R')$ (see Lemma 19). Recall that ϕ_i is contracting on I' (see Lemma 19), hence

$$|f(R) - f(R')| = |\phi_i(\hat{f}(R)) - \phi_i(\hat{f}(R'))| \leq C|\hat{f}(R) - \hat{f}(R')| \leq C' C^t. \quad (49)$$

We have

$$g(R, M) = 1 - \omega(f(R))(\hat{g}(f(R), M) + M_i - 1) \quad \text{and} \quad g(R', M') = 1 - \omega(f(R'))(\hat{g}(f(R'), M') + M_i - 1).$$

Hence

$$\begin{aligned} g(R, M) - g(R', M') &= (\omega(f(R')) - \omega(f(R)))(M_i - 1 + \hat{g}(f(R'), M')) \\ &\quad + \omega(f(R))(\hat{g}(f(R'), M') - \hat{g}(f(R), M)). \end{aligned} \quad (50)$$

We will use (49) to bound the first term in the final expression of (50); the second term will be bounded by induction. Note that the function $\omega(x)$ on I' has a bounded derivative. Further note that M_i and $\hat{g}(f(R'), M')$ are in J' . Hence there exists a constant K such that

$$\|(\omega(f(R')) - \omega(f(R)))(M_i - 1 + \hat{g}(f(R'), M'))\| \leq K|f(R') - f(R)| \leq K C' C^t. \quad (51)$$

We can now take $D > \max\{C, C_{\max}\}$ and $D' = K C' D / (D - C_{\max})$. Then we have

$$|g(R, M) - g(R', M')| \leq K C' C^t + Q D' D^{t-1} \leq D' D^t. \quad \square$$

We can now prove Lemma 21.

Proof of Lemma 21. The lemma follows from (47) and the fact that $x \in f(I)$. $\square$

Proof of Lemma 24. Let L_x be the intersection of $g(I' \times J)$ for all $(f, g) \in \mathcal{F}_{x,t}$ for all $t \geq 0$. By Helly's theorem L_x is non-empty. By Lemma 32 we have that L consists of a single point ℓ . Let $F(x) = \ell$. Lemma 32 also implies (12) for every $R \in I'$ and every $M \in J$ and t large enough. We have the following observation about the function F . If $\phi_i(y) = x$ and $x, y \in I$ then

$$F(x) = 1 - \omega(x)(F(y) + M_i - 1). \quad (52)$$

To prove (52) note that if we take $(f, g) \in \mathcal{F}_{y,t}$ then we have $(f_1, g_1) \in \mathcal{F}_{x,t+1}$ where

$$g_1(R, M) = \psi_i(f(R), g(R, M)) = 1 - \omega(\phi_i(f(R)))(g(R, M) + M_i - 1).$$

Note that as t goes to infinity we have that $f(R)$ goes to y , $\phi_i(f(R))$ goes to x , $g(R, M)$ goes to $F(y)$ and $g_1(R, M)$ goes to $F(x)$ (for any $R \in I'$ and $M \in J$) and hence (52) holds.

Now we argue continuity of F . We will show that for $x_1, x_2 \in I$ we have for some $c \in (0, 1)$

$$|F(x_1) - F(x_2)| \leq C|x_1 - x_2|^c. \quad (53)$$

We will prove (53) by induction on $\lfloor \log_{C_{\max}} |x_1 - x_2| \rfloor$, where $C_{\max}$ is the constant from Lemma 19. The base case is when $|x_1 - x_2| \geq |\omega|\tau\delta/2$. Since $F(x_1), F(x_2) \in J$ we have

$$|F(x_1) - F(x_2)| \leq 2T \leq C|\omega|\tau\delta/2 \leq C(|\omega|\tau\delta/2)^c \leq C|x_1 - x_2|^c,$$

where we choose $C \geq 4T/(|\omega|\tau\delta)$.

Now we proceed to the induction step. Let x_1, x_2 be such that $|x_1 - x_2| \leq |\omega|\tau\delta/2$. By Lemma 20 there exists ϕ_i such that both x_1, x_2 are in $\phi_i(I)$. Let $y_1 = \phi_i^{-1}(x_1)$ and $y_2 = \phi_i^{-1}(x_2)$. By Lemma 19 we have

$$(1/C_{\max})|x_1 - x_2| \leq |y_1 - y_2| \leq (1/C_{\min})|x_1 - x_2|.$$

Note that

$$\lfloor \log_{C_{\max}} |y_1 - y_2| \rfloor < \lfloor \log_{C_{\max}} |x_1 - x_2| \rfloor.$$

From equation (52) we have

$$F(x_1) - F(x_2) = (F(y_1) + M_i - 1)(\omega(x_2) - \omega(x_1)) - \omega(x_2)(F(y_1) - F(y_2)). \quad (54)$$

Note that the function $\omega(x)$ on I' has a bounded derivative. Further note that M_i and $F(y_1)$ are in J' . Hence there exists a constant K such that

$$|(F(y_1) + M_i - 1)(\omega(x_2) - \omega(x_1))| \leq K|x_2 - x_1|. \quad (55)$$

Hence

$$|F(x_1) - F(x_2)| \leq K|x_2 - x_1| + C_{\max}C|y_2 - y_1|^c \leq \left(K + C\frac{C_{\max}}{C_{\min}^c}\right)|x_2 - x_1|^c. \quad (56)$$

We will choose $c \in (0, 1)$ such that $C_{\min}^c > C_{\max}$ (such $c \in (0, 1)$ exists since $C_{\min}, C_{\max} \in (0, 1)$). We then choose C such that $C(\frac{C_{\max}}{C_{\min}^c} - 1) + K \leq 0$. For our choice of $c \in (0, 1)$ and C we can bound (56) by $C|x_2 - x_1|^c$ completing the induction step. $\square$

4.5 Obtaining an infinite sequence of pairs of field gadgets

In this section, we prove Lemma 16.

