

DECOUPLING OF CLUSTERS IN INDEPENDENT SETS IN A PERCOLATED HYPERCUBE

MRIGANKA BASU ROY CHOWDHURY, SHIRSHENDU GANGULY AND VILAS WINSTEIN

ABSTRACT. Independent sets are sets of vertices in a graph which contain no neighbors, and are well-studied in statistical mechanics as well as computer science, owing to their antiferromagnetic nature as well as their ability to encode problems involving hard constraints. Of particular interest is the setting of the nearest-neighbor hypercube $Q_d = \{0, 1\}^d$, where understanding the number of independent sets was the original motivator for Sapozhenko's famous *graph container method* [15]. A modern perspective on constraint satisfaction problems and Gibbs measures, both of which feature independent sets as a prime example, is to consider the effect of disorder. This was first done by Kronenberg and Spinka [13] who initiated the study of independent sets in a *percolated hypercube* $Q_{d,p}$ obtained by removing each edge from Q_d with probability $1 - p$ independently.

Recent work by the authors [6] resolved multiple open questions posed by [13] regarding the distribution of the number of independent sets in $Q_{d,p}$ for $p \geq \frac{2}{3}$, and also elucidated the structure of a typical independent set via a sampling algorithm. These results relied heavily on the relatively simple behavior of the model for large enough p , but as p decreases, typical independent sets become larger and feature more intricate clustering behavior. In the present article we overcome many of the challenges presented by this phenomenon and analyze the model for a wider range of p . In particular, for all $p > 0.465$ we obtain a sharp in-probability approximation for the number of independent sets in $Q_{d,p}$ in terms of explicit random variables, as well as provide a sampling algorithm. Note that this shows, curiously, that $p = \frac{1}{2}$ does not pose as a natural barrier for this problem unlike in many other problems such as the existence of isolated vertices, existence of a perfect matching, and Hamiltonicity where it appears as a point of a phase transition.

A key contribution of this work is the introduction of a new probabilistic framework to handle the clustering behavior for these low values of p . Although our analysis is restricted to $p > 0.465$, our arguments are expected to be helpful for studying this model at even lower values of p , and possibly for other related problems. In addition, although we only consider the total number of independent sets (corresponding to sampling independent sets uniformly) our techniques should readily generalize to the *partition function* of the hard-core model.

CONTENTS

1. Introduction	2
2. Idea and proof of main result	8
3. Polymer decomposition	14
4. Bounding the size of polymers	20
5. The regime with no dimers	24
6. When dimers begin to appear	27
7. Distributional limits	46
8. Sampling	51
Appendix A. Moment calculations	57
Appendix B. Entropy approximations for binomials	62
References	65

1. INTRODUCTION

This article is our second work investigating independent sets in a disordered hypercube. Referring the reader to the first paper [6] for a more detailed account of the relevant background, in this introduction we will be brief, simply introducing the pertinent objects and giving a short record of immediately adjacent results.

We begin with the notion of independent sets in a graph $G = (V, E)$ which are sets of vertices containing no neighbors. Independent sets are of interest in statistical mechanics where they arise in models of hard-core lattice gases, and the constraint that no two vertices can be neighbors can be viewed as *antiferromagnetic*, leading to connections with other models of interest such as the antiferromagnetic Ising model. As is common in statistical mechanics, it is useful to consider a suitable *Gibbs measure* on the collection of independent sets in a graph. This is known as the *hard-core model*, with a parameter $\lambda > 0$ called the *fugacity*. Given G , the hard-core model assigns to any independent set $I \subseteq V$ the probability

$$\frac{\lambda^{|I|}}{\mathcal{Z}},$$

where $|I|$ denotes the size of I and $\mathcal{Z}$, the normalizing constant, is called the partition function.

While the problem of independent sets can be studied for any graph, a particularly interesting case is when the base graph is the high dimensional lattice given by the hypercube $Q_d := \{0, 1\}^d$, endowed with the nearest-neighbor graph structure, where two vertices u and v are neighbors if and only if they differ at exactly one bit. Independent sets in Q_d correspond to error-detecting binary codes of distance 2, and this was the initial motivation for studying such sets in the classic work of Korshunov and Sapozhenko [12]. The largest independent sets are Even and Odd, the sets of vertices with an even or odd number of 1s in their binary-string representation. Indeed, Q_d is bipartite, and Even and Odd form a bipartition. These maximal independent sets correspond to the *parity bit* code used at the hardware level in most modern computers.

More recently, the study of independent sets in bipartite graphs like Q_d or subgraphs thereof has become especially popular in theoretical computer science, due to the well-known #BIS problem of counting independent sets in bipartite graphs. This problem is known to be as difficult to solve exactly as the corresponding problem for general graphs (both are #P-complete [18]), but whereas the general problem is known also to be intractible to approximation via randomized algorithms [17], it is still unknown whether or not an efficient randomized approximation scheme exists for #BIS [5].

Our goal in the present work is to estimate the partition function $\mathcal{Z}$ of the hard-core model on *random subgraphs* of Q_d , obtained via percolation. As is often the case in such settings, our results concerning the approximation of the partition function will also allow us to elucidate the structure of a typical independent set in the hard-core model on a percolated hypercube.

To keep the presentation focused, we only consider the case $\lambda = 1$ in this article, although we expect our results to generalize readily to other values of λ , as was done in [6]. When $\lambda = 1$, the hard-core model on a graph G is the uniform distribution on all independent sets in G , and thus the partition function is the total number of such independent sets; this number is denoted by $i(G)$. We will denote a hypercube that has been percolated with parameter p (meaning that each edge is deleted independently with probability $1 - p$) by $Q_{d,p}$. So our main focus will be to approximate $i(Q_{d,p})$.

1.1. Background. Broadly speaking, as p decreases and more edges are removed, there are fewer constraints for a set to be independent and thus typical independent sets become larger. When $p = 1$ so that $Q_{d,p} = Q_d$, much is known, going back to the work of Korshunov and Sapozhenko [12] who discovered the leading-order behavior of the number of independent sets:

$$i(Q_d) \sim 2^{2^{d-1}} \cdot 2 \cdot \sqrt{e}. \quad (1.1)$$

Intuitively, this approximation holds because most independent sets in Q_d are very close to simply being a subset of either the even or odd side (leading to the factor of $2^{2^{d-1}} \cdot 2$), with a small number of *defect vertices* in the opposite side. Each included defect vertex v precludes d non-defect vertices and thus eliminates d binary choices (whether or not to include each neighbor of v in the independent set), effectively reducing the number of possible independent sets by a factor of roughly 2^d . This means the probability that a uniformly random independent set includes a *particular* defect vertex v is roughly 2^{-d} . Using the powerful *graph container* method, it can be shown that the 2^{d-1} possible defect vertices do not interact and behave roughly independently, and thus that the number of defect vertices in a uniformly random independent set is approximately $\text{Bin}(2^{d-1}, 2^{-d}) \approx \text{Poisson}(\frac{1}{2})$, leading to the factor of $\sqrt{e}$.

The approximation (1.1) was sharpened by Jenssen and Perkins [11] who provided an algorithm which computes arbitrary lower-order terms in the approximation. This is based on a more careful analysis of the defects, which can be formalized via the language of *abstract polymer models* and which can subsequently be examined via the *cluster expansion*. A similar technique was used by Kronenberg and Spinka [13] who initiated the study of $i(Q_{d,p})$ for $p < 1$. They adapted the cluster expansion technique to an *annealed* version of the hard-core model on $Q_{d,p}$, and were able to exhibit arbitrarily sharp approximations of the moments $\mathbb{E}[i(Q_{d,p})^k]$ via an algorithm. These moment calculations revealed a possible phase transition at $p = \frac{2}{3}$, and they conjectured a central limit theorem for $i(Q_{d,p})$ when $p > \frac{2}{3}$ (with a small variance leading to concentration), as well as a lognormal limit theorem for $p = \frac{2}{3}$, with anticoncentration for $p < \frac{2}{3}$. Further, their moment calculations allowed them to prove a central limit theorem for $i(Q_{d,p})$ as long as $p = p_d$ tended to 1 with d at a certain rate. The behavior for a fixed $p < 1$ remained completely open.

In [6, Theorem 1.1], we addressed the conjectures of [13], exhibiting explicit deterministic quantities $M_{d,p}$ and $S_{d,p}$ such that, for fixed $p \in (\frac{2}{3}, 1)$,

$$\frac{i(Q_{d,p}) - M_{d,p}}{S_{d,p}} \xrightarrow{w} \mathcal{N}(0, 1), \quad (1.2)$$

and for $p = \frac{2}{3}$,

$$\frac{i(Q_{d,p})}{M_{d,\frac{2}{3}}} \xrightarrow{w} e^{W^{\text{Even}}/\sqrt{2}} + e^{W^{\text{Odd}}/\sqrt{2}}, \quad (1.3)$$

where $W^{\text{Even}}, W^{\text{Odd}} \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1)$. Here and in the remainder of this article, $\xrightarrow{w}$ denotes convergence in distribution as $d \rightarrow \infty$.

Both (1.2) and (1.3) derive from a more fundamental result of [6] which gives an in-probability approximation for $i(Q_{d,p})$ as follows, for $p \geq \frac{2}{3}$:

$$\frac{i(Q_{d,p})}{2^{2^{d-1}}} \stackrel{\mathbb{P}}{\sim} e^{\Psi^{\text{Even}}} + e^{\Psi^{\text{Odd}}}, \quad (1.4)$$

where Ψ^{Even} and Ψ^{Odd} are random variables which can be written explicitly in terms of the degrees of vertices in $Q_{d,p}$, with each vertex contributing once in a sum of independent terms. Here and in the rest of the article, $A \stackrel{\mathbb{P}}{\sim} B$ means that $\frac{A}{B} \xrightarrow{\mathbb{P}} 1$, where $\xrightarrow{\mathbb{P}}$ denotes convergence in probability with respect to the randomness of the percolation configuration, as $d \rightarrow \infty$. Moreover, Ψ^{Even} and Ψ^{Odd} satisfy a joint central limit theorem, converging after proper rescaling to a pair of independent normal random variables. The phase transition at $p = \frac{2}{3}$ is explained by the variance of Ψ^{Even} and Ψ^{Odd} (which can easily be explicitly computed) decreasing to 0 for $p > \frac{2}{3}$, staying constant for $p = \frac{2}{3}$, and increasing to ∞ for $p < \frac{2}{3}$.

The approximation (1.4) should be seen as an extension of (1.1) to $p < 1$, and at its core is the same fundamental mechanism, where independent sets are close to the *ground states* (uniformly random subsets of the even or odd side) with a small number of defects. The regime $p > \frac{2}{3}$ is exactly where the defects are still sparse enough that they can be effectively approximated by sampling them *with replacement*, as no repeat samples will arise (via an analysis akin to the classical *birthday problem*). When $p \leq \frac{2}{3}$, this sampling scheme fails due to the presence of defect vertices with outsized probability of occurring, which also leads to the failure of concentration and the central limit theorem (1.2).

In addition to results about the *number* of independent sets, the approximation (1.4) allowed to obtain various structural results about typical independent sets in $Q_{d,p}$, sampled uniformly [6, Theorem 1.2 and Proposition 1.4]. For $p \geq \frac{2}{3}$ this approximation scheme allowed us to conclude that the number of defect vertices in a uniformly random independent set in $Q_{d,p}$ is approximately $\text{Poisson}(\frac{1}{2}(2-p)^d)$.

1.2. Results. In our previous work [6], we focused on the case $p \geq \frac{2}{3}$ due to various technicalities that arise as p decreases and the size of a typical independent set increases. However, one very interesting question asked by [13] was about the behavior of this model near $p = \frac{1}{2}$ in light of the natural emergence of the latter as an important point of phase transition for many other features of the percolated hypercube $Q_{d,p}$ such as the existence of isolated vertices, existence of a perfect matching, and Hamiltonicity [2, 3, 4, 7, 9].

Motivated by the above, in the present article we introduce many new techniques which are organized via a convenient probabilistic framework for studying this problem. We expect that this will be useful for continued study of this model for lower values of p , and that it may also provide insights into new ways to study other related problems as well.

Notably, as a consequence of our results, *a posteriori*, it turns out that $p = \frac{1}{2}$ does not appear to be a point of phase transition for any qualitative behavior about independent sets in $Q_{d,p}$, and we obtain essentially the same result just above $p = \frac{1}{2}$ as just below. Instead, $p = \frac{1}{2}$ only corresponds to a change in the way we may present the formulas involved in our result (see Remark 7.3). Of more interest and relevance for the qualitative behavior of this model, as evidenced by our results to be presented below, is the phase transition at $p = 2 - \sqrt{2} \approx 0.586$, below which the defects start exhibiting nontrivial clustering behavior. This leads to a loss of independence between defects, and as such, our analysis in this article depends crucially on the interactions between nearby vertices in $Q_{d,p}$.

We are now in a position to state our first main result. We remind the reader that here and throughout the rest of the article, $A \stackrel{\mathbb{P}}{\sim} B$ means that $\frac{A}{B} \xrightarrow{\mathbb{P}} 1$, where $\xrightarrow{\mathbb{P}}$ denotes convergence in probability (with respect to randomness of the percolation $Q_{d,p}$ of the hypercube) as $d \rightarrow \infty$. In addition, we will occasionally use phrases like “ $o(1)$ -in-probability” to describe a random variable A with $A \xrightarrow{\mathbb{P}} 0$. Further, when we say that

a collection of random variables $\{A_S\}_{S \in \mathcal{S}}$ indexed by some collection $\mathcal{S}$ is “uniformly $o(1)$ -in-probability” if there is some random variable $M \xrightarrow{\mathbb{P}} 0$ such that $|A_S| \leq M$ for all $S \in \mathcal{S}$. Similarly, we will say that $A_S \stackrel{\mathbb{P}}{\sim} B_S$ uniformly if $\frac{A_S}{B_S} - 1$ is uniformly $o(1)$ -in-probability.

Theorem 1.1. *There is some $\gamma > 0$ such that for $p > \frac{1}{2} - \gamma$ the following holds. There are random variables $\Psi^{\text{Even}}, \Psi^{\text{Odd}}$, which can be expressed explicitly as functions of the percolation configuration, such that*

$$\frac{i(Q_{d,p})}{2^{2^{d-1}}} \stackrel{\mathbb{P}}{\sim} e^{\Psi^{\text{Even}}} + e^{\Psi^{\text{Odd}}}.$$

Moreover, there are explicit deterministic quantities $\mu = \mu_{d,p}$ and $\sigma = \sigma_{d,p}$ (defined below) such that

$$\left(\frac{\Psi^{\text{Even}} - \mu}{\sigma}, \frac{\Psi^{\text{Odd}} - \mu}{\sigma} \right) \xrightarrow{w} \mathcal{N}(0, 1) \otimes \mathcal{N}(0, 1).$$

Let us quickly see some consequences of this theorem, for which it is helpful to know that the value of σ in the statement is given by

$$\sigma^2 := \frac{1}{2} \left(2 - \frac{3}{2}p \right)^d. \quad (1.5)$$

Notice that for $p > \frac{2}{3}$ we have $\sigma \rightarrow 0$, while if $p < \frac{2}{3}$ we have $\sigma \rightarrow \infty$. The above suggests that for $p < \frac{2}{3}$, as is the case for lognormal variables with large fluctuations in the exponent, the typical value of $i(Q_{d,p})$ will be much smaller than its mean. This can indeed be made precise using sharp estimates about the mean from [13]. In addition, for much the same reason, there is no way to shift and scale $i(Q_{d,p})$ by deterministic quantities to obtain a scaling limit for $p < \frac{2}{3}$. However, we do obtain a scaling limit for the *free energy*, which is defined as $\log i(Q_{d,p})$: from Theorem 1.1 and (1.5) it is easy to see that

$$\frac{\log i(Q_{d,p}) - \log 2^{2^{d-1}} - \mu}{\sigma} \xrightarrow{w} \max\{W^{\text{Even}}, W^{\text{Odd}}\}$$

for $\frac{1}{2} - \gamma < p < \frac{2}{3}$, where $W^{\text{Even}}, W^{\text{Odd}} \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1)$.

We now briefly describe the variables Ψ^{Even} and Ψ^{Odd} appearing in the statement of Theorem 1.1, as well as the deterministic quantity μ (the quantity σ having just been defined in (1.5) above). Further elaborations on the form of these quantities arising from relatively simple moment calculations are provided later in Section 7.

Recall from the previous subsection that $N_p(v)$ denotes the (random) number of neighbors of a single vertex v in $Q_{d,p}$. We will need to consider sets containing two distinct vertices which are “adjacent” in the sense that they share a neighbor in the opposite side of the (unpercolated) hypercube. We will call a set of vertices *2-linked* if they are connected through the above adjacency structure.

Note that if two vertices share a neighbor in Q_d , they must share exactly two neighbors as demonstrated by the following diagram.



We call such subsets (e.g. the pair of shaded nodes in the top row above) *dimers*, and denote by $\text{Dimers}^{\text{Even}}$ the collection of such subsets of the even side, and similarly for $\text{Dimers}^{\text{Odd}}$. For a dimer $\{u, v\}$, we

write $N_p(\{u, v\})$ for the (random) number of vertices which neighbor either u or v in $Q_{d,p}$; note that $N_p(\{u, v\}) \neq N_p(u) + N_p(v)$ in general. With these definitions, we can write down the formula for Ψ^{Even} , the formula for Ψ^{Odd} being identical but with Even replaced by Odd throughout.

$$\Psi^{\text{Even}} := \sum_{v \in \text{Even}} \log(1 + 2^{-N_p(v)}) + \sum_{\{u,v\} \in \text{Dimers}^{\text{Even}}} \left(2^{-N_p(\{u,v\})} - 2^{-N_p(u) - N_p(v)} \right). \quad (1.7)$$

Next, we turn to the deterministic quantity μ . Specifically, we first define

$$\mu_1 := \frac{1}{2}(2-p)^d - \frac{1}{4} \left(2 - \frac{3}{2}p \right)^d + \frac{1}{6} \left(2 - \frac{7}{4}p \right)^d. \quad (1.8)$$

This is a proxy for the mean of the first sum in (1.7). Additionally define

$$\mu_2 := \frac{d(d-1)}{4} \left(\frac{(2-p)^2}{2} \right)^d \left(\frac{1 + (1-p)^2}{2-p} \right)^2, \quad (1.9)$$

$$\tilde{\mu}_2 := \frac{d(d-1)}{4} \left(\frac{(2-p)^2}{2} \right)^d. \quad (1.10)$$

Then $\mu_2 - \tilde{\mu}_2$ is the mean of the second sum in (1.7). So, we set

$$\mu := \mu_1 + \mu_2 - \tilde{\mu}_2. \quad (1.11)$$

A more complete understanding of these quantities is presented later in Section 7.

As mentioned already, our arguments also lead to a sampling scheme which is presented next. For convenience, we allow the algorithm to fail at various points, but we will show that it fails with very low probability. Additionally, here and throughout the paper we blur the lines between a collection of singletons and dimers considered as separate polymers, and the union of such a collection (which is just a set of vertices), as this should not present a scope for confusion.

Definition 1.2 (ApproxSampler). *Define the random independent set $\hat{I}$ in $Q_{d,p}$ via the following procedure.*

- (1) Let $\mathcal{H}$ be Even or Odd, with probabilities proportional to $e^{\Psi^{\text{Even}}}$ and $e^{\Psi^{\text{Odd}}}$ respectively.
- (2) Let $\tilde{S}$ denote a subset of $\mathcal{H}$ which includes each $v \in \mathcal{H}$ independently with probability $\frac{2^{-N_p(v)}}{1+2^{-N_p(v)}}$.
- (3) If $\tilde{S}$ contains any triple of vertices which is 2-linked, terminate the algorithm with a result of failure. Otherwise, remove all dimers from $\tilde{S}$ and let S_1 denote the resulting set of well-separated singletons (which share no neighbors).
- (4) Let S_2 denote a collection of dimers constructed by including each $\{u, v\} \in \text{Dimers}^{\mathcal{H}}$ independently with probability $\frac{2^{-N_p(\{u,v\})}}{1+2^{-N_p(\{u,v\})}}$.
- (5) If any neighbor of a dimer in S_2 also neighbors a vertex in S_1 , or if any pair of dimers in S_2 overlap or share a neighbor, terminate the algorithm with a result of failure. Otherwise, let $\hat{S} = S_1 \cup S_2$.
- (6) Let $\hat{I}$ denote the independent set with $\hat{I} \cap \mathcal{H} = \hat{S}$ and with every vertex in $Q_d \setminus \mathcal{H}$ which does not neighbor $\hat{S}$ included independently with probability $\frac{1}{2}$ each.

Note that there are two levels of randomness here as ApproxSampler is a *random algorithm* even when the disorder of the percolation configuration $Q_{d,p}$ is fixed, which is often termed as quenching the disorder. For simplicity of notation, we will refer to both levels of randomness using $\mathbb{P}$, but which measure is being referred

to will be clear from context. This convention is in full force in the following statement of our second main result, which states that ApproxSampler gives a good approximation for a uniformly random independent set in $Q_{d,p}$.

Theorem 1.3. *Let $p > \frac{1}{2} - \gamma$, where $\gamma > 0$ is as in Theorem 1.1. Then*


$$\mathbb{P}[\text{ApproxSampler fails}] \xrightarrow{\mathbb{P}} 0.$$

Additionally, if $\hat{I}$ denotes the output of ApproxSampler conditioned on not failing and I denotes a uniformly random independent set in $Q_{d,p}$, then I and $\hat{I}$ may be coupled so that

$$\mathbb{P}[I \neq \hat{I}] \xrightarrow{\mathbb{P}} 0.$$

1.3. Future directions. As evidenced from the statement of our results in the previous subsection, the relevant “defects” of independent sets in the parameter regime we consider are all either single vertices or dimers, consisting of two nearby vertices (which share a neighbor in the non-defect side). However, as p decreases even further and typical independent sets get larger, the defect side will contain more and more *polymers* of higher order, which are clusters of defect vertices that are *2-linked*. Specifically, for $p \leq 2 - 2^{2/3} \approx 0.413$, size-three polymers (trimers) begin to appear, and in general order- k polymers begin to appear for $p \leq 2 - 2^{(k-1)/k}$.

Broadly speaking, we introduce a framework for separating out the effect of singletons and dimers, and for analyzing their interactions. This strategy of separating polymers of different sizes and estimating their interactions ought to yield dividends even while studying lower values of p where higher-order polymers are relevant. However, as the order of polymers increases, so does the number of possible shapes that the polymers form. For instance, while every dimer looks essentially the same, geometrically speaking, as the pair of shaded vertices in top row of diagram (1.6), there are already two different nonisomorphic types of trimers, which look like the shaded vertices in the top row of either



i.e. the induced adjacency structure of the “sharing a neighbor” relation is either a triangle or a path. For larger polymers there are even more different nonisomorphic types with increasingly intricate adjacency structures between the vertices. As may be predicted, then, as p decreases to 0 the number of different relevant terms in sums of the form (1.7) will increase rapidly.

In addition, even before trimers appear (which occurs for $p \leq 2 - 2^{2/3}$), the interactions between singletons and dimers become relevant. Indeed, our proof in the present work aims to cover the complete range of p for which the singletons and dimers behave roughly independently, but not the full range of $p > 2 - 2^{2/3}$. To be more specific, although the threshold in our theorem is only written as $\frac{1}{2} - \gamma$ for some $\gamma > 0$, one may check, numerically or otherwise, that $\frac{1}{2} - \gamma \approx 0.465$, which is strictly greater than $2 - 2^{2/3} \approx 0.413$. Although we do not prove it, we do expect that this range of p is optimal, in the sense that for $p \in (2 - 2^{2/3}, \frac{1}{2} - \gamma]$, the singletons and dimers will fail to be independent. It remains an interesting problem to capture the precise nature of the interaction.

1.4. Acknowledgements. SG was partially supported by NSF Career grant-1945172. VW was partially supported by the NSF Graduate Research Fellowship grant DGE 2146752. SG learnt about this problem from Gal Kronenberg’s talk on [13] while attending the workshop titled ‘Bootstrap Percolation and its Applications’ at Banff in April 2024. He thanks the speaker as well as the organizers.

2. IDEA AND PROOF OF MAIN RESULT

In this section we provide a brief overview of the main ideas.

As has already been emphasized a few times, a key observation going back to [12] is that each independent set can be separated into its intersections with the even side and odd side of the hypercube, yielding a convenient representation for $i(Q_{d,p})$.

Since there are no edges between even vertices, every subset $S \subseteq \text{Even}$ is a valid choice for the even side of an independent set. For any $S \subseteq \text{Even}$, the number of independent sets $I \subseteq Q_{d,p}$ with $I \cap \text{Even} = S$ can be written explicitly as

$$2^{2^{d-1} - N_p(S)}, \quad (2.1)$$

where $N_p(S)$ denotes the number of vertices in Odd which are neighbors of S in $Q_{d,p}$. Indeed, any non-neighbor of S can be included in such an independent set, and since there are no edges between odd vertices, any subset T of the odd non-neighbors of S yields a valid independent set $I = S \cup T$. There are $2^{d-1} - N_p(S)$ such odd non-neighbors of S , yielding the formula (2.1). This yields the exact representation

$$\frac{i(Q_{d,p})}{2^{2^{d-1}}} = \sum_{S \subseteq \text{Even}} 2^{-N_p(S)} \quad (2.2)$$

for the total number of independent sets in $Q_{d,p}$.

Now, for well-separated sets $S \subseteq \text{Even}$, where no vertices of S share a neighbor in Odd (even in the unpercolated hypercube Q_d), the expression $2^{-N_p(S)}$ factorizes as a product of contributions from each individual vertex:

$$2^{-N_p(S)} = \prod_{v \in S} 2^{-N_p(v)}, \quad (2.3)$$

since the neighborhoods of $v \in S$ do not overlap. Moreover, the variables $2^{-N_p(v)}$ depend on disjoint subsets of the edges in the percolation configuration, meaning that the factors of (2.3) are independent (in fact, the entire family of random variables $\{N_p(v) : v \in \text{Even}\}$ is mutually independent, and similarly for Odd).

Of course, there are many sets $S \subseteq \text{Even}$ which are *not well-separated*, meaning that vertices of S share neighbors in Odd (before percolating). For general sets S , one could instead form a factorization like (2.3) but where the factors correspond to *subsets* of S which are well-separated from each other. However, the complexity of such factorizations quickly spirals out of control when S is large, since the form of each factor depends on the internal structure of the corresponding subset.

A key feature of this model which allows us to overcome this issue is that when p is bounded away from 0, the model is effectively in a *very low-temperature regime*. More precisely, most independent sets I are very close to one of the two “ground states” meaning that *exactly one of $I \cap \text{Even}$ and $I \cap \text{Odd}$ is large*, and the other is very small in comparison.

This allows us to work separately with two classes of independent sets, those with small even side and those with small odd side. If $S \subseteq \text{Even}$ is small, we can decompose it as a union of *small* subsets which are well-separated from each other, which are easier to categorize. Our approach then relies on analyzing the contributions of these small separated subsets.

In the next few sections, we briefly describe the ingredients and then combine them to prove Theorem 1.1.

2.1. Polymer models. A convenient framework for reasoning about a set S as a union of small subsets which are well-separated from each other is provided by the language of *polymer models*. Broadly speaking, a polymer model is a probability distribution on *polymer configurations*, which are collections of small units called *polymers* satisfying a compatibility condition. The defining feature is that the probability of a polymer configuration is proportional to a product of given *weights*, one for each polymer in the configuration.

In our case, polymers will be subsets $\mathbf{p} \subseteq \text{Even}$ (or Odd) which are *2-linked*, recalling that it means that they are connected under the relation of sharing a common neighbor in Q_d . We also require that polymers are *small* in a sense to be defined later, in Section 3. Two polymers are said to be *compatible* if their neighborhoods in Q_d do not intersect. We denote by PD^{Even} the set of configurations of pairwise-compatible polymers $\mathbf{p} \subseteq \text{Even}$, and similarly for Odd (PD stands for polymer-decomposable). The weight of each polymer $\mathbf{p}$ is given by

$$\varphi_{\mathbf{p}} := 2^{-N_p(\mathbf{p})}, \quad (2.4)$$

so that the weight of a configuration of compatible polymers $S \in \text{PD}^{\text{Even}}$ is

$$\prod_{\mathbf{p} \in S} \varphi_{\mathbf{p}} = \prod_{\mathbf{p} \in S} 2^{-N_p(\mathbf{p})},$$

which is the same as the weight of S in the sum (2.2), when it is considered as a *union* of its polymers (the subtle distinction between a *configuration (set) of polymers* and a set of vertices as a *union of polymers* will be ignored as it causes no confusion). The partition function (normalization constant) for this polymer model is

$$\mathcal{Z}^{\text{Even}} := \sum_{S \in \text{PD}^{\text{Even}}} \prod_{\mathbf{p} \in S} \varphi_{\mathbf{p}}, \quad (2.5)$$

and similarly there is a partition function $\mathcal{Z}^{\text{Odd}}$ on the odd side. The first step of our proof, as in various previous works on independent sets [6, 11, 13], is to reduce the study of $i(Q_{d,p})$ to the study of these polymer model partition functions.

Proposition 2.1. *For $p \in (0, 1)$,*

$$\frac{i(Q_{d,p})}{2^{2^d-1}} \stackrel{\mathbb{P}}{\sim} \mathcal{Z}^{\text{Even}} + \mathcal{Z}^{\text{Odd}}.$$

This proposition is a straightforward consequence of one the main conclusions of [13] and was already presented in our previous work [6, Lemma 4.8], but in the present article, specifically as Proposition 3.4 in Section 3 below, we present a somewhat more streamlined proof than has previously appeared. The key input is an *annealed version* of the statement from [13], where it (along with various extensions) was used to derive arbitrarily accurate estimates of the *moments* of $i(Q_{d,p})$ via the *cluster expansion* of polymer models, which will be discussed briefly in Section 3. However, it is worth emphasizing that the present work does

not make use of the cluster expansion beyond this annealed input; instead we analyze the quenched polymer model partition functions directly.

