Statistical inference on random factor graphs

Dissertation

Zur Erlangung des Doktorgrades

der Naturwissenschaften

vorgelegt beim Fachbereich Informatik und Mathematik

der Johann Wolfgang Goethe-Universität

in Frankfurt am Main

von

Philipp Johannes Michael Loick

aus Münster

Frankfurt 2021

(D 30)

vom Fachbereich Informatik und Mathematik der

Johann Wolfgang Goethe-Universität als Dissertation angenommen.

Dekan:

Prof. Dr. Lars Hedrich

Gutachter:

Prof. Dr. Amin Coja-Oghlan

Prof. Dr. Oliver Johnson

Datum der Disputation:

STATISTICAL INFERENCE ON RANDOM FACTOR GRAPHS

CONTENTS

1. Acknowledgements

First and foremost, I would like to thank Amin Coja-Oghlan for giving me the chance to pursue a PhD in mathematics. Not only did he take the risk of accepting me as a PhD student without us having worked together. He also invested significant time to help me through the administrative hurdles of a German university. Looking back at the past three years, I am also thankful for what I was able to learn from him about mathematics and computer science. Amin continuously supported us and provided invaluable input for solving the research questions that this dissertation is devoted to. He also helped us to turn vague ideas into concise mathematical proofs. Without him, this dissertation would not have been possible. At this point, I would also like to thank Gregory Sorkin from the London School of Economics who supervised me during my Master thesis and set up the initial contact to Amin. I was also pleased to be part of Amin's research team and would like to say thanks for an enjoyable work atmosphere to Oliver Gebhard, Max Hahn-Klimroth, Joon Lee, Noela Müller, Jean Ravelomanana and Maurice Rolvien. Finally, I am thankful to my wife Sophia who encouraged me to start the PhD in the first place and was a continuous source of motivation throughout the last three years.

2. Introduction

2.1. **Motivation.** Many fundamental questions in mathematics and computer science can be cast as statistical inference problems on random factor graphs. Consider the well-known problem of sparse high-dimensional linear regression where we observe m noisy measurements of the form

$$y = X\beta + w \tag{2.1}$$

with $y \in \mathbb{R}^m, X \in \mathbb{R}^{mxn}$ having iid $N(0,1)$ entries, a k-sparse parameter vector $\beta \in \{0,1\}^n$ and a noise vector $w \in \mathbb{R}^m$ with iid $N(0,\sigma^2)$ entries [62]. This setup can readily be envisioned as a bipartite graph where on the one side, we have n variables nodes representing the entries of the parameter vector β and m factor nodes representing the noisy observations y. Edges between variable and factor nodes encode the values of the input matrix X. Due to the randomness of the graph structure, we call such a construction a *random factor graph* [57]. Given knowledge of the graph structure and the noisy measurements residing on the factor nodes, the goal is to approximately recover the k-sparse parameter vector β with as few measurements m as possible.

Graph clustering serves as a second fundamental example. The problem can be best introduced using the instructive and intensely-studied *stochastic block model* [1, 43]. In its most basic form, we have a vertex set $V_n = \{x_1, \dots, x_n\}$ and pick a label ± 1 for each vertex uniformly at random giving rise to a *planted* partition of the vertex set into two clusters encoded by $\boldsymbol{\sigma} \in \{\pm 1\}^{V_n}$. Next, we connect any two vertices by an edge with probability p, if the two vertices have the same label, and with probability q otherwise. If $p > q$, edges between vertices of the same label are preferred - the assortative stochastic block model. The case $p < q$ is called the disassortative stochastic block model. Upon insertion of the edges, we remove the original labelling on the vertices. The key question is whether it is possible to recover a non-trivial approximation of the original clusters from only observing the graph structure. Put differently, what is the minimum difference between p and q such that recovery is possible. Again, this problem can be readily understood as an inference problem on a random factor graph where the vertex set V_n serves as the variable nodes and the set of edges as factor nodes. There is an edge between a factor node and a variable node if the relevant vertex is adjacent to the edge that the factor node represents.

Other prominent examples that can be described in terms of statistical inference problems on random factor graphs include principal component analysis [11], the planted clique problem [36] or constraint satisfaction [13, 26]. This list can be continued indefinitely. In this dissertation, we explore techniques for determining information-theoretic and algorithmic thresholds for such inference problems on the example of three prime inference problems: binary group testing, quantitative group testing and the Ising antiferromagnet and max cut on random regular graphs. Before we proceed to our results, let us first lay the foundation.

2.2. **Random factor graphs.** We introduced random factor graphs above for sparse high-dimensional linear regression and graph clustering. Let us abstract from these specific examples and discuss a general framework for factor graphs that is quite powerful and can express many fundamental problems in combinatorics, computer science and physics [57, 73]. A factor graph G is a bipartite graph consisting of a set of variables nodes V_n and factor nodes F_m. We assume the variables to range over a finite set Ω of size $q = |\Omega| \geq 2$. In the following, an assignment of labels to vertices will be denoted a configuration $\sigma \in \Omega^{V_n}$ which is sampled from some prior distribution. While the variable nodes represent the variables of the inference problem such as the parameter vector β in high-dimensional regression or the labels of the vertices in graph clustering, the factor nodes describe the interaction of the variable

nodes. This interaction can take the form of a linear combination as for high-dimensional regression or simply describe the edges between vertices in graph clustering. Each factor node $a \in F_m$ is associated with a function $\psi_a : \Omega^{\partial a} \to (0,\infty)$ that assigns a positive weight to each factor from the adjacent variables' value combinations. The factor graph induces a probability distribution over the label configurations via

$$\mu_G(\sigma) = \frac{\psi_G(\sigma)}{Z_G} \quad \text{with} \quad \psi_G(\sigma) = \prod_{a \in F_m} \psi_a(\sigma_{\partial a}) \quad \text{and} \quad Z_G = \sum_{\tau \in \Omega^{V_n}} \psi_G(\tau) \tag{2.2}$$

for each $\sigma \in \Omega^{V_n}$. The normalizing term Z_G is the partition function. It is intimately related to the information-theoretic threshold of many inference problems and thus of fundamental importance. We will discuss this aspect in further depth below.

