



ARTICLE

Satisfiability thresholds for regular occupation problems

Konstantinos Panagiotou  and Matija Pasch 

LMU München, Munich, Germany

Corresponding author: Konstantinos Panagiotou; Email: kpanagio@math.lmu.de

(Received 15 February 2023; revised 25 June 2024; accepted 16 December 2024; first published online 4 February 2025)

Abstract

In the last two decades the study of random instances of constraint satisfaction problems (CSPs) has flourished across several disciplines, including computer science, mathematics and physics. The diversity of the developed methods, on the rigorous and non-rigorous side, has led to major advances regarding both the theoretical as well as the applied viewpoints. Based on a *ceteris paribus* approach in terms of the density evolution equations known from statistical physics, we focus on a specific prominent class of regular CSPs, the so-called *occupation problems*, and in particular on r -in- k occupation problems. By now, out of these CSPs only the satisfiability threshold – the largest degree for which the problem admits asymptotically a solution – for the 1-in- k occupation problem has been rigorously established. Here we determine the satisfiability threshold of the 2-in- k occupation problem for all k . In the proof we exploit the connection of an associated optimization problem regarding the overlap of satisfying assignments to a fixed point problem inspired by belief propagation, a message passing algorithm developed for solving such CSPs.

Keywords: Occupation problems; satisfiability thresholds; second moment method; small subgraph conditioning; contraction coefficient; configuration model

2020 MSC Codes: Primary: 05C80

1. Introduction

Inspired by the pioneering work [16] of Erdős and Rényi in 1960, random discrete structures have been systematically studied in literally thousands of contributions. One of the initial motivations of this research was to study open problems in graph theory and combinatorics. In the following decades, however, the application of such models proved useful as a unified approach to treat a variety of problems in several fields. To mention just a few, random graphs turned out to be valuable in solving fundamental theoretical and practical problems, such as the development of error correcting codes [29], the study of statistical inference through the stochastic block model [1], and the establishment of lower bounds in complexity theory [18, 21].

The results of the past years of research suggest the existence of *phase transitions* in many classes of random discrete structures, i.e. a specific value of a given model parameter at which the properties of the system in question change dramatically. Random constraint satisfaction problems are one specific type of such structures that tend to exhibit this remarkable property and that are of particular interest in too many areas to mention, covering complexity theory, combinatorics, statistical mechanics, artificial intelligence, biology, engineering and economics. An instance of a CSP is defined by a set of variables that take values in – typically finite – domains and a set of constraints, where each constraint is satisfied for specific assignments of the subset of variables it involves. A major computational challenge is to determine whether such an



instance is satisfiable, i.e. to determine if there is an assignment of all variables that satisfies all constraints. Since the 1980s non-rigorous methods have been introduced in statistical physics that are targeted at the analysis of phase transitions in random CSPs [28, 31, 32]. Within this line of research, a variety of exciting and unexpected phenomena were discovered, for example the existence of multiple phase transitions with respect to the structure of the solution space; these transitions may have a significant impact on the hardness of the underlying instances. Since then these methods and the description of the conjectured regimes have been heavily supported by several findings, including the astounding empirical success of randomized algorithms like *belief* and *survey propagation* [5], as well as rigorous verifications, most prominently the phase transition in k -SAT [13] (for large k) and the condensation phase transition in many important models [9]. However, a complete rigorous study is still a big challenge for computer science and mathematics.

Usually, the relevant model parameter of a random CSP is a certain problem specific *density* as illustrated below. The main focus of research is to study the occurrence of phase transitions in the solution space structure and in particular the existence of (*sharp*) *satisfiability thresholds*, i.e. critical values of the density such that the probability that a random CSP admits a solution tends to one as the number of variables tends to infinity for densities below the threshold, while this limiting probability tends to zero for densities above the threshold.

1.1 Random CSPs

Two important types of random CSPs are Erdős-Rényi (ER) type and random regular CSPs. In both cases the number n of variables and the number k of variables involved in each constraint is fixed. In ER type CSPs we further fix the number m of constraints and thereby the *density* $\alpha = m/n$, i.e. the average number of constraints that a variable is involved in. In the regular case we only consider instances where each variable is involved in the same number d of constraints, which fixes the *density* d as well as the number $m = dn/k$ of constraints. In a second step we randomly choose the sets of satisfying assignments for each constraint depending on the problem. For example, in the prominent k -SAT problem one forbidden assignment is chosen uniformly at random from all possible assignments of the involved binary variables for each constraint independently. Another example is the colouring of hypergraphs, where the constraints are attached to the hyperedges and the variables to the vertices, i.e. the variables involved in a constraint correspond to the vertices incident to a hyperedge. In this case a constraint is violated iff all involved vertices take the same colour.

In our work we focus on a class of random regular CSPs in which the choice of satisfying assignments per constraint is fixed in advance, i.e. a class that contains the aforementioned colouring of (d -regular k -uniform) hypergraphs and occupation problems amongst others, but that does not include problems with further randomness in the constraints, like k -SAT and k -XORSAT. The lack of randomness on the level of constraints makes this class particularly accessible for an analysis of the asymptotic solution space structure and significantly simplifies simulations based on the well-known *population dynamics*. Using such simulations, non-rigorous results for this class have been mostly established for the case where the variables are binary valued, so-called occupation problems, or restricted to variants of hypergraph colouring for non-binary variables. Besides the extensive studies on the colouring of simple graphs, i.e. $k = 2$, the only rigorous results derived so far consider the arguably most simple type of occupation problems where each constraint is satisfied if exactly one involved variable evaluates to true, which we refer to as d -regular 1-in- k occupation problem. In our current work we strive to extend these results to general d -regular r -in- k occupation problems, i.e. problems where each constraint is satisfied if r out of the k involved variables evaluate to true.




Figure 1. On the left we see a solution of the 4-regular 2-in-3 occupation problem on a 4-regular 3-factor graph, where the rectangles and circles depict the constraints (factors) and variables (filled if they take the value one in the solution). The figure on the right shows a 2-factor in a 3-regular 4-uniform hypergraph, where the circles, solid and dashed shapes represent the vertices, hyperedges in the 2-factor and the other hyperedges respectively.

1.2 Occupation problems

We continue with the formal definition of the class of problems we consider. Let $k, d \in \mathbb{Z}_{\geq 2}$ and $r \in [k-1] := \{1, \dots, k-1\}$ be fixed. Additionally, we are given non-empty sets V of variables and constraints F . We will use the convention to index elements of V with the letter i and elements of F with the letter a (and subsequent letters) in the remainder. Then an instance o of the d -regular r -in- k occupation problem is specified by a sequence $o = (v(a))_{a \in F}$ of $m = |F|$ subsets $v(a) \subseteq V$ of size k such that each of the $n = |V|$ variables is contained in d of the subsets. In graph theory the instance o has a natural interpretation as a (d, k) -biregular graph (or d -regular k -factor graph) with disjoint node sets $V \dot{\cup} F$ and edges $\{i, a\}$ if $i \in v(a)$. By the handshaking lemma, such objects only exist if $dn = km$, which we assume in the following.

Given an instance o as just described, an assignment $x \in \{0, 1\}^V$ satisfies a constraint $a \in F$ if $\sum_{i \in v(a)} x_i = r$, otherwise x violates a . If x satisfies all constraints $a \in F$, then x is a solution of o . Notice that d times the number of 1's in x matches the total number $rm = rdn/k$ of 1's observed on the factor side, so k has to divide rn , which we also assume in the following. We write $z(o)$ for the number of solutions of o . Figure 1a shows an example of a 4-regular 2-in-3 occupation problem.

Further, for given $m, n \in \mathbb{Z}_{>0}$ let \mathcal{O} denote the set of all instances o with variables $V = [n]$ and constraints $F = [m]$. If \mathcal{O} is not empty, then the random d -regular r -in- k occupation problem \mathcal{O} is uniform on \mathcal{O} and $Z = z(\mathcal{O})$ the number of solutions of \mathcal{O} .

1.3 Examples and related problems

A problem that is closely related and can be reduced to the d -regular r -in- k occupation problem is the d -regular positive r -in- k SAT problem, a variant of k -SAT. We consider a Boolean formula

$$f = \bigwedge_{a \in F} c_a, \quad c_a = \bigvee_{i \in v(a)} i, \quad a \in F,$$

in conjunctive normal form with m clauses over n variables $i \in V$, such that no literal appears negated (hence positive r -in- k SAT), and where each clause c_a is the disjunction of k literals and each variable appears in exactly d clauses (hence d -regular). The decision problem is to determine if there exists an assignment x such that exactly r literals in each clause evaluate to true (hence r -in- k SAT). In [34] the satisfiability threshold for this problem was determined for $r = 1$, i.e.

the case where exactly one literal in each clause evaluates to true. Our Theorem 1.1 solves this problem for $r = 2$ and $k \in \mathbb{Z}_{\geq 4}$.

Our second example deals with a problem from graph theory. A k -regular d -uniform hypergraph h is a pair $h = (F, E)$ with $m = |F|$ vertices and $n = |E|$ (hyper-)edges such that each edge contains d vertices and the degree of each vertex is k . An r -factor E' is a subset of the hyperedges such that each vertex $a \in F$ is incident to r hyperedges $e_i \in E'$. In this case the problem is to determine if h has an r -factor. For example, the case $r = 1$ is the well-known perfect matching problem and the threshold was determined in [11]. An example of a 2-factor in a hypergraph is shown in Figure 1b. Theorem 1.1 solves also this problem for $r = 2$ and $k \in \mathbb{Z}_{\geq 4}$.

Several other problems in complexity and graph theory are closely related to the examples above. The satisfiability threshold in Theorem 1.1 also applies to a variant of the vertex cover problem (or hitting set problem from set theory perspective), where we choose a subset of the vertices (variables with value one) in a d -regular k -uniform hypergraph such that each hyperedge is incident to exactly two vertices in the subset. Analogously, Theorem 1.1 also establishes the threshold for a variant of the set cover problem in set theory corresponding to 2-factors in hypergraphs, i.e. given a family of d -subsets (hyperedges) and a universe (vertices) with each element contained in k subsets, the problem is to find a subfamily of the subsets such that each element of the universe is contained in exactly two subsets of the subfamily. Further, Theorem 1.1 can, e.g. also be used to give sufficient conditions for the (asymptotic) existence of Euler families in regular uniform hypergraphs as discussed in [2].

1.4 Main results

The satisfiability threshold for the d -regular 1-in- k occupation problem has been established in [11, 34], which also covers the d -regular 2-in-3 occupation problem due to colour symmetry. Our main result pins down the location of the satisfiability threshold of the random d -regular 2-in- k occupation problem for $k \in \mathbb{Z}_{\geq 4}$. For this purpose let

$$d^* = d^*(k) = \frac{kH(2/k)}{kH(2/k) - \ln \binom{k}{2}}, \quad k \in \mathbb{Z}_{\geq 4}, \quad (1)$$

where $H(p) = -p \ln(p) - (1-p) \ln(1-p)$ is the entropy of $p \in [0, 1]$. The following theorem establishes the location of the threshold at d^* .

Theorem 1.1 (2-in- k Occupation Satisfiability Threshold). *Let $k \in \mathbb{Z}_{\geq 4}$, $d \in \mathbb{Z}_{\geq 2}$, and let Z be the number of solutions from Section 1.1. There exists a sharp satisfiability threshold at d^* , i.e. for any increasing sequence $(n_i)_{i \in \mathbb{Z}_{>0}} \subseteq \mathcal{N} = \{n: dn, 2n \in k\mathbb{Z}_{>0}\}$ and $m_i = dn_i/k$ we have*

$$\lim_{i \rightarrow \infty} \mathbb{P}(Z > 0) = \begin{cases} 1, & d < d^* \\ 0, & d \geq d^* \end{cases}.$$

We provide a self-contained proof for Theorem 1.1 using the first and second moment method with small subgraph conditioning for Z . In particular, a main technical contribution in proving Theorem 1.1 is the optimization of a certain multivariate function that appears in the computation of the second moment, which encodes the interplay between the ‘similarity’ of various assignments and the change in the corresponding probability of being satisfying that they induce. A direct corollary of this optimization step at the threshold d^* is the confirmation of the conjecture by the authors in [36]. Among other things, at the core of our contribution we take a novel and rather different approach to tackle the optimization, inspired by [37] and [41] as well as other works relating the fixed points of belief propagation to the stationary points of the Bethe free entropy, respectively to the computation of the annealed free entropy density; see Section 5.6 for details.

Finally, we show that d^* is not an integer in Lemma 3.1 below, so as opposed to the case $r = 1$ [34], for $r = 2$ there is no need for a dedicated analysis at criticality.

1.5 Related work

The regular version of the random 1-in- k occupation problem (and related problems) has been studied in [11, 34] using the first and second moment method with small subgraph conditioning. The paper [37] shows that $\lim_{i \rightarrow \infty} \mathbb{P}(Z > 0) = 1$ for $d = 2$ and $k \in \mathbb{Z}_{\geq 2}$ in the d -regular 2-in- k occupation problem, i.e. the existence of 2-factors in k -regular simple graphs. A recent discussion of 2-factors (and the related Euler families) that does not rely on the probabilistic method is presented in [2]. Further, randomized polynomial time algorithms for the generation and approximate counting of 2-factors in random regular graphs have been developed in [19].

The study of Erdős-Rényi (hyper-)graphs was initiated by the groundbreaking paper [16] in 1960 and turned into a fruitful field of research with many applications, including early results on 1-factors in simple graphs [17]. On the contrary, results for the random d -regular k -uniform (hyper-)graph ensemble were rare before the introduction of the configuration (or pairing) model by Bollobás [4] and the development of the small subgraph conditioning method [23, 24]. While the proof scheme facilitated rigorous arguments to establish the existence and location of satisfiability thresholds of random regular CSPs [3, 7, 10, 14, 15, 27, 33], the problems are treated on a case by case basis, while results on entire classes of random regular CSPs are still outstanding.

One of the main reasons responsible for the complexity of a rigorous analysis of random (regular) CSPs seems to be a conjectured structural change of the solution space for increasing densities. This hypothesis has been put forward by physicists, verified in parts and mostly for ER ensembles, further led to new rigorous proof techniques [8, 10, 13] and to randomized algorithms [5, 30] for NP-hard problems that are not only of great value in practice, but can also be employed for precise numerical (though non-rigorous) estimates of satisfiability thresholds. An excellent introduction to this *replica theory* can be found in [28, 31, 40]. Specifically, numerical results indicating the satisfiability thresholds for d -regular r -in- k occupation problems (more general variants, and for ER type hypergraphs) based on this conjecture were discussed in various publications [6, 12, 20, 22, 39, 42, 43], where *occupation problems* were introduced for the first time in [35].

Another fundamental obstacle in the rigorous analysis is of a very technical nature and directly related to the second moment method as discussed in detail in our current work. In the case of regular 2-in- k occupation problems (amongst others) this optimization problem can be solved by exploiting a connection to the fixed points of belief propagation. This well-studied message passing algorithm is thoroughly discussed in [31].

1.6 Open problems

In this work we rigorously establish the threshold for $r = 2$ and $k \in \mathbb{Z}_{\geq 4}$ for the random regular r -in- k occupation problem. A rigorous proof for general r (and k) seems to be involved, but further assumptions may significantly simplify the analysis. For example, as an extension of the current work one may focus on r -in- $2r$ occupation problems, where the constraints are symmetric in the colours. As can be seen from our proof, this yields useful symmetry properties. Further, as suggested by the literature [8, 9] such *balanced* problems [42, 43] are usually more accessible to a rigorous study. On the other hand, the optimization usually also significantly simplifies if only carried out for $k \geq k_0(r)$ for some (large) $k_0(r)$.

Apart from the generalizations discussed above, results for the general r -in- k occupation problems are also still outstanding for Erdős-Rényi type CSPs, the only exception being the satisfiability threshold for perfect matchings which was recently established by Kahn [25]. Further, there only exist bounds for the exact cover problem [26] on 3-uniform hypergraphs, i.e. $r = 1$ and $k = 3$.

1.7 Outline of the Proofs

In Section 2 we present the proof strategy on a high level. Then, we turn to the notation and do some groundwork, in particular the analysis of $d^*(k)$, in Section 3. The easy part of the main result is established in Section 4 using the first moment method. The remainder is devoted to the proof that solutions exist below the threshold with probability tending to one, starting with the second moment method in Section 5. Most of the twenty pages in this section are devoted to the solution of the optimization problem and related conjecture from [36] using a belief propagation inspired approach.

Finally, we complete the small subgraph conditioning method in Section 6, using the proof of Lemma 2.8 in Appendix A as a blueprint.

2. Proof techniques

In this section we give a high-level overview of our proof. We make heavy use of the so-called *configuration model* for the generation of random instances in the form used by Moore [34].

2.1 The configuration model

Working with the uniform distribution on d -regular k -uniform hypergraphs directly is challenging. Instead, we show Theorem 1.1 for occupation problems on so-called *configurations*. A d -regular k -configuration is a bijection $g: [n] \times [d] \rightarrow [m] \times [k]$, where the v -edges $(i, h') \in [n] \times [d]$ represent pairs of variables $i \in [n]$ and so-called i -edges, i.e. half-edge indices $h' \in [d]$. The image $(a, h) = g(i, h')$ is an f -edge, i.e. a pair of a constraint (factor) $a \in [m]$ and an a -edge (or half-edge) $h \in [k]$, indicating that the i -edge h' of the variable i is wired to the a -edge h of a and thereby suggesting that i is connected to a in the corresponding d -regular k -factor graph. Notice that we can represent g by an equivalent, four-partite, graph with (disjoint) vertex sets given by the variables $V = [n]$, constraints (factors) $F = [m]$, v -edges $H' = [n] \times [d]$ and f -edges $H = [m] \times [k]$, where each variable $i \in [n]$ connects to all its v -edges $(i, h') \in H'$, each constraint $a \in [m]$ to all its f -edges $(a, h) \in H$ and a v -edge (i, h') connects to an f -edge (a, h) if $g(i, h') = (a, h)$.

Let \mathcal{G} be the set of all d -regular k -configurations on n variables, and notice that $|\mathcal{G}| = \emptyset$ iff $dn \neq km$ and $|\mathcal{G}| = (dn)! = (km)!$ for $m = dn/k \in \mathbb{Z}$, which we assume from here on. Further, the occupation problem on factor graphs directly translates to configurations, i.e. an assignment $x \in \{0, 1\}^n$ is a solution of $g \in \mathcal{G}$ if for each constraint $a \in [m]$ there exist exactly two distinct a -edges $h, h' \in [k]$ such that $x_{i(a,h)} = x_{i(a,h')} = 1$, where $i(a, h) = (g^{-1}(a, h))_1$ denotes the h -th neighbour of a . Say, the occupation problem on a configuration corresponding to the example in Figure 1a is shown in Figure 2a.

Let $Z(g)$ be the number of solutions of $g \in \mathcal{G}$. As before, $Z = 0$ almost surely unless $2n \in k\mathbb{Z}$. Theorem 1.1 is a straightforward consequence of the following result.

Theorem 2.1 (Satisfiability Threshold for Configurations). *Theorem 1.1 also holds for Z as defined in this section.*

Theorem 2.1 translates to the models in Section 1.1 and Section 1.2 using standard arguments, namely symmetry, the discussion of parallel edges, contiguity, and the fact that both the variable and the factor neighbourhoods are unique with probability tending to one. Hence, in the remainder of this contribution we exclusively consider the occupation problem on configurations.

2.2 The first moment method

In the first step we apply the first moment method to the occupation problem on configurations, yielding the following result.

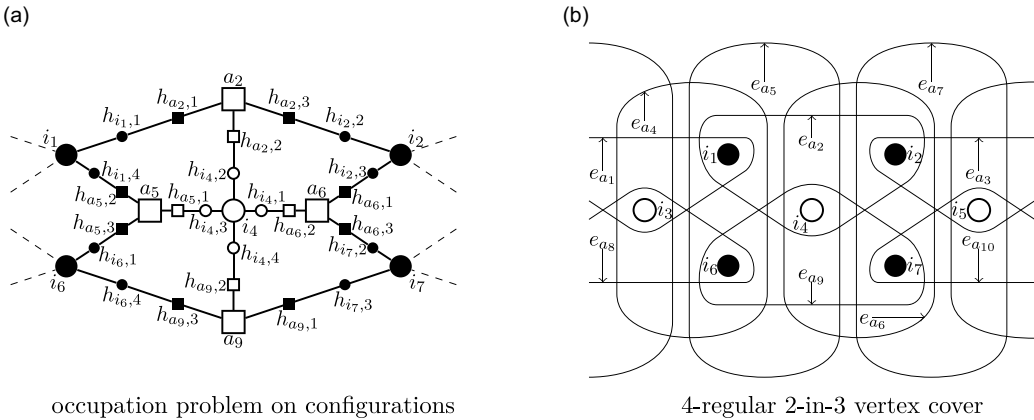


Figure 2. The figure on the left shows the solution on a configuration corresponding to the solution in Figure 1. We only denoted a -edges (small boxes, filled if they the a -edge takes the value one) and i -edges (small circles, filled if the i -edge takes the value one) instead of f -edges and v -edges for brevity (e.g. $h_{a1,1}$ instead of $(a_1, h_{a1,1})$). The figure on the right illustrates the corresponding 2-in-3 vertex cover (given by the filled circles).

Lemma 2.2 (First Moment Method). *Let $k \in \mathbb{Z}_{\geq 4}$, $d \in \mathbb{Z}_{\geq 2}$. For $n \in \mathcal{N}$ tending to infinity*

$$\mathbb{E}[Z] \sim \sqrt{d} e^{n\phi_1}, \quad \text{where} \quad \phi_1 = \frac{d}{k} \ln \binom{k}{2} - (d-1)H\left(\frac{2}{k}\right).$$

In particular, $\mathbb{E}[Z] \rightarrow \infty$ for $d < d^*$ and $\mathbb{E}[Z] \rightarrow 0$ for $d > d^*$ with d^* as in (1). So, Markov's inequality implies $\mathbb{P}(Z > 0) \rightarrow 0$ for $d > d^*$. The map ϕ_1 is known as *annealed free entropy density*.

2.3 The second moment method

Let $k \in \mathbb{Z}_{\geq 4}$ and $d \in \mathbb{Z}_{\geq 2}$. We denote the set of distributions on a finite set S by $\mathcal{P}(S)$ and identify $p \in \mathcal{P}(S)$ with its probability mass function, meaning $\mathcal{P}(S) = \{p \in [0, 1]^S : \sum_{x \in S} p(x) = 1\}$. Further, let $\mathcal{P}_\ell(S) = \{p \in \mathcal{P}(S) : \ell p \in \mathbb{Z}^S\}$ be the empirical distributions over $\ell \in \mathbb{Z}_{>0}$ trials.

