

The complexity of approximately counting in 2-spin systems on k -uniform bounded-degree hypergraphs*

Andreas Galanis[†]

Leslie Ann Goldberg[†]

Abstract

One of the most important recent developments in the complexity of approximate counting is the classification of the complexity of approximating the partition functions of antiferromagnetic 2-spin systems on bounded-degree graphs. This classification is based on a beautiful connection to the so-called uniqueness phase transition from statistical physics on the infinite Δ -regular tree. Our objective is to study the impact of this classification on unweighted 2-spin models on k -uniform hypergraphs. As has already been indicated by Yin and Zhao, the connection between the uniqueness phase transition and the complexity of approximate counting breaks down in the hypergraph setting. Nevertheless, we show that for every non-trivial symmetric k -ary Boolean function f there exists a degree bound Δ_0 so that for all $\Delta \geq \Delta_0$ the following problem is NP-hard: given a k -uniform hypergraph with maximum degree at most Δ , approximate the partition function of the hypergraph 2-spin model associated with f . It is NP-hard to approximate this partition function even within an exponential factor. By contrast, if f is a trivial symmetric Boolean function (e.g., any function f that is excluded from our result), then the partition function of the corresponding hypergraph 2-spin model can be computed exactly in polynomial time.

The full version of this paper is available at arxiv.org/abs/1505.06146. For the convenience of the reader, we have also included in the theorem statements the number of the respective theorem in the full version.

1 Introduction

One of the most important recent developments in the complexity of approximate counting is the classification of the complexity of approximating the partition func-

tions of antiferromagnetic 2-spin systems on bounded-degree graphs [10, 17]. This classification is based on a beautiful connection to the so-called uniqueness phase transition from statistical physics on the infinite Δ -regular tree, which was first established in the context of the hard-core model in the works of [18, 16] (see also [6, 13] for related results) and later developed [15, 8, 17, 10] in the more general framework of antiferromagnetic 2-spin systems.

Our objective is to study the impact of this classification on *unweighted* 2-spin models on k -uniform *hypergraphs*. A k -uniform hypergraph $H = (V, \mathcal{F})$ consists of a vertex set V and a set \mathcal{F} of arity- k hyperedges which are k -element subsets of V . Thus, a 2-uniform hypergraph is the same as a graph. The degree of a vertex $v \in V$ is the number of edges that contain v , namely $|\{e \in \mathcal{F} \mid v \in e\}|$. The maximum degree of H is (naturally) the maximum degree of the vertices of H .

A 2-spin model on the class of k -uniform hypergraphs is specified by a symmetric function $f : \{0, 1\}^k \rightarrow \mathbb{R}_+$. The 2-spin model is *unweighted* if the function f is *Boolean*, meaning that its range is a subset of the two-element set $\{0, 1\}$. Given a k -uniform hypergraph $H = (V, \mathcal{F})$, each assignment $\sigma : V \rightarrow \{0, 1\}$ induces a weight

$$w_{f;H}(\sigma) := \prod_{\{v_1, \dots, v_k\} \in \mathcal{F}} f(\sigma(v_1), \dots, \sigma(v_k)).$$

The assignment σ is sometimes referred to as a *configuration*. The *partition function* $Z_{f;H}$ corresponding to f and H is defined as follows.

$$\begin{aligned} Z_{f;H} &:= \sum_{\sigma: V \rightarrow \{0,1\}} w_{f;H}(\sigma) \\ &= \sum_{\sigma: V \rightarrow \{0,1\}} \prod_{\{v_1, \dots, v_k\} \in \mathcal{F}} f(\sigma(v_1), \dots, \sigma(v_k)). \end{aligned}$$

Given a symmetric function $f : \{0, 1\}^k \rightarrow \mathbb{R}_+$ and a hypergraph $H = (V, \mathcal{F})$ we will use $\mu_{f,H}(\cdot)$ to denote the distribution on configurations $\sigma : V \rightarrow \{0, 1\}$ in which the probability of configuration σ is proportional to its weight so $\mu_{f;H}(\sigma) \propto w(\sigma)$. The distribution

*The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) ERC grant agreement no. 334828. The paper reflects only the authors' views and not the views of the ERC or the European Commission. The European Union is not liable for any use that may be made of the information contained therein.

[†]Department of Computer Science, University of Oxford, Wolfson Building, Parks Road, Oxford, OX1 3QD, UK. andreas.galanis@cs.ox.ac.uk, leslie.goldberg@cs.ox.ac.uk

$\mu_{f,H}(\cdot)$ is called the *Gibbs distribution* associated with the partition function $Z_{f,H}$.

The computational problem that we study is the problem of approximating $Z_{f,H}$, given H as input. Formally, this problem has three parameters — a symmetric arity- k Boolean function f , a degree bound Δ , and a value $c > 1$ which specifies the desired accuracy of the approximation. The problem is defined as follows.

Name #Hyper2Spin(f, Δ, c).

Instance An n -vertex k -uniform hypergraph H with maximum degree at most Δ .

Output A number \hat{Z} such that $c^{-n}Z_{f,H} \leq \hat{Z} \leq c^n Z_{f,H}$.

The most well-known example of an unweighted 2-spin model is the independent set model on graphs. In this case $k = 2$, and f is the function given by $f(0,0) = f(0,1) = f(1,0) = 1$ and $f(1,1) = 0$. Independent sets are in one-to-one correspondence with configurations in the model — vertices that are in a given independent set are assigned spin 1 by the corresponding configuration σ . The partition function $Z_{f,H}$ is simply the number of independent sets of the graph H .

Let us now consider larger arity. A (weak) independent set in a hypergraph is a subset of vertices that does not contain a hyperedge as a subset. Weak independent sets correspond to configurations in the unweighted 2-spin model in which f is the function $f: \{0,1\}^k \rightarrow \{0,1\}$ where $f(s_1, \dots, s_k) = 1$ iff at least one of s_1, \dots, s_k is 0. A strong independent set in a hypergraph is a subset of vertices that does not contain more than one vertex of any given hyperedge. Strong independent sets correspond to the unweighted 2-spin model in which $f(s_1, \dots, s_k) = 1$ iff at most one of s_1, \dots, s_k is 1. Note that the two notions of hypergraph independent set coincide in the case $k = 2$, which is the graph case that we have already considered.

The main motivation for our work is the following striking result about the complexity of approximating the partition function of the independent set model on bounded-degree graphs: (i) There exists an FPRAS for the number of independent sets in graphs of maximum degree at most 5 [18]; (ii) There is no FPRAS for the number of independent sets in graphs of maximum degree at most 6 [16] (unless NP=RP). This computational transition was proved using insights from phase transitions and, in fact, the transition coincides with the so-called uniqueness threshold of the independent set model on the infinite Δ -regular tree.

The question that we seek to address in this work is whether a similar computational transition occurs

for the complexity of approximating the partition function of (unweighted) 2-spin models on k -uniform hypergraphs, in terms of the maximum degree Δ . While the case $k = 2$ is completely covered by the results in the previous paragraph, the picture for $k \geq 3$ appears to be much more intricate and the complexity threshold may differ from the uniqueness threshold.