Given information on the interaction between variable and factor nodes, we can now easily construct a factor graph from the variable nodes V_n and factor nodes F_m. In many interesting cases, this graph construction will exhibit some kind of randomness such as a random sequence of variable and factor degrees or an edge probability between any variable and factor node. We will denote such a random factor graph model $\mathbb{G}$ as the null model. Note that the construction of $\mathbb{G}$ in this definition is independent of the weights of the factor nodes. At the same time, $\mathbb{G}$ induces a reweighted graph distribution $\hat{\mathbb{G}}$ and given $\sigma \in \Omega^{V_n}$ a planted model $\mathbb{G}^*$ [14, 73] which for any event $\mathscr{A}$ are defined by

$$\mathbb{P}\left[\hat{\mathbb{G}} \in \mathscr{A}\right] = \frac{\mathbb{E}\left[Z_{\mathbb{G}} \mathbf{1}\left\{\mathbb{G} \in \mathscr{A}\right\}\right]}{\mathbb{E}\left[Z_{\mathbb{G}}\right]} \qquad \text{and} \qquad \mathbb{P}\left[\mathbb{G}^*(\sigma) \in \mathscr{A}\right] = \frac{\mathbb{E}\left[\psi_{\mathbb{G}}(\sigma)\mathbf{1}\left\{\mathbb{G} \in \mathscr{A}\right\}\right]}{\mathbb{E}\left[\psi_{\mathbb{G}}(\sigma)\right]}. \tag{2.3}$$

Along the same lines, we can define a distribution $\hat{\boldsymbol{\sigma}} \in \Omega^{V_n}$ on label configurations by

$$\mathbb{P}\left[\hat{\boldsymbol{\sigma}} = \sigma\right] = \frac{\mathbb{E}\left[\psi_{\mathbb{G}}(\sigma)\right]}{\sum_{\tau \in \Omega^{V_n}} \mathbb{E}\left[\psi_{\mathbb{G}}(\tau)\right]}. \tag{2.4}$$

Finally, let $\boldsymbol{\sigma} \in \Omega^{V_n}$ be a label configuration sampled uniformly at random. $\hat{\mathbb{G}}, \mathbb{G}^*, \hat{\boldsymbol{\sigma}}, \boldsymbol{\sigma}$ and the distribution from (2.2) are connected with the well-known Nishimori property.

Fact 2.1 (Proposition 3.2 in [22])**.** *For any graph G and spin configuration $\sigma \in \Omega^{V_n}$ we have*

$$\mathbb{P}\left[\hat{\mathbb{G}} = G\right] \mu_G(\sigma) = \mathbb{P}\left(\hat{\boldsymbol{\sigma}} = \sigma\right) \mathbb{P}\left(\mathbb{G}^* = G \mid \boldsymbol{\sigma} = \sigma\right).$$

Let us put these definitions into perspective by considering the concrete example of the disassortative stochastic block model. For some real parameter $\beta > 0$ and $d \geq 3$, we assign a label $\{\pm 1\}$ to each vertex uniformly at random and assume the probabilities of intra- and inter-community edges to be

$$p = \frac{2de^{-\beta}}{n\left(1 + e^{-\beta}\right)} \qquad \text{and} \qquad q = \frac{2d}{n\left(1 + e^{-\beta}\right)}$$

such that the resulting graph is sparse and vertices asymptotically have average degree d as $n \to \infty$ [14]. In the terminology of above, the resulting graph is the *planted* model. The corresponding null model $\mathbb{G}$ is a graph with $p = q = d/n$, i.e. a plain Erdős-Rényi graph where two vertices are connected by an edge with some fixed probability irrespective of their label. Let $E(G)$ denote the set of edges in a graph G. Along the lines of (2.2) we can define a

distribution over label configurations $\sigma \in \{\pm 1\}^{V_n}$ for any graph G given by

$$\mu_{G,\beta}(\sigma) = \frac{\psi_G(\sigma)}{Z_{G,\beta}} = \frac{\exp\left(-\beta \mathscr{H}_G(\sigma)\right)}{Z_{G,\beta}} \tag{2.5}$$
$$\text{where} \quad \mathscr{H}_G(\sigma) = \sum_{(v,w)\in E(G)} \mathbf{1}\{\sigma_v = \sigma_w\}$$
$$\text{and} \quad Z_{G,\beta} = \sum_{\tau\in\{\pm 1\}^{V_n}} \exp\left(-\beta \mathscr{H}_G(\tau)\right).$$

The term $\mathscr{H}_G(\sigma)$ counting the number of monochromatic edges in G is typically referred to as the Hamiltonian and $\mu_{G,\beta}$ the Boltzmann distribution for this problem. We can equivalently define $\hat{\mathbb{G}}$ and $\hat{\boldsymbol{\sigma}}$. Note that the real parameter β governs the graph structure in the planted model and the probability of observing a configuration under distribution $\mu_{G,\beta}$. For $\beta = 0$, no penalty term is imposed for monochromatic edges and the planted model $\mathbb{G}^*$ coincides with the plain Erdős-Rényi model $\mathbb{G}$. In this case, $\mu_{G,\beta}$ is simply the uniform distribution over all configurations $\sigma \in \{\pm 1\}^{V_n}$. As β increases, fewer and fewer edges between vertices of the same label will be observed under $\mathbb{G}^*$ and the distribution $\mu_{G,\beta}$ concentrates on configurations with few monochromatic edges.