Due to the structural similarity between Proposition 2.1 and our goal, which is Theorem 1.1, it may be tempting to simply define $\Psi^{\text{Even}} = \log \mathcal{Z}^{\text{Even}}$ and similarly for Odd. However, these quantities are not easy to understand a priori, letting alone their joint distribution, and we will instead obtain a good approximation of $\log \mathcal{Z}^{\text{Even}}$ by a simpler and more explicit quantity (already spelled out in (1.7)) which is more amenable to analysis.

2.2. Main step: reducing to a simpler polymer model. In this subsection we do not discuss the joint behavior of $\mathcal{Z}^{\text{Even}}$ and $\mathcal{Z}^{\text{Odd}}$, deciding to focus only on $\mathcal{Z}^{\text{Even}}$. Everything in this section also goes through with Even replaced by Odd.

In order to analyze $\mathcal{Z}^{\text{Even}}$, a central observable will be the partition function of a simpler polymer model under which *each vertex behaves independently*. It is worth pointing out that this simpler partition function already serves as a strong proxy for $\mathcal{Z}^{\text{Even}}$ if $p > 2 - \sqrt{2} \approx 0.586$ but not beyond. Thus the simpler object does not by itself suffice to approximate $\mathcal{Z}^{\text{Even}}$. Rather, it serves as an important base model to compare to. This quantitative comparison for $p \leq 2 - \sqrt{2}$ is conceptually one of the major contributions of the present work.

To define the simpler model, we replace $\varphi_{\mathbf{p}} = 2^{-N_p(\mathbf{p})}$ by

$$\tilde{\varphi}_{\mathbf{p}} := \prod_{v \in \mathbf{p}} \varphi_v, \quad (2.6)$$

where again $\varphi_v = 2^{-N_p(v)}$. This means the weight of a polymer configuration S is simply

$$\prod_{\mathbf{p} \in S} \tilde{\varphi}_{\mathbf{p}} = \prod_{\mathbf{p} \in S} \prod_{v \in \mathbf{p}} \varphi_v = \prod_{v \in S} \varphi_v$$

(recall from the previous subsection our identification of a *set* of compatible polymers and the *union* of those same polymers). We also remove the restriction that polymers be small in this polymer model, so that *any* set $S \subseteq \text{Even}$ is given the above weight. The partition function of this model is thus

$$\tilde{\mathcal{Z}}^{\text{Even}} := \sum_{S \subseteq \text{Even}} \prod_{v \in S} \varphi_v = \prod_{v \in \text{Even}} (1 + \varphi_v). \quad (2.7)$$

In other words, to sample a set $S \subseteq \text{Even}$ from this polymer model, one includes each vertex $v \in \text{Even}$ independently with probability $\frac{\varphi_v}{1 + \varphi_v}$. The simplicity of the formula for $\tilde{\mathcal{Z}}^{\text{Even}}$ makes it amenable to analysis.

In light of the above, the key innovation in this paper is to develop a characterization of the ratio $\mathcal{Z}^{\text{Even}} / \tilde{\mathcal{Z}}^{\text{Even}}$.

The first step in this analysis is to reduce both $\mathcal{Z}^{\text{Even}}$ and $\tilde{\mathcal{Z}}^{\text{Even}}$ by eliminating terms with polymers which are too large. To this end, for each $k \in \mathbb{N}$ define $\text{PD}_{\leq k}^{\text{Even}}$ to be the set of configurations of pairwise-compatible polymers $\mathbf{p} \subseteq \text{Even}$ with at most k vertices, and define

$$\mathcal{Z}_{\leq k}^{\text{Even}} := \sum_{S \in \text{PD}_{\leq k}^{\text{Even}}} \prod_{\mathbf{p} \in S} \varphi_{\mathbf{p}}, \quad \text{and} \quad \tilde{\mathcal{Z}}_{\leq k}^{\text{Even}} := \sum_{S \in \text{PD}_{\leq k}^{\text{Even}}} \prod_{\mathbf{p} \in S} \tilde{\varphi}_{\mathbf{p}}. \quad (2.8)$$

As has already been indicated, in the regime of p addressed in this article, polymers of sizes more than 2 do not appear, which is formalized by the following proposition.

Proposition 2.2. *For $p > 2 - 2^{2/3} \approx 0.413$,*

$$\mathcal{Z}^{\text{Even}} \stackrel{\mathbb{P}}{\sim} \mathcal{Z}_{\leq 2}^{\text{Even}} \quad \text{and} \quad \tilde{\mathcal{Z}}^{\text{Even}} \stackrel{\mathbb{P}}{\sim} \tilde{\mathcal{Z}}_{\leq 2}^{\text{Even}}.$$

This will be proved in two parts, as Propositions 4.1 and 4.2 in Section 4 below. With Proposition 2.2 in place, we can think of both polymer models as having two parts: first, a collection S_1 of well-separated singletons, and second, a collection S_2 of well-separated dimers *which are also compatible with S_1* .

While in principle, the law of S_2 given S_1 depends on the instantiation of S_1 , one of the key observations in this paper is that there is some threshold $\frac{1}{2} - \gamma$ strictly less than $\frac{1}{2}$ such that

when $p > \frac{1}{2} - \gamma$, the distributions of S_1 and S_2 are in fact approximately independent

under both polymer models. Numerically, this threshold is around 0.465.

In order to state this approximate independence result, we define for each $k \in \mathbb{N}$

$$\mathcal{Z}_k^{\text{Even}} := \sum_{S \in \text{PD}_k^{\text{Even}}} \prod_{\mathbf{p} \in S} \varphi_{\mathbf{p}} \quad \text{and} \quad \tilde{\mathcal{Z}}_k^{\text{Even}} := \sum_{S \in \text{PD}_k^{\text{Even}}} \prod_{\mathbf{p} \in S} \tilde{\varphi}_{\mathbf{p}}, \quad (2.9)$$

where $\text{PD}_k^{\text{Even}}$ denotes the set of configurations of pairwise-compatible k -polymers. Note that we have $\tilde{\mathcal{Z}}_1^{\text{Even}} = \mathcal{Z}_1^{\text{Even}}$ since $\varphi_v = \tilde{\varphi}_v$ for singletons, but this does not extend to $k > 1$. The following proposition exhibits the aforementioned asymptotic independence of the singletons and dimers.

Proposition 2.3. *There is some $\gamma > 0$ such that for all $p > \frac{1}{2} - \gamma$, we have*

$$\mathcal{Z}_{\leq 2}^{\text{Even}} \stackrel{\mathbb{P}}{\sim} \mathcal{Z}_1^{\text{Even}} \cdot \mathcal{Z}_2^{\text{Even}} \quad \text{and} \quad \tilde{\mathcal{Z}}_{\leq 2}^{\text{Even}} \stackrel{\mathbb{P}}{\sim} \mathcal{Z}_1^{\text{Even}} \cdot \tilde{\mathcal{Z}}_2^{\text{Even}}.$$

To show this, we must prove that the restriction that the dimer part S_2 be well-separated from the singleton part S_1 is superfluous. In other words, when one first samples the singleton set S_1 from the polymer model corresponding to $\mathcal{Z}_1^{\text{Even}}$, then by simply sampling the collection of dimers S_2 from the polymer model with partition function $\mathcal{Z}_2^{\text{Even}}$, one obtains S_2 which is well-separated from S_1 with high probability. This follows by developing an understanding of the probabilistic behavior of the *set of dimers adjacent to S_1* reliant on certain large deviation estimates.

Finally, the dimer part S_2 can itself be approximated by independently including every possible dimer, with the appropriate probability. In other words, the requirement that S_2 be well-separated is also superfluous; in formulas, this means

$$\mathcal{Z}_2^{\text{Even}} \stackrel{\mathbb{P}}{\sim} \sum_{S_2 \subseteq \text{Dimers}^{\text{Even}}} \prod_{\mathbf{d} \in S_2} \varphi_{\mathbf{d}} = \prod_{\mathbf{d} \in \text{Dimers}^{\text{Even}}} (1 + \varphi_{\mathbf{d}}). \quad (2.10)$$

Here we recall from just above (1.7) the notation $\text{Dimers}^{\text{Even}}$ for the collection of dimers (i.e. size-2 polymers) in the even side (note that we typically use $\mathbf{d}$ to denote a dimer, instead of $\mathbf{p}$ which denotes an arbitrary polymer). Now, we can go one step further than (2.10) and approximate each factor of $(1 + \varphi_{\mathbf{d}})$ by $e^{\varphi_{\mathbf{d}}}$. For the parameter regime we consider, the error in this approximation is small enough that it does not accumulate

across the product over dimers. Moreover, the same analysis goes through for $\tilde{\mathcal{Z}}_2^{\text{Even}}$. Thus, defining

$$\Delta^{\text{Even}} = \sum_{\mathbf{d} \in \text{Dimers}^{\text{Even}}} \varphi_{\mathbf{d}} \quad \text{and} \quad \tilde{\Delta}^{\text{Even}} = \sum_{\mathbf{d} \in \text{Dimers}^{\text{Even}}} \tilde{\varphi}_{\mathbf{d}}, \quad (2.11)$$

we arrive at the following lemma.

Lemma 2.4. *For $p > 0.391$,*

$$\mathcal{Z}_2^{\text{Even}} \stackrel{\mathbb{P}}{\sim} e^{\Delta^{\text{Even}}} \quad \text{and} \quad \tilde{\mathcal{Z}}_2^{\text{Even}} \stackrel{\mathbb{P}}{\sim} e^{\tilde{\Delta}^{\text{Even}}}.$$

To summarize at a very high level, Lemma 2.4 and Proposition 2.3 capture the phenomenon that if there are not many dimers present, then with high probability none of them will “collide” with each other or with any of the singletons present in the defect side of a uniformly-sampled independent set, allowing for a sampling strategy which first produces only the singletons, then afterwards produces the dimers. We expect this broad strategy of hierarchically sampling larger and larger polymers to be useful to analyze the model even for values of p below those addressed in this paper.

Remark 2.5. It is worth emphasizing that the strategy of hierarchically sampling the dimers after the singletons does not work for the *worst-case* singleton sets that could appear in the first round of sampling, at least not for all p under consideration in the present article. Indeed, as will be discussed further in Section 6.4, the probability that the sampled dimer set avoids the worst-case singleton set converges to 0 only when p is bigger than approximately 0.549. To cover the full regime of $p > \frac{1}{2} - \gamma$, we instead present an *average-case* analysis in Sections 6.5 and 6.6, averaging over the sampled singleton set.

In any case, combining Propositions 2.2 and 2.3 as well as Lemma 2.4, we obtain our final in-probability approximation of each polymer partition function.

Corollary 2.6. *There is some $\gamma > 0$ such that for all $p > \frac{1}{2} - \gamma$, we have*

$$\mathcal{Z}^{\text{Even}} \stackrel{\mathbb{P}}{\sim} \tilde{\mathcal{Z}}^{\text{Even}} \cdot e^{\Delta^{\text{Even}} - \tilde{\Delta}^{\text{Even}}}.$$

Of course, the same holds with Even replaced by Odd throughout.

Before proceeding, we remark that our actual proof strategy slightly reorganizes Proposition 2.3 and Lemma 2.4 and these particular statements do not appear in the body of the work. However, the main important result is actually Corollary 2.6 itself, which will be proved directly via Proposition 6.1 below. If the reader wishes to use the statements of Proposition 2.3 or Lemma 2.4 directly, these statements may be proven by combining the statement of Proposition 6.1 with that of Lemma 6.4.

2.3. Distributional limits. With Corollary 2.6 in hand, we define

$$\Psi^{\text{Even}} = \log \tilde{\mathcal{Z}}^{\text{Even}} + \Delta^{\text{Even}} - \tilde{\Delta}^{\text{Even}}, \quad (2.12)$$

and similarly for Ψ^{Odd} ; note that this definition agrees with (1.7). Also, let μ and σ be as defined previously in (1.11) and (1.5); we remind the reader that the justification for these definitions can be found in Section 7. With these definitions, we have a joint central limit theorem for Ψ^{Even} and Ψ^{Odd} , with the limit being a pair of independent normal random variables.

Proposition 2.7. For $p > 0.455$,

$$\left(\frac{\Psi^{\text{Even}} - \mu}{\sigma}, \frac{\Psi^{\text{Odd}} - \mu}{\sigma} \right) \xRightarrow{w} \mathcal{N}(0, 1) \otimes \mathcal{N}(0, 1).$$

This will be proved as Proposition 7.6 in Section 7 below. Note that the threshold 0.455 appearing in this result is unrelated to the threshold $\frac{1}{2} - \gamma$ of Corollary 2.6. Instead, it is related to the specific form of the constant μ defined in (1.11), which is a good approximation for the true mean of Ψ^{Even} if $p > 0.455$.

We also remark that the fluctuations in this limit theorem are solely due to the $\log \tilde{\mathcal{Z}}^{\text{Even}}$ and $\log \tilde{\mathcal{Z}}^{\text{Odd}}$ terms in (2.12). These two variables themselves exhibit a joint CLT, and the variables $\Delta^{\text{Even}}, \Delta^{\text{Odd}}, \tilde{\Delta}^{\text{Even}}$, and $\tilde{\Delta}^{\text{Odd}}$ each concentrate around deterministic quantities.

The variables contributing to the fluctuations can be written as sums of independent variables, namely

$$\log \tilde{\mathcal{Z}}^{\text{Even}} = \sum_{v \in \text{Even}} \log(1 + \varphi_v) \quad \text{and} \quad \log \tilde{\mathcal{Z}}^{\text{Odd}} = \sum_{v \in \text{Odd}} \log(1 + \varphi_v).$$

Note however that *the two sums are not independent* since each edge's state in the percolation configuration contributes to one summand in each sum. However this dependence is rather weak and as in [6, Lemma 3.1], a dependency-neighborhood version of Stein's method suffices to conclude the joint CLT to a pair of normal random variables which are independent in the limit.

2.4. Proof of first main theorem. With all of the above statements in place, we can easily conclude the proof of our first main theorem, yielding an in-probability approximation of $i(Q_{d,p})$ as a simple function of two explicit random variables which exhibit a joint central limit to a pair of independent standard normals.

Proof of Theorem 1.1. By Proposition 2.1 and Corollary 2.6, we have

$$\begin{aligned} \frac{i(Q_{d,p})}{2^{2^d-1}} &\stackrel{\mathbb{P}}{\sim} \tilde{\mathcal{Z}}^{\text{Even}} \cdot e^{\Delta^{\text{Even}} - \tilde{\Delta}^{\text{Even}}} + \tilde{\mathcal{Z}}^{\text{Odd}} \cdot e^{\Delta^{\text{Odd}} - \tilde{\Delta}^{\text{Odd}}} \\ &= e^{\Psi^{\text{Even}}} + e^{\Psi^{\text{Odd}}}. \end{aligned}$$

So Proposition 2.7 finishes the proof. □

2.5. Sampling independent sets. As is amply clear from the above few sections, there is a deep connection between approximating a partition function and understanding the structure of samples from the relevant distribution. Indeed, the latter can be leveraged to prove our second main theorem, Theorem 1.3, which states that the approximate sampling algorithm ApproxSampler given in Definition 1.2 gives a good approximation to the uniform distribution on all independent sets. Recalling the definition of ApproxSampler, we give some brief intuition regarding this result.

First of all, the relative sizes of $\mathcal{Z}^{\text{Even}}$ and $\mathcal{Z}^{\text{Odd}}$ are the right quantities to use in deciding whether the constructed independent set will be mostly contained in the Odd side or the Even side, respectively. Next, supposing we've chosen Even to be the defect side, the statement of Corollary 2.6 can be viewed as a mnemonic for the sampling algorithm ApproxSampler as follows: To sample from the distribution with partition function $\mathcal{Z}^{\text{Even}}$, first sample vertices independently (which has partition function $\tilde{\mathcal{Z}}^{\text{Even}}$). This gives the wrong weight to the dimers which will inevitably appear when $p \leq 2 - \sqrt{2}$, so to correct this we

remove all such dimers (which corresponds to dividing the partition function by $e^{\tilde{\Delta}}$) and then resample the dimers according to the correct distribution (which corresponds to multiplying the partition function by e^{Δ}).

Of course, to make this rigorous we must carefully extract the relevant parts of the various propositions we prove, and so we cannot yet provide a proof of Theorem 1.3 as we did with Theorem 1.1 above. Instead, the proof of Theorem 1.3 will be presented in Section 8.

2.6. Organization of the article. We begin in Section 3 with a brief review of polymer models and a proof of Proposition 2.1, which is restated as Proposition 3.4.

Next, in Section 4, we prove that only singletons and dimers are relevant for the regime we are considering, i.e. we prove Proposition 2.2, which is split into Propositions 4.1 and 4.2.

In Section 5, we survey the case $p > 2 - \sqrt{2}$, where in fact the dimers are also irrelevant and an analysis of the singletons suffices. Much of this case was already covered in our previous work [6], so this section is mostly an informal overview which gives some context for the next section.

Section 6 contains our analysis of the joint behavior of the singletons and dimers, and contains the bulk of the technical work done in this article, as well as the main novel conceptual contributions. We prove Corollary 2.6 via Proposition 6.1, showing that, for $p > \frac{1}{2} - \gamma$, the dimers do not interact with the singletons or each other when they are sampled independently.

Next, in Section 7, we prove the distributional limit statement, Proposition 2.7, restated as Proposition 7.6 below, showing that the two quantities Ψ^{Even} and Ψ^{Odd} in the exponents of our representation for $i(Q_{d,p})$ have independent Gaussian limits after proper rescaling.

Finally, in Section 8, we prove Theorem 1.3, showing that the algorithm ApproxSampler of Definition 1.2 yields a good approximation to a uniformly random independent set in $Q_{d,p}$.

This article contains two appendices: Appendix A contains various moment calculations which are used throughout, and Appendix B contains mostly standard lemmas relating binomial probabilities and entropy functions, which are used primarily in the proof of Proposition 6.1.

3. POLYMER DECOMPOSITION

In this section, we set up the polymer framework which makes precise the intuition that a typical independent set in $Q_{d,p}$ is only a small perturbation away from one of the two *ground states*, which consist of independent sets entirely contained within one of the two halves of the hypercube. Moreover, the perturbation can be represented by a configuration of *polymers*, which are defects in the opposite side of the hypercube from the ground state. In general, this strategy is made rigorous via *polymer models*, which can translate a variety of low-temperature statistical mechanical models into an associated high-temperature hard-core model on a different graph.

One key appeal of polymer models is the *cluster expansion* which represents their partition functions in a convenient form which can often yield insights into the model in question. However, as previously mentioned in Section 2.1, it should be noted that the present work *does not* use the cluster expansion beyond one key

annealed input to be discussed below. Instead, we work with the partition functions of (random) polymer models directly as random variables.

We remark that many of the results of this section already appeared, implicitly or otherwise, in our previous work [6]: see Section 4 of that article, and in particular Lemma 4.8 therein. Nevertheless, we provide a complete exposition in the present article for the reader's convenience, and the proofs we present here have been reworked for improved clarity.

3.1. Review of polymer models and the cluster expansion. To discuss general polymer models briefly, let us consider a general set Polymers of possible polymers $\mathbf{p}$ (while the precise definition will appear shortly, the reader can think of them as small connected sets for now). Suppose that there is a pairwise compatibility relation on Polymers , and say that a collection $S \subseteq \text{Polymers}$ is compatible if every pair of distinct $\mathbf{p}_1, \mathbf{p}_2 \in S$ are compatible. Additionally suppose that there is a weight $\omega(\mathbf{p})$ for each polymer $\mathbf{p}$. This data defines a polymer model with a partition function given by

$$\sum_{\substack{S \subseteq \text{Polymers} \\ \text{compatible}}} \prod_{\mathbf{p} \in S} \omega(\mathbf{p}).$$

The first step in using polymer models is to show that the partition function of a model of interest is close to the partition function of a particular polymer model, or perhaps a sum of multiple such partition functions. In the present article, for instance, this is the content of Proposition 2.1, which will be proven below as Proposition 3.4 after we introduce the details of our polymer model.

Often, to extract useful information or efficiently compute approximations to these polymer model partition functions, one can use the cluster expansion. Briefly, a *cluster* of polymers is a multiset Γ of polymers which are mutually *incompatible*, in the sense that the *incompatibility graph* H_Γ , which has vertex set Γ and an edge between incompatible polymers, is connected. The weight of a cluster Γ is given by

$$\omega(\Gamma) := \phi(H_\Gamma) \prod_{\mathbf{p} \in \Gamma} \omega(\mathbf{p}),$$

where $\phi(H)$ is the *Ursell function* of a graph H , defined by

$$\phi(H) := \frac{1}{|V(H)|!} \sum_{\substack{A \subseteq E(H) \\ \text{spanning} \\ \text{connected}}} (-1)^{|A|}.$$

The precise details here are not so important for our purposes, but the key point is the following formal equality, which is known as the *cluster expansion*:

$$\sum_{\substack{S \subseteq \text{Polymers} \\ \text{compatible}}} \prod_{\mathbf{p} \in S} \omega(\mathbf{p}) = \exp \left(\sum_{\Gamma \text{ cluster}} \omega(\Gamma) \right). \quad (3.1)$$

Note that the above is a priori only a formal equality; since a cluster is a multiset, even if there are only finitely many possible polymers, there are infinitely many possible clusters. So an important step in using the cluster expansion is to prove that the sum inside the exponential actually converges.

3.2. The cluster expansion of an annealed model. In the present article, we consider a *random* polymer model, i.e. one where the weights are all random variables, determined by the percolation configuration. This complicates matters and it is not a priori clear how to prove directly that this polymer model gives a good approximation for $i(Q_{d,p})$, as stated in Proposition 2.1. However, [13] applied the technique outlined above to an *annealed* version of the model, where a set of vertices in Q_d (not necessarily an independent set) is given a weight equal to the probability that it is independent in $Q_{d,p}$. The partition function of this model is then exactly equal to $\mathbb{E}[i(Q_{d,p})]$. We will use this input and a careful application of Markov's inequality to obtain Proposition 2.1.

We remark that [13] actually used a similar method to compute *all* moments of $i(Q_{d,p})$, by constructing a deterministic model for each $k \in \mathbb{N}$ with partition functions $\mathbb{E}[i(Q_{d,p})^k]$. The main contribution of [13] is to use the cluster expansion of associated polymer models in order to get accurate estimates of the moments of $i(Q_{d,p})$, which can be computed via an algorithm. Unfortunately, the moments of $i(Q_{d,p})$ themselves are not descriptive of the typical behavior of $i(Q_{d,p})$ unless p is very close to 1, but it turns out that some of the results of [13] will be useful for our purposes.

We will henceforth only focus on the $k = 1$ polymer model with partition function $\mathbb{E}[i(Q_{d,p})]$, and we now briefly describe this polymer model, introduced by [13].

Definition 3.1 (Language relating to polymer models from [13]). *In these definitions, we use the notation $N(A)$ to denote the set of neighbors of A in Q_d , rather than the number of neighbors. Additionally, $\mathcal{H}$ denotes either Even or Odd.*

- A subset $A \subseteq \mathcal{H}$ is said to be 2-linked if there is no decomposition $A = A_1 \cup A_2$ into nonempty sets A_1 and A_2 which satisfy $N(A_1) \cap N(A_2) = \emptyset$.
- For a subset $A \subseteq \mathcal{H}$, we denote by $[A]$ the closure of A , defined to be the largest subset of $\mathcal{H}$ with the same neighborhood set as A . In other words, $[A] = \{v \in \mathcal{H} : N(v) \subseteq N(A)\}$.
- A 2-linked set $\mathbf{p} \subseteq \mathcal{H}$ with $|\mathbf{p}| \leq \frac{3}{4}2^{d-1}$ is called a polymer.
- Polymers $\mathbf{p}_1$ and $\mathbf{p}_2$ are said to be compatible if $\mathbf{p}_1 \cup \mathbf{p}_2$ is not 2-linked.
- We denote by $\text{PD}^{\mathcal{H}}$ the set of pairwise compatible collections of polymers.
- The weight of a polymer $\mathbf{p}$ is $\omega(\mathbf{p}) = \mathbb{E}[\varphi_{\mathbf{p}}]$, where $\varphi_{\mathbf{p}} = 2^{-N_{\mathbf{p}}(\mathbf{p})}$. We remind the reader that $N_{\mathbf{p}}(\mathbf{p})$ denotes the (random) number of vertices neighboring $\mathbf{p}$ in the percolated graph $Q_{d,p}$.
- We define clusters and their weights as in the previous section, which is a general construction that works for all polymer models. The set of clusters will be denoted by $\mathcal{C}^{\mathcal{H}}$.

A few remarks are in order. First, the number $\frac{3}{4}$ which appears in the definition of a polymer is a somewhat arbitrary choice, and, as remarked in [13], any constant between $\frac{1}{2}$ and 1 can be used, for reasons which we will not explore in this article. Second, the definition given in [13, Equation 14] for the weight $\omega(\mathbf{p})$ of a polymer differs from the expression presented here, but it is easy to see that the two expressions are equal. Finally, note that the two polymer models, one for each half $\mathcal{H} \in \{\text{Even}, \text{Odd}\}$, actually have identical partition functions, but we will nonetheless keep these separate.

With these definitions in hand, we can state one of the main results of [13], which expands $\mathbb{E}[i(Q_{d,p})]$ in a sum of two convergent cluster expansions for these two polymer models.

Theorem 3.2 (Theorem 3.1 of [13], with $k = 1$). *For $p \in (0, 1)$,*

$$\mathbb{E}[i(Q_{d,p})] = 2^{2^{d-1}} \left(\exp \left(\sum_{\Gamma \in \mathcal{C}^{\text{Even}}} \omega(\Gamma) \right) + \exp \left(\sum_{\Gamma \in \mathcal{C}^{\text{Odd}}} \omega(\Gamma) \right) \right) \cdot \left(1 + \exp \left(-\Omega(2^d/d^4) \right) \right),$$

where the cluster expansion series inside the exponentials converge absolutely.

We remark here that the series in the above exponential expressions may be evaluated via bounds (arising from the graph container method) on the expectations of terms corresponding to higher order clusters of polymers, and one can obtain

$$\mathbb{E}[i(Q_{d,p})] = 2^{2^{d-1}} \cdot 2 \cdot \exp \left(\frac{1}{2} (2-p)^d \cdot (1 + o(1)) \right), \quad (3.2)$$

which is a simplified version of [13, Theorem 1.1].

Now, by the formal cluster expansion (3.1), which converges absolutely as stated in Theorem 3.2, for each $\mathcal{H} \in \{\text{Even}, \text{Odd}\}$, we have

$$\begin{aligned} \exp \left(\sum_{\Gamma \in \mathcal{C}^{\mathcal{H}}} \omega(\Gamma) \right) &= \sum_{S \in \text{PD}^{\mathcal{H}}} \prod_{\mathbf{p} \in S} \omega(\mathbf{p}) \\ &= \sum_{S \in \text{PD}^{\mathcal{H}}} \prod_{\mathbf{p} \in S} \mathbb{E}[\varphi_{\mathbf{p}}] \\ &= \mathbb{E} \left[\sum_{S \in \text{PD}^{\mathcal{H}}} \prod_{\mathbf{p} \in S} \varphi_{\mathbf{p}} \right] = \mathbb{E}[\mathcal{Z}^{\mathcal{H}}], \end{aligned}$$

using at the second-to-last equality the fact that the $\varphi_{\mathbf{p}}$ are independent for $\mathbf{p} \in S$ when $S \in \text{PD}^{\mathcal{H}}$, as there are no shared vertices between any $\mathbf{p} \in S$. The last equality is simply the definition of $\mathcal{Z}^{\mathcal{H}}$ from (2.5). So, as an immediate corollary of Theorem 3.2, we obtain the following annealed version of Proposition 2.1 providing an approximation of the expected number of independent sets in $Q_{d,p}$.

Corollary 3.3. *For $p \in (0, 1)$,*

$$\mathbb{E} \left[\frac{i(Q_{d,p})}{2^{2^{d-1}}} \right] = \mathbb{E}[\mathcal{Z}^{\text{Even}} + \mathcal{Z}^{\text{Odd}}] \cdot \left(1 + \exp \left(-\Omega(2^d/d^4) \right) \right).$$

3.3. Polymer decomposition of the quenched partition function. Using Corollary 3.3 as well as another input from [13] to be mentioned below, we can now obtain a quenched approximation for the number of independent sets in $Q_{d,p}$. This takes the form of a sum of two polymer model partition functions, which we will analyze directly in the sequel. The reader familiar with polymer models should note that we will *not* apply the cluster expansion to these partition functions at this point. The following proposition was previously stated as Proposition 2.1, and we remind the reader that $A \stackrel{\mathbb{P}}{\sim} B$ means that $\frac{A}{B} \xrightarrow{\mathbb{P}} 1$ as $d \rightarrow \infty$.

Proposition 3.4. *For $p \in (0, 1)$,*

$$\frac{i(Q_{d,p})}{2^{2^{d-1}}} \stackrel{\mathbb{P}}{\sim} \mathcal{Z}^{\text{Even}} + \mathcal{Z}^{\text{Odd}}.$$

where the terms on the RHS were defined in (2.5).

The proof is an application of Markov's inequality; however, since Corollary 3.3 only says that the expectations of the two sides in Proposition 3.4 are close, we will also need some form of monotonicity as well as an a priori bound on $i(Q_{d,p})$ to make the argument work.

As previously mentioned in Section 2, specifically in (2.2), we have

$$\frac{i(Q_{d,p})}{2^{2^{d-1}}} = \sum_{T \subseteq \text{Even}} 2^{-N_p(T)}.$$

Note that, if T can be represented as a union of well-separated polymers $T = \bigcup_{\mathbf{p} \in S} \mathbf{p}$ for $S \in \text{PD}^{\text{Even}}$, then

$$2^{-N_p(T)} = \prod_{\mathbf{p} \in S} \varphi_{\mathbf{p}},$$

and so the quantity $2^{-N_p(T)}$ appears as a term inside the expression $\mathcal{Z}^{\text{Even}}$ on the right-hand side of Proposition 3.4. In other words, every independent set I with an even side which is representable via a polymer configuration is accounted for in $\mathcal{Z}^{\text{Even}}$, and the number of such sets is exactly $2^{2^{d-1}} \mathcal{Z}^{\text{Even}}$. Similarly, every independent set I with an odd side which is representable via a polymer configuration is accounted for in $\mathcal{Z}^{\text{Odd}}$ and the number of such sets is exactly $2^{2^{d-1}} \mathcal{Z}^{\text{Odd}}$.