In order to apply the second moment method we will consider a (new) CSP with m factors on n variables with the larger domain $\{0, 1\}^2$, and where the constraint $a \in [m]$ is satisfied by an assignment $x \in (\{0, 1\}^2)^n$ if $\sum_{i \in v(a)} x_{i,1} = \sum_{i \in v(a)} x_{i,2} = 2$. Here, there are qualitatively three types of satisfying assignments for the constraints, namely with 0, 1 or 2 overlapping ones. We will analyse the empirical *overlap distributions* $p \in \mathcal{P}_m(\{0, 1, 2\})$ of assignments satisfying all constraints, which determine the empirical distributions $p_e \in \mathcal{P}_{km}(\{0, 1\}^2)$ of the values $\{0, 1\}^2$ over the km edges, given by

$$p_e(11) = \frac{1}{k}p(1) + \frac{2}{k}p(2) \quad \text{and} \quad p_e(10) = p_e(01) = \frac{1}{k}p(1) + \frac{2}{k}p(0).$$

So, if $p \in \mathcal{P}_m(\{0, 1, 2\})$ is an *achievable* empirical overlap distribution on the m factors, then p_e is necessarily an empirical distribution on the n variables; thus the achievable overlap distributions are contained in $\mathcal{P}_n = \{p \in \mathcal{P}_m(\{0, 1, 2\}) : p_e \in \mathcal{P}_n(\{0, 1\}^2)\}$.

In the first – combinatorial – part we establish that the second moment can be written as a sum of all contributions over all achievable overlap distributions.

Lemma 2.3 (Second Moment Combinatorics). *For any $n \in \mathcal{N}$ we have*

$$\mathbb{E}[Z^2] = \sum_{p \in \mathcal{P}_n} E(p), \quad \text{where} \quad E(p) = \binom{m}{mp} \prod_{s \in \{0,1,2\}} \binom{k}{s, 2-s, 2-s, k-4+s}^{mp(s)} \binom{n}{np_e} \left(\frac{dn}{np_e}\right)^{-1}.$$

Here, we use the notation $\binom{m}{mp}$, $p \in \mathcal{P}_m(\{0, 1, 2\})$, for multinomial coefficients.

To study further the second moment in Lemma 2.3, we identify the maximal contributions. For this purpose, let $p^* \in \mathcal{P}(\{0, 1, 2\})$ be the hypergeometric distribution with

$$p^*(s) = \frac{\binom{2}{s} \binom{k-2}{2-s}}{\binom{k}{2}} \text{ for } s \in \{0, 1, 2\}, \quad \text{and} \quad p_e^*(1, 1) = \frac{4}{k^2}, \quad p_e^*(1, 0) = \frac{2(k-2)}{k^2}. \quad (2)$$

The overlap distribution p^* is a natural candidate for maximizing $E(p)$. Indeed, we obtain p^* when we consider two independent uniformly random assignments in $\{0, 1\}^k$ with 2 ones each, and p_e^* is exactly the marginal probability if we jointly consider two independent uniformly random assignments in $\{0, 1\}^n$ to the variables with $2n/k$ ones each. In the next step, we derive the limits of the log-densities $\frac{1}{n} \ln(E(p))$. Recall that the K(ullback)-L(eibler) divergence $D_{\text{KL}}(p \parallel q)$ of two distributions $p, q \in \mathcal{P}(S)$, such that p is absolutely continuous with respect to q , is

$$D_{\text{KL}}(p \parallel q) = \sum_{x \in S} p(x) \ln \left(\frac{p(x)}{q(x)} \right).$$

Lemma 2.4 (Second Moment Asymptotics). *For any fully supported $p \in \mathcal{P}(\{0, 1, 2\})$ and any sequence $(p_n)_{n \in \mathcal{N}} \subseteq \mathcal{P}_n$ with $\lim_{n \rightarrow \infty} p_n = p$ we have $\lim_{n \rightarrow \infty} \frac{1}{n} \ln(E(p_n)) = \phi_2(p)$, where*

$$\phi_2(p) = 2\phi_1 - \frac{d}{k} \Delta_d(p) \quad \text{and} \quad \Delta_d(p) = D_{\text{KL}}(p \parallel p^*) - \frac{(d-1)k}{d} D_{\text{KL}}(p_e \parallel p_e^*).$$

The following proposition is the main contribution of this work.

Proposition 2.5 (Second Moment Minimizers). *For $k = 4$ the global minimizers of $\Delta_{d^*(4)}$ are p^* , $p^{(0)}$ given by $p^{(0)}(0) = 1$ and $p^{(2)}$ given by $p^{(2)}(2) = 1$. For $k \in \mathbb{Z}_{\geq 5}$ the global minimizers of $\Delta_{d^*(k)}$ are p^* and $p^{(2)}$.*

With Proposition 2.5, we easily verify that p^* is the unique minimizer of Δ_d for any $d < d^*(k)$, since the KL divergence is minimized by its unique root and $(d-1)k/d$ is increasing in d . This conclusion then allows us to compute the limit of the scaled second moment using Laplace's method for sums. More than that, we confirm the conjecture by the authors in [36] as an immediate corollary.

Proposition 2.6 (Second Moment Limit). *For any $k \in \mathbb{Z}_{\geq 4}$ and $d < d^*(k)$*

$$\frac{\mathbb{E}[Z^2]}{\mathbb{E}[Z]^2} \sim \sqrt{\frac{k-1}{k-d}}, \quad \text{as } n \in \mathcal{N} \text{ tends to infinity.}$$

Proposition 2.6 and the Paley-Zygmund inequality yield $\liminf_{n \rightarrow \infty} \mathbb{P}(Z > 0) \geq \sqrt{\frac{k-d}{k-1}}$. While this bound suggests that a threshold exists, we need to show that the threshold at d^* is sharp.

2.4 Small subgraph conditioning

We complete the proof of Theorem 2.1 using the small subgraph conditioning method. For this purpose let $a^b = \prod_{c=0}^{b-1} (a-c)$ denote the falling factorial.

Theorem 2.7 (Small Subgraph Conditioning, [34, Theorem 2]). *Let Z_n and $X_{n,1}, X_{n,2}, \dots$ be non-negative integer-valued random variables. Suppose that $\mathbb{E}[Z_n] > 0$ and that for each $\ell \in \mathbb{Z}_{>0}$ there are $\lambda_\ell \in \mathbb{R}_{>0}$, $\delta_\ell \in \mathbb{R}_{>-1}$ such that for any $L \in \mathbb{Z}_{>0}$*

- a) the variables $X_{n,1}, \dots, X_{n,L}$ are asymptotically independent and Poisson with $\mathbb{E}[X_{n,\ell}] \sim \lambda_\ell$,
 b) for any sequence r_1, \dots, r_L of non-negative integers,

$$\frac{\mathbb{E}\left[Z_n \prod_{\ell=1}^L X_{n,\ell}^{r_\ell}\right]}{\mathbb{E}[Z_n]} \sim \prod_{\ell=1}^L [\lambda_\ell(1 + \delta_\ell)]^{r_\ell},$$

- c) we explain the variance, i.e.

$$\frac{\mathbb{E}[Z_n^2]}{\mathbb{E}[Z_n]^2} \sim \exp\left(\sum_{\ell \geq 1} \lambda_\ell \delta_\ell^2\right) \quad \text{and} \quad \sum_{\ell \geq 1} \lambda_\ell \delta_\ell^2 < \infty.$$

Then $\lim_{n \rightarrow \infty} \mathbb{P}(Z_n > 0) = 1$.

We will apply Theorem 2.7 to the number Z of solutions from Section 2.1 and the numbers X_ℓ of small cycles in the configuration G . In order to understand what a cycle in a configuration is, we recall the representation of a configuration g as a four-partite graph from Section 2.1.

Since we are mostly interested in the factor graph associated with a configuration we divide the lengths of paths by three, e.g. what we call a cycle of length four in the bijection, is actually a cycle of length twelve in its equivalent four-partite graph representation. Figures 1a and 2a show an example of a factor graph and the corresponding configuration in its graph representation. Showing the following statement, which establishes Assumption 2.7a), is rather routine.

Lemma 2.8 (Small Cycles). For $\ell \in \mathbb{Z}_{>0}$ let X_ℓ be the number of 2ℓ -cycles in G , and set

$$\lambda_\ell = \frac{[(k-1)(d-1)]^\ell}{2^\ell}.$$

Then X_1, \dots, X_L are asymptotically independent and Poisson with $\mathbb{E}[X_\ell] \sim \lambda_\ell$ for all $L \in \mathbb{Z}_{>0}$.

We give a self-contained proof of Lemma 2.8 in the appendix, which we build upon to argue that Assumption 2.7b) in Theorem 2.7 holds. With Lemma 2.8 in place, we consider the base case in Assumption 2.7b), i.e. for $\ell \in \mathbb{Z}_{>0}$ we let $r_\ell = 1$ and $r_{\ell'} = 0$ otherwise, to determine $\delta_\ell = (1-k)^{-\ell}$. We easily verify that $\sum_{\ell \geq 1} \lambda_\ell \delta_\ell^2 = \frac{1}{2} \ln\left(\frac{k-1}{k-d}\right)$ and thereby establish Assumption 2.7c) using Proposition 2.6. Finally, we follow the proof of Lemma 2.8 to complete the verification of Assumption 2.7b) and thereby complete the proof of Theorem 1.1.

3. Preliminaries and notation

After introducing notation in Section 3.1, we establish a few basic facts in Section 3.2.

3.1 Notation

We use the notation $[n] = \{1, \dots, n\}$ and $[n]_0 = \{0\} \cup [n]$ for $n \in \mathbb{Z}_{>0}$, denote the falling factorial with $n^{\underline{k}}$ for $n, k \in \mathbb{Z}_{\geq 0}$, $k \leq n$, and multinomial coefficients with $\binom{n}{k}$ for $n \in \mathbb{Z}_{\geq 0}$ and $k \in \mathbb{Z}_{\geq 0}^d$, $d \in \mathbb{Z}_{>1}$, such that $\sum_{i \in [d]} k_i = n$. For functions f, g on integers with $\lim_{n \rightarrow \infty} f(n)/g(n) = 1$ we write $f(n) \sim g(n)$. We make heavy use of Stirling's formula [38], i.e.

$$\sqrt{2\pi n} \left(\frac{n}{e}\right)^n e^{\frac{1}{12n+1}} \leq n! \leq \sqrt{2\pi n} \left(\frac{n}{e}\right)^n e^{\frac{1}{12n}}, \quad n \in \mathbb{Z}_{>0},$$

and in particular $n! \sim \sqrt{2\pi n} \left(\frac{n}{e}\right)^n$. If a random variable X has law P we write $X \sim P$ and use $\text{Po}(\lambda)$ to denote the Poisson distribution with parameter λ . Distributions $p \in \mathcal{P}(\mathcal{S})$ in the convex polytope $\mathcal{P}(\mathcal{S})$ of distributions with finite support \mathcal{S} are identified with their probability mass

functions $p \in [0, 1]^S$. Further, $\mathcal{P}_n(\mathcal{S}) = \{p \in \mathcal{P}(\mathcal{S}) : np \in [n]_0\}$ denotes the set of empirical distributions obtained from $n \in \mathbb{Z}_{\geq 1}$ trials. Let v^t denote the transpose of a vector v . Finally, we use ‘iff’ for ‘if and only if’.

3.2 Basic observations

We briefly establish the claims in Section 1.1 for the configuration version, and the claim that d^* is not an integer.

Lemma 3.1. *The set \mathcal{G} is empty iff $dn \neq km$, so let $dn = km$. Then, we have $Z = 0$ almost surely if $n_1 = 2n/k \notin \mathbb{Z}$. Finally, $d^* \in (1, \infty) \setminus \mathbb{Z}$.*

Proof. Since $g \in \mathcal{G}$ is a bijection $g: [n] \times [d] \rightarrow [m] \times [k]$, the set \mathcal{G} is empty for $dn \neq km$ and $|\mathcal{G}| = (dn)! = (km)!$ otherwise, which proves the first assertion. Next, we fix a solution $x \in \{0, 1\}^n$ of $g \in \mathcal{G}$ with n'_1 ones. Then two a -edges h have to take the value one, i.e. $x_{i(a,h)} = 1$, for each $a \in [m]$ and hence $2m$ f -edges $(a, h) \in [m] \times [k]$ in total. On the other hand, there are dn'_1 v -edges $(i, h) \in [n] \times [d]$ that take the value one. Since g is a bijection, $dn'_1 = 2m$, so $n_1 = n'_1 \in \mathbb{Z}$.

For the last assertion, we first focus on the denominator of d^* , i.e.

$$kH\left(\frac{2}{k}\right) - \ln\left(\frac{k}{2}\right) = -\ln\left(\frac{\binom{k}{2}}{\binom{2}{2}} \left(\frac{2}{k}\right)^2 \left(\frac{k-2}{k}\right)^{k-2}\right) > 0,$$

so $d^* > 0$ for $k \in \mathbb{Z}_{\geq 3}$. Next, notice that d^* is a solution of $f(d) = 1$ with

$$f(d) = e^{(d-1)(kH(2/k) - \ln(\frac{k}{2})) - \ln(\frac{k}{2})} = \frac{2}{k(k-1)} \left(\frac{k^{k-1}}{2(k-2)^{k-2}(k-1)} \right)^{d-1},$$

which directly implies that $d^* > 1$ and further, since $\gcd(k, k-1) = 1$, that $d^* \in (1, \infty) \setminus \mathbb{Z}$. \square

4. The first moment method – proof of Lemma 2.2

This short section is dedicated to the proof of Lemma 2.2. We write the expectation in terms of the number $|\mathcal{E}|$ of pairs $(g, x) \in \mathcal{E}$ such that $x \in \{0, 1\}^n$ satisfies $g \in \mathcal{G}$, i.e.

$$\mathbb{E}[Z] = \frac{|\mathcal{E}|}{|\mathcal{G}|} = \frac{1}{(dn)!} \binom{n}{n_1} \binom{k}{2}^m (2m)!(dn - 2m)!,$$

with $n_1 = 2n/k$ and for the following reasons. First, we choose the n_1 variables with value one in x , then we choose the two a -edges for each constraint $a \in [m]$ with value one, wire the v -edges and f -edges with value one and finally wire the edges with value zero. In particular, this implies that $\mathbb{E}[Z] > 0$ for all $n \in \mathcal{N}$.

Using Stirling’s formula, the asymptotics are given by

$$\mathbb{E}[Z] = \frac{\binom{n}{n_1} \binom{k}{2}^m}{\binom{km}{2m}} \sim \sqrt{\frac{2\pi km \frac{2}{k} (1 - \frac{2}{k})}{2\pi n \frac{2}{k} (1 - \frac{2}{k})}} \exp\left(nH\left(\frac{2}{k}\right) - kmH\left(\frac{2}{k}\right) + m \ln\left(\binom{k}{2}\right)\right) = \sqrt{d} e^{n\phi_1}.$$

5. The second moment method

In this section we consider the case $d < d^*$. We prove Lemma 2.3, Lemma 2.4, Proposition 2.5 and Proposition 2.6, the main contribution of this work.

5.1 How to square a constraint satisfaction problem

In order to facilitate the presentation we introduce the *squared* d -regular 2-in- k occupation problem. As before, an instance of this problem is given by a bijection $g: [n] \times [d] \rightarrow [m] \times [k]$. Now, for an assignment $x \in (\{0, 1\}^2)^n$ let $y_{g,x} = (x_{i(a,h)})_{a \in [m], h \in [k]}$ be the corresponding f -edge assignment under g , where we recall from Section 2.1 that $i(a, h) = (g^{-1}(a, h))_1 \in [n]$ is the variable $i(a, h)$ wired to the f -edge (a, h) under g . A constraint $a \in [m]$ is satisfied by a constraint assignment $x \in (\{0, 1\}^2)^k$ iff $x \in \mathcal{S}^{(2)}$, where

$$\mathcal{S}^{(2)} = \left\{ x \in (\{0, 1\}^2)^k : \sum_{h \in [k]} x_{h,1} = \sum_{h \in [k]} x_{h,2} = 2 \right\}.$$

An f -edge assignment $x \in (\{0, 1\}^2)^{m \times k}$ is satisfying if $x_a = (x_{a,h})_{h \in [k]}$ satisfies a for all $a \in [m]$. Finally, an assignment $x \in (\{0, 1\}^2)^n$ is a solution of g if $y_{g,x}$ is satisfying. Notice that the pairs of solutions $x, x' \in \{0, 1\}^n$ of the standard problem on g are in one to one correspondence with the solutions $y \in (\{0, 1\}^2)^n$ of the squared problem on g via $y = (x_i, x'_i)_{i \in [n]}$. So, $z^{(2)}(g) = z(g)^2$ for the number $z^{(2)}(g)$ of solutions of the squared problem, hence $Z^{(2)} = Z^2$ for $Z^{(2)} = z^{(2)}(G)$ and in particular $\mathbb{E}[Z^{(2)}] = \mathbb{E}[Z^2]$. This equivalence allows us to entirely focus on the squared problem.

5.2 Proof of Lemma 2.3

As before, we can write $\mathbb{E}[Z^{(2)}] = \frac{1}{(dn)!} |\mathcal{E}|$, where $|\mathcal{E}|$ is the number of pairs $(g, x) \in \mathcal{E}$ such that $x \in (\{0, 1\}^2)^n$ solves g . Set

$$\mathcal{Y} = \left\{ y \in (\{0, 1\}^2)^{m \times k} : y_a \in \mathcal{S}^{(2)} \text{ for all } a \in [m] \right\}.$$

For $y \in \mathcal{Y}$ let the *overlap distribution* $p_y \in \mathcal{P}_m(\{0, 1, 2\})$ be given by

$$p_y(s) = \frac{1}{m} |\{a \in [m] : |y_a^{-1}(1, 1)| = s\}|, \quad s \in \{0, 1, 2\}.$$

Further, let the *edge distribution* $q_y \in \mathcal{P}_{km}(\{0, 1\}^2)$ be given by

$$q_y(x) = \frac{1}{km} |\{(a, h) \in [m] \times [k] : y_{a,h} = x\}| = \frac{1}{km} |y^{-1}(x)|, \quad x \in \{0, 1\}^2.$$

Using that $|y_a^{-1}(1, 0)| = |y_a^{-1}(0, 1)| = 2 - |y_a^{-1}(1, 1)|$ and hence $|y^{-1}(0, 0)| = k - 4 + |y^{-1}(1, 1)|$ we directly get

$$\begin{aligned} q_y(1, 1) &= \frac{1}{km} \sum_{a \in [m]} |y_a^{-1}(1, 1)| = \frac{1}{km} \sum_{s \in \{0, 1, 2\}} s |\{a \in [m] : |y_a^{-1}(1, 1)| = s\}| = \sum_{s \in \{0, 1, 2\}} \frac{s}{k} p_y(s), \\ q_y(1, 0) &= q_y(0, 1) = \frac{1}{km} \sum_{s \in \{0, 1, 2\}} (2 - s) |\{a \in [m] : |y_a^{-1}(1, 1)| = s\}| = \sum_{s \in \{0, 1, 2\}} \frac{2 - s}{k} p_y(s), \\ q_y(0, 0) &= \sum_{s \in \{0, 1, 2\}} \frac{k - 4 + s}{k} p_y(s). \end{aligned}$$

Hence, let $p_e = Wp \in \mathcal{P}(\{0, 1\}^2)$ denote the edge distribution of any (not necessarily empirical) overlap distribution $p \in \mathcal{P}(\{0, 1, 2\})$, where $W \in [0, 1]^{\{0, 1\}^2 \times \{0, 1, 2\}}$ is given by

$$W_{11,s} = \frac{s}{k}, \quad W_{10,s} = W_{01,s} = \frac{2-s}{k} \quad \text{and} \quad W_{00,s} = \frac{k-4+s}{k}, \quad s \in \{0, 1, 2\}. \quad (3)$$

Now, notice that for any $(g, x) \in \mathcal{E}$ we have $y_{g,x,a,h} = x_{i(a,h)}$ for all $a \in [m], h \in [k]$, hence $g(x^{-1}(z) \times [d]) = y_{g,x}^{-1}(z)$ and by that

$$q_{y_{g,x}}(z) = \frac{|y_{g,x}^{-1}(z)|}{km} = \frac{d|x^{-1}(z)|}{km} = \frac{1}{n}|x^{-1}(z)| = q_x(z) \text{ for } z \in \{0, 1\}^2,$$

i.e. the relative frequencies of the values in the f -edge assignment $y_{g,x}$ coincide with the relative frequencies $q_x \in \mathcal{P}_n(\{0, 1\}^2)$ of the values in the variable assignment x . In particular, this shows that a satisfying f -edge assignment $y \in \mathcal{Y}$ is only *attainable* if $q_y \in \mathcal{P}_n(\{0, 1\}^2)$, and thereby

$$\mathbb{E}[Z^{(2)}] = \frac{1}{(dn)!} \sum_{p \in \mathcal{P}_n} |\{(g, x) \in \mathcal{E} : p_{y_{g,x}} = p\}|.$$

Now, fix an attainable satisfying f -edge assignment $y \in \mathcal{Y}$ and an assignment $x \in (\{0, 1\}^2)^n$ with $q_y = q_x$, i.e. $|x^{-1}(z) \times [d]| = |y^{-1}(z)|$ for all $z \in \{0, 1\}^2$. Any bijection g with $y = y_{g,x}$ needs to respect $g(x^{-1}(z) \times [d]) = y_{g,x}^{-1}(z)$ for $z \in \{0, 1\}^2$ and can hence be uniquely decomposed into its restrictions $g_z : x^{-1}(z) \times [d] \rightarrow y_{g,x}^{-1}(z)$. On the other hand, any choice of such restrictions g_z gives a bijection g with $y = y_{g,x}$, and so

$$|\mathcal{E}_{x,y}| = \prod_{z \in \{0,1\}^2} (dnq_x(z))! = \prod_{z \in \{0,1\}^2} (dnp_{y_e(z)})!, \text{ where } \mathcal{E}_{x,y} = \{(g, x) \in \mathcal{E} : y_{g,x} = y\}.$$

Notice that $\mathcal{E}_{x,y} \cap \mathcal{E}_{x',y'} = \emptyset$ for any $(x, y) \neq (x', y')$, which is obvious for $x \neq x'$, and also for $y \neq y'$, since $y_{g,x} = y \neq y' = y_{g',x}$ implies that $g \neq g'$. But since $|\mathcal{E}_{x,y}|$ only depends on p_y (actually only on p_{y_e}) this completes the proof, because for any fixed attainable overlap distribution $p \in \mathcal{P}_n$, we can now independently choose the satisfying f -edge assignment y and variable assignment x , subject to $q_x = p_e$ and $p_y = p$ (which implies $q_y = q_x$). So we have $\mathbb{E}[Z^{(2)}] = \sum_p E(p)$ with $p \in \mathcal{P}_n$ and

$$E(p) = \frac{1}{(dn)!} \binom{n}{np_e} \binom{m}{mp} \prod_{s \in \{0,1,2\}} \binom{k}{s, 2-s, 2-s, k-4+s}^{mp(s)} \prod_{x \in \{0,1\}^2} (dnp_e(z))!,$$

where we choose a variable assignment x with $q_x = p_e$, an f -edge assignment y with $p_y = p$ by first choosing one of the $\binom{m}{mp}$ options for $(|y_a^{-1}(1, 1)|)_{a \in [m]}$ and then independently one of the $\binom{k}{s, 2-s, 2-s, k-4+s}$ satisfying constraint assignments for each of the $mp(s)$ constraints with overlap $s \in \{0, 1, 2\}$, and finally choosing a bijection g with $(g, x) \in \mathcal{E}_{x,y}$.