This issue has been discussed in [19] in the special case of approximately counting the strong independent sets of a hypergraph. While the full picture is still incomplete, it is useful to see why the complexity threshold may differ from the uniqueness threshold in this particular model for $k = 3$. As is implicit in [11] and was spelled out explicitly in [19], uniqueness holds for this model on the infinite Δ -regular 3-uniform hypertree if and only if $\Delta \leq 3$. For $\Delta \leq 3$, the results of [11, 19] establish that a (non-trivial) analogue of Weitz’s self-avoiding walk computational-tree approach yields an efficient approximation scheme for the partition function by (implicitly) establishing a strong spatial mixing result. Strong spatial mixing does not hold when $\Delta \geq 4$ because the infinite Δ -regular 3-uniform hypertree is in non-uniqueness. While it is known that it is hard to approximate the partition function for $\Delta \geq 6$, Yin and Zhao [19] show that natural gadgets cannot be used to show hardness for $\Delta = 4, 5$ and these cases remain open.

Generally, as the results of [11, 19] suggest, one would expect that, for “natural” functions f , an FPRAS should exist up to the strong spatial mixing threshold, but this is (in general) below the uniqueness threshold of the Δ -regular k -uniform hypertree.

Above the uniqueness threshold, approximating the partition function may be hard, but this is not known in general, even for the special case of strong independent sets. Thus, it is not clear from the literature that there is a computational threshold where approximating the partition function on hypergraphs of maximum degree Δ becomes intractable and it is not clear whether this threshold, if it exists, coincides with the uniqueness threshold.

The main contribution of this paper is showing that, for every function f (apart from seven special “easy” functions), there is indeed a “barrier” value Δ_0 such that for all $\Delta \geq \Delta_0$, it is NP-hard to approximate the partition function.

DEFINITION 1.1. (DEFINITION 1) *For $k \geq 2$, let $\text{EASY}(k)$ be the set containing the following seven functions.*

$$\begin{aligned} f_{\text{zero}}^{(k)}(x_1, \dots, x_k) &= 0, \\ f_{\text{one}}^{(k)}(x_1, \dots, x_k) &= 1, \end{aligned}$$

$$\begin{aligned}
f_{\text{allzero}}^{(k)}(x_1, \dots, x_k) &= \mathbf{1}\{x_1 = \dots = x_k = 0\}, \\
f_{\text{allone}}^{(k)}(x_1, \dots, x_k) &= \mathbf{1}\{x_1 = \dots = x_k = 1\}, \\
f_{\text{EQ}}^{(k)}(x_1, \dots, x_k) &= \mathbf{1}\{x_1 = \dots = x_k\}, \\
f_{\text{even}}^{(k)}(x_1, \dots, x_k) &= \mathbf{1}\{x_1 \oplus \dots \oplus x_k = 0\}, \\
f_{\text{odd}}^{(k)}(x_1, \dots, x_k) &= \mathbf{1}\{x_1 \oplus \dots \oplus x_k = 1\}.
\end{aligned}$$

Considering each of the functions in $\text{EASY}(k)$, we obtain the following observation.

OBSERVATION 1.1. (OBSERVATION 2) *Let $k \geq 2$ and $f \in \text{EASY}(k)$. Then, the problem of (exactly) computing $Z_{f,H}$, given as input a k -uniform hypergraph H , can be solved in time polynomial in the size of H .*

Our main theorem is a contrasting hardness result.

THEOREM 1.1. (THEOREM 3) *Let $k \geq 2$ and let $f : \{0,1\}^k \rightarrow \{0,1\}$ be a symmetric Boolean function such that $f \notin \text{EASY}(k)$. Then, there exists Δ_0 such that for all $\Delta \geq \Delta_0$, there exists $c > 1$ such that $\#\text{Hyper2Spin}(f, \Delta, c)$ is NP-hard.*

Thus we show that for all $k \geq 2$, for all non-trivial symmetric Boolean functions f , for all sufficiently large Δ , it is NP-hard to approximate $Z_{f,H}$, even within an exponential factor, given a k -uniform hypergraph H of maximum degree at most Δ . We do not pursue the task of obtaining an explicit bound on Δ , since this would require heavy numerical work (depending on the function f) and we do not expect that it would yield the exact value of the threshold, even if such a threshold exists.

1.1 Counting Constraint Satisfaction and Related Results Suppose that Γ is a set of Boolean functions of different arities. Thus, an arity- k function in Γ is a function from $\{0,1\}^k$ to $\{0,1\}$. The counting constraint satisfaction problem $\#\text{CSP}(\Gamma)$ is the problem of computing the CSP partition function $Z_{\Gamma,I}$ where I is a CSP instance consisting of a set V of variables and a set \mathcal{S}^1 of constraints, where each constraint $C = (v_1, \dots, v_k, f) \in \mathcal{S}$ constrains the variables v_1, \dots, v_k by applying a particular k -ary function $f \in \Gamma$. The value of the partition function is given by

$$Z_{\Gamma,I} := \sum_{\sigma: V \rightarrow \{0,1\}} \prod_{(v_1, \dots, v_k, f) \in \mathcal{S}} f(\sigma(v_1), \dots, \sigma(v_k)).$$

¹The reader who is familiar with weighted counting CSP may have expected \mathcal{S} to be a multiset rather than a set, but this is not necessary here since the functions in Γ have range $\{0,1\}$. Restricting \mathcal{S} to be a set allows some technical simplifications later.

The constraint $C = (v_1, \dots, v_k, f)$ could use a particular variable more than once. For example, it is possible that v_1 and v_2 are both the same variable. The problem $\#\text{CSP}_{\Delta}(\Gamma)$ is the problem of computing $Z_{\Gamma,I}$ given an instance I in which each variable is used at most Δ times. We can also define a related approximation problem, similar to $\#\text{Hyper2Spin}(f, \Delta, c)$.

Name $\#\text{CSP}_{\Delta,c}(\Gamma)$.

Instance An n -variable instance I of a CSP in which all constraints apply functions from Γ and each variable is used at most Δ times.

Output A number \hat{Z} such that $c^{-n} Z_{\Gamma,I} \leq \hat{Z} \leq c^n Z_{\Gamma,I}$.

It is clear that our problem $\#\text{Hyper2Spin}(f, \Delta, c)$ is closely related to the problem $\#\text{CSP}_{\Delta,c}(\{f\})$. In particular, $\#\text{Hyper2Spin}(f, \Delta, c)$ is the special case of $\#\text{CSP}_{\Delta,c}(\{f\})$ in which constraints are not allowed to re-use variables. Thus, Theorem 1.1 has the following immediate corollary.