Proof of Lemma 16. For $j = 0$, we set $\mathcal{T}_{1,0} = \mathcal{T}_1, \mathcal{T}_{2,0} = \mathcal{T}_2$. For $j \geq 1$, we proceed inductively. For $i \in \{1, 2\}$, we construct $\mathcal{T}_{i,j+1}$ by just combining $\mathcal{T}_{i,j}$ with the degenerate tree of Example 10 according to Lemma 17; we then have that

$$R_{i,j+1} = \frac{1 + \gamma\lambda R_{i,j}}{\beta + \lambda R_{i,j}}, \quad M_{i,j+1} = 1 - \omega(R_{i,j})M_{i,j}.$$

Since $R_{1,j} = R_{2,j}$ and $M_{1,j} \neq M_{2,j}$, we therefore have that $R_{1,j+1} = R_{2,j+1}$ and $M_{1,j+1} \neq M_{2,j+1}$ as well. The function $f(x) = \frac{1 + \gamma\lambda x}{\beta + \lambda x}$ has a unique fixpoint x^* in the interval $(\beta, 1/\gamma)$ and for any starting point $x_0 \in (\beta, 1/\gamma)$ with $x_0 \neq x^*$, repeated application of f produces x_i 's that are getting strictly closer to x^* . It follows that the sequence $\{R_{1,j}\}$ constructed above, and hence $\{R_{2,j}\}$ as well, has pairwise distinct elements except for the case that $R_1 = R_2 = x^*$. In this case, we let $\mathcal{T}_{i,0}$ to be the field gadget obtained by combining two copies of $\mathcal{T}_i$ according to Lemma 17 and proceed as before. Note that whenever $x^* \neq 1$, we also have that $R_{i,0} \neq x^*$ as needed since

$$R_{i,0} = \frac{1 + \gamma\lambda(x^*)^2}{\beta + \lambda(x^*)^2} \neq f(x^*) = x^*,$$

where the disequality holds because $x^* \neq 1$. We also have that $R_{1,0} = R_{2,0}$ and $M_{1,0} \neq M_{2,0}$ since, for $i \in \{1, 2\}$, $M_{i,0} = 1 - \omega(R_{i,0})(2M_i - 1)$ and $M_1 \neq M_2$. Finally, in the case that $x^* = 1$, which corresponds to $\lambda = \frac{1-\beta}{1-\gamma}$ for $\beta, \gamma \in (0, 1)$ and $\beta \neq \gamma$, we let $\mathcal{T}_{i,0}$ to be the field gadget obtained by combining, according to Lemma 17, $\mathcal{T}_i$ and the field gadget $\mathcal{T}$ of Lemma 30. Then, we have that $R_{i,0} \neq x^*$ as needed since

$$R_{i,0} = \frac{1 + \gamma\lambda R}{\beta + \lambda R} \neq 1,$$

since the effective field R of $\mathcal{T}$ satisfies $R \neq 1$ by Lemma 30. We also have that $R_{1,0} = R_{2,0}$ and $M_{1,0} \neq M_{2,0}$ since, for $i \in \{1, 2\}$, $M_{i,0} = 1 - \omega(R_{i,0})(M_i + M - 1)$ and $M_1 \neq M_2$, where M is the magnetization gap of $\mathcal{T}$.

This finishes the proof of Lemma 16. $\square$

5 Proofs of the remaining inapproximability results

With all the field-gadget constructions in place, we now give/finish the proofs of the remaining inapproximability results.

5.1 The bipartite “phase-gadget” of Sly and Sun

Here, we state more precisely the properties of the bipartite phase-gadgets that we described in Section 2.3. Recall that a graph $G \in \mathcal{G}_n^\ell$ is an almost- Δ -regular bipartite graph with n vertices on each side and ℓ “ports” of degree $\Delta - 1$, which will be used to connect distinct copies of gadgets. Recall also that for $\sigma : U^+ \cup U^- \rightarrow \{0, 1\}$, we define the *phase* $\mathcal{Y}(\sigma)$ of the configuration σ as the side which has the most occupied vertices under σ , i.e.,

$$\mathcal{Y}(\sigma) = \begin{cases} + & \text{if } |\sigma_{U^+}| \geq |\sigma_{U^-}|, \\ - & \text{otherwise.} \end{cases}$$

Lemma 33 ([16, Section 4.1]). *Let $\Delta \geq 3$ and $(\beta, \gamma, \lambda) \in \mathcal{U}_\Delta$. There exist real numbers $q^+, q^- \in (0, 1)$ with $q^- < q^+$ such that the following hold for every integer $\ell \geq 0$, real $\epsilon > 0$, and all sufficiently large integers n . There exists a graph $G \in \mathcal{G}_n^\ell$ such that the Gibbs distribution $\mu = \mu_{G;\beta,\gamma,\lambda}$ satisfies the following:*

1. $\frac{1-\epsilon}{2} \leq \mu(\mathcal{Y}(\sigma) = \pm) \leq \frac{1+\epsilon}{2}$.
2. for every $\tau : W \rightarrow \{0, 1\}$, we have that $(1-\epsilon)Q_W^\pm(\tau) \leq \mu(\sigma_W = \tau \mid \mathcal{Y}(\sigma) = \pm) \leq (1+\epsilon)Q_W^\pm(\tau)$, where $Q_W^+(\cdot), Q_W^-(\cdot)$ are product distribution defined by

$$Q_W^\pm(\tau) = (q^\pm)^{|\tau^{-1}(1) \cap W^+|} (1 - q^\pm)^{|\tau^{-1}(0) \cap W^+|} (q^\mp)^{|\tau^{-1}(1) \cap W^-|} (1 - q^\mp)^{|\tau^{-1}(0) \cap W^-|}. \quad (57)$$

Remark 34. We remark that for any fixed $\epsilon > 0$ and $\ell \geq 1$, we can find n and $G \in \mathcal{G}_n^\ell$ satisfying Items 1 and 2 by brute force in constant time. Namely, the values of q^+, q^- can be obtained by solving the system of equations $x = \frac{1}{\lambda} \left(\frac{\beta y + 1}{y + \gamma} \right)^{\Delta-1}$, $y = \frac{1}{\lambda} \left(\frac{\beta x + 1}{x + \gamma} \right)^{\Delta-1}$ for $x, y > 0$; it can be shown that for $(\beta, \gamma, \lambda) \in \mathcal{U}_\Delta$ this system has a unique solution with $x < y$ (see, e.g., [5, Lemma 7], or [8, Section 6.2]). Using these values, for any constants n, ℓ and any graph $G \in \mathcal{G}_n^\ell$ we can check whether Items 1 and 2 hold in constant time, and therefore find G satisfying Lemma 33.