5.3 Empirical overlap distributions

This section is dedicated to deriving properties of the set \mathcal{P}_n for $n \in \mathcal{N}$. In the following we will use the canonical ascending order on $\{0, 1, 2\}$ to denote points in $\mathbb{R}^{\{0,1,2\}}$ and the ascending lexicographical order on $\{0, 1\}^2$ to denote points in $\mathbb{R}^{\{0,1\}^2}$. Recall that $p^{(s)} \in \mathcal{P}(\{0, 1, 2\})$ given by $p^{(s)}(s) = 1$ for $s \in \{0, 1, 2\}$ denote the corners of the convex polytope $\mathcal{P}(\{0, 1, 2\})$ and further consider the vectors in $\mathbb{R}^{\{0,1,2\}}$

$$b_1 = (-d, d, 0)^t, \quad b_2 = (1, -2, 1)^t. \quad (4)$$

Finally, let

$$\mathcal{X} = \{x \in \mathbb{R}^2 : (b_1, b_2)x \geq -p^{(0)}\}, \quad \mathcal{X}_n = \mathcal{X} \cap (m^{-1}\mathbb{Z})^2.$$

Lemma 5.1. *The map $\iota_n : \mathcal{X}_n \rightarrow \mathcal{P}_n, x \mapsto p^{(0)} + (b_1, b_2)x$ is a bijection.*

Proof. We use the shorthands $1_{\{0,1,2\}} = (1)_{s \in \{0,1,2\}}$ and

$$1_{\{0,1,2\}}^\perp = \left\{ x \in \mathbb{R}^{\{0,1,2\}} : 1_{\{0,1,2\}}^\top x = 0 \right\} = \left\{ x \in \mathbb{R}^{\{0,1,2\}} : \sum_{s \in \{0,1,2\}} x_s = 0 \right\}.$$

Note that $\mathcal{P}(\{0, 1, 2\}) \subseteq p^{(0)} + 1_{\{0,1,2\}}^\perp = \{p^{(0)} + x : x \in 1_{\{0,1,2\}}^\perp\}$. On the other hand, (b_1, b_2) is a basis of $1_{\{0,1,2\}}^\perp$, and hence

$$\iota : \mathbb{R}^2 \rightarrow p^{(0)} + 1_{\{0,1,2\}}^\perp, \quad x \mapsto p^{(0)} + (b_1, b_2)x, \quad (5)$$

is bijective. This gives that $\iota(\mathcal{X}) = \mathcal{P}(\{0, 1, 2\})$ and that ι_n is the restriction of ι to \mathcal{X}_n , so ι_n is a bijection from \mathcal{X}_n to $\mathcal{P}(\{0, 1, 2\}) \cap \iota((m^{-1}\mathbb{Z})^2)$. Consequently, it remains to show that $\mathcal{P}_n = \mathcal{P}(\{0, 1, 2\}) \cap \iota((m^{-1}\mathbb{Z})^2)$, where

$$\iota((m^{-1}\mathbb{Z})^2) = \left\{ p^{(0)} + \frac{i_1}{m}b_1 + \frac{i_2}{m}b_2 : i \in \mathbb{Z}^2 \right\}$$

is a grid anchored at $p^{(0)}$ and spanned by $m^{-1}b_1$ and $m^{-1}b_2$. Note that

$$p_e^{(0)} = \left(\frac{k-4}{k}, \frac{2}{k}, \frac{2}{k}, 0 \right)^\top,$$

so $np_e^{(0)} \in \mathbb{Z}^{\{0,1\}^2}$ since $n \in \mathcal{N}$, and hence $p^{(0)} \in \mathcal{P}_n$ by the definition of \mathcal{P}_n . Next, we show that \mathcal{P}_n is on the grid, i.e. $\mathcal{P}_n \subseteq \iota((m^{-1}\mathbb{Z})^2)$. For this purpose fix $p \in \mathcal{P}_n$ and let $x = \iota^{-1}(p)$, i.e. $mp \in \mathbb{Z}^{\{0,1,2\}}$, $n(Wp) \in \mathbb{Z}^{\{0,1\}^2}$ and $p = p^{(0)} + x_1b_1 + x_2b_2$. This directly gives $mx_2 = mp(2) \in \mathbb{Z}$. Further, we notice that b_2 is in the kernel of W from Equation (3), i.e. $Wb_2 = 0_{\{0,1\}^2}$, and $Wb_1 = \frac{d}{k}w$ with $w = (1, -1, -1, 1)^\top$. This directly gives $p_e(1, 1) = 0 + \frac{d}{k}x_1 + 0$ and hence $mx_1 = np_e(1, 1) \in \mathbb{Z}$, i.e. $x \in (m^{-1}\mathbb{Z})^2$ and hence $p = \iota(x) \in \iota((m^{-1}\mathbb{Z})^2)$. Conversely, for any $x \in \mathcal{X}_n$ and with $p = \iota(x)$ we have $p \in \mathcal{P}(\{0, 1, 2\})$ since $x \in \mathcal{X}$, further $mp = mp^{(0)} + (b_1, b_2)(mx) \in \mathbb{Z}^{\{0,1,2\}}$ since $mx \in \mathbb{Z}^2$ and the other terms on the right-hand side are integer valued by definition, and finally $np_e = np_e^{(0)} + mx_1w \in \mathbb{Z}^{\{0,1\}^2}$. \square

Using Lemma 5.1 we have $\mathbb{E}[Z^{(2)}] = \sum_{x \in \mathcal{X}_n} E(\iota_n(x))$, where $\mathcal{X}_n \subseteq \mathbb{R}^2$ may be considered as a normalization of the grid $\mathcal{P}_n \subseteq p^{(0)} + 1_{\{0,1,2\}}^\perp$. In order to prepare the upcoming asymptotics of the second moment, we give a complete characterization of the convex polytope \mathcal{X} and the image of \mathcal{X} under $W(b_1, b_2)$, i.e. the image $p_e = Wp$ of $p \in \mathcal{P}(\{0, 1, 2\})$ under W from Equation (3). Let $w = (1, -1, -1, 1)^\top$ from the proof of Lemma 5.1, and set

$$\mathcal{W} = \left\{ p_e^{(0)} + yw : y \in [0, 2/k] \right\} \subseteq \mathcal{P}(\{0, 1\}^2), \quad \mathcal{X}_p = \left\{ x \in \mathcal{X} : x_1 = \frac{k}{d}p(1, 1) \right\} \text{ for } p \in \mathcal{W}.$$

Moreover, recall the definition of p^* from (2) and the bijection ι from (5), and let

$$x^{(0)} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \quad x^{(1)} = d^{-1} \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad x^{(2)} = d^{-1} \begin{pmatrix} 2 \\ d \end{pmatrix} \in \mathbb{R}^2 \text{ and } x^* = \iota^{-1}(p^*).$$

Lemma 5.2. *The set \mathcal{X} is a two-dimensional convex polytope with corners $x^{(0)}, x^{(1)}, x^{(2)}$, and x^* is in the interior of \mathcal{X} . The image of \mathcal{X} under $W(b_1, b_2)$ is the one-dimensional convex polytope \mathcal{W} with corners $p_e^{(0)}$ and $p_e^{(2)}$. Further, the preimage of $p \in \mathcal{W}$ under $W(b_1, b_2)$ is \mathcal{X}_p , where $\mathcal{X}_{p_e^{(s)}} = \{p^{(s)}\}$ for $s \in \{0, 2\}$ and the intersection of \mathcal{X}_p with the interior of \mathcal{X} is non-empty otherwise.*

Proof. Notice that $\iota(x^{(s)}) = p^{(s)}$ for $s \in \{0, 1, 2\}$, so since $\mathcal{P}(\{0, 1, 2\})$ is the convex hull of its corners $p^{(s)}$, $s \in \{0, 1, 2\}$, we have that \mathcal{X} is the convex hull of $x^{(s)}$, $s \in \{0, 1, 2\}$, i.e. a two-dimensional convex polytope with corners $x^{(s)}$, since ι is an affine transformation. In particular this also directly yields that x^* is in the interior of \mathcal{X} . Further, this shows that for any $x \in \mathcal{X}$ we have $x_1 \geq 0$ with equality iff $x = x^{(0)}$ and further $x_1 \leq \frac{2}{d}$ with equality iff $x = x^{(2)}$. Using $Wb_2 = 0_{\{0,1\}^2}$ and $Wb_1 = \frac{d}{k}w$ from the proof of Lemma 5.1 we directly get that

$$W(b_1, b_2)x = p_e^{(0)} + \frac{d}{k}x_1w \text{ with } \frac{d}{k}x_1 \in [0, 2/k],$$

hence the image of \mathcal{X} under $W(b_1, b_2)$ is a subset of \mathcal{W} . Conversely, for $y \in [0, 2/k]$ and $x = \frac{k}{2}yx^{(2)} \in \mathcal{X}$ we have $W(b_1, b_2)x = p_e^{(0)} + yw$, which shows that \mathcal{W} is the image of \mathcal{X} under $W(b_1, b_2)$. This also shows that \mathcal{X}_p is the preimage of $p \in \mathcal{W}$, since for $y \in [0, 2/k]$ and $p = p_e^{(0)} + yw$ we have $p(1, 1) = y$. This in turn directly yields that $\mathcal{X}_{p_e^{(s)}} = \{p^{(s)}\}$ for $s \in \{0, 2\}$. To see that \mathcal{X}_p contains interior points of \mathcal{X} otherwise, we can consider non-trivial convex combinations of x^* and $x^{(0)}$ for $\frac{k}{d}p(1, 1) < x_1^*$ and non-trivial convex combinations of x^* and $x^{(2)}$ for $\frac{k}{d}p(1, 1) > x_1^*$, which are points in the interior of \mathcal{X} . \square

Notice that in the two-dimensional case at hand, the proof of Lemma 5.2 is overly formal. The set \mathcal{X} is simply (the convex hull of) the triangle given by $x^{(s)}$, $s \in \{0, 1, 2\}$, with \mathcal{X}_p given by the vertical lines in \mathcal{X} with $x_1 = \frac{d}{k}p(1, 1)$. Further, the set \mathcal{X}_n is a canonical discretization of \mathcal{X} in that it is given by the points of the grid $(m^{-1}\mathbb{Z})^2$ contained in the triangle \mathcal{X} .

5.4 Proof of Lemma 2.4

We derive Lemma 2.4 from the following stronger assertion.

Lemma 5.3. *Let $\mathcal{U} \subseteq \mathcal{P}(\{0, 1, 2\})$ be a subset with non-empty interior and such that the closure of \mathcal{U} is contained in the interior of $\mathcal{P}(\{0, 1, 2\})$. Then there exists a constant $c = c(\mathcal{U}) \in \mathbb{R}_{>0}$ such that for all $n \in \mathcal{N}$ and all $p \in \mathcal{P}_n \cap \mathcal{U}$ we have $\tilde{E}(p)e^{-c/n} \leq E(p) \leq \tilde{E}(p)e^{c/n}$, where*

$$\tilde{E}(p) = \sqrt{\frac{d^3}{(2\pi)^2 m^2 \prod_s p(s)}} e^{n\phi_2(p)}.$$

Proof. Let \mathcal{C} denote the closure of \mathcal{U} and $\pi_s : \mathcal{C} \rightarrow [0, 1]$, $p \mapsto p(s)$ the projection for $s \in \{0, 1, 2\}$. Since \mathcal{C} is compact, the continuous map π_s attains its maximum $p_+(s)$ and its minimum $p_-(s)$, which directly gives $0 < p_-(s) < p_+(s) < 1$ since all $p \in \mathcal{C}$ are fully supported and the interior of \mathcal{C} is non-empty (that gives the second inequality). Using Lemma 5.2, the continuous map $\pi : \mathcal{C} \rightarrow [0, 2/k]$, $p \mapsto p_e(1, 1)$, and the same reasoning as above we obtain the maximum $p_{e,+}(1, 1)$ and minimum $p_{e,-}(1, 1)$ of π with $0 < p_{e,-}(1, 1) < p_{e,+}(1, 1) < 2/k$, which directly give the bounds $p_{e,-}(x), p_{e,+}(x) > 0$ for $x \in \{0, 1\}^2$ as functions of $p_{e,+}(1, 1)$ and $p_{e,-}(1, 1)$. Now, we can use these bounds with the Stirling bound to obtain a constant $c \in \mathbb{R}_{>0}$ such that for all $n \in \mathcal{N}$ and $p \in \mathcal{P}_n \cap \mathcal{C}$ we have $E'(p)e^{-c/n} \leq E(p) \leq E'(p)e^{c/n}$, where

$$\begin{aligned} E'(p) &= \sqrt{\frac{2\pi m d^3}{\prod_s (2\pi m p(s))}} \prod_{s \in \{0,1,2\}} \binom{k}{s, 2-s, 2-s, k-4+s}^{mp(s)} e^{mH(p) - (d-1)nH(p_e)} \\ &= \sqrt{\frac{d^3}{(2\pi)^2 m^2 \prod_s p(s)}} e^{2m \ln \binom{k}{2} - mD_{\text{KL}}(p \| p_e^*) - (d-1)nH(p_e)}. \end{aligned}$$

To see that $E'(p) = \tilde{E}(p)$, we observe that $D_{\text{KL}}(p_e \| p_e^*) = 2H(2/k) - H(p_e)$, since $H(p_e^*) = 2H(2/k)$, $p_e(1, 0) = p_e(0, 1)$, and $p_e(1, 1) + p_e(1, 0) = 2/k$ for any $p \in \mathcal{P}(\{0, 1, 2\})$. \square

Now, Lemma 2.4 is an immediate corollary. To see this, fix a fully supported overlap distribution $p \in \mathcal{P}(\{0, 1, 2\})$ and a sequence $(p_n)_{n \in \mathcal{N}} \subseteq \mathcal{P}_n$ converging to p , e.g. $p_n = \iota(m^{-1} \lfloor mx_1 \rfloor, m^{-1} \lfloor mx_2 \rfloor)$ with $\iota(x) = p$ and n sufficiently large. Further, fix a neighbourhood \mathcal{U} of p as described in Lemma 5.3, which is possible since p is fully supported. Then we have $p_n \in \mathcal{U}$ for sufficiently large n , hence with the continuity of ϕ_2 we have

$$\lim_{n \rightarrow \infty} \frac{1}{n} \ln(E(p_n)) = \lim_{n \rightarrow \infty} \frac{1}{n} \ln(\tilde{E}(p_n)) = \lim_{n \rightarrow \infty} \phi_2(p_n) = \phi_2(p).$$

5.5 Proof of Proposition 2.6

We postpone the proof of Proposition 2.5 and continue with Laplace's method for sums using the result. We obtain that

$$\frac{\mathbb{E}[Z^2]}{\mathbb{E}[Z]^2} = \sum_p e(p), \quad e(p) = \frac{E(p)}{\mathbb{E}[Z]^2} = \frac{\binom{2n/k}{np_e(1,1)} \binom{(k-2)n/k}{np_e(0,1)} \binom{dn}{2dn/k} \binom{m}{mp} \prod_s p^*(s)^{mp(s)}}{\binom{2dn/k}{dnp_e(1,1)} \binom{(k-2)dn/k}{dnp_e(0,1)} \binom{n}{2n/k}},$$

where the sum is over $p \in \mathcal{P}_n$. First, we use Proposition 2.5 to show that Laplace's method of sums is applicable. While we have already established that Δ_{d^*} is non-negative, we still need to ensure that p^* is the unique minimizer of Δ_d for $d < d^*$ and that the Hessian at p^* is positive definite. We will need the second order Taylor approximation of the KL divergence. To be specific, let μ^* have finite non-trivial support \mathcal{S} and let $f: \mathcal{P}(\mathcal{S}) \rightarrow \mathbb{R}_{\geq 0}$, $\mu \mapsto D_{\text{KL}}(\mu \parallel \mu^*)$, be the corresponding KL divergence. Then

$$f^{(2)}: \mathcal{P}(\mathcal{S}) \rightarrow \mathbb{R}_{\geq 0}, \quad \mu \mapsto \frac{1}{2} D_{\chi^2}(\mu \parallel \mu^*) = \frac{1}{2} \sum_s \frac{(\mu(s) - \mu^*(s))^2}{\mu^*(s)} = \frac{1}{2} (\mu - \mu^*)^t D_{\mu^*}^{-1} (\mu - \mu^*),$$

is the second order Taylor approximation of f at μ^* , where $D_{\chi^2}(\mu \parallel \mu^*)$ denotes Pearson's χ^2 divergence, $D_{\mu^*} = (\delta_{ij} \mu^*(i))_{i,j \in \mathcal{S}}$ the matrix with μ^* on the diagonal, and $\delta_{ij} = 1$ if $i = j$ and 0 otherwise; this can be easily seen by considering the extension of f to $\mathbb{R}_{\geq 0}^{\mathcal{S}}$. On the other hand, we would like to consider Δ_d as a function over the suitable domain \mathcal{X} from Section 5.3, however relative to the base point p^* . Hence, let $\mathcal{X}^* = \{x - x^*: x \in \mathcal{X}\}$ be the triangle \mathcal{X} centred at x^* instead of $x^{(0)}$, and $\iota^*: \mathcal{X}^* \rightarrow \mathcal{P}(\{0, 1, 2\})$ the bijection given by

$$\iota^*(x) = \iota(x + x^*) = p^{(0)} + (b_1, b_2)x + (b_1, b_2)x^* = \iota(x^*) + (b_1, b_2)x = p^* + (b_1, b_2)x$$

for $x \in \mathcal{X}^*$, with b_1, b_2 from Equation (4). Now, let $\gamma_d: \mathcal{X}^* \rightarrow \mathbb{R}_{\geq 0}$, $x \mapsto \Delta_d(\iota^*(x))$, denote the corresponding parametrization of Δ_d . Then, using the chain rule for multivariate calculus as indicated above for both (b_1, b_2) and W from (3), we derive the Hessian

$$H_d = (b_1, b_2)^t \left(D_{p^*}^{-1} - \frac{(d-1)k}{d} W^t D_{p_e^*}^{-1} W \right) (b_1, b_2) \quad (6)$$

of γ_d at $0_{[2]} \in \mathbb{R}^2$, using the shorthand $D_{\mu^*} = (\delta_{ij} \mu^*(i))_{i,j}$. The properties of the KL divergence imply that $\gamma_d(0_{[2]}) = 0$ and γ_d has a stationary point at $0_{[2]}$. Now, the second order Taylor approximation $\gamma_d^{(2)}: \mathcal{X}^* \rightarrow \mathbb{R}$, $x \mapsto \frac{1}{2} x^t H_d x$, of γ_d at $0_{[2]}$ can be written as $\gamma_d^{(2)} = \Delta_d^{(2)} \circ \iota^*$ with

$$\Delta_d^{(2)}(p) = \frac{1}{2} \left[D_{\chi^2}(p \parallel p^*) - \frac{(d-1)k}{d} D_{\chi^2}(p_e \parallel p_e^*) \right]. \quad (7)$$

Further, for any neighbourhood \mathcal{U} of $0_{[2]}$ such that the closure of \mathcal{U} is contained in the interior of \mathcal{X}^* , Taylor's theorem yields a constant $c \in \mathbb{R}_{>0}$ such that

$$\gamma_d^{(2)}(x) - c \|x\|_2^3 \leq \gamma_d(x) \leq \gamma_d^{(2)}(x) + c \|x\|_2^3 \quad (8)$$

for all $x \in \mathcal{U}$. Since H_d is symmetric, let $\lambda_1, \lambda_2 \in \mathbb{R}$ with $\lambda_1 \leq \lambda_2$ denote its eigenvalues and fix a corresponding orthonormal basis of eigenvectors $v_1, v_2 \in \mathbb{R}^2$, i.e. $H_d v_1 = \lambda_1 v_1$ and $H_d v_2 = \lambda_2 v_2$.

Formally and analogously to the KL divergence we will take the liberty to identify Δ_d and $\Delta_d^{(2)}$ with their extensions to the maximal domains $\mathcal{D} \subseteq \mathbb{R}^{\{0,1,2\}}$ and $\mathcal{D}^{(2)} = \mathbb{R}^{\{0,1,2\}}$ respectively. In particular, Lemma 5.2 shows that for any fully supported $p \in \mathcal{P}(\{0, 1, 2\})$ the edge distribution p_e also has full support, hence we can use the Lipschitz continuity of W on $\mathbb{R}^{\{0,1,2\}}$ to find $\varepsilon \in \mathbb{R}_{>0}$ such that both $p' > 0$ and $Wp' > 0$ for any $p' \in \mathcal{B}_\varepsilon(p) \subseteq \mathbb{R}_{>0}^{\{0,1,2\}}$ and thereby Δ_d is well-defined and smooth on $\mathcal{B}_\varepsilon(p)$.

Lemma 5.4. *Let $k \in \mathbb{Z}_{\geq 4}$ and $d \in (0, d^*)$. Then the unique minimizer of γ_d is $0_{[2]}$ and H_d is positive definite.*

Proof. Using Proposition 2.5 we know that H_{d^*} is positive semidefinite since $0_{[2]}$ is a global minimum of γ_{d^*} . This in turn yields that $\gamma_{d^*}^{(2)} \geq 0$ or equivalently $\Delta_{d^*}^{(2)} \geq 0$. Now, for any $d < d^*$ the unique minimizer of Δ_d is p^* since $\Delta_d(p^*) = 0$, further $\Delta_d(p) > 0$ for any $p \neq p^*$ with $p_e = p_e^*$ and

$$\Delta_d(p) = D_{\text{KL}}(p \parallel p^*) - \left(1 - \frac{1}{d}\right) k D_{\text{KL}}(p_e \parallel p_e^*) > \Delta_{d^*}(p) \geq 0$$

for any p with $p_e \neq p_e^*$. But the same argumentation shows that p^* is the unique minimizer of $\Delta_d^{(2)}$, since $D_{\chi^2}(\mu \parallel \mu^*)$ is also minimal with value 0 iff $\mu = \mu^*$. This in turn shows that $\gamma_d^{(2)}$ is uniquely minimized at $0_{[2]}$ and hence H_d is positive definite. \square

Let $\eta_{\text{KL}} = \sup_{p \neq p^*} \frac{D_{\text{KL}}(p_e \parallel p_e^*)}{D_{\text{KL}}(p \parallel p^*)}$ denote the contraction coefficient with respect to the KL divergence. Notice that by Proposition 2.5 we have $\frac{d^*}{(d^*-1)k} \geq \frac{D_{\text{KL}}(p_e \parallel p_e^*)}{D_{\text{KL}}(p \parallel p^*)}$ for all $p \neq p^*$ with equality for $p = p^{(2)}$, hence $\eta_{\text{KL}} = \frac{d^*}{(d^*-1)k}$ (so Proposition 2.5 indeed confirms the conjecture by the authors in [36]). Further, let $\eta_{\chi^2} = \sup_{p \neq p^*} \frac{D_{\chi^2}(p_e \parallel p_e^*)}{D_{\chi^2}(p \parallel p^*)}$ denote the contraction coefficient with respect to the χ^2 divergence. The proof of Lemma 5.4 suggests that $\eta_{\chi^2} \leq \eta_{\text{KL}}$, a result known from the literature.

In the rest of this section we discuss the straightforward (but cumbersome) application of Laplace's method for sums. For convenience, we first show that the boundaries can be neglected and derive the asymptotics of the sum on the interior using the uniform convergence established in Lemma 5.3.