However, there may be independent sets for which neither side is representable via a polymer configuration, which are not counted in the right-hand side of Proposition 3.4. Moreover, there are also independent sets which are counted in both $\mathcal{Z}^{\text{Even}}$ and $\mathcal{Z}^{\text{Odd}}$, since both sides can be decomposed into well-separated polymers. This means that the two sides of Proposition 3.4 are not immediately comparable via an inequality. To remedy this, we use the following lemma from [13] which gives another annealed bound on the total number of independent sets for which both sides are not representable by a polymer configuration.

Lemma 3.5 (Lemma 3.3 of [13], with $k = 1$). *For $p \in (0, 1)$,*

$$\sum_{\substack{I \subseteq Q_d \\ |[I \cap \text{Even}]|, |[I \cap \text{Odd}]| > \frac{3}{4} 2^{d-1}}} \mathbb{P}[I \text{ is independent in } Q_{d,p}] \leq \mathbb{E}[i(Q_{d,p})] \cdot \exp\left(-\Omega(2^d/d)\right).$$

Recall from Definition 3.1 that $[A]$ is the closure of a set, and that $||[\mathbf{p}]|| \leq \frac{3}{4} 2^{d-1}$ is the only requirement, beyond being 2-linked, for $\mathbf{p}$ to be a polymer under Definition 3.1. So if $|[I \cap \text{Even}]| \leq \frac{3}{4} 2^{d-1}$, then $I \cap \text{Even}$ can be represented by a polymer configuration, and similarly for $I \cap \text{Odd}$. Thus we immediately obtain the following corollary.

Corollary 3.6. *Let TooBig denote the number of independent sets $I \subseteq Q_d$ for which neither $I \cap \text{Even}$ nor $I \cap \text{Odd}$ are representable via a polymer configuration. Then, for all $p \in (0, 1)$,*

$$\mathbb{E}[\text{TooBig}] \leq \mathbb{E}[i(Q_{d,p})] \cdot \exp\left(-\Omega(2^d/d)\right).$$

With this in hand, we can now prove Proposition 3.4.

Proof of Proposition 3.4. Define

$$\mathcal{X} := \frac{i(Q_{d,p})}{2^{2^{d-1}}} - \mathcal{Z}^{\text{Even}} - \mathcal{Z}^{\text{Odd}},$$

so that our goal is to show that

$$\frac{\mathcal{X}}{i(Q_{d,p})/2^{2^{d-1}}} \xrightarrow{\mathbb{P}} 0. \quad (3.3)$$

An initial idea might be to use Markov's inequality, since the expectation of $\mathcal{X}$ is small; indeed Corollary 3.3 says that

$$|\mathbb{E}[\mathcal{X}]| \leq \mathbb{E}\left[\frac{i(Q_{d,p})}{2^{2^{d-1}}}\right] \cdot \exp\left(-\Omega(2^d/d^4)\right). \quad (3.4)$$

However, $\mathcal{X}$ is not positive (note that the absolute value is on the outside of the expectation above), and moreover the denominator of (3.3) is random. Thus we need to be a bit more careful.

First we will show that

$$\frac{\text{TooBig}}{i(Q_{d,p})} \xrightarrow{\mathbb{P}} 0, \quad (3.5)$$

where TooBig is defined as in Corollary 3.6. In order to apply Markov's inequality with that corollary, we will need an a priori lower bound on $i(Q_{d,p})$. Since $i(Q_{d,p}) \geq 2 \cdot 2^{2^{d-1}}$ deterministically, by (3.2) we have

$$i(Q_{d,p}) \geq \exp\left(-\frac{1}{2}(2-p)^d(1+o(1))\right) \cdot \mathbb{E}[i(Q_{d,p})]. \quad (3.6)$$

So for any $\varepsilon > 0$, we have

$$\begin{aligned} \mathbb{P}[\text{TooBig} \geq \varepsilon \cdot i(Q_{d,p})] &\leq \mathbb{P}\left[\text{TooBig} \geq \varepsilon \cdot \exp\left(-\frac{1}{2}(2-p)^d(1+o(1))\right) \mathbb{E}[i(Q_{d,p})]\right] \\ &\leq \frac{\exp(-\Omega(2^d/d))}{\varepsilon \cdot \exp(-\frac{1}{2}(2-p)^d(1+o(1)))}, \end{aligned}$$

which tends to zero as $d \rightarrow \infty$ since $p > 0$, proving (3.5).

Next we turn to (3.3). By the definitions of $\mathcal{X}$ and TooBig, as well as the discussion (between the statements of Proposition 3.4 and Lemma 3.5 above) of how the polymer partition functions relate to the number of polymer-decomposable independent sets,

$$\text{TooBig} - 2^{2^{d-1}} \mathcal{X} = \underbrace{2^{2^{d-1}} \cdot \mathcal{Z}^{\text{Even}}}_{\substack{\text{number of independent sets} \\ \text{polymer-decomposable} \\ \text{on even side}}} + \underbrace{2^{2^{d-1}} \cdot \mathcal{Z}^{\text{Odd}}}_{\substack{\text{number of independent sets} \\ \text{polymer-decomposable} \\ \text{on odd side}}} - \underbrace{(i(Q_{d,p}) - \text{TooBig})}_{\substack{\text{number of independent sets} \\ \text{polymer-decomposable} \\ \text{on either side}}} \geq 0,$$

since the above quantity is exactly the number of independent sets for which *both sides* have representations as polymer configurations. So, using this nonnegativity, we now have

$$\begin{aligned}
& \mathbb{P}[\text{TooBig} - 2^{2^{d-1}} \mathcal{X} \geq \varepsilon \cdot i(Q_{d,p})] \\
& \leq \mathbb{P}\left[\text{TooBig} - 2^{2^{d-1}} \mathcal{X} \geq \varepsilon \cdot \exp\left(-\frac{1}{2}(2-p)^d(1+o(1))\right) \cdot \mathbb{E}[i(Q_{d,p})]\right] \quad (\text{by (3.6)}) \\
& \leq \frac{\mathbb{E}[\text{TooBig} - 2^{2^{d-1}} \mathcal{X}]}{\varepsilon \cdot \exp\left(-\frac{1}{2}(2-p)^d(1+o(1))\right) \cdot \mathbb{E}[i(Q_{d,p})]} \quad (\text{by Markov's inequality}) \\
& \leq \frac{\mathbb{E}[\text{TooBig}] + 2^{2^{d-1}} |\mathbb{E}[\mathcal{X}]|}{\varepsilon \cdot \exp\left(-\frac{1}{2}(2-p)^d(1+o(1))\right) \cdot \mathbb{E}[i(Q_{d,p})]} \\
& \leq \frac{\exp(-\Omega(2^d/d)) + \exp(-\Omega(2^d/d^4))}{\varepsilon \cdot \exp\left(-\frac{1}{2}(2-p)^d(1+o(1))\right)}. \quad (\text{by (3.4) and Corollary 3.6})
\end{aligned}$$

The above tends to zero since $p > 0$, proving that

$$\frac{\text{TooBig}}{i(Q_{d,p})} - \frac{\mathcal{X}}{i(Q_{d,p})/2^{2^{d-1}}} \xrightarrow{\mathbb{P}} 0.$$

Combining this with (3.5) completes the proof. $\square$

4. BOUNDING THE SIZE OF POLYMERS

In principle, a polymer configuration S can contain polymers of any size, and there are increasingly many possibilities for a polymer of size k as k grows, which would be somewhat cumbersome to deal with. However, it turns out that for p bounded away from zero, only finitely many types of polymers are relevant in the decomposition. This is formalized in Proposition 4.1 below, which is a more general version of Proposition 2.2 above.

In the sequel, we only consider p large enough that in fact only polymers of size ≤ 2 (i.e. singletons and dimers) are relevant. Nevertheless, we state the results in this section in greater generality in the hopes that they may be useful for future work investigating smaller values of p . In addition, we remark that a version of these results (covering the case where only singletons are relevant) already appeared in our previous work [6, Lemma 4.9]. Thus, as with Section 3 above, this section is provided more for completeness, in addition to providing streamlined proofs, rather than presenting new ideas.

Note that in this section we drop the superscript $\mathcal{H}$ denoting the side of the hypercube; for instance, $\mathcal{Z}^{\mathcal{H}}$ will simply become $\mathcal{Z}$. This is because all results in this section focus on only one side at a time, and hold equally for the even and the odd sides.

Recall from (2.5) and (2.8) the notation

$$\mathcal{Z} = \sum_{S \in \text{PD}} \prod_{\mathbf{p} \in S} \varphi_{\mathbf{p}} \quad \text{and} \quad \mathcal{Z}_{\leq k} = \sum_{S \in \text{PD}_{\leq k}} \prod_{\mathbf{p} \in S} \varphi_{\mathbf{p}}.$$

Here, and in the sequel, $\text{PD} = \text{PD}^{\mathcal{H}}$ denotes the collection of separated polymer configurations on side $\mathcal{H}$ of the hypercube, and $\text{PD}_{\leq k}$ denotes the subcollection where all the polymers are constrained to have size at most k .

The following proposition gives the threshold in p for the irrelevance of polymers of size larger than k , for each $k \in \mathbb{N}$. This proposition will be proved in Section 4.2 below.

Proposition 4.1. *For each $k \geq 1$, if $p > 2 - 2^{k/(k+1)}$, then we have $\mathcal{Z} \stackrel{\mathbb{P}}{\sim} \mathcal{Z}_{\leq k}$.*

4.1. The independent singletons model. In this section, as a warm-up which will also be useful in the overall proof of our main theorem, we prove that a random set drawn according to a simpler product measure (which was alluded to in the proof ideas presented in Section 2, and which will appear repeatedly in our arguments) also satisfies a polymer decomposition with small polymers when p is bounded away from zero. Specifically, we consider the model which samples $S \subseteq \mathcal{H}$ by choosing each vertex v to be included in S independently with probability $\frac{\varphi_v}{1+\varphi_v}$, where we recall from (2.4) that $\varphi_v = 2^{-N_p(v)}$. If the weight of a set S is written simply as $\prod_{v \in S} \varphi_v$, then the partition function of this model is

$$\prod_{v \in \mathcal{H}} (1 + \varphi_v).$$

Recall from (2.6) the notation $\tilde{\varphi}_{\mathbf{p}} = \prod_{v \in \mathbf{p}} \varphi_v$. Then in expanding the above product, each term can be written as a product of $\tilde{\varphi}_{\mathbf{p}}$ for some polymers $\mathbf{p}$ in a valid polymer configuration, as follows

$$\prod_{v \in \mathcal{H}} (1 + \varphi_v) = \sum_{S \in \widetilde{\text{PD}}} \prod_{\mathbf{p} \in S} \tilde{\varphi}_{\mathbf{p}} = \tilde{\mathcal{Z}},$$

recalling the definition of $\tilde{\mathcal{Z}}$ from (2.7). Note that in defining $\tilde{\mathcal{Z}}$ we use a slightly relaxed notion of polymers than that which appears in [13] and in Definition 3.1; specifically, we do not impose the restriction that $||[\mathbf{p}]|| \leq \frac{3}{4}2^{d-1}$. We denote the collection of these polymers by $\widetilde{\text{Polymers}}$, and define $\widetilde{\text{PD}}$ and $\widetilde{\text{PD}}_{\leq k}$ analogously to PD and $\text{PD}_{\leq k}$. However we remark that for large enough d and for any fixed k , we have $\widetilde{\text{PD}}_{\leq k} = \text{PD}_{\leq k}$, and so we will simply use the latter notation in the sequel, where we will only work with polymers of bounded size.

The following proposition is the analogue of Proposition 4.1 for this model. It says that for large enough p , only small enough polymers are relevant. Recall from (2.8) the notation

$$\tilde{\mathcal{Z}}_{\leq k} = \sum_{S \in \text{PD}_{\leq k}} \prod_{\mathbf{p} \in S} \tilde{\varphi}_{\mathbf{p}}.$$

Proposition 4.2. *For each $k \geq 1$, if $p > 2 - 2^{k/(k+1)}$, then we have $\tilde{\mathcal{Z}} \stackrel{\mathbb{P}}{\sim} \tilde{\mathcal{Z}}_{\leq k}$.*

The key lemma is the following, which bounds the expected total weight of polymers of size at least k . We denote the collection of relaxed polymers of size exactly j by $\widetilde{\text{Polymers}}_j$.

Lemma 4.3. *For every $p \in (0, 1)$ and every fixed $k \geq 0$,*

$$\sum_{j > k} \sum_{\mathbf{p} \in \widetilde{\text{Polymers}}_j} \mathbb{E}[\tilde{\varphi}_{\mathbf{p}}] \leq \text{poly}(d) \cdot \left(\frac{(2-p)^{k+1}}{2^k} \right)^d.$$

Proof of Lemma 4.3. The number of polymers of size j is at most $2^d(ed^2)^j$; to see why, consider the following method of constructing a polymer. First, choose one of the j^{j-2} labeled trees of size j . Next, let choose any of the 2^{d-1} vertices of $\mathcal{H}$ to be the first vertex, with label 1. Perform breadth-first or depth-first search in the

tree, always choosing the vertex with minimum label among possible choices; when a vertex i is selected, attach it to its parent k (with the root having label 1) by selecting one of the $\leq d^2$ vertices sharing a neighbor with the vertex labeled k , and giving it the label i .

Every polymer of size j is formed by an instance of this construction, since the 2-linked graph structure of any polymer has at least one spanning tree. Moreover, each polymer is counted at least $j!$ times since we can permute the labels of the tree arbitrarily, and in any permutation of the labels it is still possible to generate the same polymer. So the number of polymers of size j is at most

$$j^{j-2} \cdot 2^{d-1} (d^2)^{j-1} \cdot \frac{1}{j!} \leq 2^d (d^2)^j \frac{j^j}{j!} \leq 2^d (ed^2)^j$$

as required.

Now for each polymer $\mathbf{p}$ with size j , we have $\tilde{\varphi}_{\mathbf{p}} \stackrel{d}{=} 2^{-\text{Bin}(jd, p)}$. So by a simple calculation like the one in presented in Lemma A.1 (a), we have

$$\mathbb{E}[\tilde{\varphi}_{\mathbf{p}}] = \left(\frac{2-p}{2} \right)^{jd},$$

and so

$$\sum_{j>k} \sum_{\mathbf{p} \in \widetilde{\text{Polymers}}_j} \mathbb{E}[\tilde{\varphi}_{\mathbf{p}}] \leq 2^d \sum_{j>k} \left(ed^2 \left(\frac{2-p}{2} \right)^d \right)^j.$$

For large enough d , we have $ed^2 \left(\frac{2-p}{2} \right)^d < \frac{1}{2}$, and so the geometric sum is dominated by 2 times its first term, which is the term corresponding to $j = k+1$, i.e. we have

$$\sum_{j>k} \sum_{\mathbf{p} \in \widetilde{\text{Polymers}}_j} \mathbb{E}[\tilde{\varphi}_{\mathbf{p}}] \leq 2^d \cdot 2 \cdot \left(ed^2 \left(\frac{2-p}{2} \right)^d \right)^{k+1},$$

which is the desired bound. $\square$

With this lemma in hand, we can now prove Proposition 4.2.

Proof of Proposition 4.2. Consider the quantity

$$\exp \left(\sum_{j>k} \sum_{\mathbf{p} \in \widetilde{\text{Polymers}}_j} \tilde{\varphi}_{\mathbf{p}} \right) \cdot \sum_{S \in \text{PD}_{\leq k}} \prod_{\mathbf{p} \in S} \tilde{\varphi}_{\mathbf{p}}.$$

By Taylor expanding the exponential of the (finite) sum over polymers of size larger than k , one obtains various products of terms in the sum, and those corresponding to collections of m distinct polymers appear with a coefficient of 1, since they appear $m!$ times, each with a coefficient of $\frac{1}{m!}$ from the Taylor series for the exponential function. Moreover, since all terms in the sum are positive, we arrive at

$$\exp \left(\sum_{j>k} \sum_{\mathbf{p} \in \widetilde{\text{Polymers}}_j} \tilde{\varphi}_{\mathbf{p}} \right) \cdot \sum_{S \in \text{PD}_{\leq k}} \prod_{\mathbf{p} \in S} \tilde{\varphi}_{\mathbf{p}} \geq \sum_{T \in \widetilde{\text{PD}}} \prod_{\mathbf{p} \in T} \tilde{\varphi}_{\mathbf{p}} = \tilde{\mathcal{Z}},$$

as soon as d is large enough so that $\widetilde{\text{PD}}_{\leq k} = \text{PD}_{\leq k}$. On the other hand, by definition we have

$$\sum_{S \in \text{PD}_{\leq k}} \prod_{\mathbf{p} \in S} \tilde{\varphi}_{\mathbf{p}} = \tilde{\mathcal{Z}}_{\leq k},$$

and we thus obtain

$$\left(\exp \left(\sum_{j>k} \sum_{\mathbf{p} \in \widetilde{\text{Polymers}}_j} \tilde{\varphi}_{\mathbf{p}} \right) - 1 \right) \cdot \sum_{S \in \text{PD}_{\leq k}} \prod_{\mathbf{p} \in S} \tilde{\varphi}_{\mathbf{p}} \geq \tilde{\mathcal{Z}} - \tilde{\mathcal{Z}}_{\leq k} \geq 0,$$

which, upon rearranging, yields

$$\exp \left(\sum_{j>k} \sum_{\mathbf{p} \in \widetilde{\text{Polymers}}_j} \tilde{\varphi}_{\mathbf{p}} \right) \geq \frac{\tilde{\mathcal{Z}}}{\tilde{\mathcal{Z}}_{\leq k}} \geq 1.$$

So the proof will be finished if we can show that

$$\sum_{j>k} \sum_{\mathbf{p} \in \widetilde{\text{Polymers}}_j} \tilde{\varphi}_{\mathbf{p}} \xrightarrow{\mathbb{P}} 0.$$

But this follows from Lemma 4.3 and Markov's inequality, since when $p > 2 - 2^{k/(k+1)}$ the upper bound in that lemma tends to 0 as $d \rightarrow \infty$. $\square$

4.2. The true independent set model. We now turn to the proof of Proposition 4.1, which is the same statement as Proposition 4.2 which was just proved, but with

$$\tilde{\varphi}_{\mathbf{p}} = \prod_{v \in \mathbf{p}} 2^{-N_p(v)} \quad \text{replaced by} \quad \varphi_{\mathbf{p}} = 2^{-N_p(\mathbf{p})}.$$

The key lemma in the previous proof, Lemma 4.3, crucially leveraged the independence of $\varphi_v = 2^{-N_p(v)}$ for distinct vertices to give an upper bound on the sum of $\mathbb{E}[\tilde{\varphi}_{\mathbf{p}}]$ over *all* large polymers $\mathbf{p}$. While it may be expected that a similar bound holds for $\mathbb{E}[\varphi_{\mathbf{p}}]$ for *small enough* polymers, it is far from obvious that the sum over all large polymers may be bounded in the same way, since $\varphi_{\mathbf{p}}$ may be much larger than $\tilde{\varphi}_{\mathbf{p}}$ for large polymers owing to the large number of overlapping neighborhoods of vertices in $\mathbf{p}$.

The *graph container method* is a powerful tool that allows one to bound the number of such large sets based on their neighborhood sizes, and it yields the estimates we need. This calculation was carried out in [13], using graph container results presented in [10] that follow the work of [16] closely, and the result is the following lemma, which will replace Lemma 4.3 in the proof of Proposition 4.1. Note that we are now restricting to polymers as defined by [13], which are the same polymers as those appearing in the decomposition given by Proposition 3.4. Specifically, these include the restriction that $|\mathbf{p}| \leq \frac{3}{4}2^{d-1}$, and we denote the collection of such polymers with size exactly j by Polymers_j .

Lemma 4.4 (Adapted from Lemma 4.1 of [13]). *For every $p \in (0, 1)$ and every fixed $k \geq 1$,*

$$\sum_{j>k} \sum_{\mathbf{p} \in \text{Polymers}_j} \mathbb{E}[\varphi_{\mathbf{p}}] \leq \text{poly}(d) \cdot \left(\frac{(2-p)^{k+1}}{2^k} \right)^d$$

We remark that this statement is weaker than the original version of Lemma 4.1 of [13], but will suffice for our purposes. For the reader who wishes to compare the original statement to the above statement, we remark that the original statement is about the sum over all *clusters* of polymers, and each term in the sum includes an extra exponential factor not written above. The above statement follows from the original statement by restricting the sum to only clusters of size 1 (which are the same as individual polymers), using the fact that every term in the sum is nonnegative. Moreover, the aforementioned exponential factors are all ≥ 1 , so we can simply erase them and still retain the same bound. With Lemma 4.4 in hand, Proposition 4.1 follows by the same proof as for Proposition 4.2, but with $\tilde{\varphi}_p$ replaced by φ_p .

5. THE REGIME WITH NO DIMERS

The main new technical and conceptual contribution of the present article is a sharp understanding of $i(Q_{d,p})$ in the regime of parameters where dimers are relevant. But, before pursuing that in Section 6 below, in the present section we highlight how our preparation up to this point yields a quick analysis of the regime $p > 2 - \sqrt{2} \approx 0.586$, where no dimers appear in the defect side of typical independent sets. In particular, with a bit of extra care in tracking error terms through our proofs of Propositions 3.4, 4.1, and 4.2, one can recover the results of our previous work [6], which were proven therein for $p \geq \frac{2}{3}$. We do not spell out all of the details presently as these results have already appeared in our previous article.

5.1. In-probability approximation. By Propositions 3.4 and 4.1, when $p > 2 - \sqrt{2}$, we have

$$\frac{i(Q_{d,p})}{2^{d-1}} \stackrel{\mathbb{P}}{\sim} \mathcal{Z}_1^{\text{Even}} + \mathcal{Z}_1^{\text{Odd}}.$$

We remind the reader that

$$\mathcal{Z}_1^{\mathcal{H}} = \sum_{S \in \text{PD}_1^{\mathcal{H}}} \prod_{v \in S} \varphi_v,$$

where $\text{PD}_1^{\mathcal{H}}$ denotes the collection of well-separated subsets of $\mathcal{H}$. Now, by Proposition 4.2, since $p > 2 - \sqrt{2}$, we have

$$\mathcal{Z}_1^{\mathcal{H}} \stackrel{\mathbb{P}}{\sim} \tilde{\mathcal{Z}}^{\mathcal{H}} = \prod_{v \in \mathcal{H}} (1 + \varphi_v)$$

for each $\mathcal{H}$, using the fact that $\mathcal{Z}_1^{\mathcal{H}} = \tilde{\mathcal{Z}}^{\mathcal{H}}$ since $\varphi_v = \tilde{\varphi}_v$ for singletons. Thus we have

$$\frac{i(Q_{d,p})}{2^{d-1}} \stackrel{\mathbb{P}}{\sim} \exp(\Phi_{\log}^{\text{Even}}) + \exp(\Phi_{\log}^{\text{Odd}}), \quad (5.1)$$

where

$$\Phi_{\log}^{\mathcal{H}} := \log \tilde{\mathcal{Z}}^{\mathcal{H}} = \sum_{v \in \mathcal{H}} \log(1 + \varphi_v). \quad (5.2)$$

This completes the first part of our program, namely exhibiting an explicit in-probability approximation for $i(Q_{d,p})$. In the regime $p > 2 - \sqrt{2}$, note that the quantities $\Delta^{\mathcal{H}}$ and $\tilde{\Delta}^{\mathcal{H}}$ for $\mathcal{H} \in \{\text{Even}, \text{Odd}\}$, as defined in (2.11) via sums of contributions coming from dimers, are not relevant for this approximation.

5.2. Distributional limits. Next, as will be proved in Section 7 below (specifically as Corollary 7.4), $\Phi_{\log}^{\text{Even}}$ and $\Phi_{\log}^{\text{Odd}}$ satisfy a joint central limit theorem with explicit recentering and rescaling, i.e. we have

$$\left(\frac{\Phi_{\log}^{\text{Even}} - \mu_1}{\sigma}, \frac{\Phi_{\log}^{\text{Odd}} - \mu_1}{\sigma} \right) \xrightarrow{w} \mathcal{N}(0, 1) \otimes \mathcal{N}(0, 1),$$

where, as previously defined in (1.8) and (1.5),

$$\begin{aligned} \mu_1 &= \frac{1}{2}(2-p)^d - \frac{1}{4}\left(2 - \frac{3}{2}p\right)^d + \frac{1}{6}\left(2 - \frac{7}{4}p\right)^d, \\ \sigma^2 &= \frac{1}{2}\left(2 - \frac{3}{2}p\right)^d. \end{aligned} \tag{5.3}$$

In fact, we may simplify the centering quantity in the regime $p > 2 - \sqrt{2}$. Specifically, in this regime the third term of μ_1 is $\ll \sigma$, which follows from the fact that

$$\left(2 - \frac{7}{4}p\right)^2 < 2 - \frac{3}{2}p$$

for such p . To verify the above inequality, note that the difference between the left and right-hand sides is

$$\frac{49}{16}p^2 - \frac{11}{2}p + 2,$$

which has roots

$$\frac{11 \pm \sqrt{11^2 - 2 \cdot 49}}{49/8} \approx 0.506, 1.289.$$

Therefore, since $2 - \sqrt{2} \approx 0.586$, the difference is negative on the interval $(2 - \sqrt{2}, 1)$. All this is to say that if we define

$$\mu' = \frac{1}{2}(2-p)^d - \frac{\sigma^2}{2},$$

then we in fact have

$$\left(\frac{\Phi_{\log}^{\text{Even}} - \mu'}{\sigma}, \frac{\Phi_{\log}^{\text{Odd}} - \mu'}{\sigma} \right) \xrightarrow{w} \mathcal{N}(0, 1) \otimes \mathcal{N}(0, 1) \tag{5.4}$$

when $p > 2 - \sqrt{2}$.

5.3. Recovery of previous results. With a bit of extra care in tracking the errors through our proofs, our present results may be used to witness a previously-discovered phase transition at $p = \frac{2}{3}$, whereby the quantity $i(Q_{d,p})$ itself exhibits a Gaussian scaling limit for $p > \frac{2}{3}$, a lognormal limit for $p = \frac{2}{3}$, and no scaling limit for $p < \frac{2}{3}$. As these distributional limit results were already proved in [6], we do not provide full details presently, instead being satisfied with an informal discussion.

In the statements of Propositions 3.4, 4.1, and 4.2, various quantities are related via $\stackrel{\mathbb{P}}{\sim}$, meaning there is a multiplicative $(1 + o(1))$ -in-probability error. By carefully following the proofs this may be improved to a multiplicative error of $(1 + o(\sigma))$ -in-probability when $p > \frac{2}{3}$, which is the regime for which $\sigma \rightarrow 0$ (recall the formula for σ^2 in (5.3)). Additionally, $\sigma^2 \ll \sigma$ in this regime, so if we define

$$\hat{\mu} = \frac{1}{2}(2-p)^d,$$

then the improved version of (5.1) combined with (5.4) yields

$$\begin{aligned} \frac{i(Q_{d,p})}{2^{2^{d-1}} e^{\hat{\mu}}} &= \left(\exp\left(\Phi_{\log}^{\text{Even}} - \hat{\mu}\right) + \exp\left(\Phi_{\log}^{\text{Odd}} - \hat{\mu}\right) \right) (1 + o(\sigma)) \\ &= \left(2 + \Phi_{\log}^{\text{Even}} + \Phi_{\log}^{\text{Odd}} - 2\hat{\mu} \right) (1 + o(\sigma)). \end{aligned}$$

Here the $o(\sigma)$ error term satisfies $\frac{o(\sigma)}{\sigma} \xrightarrow{\mathbb{P}} 0$ and we have used the fact that the (5.4) implies that the error in approximating the exponentials by linear functions is $o(\sigma)$, since $(\Phi_{\log}^{\text{Even}} - \hat{\mu})^2 = O(\sigma^2) = o(\sigma)$ because $\sigma \rightarrow 0$ (here again the $O(\sigma^2)$ means in-probability). The central limit theorem (5.4) also shows that

$$\frac{\Phi_{\log}^{\text{Even}} + \Phi_{\log}^{\text{Odd}} - 2\hat{\mu}}{\sigma} \xrightarrow{w} \mathcal{N}(0, 2),$$

and so, since the approximation is $(1 + o(\sigma))$ -in-probability, we obtain a central limit theorem for $i(Q_{d,p})$. Spelling it out with all quantities written explicitly in terms of d and p , we get

$$\frac{i(Q_{d,p}) - 2 \cdot 2^{2^{d-1}} \exp\left(\frac{1}{2}(2-p)^d\right)}{2^{2^{d-1}} \exp\left(\frac{1}{2}(2-p)^d\right) (2 - \frac{3}{2}p)^{d/2}} \xrightarrow{w} \mathcal{N}(0, 1).$$

This recovers the first part of [6, Theorem 1.1], which extends [13, Theorem 1.3] from a regime of nonconstant $p = p_n \rightarrow 1$ as $d \rightarrow \infty$ to all constant $p \in (\frac{2}{3}, 1)$.

We also recover the second part of [6, Theorem 1.1], which covers the case $p = \frac{2}{3}$. In this case $\sigma^2 = \frac{1}{2}$, and (5.1) thus gives

$$\frac{i(Q_{d,2/3})}{2^{2^{d-1}} e^{\mu'}} \stackrel{\mathbb{P}}{\sim} \left(\exp\left(\frac{1}{\sqrt{2}} \frac{\Phi_{\log}^{\text{Even}} - \mu'}{\sigma}\right) + \exp\left(\frac{1}{\sqrt{2}} \frac{\Phi_{\log}^{\text{Odd}} - \mu'}{\sigma}\right) \right).$$

So, after plugging in the value for μ' in terms of d with $p = \frac{2}{3}$, (5.4) gives

$$\frac{i(Q_{d,2/3})}{2^{2^{d-1}} \exp\left(\frac{1}{2}(2-p)^d - \frac{1}{4}\right)} \xrightarrow{w} e^{W^{\text{Even}}/\sqrt{2}} + e^{W^{\text{Odd}}/\sqrt{2}},$$

where $W^{\text{Even}}, W^{\text{Odd}} \sim \mathcal{N}(0, 1)$ are independent.

