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

Andreas Galanis†   Leslie Ann Goldberg†

20 June 2016

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.

Keywords: Approximate counting, bounded-degree hypergraphs, 2-spin systems, counting constraint satisfaction.

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 functions of antiferromagnetic 2-spin systems on bounded-degree graphs [14, 21]. 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 [22, 20] (see also [10, 17] for related results) and later developed [19, 12, 21, 14] 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 $f : \{0,1\}^k \to \mathbb{R}_+$. 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 $\sigma: V \to \{0, 1\}$ induces a weight

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

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.

$$Z_{f;H} := \sum_{\sigma: V \to \{0,1\}} w_{f;H}(\sigma) = \sum_{\sigma: V \to \{0,1\}} \prod_{\{v_1, \ldots, v_k\} \in \mathcal{F}} f(\sigma(v_1), \ldots, \sigma(v_k)).$$

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

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 $f: \{0,1\}^k \to \{0,1\}$ where $f(s_1, \ldots, s_k) = 1$ iff at least one of $s_1, \ldots, s_k$ 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 $f(s_1, \ldots, s_k) = 1$ iff at most one of $s_1, \ldots, s_k$ 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 [22]; (ii) There is no FPRAS for the number of independent sets in graphs of maximum degree at most 6 [20] (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 [23] 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 [15] and was spelled out explicitly in [23], 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 [15, 23] 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 [23] show that natural gadgets cannot be used to show hardness for $\Delta = 4, 5$ and these cases remain open.

Generally, as the results of [15, 23] 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. *For $k \geq 2$, let* EASY(k) *be the set containing the following seven functions.*

$$f^{(k)}_{\text{zero}}(x_1, \ldots, x_k) = 0, \quad f^{(k)}_{\text{one}}(x_1, \ldots, x_k) = 1, \quad f^{(k)}_{\text{allzero}}(x_1, \ldots, x_k) = \mathbf{1}\{x_1 = \ldots = x_k = 0\},$$

$$f^{(k)}_{\text{allone}}(x_1, \ldots, x_k) = \mathbf{1}\{x_1 = \ldots = x_k = 1\}, \quad f^{(k)}_{\text{EQ}}(x_1, \ldots, x_k) = \mathbf{1}\{x_1 = \ldots = x_k\},$$

$$f^{(k)}_{\text{even}}(x_1, \ldots, x_k) = \mathbf{1}\{x_1 \oplus \cdots \oplus x_k = 0\}, \quad f^{(k)}_{\text{odd}}(x_1, \ldots, x_k) = \mathbf{1}\{x_1 \oplus \cdots \oplus x_k = 1\}.$$

Considering each of the functions in EASY(k), we obtain the following observation.

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

Note added in final version: Subsequent to this paper, the authors, together with Bezáková, Guo and Štefankovič [1], have studied the issue of a computational transition specifically for the problem of counting the weak independent sets of a hypergraph. They found further evidence, for this problem, that the complexity threshold, if it exists, may differ from the uniqueness threshold.

Particularly, for the weak independent set model, they gave an FPTAS which works even beyond the strong spatial mixing threshold and they showed inapproximability even below the uniqueness threshold.

Note, however that, while Theorem 3 does not guarantee the existence of a complexity threshold, it does at least show inapproximability if the degree bound is sufficiently large. That is, for every non-trivial f, it shows that for all sufficiently large Δ, it is NP-hard to approximate the partition function $Z_{f;H}$, given a k-uniform hypergraph H of maximum degree at most Δ.[1]

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 $\#\mathsf{CSP}(\Gamma)$ 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 $\mathcal{S}$[2] of constraints, where each constraint $C = (v_1, \ldots, v_k, f) \in \mathcal{S}$ constrains the variables $v_1, \ldots, v_k$ by applying a particular k-ary function $f \in \Gamma$. The value of the partition function is given by

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

The constraint $C = (v_1, \ldots, v_k, f)$ 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 $\#\mathsf{CSP}_\Delta(\Gamma)$ 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 $\#\mathsf{Hyper2Spin}(f, \Delta, c)$.

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

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 $\#\mathsf{Hyper2Spin}(f, \Delta, c)$ is closely related to the problem $\#\mathsf{CSP}_{\Delta,c}(\{f\})$. In particular, $\#\mathsf{Hyper2Spin}(f, \Delta, c)$ is the special case of $\#\mathsf{CSP}_{\Delta,c}(\{f\})$ in which constraints are not allowed to re-use variables. Thus, Theorem 3 has the following immediate corollary.

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

The combined results of [7] and [5] 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 [9] and bounded-degree approximate counting CSP [11] 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 [11] allows us to deduce (see Observation 27) that for every $\Delta \geq 6$ and $k \geq 2$ and every symmetric k-ary Boolean function $f \notin \mathsf{EASY}(k)$, there is no FPRAS for $\#\mathsf{CSP}(\{f, \delta_0, \delta_1\})$ unless NP = RP.[3] This hardness result extends from the uniqueness phase transition of the independent set

---

[1] It is an open question whether Theorem 3 continues to hold if the input H is further restricted to be a Δ-regular k-uniform hypergraph. Our result does not apply to this input restriction because our gadgets are not regular hypergraphs.

[2] The 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.

[3] We remark here that [11] also gave a partial classification for $\Delta = 3, 4, 5$, the remaining cases were (partially) resolved in [15].

model (which occurs at $\Delta = 6$) because pinning allows constructions which realise arbitrarily bad configurations.

The result of [11] 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 [3] 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 $f(s_1, \ldots, s_k) = 1$ iff at least one of $s_1, \ldots, s_k$ is 0) is not in $\mathsf{EASY}(k)$ for any $k \geq 2$, the result of Bordewich et al. [3] 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 3 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 Techniques

In order to prove Theorem 3, we will construct a k-uniform hypergraph H such that the spin-system induced by f on H induces an anti-ferromagentic 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 (Theorem 20) 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 [11], 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 [11], 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 3. 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 [11] to prove Theorem 3. The proof would even be straightforward if perfect reaslisations 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 [11]. 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 5. *Let $f : \{0, 1\}^k \to \mathbb{R}_+$ 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 positive ε'.

This can be done via standard powering (see the upcoming Lemma 12): 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 Lemma 12, the maximum degree of H' is proportional to $\log(1/\varepsilon')$. Nevertheless, the possibility of powering motivates the following definition.

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 $f \colon \{0,1\}^k \to \{0,1\}$ be the weak independent set function where $f(s_1, \ldots, s_k) = 1$ if and only if at least one of $s_1, \ldots, s_k$ 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 arbitrary polynomial precision using bounded-degree hypergraphs. Note that this example already shows that it is impossible to prove Theorem 3 by approximating the pinning functions δ_0 and δ_1 and then applying the result of [11].

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 13 and Lemma 14) 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 3 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) (Sections 4.1 and 4.3, respectively) 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 3.1) antiferromagnetic binary 2-spin models, showing (Lemma 22) 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 4.2), 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, which are explained in detail

6

in Section 4.2.1. 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 28). The proof for this class of self-dual functions f is given in Section 4.2.2.

For those self-dual functions f where the associated decision problem is not hard (Section 4.2.3), we show that one can realise approximate equality in the following sense.

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 exist distinct vertices $v_1, \ldots, v_t$ such that for each $s \in \{0, 1\}$,*

$$\mu_{f;H}(\sigma_{v_1} = \ldots = \sigma_{v_t} = s) \geq (1 - \varepsilon)/2.$$

We will refer to $v_1, \ldots, v_t$ as the terminals of H.

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 10 and 13, 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)).

1.3   Notation

We conclude this section with a piece of notation. Given a configuration σ and a subset $T \subseteq V$, we will use the notation σ_T to denote the restriction of σ to vertices in T. For a vertex $v \in V$, we will also use σ_v to denote the spin $\sigma(v)$ of vertex v in σ. Given a hyperedge $e \in \mathcal{F}$, we will denote by $H \setminus e$ the hypergraph $(V, \mathcal{F} \setminus e)$.

2   Properties of non-negative symmetric functions with domain $\{0, 1\}^k$

In this section, we study the concepts of pinning and equality that we will use for the proof of Theorem 3. 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 $f : \{0, 1\}^k \to \mathbb{R}_+$. Since f is symmetric, there are values $w_0, w_1, \ldots, w_k \in \mathbb{R}_+$ such that $f(x_1, \ldots, x_k) = w_\ell$ whenever $x_1 + \ldots + x_k = \ell$. We will refer to f and to the values w_i in the definitions and proofs in this section.

2.1   Pinning and equality

We start with the following remark, which follows from Defintion 7 (and makes the definition easier to apply).

Remark 9. *If H is an ε-realisation of t-equality and $v_1, \ldots, v_t$ are the terminals of H, then it also holds that $\mu_{f;H}(\sigma_{v_1} = \ldots = \sigma_{v_t} = s) \leq (1 + \varepsilon)/2$ for each $s \in \{0, 1\}$. Further, we have that $\mu_{f;H}(\exists i, j : \sigma_{v_i} \neq \sigma_{v_j}) \leq \varepsilon$.*

Next, we give a straightforward extension to the notion of supporting t-equality (see Definition 8).

Lemma 10. *Let $t \geq 2$ be an integer. The function f supports t-equality iff f supports 2-equality.*

Proof. It is immediate that if f supports t-equality for some $t \geq 2$ then it supports 2-equality (terminals $v_3, \ldots, v_t$ can simply be ignored).

We will now suppose that f supports 2-equality and show that it supports t-equality for a given $t > 2$. Consider $\varepsilon > 0$. Choose $\delta > 0$ to be sufficiently small (with respect to ε and t) so that $\delta \leq \varepsilon 2^{-(t+2)}$ and

$$\left(\frac{1-\delta}{1+\delta}\right)^{\binom{t}{2}} \geq \max\left\{\frac{1-\varepsilon/2}{1+\varepsilon/2}, \frac{1}{2}\right\}.$$

Suppose that H is a δ-realisation of 2-equality so it has terminals x and y so that for each $s \in \{0,1\}$

$$\frac{1-\delta}{2} \leq \mu_{f;H}(\sigma_x = \sigma_y = s) \leq \frac{1+\delta}{2} \text{ and } \mu_{f;H}(\sigma_x = s; \sigma_y = s \oplus 1) \leq \delta.$$

Let H' be the hypergraph constructed as follows. Let $T = \{v_1, \ldots, v_t\}$ be a set of t vertices which will be the terminals of H'. For each $1 \leq i < j \leq t$, let H_{ij} be a new copy of H but identify the terminal x of H_{ij} with v_i and the terminal y of H_{ij} with v_j. Let H' be the resulting hypergraph. Now for any $\tau : \{v_1, \ldots, v_t\} \to \{0,1\}$ that does *not* satisfy $\tau(v_1) = \cdots = \tau(v_t)$, the contribution to $Z_{f;H'}$ from configurations σ with $\sigma_T = \tau$ is at most

$$\delta Z_{f;H}\left(\left(\frac{1+\delta}{2}\right) Z_{f;H}\right)^{\binom{t}{2}-1} = \left(\frac{\delta}{1+\delta}\right)(1+\delta)^{\binom{t}{2}}\frac{Z_{f;H}^{\binom{t}{2}}}{2^{\binom{t}{2}-1}}$$

On the other hand, by considering the contribution from configurations σ with $\sigma(v_1) = \cdots = \sigma(v_t)$, we obtain that $Z_{f;H'}$ is at least

$$2\left(\left(\frac{1-\delta}{2}\right) Z_{f;H}\right)^{\binom{t}{2}} = (1-\delta)^{\binom{t}{2}}\frac{Z_{f;H}^{\binom{t}{2}}}{2^{\binom{t}{2}-1}},$$

Thus

$$\mu_{f;H'}(\sigma_T = \tau) \leq \frac{\left(\frac{\delta}{1+\delta}\right)(1+\delta)^{\binom{t}{2}}}{(1-\delta)^{\binom{t}{2}}} \leq \frac{2\delta}{1+\delta} \leq 2\delta$$

so, since $2^t(2\delta) \leq \varepsilon/2$, $\mu_{f;H'}(\exists i,j : \sigma_{v_i} \neq \sigma_{v_j}) \leq \varepsilon/2$. Furthermore, for any $s \in \{0,1\}$,

$$\frac{\mu_{f;H'}(\sigma(v_1) = \cdots = \sigma(v_t) = s)}{\mu_{f;H'}(\sigma(v_1) = \cdots = \sigma(v_t) = s \oplus 1)} \geq \left(\frac{1-\delta}{1+\delta}\right)^{\binom{t}{2}} \geq \frac{1-\varepsilon/2}{1+\varepsilon/2}.$$

It follows that

$$(1+\varepsilon/2)\mu_{f;H'}(\sigma(v_1) = \cdots = \sigma(v_t) = s) \geq (1-\varepsilon/2)\mu_{f;H'}(\sigma(v_1) = \cdots = \sigma(v_t) = s \oplus 1)$$
$$\geq (1-\varepsilon/2)(1 - \mu_{f;H'}(\sigma(v_1) = \cdots = \sigma(v_t) = s) - \varepsilon/2),$$

so $\mu_{f;H'}(\sigma(v_1) = \cdots = \sigma(v_t) = s) \geq (1-\varepsilon)/2$. Thus, H' is an ε-realisation of t-equality. $\qquad\square$

Lemma 10 motivates the following definition.

Definition 11. *A function f supports equality if, for some $t \geq 2$, it supports t-equality. (In this case, Lemma 10 shows that f supports t-equality for every $t \geq 2$.)*

The following lemma gives sufficient conditions for pinning-to-0, pinning-to-1 and 2-equality.

Lemma 12. *Let H be a hypergraph and $f : \{0,1\}^k \to \mathbb{R}_+$ be symmetric.*

1. *If there is a vertex v in H such that $\mu_{f;H}(\sigma_v = 0) > \mu_{f;H}(\sigma_v = 1)$, then f supports pinning-to-0.*

2. *If there is a vertex v in H such that $\mu_{f;H}(\sigma_v = 1) > \mu_{f;H}(\sigma_v = 0)$, then f supports pinning-to-1.*

3. *If there are vertices x, y in H such that $\mu_{f;H}(\sigma_x = \sigma_y = 0) = \mu_{f;H}(\sigma_x = \sigma_y = 1)$ and $\mu_{f;H}(\sigma_x = \sigma_y) > \mu_{f;H}(\sigma_x \neq \sigma_y)$, then f supports 2-equality.*
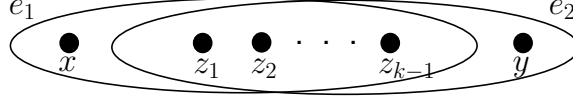
8

Figure 1: The hypergraph H used in the proof of Lemma 13. The hypergraph has two hyperedges e_1, e_2 such that $e_1 \cap e_2 = \{z_1, \ldots, z_{k-1}\}$, $e_1 \backslash e_2 = \{x\}$ and $e_2 \backslash e_1 = \{y\}$. We focus on the spins s_1, s_2 of the vertices x, y, respectively, i.e., for a configuration σ on H, $s_1 = \sigma_x$ and $s_2 = \sigma_y$.

Proof. We start with Item 3, which is the most difficult. Given $\varepsilon > 0$, we will use H to construct a hypergraph H' which is an ε-realisation of 2-equality. We start by constructing a hypergraph H'' with terminals v_1 and v_2. We construct H'' by taking two copies of H. In the first copy, we identify the vertex x with the terminal v_1 and the vertex y with the terminal v_2. The second copy is disjoint from the first one, except that we identify the vertex x of the second copy with the terminal v_2 and the vertex y of the second copy with the vertex v_1.

Let $\mu := \mu_{f;H}$. Then let $p = \mu(\sigma_x = \sigma_y = 0)^2 = \mu(\sigma_x = \sigma_y = 1)^2$ and $q = \mu(\sigma_x = 0, \sigma_y = 1)\mu(\sigma_x = 1, \sigma_y = 0)$.

If $q = 0$ then we can take H' to be H''. Then $\mu_{f;H'}(\sigma_{v_1} \neq \sigma_{v_2}) = 0$. However, $\mu_{f;H'}(\sigma_{v_1} = \sigma_{v_2} = 0) = \mu_{f;H'}(\sigma_{v_1} = \sigma_{v_2} = 1)$ so for $s \in \{0,1\}$, $\mu_{f;H'}(\sigma_{v_1} = \sigma_{v_2} = s) = 1/2$ and H' is a 0-realisation of 2-equality.

So suppose $q > 0$. Contruct H' by taking $r = 1 + \lceil \ln \varepsilon / \ln(q/p) \rceil$ disjoint copies of H'', identifying all terminals v_1 and all terminals v_2. Let $\mu' := \mu_{f;H'}$. We have

$$\mu'(\sigma_{v_1} = \sigma_{v_2} = 0) \propto p^r, \mu'(\sigma_{v_1} = \sigma_{v_2} = 1) \propto p^r, \mu'(\sigma_{v_1} = 0, \sigma_{v_2} = 1) = \mu'(\sigma_{v_1} = 1, \sigma_{v_2} = 0) \propto q^r.$$

Our choice of r ensures that $\mu'(\sigma_{v_1} = 0, \sigma_{v_2} = 1)/\mu'(\sigma_{v_1} = \sigma_{v_2} = 0) < \varepsilon$, so H' is an ε-realisation of 2-equality.

The proofs for Items 1 and 2 are similar but simpler. We will do Item 1. Let $p = \mu(\sigma_v = 0)$ and $q = \mu(\sigma_v = 1)$. Construct H' by taking $r = 1 + \lceil \ln \varepsilon / \ln(q/p) \rceil$ disjoint copies of H, identifying vertex v in all copies. Then $\mu_{f;H'}(\sigma_v = 1) = q^r/(q^r + p^r) \leq (q/p)^r \leq \varepsilon$ so H' is an ε-realisation of pinning-to-0. □

2.2 Classifying functions with respect to pinning and equality

The following lemma will be used in our classification.

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

Proof. 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 $\{x, y, z_1, \ldots, z_{k-1}\}$ and hyperedge set $\mathcal{F} = \{e_1, e_2\}$, where $e_1 = \{x, z_1, \ldots, z_{k-1}\}$ and $e_2 = \{y, z_1, \ldots, z_{k-1}\}$ (see Figure 1). Let $\mu := \mu_{f;H}$. We have that

$$\mu(\sigma_x = s_1, \sigma_y = s_2) \propto Z_{s_1 s_2} \text{ for } s_1, s_2 \in \{0,1\},$$

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 $\mu(\sigma_x = 0) \propto Z_{01} + Z_{00}$ and $\mu(\sigma_x = 1) \propto Z_{10} + Z_{11}$. Since $Z_{01} = Z_{10}$, $Z_{00} \neq Z_{11}$ would imply $\mu(\sigma_x = 0) \neq \mu(\sigma_x = 1)$, contradicting that f does not support pinning-to-0 or pinning-to-1 (by Lemma 12).

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

$$\left[\sum_{\ell=0}^{k-1} \binom{k-1}{\ell} w_\ell^2 \right] \left[\sum_{\ell=0}^{k-1} \binom{k-1}{\ell} w_{\ell+1}^2 \right] \geq \left[\sum_{\ell=0}^{k-1} \binom{k-1}{\ell} w_\ell w_{\ell+1} \right]^2 \tag{1}$$
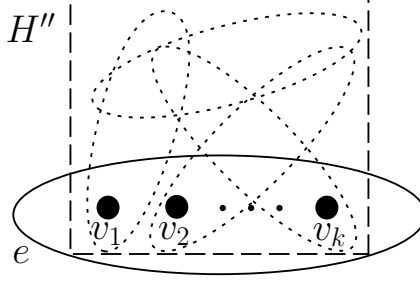
Figure 2: The hypergraph H' used to show that $w_0 = w_k$ in Lemma 14. It consists of a hypergraph H'' (dashed in the figure) which ε-realises k-equality among its terminals $v_1, \ldots, v_k$ and the hyperedge $e = \{v_1, \ldots, v_k\}$. We show that if $w_0 > w_k$, then H' realises pinning-to-0.