5.2 Preliminaries

Let $G = (V, E)$ be a graph and (β, γ) be antiferromagnetic. It will sometimes be convenient to consider models where instead of a uniform field λ over vertices of G , we allow vertices to have their “own” field. More precisely, for a field vector $\boldsymbol{\lambda} = \{\lambda_v\}_{v \in V}$, we define

$$\mu_{G;\beta,\gamma,\boldsymbol{\lambda}}(\sigma) = \frac{w_{G;\beta,\gamma,\boldsymbol{\lambda}}(\sigma)}{Z_{G;\beta,\gamma,\boldsymbol{\lambda}}}, \text{ where } w_{G;\beta,\gamma,\boldsymbol{\lambda}}(\sigma) := \beta^{m_0(\sigma)} \gamma^{m_1(\sigma)} \prod_{v \in V; \sigma(v)=1} \lambda_v. \quad (58)$$

Lemma 35. *Let (β, γ) be antiferromagnetic and $\lambda, \lambda_1, \lambda_2 > 0$. Let $G = (V, E)$ be a graph and $S \subseteq V$. For $i \in \{1, 2\}$, let $\boldsymbol{\lambda}_i$ be the field vector on V , where every $v \in S$ has field λ_i , whereas every $v \in V \setminus S$ has field λ . Let μ_i be the Gibbs distributions on G , respectively, with parameters $\beta, \gamma, \boldsymbol{\lambda}_i$. Then, for every $v \in V$, it holds that*

$$|\mathbf{E}_{\sigma \sim \mu_2}[\sigma(v)] - \mathbf{E}_{\sigma \sim \mu_1}[\sigma(v)]| \leq 2|S| \left| \frac{\lambda_2}{\lambda_1} - 1 \right|.$$

Proof. Assume first that $\lambda_1 \geq \lambda_2$. For $v \in V$, let

$$Z_{G,v}^{\text{in}}(\boldsymbol{\lambda}_i) := \sum_{\sigma: V \rightarrow \{0,1\}; \sigma(v)=1} w_{G;\beta,\gamma,\boldsymbol{\lambda}_i}(\sigma), \quad Z_G(\boldsymbol{\lambda}_i) := Z_{G;\beta,\gamma,\boldsymbol{\lambda}_i}^{\text{in}},$$

so that $\mathbf{E}_{\sigma \sim \mu_i}[\sigma(v)] = \frac{Z_{G,v}^{\text{in}}(\boldsymbol{\lambda}_i)}{Z_G(\boldsymbol{\lambda}_i)}$. Therefore,

$$\mathbf{E}_{\sigma \sim \mu_1}[\sigma(v)] - \mathbf{E}_{\sigma \sim \mu_2}[\sigma(v)] = \frac{Z_{G,v}^{\text{in}}(\boldsymbol{\lambda}_1) - Z_{G,v}^{\text{in}}(\boldsymbol{\lambda}_2)}{Z_G(\boldsymbol{\lambda}_1)} + \frac{Z_{G,v}^{\text{in}}(\boldsymbol{\lambda}_2)}{Z_G(\boldsymbol{\lambda}_2)} \frac{Z_G(\boldsymbol{\lambda}_2) - Z_G(\boldsymbol{\lambda}_1)}{Z_G(\boldsymbol{\lambda}_1)} \quad (59)$$

For arbitrary $\sigma : V \rightarrow \{0, 1\}$, we have that

$$\frac{w_{G;\beta,\gamma,\boldsymbol{\lambda}_1}(\sigma) - w_{G;\beta,\gamma,\boldsymbol{\lambda}_2}(\sigma)}{w_{G;\beta,\gamma,\boldsymbol{\lambda}_1}(\sigma)} = \frac{\lambda_1^{|\sigma_S|} - \lambda_2^{|\sigma_S|}}{\lambda_1^{|\sigma_S|}} \leq |S| \left(1 - \frac{\lambda_2}{\lambda_1} \right).$$

From this, it follows that

$$\left| \frac{Z_{G,v}^{\text{in}}(\lambda_1) - Z_{G,v}^{\text{in}}(\lambda_2)}{Z_G(\lambda_1)} \right| \leq |S| \left(1 - \frac{\lambda_2}{\lambda_1} \right), \quad \left| \frac{Z_G(\lambda_2) - Z_G(\lambda_1)}{Z_G(\lambda_1)} \right| \leq |S| \left(1 - \frac{\lambda_2}{\lambda_1} \right).$$

Hence, by taking absolute values in (59) and using $\frac{Z_{G,v}^{\text{in}}(\lambda_2)}{Z_G(\lambda_2)} \leq 1$ and the triangle inequality, we obtain

$$|\mathbf{E}_{\sigma \sim \mu_1}[\sigma(v)] - \mathbf{E}_{\sigma \sim \mu_2}[\sigma(v)]| \leq 2|S| \left(1 - \frac{\lambda_2}{\lambda_1} \right),$$

finishing the proof for $\lambda_1 \geq \lambda_2$. The case $\lambda_1 < \lambda_2$ follows from the previous case since

$$|\mathbf{E}_{\sigma \sim \mu_1}[\sigma(v)] - \mathbf{E}_{\sigma \sim \mu_2}[\sigma(v)]| \leq 2|S| \left(1 - \frac{\lambda_1}{\lambda_2} \right) \leq 2|S| \left(\frac{\lambda_2}{\lambda_1} - 1 \right). \quad \square$$

5.3 Proof of Lemma 15

We start with the proof of Lemma 15.

Proof of Lemma 15. We first prove the bounds on $\mathcal{M}_{\beta, \gamma, \lambda}(H_{G, \mathcal{T}}^k)$.

For an edge $e = (u, v) \in E_H$ and $i \in [k]$, let $P_e^{i,+}, P_e^{i,-}$ be the i -th paths connecting the $+/+$ and $-/-$ sides of the gadgets G_u, G_v , respectively. Let the (two) internal vertices of the path $P_e^{i,\pm}$ be $t_{e,1}^{i,\pm}$ and $t_{e,2}^{i,\pm}$. For $j = 1, 2$ denote by $\mathcal{T}_{e,j}^{i,\pm}$ the copy of $\mathcal{T}$ that was appended to $t_{e,j}^{i,\pm}$.

Note that the vertex set of $H_{G, \mathcal{T}}^k$ is the disjoint union of $V(\hat{H}_G^k)$ and $V(\mathcal{T}_{e,j}^{i,s})$ for $e \in E(H), s \in \{+, -\}, i \in [k], j \in \{1, 2\}$, so we have

$$\mathcal{M}_{\beta, \gamma, \lambda}(H_{G, \mathcal{T}}^k) = \mathbf{E}_{\sigma \sim \mu} [|\sigma_{V(\hat{H}_G^k)}|] + \sum_{e \in E(H)} \sum_{s \in \{+, -\}} \sum_{i \in [k]} \sum_{j \in \{1, 2\}} \mathbf{E}_{\sigma \sim \mu} [|\sigma_{V(\mathcal{T}_{e,j}^{i,s})}|]. \quad (60)$$