Lemma 5.5. *Let $d \in (0, d^*)$ and let \mathcal{U} be a neighbourhood of p^* such that its closure is contained in the interior of $\mathcal{P}(\{0, 1, 2\})$. Then*

$$\frac{\mathbb{E}[Z^2]}{\mathbb{E}[Z]^2} = \sum_{p \in \mathcal{P}_n} e(p) \sim \sum_{p \in \mathcal{P}_n \cap \mathcal{U}} \sqrt{\frac{d}{(2\pi)^2 m^2 \prod_s p(s)}} e^{-m\Delta_d(p)}.$$

Proof. Let $\Delta_{\min} > 0$ denote the global minimum of Δ_d on $\mathcal{P}(\{0, 1, 2\}) \setminus \mathcal{U}$. Now, we can use the well-known bounds $\frac{1}{a+1} \exp(aH(\frac{b}{a})) \leq \binom{a}{b} \leq \exp(aH(\frac{b}{a}))$ for binomial coefficients and the corresponding upper bound for multinomial coefficients (using the entropy of the distribution determined by the weights $\frac{b_i}{a}$) to derive

$$\sum_{p \notin \mathcal{U}} e(p) \leq \rho(n) \sum_{p \notin \mathcal{U}} e^{-m\Delta_d(p)} \leq \rho(n) e^{-m\Delta_{\min}} |\mathcal{P}_m(\{0, 1, 2\})| = \rho(n) \binom{m+1}{2} e^{-m\Delta_{\min}}, \text{ where}$$

$$\rho(n) = (n+1) \left(\frac{2dn}{k} + 1 \right) \left(\frac{(k-2)dn}{k} + 1 \right).$$

Here, we used the form of $e(p)$ introduced at the beginning of this section and further notice that the bounds used are tight for the log-densities, i.e. the exponent is $\Delta_d(p)$ by the computations in Section 5.4. The right-hand side vanishes for n tending to infinity, hence we have

$$\frac{\mathbb{E}[Z^2]}{\mathbb{E}[Z]^2} \sim \sum_{p \in \mathcal{U}} e(p).$$

Now, the result directly follows using Lemma 5.3 and Lemma 2.2. \square

Lemma 5.5 shows that the overlap distributions p with material contributions $e(p)$ to the second moment are concentrated around p^* . Hence, instead of considering a fixed neighbourhood \mathcal{U} of p^* we consider a sequence $(\mathcal{U}_n)_{n \in \mathcal{N}}$ of decreasing neighbourhoods. First, we choose a scaling that improves the assertion of Lemma 5.5 and further allows to simplify the asymptotics of the right-hand side, in the sense that the leading factor collapses to a constant and $\gamma_d(x) = \Delta_d(\iota^*(x))$ can be replaced by its second order Taylor approximation $\gamma_d^{(2)}(x) = \Delta_d^{(2)}(\iota^*(x)) = \frac{1}{2}x^t H_d x$ from above (7). For this purpose let $\mathcal{U}^* \subseteq \mathcal{X}^*$ be a sufficiently small neighbourhood of $0_{[2]}$ (in particular bounded away from the boundary of \mathcal{X}^*), further

$$\mathcal{U}_n^* = \left\{ x \in \mathcal{X}^* : \|x\|_2 < \frac{\ln(m)}{\sqrt{m}} \right\} \text{ and } \mathcal{X}_n^* = \{x - x^* : x \in \mathcal{X}_n\} \cap \mathcal{U}_n^* \text{ for } n \in \mathcal{N}.$$

In the following we restrict to $n \geq n_0$ where $n_0 \in \mathcal{N}$ is such that $\mathcal{U}_{n_0}^* \subseteq \mathcal{U}^*$.

Lemma 5.6. For $d \in (0, d^*)$ we have

$$\frac{\mathbb{E}[Z^2]}{\mathbb{E}[Z]^2} \sim \sqrt{\frac{d}{(2\pi)^2 m^2 \prod_s p^*(s)}} \sum_{x \in \mathcal{X}_n^*} e^{-\frac{m}{2} x^t H_d x}.$$

Proof. First, notice that we can apply Lemma 5.5 to $\iota^*(\mathcal{U}^*)$. So, we need to show that the sum over $\mathcal{U}^* \setminus \mathcal{U}_n^*$ is negligible. Then we proceed to derive the asymptotics of the sum over \mathcal{U}_n^* . Obviously, we have $\gamma_{\min}(n) \rightarrow 0$ for $n \rightarrow \infty$ with $\gamma_{\min}(n) = \min_{x \notin \mathcal{U}_n^*} \gamma_d(x) > 0$, since $\gamma_d(x) = \Delta_d(\iota^*(x))$ is continuous and $\gamma_d(0_{[2]}) = 0$. The main objective of the proof is to show that $\gamma_{\min}(n)$ converges to zero sufficiently slow. But with $\gamma_d^{(2)}(x) = \frac{1}{2}x^t H_d x$ from above (7) and for any $x \in \mathbb{R}^2$ we have

$$\gamma_d^{(2)}((v_1, v_2)x) = \frac{1}{2}(\lambda_1 x_1^2 + \lambda_2 x_2^2) \geq \frac{\lambda_1}{2} \|x\|_2^2 = \frac{\lambda_1}{2} \|(v_1, v_2)x\|_2^2$$

since (v_1, v_2) is an orthonormal basis, so $\gamma_d^{(2)}(x) \geq \frac{\lambda_1}{2} \|x\|_2^2$ for all $x \in \mathbb{R}^2$. Now, for any sufficiently small $\varepsilon \in (0, 1)$ let $c \in \mathbb{R}_{>0}$ be the constant for $\mathcal{B}_\varepsilon(0_{[2]})$ from Taylor's theorem applied to γ_d at $0_{[2]}$, then for any $x \in \mathcal{U}^* = \mathcal{B}_\varepsilon(0_{[2]}) \cap \mathcal{B}_\delta(0_{[2]})$, with $\delta = \frac{\varepsilon \lambda_1}{2c}$, we have $\gamma_d(x) \geq (1 - \varepsilon)\gamma_d^{(2)}(x)$ since

$$\gamma_d(x) - (1 - \varepsilon)\gamma_d^{(2)}(x) \geq \varepsilon\gamma_d^{(2)}(x) - c\|x\|_2^3 \geq \left(\frac{\varepsilon \lambda_1}{2} - c\delta\right) \|x\|_2^2 = 0.$$

In combination we have $\gamma_d(x) \geq \frac{(1-\varepsilon)\lambda_1}{2} \|x\|_2^2$ and using $p = \iota^*(x)$ hence

$$\lim_{n \rightarrow \infty} \sum_{x \notin \mathcal{U}_n^*} e(p) = \lim_{n \rightarrow \infty} \sum_{x \in \mathcal{U}^* \setminus \mathcal{U}_n^*} \sqrt{\frac{d}{(2\pi)^2 m^2 \prod_s p(s)}} e^{-m\gamma_d(x)} \leq \lim_{n \rightarrow \infty} C m e^{-\frac{(1-\varepsilon)\lambda_1}{2} \ln(m)^2} = 0,$$

by using (the proof of) Lemma 5.5 and some sufficiently large constant C . With this we have

$$\frac{\mathbb{E}[Z^2]}{\mathbb{E}[Z]^2} \sim \sum_{x \in \mathcal{X}_n^*} e(\iota^*(x)) \sim \sqrt{\frac{d}{(2\pi)^2 m^2 \prod_s p^*(s)}} \sum_{x \in \mathcal{X}_n^*} e^{-m\gamma_d^{(2)}(x)},$$

where the last equivalence follows from the fact that the leading factor converges to the respective constant uniformly on \mathcal{U}_n^* and by (8) on \mathcal{U}^* . \square

Lemma 5.6 completes the analytical part of the proof. For the last, measure theoretic, part we recall the bijection ι_n from Lemma 5.1. The translation of the sum on the right-hand side of Lemma 5.6 into a Riemann sum and further into the integral $\int g_\infty(x)dx$, where

$$g_\infty: \mathbb{R}^2 \rightarrow \mathbb{R}_{>0}, \quad y \mapsto \sqrt{\frac{d}{(2\pi)^2 \prod_s p^*(s)}} \exp\left(-\frac{1}{2}y^t H_d y\right),$$

is essentially given by the grid $\mathcal{X}_n \subseteq (m^{-1}\mathbb{Z})^2 \subseteq \mathbb{R}^2$. We make this rigorous in the following.

Lemma 5.7. *We have*

$$\frac{\mathbb{E}[Z^2]}{\mathbb{E}[Z]^2} \sim \int g_\infty(x)dx.$$

Proof. We start with the partition of \mathbb{R}^2 into the squares

$$\mathcal{Q}_{n,x} = \left\{ x + \alpha_1 \begin{pmatrix} 1 \\ 0 \end{pmatrix} + \alpha_2 \begin{pmatrix} 0 \\ 1 \end{pmatrix} : \alpha \in \left[-\frac{1}{2m}, \frac{1}{2m} \right]^2 \right\}, x \in (m^{-1}\mathbb{Z})^2.$$

Next, we need a suitable selection of squares to cover the disc

$$x^* + \mathcal{U}_n^* = \left\{ x^* + x : x \in \mathbb{R}^2, \|x\|_2 < \frac{\ln(m)}{\sqrt{m}} \right\} \subseteq \mathbb{R}^2$$

corresponding to the disc \mathcal{U}_n^* . For this purpose let $x_{\min}, x_{\max} \in (m^{-1}\mathbb{Z})^2$ be given by

$$\begin{aligned} x_{\min,1} &= m^{-1} \left\lfloor m \left(x_1^* - \frac{\ln(m)}{\sqrt{m}} \right) \right\rfloor, x_{\min,2} = m^{-1} \left\lfloor m \left(x_2^* - \frac{\ln(m)}{\sqrt{m}} \right) \right\rfloor, \\ x_{\max,1} &= m^{-1} \left\lceil m \left(x_1^* + \frac{\ln(m)}{\sqrt{m}} \right) \right\rceil, x_{\max,2} = m^{-1} \left\lceil m \left(x_2^* + \frac{\ln(m)}{\sqrt{m}} \right) \right\rceil. \end{aligned}$$

Further, let $\mathcal{G}_n = (m^{-1}\mathbb{Z})^2 \cap ([x_{\min,1}, x_{\max,1}] \times [x_{\min,2}, x_{\max,2}])$. By the definition of x_{\min} and x_{\max} the points on the boundary are not in $x^* + \mathcal{U}_n^*$, which ensures that $x^* + \mathcal{U}_n^* \subseteq \mathcal{Q}_n$ with $\mathcal{Q}_n = \bigcup_{x \in \mathcal{G}_n} \mathcal{Q}_{n,x}$. Further, we have $\mathcal{Q}_- \subseteq \mathcal{Q}_n \subseteq \mathcal{Q}_+$ with

$$\mathcal{Q}_- = \left\{ x \in \mathbb{R}^2 : \|x - x^*\|_\infty \leq \frac{\ln(m)}{\sqrt{m}} \right\}, \mathcal{Q}_+ = \left\{ x \in \mathbb{R}^2 : \|x - x^*\|_\infty \leq \frac{\ln(m)}{\sqrt{m}} + \frac{3}{2m} \right\},$$

which ensures that $\mathcal{Q}_n \subseteq \mathcal{X}$ for $n \in \mathcal{N}$ sufficiently large. Now, we translate the notions back to \mathcal{X}^* using the bijection $\tau: \mathbb{R}^2 \rightarrow \mathbb{R}^2, x \mapsto x - x^*$, i.e. let $\mathcal{G}_n^* = \tau(\mathcal{G}_n)$, $\mathcal{Q}_{n,x}^* = \tau(\mathcal{Q}_{n,\tau^{-1}(x)})$ for $x \in \mathcal{G}_n^*$, $\mathcal{Q}_n^* = \tau(\mathcal{Q}_n)$, $\mathcal{Q}_-^* = \tau(\mathcal{Q}_-)$ and $\mathcal{Q}_+^* = \tau(\mathcal{Q}_+)$. This directly gives

$$\begin{aligned} \mathcal{Q}_{n,x}^* &= \left\{ x + \alpha_1 \begin{pmatrix} 1 \\ 0 \end{pmatrix} + \alpha_2 \begin{pmatrix} 0 \\ 1 \end{pmatrix} : \alpha \in \left[-\frac{1}{2m}, \frac{1}{2m} \right]^2 \right\}, x \in \mathcal{G}_n^*, \mathcal{Q}_n^* = \bigcup_{x \in \mathcal{G}_n^*} \mathcal{Q}_{n,x}^*, \\ \mathcal{Q}_-^* &= \left\{ x \in \mathbb{R}^2 : \|x\|_\infty \leq \frac{\ln(m)}{\sqrt{m}} \right\}, \mathcal{Q}_+^* = \left\{ x \in \mathbb{R}^2 : \|x\|_\infty \leq \frac{\ln(m)}{\sqrt{m}} + \frac{3}{2m} \right\}, \end{aligned}$$

and $\mathcal{U}_n^* \subseteq \mathcal{Q}_-^* \subseteq \mathcal{Q}_n^* \subseteq \mathcal{Q}_+^* \subseteq \mathcal{X}^*$ for $n \in \mathcal{N}$ sufficiently large. Further, with Lemma 5.6 and the definition of $\gamma_d^{(2)}$ we now have

$$\frac{\mathbb{E}[Z^2]}{\mathbb{E}[Z]^2} \sim \sum_{x \in \mathcal{G}_n^*} \sqrt{\frac{d}{(2\pi)^2 m^2 \prod_s p^*(s)}} \exp\left(-\frac{1}{2} m x^t H_d x\right).$$

Finally, we need to adjust the scaling to turn the sum on the right-hand side into a Riemann sum. For this purpose let $\sigma: \mathbb{R}^2 \rightarrow \mathbb{R}^2$, $x \mapsto \sqrt{m}x$, further $\mathcal{G}'_n = \sigma(\mathcal{G}_n^*)$, $\mathcal{Q}'_{n,x} = \sigma(\mathcal{Q}_{n,\sigma^{-1}(x)}^*)$ for $x \in \mathcal{G}'_n$, $\mathcal{Q}'_- = \sigma(\mathcal{Q}_-^*)$, $\mathcal{Q}'_+ = \sigma(\mathcal{Q}_+^*)$. This directly gives

$$\mathcal{Q}'_{n,x} = \left\{ x + \alpha_1 \begin{pmatrix} 1 \\ 0 \end{pmatrix} + \alpha_2 \begin{pmatrix} 0 \\ 1 \end{pmatrix} : \alpha \in \left[-\frac{1}{2\sqrt{m}}, \frac{1}{2\sqrt{m}} \right]^2 \right\}, x \in \mathcal{G}'_n, \mathcal{Q}'_n = \bigcup_{x \in \mathcal{G}'_n} \mathcal{Q}'_{n,x},$$

$$\mathcal{Q}'_- = \{x \in \mathbb{R}^2 : \|x\|_\infty \leq \ln(m)\}, \mathcal{Q}'_+ = \left\{ x \in \mathbb{R}^2 : \|x\|_\infty \leq \ln(m) + \frac{3}{2\sqrt{m}} \right\},$$

and $\mathcal{Q}'_- \subseteq \mathcal{Q}'_n \subseteq \mathcal{Q}'_+$. Using that $m x^t H_d x = \sigma(x)^t H_d \sigma(x)$ for all $x \in \mathcal{G}_n^*$ and further that the area of $\mathcal{Q}'_{n,x}$ is m^{-1} for all $x \in \mathcal{G}'_n$ we have

$$\frac{\mathbb{E}[Z^2]}{\mathbb{E}[Z]^2} \sim \sum_{x \in \mathcal{G}'_n} \sqrt{\frac{d}{(2\pi)^2 m^2 \prod_s p^*(s)}} \exp\left(-\frac{1}{2} x^t H_d x\right) = \int g_n(y) dy,$$

$$g_n(y) = \sum_{x \in \mathcal{G}'_n} \mathbb{1}\{y \in \mathcal{Q}'_{n,x}\} \sqrt{\frac{d}{(2\pi)^2 \prod_s p^*(s)}} \exp\left(-\frac{1}{2} x^t H_d x\right), y \in \mathbb{R}^2.$$

In order to show that $\int g_n(y) dy$ converges to $\int g_\infty(y) dy$ we recall from Lemma 5.4 that H_d is positive definite, which ensures that $\int g_\infty(y) dy$ exists and is finite. Now, using Taylor's theorem with order 0 and the Lagrange form of the first order remainder with the fact that the absolutes of the first derivatives of g_∞ are bounded from above yields a constant $c \in \mathbb{R}_{>0}$ such that for all $n \in \mathcal{N}$ and all $y \in \mathcal{Q}'_n$, with $x \in \mathcal{G}'_n$ such that $y \in \mathcal{Q}'_{n,x}$, we have

$$\|g_\infty(y) - g_n(y)\|_\infty = \|g_\infty(y) - g_\infty(x)\|_\infty \leq \frac{c}{\sqrt{m}}.$$

This bound directly suggests that

$$\left| \int \mathbb{1}\{y \in \mathcal{Q}'_{n,x}\} g_\infty(y) dy - \int \mathbb{1}\{y \in \mathcal{Q}'_{n,x}\} g_n(y) dy \right| \leq c m^{-\frac{3}{2}},$$

$$\left| \int \mathbb{1}\{y \in \mathcal{Q}'_n\} g_\infty(y) dy - \int \mathbb{1}\{y \in \mathcal{Q}'_n\} g_n(y) dy \right| \leq \frac{c}{\sqrt{m}} \left(2 \ln(m) + \frac{3}{\sqrt{m}} \right)^2 \text{ and}$$

$$\left| \int g_\infty(y) dy - \int g_n(y) dy \right| \leq \frac{c}{\sqrt{m}} \left(2 \ln(m) + \frac{3}{\sqrt{m}} \right)^2 + \int \mathbb{1}\{y \notin \mathcal{Q}'_n\} g_\infty(y) dy.$$

In particular the last bound suggests that $\int g_n(y) dy \rightarrow \int g_\infty(y) dy$ since the error on the right-hand side tends to zero as n tends to infinity. \square

The only remaining part of the proof is to compute $\int g_\infty(x) dx$.

Lemma 5.8. We have $\int g_\infty(x) dx = \sqrt{\frac{k-1}{k-d}}$.

Proof. The Gaussian integral gives

$$\int g_{\infty}(x)dx = \sqrt{\frac{d}{(2\pi)^2 \prod_s p^*(s)}} \sqrt{\frac{(2\pi)^2}{\det(H_d)}} = \sqrt{\frac{d}{\det(H_d) \prod_s p^*(s)}}.$$

Recall that the Hessian $H_d \in \mathbb{R}^{2 \times 2}$ from (6) is a 2×2 matrix and all entries are given explicitly, so a straightforward calculation asserts that $\det(H_d) = (k-d)d/[(k-1) \prod_s p^*(s)]$. \square

Finally, combining Lemma 5.7 with Lemma 5.8 completes the proof of Proposition 2.6.

5.6 Proof of Proposition 2.5

This section is dedicated to the identification of the minimizers of Δ_d . First, we show that the stationary points of Δ_d including their characterization are in one to one correspondence with the fixed points of a ratio of belief propagation messages for the constraint satisfaction problem defined in Section 5.1, cf. Lemmas 5.9 to 5.12 below. Considering the ratio turns out to be sufficient and allows to consider a one-dimensional problem. For more background on the correspondence of stationary points of Δ_d and fixed points of belief propagation we refer the reader to [31, 41].

The major advantage of this equivalent fixed point problem, cf. Equation (9) below, is that the base function ι_{BP}^* is a simple rational function and the target function ι_{BP} is obtained from ι_{BP}^* by a single exponentiation. Thus, in the second step we discuss ι_{BP}^* and introduce a piecewise linear bound on ι_{BP}^* on $[1, k]$ given by g_1^* and g_k^* , cf. Lemma 5.13 below. The piecewise function given by g_1^* and g_k^* on $[1, k]$ and ι_{BP}^* on $[k, \infty)$ is increasing and piecewise convex, which allows to establish the unique fixed point of ι_{BP} on $(1, \infty)$ (and for reasonable choices of d) in Lemmas 5.16 and 5.17. For symmetry reasons this also completes the discussion for $k = 4$.

In the last step, we establish that ι_{BP} has no fixed points on $(0, 1)$, for reasonable d and $k > 4$, cf. Lemma 5.19 below. By combining the results we obtain that ι_{BP} has two fixed points (three for $k = 4$), one at 1 corresponding to the desired interior minimum of Δ_d and one in $[k, \infty)$ corresponding to a saddle point of Δ_d (two for $k = 4$, and for reasonable d). This allows to identify the two minimizers of Δ_{d^*} (three for $k = 4$) and completes the proof.

We begin with characterizing the stationary points of Δ_d for any $d \in \mathbb{R}_{>0}$.

In order to pin them down we first determine the stationary points of the restriction of Δ_d to overlap distributions with the same fixed edge distribution. For this purpose, recall the line $\mathcal{W} \subseteq \mathcal{P}(\{0, 1\}^2)$ of attainable edge distributions and the lines $\mathcal{P}_q = \iota(\mathcal{X}_q) = \{p \in \mathcal{P}(\{0, 1, 2\}) : p_e = q\}$ of overlap distributions with fixed edge distribution $q \in \mathcal{W}$ from Lemma 5.2. Further, let $\Delta_{d,q} : \mathcal{P}_q \rightarrow \mathbb{R}$ denote the restriction of Δ_d to \mathcal{P}_q . For $x \in \mathbb{R}_{>0}$ let $p_x \in \mathcal{P}(\{0, 1, 2\})$ be given by $p_x(s) = p^*(s)x^s / \sum_s p^*(s)x^s$, $s \in \{0, 1, 2\}$, further let $p_0 = p^{(0)}$, $p_{\infty} = p^{(2)}$, and $\mathcal{P}_{\min} = \{p_x : x \in [0, \infty]\}$. Finally, let $\iota_{\text{rp}} : [0, \infty] \rightarrow \mathcal{P}_{\min}$, $x \mapsto p_x$, denote the induced map and $\iota_{\text{pe}} : \mathcal{P}_{\min} \rightarrow \mathcal{W}$, $p \mapsto p_e$, the corresponding edge distributions.

Lemma 5.9. *For all $q \in \mathcal{W} \setminus \{p_e^{(0)}, p_e^{(2)}\}$ the map $\Delta_{d,q}$ has a unique stationary point $p_q \in \mathcal{P}_q$ that is a global minimum. The unique global minimizer of $\Delta_{d,p_e^{(s)}}$ is $p_{p_e^{(s)}} = p^{(s)}$ for $s \in \{0, 2\}$. Further, we have $\mathcal{P}_{\min} = \{p_q : q \in \mathcal{W}\}$ and the maps ι_{rp} , ι_{pe} are bijections.*

Proof. Recall from Lemma 5.2 that \mathcal{P}_q is one-dimensional for $q \in \mathcal{W} \setminus \{p_e^{(0)}, p_e^{(2)}\}$. Further, the map $\Delta_{d,q}$ is strictly convex since the KL divergence $D_{\text{KL}}(p \parallel p^*)$ (respectively $x \ln(x)$) is and further $D_{\text{KL}}(p_e \parallel p_e^*) = D_{\text{KL}}(q \parallel p_e^*)$ is constant. Now, fix an interior point $p_o \in \mathcal{P}_q$ and let a boundary point $p_b \in \mathcal{P}_q$ be given. Then p_b is not fully supported since it is on the boundary of $\mathcal{P}(\{0, 1, 2\})$ and hence the derivative of $D_{\text{KL}}(\alpha p_o + (1-\alpha)p_b \parallel p^*)$ tends to $-\infty$ as α tends to 0, which shows that $\Delta_{d,q}$ is not minimized on the boundary. Hence, we know that there exists exactly one stationary point $p_q \in \mathcal{P}_q$ and that $\Delta_{d,q}(p)$ is minimal iff $p = p_q$. As discussed in Lemma 5.2 we have

$\mathcal{P}_q = \{p^{(s)}\}$ for $q = p_e^{(s)}$ and $s \in \{0, 2\}$, so $p_q = p^{(s)}$ is obviously the unique global minimizer of $\Delta_{d,q}$ in this case and further $\Delta_{d,q}$ has no stationary points (since \mathcal{P}_q has empty interior). This shows that the map $q \mapsto p_q$ for $q \in \mathcal{W}$ is a bijection.