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 (1) holds only if there exists $\alpha \geq 0$ such that $w_{\ell+1} = \alpha w_\ell$ for every $\ell = 0, \ldots, k-1$, which yields $w_\ell = \alpha^\ell w_0$ for $\ell = 0, \ldots, k$. 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 Item 3 of Lemma 12 yields that f supports 2-equality. □

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 $\ell = 0, \ldots, k$.*

Proof. Let $k \geq 2$ and suppose that f supports 2-equality but neither pinning-to-0 nor pinning-to-1. Let H be the hypergraph with the vertex set $\{v_1, \ldots, v_k\}$ and the single hyperedge $e = \{v_1, \ldots, v_k\}$. We may assume that f is not a constant function since a constant function does not support 2-equality. Thus, at least one of the w_ℓ's is non-zero and hence $Z_{f;H} > 0$.

We will start by establishing the claim for $\ell = 0$, by showing that $w_0 = w_k$. Assume for contradiction that $w_0 \neq w_k$. W.l.o.g we may assume that $w_0 > w_k$ (if it is the other way around, then we can swap 0's and 1's in the following argument). Since f supports 2-equality, it also supports k-equality by Lemma 10. Choose $\varepsilon > 0$ sufficiently small so that

$$w_0 \left(\frac{1-\varepsilon}{2} \right) > w_k \left(\frac{1+\varepsilon}{2} \right) + 2^k \varepsilon.$$

Construct H' by taking H and a distinct ε-realisation H'' of k-equality and identifying the vertices of H with the terminals of H'' (see Figure 2). Then the contribution to $Z_{f,H'}$ from configurations with $\sigma(v_k) = 0$ is at least the contribution from all configurations that assign spin 0 to all terminals (giving a contribution of at least $w_0 \left(\frac{1-\varepsilon}{2} \right) Z_{f,H''}$ since H'' is an ε-realisation of k-equality) so we get

$$\mu_{f,H'}(\sigma(v_k) = 0) Z_{f,H'} \geq w_0 \left(\frac{1-\varepsilon}{2} \right) Z_{f,H''}.$$

Now consider the contribution to $Z_{f,H'}$ from configurations with $\sigma(v_k) = 1$. The contribution from configurations which map all terminals to spin 1 is at most $w_k \left(\frac{1+\varepsilon}{2} \right) Z_{f,H''}$ since H'' is an ε-realisation of k-equality. In addition, configurations which do not make the spins at the terminals equal contribute at most $2^k \varepsilon Z_{f,H''}$. So we get

$$\mu_{f;H'}(\sigma(v_k) = 1) Z_{f,H'} \leq w_k \left(\frac{1+\varepsilon}{2} \right) Z_{f,H''} + 2^k \varepsilon Z_{f;H''}.$$

It follows by the choice of ε that $\mu_{f;H'}(\sigma(v_k) = 0) > \mu_{f,H'}(\sigma(v_k) = 1)$, so Lemma 12 shows that f supports pinning-to-0, contrary to the statement of the lemma. Thus, we have shown that $w_0 = w_k$.

For ease of notation, we let $Z_{0,k} = w_0$ and $Z_{1,k} = w_k$. Then, for $t \in \{1, \ldots, k-1\}$ let $V_t = \{v_1, \ldots, v_t\}$. For $s \in \{0, 1\}$ let $Z'_{s,t} = \sum_{\ell=0}^{k-1-t} \binom{k-1-t}{\ell} w_{\ell+s}$ and $Z''_{s,t} = \sum_{\ell=0}^{k-1-t} \binom{k-1-t}{\ell} w_{\ell+s+t}$. Let $Z_{s,t} = Z'_{s,t} + Z''_{s,t}$. We will establish the following system of equalities.

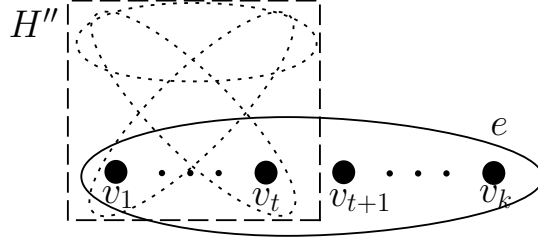$$\text{For } t \in \{1, \ldots, k\}, Z_{0,t} = Z_{1,t}. \tag{2}$$

Figure 3: The hypergraph H' used to show that $Z_{0,t} = Z_{1,t}$ for $t \in \{2, \ldots, k-1\}$. It consists of a hypergraph H'' (dashed in the figure) which ε-realises t-equality among its terminals $v_1, \ldots, v_t$ and the hyperedge $e = \{v_1, \ldots, v_t, v_{t+1}, \ldots, v_k\}$. We show that if $Z_{0,t} > Z_{1,t}$, then H' realises pinning-to-0.

We have already dealt with the case $t = k$. Next, consider $t = 1$. Note that $\mu_{f;H}(\sigma(v_k) = 0) = Z_{0,1}/Z_{f;H}$ and $\mu_{f;H}(\sigma(v_k) = 1) = Z_{1,1}/Z_{f,H}$ (to see these, observe that $Z'_{s,1}$ is the contribution from configurations where, say, vertex v_1 has spin 0 and v_k has spin s; similarly, $Z''_{s,1}$ is the contribution where v_1 has spin 1 and v_k has spin s). Thus, if $Z_{0,1} \neq Z_{1,1}$ then $\mu_{f;H}(\sigma(v_k) = 0) \neq \mu_{f;H}(\sigma(v_k) = 1)$ so Lemma 12 shows that f supports pinning-to-0 or pinning-to-1, contrary to the statement of the lemma.

We will now show that Equation (2) holds for $t \in \{2, \ldots, k-1\}$. The proof is similar to the proof that $w_0 = w_k$ above. Consider some $t \in \{2, \ldots, k-1\}$, and suppose for contradiction that $Z_{0,t} > Z_{1,t}$. Let $\rho = Z_{1,t} \geq 0$ and $\delta = Z_{0,t} - Z_{1,t} > 0$. Again, since f supports 2-equality, it also supports t-equality by Lemma 10. Choose ε sufficiently small so that

$$(\rho + \delta)\left(\frac{1 - \varepsilon}{2}\right) > \rho\left(\frac{1 + \varepsilon}{2}\right) + 2^k \varepsilon$$

Construct H' by taking H and a distinct ε-realisation H'' of t-equality and identifying the terminals $v_1, \ldots, v_t$ in H and H'' (see Figure 3). Note that vertex v_k is not a vertex of H''. So the contribution to $Z_{f,H'}$ from configurations with $\sigma(v_k) = 0$ is at least the contribution from such configurations which also satisfy $\sigma(v_1) = \cdots = \sigma(v_t) = 0$ (giving a contribution of at least $Z'_{0,t}\left(\frac{1-\varepsilon}{2}\right)Z_{f,H''}$) and a similar contribution from configurations which also satisfy $\sigma(v_1) = \cdots = \sigma(v_t) = 1$ (giving a contribution of at least $Z''_{0,t}\left(\frac{1-\varepsilon}{2}\right)Z_{f,H''}$) so we get

$$\mu_{f,H'}(\sigma(v_k) = 0)Z_{f,H'} \geq Z'_{0,t}\left(\tfrac{1-\varepsilon}{2}\right)Z_{f,H''} + Z''_{0,t}\left(\tfrac{1-\varepsilon}{2}\right)Z_{f,H''} = Z_{0,t}\left(\tfrac{1-\varepsilon}{2}\right)Z_{f,H''}$$
$$\geq (\rho + \delta)\left(\tfrac{1-\varepsilon}{2}\right)Z_{f,H''}.$$

Now consider the contribution to $Z_{f,H'}$ from configurations with $\sigma(v_k) = 1$. This is at most the contributions which also satisfy $\sigma(v_1) = \cdots = \sigma(v_t) = 0$ (giving a contribution at most $Z'_{1,t}\left(\frac{1+\varepsilon}{2}\right)Z_{f,H''}$) and a similar term $Z''_{1,t}\left(\frac{1+\varepsilon}{2}\right)Z_{f,H''}$ from the contributions which also satisfy $\sigma(v_1) = \cdots = \sigma(v_t) = 1$. In addition, configurations which do not satisfy $\sigma(v_1) = \cdots = \sigma(v_t)$ contribute at most $2^k \varepsilon Z_{f,H''}$. So we get

$$\mu_{f;H'}(\sigma(v_k) = 1)Z_{f,H'} \leq \rho\left(\tfrac{1+\varepsilon}{2}\right)Z_{f,H''} + 2^k \varepsilon Z_{f;H''}.$$

Again by the choice of ε, we have $\mu_{f,H'}(\sigma(v_k) = 0) > \mu_{f,H'}(\sigma(v_k) = 1)$, so Lemma 12 shows that f supports pinning-to-0, contrary to the statement of the lemma. So we have now established the system of equations (2).

We view (2) as a system of equations over the $k + 1$ (real) variables $w_0, \ldots, w_k$. We refer to the equation $Z_{0,t} = Z_{1,t}$ as "Equation t" or "the t'th equation". In the remainder of the proof we will show that the solutions to this system of equations are exactly the assignments of real values to the variables satisfying

$$w_\ell = w_{k-\ell} \text{ for } \ell = 0, \ldots, k. \tag{3}$$

We first show that any solution satisfying (3) satisfies (2). The case $t = k$ is obvious. For $t \in \{1, \ldots, k-1\}$, we will show that any solution satisfying (3) satisfies $Z'_{0,t} = Z''_{1,t}$ and $Z'_{1,t} = Z''_{0,t}$.

11

To see the first of these, consider the coefficient $\binom{k-1-t}{\ell'}$ of $w_{\ell'+1+t}$ in $Z''_{1,t}$. Now let $\ell = k - (\ell' + 1 + t)$ so $w_\ell = w_{\ell'+1+t}$. The coefficient of w_ℓ in $Z'_{0,t}$ is $\binom{k-1-t}{\ell} = \binom{k-1-t}{k-1-t-\ell} = \binom{k-1-t}{\ell'}$. This establishes $Z'_{0,t} = Z''_{1,t}$. The proof that $Z'_{1,t} = Z''_{0,t}$ is similar. We will now show that the solution in (3) has the appropriate dimension, so there are no other solutions of (2). It is simplest to split the calculation into two cases, depending on the parity of k.

First, suppose $k = 2r + 1$ is odd. Consider $1 \le t \le t' \le k - 1$. We will compute the coefficients of the variable w_t in the quantities $Z'_{0,t'}$, $Z''_{0,t'}$, $Z''_{1,t'}$ and $Z''_{1,t'}$ using the convention that $\binom{a}{b} = 0$ if $b \notin \{0, \ldots, a\}$. These are $\binom{2r-t'}{t}$, $\binom{2r-t'}{t-t'}$, $\binom{2r-t'}{t-1}$ and $\binom{2r-t'}{t-t'-1}$, respectively. Consider the case where $t \ge r + 1$. Then $2r - t' < t - 1 < t$ so these can be simplified to 0, $\binom{2r-t'}{t-t'}$, 0 and $\binom{2r-t'}{t-t'-1}$, respectively. If $t' > t$ then all four coefficients are 0 so w_t is not in Equation t'. If $t' = t$ then the final coefficient is 0, but the 2nd of these coefficients is 1, so w_t has a non-zero coefficient in Equation t. We conclude for $t \ge r + 1$ that w_t is not in Equations $t+1, \ldots, k-1$, but it has a non-zero coefficient in Equation t.

Thus, Equations $k, k-1, \ldots, r+1$ give us $k - r = r + 1$ linearly independent equations. Thus the solution space of the system over our $k + 1$ variables has dimension at most $k + 1 - (r+1) = r + 1$. But we have already shown that Equation (3) gives a solution, and the dimension of this solution is $r + 1$.

Similarly, suppose that $k = 2r$ is even. Once again consider $1 \le t \le t' \le k - 1$ and compute the coefficients of the variable w_t in the quantities $Z'_{0,t'}$, $Z''_{0,t'}$, $Z''_{1,t'}$ and $Z''_{1,t'}$. These are $\binom{2r-1-t'}{t}$, $\binom{2r-1-t'}{t-t'}$, $\binom{2r-1-t'}{t-1}$ and $\binom{2r-1-t'}{t-t'-1}$, respectively. Consider the case where $t \ge r + 1$. Then $2r - 1 - t' < t - 1 < t$ so these can again be simplified to 0, $\binom{2r-1-t'}{t-t'}$, 0 and $\binom{2r-1-t'}{t-t'-1}$, respectively. Once again, if $t' > t$ then all four coefficients are 0 so w_t is not in Equation t'. However, if $t' = t$ then the final coefficient is 0, but the 2nd of these coefficients is 1, so w_t has a non-zero coefficient in Equation t. As before, we conclude that the equations $k, k-1, \ldots, r+1$ give us r linearly independent equations, so the solution space of the system has dimension at most $(k+1) - (r) = r + 1$. But $w_\ell = w_{k-\ell}$ is a solution whose dimension is $r + 1$. □

2.3 Realising conditional distributions induced by pinning and equality

We will use pinnings or equality to construct and analyze gadgets; the upcoming Lemma 15 is a first step in doing this effortlessly. It asserts that when f supports one of the properties pinning-to-0, pinning-to-1 or equality we can consider appropriate conditional distributions (depending on the property); these conditional distributions can then be realised using appropriately constructed hypergraphs.

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).

Lemma 15. *Let $f : \{0,1\}^k \to \mathbb{R}_+$ be symmetric. Let $H = (V, \mathcal{F})$ be a hypergraph and $S \subseteq V$. Let $\varepsilon > 0$.*

1. *Suppose that f supports pinning-to-s for some $s \in \{0, 1\}$. Suppose that $u \in V$ satisfies $\mu_{f;H}(\sigma_u = s) > 0$. Then there is a hypergraph $H' = (V', \mathcal{F}')$ with $V \subseteq V'$ and $\mathcal{F} \subseteq \mathcal{F}'$ such that for every $\tau : S \to \{0, 1\}$ it holds that*

$$\left| \mu_{f;H'}(\sigma_S = \tau) - \mu_{f;H}(\sigma_S = \tau \mid \sigma_u = s) \right| \le \varepsilon.$$

2. *Suppose that f supports equality and that $R \subseteq V$ satisfies $\mu_{f;H}(\sigma_R^{\mathsf{eq}}) > 0$. Then there is a hypergraph $H' = (V', \mathcal{F}')$ with $V \subseteq V'$ and $\mathcal{F} \subseteq \mathcal{F}'$ such that for every $\tau : S \to \{0, 1\}$ it holds that*

$$\left| \mu_{f;H'}(\sigma_S = \tau) - \mu_{f;H}(\sigma_S = \tau \mid \sigma_R^{\mathsf{eq}}) \right| \le \varepsilon,$$
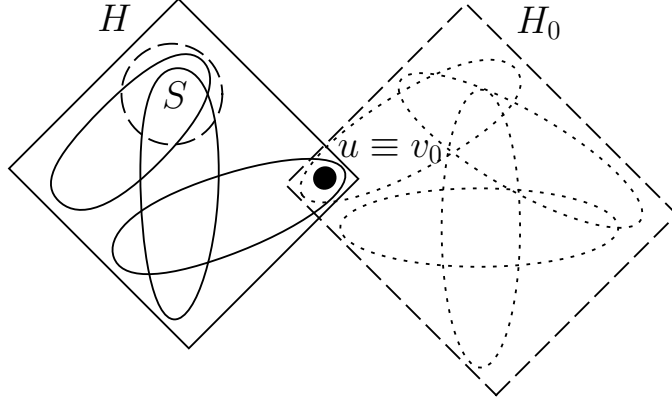
Figure 4: The hypergraph H' used in the proof of Item 1 in Lemma 15. It consists of the hypergraph H and a hypergraph H_0 with terminal v_0 which ε'-realises pinning-to-0. We identify the vertex u of H with the terminal vertex v_0 of H_0 keeping otherwise the two hypergraphs H, H_0 disjoint. Then, for any subset S of vertices in H (which may in general include u as well), the distribution $\mu_{f;H'}(\sigma_S = \cdot)$ is well approximated by the conditional distribution $\mu_{f,H}(\sigma_S = \cdot \mid \sigma_u = 0)$.

Proof. Denote $\mu_{f;H}$ by μ. We begin with the proof of Item 1. We will take $s = 0$. The proof for $s = 1$ is similar, swapping 0's and 1's. Let $\varepsilon' > 0$ be sufficiently small so that

$$\left(\frac{\varepsilon'}{1 - \varepsilon'}\right)\left(\frac{\mu(\sigma_u = 1)}{\mu(\sigma_u = 0)}\right) \leq \frac{\varepsilon}{1 - \varepsilon}. \tag{4}$$

Since f supports pinning-to-0, there exists a hypergraph H_0 with vertex set V_0 which ε'-realises pinning-to-0. Denote by v_0 the terminal of H_0. For $i \in \{0, 1\}$, let $p_i = \mu_{f;H_0}(\sigma_{v_0} = i)$, so that $p_0 \geq 1 - \varepsilon'$, $p_1 \leq \varepsilon'$. To construct the hypergraph H' (see Figure 4), take a (distinct) copy of H_0 and identify v_0 with u (the copy of H_0 is otherwise disjoint from the rest of H). Let $\mu' := \mu_{f;H'}$. Note that $\mu'(\sigma_u = s) \propto p_s \mu(\sigma_u = s)$ for $s \in \{0, 1\}$, so that

$$\frac{\mu'(\sigma_u = 1)}{\mu'(\sigma_u = 0)} = \left(\frac{p_1}{p_0}\right)\left(\frac{\mu(\sigma_u = 1)}{\mu(\sigma_u = 0)}\right).$$

Equation (4) yields $\mu'(\sigma_u = 1) \leq \varepsilon$.

Let $\tau : S \to \{0, 1\}$. Note that

$$\mu'(\sigma_S = \tau) = \mu'(\sigma_u = 0)\mu'(\sigma_S = \tau \mid \sigma_u = 0) + \mu'(\sigma_S = \tau, \sigma_u = 1)$$
$$= \mu'(\sigma_u = 0)\,\mu(\sigma_S = \tau \mid \sigma_u = 0) + \mu'(\sigma_S = \tau, \sigma_u = 1),$$

where in the second equality we used that conditioned on the spin of u, σ_S is independent of the configuration $\sigma_{V_0 \setminus \{u\}}$. It follows that

$$|\mu'(\sigma_S = \tau) - \mu(\sigma_S = \tau \mid \sigma_u = 0)| \leq \mu'(\sigma_u = 1) \leq \varepsilon$$

completing the proof of Item 1.