2.3. **Teacher-student model.** With the definition of a random factor graph at hand, the teacher-student model provides an instructive analogy to understand statistical inference on random factor graphs [73]. Suppose we have a teacher who selects a configuration $\boldsymbol{\sigma} \in \Omega^{V_n}$ from some prior distribution. We will refer to this planted configuration as the ground truth. For high-dimensional linear regression, the ground truth is the parameter vector β, while for graph clustering it is the partition of the vertex set into two classes. Thereafter, the planted model $\mathbb{G}^*$ is constructed given the ground truth according to (2.3), i.e., noisy measurements are calculated according to (2.1) for high-dimensional linear regression and vertices are connected with intra- and inter-cluster probabilities for graph clustering. Thereafter, the teacher hands $\mathbb{G}^*$ to the student without providing the ground truth $\boldsymbol{\sigma}$. The student might have access to the prior distribution of $\boldsymbol{\sigma}$ and the distribution of $\mathbb{G}^*$ given $\boldsymbol{\sigma}$ which is called the Bayes optimal case [73]. In other problems, there even might be a mismatching prior and/or model [73]. Now the key question is how much of a mark does $\boldsymbol{\sigma}$ leave on $\mathbb{G}^*$? Put differently, can the student (approximately) recover the ground truth $\boldsymbol{\sigma}$ just using the planted model $\mathbb{G}^*$ and information on the prior of $\boldsymbol{\sigma}$ and the distribution of $\mathbb{G}^*$ given $\boldsymbol{\sigma}$?

It should be clear that the answer depends on the signal-to-noise ratio in the underlying problem. For high-dimensional linear regression, the variance σ^2 of the noise vector clearly impedes the ability of the student to recover β. Similarly, the number of measurement m can be seen as the signal - the more noisy measurements are conveyed, the easier the inference task [62]. For graph clustering, we can think of $\exp(-\beta)$ as the noise parameter. Similarly, the larger the average degree d, the more is revealed about the ground truth $\boldsymbol{\sigma}$ [14]. Thus, the fundamental task in statistical inference problems is to determine signal-to-noise thresholds from which (approximate) recovery of the ground truth is possible.

2.4. **Presumed computational hardness.** So far, we avoided the question of what we mean when we say that (approximate) recovery is possible. Let us now close this gap. To get started, note that how we measure recovery is often problem-specific. For graph clustering, it might be a configuration estimate that has a non-trivial overlap with the ground truth. For other problems such as group testing which we discuss below we need to infer the correct configuration with high probability over the randomness of the graph [1].

[1] We say that a sequence of events $\mathscr{E}_n$ holds with high probability (w.h.p.), if $\lim_{n\to\infty} \mathbb{P}[\mathscr{E}_n] = 1$.

While how we measure recovery is often problem-specific, the question of finding thresholds from which recovery is possible typically comes in two instalments [73]. First, there is the *information-theoretic* view which asks for the minimum signal-to-noise ratio such that recovery is in principle possible. Put differently, the information-theoretic perspective is interested in the threshold from which sufficient information is contained in $\mathbb{G}^*$ to recover $\boldsymbol{\sigma}$ irrespective of computational resources. Second, the *algorithmic* threshold is concerned with the threshold from which efficient recovery is possible, i.e., a polynomial-time algorithm is able to recover $\boldsymbol{\sigma}$ from $\mathbb{G}^*$. One might be tempted to think that the ability to recover $\boldsymbol{\sigma}$ increases continuously with the signal-to-noise ratio, such as with the number of measurements for high-dimensional linear regression. Indeed, many statistical inference problems exhibit phase transitions where all the way up to a threshold inference is impossible just to become possible once we cross the threshold [57]. Moreover, one might be tempted to think that the information-theoretic and algorithmic thresholds for inference problems always coincide. Interestingly, many statistical inference problems appear to undergo an impossible-hard-easy transition where below some threshold m_{inf} the problem cannot even be solved information-theoretically (*impossible* regime). Conversely, above some threshold m_{alg} efficient algorithms are known to solve the problem (*easy* regime). When these two bounds do not coincide, one is faced with a regime where the problem in principle contains sufficient information to recover the ground truth, but no efficient algorithms are known to succeed (*hard* regime). Prominent examples are the planted clique problem [36], code division multiple access (CDMA) [73], the pooled data problem [3], sparse principal component analysis (PCA) [11] or sparse high-dimensional regression [37]. In recent years, increasing evidence has been put forward that such computational gaps might indeed be due to genuine computational hardness in the underlying problem. In other inference problems such as the stochastic block model [1] or compressed sensing [29], we do not find such a gap. The phenomenon of computational hardness is still far from being understood and continues to be a field of active research of fundamental importance in mathematics and computer science. A key building block towards this end is to derive sharp information-theoretic and algorithmic thresholds in inference problems.

Based on the celebrated, yet non-rigorous *cavity* method, physicists have put forward a number of striking predictions for such phase transitions [57]. In recent years, mathematicians made progress to turn heuristic arguments into proofs and develop techniques that allow the rigorous vindications of some of these predictions (see i.e., [2, 8, 34, 35, 42, 70] - to name a few). While many gaps still remain, the versatility of these novel techniques allows the application to a wide range of problems. A significant body of this dissertation is devoted to adapt techniques from mathematical physics and put them to use in classical problems in combinatorics and information theory. Let us dive into these techniques that mathematicians today have at their disposal to derive information-theoretic and algorithmic thresholds in inference problems.