Finally, for $2 - \sqrt{2} < p < \frac{2}{3}$, (5.1) yields

$$\frac{i(Q_{d,p})}{2^{2^{d-1}} e^{\mu'}} \stackrel{\mathbb{P}}{\sim} \exp\left(\sigma \cdot \frac{\Phi_{\log}^{\text{Even}} - \mu'}{\sigma}\right) + \exp\left(\sigma \cdot \frac{\Phi_{\log}^{\text{Odd}} - \mu'}{\sigma}\right).$$

Since, by (5.4) the quantities inside the exponentials fluctuate at scale $\sigma \rightarrow \infty$, it may easily be checked that this has no scaling limit under any deterministic recentering and rescaling.

5.4. Sampling with replacement. We would like to briefly remark that the phase transition at $p = \frac{2}{3}$ corresponds to the following qualitative behavior of typical independent sets in $Q_{d,p}$, which was already mentioned in [6] as well as in Sections 1 and 2 above.

When $p > \frac{2}{3}$, one way to sample the defect side of an independent set is to first determine the (random) number of singletons one must include, and then sample them *with replacement* from a particular probability distribution which depends on the percolation configuration and weights each singleton v proportionally to φ_v . This works because when $p > \frac{2}{3}$, the typical number of singletons is small enough that there will be

no repeats in the sample with high probability, and in effect we get a sample of a *subset* with the correct distribution. For reasons which we do not delve into presently but which were thoroughly explored in [6], the exponential form of the solution arises from this scheme of sampling with replacement.

However, when $p \leq \frac{2}{3}$, the number of singletons that must be sampled is large enough that when sampling with replacement one begins to see repeated singletons. This situation was analyzed in [6] as a non-uniform random birthday problem, and the outcome of this analysis is that a correction of $-\frac{\sigma^2}{2}$ in μ' is necessary when using the exponential formula arising from sampling with replacement.

In sampling with replacement, there are still no dimers for $p \in (2 - \sqrt{2}, \frac{2}{3})$, meaning that it is very unlikely that two *neighboring vertices* will both be sampled. This leads to the interpretation that the transition happens when certain *individual singletons* gain an outsized probability in the aforementioned random probability distribution. In effect, even when there are no dimers, this transition marks the point where the influence of singletons becomes lopsided and a small number of high-relevance singletons begin to make a meaningful contribution beyond the aggregate behavior.

6. WHEN DIMERS BEGIN TO APPEAR

We now arrive at our main contribution in this paper, which is an analysis of model in the regime of parameters where typical independent sets will contain dimers in their defect sides. This begins to happen when $p \leq 2 - \sqrt{2} \approx 0.586$, and our goal in this section is to prove that the set of singletons and the set of dimers in the defect side of a uniformly random independent set behave roughly independently from each other for a large range of p below this value, extending all the way below $p = \frac{1}{2}$. This will be formalized in Proposition 6.1 below. As discussed in Section 1.2, this allows us to resolve the question posed by [13] about the existence of a phase transition at $p = \frac{1}{2}$, which is a natural point of phase transition for many other features of the percolated hypercube. We reiterate that, interestingly, our results show that there is actually *no* phase transition at $p = \frac{1}{2}$ for the behavior of independent sets in $Q_{d,p}$.


However, our analysis does uncover a *different* point of potential phase transition below $p = \frac{1}{2}$. To explain this, first note that in this section we always restrict to the regime of $p > 2 - 2^{2/3} \approx 0.413$, so that there are no trimers in the defect side (by Proposition 4.1). But even for certain p above this value where more complex clustering behavior arises as *adjacent* singletons and dimers merge to form trimers, it seems that the singletons and dimers will start to interact in a more subtle way. Specifically, if one tried to sample the singleton set and the dimer set independently from one another, we expect that there will be some singleton and some dimer which contain *the same* vertex. This is analogous to the phase transition at $p = \frac{2}{3}$ discussed in Section 5.4 above, where in sampling singletons with replacement, one vertex will be sampled multiple times if $p \leq \frac{2}{3}$, even though there are no *adjacent* vertices sampled until $p \leq 2 - \sqrt{2} \approx 0.586$.

Our proof aims to cover the full regime of parameters p for which the above “overlapping samples” behavior does not occur, and it is valid for all p above a certain threshold which we denote by $\frac{1}{2} - \gamma$ for some $\gamma > 0$. Although we do not prove this, one might guess after observing the structure of our argument that the threshold we obtain is indeed sharp, in the sense that the singleton set and dimer set will not be independent from one another below it. One should also note that while we will simply be satisfied herein with proving

that this threshold is strictly less than $\frac{1}{2}$, one may easily numerically approximate the threshold as about 0.465; in particular, it is strictly larger than $2 - 2^{2/3} \approx 0.413$.

Another interesting point that was mentioned briefly in Remark 2.5 is that a *worst-case* analysis does not suffice for the full regime of independence. To be specific, if we sampled the singleton set S_1 and then the dimer set S_2 independently, then there are certain worst-case choices of S_1 for which the independently sampled S_2 will actually overlap with S_1 , if p is low enough (below roughly 0.548). To overcome this, we must perform an *average-case* analysis (in Sections 6.5 and 6.6), which will allow us to prove that the singleton set and dimer set do not interact for $p > \frac{1}{2} - \gamma$ by averaging over the singleton set S_1 . However, we do also provide a short non-rigorous overview of the worst-case analysis in Section 6.4 as a warm-up to the more technical later sections, as many of the relevant tools can be introduced in this easier setting.

Before stating the main result of this section, Proposition 6.1 below, we summarize the above few paragraphs with the following diagram.



We also emphasize that, although we do not pursue this presently, we expect that the techniques developed in this section will prove fruitful for analyzing the regime of non-independence with values of p down to $2 - 2^{2/3}$ and even beyond, when trimers and other higher-order polymers begin to appear.

6.1. Statement of main proposition. In order to state our main proposition, we remind the reader of our running notation, introduced first in Section 2. As in Section 4 above, to avoid notational overload we drop the superscript $\mathcal{H}$ as every statement in this section relates two quantities which correspond to the *same side*. So, for instance, $\text{PD}_{\leq 2}$ means $\text{PD}_{\leq 2}^{\mathcal{H}}$ here. Recall that

$$\begin{aligned} \mathcal{Z} &= \sum_{S \in \text{PD}} \prod_{\mathbf{p} \in S} \varphi_{\mathbf{p}}, \\ \mathcal{Z}_{\leq 2} &= \sum_{S \in \text{PD}_{\leq 2}} \prod_{\mathbf{p} \in S} \varphi_{\mathbf{p}}, \\ \mathcal{Z}_1 &= \sum_{S \in \text{PD}_1} \prod_{v \in S} \varphi_v, \\ \Delta &= \sum_{\mathbf{d} \in \text{Dimers}} \varphi_{\mathbf{d}}, \end{aligned}$$

as well as the corresponding notions of $\tilde{\mathcal{Z}}$, $\tilde{\mathcal{Z}}_{\leq 2}$, $\tilde{\mathcal{Z}}_1$, and $\tilde{\Delta}$ which are obtained by replacing each $\varphi_{\mathbf{p}}$ by $\tilde{\varphi}_{\mathbf{p}}$, et cetera. We remind the reader that $\tilde{\mathcal{Z}}_1 = \mathcal{Z}_1$ since $\varphi_v = \tilde{\varphi}_v$ for singletons v , but that $\tilde{\mathcal{Z}}_{\leq 2} \neq \mathcal{Z}_{\leq 2}$ and $\Delta \neq \tilde{\Delta}$ since $\tilde{\varphi}_{\mathbf{d}} \neq \varphi_{\mathbf{d}}$ in general for dimers $\mathbf{d}$.

The following proposition is the main result of this section, and holds for both $\mathcal{H} = \text{Even}$ and $\mathcal{H} = \text{Odd}$ (although the $\mathcal{H}$ superscript is only present implicitly in the statement).

Proposition 6.1. *There exists a universal constant $\gamma > 0$ such that for $p > \frac{1}{2} - \gamma$,*

$$\mathcal{Z}_{\leq 2} \stackrel{\mathbb{P}}{\sim} \mathcal{Z}_1 \cdot e^\Delta \quad \text{and} \quad \tilde{\mathcal{Z}}_{\leq 2} \stackrel{\mathbb{P}}{\sim} \mathcal{Z}_1 \cdot e^{\tilde{\Delta}}.$$

Together with Propositions 4.1 and 4.2 which state that for $p > 2 - 2^{2/3}$ we have

$$\mathcal{Z} \stackrel{\mathbb{P}}{\sim} \mathcal{Z}_{\leq 2} \quad \text{and} \quad \tilde{\mathcal{Z}} \stackrel{\mathbb{P}}{\sim} \tilde{\mathcal{Z}}_{\leq 2},$$

Proposition 6.1 implies Corollary 2.6, which states that

$$\mathcal{Z} \stackrel{\mathbb{P}}{\sim} \tilde{\mathcal{Z}} \cdot e^{\Delta - \tilde{\Delta}}$$

for all p larger than some threshold strictly below $\frac{1}{2}$. Combining this with Proposition 3.4, namely that

$$i(Q_{d,p}) \stackrel{\mathbb{P}}{\sim} \mathcal{Z}^{\text{Even}} + \mathcal{Z}^{\text{Odd}},$$

completes our in-probability approximation of $i(Q_{d,p})$, as discussed in Section 2.

Note that we will focus on proving the first of the two statements in Proposition 6.1, with no tildes. As will be expanded upon after the statement of Lemma 6.6 below, our proof of the tilde-free statement will also easily yield a proof of the tilde statement, simply because $\tilde{\varphi}_{\mathbf{p}} \leq \varphi_{\mathbf{p}}$.

6.2. Proof overview and section outline. We will reinterpret the statement of Proposition 6.1 probabilistically as follows. We would like to understand the distribution of $S_{\leq 2} \in \text{PD}_{\leq 2}$ sampled from the model with partition function $\mathcal{Z}_{\leq 2}$, which contains both singletons and dimers, and which is sampled with probability proportional to

$$\prod_{\text{singleton } v \in S_{\leq 2}} \varphi_v \cdot \prod_{\text{dimer } \mathbf{d} \in S_{\leq 2}} \varphi_{\mathbf{d}}.$$

We will compare this to the alternate sampling scheme whereby we sample the singleton part and the dimer part independently. First sample a well-separated singleton set $S_1 \in \text{PD}_1$ from the model with partition function $\mathcal{Z}_1$, i.e. with probability proportional to $\prod_{v \in S_1} \varphi_v$ (but note that we are restricting to only $S_1 \in \text{PD}_1$). Then we sample a dimer set S_2 by simply including each dimer independently with probability $\frac{\varphi_{\mathbf{d}}}{1 + \varphi_{\mathbf{d}}}$, so that the probability of a particular dimer set is proportional to $\prod_{\mathbf{d} \in S_2} \varphi_{\mathbf{d}}$. For the dimer set S_2 , unlike the singleton set S_1 , we have no separation constraint, but the set S_2 we sample will turn out to naturally be well-separated with high probability. Proposition 6.1 amounts to showing that the distribution of $S_{\leq 2}$ is close to that of $S_1 \cup S_2$.

The key difference, a priori, between $S_{\leq 2}$ and $S_1 \cup S_2$ is that we force the singleton and dimer parts of $S_{\leq 2}$ to not interact, meaning that no singletons are adjacent to any dimers in $S_{\leq 2}$. On the other hand, there is no such restriction in $S_1 \cup S_2$ as the two parts are sampled independently. In order to prove Proposition 6.1, we will introduce and analyze an auxiliary random variable associated to the singleton set S_1 , denoted by $\text{Adj}(S_1)$, which controls the probability that the independently sampled dimer set S_2 will interact nontrivially with S_1 .

Definition 6.2. *For a set $S \subseteq \mathcal{H}$ and a dimer $\mathbf{d}$, we write $\mathbf{d} \sim S$ if $\mathbf{d}$ shares a neighbor with a vertex in S . Note that this includes the case where S contains a vertex of $\mathbf{d}$ itself. Now for $S \subseteq \mathcal{H}$, define*

$$\text{Adj}(S) := \sum_{\mathbf{d} \sim S} \varphi_{\mathbf{d}}.$$

For future reference, we also define $\widetilde{\text{Adj}}(S)$ in the same way but with $\varphi_{\mathbf{d}}$ replaced by $\widetilde{\varphi}_{\mathbf{d}}$.

In Section 6.3 below, we will prove a convenient representation for the dimer part of the partition function $\mathcal{Z}_{\leq 2}$. This will motivate the definition of $\text{Adj}(S)$ and it also shows that the dimers do not interact with themselves. As will be explained further in that section, showing that singletons and dimers behave independently amounts to showing that $\text{Adj}(S_1)$ is small, for typical samples of the well-separated singleton set S_1 .

An initial idea for proving this is to show that in fact, for *any* well-separated singleton set S_1 that could be sampled, $\text{Adj}(S_1)$ is uniformly small. This argument is outlined briefly in Section 6.4 below. However, since this *worst-case* strategy does not work for $p < 0.548$, that section is presented somewhat non-rigorously and should be considered as a warm-up for later sections which cover the full range of p in Proposition 6.1.

In Section 6.5, we present the setup for the *average-case* argument, where we prove that $\text{Adj}(S_1)$ is small for *most* well-separated singleton sets $S_1 \in \text{PD}_1$ that are sampled from the model with partition function $\mathcal{Z}_1$. For this, we will introduce the auxiliary (random) measure π on subsets of $\mathcal{H}$, which samples S with probability

$$\pi[S] = \frac{1}{\mathcal{Z}} \prod_{v \in S} \varphi_v, \quad (6.1)$$

or in other words includes each vertex $v \in \mathcal{H}$ independently with probability $\frac{\varphi_v}{1+\varphi_v}$. With this definition in hand, our result will turn out to reduce to the following statement:

$$\pi[\text{Adj}(S) > \varepsilon | S \in \text{PD}_1] \xrightarrow{\mathbb{P}} 0 \quad (6.2)$$

for some choice of $\varepsilon = \varepsilon_d \rightarrow 0$.

To prove (6.2), in Section 6.6 we will decompose the sum defining $\text{Adj}(S)$ into polynomially many pieces, such that each piece consists of variables which are independent under π . This will allow us to apply various standard estimates of binomial probabilities via the relative entropy function, although much care will be needed as all of the binomial random variables in question have *random* specifications in terms of the percolation configuration. As such, Section 6.6 is the most technical part of this entire paper.

6.3. A convenient representation of the partition function. In this section we show that the dimer part of the polymer configuration $S_{\leq 2}$ is quite well-behaved for the regime of parameters we consider. Specifically, it behaves similarly to sampling dimers independently, and we may rewrite the dimer part of the partition function $\mathcal{Z}_{\leq 2}$ in a convenient exponential form involving the quantity $\text{Adj}(S)$ introduced in Definition 6.2.

This should be compared to the exponential form obtained in our previous work [6] for the singleton partition function when $p > \frac{2}{3}$ (which was re-proven in Section 5.1 above). Our result here for dimers has an interpretation analogous to the one discussed in Section 5.4, where one can consider instead sampling the dimers *with replacement* after choosing the correct number of dimers to sample, and no individual dimer will have a weight large enough to be sampled twice.

Let us begin by recording a definition.

Definition 6.3. For any $S_1 \in \text{PD}_1$, we denote by $\text{Compat}_2(S_1)$ the collection of dimers $\mathbf{d}$ with $\mathbf{d} \not\sim S_1$, meaning that $S_1 \cup \{\mathbf{d}\}$ remains a valid polymer configuration.

With this notation, we have the following representation for the singleton-and-dimer partition function.

$$\begin{aligned} \mathcal{Z}_{\leq 2} &= \sum_{S_{\leq 2} \in \text{PD}_{\leq 2}} \prod_{\mathbf{p} \in S_{\leq 2}} \varphi_{\mathbf{p}} \\ &= \sum_{S_1 \in \text{PD}_1} \left(\prod_{v \in S_1} \varphi_v \cdot \sum_{\substack{S_2 \in \text{PD}_2 \\ S_2 \subseteq \text{Compat}_2(S_1)}} \prod_{\mathbf{d} \in S_2} \varphi_{\mathbf{d}} \right). \end{aligned} \quad (6.3)$$

We will first show that the constraint $S_2 \in \text{PD}_2$ is superfluous, i.e. that

$$\sum_{\substack{S_2 \in \text{PD}_2 \\ S_2 \subseteq \text{Compat}_2(S_1)}} \prod_{\mathbf{d} \in S_2} \varphi_{\mathbf{d}} \stackrel{\mathbb{P}}{\sim} \sum_{S_2 \subseteq \text{Compat}_2(S_1)} \prod_{\mathbf{d} \in S_2} \varphi_{\mathbf{d}}, \quad (6.4)$$

with the approximation being uniform in $S_1 \in \text{PD}_1$. Since

$$\sum_{S_2 \subseteq \text{Compat}_2(S_1)} \prod_{\mathbf{d} \in S_2} \varphi_{\mathbf{d}} = \prod_{\mathbf{d} \not\sim S_1} (1 + \varphi_{\mathbf{d}}), \quad (6.5)$$

we can interpret the approximation (6.4) as saying that, in order to sample a collection S_2 of dimers compatible with any singleton set $S_1 \in \text{PD}_1$, with the correct distribution as given by (6.3), we may independently include each dimer $\mathbf{d} \in \text{Compat}_2(S_1)$ in the set S_2 with probability $\frac{\varphi_{\mathbf{d}}}{1 + \varphi_{\mathbf{d}}}$. Thus, this is a statement about non-collision (with one another) of independently drawn dimers.

We will also apply the approximation $1 + \varphi_{\mathbf{d}} \approx e^{\varphi_{\mathbf{d}}}$ to the right-hand side of (6.5), and the smallness of $\varphi_{\mathbf{d}}$ for a certain range of p which suffices for our purposes will ensure that the error does not accumulate as the product over all dimers $\mathbf{d} \not\sim S_1$ is taken. We combine this approximation with (6.4) in the following lemma.

Lemma 6.4. *For $p > 0.391$, we have*

$$\sum_{\substack{S_2 \in \text{PD}_2 \\ S_2 \subseteq \text{Compat}_2(S_1)}} \prod_{\mathbf{d} \in S_2} \varphi_{\mathbf{d}} \stackrel{\mathbb{P}}{\sim} \exp \left(\sum_{\mathbf{d} \not\sim S_1} \varphi_{\mathbf{d}} \right) = \exp(\Delta - \text{Adj}(S_1)),$$

where the approximation is uniform among all $S_1 \subseteq \mathcal{H}$ (recall what this means from above the statement of Theorem 1.1). The same statement also holds with $\varphi_{\mathbf{d}}$ replaced by $\tilde{\varphi}_{\mathbf{d}}$, Δ replaced by $\tilde{\Delta}$, and Adj replaced by $\tilde{\text{Adj}}$.

We remark that the threshold $p > 0.391$ arises from a simple moment calculation and is not related to the threshold $\frac{1}{2} - \gamma$ of our main result which arises from other considerations. Additionally, the equality in the statement follows directly from the definition of Δ in (2.11) and of Adj in Definition 6.2, so it only remains to prove the uniform in-probability approximation.

Proof of Lemma 6.4. We first derive an upper bound on the difference between the two sides of (6.4). Note that if $S_2 \subseteq \text{Compat}_2(S_1)$ but $S_2 \notin \text{PD}_2$, then there are two dimers in S_2 which are adjacent or overlapping,

i.e. there are distinct $\mathbf{d}_1, \mathbf{d}_2 \in S_2$ with $\mathbf{d}_1 \sim \mathbf{d}_2$. This implies that

$$0 \leq \sum_{\substack{S_2 \not\subseteq \text{PD}_2 \\ S_2 \subseteq \text{Compat}_2(S_1)}} \prod_{\mathbf{d} \in S_2} \varphi_{\mathbf{d}} \leq \left(\sum_{\substack{\mathbf{d}_1 \neq \mathbf{d}_2 \\ \mathbf{d}_1 \sim \mathbf{d}_2}} \varphi_{\mathbf{d}_1} \varphi_{\mathbf{d}_2} \right) \cdot \exp \left(\sum_{\mathbf{d} \not\sim S_1} \varphi_{\mathbf{d}} \right).$$

To see this, note that in the Taylor expansion of the exponential, every collection of distinct dimers $\mathbf{d}$ will lead to a term which is the product of the relevant $\varphi_{\mathbf{d}}$, with a coefficient of 1 (this is similar to the strategy used in the proof of Proposition 4.2 above). Note that the prefactor above does not depend on S_1 .

Now consider the right-hand side of (6.4) which is equal to the right-hand side of (6.5). Since

$$e^{x-x^2/2} \leq 1+x \leq e^x$$

for all $x \geq 0$, we have

$$\exp \left(\sum_{\mathbf{d} \not\sim S_1} \varphi_{\mathbf{d}} - \frac{1}{2} \sum_{\mathbf{d} \not\sim S_1} \varphi_{\mathbf{d}}^2 \right) \leq \sum_{S_2 \subseteq \text{Compat}_2(S_1)} \prod_{\mathbf{d} \in S_2} \varphi_{\mathbf{d}} \leq \exp \left(\sum_{\mathbf{d} \not\sim S_1} \varphi_{\mathbf{d}} \right).$$

Moreover, since every $\varphi_{\mathbf{d}}$ is positive, we can drop the restriction that $\mathbf{d} \not\sim S_1$ in the second sum in the first exponential above, obtaining

$$\exp \left(-\frac{1}{2} \sum_{\mathbf{d} \in \text{Dimers}} \varphi_{\mathbf{d}}^2 \right) \cdot \exp \left(\sum_{\mathbf{d} \not\sim S_1} \varphi_{\mathbf{d}} \right) \leq \sum_{S_2 \subseteq \text{Compat}_2(S_1)} \prod_{\mathbf{d} \in S_2} \varphi_{\mathbf{d}} \leq \exp \left(\sum_{\mathbf{d} \not\sim S_1} \varphi_{\mathbf{d}} \right),$$

wherein again the pre-factor in the left-hand side does not depend on S . Thus the non-tilde case of the lemma will be proved if we can show that both

$$\sum_{\substack{\mathbf{d}_1 \neq \mathbf{d}_2 \\ \mathbf{d}_1 \sim \mathbf{d}_2}} \varphi_{\mathbf{d}_1} \varphi_{\mathbf{d}_2} \xrightarrow{\mathbb{P}} 0 \quad \text{and} \quad \sum_{\mathbf{d} \in \text{Dimers}} \varphi_{\mathbf{d}}^2 \xrightarrow{\mathbb{P}} 0,$$

which, by positivity of $\varphi_{\mathbf{d}}$ and the fact that $\mathbf{d} \sim \mathbf{d}$ for any $\mathbf{d} \in \text{Dimers}$, is implied by

$$\sum_{\mathbf{d}_1 \sim \mathbf{d}_2} \varphi_{\mathbf{d}_1} \varphi_{\mathbf{d}_2} \xrightarrow{\mathbb{P}} 0. \tag{6.6}$$

This follows from a simple moment calculation which is relegated to the appendix. Namely, Lemma A.4 (b) states that the expectation of the above quantity converges to 0 whenever $p > 0.391$. Moreover, since the same algebra goes through with $\varphi_{\mathbf{d}}$ replaced by $\tilde{\varphi}_{\mathbf{d}}$, and since $\tilde{\varphi}_{\mathbf{d}} \leq \varphi_{\mathbf{d}}$, (6.6) also implies the tilde case of the lemma. $\square$

With Lemma 6.4 in hand, we have effectively relegated the interaction between the singleton part S_1 and the dimer part S_2 to the study of the quantity $\text{Adj}(S_1)$, and as such

for the rest of this section, we will use S instead of S_1 to denote the singleton set,

to reduce notational clutter.

For instance, recalling the representation (6.3) for $\mathcal{Z}_{\leq 2}$, Lemma 6.4 implies that

$$\mathcal{Z}_{\leq 2} \stackrel{\mathbb{P}}{\sim} \sum_{S \in \text{PD}_1} \left(\prod_{v \in S} \varphi_v \cdot \exp(\Delta - \text{Adj}(S)) \right),$$

crucially using the fact that the approximation in Lemma 6.4 is uniform over all singleton sets. Recalling the definition of π from (6.1), we obtain

$$\mathcal{Z}_{\leq 2} \stackrel{\mathbb{P}}{\sim} e^\Delta \cdot \tilde{\mathcal{Z}} \cdot \mathbb{E}_\pi \left[\mathbf{1}_{\{S \in \text{PD}_1\}} \cdot e^{-\text{Adj}(S)} \right].$$

Now recalling the definitions of $\tilde{\mathcal{Z}}$ and $\mathcal{Z}_1$ from (2.7) and (2.9), we obtain

$$\mathcal{Z}_{\leq 2} \stackrel{\mathbb{P}}{\sim} e^\Delta \cdot \mathcal{Z}_1 \cdot \mathbb{E}_\pi \left[e^{-\text{Adj}(S)} \middle| S \in \text{PD}_1 \right]. \quad (6.7)$$

So, to prove the first approximation in Proposition 6.1, which states that $\mathcal{Z}_{\leq 2} \stackrel{\mathbb{P}}{\sim} \mathcal{Z}_1 \cdot e^\Delta$, it is equivalent to show that

$$\mathbb{E}_\pi \left[e^{-\text{Adj}(S)} \middle| S \in \text{PD}_1 \right] \xrightarrow{\mathbb{P}} 1. \quad (6.8)$$

First of all, $\text{Adj}(S) \geq 0$, and so the above expectation is always at most 1. To show that it converges to 1, we need to prove some form of a corresponding upper bound on $\text{Adj}(S)$. This task will occupy us for the remainder of this section of the paper.

Note also that all of the above manipulations can be done with φ_d replaced by $\tilde{\varphi}_d$ everywhere, and we will obtain

$$\tilde{\mathcal{Z}}_{\leq 2} \stackrel{\mathbb{P}}{\sim} e^{\tilde{\Delta}} \cdot \mathcal{Z}_1 \cdot \mathbb{E}_\pi \left[e^{-\tilde{\text{Adj}}(S)} \middle| S \in \text{PD}_1 \right].$$

Once we have a workable upper bound for $\text{Adj}(S)$, the same bound will also follow for $\tilde{\text{Adj}}(S)$, since $\tilde{\varphi}_d \leq \varphi_d$ deterministically. So the tilde case of Proposition 6.1 follows from the same analysis as the non-tilde case, which we focus on henceforth.

6.4. A worst-case argument to bound the contributions of adjacent dimers. In this section we present a sketch of a “worst-case” type argument to prove (6.8) by showing that $\text{Adj}(S)$ is small for all sets $S \subseteq \mathcal{H}$, assuming a simple *a priori* bound on the size of S . We will see what this size constraint should be shortly, and then find the worst-case S , which maximizes $\text{Adj}(S)$ while satisfying this constraint. If this worst-case S has $\text{Adj}(S) \xrightarrow{\mathbb{P}} 0$, then the proof of (6.8) will be finished.

Unfortunately, this strategy does not work for the full range of p which we wish to consider, and in later sections we will present an “average-case” argument which proves (6.8) in what we believe to be the full regime of parameters for which it holds. Nevertheless, some ideas used later will already surface in this easier worst-case analysis, so this somewhat non-rigorous section should be viewed as a warm-up.

6.4.1. Size constraint. First we remark that, as can be proven via simple moment calculations such as those done in Lemma A.1, the typical size of a set S sampled from π is approximately $\frac{1}{2}(2-p)^d$. Since vertices are included in S independently from one another, it is easy to show (via a Chernoff bound for instance) that

$$\pi \left[|S| > C(2-p)^d \right] \leq e^{-c(2-p)^d} (1 + o(1)), \quad (6.9)$$

for some constants $C, c > 0$, where the $o(1)$ term converges to 0 in probability.

On the other hand, by the definitions of $\tilde{\mathcal{Z}}$ from (2.7) and $\mathcal{Z}_1$ from (2.9), as well as recalling the definition of π from (6.1), we have

$$\pi[S \in \text{PD}_1] = \frac{\mathcal{Z}_1}{\tilde{\mathcal{Z}}} \stackrel{\mathbb{P}}{\sim} \frac{\mathcal{Z}_1}{\tilde{\mathcal{Z}}_{\leq 2}},$$

using Proposition 4.2 for the approximation of $\tilde{\mathcal{Z}}$ by $\tilde{\mathcal{Z}}_{\leq 2}$. Now, by similar manipulations as just done in the previous section to derive (6.7), we may obtain

$$\begin{aligned} \tilde{\mathcal{Z}}_{\leq 2} &\stackrel{\mathbb{P}}{\sim} e^{\tilde{\Delta}} \cdot \mathcal{Z}_1 \cdot \mathbb{E}_\pi \left[e^{-\widetilde{\text{Adj}}(S)} \middle| S \in \text{PD}_1 \right] \\ &\leq e^{\tilde{\Delta}} \cdot \mathcal{Z}_1, \end{aligned}$$

using simply the fact that $\widetilde{\text{Adj}}(S) \geq 0$ in the last inequality. Thus we obtain

$$\pi[S \in \text{PD}_1] \geq e^{-\tilde{\Delta}}(1 - o(1)), \quad (6.10)$$

where again the $o(1)$ term converges to 0 in probability.

Now, combining (6.9) and (6.10), we find that

$$\pi \left[|S| > C(2-p)^d \middle| S \in \text{PD}_1 \right] \leq \exp \left(-c(2-p)^d + \tilde{\Delta} \right) \cdot (1 + o(1)), \quad (6.11)$$

where again $o(1) \xrightarrow{\mathbb{P}} 0$. Now, moment calculations which can be found in Lemma A.4 show that $\tilde{\Delta}$ concentrates at constant scale around its mean, which is

$$\tilde{\mu}_2 = \frac{d(d-1)}{4} \left(\frac{(2-p)^2}{2} \right)^d$$

as defined first in (1.10). Since $2-p < 2$ we have $(2-p)^d \gg \left(\frac{(2-p)^2}{2} \right)^d$, which means that the bound (6.11) tends to 0. Thus, to prove (6.8), it suffices to show that

$$\max \left\{ \text{Adj}(S) : |S| \leq C(2-p)^d \right\} \xrightarrow{\mathbb{P}} 0,$$

where C is as in (6.9).