The proof of Item 2 is completely analogous, though slightly more technical. Let $r = |R|$. As before, let $\varepsilon' > 0$ be sufficiently small, to be picked later. Since f supports equality, there exists a hypergraph H_1 with vertex set V_1 which ε'-realises r-equality. Denote by $T = \{v_1, \dots, v_r\}$ the set of terminals of H_1. For $\eta : T \to \{0, 1\}$, let $p_\eta := \mu_{f;H_1}(\sigma_R = \eta)$. Since H_1 is an ε'-realisation of r-equality, we have that

$$(1 - \varepsilon')/2 \leq p_{\mathbf{0}}, p_{\mathbf{1}} \leq (1 + \varepsilon')/2, \quad \sum_{\eta \neq \mathbf{0}, \mathbf{1}} p_\eta \leq \varepsilon'.$$

13

To construct the hypergraph H', take a (distinct) copy of H_1 and identify (in an arbitrary way) vertices in the set R with terminals in the set T (the copy of H_1 is otherwise disjoint from the rest of H), we will keep the notation R for the merged vertices in H' (see Figure 5). Let $\mu' := \mu_{f;H'}$. For $\eta : R \to \{0,1\}$, we $\mu'(\sigma_R = \eta) \propto p_\eta \mu(\sigma_R = \eta)$, so that

$$\mu'(\sigma_R = \eta) = \frac{p_\eta \mu(\sigma_R = \eta)}{p_{\mathbf{0}} \mu(\sigma_R = \mathbf{0}) + p_{\mathbf{1}} \mu(\sigma_R = \mathbf{1}) + \sum_{\eta' \neq \mathbf{0}, \mathbf{1}} p_{\eta'} \mu(\sigma_R = \eta')}.$$

For $\mathbf{s} \in \{\mathbf{0}, \mathbf{1}\}$, we use the upper and lower bounds on $p_{\mathbf{s}}$ to obtain

$$\mu'(\sigma_R = \mathbf{s}) \leq \frac{p_{\mathbf{s}} \mu(\sigma_R = \mathbf{s})}{p_{\mathbf{0}} \mu(\sigma_R = \mathbf{0}) + p_{\mathbf{1}} \mu(\sigma_R = \mathbf{1})} \leq \left(\frac{1 + \varepsilon'}{1 - \varepsilon'}\right) \left(\frac{\mu(\sigma_R = \mathbf{s})}{\mu(\sigma_R = \mathbf{0}) + \mu(\sigma_R = \mathbf{1})}\right), \tag{5}$$

$$\mu'(\sigma_R = \mathbf{s}) \geq \frac{p_{\mathbf{s}} \mu(\sigma_R = \mathbf{s})}{p_{\mathbf{0}} \mu(\sigma_R = \mathbf{0}) + p_{\mathbf{1}} \mu(\sigma_R = \mathbf{1}) + \varepsilon'} \geq \left(\frac{1 - \varepsilon'}{1 + \varepsilon'}\right) \left(\frac{\mu(\sigma_R = \mathbf{s})}{\mu(\sigma_R = \mathbf{0}) + \mu(\sigma_R = \mathbf{1}) + \frac{2\varepsilon'}{1 + \varepsilon'}}\right). \tag{6}$$

From (5) and (6), we obtain that for $\mathbf{s} \in \{\mathbf{0}, \mathbf{1}\}$, as $\varepsilon' \downarrow 0$, it holds that

$$\mu'(\sigma_R = \mathbf{s}) \to \frac{\mu(\sigma_R = \mathbf{s})}{\mu(\sigma_R = \mathbf{0}) + \mu(\sigma_R = \mathbf{1})} = \mu(\sigma_R = \mathbf{s} \mid \sigma_R^{\mathsf{eq}}), \quad \text{and thus } \mu'(\neg \sigma_R^{\mathsf{eq}}) \to 0, \tag{7}$$

where the latter limit follows by observing that $\mu'(\neg \sigma_R^{\mathsf{eq}}) = 1 - \mu'(\sigma_R = \mathbf{0}) - \mu'(\sigma_R = \mathbf{1})$.

Let $\tau : S \to \{0,1\}$. Note that

$$\begin{aligned}
\mu'(\sigma_S = \tau) &= \mu'(\sigma_R = \mathbf{0})\mu'(\sigma_S = \tau \mid \sigma_R = \mathbf{0}) + \\
&\qquad \mu'(\sigma_R = \mathbf{1})\mu'(\sigma_S = \tau \mid \sigma_R = \mathbf{1}) + \mu'(\sigma_S = \tau, \neg\sigma_R^{\mathsf{eq}}) \\
&= \mu'(\sigma_R = \mathbf{0})\mu(\sigma_S = \tau \mid \sigma_R = \mathbf{0}) + \\
&\qquad \mu'(\sigma_R = \mathbf{1})\mu(\sigma_S = \tau \mid \sigma_R = \mathbf{1}) + \mu'(\sigma_S = \tau, \neg\sigma_R^{\mathsf{eq}}),
\end{aligned} \tag{8}$$

where again in the second equality we used that conditioned on σ_R, σ_S is independent of the configuration $\sigma_{V_1 \setminus R}$. Using the limits in (7) and the equality (8), it is not hard to see that as $\varepsilon' \downarrow 0$, it holds that

$$\mu'(\sigma_S = \tau) \to \mu(\sigma_R = \mathbf{0} \mid \sigma_R^{\mathsf{eq}}) \mu(\sigma_S = \tau \mid \sigma_R = \mathbf{0}) + \mu(\sigma_R = \mathbf{1} \mid \sigma_R^{\mathsf{eq}}) \mu(\sigma_S = \tau \mid \sigma_R = \mathbf{1}). \tag{9}$$

The right-hand side in (9) is equal to $\mu(\sigma_S = \tau \mid \sigma_R^{\mathsf{eq}})$, from where it follows that by choosing small ε', the hypergraph H' satisfies $\left| \mu'(\sigma_S = \tau) - \mu\left(\sigma_S = \tau \mid \sigma_R^{\mathsf{eq}}\right) \right| \leq \varepsilon$, as wanted.

This concludes the proof of Lemma 15. $\qquad\square$

Typically, when f satisfies, say, pinning-to-0 we will be interested in "pinning" more than one vertex to zero, while Item 1 from Lemma 15 accomodates only one vertex. We will also be interested in cases where f satisfies multiple properties amongst pinning-to-0, pinning-to-1 and equality. We extend Lemma 15 to address this more general framework. Prior to that, it will be useful for our applications to set up a convenient notation.

Definition 16. *Let $f : \{0,1\}^k \to \mathbb{R}_+$ be symmetric. Let $H = (V, \mathcal{F})$ be a hypergraph and assume that $\mathcal{V} := (V_0, V_1, V_2, \ldots, V_r)$ is a labelled collection of disjoint subsets of V such that: (i) $V_0 = \emptyset$ if f does not support pinning-to-0, (ii) $V_1 = \emptyset$ if f does not support pinning-to-1, (iii) $V_2 = \ldots = V_r = \emptyset$ if f does not support equality, (iv) it holds that $\mu_{f;H}(\sigma_{V_0} = \mathbf{0}, \sigma_{V_1} = \mathbf{1}, \sigma_{V_2}^{\mathsf{eq}}, \ldots, \sigma_{V_r}^{\mathsf{eq}}) > 0$. We will then say that $\mathcal{V}$ is* admissible *for the hypergraph H (with respect to f) and denote by $\mu_{f;H}^{\mathrm{cond}(\mathcal{V})}$ the probability distribution $\mu_{f;H}(\cdot \mid \sigma_{V_0} = \mathbf{0}, \sigma_{V_1} = \mathbf{1}, \sigma_{V_2}^{\mathsf{eq}}, \ldots, \sigma_{V_r}^{\mathsf{eq}})$.*

Definition 16 provides the framework to prove the following generalization of Lemma 15. In simple words, we will now be able to combine the conditional distributions that we can realize via hypergraphs using f.
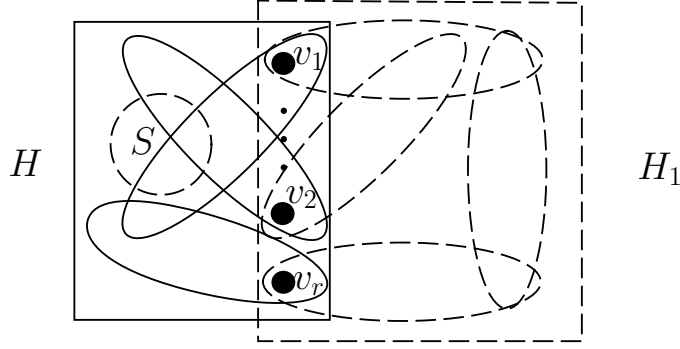
Figure 5: The hypergraph H' used in the proof of Item 2 in Lemma 15. It consists of the hypergraph H and a hypergraph H_1 which realises r-equality among its terminals $v_1, \ldots, v_r$. We identify the vertices in R with the terminal vertices $v_1, \ldots, v_r$ of H_1 keeping otherwise the two hypergraphs H, H_1 disjoint. Then, for any subset S of vertices in H (which may in general include any of the vertices $v_1, \ldots, v_r$), the distribution $\mu_{f;H'}(\sigma_S = \cdot)$ is well approximated by the conditional distribution $\mu_{f,H}(\sigma_S = \cdot \mid \sigma_R^{\mathsf{eq}})$.

Lemma 17. *Let $f : \{0,1\}^k \to \mathbb{R}_+$ be a symmetric function. Let H be a hypergraph with vertex set V and let S be a subset of V. Let $\mathcal{V}$ an admissible collection of subsets of V with respect to H. Then, for every $\varepsilon > 0$, there is a hypergraph $H' = (V', \mathcal{F}')$ with $V \subseteq V'$ and $\mathcal{F} \subseteq \mathcal{F}'$ such that, for every $\tau : S \to \{0,1\}$, it holds that*

$$\left| \mu_{f;H'}(\sigma_S = \tau) - \mu_{f;H}^{\mathrm{cond}(\mathcal{V})}(\sigma_S = \tau) \right| \leq \varepsilon,$$

where $\mu_{f;H}^{\mathrm{cond}(\mathcal{V})}(\cdot)$ is as in Definition 16.

Proof. To treat the different cases which arise as uniformly as possible, it will be convenient to use the following notation for the purposes of this proof. Let $\sigma : V \to \{0,1\}$ and $X \subseteq V$. For $\mathsf{s} \in \{\mathbf{0}, \mathbf{1}\}$, we will write σ_X^{s} as an alternative notation to $\sigma_X = \mathsf{s}$.

We are now set to prove the claim. Suppose that $V = (V_0, V_1, V_2, \ldots, V_r)$. We proceed by induction on $W := |V_0| + |V_1| + (r - 1)$. The base case $W = 0$ is trivial — we can take $H' = H$. To carry out the induction step, assume that the claim holds when $W = t$, we show it when $W = t + 1$. We have that at least one of the following holds: (i) $|V_0| \geq 1$, (ii) $|V_1| \geq 1$, (iii) $r \geq 2$. For each of these cases, we will have a collection of sets $\overline{\mathcal{V}}$ for which we will invoke the inductive hypothesis and a set X which we wish to add to $\overline{\mathcal{V}}$ to conclude the claim for $\mathcal{V}$. Also, we want to be able to condition on the event σ_X^{s} for some $\mathsf{s} \in \{\mathbf{0}, \mathbf{1}, \mathsf{eq}\}$ (the value of s will depend on the case that we consider). More precisely, the set X, the value of $\mathsf{s} \in \{\mathbf{0}, \mathbf{1}, \mathsf{eq}\}$ and the collection of sets $\overline{\mathcal{V}}$ are specified as follows for the respective cases:

- $X = \{v\}$ for some $v \in V_0$, $\mathsf{s} = \mathbf{0}$, and $\overline{\mathcal{V}} = (V_0 \setminus X, V_1, V_2, \ldots, V_r)$,
- $X = \{v\}$ for some $v \in V_1$, $\mathsf{s} = \mathbf{1}$, and $\overline{\mathcal{V}} = (V_0, V_1 \setminus X, V_2, \ldots, V_r)$,
- $X = V_r$, $\mathsf{s} = \mathsf{eq}$, and $\overline{\mathcal{V}} = (V_0, V_1, V_2, \ldots, V_{r-1})$.

Note that since $\mathcal{V}$ is admissible for H under f, the same is trivially true for $\overline{\mathcal{V}}$ as well.

Let S' be an arbitrary subset of V (which may be our target set S). By induction, for the collection of subsets $\overline{\mathcal{V}}$, we have that for every $\varepsilon > 0$, there exists a hypergraph H' with $V \subseteq V'$ and $\mathcal{F} \subseteq \mathcal{F}'$ such that, for every $\tau' : S' \to \{0,1\}$, it holds that

$$\left| \mu_{f;H'}(\sigma_{S'} = \tau') - \mu_{f;H}^{\mathrm{cond}(\overline{\mathcal{V}})}(\sigma_{S'} = \tau') \right| \leq \varepsilon, \tag{10}$$

Since $\mu(\sigma_{V_0}^0, \sigma_{V_1}^1, \sigma_{V_2}^{\mathsf{eq}}, \ldots, \sigma_{V_r}^{\mathsf{eq}}) > 0$ (by the assumption that $\mathcal{V}$ is admissible), we also have that $M := \mu_{f;H}^{\mathrm{cond}(\overline{\mathcal{V}})}(\sigma_X^{\mathsf{s}}) > 0$, where $X, \mathsf{s}, \overline{\mathcal{V}}$ were defined above. Let

$$\varepsilon_1 := \varepsilon/2, \quad \varepsilon_2 := \varepsilon_1 M^2/4, \quad \varepsilon_3 := \varepsilon_2/2^{|S|}.$$

15

Now consider $\tau : S \to \{0,1\}$. if $\mu_{f;H}(\sigma_S = \tau, \sigma_X^{\mathbf{s}}) = 0$ then $\mu_{f;H}^{\mathrm{cond}(\overline{\mathcal{V}})}(\sigma_S = \tau, \sigma_X^{\mathbf{s}}) = 0$. Also, since $V \subseteq V'$ and $\mathcal{F} \subseteq \mathcal{F}'$. $\mu_{f;H'}(\sigma_S = \tau, \sigma_X^{\mathbf{s}}) = 0$.

Otherwise, $\mu_{f;H}(\sigma_S = \tau, \sigma_X^{\mathbf{s}}) > 0$. In this case, applying (10) with $S' = S \cup X$ we find that there is a hypergraph $H' = (V', \mathcal{F}')$ with vertex set $V \subseteq V'$ and $\mathcal{F} \subseteq \mathcal{F}'$ such that for each assignment $\tau' : S' \to \{0,1\}$ that is consistent with $\tau'_S = \tau$ and $\tau'^{\mathbf{s}}_X$, we have

$$\left| \mu_{f;H'}(\sigma_{S'} = \tau') - \mu_{f;H}^{\mathrm{cond}(\overline{\mathcal{V}})}(\sigma_{S'} = \tau') \right| \le \varepsilon.$$

Summing over all $\tau' : S' \to \{0,1\}$ that satisfy both $\tau'_S = \tau$ and $\tau'^{\mathbf{s}}_X$, we conclude that

$$\left| \mu_{f;H'}(\sigma_S = \tau, \sigma_X^{\mathbf{s}}) - \mu_{f;H}^{\mathrm{cond}(\overline{\mathcal{V}})}(\sigma_S = \tau, \sigma_X^{\mathbf{s}}) \right| \le \varepsilon_3. \tag{11}$$

So now we have established (11) for all $\tau : S \to \{0,1\}$ (whether $\mu_{f;H}(\sigma_S = \tau, \sigma_X^{\mathbf{x}})$ is 0 or not). Summing over $\tau : S \to \{0,1\}$, we obtain

$$\left| \mu_{f;H'}(\sigma_X^{\mathbf{s}}) - \mu_{f;H}^{\mathrm{cond}(\overline{\mathcal{V}})}(\sigma_X^{\mathbf{s}}) \right| \le 2^{|S|} \varepsilon_3 = \varepsilon_2. \tag{12}$$

This implies $\mu_{f;H'}(\sigma_X^{\mathbf{s}}) \ge M - \varepsilon_2 > M/2 > 0$.

Now to simplify the notation define $A := \mu_{f;H}^{\mathrm{cond}(\overline{\mathcal{V}})}(\sigma_S = \tau, \sigma_X^{\mathbf{s}})$. Now note that

$$\mu_{f;H}^{\mathrm{cond}(\mathcal{V})}(\sigma_S = \tau) = A/M.$$

Let $A' = \mu_{f;H'}(\sigma_S = \tau, \sigma_X^{\mathbf{s}})$ and let $M' = \mu_{f;H'}(\sigma_X^{\mathbf{s}})$. Equation (11) (together with $\varepsilon_3 \le \varepsilon_2$) shows that $A' - \varepsilon_2 \le A \le A' + \varepsilon_2$. Also, Equation (12) shows $M' - \varepsilon_2 \le M \le M' + \varepsilon_2$. Then we will use the bound

$$\max\left(\frac{A + \varepsilon_2}{M - \varepsilon_2} - \frac{A}{M}, \ \frac{A}{M} - \frac{A - \varepsilon_2}{M + \varepsilon_2} \right) \le \frac{\varepsilon_2(M + A)}{M(M - \varepsilon_2)} \le \frac{2\varepsilon_2}{M(M - \varepsilon_2)} \le 4\varepsilon_2/M^2 = \varepsilon_1$$

to conclude that

$$\left| \frac{A'}{M'} - \frac{A}{M} \right| = \left| \mu_{f;H'}(\sigma_S = \tau \mid \sigma_X^{\mathbf{s}}) - \mu_{f;H}^{\mathrm{cond}(\mathcal{V})}(\sigma_S = \tau) \right| \le \varepsilon_1. \tag{13}$$

We will now apply Lemma 15 to H' and S with error parameter ε_1. To apply the lemma, we need the fact that f supports pinning-to-s if $\mathbf{s}$ is $\mathbf{0}$ or $\mathbf{1}$. If $\mathbf{s} = \mathsf{eq}$ then we need the fact that f supports equality. Both of these follow from the admissibility of $\mathcal{V}$ and the construction of X and $\mathbf{s}$. We also need the fact that $\mu_{f;H'}(\sigma_X^{\mathbf{s}}) > 0$, which we have established above. Then Lemma 15 shows that there exists a hypergraph $H'' = (V'', \mathcal{F}'')$ with $V' \subseteq V''$ and $\mathcal{F}' \subseteq \mathcal{F}''$ such that for every $\tau : S \to \{0,1\}$, it holds that

$$\left| \mu_{f;H''}(\sigma_S = \tau) - \mu_{f;H'}(\sigma_S = \tau \mid \sigma_X^{\mathbf{s}}) \right| \le \varepsilon_1. \tag{14}$$

It follows by (13) and (14) that

$$\left| \mu_{f;H''}(\sigma_S = \tau) - \mu_{f;H}^{\mathrm{cond}(\mathcal{V})}(\sigma_S = \tau) \right| \le 2\varepsilon_1 = \varepsilon.$$

This completes the induction. □

3 A general inapproximability lemma

The purpose of this section is to prove the following lemma, which will allow us to exploit our study of pinning-to-0, pinning-to-1 and equality.

Lemma 18. *Let $f : \{0,1\}^k \to \mathbb{R}_+$ 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*

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

Suppose that all of the following hold:

$$\mu_{00} + \mu_{11} > 0, \ \min\{\mu_{00}, \mu_{11}\} < \sqrt{\mu_{01}\mu_{10}}, \ \max\{\mu_{00}, \mu_{11}\} \le \sqrt{\mu_{01}\mu_{10}}.$$

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

The proof of Lemma 18 uses inapproximability results for antiferromagnetic 2-spin systems on bounded-degree graphs. Thus, before giving its proof, it will be helpful to make a detour to extract the results that will be useful in the proof of Lemma 18.

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).*

3.1 Inapproximability for antiferromagnetic 2-spin systems

We review inapproximability results for the partition function of antiferromagnetic 2-spin models on graphs. We start with a few relevant definitions following [14]. A 2-spin model on a graph is specified by three parameters $\beta, \gamma \ge 0$ and $\lambda > 0$. For a graph $G = (V, E)$, configurations of the model are all possible assignments $\sigma : V \to \{0,1\}$ and the partition function is given by

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

where we adopt the convention that $0^0 \equiv 1$ when one of the parameters β, γ is equal to zero. The case $\beta = \gamma$ corresponds to the Ising model, while the case $\beta = 0$ and $\gamma = 1$ corresponds to the hard-core model.

The 2-spin system with parameters β, γ, λ is called *antiferromagnetic* if $\beta\gamma < 1$. In [21], it was shown that the computational hardness of approximating the partition function in antiferromagnetic 2-spin systems on Δ-regular graphs is captured by the so-called uniqueness threshold on the infinite Δ-regular tree. More precisely, we have the following.

Theorem 20 ([21, Theorems 2 & 3])**.** *Consider the 2-spin systems specified by the following parameters β, γ, λ: (i) $0 < \beta = \gamma < 1$, $\lambda > 0$ (antiferromagnetic Ising model), (ii) $\beta = 0, \gamma = 1$, $\lambda > 0$ (hard-core model). If the 2-spin system specified by the parameters β, γ, λ is in the non-uniqueness regime of the infinite Δ-regular tree for $\Delta \ge 3$, 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 the class of Δ-regular graphs G.*

Note that it is not important for us that Theorem 20 shows hardness all the way to the uniqueness threshold. It would suffice to have a weaker bound for the antiferrogmanetic Ising model and the hard-core model. Luby and Vigoda [16] provide such a result for the hard-core model. Rather than explicitly deriving such a bound for the antiferromagnetic Ising model, we use Theorem 20. Cai et al [4, Theorems 1 & 2] give similar results that apply in some of the relevant parameter space, however it is simpler to work with the later paper [21], especially since the latter shows that the partition function is hard to approximate even within an exponential factor.