2.5. **Information-theoretic techniques.** In the teacher-student analogy, the information-theoretic view is concerned with the imprint that the ground truth $\boldsymbol{\sigma}$ leaves on the factor graph $\mathbb{G}^*$ that is conveyed to the student. The key quantity in this regard is the mutual information between $\boldsymbol{\sigma}$ and $\mathbb{G}^*$

$$I(\boldsymbol{\sigma}, \mathbb{G}^*) = \sum_{\sigma \in \Omega^{Vn}, G} \mathbb{P}\left[\boldsymbol{\sigma} = \sigma, \mathbb{G}^* = G\right] \log \frac{\mathbb{P}\left[\boldsymbol{\sigma} = \sigma, \mathbb{G}^* = G\right]}{\mathbb{P}\left[\boldsymbol{\sigma} = \sigma\right] \mathbb{P}\left[\mathbb{G}^* = G\right]}.$$

A easy-to-derive, yet striking insight is that the mutual information is intimately related to the partition function of $\mathbb{G}^*$. Assuming a fixed factor node degree k and a fluctuating variable node degree $\boldsymbol{d}$ and letting $\Lambda(x) = x\log(x)$ we obtain the following relationship.

Lemma 2.2 (Proposition 3.1 in [22])**.** *Under certain symmetry conditions we have w.h.p.*

$$I(\boldsymbol{\sigma}, \mathbb{G}^*)/n = \log q + \frac{\mathbb{E}[\boldsymbol{d}]}{k} q^{-k} \sum_{\tau \in \Omega^k} \Lambda(\psi(\tau)) - \mathbb{E}\left[\log Z_{\mathbb{G}^*}\right]/n + o(1).$$

Thus, calculating the mutual information between $\boldsymbol{\sigma}$ and $\mathbb{G}^*$ boils down to calculating $\mathbb{E}\left[\log Z_{\mathbb{G}^*}\right]$ - the free energy in physics jargon. However, since the expectation is outside the logarithm, this task is far from trivial and turns out to be a considerable challenge. By Jensen's inequality and standard results [57] we know that

$$\log \mathbb{E}[Z_{\mathbb{G}}] \leq \mathbb{E}\left[\log Z_{\mathbb{G}^*}\right] \leq \log \mathbb{E}[Z_{\mathbb{G}^*}] \tag{2.6}$$

In comparison with the free energy, calculating the logarithm of the expectation of the partition function is much more amenable. For certain regimes in many inference problems, we encounter the fortunate situation where (2.6) holds with equality and puts us in the convenient position that we can calculate the first moment and derive sharp information-theoretic results. Both the binary and the quantitative group testing problem that we consider in this dissertation fall into this category. In those situations, we calculate the first moment of the partition function in the planted model, i.e., the number of alternative configurations $\sigma \in \Omega^{V_n}$ that are consistent with the factor graph generated by the ground truth. To facilitate this calculation, the first moment is typically calculated in two steps. First, we consider configurations exhibiting a small overlap with the ground truth and show that their number is vanishing above the information-theoretic threshold (see also [4]). Second, local stability arguments are employed to rule out any alternative configurations with a large overlap with the true configuration. We will dive further into this point in subsequent sections.

Unfortunately, we often do not have the luxury that for the interesting regime (2.6) holds with equality. The maximum cut on random regular graphs that we also consider in this dissertation is a case in point. In the relevant regime, the bounds from (2.6) are not tight and we need to resort to different methods. Two techniques have emerged in recent years that allow tackling the free energy. While quite elaborate and challenging, both essentially boil down to local calculations that trace the impact on the partition function from a small number of local changes. The first is a coupling technique known as the Aizenman-Sims-Starr scheme which provides an upper bound on the free energy [2]. The second is known as the interpolation method which provides a lower bound [34, 42]. In many inference problems, the bounds derived from both techniques coincide and thus provide a tight expression for the free energy. It should be noted that the resulting expression typically comes as a variational formula that requires optimisation over a functional called the Bethe functional. We will encounter this exact situation when deriving a bound on the maximum cut size on random regular graphs. Fortunately, in our case we can throw a bridge between the specific functional and random walks to derive an explicit bound. However, in most situations one cannot hope for a simpler formulation than the general optimisation problem which precisely describes the intricate dependencies between variable and factor nodes.

2.6. **Algorithmic techniques.** Having established the information-theoretic thresholds of a problem, the key question is whether efficient algorithms exist that attain these bounds resulting in a simple impossible-easy phase transition with no regime of presumed computational hardness. Algorithmic approaches to inference problems are numerous and often problem-specific, so we will focus here on a powerful message-passing algorithm called belief propagation that is well compatible with the notion of factor graphs [57]. The significance of belief propagation reveals itself in the fact that for many inference problems it works all the way down to the information-theoretic threshold (see i.e., [29] in the case of

compressed sensing). Even if it falls short of the information-theoretic bound, there are typically no other efficient algorithms known that improve beyond the algorithmic threshold of belief propagation.

Following along the lines of [57], let $\mathscr{P}(\Omega)$ denote the set of all probability distributions over Ω. We take a factor graph G and define a message space $\mathscr{M}(G)$ consisting of all sets of messages $\nu = (\nu_{x\to a}, \nu_{a\to x})_{x\in V, a\in F, x\in\partial a}$ where $\nu_{x\to a}, \nu_{a\to x} \in \mathscr{P}(\Omega)$. The belief propagation operator now maps some $\nu \in \mathscr{M}(G)$ to $\hat{\nu} \in \mathscr{M}(G)$ according to

$$\hat{\nu}_{x\to a}(\sigma) = \frac{\prod_{b\in\partial x\setminus a} \hat{\nu}_{b\to x}(\sigma)}{\sum_{\tau\in\Omega}\prod_{b\in\partial x\setminus a}\hat{\nu}_{b\to x}(\tau)} \qquad \text{and}$$
$$\hat{\nu}_{a\to x}(\sigma) = \frac{\sum_{\tau\in\Omega^{\partial a}} \mathbf{1}\{\tau_x = \sigma\}\psi_a(\tau)\prod_{y\in\partial a\setminus x}\nu_{y\to a}(\tau_y)}{\sum_{\tau\in\Omega^{\partial a}}\psi_a(\tau)\prod_{y\in\partial a\setminus x}\nu_{y\to a}(\tau_y)}$$