6.4.2. Worst-case sum. Given that $|S| \leq C(2-p)^d$, since each vertex is only adjacent to at most d^4 distinct dimers, the number of dimers $\mathbf{d} \sim S$ is at most $\text{poly}(d) \cdot (2-p)^d$. Thus, to bound

$$\text{Adj}(S) = \sum_{\mathbf{d} \sim S} \varphi_{\mathbf{d}}$$

for S with this size constraint, it suffices to bound the sum of the *largest* $\text{poly}(d) \cdot (2-p)^d$ variables $\varphi_{\mathbf{d}}$. To do this, we need to get a better understanding of how many dimers $\mathbf{d}$ have each possible value of $\varphi_{\mathbf{d}}$.

Since we are just looking for an upper bound, let us first use the fact that if $\mathbf{d} = \{u, v\}$ then

$$\varphi_{\mathbf{d}} \leq 4 \cdot 2^{-N_p(u) - N_p(v)},$$

because there are exactly two vertices which neighbor both u and v in the unpercolated hypercube Q_d , so $N_p(\{u, v\}) \geq N_p(u) + N_p(v) - 2$. Now we will use the fact that $N_p(u)$ are independent $\text{Bin}(d, p)$ variables (with respect to the randomness of the percolation $Q_{d,p}$) to bound the number of dimers $\mathbf{d} = \{u, v\}$ with various (low) values of $N_p(u) + N_p(v)$.

Using standard entropy estimates on binomial probabilities, many of which are included in Appendix B, we have

$$\mathbb{P}[N_p(u) = a] = \text{poly}(d) \cdot 2^{-dH_p(r)},$$

where $r = \frac{a}{d}$ and H_p is the binary relative entropy function

$$H_p(q) := q \log_2 \frac{q}{p} + (1 - q) \log_2 \frac{1 - q}{1 - p}. \quad (6.12)$$

Now since there are $\text{poly}(d) \cdot 2^d$ dimers in total and $N_p(u)$ are independent for $u \in \mathcal{H}$, the number of dimers $\mathbf{d} = \{u, v\}$ which have $N_p(u) = a$ and $N_p(v) = b$ is thus within a (fixed) polynomial factor of

$$2^{d(1-H_p(\frac{a}{d})-H_p(\frac{b}{d}))}$$

with high probability. Now by convexity of relative entropy, we have

$$1 - H_p\left(\frac{a}{d}\right) - H_p\left(\frac{b}{d}\right) \leq 1 - 2 \cdot H_p\left(\frac{a+b}{2d}\right),$$

and so the number of dimers $\mathbf{d} = \{u, v\}$ for which $N_p(u) + N_p(v) = c$ is, with high probability, at most

$$\text{poly}(d) \sum_{a+b=c} 2^{d(1-H_p(\frac{a}{d})-H_p(\frac{b}{d}))} \leq \text{poly}(d) \cdot 2^{d(1-2 \cdot H_p(\frac{c}{2d}))}. \quad (6.13)$$

In the above bound we have used the fact that the sum in the left-hand side of (6.13) has only polynomially many terms.

For notational convenience, let us define

$$h_c^* := 1 - 2 \cdot H_p\left(\frac{c}{2d}\right).$$

With this definition, we see that the sum of the $\text{poly}(d) \cdot (2 - p)^d$ largest values of $\varphi_{\mathbf{d}}$ is at most

$$\text{poly}(d) \sum_{c=0}^{c_{\max}} 2^{d \cdot h_c^*} \cdot 2^{-c}, \quad (6.14)$$

with high probability, where $c_{\max}$ is the maximal value of c for which $2^{d \cdot h_c^*} \leq \text{poly}(d) \cdot (2 - p)^d$, which ensures that the number of dimers $\mathbf{d}$ satisfying $\tilde{\varphi}_{\mathbf{d}} = 2^{-c}$ is at most $\text{poly}(d) \cdot (2 - p)^d$.

Finally, since there are only polynomially many terms in the sum (6.14), we will have $\text{Adj}(S) \xrightarrow{\mathbb{P}} 0$ uniformly for $|S| \leq C(2 - p)^d$ if we can show that

$$\text{poly}(d) \cdot 2^{dh_c^* - c} \ll 1$$

whenever c is small enough that $2^{dh_c^*} \leq \text{poly}(d) \cdot (2 - p)^d$. In other words, setting $t = \frac{c}{2d}$, we need to have

$$1 - 2 \cdot H_p(t) - 2t < 0 \quad \text{for all } t \in [0, 1] \quad \text{satisfying} \quad 1 - 2 \cdot H_p(t) < \log_2(2 - p).$$

This can be numerically checked, and it holds for $p > 0.549$, but does not hold for $p < 0.548$. Thus, in order to achieve our goal of investigating the behavior near $p = \frac{1}{2}$, we will need a more sophisticated argument which leverages the fact that (6.8) is an average and so we can ignore worst-case sets S if they are rare.

6.5. The average-case argument. In this section we lay the groundwork for the average-case argument to prove (6.8), i.e. that

$$\mathbb{E}_\pi \left[e^{-\text{Adj}(S)} \mid S \in \text{PD}_1 \right] \xrightarrow{\mathbb{P}} 1, \quad (6.15)$$

for all $p > \frac{1}{2} - \gamma$, where $\text{Adj}(S)$ is as in Definition 6.2 and π is defined in (6.1). This will finish the proof of Proposition 6.1, as discussed in Section 6.3. As mentioned at the beginning of Section 6, we expect that the strategy we employ in this section and the next actually works for all p such that the result does indeed hold, although we do not prove this. The heuristic for this belief was outlined at the beginning of Section 6, and is related to a phase transition similar in spirit to the one which occurs at $p = \frac{2}{3}$, which was discussed in Section 5.4.

Note that since $\text{Adj}(S) \geq 0$, to prove (6.15) it suffices to show that

$$\pi[\text{Adj}(S) > \varepsilon \mid S \in \text{PD}_1] \xrightarrow{\mathbb{P}} 0 \quad (6.16)$$

for some choice of $\varepsilon = \varepsilon_d \rightarrow 0$ as $d \rightarrow \infty$. In other words, instead of attempting to fully rule out sets with large $\text{Adj}(S)$ as in the worst-case argument of the previous section, we will show that such sets are rare under π conditioned on $S \in \text{PD}_1$. In fact, we will show a stronger statement, that

$$\pi[\text{Adj}(S) > \varepsilon] \ll \pi[S \in \text{PD}_1] \cdot (1 + o(1)) \quad (6.17)$$

for some choice of $\varepsilon = \varepsilon_d$ decreasing to zero polynomially as a function of d , where the $o(1)$ term converges to 0 in probability. This will be achieved by the following two lemmas, which give a lower bound on the right-hand side and an upper bound on the left-hand side.

Lemma 6.5. *For all $p \in (0, 1)$,*

$$\pi[S \in \text{PD}_1] \geq \exp \left(-\text{poly}(d) \cdot \left(\frac{(2-p)^2}{2} \right)^d \right) \cdot (1 - o(1)),$$

where the $o(1)$ term tends to 0 in probability.

Lemma 6.6. *There exists a universal constant $\gamma > 0$ such that for $p > \frac{1}{2} - \gamma$, there is some constant $\alpha_p > \frac{(2-p)^2}{2}$ such that for any $\varepsilon \geq \frac{1}{\text{poly}(d)}$,*

$$\pi[\text{Adj}(S) > \varepsilon] \leq \exp \left(-\frac{1}{\text{poly}(d)} \cdot \alpha_p^d \right) \cdot (1 + o(1)),$$

where the $o(1)$ term tends to 0 in probability.

These two lemmas combine to prove (6.17) which proves (6.15) and subsequently, by the discussion in Section 6.3, finishes the proof of Proposition 6.1 in the first (non-tilde) case. As mentioned in Section 6.1, the tilde case follows from the non-tilde case simply because $\widetilde{\varphi}_d \leq \varphi_d$, and so $\widetilde{\text{Adj}}(S) \leq \text{Adj}(S)$, meaning that the bound (6.17) with $\text{Adj}(S)$ replaced by $\widetilde{\text{Adj}}(S)$ follows immediately from the original bound. Thus, Lemmas 6.5 and 6.6 complete the full proof of Proposition 6.1.

We now prove Lemma 6.5, with the proof of Lemma 6.6, which is one of the key novelties of this article, deferred to the next subsection.

Proof of Lemma 6.5. First note that

$$\begin{aligned}
\tilde{\mathcal{Z}} &\stackrel{\mathbb{P}}{\sim} \tilde{\mathcal{Z}}_{\leq 2} && \text{(by Proposition 4.2)} \\
&= \sum_{S \in \text{PD}_1} \left(\prod_{v \in S} \varphi_v \cdot \sum_{\substack{S_2 \in \text{PD}_2 \\ S_2 \subseteq \text{Compat}_2(S)}} \prod_{\mathbf{d} \in S_2} \tilde{\varphi}_{\mathbf{d}} \right) && (6.18) \\
&\stackrel{\mathbb{P}}{\sim} \sum_{S \in \text{PD}_1} \left(\prod_{v \in S} \varphi_v \cdot \exp(\tilde{\Delta} - \widetilde{\text{Adj}}(S)) \right) && \text{(by Lemma 6.4)} \\
&\leq e^{\tilde{\Delta}} \cdot \sum_{S \in \text{PD}_1} \prod_{v \in S} \varphi_v. && \text{(since } \widetilde{\text{Adj}}(S) \geq 0)
\end{aligned}$$

Note that at step (6.18), we used the manipulation done in (6.3) but with $\varphi_{\mathbf{d}}$ replaced by $\tilde{\varphi}_{\mathbf{d}}$. Additionally, in the application of Lemma 6.4 above, it was crucial that the approximation in that lemma is uniform over S .

Now, recalling the definition of π from (6.1), rearranging the above display shows that

$$\begin{aligned}
\pi[S \in \text{PD}_1] &= \frac{1}{\tilde{\mathcal{Z}}} \sum_{S \in \text{PD}_1} \prod_{v \in S} \varphi_v \\
&\geq \exp(-\tilde{\Delta}) \cdot (1 - o(1)),
\end{aligned}$$

where $o(1) \xrightarrow{\mathbb{P}} 0$. Now, as it turns out, $\tilde{\Delta}$ concentrates at around its mean

$$\tilde{\mu}_2 = \frac{d(d-1)}{4} \left(\frac{(2-p)^2}{2} \right)^d,$$

this notation being initially defined in (1.10). This can be seen by moment calculations in the appendix, specifically in Lemma A.4. The concentration result is also written explicitly as Lemma 7.5 below, which states that

$$\tilde{\Delta} - \tilde{\mu}_2 \xrightarrow{\mathbb{P}} 0.$$

This finishes the proof. □

6.6. Bounding the contribution of singleton sets with large dimer adjacency. In this section we prove Lemma 6.6. For convenience, we will replace the distribution π with one which chooses each singleton v independently with probability φ_v instead of the unwieldy ratio $\frac{\varphi_v}{1+\varphi_v}$ (note that this makes sense since $\varphi_v \leq 1$). This new measure stochastically dominates the old one (via inclusion of sets) since $\frac{\varphi_v}{1+\varphi_v} \leq \varphi_v$. Since $\text{Adj}(S)$ is an increasing function of S (recall from Definition 6.2 that it is the sum of $\varphi_{\mathbf{d}}$ over all $\mathbf{d} \sim S$), if we can show that the event $\{\text{Adj}(S) > \varepsilon\}$ has small probability under this new measure, then it will also have small probability under the old measure.

Henceforth we denote by π this new measure, which should not cause confusion

as we will not return to the old definition of π for the rest of this section.

Now notice that, by rewriting the definition of $\text{Adj}(S)$ from Definition 6.2, we have

$$\text{Adj}(S) = \sum_{\mathbf{d} \in \text{Dimers}} \varphi_{\mathbf{d}} \mathbf{1}_{\{\mathbf{d} \sim S\}},$$

and so broadly speaking our goal is to bound the tail of a sum of bounded random variables. Note that if the variables were i.i.d., one may have applied a version of Hoeffding's inequality or a standard large deviations bound for sums of i.i.d. variables to get a bound on this tail. However, the variables $\{\varphi_{\mathbf{d}} \mathbf{1}_{\{\mathbf{d} \sim S\}}\}_{\mathbf{d} \in \text{Dimers}}$ are not independent under π since many different dimers $\mathbf{d}$ can be adjacent to one single vertex v , meaning that many different indicators $\mathbf{1}_{\{\mathbf{d} \sim S\}}$ may be set to 1 as a result of a single vertex being added to S . Furthermore they are not identically distributed since the distribution of $\varphi_{\mathbf{d}} \mathbf{1}_{\{\mathbf{d} \sim S\}}$ under π (which samples vertices v to include in S based on their degree in $Q_{d,p}$) depends on the percolation configuration near $\mathbf{d}$.

Nonetheless, the dependence is local in nature and if the dimers are separated enough from each other, then indeed they correspond to independent variables $\varphi_{\mathbf{d}} \mathbf{1}_{\{\mathbf{d} \sim S\}}$. Specifically, the variables $\mathbf{d}_1$ and $\mathbf{d}_2$ are independent under π when there is no vertex v whose inclusion in S can cause both $\mathbf{d}_1 \sim S$ and $\mathbf{d}_2 \sim S$, since the vertices are included in S independently. This is the case when the dimers have distance at least 3 in the shared-neighbor graph, meaning that there is no vertex v which has a shared neighbor with $\mathbf{d}_1$ and a shared neighbor with $\mathbf{d}_2$, in the unpercolated hypercube Q_d .

Hence to remedy the non-independence problem, we will partition the set of dimers into $\text{poly}(d)$ layers $\{\mathcal{L}_\ell\}_\ell$ such that for each ℓ , the dimers in $\mathcal{L}_\ell$ are at shared-neighbor-distance at least 3 from each other. One particularly quick way to see that this is possible is to recall that the chromatic number of a graph is at most the maximum degree plus one, and each dimer is within distance 3 of at most $\text{poly}(d)$ other dimers. Note that the classes $\mathcal{L}_\ell$ are deterministic, i.e. they do not depend on the percolation configuration.

It would be useful to have the following definition.

Definition 6.7. *The layer of a dimer $\mathbf{d}$ is the unique integer ℓ such that $\mathbf{d} \in \mathcal{L}_\ell$.*

We will use the independence between dimers within each layer to get a good upper bound on the tail for the sum over each layer, and then use a union bound over the only $\text{poly}(d)$ many layers. However, different dimers in the same layer may still have different distributions with different typical magnitudes, and it is a priori unclear which types of dimers yield the dominant contribution to the sum. This necessitates another large deviation analysis. Towards this, we further decompose each layer in terms of the value of $\varphi_{\mathbf{d}}$, or, more precisely, by the number of neighbors in $Q_{d,p}$ each of the constituent vertices have. To make this precise, we fix an arbitrary ordering $\prec$ of the vertices and equate a dimer $\mathbf{d}$ with its pair (u, v) of vertices, where $u \prec v$.

Definition 6.8. *The type of a dimer $\mathbf{d} = (u, v)$ is the ordered pair (a, b) where $N_p(u) = a$ and $N_p(v) = b$. We will also say that a singleton v has type a if $N_p(v) = a$.*

Note that there are only $(d+1)^2$ different possible types for a dimer. For a dimer $\mathbf{d} = (u, v)$ of type (a, b) , we have

$$\varphi_{\mathbf{d}} = 2^{-N_p(\{u, v\})} \leq 2^{-N_p(u) - N_p(v) + 2} = 4 \cdot 2^{-a-b}, \quad (6.19)$$

because two vertices in the same side sharing a neighbor share exactly two neighbors in Q_d , before percolation (recall for instance the diagram (1.6)).

Finally, due to the differences in the behaviors of dimers *adjacent* to a vertex in S and dimers *containing* a vertex in S , we separate the dimers by this distinction as well.

Definition 6.9. *The status of a dimer $\mathbf{d} = (u, v)$ with respect to S is an element $\sigma \in \{O_1, O_2, N\}$ with $\sigma = O_1$ if $u \in S$, $\sigma = O_2$ if $u \notin S, v \in S$ (O stands for “overlap”), and $\sigma = N$ if $u, v \notin S$ but one of them is adjacent to a singleton in S (N stands for “neighbor”).*

We remark that *dimers with the same type and status may still not have the same distribution* under π due to the different possible values of $N_p(\{u, v\})$ arising from the different ways in which the neighborhoods of u and v may overlap. However, the relationship (6.19) will be enough for our purposes. Broadly speaking, the rest of the proof is to bound the tail of each sum of variables corresponding to dimers of the same layer, type, and status; we will then conclude with a union bound.

To start making this precise, let us denote by $\mathcal{C}(a, b; \ell; \sigma; S)$ the number of dimers of type (a, b) , layer ℓ , and status σ with respect to S . So, by (6.19), we have

$$\text{Adj}(S) \lesssim \sum_{a,b,\ell,\sigma} \mathcal{C}(a, b; \ell; \sigma; S) \cdot 2^{-a-b}.$$

There are at most $\text{poly}(d)$ terms in this sum, so $\text{Adj}(S) > \varepsilon$ implies that for some choice of a, b, ℓ, σ we have

$$\mathcal{C}(a, b; \ell; \sigma; S) \cdot 2^{-a-b} > \frac{\varepsilon}{\text{poly}(d)}, \quad (6.20)$$

which is still at least $\frac{1}{\text{poly}(d)}$ if $\varepsilon > \frac{1}{\text{poly}(d)}$. So Lemma 6.6 will follow from the following lemma, which states that, with high probability in terms of the percolation configuration, (6.20) has low probability when S is sampled from π .

Lemma 6.10. *There exists a universal constant $\gamma > 0$ such that for $p > \frac{1}{2} - \gamma$, there is a constant $\alpha_p > \frac{(2-p)^2}{2}$ such that, for all $\varepsilon = \varepsilon_d > \text{poly}(d)^{-1}$, there is a function $f(d) > \text{poly}(d)^{-1}$ (which may depend on ε_d) such that the following event (determined by the percolation configuration) holds with probability tending to 1.*

$$\left\{ \text{For all choices of } a, b, \ell, \sigma, \text{ we have } \pi \left[\mathcal{C}(a, b; \ell; \sigma; S) > \varepsilon \cdot 2^{a+b} \right] \leq \exp \left(-f(d) \cdot \alpha_p^d \right) \right\}.$$

Proof of Lemma 6.10. The cases $\sigma \in \{O_1, O_2\}$ and $\sigma = N$ are similar but with some minor differences. We tackle the former case first.

Case 1: $\sigma = O_1$. We only show the O_1 case as the O_2 case is identical.

For each a, b, ℓ , let $\mathcal{N}(a, b; \ell)$ denote the number of dimers of type (a, b) in layer ℓ ; note that this only depends on the percolation configuration and so is deterministic with respect to the measure π . So when S is sampled from π the count $\mathcal{C}(a, b; \ell; O_1; S)$ is a binomial random variable with parameters $\mathcal{N}(a, b; \ell)$ and 2^{-a} , recalling that we are using a simplified version of π which chooses a singleton of type a with probability 2^{-a} , and also that dimers in a particular layer are adjacent to S independently due to their separation. This implies that

$$\pi \left[\mathcal{C}(a, b; \ell; O_1; S) > \varepsilon \cdot 2^{a+b} \right] = \mathbb{P} \left[\text{Bin}(\mathcal{N}(a, b; \ell), 2^{-a}) > \varepsilon \cdot 2^{a+b} \right]. \quad (6.21)$$

In order to estimate the right-hand side we will make use of various standard estimates for binomial probabilities, which are stated and proved in the appendix.

First we will estimate $\mathcal{N}(a, b; \ell)$. Due to the separation imposed by the layer constraint, when considered as a random variable which depends on the random percolation configuration, $\mathcal{N}(a, b; \ell)$ itself is a sum of independent indicators, one for each $\mathbf{d} \in \mathfrak{L}_\ell$, of the event that $\mathbf{d}$ has type (a, b) . In other words, $\mathcal{N}(a, b; \ell)$ is itself a binomial random variable, measurable with respect to the percolation configuration. A fixed dimer $\mathbf{d} = (u, v)$ in layer ℓ is included in $\mathcal{N}(a, b; \ell)$ with probability

$$\mathbb{P}[N_p(u) = a, N_p(v) = b] \leq 2^{-d(H_p(\frac{a}{d}) + H_p(\frac{b}{d}))} =: p_{a,b}, \quad (6.22)$$

via the standard Lemma B.1, which bounds binomial probabilities using the entropy function H_p defined in (6.12). Thus $\mathcal{N}(a, b; \ell)$ is stochastically dominated by a $\text{Bin}(2^d d^2, p_{a,b})$ random variable. Therefore, for each polynomial $f(d)$ and for each choice of a, b, ℓ , we have

$$\mathbb{P}[\mathcal{N}(a, b; \ell) \geq f(d) \cdot 2^d d^2 \cdot p_{a,b}] \leq \frac{1}{f(d)}$$

by Markov's inequality. Since the number of labels (a, b, ℓ) is at most $\text{poly}(d)$, by a union bound, choosing the correct $f(d)$, there is a polynomial $g(d)$ such that with probability $\geq 1 - \frac{1}{d}$,

$$\mathcal{N}(a, b; \ell) \leq 2^d \cdot g(d) \cdot p_{a,b}, \quad \text{for all } a, b, \ell. \quad (6.23)$$

Now if $\mathcal{N}(a, b; \ell) \leq \varepsilon \cdot 2^{a+b}$ then the probability on the right-hand side of (6.21) is zero. Thus, by considering the values of $p_{a,b}$, (6.23) will allow us to rule out certain values of a and b a priori, by showing that those values have $\mathcal{N}(a, b; \ell) \leq \varepsilon \cdot 2^{a+b}$ with high probability. The other values of a and b not ruled out here will be handled afterwards.

We will derive an entropy inequality for the values of a and b *not* ruled out by this idea. First, by convexity, we have

$$H_p\left(\frac{a}{d}\right) + H_p\left(\frac{b}{d}\right) \geq 2H_p\left(\frac{a+b}{2d}\right),$$

and so, defining

$$p_{a+b} := 2^{-2dH_p(\frac{a+b}{2d})}.$$

we obtain $p_{a,b} \leq p_{a+b}$. Thus $\mathcal{N}(a, b; \ell)$ is stochastically dominated by $\text{Bin}(2^d d^2, p_{a+b})$, and so

$$\begin{aligned} \mathbb{P}[\mathcal{N}(a, b; \ell) > \varepsilon \cdot 2^{a+b}] &\leq \mathbb{P}[\text{Bin}(2^d d^2, p_{a+b}) > \varepsilon \cdot 2^{a+b}] \\ &\leq 2^{d(1-2H_p(t)-2t)} \cdot \text{poly}(d), \end{aligned}$$

for $t = \frac{a+b}{2d}$, via Markov's inequality. The bound vanishes exponentially fast in d whenever

$$1 - 2H_p(t) - 2t < -\delta \quad (6.24)$$

for some fixed absolute constant $\delta > 0$. In other words, when we consider values of a, b for which (6.24) holds with $t = \frac{a+b}{2d}$, the right-hand side of (6.21) is zero with high probability.

This means we only need to analyze the values of a and b for which (6.24) fails, i.e. for which

$$2H_p(t) + 2t \leq 1 + \delta, \quad (6.25)$$

for some fixed absolute constant $\delta > 0$, where $t = \frac{a+b}{2d}$. So for the remainder of the proof we assume this condition holds and also that $\mathcal{N}(a, b; \ell) > \varepsilon \cdot 2^{a+b}$, so that we may effectively use the binomial entropy bounds collected in the appendix (again, if this condition fails then the probability (6.21) is simply zero).

Let us return to (6.21) with these assumptions. We wish to apply Lemma B.1, which states that

$$\mathbb{P}[\text{Bin}(n, q_1) > nq_2] \leq 2^{-n \cdot H_{q_1}(q_2)}$$

when $q_1 \leq q_2$. We will then apply Lemma B.2, which states that

$$H_{q_1}(q_2) \geq \frac{1}{2} q_2 \cdot \log_2 \left(\frac{q_2}{q_1} \right)$$

whenever $q_2 \geq 10q_1$. To use these two lemmas for (6.21), where $n = \mathcal{N}(a, b; \ell)$, $q_1 = 2^{-a}$, and $nq_2 = \varepsilon \cdot 2^{a+b}$, we need

$$10 \cdot \mathcal{N}(a, b; \ell) \cdot 2^{-a} \leq \varepsilon \cdot 2^{a+b}, \quad \text{i.e.} \quad \mathcal{N}(a, b; \ell) \leq \frac{\varepsilon}{10} 2^{2a+b}. \quad (6.26)$$

Under (6.23), which holds with probability $\geq 1 - \frac{1}{d}$, (6.26) holds as long as

$$2^d \cdot g(d) \cdot p_{a,b} \leq \frac{\varepsilon}{10} 2^{2a+b}.$$

Recalling the definition of $p_{a,b}$ from (6.22), this is equivalent to

$$2^{-d(2s+H_p(s)+r+H_p(r))} \leq 2^{-d} \frac{\varepsilon}{10g(d)},$$

where we have defined $s = \frac{a}{d}$ and $r = \frac{b}{d}$. Since $\frac{\varepsilon}{10g(d)}$ is only polynomially small, the above is implied for d large enough whenever

$$2s + H_p(s) + r + H_p(r) > 1 + \delta_1 \quad (6.27)$$

for some universal $\delta_1 > 0$. To understand the above inequality, we turn to Lemma B.3 which states that for any $m \geq 1$, the function

$$f_m(p) := \inf\{ms + H_p(s) : s \in [0, 1]\}$$

is continuous, strictly increasing in $p \in (0, 1)$, and satisfies

$$f_m\left(\frac{1}{2}\right) = m + 1 - \log_2(2^m + 1).$$

With $m = 1$ and $m = 2$ respectively, this yields

$$r + H_{\frac{1}{2}}(r) \geq 2 - \log_2 3 \approx 0.415, \quad (6.28)$$

$$2s + H_{\frac{1}{2}}(s) \geq 3 - \log_2 5 \approx 0.678$$

for any $r, s \in [0, 1]$. The sum of these is strictly greater than 1, which means, by the continuity and increasing nature of f_m as a function of p , that there are some universal δ_1, γ_1 such that (6.27) holds for all $r, s \in [0, 1]$ and all $p > \frac{1}{2} - \gamma_1$.

All of this is to say that we may in fact apply Lemmas B.1 (which bounds binomial probabilities in terms of entropy $H_p(q)$) and B.2 (which gives a helpful estimate on $H_p(q)$) to (6.21), and we obtain

$$\pi\left[\mathcal{C}(a, b; \ell; \mathbf{O}_1; S) \geq \varepsilon \cdot 2^{a+b}\right] = \mathbb{P}\left[\text{Bin}(\mathcal{N}(a, b; \ell), 2^{-a}) \geq \varepsilon \cdot 2^{a+b}\right] \leq 2^{-\text{exponent}} \quad (6.29)$$

where

$$\text{exponent} := \mathcal{N}(a, b; \ell) \cdot H_{2^{-a}} \left(\frac{\varepsilon \cdot 2^{a+b}}{\mathcal{N}(a, b; \ell)} \right) \quad (\text{Lemma B.1})$$

$$\geq \mathcal{N}(a, b; \ell) \cdot \frac{1}{2} \cdot \left(\frac{\varepsilon \cdot 2^{a+b}}{\mathcal{N}(a, b; \ell)} \right) \log_2 \left(\frac{\varepsilon \cdot 2^{a+b}}{\mathcal{N}(a, b; \ell) \cdot 2^{-a}} \right) \quad (\text{Lemma B.2})$$

$$\geq \frac{\varepsilon}{2} \log_2(10) \cdot 2^{a+b} \quad (\text{Equation (6.26)})$$

$$= \frac{\varepsilon}{2} \log_2(10) \cdot 2^{2dt}. \quad (\text{Definition of } t)$$

Note that since (6.26) holds for all a, b, ℓ simultaneously with probability at least $1 - \frac{1}{d}$ (since it is a consequence of (6.23)), the above inequality for the exponent also holds simultaneously for all a, b, ℓ with the same probability. So we find that

$$\pi \left[\mathcal{C}(a, b; \ell, \mathbf{O}_1; S) > \varepsilon \cdot 2^{a+b} \right] \leq \exp \left(-\frac{\varepsilon}{2} \log_2(10) \cdot 2^{2dt^*} \right) \quad (6.30)$$

simultaneously for all a, b, ℓ with probability at least $1 - \frac{1}{d}$, where, recalling condition (6.25), we define

$$t^* := \inf \{ t : 2H_p(t) + 2t \leq 1 + \delta \}. \quad (6.31)$$

Eventually, we will take $\alpha_p = 2^{2t^*}$, and we will analyze the above optimization problem to show that $\alpha_p > \frac{(2-p)^2}{2}$, as stated in the lemma.

As mentioned earlier, the same upper bound holds for the $\mathbf{O}_2$ case with the same proof. At this point we will derive the same upper bound for the $\mathbf{N}$ case, after which we will analyze the optimization problem just obtained to finish the proof.

Case 2: $\sigma = \mathbf{N}$.

In this case, the dimer $\mathbf{d}$ does not overlap (share a vertex) with S , but there is a singleton $v \in S$ such that $v \sim \mathbf{d}$ (choose arbitrarily if there are multiple such singletons). The arguments will be very similar to the overlap case, with some extra work to keep track of the properties of v .