COROLLARY 1.1. (COROLLARY 4) *Let $k \geq 2$ and let $f : \{0,1\}^k \rightarrow \{0,1\}$ be a symmetric Boolean function such that $f \notin \text{EASY}(k)$. Then, there exists Δ_0 such that for all $\Delta \geq \Delta_0$, there exists $c > 1$ such that $\#\text{CSP}_{\Delta,c}(\{f\})$ is NP-hard.*

The combined results of [4] and [2] show that for (exact) counting CSPs, adding a degree bound $\Delta \geq 3$ does not change the complexity of the problem. The situation is less clear for decision and approximate counting. Previous work on bounded-degree decision CSP [5] and bounded-degree approximate counting CSP [7] considered only the so-called conservative model where intractability arises more easily. In this model, δ_0 is the unary pinning-to-0 function which is defined by $\delta_0(0) = 1$ and $\delta_0(1) = 0$. Also, δ_1 is the unary pinning-to-1 function which is defined by $\delta_1(0) = 0$ and $\delta_1(1) = 1$.

Theorem 24 of [7] allows us to deduce (see Observation 27 in the full version) that for every $\Delta \geq 6$, $k \geq 2$ and symmetric k -ary Boolean function $f \notin \text{EASY}(k)$, there is no FPRAS for $\#\text{CSP}(\{f, \delta_0, \delta_1\})$ unless $\text{NP} = \text{RP}$. This hardness result extends from the independent set uniqueness phase transition at $\Delta = 6$ because pinning allows constructions which realise arbitrarily bad configurations.

The result of [7] does not apply to our hypergraph 2-spin context where the pinning functions δ_0 and δ_1 are not present. To see this, consider the following contrasting positive result of [1] which is proved via the MCMC method: an FPRAS exists for approximating the number of (weak) independent sets in a k -uniform hypergraph of maximum degree Δ whenever $k \geq 2\Delta + 1$.

Thus, even though the weak independent function f (given by $f(s_1, \dots, s_k) = 1$ iff at least one of s_1, \dots, s_k is 0) is not in $\text{EASY}(k)$ for any $k \geq 2$, the result of Bordewich et al. [1] shows that for every $\Delta \leq (k-1)/2$, there is an FPRAS for the partition function $Z_{f,H}$ on the class of k -uniform hypergraphs H with maximum degree at most Δ .

Thus, it is clear that $\Delta = 6$ cannot always be a computational threshold in the hypergraph 2-spin framework (where there is no pinning). However, our Theorem 1.1 shows that, for every non-trivial symmetric Boolean function f , there is degree-bound Δ_0 such that approximating the partition function is intractable beyond this degree bound.

1.2 Proof Overview In order to prove Theorem 1.1, we will construct a k -uniform hypergraph H such that the spin-system induced by f on H induces an anti-ferromagnetic binary 2-spin model that is in the non-uniqueness region of the corresponding Δ -regular tree. It follows from a result of Sly and Sun that approximating the partition function of the binary model is intractable, and we will use this to show that approximating the partition function of the k -ary model is also intractable.

While this high-level approach is analogous to the one followed in [7], in our setting where the pinning functions δ_0 and δ_1 are not available, we have to tackle several obstacles. A first indicator of the difficulties that arise is that, in [7], the target 2-spin model is always the independent set model (largely due to the availability of the pinning functions δ_0 and δ_1). In contrast, our target binary 2-spin model will be weighted and depend on the function f . In fact, we will only know the parameters of the binary 2-spin model only approximately which, as we shall discuss later in detail, poses difficulties in showing that it is intractable.

To explain the argument in more detail, let us backtrack and discuss a natural approach that one might hope would lead to proof of Theorem 1.1. First, if one were able to construct hypergraphs to “realise” the pinning functions δ_0 and δ_1 then these hypergraphs could be combined with the reduction in [7] to prove Theorem 1.1. The proof would even be straightforward if perfect realisations could be found. For example, to realise δ_0 perfectly we would need a hypergraph H whose partition function is non-zero which has a vertex v such that every configuration σ with $w_{f,H}(\sigma) > 0$ satisfies $\sigma(v) = 0$. More realistically, one might hope that even “approximate” versions of the pinning functions δ_0 or δ_1 would suffice to simulate the reduction in [7]. Unfortunately, this fails rather formidably: first, as we shall see below, there are functions f which simply

cannot realise (approximate) pinning, and, second, even for those functions f which do support pinning, the bounded-degree assumption poses strict limits on the accuracy of the approximations that can be achieved.

Despite the failure of the above approach, it does turn out to be useful to explore the extent to which the pinning functions δ_0 and δ_1 can be simulated using hypergraphs. We know from the binary case (where the uniqueness phase transition coincides with the computational transition) that the achievable “boundary conditions” play an important role. Understanding the pinnings that can be (approximately) achieved gives us the relevant boundary information for the higher-arity case. To make the following discussion concrete, consider the following definition (stated more generally for weighted functions f).

DEFINITION 1.2. (DEFINITION 5) *Let $f : \{0,1\}^k \rightarrow \mathbb{R}_+$ be symmetric. Suppose that $\varepsilon \geq 0$ and $s \in \{0,1\}$. The hypergraph H is an ε -realisation of pinning-to- s if there exists a vertex v in H such that $\mu_{f,H}(\sigma_v = s) \geq 1 - \varepsilon$. We will refer to v as the terminal of H .*

Note that the perfect realisation discussed earlier corresponds to taking $\varepsilon = 0$. Suppose that we have an ε -realisation of pinning-to- s but we want an ε' -realisation for some very small $\varepsilon' > 0$. This can be done via standard powering (see Lemma 12 in the full version): Given a hypergraph H which ε -realises pinning-to- s for some $\varepsilon < 1/2$, one can construct a hypergraph H' which ε' -realises pinning-to- s . Note, however, that the size of H' may depend on ε' . For example, in the construction of Lemma 12, the maximum degree of H' is proportional to $\log(1/\varepsilon')$. Nevertheless, the possibility of powering motivates the following definition.

DEFINITION 1.3. (DEFINITION 6) *Let $s \in \{0,1\}$. We say that f supports pinning-to- s if for every $\varepsilon > 0$, there is a (finite) hypergraph H which is an ε -realisation of pinning-to- s .*

We will next consider an example which demonstrates the limits of what can be achieved. Let $f : \{0,1\}^k \rightarrow \{0,1\}$ be the weak independent set function where $f(s_1, \dots, s_k) = 1$ if and only if at least one of s_1, \dots, s_k is 0. First, note that f does not support pinning-to-1 since for every hypergraph H and every vertex v in H it holds that $\mu_{f,H}(\sigma(v) = 1) \leq 1/2$. The function f does support pinning-to-0 but there is still a limit on how small ε can be. In particular, for every k -uniform hypergraph H with maximum degree Δ , and every vertex v of H we can obtain the crude bound $\mu_{f,H}(\sigma(v) = 0) \leq 1 - 1/2^{k\Delta}$. This shows that we cannot hope to pin the spin of a vertex to 0 with ar-

bitrary polynomial precision using bounded-degree hypergraphs. Note that this example already shows that it is impossible to prove Theorem 1.1 by approximating the pinning functions δ_0 and δ_1 and then applying the result of [7].