For future use, we point out the following characterisation of the uniqueness regime on the infinite Δ-regular tree. For a 2-spin system with parameters β, γ, λ, let $h(x) := \lambda\left(\frac{\beta x+1}{x+\gamma}\right)^{\Delta-1}$ and let x^* be the (unique) positive solution of $x^* = h(x^*)$. Then, uniqueness holds on the infinite Δ-regular tree iff

$|h'(x^*)| \leq 1$, i.e., the absolute value of the derivative of h evaluated at x^* is less than or equal than 1. When, instead, $|h'(x^*)| > 1$, non-uniqueness holds on the infinite Δ-regular tree. Equivalently, one can derive the following equivalent criterion: non-uniqueness on the infinite Δ-regular tree holds iff the system of equations

$$x = \lambda\Big(\frac{\beta y + 1}{y + \gamma}\Big)^{\Delta - 1}, \quad y = \lambda\Big(\frac{\beta x + 1}{x + \gamma}\Big)^{\Delta - 1} \tag{15}$$

has multiple (i.e., more than one) positive solutions (x, y).

It is well-known (see, e.g., [21]) that antiferromagnetic 2-spin systems on Δ-regular graphs can be expressed in terms of either the Ising model or the hard-core model. Theorem 20 thus also gives the regime where general antiferromagnetic 2-spin systems are hard, albeit somewhat implicitly. For the sake of completeness we do this explicitly in the following simple corollary of Theorem 20 (which is nevertheless lengthy to prove).

Corollary 21. *Let $\beta, \gamma \geq 0$ with $\beta\gamma < 1$, $\gamma > 0$, $\lambda > 0$ and $\Delta \geq 3$. If the 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 the class of Δ-regular graphs G.*

Proof. For $\beta, \gamma, \lambda, \Delta$ as in the statement of the lemma, consider the following map:

$$R(\beta, \gamma, \lambda) = \begin{cases} \big(\sqrt{\beta\gamma}, \sqrt{\beta\gamma}, \lambda(\beta/\gamma)^{\Delta/2}\big), & \text{if } \beta > 0, \\ (0, 1, \lambda/\gamma^{\Delta}), & \text{if } \beta = 0. \end{cases}$$

To prove the claim, it suffices to show the following two facts and then to use Theorem 20.

Fact 1. For a Δ-regular graph G, a multiplicative approximation of $Z_{\beta, \gamma, \lambda; G}$ within a factor C yields a multiplicative approximation of $Z_{R(\beta, \gamma, \lambda); G}$ within a factor C.

Fact 2. The 2-spin system with parameters β, γ, λ is in the non-uniqueness regime of the infinite Δ-regular tree iff the 2-spin system with parameters $R(\beta, \gamma, \lambda)$ is in the non-uniqueness regime of the infinite Δ-regular tree.

We consider first the case $\beta > 0$. Let λ' be defined from $(\lambda')^{1/\Delta} = \lambda^{1/\Delta}\sqrt{\beta}/\sqrt{\gamma}$ and let $\beta' = \sqrt{\beta\gamma}$. Note that $R(\beta, \gamma, \lambda) = (\beta', \beta', \lambda')$.

We first show Fact 1. Let $G = (V, E)$ be a Δ-regular graph. Observe that

$$\begin{aligned} Z_{\beta, \gamma, \lambda; G} &= \sum_{\sigma: V \to \{0,1\}} \lambda^{|\sigma^{-1}(0)|} \prod_{(u,v) \in E} \beta^{\mathbf{1}\{\sigma(u)=\sigma(v)=0\}} \gamma^{\mathbf{1}\{\sigma(u)=\sigma(v)=1\}} \\ &= \sum_{\sigma: V \to \{0,1\}} \prod_{(u,v) \in E} (\beta\lambda^{2/\Delta})^{\mathbf{1}\{\sigma(u)=\sigma(v)=0\}} (\lambda^{1/\Delta})^{\mathbf{1}\{\sigma(u)\neq\sigma(v)\}} \gamma^{\mathbf{1}\{\sigma(u)=\sigma(v)=1\}} \\ &= \Big(\frac{\sqrt{\gamma}}{\sqrt{\beta}}\Big)^{|E|} \sum_{\sigma: V \to \{0,1\}} \prod_{(u,v) \in E} (\beta'(\lambda')^{2/\Delta})^{\mathbf{1}\{\sigma(u)=\sigma(v)=0\}} ((\lambda')^{1/\Delta})^{\mathbf{1}\{\sigma(u)\neq\sigma(v)\}} (\beta')^{\mathbf{1}\{\sigma(u)=\sigma(v)=1\}} \\ &= \Big(\frac{\sqrt{\gamma}}{\sqrt{\beta}}\Big)^{|E|} \sum_{\sigma: V \to \{0,1\}} (\lambda')^{|\sigma^{-1}(0)|} \prod_{(u,v) \in E} (\beta')^{\mathbf{1}\{\sigma(u)=\sigma(v)\}} \\ &= \Big(\frac{\sqrt{\gamma}}{\sqrt{\beta}}\Big)^{|E|} Z_{\beta', \beta', \lambda'; G}, \end{aligned}$$

which clearly yields the desired fact.

We next show Fact 2. Note that for positive x, y, (15) is equivalent to

$$x\Big(\frac{\beta y + 1}{y + \gamma}\Big) = \lambda\Big(\frac{\beta y + 1}{y + \gamma}\Big)^{\Delta}, \quad y\Big(\frac{\beta x + 1}{x + \gamma}\Big) = \lambda\Big(\frac{\beta x + 1}{x + \gamma}\Big)^{\Delta} \tag{16}$$

18

Thus, it suffices to show that the positive solutions (x, y) of (16) are in one-to-one correspondence with positive solutions (x', y') to

$$x'\Big(\frac{\beta'y'+1}{y'+\beta'}\Big) = \lambda'\Big(\frac{\beta'y'+1}{y'+\beta'}\Big)^\Delta, \quad y'\Big(\frac{\beta'x'+1}{x'+\beta'}\Big) = \lambda'\Big(\frac{\beta'x'+1}{x'+\beta'}\Big)^\Delta, \tag{17}$$

where $\beta' = \sqrt{\beta\gamma}$ and $\lambda' = \lambda(\sqrt{\beta}/\sqrt{\gamma})^\Delta$ are as before.

Let $x' = (\sqrt{\beta}/\sqrt{\gamma})x$ and $y' = (\sqrt{\beta}/\sqrt{\gamma})y$. Note that for $\beta > 0$ we have that x, y are positive iff x', y' are positive and that x, y are in one-to-one correspondence with x', y'. It is also simple to verify that x, y, x', y' satisfy the equations

$$\frac{\beta x + 1}{x + \gamma} = \frac{\sqrt{\beta}}{\sqrt{\gamma}} \cdot \frac{\sqrt{\beta\gamma}x' + 1}{x' + \sqrt{\beta\gamma}}, \quad \frac{\beta y + 1}{y + \gamma} = \frac{\sqrt{\beta}}{\sqrt{\gamma}} \cdot \frac{\sqrt{\beta\gamma}y' + 1}{y' + \sqrt{\beta\gamma}}, \tag{18}$$

and

$$x\Big(\frac{\beta y + 1}{y + \gamma}\Big) = x'\Big(\frac{\sqrt{\beta\gamma}y' + 1}{y' + \sqrt{\beta\gamma}}\Big), \quad y\Big(\frac{\beta x + 1}{x + \gamma}\Big) = y'\Big(\frac{\sqrt{\beta\gamma}x' + 1}{x' + \sqrt{\beta\gamma}}\Big). \tag{19}$$

Using (18) and (19), one can easily check that x, y satisfy (16) iff x', y' satisfy (17).

Next consider the case $\beta = 0$. The arguments are completely analogous to the case $\beta > 0$, up to minor technical details. Let λ' be defined from $(\lambda')^{1/\Delta} = \lambda^{1/\Delta}/\gamma$ and note that $R(\beta, \gamma, \lambda) = (0, 1, \lambda')$.

For Fact 1, we have that

$$
\begin{aligned}
Z_{\beta,\gamma,\lambda;G} &= \sum_{\sigma:V\to\{0,1\}} \lambda^{|\sigma^{-1}(0)|} \prod_{(u,v)\in E} \beta^{\mathbf{1}\{\sigma(u)=\sigma(v)=0\}} \gamma^{\mathbf{1}\{\sigma(u)=\sigma(v)=1\}} \\
&= \sum_{\sigma:V\to\{0,1\}} \prod_{(u,v)\in E} (\beta\lambda^{2/\Delta})^{\mathbf{1}\{\sigma(u)=\sigma(v)=0\}} (\lambda^{1/\Delta})^{\mathbf{1}\{\sigma(u)\neq\sigma(v)\}} \gamma^{\mathbf{1}\{\sigma(u)=\sigma(v)=1\}} \\
&= \gamma^{|E|} \sum_{\sigma:V\to\{0,1\}} \prod_{(u,v)\in E} \big(\beta\gamma(\lambda')^{2/\Delta}\big)^{\mathbf{1}\{\sigma(u)=\sigma(v)=0\}} \big((\lambda')^{1/\Delta}\big)^{\mathbf{1}\{\sigma(u)\neq\sigma(v)\}} \\
&= \gamma^{|E|} \sum_{\sigma:V\to\{0,1\}} (\lambda')^{|\sigma^{-1}(0)|} \prod_{(u,v)\in E} \big(\beta\gamma\big)^{\mathbf{1}\{\sigma(u)=\sigma(v)=0\}} \\
&= \gamma^{|E|} Z_{0,1,\lambda';G},
\end{aligned}
$$

which again yields the desired fact.

For Fact 2, set $x = \gamma x'$ and $y = \gamma y'$. The following equivalence is easy to see by inspection: x, y satisfy

$$\frac{x}{y+\gamma} = \lambda\Big(\frac{1}{y+\gamma}\Big)^\Delta, \quad \frac{y}{x+\gamma} = \lambda\Big(\frac{1}{x+\gamma}\Big)^\Delta \tag{20}$$

iff x', y' satisfy

$$\frac{x'}{y'+1} = \lambda'\Big(\frac{1}{y'+1}\Big)^\Delta, \quad \frac{y'}{x'+1} = \lambda'\Big(\frac{1}{x'+1}\Big)^\Delta. \tag{21}$$

Note that (20) and (21) correspond to equation (16) for the 2-spin systems with parameters $\beta = 0, \gamma, \lambda$ and $\beta = 0, \gamma = 1, \lambda'$, respectively.

This concludes the proof of the corollary. $\qquad\square$

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 [14, 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 $0 \le \beta < 1$ and $0 < \gamma \le 1$, non-uniqueness holds on the infinite Δ-regular tree for all sufficiently large Δ.

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

In order to prove Theorem 3, 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 Theorem 20. 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) 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 will 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 will show that for β_0 strictly less than 1 and $\gamma_0 \le 1$, 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).[4] We will prove a slightly stronger statement by giving a bound on the required accuracy ε in terms of the degree Δ, which will allow us to switch the order of quantifiers. Also, our result will be monotone in the degree-bound Δ (as in Item 1 above).

Lemma 22. *Suppose $0 \le \beta_0 < 1$. Then, for all sufficiently large Δ, for $\varepsilon = 1/\Delta$, for all β, γ which satisfy*

$$\max\{\beta_0 - \varepsilon, 0\} \le \beta < \beta_0 + \varepsilon \text{ and } 0 < \gamma < 1 + \varepsilon, \tag{22}$$

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. By choosing Δ sufficiently large, for all β, γ which satisfy (22), it clearly holds that $\beta\gamma < 1$ and thus the corresponding 2-spin system is antiferromagnetic. We next show that for all Δ sufficiently large the 2-spin system is also in the non-uniqueness regime of the infinite Δ-regular tree.

We first consider the "soft-constrained" case $\beta_0 > 0$, where we will assume throughout that $\Delta > 1/\beta_0$, so that for all β, γ satisfying (22) it holds that $\beta\gamma > 0$ and $\beta < 1$. We first recall basic facts about the uniqueness regime of soft-contrained antiferromagnetic 2-spin systems on the infinite Δ-regular tree for $\Delta \ge 3$. The reader is referred to, e.g., [14, Lemma 21] for more details.

We have already seen that non-uniqueness holds on the infinite Δ-regular tree when the system of equations (15) has multiple positive solutions. We also saw that this corresponds to the case where $|h'(x^*)| > 1$ where x^* is the unique positive solution of $x^* = h(x^*)$ for a function $h(x)$ defined shortly before Equation (15). In [14, Lemma 21], it is shown that there exist values $\lambda_1 := \lambda_1(\beta, \gamma, \Delta)$ and $\lambda_2 := \lambda_2(\beta, \gamma, \Delta)$ such that the condition $|h'(x^*)| < 1$ holds either when (i) $\sqrt{\beta\gamma} > (\Delta - 2)/\Delta$, or (ii) $\sqrt{\beta\gamma} \le (\Delta - 2)/\Delta$ and $\lambda < \lambda_1$ or $\lambda > \lambda_2$. Adapting the proof [14, Proof of Lemma 21, Item 7] it is not hard to see that the condition $|h'(x^*)| > 1$ which we are interested in holds iff $\sqrt{\beta\gamma} < (\Delta - 2)/\Delta$ and $\lambda \in (\lambda_1, \lambda_2)$, where λ_1, λ_2 are as in [14, Lemma 21, Item 7]. Thus, for $\beta_0 > 0$, our goal is to show that, for all sufficiently large Δ, for all β, γ satisfying (22) with $\varepsilon = 1/\Delta$, it holds that $\lambda_1 < 1 < \lambda_2$.

---

[4]Of course, other approaches may make it possible to directly construct "strictly antiferromagnetic" gadgets, without relying on Lemma 22.

To show the inequalities for λ_1 and λ_2, we next describe explicitly the values of λ_1 and λ_2. For fixed $\beta, \gamma > 0$ and $\Delta \geq 3$ with $\sqrt{\beta\gamma} < (\Delta - 2)/\Delta$, the values of λ_1, λ_2 can be obtained as follows (see [14, Lemma 22]). To align with the setting in [14], we denote $d := \Delta - 1$. When $\sqrt{\beta\gamma} < (d-1)/(d+1)$, it is not hard to show that the equation

$$\frac{d(1-\beta\gamma)x}{(\beta x + 1)(x+\gamma)} = 1, \text{ which is equivalent to } \beta x^2 + \big((d+1)\beta\gamma - (d-1)\big)x + \gamma = 0, \quad (23)$$

has two distinct positive solutions x_1, x_2. Without loss of generality, we may assume that $x_1 < x_2$. For future use, we remark that

$$x_1 = \Big((d-1) - (d+1)\beta\gamma - \sqrt{\big((d-1) - (d+1)\beta\gamma\big)^2 - 4\beta\gamma}\Big)/2\beta$$

$$= 2\gamma/\Big((d-1) - (d+1)\beta\gamma + \sqrt{\big((d-1) - (d+1)\beta\gamma\big)^2 - 4\beta\gamma}\Big), \quad (24)$$

where the latter expression follows by taking the conjugate expresion. We thus obtain the crude bounds

$$x_1 < \frac{2\gamma}{T}, \quad \frac{T}{2\beta} < x_2, \text{ where } T := (d-1) - (d+1)\beta\gamma. \quad (25)$$

(The upper bound for x_1 is obtained by ignoring the square root and the lower bound for x_2 is obtained from the upper bound for x_1 and noticing, from (23), that $x_1 x_2 = \gamma/\beta$.)

The values of λ_1, λ_2 in terms of x_1, x_2 are given by

$$\lambda_1 = \lambda_1(\beta, \gamma, d) := x_1\left(\frac{x_1 + \gamma}{\beta x_1 + 1}\right)^d, \quad \lambda_2 = \lambda_2(\beta, \gamma, d) := x_2\left(\frac{x_2 + \gamma}{\beta x_2 + 1}\right)^d. \quad (26)$$

For future use, note that $(x_1 + \gamma)/(\beta x_1 + 1) < x_1 + \gamma$. Also, since $\beta < 1$ and $\gamma > 0$, $(x_2 + \gamma)/(\beta x_2 + 1) > x_2/(x_2 + 1) > (x_2 - 1)/x_2$. So, using also the bounds for x_1 and x_2 from (25), the expressions in (26) yield the bounds

$$\lambda_1 < x_1(x_1 + \gamma)^d < \frac{2\gamma^{d+1}}{T}(1 + 2/T)^d < \frac{6}{T}(1 + 2/T)^d,$$

$$\lambda_2 > x_2(1 - 1/x_2)^d > \frac{T}{2\beta}(1 - 2\beta/T)^d > \frac{T}{2}(1 - 2/T)^d, \quad (27)$$

where in the rightmost inequalities we used the bounds $\gamma^{d+1} \leq \big(1 + 1/(d+1)\big)^{d+1} < 3$ for $d \geq 2$ (since $\gamma \leq 1 + \varepsilon$ from (22)) and $\beta < 1$ (since $\beta_0 < 1$ and we can use our initial assumption that d is sufficiently large and hence ε in (22) small).

For $\beta_0 > 0$ and β, γ satisfying (22), we have the bound $(d+1)\beta\gamma < (d+2)\beta_0 + 2$, so for all sufficiently large d (depending only on β_0), it holds that $T > d(1 - \beta_0)/2$. Hence, the bounds in (27) yield that, for all sufficiently large d,

$$\lambda_1 < \frac{12}{d(1-\beta_0)}\Big(1 + \frac{4}{d(1-\beta_0)}\Big)^d, \quad \lambda_2 > \frac{d(1-\beta_0)}{4}\Big(1 - \frac{4}{d(1-\beta_0)}\Big)^d.$$

Since $1 > \beta_0 > 0$, we clearly obtain that for all sufficiently large d (depending only on β_0), it holds that $\lambda_1 < 1 < \lambda_2$, as wanted. This completes the proof of the lemma for the case $\beta_0 > 0$.

We next consider the case $\beta_0 = 0$. For $\gamma > 0$, the 2-spin system specified by $\beta = 0, \gamma$ and $\lambda = 1$ is in the uniqueness regime of the infinite $(d+1)$-regular tree iff $1 \leq \lambda_c(\gamma, d) := \frac{\gamma^{d+1}d^d}{(d-1)^{d+1}}$ (see, for example, [14, Proof of Item 5 in Lemma 21]). Since $\lambda_c(\gamma, d)$ is an increasing function of γ, it suffices to show that $\lambda_c(1 + \varepsilon, d) < 1$ for all sufficiently large $d \geq 2$. Note that $(1 + \varepsilon)^{d+1} < 3$ for all $d \geq 2$ and, for $d \geq 10$, we have $\frac{d^d}{(d-1)^{d+1}} \leq 1/3$. It follows that for all $d \geq 10$, it holds that $\lambda_c(1 + \varepsilon, d) < 1$, as needed.

To complete the proof for the case $\beta_0 = 0$, we need to argue that for all sufficiently large d, for all $(\beta, \gamma) \in (0, \varepsilon) \times (0, 1 + \varepsilon)$ with $\varepsilon = 1/(d+1)$, the 2-spin system with parameters $\beta, \gamma, \lambda = 1$, is in the non-uniqueness regime of the infinite $(d+1)$-regular tree. For $d \geq 10$ and this range of the parameters β, γ we have that $\beta\gamma > 0$ and $\sqrt{\beta\gamma} \leq (d-1)/(d+1)$. Thus it suffices to establish that for all sufficiently large d it holds that $\lambda_1(\beta, \gamma, d) < 1 < \lambda_2(\beta, \gamma, d)$, where λ_1, λ_2 are as in (26). In fact, we can use the bounds in (27), so we only need to provide a lower bound on T since our derivation previously was to account for the case $\beta_0 > 0$. For this range of the parameters β, γ, we have that $(d+1)\beta\gamma < 2$. It follows that $T > d - 3$. Thus, the bounds in (27) yield

$$\lambda_1 < \frac{6}{d-3}\Big(1 + \frac{2}{d-3}\Big)^d, \quad \lambda_2 > \frac{d-3}{2}\Big(1 - \frac{2}{d-3}\Big)^d,$$

so that $\lambda_1 < 1 < \lambda_2$ for all sufficiently large d. This concludes the proof of the lemma. $\qquad\square$

3.2   Proof of Lemma 18

Using Lemmas 17 and 22, we now give the proof of Lemma 18. The idea is to use Lemma 17 to obtain a hypergraph that realises the conditional distribution with sufficient accuracy ε. The resulting hypergraph can be used to simulate an antiferromagnetic 2-spin system which, by Lemma 22 and Corollary 21, will be hard to approximate on Δ-regular graphs (for large Δ). The formal proof is as follows.