Let us further subdivide $\mathcal{C}(a, b; \ell; \mathbf{N}; S)$ depending on the type of the singleton v (i.e. the value of $N_p(v)$). In other words, let $\mathcal{C}(a, b, c; \ell; \mathbf{N}; S)$ be the count of pairs $(\mathbf{d}, v)$, such that $\mathbf{d}$ is in layer ℓ , has type (a, b) , status $\mathbf{N}$ with respect to S , and $\mathbf{d} \sim v \in S$ with type $N_p(v) = c$. Then,

$$\mathcal{C}(a, b; \ell; \mathbf{N}; S) \leq \sum_{c=0}^d \mathcal{C}(a, b, c; \ell; \mathbf{N}; S)$$

(this is not an equality only because some dimers will be counted in multiple summands if they have multiple neighbors v in S). If this is at least $\varepsilon \cdot 2^{a+b}$, we must then have some choice of c such that

$$\mathcal{C}(a, b, c; \ell; \mathbf{N}; S) \geq \frac{\varepsilon}{d+1} \cdot 2^{a+b} =: \varepsilon' \cdot 2^{a+b},$$

where $\varepsilon' = \varepsilon / \text{poly}(d)$.

Let $\mathcal{N}(a, b, c; \ell)$ be the count of pairs $(\mathbf{d}, v)$ of the kind counted in $\mathcal{C}(a, b, c; \ell; \mathbf{N}; S)$ but excluding the status- $\mathbf{N}$ condition and the $v \in S$ condition (note that this count only depends on the percolation configuration).

Now for every pair $(\mathbf{d}, v)$ counted in $\mathcal{N}(a, b, c; \ell)$, it is counted in $\mathcal{C}(a, b, c; \ell; \mathbf{N}; S)$ if and only if $v \in S$ and $\mathbf{d} \cap S = \emptyset$, which happens with probability at most 2^{-c} (under π). Further, no two pairs $(\mathbf{d}, v), (\mathbf{d}', v')$ counted in $\mathcal{N}(a, b, c; \ell)$ satisfy $v = v'$ (unless $\mathbf{d} = \mathbf{d}'$), since otherwise this would imply $\mathbf{d} \sim v = v' \sim \mathbf{d}'$ and thus $\mathbf{d}$ and $\mathbf{d}'$ are within a distance of 2 in the shared-neighbor graph, which is ruled out by construction of the layers $\mathfrak{L}_\ell$. Consequently, under π , $\mathcal{C}(a, b, c; \ell; \mathbf{N}; S)$ is stochastically dominated by a $\text{Bin}(\mathcal{N}(a, b, c; \ell), 2^{-c})$ random variable and thus

$$\pi\left[\mathcal{C}(a, b, c; \ell; \mathbf{N}; S) > \varepsilon' \cdot 2^{a+b}\right] \leq \mathbb{P}\left[\text{Bin}(\mathcal{N}(a, b, c; \ell), 2^{-c}) > \varepsilon' \cdot 2^{a+b}\right]. \quad (6.32)$$

As in the overlap case, we will now estimate $\mathcal{N}(a, b, c; \ell)$. By definition we have

$$\begin{aligned} \mathcal{N}(a, b, c; \ell) &= \sum_{\substack{\mathbf{d} \in \mathfrak{L}_\ell, \\ v \sim \mathbf{d}}} \mathbf{1}_{\{N_p(v)=c\}} \cdot \mathbf{1}_{\{\mathbf{d} \text{ has type } (a,b)\}} \\ &\leq 2d^2 \sum_{\mathbf{d} \in \mathfrak{L}_\ell} \mathbf{1}_{E_{\mathbf{d}}}, \end{aligned}$$

since there are at most $2d^2$ neighbors v of each $\mathbf{d}$, where we set

$$E_{\mathbf{d}} := \{\exists v \sim \mathbf{d} : N_p(v) = c\} \cap \{\mathbf{d} \text{ has type } (a, b)\}.$$

By a union bound over all possible neighbors v of $\mathbf{d} = (u, w)$, we have

$$\begin{aligned} \mathbb{P}[E_{\mathbf{d}}] &\leq 2d^2 \cdot \mathbb{P}[N_p(v) = c, N_p(u) = a, N_p(w) = b] \\ &\leq 2d^2 \cdot \mathbb{P}[N_p(v) + N_p(u) + N_p(w) = a + b + c] \\ &\leq 2d^2 \cdot 2^{-3dH_p(\frac{a+b+c}{3d})} =: p_{a+b+c}, \end{aligned}$$

using the fact that $N_p(v) + N_p(u) + N_p(w)$ is a $\text{Bin}(3d, p)$ random variable and applying Lemma B.1. Note also that the indicators $\mathbf{1}_{E_{\mathbf{d}}}$ for $\mathbf{d} \in \mathfrak{L}_\ell$ are independent (with respect to the percolation measure) by the separation within a layer. Thus $\mathcal{N}(a, b, c; \ell)$ is stochastically dominated by $2d^2 \cdot \text{Bin}(2^d d^2, p_{a+b+c})$. By the same argument as was used for (6.23), there is some polynomial $g(d)$ such that with probability $\geq 1 - \frac{1}{d}$, we have

$$\mathcal{N}(a, b, c; \ell) \leq 2^d \cdot g(d) \cdot p_{a+b+c}, \quad \text{for all } a, b, c, \ell. \quad (6.33)$$

Once again, we know that the right-hand side of (6.32) is zero if

$$\mathcal{N}(a, b, c; \ell) \leq \varepsilon' \cdot 2^{a+b}, \quad (6.34)$$

and we can use (6.33) to determine which values of a, b, c satisfy (6.34) with high probability. This time, since

$$\mathcal{N}(a, b, c; \ell) \leq 2d^2 \mathcal{N}(a, b; \ell),$$

the condition (6.34) is implied by

$$\mathcal{N}(a, b; \ell) \leq \varepsilon'' \cdot 2^{a+b}$$

for some $\varepsilon'' \geq \frac{\varepsilon'}{\text{poly}(d)}$. So by the same argument as in the previous case, with high probability we have $\mathcal{N}(a, b, c; \ell) \leq \varepsilon' \cdot 2^{a+b}$ simultaneously for all a, b, c, ℓ which do not satisfy

$$2H_p(t) + 2t \leq 1 + \delta \quad (6.35)$$

for some fixed $\delta > 0$, where $t = \frac{a+b}{2d}$ as before. For the rest of the proof we assume this condition holds and that $\mathcal{N}(a, b, c; \ell) > \varepsilon' \cdot 2^{a+b}$, otherwise the upper bound is just zero.

With these assumptions we return to (6.32). Recall that in the previous case, we bounded the analogous probability (6.21) by applying Lemmas B.1 and B.2 (which bound binomial probabilities in terms of entropy, and give a convenient estimate on the entropy, respectively) to obtain (6.29). We will apply the same strategy here, and, analogous to the conditions (6.26) required in the previous case, we will require the following conditions to apply Lemmas B.1 and B.2.

$$10 \cdot \mathcal{N}(a, b, c; \ell) \cdot 2^{-c} \leq \varepsilon' \cdot 2^{a+b}, \quad \text{i.e.} \quad \mathcal{N}(a, b, c; \ell) \leq \frac{\varepsilon'}{10} 2^{a+b+c}. \quad (6.36)$$

As before, if (6.33) holds (which happens with probability at least $1 - \frac{1}{d}$), then (6.36) holds as long as

$$2^d \cdot g(d) \cdot p_{a+b+c} \leq \frac{\varepsilon'}{10} 2^{a+b+c}.$$

This is equivalent to

$$2^{3d(s+H_p(s))} \geq 2^d \frac{10g(d)}{\varepsilon'},$$

where we have defined $s = \frac{a+b+c}{3d}$. Since $\frac{10g(d)}{\varepsilon'}$ is polynomial, this holds for high enough d whenever

$$s + H_p(s) > \frac{1}{3} + \delta_2 \quad (6.37)$$

for some universal $\delta_2 > 0$. As in the previous case, specifically in (6.28), we now apply Lemma B.3 with $m = 1$ and we find that

$$s + H_{\frac{1}{2}}(s) \geq 2 - \log_2 3 \approx 0.415 > \frac{1}{3}$$

for all $s \in [0, 1]$. Again by the continuity and increasing nature of the optimization problem, afforded by Lemma B.3, there are some universal δ_2, γ_2 such that (6.37) holds for all $s \in [0, 1]$ and all $p > \frac{1}{2} - \gamma_2$. In other words, we may indeed apply Lemmas B.1 and B.2 to (6.32), obtaining

$$\pi \left[\mathcal{C}(a, b, c; \ell; \mathbf{N}; S) \geq \varepsilon' 2^{a+b} \right] \leq \mathbb{P} \left[\text{Bin}(\mathcal{N}(a, b, c; \ell), 2^{-c}) \geq \varepsilon' \cdot 2^{a+b} \right] \leq 2^{-\text{exponent}},$$

where

$$\text{exponent} := \mathcal{N}(a, b; \ell) \cdot H_{2^{-c}} \left(\frac{\varepsilon' \cdot 2^{a+b}}{\mathcal{N}(a, b, c; \ell)} \right) \quad (\text{Lemma B.1})$$

$$\geq \mathcal{N}(a, b, c; \ell) \cdot \frac{1}{2} \cdot \left(\frac{\varepsilon' \cdot 2^{a+b}}{\mathcal{N}(a, b, c; \ell)} \right) \log_2 \left(\frac{\varepsilon' \cdot 2^{a+b}}{\mathcal{N}(a, b, c; \ell) \cdot 2^{-c}} \right) \quad (\text{Lemma B.2})$$

$$\geq \frac{\varepsilon'}{2} \log_2(10) \cdot 2^{a+b} \quad (\text{equation (6.36)})$$

$$= \frac{\varepsilon'}{2} \log_2(10) \cdot 2^{2dt}. \quad (\text{Definition of } t)$$

Thus we obtain the same upper bound as in (6.30) with an exponential rate given in terms of the same optimization problem (6.31) over t , with the same condition (6.35) on t as in the $\mathcal{O}_1$ case (recall (6.25)). Note that the polynomial prefactor inside the exponential in this case is potentially different due to the union bound over c , but it is still uniform over a, b, ℓ .

Final analysis. In all cases $\sigma = \{O_1, O_2, N\}$ above, we arrive at the same upper bound (6.30) with the same exponential rate given in terms of the same optimization problem (6.31). Thus, by a union bound, there is some function $f(d) > \text{poly}(d)^{-1}$ depending on the choice of $\varepsilon > \text{poly}(d)^{-1}$ such that, with probability tending to 1 we have the following inequality simultaneously for all a, b, ℓ, σ :

$$\pi \left[\mathcal{C}(a, b; \ell; \sigma; S) > \varepsilon \cdot 2^{a+b} \right] \leq \exp \left(-f(d) \cdot 2^{2dt^*} \right),$$

where t^* solves the optimization problem (6.31), restated as follows:

$$t^* = \inf \{ t : 2H_p(t) + 2t \leq 1 + \delta \}$$

for some small constant $\delta > 0$ which we may choose. As mentioned before, we will take $\alpha_p = 2^{2t^*}$. So, to finish the proof of the lemma, we must simply show that $\alpha_p > \frac{(2-p)^2}{2}$.

For this, a calculation relegated to Lemma B.4 shows that

$$\frac{(2-p)^2}{2} = 2^{1-2s^*-2H_p(s^*)},$$

where s^* maximizes $1 - 2s - 2H_p(s)$ over $s \in [0, 1]$. So to prove that $\alpha_p > \frac{(2-p)^2}{2}$, it suffices to show that

$$\sup \{ 1 - 2s - 2H_p(s) : s \in [0, 1] \} < 2 \cdot \inf \{ t : 2H_p(t) + 2t \leq 1 + \delta \},$$

or in other words that

$$1 - 2s - 2H_p(s) < 2t \tag{6.38}$$

for all $s \in [0, 1]$ and all t satisfying $2H_p(t) + 2t \leq 1 + \delta$. By this condition on t , (6.38) would hold if we had

$$1 - 2s - 2H_p(s) + (1 + \delta) < 2t + (2t + 2H_p(t)) \tag{6.39}$$

for all $t, s \in [0, 1]$. Indeed, (6.38) is obtained by adding the inequality

$$-(1 + \delta) \leq -(2H_p(t) + 2t)$$

to (6.39). Rearranging (6.39) yields

$$2(s + H_p(s)) + 2(2t + H_p(t)) > 2 + \delta,$$

and so we may apply Lemma B.3 again with $m = 1$ and $m = 2$. This yields

$$2(2t + H_{\frac{1}{2}}(t)) \geq 2(3 - \log_2(5)) > 1.35$$

$$2(s + H_{\frac{1}{2}}(s)) \geq 2(2 - \log_2(3)) > 0.83,$$

for all $s, t \in [0, 1]$, and $1.35 + 0.83 = 2.18$ is strictly larger than 2. By the continuity and monotonicity from Lemma B.3, there is some $\gamma_3 > 0$ such that for $p > \frac{1}{2} - \gamma_3$, the sum of the relevant quantities is still at least $2 + \delta$, for our fixed universal choice of δ , as long as it is small enough. Thus, taking $\gamma = \min\{\gamma_1, \gamma_2, \gamma_3\}$, the proof is concluded. $\square$

7. DISTRIBUTIONAL LIMITS

In this section we prove Proposition 2.7, which states that

$$\left(\frac{\Psi^{\text{Even}} - \mu}{\sigma}, \frac{\Psi^{\text{Odd}} - \mu}{\sigma} \right) \xrightarrow{w} \mathcal{N}(0, 1) \otimes \mathcal{N}(0, 1)$$

for some explicit deterministic quantities $\mu = \mu_{d,p}$ and $\sigma = \sigma_{d,p}$ (initially defined in (1.11) and (1.5)), and for all p larger than some specific threshold 0.455. It should be noted that the threshold 0.455 is not a threshold for any qualitative change, but rather the point above which the *specific approximate mean* $\mu_{d,p}$ we use is a good enough approximation. This threshold is unrelated to the qualitatively meaningful threshold $\frac{1}{2} - \gamma$ appearing in Corollary 2.6 (which is a consequence of Proposition 6.1), aside from the fact that we have chosen the specific approximate mean in (1.11) to be valid for all possible values of $p > \frac{1}{2} - \gamma \approx 0.465$.

Recall that in the previous few sections, we have finished the proof of Corollary 2.6, deriving an in-probability approximation for the number of independent sets in $Q_{d,p}$, namely that

$$i(Q_{d,p}) \stackrel{\mathbb{P}}{\sim} 2^{2^{d-1}} \cdot (e^{\Psi^{\text{Even}}} + e^{\Psi^{\text{Odd}}}),$$

where the random variables Ψ^{Even} and Ψ^{Odd} can be written explicitly in terms of the percolation configuration. We remind the reader that, as initially defined in (1.7), we have

$$\Psi^{\mathcal{H}} = \log \tilde{\mathcal{Z}}^{\mathcal{H}} + \Delta^{\mathcal{H}} - \tilde{\Delta}^{\mathcal{H}},$$

for $\mathcal{H} \in \{\text{Even}, \text{Odd}\}$, where, as defined in (2.7) and (2.11),

$$\tilde{\mathcal{Z}}^{\mathcal{H}} = \prod_{v \in \mathcal{H}} (1 + \varphi_v), \quad \Delta^{\mathcal{H}} = \sum_{\mathbf{d} \in \text{Dimers}^{\mathcal{H}}} \varphi_{\mathbf{d}}, \quad \text{and} \quad \tilde{\Delta}^{\mathcal{H}} = \sum_{\mathbf{d} \in \text{Dimers}^{\mathcal{H}}} \tilde{\varphi}_{\mathbf{d}}.$$

It is also worth remarking that the requirement that p is bigger than some threshold in the statement is not necessary for the central limit theorem itself to hold, when Ψ^{Even} and Ψ^{Odd} are re-centered by their exact mean and rescaled by their exact standard deviation. However, the exact mean and standard deviation do not have explicit closed-form expressions in general. So we introduce the simpler closed-form quantities μ and σ which work as proxies for the mean and standard deviation, but only for p large enough. The threshold for the validity of these approximations is below the threshold $\frac{1}{2} - \gamma$ in Corollary 2.6/Proposition 6.1, so these approximations do not restrict the range of p for which our main theorem holds.

The first step towards proving the CLT will show that the quantities $\log \tilde{\mathcal{Z}}^{\text{Even}}$ and $\log \tilde{\mathcal{Z}}^{\text{Odd}}$ satisfy a joint central limit theorem. We then show that the quantities $\Delta^{\mathcal{H}}$ and $\tilde{\Delta}^{\mathcal{H}}$ both concentrate around deterministic values for each $\mathcal{H} \in \{\text{Even}, \text{Odd}\}$. In other words, the fluctuations of $(\Psi^{\text{Even}}, \Psi^{\text{Odd}})$ are driven purely by those of $(\log \tilde{\mathcal{Z}}^{\text{Even}}, \log \tilde{\mathcal{Z}}^{\text{Odd}})$.

Before moving on, we remark that the proof of the central limit theorem for $(\log \tilde{\mathcal{Z}}^{\text{Even}}, \log \tilde{\mathcal{Z}}^{\text{Odd}})$ uses the same techniques as our previous proof of a similar central limit theorem for a simplification of these variables in [6, Lemma 3.1]. Only a minor modification of that proof is required for the present setting, and the concentration for $\Delta^{\mathcal{H}}$ and $\tilde{\Delta}^{\mathcal{H}}$ only requires a few novel but relatively routine moment calculations. Nevertheless, we include complete proofs for the reader's convenience.

7.1. Joint central limit theorem for sums of logarithmic expressions. For notational convenience, as in (5.2), we set

$$\Phi_{\log}^{\mathcal{H}} = \log \tilde{\mathcal{Z}}^{\mathcal{H}} = \sum_{v \in \mathcal{H}} \log(1 + \varphi_v),$$

for $\mathcal{H} \in \{\text{Even}, \text{Odd}\}$. When appropriate, we will drop the superscript $\mathcal{H}$; for instance, the mean and variance of these sums are identical for both sides, and we will denote these simply by $\mathbb{E}[\Phi_{\log}]$ and $\text{Var}(\Phi_{\log})$. Note that, while both $\Phi_{\log}^{\text{Even}}$ and $\Phi_{\log}^{\text{Odd}}$ are sums of independent random variables, the two sums themselves are not independent, because deleting one edge affects the neighborhood of one odd vertex and one even vertex. However, their dependence is sufficiently weak that we will be able to prove the following joint CLT to a pair of i.i.d. standard normal random variables.

Lemma 7.1. *For $p \in (0, 1)$,*

$$\left(\frac{\Phi_{\log}^{\text{Even}} - \mathbb{E}[\Phi_{\log}]}{\sqrt{\text{Var}(\Phi_{\log})}}, \frac{\Phi_{\log}^{\text{Odd}} - \mathbb{E}[\Phi_{\log}]}{\sqrt{\text{Var}(\Phi_{\log})}} \right) \xrightarrow{w} \mathcal{N}(0, 1) \otimes \mathcal{N}(0, 1).$$

We will invoke Stein's method to prove CLT for random variables with a sparse dependency graph (a graph with the random variables as the vertex set and an edge set such that any vertex is independent of the set of all variables outside its neighborhood). The precise formulation we will rely on is [14, Theorem 3.6] which was adapted from the main result of [1].

Theorem 7.2. [14, Theorem 3.6] *Let $X_1, \dots, X_n$ be random variables with $\mathbb{E}[X_i^4] < \infty$, $\mathbb{E}[X_i] = 0$, $\sigma^2 = \text{Var}(\sum_{i=1}^n X_i)$, and define $\tilde{W} = \sum_{i=1}^n X_i / \sigma$. Let the collection $(X_i, \dots, X_i)$ have dependency neighborhoods N_i , $i = 1, \dots, n$, with $D := \max_{1 \leq i \leq n} |N_i|$. Then*

$$d_{\text{Was}}(\tilde{W}, W) \lesssim \frac{D^2}{\sigma^3} \sum_{i=1}^n \mathbb{E}|X_i|^3 + \frac{D^{3/2}}{\sigma^2} \sqrt{\sum_{i=1}^n \mathbb{E}[X_i^4]}, \quad (7.1)$$

where W is a standard normal random variable and d_{Was} is the 1-Wasserstein distance.

In the above statement, the *dependency neighborhood* of a variable X_i in the collection $\{X_1, \dots, X_n\}$ is the set of indices j for which X_i and X_j are not independent.

Presently we will not be concerned with the precise rate of convergence under the 1-Wasserstein distance, and so we forgo a precise definition of this distance. Suffice it to say, $\tilde{W}$ converges in distribution to W if $d_{\text{Was}}(\tilde{W}, W) \rightarrow 0$, so we will be satisfied with showing that the right-hand side of (7.1) converges to 0.

Proof of Lemma 7.1. Fix $c = (c_{\text{Even}}, c_{\text{Odd}})$ in the unit circle in $\mathbb{R}^2$, and set

$$\Phi^c = c_{\text{Even}} \left(\Phi_{\log}^{\text{Even}} - \mathbb{E}[\Phi_{\log}] \right) + c_{\text{Odd}} \left(\Phi_{\log}^{\text{Odd}} - \mathbb{E}[\Phi_{\log}] \right).$$

We can write Φ^c as a sum of the following 2^d variables indexed by vertices v of Q_d :

$$X_v := c_{\mathcal{H}} (\log(1 + \varphi_v) - \mathbb{E}[\log(1 + \varphi_v)]),$$

where we take $\mathcal{H} = \text{Even}$ if $v \in \text{Even}$ and $\mathcal{H} = \text{Odd}$ if $v \in \text{Odd}$.

We will apply Theorem 7.2 to conclude that

$$\frac{\Phi^c}{\sqrt{\mathbf{Var}(\Phi^c)}} \xrightarrow{w} \mathcal{N}(0, 1).$$

Since this holds for an arbitrary c in the unit circle, the Cramer-Wold theorem then implies that

$$\left(\frac{\Phi_{\log}^{\text{Even}} - \mathbb{E}[\Phi_{\log}]}{\sqrt{\mathbf{Var}(\Phi^c)}}, \frac{\Phi_{\log}^{\text{Odd}} - \mathbb{E}[\Phi_{\log}]}{\sqrt{\mathbf{Var}(\Phi^c)}} \right) \xrightarrow{w} \mathcal{N}(0, 1) \otimes \mathcal{N}(0, 1). \quad (7.2)$$

Now, by Lemma A.3 we have

$$\begin{aligned} \mathbf{Var}(\Phi^c) &= c_{\text{Even}}^2 \mathbf{Var}(\Phi_{\log}^{\text{Even}}) + c_{\text{Odd}}^2 \mathbf{Var}(\Phi_{\log}^{\text{Odd}}) + c_{\text{Even}} c_{\text{Odd}} \mathbf{Cov}(\Phi_{\log}^{\text{Even}}, \Phi_{\log}^{\text{Odd}}) \\ &= \mathbf{Var}(\Phi_{\log})(1 + o(1)), \end{aligned}$$

and so (7.2) implies the conclusion of the lemma.

So it just remains to verify the conditions of Theorem 7.2 and show that the right-hand side of (7.1) converges to 0. Of course by definition we have $\mathbb{E}[X_v] = 0$. To bound the moments of the X_v , notice that

$$\begin{aligned} \mathbb{E}[X_v^k] &\leq \mathbb{E}[(\log(1 + \varphi_v) - \mathbb{E}[\log(1 + \varphi_v)])^k] \\ &\leq \sum_{j=0}^k \binom{k}{j} \mathbb{E}[(\log(1 + \varphi_v))^j] \mathbb{E}[\log(1 + \varphi_v)]^{k-j} \\ &\leq \sum_{j=0}^k \binom{k}{j} \mathbb{E}[\varphi_v^j] \mathbb{E}[\varphi_v]^{k-j} \\ &\leq \sum_{j=0}^k \binom{k}{j} \mathbb{E}[\varphi_v^k]^{j/k} \mathbb{E}[\varphi_v^k]^{(k-j)/k} \\ &= 2^k \cdot \mathbb{E}[\varphi_v^k], \end{aligned}$$

where we used the fact that $\log(1 + \varphi_v) \leq \varphi_v$ (recall that $\varphi_v > 0$) and then applied Jensen's inequality. So we have

$$\mathbb{E}[X_v^3] \lesssim \mathbb{E}[\varphi_v^3] = \left(\frac{8-7p}{8} \right)^d,$$

and

$$\mathbb{E}[X_v^4] \lesssim \mathbb{E}[\varphi_v^4] = \left(\frac{16-15p}{16} \right)^d,$$

where these moments are computed in the appendix, in Lemma A.1 (a).

Now, for the size of the dependency neighborhoods, each X_v is independent of all but d other variables, which are exactly the ones corresponding to neighbors of v . So by (7.1), the 1-Wasserstein distance between $\mathbf{Var}(\Phi^c)^{-1/2} \Phi^c$ and a standard normal random variable is, up to constants, at most

$$\frac{d^2 2^d \left(\frac{8-7p}{8} \right)^d}{\left(\frac{4-3p}{2} \right)^{3d/2}} + \frac{d^{3/2} \sqrt{2^d \left(\frac{16-15p}{16} \right)^d}}{\left(\frac{4-3p}{2} \right)^d} = d^2 \left(\frac{(8-7p)^2}{2(4-3p)^3} \right)^{d/2} + d^{3/2} \left(\frac{16-15p}{2(4-3p)^2} \right)^{d/2}.$$

The proof will thus be complete if we can show that both

$$\frac{(8-7p)^2}{2(4-3p)^3} < 1 \quad \text{and} \quad \frac{16-15p}{2(4-3p)^2} < 1.$$

First notice that the second bound implies the first, since

$$(8-7p)^2 \leq (16-15p)(4-3p).$$

To verify the above inequality, note that the difference

$$(8-7p)^2 - (16-15p)(4-3p)$$

is a quadratic with a positive coefficient of p^2 (namely $49 - 45$) which vanishes at $p = 0$ and $p = 1$, meaning that it must be negative for $0 < p < 1$. Now, to actually prove the second inequality, note that

$$2(4-3p)^2 - (16-15p) = 18p^2 - 33p + 16$$

is another quadratic with a positive coefficient of p^2 , and with discriminant

$$33^2 - 4 \cdot 16 \cdot 18 = -63,$$

so it is positive for all p , finishing the proof. $\square$

7.2. Approximations of the expectation and variance. Lemma 7.1 exhibits a joint central limit theorem for $\Phi_{\log}^{\text{Even}}$ and $\Phi_{\log}^{\text{Odd}}$ upon centering by $\mathbb{E}[\Phi_{\log}]$ and rescaling by $\sqrt{\text{Var}(\Phi_{\log})}$. However, there is no simple closed form for these quantities in general, and this is particularly an issue when $p \leq \frac{1}{2}$, as will be discussed in Remark 7.3 below. In this section, we construct quantities $\mu = \mu_{d,p}$ and $\sigma = \sigma_{d,p}$ with explicit closed-form expressions that can be used as a replacement for the mean and standard deviation of $\Phi_{\log}$ in the central limit theorem. Specifically, these quantities will satisfy

$$\sigma^2 = \text{Var}(\Phi_{\log})(1 + o(1)) \quad \text{and} \quad \mu = \mathbb{E}[\Phi_{\log}] + o(\sigma),$$

which suffice for our purposes.

For σ , a simple moment calculation which has been relegated to the appendix as Lemma A.2 (a) shows that

$$\text{Var}(\Phi_{\log}) = \frac{1}{2} \left(2 - \frac{3}{2}p \right)^d (1 + o(1))$$

for all $p \in (0, 1)$. So, as previously defined in (1.5), we set

$$\sigma^2 = \frac{1}{2} \left(2 - \frac{3}{2}p \right)^d.$$

Next we turn to μ . An initial strategy for approximating $\mathbb{E}[\Phi_{\log}]$ might be to expand $\log(1 + \varphi_v)$ in a Taylor series and then interchange the expectation and summation; since the moments of φ_v have explicit expressions given in Lemma A.1, this would in principle lead to a series expansion for $\mathbb{E}[\Phi_{\log}]$ which can be truncated to obtain a closed-form expression. This does indeed work for $p > \frac{1}{2}$, and leads to an approximation of $\mathbb{E}[\Phi_{\log}]$ with additive $o(1)$ error. However, for $p \leq \frac{1}{2}$, every term in the series grows exponentially with d , so there is no such approximation of $\mathbb{E}[\Phi_{\log}]$ with only $o(1)$ additive error.

Remark 7.3. The ability or inability to accurately approximate $\mathbb{E}[\Phi_{\log}]$ by a closed-form expression, or equivalently to approximate $\Phi_{\log}$ itself by a polynomial in the variables φ_v , appears to be the extent of the changes in this model at $p = \frac{1}{2}$. So, unlike the many other features of $Q_{d,p}$ which exhibit a qualitative phase transition at $p = \frac{1}{2}$, the behavior of independent sets does not change sharply at this value.

In any case, for such values of $p \leq \frac{1}{2}$, the standard deviation proxy σ itself also grows exponentially with d , and only finitely many of the terms in the series $\Phi_{\log}$ grow faster, allowing us to obtain an approximation with $o(\sigma)$ additive error, meaning this error is dominated in the CLT by the normalization by σ . Specifically, for the range of p that we consider, we only need to include three terms in the series.

Namely, for $k \geq 0$ we define

$$\mu_1^k := 2^{d-1} \mathbb{E}[\varphi_v^k] = \frac{1}{2} \left(2 - \frac{2^k - 1}{2^{k-1}} p \right)^d, \quad (7.3)$$

where the moment calculation here follows from Lemma A.1 (b). Then, Lemma A.2 (b) states that

$$\mathbb{E}[\Phi_{\log}] = \mu_1 - \frac{\mu_1^2}{2} + \frac{\mu_1^3}{3} + o(\sigma)$$

for all $p > 0.455$. It's worth noting that this range of p is simply where $\mu_1^4 \ll \sigma$, and is not related to the threshold $\frac{1}{2} - \gamma$ of our main theorem. We also remark again that numerical calculations show that the threshold $\frac{1}{2} - \gamma$ can be taken as low as approximately 0.465, although we do not prove this rigorously. So in principle this three-term approximation holds for the full possible range of our theorem, and regardless it holds for values of p strictly below $\frac{1}{2}$. Thus we set

$$\mu_1 = \mu_1 - \frac{\mu_1^2}{2} + \frac{\mu_1^3}{3}$$

as in (1.8), and the above discussion leads to the following corollary of Lemma 7.1.