for $\sigma \in \Omega$. The belief propagation algorithm consists of iteratively applying the above operator. The hope is that upon convergence, the messages from factor to variable nodes defined on the probability distribution over Ω are reminiscent of the ground truth. To this end, we define the marginal

$$\nu_x(\sigma) = \frac{\prod_{b\in\partial x}\hat{\nu}_{b\to x}(\sigma)}{\sum_{\tau\in\Omega}\prod_{b\in\partial x}\hat{\nu}_{b\to x}(\tau)} \qquad (\sigma\in\Omega).$$

It is straightforward to prove that belief propagation converges on factors graphs without cycles [57]. Remarkably, it seems to also perform well on factor graphs that do contain cycles but are locally tree-like. While we still lack a comprehensive general understanding of belief propagation on cyclic factor graphs, for some specific models with a locally tree-like structure it could be shown that belief propagation converges - see [19] for an example. For other problems, it is possible to retract to problem-specific algorithms that are inspired by belief propagation, but are easier to analyse. In fact, we will present such an algorithm for the binary group testing problem that essentially performs the first update round of belief propagation. With an additional clean-up step, we thereby find a combinatorially meaningful algorithm that solves the binary group testing problem down to the information-theoretic threshold. In certain situations when the inference problem is dense rather than sparse, physics intuition suggests that we can meaningfully enhance the efficiency of belief propagation while maintaining its reliability [9]. The crucial insight which is far from being rigorously established is that for some dense inference problems including the recipient node in the message update only adds a negligible error term. Thus, we merely need to calculate $|V_n| + |F_m|$ messages in each round rather than $|V_n| \times |F_m|$. The resulting algorithm trades under the name of *approximate message passing* and has been, among others, applied to the quantitative group testing problem which we will discuss in due course [3].

An important question that we have left out in the above section is the starting point of belief propagation. The default approach is to initialise belief propagation with uniform messages or messages incorporating information about the prior distribution. However, it turns out that in many inference problems, we do not just have one fix point to which belief propagation converges, but two - a trivial one which does not contain any meaningful information about the ground truth and a non-trivial one that is close to the desired ground truth [57]. If we initialise belief propagation with uniform messages, we often find ourselves trapped in this trivial fix point without any hope of getting to the non-trivial desirable fix point. Conversely, if we had some means to initialise the messages with the ground truth, belief propagation would take us to the non-trivial fix point. However, for inference problems we do not have access to the ground truth - that would defy the entire point. Fortunately, it turns out that belief propagation often already takes us to the non-trivial fix point

if we have some clue about the ground truth. If we have control over the graph structure like in compressed sensing or group testing, the notion of spatial coupling provides redemption. Pioneered in the field of coding theory [48, 49, 50], spatial coupling is concerned with enforcing a geometric structure on the graph $\mathbb{G}^*$. Instead of trying to infer the entire ground truth in one attempt, we proceed step-wise and divide the vertex and factor nodes up into compartments where we only allow edges between variables and factor nodes in nearby compartments. In effect, if we describe the factor graph by means of an adjacency matrix, this geometric structure gives rise to a band matrix. For the first couple of compartments we provide extra measurements that are negligible in the grand scheme of things, but which allow us to easily recover the ground truth for these variable nodes. Then, we proceed from compartment to compartment inferring the ground truth with knowledge of the ground truth of previous compartments. The (approximate) knowledge of the ground truth in prior departments provides us with a starting point that is reminiscent of the overall ground truth and thus facilitates inference. In groundbreaking work, this idea in conjunction with belief propagation has been put to use to show that the compressed sensing problem can be efficiently solved above the information-theoretic threshold [29]. In this dissertation, we provide a second example where the notion of spatial coupling allows us to close the gap between the information-theoretic and algorithmic thresholds, thereby solving a long-standing open problem.

2.7. **Contribution and outlook.** As laid out above, we will consider three classical inference problems and provide novel results on their information-theoretic and algorithmic thresholds. Section 3 will be devoted to the binary group testing problem where the goal is to recover a small set of infected individuals in a large population by a pooled testing scheme. In the first paper

Information-theoretic and algorithmic threshold for group testing

by Amin Coja-Oghlan, Oliver Gebhard, Max Hahn-Klimroth and Philipp Loick, we establish a matching information-theoretic lower and upper bound for the test design prevailing in the literature and analyse two prominent non-adaptive algorithms. In the second paper

Optimal group testing

by Amin Coja-Oghlan, Oliver Gebhard, Max Hahn-Klimroth and Philipp Loick, we derive an information-theoretic lower bound for any non-adaptive test design. Moreover, we present a novel test design inspired by the idea of spatial coupling and an efficient non-adaptive algorithm that attains this information-theoretic lower bound. As a corollary, we find a two-stage algorithm that attains the universal information-theoretic lower bound. Thereby, we show that group testing undergoes a plain impossible-easy transition. Finally, in the paper

Improved bounds for noisy group testing with constant tests per item

by Oliver Gebhard, Oliver Johnson, Philipp Loick and Maurice Rolvien, we consider the problem of binary group testing with noisy measurements and adapt popular non-adaptive algorithms from the noiseless setting to derive algorithmic upper bounds. Finally, we consider the practical implications of group testing in the empirical work

Efficient and accurate group testing via Belief Propagation: an empirical study

by Amin Coja-Oghlan, Max Hahn-Klimroth, Philipp Loick and Manuel Penschuck, where we run the belief propagation algorithm on moderately small instances for the noiseless and noisy setting. Thereafter, we will discuss the problem of quantitative group testing based on the paper