Further, for q in the interior of \mathcal{W} the stationary point p_q is fully supported and the unique root of the first derivative of $\Delta_{d,q}$ in the direction b_2 from (4), i.e.

$$\ln \left(\frac{p_q(0)}{p^*(0)} \right) + \ln \left(\frac{p_q(2)}{p^*(2)} \right) = 2 \ln \left(\frac{p_q(1)}{p^*(1)} \right) \text{ or equivalently } \frac{p_q(2)/p^*(2)}{p_q(1)/p^*(1)} = \frac{p_q(1)/p^*(1)}{p_q(0)/p^*(0)}.$$

Let \mathcal{P}'_{\min} denote the set of all fully supported $p \in \mathcal{P}(\{0, 1, 2\})$ satisfying $\frac{p(2)/p^*(2)}{p(1)/p^*(1)} = \frac{p(1)/p^*(1)}{p(0)/p^*(0)}$, i.e. our set of candidates for stationary points. Now, for $p \in \mathcal{P}'_{\min}$ let $q = p_e$, then we obviously have $p \in \mathcal{P}_q$ and p is a root of the first derivative of $\Delta_{d,q}$ in the direction b_2 , so p is the unique root and $p = p_q$. Hence, the map $\iota'_{pe} : \mathcal{P}'_{\min} \rightarrow \mathcal{W}$, $p \mapsto p_e$, is a bijection (up to the corners of \mathcal{W}) with inverse $q \mapsto p_q$. Now, let $\iota_{pr} : \mathcal{P}'_{\min} \rightarrow \mathbb{R}_{>0}$, $p \mapsto x_p$, with $x_p = \frac{p(1)p^*(0)}{p^*(1)p(0)}$. Notice that ι_{pr} is surjective since for any $x \in \mathbb{R}_{>0}$ we have

$$\frac{p_x(2)/p^*(2)}{p_x(1)/p^*(1)} = \frac{x^2}{x} = x = \frac{p_x(1)/p^*(1)}{p_x(0)/p^*(0)}$$

and hence $p_x \in \mathcal{P}'_{\min}$. To show that ι_{pr} is injective let $p \in \mathcal{P}'_{\min}$ and $x = x_p$. Using the definition of x_p and the defining property of \mathcal{P}'_{\min} we get

$$p(0) = p^*(0) \frac{p(0)}{p^*(0)}, p(1) = p^*(1)x \frac{p(0)}{p^*(0)}, p(2) = p^*(2)x \frac{p(1)}{p^*(1)} = p^*(2)x^2 \frac{p(0)}{p^*(0)}, \text{ so}$$

$$p(s) = \frac{p(s)}{p(0) + p(1) + p(2)} = \frac{p^*(s)x^s}{\sum_s p^*(s)x^s} = p_x(s), s \in \{0, 1, 2\}.$$

This shows that $\mathcal{P}_{\min} = \mathcal{P}'_{\min} \cup \{p_0, p_\infty\}$, that ι_{rp} is a bijection with inverse ι_{pr} (canonically extended to the endpoints), and finally that $\iota_{pe} = \iota'_{pe}$ is a bijection as well. \square

Lemma 5.9 has a few immediate consequences. For one, the only minimizers of Δ_d in the direction b_2 , from (4), on the boundary are $p^{(0)}$ and $p^{(2)}$, while all other boundary points are maximizers in the direction b_2 , hence if p is a global minimizer of Δ_d on the boundary, we have $p \in \{p^{(0)}, p^{(2)}\}$. Further, all stationary points of Δ_d are either local minima or saddle points. Finally, we have $p \in \mathcal{P}_{\min}$ for any stationary point $p \in \mathcal{P}(\{0, 1, 2\})$ of Δ_d since then also the derivative in the direction of b_2 vanishes.

For the upcoming characterization of the stationary points of Δ_d let

$$\iota_{rr} : \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}, \quad \iota_{rr}(x) = \iota_{rr}^*(x)^{\frac{d-1}{d}}, \quad \iota_{rr}^*(x) = \frac{p_{x,e}(1, 1)p_{x,e}(0, 0)}{p_{x,e}(1, 0)p_{x,e}(0, 1)}.$$

Notice that $\iota_{rr}(x) \in \mathbb{R}_{>0}$ for $x \in \mathbb{R}_{>0}$ since then p_x is fully supported and hence $p_{x,e}$ is fully supported by Lemma 5.2. Finally, let $\mathcal{X}_{st} = \{x \in \mathbb{R}_{>0} : \iota_{rr}(x) = x\}$ denote the fixed points of ι_{rr} and $\mathcal{P}_{st} = \{p_x : x \in \mathcal{X}_{st}\}$ the corresponding distributions. Notice that $p_x = p^*$ for $x = 1$ and further $\iota_{rr}^*(1) = 1$, i.e. $\iota_{rr}(1) = 1$ for all $d \in \mathbb{R}_{>0}$, hence $1 \in \mathcal{X}_{st}$ and $p^* \in \mathcal{P}_{st}$ for all $d \in \mathbb{R}_{>0}$.

Lemma 5.10. *The stationary points of Δ_d are given by \mathcal{P}_{st} .*

Proof. Using Lemma 5.9, a fully supported distribution $p \in \mathcal{P}(\{0, 1, 2\})$ is a stationary point of Δ_d iff there exists $x \in \mathbb{R}_{>0}$ such that $p = p_x$ and the derivative of Δ_d at p_x in the direction b_1 vanishes, i.e. p_x is a solution of

$$0 = \left(\left(\ln \left(\frac{p_x(s)}{p^*(s)} \right) \right)_{s \in \{0,1,2\}}^t - \frac{(d-1)k}{d} \left(\ln \left(\frac{p_{x,e}(y)}{p_e^*(y)} \right) \right)_{y \in \{0,1\}^2}^t W \right) b_1,$$

where we used the chain rule for multivariate calculus, that W is column stochastic and that $b_1 \in 1_{\{0,1,2\}}^\perp$. Recall from Section 5.3, e.g. from the proof of Lemma 5.1, that $Wb_1 = \frac{d}{k}w$, hence computing the dot product with b_1 gives

$$0 = d \ln(x) - (d-1) \ln \left(\frac{p_{x,e}(1,1)p_{x,e}(0,0)}{p_{x,e}(1,0)p_{x,e}(0,1)} \right).$$

Obviously, equality holds if and only if $x \in \mathcal{X}_{\text{st}}$, hence p is a stationary point of Δ_d iff $p \in \mathcal{P}_{\text{st}}$. \square

Lemma 5.10 does not only allow to translate the stationary points of Δ_d into fixed points of ι_{rr} , it also allows to translate the types as follows.

Lemma 5.11. *Fix $x \in \mathbb{R}_{>0}$. Then we have $\iota_{\text{rr}}(x) < x$ iff $(\Delta_d \circ \iota_{\text{rp}})'(x) > 0$, $\iota_{\text{rr}}(x) > x$ iff $(\Delta_d \circ \iota_{\text{rp}})'(x) < 0$, and $\iota_{\text{rr}}(x) = x$ iff $(\Delta_d \circ \iota_{\text{rp}})'(x) = 0$.*

Proof. Fix $x \in \mathbb{R}_{>0}$. The proof of Lemma 5.10 directly suggests that the first derivative of Δ_d at p_x in the direction b_1 is strictly positive iff

$$0 < \ln(x) - \frac{d-1}{d} \ln(\iota_{\text{rr}}^*(x)),$$

which holds iff $\iota_{\text{rr}}(x) < x$. We're left to establish that the direction of ι_{rp} is consistent with b_1 . Intuitively, using Lemma 5.2 and Lemma 5.9 we can argue that $x \mapsto p_{x,e}(1,1)$ is a bijection and hence either increasing or decreasing. Taking the limits $x \rightarrow 0$ and $x \rightarrow \infty$ suggests that it is increasing, hence with $c \in \mathbb{R}^2$ given by $\iota'_{\text{rp}}(x) = (b_1, b_2)c$, we know that $c_1 \geq 0$.

Formally, we quantify the direction of ι_{rp} . For this purpose we compute the derivative of ι_{rp} at $x \in \mathbb{R}_{>0}$, given by

$$\iota'_{\text{rp}}(x) = \left(\frac{sp^*(s)x^{s-1} \sum_{s' \in \{0,1,2\}} p^*(s')x^{s'} - p^*(s)x^s \sum_{s' \in \{0,1,2\}} s'p^*(s')x^{s'-1}}{\left(\sum_{s' \in \{0,1,2\}} p^*(s')x^{s'} \right)^2} \right)_{s \in \{0,1,2\}}.$$

Notice that $v = \iota'_{\text{rp}}(x) \in 1_{\{0,1,2\}}^\perp$, since $\iota_{\text{rp}}(\mathbb{R}_{>0}) \subseteq \mathcal{P}(\{0,1,2\})$ or by computing $\sum_s v_s = 0$ directly. Now, let $c \in \mathbb{R}^2$ be given by $v = (b_1, b_2)c$. This directly gives $c_2 = v_2$ and hence $c_1 = d^{-1}(v_1 + 2v_2) = d^{-1} \sum_s sv_s$. Now, notice that

$$\begin{aligned} S = dx c_1 &= \sum_{s,s' \in \{0,1,2\}} p_x(s)p_x(s')s(s-s') \\ &= \sum_{s>s'} p_x(s)p_x(s')s(s-s') - \sum_{s>s'} p_x(s)p_x(s')s'(s-s') = \sum_{s>s'} p_x(s)p_x(s')(s-s')^2 > 0, \end{aligned}$$

which directly gives $c_1 = \frac{S}{dx} \in \mathbb{R}_{>0}$. Now, with $\nabla = \left(\frac{\partial \Delta_d}{\partial p(s)}(p_x) \right)_{s \in \{0,1,2\}} \in \mathbb{R}^{\{0,1,2\}}$ denoting the partial derivatives of Δ_d at p_x and using the chain rule we have

$$(\Delta_d \circ \iota_{\text{rp}})'(x) = \nabla^t \iota'_{\text{rp}}(x) = c_1 \nabla^t b_1 + c_2 \nabla^t b_2 = c_1 \nabla^t b_1,$$

since the derivative $\nabla^t b_2$ of Δ_d at p_x in the direction b_2 is zero, hence we have $(\Delta_d \circ \iota_{\text{rp}})'(x) > 0$ iff the derivative $\nabla^t b_1$ of Δ_d at p_x in the direction b_1 is strictly positive, which is the case iff $\iota_{\text{rr}}(x) < x$. \square

Lemma 5.11 with Lemma 5.10 shows that control over ι_{rr} gives complete control over the location and characterization of the stationary points of Δ_d . However, instead of solving the fixed point equation given by ι_{rr} directly, we use a slight modification inspired by the *belief propagation* algorithm applied to the constraint satisfaction discussed in Section 5.1 and initialized with uniform messages.

For this purpose let $N \in \mathbb{Z}_{\geq 0}$, further $N_1, N_2 \in [N]_0$ and the hypergeometric distribution $p_{N,N_1,N_2} \in \mathcal{P}(\mathbb{Z})$ be given by

$$p_{N,N_1,N_2}(s) = \frac{\binom{N_1}{s} \binom{N-N_1}{N_2-s}}{\binom{N}{N_2}} = \frac{\binom{N}{N-N_1-N_2+s, N_1-s, N_2-s, s}}{\binom{N}{N_1} \binom{N}{N_2}} \text{ for } s \in \mathbb{Z}.$$

The latter form directly shows that $p_{N,N_1,N_2} = p_{N,N_2,N_1}$. Now, for $y \in \{0, 1\}^2$ let $p_y^* \in \mathcal{P}(\{0, 1, 2\})$ be given by $p_{(1,1)}^*(s) = p_{k-1,1,1}(s-1)$, $p_{(1,0)}^*(s) = p_{k-1,1,2}(s)$, $p_{(0,1)}^*(s) = p_{k-1,2,1}(s)$, $p_{(0,0)}^*(s) = p_{k-1,2,2}(s)$ for $s \in \{0, 1, 2\}$. On the other hand, for $y \in \{0, 1\}^2$ let $p_y' \in \mathcal{P}(\{0, 1, 2\})$ be given by $p_{(1,1)}'(s) = p_{k-1,1,1}(s)$ for $s \in \{0, 1, 2\}$ and further $p_y' = p_y^*$ for $y \in \{0, 1\}^2 \setminus \{(1, 1)\}$. Finally, for a distribution $p \in \mathcal{P}(\{0, 1, 2\})$ let $f_p: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$, $x \mapsto \sum_{s \in \{0,1,2\}} p(s)x^s$, be its probability generating function, and further $\iota_{BP}: \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$ given by

$$\iota_{BP}(x) = \iota_{BP}^*(x)^{d-1}, \quad \iota_{BP}^*(x) = \frac{f_{p_{(1,1)}'}(x)f_{p_{(0,0)}'}(x)}{f_{p_{(1,0)}'}(x)f_{p_{(0,1)}'}(x)}, \text{ for } x \in \mathbb{R}_{>0}. \quad (9)$$

Lemma 5.12. Fix $x \in \mathbb{R}_{>0}$. Then we have $\iota_{rr}(x) < x$ iff $\iota_{BP}(x) < x$, $\iota_{rr}(x) > x$ iff $\iota_{BP}(x) > x$, and $\iota_{rr}(x) = x$ iff $\iota_{BP}(x) = x$.

Proof. First, notice that the normalization constant of p_x cancels out in ι_{rr}^* , as does the normalization constant $\binom{k}{2}^2$ of p^* . Further, with $v = (k-4+s, 2-s, 2-s, s)^t \in \mathbb{R}^{[0,1]^2}$ we have $W_{y,s}(\binom{k}{v}) = \frac{v_y}{k} \binom{k}{v} = \binom{k-1}{v-(\delta_{y,z})_z}$ for $y \in \{0, 1\}^2$, $s \in \{0, 1, 2\}$, and thereby

$$\iota_{rr}^*(x) = \frac{\left(\sum_s \binom{k-1}{k-4+s, 2-s, 2-s, s-1} x^s\right) \left(\sum_s \binom{k-1}{k-5+s, 2-s, 2-s, s} x^s\right)}{\left(\sum_s \binom{k-1}{k-4+s, 2-s, 1-s, s} x^s\right) \left(\sum_s \binom{k-1}{k-4+s, 1-s, 2-s, s} x^s\right)} \text{ for } x \in \mathbb{R}_{>0}.$$

Now, since the normalization constants cancel out in total, this directly gives

$$\iota_{rr}^*(x) = \frac{f_{p_{(1,1)}^*}(x)f_{p_{(0,0)}^*}(x)}{f_{p_{(1,0)}^*}(x)f_{p_{(0,1)}^*}(x)} = x\iota_{BP}^*(x) \text{ for } x \in \mathbb{R}_{>0},$$

using that $p_{(1,1)}^*(s) = p_{(1,1)}'(s-1)$ for $s \in \{0, 1, 2\}$, hence $f_{p_{(1,1)}^*}(x) = xf_{p_{(1,1)}'}(x)$, and $p_y^*(s) = p_y'(s)$ for $y \neq (1, 1)$. Now, we have $x = \iota_{rr}(x)$ iff $x = x^{\frac{d-1}{d}} \iota_{BP}(x)^{\frac{1}{d}}$, which holds iff $x^{\frac{1}{d}} = \iota_{BP}(x)^{\frac{1}{d}}$, which then again is equivalent to $x = \iota_{BP}(x)$. Equivalence of the inequalities follows analogously. \square

The following part is dedicated to the identification of the fixed points of ι_{BP} . We start with a discussion of ι_{BP}^* . For this purpose let $g_1^*, g_k^*: \mathbb{R}_{\geq 1} \rightarrow \mathbb{R}_{\geq 1}$ be given by

$$g_1^*(x) = \frac{1}{k-1}(x-1) + 1 \text{ and } g_k^*(x) = \frac{13k-12}{27(k-1)(k-2)}(x-k) + \frac{2(7k-12)}{9(k-2)}, x \in \mathbb{R}_{\geq 1}.$$

Lemma 5.13. For any $k \in \mathbb{Z}_{\geq 4}$ we have

$$\lim_{x \rightarrow 0} \iota_{BP}^*(x) = \frac{k-4}{k-3}, \quad \iota_{BP}^*(1) = 1 = g_1^*(1) \quad \text{and} \quad \iota_{BP}^*(k) = g_k^*(k).$$

For the first derivative we have

$$\iota_{BP}'^*(1) = g_1'^*(1), \quad \iota_{BP}'^*(k) = g_k'^*(k) \quad \text{and} \quad \lim_{x \rightarrow \infty} \iota_{BP}'^*(x) = \frac{1}{2(k-2)}.$$

Moreover, for the second derivative we have

$$\iota_{\text{BP}}^{*''}(x) < 0 \text{ for } x \in (0, k), \quad \iota_{\text{BP}}^{*''}(k) = 0, \quad \iota_{\text{BP}}^{*''}(x) > 0 \text{ for } x \in \mathbb{R}_{>k}.$$

For $k = 4$ and $x \in \mathbb{R}_{>0}$ we have $\iota_{\text{BP}}^*(x^{-1}) = (\iota_{\text{BP}}^*(x))^{-1}$.

Proof. Using $f_p(1) = 1$ for the moment generating function of any finitely supported law p , we have $\iota_{\text{BP}}^*(1) = 1$. Further, using that the first moment of a hypergeometric distribution p_{N,N_1,N_2} is $\frac{N_1 N_2}{N}$ and that $f_p'(1)$ is the first moment of p , we have $\iota_{\text{BP}}^{*'}(1) = \frac{1}{k-1} + \frac{4}{k-1} - 2 \cdot \frac{2}{k-1} = \frac{1}{k-1}$. The symmetry of ι_{BP}^* for the special case $k = 4$ can be seen as follows. First, recall that $p_{N,N_1,N_2}(s) = p_{N,N-N_1,N_2}(N_2 - s)$ for any hypergeometric distribution p_{N,N_1,N_2} . For $s \in \{0, 1, 2\}$ this gives

$$\begin{aligned} p'_{(0,0)}(s) &= p_{3,2,2}(s) = p_{3,1,2}(2-s) = p'_{(1,0)}(2-s) \\ &= p'_{(0,1)}(2-s) = p_{3,2,1}(2-s) = p_{3,1,1}(s-1) = p'_{(1,1)}(s-1). \end{aligned}$$

These relations can be directly translated to the moment generating functions, i.e.

$$f_{p'_{(0,0)}}(x) = x^2 f_{p'_{(1,0)}}(x^{-1}) = x^2 f_{p'_{(0,1)}}(x^{-1}) = x f_{p'_{(1,1)}}(x)$$

for $x \in \mathbb{R}_{>0}$. Using these transformations we have

$$\iota_{\text{BP}}^*(x^{-1}) = \frac{f_{p'_{(1,1)}}(x^{-1}) f_{p'_{(0,0)}}(x^{-1})}{f_{p'_{(1,0)}}(x^{-1}) f_{p'_{(0,1)}}(x^{-1})} = \frac{x^{-1} f_{p'_{(1,0)}}(x) x^{-2} f_{p'_{(0,1)}}(x)}{x^{-1} f_{p'_{(1,1)}}(x) x^{-2} f_{p'_{(0,0)}}(x)} = (\iota_{\text{BP}}^*(x))^{-1}.$$

For $k \in \mathbb{Z}_{\geq 4}$ and $x \in \mathbb{R}_{>0}$ direct computation gives

$$\begin{aligned} \iota_{\text{BP}}^*(x) &= \frac{1}{2(k-2)}x + \frac{2k-5}{2(k-2)} + \frac{(k-1)(k-3)(x-1)}{2(k-2)(2x+k-3)^2}, \\ \iota_{\text{BP}}^{*'}(x) &= \frac{1}{2(k-2)} + \frac{(k-1)(k-3)(-2x+k+1)}{2(k-2)(2x+k-3)^3}, \\ \iota_{\text{BP}}^{*''}(x) &= \frac{4(k-1)(k-3)(x-k)}{(k-2)(2x+k-3)^4}. \end{aligned}$$

The remaining assertions follow immediately or with routine computations. \square

Lemma 5.13 has the following immediate consequences.

Corollary 5.14. For any $d \in (0, 2]$ we have $\iota_{\text{BP}}^* \in (x, 1)$ for $x \in (0, 1)$ and $\iota_{\text{BP}}^*(x) \in (1, x)$ for $x \in \mathbb{R}_{>1}$. In particular, p^* is the unique minimizer of Δ_d .

Proof. Using Lemma 5.13 we notice that $\iota_{\text{BP}}^{*'}(x) \in [\iota_{\text{BP}}^{*'}(k), \frac{1}{k-1}] \subset (0, 1)$ for $x \in \mathbb{R}_{\geq 1}$ and $\iota_{\text{BP}}^*(x) = x$ for $x = 1$, hence we have $\iota_{\text{BP}}^*(x) \in (1, x)$ for $x \in \mathbb{R}_{>1}$. For $k = 4$ this gives $\iota_{\text{BP}}^*(x) \in (x, 1)$ using the symmetry result. For $k \in \mathbb{Z}_{>4}$ we have $\lim_{x \rightarrow 0} \iota_{\text{BP}}^*(x) > 0$, which gives $x^* = \inf\{x \in \mathbb{R}_{>0} : \iota_{\text{BP}}^*(x) \leq x\} \in (0, 1]$ using $\iota_{\text{BP}}^*(1) = 1$. Assume that $x^* < 1$, then using the continuity of $x - \iota_{\text{BP}}^*(x)$ we directly get $\iota_{\text{BP}}^*(x^*) = x^*$, and further $\iota_{\text{BP}}^{*'}(x^*) \leq 1$ since $\iota_{\text{BP}}^*(x) > x$ for $x \in (0, x^*)$. But then, since $\iota_{\text{BP}}^{*''}(x) < 0$ for $x \in (0, k)$, this implies that $\iota_{\text{BP}}^{*'}(x) < 1$ for $x \in (x^*, 1]$ which yields that $\iota_{\text{BP}}^*(1) < 1$ and hence a contradiction. This shows that $\iota_{\text{BP}}^*(x) \in (x, 1)$ for $x \in (0, 1)$. Now, for any $d \in (0, 2]$ we have $\iota_{\text{BP}}(x) \geq \iota_{\text{BP}}^*(x) \in (x, 1)$ for $x \in (0, 1)$ and $\iota_{\text{BP}}(x) \leq \iota_{\text{BP}}^*(x) \in (1, x)$ for $x \in \mathbb{R}_{>1}$. Hence, using Lemma 5.11 and Lemma 5.12 we immediately get that $p^* = p_1$ is the unique minimizer of Δ_d . \square

Corollary 5.14 covers the case of simple graphs that was discussed in [37]. On the other hand, Corollary 5.14 suggests that Proposition 2.5 can only hold if $d^* > 2$.