Nevertheless, pinning-to-0 and pinning-to-1 will be important for us since, whenever a function f supports one (or both) of these notions, we will be able to use them to decrease the arity of the function f . This is particularly useful since, recall, our ultimate goal is to obtain an intractable binary 2-spin model. Intriguingly, there are functions f which do not support either pinning-to-0 or pinning-to-1. For example, consider the function f which is induced by the “not-all-equal” constraint. Then, for every hypergraph H with $Z_{f,H} > 0$ and every vertex v of H , it is easy to see that

$$\mu_{f,H}(\sigma(v) = 1) = \mu_{f,H}(\sigma(v) = 0) = 1/2.$$

More generally, the same phenomenon holds for any function f whose value is unchanged when the argument is complemented; such functions are called “self-dual”.

The first point that we address in this work is a complete characterisation of the functions f which support either the notion of pinning-to-0 or pinning-to-1. We show (Lemma 2.1 and Lemma 2.2) that any function f other than those that are self-dual do support either pinning-to-0 or pinning-to-1 (but perhaps not both).

We show this classification even for weighted functions f , see Section 2 for more details. The classification allows us to split the proof of Theorem 1.1 into three cases: (i) f supports both pinning-to-0 and pinning-to-1, (ii) f is self-dual, and (iii) f supports exactly one of pinning-to-0 and pinning-to-1.

In cases (i) and (iii) where pinning is available we show how to use the approximate pinning to simulate binary antiferromagnetic 2-spin models that are intractable. A difficulty that arises in the proof is that not every anti-ferromagnetic binary 2-spin model is in the non-uniqueness region. In fact, there are relevant values of the parameters for which the corresponding binary 2-spin model is actually in the uniqueness region for all sufficiently large Δ . To make matters worse, we will not be able to control the parameters of the resulting binary model with perfect accuracy. In particular, to analyse the k -ary gadgets, we will use ε -realisations of pinnings via hypergraphs for some small $\varepsilon > 0$. Thus, we are faced with the possibility that the idealised binary 2-spin model (i.e., the one corresponding to $\varepsilon = 0$) may be in the non-uniqueness region, but we need to prove that the approximate version that we actually achieve is also in the non-uniqueness region. In fact, the idealised binary 2-spin model will sometimes even be on

the boundary of the region where intractability holds for sufficiently large Δ , which makes our task harder.

Our approach to this is to revisit (Section 3) antiferromagnetic binary 2-spin models, showing (Lemma 3.2) that there is a sufficiently-wide strip outside of the natural square where the parameters are at most 1 where the system is in the non-uniqueness region. We will then carefully ensure that all of the idealised systems are inside this strip, so that even the approximations are still in non-uniqueness.

In case (ii) (Section 4.1), where the function f is self-dual and hence no pinning is possible, we first classify those self-dual functions f where the related decision problem is NP-hard. In order to do so, we use techniques (polymorphisms) from constraint satisfaction. While this hardness is not for the bounded-degree setting, we show how to lift the results to bounded-degree hypergraphs by showing that one can force the spins of two vertices to be equal (Lemma 4.2). The proof for this class of self-dual functions f is given in Section 4.1.1. For those self-dual functions f where the associated decision problem is not hard (Section 4.1.2), we show that one can realise approximate equality in the following sense.

DEFINITION 1.4. (DEFINITION 7) *Let $\varepsilon \geq 0$ and $t \geq 2$ be an integer. The hypergraph H is an ε -realisation of t -equality if there are distinct vertices v_1, \dots, v_t such that for $s \in \{0, 1\}$, $\mu_{f,H}(\sigma_{v_1} = \dots = \sigma_{v_t} = s) \geq (1 - \varepsilon)/2$.*

DEFINITION 1.5. (DEFINITION 8) *A function f supports t -equality if for every $\varepsilon > 0$, there is a (finite) hypergraph H which is an ε -realisation of t -equality.*

Using the upcoming Lemmas 2.3 and 2.1, we show that a self-dual function f supports t -equality for every integer $t \geq 2$. Roughly, this allows us to decrease the arity of the function by carefully using (approximate) equality to obtain an anti-ferromagnetic binary 2-spin model which is intractable (note that we again have to deal with the approximation issue that we described for cases (i) and (iii)).

2 Classifying functions with respect to pinning and equality

In this section, we study the concepts of pinning and equality that we will use for the proof of Theorem 1.1. While our primary interest is in symmetric *Boolean* functions f , the results of this section extend effortlessly to non-negative symmetric functions f with domain $\{0, 1\}^k$ and range \mathbb{R}_+ . For the remainder of this section, we consider a symmetric function $f : \{0, 1\}^k \rightarrow \mathbb{R}_+$. Since f is symmetric, there are values $w_0, w_1, \dots, w_k \in \mathbb{R}_+$ such that

$$f(x_1, \dots, x_k) = w_\ell$$

whenever

$$x_1 + \dots + x_k = \ell.$$

We will refer to f and to the values w_ℓ in the proofs.

The following lemma will be used in our classification.

LEMMA 2.1. (LEMMA 13) *Let $k \geq 2$. For all $f : \{0, 1\}^k \rightarrow \mathbb{R}_+$ which are not constant, it holds that f supports at least one of pinning-to-0, pinning-to-1 and 2-equality.*

Proof. (Sketch) Assume that f does not support pinning-to-0 or pinning-to-1. We will show that f supports 2-equality. Let H be the hypergraph with vertex set $\{x, y, z_1, \dots, z_{k-1}\}$ and hyperedge set $\mathcal{F} = \{e_1, e_2\}$, where $e_1 = \{x, z_1, \dots, z_{k-1}\}$ and $e_2 = \{y, z_1, \dots, z_{k-1}\}$. For $s_1, s_2 \in \{0, 1\}$, we have that

$$\mu_{f;H}(\sigma_x = s_1, \sigma_y = s_2) \propto Z_{s_1 s_2},$$

where

$$Z_{s_1 s_2} = \sum_{\ell=0}^{k-1} \binom{k-1}{\ell} w_{\ell+s_1} w_{\ell+s_2}.$$

Note that $Z_{01} = Z_{10}$.

We first show that $Z_{00} = Z_{11}$. Assume otherwise. Note that

$$\mu_{f;H}(\sigma_x = 0) \propto Z_{01} + Z_{00},$$

$$\mu_{f;H}(\sigma_x = 1) \propto Z_{10} + Z_{11}.$$

Since $Z_{01} = Z_{10}$, $Z_{00} \neq Z_{11}$ would imply $\mu_{f;H}(\sigma_x = 0) \neq \mu_{f;H}(\sigma_x = 1)$, contradicting that f does not support pinning-to-0 or pinning-to-1 (see Items 1 and 2 of Lemma 12 in the full version).