Quantitative group testing in the sublinear regime

by Oliver Gebhard, Max Hahn-Klimroth, Dominik Kaaser and Philipp Loick in which we derive an information-theoretic upper bound matching the known lower bound and analyse a greedy algorithm inspired by the first stage of belief propagation. Finally, we will deal with the Ising antiferromagnet in the paper

The Ising antiferromagnet and max cut on random regular graphs

by Amin Coja-Oghlan, Philipp Loick, Balazs Mezei and Gregory Sorkin. First, we pinpoint the replica symmetry breaking phase transition of the Ising antiferromagnet at the combinatorially meaningful Kesten-Stigum bound. Second, we use the interpolation method to establish an upper bound on the maximum cut size in a random regular graph, thereby vindicating a prediction from statistical physics. In the follow-up work

The Ising antiferromagnet in the replica symmetric phase

by Christian Fabian and Philipp Loick we characterize the limiting distribution of the partition function in the replica symmetric regime. The proof is based on a combination of the method of moments, spatial mixing arguments and small subgraph conditioning.

3. Binary group testing

3.1. **Setting & notation.** Binary group testing is a prime example of a statistical inference problem. Suppose we have a large population of n individuals, out of which k suffer from some rare disease with k being small in comparison to n. We employ the common assumption in the literature that $k \sim n^\theta$ for some $\theta \in (0,1)$. Rather than testing every individual separately, we have access to a pooled binary test scheme that can test groups of individuals. A test result is positive if and only if there is at least one infected individual included in the tested group and negative otherwise. The key insight is that the sparsity of the problem allows us to get by with significantly fewer tests than individual testing. To see why, consider the original two-stage test scheme proposed by Dorfman in 1943 to test soldiers in the US army for Syphilis [30]. In pioneering work, Dorfman proposed to test groups of individuals with a potential follow-up round of individual tests. To be precise, if a pool test returned positive, the individuals would be tested individually and the original pool test would have been wasted. However, if a test returned negative, we can be sure that every individual in the group is healthy and thus save a considerable number of individual tests. While a good starting point, this design is clearly not optimal. A considerable body of mathematical research has been devoted to group testing since then. Particularly, since the early 2000s the group testing problem has regained attention and today is used for DNA sequencing [51, 61], protein interaction experiments [59, 71] or the current COVID-19 pandemic [68].

Binary group testing can be readily described as an inference problem on random factor graphs. On the one side, we have a set $V_n = \{x_1, \dots, x_n\}$ of variable nodes representing the individuals and a set of $F_m = \{a_1, \dots, a_m\}$ factor nodes for the pooled tests. The edges between variable and factor nodes in the graph $\mathbb{G}$ encode the assignment of individuals to tests. In the following, let $(\Delta_x)_{x \in V_n}$ denote the variable degrees, i.e. the number of tests that an individual is assigned to. Similarly, let $(\Gamma_a)_{a \in \Gamma_m}$ be the sequence of test degrees. In line with standard graph notation, we write ∂x and ∂a to denote the neighbouring tests and individuals for an individual x and a test a, respectively. Note that in Dorfman's original two-stage test scheme, we have $\Delta_x = 1$ for all $x \in V_n$ in the first stage. In due course, we will see alternative test designs with different values of Δ_x. To proceed in the teacher-student analogy, we assume that a ground truth $\boldsymbol{\sigma} \in \{0,1\}^{V_n}$ is sampled uniformly at random among all configurations with Hamming weight k. Note that the construction of the factor graph $\mathbb{G}$ is independent of $\boldsymbol{\sigma}$. One peculiarity of the group testing problem is that we have control over the construction of $\mathbb{G}$ which we will exercise deliberately to best facilitate inference of the ground truth. Given

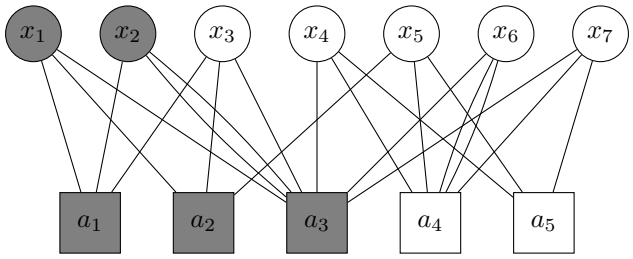

FIGURE 1. Illustration of a random regular test design for non-adaptive group testing. The figure is adopted from [17].

the ground truth $\boldsymbol{\sigma}$ and the graph $\mathbb{G}$, we can calculate the (noiseless) test result $\hat{\boldsymbol{\sigma}}_a$ for any test $a \in F_m$ as

$$\hat{\boldsymbol{\sigma}}_a = \hat{\boldsymbol{\sigma}}_a(\mathbb{G}, \boldsymbol{\sigma}) = \mathbf{1}\left\{\sum_{x \in \partial a} \boldsymbol{\sigma}_x > 0\right\}.$$

Now the teacher hands the graph $\mathbb{G}$ together with the test results $\hat{\boldsymbol{\sigma}} = (\hat{\boldsymbol{\sigma}}_a)_{a \in F_m}$ to the students whose task it is to recover $\boldsymbol{\sigma}$ with high probability. As for most inference problems, this inference task comes with two guiding questions. First, what is the minimum number of tests $m_{\mathrm{inf}}(n,\theta)$ regardless of computational resources to infer $\boldsymbol{\sigma}$ w.h.p. from $\mathbb{G}$ and $\hat{\boldsymbol{\sigma}}$? Second, what is the minimum number of tests $m_{\mathrm{alg}}(n,\theta)$ such that a polynomial-time algorithm can infer $\boldsymbol{\sigma}$ w.h.p.?