Corollary 5.15. For all $k \in \mathbb{Z}_{\geq 4}$ we have $d^* \in \mathbb{R}_{>2}$.

Proof. For $d \in \mathbb{R}_{>0}$ let $f(d) = (\Delta_d \circ \iota_{\text{rp}})(\infty)$ and notice that $f(d) = \frac{k}{d}\phi_1$. Corollary 5.14 shows that $f(d) > 0$ for all $d \in (0, 2]$. On the other hand, as derived in Section 4 we know that d^* is the unique root of f , which shows that $d^* > 2$. \square

Based on Corollary 5.14 we can restrict to $d \in \mathbb{R}_{>2}$, while Corollary 5.15 motivates the discussion of this interval. Further, the restriction $d \in \mathbb{R}_{>2}$ ensures that x^{d-1} is increasing and convex on $\mathbb{R}_{>0}$. In the following we will consider the intervals $\mathbb{R}_{\geq k}$, $[\bar{x}, k]$, $[1, \bar{x}]$ and $(0, 1]$ independently, where $\bar{x} = \frac{1}{7}(k+6) \in (1, k)$ is the intersection of g_1 and g_k with $g_1(\bar{x}) = g_k(\bar{x}) = \frac{8}{7}$.

Lemma 5.16. For $d \in (2, d_k)$, with $d_k = \frac{\ln(k)}{\ln(\iota_{\text{BP}}^*(k))} + 1$, there exists $x_{\max} \in \mathbb{R}_{>k}$ such that $\iota_{\text{BP}}(x) < x$ for $x \in [k, x_{\max})$, $\iota_{\text{BP}}(x_{\max}) = x_{\max}$ and $\iota_{\text{BP}}(x) > x$ for $x \in \mathbb{R}_{>x_{\max}}$.

Proof. Let $f(d) = (\iota_{\text{BP}}^*(k))^{d-1}$ for $d \in \mathbb{R}_{\geq 2}$, i.e. $f(d) = \iota_{\text{BP}}(k)$ is the value of ι_{BP} at k under a variation of d . We know from Lemma 5.13 that $\iota_{\text{BP}}^*(k) \in \mathbb{R}_{>1}$, hence $f(d)$ is strictly increasing, and further direct computation gives $f(d_k) = k$, so we have $\iota_{\text{BP}}(k) < k$ for any $d \in (2, d_k)$. Now, since ι_{BP}^* is strictly increasing and convex on $\mathbb{R}_{>k}$ by Lemma 5.13 and further the function x^{d-1} is increasing and convex for $d \in (2, d_k)$, we know that ι_{BP} is convex and increasing on $\mathbb{R}_{>k}$, or formally

$$\begin{aligned}\iota'_{\text{BP}}(x) &= (d-1)\iota_{\text{BP}}^*(x)^{d-2}\iota_{\text{BP}}^{*\prime}(x) = (d-1)\iota_{\text{BP}}(x)\frac{\iota_{\text{BP}}^{*\prime}(x)}{\iota_{\text{BP}}^*(x)} > 0, \\ \iota''_{\text{BP}}(x) &= (d-1)\iota_{\text{BP}}(x)\left[(d-2)\left(\frac{\iota_{\text{BP}}^{*\prime}(x)}{\iota_{\text{BP}}^*(x)}\right)^2 + \frac{\iota_{\text{BP}}^{*\prime\prime}(x)}{\iota_{\text{BP}}^*(x)}\right] > 0.\end{aligned}$$

Using Lemma 5.13 and for $x \rightarrow \infty$ we have $\iota_{\text{BP}}^*(x) \rightarrow \infty$ since $\iota_{\text{BP}}^{*\prime}(x) \rightarrow \frac{1}{2(k-2)}$ and hence $\iota'_{\text{BP}}(x) \rightarrow \infty$ by the above, i.e. $\iota_{\text{BP}}(x) - x \rightarrow \infty$, which suggests the existence of $x_{\max} \in \mathbb{R}_{>k}$ with $\iota_{\text{BP}}(x_{\max}) = x_{\max}$ since $\iota_{\text{BP}}(k) < k$. Now, let $x_+ = \inf\{x \in \mathbb{R}_{\geq k} : \iota_{\text{BP}}(x) \geq x\}$, then we have $x_+ \in (k, x_{\max}]$. Since $\iota_{\text{BP}}(x) < x$ for $x \in [k, x_+)$ we need $\iota'_{\text{BP}}(x_+) \geq 1$, which gives $\iota'_{\text{BP}}(x) > 1$ for $x > x_+$ since $\iota'_{\text{BP}}(x) > 0$, hence $\iota_{\text{BP}}(x) > x$, thereby $x_+ = x_{\max}$, and in summary $\iota_{\text{BP}}(x) < x$ for $x \in [k, x_{\max})$, $\iota_{\text{BP}}(x_{\max}) = x_{\max}$ and $\iota_{\text{BP}}(x) > x$ for $x \in \mathbb{R}_{>x_{\max}}$. \square

The proof of Lemma 5.16 serves as a blueprint for the next two cases, where we do not consider ι_{BP} directly since $\iota_{\text{BP}}^{*\prime\prime}(x) < 0$ on $(1, k)$, but work with $g_k(x) = g_k^*(x)^{d-1}$ and $g_1(x) = g_1^*(x)^{d-1}$ instead, which are convex, increasing and upper bounds for ι_{BP} on $[1, k]$ since $\iota_{\text{BP}}^{*\prime\prime}(x) < 0$ on $(1, k)$. In the spirit of Lemma 5.16 we continue to consider the maximal domain for $d \in \mathbb{R}_{>2}$. Let

$$d_{\bar{x}} = \frac{\ln(\bar{x})}{\ln(g_1(\bar{x}))} + 1 \text{ and } d_{\max} = \min(d_{\bar{x}}, d_k).$$

We postpone the proof that $d^* \leq d_{\max}$, instead we focus on the interval $(1, k)$.

Lemma 5.17. For any $d \in (2, d_{\max}) \subseteq (2, k)$ and all $x \in (1, k]$ we have $\iota_{\text{BP}}(x) < x$.

Proof. Fix $d \in (2, d_{\max})$. Since $\iota_{\text{BP}}^{*\prime\prime}(x) < 0$ for $x \in [1, k]$, we know that $\iota_{\text{BP}}^*(x) \leq g_k^*(x)$ for $x \in [\bar{x}, k]$ and $\iota_{\text{BP}}^*(x) \leq g_1^*(x)$ for $x \in [1, \bar{x}]$, so using that x^{d-1} is increasing we have that $\iota_{\text{BP}}(x) \leq g_k(x)$ for $x \in [\bar{x}, k]$ and $\iota_{\text{BP}}(x) \leq g_1(x)$ for $x \in [1, \bar{x}]$. Analogous to Lemma 5.16 we notice that $g_k(k) = \iota_{\text{BP}}(k) < k$ since $d < d_k$, that $g_k(\bar{x}) = g_1(\bar{x}) < \bar{x}$ since $d < d_{\bar{x}}$ and that $g_1(1) = 1$. Further, since g_1^* , g_k^* are increasing and convex, using that x^{d-1} is increasing and convex yields that g_1 , g_k are increasing and convex. In particular, we can upper bound g_k with the line $l_k: [\bar{x}, k] \rightarrow [g_k(\bar{x}), g_k(k)]$ connecting $(\bar{x}, g_k(\bar{x}))$ and $(k, g_k(k))$, which is entirely and strictly under the diagonal. Analogously, we can upper bound g_1 with the line $l_1: [1, \bar{x}] \rightarrow [1, g_1(\bar{x})]$ connecting $(1, 1)$ and $(\bar{x}, g_1(\bar{x}))$, which is also entirely and strictly under the diagonal except for $(1, 1)$ where the two lines intersect. In total,

$\iota_{\text{BP}}(x) \leq \min(g_1(x), g_k(x)) \leq \min(l_1(x), l_k(x)) < x$ for all $x \in (1, k]$. Finally, another implication is that $\frac{d-1}{k-1} = g'_1(1) \leq l'_1(1) < 1$ since l_1 is below the diagonal, which suggests that $d_{\max} \leq k$. \square

Combining Lemma 5.16 and Lemma 5.17 shows for any $d \in (2, d_{\max})$ that $\Delta_d \circ \iota_{\text{rp}}$ has exactly one stationary point x_{\max} on $\mathbb{R}_{>1}$ which is the unique maximizer of $\Delta_d \circ \iota_{\text{rp}}$ on this interval. Further, since $d_{\max} \leq k$ and using $\iota'_{\text{BP}}(1) = \frac{d-1}{k-1}$ we also know that $x = 1$ is an isolated minimizer of $\Delta_d \circ \iota_{\text{rp}}$. Aside, notice that this argumentation can be used to show that H_d as defined in Section 5.5 is positive semidefinite for all $d < k$, hence with the arguments from the proof of Lemma 5.4 we see that H_d is positive definite for all $d < k$ and finally with Lemma 5.8 that $\eta_{\chi^2} = \frac{1}{k-1}$.

For the low overlap region $x \in (0, 1)$ we need a significantly different approach, since ι_{BP}^* is increasing and concave, but $\iota_{\text{BP}}(x) < \iota_{\text{BP}}^*(x)$ and we need to show that $\iota_{\text{BP}}(x) > x$. This means that first order approximations as used for $(1, k)$ are useless since they are upper bounds to ι_{BP}^* and there are no immediate implications for ι''_{BP} as was the case for $\mathbb{R}_{>k}$. However, the symmetric case $k = 4$ can be discussed easily.

Corollary 5.18. *For $k = 4$ and $d \in (2, d_{\max})$ we have $\iota_{\text{BP}}(x) < x$ for $x \in (0, x_{\max}^{-1})$, $\iota_{\text{BP}}(x_{\max}^{-1}) = x_{\max}^{-1}$, and $\iota_{\text{BP}}(x) > x$ for $x \in (x_{\max}^{-1}, 1)$.*

Proof. Combining Lemma 5.17 and Lemma 5.16 we have $\iota_{\text{BP}}(x) < x$ for $x \in (1, x_{\max})$ and $\iota_{\text{BP}}(x) > x$ for $x \in (x_{\max}, \infty)$, hence using the symmetry from Lemma 5.13 directly gives the result. \square

Corollary 5.18 allows to restrict to $k \in \mathbb{Z}_{>4}$ in the remainder. Now, we basically reverse the method used for the interval $(1, k)$, i.e. instead of using tangents g_1^* , g_k^* to ι_{BP}^* and scaling them with $(d-1)$, we scale ι_{BP}^* such that the diagonal is a tangent, meaning we consider $\iota_k = \iota_{\text{BP}}$ for $d = k$ since $\iota'_{\text{BP}}(1) = \frac{d-1}{k-1}$, and show that ι_k is sufficiently convex to ensure $\iota_k(x) > x$ for $x \in (0, 1)$. The next lemma ensures that this approach is applicable for all $k \geq 5$.

Lemma 5.19. *For any $k \in \mathbb{Z}_{\geq 5}$, $d \in (2, k]$ and all $x \in (0, 1)$ we have $\iota_{\text{BP}}(x) > x$.*

Proof. Let $k \in \mathbb{Z}_{\geq 5}$. As derived in the proof of Lemma 5.17 we have $\iota_k(1) = 1$ and $\iota'_k(1) = 1$, i.e. the diagonal is a tangent to ι_k at $x = 1$. Further, as discussed in the proof of Lemma 5.16,

$$\iota''_k(x) = (k-1) \frac{\iota_k(x)}{\iota_{\text{BP}}^*(x)^2} \left[(k-2) \iota_{\text{BP}}^{*'}(x)^2 + \iota_{\text{BP}}^*(x) \iota_{\text{BP}}^{*''}(x) \right] \text{ for } x \in (0, 1).$$

Since the leading factor is clearly strictly positive, we may focus on the term in the square brackets. Further, using that moment generating functions are strictly positive for strictly positive real numbers we can extract the strictly positive denominator of the term and normalize to get

$$f(x) = \frac{\iota_{\text{BP}}^*(x)^2 \left[(k-1) f_{p(1,0)}'(x) \right]^6 (k-2)^2 \iota_k''(x)}{(k-1) \iota_k(x)} = \sum_{i \in [6]_0} a_i x^i,$$

where

$$\begin{aligned} a_6 &= 16(k-2), \\ a_5 &= 48(k-2)(k-3), \\ a_4 &= 4(k-3) \left[9(k-2)(k-3) + 4(k-2)(k-4) + 2(k-1) \right], \\ a_3 &= 8(k-3) \left[3(k-2)^3 + (k-2)(k^2 - 3k + 4) + 2(k-1)(k-4) \right], \\ a_2 &= 4(k-3)b_2, \quad b_2 = (k-2)^2 \left[(k-4)^2 + 3(k^2 - 3k + 4) \right] - (k-1)(k^2 + 11k - 36), \\ a_1 &= 4(k-3)^2 b_1, \quad b_1 = (k-2)(k-4)(k^2 - 3k + 4) - 2(k-1)(2k^2 - 3k - 4), \\ a_0 &= (k-2)(k-3)^2 b_0, \quad b_0 = (k^2 - 3k + 4)^2 - 4k(k-1)(k-4). \end{aligned}$$

We can easily verify that $a_i > 0$ for $3 \leq i \leq 6$ using that $x^2 - 3x + 4 > 0$ for all $x \in \mathbb{R}$. Viewing the b_i , $i \in \{0, 1, 2\}$, as polynomials $b_i(x)$, $x \in \mathbb{R}$, of degree 4 and evaluated at $x = k$, we have

$$b_2''(x) = 48x^2 - 228x + 254, b_2(5) = 82, b_2'(5) = 225,$$

hence $b_2''(x) > 0$ for $x > x_2$ with $x_2 = \frac{19}{8} + \frac{\sqrt{201}}{24} < 3$ and by that $b_2(x) > 0$ for all $x \in \mathbb{R}_{\geq 5}$, so in particular $b_2 = b_2(k) > 0$ since $k \in \mathbb{Z}_{\geq 5}$ and thereby $a_2 > 0$. Using the same technique for the degree four polynomials $b_1(x)$ we obtain that $b_1''(x) > 0$ for $x > x_1$ with $x_1 = \frac{13}{4} + \frac{\sqrt{561}}{12} < 6$, $b_1(10) = 564$, $b_1'(10) = 854$, and hence that $b_1 > 0$ if $k \geq 10$. For $b_0(x)$ we have $b_0''(x) > 0$ for $x > x_0$ with $x_0 = \frac{5}{2} + \frac{\sqrt{3}}{6} < 3$, $b_0'(3) = 20$, $b_0(3) = 40$, so $b_0 > 0$ for all $k \in \mathbb{Z}_{\geq 5}$.

Hence, for $k \in \mathbb{Z}_{\geq 10}$ we know that $a_i > 0$ for all $i \in [6]_0$, which directly implies that $f(x) > 0$ for all $x \in \mathbb{R}_{>0}$, so $\iota_k''(x) > 0$, further $\iota_k'(x) < 1$ for $x \in (0, 1)$, $\iota_k'(x) > 1$ for $x \in \mathbb{R}_{>1}$ and thereby $\iota_k(x) > x$ for $x \in \mathbb{R}_{>0} \setminus \{1\}$.

For $5 \leq k \leq 9$ we still have $a_i > 0$ for $i \in [6]_0 \setminus \{1\}$. For $6 \leq k \leq 9$ we consider the quadratic function $g_k(x) = \sum_{i \in \{0,1,2\}} a_i x^i$ explicitly, given by

$$\begin{aligned} g_6(x) &= 6120x^2 - 11664x + 8784, g_7(x) = 24960x^2 - 25344x + 41600, \\ g_8(x) &= 72560x^2 - 34400x + 156000, g_9(x) = 172944x^2 - 9504x + 484848. \end{aligned}$$

It turns out that $g_k(x) > 0$ for all $x \in \mathbb{R}$ and $6 \leq k \leq 9$, which in particular yields $f(x) > 0$ for all $x \in \mathbb{R}_{>0}$. Using the same argumentation as for $k \in \mathbb{Z}_{\geq 10}$ shows that $\iota_k(x) > x$ for all $x \in \mathbb{R}_{>0} \setminus \{1\}$ and $k \in \mathbb{Z}_{\geq 6}$.

As opposed to the previous cases the function ι_k is not convex for $k = 5$, and while this slightly complicates the computation, we will show that this does not affect the overall picture. Now, we consider the complete sixth order polynomial

$$f(x) = 48x^6 + 288x^5 + 592x^4 + 1088x^3 + 656x^2 - 3296x + 1392.$$

We notice that $f''(x) > 0$ for all $x \in \mathbb{R}_{\geq 0}$ since $a_i > 0$ for $i \in [6] \setminus \{1\}$, hence $f'(x)$ is strictly increasing for $x \in \mathbb{R}_{\geq 0}$, which shows the existence of a unique root $x_{\min} \in (0, 1)$, using $f'(0) < 0$ and $f'(1) > 0$, i.e. x_{\min} is the unique minimizer of f on $\mathbb{R}_{\geq 0}$. Computing $f(x_{1-}) > 0$, $f(x_{1+}) < 0$ and $f(1) > 0$ with $x_{1-} = 0.581$ and $x_{1+} = 0.582$ ensures the existence of exactly two roots $x_1 \in (x_{1-}, x_{1+})$ and $x_2 \in (x_{1+}, 1)$ of f , hence ι_k is convex on $(0, x_1)$, concave on (x_1, x_2) and convex on $\mathbb{R}_{>x_2}$. Let $g: \mathbb{R} \rightarrow \mathbb{R}$, $x \mapsto \iota_k'(x_{1-})(x - x_{1-}) + \iota_k(x_{1-})$ denote the tangent of ι_k at x_{1-} . Since we can write both ι_k and ι_k' as the ratio of polynomials with integer coefficients, we can compute $\iota_k(x_{1-}) \in (0.584, 0.585)$, $\iota_k'(x_{1-}) \in (0.99, 0.991)$ and $g(x_{1+}) \in (0.585, 0.586)$ exactly. Using the convexity of ι_k on $(0, x_1)$, $g' = \iota_k'(x_{1-}) < 1$ and $g(x_{1+}) > x_{1+}$ immediately gives that $\iota_k(x) \geq g(x) > x$ for $x \in (0, x_1]$. The fact that $\iota_k(x) > x$ for $x \in (x_2, 1)$ immediately follows from the convexity of ι_k on $(x_2, 1)$ and that the diagonal is a tangent to ι_k at $x = 1$. But now, since $(x_1, \iota_k(x_1))$ and $(x_2, \iota_k(x_2))$ are above the diagonal, so is the line connecting the two, which is a lower bound to ι_k on (x_1, x_2) since ι_k is concave on this interval. By that we have finally showed that the overall picture is also the same for $k = 5$, i.e. $\iota_k(x) > x$ for $x \in \mathbb{R}_{>0} \setminus \{1\}$.

The fact that $\iota_k(x) > x$, i.e. $\iota_{\text{BP}}(x) > x$ with $d = k$, for $x \in \mathbb{R}_{>0} \setminus \{1\}$ shows that the only stationary point of Δ_k is a saddle at p^* (respectively a maximum of $\Delta_k \circ \iota_{\text{rp}}$ at $x = 1$). But more importantly, since x^{d-1} is decreasing in d for $x \in (0, 1)$ and $\iota_{\text{BP}}^*(x) \in (0, 1)$, we have $\iota_{\text{BP}}(x) \geq \iota_k(x) > x$ for all $d \in (2, k]$. \square

The combination of Lemma 5.16, Lemma 5.17 and Lemma 5.18 shows that for $k = 4$ and all $d \in (2, d_{\max})$ there exist exactly three fixed points $x_- < x_0 < x_+$ of ι_{BP} , with $x_0 = 1$, $x_+ \in (k, \infty)$ and $x_- = x_+^{-1}$, hence Lemma 5.12 and Lemma 5.11 suggest that x_0 is a minimizer of $\Delta_d \circ \iota_{\text{rp}}$ while x_- and x_+ are maximizers. In particular, we have the three minimizers $\{0, 1, \infty\}$ of $\Delta_d \circ \iota_{\text{rp}}$ in total.

The combination of Lemma 5.16, Lemma 5.17 and Lemma 5.19 shows that for $k \in \mathbb{Z}_{\geq 5}$ and all $d \in (2, d_{\max}) \subseteq (0, k)$ there exist exactly two fixed points $x_0 < x_+$ of ι_{BP} , with $x_0 = 1$ and $x_+ \in (k, \infty)$, hence Lemma 5.12 and Lemma 5.11 suggest that x_0 is a minimizer of $\Delta_d \circ \iota_{rp}$ while x_+ is a maximizer. In particular, we have the two minimizers $\{1, \infty\}$ of $\Delta_d \circ \iota_{rp}$ in total, while $x = 0$ is a maximizer in these cases.

The last step is to show that we have $d^* \in (2, d_{\max})$, which then directly establishes that the unique minimizers of $\Delta_{d^*} \circ \iota_{rp}$ are given by $\{0, 1, \infty\}$ for $k = 4$ and $\{1, \infty\}$ for $k \in \mathbb{Z}_{\geq 5}$ as required. On the other hand, direct computation as in the proof of Corollary 5.15 shows that all minimizers are roots of $\Delta_{d^*} \circ \iota_{rp}$, i.e. all minimizers are global minimizers and the global minimum of $\Delta_{d^*} \circ \iota_{rp}$ is 0. Lemma 5.9 directly suggests that the global minimizers of $\Delta_{d^*} \circ \iota_{rp}$ are in one to one correspondence with the global minimizers of Δ_{d^*} via $x \mapsto p_x$, which then completes the proof of Proposition 2.5.

Lemma 5.20. *For all $k \in \mathbb{Z}_{\geq 4}$ we have $2 < d^* < d_k < d_{\bar{x}}$.*

Proof. Recall from Corollary 5.15 that $d^* > 2$. For convenience, we consider the extensions of $d^* - 1$, $d_{\bar{x}} - 1$ and $d_k - 1$ to the real line, i.e. for $x \in \mathbb{R}_{\geq 3}$ let

$$f_0(x) = \frac{\ln(x(x-1)/2)}{xH(2/x) - \ln(x(x-1)/2)}, f_1(x) = \frac{\ln((x+6)/7)}{\ln(8/7)}, f_2(x) = \frac{\ln(x)}{\ln\left(\frac{2(7x-12)}{9(x-2)}\right)},$$

i.e. $d^* = f_0(k) + 1$, $d_{\bar{x}} = f_1(k) + 1$ and $d_k = f_2(k) + 1$ for all $k \in \mathbb{Z}_{\geq 4}$. We start with the asymptotic comparison of f_1 and f_2 . The corresponding rearrangement gives

$$f_1(x) = m_1 \ln(x) + t_1(x), m_1 = \frac{1}{\ln(8/7)}, t_1(x) = -\frac{\ln(7)}{\ln(8/7)} + \frac{\ln\left(1 + \frac{6}{x}\right)}{\ln(8/7)},$$

$$f_2(x) = m_2(x) \ln(x), m_2(x) = \frac{1}{\ln(14/9) + \ln\left(1 + \frac{2}{7(x-2)}\right)}.$$

Notice that $\ln(x) > 0$ since $x \geq 3$, further $t_1(x)$ is decreasing while $m_2(x)$ is increasing, and thereby we have $f_1(x) \geq f_{1\infty}(x)$ and $f_2(x) \leq f_{2\infty}(x)$ with

$$f_{1\infty}(x) = m_{1\infty} \ln(x) + t_{1\infty}, m_{1\infty} = m_1, t_{1\infty} = -\frac{\ln(7)}{\ln(8/7)},$$

$$f_{2\infty}(x) = m_{2\infty} \ln(x), m_{2\infty} = \frac{1}{\ln(14/9)}.$$

Notice that $m_{1\infty} > m_{2\infty}$, $t_{1\infty} < 0$ and further $f_{1\infty}(x) > f_{2\infty}(x)$ iff $x > x_{12}$ with $x_{12} = \exp\left(\frac{-t_{1\infty}}{m_{1\infty} - m_{2\infty}}\right) \in (16, 17)$, so $f_1(x) > f_2(x)$ for all $x \in \mathbb{R}_{\geq 17}$ and hence $d_k < d_{\bar{x}}$ for $k \in \mathbb{Z}_{\geq 17}$. We check by hand that $d_k < d_{\bar{x}}$ also holds for $4 \leq k \leq 16$.