Further, we have that $Z_{00} Z_{11} \geq Z_{01}^2$, since

$$(2.1) \quad \left[\sum_{\ell=0}^{k-1} \binom{k-1}{\ell} w_\ell^2 \right] \left[\sum_{\ell=0}^{k-1} \binom{k-1}{\ell} w_{\ell+1}^2 \right] \geq \left[\sum_{\ell=0}^{k-1} \binom{k-1}{\ell} w_\ell w_{\ell+1} \right]^2$$

holds as an immediate consequence of the Cauchy-Schwartz inequality. From $Z_{00} = Z_{11}$, we thus obtain that $Z_{00} \geq Z_{01}$. Equality in (2.1) holds only if there exists $\alpha \geq 0$ such that $w_{\ell+1} = \alpha w_\ell$ for every $\ell = 0, \dots, k-1$, which yields $w_\ell = \alpha^\ell w_0$ for $\ell = 0, \dots, k$. This gives $Z_{11} = \alpha^2 Z_{00}$, so $Z_{00} = Z_{11}$ leaves only the possibility $\alpha = 1$, which in turn yields that f is a constant function.

Thus, it holds that $Z_{00} = Z_{11} > Z_{01} = Z_{10}$, so by powering-up arguments f supports 2-equality (see Item 3 of Lemma 12 in the full version for details)

Recall that self-dual functions do not support pinning-to-0 or pinning-to-1, so Lemma 2.1 yields that self-dual functions support 2-equality. We show that the converse holds as well, in the following sense.

LEMMA 2.2. (LEMMA 14) *Let $k \geq 2$. If f supports 2-equality but neither pinning-to-0 nor pinning-to-1, then it holds that $w_\ell = w_{k-\ell}$ for all $\ell = 0, \dots, k$.*

To prove this, we will use the following (simple) lemma which allows us to say that f “supports equality” iff it supports 2-equality (hence t -equality for every $t \geq 2$).

LEMMA 2.3. (LEMMA 10) *Let $t \geq 2$ be an integer. The function f supports t -equality iff f supports 2-equality.*

Proof. (Proof of Lemma 2.2 — Sketch) Consider the hypergraph with the single hyperedge $\{v_1, \dots, v_k\}$. For $t = 1, \dots, k$ derive linear equations on the variables w_0, \dots, w_k by considering gadgets which force equality among the vertices $\{v_1, \dots, v_t\}$ (such gadgets exist since f supports t -equality by Lemma 2.3). Since f does not support pinning-to-0 or pinning-to-1, it must be the case that $\mu_{f;H}(\sigma(v_k) = 1) = 1/2$ (otherwise, one can use powering to show that f supports either pinning-to-0 or pinning-to-1). This gives an additional linear equation. We then show that these linear equations are satisfied iff $w_\ell = w_{k-\ell}$ for all $\ell = 0, \dots, k$.

3 A general inapproximability lemma

The purpose of this section is to prove Lemma 3.1, which will allow us to exploit our study of pinning-to-0, pinning-to-1 and equality. Given a set S of vertices, it will be convenient to write $\sigma_S = \mathbf{0}$ to denote the event that all vertices in S are assigned the spin 0 under the assignment σ . We will similarly write $\sigma_S = \mathbf{1}$. We will also use σ_S^{eq} to denote the event that all vertices in S have the same spin under σ (the spin could be 0 or 1). We use the following technical definition.

DEFINITION 3.1. (DEFINITION 16) *Let $f : \{0, 1\}^k \rightarrow \mathbb{R}_+$ be symmetric. Let $H = (V, \mathcal{F})$ be a hypergraph and assume that $\mathcal{V} := (V_0, V_1, V_2, \dots, V_r)$ is a labelled collection of disjoint subsets of V such that:*

1. $V_0 = \emptyset$ if f does not support pinning-to-0,
2. $V_1 = \emptyset$ if f does not support pinning-to-1,
3. $V_2 = \dots = V_r = \emptyset$ if f does not support equality,
4. $\mu_{f;H}(\sigma_{V_0} = \mathbf{0}, \sigma_{V_1} = \mathbf{1}, \sigma_{V_2}^{\text{eq}}, \dots, \sigma_{V_r}^{\text{eq}}) > 0$.

We will then say that \mathcal{V} is admissible for the hypergraph H (with respect to f) and denote by $\mu_{f;H}^{\text{cond}(\mathcal{V})}$ the probability distribution

$$\mu_{f;H}(\cdot \mid \sigma_{V_0} = \mathbf{0}, \sigma_{V_1} = \mathbf{1}, \sigma_{V_2}^{\text{eq}}, \dots, \sigma_{V_r}^{\text{eq}}).$$

LEMMA 3.1. (LEMMA 18) *Let $f : \{0,1\}^k \rightarrow \mathbb{R}_+$ be symmetric. Let H be a hypergraph, let \mathcal{V} be admissible for H and let x and y be vertices of H . For $s_1, s_2 \in \{0,1\}$, define $\mu_{s_1 s_2}$ by*

$$\mu_{s_1 s_2} := \mu_{f;H}^{\text{cond}(\mathcal{V})}(\sigma(x) = s_1, \sigma(y) = s_2).$$

Suppose that

$$\begin{aligned} \mu_{00} + \mu_{11} &> 0, \quad \min\{\mu_{00}, \mu_{11}\} < \sqrt{\mu_{01}\mu_{10}}, \\ \max\{\mu_{00}, \mu_{11}\} &\leq \sqrt{\mu_{01}\mu_{10}}. \end{aligned}$$

Then, for all sufficiently large Δ , there exists $c > 1$ such that $\#\text{Hyper2Spin}(f, \Delta, c)$ is NP-hard.

The proof of Lemma 3.1 uses inapproximability results for antiferromagnetic 2-spin systems on bounded-degree graphs. Thus, before giving its proof, we detour to extract the ingredients that we need.

REMARK 3.1. (REMARK 19) *Note that the inequalities for the μ_{ij} 's are stronger than the standard antiferromagnetic condition $\mu_{00}\mu_{11} < \mu_{01}\mu_{10}$ for 2-spin models on graphs. This is to ensure that the corresponding 2-spin system lies in the non-uniqueness region for all sufficiently large Δ (and hence is intractable). In fact, if $\max\{\mu_{00}, \mu_{11}\} > \sqrt{\mu_{01}\mu_{10}}$, for the corresponding binary 2-spin system (even if it is antiferromagnetic), approximating its partition function may be tractable for all graphs (when the external field is fixed).*

We review inapproximability results for the partition function of antiferromagnetic 2-spin models on graphs. We start with a few relevant definitions following [10]. A 2-spin model on a graph is specified by three parameters $\beta, \gamma \geq 0$ and $\lambda > 0$. For a graph $G = (V, E)$, configurations are assignments $\sigma : V \rightarrow \{0,1\}$ and the partition function is given by

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

The 2-spin system with parameters β, γ, λ is called *antiferromagnetic* if $\beta\gamma < 1$. Sly and Sun [17] showed that if an anti-ferromagnetic 2-spin system specified by the parameters β, γ, λ is in the non-uniqueness regime of the infinite Δ -regular tree, then there is a $c > 1$ such that it is NP-hard to approximate $Z_{\beta, \gamma, \lambda; G}$ within a factor of c^n on Δ -regular graphs G . (We will simply cite Sly and Sun when we use this result below, without being more explicit. Details are in the full version.)