There are two settings for binary group testing that are worth mentioning at this point. On the one hand, we might consider *adaptive* group testing settings where testing is performed over several stages and the design of later stages might depend on the results of earlier stages. Dorfman's original two-stage design is a case in point where the decision for individual testing in the second stage depends on the pooled test result of the first stage. On the other hand, there are *non-adaptive* group testing settings where all tests must be specified upfront and inference of $\boldsymbol{\sigma}$ is only admissible based on these tests. A classical non-adaptive test design prevalent in the literature would be a random regular factor graph where we fix either the variable degrees or both the variable and factor degrees. An illustration of the latter approach is provided in Figure 1 While both approaches are important, research in recent years has predominantly focused on non-adaptive test designs in the interest of speed and the possibility for automation.

3.2. **Prior research.** Let us review the state of research prior to our work on binary group testing and start with a simple counting argument that reveals the universal information-theoretic lower bound for any group testing design - adaptive or non-adaptive. A necessary condition for us to infer $\boldsymbol{\sigma}$ is that the number of possible test outcomes 2^m must exceed the number of possible configuration $\binom{n}{k}$. A straightforward application of Stirling's formula thus yields

$$m_{\mathrm{ad}} = \frac{1}{\log 2} k \log(n/k) \qquad \left(k \sim n^{\theta}\right) \tag{3.1}$$

For the random regular factor graph model described above as an example of a *non-adaptive* setting, [45] derived the non-adaptive information-theoretic lower bound

$$m_{\mathrm{inf}} = \max\left\{\frac{1}{\log 2}, \frac{\theta}{(1-\theta)\log^2 2}\right\} k \log(n/k)$$

When it comes to efficient non-adaptive algorithms, the so-called `COMP` and `DD` algorithm have emerged as plain, yet prominent algorithms. Despite their simplicity, they perform

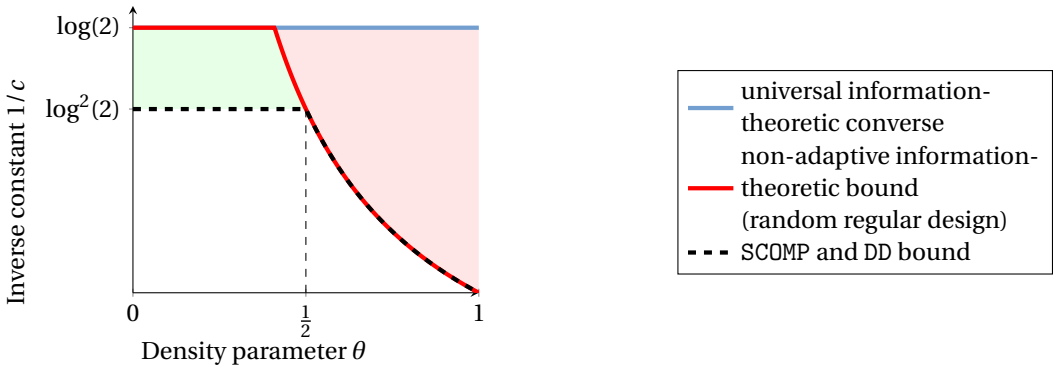

FIGURE 2. Phase diagram for binary group testing prior to our work *Optimal group testing*. In the red area, inference is information-theoretically impossible for a random regular design. It remained open whether there exists a different non-adaptive design that could attain the universal information-theoretic lower bound. The green area marks the regime for which inference is information-theoretically possible but where all previously known algorithms failed.

reasonably well. All that the COMP algorithm does is to classify any individual included in at least one negative test as healthy. All other individuals are classified as infected. Clearly, this algorithm only succeeds if every healthy individual is included in at least one negative test. The DD algorithm moves one step beyond COMP. As the COMP algorithm, it first classifies any individuals in negative tests as healthy and removes them from the test design. Thereafter, it searches for positive tests only featuring one remaining individual (after the removal of definitely healthy individuals) and classifies those individuals as infected. All other individuals are deemed to be healthy. While the COMP algorithm produces an estimate of $\boldsymbol{\sigma}$ that is at least consistent with the test results, it is important to note that the DD algorithm might yield an estimate that is not even consistent with the test results. Clearly, the algorithmic upper bound of DD algorithm is at least as good as the bound of COMP and reads as follows

$$m_{\mathrm{alg}} = \max\left\{\frac{1}{\log^2 2}, \frac{\theta}{(1-\theta)\log^2 2}\right\} k\log(n/k).$$

For adaptive designs, [65] proposed an adaptive three-stage algorithm attaining the universal information-theoretic lower-bound m_{ad}. These bounds are visualised in the phase diagram in Figure 2.

Significant research effort was devoted to resolving the open questions (see i.e. [10, 56, 58, 7]) in the phase diagram which can be summarized as follows

- Is the algorithmic threshold m_{alg} tight for the DD algorithm? In other words, does DD fail to succeed once we have $m < m_{\mathrm{alg}}$?
- Is the non-adaptive information-theoretic lower bound m_{inf} information-theoretically achievable by a random regular design?
- Is the non-adaptive information-theoretic lower bound m_{inf} achievable by an efficient algorithm?
- Are there alternative non-adaptive test designs that attain the universal information-theoretic lower bound m_{ad}?
- If no such non-adaptive test designs exist, does there exist an efficient two-stage (rather than three-stage) algorithm attaining the information-theoretic lower bound?

In the first two papers on binary group testing, we will address the above questions and entirely resolve the phase diagram in Figure 2.