Fix arbitrary $e \in E(H), s \in \{+, -\}, i \in [k], j \in \{1, 2\}$, and for convenience let $\mu_{e,j}^{i,s}$ denote the Gibbs distribution on $\mathcal{T}_{e,j}^{i,s}$ with parameters β, γ, λ . Note that, conditioned on the spin of $t_{e,j}^{i,s}$, the distribution μ factorizes; more precisely, for $\tau : V(\mathcal{T}_{e,j}^{i,s}) \rightarrow \{0, 1\}$ and $w \in \{0, 1\}$, we have

$$\mu \left(\sigma_{V(\mathcal{T}_{e,j}^{i,s})} = \tau \mid \sigma(t_{e,j}^{i,s}) = w \right) = \mu_{e,j}^{i,s}(\tau \mid \tau(t_{e,j}^{i,s}) = w).$$

Using that the magnetization gap of $\mathcal{T}_{e,j}^{i,s}$ is M , we therefore obtain that

$$\begin{aligned} \mathbf{E}_{\sigma \sim \mu} [|\sigma_{V(\mathcal{T}_{e,j}^{i,s})}|] &= \sum_{w \in \{0, 1\}} \mu(\sigma(t_{e,j}^{i,s}) = w) \mathbf{E}_{\tau \sim \mu_{e,j}^{i,s}} [|\tau| \mid \tau(t_{e,j}^{i,s}) = w] \\ &= M \mathbf{E}_{\sigma \sim \mu} [\sigma(t_{e,j}^{i,s})] + \mathcal{A}', \end{aligned}$$

where the last equality follows by writing $\mu(\sigma(t_{e,j}^{i,s}) = 0) = 1 - \mu(\sigma(t_{e,j}^{i,s}) = 1) = 1 - \mathbf{E}_{\sigma \sim \mu} [\sigma(t_{e,j}^{i,s})]$ and noting that $\mathcal{A}' = \mathbf{E}_{\tau \sim \mu_{e,j}^{i,s}} [|\tau| \mid \tau(t_{e,j}^{i,s}) = 0]$. Plugging into (60), we obtain

$$\mathcal{M}_{\beta, \gamma, \lambda}(H_{G, \mathcal{T}}^k) = 4k\mathcal{A}'|E(H)| + \mathbf{E}_{\sigma \sim \mu} [|\sigma_{V(\hat{H}_G^k)}|] + M \sum_{e \in E(H)} \sum_{s \in \{+, -\}} \sum_{i \in [k]} M_e^{i,s}, \quad (61)$$

where, for $e \in E(H), i \in [k]$ and $s \in \{+, -\}$,

$$M_e^{i,s} := \mathbf{E}_{\sigma \sim \mu} [\sigma(t_{e,1}^{i,s}) + \sigma(t_{e,2}^{i,s})].$$

We can expand $M_e^{i,s}$ according to the configuration on the ports $W_e^{i,s} := \{w_u^{i,s}, w_v^{i,s}\}$, i.e.,

$$M_e^{i,s} = \sum_{\tau: W_e^{i,s} \rightarrow \{0,1\}} \mu(\sigma_{W_e^{i,s}} = \tau) \mathbf{E}_{\sigma \sim \mu}[\sigma(t_{e,1}^{i,s}) + \sigma(t_{e,2}^{i,s}) \mid \sigma_{W_e^{i,s}} = \tau]. \quad (62)$$

For convenience, let $\overline{W}_e^{i,s}$ be the set $(W_u \cup W_v) \setminus W_e^{i,s}$, i.e., the set of all ports in the gadgets G_u, G_v other than those in $W_e^{i,s}$. We have

$$\mu(\sigma_{W_e^{i,s}} = \tau) = \sum_{\substack{Y: V(H) \rightarrow \{+, -\}; \\ \eta: \overline{W}_e^{i,s} \rightarrow \{0,1\}}} \mu(\sigma_{W_e^{i,s}} = \tau \mid \widehat{\mathcal{Y}}(\sigma) = Y, \sigma_{\overline{W}_e^{i,s}} = \eta) \mu(\widehat{\mathcal{Y}}(\sigma) = Y, \sigma_{\overline{W}_e^{i,s}} = \eta). \quad (63)$$

Now, observe that $\overline{W}_e^{i,s}$ disconnects the vertices in $U_u \cup U_v$ from the rest of the graph $H_{G,\mathcal{T}}^k$, so, conditioned on $\eta: \overline{W}_e^{i,s} \rightarrow \{0,1\}$, the configuration on $W_e^{i,s}$ is distributed according to μ_e , where μ_e is the Gibbs distribution with parameters (β, γ, λ) on the subgraph of $H_{G,\mathcal{T}}^k$ induced by $U_u \cup U_v$. Therefore, for any $Y: V(H) \rightarrow \{+, -\}$ and $\eta: \overline{W}_e^{i,s} \rightarrow \{0,1\}$, we have that

$$\mu(\sigma_{W_e^{i,s}} = \tau \mid \widehat{\mathcal{Y}}(\sigma) = Y, \sigma_{\overline{W}_e^{i,s}} = \eta) = \mu_e(\sigma_{W_e^{i,s}} = \tau \mid \mathcal{Y}(\sigma_{U_u}) = Y_u, \mathcal{Y}(\sigma_{U_v}) = Y_v, \sigma_{\overline{W}_e^{i,s}} = \eta). \quad (64)$$

Since G_u, G_v are copies of G , and G satisfies Lemma 33, by Item 2 we obtain that $\mu_e(\sigma_{W_e^{i,s}} = \tau \mid \mathcal{Y}(\sigma_{U_u}) = Y_u, \mathcal{Y}(\sigma_{U_v}) = Y_v, \sigma_{\overline{W}_e^{i,s}} = \eta)$ is approximately independent of η and within a factor of $(1 \pm 8\epsilon)$ from the Gibbs distribution on $P_e^{i,\pm}$ with field vector $\boldsymbol{\lambda}_\pm(Y_u, Y_v)$ given by

$$\lambda_{w_u^{i,\pm}} = \frac{q^{\pm Y_u}}{1 - q^{\pm Y_u}}, \quad \lambda_{t_{e,1}^{i,\pm}} = \lambda_{t_{e,2}^{i,\pm}} = \lambda R, \quad \lambda_{w_v^{i,\pm}} = \frac{q^{\pm Y_v}}{1 - q^{\pm Y_v}}, \quad \text{with the notation } \pm\pm \equiv \pm \text{ and } \pm\mp \equiv \mp.$$

More precisely, we have

$$\mu_e(\sigma_{W_e^{i,s}} = \tau \mid \mathcal{Y}(\sigma_{U_u}) = Y_u, \mathcal{Y}(\sigma_{U_v}) = Y_v, \sigma_{\overline{W}_e^{i,s}} = \eta) = (1 \pm 8\epsilon) \mu_{P_e^{i,s}; \beta, \gamma, \boldsymbol{\lambda}_s(Y_u, Y_v)}(\sigma_{W_e^{i,s}} = \tau). \quad (65)$$