For us, the case $\lambda = 1$ (which is usually referred to as the case without an external field) will be especially important. To motivate what follows, the reader

should first bear in mind the following two facts (see [10, Lemma 21]) about the uniqueness regime for antiferromagnetic 2-spin systems. The two cases correspond to whether or not one of the parameters β, γ is larger than 1. These parameters cannot both be larger than 1 because of the antiferromagnetic condition $\beta\gamma < 1$.

1. when β and γ satisfy $0 \leq \beta < 1$ and $0 < \gamma \leq 1$, non-uniqueness holds on the infinite Δ -regular tree for all sufficiently large Δ .
2. when β and γ satisfy $0 \leq \beta < 1$ and $\gamma > 1$ then uniqueness holds on the infinite Δ -regular tree for all sufficiently large Δ .

In order to prove Theorem 1.1, we will construct a family of k -uniform hypergraphs so that the 2-spin model that f induces on these hypergraphs simulates an anti-ferromagnetic binary 2-spin model. Thus, the constructed hypergraphs will be viewed as binary gadgets. It will be important that the induced binary 2-spin model is in the non-uniqueness region so that we can prove hardness using the result of Sly and Sun [17]. Our constructions will use the conditional distributions induced by pinning or equality to simplify the analysis of the gadgets.

The conditional distribution will yield an idealised antiferromagnetic 2-spin system with parameters β_0 and γ_0 , say. The delicate issue that arises is that the hypergraphs that we can construct to simulate these conditional distributions (see Lemmas 15 and 17 in the full version) are imperfect. There is always a small error ε . So even if the ideal antiferromagnetic spin-system given by β_0 and γ_0 is in the non-uniqueness region, we will have constructed some nearby binary spin-system given by (say) parameters β and γ and we will have to prove that the spin-system given by β and γ is also an anti-ferromagnetic spin system in the non-uniqueness region.

In general, the error bound (that we get from Lemma 17) will tell us that for some small constant ε , $|\beta - \beta_0| < \varepsilon$ and $|\gamma - \gamma_0| < \varepsilon$. The most difficult case will be when γ_0 is close to 1 (including the case where γ_0 is actually 1). In this case, we might have γ slightly larger than 1 and we will thus need to exclude Item 2 above.

In order to overcome these obstacles, we rely on making the error ε very small, at the expense, of course, of potentially increasing the degree bound Δ . By a continuity-type of argument, we show that for β_0 strictly less than 1 and $\gamma_0 \leq 1$, for all β, γ which are sufficiently close to β_0, γ_0 , there exists a Δ such that the 2-spin system with parameters β, γ is in the non-uniqueness regime of the infinite Δ -regular tree (which can then be used to derive hardness). In Lemma 3.2, which is proved

in the full version, we actually prove a slightly stronger statement by giving a bound on the required accuracy ε in terms of the degree Δ , which allows us to switch the order of quantifiers. Also, our result is monotone in the degree-bound Δ (as in Item 1 above).

LEMMA 3.2. (LEMMA 22) *Suppose $0 \leq \beta_0 < 1$. Then, for all sufficiently large Δ , for $\varepsilon = 1/\Delta$, for all β, γ which satisfy $\max\{\beta_0 - \varepsilon, 0\} \leq \beta < \beta_0 + \varepsilon$ and $0 < \gamma < 1 + \varepsilon$, the 2-spin system with parameters β, γ and $\lambda = 1$ (no external field) is antiferromagnetic and in the non-uniqueness regime of the infinite Δ -regular tree.*

Proof. [Proof of Lemma 3.1 (Sketch)] The idea is to first obtain a hypergraph that realises the conditional distribution with sufficient small accuracy ε (see Lemma 17). The resulting hypergraph can be used to simulate an antiferromagnetic 2-spin system which, by Lemma 3.2 and the result of Sly and Sun [17], will be hard to approximate on Δ -regular graphs (for large Δ).

4 Proof of Theorem 1.1

In this section, we discuss the proof of Theorem 1.1. Details are in the full version. Let $k \geq 2$ and $f : \{0, 1\}^k \rightarrow \{0, 1\}$ be a *symmetric Boolean function* with $f \notin \text{EASY}(k)$. Our goal is to show that there exists Δ_0 such that for all $\Delta \geq \Delta_0$, there exists $c > 1$ such that $\#\text{Hyper2Spin}(f, \Delta, c)$ is NP-hard.

For $k = 2$ this goal is easily reduced to a similar question regarding the independent set model on graphs, which is already understood. Thus, we restrict our attention here to the case $k \geq 3$. We may further assume that at least one of w_0, \dots, w_k is 0 (otherwise f is the constant function $f_{\text{one}}^{(k)}$) and at least one is 1 (otherwise f is the constant function $f_{\text{zero}}^{(k)}$).

By Lemma 2.1, to prove Theorem 1.1 we may split the analysis into the following cases:

1. f supports both pinning-to-0 and pinning-to-1.
2. f supports 2-equality (but neither pinning-to-0 nor pinning-to-1).
3. f supports pinning-to-0 or pinning-to-1 (but not both). (By swapping 0 and 1, it would be identical to assume that f supports pinning-to-0 but not pinning-to-1.)

Here we will discuss the proof for Case 2 — the proofs for the remaining cases are in the full version.

4.1 Hardness for Self-Dual functions (Case II)

Recall that the w_ℓ 's are such that $f(x_1, \dots, x_k) = w_\ell$ whenever $x_1 + \dots + x_k = \ell$. The only difference between the notation here and that of Section 2 is that the w_ℓ 's

are now restricted to be unweighted, i.e., either 0 or 1. By Lemma 2.2 we conclude that f is self-dual, meaning that $w_\ell = w_{k-\ell}$ for all $\ell \in \{0, \dots, k\}$. We deal with Case II by splitting the analysis into two cases — the case where $w_0 = 0$ (Section 4.1.1) and the case where $w_0 = 1$ (Section 4.1.2).

4.1.1 The case $w_0 = 0$ We will first establish the NP-hardness of a related CSP decision problem. This will immediately imply that the corresponding approximate counting CSP problem is also NP-hard. We will need to extend this both by introducing degree bounds and by moving to the more restricted hypergraph 2-spin model where repeated variables are not allowed in constraints. To do both of these, we will prove that f supports *perfect equality*.

Most of the details are deferred to the full version, but here is a sketch. Let $\text{CSP}(\Gamma)$ be the problem of determining whether the partition function $Z_{\Gamma, I}$ is non-zero, given an instance I of a CSP in which all constraints are from the set Γ .

LEMMA 4.1. (LEMMA 25) *Suppose $k > 2$. Let $f \neq f_{\text{zero}}^{(k)}, f_{\text{odd}}^{(k)}$ be an arity- k symmetric Boolean formula that is self-dual and satisfies $w_0 = 0$. Then $\text{CSP}(\{f\})$ is NP-hard.*

Proof. (Sketch) We use Schaefer's famous dichotomy theorem [14] which classifies the more general problem $\text{CSP}(\Gamma)$ determining for which Γ it is tractable. We use an algebraic formulation of Chen [3, Theorem 3.21], which enables us to classify $\Gamma = \{f\}$ by examining algebraic identities called polymorphisms of f . The actual application of Schaefer's theorem is quite lengthy (and can be found in the full version).