Lemma 18. *Let $f : \{0, 1\}^k \to \mathbb{R}_+$ 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*

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

Suppose that all of the following hold:

$$\mu_{00} + \mu_{11} > 0, \ \min\{\mu_{00}, \mu_{11}\} < \sqrt{\mu_{01}\mu_{10}}, \ \max\{\mu_{00}, \mu_{11}\} \leq \sqrt{\mu_{01}\mu_{10}}. \tag{28}$$

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

Proof. Let $H = (V, \mathcal{F})$. We start by applying Lemma 17 with $S = \{x, y\}$. For every $\varepsilon' > 0$ and every $s_1, s_2 \in \{0, 1\}$, the lemma shows that there is a hypergraph $H' = (V', \mathcal{F}')$ with $V \subseteq V'$ and $\mathcal{F} \subseteq \mathcal{F}'$ so that

$$\left|\mu_{f;H'}(\sigma_x = s_1, \sigma_y = s_2) - \mu_{f;H}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = s_1, \sigma_y = s_2)\right| \leq \varepsilon'. \tag{29}$$

For $s_1, s_2 \in \{0, 1\}$, let $\mu'_{s_1 s_2} = \mu_{f;H'}(\sigma_x = s_1, \sigma_y = s_2)$. Thus, (29) becomes

$$\left|\mu'_{s_1 s_2} - \mu_{s_1 s_2}\right| \leq \varepsilon'. \tag{30}$$

The conditions in (28) guarantee that μ_{01} and μ_{10} are positive, so by choosing ε' sufficiently small, we can also guarantee that μ'_{01} and μ'_{10} are positive.

Assume without loss of generality that $\mu_{00} \leq \mu_{11}$. (Otherwise, we will swap the role of the spins 0 and 1.) Let $\beta_0 = \mu_{00}^2/\mu_{01}\mu_{10}$ and let $\gamma_0 = \mu_{11}^2/\mu_{01}\mu_{10}$. By (28), $0 \leq \beta_0 < 1$ and $0 < \gamma_0 \leq 1$.

Next, we "symmetrise" the hypergraph H' to obtain a hypergraph H'' (an analogous argument was used previously in the proof of Lemma 12). To do this, take two disjoint copies of H' which we denote by H'_1, H'_2. For $i = 1, 2$, denote by x_i, y_i the images of the vertices x, y in H'_i. Now identify vertices x_1 and y_2 into a single vertex x, and similarly identify vertices x_2 and y_1 into a single vertex y. Let H'' be the final hypergraph and let $\mu''_{s_1 s_2}$ denote $\mu_{f;H''}(\sigma(x) = s_1, \sigma(y) = s_2)$. Then for all $s_1, s_2 \in \{0, 1\}$, we have

$$\mu''_{s_1 s_2} = \frac{\mu'_{s_1 s_2}\mu'_{s_2 s_1}}{\sum_{t_1, t_2 \in \{0, 1\}} \mu'_{t_1 t_2}\mu'_{t_2 t_1}}.$$

Note that $\mu''_{01} = \mu''_{10} > 0$.

Consider the 2-spin system with parameters $\beta := \mu''_{00}/\mu''_{01}$, $\gamma := \mu''_{11}/\mu''_{01}$, and $\lambda = 1$. Using the definitions of β_0 and γ_0 and Equation (30) we find that for every $\Delta' \geq 3$, and for every sufficiently small $\varepsilon' > 0$, we have

$$|\beta - \beta_0| < 1/\Delta' \text{ and } |\gamma - \gamma_0| < 1/\Delta'. \tag{31}$$

By Lemma 22, there is a $\Delta'_0 \geq 3$ such that if $\Delta' \geq \Delta'_0$ then the spin system with parameters $\beta, \gamma, \lambda = 1$ is in the non-uniqueness regime of the infinite Δ'-regular tree. Thus, by Corollary 21, there is a $c > 1$ such that approximating $Z_{\beta,\gamma,1;G}$ within a factor of c^n is NP-hard on the class of Δ'-regular n-vertex graphs G.

Let Δ'' be the maximum degree of the hypergraph H''. We will show the lemma for all $\Delta \geq \Delta'\Delta''$.

Let $G = (V, E)$ be a Δ'-regular graph for which we want to compute $Z_{\beta,\gamma,1;G}$. We construct the hypergraph H''' by replacing each edge of G with a copy of the hypergraph H'' as follows. For each edge $(u, v) \in E$, take a (distinct) copy $H_{uv} = (V_{uv}, \mathcal{F}_{uv})$ of the hypergraph H''. Denote by x_{uv}, y_{uv} the images of the vertices x, y in H_{uv}. Now for each $u \in V$ identify the vertices $x_{uv_1}, \ldots, x_{uv_{\Delta'}}$ into a single vertex u, where $v_1, \ldots, v_{\Delta'}$ denote the neighbors of u in G. It is clear that the hypergraph H''' has maximum degree Δ.

Let $V''', \mathcal{F}'''$ denote the vertex and hyperedge sets of H''', respectively. For $\sigma : V \to \{0, 1\}$, let $\Sigma(\sigma) = \{\tau : V''' \to \{0, 1\} \mid \tau_V = \sigma\}$. The total contribution to the partition function $Z_{f;H'''}$ from configurations in $\Sigma(\sigma)$ is exactly $\prod_{(u,v)\in E}(\mu''_{\sigma(u)\sigma(v)}Z_{f;H''})$. Thus, we obtain

$$Z_{f;H'''} = (\mu''_{01}Z_{f;H''})^{|E|} \sum_{\sigma:V\to\{0,1\}} \prod_{(u,v)\in E} \frac{\mu''_{\sigma(u)\sigma(v)}}{\mu''_{01}} = (\mu''_{01}Z_{f;H''})^{|E|}\, Z_{\beta,\gamma,1;G}. \tag{32}$$

Since $\mu''_{01}Z_{f;H''}$ is an explicitly computable constant, it follows that an approximation to the partition function $Z_{f;H'''}$ within a factor of c^n yields an approximation to $Z_{\beta,\gamma,1;G}$ within a factor of c^n. This completes the proof since Δ' and H'' are fixed, which guarantees that the number of vertices of H''' is a constant multiple of the number of vertices of G. $\qquad\square$

4   Proof of Theorem 3

In this section, we give the proof of Theorem 3. Let $k \geq 2$ and $f : \{0,1\}^k \to \{0,1\}$ be a *symmetric Boolean* function with $f \notin \mathsf{EASY}(k)$. Our goal is to show that there exists Δ_0 such that for all $\Delta \geq \Delta_0$, there exists $c > 1$ such that $\#\mathsf{Hyper2Spin}(f, \Delta, c)$ is NP-hard.

The case $k = 2$ corresponds to approximating the partition function of *unweighted* 2-spin systems in graphs. It is not hard to see that the only symmetric arity-2 Boolean functions with $f \notin \mathsf{EASY}(k)$ are the function which is 1 if at least one of x_1 and x_2 is 0 and the function which is 1 if at least one of x_1 and x_2 is 1. The partition function in both of these cases corresponds to counting the number of the independent sets, or equivalently to the 2-spin system with $\beta = 0$ and $\gamma = 1$ (and no external field). For this model, it is well known that non-uniqueness holds in the infinite Δ-regular tree when $\Delta \geq 6$, which in conjuction with Theorem 21 completes the proof of Theorem 3 in the special case $k = 2$ (alternatively, one may use Lemma 22 to argue that non-uniqueness holds for all sufficiently large Δ).

Thus, for the rest of the proof, we will assume that $k \geq 3$. We may further assume that at least one of $w_0, \ldots, w_k$ is 0 (otherwise f is the constant function $f^{(k)}_{\mathsf{one}}$) and at least one is 1 (otherwise f is the constant function $f^{(k)}_{\mathsf{zero}}$).

By Lemma 13, to prove Theorem 3 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.)

In each case, the goal is to show that there exists Δ_0 such that for all $\Delta \geq \Delta_0$, there exists $c > 1$ such that $\#\mathsf{Hyper2Spin}(f, \Delta, c)$ is NP-hard.

For example, when f is the function corresponding to weak independent sets, f supports pinning-to-0 but not pinning-to-1, so f is in Case 3. The same is true when f is the function corresponding to strong independent sets. On the other hand, if f is the "not-all-equal" function, then it supports 2-equality, but neither pinning-to-0 nor pinning-to-1, so it is in Case 2.

Before delving into the proofs for each of this cases, we give a piece of terminology which will simplify the exposition. The reader may wish to recall Definition 16. We will typically invoke Lemma 18 for a hypergraph H and an admissible collection of sets $\mathcal{V} = (V_0, V_1, V_2, \ldots, V_r)$. Rather than formally defining $\mathcal{V}$ in each such application of Lemma 18, it will be convenient (and more instructive) to say, e.g., pin vertices x_1, x_2 to 0 (instead of specifying V_0 as $V_0 = \{x_1, x_2\}$), pin vertices x_3, x_4 to 1 (instead of specifying V_1 as $V_1 = \{x_3, x_4\}$), and force equality among x_5, x_6, x_7 (instead of specifying V_2 as $V_2 = \{x_5, x_6, x_7\}$).

4.1 Case I

In this section, we assume that the function f supports both pinning-to-0 and pinning-to-1. In this case, Lemma 18 will always be applied to the hypergraph with a single edge $e := \{x_1, x_2, \ldots, x_k\}$. For explicitness and with a slight abuse of notation we will denote by e this hypergraph.

Case Ia Suppose first that exactly one of $w_0, \ldots, w_k$ is equal to 1, say $w_j = 1$. We may assume that $j \neq 0$ and $j \neq k$, otherwise $f = f_{\mathsf{allzero}}^{(k)}$ or $f = f_{\mathsf{allone}}^{(k)}$, respectively. We will consider separately the cases $j = 1$ and $k > j \geq 2$.

Suppose first that $k > j \geq 2$. Pin $x_1, \ldots, x_{j-2}$ to 1 (if $j = 2$, no vertex is pinned to one), pin $x_{j+2}, \ldots, x_k$ to 0 (if $j + 2 > k$, no vertex is pinned to zero) and set $x := x_{j-1}$, $y := x_j$ (since $j < k$, note that vertex x_{j+1} is "free"). We have:

$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 0) \propto w_{j-2} + w_{j-1} = 0,$$
$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = 0, \sigma_y = 1) \propto w_{j-1} + w_j = 1,$$
$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 1) \propto w_j + w_{j+1} = 1.$$

Also, by symmetry, $\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = 0, \sigma_y = 1) = \mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = 1, \sigma_y = 0)$. (We will use similar symmetry arguments in the rest of this proof without pointing them out explicitly). So , from Lemma 18, for all sufficiently large Δ, there exists $c > 1$ such that $\#\mathsf{Hyper2Spin}(f, \Delta, c)$ is NP-hard.

Suppose next that $j = 1$. Pin $x_4, \ldots, x_k$ to 0 and set $x := x_1$, $y := x_2$ (since $k \geq 3$, note that vertex x_3 is "free"). We have:

$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 0) \propto w_0 + w_1 = 1,$$
$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = 0, \sigma_y = 1) \propto w_1 + w_2 = 1,$$
$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 1) \propto w_2 + w_3 = 0,$$

so, from Lemma 18, for all sufficiently large Δ, there exists $c > 1$ such that $\#\mathsf{Hyper2Spin}(f, \Delta, c)$ is NP-hard.

Case Ib. Suppose next that at least two of $w_0, \ldots, w_k$ are equal to 1. Let i, j be two indices with $i < j$ such that $w_i = w_j = 1$ and $w_{i+1} = \ldots = w_{j-1} = 0$. We will first prove that for any two such indices, it holds that $j = i + 2$ (otherwise, we will show that $\#\mathsf{Hyper2Spin}(f, \Delta, c)$ is NP-hard). Since $f \neq f_{\mathsf{EQ}}^{(k)}$, we may assume that either $i > 0$ or $j < k$. Without loss of generality, we assume that $j < k$ (otherwise we may swap the spins 0 and 1).

1. If $j > i + 2$, we consider cases whether $w_{j+1} = 0$ or 1. If $w_{j+1} = 0$, pin $x_1, \ldots, x_{j-2}$ to 1, pin $x_{j+2}, \ldots, x_k$ to 0 and set $x := x_{j-1}$, $y := x_j$ (since $j < k$, note that vertex x_{j+1} is "free"). We have:

$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 0) \propto w_{j-2} + w_{j-1} = 0,$$

24

$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = 0, \sigma_y = 1) \propto w_{j-1} + w_j = 1,$$
$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 1) \propto w_j + w_{j+1} = 1,$$

so, from Lemma 18, for all sufficiently large Δ, there exists $c > 1$ such that $\#\mathsf{Hyper2Spin}(f, \Delta, c)$ is NP-hard. If $w_{j+1} = 1$, pin $x_1, \ldots, x_{j-1}$ to 1, pin $x_{j+2}, \ldots, x_k$ to 0 and set $x := x_j$, $y := x_{j+1}$. We have:

$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 0) \propto w_{j-1} = 0,$$
$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = 0, \sigma_y = 1) \propto w_j = 1,$$
$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 1) \propto w_{j+1} = 1,$$

so, from Lemma 18, for all sufficiently large Δ, there exists $c > 1$ such that $\#\mathsf{Hyper2Spin}(f, \Delta, c)$ is NP-hard.

2. Assume now that $j = i + 1$. Suppose there exists j' such that $w_i = w_{i+1} = \ldots = w_{j'} = 1$ and $w_{j'+1} = 0$. Pin $x_1, \ldots, x_{j'-1}$ to 1, pin $x_{j'+2}, \ldots, x_k$ to 0 and set $x := x_{j'}$, $y := x_{j'+1}$. We have:

$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 0) \propto w_{j'-1} = 1,$$
$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = 0, \sigma_y = 1) \propto w_{j'} = 1,$$
$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 1) \propto w_{j'+1} = 0,$$

so, from Lemma 18, for all sufficiently large Δ, there exists $c > 1$ such that $\#\mathsf{Hyper2Spin}(f, \Delta, c)$ is NP-hard.

If such a j' does not exist, then it holds that $w_i = w_{i+1} = \ldots = w_k = 1$, so there exists $i' > 0$ such that $w_{i'} = \ldots = w_k = 1$ and $w_{i'-1} = 0$ (otherwise $f = f_{\mathsf{one}}^{(k)}$). Pin $x_{i'+2}, \ldots, x_k$ to 0, pin $x_1, \ldots, x_{i'-1}$ to 1 and set $x := x_{i'}$, $y := x_{i'+1}$. We have:

$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 0) \propto w_{i'-1} = 0,$$
$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = 0, \sigma_y = 1) \propto w_{i'} = 1,$$
$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 1) \propto w_{i'+1} = 1,$$

so again from Lemma 18, for all sufficiently large Δ, there exists $c > 1$ such that $\#\mathsf{Hyper2Spin}(f, \Delta, c)$ is NP-hard.

It follows that for every two indices with $i < j$ such that $w_i = w_j = 1$ and $w_{i+1} = \ldots = w_{j-1} = 0$, it holds that $j = i + 2$. Let i' be the minimum integer such that $w_{i'} = 1$. We have that $w_0 = w_1 = \ldots = w_{i'-1} = 0$. Let j' be the maximum integer such that $w_{i'} = w_{i'+2} = \ldots = w_{i'+2j'} = 1$. By assumption, at least two w_i's are equal to 1, so we have that $j' \geq 1$. We also have that $w_{i'+1} = w_{i'+3} = \ldots = w_{i'+2j'-1} = 0$, $w_{i'+2j'+1} = \ldots = w_k = 0$.

We may assume that either $i' \notin \{0, 1\}$ or $i' + 2j' \notin \{k - 1, k\}$ (otherwise either $f = f_{\mathsf{even}}^{(k)}$ or $f = f_{\mathsf{odd}}^{(k)}$). Let us assume first that $i' + 2j' \notin \{k - 1, k\}$, i.e., $i' + 2j' \leq k - 2$, so that $w_{i'+2j'+2} = 0$. Pin $x_1, \ldots, x_{i'}$ to 1, set $x := x_{i'+1}, y := x_{i'+2}$ and pin $x_{i'+2j'+3}, \ldots, x_k$ to 0. We have:

$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 0) \propto \sum_{\ell=0}^{2j'} \binom{2j'}{\ell} w_{i'+\ell} = 2^{2j'-1},$$

$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = 0, \sigma_y = 1) \propto \sum_{\ell=0}^{2j'} \binom{2j'}{\ell} w_{i'+1+\ell} = 2^{2j'-1},$$

$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 1) \propto \sum_{\ell=0}^{2j'} \binom{2j'}{\ell} w_{i'+2+\ell} = 2^{2j'-1} - 1,$$

so from Lemma 18, for all sufficiently large Δ, there exists $c > 1$ such that $\#\mathsf{Hyper2Spin}(f, \Delta, c)$ is NP-hard.

The case $i' \notin \{0, 1\}$ can be covered by an analogous argument, the only difference being that now we pin $x_1, \ldots, x_{i'-1}$ to 1, set $x := x_{i'}, y := x_{i'+1}$ and pin $x_{i'+2j'+1}, \ldots, x_k$ to 0. Other than that, the previous calculations may be easily modified to obtain

$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 0) \propto 2^{2j'-1} - 1,$$
$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = 0, \sigma_y = 1) \propto 2^{2j'-1},$$
$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 1) \propto 2^{2j'-1},$$

yielding, by Lemma 18, that for all sufficiently large Δ, there exists $c > 1$ such that $\#\mathsf{Hyper2Spin}(f, \Delta, c)$ is NP-hard. This concludes the proof in the case where f supports both pinning-to-0 and pinning-to-1.

4.2 Case II

In Case II we assume that $k \geq 3$ and that f is a symmetric arity-k Boolean function that is not in $\mathsf{EASY}(k)$. The function f supports 2-equality but does not support pinning-to-0 or pinning-to-1.

By Lemma 14 we conclude that f is self-dual, meaning that $w_\ell = w_{k-\ell}$ for all $\ell \in \{0, \ldots, k\}$. By Lemma 10, we conclude that f supports t-equality for all $t \geq 2$. Our goal is to show that for all sufficiently large Δ, there exists $c > 1$ such that the approximation problem $\#\mathsf{Hyper2Spin}(f, \Delta, c)$ is NP-hard. We prove this by considering two cases, depending on whether $w_0 = 0$ or $w_0 = 1$. We will use the following lemma in both cases.

Lemma 23. *Let f be an arity-k symmetric Boolean formula that is self-dual. Let H be a hypergraph with vertex set V and denote by $\Sigma := \{\sigma \mid \sigma : V \to \{0, 1\}\}$ the set of all $\{0, 1\}$ assignments on V. Let $Q : \Sigma \to \Sigma$ be the map which maps an assignment σ to its complement $\bar{\sigma}$, i.e., $\bar{\sigma}$ is defined by $\bar{\sigma}_v = 1 - \sigma_v$ for all $v \in V$. Then for every $\Sigma' \subseteq \Sigma$, it holds that $\mu_{f;H}(\Sigma') = \mu_{f;H}(Q(\Sigma'))$.*

Proof. For every $\sigma \in \Sigma$, self-duality gives that $w_{f;H}(\sigma) = w_{f;H}(\bar{\sigma})$. Summing this equality over all assignments σ in the subset Σ' yields the result. $\qed$

We will split the analysis into two cases – the case where $w_0 = 0$ (Section 4.2.2) and the case where $w_0 = 1$ (Section 4.2.3). Before these two sections, we make a digression into Constraint Satistfaction Problems (CSP). The digression will introduce and prove a lemma that we will need for the $w_0 = 0$ case. In addition, it will provide some missing detail which we used in the Introduction to explain the context of existing work.