Combining (64) and (65), and plugging back into (63), we obtain

$$\mu(\sigma_{W_e^{i,s}} = \tau) = (1 \pm 8\epsilon) \sum_{Y: V(H) \rightarrow \{+, -\}} \mu(\widehat{\mathcal{Y}}(\sigma) = Y) \mu_{P_e^{i,s}; \beta, \gamma, \boldsymbol{\lambda}_s(Y_u, Y_v)}(\sigma_{W_e^{i,s}} = \tau)$$

In turn, plugging this back into (62) yields that

$$M_e^{i,s} = (1 \pm 8\epsilon) \sum_{Y: V(H) \rightarrow \{+, -\}} \mu(\widehat{\mathcal{Y}}(\sigma) = Y) \mathbf{E}_{\sigma \sim \mu_{P_e^{i,s}; \beta, \gamma, \boldsymbol{\lambda}_s(Y_u, Y_v)}}[\sigma(t_{e,1}^{i,s}) + \sigma(t_{e,2}^{i,s})]. \quad (66)$$

For $s_1, s_2, s_3 \in \{+, -\}$, let

$$A_{s_2 s_3}^{s_1} := \mathbf{E}_{\sigma \sim \mu_{P_e^{i,s_1}; \beta, \gamma, \boldsymbol{\lambda}_s(s_2, s_3)}}[\sigma(t_{e,1}^{i,s_1}) + \sigma(t_{e,2}^{i,s_1})].$$

To give expressions for $A_{s_2 s_3}^{s_1}$, it will be convenient to consider the following vectors/matrices:

$$\mathbf{q}^\pm := \begin{bmatrix} 1 - q^\pm \\ q^\pm \end{bmatrix}, \quad \mathbf{L}(x) := \begin{bmatrix} \beta^3 + 2\beta x + \gamma x^2 & \beta^2 + x(1 + \beta\gamma) + \gamma^2 x^2 \\ \beta^2 + x(1 + \beta\gamma) + \gamma^2 x^2 & \beta + 2\gamma x + \gamma^3 x^2 \end{bmatrix}, \quad \begin{aligned} f_{\pm\pm}(x) &:= (\mathbf{q}^\pm)^\top \mathbf{L}(x) \mathbf{q}^\pm \\ f_{\pm\mp}(x) &:= (\mathbf{q}^\pm)^\top \mathbf{L}(x) \mathbf{q}^\mp \end{aligned} \quad (67)$$

where q^+, q^- are the constants in Lemma 33, and recall that $q^+ \neq q^-$. Roughly, the matrix $\mathbf{L}(x)$ corresponds to the interaction matrix between the endpoints of a path of length 4, whose vertices

have vertex fields $1, x, x, 1$ respectively. Then, the functions $f_{\pm\pm}(x)$ capture the total weight of configurations on the path, when the endpoints are occupied with probabilities $q^\pm, q^\pm$, respectively; similarly for the functions $f_{\pm\mp}(x)$. The magnetization of the two middle vertices is given by the log-derivative with respect to x and multiplying by x (this formula can be verified using the definition of magnetization), yielding the following expressions for the A 's:

$$\begin{aligned} A_{++}^+ &= A_{--}^- = \lambda R \cdot \frac{f'_{++}(\lambda R)}{f_{++}(\lambda R)}, & A_{+-}^- &= A_{-+}^+ = \lambda R \cdot \frac{f'_{--}(\lambda R)}{f_{--}(\lambda R)}, \\ A_{+-}^+ &= A_{-+}^- = A_{-+}^- = A_{+-}^+ = \lambda R \cdot \frac{f'_{+-}(\lambda R)}{f_{+-}(\lambda R)} = \lambda R \cdot \frac{f'_{-+}(\lambda R)}{f_{-+}(\lambda R)}. \end{aligned}$$

Now, let's fix a phase vector $Y : V(H) \rightarrow \{+, -\}$, and compute

$$\sum_{e=(u,v) \in E(H)} \sum_{s \in \{+, -\}} \sum_{i \in [k]} \mathbf{E}_{\sigma \sim \mu_{P_e^{i,s}; \beta, \gamma, \lambda s(Y_u, Y_v)}} [\sigma(t_{e,1}^{i,s}) + \sigma(t_{e,2}^{i,s})]. \quad (68)$$

where the functions $f_{\pm\pm}, f_{\pm\mp}$ are as in (67). We have that the contribution from an edge $e = (u, v)$ with $Y_u \neq Y_v$ is kB , while the contribution from an edge $e = (u, v)$ with $Y_u = Y_v$ is kC , where

$$\begin{aligned} B &:= A_{+-}^+ + A_{-+}^- = A_{-+}^+ + A_{+-}^- = \lambda R \left(\frac{f'_{+-}(\lambda R)}{f_{+-}(\lambda R)} + \frac{f'_{-+}(\lambda R)}{f_{-+}(\lambda R)} \right), \\ C &:= A_{++}^+ + A_{--}^- = A_{--}^+ + A_{++}^- = \lambda R \left(\frac{f'_{++}(\lambda R)}{f_{++}(\lambda R)} + \frac{f'_{--}(\lambda R)}{f_{--}(\lambda R)} \right). \end{aligned}$$

It follows that the value of the sum in (68) is

$$kB \text{Cut}_H(Y) + kC(|E(H)| - \text{Cut}_H(Y)) = k(B - C)\text{Cut}_H(Y) + kC|E(H)|. \quad (69)$$

Combining (66) with (61) and using the value of the sum in (68) as obtained in (69), finishes the proof of the bounds on $\mathcal{M}_{\beta, \gamma, \lambda}(H_{G, \mathcal{T}}^k)$.