The following technical lemma is inspired by techniques from [9] and shows that in the case $w_0 = 0$ we can simulate perfect equality.

LEMMA 4.2. (LEMMA 28) *Suppose $k > 2$. Let $f \neq f_{\text{zero}}^{(k)}, f_{\text{odd}}^{(k)}$ be an arity- k symmetric Boolean formula that is self-dual and satisfies $w_0 = 0$. Then there is a hypergraph H with $Z_{f, H} > 0$ which contains vertices x and y such that for any configuration $\sigma : V(H) \rightarrow \{0, 1\}$ with $w_{f, H}(\sigma) > 0$, we have $\sigma(x) = \sigma(y)$.*

Proof. (Sketch) We first prove that there exists a hypergraph $H_0 = (V_0, \mathcal{F}_0)$ such that $Z_{f, H_0} = 0$ (H_0 is the complete k -uniform hypergraph on $2k - 1$ vertices). We then remove hyperedges of H_0 successively to obtain a hypergraph $H' = (V', \mathcal{F}')$ such that $Z_{f, H'} = 0$ and for every $e \in \mathcal{F}'$, it holds that $Z_{f, H' \setminus e} > 0$. We then choose $e \in \mathcal{F}'$ and we replace successively the vertices in e by

new vertices (without changing the instances of these vertices in other hyperedges). We stop as soon as the hypergraph has a non-zero partition function. We focus on the vertex that caused the switch from zero to non-zero. By copying the (modified) hyperedge e appropriately and adding this to the instance, we ensure equality between two vertices in the resulting hypergraph.

We will now combine Lemmas 4.1 and 4.2 to obtain the $w_0 = 0$ part of Case II.

LEMMA 4.3. (LEMMA 29) *Suppose $k > 2$. Let $f \neq f_{\text{zero}}^{(k)}, f_{\text{odd}}^{(k)}$ be an arity- k symmetric Boolean formula that is self-dual and satisfies $w_0 = 0$. Then there is a Δ_0 such that for every $\Delta \geq \Delta_0$, there exists $c > 1$ such that $\#\text{Hyper2Spin}(f, \Delta, c)$ is NP-hard.*

Proof. (Sketch) Let $\Gamma = \{f\}$. Lemma 4.1 shows that it is NP-hard to determine whether $Z_{\Gamma, I} = 0$ given a CSP(Γ) instance I . From I we show how to (efficiently) construct a k -uniform hypergraph H' with degree at most Δ_0 so that $Z_{\Gamma, I} = 0$ if and only if $Z_{f; H'} = 0$. The construction uses the perfect equality gadget from Lemma 4.2 to reduce the degree and to avoid the use of repeated variables in constraints (hyperedges).

4.1.2 The case $w_0 = 1$ The $w_0 = 1$ case uses the following technical lemma.

LEMMA 4.4. (LEMMA 31) *Let $f : \{0, 1\}^k \rightarrow \{0, 1\}$ be a self-dual symmetric function. Let t_1, t_2 be integers such that $t_1 \geq 1, t_2 \geq 0, 2t_1 + t_2 \leq k$. Suppose that*

$$(4.2) \quad 0 < w_0 + w_{t_2} + w_{2t_1} + w_{2t_1+t_2} < 2(w_{t_1} + w_{t_1+t_2}).$$

Then there exists Δ_0 such that for every $\Delta \geq \Delta_0$, there exists $c > 1$ such that $\#\text{Hyper2Spin}(f, \Delta, c)$ is NP-hard.

Proof. (Sketch) Let $e = \{x_1, \dots, x_k\}$. Let \mathcal{V} force equality on the sets of vertices $\{x_1, \dots, x_{t_1}\}, \{x_{t_1+1}, \dots, x_{2t_1}\}, \{x_{2t_1+1}, \dots, x_{2t_1+t_2}\}$ and, whenever $2t_1+t_2 < k$, on $\{x_{2t_1+t_2+1}, \dots, x_k\}$. Let $t_3 := k - 2t_1 - t_2$ and set $x := x_1$ and $y := x_{t_1+1}$. By the assumptions, we have that $t_3 \geq 0$. We will consider here the case where $t_2 > 0$ and $t_3 > 0$ (note that the condition $2t_1 + t_2 < k$ is equivalent to $t_3 > 0$).

Define

$$\begin{aligned} Z_{00}^{\text{ex}} &= w_0 + w_{t_2} + w_{2t_1} + w_{2t_1+t_2}, \\ Z_{01}^{\text{ex}} &= 2(w_{t_1} + w_{t_1+t_2}), \end{aligned}$$

and note that (4.2) is equivalent to

$$0 < Z_{00}^{\text{ex}} < Z_{01}^{\text{ex}}.$$

In this case, we have that

$$\begin{aligned} \mu_{f;e}^{\text{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 0) &\propto w_0 + w_{t_2} + w_{t_3} + w_{t_2+t_3} = Z_{00}^{\text{ex}}, \\ \mu_{f;e}^{\text{cond}(\mathcal{V})}(\sigma_x = 0, \sigma_y = 1) &\propto w_{t_1} + w_{t_1+t_2} + w_{t_1+t_3} + w_{t_1+t_2+t_3} = Z_{01}^{\text{ex}}, \end{aligned}$$

where the equalities follow by self-duality.

It will then turn out that (4.2) is equivalent to the condition in Lemma 3.1, from which the result follows. The proofs for the remaining cases for the values of t_2, t_3 are analogous and are included in the full version.

We now sketch the remainder of the argument for the $w_0 = 1$ part of Case II. Recall that $k \geq 3$ and that f is a symmetric self-dual arity- k Boolean function that is not in $\text{EASY}(k)$. The function f supports 2-equality but does not support pinning-to-0 or pinning-to-1. Also, $w_0 = 1$. Our goal is to show that for all sufficiently large Δ , there exists $c > 1$ such that the approximation problem $\#\text{Hyper2Spin}(f, \Delta, c)$ is NP-hard.

Let $0 < i \leq k/2$ be the smallest positive index with $w_i = 1$. Clearly, we may assume that such an index i exists (otherwise, by the self-duality of f , we have $f = f_{\text{EQ}}^{(k)}$). The following two claims follow from applying Lemma 4.4 for appropriate values of t_1 and t_2 . Their proofs can be found in the full version.