4.2.1 A digression regarding Constraint Satisfaction Problems

Recall the CSP definitions from Section 1.1. Let $\mathsf{CSP}(\Gamma)$ 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.

We will use the following CSP terminology. Let f be an arity-k Boolean function. For some positive integer m, let g be a function $g : \{0, 1\}^m \to \{0, 1\}$. Suppose that $\overline{x}_1, \ldots, \overline{x}_m$ are m Boolean k-tuples so for $i \in \{1, \ldots, m\}$ we can write $\overline{x}_i$ as a tuple $\overline{x}_i = (x_{i,1}, \ldots, x_{i,k})$ in $\{0, 1\}^k$. We will let $\overline{y}_g(\overline{x}_1, \ldots, \overline{x}_m) = (y_1, \ldots, y_k)$ be the Boolean k-tuple constructed from g and from $\overline{x}_1, \ldots, \overline{x}_m$ as follows. For each $j \in \{1, \ldots, k\}$, y_j is obtained by applying g to $x_{1,j}, \ldots, x_{m,j}$ so $y_j = g(x_{1,j}, \ldots, x_{m,j})$. The function g is said to be a *polymorphism* of f if, for any choice of m tuples $\overline{x}_1, \ldots, \overline{x}_m$ satisfying $f(\overline{x}_1) = \cdots = f(\overline{x}_m) = 1$, we also have $f(\overline{y}_g(\overline{x}_1, \ldots, \overline{x}_m)) = 1$.

We will use the following algebraic formulation of Chen [6, Theorem 3.21] of Schaefer's famous dichotomy theorem [18].

Theorem 24. *(Schaefer) Let Γ be a finite Boolean constraint language. The problem $\mathsf{CSP}(\Gamma)$ is polynomial-tractable if one of the following six functions is a polymorphism of every function $f \in \Gamma$.*

1. *g is the unary function g_0 with $g_0(0) = g_0(1) = 0$.*

2. *g is the unary function g_1 with $g_1(0) = g_1(1) = 1$.*

3. *g is the arity-2 Boolean function AND.*

4. *g is the arity-2 Boolean function OR.*

5. *g is the ternary majority function Maj defined by $\mathsf{Maj}(a, b, c) = (a \wedge b) \vee (a \wedge c) \vee (b \wedge c)$.*

6. *g is the ternary minority function $\mathsf{Minority}$ defined by $\mathsf{Minority}(a, b, c) = a \oplus b \oplus c$.*

Otherwise, $\mathsf{CSP}(\Gamma)$ is NP-complete.

We start by showing that the decision problem $\mathsf{CSP}(\{f\})$ is NP-hard when f is a non-trivial arity-k symmetric Boolean formula that is self-dual and satisfies $w_0 = 0$. In Section 4.2.2 we will use this fact and the fact f supports equality to show that $\#\mathsf{Hyper2Spin}(f, \Delta, c)$ is NP-hard.

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

Proof. We will show that f does not satisfy any of the tractable cases in Schaefer's dichotomy theorem (Lemma 24). Since f is not $f_{\mathsf{zero}}^{(k)}$, it is not identically zero. Thus, there is a j in the range $1 \leq j \leq \lceil k/2 \rceil$ such that $w_j = 1$. We will use this value j in the cases below.

Cases 1 and 2: We first show that g_0 is not a polymorphism of f. To see this, let $\overline{x}_1$ be any k-tuple with j ones so that $f(\overline{x}_1) = 1$. Applying the function g_0 position-wise, we get $\overline{y}_{g_0}(\overline{x}_1, \overline{x}_2) = (y_1, \ldots, y_k) = (g_0(x_{1,1}), \ldots, g_0(x_{1,k})) = (0, \ldots, 0)$. But since $w_0 = 0$, $f(0, \ldots, 0) = 0$, contrary to the fact that $f(y_1, \ldots, y_k)$ would have to be 1 if g_0 were a polymorphism of f. Similarly, applying g_1 componentwise to $\overline{x}_1$ we get $\overline{y}_{g_1}(\overline{x}_1, \overline{x}_2) = (1, \ldots, 1)$. Since $w_k = 0$ (by self-duality), $f(1, \ldots, 1) = 0$ so g_1 is not a polymorphism of f.

Cases 3 and 4: Instead of doing both cases, we first use self-duality to argue that if OR is a polymorphism of f then so is AND. For this, suppose that OR is a polymorphism of f. Let $\overline{x}_1$ and $\overline{x}_2$ be two tuples with $f(\overline{x}_1) = f(\overline{x}_2) = 1$. For $i \in \{1, 2\}$, let $\neg\overline{x}_i$ be the position-wise Boolean complement of $\overline{x}_i$. By self-duality, $f(\neg\overline{x}_1) = f(\neg\overline{x}_2) = 1$ so $f(y_{\mathsf{OR}}(\neg\overline{x}_1, \neg\overline{x}_2)) = 1$. But $y_{\mathsf{OR}}(\neg\overline{x}_1, \neg\overline{x}_2)$ is the position-wise Boolean complement of $y_{\mathsf{AND}}(\overline{x}_1, \overline{x}_2)$, so by self-duality, we also have $f(y_{\mathsf{AND}}(\overline{x}_1, \overline{x}_2)) = 1$, establishing that AND is also a polymorphism of f. Thus, we can complete both cases by just showing that AND is actually not a polymorphism of f.

Recall the value $j \leq \lceil k/2 \rceil$ from above. We first deal with the simplest case where there is a $j \leq k/2$ with $w_j = 1$. Consider two tuples $\overline{x}_1$ and $\overline{x}_2$, each with j ones, chosen so that there is no position ℓ with $x_{1,\ell} = x_{2,\ell} = 1$. This is possible since $j \leq k/2$. Then $f(\overline{x}_1) = f(\overline{x}_2) = 1$ but $\overline{y}_{\mathsf{AND}}(\overline{x}_1, \overline{x}_2) = (0, \ldots, 0)$. Now $w_0 = 0$, so $f(\overline{y}_{\mathsf{AND}}(\overline{x}_1, \overline{x}_2)) = 0$ and we have shown that AND is not a polymorphism of f.

We now deal with the remaining case. We have $j = \lceil k/2 \rceil$ and $w_j = 1$ and every $\ell \neq j$ has $w_\ell = 0$. In this case we can consider any distinct tuples $\overline{x}_1$ and $\overline{x}_2$ with exactly j ones. Then $\overline{y}_{\mathsf{AND}}(\overline{x}_1, \overline{x}_2)$ has fewer than j ones so $f(\overline{y}_{\mathsf{AND}}(\overline{x}_1, \overline{x}_2)) = 0$, so AND is not a polymorphism of f.

Case 5: Suppose that Maj is a polymorphism of f. We will derive a contradiction. In order to simplify the tedious special cases arising from floors and ceilings we write k as $k = 6r + 3a + b$ where r is a non-negative integer, $a \in \{0, 1\}$ and $b \in \{0, 1, 2\}$.

First, the fact that Maj is a polymorphism of f implies that for every $\ell \leq 2r + a$, we have $w_\ell = 0$. To see this, let $\overline{x}_1$, $\overline{x}_2$ and $\overline{x}_3$ be three k-tuples, each with ℓ ones, such that there is no position p with more than a single one amongst $x_{1,p}$, $x_{2,p}$ and $x_{3,p}$. This is possible since $\ell \leq k/3$. Then $\overline{y}_{\mathsf{Maj}}(\overline{x}_1, \overline{x}_2, \overline{x}_3) = (0, \ldots, 0)$. So if Maj is a polymorphism of f we must have that one of $f(\overline{x}_1)$, $f(\overline{x}_2)$ and $f(\overline{x}_3)$ is 0 (which means, by symmetry of f, that all of them are 0), so $w_\ell = 0$.

Now consider any integer ℓ in the range $2r + a < \ell \leq 3r + a$. Specifically, for an integer s in the range $1 \leq s \leq r$, let $\ell = s + 2r + a$. Consider three k-tuples $\overline{x}_1$, $\overline{x}_2$ and $\overline{x}_3$, each with ℓ ones, such that s positions have ones in all three tuples, s positions have ones in tuples $\overline{x}_1$ and $\overline{x}_2$, and the remaining

27

positions have a one in exactly one tuple. Then $\overline{y}_{\mathsf{Maj}}(\overline{x}_1, \overline{x}_2, \overline{x}_3)$ has $2s$ ones. Since $2s \leq 2r$, we have $w_{2s} = 0$. So if Maj is a polymorphism we must have $f(\overline{x}_1) = f(\overline{x}_2) = f(\overline{z}_3) = 0$ so $w_\ell = 0$.

So the only possible values of j where we could have $w_j = 1$ satisfy $j > 3r + a$. By self-duality, they also satisfy $k - j > 3r + a$ so $3r + a + 1 \leq j < 3r + 2a + b$. Since $a + 1 < 2a + b$, the pair (a, b) is in the set $\{(0, 2), (1, 1), (1, 2)\}$. So the three possibilities are

- $(a, b) = (0, 2)$ so $k = 6r + 2$ is even and the only j with $w_j = 1$ satisfies $3r + 1 \leq j < 3r + 2$ so $j = 3r + 1 = k/2$.

- $(a, b) = (1, 1)$ so $k = 6r + 3 + 1$ is even and the only j with $w_j = 1$ satisfies $3r + 2 \leq j < 3r + 3$ so $j = 3r + 2 = k/2$.

- $(a, b) = (1, 2)$ so $k = 6r + 3 + 2$ is odd and the only j with $w_j = 1$ satisfy $3r + 2 \leq j < 3r + 4$ so by self-duality, there exactly two values w_j that are non-zero, and these are $j = 3r + 2 = \lfloor k/2 \rfloor$ and $j = 3r + 3 = \lceil k/2 \rceil$.

There must be a j with $w_j \neq 0$ since f is not the constant zero function. We show that in all three cases Maj is not a polymorphism. We take the first two cases together, so suppose that k is even, and that there is exactly one positive w_j which is $w_{k/2}$. Since k is even and greater than 2, it is at least 4. Choose $\overline{x}_1$ with 1's in positions $1, \ldots, k/2$ and $\overline{x}_2$ with 1's in positions $2, \ldots, k/2 + 1$. Choose $\overline{x}_3$ with 1's in positions $k/2 + 1, \ldots, k - 1$ and 1. Then $\overline{y}_{\mathsf{Maj}}(\overline{x}_1, \overline{x}_2, \overline{x}_3)$ has $(k/2) + 1$ ones (in positions $1, \ldots, k/2 + 1$) so $f(\overline{y}_{\mathsf{Maj}}(\overline{x}_1, \overline{x}_2, \overline{x}_3)) = 0$ and Maj is not a polymorphism.

The final case is similar. Suppose that $k = 2t + 1$ and that there are exactly two positive w_j's which are w_t and w_{t+1}. Choose $\overline{x}_1$ with 1's in positions $1, \ldots, t + 1$. Choose $\overline{x}_2$ with 1's in positions $1, \ldots, t$ and $t + 2$. Choose $\overline{x}_3$ with 1's in positions $t + 1, \ldots, 2t + 1$. Then $\overline{y}_{\mathsf{Maj}}(\overline{x}_1, \overline{x}_2, \overline{x}_3)$ has $t + 2$ ones (in positions $1, \ldots, t + 2$) so $f(\overline{y}_{\mathsf{Maj}}(\overline{x}_1, \overline{x}_2, \overline{x}_3)) = 0$ and Maj is not a polymorphism.

Case 6: Suppose that $\mathsf{Minority}$ is a polymorphism of f. We will derive some consequences about the w_ℓ values which will give us a contradiction.

First, there is no index ℓ such that $w_\ell = 1$ and $w_{\ell-2} = 0$. Clearly this is not the case for $\ell = k$ since $w_k = 0$. Suppose for contradiction that it is true for some $2 \leq \ell < k$. Construct tuples $\overline{x}_1$, $\overline{x}_2$ and $\overline{x}_3$ each with ℓ ones, such that the first $\ell - 2$ positions have ones in all three tuples, and in each of the next three positions there is exactly one zero. Any remaining positions are all zero. This is possible since $\ell + 1 \leq k$. Then $\overline{y}_{\mathsf{Minority}}(\overline{x}_1, \overline{x}_2, \overline{x}_3)$ has $\ell - 2$ ones so if $\mathsf{Minority}$ is a polymorphism of f then $f(\overline{y}_{\mathsf{Minority}}(\overline{x}_1, \overline{x}_2, \overline{x}_3)) = 1$ so $w_{\ell-2} = 1$, contradicting the assumption.

Next, there is no odd index $3 \leq j \leq k/2$ such that $w_j = 0$ and $w_{j-2} = 1$. Suppose for contradiction that this is true for some $j = 2r + 1$. Construct tuples $\overline{x}_1$, $\overline{x}_2$ and $\overline{x}_3$ as follows.

- The first $r - 1$ positions have ones in tuples $\overline{x}_1$ and $\overline{x}_3$.

- The next $r - 1$ positions have ones in tuples $\overline{x}_2$ and $\overline{x}_3$.

- The next r positions have ones in tuple $\overline{x}_1$ only.

- The next r positions have ones in tuple $\overline{x}_2$ only.

- The next position has a one in tuple $\overline{x}_3$ only.

- Any remaining positions are all zero.

Each tuple has $2r - 1 = j - 2$ ones. The construction is possible since $2(r - 1) + 2r + 1 = 4r - 1 \leq 4r + 2 = 2j \leq k$. Then $\overline{y}_{\mathsf{Minority}}(\overline{x}_1, \overline{x}_2, \overline{x}_3)$ has $2r + 1 = j$ ones so and if $\mathsf{Minority}$ is a polymorphism of f then $f(\overline{y}_{\mathsf{Minority}}(\overline{x}_1, \overline{x}_2, \overline{x}_3)) = 1$ so $w_j = 1$, contradicting the assumption.

The first fact rules out the possibility that there is an even index w_ℓ with $w_\ell = 1$. This is ruled out because $w_0 = 0$ and we can derive a contradiction by considering the smallest even ℓ such that $w_\ell = 1$. This also tells us that k is even. This follows because there has to be some j with $w_j = 1$ and from the above, j has to be odd. But if k is odd then $k - j$ is even, yet self-duality would imply that $w_{k-j} = 1$.

So since f is not the trivial all-zero function $f_{\mathsf{zero}}^{(k)}$, there is some (odd) $j \leq k/2$ with $w_j = 1$. Take j as large as possible. The first fact tells us that for all odd $\ell < j$, $w_\ell = 1$. The second fact tells us

that for all odd ℓ between j and $k/2$, $w_\ell = 1$. Thus, all odd ℓ have $w_\ell = 1$. This implies that $f = f_{\mathsf{odd}}^{(k)}$, contrary to the statement of the lemma. $\qquad\square$

Before returning to our main proof, we present one more lemma that uses the language of polymorphisms, and we use this lemma to prove Observation 27, which supports our interpretation of existing literature in the introduction to this paper.

Lemma 26. *Suppose that $k \geq 3$. Let f be a symmetric k-ary Boolean function that is not in $\mathsf{EASY}(k)$. Then either the arity-2 Boolean function AND is not a polymorphism of f or the arity-2 Boolean function OR is not a polymorphism of f (or both).*

Proof. We break the analysis into two cases.

Case 1. There is an index j in the range $1 \leq j \leq k-1$ such that $w_j = 1$ and $w_{j+1} = 0$:
Let $\overline{x}_1$ have ones in the first j positions (and only in those positions) and let $\overline{x}_2$ have ones in positions $2, \ldots, j+1$ (only). Then $f(\overline{x}_1) = f(\overline{x}_2) = 1$. But $\overline{y}_{\mathsf{OR}}(\overline{x}_1, \overline{x}_2))$ has $j+1$ ones so $f(\overline{y}_{\mathsf{OR}}(\overline{x}_1, \overline{x}_2)) = 0$ and OR is not a polymorphism of f.

Case 2. There is an index j in the range $0 \leq j \leq k-2$ such that $w_j = 0$ and $w_{j+1} = 1$:
Let $\overline{x}_1$ have ones in the first $j+1$ positions (and only in those positions) and let $\overline{x}_2$ have ones in positions $2, \ldots, j+2$ (only). Then $f(\overline{x}_1) = f(\overline{x}_2) = 1$. But $\overline{y}_{\mathsf{AND}}(\overline{x}_1, \overline{x}_2))$ has j ones so $f(\overline{y}_{\mathsf{AND}}(\overline{x}_1, \overline{x}_2)) = 0$ and AND is not a polymorphism of f.

If neither Case 1 nor Case 2 applies then f must be one of the four functions $f_{\mathsf{zero}}^{(k)}$, $f_{\mathsf{one}}^{(k)}$, $f_{\mathsf{allzero}}^{(k)}$ and $f_{\mathsf{allone}}^{(k)}$ defined by $f_{\mathsf{zero}}^{(k)}(x_1, \ldots, x_k) = 0$, $f_{\mathsf{one}}^{(k)}(x_1, \ldots, x_k) = 1$,

$$f_{\mathsf{allzero}}^{(k)}(x_1, \ldots, x_k) = \mathbf{1}\{x_1 = \ldots = x_k = 0\},$$

and

$$f_{\mathsf{allone}}^{(k)}(x_1, \ldots, x_k) = \mathbf{1}\{x_1 = \ldots = x_k = 1\}.$$

All four of these functions are in $\mathsf{EASY}(k)$. $\qquad\square$

Observation 27. *Suppose that $\Delta \geq 6$ and $k \geq 3$ and that f is a symmetric k-ary Boolean function that is not in $\mathsf{EASY}(k)$. Then there is no FPRAS for $\#\mathsf{CSP}(\{f, \delta_0, \delta_1\})$ unless $\mathsf{NP} = \mathsf{RP}$.*

Proof. Theorem 24 of [11] gives the result unless (i) f is affine (given by a linear equation over GF_2), or (ii) f is in a set of functions called IM-conj. All symmetric affine functions are in $\mathsf{EASY}(k)$. IM-conj is the same as the class IM_2 studied in [2, 8]. As described in these works, the Galois correspondence between Post's lattice and its dual shows that f is in IM_2 if and only if it has both AND and OR as polymorhisms. Thus, by Lemma 26 this case does not arise. $\qquad\square$

4.2.2 The case $w_0 = 0$

We now return to our proof. By assumption, $k \geq 3$ and f is a symmetric arity-k Boolean function that is not in $\mathsf{EASY}(k)$. Since we are in Case II, the function f supports 2-equality but it does not support pinning-to-0 or pinning-to-1. By Lemma 14, we know that f is self-dual. We are interested in the case $w_0 = 0$ and we know from Lemma 25 that a related decision CSP problem is NP-hard. We wish to use the hardness of the CSP decision problem to show hardness of our bounded-degree counting problem. We will use the fact that f supports equality to introduce degree bounds and also to move to the more restricted hypergraph 2-spin model where repeated variables are not allowed. The following technical lemma is inspired by techniques from [13].

Lemma 28. *Suppose $k > 2$. Let $f \neq f_{\mathsf{zero}}^{(k)}, f_{\mathsf{odd}}^{(k)}$ 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 $\sigma : V(H) \to \{0, 1\}$ with $w_{f;H}(\sigma) > 0$, we have $\sigma(x) = \sigma(y)$.*

Proof. We first prove that there exists a hypergraph $H_0 = (V_0, \mathcal{F}_0)$ such that $Z_{f;H_0} = 0$. To see this, consider the complete k-uniform hypergraph on $2k - 1$ vertices, i.e., $V_0 = \{1, \ldots, 2k - 1\}$ and $\mathcal{F}_0$ is the set of all k-element subsets of V_0. Consider an arbitrary assignment $\sigma : V_0 \to \{0, 1\}$. Under σ, there exist k vertices which have the same spin, w.l.o.g. assume that these vertices are $1, \ldots, k$. Since $w_0 = w_k = 0$ (note that $w_k = 0$ by self-duality), it follows that the hyperedge $\{1, \ldots, k\} \in \mathcal{F}_0$ is not satisfied under σ. Since σ was arbitrary, this proves that $Z_{f;H_0} = 0$. By removing hyperedges of H_0 successively, we obtain a hypergraph $H' = (V', \mathcal{F}')$ such that $Z_{f;H'} = 0$ and for every $e \in \mathcal{F}'$, it holds that $Z_{f;H' \setminus e} > 0$. Since f is not the all-zero function $f_{\mathsf{zero}}^{(k)}$, we can conclude that H' has at least one hyperedge.