We next prove the bounds on $\text{AVG-CUT}_\mu(H)$. The upper bound is trivial, so we focus on the lower bound, which is along the same lines as in [16, Lemma 4.2]. Let $\hat{\mu}$ be the Gibbs distribution on $\hat{H}_G^k$ and note that this is a product distribution over the gadgets. For convenience, let $\mathbf{L} = \{L_{ij}\}_{i,j \in \{0,1\}}$ denote the (entries of the) matrix $\mathbf{L}(\lambda R)$. For phase configurations $Y : V(H) \rightarrow \{+, -\}$, we have that

$$\mu(\hat{\mathcal{Y}}(\sigma) = Y) \propto \hat{\mu}(\hat{\mathcal{Y}}(\sigma) = Y) \sum_{\tau: W \rightarrow \{0,1\}} \hat{\mu}(\sigma_W = \tau \mid \hat{\mathcal{Y}}(\sigma) = Y) \prod_{e=(u,v) \in E(H)} \prod_{s \in \{+, -\}} \prod_{i \in [k]} L_{\tau(w_u^{i,s}), \tau(w_v^{i,s})}$$

By Items 1 and 2 of Lemma 33, for all $Y : V(H) \rightarrow \{+, -\}$ and $\tau : W \rightarrow \{0, 1\}$, it holds that

$$\hat{\mu}(\hat{\mathcal{Y}}(\sigma) = Y) = \left(\frac{1 \pm \epsilon}{2} \right)^{|V(H)|}, \quad \hat{\mu}(\sigma_W = \tau \mid \hat{\mathcal{Y}}(\sigma) = Y) = (1 \pm \epsilon)^{|V(H)|} \prod_{v \in V} Q_{W_v}^{Y_v}(\tau_{W_v}),$$

where, for $v \in V(H)$, $Q_{W_v}^\pm$ is the product distribution defined in (57) for the gadget G_v . Using that

$$\mu(\hat{\mathcal{Y}}(\sigma) = Y) \propto \left(\frac{1 \pm \epsilon}{2} \right)^{|V(H)|} \prod_{e=(u,v) \in E(H)} \prod_{s \in \{+, -\}} \prod_{i \in [k]} \sum_{\tau: W_e^{i,s} \rightarrow \{0,1\}} Q^{Y_u}(\tau(w_u^{i,s})) L_{\tau(w_u^{i,s}), \tau(w_v^{i,s})} Q_{W_v}^{Y_v}(\tau(w_v^{i,s})).$$

Observe now that $\sum_{\tau: W_e^{i,s} \rightarrow \{0,1\}} Q_{W_v}^{Y_u}(\tau(w_u^{i,s})) L_{\tau(w_u^{i,s}), \tau(w_v^{i,s})} Q^{Y_v}(\tau(w_v^{i,s})) = (\mathbf{q}^{Y_u})^\top \mathbf{L} \mathbf{q}^{Y_v}$, so

$$\mu(\hat{\mathcal{Y}}(\sigma) = Y) \propto \left(\frac{1 \pm \epsilon}{2} \right)^{|V(H)|} \prod_{(u,v) \in E(H); Y_u \neq Y_v} (f_{+-}(\lambda R) f_{-+}(\lambda R))^k \prod_{(u,v) \in E(H); Y_u = Y_v} (f_{++}(\lambda R) f_{--}(\lambda R))^k.$$

It follows that for

$$A := A(R) = \frac{f_{+-}(\lambda R) \cdot f_{-+}(\lambda R)}{f_{--}(\lambda R) \cdot f_{++}(\lambda R)},$$

we have that

$$\mu(\hat{\mathcal{Y}}(\sigma) = Y) \propto \left(\frac{1}{2} \pm \epsilon\right)^{|V(H)|} A^{k \text{Cut}_H(Y)}.$$

Now note that for $R > 0$, we have $A > 1$ since $q^+ \neq q^-$ and for all x we have that

$$f_{+-}(x) \cdot f_{-+}(x) - f_{++}(x) \cdot f_{--}(x) = (1 - \beta\gamma)^3 x^2 (q^+ - q^-)^2.$$

Note also that $A(0) = 1$, while $B(R) - C(R) = (\log A(R))'$ thus proving that $A(R), B(R), C(R)$ are rational functions of R satisfying (1).

Let $\hat{K} > 1$ be such that $1 - \frac{1}{\hat{K}} \leq \frac{4}{k \log A}$. From the estimates above, we have

$$\mu\left(\frac{\text{Cut}_H(\hat{\mathcal{Y}}(\sigma))}{\text{MAX-CUT}(H)} \leq \frac{1}{\hat{K}}\right) \leq \frac{\left(\frac{1}{2} + \epsilon\right)^{|V(H)|} 2^{|V(H)|} A^{\frac{k}{\hat{K}} \text{MAX-CUT}(H)}}{\left(\frac{1}{2} - \epsilon\right)^{|V(H)|} A^{k \text{MAX-CUT}(H)}} \leq \left(\frac{1}{2}\right)^{|V(H)|}, \quad (70)$$

where in the last inequality we used that

$$A, \hat{K} > 1, \quad \text{MAX-CUT}(H) \geq \frac{1}{2}|E(H)| = \frac{3}{4}|V(H)|, \quad \text{and } 2\left(\frac{1}{2} + \epsilon\right)/\left(\frac{1}{2} - \epsilon\right) \leq 4.$$

Since $\text{AVG-CUT}_\mu(H) := \sum_{Y: V(H) \rightarrow \{+, -\}} \mu(\hat{\mathcal{Y}}(\sigma) = Y) |\text{Cut}_H(Y)|$, from (70), and using that $k \geq 10/\log A$, for $|V(H)| \geq 2kA$, we obtain

$$\text{AVG-CUT}_\mu(H) \geq \left(1 - (1/2)^{|V(H)|}\right) \frac{\text{MAX-CUT}(H)}{\hat{K}} \geq \frac{\text{MAX-CUT}(H)}{K},$$

where $K = 1 + \frac{6}{k \log A}$ is as in the statement of the lemma, as wanted.

This finishes the proof of Lemma 15. $\square$

5.4 Constant-factor inapproximability — Proof of Theorems 3, 5, 8

We now give the relatively straightforward proofs of Theorems 3, 5, 8.

Proof of Theorems 3, 5, 8. Analogously to the proof of Theorem 1, these theorems can be obtained as immediate corollaries of Theorem 8 using the field-gadget constructions of Theorem 14. The only extra fact that we need to note is that in the case of the hard-core model ($\beta = 1, \gamma = 0$), we have that $(\beta, \gamma, \lambda) \in \mathcal{U}_\Delta^*$ iff $\lambda > \lambda_c(\Delta)$, for all $\Delta \geq 3$. Similarly in the case of the antiferromagnetic Ising model ($0 < \beta = \gamma < 1$), we have that $(\beta, \gamma, \lambda) \in \mathcal{U}_\Delta^*$ iff $\lambda \in (\frac{1}{\lambda_c}, \lambda_c)$, where $\lambda_c = \lambda_c^{\text{Ising}}(\Delta, \beta)$, for all $\Delta \geq 3$. $\square$

5.5 Proof of Theorems 2, 4, and 7

We next give the proof of Theorem 7, using Theorem 28.

Proof of Theorem 7. Suppose for the sake of contradiction that, for arbitrarily small $\kappa > 0$, there is a polynomial-time algorithm that, on input a graph G of maximum degree Δ , produces a $(1 + \frac{\kappa}{\log |V(G)|})$ -approximation of the magnetization $\mathcal{M}_{G; \beta, \gamma, \lambda}$. We will show that we can approximate MAX-CUT on 3-regular graphs within a constant factor arbitrarily close to 1, contradicting the inapproximability result of [1].