We are left to show that $1 < f_0(x) < f_2(x)$ for $x \in \mathbb{Z}_{\geq 4}$. Again, we start with the asymptotic comparison, where the corresponding rearrangement $f_0(x) = m_0(x) \ln(x) + t_0(x)$ is given by

$$m_0(x) = \frac{1 + \ln\left(1 - \frac{1}{x}\right)}{n_0(x)}, t_0(x) = \frac{-\ln(2)}{n_0(x)}, n_0(x) = -(x-2) \ln\left(1 - \frac{2}{x}\right) - \ln\left(1 - \frac{1}{x}\right) - \ln(2).$$

Recall that for given $c \in \mathbb{R}$ we have $\ln\left(1 + \frac{c}{x}\right) \sim \frac{c}{x}$ and $\ln\left(1 + \frac{c}{x}\right) \leq \frac{c}{x}$ for all $x \in \mathbb{R}_{>|c|}$, since $\ln(x)$ is concave and the tangent at 1 is $x - 1$. Hence, for all $x \in \mathbb{R}_{\geq 3}$ we have $n_0(x) \geq n_+(x) > 0$ since $x > 2$ and $x > x_1$, where

$$n_+(x) = \ln(e^2/2) - \frac{3}{x} \text{ and } x_1 = \frac{3}{\ln(e^2/2)} \in (2, 3).$$

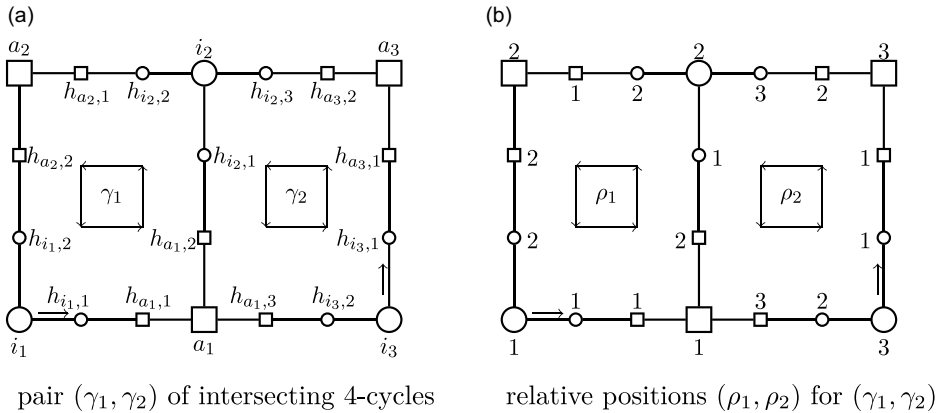


Figure 3. The left figure shows a sequence $\gamma = (\gamma_1, \gamma_2)$ of two directed (intersecting) four-cycles with base variables i_1 and i_3 and directions indicated by the arrows respectively. Analogously to Figure 2a we only denoted the i -edges and a -edges instead of the v -edges and f -edges. The relative positions $\rho = (\rho_1, \rho_2)$ corresponding to γ are depicted in the right figure. Here, the variables, constraints, i -edges and a -edges are labelled according to the order of first traversal (where γ_1 is traversed before γ_2). The numbers $n(\rho) = 3$, $m(\rho) = 3$, $e(\rho) = 7$ of variables, constraints and edges in ρ are equal to the corresponding numbers in γ , further the degree $d_j(\rho)$ of the variable $j \in [3]$ equals the degree of j in γ , and analogously for the degrees $k_b(\rho)$ of the constraints $b \in [3]$ in ρ . The absolute values $\alpha = (\alpha_v, \alpha_f, (\alpha_{v,j})_{j \in [3]}, (\alpha_{f,b})_{b \in [3]})$ are given by $\alpha_v = (i_j)_{j \in [3]}$, $\alpha_f = (a_b)_{b \in [3]}$, $\alpha_{v,j} = (h_{i_j,e})_{e \in [d_j(\rho)]}$, $j \in [3]$, and $\alpha_{f,b} = (h_{a_b,e})_{e \in [k_b(\rho)]}$, $b \in [3]$, i.e. they store the (initial) labels of γ corresponding to the labels of ρ .

Since $t_0(x) < 0$ and $m_0(x) \leq m_+(x)$ we have $f_0(x) \leq f_+(x) = m_+(x) \ln(x)$ with

$$m_+(x) = \left(\ln(e^2/2) - \frac{3}{x} \right)^{-1}.$$

Now, since $m_+(x)$ is decreasing in x and $m_2(x)$ is increasing in x , we numerically determine $x^* \in \mathbb{R}_{>0}$ such that $m_+(x^*) = m_2(x^*)$ and find that $x^* \in (4, 5)$. In particular, we have $f_0(x) \leq f_+(x) < f_2(x)$ for $x \in \mathbb{R}_{\geq 5}$, and check that $d^* < d_k$ for $k = 4$ by hand. \square

Lemma 5.20 concludes the proof of Proposition 2.5 as discussed before.

6. Small subgraph conditioning

In this section we prove the remaining parts of Theorem 2.7, thereby establishing Theorem 2.1. The first part of the proof heavily relies on Section A and illustrates the correspondences. We start with the derivation of δ_ℓ by computing $\mathbb{E}[ZX_\ell]$. For this purpose we fix $\ell \in \mathbb{Z}_{>0}$, $n \in \mathcal{N}$ sufficiently large, and let \bar{c}_ℓ denote the canonical 2ℓ -cycle, i.e. the cycle with variables i , constraints a in $[\ell]$ and i -edges, a -edges in $\{1, 2\}$ with labels ordered by first traversal, see e.g. the left cycle in Figure 3b. Analogous to the previous sections we rewrite the expectation and count the number $|\mathcal{E}|$ of triplets $(g, c, x) \in \mathcal{E}$ such that c is a 2ℓ -cycle and x a solution in g , i.e.

$$\mathbb{E}[ZX_\ell] = \frac{|\mathcal{E}|}{|\mathcal{G}|} = \sum_{y \in \{0,1\}^\ell} \frac{e_1 e_2 e_3}{2^\ell (dn)!},$$

where

$$\begin{aligned} e_1 &= e_1(y) = \binom{n}{n_1} (n_1)_{r_1} (n - n_1)^{\ell - r_1} (d(d-1))^\ell, \\ e_2 &= e_2(y) = \binom{k}{2}^m m^\ell 2^{r_2} (2(k-2))^{2(r_1-r_2)} ((k-2)(k-3))^{\ell-2r_1+r_2}, \\ e_3 &= e_3(y) = (dn_1 - 2r_1)! (d(n - n_1) - 2(\ell - r_1))!, \end{aligned}$$

and $r = r(y) = (r_1, r_2)$ is defined as follows. For $y \in \{0, 1\}^\ell$ we let $r_1 = r_1(y)$ denote the number of ones in y . Further, let $r_2 = r_2(y)$ denote the number of constraints $b \in [\ell]$ in \bar{c}_ℓ such that both b -edges take the value one under the assignment y of the variables $j \in [\ell]$ in \bar{c}_ℓ . With y fixed we can compute the number of suitable triplets (g, c, x) as follows. The denominator in the first line reflects $|\mathcal{G}|^{-1}$ and the compensation 2ℓ as we will count directed cycles γ in g . The sum over $y \in \{0, 1\}^\ell$ implements the choice of the assignment of the variables visited by γ such that the variables i_1, \dots, i_ℓ traversed by γ correspond to the variables $1, \dots, \ell$ in \bar{c}_ℓ in this order, i.e. $x_{i_1} = y_1, \dots, x_{i_\ell} = y_\ell$. The first term in e_1 chooses the variables that take the value one under the solution x . Then we choose the r_1 variables out of the n_1 -variables that participate in the directed cycle γ and take the value one consistent with y (hence an ordered choice). Analogously, we then choose the variables in γ taking zero under x . Finally, we choose the two i -edges traversed by γ for each of the ℓ variables i in the cycle.

The first term in e_2 is the usual choice of the two a -edges taking one under x for each $a \in [m]$. Then we choose the constraints visited by γ . The remaining terms account for the ordered choice of the two a -edges that are traversed by γ and that is consistent with the assignments y and x in the following sense. The (already chosen) variables i_1, \dots, i_ℓ and constraints a_1, \dots, a_ℓ traversed by γ correspond to the variables $1, \dots, \ell$ and constraints $1, \dots, \ell$ in \bar{c}_ℓ in this order respectively. Further, the assignment of these variables is already fixed by y and the a -edges taking the value one for each of these constraints are also fixed by our previous choice. Hence, if $y_1 = y_2 = 1$, then we have only two choices for the a_1 -edge connecting to i_1 , while the a_1 -edge connecting to i_2 is fixed afterwards. For $y_1 = 1$ and $y_2 = 0$ we have two choices for the a_1 -edge connecting to i_1 and $(k-2)$ choices for the a_1 -edge connecting to i_2 . The case $y_1 = 0, y_2 = 1$ is symmetric and we see that we have $(k-2)$ and $(k-3)$ choices for the remaining case $y_1 = y_2 = 0$ analogously. To derive the number of constraints for each of the cases above we recall that we have $r_1(y)$ ones in total and $r_2(y)$ ones whose successor is one (i.e. the constraint a between the two ones takes the value one on both a -edges, and where the successor of y_ℓ is y_1). But then $(r_1 - r_2)$ ones in y do not have the successor one, i.e. they have the successor zero. Complementarily we see that since r_2 ones are succeeded by a one there are r_2 ones that are preceded by a one, hence there are $(r_1 - r_2)$ ones that are preceded by zero. Then again, this means that there are $(r_1 - r_2)$ zeros that are succeeded by a one, hence the remaining $(\ell - 2r_1 + r_2)$ zeros out of the $(\ell - r_1)$ zeros are succeeded by a zero. This fixes γ , so in particular $2r_1$ v -edges that take the value one and $2(\ell - r_1)$ v -edges that take the value zero. The two terms in e_3 then wire the remaining edges.

We divide by $\mathbb{E}[Z]$ to match the left hand side of Theorem 2.7b), i.e.

$$\frac{\mathbb{E}[ZX_\ell]}{\mathbb{E}[Z]} = \sum_{y \in \{0,1\}^\ell} \frac{e_1 e_2 e_3}{2\ell(2m)!(dn - 2m)!}, \text{ where}$$

$$e_1 = e_1(y) = n_1^{r_1} (n - n_1)^{\ell - r_1} (d(d-1))^\ell,$$

$$e_2 = e_2(y) = m^\ell 2^{r_2} (2(k-2))^{2(r_1 - r_2)} ((k-2)(k-3))^{\ell - 2r_1 + r_2},$$

$$e_3 = e_3(y) = (dn_1 - 2r_1)!(d(n - n_1) - 2(\ell - r_1))!,$$

and using Stirling's formula we easily derive that

$$\frac{\mathbb{E}[ZX_\ell]}{\mathbb{E}[Z]} \sim \lambda_\ell \sum_{y \in \{0,1\}^\ell} M_{11}^{r_2} M_{01}^{r_1 - r_2} M_{10}^{r_1 - r_2} M_{00}^{\ell - 2r_1 + r_2} = \lambda_\ell (1 + \delta_\ell), \quad M = \begin{pmatrix} 1 - \frac{2}{k-1} & 1 - \frac{1}{k-1} \\ \frac{2}{k-1} & \frac{1}{k-1} \end{pmatrix}.$$

The matrix M has a nice interpretation as a (column stochastic) transition probability matrix in a two state Markov process, with

$$1 + \delta_\ell = \sum_{\substack{y \in \{0,1\}^\ell \\ y_1=0}} M_{11}^{r_2} M_{01}^{r_1-r_2} M_{10}^{r_1-r_2} M_{00}^{\ell-2r_1+r_2} + \sum_{\substack{y \in \{0,1\}^\ell \\ y_1=1}} M_{11}^{r_2} M_{01}^{r_1-r_2} M_{10}^{r_1-r_2} M_{00}^{\ell-2r_1+r_2}$$

reflecting the probabilities that we return to the starting point given that the starting point is zero and one respectively. Let us consider the first partial sum restricted to sequences y (of Markov states) such that $y_1 = 0$, i.e. we start in the state zero. Then M_{0y_2} reflects the probability that we move from the initial state zero to the state y_2 given that we are in state zero (which is the case because we know that $y_1 = 0$). As discussed above we will move from a one to a one in y exactly r_2 times, from a one to a zero ($r_1 - r_2$) times, from a zero to a one ($r_1 - r_2$) times and from a zero to a zero ($\ell - 2r_1 + r_2$) times. Hence the contribution to the first partial sum for given y exactly reflects the probability that we start in the state zero and (with this given) return to the state zero after ℓ steps (since the successor of y_ℓ is $y_1 = 0$). Since we sum over all such sequences y the first sum reflects the probability that we reach state zero after ℓ steps given that we start in the state zero. The discussion of the second sum is completely analogous. This directly yields

$$1 + \delta_\ell = (M^\ell)_{00} + (M^\ell)_{11} = \text{Tr}(M^\ell) = \lambda'_1 + \lambda'_2 = \lambda_1^\ell + \lambda_2^\ell, \quad \lambda_1 = 1, \lambda_2 = -\frac{1}{k-1},$$

where we used the Kolmogorov-Chapman equalities in the first step, i.e. that the ℓ -step transition probability matrix is the ℓ -th power of the one step transition probability matrix, which allow to translate the first sum into the transition probability $(M^\ell)_{00}$ that we reach the state zero after ℓ steps given that we start in the state zero and analogously for the second sum. In the second step we use the definition of the trace, while in the third step we use that the trace is the sum of the eigenvalues λ'_1, λ'_2 of M^ℓ . In the next step we use that the eigenvalues λ'_1, λ'_2 of the ℓ -th power M^ℓ of the matrix M are the ℓ -th powers of the eigenvalues λ_1, λ_2 of M . In particular this also yields that $\delta_\ell > -1$ for all $k > 3$ and establishes $\delta_\ell = (1 - k)^{-\ell}$.

Following the strategy of Section A we turn to the case of disjoint cycles. Similarly, the present case is a canonical extension of the single cycle case discussed above. We fix $L \in \mathbb{Z}_{>0}$, $r \in \mathbb{Z}_{\geq 0}^L$ and $n \in \mathcal{N}$ sufficiently large. Further, as in the previous sections we rewrite the expectation and count the number $|\mathcal{E}|$ of triplets $(g, c, x) \in \mathcal{E}$ such that $c = (c_s)_{s \in [\bar{r}]}$ is a sequence of $\bar{r} = \sum_{\ell \in [L]} r_\ell$ distinct $2\ell_s$ -cycles c_s in the configuration g sorted by their length ℓ_s in ascending order (as described in Section A) and x is a solution of g . This yields

$$\mathbb{E} \left[Z \prod_{\ell \in [L]} (X_\ell)_{r_\ell} \right] = \frac{|\mathcal{E}|}{|\mathcal{G}|} = \frac{|\mathcal{E}_0|}{|\mathcal{G}|} + \frac{|\mathcal{E}_1|}{|\mathcal{G}|},$$

where $\mathcal{E}_0 \subseteq \mathcal{E}$ is the set over all triplets $(g, c, x) \in \mathcal{E}$ involving sequences c of disjoint cycles and $\mathcal{E}_1 = \mathcal{E} \setminus \mathcal{E}_0$. We begin with the first contribution, which can be regarded as a combination of the discussion of disjoint cycles in Section A and the single cycle case above, i.e.

$$\begin{aligned} \frac{|\mathcal{E}_0|}{|\mathcal{G}|} &= \sum_{y \in \{0,1\}^l} \frac{e_1 e_2 e_3}{(dn)! \prod_{s \in [\bar{r}]} (2\ell_s)^{l_s}}, \\ e_1 &= e_1(y) = \binom{n}{n_1} n_1^{r_1} (n - n_1)^{l - r_1} (d(d-1))^{l_1}, \\ e_2 &= e_2(y) = \binom{k}{2}^m m^{l_2} 2^{r_2} (2(k-2))^{2(r_1-r_2)} ((k-2)(k-3))^{l - 2r_1 + r_2}, \\ e_3 &= e_3(y) = (dn_1 - 2r_1)! (d(n - n_1) - 2(l - r_1))!, \\ l &= \sum_{s \in [\bar{r}]} \ell_s, r_i = \sum_{s \in [\bar{r}]} r_i(y_s), i \in [2], \end{aligned}$$

where $y = (y_s)_{s \in [\bar{r}]}$ is the subdivision of y corresponding to the definition of c , and r_1, r_2 are the notions defined above. The combinatorial arguments are now fairly self-explanatory, e.g. we make an ordered choice of the $r_1(y_1)$ variables taking one for γ_1 , then an ordered choice of $r_1(y_2)$ variables taking one for γ_2 out of the remaining $n_1 - r_1(y_1)$ variables taking one and so on.

The asymptotics are also completely analogous to the single cycle case and Section A. First, we notice that the sum is still bounded, i.e. we can also use the asymptotic equivalences for the corresponding ratio here. Then, the sum can be decomposed into the product of the \bar{r} factors that correspond to the single cycle case above, analogously to Section A, which yields

$$\frac{|\mathcal{E}_0|}{|\mathcal{G}|\mathbb{E}[Z]} \sim \prod_{\ell \in [L]} \lambda_\ell^{r_\ell} (1 + \delta_\ell)^{r_\ell}.$$

Now we turn to the proof that the second contribution involving \mathcal{E}_1 is negligible, which is a combination of the above and the discussion of intersecting cycles in Section A. We let

$$\begin{aligned} \mathcal{E}_2 &= \{(g, \gamma, x) : (g, c(\gamma), x) \in \mathcal{E}_1\}, \mathcal{R} = \{\rho(\gamma) : (g, \gamma, x) \in \mathcal{E}_2\} \text{ and} \\ \mathcal{E}_\rho &= \{(g, \gamma, x) \in \mathcal{E}_2 : \rho(\gamma) = \rho\} \text{ for } \rho \in \mathcal{R} \end{aligned}$$

denote the sets that match the corresponding sets in Section A. For relative positions $\rho \in \mathcal{R}$ we consider an assignment $y \in \{0, 1\}^{n(\rho)}$ of the variables $V = [n(\rho)]$ in the corresponding union of cycles $c = c(\rho)$ and let

$$\begin{aligned} \tau_1 &= \tau_1(\rho, y) = |\{j \in V : y_j = 1\}|, \\ \mathfrak{o}(b) &= \mathfrak{o}_{\rho, y}(b) = |\{h \in [k_b(\rho)] : y_{i_c(b, h)} = 1\}| \text{ for } b \in [m(\rho)] \text{ and} \\ \mathfrak{o} &= \mathfrak{o}(\rho, y) = \sum_{b \in [m(\rho)]} \mathfrak{o}(b) \end{aligned}$$

denote the number of variables $j \in V$ in c that take the value one under y , the number of b -edges for a constraint $b \in [m(\rho)]$ in c that take the value one under y and the number of f -edges in c that take the value one under y respectively. Since c is a configuration the number of v -edges in c that take the value one under y is also \mathfrak{o} . We are particularly interested in the assignments

$$y \in \mathcal{Y} = \mathcal{Y}(\rho) = \{z \in \{0, 1\}^{n(\rho)} : \forall b \in [m(\rho)] \mathfrak{o}(b) \in [2 + k_b - k, 2]\}$$

that do not directly violate a constraint $b \in [m(\rho)]$ in $c(\rho)$ in the sense that $\mathfrak{o}(b) \leq 2$ and also do not indirectly violate b in that $2 - \mathfrak{o}(b) \leq k - k_b$, i.e. there are sufficiently many b -edges left to take the remaining $(2 - \mathfrak{o}(b))$ ones. With this slight extension of our machinery we can derive

$$\begin{aligned} \frac{|\mathcal{E}_1|}{|\mathcal{G}|} &= \sum_{\rho \in \mathcal{R}} \frac{|\mathcal{E}_\rho|}{(dn)! \prod_{s \in [\tau]} (2\ell_s)}, |\mathcal{E}_\rho| = \sum_{y \in \mathcal{Y}} e_1 e_2 e_3, \\ e_1 &= e_1(\rho, y) = \binom{n}{n_1} n_1^{\tau_1} (n - n_1)^{n(\rho) - \tau_1} \prod_{j \in [n(\rho)]} d_j^{d_j(\rho)}, \\ e_2 &= e_2(\rho, y) = \binom{k}{2}^m m^{m(\rho)} \prod_{b \in [m(\rho)]} (2^{\mathfrak{o}(b)} (k - 2)^{k_b(\rho) - \mathfrak{o}(b)}), \\ e_3 &= e_3(\rho, y) = (dn_1 - \mathfrak{o})! (dn - n_1 - (e(\rho) - \mathfrak{o}))!, \end{aligned}$$

for the following reasons. With $\rho \in \mathcal{R}$ and $y \in \mathcal{Y}(\rho)$ fixed we choose the n_1 variables out of the n variables in the configuration g that should take the value one under x . Out of these n_1 variables we choose the τ_1 variables (ordered by first traversal) that take the value one in the directed cycles γ under x , corresponding to the τ_1 variables in ρ that take one under y (more precisely we choose the values $i \in [n]$ of the absolute values α_v for the τ_1 variables $j \in [n(\rho)]$ in ρ that take the value

one under y) and analogously for the variables that take zero. Then, for each variable $j \in [n(\rho)]$ in ρ and corresponding variable $i = \alpha_v(j)$ in γ we choose the i -edges that participate in γ (meaning that we choose $\alpha_{v,j}$). On the constraint side we first choose the two a -edges that take the value one under x in g for each $a \in [m]$. Then we select the $m(\rho)$ constraints that participate in γ (i.e. we fix α_f). Further, for each constraint $b \in [m(\rho)]$ in ρ and its corresponding constraint $a = \alpha_f(b)$ in γ we choose the $\alpha(b)$ a -edges that take the value one in γ under x consistent with ρ and y out of the two a -edges that take the value one in g under x and analogously for the a -edges that take the value zero (which means that we fix $\alpha_{f,b}$ for $b \in [m(\rho)]$ consistent with the choice of y and the choice of the two a -edges that take the value one for each $a \in [m]$). This fixes the sequence of the directed cycles (i.e. the isomorphism α and further γ). The remaining terms wire the $(dn_1 - \alpha)$ remaining v -edges that take the value one and the v -edges taking zero respectively.