Corollary 7.4. For $p > 0.455$,

$$\left(\frac{\Phi_{\log}^{\text{Even}} - \mu_1}{\sigma}, \frac{\Phi_{\log}^{\text{Odd}} - \mu_1}{\sigma} \right) \xrightarrow{w} \mathcal{N}(0, 1) \otimes \mathcal{N}(0, 1).$$

7.3. Concentration of the dimer sums. We now show that the quantities $\Delta^{\mathcal{H}}$ and $\tilde{\Delta}^{\mathcal{H}}$ concentrate at $o(1)$ scale around deterministic values. As the joint behavior of these quantities across both sides is not relevant, we drop the superscript $\mathcal{H}$ in this section for readability.

In Lemma A.4 (a), a short calculation leads to the following exact expressions for the means of Δ and $\tilde{\Delta}$:

$$\mathbb{E}[\Delta] = \mu_2, \quad \mathbb{E}[\tilde{\Delta}] = \tilde{\mu}_2,$$

where

$$\begin{aligned} \mu_2 &= \frac{d(d-1)}{4} \left(\frac{(2-p)^2}{2} \right)^d \cdot \left(\frac{1 + (1-p)^2}{2-p} \right)^2, \\ \tilde{\mu}_2 &= \frac{d(d-1)}{4} \left(\frac{(2-p)^2}{2} \right)^d, \end{aligned}$$

as previously defined in (1.9) and (1.10). Additionally, Lemma A.4 (b) calculates the second moments of φ_d and $\tilde{\varphi}_d$ and shows that for $p > 0.391$, we have $\text{Var}(\Delta), \text{Var}(\tilde{\Delta}) \ll 1$. Again, this number arises simply as the range of p for which $2^d \mathbb{E}[\varphi_d^2] \ll 1$ and has nothing to do with the threshold $\frac{1}{2} - \gamma$ of our main theorem. This variance calculation yields the following lemma, giving the $o(1)$ concentration of Δ and $\tilde{\Delta}$ around their means.

Lemma 7.5. *For $p > 0.391$ we have*

$$\Delta - \mu_2 \xrightarrow{\mathbb{P}} 0 \quad \text{and} \quad \tilde{\Delta} - \tilde{\mu}_2 \xrightarrow{\mathbb{P}} 0.$$

Now recall from (2.12) that

$$\Psi^{\mathcal{H}} = \Phi_{\log}^{\mathcal{H}} + \Delta^{\mathcal{H}} - \tilde{\Delta}^{\mathcal{H}}$$

for $\mathcal{H} \in \{\text{Even}, \text{Odd}\}$, and set

$$\mu = \mu_1 + \mu_2 - \tilde{\mu}_2,$$

as originally defined in (1.11). Then Corollary 7.4 and Lemma 7.5 yield the following, which was previously stated as Proposition 2.7.

Proposition 7.6. *For $p > 0.455$, we have*

$$\left(\frac{\Psi^{\text{Even}} - \mu}{\sigma}, \frac{\Psi^{\text{Odd}} - \mu}{\sigma} \right) \xRightarrow{w} \mathcal{N}(0, 1) \otimes \mathcal{N}(0, 1).$$

We finally remark that, as can be checked numerically, the quantities μ_2 and $\tilde{\mu}_2$ are indeed of larger order than σ when p is less than 0.508. This means that, to ensure the validity of our main theorem which holds for values of p below $\frac{1}{2}$, these terms must be included in the formula for μ , and cannot be absorbed into an $o(\sigma)$ error term like the tail of the infinite series discussed in Section 7.2.

8. SAMPLING

In this section we prove Theorem 1.3, exhibiting an approximate sampling algorithm which, given a quenched realization of $Q_{d,p}$, yields an independent set with distribution very close to that of a uniformly random independent set in $Q_{d,p}$.

We remind the reader of the structure of this algorithm (initially presented in Definition 1.2), now using notation developed in the body of the work. In order to present the algorithm, let us first introduce the probabilities q^{Even} and q^{Odd} defined by

$$q^{\mathcal{H}} \propto e^{\Psi^{\mathcal{H}}} = \tilde{\mathcal{Z}}^{\mathcal{H}} \cdot \exp(\Delta^{\mathcal{H}} - \tilde{\Delta}^{\mathcal{H}}), \quad q^{\text{Even}} + q^{\text{Odd}} = 1, \quad (8.1)$$

recalling the definition of $\tilde{\mathcal{Z}}^{\mathcal{H}}$ from (2.7) and the definitions of $\Delta^{\mathcal{H}}$ and $\tilde{\Delta}^{\mathcal{H}}$ from (2.11). As we will see shortly in Lemma 8.2 below, q^{Even} and q^{Odd} approximate the probabilities that a uniformly random independent set I in $Q_{d,p}$ has the even side or the odd side as its defect side.

Additionally we remind the reader that the the algorithm may fail (in particular at steps 3 and 5 below), and we enforce our usual convention of not distinguishing between a set of vertices and the corresponding set of polymers, as mentioned in Section 2.1.

Definition 8.1 (ApproxSampler). *Define the random independent set $\hat{I}$ in $Q_{d,p}$ via the following procedure.*

- (1) *Let $\mathcal{H}$ be Even or Odd, with probabilities q^{Even} and q^{Odd} respectively.*
- (2) *Let $\tilde{S}$ denote a subset of $\mathcal{H}$ constructed by including each $v \in \mathcal{H}$ independently with probability $\frac{\varphi_v}{1+\varphi_v}$. In other words, sample $\tilde{S} \sim \pi^{\mathcal{H}}$, where $\pi^{\mathcal{H}}$ was defined in (6.1) (with the superscript $\mathcal{H}$ being implicit in that definition).*
- (3) *If $\tilde{S} \notin \text{PD}_{\leq 2}$ (recall this notation from just above (2.8)), terminate the algorithm with a result of failure. Otherwise, remove all dimers from $\tilde{S}$ and let S_1 denote the resulting set of well-separated singletons.*
- (4) *Let S_2 denote a collection of dimers constructed by including each $\{u, v\} \in \text{Dimers}^{\mathcal{H}}$ independently with probability $\frac{\varphi_d}{1+\varphi_d}$.*
- (5) *If $S_1 \cup S_2 \notin \text{PD}_{\leq 2}$, terminate the algorithm with a result of failure. Otherwise, let $\hat{S} = S_1 \cup S_2$.*
- (6) *Let $\hat{I}$ denote the independent set with $\hat{I} \cap \mathcal{H} = \hat{S}$ and with every vertex in $Q_d \setminus \mathcal{H}$ which does not neighbor $\hat{S}$ included independently with probability $\frac{1}{2}$ each.*

There are three parts to the proof of Theorem 1.3. First, we must show that the defect side of a uniformly random independent set I in $Q_{d,p}$ is $\mathcal{H}$ with probability approximately $q^{\mathcal{H}}$. Next, we must show that ApproxSampler fails (at steps 3 or 5) with low probability. Finally, we must show that if ApproxSampler does not fail, then the set $\hat{S}$ constructed in step 5 is close in distribution to the defect side of a uniformly random independent set with defects on side $\mathcal{H}$, and that this implies that the total variation distance between a uniformly random independent set I and the independent set $\hat{I}$ constructed via ApproxSampler is small.

8.1. The defect side of a uniformly random independent set. Let us begin with a lemma detailing some properties of the defect side of a uniformly random independent set; these statements are simply reinterpretations of various previous statements throughout the article.

Lemma 8.2. *Let I denote a uniformly random independent set in $Q_{d,p}$.*

- (1) *For $p > 0$,*

$$\mathbb{P}\left[I \cap \text{Even} \in \text{PD}^{\text{Even}} \text{ or } I \cap \text{Odd} \in \text{PD}^{\text{Odd}}, \text{ but not both}\right] \xrightarrow{\mathbb{P}} 1.$$

- (2) *For $p > 0$ and any $S \in \text{PD}^{\text{Even}}$ (uniformly),*

$$\mathbb{P}\left[I \cap \text{Even} = S \mid I \cap \text{Even} \in \text{PD}^{\text{Even}}\right] \stackrel{\mathbb{P}}{\sim} \frac{\prod_{\mathbf{p} \in S} \varphi_{\mathbf{p}}}{\mathcal{Z}^{\text{Even}}},$$

and similarly for Even replaced by Odd.

- (3) *For $p > \frac{1}{2} - \gamma$ (where $\gamma > 0$ is as in Proposition 6.1),*

$$\mathbb{P}\left[I \cap \text{Even} \in \text{PD}^{\text{Even}}\right] \stackrel{\mathbb{P}}{\sim} q^{\text{Even}},$$

and similarly for Even replaced by Odd.

Proof of Lemma 8.2. Recall from the discussion in Section 3.3 that $2^{2^{d-1}} \mathcal{Z}^{\text{Even}}$ is the number of independent sets in $Q_{d,p}$ whose even side is polymer-decomposable, and similarly for $2^{2^{d-1}} \mathcal{Z}^{\text{Odd}}$. Now Corollary 3.6

states that the number of independent sets that are *not accounted for* in

$$2^{2^{d-1}} \left(\mathcal{Z}^{\text{Even}} + \mathcal{Z}^{\text{Odd}} \right) \quad (8.2)$$

is an $o(1)$ -in-probability fraction of all independent sets. Further, by Proposition 3.4, which states that $i(Q_{d,p})$ is equal to (8.2) times a $(1 + o(1))$ -in-probability factor, we observe that the number of independent sets which are *double-counted* in (8.2) is also an $o(1)$ -in-probability fraction of all independent sets. In other words, all but an $o(1)$ -in-probability fraction of independent sets are polymer decomposable on exactly one of the two sides, which proves part (1).

Of these sets, there are $2^{2^{d-1}} \mathcal{Z}^{\text{Even}} (1 - o(1))$ which are decomposable on the even side and similarly for the odd side. And, as mentioned at the beginning of Section 2 (in particular see (2.1)), the number of independent sets which have even side S is $2^{2^{d-1} - N_p(S)}$, which proves part (2) for the even side; the odd case follows by the same reasoning.

Finally, by Corollary 2.6, for $p > \frac{1}{2} - \gamma$ we have

$$\mathcal{Z}^{\text{Even}} \stackrel{\mathbb{P}}{\sim} \tilde{\mathcal{Z}}^{\text{Even}} \cdot \exp\left(\Delta^{\text{Even}} - \tilde{\Delta}^{\text{Even}}\right),$$

and similarly for Odd, which yields part (3), recalling the definition (8.1) of q^{Even} and q^{Odd} . $\square$

Now let us establish that total variation closeness of defect sides is sufficient to yield total variation closeness of the full independent sets.

Lemma 8.3. *Let I denote a uniformly random independent set in $Q_{d,p}$ and let $\hat{I}$ denote the output of ApproxSampler, given that ApproxSampler does not fail. Let S denote the random subset of Even given by $I \cap \text{Even}$, conditioned on $I \cap \text{Even} \in \text{PD}$. Let $\hat{S}$ denote the random subset of Even given in step 5 of ApproxSampler, conditioned on choosing $\mathcal{H} = \text{Even}$ in step 1 (and on ApproxSampler not failing). If $\text{TV}(S, \hat{S}) \xrightarrow{\mathbb{P}} 0$ and the analogous statement holds for Odd, then $\text{TV}(I, \hat{I}) \xrightarrow{\mathbb{P}} 0$ as well.*

Proof of Lemma 8.3. We need to exhibit a coupling between I and $\hat{I}$ such that $\mathbb{P}[I \neq \hat{I}] \xrightarrow{\mathbb{P}} 0$. By Lemma 8.2 (1), we may couple the choice of $\mathcal{H}$ in step 1 of ApproxSampler to agree with the polymer-decomposable defect side of I with probability $\xrightarrow{\mathbb{P}} 1$ (we may simply ignore the event where both $I \cap \text{Even}$ and $I \cap \text{Odd}$ are polymer-decomposable, since this happens with probability $\xrightarrow{\mathbb{P}} 0$). The hypothesis of the lemma is that we may couple the defect sides of I and $\hat{I}$ with high probability, and given that I has defect side S , the distribution of the non-defect side is a uniformly random subset of the vertices which do not neighbor S . This is exactly what is achieved in step 6 of ApproxSampler, meaning that on the event that the defect sides S and $\hat{S}$ are coupled, the entire independent sets I and $\hat{I}$ may be coupled perfectly, finishing the proof. $\square$

Thus, to prove Theorem 1.3, it suffices to show that ApproxSampler fails with probability $\xrightarrow{\mathbb{P}} 0$, and that, given that ApproxSampler does not fail, the defect side constructed by the algorithm is close in total variation distance to the defect side of a uniformly random set. These are captured in the following two lemmas, which together finish the proof via an application of Lemma 8.3.

Lemma 8.4. *For all $p > \frac{1}{2} - \gamma$, where $\gamma > 0$ is defined as in Proposition 6.1,*

$$\mathbb{P}[\text{ApproxSampler fails}] \xrightarrow{\mathbb{P}} 0.$$

Lemma 8.5. *Let S denote $I \cap \text{Even}$, where I is a uniformly random independent set conditioned on the event $I \cap \text{Even} \in \text{PD}^{\text{Even}}$. Let $\hat{S}$ denote the result of step 5 in ApproxSampler, conditioned on choosing $\mathcal{H} = \text{Even}$ in step 1 and on the algorithm not failing at steps 3 or 5. For all $p > \frac{1}{2} - \gamma$, where $\gamma > 0$ is defined as in Proposition 6.1, we have*

$$\text{TV}(S, \hat{S}) \xrightarrow{\mathbb{P}} 0.$$

Additionally, the analogous statement with Even replaced by Odd also holds.

We will prove Lemma 8.4 in Section 8.2, and we will prove Lemma 8.5 in Section 8.3 below.

8.2. The algorithm rarely fails. Let us assume that in step 1, we chose $\mathcal{H} = \text{Even}$. So we will drop all superscripts of $\mathcal{H}$ to simplify the notation. The case where $\mathcal{H} = \text{Odd}$ follows by identical reasoning. For ApproxSampler to fail at all, it must fail at step 3 or step 5, so Lemma 8.4 follows from the following pair of lemmas, which we prove in turn.

Lemma 8.6. *For all $p > 2 - 2^{2/3} \approx 0.413$,*

$$\mathbb{P}[\text{ApproxSampler fails at step 3}] \xrightarrow{\mathbb{P}} 0.$$

Lemma 8.7. *For all $p > \frac{1}{2} - \gamma$, where $\gamma > 0$ is defined as in Proposition 6.1,*

$$\mathbb{P}[\text{ApproxSampler fails at step 5} | \text{ApproxSampler does not fail at step 3}] \xrightarrow{\mathbb{P}} 0$$

Proof of Lemma 8.6. The probability that ApproxSampler fails at step 3 is the same as the probability that a sample from π is not in $\text{PD}_{\leq 2}$. Recalling the definition of π from (6.1) as well as the definitions of $\tilde{\mathcal{Z}}$ and $\tilde{\mathcal{Z}}_{\leq 2}$ from (2.7) and (2.8) respectively, we see that

$$\mathbb{P}[\text{ApproxSampler fails at step 3}] = \frac{\tilde{\mathcal{Z}} - \tilde{\mathcal{Z}}_{\leq 2}}{\tilde{\mathcal{Z}}} \xrightarrow{\mathbb{P}} 0, \quad (8.3)$$

invoking Proposition 4.2 which states that $\tilde{\mathcal{Z}} \sim \tilde{\mathcal{Z}}_{\leq 2}$ for all $p > 2 - 2^{2/3}$. $\square$

Proof of Lemma 8.7. ApproxSampler could fail in two different ways at step 5: $S_1 \cup S_2 \notin \text{PD}_{\leq 2}$ only if either $S_2 \notin \text{PD}_2$ or $S_2 \not\subseteq \text{Compat}_2(S_1)$.

Let us consider the case $S_2 \notin \text{PD}_2$ first. By a union bound,

$$\mathbb{P}[S_2 \notin \text{PD}_2] \leq \sum_{\mathbf{d}_1 \sim \mathbf{d}_2} \varphi_{\mathbf{d}_1} \varphi_{\mathbf{d}_2}; \quad (8.4)$$

here we have also used the bound $\frac{\varphi_{\mathbf{d}}}{1+\varphi_{\mathbf{d}}} \leq \varphi_{\mathbf{d}}$. Now, by a simple moment calculation done in Lemma A.4 (b) below, the right-hand side in (8.4) has expectation (in terms of the percolation configuration) tending to 0 for $p > 0.391$; this shows immediately that

$$\mathbb{P}[S_2 \notin \text{PD}_2] \xrightarrow{\mathbb{P}} 0. \quad (8.5)$$

Let us next consider the failure via $S_2 \not\subseteq \text{Compat}_2(S_1)$. Using the bound $\frac{\varphi_{\mathbf{d}}}{1+\varphi_{\mathbf{d}}} \leq \varphi_{\mathbf{d}}$ again, as well as another union bound, conditioning on the value of the singleton set S_1 constructed in step 3, we have

$$\mathbb{P}[S_2 \not\subseteq \text{Compat}_2(S_1) | \text{step 3 yields } S_1] \leq \sum_{\mathbf{d} \sim S_1} \varphi_{\mathbf{d}} = \text{Adj}(S_1), \quad (8.6)$$

recalling the notation Adj from Definition 6.2. Since probabilities are at most 1, we may replace this upper bound by $\min\{1, \text{Adj}(S_1)\}$, which will be important in the sequel for turning smallness in probability into smallness in expectation.

The upper bound of $\text{Adj}(S_1)$ given in (8.6) is actually *not* small under the worst-case choice of S_1 , as first alluded to in Remark 2.5 and discussed in further detail in Section 6.4. Instead, we need to use our average-case analysis and leverage the fact that $\text{Adj}(S_1)$ is small with high probability for $S_1 \sim \pi$ conditioned to be in PD_1 , as discussed in Section 6.5; in particular, we will make use of (6.16), which is a consequence of Lemmas 6.5 and 6.6 and which states that, for all $p > \frac{1}{2} - \gamma$,

$$\pi[\text{Adj}(S_1) > \varepsilon | S_1 \in \text{PD}_1] \xrightarrow{\mathbb{P}} 0 \quad (8.7)$$

for some choice of $\varepsilon = \varepsilon_d \rightarrow 0$ as $d \rightarrow \infty$. Note that the same statement holds for $\widetilde{\text{Adj}}$ as well, which is the same as Adj but with each φ_d replaced by $\widetilde{\varphi}_d$; this is simply because $\widetilde{\varphi}_d \leq \varphi_d$. These facts will also be relevant in Section 8.3 below, where we prove Lemma 8.7, showing that the defect side generated by the sampler is close to the defect side of a uniform independent set.

To make use of (8.7), we will need to understand the distribution of the set S_1 constructed in step 3 of ApproxSampler. Since ApproxSampler fails if the set $\widetilde{S} \sim \pi$ is not in $\text{PD}_{\leq 2}$, in the case where the algorithm does not fail, $\widetilde{S}$ is conditioned to be in $\text{PD}_{\leq 2}$, meaning that the probability of a particular value of S_1 constructed in step 3 is

$$\mathbb{P}[\text{step 3 yields } S_1 | \text{step 3 does not fail}] = \frac{1}{\widetilde{\mathcal{Z}}_{\leq 2}} \prod_{v \in S_1} \varphi_v \sum_{\substack{S_2 \in \text{PD}_2 \\ S_2 \subseteq \text{Compat}_2(S_1)}} \prod_{d \in S_2} \widetilde{\varphi}_d.$$

By Lemma 6.4, the inner sum over S_2 is approximated in probability by $e^{\widetilde{\Delta} - \widetilde{\text{Adj}}(S_1)}$, and this approximation is *uniform* in the choice of S_1 . So

$$\mathbb{P}[\text{step 3 yields } S_1 | \text{step 3 does not fail}] \stackrel{\mathbb{P}}{\sim} \frac{\prod_{v \in S_1} \varphi_v}{\widetilde{\mathcal{Z}}_{\leq 2}} \cdot \exp\left(\widetilde{\Delta} - \widetilde{\text{Adj}}(S_1)\right).$$

Now, by Proposition 6.1 which states that $\widetilde{\mathcal{Z}}_{\leq 2} \stackrel{\mathbb{P}}{\sim} \mathcal{Z}_1 \cdot e^{\widetilde{\Delta}}$ for $p > \frac{1}{2} - \gamma$, we have

$$\mathbb{P}[\text{step 3 yields } S_1 | \text{step 3 does not fail}] \stackrel{\mathbb{P}}{\sim} \frac{\prod_{v \in S_1} \varphi_v}{\mathcal{Z}_1} \cdot e^{-\widetilde{\text{Adj}}(S_1)} \quad (8.8)$$

uniformly for any $S_1 \in \text{PD}_1$ (sets $S_1 \notin \text{PD}_1$ are not possible at this stage). Combining (8.8) with (8.6) and recalling the definition of π from (6.1), we obtain

$$\begin{aligned} \mathbb{P}[S_2 \not\subseteq \text{Compat}_2(S_1) | \text{step 3 does not fail}] &= \mathbb{E}[\mathbb{P}[S_2 \not\subseteq \text{Compat}_2(S_1) | \text{step 3 yields } S_1]] \\ &\leq \mathbb{E}[\min\{1, \text{Adj}(S_1)\}] \\ &\stackrel{\mathbb{P}}{\sim} \sum_{S_1 \in \text{PD}_1} \frac{\prod_{v \in S_1} \varphi_v}{\mathcal{Z}_1} \cdot e^{-\widetilde{\text{Adj}}(S_1)} \cdot \min\{1, \text{Adj}(S_1)\} \\ &\leq \mathbb{E}_{S_1 \sim \pi}[\min\{1, \text{Adj}(S_1)\} | S_1 \in \text{PD}_1], \end{aligned}$$

where the last inequality follows simply from the fact that $\widetilde{\text{Adj}}(S_1) \geq 0$. Now, using (8.7), we find that the right-hand side converges to 0 in probability. Combining this with (8.5) and (8.3) finishes the proof. $\square$

8.3. The defect distributions are close. Now we turn to the proof of Lemma 8.5, showing that the defect side from the sampler is close to the defect side of a uniform independent set, in total variation distance. Recall that in the context of that lemma, S denotes $I \cap \text{Even}$, where I is a uniformly random independent set conditioned on the event $I \cap \text{Even} \in \text{PD}^{\text{Even}}$. Additionally, $\hat{S}$ denote the result of step 5 in ApproxSampler, conditioned on choosing $\mathcal{H} = \text{Even}$ in step 1 and on the algorithm not failing at steps 3 or 5. We only focus on the Even case in this section as the Odd case is identical, and henceforth we drop all Even superscripts.

Proof of Lemma 8.5. Given that ApproxSampler does not fail, the defect side $\hat{S}$ obtained in step 5 is guaranteed to be in $\text{PD}_{\leq 2}$, meaning that it can be decomposed as $S_1 \in \text{PD}_1$ and $S_2 \in \text{PD}_2$ with $S_2 \subseteq \text{Compat}_2(S_1)$. Additionally, since S is guaranteed to be in PD by our conditioning assumption, we have

$$\text{TV}(S, \hat{S}) = \sum_{S_1 \in \text{PD}_1} \sum_{\substack{S_2 \in \text{PD}_2 \\ S_2 \subseteq \text{Compat}_2(S_1)}} \left| \mathbb{P}[S = S_1 \cup S_2] - \mathbb{P}[\hat{S} = S_1 \cup S_2] \right| + \sum_{S' \in \text{PD} \setminus \text{PD}_{\leq 2}} \mathbb{P}[S = S']. \quad (8.9)$$

By Lemma 8.2 (2), for any $S' \in \text{PD}$,

$$\mathbb{P}[S = S'] \stackrel{\mathbb{P}}{\sim} \frac{\prod_{\mathbf{p} \in S'} \varphi_{\mathbf{p}}}{\mathcal{Z}}, \quad (8.10)$$

and this approximation is uniform among $S' \in \text{PD}$. This implies that the second sum in (8.9) is

$$\sum_{S' \in \text{PD} \setminus \text{PD}_{\leq 2}} \mathbb{P}[S = S'] \stackrel{\mathbb{P}}{\sim} \frac{\mathcal{Z} - \mathcal{Z}_{\leq 2}}{\mathcal{Z}} \xrightarrow{\mathbb{P}} 0, \quad (8.11)$$

recalling the definitions of $\mathcal{Z}$ and $\mathcal{Z}_{\leq 2}$ from (2.5) and (2.8) and invoking Proposition 4.1 which states that $\mathcal{Z} \stackrel{\mathbb{P}}{\sim} \mathcal{Z}_{\leq 2}$ for $p > 2 - 2^{2/3}$.

Now let us turn to the first sum. For S the defect side of the uniform independent set, (8.10) yields

$$\mathbb{P}[S = S_1 \cup S_2] \stackrel{\mathbb{P}}{\sim} \frac{\prod_{v \in S_1} \varphi_v \prod_{\mathbf{d} \in S_2} \varphi_{\mathbf{d}}}{\mathcal{Z}} \stackrel{\mathbb{P}}{\sim} \frac{\prod_{v \in S_1} \varphi_v \prod_{\mathbf{d} \in S_2} \varphi_{\mathbf{d}}}{\mathcal{Z}_1 \cdot e^{\Delta}}, \quad (8.12)$$

where we have invoked Proposition 4.1 (which states that $\mathcal{Z} \stackrel{\mathbb{P}}{\sim} \mathcal{Z}_{\leq 2}$) and Proposition 6.1 (which states that $\mathcal{Z}_{\leq 2} \stackrel{\mathbb{P}}{\sim} \mathcal{Z}_1 \cdot e^{\Delta}$) to approximate $\mathcal{Z}$ by $\mathcal{Z}_1 \cdot e^{\Delta}$. It's crucial to note that all of these approximations are uniform in S_1 and S_2 , in the sense that the error terms are bounded by quantities tending (in probability) to 0 which do not depend on S_1 or S_2 .

Now let us examine the corresponding probability for $\hat{S}$ constructed in step 5 of ApproxSampler. Recall from (8.8) that

$$\mathbb{P}[\text{step 3 yields } S_1 | \text{step 3 does not fail}] \stackrel{\mathbb{P}}{\sim} \frac{\prod_{v \in S_1} \varphi_v}{\mathcal{Z}_1} \cdot e^{-\widetilde{\text{Adj}}(S_1)}. \quad (8.13)$$

Now, given that we obtain S_1 in step 3 and that ApproxSampler does not fail at step 5, the probability of obtaining a particular value of S_2 in that step (which must satisfy $S_2 \in \text{PD}_2$ and $S_2 \subseteq \text{Compat}_2(S_1)$) is

$$\mathbb{P}[\text{step 5 yields } S_2 | \text{step 3 yields } S_1] = \frac{\prod_{\mathbf{d} \in S_2} \varphi_{\mathbf{d}}}{\sum_{\substack{S'_2 \in \text{PD}_2 \\ S'_2 \subseteq \text{Compat}_2(S_1)}} \prod_{\mathbf{d} \in S'_2} \varphi_{\mathbf{d}}} \stackrel{\mathbb{P}}{\sim} \frac{\prod_{\mathbf{d} \in S_2} \varphi_{\mathbf{d}}}{e^{\Delta}} e^{\text{Adj}(S_1)}, \quad (8.14)$$

where we have applied Lemma 6.4 which approximates the big sum in the denominator by $e^{\Delta - \text{Adj}(S_1)}$ and the approximation here is again uniform in S_1 .

Combining (8.12), (8.13) and (8.14), we find that the first sum in (8.9) satisfies

$$\begin{aligned}
& \sum_{S_1 \in \text{PD}_1} \sum_{\substack{S_2 \in \text{PD}_2 \\ S_2 \subseteq \text{Compat}_2(S_1)}} \left| \mathbb{P}[S = S_1 \cup S_2] - \mathbb{P}[\widehat{S} = S_1 \cup S_2] \right| \\
& \quad \stackrel{\mathbb{P}}{\sim} \sum_{S_1 \in \text{PD}_1} \sum_{\substack{S_2 \in \text{PD}_2 \\ S_2 \subseteq \text{Compat}_2(S_1)}} \frac{\prod_{v \in S_1} \varphi_v \prod_{\mathbf{d} \in S_2} \varphi_{\mathbf{d}}}{\mathcal{Z}_1 \cdot e^{\Delta}} \left| 1 - e^{\text{Adj}(S_1) - \widetilde{\text{Adj}}(S_1)} \right| \\
& \quad = \sum_{S_1 \in \text{PD}_1} \frac{\prod_{v \in S_1} \varphi_v}{\mathcal{Z}_1} \left| 1 - e^{\text{Adj}(S_1) - \widetilde{\text{Adj}}(S_1)} \right| \cdot \underbrace{\frac{\sum_{\substack{S_2 \in \text{PD}_2 \\ S_2 \subseteq \text{Compat}_2(S_1)}} \prod_{\mathbf{d} \in S_2} \varphi_{\mathbf{d}}}{e^{\Delta}}}_{(*)}
\end{aligned}$$

To handle the factor labeled $(*)$, we apply Lemma 6.4, which states that

$$\sum_{\substack{S_2 \in \text{PD}_2 \\ S_2 \subseteq \text{Compat}_2(S_1)}} \prod_{\mathbf{d} \in S_2} \varphi_{\mathbf{d}} \stackrel{\mathbb{P}}{\sim} \exp \left(\sum_{\mathbf{d} \not\sim S_1} \varphi_{\mathbf{d}} \right) = \exp(\Delta - \text{Adj}(S_1)),$$

where the approximation is uniform in S_1 . So, continuing from the previous display, the first sum in (8.9) is

$$\stackrel{\mathbb{P}}{\sim} \sum_{S_1 \in \text{PD}_1} \frac{\prod_{v \in S_1} \varphi_v}{\mathcal{Z}_1} \left| 1 - e^{\text{Adj}(S_1) - \widetilde{\text{Adj}}(S_1)} \right| e^{-\text{Adj}(S_1)} = \mathbb{E}_{\pi} \left[\left| e^{-\text{Adj}(S_1)} - e^{-\widetilde{\text{Adj}}(S_1)} \right| \middle| S_1 \in \text{PD}_1 \right].$$

Now, since $\text{Adj}(S_1), \widetilde{\text{Adj}}(S_1) \geq 0$ and the derivative of $t \mapsto e^{-t}$ is bounded by 1 for $t \geq 0$, the right-hand side above is at most

$$\mathbb{E}_{\pi} \left[\left| \text{Adj}(S_1) - \widetilde{\text{Adj}}(S_1) \right| \middle| S_1 \in \text{PD}_1 \right],$$

which tends to 0 in probability by (8.7) as well as the corresponding fact for $\widetilde{\text{Adj}}$. In combination with (8.11), this finishes the proof of the lemma. $\square$

APPENDIX A. MOMENT CALCULATIONS

Recall from (2.4) and (2.6) that for each polymer $\mathbf{p}$, we have defined $\varphi_{\mathbf{p}} = 2^{-N_p(\mathbf{p})}$, and $\widetilde{\varphi}_{\mathbf{p}} = \prod_{v \in \mathbf{p}} 2^{-N_p(v)}$. Additionally we introduce the notation

$$\Phi_k = \sum_{v \in \mathcal{H}} \varphi_v^k$$

for each $k \geq 0$, and we recall from (5.2) and (2.11) the notation

$$\Phi_{\log} = \sum_{v \in \mathcal{H}} \log(1 + \varphi_v), \quad \Delta = \sum_{\mathbf{d} \in \text{Dimers}^{\mathcal{H}}} \varphi_{\mathbf{d}} \quad \text{and} \quad \widetilde{\Delta} = \sum_{\mathbf{d} \in \text{Dimers}^{\mathcal{H}}} \widetilde{\varphi}_{\mathbf{d}}.$$

Note that these quantities depend on the side $\mathcal{H} \in \{\text{Even}, \text{Odd}\}$, but for most of the following moment calculations, the result will be identical for both sides and for notational convenience we will not mention the side explicitly. The exception will be the covariance of $\Phi_{\log}$ for the even and odd side, and there we will use the notation $\Phi_{\log}^{\text{Even}}$ and $\Phi_{\log}^{\text{Odd}}$ to distinguish between the two sides.