Choose $e \in \mathcal{F}'$ and let $S \subseteq e$ be the set of vertices of e that have non-zero degree in $H' \setminus e$. By the minimality of H', we have $S \neq \emptyset$. Let $i = |S|$. Denote the vertices in S by $v_1, \ldots, v_i$ and the vertices in $e \setminus S$ by $v_{i+1}, \ldots, v_k$. Consider i new vertices $u_1, \ldots, u_i \notin V'$ and for each $t \in \{0, \ldots, i\}$, let $e_t = \{u_1, \ldots, u_t, v_{t+1}, \ldots, v_i\} \cup (e \setminus S)$. Further, consider the hypergraphs $H_t = (V_t, \mathcal{F}_t)$ where $V_t = V' \cup \{u_1, \ldots, u_t\}$ and $\mathcal{F}_t = (\mathcal{F}' \setminus \{e\}) \cup \{e_t\}$. Note that $e_0 = e$ and e_i has no vertices that are in other hyperedges of H'. By the minimality of H', we can conclude that $Z_{f;H_i} > 0$. Let j be the smallest integer such that $Z_{f;H_j} > 0$. Then $1 \leq j \leq k$. Also, for every $\sigma : V_j \to \{0, 1\}$ with $w_{f;H_j}(\sigma) > 0$, it must hold that $\sigma(u_j) \neq \sigma(v_j)$ (otherwise, we would have $w_{f;H_{j-1}}(\sigma|_{V_{j-1}}) > 0$, contradicting the fact that $Z_{f;H_{j-1}} = 0$).

Let $H = (V_j \cup \{u'_j\}, \mathcal{F}_j \cup \{e'_j\})$, where $e'_j = \{u_1, \ldots, u_{j-1}, u'_j, v_{j+1}, \ldots, v_i\} \cup (e \setminus S)$. As above, we conclude that $Z_{f;H} > 0$. Also, for every $\sigma : V_j \cup \{u'_j\} \to \{0, 1\}$ with $w_{f;H}(\sigma) > 0$ it must hold that $\sigma(u_j) \neq \sigma(v_j)$ and $\sigma(u'_j) \neq \sigma(v_j)$, so $\sigma(u_j) = \sigma(u'_j)$. So the vertices x and y in the statement of the lemma can be taken to be u_j and u'_j. □

We will now combine Lemmas 25 and 28 to conclude the following.

Lemma 29. *Suppose $k > 2$. Let $f \neq f_{\mathsf{zero}}^{(k)}, f_{\mathsf{odd}}^{(k)}$ 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 $\#\mathsf{Hyper2Spin}(f, \Delta, c)$ is NP-hard.*

Proof. From Lemma 28 we know there is a hypergraph H'' with $Z_{f;H''} > 0$ which has vertices x and y such that for any configuration $\sigma : V(H'') \to \{0, 1\}$ with $w_{f;H''}(\sigma) > 0$, we have $\sigma(x) = \sigma(y)$. Let Z_0 be the contribution to $Z_{f;H''}$ from configurations with $\sigma(x) = \sigma(y) = 0$ and let Z_1 be the contribution to $Z_{f;H''}$ from configurations with $\sigma(x) = \sigma(y) = 1$. By Lemma 23, we have $Z_0 = Z_1$ so, since $Z_0 + Z_1 = Z_{f;H''}$, which is positive, both Z_0 and Z_1 are positive. Let Δ'' be the maximum degree of H'' and let $\Delta_0 = 2\Delta'' + 1$.

Let $\Gamma = \{f\}$. We know from Lemma 25 that it is NP-hard to determine whether $Z_{\Gamma, I} = 0$ given a CSP instance I.

Consider a CSP instance I with variable set V and constraint set $\mathcal{S}$. We will 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 $Z_{f;H'} = 0$. This will imply that determining whether $Z_{f;H'} = 0$ is NP-hard given a k-uniform hypergraph H' with degree at most Δ_0. Hence, for every $\Delta \geq \Delta_0$, $\#\mathsf{Hyper2Spin}(f, \Delta, c)$ is also NP-hard.

The construction is straightforward, apart from the notation. For each $v \in V$, let $n(v)$ be the number of times that variable v is used (taking all of the constraints in $\mathcal{S}$ together). Let $V' = \{(v, j) \mid v \in V, 1 \leq j \leq n(v)\}$. Let $\mathcal{S}'$ be a set of constraints that is identical to $\mathcal{S}$ except that, for each $v \in V$ and $1 \leq j \leq n(v)$, the j'th use of vertex v is replaced with (v, j). Note that the constraints in $\mathcal{S}'$ use each variable in V' exactly once. Thus, they can be viewed as hyperedges of a k-uniform hypergraph.

To build H' we will take the vertices in V' and the hyperedges in $\mathcal{S}'$ but we will add some additional vertices and hyperedges. In particular, for each $v \in V$ and $j \in \{1, \ldots, n(v) - 1\}$ we will take a new copy $H''_{v,j}$ of H''. We will identify the vertex x of $H''_{v,j}$ with (v, j) and the vertex y of $H''_{v,j}$ with $(v, j + 1)$. This completes the construction of the hypergraph H'. Note that the degree of each vertex of H' is at most $2\Delta'' + 1 = \Delta_0$. Also, H' can be efficiently constructed given I.

Now, by the properties of H'', every configuration $\sigma' : V(H') \to \{0, 1\}$ with $w_{f;H'}(\sigma) > 0$ has $\sigma((v, j)) = \sigma((v, j'))$ for every $v \in V$ and every $1 \leq j, j' \leq n(v)$. Also, σ' induces a configuration

$\sigma : V \to \{0,1\}$ by taking $\sigma(v) = \sigma'(v,1) = \cdots = \sigma'(v,n(v))$. Now note that

$$Z_{f;H'} = \sum_{\sigma : V \to \{0,1\}} \left(\prod_{(v_1,\dots,v_k,f) \in \mathcal{S}} f(\sigma(v_1),\dots,\sigma(v_k)) \right) \left(\prod_{v \in V} Z_{\sigma(v)}^{n(v)-1} \right).$$

This is identical to the partition function $Z_{f;I}$ apart from the factors of Z_0 and Z_1, which are both positive. We conclude that $Z_{f;H'}$ is positive if and only if $Z_{f;I}$ is positive. $\qquad\square$

4.2.3   The case $w_0 = 1$

We now consider the case where $w_0 = 1$. We will use the following corollary of Lemma 18 (tailored to the case where f is self-dual).

Corollary 30. *Let f be an arity-k symmetric Boolean formula that is self-dual. Let H be a hypergraph, let $\mathcal{V}$ be admissible for H (with $V_0 = V_1 = \emptyset$) and let x and y be vertices of H. Suppose that*

$$0 < \mu_{f;H}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 0) < \mu_{f;H}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = 0, \sigma_y = 1). \tag{33}$$

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

Proof. For $i,j \in \{0,1\}$, let $\mu_{ij} = \mu_{f;H}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = i, \sigma_y = j)$. By Lemma 23, we have $\mu_{11} = \mu_{00}$ and $\mu_{01} = \mu_{10}$. It follows that condition (28) is equivalent to condition (33), so the corollary follows by applying Lemma 18. $\qquad\square$

We use Corollary 30 to show the following lemma.

Lemma 31. *Let $f : \{0,1\}^k \to \{0,1\}$ 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*

$$0 < w_0 + w_{t_2} + w_{2t_1} + w_{2t_1+t_2} < 2(w_{t_1} + w_{t_1+t_2}). \tag{34}$$

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

Proof. We will apply Corollary 30. 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 $\{x_{2t_1+t_2+1}, \dots, x_k\}$ (see Figure 6). Let $t_3 := k - 2t_1 - t_2$ and set $x := x_1$ and $y := x_{t_1+1}$. Note that the condition $2t_1 + t_2 < k$ is equivalent to $t_3 > 0$. By the assumptions, we also have that $t_3 \geq 0$. Now define

$$Z_{00}^{\mathrm{ex}} = w_0 + w_{t_2} + w_{2t_1} + w_{2t_1+t_2},$$
$$Z_{01}^{\mathrm{ex}} = 2(w_{t_1} + w_{t_1+t_2}).$$

First, suppose $t_2 > 0$ and $t_3 > 0$. In this case, we have (using self-duality in the first equality in each line)

$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 0) \propto w_0 + w_{t_2} + w_{t_3} + w_{t_2+t_3} = w_0 + w_{t_2} + w_{2t_1+t_2} + w_{2t_1} = Z_{00}^{\mathrm{ex}},$$

$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = 0, \sigma_y = 1) \propto w_{t_1} + w_{t_1+t_2} + w_{t_1+t_3} + w_{t_1+t_2+t_3} = w_{t_1} + w_{t_1+t_2} + w_{t_1+t_2} + w_{t_1} = Z_{01}^{\mathrm{ex}}.$$

It is now immediate that inequality (34) is equivalent to condition (33) in Corollary 30, from which the result follows.

The proofs for the remaining cases for the values of t_2, t_3 are completely analogous. If $t_2 > 0$ and $t_3 = 0$ then $w_0 = w_{2t_1+t_2}$ and $w_{t_2} = w_{2t_1}$ so we have

$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 0) \propto w_0 + w_{t_2} = \tfrac{1}{2} Z_{00}^{\mathrm{ex}},$$
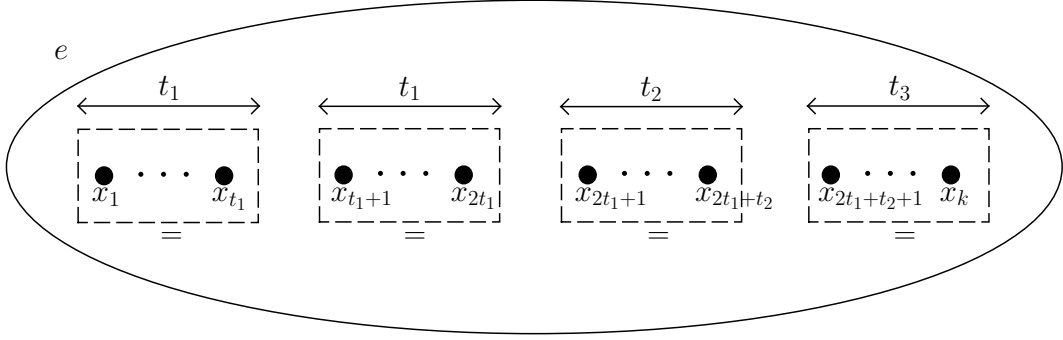
Figure 6: The hypergraph used in the proof of Lemma 31. The hypergraph has just one hyperedge $e = \{x_1, x_2, \ldots, x_k\}$. We partition the hyperedge into sets of sizes t_1, t_1, t_2, t_3 and look at the conditional distribution where the spins of the vertices in each set are equal.

$$\mu_{f;e}^{\text{cond}(\mathcal{V})}(\sigma_x = 0, \sigma_y = 1) \propto w_{t_1} + w_{t_1+t_2} = \tfrac{1}{2}Z_{01}^{\text{ex}}.$$

If $t_2 = t_3 = 0$ then also $w_{t_2} = w_0$ and $w_{t_1+t_2} = w_{t_1}$, so

$$\mu_{f;e}^{\text{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 0) \propto w_0 = \tfrac{1}{4}Z_{00}^{\text{ex}},$$
$$\mu_{f;e}^{\text{cond}(\mathcal{V})}(\sigma_x = 0, \sigma_y = 1) \propto w_{t_1} = \tfrac{1}{4}Z_{01}^{\text{ex}}.$$

Finally, if $t_2 = 0$ and $t_3 > 0$ then

$$\mu_{f;e}^{\text{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 0) \propto w_0 + w_{t_3} = w_0 + w_{2t_1} = \tfrac{1}{2}Z_{00}^{\text{ex}},$$
$$\mu_{f;e}^{\text{cond}(\mathcal{V})}(\sigma_x = 0, \sigma_y = 1) \propto w_{t_1} + w_{t_1+t_3} = 2w_{t_1} = \tfrac{1}{2}Z_{01}^{\text{ex}}.$$

Thus, in each of the above cases, the result follows by Corollary 30. □

We will also use the following inequality for binomial coefficients (which is slightly stronger than the well-known log-concavity property $\binom{n}{i-1}\binom{n}{i+1} \leq \binom{n}{i}^2$).

Lemma 32. *For all $n \geq i \geq 2$, it holds that*

$$\left[1 + \binom{n}{i}\right]\left[1 + \binom{n}{i-2}\right] \leq \binom{n}{i-1}^2.$$

Equality holds iff $n = 2, i = 2$.

Proof. For $i = 2$ the inequality becomes $n^2 - n + 2 \leq n^2$ which holds for all $n \geq 2$ (equality only if $n = 2$). The same argument applies for $i = n$. Thus, we may assume $n - 1 \geq i \geq 3$. For $i = 3$ the inequality becomes

$$\left(1 + \frac{n(n-1)(n-2)}{6}\right)(n+1) \leq \frac{n^2(n-1)^2}{4} \Leftrightarrow n^4 - 2n^3 + 5n^2 - 16n - 12 \geq 0.$$

The latter can easily be verified that it holds strictly for all $n \geq 4$. The same argument applies for $i = n - 1$. Thus, we may assume $n - 2 \geq i \geq 4$.

For all such values of n, i, the inequality follows by summing

$$1 + \binom{n}{i} + \binom{n}{i-2} < \frac{1}{i-1}\binom{n}{i}\binom{n}{i-2}, \tag{35}$$

$$\frac{i}{i-1}\binom{n}{i}\binom{n}{i-2} < \binom{n}{i-1}^2. \tag{36}$$

Inequality (36) is equivalent to $n-i+1 < n-i+2$ which is trivially true. To see (35), we first rewrite it in the equivalent form

$$\left[\binom{n}{i} - (i-1)\right]\left[\binom{n}{i-2} - (i-1)\right] > i(i-1),$$

which follows from the inequalities $\binom{n}{i} \geq 2(i+1)$ and $\binom{n}{i-2} \geq 2(i-1)$ (both are special cases of $\binom{n}{j} \geq 2(j+1)$ which holds for $j \geq 2$ and $n \geq j+2$). □

We can now do the proof for this case. Recall that $k \geq 3$ and that f is a symmetric self-dual arity-k Boolean function that is not in $\mathsf{EASY}(k)$. The function f supports 2-equality but does not support pinning-to-0 or pinning-to-1. We are assuming that $w_0 = 1$. Our goal is to show that for all sufficiently large Δ, there exists $c > 1$ such that the approximation problem $\#\mathsf{Hyper2Spin}(f, \Delta, c)$ 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 $f = f_{\mathsf{EQ}}^{(k)}$).

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 $\#\mathsf{Hyper2Spin}(f, \Delta, c)$ is NP-hard.*

Proof. Let r be the smallest positive integer with $ri \leq k$ and $w_{ri} = 0$. Since $w_i = 1$, we have $r \geq 2$. We next check that the conditions of Lemma 31 are satisfied with $t_1 = i$ and $t_2 = (r-2)i$. First, we clearly have $t_1 \geq 1$, $t_2 \geq 0$, $2t_1 + t_2 = ri \leq k$. Moreover, the left-most inequality in (34) is true since $w_0 = 1$. Finally, the right-most inequality in (34) also holds, since

$$w_{t_1} + w_{t_1+t_2} = w_i + w_{(r-1)i} = 2,$$
$$w_0 + w_{t_2} + w_{2t_1} + w_{2t_1+t_2} = w_0 + w_{(r-2)i} + w_{2i} + w_{ri} \leq 3,$$

where we used that $w_{(r-1)i} = 1$ and $w_{ri} = 0$ (by the choice of r). Applying Corollary 30 yields the claim. □

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 $\#\mathsf{Hyper2Spin}(f, \Delta, c)$ is NP-hard.*

Proof. Let r be the smallest positive integer that is less than k, is not divisible by i and has $w_r = 1$. By the choice of i, we have $r > i$. By self-duality, r is not in $\{k-i+1, \ldots, k-1\}$ so $r \leq k-i$. We next check that the conditions of Lemma 31 are satisfied with $t_1 = i$ and $t_2 = r - i$. We clearly have $t_1 \geq 1$, $t_2 > 0$, $2t_1 + t_2 = r + i \leq k$. Moreover, the left-most inequality in (34) is true since $w_0 = 1$. Finally, the right-most inequality in (34) also holds, since

$$w_{t_1} + w_{t_1+t_2} = w_i + w_r = 2,$$
$$w_0 + w_{t_2} + w_{2t_1} + w_{2t_1+t_2} = w_0 + w_{r-i} + w_{2i} + w_{r+i} \leq 3,$$

where we used that $w_r = 1$ and $w_{r-i} = 0$ (by the choice of r). Applying Corollary 30 yields the claim. □

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

- $w_0 = w_k = 1$. (We know that this is true because this is the case that we are dealing with in the current section, Section 4.2.3).

- The positive integers $\ell \in \{1, \ldots, k-1\}$ with $w_\ell = 1$ are precisely the multiples of i. This follows from Claims 33 and 34.

- k is a multiple of i. Suppose instead that $k = mi + u$ for some non-negative integer m and some integer $u \in \{1, \ldots, i-1\}$. 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 $w_u = 1$. But this contradicts the choice of i.

- $i > 2$. If $i = 1$ then f is the all-one function $f = f_{\mathsf{one}}^{(k)}$. If $i = 2$ then k is even since it is a multiple of i. Then f is the easy function $f = f_{\mathsf{even}}^{(k)}$.

- $2i \leq k$. We know that k is a multiple of i, but if k is equal to i, then f is the equality function $f = f_{\mathsf{EQ}}^{(k)}$.

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

$$
\begin{aligned}
\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 0) &\propto \sum_{\ell=0}^{2i-2} \binom{2i-2}{\ell}(w_\ell + w_{\ell+k-2i}), \\
\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = 0, \sigma_y = 1) &\propto \sum_{\ell=0}^{2i-2} \binom{2i-2}{\ell}(w_{\ell+1} + w_{\ell+k-2i+1}).
\end{aligned}
\tag{37}
$$

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, (37) becomes

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

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 $1 + \binom{2i-2}{i} < \binom{2i-2}{i-1}$, so that Corollary 30 yields that for all sufficiently large Δ, there exists $c > 1$ such that $\#\mathsf{Hyper2Spin}(f, \Delta, c)$ is NP-hard. To see that $1 + \binom{2i-2}{i} < \binom{2i-2}{i-1}$, let $n = 2i - 2$. Note that $2 < i < n$. Then Lemma 32 gives

$$
\left[1 + \binom{2i-2}{i}\right]\left[1 + \binom{2i-2}{i-2}\right] < \binom{2i-2}{i-1}^2.
$$

The desired inequality follows after observing that $\binom{2i-2}{i-2} = \binom{2i-2}{i}$ and simplifying.

4.3 Case III

Throughout this section, we will assume that f supports pinning-to-0 (the case that f supports pinning-to-1 is analogous by swapping the spins 0 and 1). Our goal is to show that there exists Δ_0 such that for all $\Delta \geq \Delta_0$, there exists $c > 1$ such that $\#\mathsf{Hyper2Spin}(f, \Delta, c)$ is NP-hard.

We have the following analogue of Lemma 12 (tailored to the case where f supports pinning-to-0).