Let $A(R), B(R), C(R)$ be the rational functions in Lemma 15 and set $D(R) = B(R) - C(R)$. From Footnote 2, we have that for all but finitely many values of R we have that $D(R) \neq 0$. Using Theorem 28, we therefore conclude that there are constants $R^*, M, \Xi, L > 0$ such that $D(R^*) \neq 0$ and an algorithm, which, on input a rational $r \in (0, 1/2)$, outputs in time $\text{poly}(\text{bits}(r))$ a pair of field gadgets $\mathcal{T}_1, \mathcal{T}_2$, each of maximum degree 3 and size $\leq L|\log r|$, such that

$$|R_1 - R^*|, |R_2 - R^*| \leq r, \quad |M_1|, |M_2| \leq \Xi, \quad |M_1 - M_2| \geq \hat{M}. \quad (71)$$

where, for $i \in \{1, 2\}$, $R_i := R_{\mathcal{T}_i}$ is the effective field of $\mathcal{T}_i$ and $M_i := M_{\mathcal{T}_i}$ is the magnetization gap of $\mathcal{T}_i$. For convenience, for $i \in \{1, 2\}$ we let $\mathcal{A}'_i, A_i, B_i, C_i, D_i$ denote $\mathcal{A}'_{\mathcal{T}_i}, A(R_i), B(R_i), C(R_i), D(R_i)$, respectively. Since the functions A, B, C, D are continuously differentiable with respect to R , $D(R^*) \neq 0$ and $A(R^*) > 1$, there are small constants $r_0, \zeta > 0$ and large constants $\eta, D^* > 1$ so that, for all $r < r_0$, (71) further implies that

$$A_1, A_2 \geq 1 + \zeta, \quad D_1, D_2 \neq 0 \text{ and } |D_1|, |D_2| \leq D^*, \quad \left| \frac{R_1}{R_2} - 1 \right| \leq \eta r, \quad |C_1 - C_2| \leq \eta r, \quad |D_1 - D_2| \leq \eta r, \quad \left| \frac{1}{\log A_1} - \frac{1}{\log A_2} \right| \leq \eta r. \quad (72)$$

Let k be a large integer satisfying $k > 1 + 10/\log(1 + \zeta)$ and $\epsilon > 0$ be a small constant (both to be specified later). Let $G \in \mathcal{G}_n^{3k}$ be a graph satisfying Lemma 33.

Let H be a 3-regular graph, an instance of the MAX-CUT problem. We will use the algorithm of Theorem 28 for $r = \min\{\frac{\epsilon}{50\eta k \Xi L D^* |V(G)| |V(H)|^3}, r_0\}$, so we obtain field gadgets $\mathcal{T}_1, \mathcal{T}_2$ that satisfy (71) and therefore (72) as well. For $i \in \{1, 2\}$, consider the graph $H_{G, \mathcal{T}_i}^k$ as in Lemma 15. For later use, note that the number of vertices in $H_{G, \mathcal{T}_i}^k$ is at most

$$|V(H)| |V(G)| + 2k|E(H)| |V(\mathcal{T}_i)| \leq 10kL|E(H)| \log |V(H)|,$$

using that $|V(\mathcal{T}_i)| \leq 4L \log |V(H)|$ and $|V(G)| \leq \log |V(H)|$ for all sufficiently large $|V(H)|$.

Let μ_i denote the Gibbs distribution on $H_{G, \mathcal{T}_i}^k$ with parameters β, γ, λ . By Lemma 15, the average magnetization of the graph $H_{G, \mathcal{T}_i}^k$ satisfies

$$\mathcal{M}_{\beta, \gamma, \lambda}(H_{G, \mathcal{T}_i}^k) = 4k\mathcal{A}'_i|E(H)| + \mathbf{E}_{\sigma \sim \mu_i} \left[|\sigma_{V(\hat{H}_G^k)}| \right] + (1 \pm 8\epsilon)kM_iQ_i. \quad (73)$$

where $Q_i = D_i \text{AVG-CUT}_{\mu_i}(H) + C_i|E(H)|$ and $\text{AVG-CUT}_{\mu_i}(H)$ satisfies

$$1/K_i \leq \frac{\text{AVG-CUT}_{\mu_i}(H)}{\text{MAX-CUT}(H)} \leq 1 \text{ for } K_i := 1 + \frac{6}{k \log A_i}. \quad (74)$$

Let

$$\mathcal{D} := \mathcal{M}_{\beta, \gamma, \lambda}(H_{G, \mathcal{T}_1}^k) - \mathcal{M}_{\beta, \gamma, \lambda}(H_{G, \mathcal{T}_2}^k), \quad \mathcal{D}' := \mathbf{E}_{\sigma \sim \mu_1} \left[|\sigma_{V(\hat{H}_G^k)}| \right] - \mathbf{E}_{\sigma \sim \mu_2} \left[|\sigma_{V(\hat{H}_G^k)}| \right],$$

so that from (73)

$$\mathcal{D} = 4k(\mathcal{A}'_1 - \mathcal{A}'_2)|E(H)| + \mathcal{D}' + (1 \pm 8\epsilon)k((M_1 - M_2)Q_1 + M_2(Q_1 - Q_2)). \quad (75)$$

We next show that for all sufficiently large $|V(H)|$ it holds that

$$\begin{aligned} \mathcal{D}' &\leq |V(H)| |V(G)| \cdot 2k|E(H)| (|V(\mathcal{T}_1)| + |V(\mathcal{T}_2)|) \eta r \leq \epsilon, \\ |M_2(Q_1 - Q_2)| &\leq \Xi(D^* + 2)|E(H)| \eta r \leq \epsilon/(10k). \end{aligned} \quad (76)$$

The first inequality follows from applying Lemma 35 to the graph H_G^k (without the field gadgets). Namely, for $i \in \{1, 2\}$, denote by ν_i the Gibbs distribution on H_G^k where every vertex in $\hat{H}_G^k$ has field λ whereas every vertex in $V(H_G^k) \setminus V(\hat{H}_G^k)$ has field λR_i . Then, we have that

$$\mathbf{E}_{\sigma \sim \mu_i} [|\sigma_{V(\hat{H}_G^k)}|] = \mathbf{E}_{\sigma \sim \nu_i} [|\sigma_{V(\hat{H}_G^k)}|] \quad (77)$$

and hence the inequality follows by applying Lemma 35 to each of the vertices in $V(\hat{H}_G^k)$ and using the bound on $|\frac{R_1}{R_2} - 1| \leq \eta r$ from (72). The second inequality in (76) follows from the fact that $|M_2|$ is bounded by Ξ , the bounds on $|C_1 - C_2|, |D_1 - D_2|, |D_1|, |D_2|$ from (72) and the bound $\text{AVG-CUT}_{\mu_i}(H) \leq \text{MAX-CUT}(H)$ from (74).