A.1. Singleton moments. Recall the following constants defined in (7.3) which turn out to be the expectations of Φ_k , or equivalently 2^{d-1} times the k th moments of φ_v .

$$\mu_1^k = \frac{1}{2} \left(2 - \frac{2^k - 1}{2^{k-1}} p \right)^d.$$

We will also use the notation

$$\sigma^2 = \mu_1^2 = \frac{1}{2} \left(2 - \frac{3}{2} p \right)^d \quad \text{and} \quad \sigma = \sqrt{\sigma^2}.$$

We first collect the following moments of singleton variables and sums of singleton variables.

Lemma A.1. *The following facts hold for all $p \in (0, 1)$.*

(a) *For each v in $\mathcal{H}$ and each $k \geq 0$, we have*

$$\begin{aligned} \mathbb{E}[\varphi_v^k] &= \left(1 - \frac{2^k - 1}{2^k} p \right)^d, \\ \mathbf{Var}(\varphi_v^k) &= \mathbb{E}[\varphi_v^{2k}] (1 - o(1)). \end{aligned}$$

(b) *For each $k \geq 0$, we have*

$$\begin{aligned} \mathbb{E}[\Phi_k] &= \mu_1^k, \\ \mathbf{Var}(\Phi_k) &= \mu_1^{2k} (1 - o(1)). \end{aligned}$$

Proof of Lemma A.1.

(a) Since $N_p(v)$ is a sum of d independent $\text{Ber}(p)$ variables,

$$\mathbb{E}[\varphi_v^k] = \left((1 - p) + 2^{-k} p \right)^d = \left(1 - \frac{2^k - 1}{2^k} p \right)^d.$$

For the second claim, it suffices to show that $\mathbb{E}[\varphi_v^k]^2 \ll \mathbb{E}[\varphi_v^{2k}]$, or in other words, that

$$\left(1 - \frac{2^k - 1}{2^k} p \right)^2 < 1 - \frac{2^{2k} - 1}{2^{2k}} p.$$

The validity of this inequality for $0 < p < 1$ can be checked by multiplying both sides by 2^{2k} and expanding the square; this gives

$$\begin{aligned} 2^{2k} - 2 \cdot 2^k (2^k - 1) p + (2^k - 1)^2 p^2 &< 2^{2k} - (2^{2k} - 1) p \\ -2 \cdot 2^k (2^k - 1) + (2^k - 1)^2 p &< -(2^k + 1)(2^k - 1) \\ (2^k - 1)^2 p &< (2^k - 1)(2 \cdot 2^k - (2^k + 1)) \\ (2^k - 1) p &< 2 \cdot 2^k - (2^k + 1) = 2^k - 1, \end{aligned}$$

which holds as $p < 1$.

(b) These claims follow from the previous part, using the fact that the number of vertices in each side is 2^{d-1} . The claim about the variance also uses the independence of φ_v amongst $v \in \mathcal{H}$. $\square$

Next we obtain approximations for the mean and variance of $\Phi_{\log}$ which are used in Section 7.

Lemma A.2. *The following approximations hold for $\Phi_{\log}$.*

(a) *For all $p \in (0, 1)$, we have*

$$\mathbf{Var}(\Phi_{\log}) = \sigma^2(1 + o(1))$$

(b) *For $p > 0.455$, we have*

$$\mathbb{E}[\Phi_{\log}] = \mu_1^1 - \frac{\mu_1^2}{2} + \frac{\mu_1^3}{3} + o(\sigma)$$

Proof of Lemma A.2.

(a) Since $x \geq \log(1+x) \geq x - \frac{x^2}{2}$ for all $x \in [0, 1]$, using Lemma A.1 (a), we have

$$\left(1 - \frac{1}{2}p\right)^d \geq \mathbb{E}[\log(1 + \varphi_v)] \geq \left(1 - \frac{1}{2}p\right)^d - \frac{1}{2}\left(1 - \frac{3}{4}p\right)^d,$$

and since $x^2 \geq \log^2(1 + \varphi_v) \geq x^2 - x^3 + \frac{1}{4}x^4$, we also have

$$\left(1 - \frac{3}{4}p\right)^d \geq \mathbb{E}[\log^2(1 + \varphi_v)] \geq \left(1 - \frac{3}{4}p\right)^d - \left(1 - \frac{7}{8}p\right)^d + \frac{1}{4}\left(1 - \frac{15}{16}p\right)^d.$$

Combining these, we obtain

$$\begin{aligned} & \left(1 - \frac{3}{4}p\right)^d - \left(1 - \frac{1}{2}p\right)^{2d} + \left(1 - \frac{1}{2}p\right)^d \left(1 - \frac{3}{4}p\right)^d - \frac{1}{4}\left(1 - \frac{3}{4}p\right)^{2d} \\ & \geq \mathbf{Var}(\log(1 + \varphi_v)) \geq \end{aligned} \quad (\text{A.1})$$

$$\left(1 - \frac{3}{4}p\right)^d - \left(1 - \frac{7}{8}p\right)^d + \frac{1}{4}\left(1 - \frac{15}{16}p\right)^d - \left(1 - \frac{1}{2}p\right)^{2d}. \quad (\text{A.2})$$

Now for both the upper bound (A.1) and the lower bound (A.2), each term after the first is exponentially smaller than the first term. This is clear by inspection for all terms other than $(1 - \frac{1}{2}p)^{2d}$, and the fact that this term is also exponentially smaller than $(1 - \frac{3}{4}p)^d$ follows from the same argument as in Lemma A.1 above. This means that

$$\mathbf{Var}(\log(1 + \varphi_v)) = \left(1 - \frac{3}{4}p\right)^d (1 + o(1)),$$

and so since $\Phi_{\log}$ is a sum of 2^{d-1} independent random variables of the form $\log(1 + \varphi_v)$, we have

$$\mathbf{Var}(\Phi_{\log}) = \sigma^2(1 + o(1)),$$

since $\sigma^2 = 2^{d-1}\left(1 - \frac{3}{4}p\right)^d$.

(b) This time we use the fact that

$$x - \frac{x^2}{2} + \frac{x^3}{3} \geq \log(1+x) \geq x - \frac{x^2}{2} + \frac{x^3}{3} - \frac{x^4}{4}$$

for $x \geq 0$. This gives

$$0 \geq \mathbb{E}[\Phi_{\log}] - \left(\mu_1^1 - \frac{\mu_1^2}{2} + \frac{\mu_1^3}{3}\right) \geq -\frac{\mu_1^4}{4}$$

We thus want to show that $\mu_1^4 \ll \sigma$, which is implied by

$$\left(2 - \frac{15}{8}p\right)^2 < \left(2 - \frac{3}{2}p\right).$$

The difference between the left-hand side and the right-hand side is

$$2 - 6p + \frac{225}{64}p^2,$$

which has roots

$$\frac{6 \pm \sqrt{36 - 225/8}}{225/32} \approx 0.4542, 1.2524.$$

Therefore, for $0.455 < p < 1$, we have $\mu_1^4 \ll \sigma$, finishing the proof. $\square$

The following lemma gives the independence input for Lemma 7.1 allowing us to show that the scaling limit of $(\Phi_{\log}^{\text{Even}}, \Phi_{\log}^{\text{Odd}})$ is a pair of independent Gaussians.

Lemma A.3. *For all $p \in (0, 1)$,*

$$\mathbf{Cov}(\Phi_{\log}^{\text{Even}}, \Phi_{\log}^{\text{Odd}}) \ll \mathbf{Var}(\Phi_{\log}).$$

Proof of Lemma A.3. First note that $\mathbf{Cov}(\Phi_{\log}^{\text{Even}}, \Phi_{\log}^{\text{Odd}}) \geq 0$ by the FKG inequality since both random variables are decreasing in the percolation configuration. Now, since φ_v and φ_w are independent when $v \not\sim w$ in Q_d , we have

$$\begin{aligned} \mathbf{Cov}(\Phi_{\log}^{\text{Even}}, \Phi_{\log}^{\text{Odd}}) &= \sum_{\substack{v \in \text{Even}, u \in \text{Odd} \\ v \sim u}} (\mathbb{E}[\log(1 + \varphi_v) \log(1 + \varphi_u)] - \mathbb{E}[\log(1 + \varphi_v)] \mathbb{E}[\log(1 + \varphi_u)]) \\ &\leq \sum_{\substack{v \in \text{Even}, u \in \text{Odd} \\ v \sim u}} \mathbb{E}[\log(1 + \varphi_v) \log(1 + \varphi_u)] \\ &\leq \sum_{\substack{v \in \text{Even}, u \in \text{Odd} \\ v \sim u}} \mathbb{E}[\varphi_v \varphi_u], \end{aligned}$$

using the inequality $\log(1 + x) \leq x$ for $x \in [0, 1]$ at the last step. Now, when $v \sim w$, there is exactly one edge connecting them in Q_d , and there are $2d - 2$ edges which are incident to only one of the two vertices. So $\varphi_v \varphi_w \sim 2^{-2 \cdot \text{Ber}(p)} 2^{-\text{Bin}(2d-2, p)}$, where the two random variables in the exponent are independent. Additionally, for each of the 2^{d-1} vertices $v \in \text{Even}$, there are d vertices $u \in \text{Odd}$ such that $v \sim u$, and so we have

$$\mathbf{Cov}(\Phi_{\log}^{\text{Even}}, \Phi_{\log}^{\text{Odd}}) = 2^{d-1} d \left(1 - \frac{3}{4}p\right) \left(1 - \frac{1}{2}p\right)^{2d-2},$$

which is $\ll \sigma^2$ since

$$2 \cdot \left(1 - \frac{1}{2}p\right)^2 = 2 - 2p + \frac{1}{2}p^2 < 2 - 2p + \frac{1}{2}p = 2 - \frac{3}{2}p.$$

The conclusion thus follows by Lemma A.2 (a). $\square$

A.2. Dimer moments. Recall the following constants defined in (1.9) and (1.10), which turn out to be the expectations of Δ and $\tilde{\Delta}$ respectively:

$$\mu_2 = \frac{d(d-1)}{4} \left(\frac{(2-p)^2}{2} \right)^d \cdot \left(\frac{1 + (1-p)^2}{2-p} \right)^2,$$

$$\tilde{\mu}_2 = \frac{d(d-1)}{4} \left(\frac{(2-p)^2}{2} \right)^d.$$

Notice that the above two expressions differ only by a multiple which does not depend on d . Additionally, the exponential orders of μ_2 and $\tilde{\mu}_2$ are the same as that of $\frac{1}{2d}(\mu_1^1)^2$. Intuitively, this is because each $\varphi_{\mathbf{d}}$ and $\tilde{\varphi}_{\mathbf{d}}$ behaves approximately like a product of two independent singleton variables φ_v , since most of the neighbors of $\mathbf{d}$ are not neighbors of both vertices in $\mathbf{d}$. Moreover, there are the same number of singletons and dimers, up to a $\text{poly}(d)$ factor.

Lemma A.4. *Recall the definitions of Δ and $\tilde{\Delta}$ from (2.11) above.*

(a) *For all $p \in (0, 1)$, we have*

$$\mathbb{E}[\Delta] = \mu_2 \quad \text{and} \quad \mathbb{E}[\tilde{\Delta}] = \tilde{\mu}_2.$$

(b) *For all $p > 0.391$, we have*

$$\text{Var}(\Delta) \ll 1 \quad \text{and} \quad \text{Var}(\tilde{\Delta}) \ll 1.$$

Additionally for $p > 0.391$, we have

$$\mathbb{E} \left[\sum_{\mathbf{d}_1 \sim \mathbf{d}_2} \varphi_{\mathbf{d}_1} \varphi_{\mathbf{d}_2} \right] \ll 1 \quad \text{and} \quad \mathbb{E} \left[\sum_{\mathbf{d}_1 \sim \mathbf{d}_2} \tilde{\varphi}_{\mathbf{d}_1} \tilde{\varphi}_{\mathbf{d}_2} \right] \ll 1,$$

where $\mathbf{d}_1 \sim \mathbf{d}_2$ if the neighborhoods of the two dimers overlap.

Proof of Lemma A.4.

(a) Let us consider $\tilde{\Delta}$ first. For $\mathbf{d} = \{u, v\}$, we have

$$\tilde{\varphi}_{\mathbf{d}} = 2^{-N_p(u) - N_p(v)} \stackrel{d}{=} 2^{-\text{Bin}(2d, p)},$$

and so by the same reasoning as in Lemma A.1 (a), we have

$$\mathbb{E}[\tilde{\varphi}_{\mathbf{d}}] = \left(1 - \frac{1}{2}p \right)^{2d}.$$

Now there are $2^{d-1} \frac{d(d-1)}{2}$ possible dimers $\mathbf{d} \in \text{Dimers}^{\mathcal{H}}$, and so we have

$$\mathbb{E}[\tilde{\Delta}] = 2^{d-1} \frac{d(d-1)}{2} \left(1 - \frac{1}{2}p \right)^{2d} = \tilde{\mu}_2.$$

As for Δ , we must consider $\varphi_{\mathbf{d}}$ instead of $\tilde{\varphi}_{\mathbf{d}}$. Now for $\mathbf{d} = \{u, v\}$, we know that u and v share a neighbor; in this case, in fact, u and v share exactly two neighbors since their binary string representations differ in exactly two bits, and swapping either one of these bits leads to a shared neighbor. Each of these common neighbors is a neighbor of $\mathbf{d}$ in $Q_{d,p}$ with probability $1 - (1-p)^2$, since it is not a neighbor exactly when both relevant edges are absent. There are also $2d - 2$ vertices

in the opposite side which are neighbors of exactly one of u or v . So we have

$$\varphi_{\mathbf{d}} \stackrel{d}{=} 2^{-\text{Bin}(2, 1-(1-p)^2) - \text{Bin}(2d-2, p)},$$

where the two Binomials are independent. So we have

$$\begin{aligned} \mathbb{E}[\varphi_{\mathbf{d}}] &= \left((1-p)^2 + \frac{1}{2}(1 - (1-p)^2) \right)^2 \left(1 - \frac{1}{2}p \right)^{2d-2} \\ &= \left(\frac{\frac{1}{2} + \frac{1}{2}(1-p)^2}{1 - \frac{1}{2}p} \right)^2 \left(1 - \frac{1}{2}p \right)^{2d} \\ &= \left(\frac{1 + (1-p)^2}{2-p} \right)^2 \left(1 - \frac{1}{2}p \right)^{2d}. \end{aligned}$$

This differs from $\mathbb{E}[\tilde{\varphi}_{\mathbf{d}}]$ by the same factor as μ_2 differs from $\tilde{\mu}_2$, and so we obtain $\mathbb{E}[\Delta] = \mu_2$.

- (b) First note that $\varphi_{\mathbf{d}_1}$ is independent of $\varphi_{\mathbf{d}_2}$ if $\mathbf{d}_1 \cap \mathbf{d}_2 = \emptyset$, and for each $\mathbf{d}_1$ there are at most $2d^2$ dimers $\mathbf{d}_2$ for which this does not hold. So we have

$$\begin{aligned} \mathbf{Var}(\Delta) &\leq \sum_{\substack{\mathbf{d}_1, \mathbf{d}_2 \in \text{Dimers} \\ \mathbf{d}_1 \cap \mathbf{d}_2 \neq \emptyset}} \mathbb{E}[\varphi_{\mathbf{d}_1} \varphi_{\mathbf{d}_2}] \\ &\leq \sum_{\substack{\mathbf{d}_1, \mathbf{d}_2 \in \text{Dimers} \\ \mathbf{d}_1 \cap \mathbf{d}_2 \neq \emptyset}} \sqrt{\mathbb{E}[\varphi_{\mathbf{d}_1}^2] \mathbb{E}[\varphi_{\mathbf{d}_2}^2]} \\ &\leq \text{poly}(d) \cdot 2^d \cdot \mathbb{E}[\varphi_{\mathbf{d}}^2]. \end{aligned}$$

Now using the expression for $\varphi_{\mathbf{d}}$ from the previous part, we see that

$$\mathbb{E}[\varphi_{\mathbf{d}}^2] = c_p \cdot \mathbb{E}\left[2^{-2 \cdot \text{Bin}(2d, p)}\right]$$

for some constant c_p depending only on p . Now, as in Lemma A.1 (a), we have

$$\mathbb{E}\left[2^{-2 \cdot \text{Bin}(2d, p)}\right] = \left(1 - \frac{3}{4}p\right)^{2d}.$$

Thus

$$\mathbf{Var}(\Delta) \lesssim \text{poly}(d) \cdot 2^d \left(1 - \frac{3}{4}p\right)^{2d},$$

and this is $\ll 1$ when $\left(1 - \frac{3}{4}p\right)^2 < \frac{1}{2}$, which happens for $p > \frac{2}{3}(2 - \sqrt{2}) \approx 0.3905$. Additionally, $\tilde{\Delta}$ is a sum of i.i.d. random variables $\tilde{\varphi}_{\mathbf{d}} \leq \varphi_{\mathbf{d}}$, and so by the same reasoning we have $\mathbf{Var}(\tilde{\Delta}) \ll 1$ as well.

The second statement, about the sums over $\mathbf{d}_1 \sim \mathbf{d}_2$, follows by the same reasoning since for any dimer $\mathbf{d}_1$ there are at most $\text{poly}(d)$ dimers $\mathbf{d}_2$ which satisfy $\mathbf{d}_1 \sim \mathbf{d}_2$. $\square$

APPENDIX B. ENTROPY APPROXIMATIONS FOR BINOMIALS

Recall the binary (relative) entropy between two Bernoulli distributions with parameters p and q , defined in (6.12) above:

$$H_p(q) = q \log_2 \frac{q}{p} + (1-q) \log_2 \frac{1-q}{1-p}.$$

The following lemma is standard; see for instance [8, Page 353].

Lemma B.1. *For $p, q \in [0, 1]$ with $q \geq p$, we have $\mathbb{P}[\text{Bin}(n, p) \geq nq] \leq 2^{-n \cdot H_p(q)}$.*

Next, we need a lemma which bounds the binary relative entropy in terms of a more easily accessible quantity.

Lemma B.2. *For any $0 < p < x < 1$ with $x \geq 10p$, we have*

$$H_p(x) \geq \frac{1}{2} \cdot x \log_2(x/p)$$

We remark that the constant 10 is not the best possible constant here, although 5 seems to be closer to being optimal.

Proof of Lemma B.2. Recall from (6.12) that

$$H_p(x) = x \log_2 \frac{x}{p} + (1-x) \log_2 \frac{1-x}{1-p}.$$

Since $x > p$, $1-x < 1-p$, and thus the second summand is negative. Further, the ratio between the (negative) second summand and the first summand is

$$\begin{aligned} \frac{(1-x) \log_2 \frac{1-p}{1-x}}{x \log_2 \frac{x}{p}} &= \frac{(1-x) \log_2 \left(1 + \frac{x-p}{1-x}\right)}{x \log_2 \frac{x}{p}} \\ &= \frac{(1-x) \ln \left(1 + \frac{x-p}{1-x}\right)}{x \ln \frac{x}{p}} \\ &\leq \frac{(1-x) \left(\frac{x-p}{1-x}\right)}{x \ln \frac{x}{p}} \quad (\ln(1+t) \leq t) \\ &= \frac{1 - \frac{p}{x}}{\ln \frac{x}{p}} = \frac{\frac{p}{x} - 1}{\ln \frac{p}{x}}. \end{aligned}$$

Now since we assumed that $\frac{p}{x} \leq \frac{1}{10}$ and $\frac{t-1}{\ln t}$ is an increasing function of $t > 0$, the right-hand side above is at most $\frac{0.1-1}{\ln 0.1} \approx 0.3908 < \frac{1}{2}$, finishing the proof. $\square$

The following lemma is used repeatedly in Section 6 to show that various optimization problems have valid solutions for a range of p which includes values slightly below $\frac{1}{2}$.

Lemma B.3. *For every fixed $m \geq 1$, let*

$$f(p) := \inf_{s \in (0,1)} ms + H_p(s), \quad p \in (0, 1).$$

Then f is continuous and strictly increasing in p , and

$$f\left(\frac{1}{2}\right) = m + 1 - \log_2(2^m + 1).$$

Proof of Lemma B.3. Let us write $g(p, s) = ms + H_p(s)$ for the objective function, and $s^*(p)$ for the minimizer. Then, to compute $s^* = s^*(p)$, note that g is convex in s , and thus it suffices to equate its

derivative to zero. The derivative of $H_p(x)$ is $\log_2(x/p) - \log_2((1-x)/(1-p))$. Setting $\partial g/\partial s$ to zero, we get

$$\begin{aligned} m + \left(\log_2 \frac{s^*}{p} - \log_2 \frac{1-s^*}{1-p} \right) = 0 &\iff \frac{s^*}{1-s^*} = 2^{-m} \frac{p}{1-p} \\ &\iff s^* = \frac{p}{(1-p)2^m + p}. \end{aligned}$$

In particular, $s^* < p$ always. Observe that continuity of f follows from the continuity of H_p and $s^*(p)$. The derivative of f is

$$f'(p) = \frac{\partial g}{\partial p} \Big|_{s=s^*(p)} + \frac{\partial g}{\partial s} \Big|_{s=s^*(p)} \cdot (s^*)'(p)$$

The second term is zero due to the optimality of s^* . Further, since $\partial H_p(x)/\partial p$ is positive when $x < p$, we have that $f'(p) > 0$ always, showing that f is strictly increasing.

Setting $p = 1/2$, we get $s^* = 1/(2^m + 1)$, and thus

$$\begin{aligned} f(1/2) &= g(1/2, s^*(1/2)) \\ &= \frac{m}{2^m + 1} + \frac{1}{2^m + 1} \log_2 \left(\frac{1}{1/2 \cdot (2^m + 1)} \right) + \frac{2^m}{2^m + 1} \log_2 \left(\frac{2^m}{1/2 \cdot (2^m + 1)} \right) \\ &= m + 1 - \log_2(2^m + 1), \end{aligned}$$

as required. $\square$

Finally, the following lemma expresses the exponential quantity appearing in the dimer moment calculations of Lemma A.4 as the solution of an optimization problem.

Lemma B.4. *For all $p \in (0, 1)$,*

$$\frac{1}{\text{poly}(d)} \cdot 2^{d(1-2s^*-2H_p(s^*))} \leq \left(\frac{(2-p)^2}{2} \right)^d \leq \text{poly}(d) \cdot 2^{d(1-2s^*-2H_p(s^*))},$$

where s^* maximizes $1 - 2s - 2H_p(s)$ over $s \in [0, 1]$. In particular, this implies that

$$\frac{(2-p)^2}{2} = 2^{1-2s^*-2H_p(s^*)}.$$

Proof of Lemma B.4. By the same argument as in Lemma A.1,

$$\frac{(2-p)^{2d}}{2^d} = 2^d \mathbb{E} \left[2^{-\text{Bin}(2d, p)} \right] = 2^d \sum_{k=0}^{2d} 2^{-k} \mathbb{P}[\text{Bin}(2d, p) = k]. \quad (\text{B.1})$$

Now a standard derivation using Stirling's approximation (see e.g. [8, Page 353]), we have

$$\frac{1}{2d+1} 2^{2dH(k/2d)} \leq \binom{2d}{k} \leq 2^{2dH(k/2d)},$$

where

$$H(q) = -q \log_2 q - (1-q) \log_2 (1-q)$$

is the standard binary entropy function. This immediately implies that

$$\frac{1}{2d+1} 2^{-2dH_p(k/2d)} \leq \mathbb{P}[\text{Bin}(2d, p) = k] \leq 2^{-2dH_p(k/2d)}.$$

Now using the fact that there are polynomially many terms in the sum (B.1), we find that

$$\frac{1}{\text{poly}(d)} \cdot 2^{d(1-2(k^*/2d)-2H_p(k^*/2d))} \leq \frac{(2-p)^{2d}}{2^d} \leq \text{poly}(d) \cdot 2^{d(1-2(k^*/2d)-2H_p(k^*/2d))} \quad (\text{B.2})$$

for some fixed polynomials in the upper and lower bounds, where $k^* \in \{0, \dots, 2d\}$ maximizes the exponent $1 - 2(k^*/2d) - 2H_p(k^*/2d)$.

Now, for any fixed $p \in (0, 1)$, the function $g(s) = 1 - 2s - 2H_p(s)$ is concave. Thus it has a unique maximum $s^* \in [0, 1]$, and the optimizer k^* is one of the two points in $\{0, \dots, 2d\}$ nearest to s^* on either side. This means that $|\frac{k^*}{2d} - s^*| \leq \frac{1}{2d}$, and so since $g'(s^*) = 0$ we must have

$$|g(s^*) - g(k^*)| \leq \frac{C}{d^2}$$

for some constant C , by Taylor expansion. This combined with (B.2) finishes the proof. $\square$

REFERENCES

- [1] Pierre Baldi, Yosef Rinott, and Charles Stein. A normal approximation for the number of local maxima of a random function on a graph. In *Probability, statistics, and mathematics*, pages 59–81. Elsevier, 1989.
- [2] Béla Bollobás. The evolution of the cube. In *North-Holland Mathematics Studies*, volume 75, pages 91–97. Elsevier, 1983.
- [3] Béla Bollobás. Complete matchings in random subgraphs of the cube. *Random Structures & Algorithms*, 1(1):95–104, 1990.
- [4] Ju D Burtin. The probability of connectedness of a random subgraph of an n-dimensional cube. *Problemy Peredaci Informacii*, 13(2):90–95, 1977.
- [5] Sarah Cannon and Will Perkins. Counting independent sets in unbalanced bipartite graphs. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1456–1466. SIAM, 2020.
- [6] Mriganka Basu Roy Chowdhury, Shirshendu Ganguly, and Vilas Winstein. Gaussian to log-normal transition for independent sets in a percolated hypercube. *arXiv preprint arXiv:2410.07080*, 2024.
- [7] Padraig Condon, Alberto Espuny Díaz, Antonio Girão, Daniela Kühn, and Deryk Osthus. Hamiltonicity of random subgraphs of the hypercube. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 889–898. SIAM, 2021.
- [8] Thomas M Cover. *Elements of information theory*. John Wiley & Sons, 1999.
- [9] Paul Erdős and Joel Spencer. Evolution of the n-cube. *Computers & Mathematics with Applications*, 5(1):33–39, 1979.
- [10] David Galvin. A threshold phenomenon for random independent sets in the discrete hypercube. *Combinatorics, probability and computing*, 20(1):27–51, 2011.
- [11] Matthew Jenssen and Will Perkins. Independent sets in the hypercube revisited. *Journal of the London Mathematical Society*, 102(2):645–669, 2020.
- [12] Aleksej D Korshunov and Alexander A Sapozhenko. The number of binary codes with distance 2. *Problemy Kibernet*, 40:111–130, 1983.
- [13] Gal Kronenberg and Yinon Spinka. Independent sets in random subgraphs of the hypercube. *arXiv preprint arXiv:2201.06127*, 2022.
- [14] Nathan Ross. Fundamentals of Stein’s method. *Probability Surveys*, 8:210–293, 2011.
- [15] Aleksandr Antonovich Sapozhenko. On the number of connected subsets with given cardinality of the boundary in bipartite graphs. *Metody Diskret. Analiz*, 45(45):42–70, 1987.
- [16] Aleksandr Antonovich Sapozhenko. On the number of connected subsets with given cardinality of the boundary in bipartite graphs. *Metody Diskret. Analiz*, 45(45):42–70, 1987.
- [17] Allan Sly. Computational transition at the uniqueness threshold. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 287–296. IEEE, 2010.
- [18] Mingji Xia, Peng Zhang, and Wenbo Zhao. Computational complexity of counting problems on 3-regular planar graphs. *Theoretical Computer Science*, 384(1):111–125, 2007.

MRIGANKA BASU ROY CHOWDHURY, UNIVERSITY OF CALIFORNIA, BERKELEY

Email address: mriganka_brc@berkeley.edu

SHIRSHENDU GANGULY, UNIVERSITY OF CALIFORNIA, BERKELEY

Email address: sganguly@berkeley.edu

VILAS WINSTEIN, UNIVERSITY OF CALIFORNIA, BERKELEY

Email address: vilas@berkeley.edu