Lemma 35. *Assume that f supports pinning-to-0. Let $H = (V, \mathcal{F})$ and let $V_0 \subseteq V$ be admissible for the hypergraph H, i.e., $\mu_{f;H}(\sigma_{V_0} = \mathbf{0}) > 0$. With $\mathcal{V} = (V_0)$, recall that $\mu_{f;H}^{\mathrm{cond}(\mathcal{V})}(\cdot) = \mu_{f;H}(\cdot \mid \sigma_{V_0} = \mathbf{0})$.*

1. *If there exists a vertex v in H such that $\mu_{f;H}^{\mathrm{cond}(\mathcal{V})}(\sigma_v = 1) > \mu_{f;H}^{\mathrm{cond}(\mathcal{V})}(\sigma_v = 0)$, then f supports pinning-to-1.*

2. *If there exists a subset S of V such that $\mu_{f;H}^{\mathrm{cond}(\mathcal{V})}(\sigma_S = \mathbf{1}) = \mu_{f;H}^{\mathrm{cond}(\mathcal{V})}(\sigma_S = \mathbf{0}) = 1/2$, then f supports equality.*

Proof. For Item 1, choose ε in the range $0 < \varepsilon < |\mu_{f;H}^{\mathrm{cond}(\mathcal{V})}(\sigma_v = 1) - \mu_{f;H}^{\mathrm{cond}(\mathcal{V})}(\sigma_v = 0)|/2$. By Lemma 17, there exists a hypergraph H' with vertex set $V' \supseteq V$, such that for $s \in \{0, 1\}$, it holds that

$$|\mu_{f;H'}(\sigma_v = s) - \mu_{f;H}^{\mathrm{cond}(\mathcal{V})}(\sigma_v = s)| \leq \varepsilon.$$

It follows that $\mu_{f;H'}(\sigma_v = 1) > \mu_{f;H'}(\sigma_v = 0)$, so by Lemma 12, f supports pinning-to-1.

For Item 2, for $\varepsilon > 0$, apply Lemma 17 to conclude that there exists a hypergraph H' which ε-realises $|S|$-equality. By Lemma 10, we obtain that f supports equality. □

We remark here that, as in the analysis for Case I in Section 4.1, rather than formally defining V_0, in each of the subcases which we consider, we will typically say, e.g., pin vertices x_1, x_2 to 0 instead of specifying V_0 as $V_0 = \{x_1, x_2\}$. Also, unless otherwise stated, we will have $\mathcal{V} = (V_0)$.

Our first application of Lemma 35 is to show that if $w_0 = 0$, then f also supports pinning-to-1. So assume that $w_0 = 0$. Let i be the smallest index such that $w_i \neq 0$ (if such an index does not exist, then $f = f_{\mathsf{zero}}^{(k)}$), so that $1 \leq i \leq k-1$ (if $i = k$ then $f = f_{\mathsf{allone}}^{(k)}$). Let $e = \{x_1, x_2, \ldots, x_k\}$, pin $x_{i+1}, \ldots, x_k$ to 0 and set $x := x_1$.

$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = 0) \propto \sum_{\ell=0}^{i-1} \binom{i-1}{\ell} w_\ell = 0, \quad \mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = 1) \propto \sum_{\ell=0}^{i-1} \binom{i-1}{\ell} w_{\ell+1} = w_i = 1.$$

Thus, if $w_0 \neq 1$, by Lemma 35, we obtain that f also supports pinning-to-1, in which case we fall back in Case I.

Thus for the rest of the proof in this section we will assume that $w_0 = 1$ (otherwise, as we showed above, we fall back in Case I, where f supports both pinning-to-0 and pinning-to-1). Let $i > 0$ be the minimum index i such that $w_i = 1$ (we may assume that such an index exists, otherwise $f = f_{\mathsf{allzero}}^{(k)}$). We consider the cases $i = 1$ and $i \geq 2$ separately.

4.3.1 The case $i \geq 2$

The setting of this section is $w_0 = w_i = 1$ and $w_1 = \ldots = w_{i-1} = 0$ (and f supports pinning-to-0). We may assume that $k > i$ (otherwise $f = f_{\mathsf{EQ}}^{(k)}$).

Lemma 36. *If $k > i \geq 2$, then the function f supports equality.*

Proof. We show that f supports i-equality, so by Lemma 10, we obtain that f supports equality.

Let $e = \{x_1, \ldots, x_k\}$ and pin vertices $x_{i+1}, \ldots, x_k$ to 0. Denote $S := \{x_1, \ldots, x_i\}$. We have

$$\mu_{f;H}^{\mathrm{cond}(\mathcal{V})}(\sigma_S = \mathbf{0}) \propto w_0 = 1, \quad \mu_{f;H}^{\mathrm{cond}(\mathcal{V})}(\sigma_S = \mathbf{1}) \propto w_i = 1, \quad \mu_{f;H}^{\mathrm{cond}(\mathcal{V})}(\neg\sigma_S^{\mathsf{eq}}) \propto 0.$$

It follows that $\mu_{f;H}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = 0) = \mu_{f;H}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = 1) = 1/2$, so the result follows from Lemma 35. □

Case IIIa There exists $j > i$ such that $w_j = 1$.

Let j be the minimum such index. We have the following lemma.

Lemma 37. *If $j \neq 2i$, the function f supports also pinning-to-1.*

Proof. Suppose first that $j < 2i$. Let $e = \{x_1, \ldots, x_k\}$, pin $x_{j+1}, \ldots, x_k$ to zero and set $x := x_1$. We have:

$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = 0) \propto \sum_{\ell=0}^{j-1} \binom{j-1}{\ell} w_\ell = w_0 + \binom{j-1}{i} w_i = 1 + \binom{j-1}{i},$$

$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = 1) \propto \sum_{\ell=0}^{j-1} \binom{j-1}{\ell} w_{\ell+1} = \binom{j-1}{i-1} w_i + w_j = 1 + \binom{j-1}{i-1}.$$

From $2i > j$, we have $i - 1 \geq (j-1)/2$, so that $\mu_{f;e}^{\text{cond}(\mathcal{V})}(\sigma_x = 1) > \mu_{f;e}^{\text{cond}(\mathcal{V})}(\sigma_x = 0)$, so Lemma 35 yields that f also supports pinning-to-1.

We next consider the more difficult case $j > 2i$. Let H be the hypergraph with vertex set

$$\{x_1, \ldots, x_k\} \cup \{y_{i+1}, \ldots, y_k\} \cup \{z_{i+1}, \ldots, z_k\}$$

and hyperedges $\{e_X, e_Y, e_Z\}$ with

$$e_X = \{x_1, \ldots, x_j, \ldots, x_k\}, \ e_Y = \{x_1, .., x_i, y_{i+1}, \ldots, y_j, \ldots, y_k\},$$
$$e_Z = \{x_{i+1}, \ldots, x_{2i}, z_{i+1}, \ldots, z_j, \ldots, z_k\}.$$

We will pin to zero the vertices in $V_0 := \{x_{j+1}, \ldots, x_k\} \cup \{y_{j+1}, \ldots, y_k\} \cup \{z_{j+1}, \ldots, z_k\}$ (thus, in the conditional distribution, each of the hyperedges e_X, e_Y, e_Z has j vertices). Also, we will force equality among the sets $V_2 := \{x_1, x_2, ..., x_{2i}\}$, $V_3 := \{x_{2i+1}, \ldots, x_j\}$, $V_4 := \{y_{i+1}, \ldots, y_j\}$ and $V_5 := \{z_{i+1}, \ldots, z_j\}$. Let $x := x_1$.

Consider $\mu_{f;H}^{\text{cond}(\mathcal{V})}(\cdot) = \mu_{f;H}(\cdot \mid \sigma_{V_0} = \mathbf{1}, \sigma_{V_2}^{\text{eq}}, \sigma_{V_3}^{\text{eq}}, \sigma_{V_4}^{\text{eq}}, \sigma_{V_5}^{\text{eq}})$. We will show that

$$\mu_{f;H}^{\text{cond}(\mathcal{V})}(\sigma_x = 1) > \mu_{f;H}^{\text{cond}(\mathcal{V})}(\sigma_x = 0), \tag{38}$$

so the result will follow by Lemma 35. To see (38), assume first that $\sigma_x = 0$. In the conditional distribution $\mu_{f;H}^{\text{cond}(\mathcal{V})}(\cdot)$, it then holds that $\sigma_{V_2} = \mathbf{0}$. From $j > 2i$, there is only one way to satisfy e_Y, e_Z by setting $\sigma_{V_4} = \sigma_{V_5} = \mathbf{0}$. Further, there are at most two ways to satisfy e_X (by setting σ_{V_3} to be $\mathbf{0}$ and possibly $\mathbf{1}$, if $w_{j-2i} = 1$). Thus, the total weight of configurations with $\sigma_x = 0$ is at most two. On the other hand, if $\sigma_x = 1$ and hence $\sigma_{V_2} = \mathbf{1}$, from $w_j = 1$, there is at least one way to satisfy to e_X (by setting $\sigma_{V_3} = \mathbf{1}$). Further, from $w_i = w_j = 1$, there are two ways to satisfy each of e_Y, e_Z (any combination of $\{\sigma_{V_4}, \sigma_{V_5}\} \in \{\mathbf{0}, \mathbf{1}\}$ works). Thus, the total weight of configurations with $\sigma_x = 1$ is at least 4. We thus obtain that $\mu_{f;H}^{\text{cond}(\mathcal{V})}(\sigma_x = 1) > \mu_{f;H}^{\text{cond}(\mathcal{V})}(\sigma_x = 0)$, as wanted.

This concludes the proof. $\qquad\square$

By Lemma 37, we may thus assume that $j = 2i$ (otherwise we fall back into Case I, since f supports both pinning-to-0 and pinning-to-1). That is, we have $w_0 = w_i = w_{2i}$ and for $j \leq 2i$ with $j \neq 0, i, 2i$, it holds that $w_j = 0$.

If $i > 2$, by pinning $k - 2i$ variables to 0, we (approximately) get a non-trivial self-dual constraint with arity $2i$ (since $w_0 = w_i = w_{2i} = 1$) and hence the proof may be completed analogously to Case II. In particular, let $e = \{x_1, \ldots, x_k\}$ and pin $x_{2i+1}, \ldots, x_k$ to 0. Denote $x := x_{2i-1}$, $y := x_{2i}$. Note that if σ is such that $\sigma_x = \sigma_y = 0$, σ has non-zero weight iff exactly 0 or i vertices from $x_1, \ldots, x_{2i-2}$ have spin 1 under σ. Similarly, if σ is such that $\sigma_x = \sigma_y = 1$, σ has non-zero weight iff exactly $i - 2$ or $2i - 2$ vertices from $x_1, \ldots, x_{2i-2}$ have spin 1 under σ. On the other hand, if σ is such that $\sigma_x = 0$ and $\sigma_y = 1$ (or vice versa), σ has non-zero weight iff exactly $i - 1$ vertices from $x_1, \ldots, x_{2i-2}$ have spin 1. It follows that for $\{s_1, s_2\} \in \{0, 1\}$, we have $\mu_{f;e}^{\text{cond}(\mathcal{V})}(\sigma_x = s_1, \sigma_y = s_2) \propto Z_{s_1 s_2}$ where

$$Z_{00} = \binom{2i-2}{0} + \binom{2i-2}{i}, \quad Z_{11} = \binom{2i-2}{i-2} + \binom{2i-2}{2i-2}, \quad Z_{01} = Z_{10} = \binom{2i-2}{i-1}.$$

Note that $Z_{00} = Z_{11}$ and $Z_{00} Z_{11} < Z_{01} Z_{10}$ by Lemma 32 (since $i > 2$), which gives that $Z_{00} = Z_{11} < Z_{01} = Z_{10}$. Lemma 18 thus implies that, for all sufficiently large Δ, there exists $c > 1$ such that #Hyper2Spin(f, Δ, c) is NP-hard.

Thus, it remains to consider the case $i = 2$, i.e., $w_0 = w_2 = w_4 = 1$ and $w_1 = w_3 = 0$. Let j' be the largest integer such that $w_0 = w_2 = \ldots = w_{2j'} = 1$ and $w_1 = \ldots = w_{2j'-1} = 0$. We may assume that $2j' < k$, otherwise $f = f_{\text{even}}^{(k)}$.

Assume first $w_{2j'+1} = 1$. Let $e = \{x_1, x_2, \ldots, x_k\}$, pin $x_{2j'+2}, \ldots, x_k$ to 0 and set $x := x_1$. We have:

$$\mu_{f;e}^{\text{cond}(\mathcal{V})}(\sigma_x = 0) \propto \sum_{\ell=0}^{2j'} \binom{2j'}{\ell} w_\ell = 2^{2j'-1}, \quad \mu_{f;e}^{\text{cond}(\mathcal{V})}(\sigma_x = 1) \propto \sum_{\ell=0}^{2j'} \binom{2j'}{\ell} w_{\ell+1} = 2^{2j'-1} + 1.$$

By Lemma 35, we obtain that f also supports pinning-to-1, in which case we fall back in Case I.

Thus, we may assume that $w_{2j'+1} = 0$. We may further assume that $2j' + 1 < k$, otherwise $f = f_{\mathsf{even}}^{(k)}$. By the choice of j', it follows that $w_{2j'+2} = 0$. As before, let $e = \{x_1, x_2, \ldots, x_k\}$, pin $x_{2j'+3}, \ldots, x_k$ to 0, set $x := x_1$, $y := x_2$. We have:

$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 0) \propto \sum_{\ell=0}^{2j'} \binom{2j'}{\ell} w_\ell = 2^{2j'-1},$$

$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = 0, \sigma_y = 1) \propto \sum_{\ell=0}^{2j'} \binom{2j'}{\ell} w_{\ell+1} = 2^{2j'-1},$$

$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 1) \propto \sum_{\ell=0}^{2j'} \binom{2j'}{\ell} w_{\ell+2} = 2^{2j'-1} - 1.$$

Once again, we obtain from Lemma 18 that for all sufficiently large Δ, there exists $c > 1$ such that $\#\mathsf{Hyper2Spin}(f, \Delta, c)$ is NP-hard.

Case IIIb There does not exist $j > i$ such that $w_j = 1$. Then we have $w_0 = w_i = 1$ and $w_j = 0$ for all $j \neq 0, i$ (recall also that $i \geq 2$). We may assume that $i < k$, otherwise $f = f_{\mathsf{EQ}}^{(k)}$. We consider separately the cases $i = 2$ and $i > 2$.

Consider first the case $i > 2$. Let $e = \{x_1, x_2, \ldots, x_k\}$, pin $x_{i+2}, \ldots, x_k$ to zero, force equality between $x_3, \ldots, x_{i+1}$ and set $x := x_1, y := x_2$. We have:

$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = 0, \sigma_y = 0) \propto w_0 + w_{i-1} = 1,$$
$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = 0, \sigma_y = 1) \propto w_1 + w_i = 1,$$
$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = 1, \sigma_y = 1) \propto w_2 + w_{i+1} = 0.$$

From Lemma 18, we obtain that for all sufficiently large Δ, there exists $c > 1$ such that $\#\mathsf{Hyper2Spin}(f, \Delta, c)$ is NP-hard.

Finally, we consider the case $i = 2$. We may assume that $k \geq 4$, otherwise $f = f_{\mathsf{even}}^{(3)}$. Let $e = \{x_1, x_2, \ldots, x_k\}$, pin $x_5, \ldots, x_k$ to zero and set $x := x_1, y := x_2$ (note that x_3, x_4 are "free"). We have:

$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = 0, \sigma_y = 0) \propto \binom{2}{0} w_0 + \binom{2}{2} w_2 = 2,$$

$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = 0, \sigma_y = 1) \propto \binom{2}{1} w_2 = 2,$$

$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = 1, \sigma_y = 1) \propto \binom{2}{0} w_2 = 1.$$

From Lemma 18, we obtain that for all sufficiently large Δ, there exists $c > 1$ such that $\#\mathsf{Hyper2Spin}(f, \Delta, c)$ is NP-hard.

4.3.2 The case $i = 1$

In this case, we begin with the assumption that $w_0 = w_1 = 1$ (and f supports pinning-to-0).

Let j be the minimum index $\ell > 1$ such that $w_\ell = 1$. Let $e = \{x_1, x_2, \ldots, x_k\}$. If $j \geq 3$ or such a j does not exist, pin $x_3, \ldots, x_k$ to 0, and set $x := x_1, y := x_2$. We have:

$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 0) \propto w_0 = 1,$$
$$\mu_{f;e}^{\mathrm{cond}(\mathcal{V})}(\sigma_x = 0, \sigma_y = 1) \propto w_1 = 1,$$

$$\mu_{f;e}^{\text{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 1) \propto w_2 = 0.$$

It follows by Lemma 18 that for all sufficiently large Δ, there exists $c > 1$ such that $\#\mathsf{Hyper2Spin}(f, \Delta, c)$ is NP-hard.

If $j = 2$, let j' be the first index $j' > j$ such that $w_{j'} = 0$. We may assume that j' exists otherwise $f = f_{\text{one}}^{(k)}$. We have $j' \geq 3$. Let $e = \{x_1, x_2, \ldots, x_k\}$. Pin $x_{j'+1}, \ldots, x_k$ to 0, set $x := x_1$, $y := x_2$. We have:

$$\mu_{f;e}^{\text{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 0) \propto \sum_{\ell=0}^{j'-2} \binom{j'-2}{\ell} w_\ell = 2^{j'-2},$$

$$\mu_{f;e}^{\text{cond}(\mathcal{V})}(\sigma_x = 0, \sigma_y = 1) \propto \sum_{\ell=0}^{j'-2} \binom{j'-2}{\ell} w_{\ell+1} = 2^{j'-2},$$

$$\mu_{f;e}^{\text{cond}(\mathcal{V})}(\sigma_x = \sigma_y = 1) \propto \sum_{\ell=0}^{j'-2} \binom{j'-2}{\ell} w_{\ell+2} = 2^{j'-2} - 1.$$

From Lemma 18, for all sufficiently large Δ, there exists $c > 1$ such that $\#\mathsf{Hyper2Spin}(f, \Delta, c)$ is NP-hard.

References

[1] I. Bezáková, A, Galanis, L.A. Goldberg, H. Guo and D, Štefankovič, Approximation via Correlation Decay when Strong Spatial Mixing Fails, *ICALP,* 2016.

[2] Elmar Böhler, Nadia Creignou, Matthias Galota, Steffen Reith, Henning Schnoor, and Heribert Vollmer. Boolean circuits as a data structure for boolean functions: Efficient algorithms and hard problems. *Logical Methods in Computer Science*, 8(3), 2010.

[3] 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.

[4] Jin-Yi Cai, Xi Chen, Heng Guo, and Pinyan Lu. Inapproximability after uniqueness phase transition in two-spin systems. In *Combinatorial Optimization and Applications - 6th International Conference, COCOA 2012, Banff, AB, Canada, August 5-9, 2012. Proceedings*, pages 336–347, 2012.

[5] Jin-Yi Cai, Pinyan Lu, and Mingji Xia. The complexity of complex weighted boolean #CSP. *J. Comput. Syst. Sci.*, 80(1):217–236, 2014.

[6] Hubie Chen. A rendezvous of logic, complexity, and algebra. *ACM Comput. Surv.*, 42(1), 2009.

[7] Nadia Creignou and Miki Hermann. Complexity of generalized satisfiability counting problems. *Inf. Comput.*, 125(1):1–12, 1996.

[8] Nadia Creignou, Phokion G. Kolaitis, and Bruno Zanuttini. Preferred representations of boolean relations. *Electronic Colloquium on Computational Complexity (ECCC)*, 119, 2005.

[9] Víctor Dalmau and Daniel K. Ford. Generalized satisfability 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.

[10] Martin E. Dyer, Alan M. Frieze, and Mark Jerrum. On counting independent sets in sparse graphs. *SIAM J. Comput.*, 31(5):1527–1541, 2002.

[11] Martin E. Dyer, Leslie Ann Goldberg, Markus Jalsenius, and David Richerby. The complexity of approximating bounded-degree boolean #CSP. *Inf. Comput.*, 220:1–14, 2012.

[12] 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.

[13] 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.

[14] 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.

[15] 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.

[16] Michael Luby and Eric Vigoda. Fast convergence of the Glauber dynamics for sampling independent sets. *Random Struct. Algorithms*, 15(3-4):229–241, 1999.

[17] 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.

[18] 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.

[19] Alistair Sinclair, Piyush Srivastava, and Marc Thurley. Approximation algorithms for two-state anti-ferromagnetic 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.

[20] Allan Sly. Computational transition at the uniqueness threshold. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 287–296, 2010.

[21] Allan Sly and Nike Sun. Counting in two-spin models on d-regular graphs. *Ann. Probab.*, 42(6):2383–2416, 2014.

[22] 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.

[23] Yitong Yin and Jinman Zhao. Counting hypergraph matchings up to uniqueness threshold. *CoRR*, abs/1503.05812, 2015.