Using the purported algorithm for the magnetizations on $H_{G, \mathcal{T}_i}^k$, and since the latter graph has at most $10kL|E(H)| \log |V(H)|$ vertices for all sufficiently large $|V(H)|$, we can compute an estimate of $\mathcal{M}_{\beta, \gamma, \lambda}(H_{G, \mathcal{T}_i}^k)$ that is off by at most an additive

$$\frac{\kappa}{\log |V(H_{G, \mathcal{T}_i}^k)|} 10kL|E(H)| \log |V(H)| \leq 10kL|E(H)|\kappa.$$

By subtracting these estimates for $i \in \{1, 2\}$, we therefore compute $\hat{\mathcal{D}}$ which satisfies

$$|\hat{\mathcal{D}} - \mathcal{D}| \leq 20kL|E(H)|\kappa. \quad (78)$$

Then, since the magnetization gaps M_1, M_2 satisfy $|M_1 - M_2| \geq \hat{M}$ and $D_1 \neq 0$ (see (71) and (72)), we can compute $\widehat{\text{MC}} = \frac{\hat{\mathcal{D}} - 4k(\mathcal{A}'_1 - \mathcal{A}'_2)|E(H)|}{k(M_1 - M_2)D_1} - \frac{C_1}{D_1}|E(H)|$ which using (75), (76) and (78) satisfies

$$\begin{aligned} |\widehat{\text{MC}} - \text{AVG-CUT}_{\mu_1}(H)| &\leq \frac{20kL|E(H)|\kappa + 8k|M_1 - M_2||Q_1|\epsilon + 2\epsilon}{k|M_1 - M_2||D_1|} \\ &\leq \left(\frac{20L}{\hat{M}|D_1|}\kappa + 8\left(1 + \frac{|C_1|}{|D_1|}\right)\epsilon + \frac{2}{\hat{M}|D_1|}\epsilon \right) |E(H)|. \end{aligned} \quad (79)$$

By choosing k sufficiently large, we have from (74) that $\text{AVG-CUT}_{\mu_1}(H)$ is within a factor arbitrarily close to 1 from $\text{MAX-CUT}(H)$. By choosing ϵ, κ to be sufficiently small positive constants we can further ensure from (79) that our approximation $\widehat{\text{MC}}$ is within a factor arbitrarily close to 1 from $\text{AVG-CUT}_{\mu_1}(H)$, and hence from $\text{MAX-CUT}(H)$ as well. This finishes the contradiction argument and completes the proof of Theorem 7. $\square$

Proof of Theorems 2 and 4. The theorems can be obtained as immediate corollaries of Theorem 7 using the field-gadget constructions of Theorem 13. As in Section 5.4, we again note that in the case of the hard-core model ($\beta = 1, \gamma = 0$), we have that $(\beta, \gamma, \lambda) \in \mathcal{U}_\Delta^*$ iff $\lambda > \lambda_c(\Delta)$, for all $\Delta \geq 3$. Similarly in the case of the antiferromagnetic Ising model ($0 < \beta = \gamma < 1$), we have that $(\beta, \gamma, \lambda) \in \mathcal{U}_\Delta^*$ iff $\lambda \in (\frac{1}{\lambda_c}, \lambda_c)$, where $\lambda_c = \lambda_c^{\text{Ising}}(\Delta, \beta)$, for all $\Delta \geq 3$. $\square$

References

- [1] Paola Alimonti and Viggo Kann. Hardness of approximating problems on cubic graphs. In *Algorithms and Complexity*, pages 288–298, 1997.
- [2] Nima Anari, Kuikui Liu, and Shayan Oveis Gharan. Spectral independence in high-dimensional expanders and applications to the hardcore model. *arXiv preprint arXiv:2001.00303*, 2020.

- [3] Ivona Bezáková, Andreas Galanis, Leslie Ann Goldberg, and Daniel Štefankovič. Inapproximability of the independent set polynomial in the complex plane. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2018, page 1234–1240, 2018.
- [4] Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. Inapproximability for antiferromagnetic spin systems in the tree nonuniqueness region. *J. ACM*, 62(6), 2015.
- [5] Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. Inapproximability of the partition function for the antiferromagnetic Ising and hard-core models. *Combinatorics, Probability and Computing*, 25(4):500–559, 2016.
- [6] Mark Jerrum and Alistair Sinclair. Polynomial-time approximation algorithms for the Ising model. *SIAM Journal on Computing*, 22(5):1087–1116, 1993.
- [7] Liang Li, Pinyan Lu, and Yitong Yin. Correlation decay up to uniqueness in spin systems. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 67–84, 2013.
- [8] Fabio Martinelli, Alistair Sinclair, and Dror Weitz. Fast mixing for independent sets, colorings, and other models on trees. *Random Structures & Algorithms*, 31(2):134–172, 2007.
- [9] Viresh Patel and Guus Regts. Deterministic polynomial-time approximation algorithms for partition functions and graph polynomials. *SIAM Journal on Computing*, 46(6):1893–1919, 2017.
- [10] Han Peters and Guus Regts. On a conjecture of Sokal concerning roots of the independence polynomial. *Michigan Math. J.*, 68(1):33–55, 2019.
- [11] Ricardo Restrepo, Daniel Štefankovič, Juan C. Vera, Eric Vigoda, and Linji Yang. Phase transition for Glauber dynamics for independent sets on regular trees. *SIAM Journal on Discrete Mathematics*, 28(2):835–861, 2014.
- [12] Leonard J. Schulman, Alistair Sinclair, and Piyush Srivastava. Symbolic integration and the complexity of computing averages. In *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 1231–1245, 2015.
- [13] Alistair Sinclair and Piyush Srivastava. Lee-Yang theorems and the complexity of computing averages. *Communications in Mathematical Physics*, 329(3):827–858, 2014.
- [14] Alistair Sinclair, Piyush Srivastava, and Marc Thurley. Approximation algorithms for two-state anti-ferromagnetic spin systems on bounded degree graphs. *J. Stat. Phys.*, 155(4):666–686, 2014.
- [15] Allan Sly. Computational transition at the uniqueness threshold. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 287–296, 2010.
- [16] Allan Sly and Nike Sun. Counting in two-spin models on d-regular graphs. *Ann. Probab.*, 42(6):2383–2416, 11 2014.
- [17] Dror Weitz. Counting independent sets up to the tree threshold. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*, pages 140–149, 2006.