As opposed to the rather demanding combinatorial part the asymptotics are still easy to derive since both sums are bounded, so the procedure analogous to Section A yields

$$\frac{|\mathcal{E}_1|}{|\mathcal{G}|\mathbb{E}[Z]} \sim \sum_{\rho \in \mathcal{R}} \sum_{y \in \mathcal{Y}} c_1(\rho, y) n^{n(\rho) + m(\rho) - e(\rho)},$$

where $c_1(\rho, y)$ is a constant compensating the bounded terms. The right hand side tends to zero by the argumentation in Section A, so this contribution is indeed negligible. This shows that $\frac{|\mathcal{E}|}{|\mathcal{G}|} \sim \frac{|\mathcal{E}_0|}{|\mathcal{G}|}$ and thereby establishes Theorem 2.7 (2).

With $d \in [1, d^*) \subseteq [1, k]$ as discussed in Lemma 5.17 and Lemma 5.20, λ_ℓ as derived in Lemma 2.8, $\delta_\ell = (1 - k)^{-\ell}$, the asymptotics of the second moment discussed in Lemma 2.6 and the Taylor series $\ln(1 - x) = -\sum_{\ell \geq 1} x^\ell / \ell$, $x \in (0, 1)$, we establish Theorem 2.7 (3) by applying our results to the sum

$$\sum_{\ell \geq 1} \lambda_\ell \delta_\ell^2 = \sum_{\ell \geq 1} \frac{1}{2\ell} \left(\frac{d-1}{k-1} \right)^\ell = -\frac{1}{2} \ln \left(1 - \frac{d-1}{k-1} \right) = \ln \left(\sqrt{\frac{k-1}{k-d}} \right).$$

This concludes the proof of Theorem 2.7 and further the proof of Theorem 1.1.

Acknowledgements

An extended abstract of a preliminary version of this paper has appeared at the proceedings of ICALP '19 [36]. The results here apply to all 2-in- k occupation problems, as opposed to only $k = 4$ in [36].

Funding statement

The research leading to these results has received funding from the European Research Council, ERC Grant Agreement 772606–PTRCSP.

References

- [1] Abbe, E. (2018) Community detection and stochastic block models: recent developments. *J Mach Learn Res* **18** 1–86.
- [2] Bahmanian, M. A. and Šajna, M. (2017) Quasi-Eulerian hypergraphs. *Electron. J. Combin* **24** 12.
- [3] Bapst, V., Coja-Oghlan, A. and Efthymiou, C. (2017) Planting colourings silently. *Combin. Probab. Comput* **26** 338–366.
- [4] Bollobás, B. (1980) A probabilistic proof of an asymptotic formula for the number of labelled regular graphs. *European J. Combin* **1** 311–316.
- [5] Braunstein, A., Mézard, M. and Zecchina, R. (2005) Survey propagation: an algorithm for satisfiability. *Random Struct. Algor.* **27** 201–226.
- [6] Castellani, T., Napolano, V., Ricci-Tersenghi, F. and Zecchina, R. (2003) Bicolouring random hypergraphs. *J. Phys. A* **36** 11037–11053.
- [7] Coja-Oghlan, A. (2016) Constraint satisfaction: random regular k -SAT. In *Statistical physics, optimization, inference, and message-passing algorithms*. Oxford: Oxford Univ. Press, pp. 231–251.

- [8] Coja-Oghlan, A., Efthymiou, C., Jaafari, N., Kang, M. and Kapetanopoulos, T. (2018) Charting the replica symmetric phase. *Comm. Math. Phys* **359** 603–698.
- [9] Coja-Oghlan, A., Krzakala, F., Perkins, W. and Zdeborova, L. (2018) Information-theoretic thresholds from the cavity method. *Adv. Math* **333** 694–795.
- [10] Coja-Oghlan, A. and Panagiotou, K. (2016) The asymptotic k -SAT threshold. *Adv. Math* **288** 985–1068.
- [11] Cooper, C., Frieze, A., Molloy, M. and Reed, B. (1996) Perfect matchings in random r -regular, s -uniform hypergraphs. *Combin. Probab. Comput* **5** 1–14.
- [12] Dall'Asta, L., Ramezani, A. and Zecchina, R. (2008) Entropy landscape and non-Gibbs solutions in constraint satisfaction problems. *Phys. Rev. E* **77** 031118, 16.
- [13] Ding, J., Sly, A. and Sun, N. (2015) Proof of the satisfiability conjecture for large k . In *STOC'15—Proceedings of the 2015 ACM Symposium on Theory of Computing*. ACM, New York, pp. 59–68.
- [14] Ding, J., Sly, A. and Sun, N. (2016) Maximum independent sets on random regular graphs. *Acta Math* **217** 263–340.
- [15] Ding, J., Sly, A. and Sun, N. (2016) Satisfiability threshold for random regular NAE-SAT. *Comm. Math. Phys* **341** 435–489.
- [16] Erdős, P. and Rényi, A. (1960) On the evolution of random graphs. *Magyar Tud. Akad. Mat. Kutató Int. Közl* **5** 17–61.
- [17] Erdős, P. and Rényi, A. (1966) On the existence of a factor of degree one of a connected random graph. *Acta Math. Acad. Sci. Hungar* **17** 359–368.
- [18] Feige, U. (2002) Relations between average case complexity and approximation complexity. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing, STOC '02*, ACM, New York, NY, USA, pp. 534–543.
- [19] Frieze, A., Jerrum, M., Molloy, M., Robinson, R. W. and Wormald, N. C. (1996) Generating and counting Hamiltonian cycles in random regular graphs. *J. Algorithms* **21** 176–198.
- [20] Gabrié, M., Dani, V., Semerjian, G. and Zdeborová, L. (2017) Phase transitions in the q -coloring of random hypergraphs. *J. Phys. A* **50** 44.
- [21] Galanis, A., Štefankovič, D. and Vigoda, E. (2014) Inapproximability for antiferromagnetic spin systems in the tree non-uniqueness region. In *Proceedings of the Forty-sixth Annual ACM Symposium on Theory of Computing, STOC '14*, New York, NY, USA, pp. 823–831.
- [22] Higuchi, S. and Mézard, M. (2010) Correlation-based decimation in constraint satisfaction problems. *J. Phys. Conf. Ser.* **233** 012003.
- [23] Janson, S. (1995) Random regular graphs: asymptotic distributions and contiguity. *Combin. Probab. Comput* **4** 369–405.
- [24] Janson, S., Łuczak, T. and Ruciński, A. (2000) *Random graphs*. Wiley-Interscience, New York, Wiley-Interscience Series in Discrete Mathematics and Optimization.
- [25] Kahn, J. (2022) Hitting times for Shamir's problem. *Trans. Amer. Math. Soc* **375** 627–668.
- [26] Kalapala, V. and Moore, C. (2008) The phase transition in exact cover. *Chic. J. Theoret. Comput. Sci.* **5** 9, pages Article.
- [27] Kemkes, G., Pérez-Giménez, X. and Wormald, N. C. (2010) On the chromatic number of random d -regular graphs. *Adv. Math* **223** 300–328.
- [28] Krzakala, F., Ricci-Tersenghi, F., Zdeborová, L., Zecchina, R., Tramel, E. W. and Cugliandolo, L. F. (2016) *Statistical physics, optimization, inference, and message-passing algorithms*, Oxford: Oxford University Press, editors.
- [29] Kudekar, S., Richardson, T. and Urbanke, R. (2013) Spatially coupled ensembles universally achieve capacity under belief propagation. *IEEE Trans. Inform. Theory* **59** 7761–7813.
- [30] Maneva, E., Mossel, E. and Wainwright, M. J. (2007) A new look at survey propagation and its generalizations. *J. ACM* **54** 17.
- [31] Mézard, M. and Montanari, A. (2009) *Information, physics, and computation*. Oxford University Press, Oxford. Oxford Graduate Texts
- [32] Mézard, M., Parisi, G. and Virasoro, M. A. (1987) Spin glass theory and beyond. In *World Scientific Lecture Notes in Physics*, vol. **9**, Teaneck, NJ: World Scientific Publishing Co., Inc.
- [33] Molloy, M., Robalewska, H. D., Robinson, R. W. and Wormald, N. C. (1997) 1-factorizations of random regular graphs. *Random Structures Algorithms* **10** 305–321.
- [34] Moore, C. (2016) The phase transition in random regular exact cover. *Ann. Inst. Henri Poincaré D* **3** 349–362.
- [35] Mora, T. (2007) *Geometry and Inference in Optimization and in Information Theory*. Theses, Université Paris Sud - Paris XI.
- [36] Panagiotou, K. and Pasch, M. (2019) Satisfiability Thresholds for Regular Occupation Problems. In *46th Int. Col. on Automata, Languages, and Programming (ICALP '19)*, volume 132 of Leibniz International Proceedings in Informatics (LIPIcs), 90:1–90:14.
- [37] Robalewska, H. D. (1996) 2-factors in random regular graphs. *J. Graph Theory* **23** 215–224.
- [38] Robbins, H. (1955) A remark on Stirling's formula. *Amer. Math. Monthly* **62** 26–29.
- [39] Schmidt, C., Guenther, N. and Zdeborová, L. (2016) Circular coloring of random graphs: statistical physics investigation. *J. Stat. Mech. Theory Exp* **2016** 083303, 28.
- [40] Talagrand, M. (2003) *Spin glasses: a challenge for mathematicians*. volume **46** of Results in Mathematics and Related Areas. 3rd Series. A Series of Modern Surveys in Mathematics, Springer-Verlag, Berlin,
- [41] Yedidia, J. S., Freeman, W. T. and Weiss, Y. (2001) Bethe free energy, kikuchi approximations, and belief propagation algorithms, Technical Report TR2001-16, MERL - Mitsubishi Electric Research Laboratories, Cambridge, MA 02139, May 2001.

- [42] Zdeborová, L. and Krzakala, F. (2011) Quiet planting in the locked constraint satisfaction problems. *SIAM J. Discrete Math* **25** 750–770.
- [43] Zdeborová, L. and Mézard, M. (2008) Constraint satisfaction problems with isolated solutions are hard. *J. Stat. Mech. Theory Exp* **2008** P12004.

Appendix

A. Proof of Lemma 2.8

We present the proof of Lemma 2.8 in detail so as to facilitate the presentation of the small subgraph conditioning method in Section 6. Lemma 2.8 can be shown by a direct application of the method of moments, which is discussed, for example, in [24] (Theorem 6.10).

Theorem A.1 (Method of Moments). *Let $L \in \mathbb{Z}_{>0}$ and $((X_{\ell,i})_{\ell \in [L]})_{i \in \mathbb{Z}_{>0}}$ be a sequence of a vector of random variables. If $\lambda \in \mathbb{R}_{\geq 0}^L$ is such that, as $i \rightarrow \infty$,*

$$\mathbb{E} \left[\prod_{\ell=1}^L (X_{\ell,i})^{r_\ell} \right] \rightarrow \prod_{\ell=1}^L \lambda_\ell^{r_\ell}$$

for every $r \in \mathbb{Z}_{\geq 0}^L$, then $(X_{\ell,i})_{\ell \in [L]}$ converges in distribution to $(Z_\ell)_{\ell \in [L]}$, where the $Z_\ell \sim \text{Po}(\lambda_\ell)$ are independent Poisson distributed random variables.

First, we notice that $G = G_{k,d,n,m}$ and further $X_\ell = X_{k,d,n,\ell}$ is only defined for $m = dn/k \in \mathbb{Z}$ as stated in Lemma 3.1, hence Lemma 2.8 only applies to such sequences of configurations. Fix $k, d \in \mathbb{Z}_{>1}$. Before we turn to the general case we consider the $\mathbb{E}[X_\ell]$ for $\ell \in \mathbb{Z}_{>0}$. For this purpose let n and $m(n)$ be sufficiently large. Let $\mathcal{C}_{\ell,g}$ be the set of all 2ℓ -cycles in $g \in \mathcal{G}$. Then

$$\mathbb{E}[X_\ell] = \sum_{g \in \mathcal{G}} \frac{X_\ell(g)}{|\mathcal{G}|} = |\mathcal{G}|^{-1} \sum_{g \in \mathcal{G}} |\mathcal{C}_{\ell,g}| = \frac{|\mathcal{E}|}{|\mathcal{G}|}, \text{ where } \mathcal{E} = \{(g, c) : g \in \mathcal{G}, c \in \mathcal{C}_{\ell,g}\}.$$

With this at hand we obtain that

$$\mathbb{E}[X_\ell] = \frac{1}{2\ell(dn)!} n^\ell m^\ell (d(d-1))^\ell (k(k-1))^\ell (dn-2\ell)!$$

using the following combinatorial arguments. Instead of counting pairs (g, c) of configurations g and 2ℓ -cycles $c \in \mathcal{C}_{\ell,g}$ we count pairs (g, γ) of configurations g and directed 2ℓ -cycles γ (based at a variable node) in g . There are exactly 2ℓ directed cycles γ corresponding to each (undirected) cycle c of length 2ℓ since we can choose the base from the ℓ variables in c and γ is then determined by one of the two possible directions. The denominator reflects the compensation for this counting next to the probability $|\mathcal{G}|^{-1}$. Further, the term n^ℓ reflects the ordered choice of the variables for the directed cycle, as does m^ℓ for the constraints. The next two terms account for the choice of the two i -edges and a -edges traversed by the cycle for each of the ℓ variables i and constraints a . This fixes the directed cycle γ and further the corresponding undirected cycle $c(\gamma)$. In particular, the 2ℓ edges of the cycle c in g are fixed, i.e. the corresponding restriction of g to c . This leaves us with $(dn-2\ell)$ half-edges in $\text{dom}(g)$ and $(km-2\ell)$ half-edges in $\text{im}(g)$ that have not been wired yet. The last term gives the number of such wirings.

Next, we turn to asymptotics. Extracting λ_ℓ and expanding the falling factorials yields

$$\mathbb{E}[X_\ell] = \lambda_\ell d^\ell k^\ell \frac{n! m! (dn-2\ell)!}{(dn)!(n-\ell)!(m-\ell)!}.$$

Using Stirling's formula we readily obtain that

$$\mathbb{E}[X_\ell] \sim \lambda_\ell d^\ell k^\ell \sqrt{\frac{nm(dn-2\ell)}{dn(n-\ell)(m-\ell)}} \frac{n^n m^m (dn-2\ell)^{dn-2\ell}}{(dn)^{dn} (n-\ell)^{n-\ell} (m-\ell)^{m-\ell}},$$

and so

$$\mathbb{E}[X_\ell] \sim \lambda_\ell d^\ell k^\ell \sqrt{\frac{(1 - \frac{2\ell}{dn})}{(1 - \frac{\ell}{n})(1 - \frac{\ell}{m})}} \frac{n^\ell m^\ell (1 - \frac{2\ell}{dn})^{dn-2\ell}}{(dn)^{2\ell} (1 - \frac{\ell}{n})^{n-\ell} (1 - \frac{\ell}{m})^{m-\ell}} \sim \lambda_\ell d^\ell k^\ell \frac{n^\ell m^\ell}{(dn)^{2\ell}}.$$

Using that $dn = km$ leads to

$$\mathbb{E}[X_\ell] \sim \lambda_\ell d^\ell k^\ell \frac{n^\ell (dk^{-1}n)^\ell}{(dn)^{2\ell}} = \lambda_\ell,$$

as claimed. We turn to the general case. For this purpose let $L \in \mathbb{Z}_{>0}$, $r \in \mathbb{Z}_{\geq 0}^L$ and let n and m be sufficiently large. Then

$$X_\ell(g)^{r_\ell} = \prod_{s=0}^{r_\ell-1} (|\mathcal{C}_{\ell,g}| - s) = |\mathcal{C}_{\ell,r_\ell,g}|, \text{ where } \mathcal{C}_{\ell,r_\ell,g} = \{c \in \mathcal{C}_{\ell,g}^{r_\ell} : \forall s \in [r_\ell] \forall s' \in [s-1] c_s \neq c_{s'}\}$$

for $g \in \mathcal{G}$, since this corresponds to an ordered choice of 2ℓ -cycles in g without repetition. The product can then be directly written as

$$\prod_{\ell=1}^L X_\ell(g)^{r_\ell} = |\mathcal{C}_{r,g}|, \text{ where } \mathcal{C}_{r,g} = \prod_{\ell=1}^L \mathcal{C}_{\ell,r_\ell,g}.$$

To avoid double indexed sequences we use the equivalent representation $c = (c_s)_{s \in [\tau]} \in \mathcal{C}_{r,g}$ where $\tau = \sum_{1 \leq \ell \leq L} r_\ell$. From the above we see that the cycles c_s are ordered by their length ℓ_s in ascending order and are pairwise distinct. We obtain that

$$\mathbb{E} \left[\prod_{\ell=1}^L X_\ell^{r_\ell} \right] = \frac{|\mathcal{E}|}{|\mathcal{G}|}, \text{ where } \mathcal{E} = \{(g, c) : g \in \mathcal{G}, c \in \mathcal{C}_{r,g}\}.$$

Since we have ℓ_s distinct variables and constraints in each cycle c_s respectively, we can have at most $\mathfrak{l} = \sum_{s \in [\tau]} \ell_s$ distinct variables and constraints in c . Specifically, we only have $|V(c)| = \mathfrak{l}$ variables and $|F(c)| = \mathfrak{l}$ constraints iff all cycles c_s are disjoint. So, let

$$\mathcal{E}_0 = \{(g, c) \in \mathcal{E} : |V(c)| = |F(c)| = \mathfrak{l}\}$$

denote the set of pairs $(g, c) \in \mathcal{E}$ with disjoint cycles and further $\mathcal{E}_1 = \mathcal{E} \setminus \mathcal{E}_0$ the remaining pairs. Then we have

$$\frac{|\mathcal{E}_0|}{|\mathcal{G}|} = \frac{1}{(dn)! \prod_{s=1}^{\tau} (2\ell_s)} n^{\mathfrak{l}} m^{\mathfrak{l}} (d(d-1)^{\mathfrak{l}} (k(k-1))^{\mathfrak{l}} (dn-2\mathfrak{l})!)$$

for the following reasons. For each cycle c_s in c counting the $2\ell_s$ directed cycles facilitates the computation, hence we find the corresponding product in the denominator. Since the variables within each directed cycle and the cycles in the sequence are ordered we have an ordered choice of all variables. Further, since the ℓ_s variables within each cycle are distinct and the cycles are pairwise disjoint we choose all variables without repetition. This explains the first falling factorial. The next term for the constraints follows analogously. But since variables and constraints are disjoint the edges are too, hence we choose two edges for each of the \mathfrak{l} variables and constraints respectively. Then we wire the remaining edges.

The asymptotics are derived analogously to the base case, i.e.

$$\frac{|\mathcal{E}_0|}{|\mathcal{G}|} \sim \frac{(d-1)^{\mathfrak{l}} (k-1)^{\mathfrak{l}}}{\prod_{s=1}^{\tau} (2\ell_s)} = \prod_{s=1}^{\tau} \lambda_{\ell_s} = \prod_{\ell=1}^L \lambda_{\ell}^{r_\ell},$$

using the definition of $c = (c_s)_{s \in [\tau]}$ in the last step. Since the contribution of the disjoint cycles already yields the desired result, we want to show that the contribution of intersecting cycles is

negligible. As before, we count directed cycles γ_s and adjust the result accordingly, so let

$$\mathcal{E}_2 = \{(g, \gamma) : (g, c(\gamma)) \in \mathcal{E}_1\}, \text{ i.e. } |\mathcal{E}_2| = |\mathcal{E}_1| \prod_{s \in [\tau]} (2\ell_s).$$

In the next step we consider the *relative position representations* (α, ρ) of sequences γ of directed cycles. Instead of a formal introduction we illustrate this concept in Figure 3. The corresponding decomposition of the contributions to the expectation according to ρ is

$$\frac{|\mathcal{E}_1|}{|\mathcal{G}|} = \sum_{\rho \in \mathcal{R}} \frac{|\mathcal{E}_\rho|}{|\mathcal{G}| \prod_{s \in [\tau]} (2\ell_s)}, \mathcal{E}_\rho = \{(g, \gamma) \in \mathcal{E}_2 : \rho(\gamma) = \rho\}, \mathcal{R} = \{\rho(\gamma) : (g, \gamma) \in \mathcal{E}_2\}.$$

For the following reasons we can then derive

$$|\mathcal{E}_\rho| = n^{\underline{n(\rho)}} m^{\underline{m(\rho)}} \prod_{j \in [n(\rho)]} d^{\underline{d_j(\rho)}} \prod_{b \in [m(\rho)]} k^{\underline{k_b(\rho)}} (dn - e(\rho))!.$$

Since ρ is fixed, we have to fix the absolute values α , thereby the directed cycle γ , and wire the remaining edges. But the first four terms exactly correspond to the number of choices for the index vectors in α . This fixes γ , further the union $c(\gamma)$ of cycles and in particular $e(\rho)$ edges. The remaining term counts the number of choices to wire the remaining edges.

For the asymptotics we notice that $n(\rho), m(\rho) \leq l$ and that also the two products are bounded in both the multiplication region and values. But this further implies that $|\mathcal{R}|$ is bounded, i.e. the summation region is also finite in the limit and hence we can consider the asymptotics of each term separately, which yields

$$\begin{aligned} \frac{|\mathcal{E}_1|}{|\mathcal{G}|} &= \sum_{\rho \in \mathcal{R}} \frac{\prod_{i \in [n(\rho)]} d^{\underline{d_i(\rho)}} \prod_{a \in [m(\rho)]} k^{\underline{k_a(\rho)}} n^{\underline{n(\rho)}} m^{\underline{m(\rho)}} (dn - e(\rho))!}{\prod_{s \in [\tau]} (2\ell_s) (dn)!} \\ &= \sum_{\rho \in \mathcal{R}} c_1(\rho) \frac{n^{\underline{n(\rho)}} m^{\underline{m(\rho)}} (dn - e(\rho))!}{(dn)!} \\ &\sim \sum_{\rho \in \mathcal{R}} c_1(\rho) \left(\frac{1}{e}\right)^{n(\rho)} \left(\frac{d}{ke}\right)^{m(\rho)} \left(\frac{e}{d}\right)^{e(\rho)} n^{n(\rho)+m(\rho)-e(\rho)} \\ &= \sum_{\rho \in \mathcal{R}} c_2(\rho) n^{n(\rho)+m(\rho)-e(\rho)}, \end{aligned}$$

where we summarized the terms that only depend on ρ into constants. Now, let $\rho \in \mathcal{R}$ and let $c = c(\rho)$ be the graph of ρ as introduced in Section 2.4. Since ρ is a sequence of directed cycles that are not all disjoint, its graph c is the union of the corresponding (undirected) cycles that are not all disjoint. But then c has more edges than vertices, i.e. $3e(\rho) > n(\rho) + m(\rho) + 2e(\rho)$, and hence

$$\frac{|\mathcal{E}_1|}{|\mathcal{G}|} \sim \sum_{\rho \in \mathcal{R}} c_2(\rho) n^{n(\rho)+m(\rho)-e(\rho)} \leq n^{-1} \sum_{\rho \in \mathcal{R}} c_2(\rho) = c_3 n^{-1},$$

which shows that this contribution is negligible. This establishes the asymptotic equivalence

$$\mathbb{E} \left[\prod_{\ell \in [L]} X_\ell^{r_\ell} \right] \sim \prod_{\ell \in [L]} \lambda_\ell^{r_\ell}$$

and allows to apply the method of moments, which directly yields Lemma 2.8.