CLAIM 4.1. (CLAIM 33) *If there is a positive integer r with $ri \leq k$ and $w_{ri} = 0$ then, for all sufficiently large Δ , there exists $c > 1$ such that $\#\text{Hyper2Spin}(f, \Delta, c)$ is NP-hard.*

CLAIM 4.2. (CLAIM 34) *If there is a positive integer $r < k$ that is not divisible by i and has $w_r = 1$ then, for all sufficiently large Δ , there exists $c > 1$ such that $\#\text{Hyper2Spin}(f, \Delta, c)$ is NP-hard.*

The remaining cases that we have to deal with are now quite constrained, satisfying the following properties:

- $w_0 = w_k = 1$ ($w_k = 1$ by self-duality),
- the positive integers $\ell \in \{1, \dots, k-1\}$ with $w_\ell = 1$ are precisely the multiples of i (by Claims 4.1 and 4.2),
- k is a multiple of i . Suppose instead that $k = mi + u$ for some non-negative integer m and some integer $u \in \{1, \dots, i-1\}$. Then $w_{mi} = 1$ since mi is either 0 (if $m = 0$) or it is a positive multiple of i which is less than k . Now $k - mi = u$ so by self-duality $w_u = 1$, contradicting the choice of i .

- $i > 2$. If $i = 1$ then f is the all-one function $f = f_{\text{one}}^{(k)}$. If $i = 2$ then k is even since it is a multiple of i —by Item (iii)—and hence f is the easy function $f = f_{\text{even}}^{(k)}$.
- $2i \leq k$. We know that k is a multiple of i , but if k is equal to i then $f = f_{\text{EQ}}^{(k)}$.

To finish the proof, we consider the hypergraph with a single edge $e = \{x_1, \dots, x_k\}$. Let x be x_{2i-1} and let y be x_{2i} . Let \mathcal{V} force equality among the vertices in $S = \{x_{2i+1}, \dots, x_k\}$. Suppose first that $k > 2i$ (so that $|S| \geq 1$). We will use ℓ to denote the number of spin-one vertices in x_1, \dots, x_{2i-1} . Then, since the assignment to vertices in S can be either the $\mathbf{0}$ or $\mathbf{1}$ assignment, we get

$$(4.3) \quad \begin{aligned} \mu_{f;e}^{\text{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 0) &\propto \sum_{\ell=0}^{2i-2} \binom{2i-2}{\ell} (w_\ell + w_{\ell+k-2i}), \\ \mu_{f;e}^{\text{cond}(\mathcal{V})}(\sigma_x = 0, \sigma_y = 1) &\propto \sum_{\ell=0}^{2i-2} \binom{2i-2}{\ell} (w_{\ell+1} + w_{\ell+k-2i+1}). \end{aligned}$$

But in the range $0 \leq \ell \leq 2i-2$, w_ℓ is only positive if $\ell \in \{0, i\}$. Similarly, by self-duality $w_{\ell+k-2i} = w_{2i-\ell}$, which is only positive if $\ell \in \{0, i\}$. Similarly, $w_{\ell+1}$ is only positive if $\ell = i-1$ and $w_{\ell+k-2i+1} = w_{2i-\ell-1}$ which is only positive if $\ell = i-1$. So, (4.3) becomes

$$\begin{aligned} \mu_{f;e}^{\text{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 0) &\propto 2 + 2 \binom{2i-2}{i}, \\ \mu_{f;e}^{\text{cond}(\mathcal{V})}(\sigma_x = 0, \sigma_y = 1) &\propto 2 \binom{2i-2}{i-1}. \end{aligned}$$

If $k = 2i$ then we get the same equations (apart from a factor of 2, which makes no difference).

To finish the argument, we need only show that $1 + \binom{2i-2}{i} < \binom{2i-2}{i-1}$. We do this using manipulation of Binomial coefficients (see Lemma 32 in the full version). Then we can use Lemma 3.1 to show that for all sufficiently large Δ , there exists $c > 1$ such that $\#\text{Hyper2Spin}(f, \Delta, c)$ is NP-hard. This completes the proof of this case.

References

- [1] Magnus Bordewich, Martin E. Dyer, and Marek Karpinski. Path coupling using stopping times and counting independent sets and colorings in hypergraphs. *Random Struct. Algorithms*, 32(3):375–399, 2008.
- [2] Jin-Yi Cai, Pinyan Lu, and Mingji Xia. The complexity of complex weighted boolean #CSP. *J. Comput. Syst. Sci.*, 80(1):217–236, 2014.
- [3] Hubie Chen. A rendezvous of logic, complexity, and algebra. *ACM Comput. Surv.*, 42(1), 2009.
- [4] Nadia Creignou and Miki Hermann. Complexity of generalized satisfiability counting problems. *Inf. Comput.*, 125(1):1–12, 1996.
- [5] Víctor Dalmau and Daniel K. Ford. Generalized satisfiability with limited occurrences per variable: A study through delta-matroid parity. In *Mathematical Foundations of Computer Science 2003, 28th International Symposium, MFCS 2003, Bratislava, Slovakia, August 25–29, 2003, Proceedings*, pages 358–367, 2003.
- [6] Martin E. Dyer, Alan M. Frieze, and Mark Jerrum. On counting independent sets in sparse graphs. *SIAM J. Comput.*, 31(5):1527–1541, 2002.
- [7] Martin E. Dyer, Leslie Ann Goldberg, Markus Jalseenius, and David Richerby. The complexity of approximating bounded-degree boolean #CSP. *Inf. Comput.*, 220:1–14, 2012.
- [8] Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. Inapproximability of the partition function for the antiferromagnetic Ising and hard-core models. *CoRR*, abs/1203.2226, 2012.
- [9] Jan Kratochvíl, Petr Savický, and Zsolt Tuza. One more occurrence of variables makes satisfiability jump from trivial to np-complete. *SIAM J. Comput.*, 22(1):203–210, 1993.
- [10] Liang Li, Pinyan Lu, and Yitong Yin. Correlation decay up to uniqueness in spin systems. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6–8, 2013*, pages 67–84, 2013.
- [11] Jingcheng Liu and Pinyan Lu. FPTAS for counting monotone CNF. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4–6, 2015*, pages 1531–1548, 2015.
- [12] Michael Luby and Eric Vigoda. Fast convergence of the Glauber dynamics for sampling independent sets. *Random Struct. Algorithms*, 15(3–4):229–241, 1999.
- [13] Elchanan Mossel, Dror Weitz, and Nicholas Wormald. On the hardness of sampling independent sets beyond the tree threshold. *Probab. Theory Related Fields*, 143(3–4):401–439, 2009.
- [14] Thomas J. Schaefer. The complexity of satisfiability problems. In *Proceedings of the 10th Annual ACM Symposium on Theory of Computing, May 1–3, 1978, San Diego, California, USA*, pages 216–226, 1978.
- [15] Alistair Sinclair, Piyush Srivastava, and Marc Thurley. Approximation algorithms for two-state antiferromagnetic spin systems on bounded degree graphs. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17–19, 2012*, pages 941–953, 2012.
- [16] Allan Sly. Computational transition at the unique-

- ness threshold. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 287–296, 2010.
- [17] Allan Sly and Nike Sun. Counting in two-spin models on d -regular graphs. *Ann. Probab.*, 42(6):2383–2416, 2014.
- [18] Dror Weitz. Counting independent sets up to the tree threshold. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing, Seattle, WA, USA, May 21-23, 2006*, pages 140–149, 2006.
- [19] Yitong Yin and Jinman Zhao. Counting hypergraph matchings up to uniqueness threshold. *CoRR*, abs/1503.05812, 2015.