3.3. **Results.** The information-theoretic and algorithmic results in the preceding section evince that for all relevant questions the number of tests is of order $\Theta(k\log(n/k))$, so we will write $m = ck\log(n/k)$ for some constant $c > 0$. Throughout the papers, we employ a model where each individual is assigned to $\Delta = cd\log(n/k)$ tests uniformly at random without replacement [2] for some constant $d > 0$. This design gives rise to fluctuating test degrees $(\Gamma_a)_{a\in F_m}$. However, we readily find that the test degrees are tightly concentrated in the sense that w.h.p.

$$d n/k - \sqrt{n}\log(n/k)\log n \le \min_{a\in F_m}\Gamma_a \le \max_{a\in F_m}\Gamma_a \le d n/k + \sqrt{n}\log(n/k)\log n.$$

Moreover, some standard techniques reveal that w.h.p. the number of negative tests $\boldsymbol{m}_0$ is concentrated at

$$\boldsymbol{m}_0 = \exp(-d)m + O\left(\sqrt{m}\log^2 m\right). \tag{3.2}$$

Thus, in order to maximize the test entropy, it seems suitable to set $d = \log(2)$. Indeed, in [17] we show that $c, d = \Theta(1)$ and this specific choice of d best facilitate inference in our model. Before we get to the results, we should introduce some types of individuals that are of fundamental importance for the analysis of algorithms in subsequent sections. To be precise, we split the set of healthy individuals V_0 into two subsets V_0^+ and V_0^-. An individual is defined to be in V_0^-, if it shows up in at least one negative test. Formally,

$$V_0^- = \{x \in V_0 : \exists a \in \partial x : \hat{\sigma}_a = 0\} \qquad \text{and} \qquad V_0^+ = V_0 \setminus V_0^-.$$

Similarly, we divide the set of infected individuals V_1 into three subsets V_1^{--}, V_1^+ and a remaining set $V_1 \setminus (V_1^{--} \cup V_1^+)$. The set V_1^{--} contains those infected individuals that are included in at least one test where all other individuals are from the set V_0^-. Conversely, the set V_1^+ contains all individuals that are only included in tests with at least one other infected individual. Formally,

$$V_1^{--} = \left\{x \in V_1 : \exists a \in \partial x : \forall y \in \partial a \setminus x : y \in V_0^-\right\} \quad \text{and} \quad V_1^+ = \left\{x \in V_1 : \forall a \in \partial x : \exists y \in \partial a : y \in V_1\right\}.$$

To throw the bridge to the DD algorithm, note that the first step of DD identifies all individuals from V_0^-, i.e. the definitely healthy individuals. Correspondingly, the second step of DD classifies the individuals from V_1^{--} as infected, i.e. the definitely infected individuals. We can now state our first result regarding the information-theoretic threshold of the random regular design.

Theorem 3.1 (Theorem 1.1 in [17])**.** *Suppose that* $0 < \theta < 1$, $k \sim n^\theta$ *and* $\varepsilon > 0$ *and let*

$$m_{inf} = m_{inf}(n,\theta) = \max\left\{\frac{1}{\log 2}, \frac{\theta}{(1-\theta)\log^2 2}\right\} k\log(n/k).$$

(1) If $m > (1+\varepsilon)m_{inf}$*, then there exists an algorithm that given* $\mathbb{G}, \hat{\boldsymbol{\sigma}}$ *outputs* $\boldsymbol{\sigma}$ *with high probability.*

(2) If $m < (1+\varepsilon)m_{inf}$*, then there does not exist any algorithm that given* $\mathbb{G}, \hat{\boldsymbol{\sigma}}, k$ *outputs* $\boldsymbol{\sigma}$ *with a non-vanishing probability.*

[2] Note that in the paper "Information-theoretic and algorithmic thresholds for group testing" we assume that individuals are assigned to tests uniformly at random *with* replacement. Despite this small difference in the replacement rule, the models are almost identical and all results carry over to the model where we sample without replacement.

We stress here that the terminology from the theorem refers to *any* algorithm, not necessarily an efficient one. Let us shed light on the proof strategy of this information-theoretic result, starting with the first statement. To this end, we will employ a technique presented earlier. It turns out that in order to derive the statement of the theorem it suffices to calculate the planted first moment, i.e. count the number of alternative configurations other than the ground truth $\boldsymbol{\sigma}$ that yield the same test result. If we find that for a certain value of m, there does not exist another configuration consistent with the test result, we know that it is information-theoretically possible to recover $\boldsymbol{\sigma}$ by using a brute-force approach. Conversely, when there are other configurations next to the ground truth yielding the test results and featuring k infected individuals, the Nishimori property informs us that we have no way to distinguish the true configuration from the alternatives. Our best guess is a random pick. The proof of the first statement of the theorem proceeds in two steps. First, we show that for $m < m_{\mathrm{inf}}$ there does not exist a second configuration consistent with the test results w.h.p. that has a small overlap with the true configuration. Second, we consider configurations that have a large overlap with the true configuration and rule out any other consistent configuration for $m < m_{\mathrm{inf}}$.

For the first step, let

$$S_k(\mathbb{G}, \hat{\boldsymbol{\sigma}}) = \left\{\sigma \in \{0,1\}^{V_n} \setminus \boldsymbol{\sigma} : \forall a \in F_m : \hat{\sigma}_a(\mathbb{G}, \sigma) = \hat{\boldsymbol{\sigma}}_a\right\} \tag{3.3}$$

be the the set of all alternative configurations σ that are consistent with the test results. Along those lines, we let

$$Z_{k,\ell} = Z_{k,\ell}(\mathbb{G}, \hat{\boldsymbol{\sigma}}) = |\{\sigma \in S_k(\mathbb{G}, \hat{\boldsymbol{\sigma}}) : \langle \boldsymbol{\sigma}, \sigma \rangle = \ell\}| \tag{3.4}$$

be the number of consistent configurations that have an overlap of ℓ in the number of infected individuals with $\boldsymbol{\sigma}$. In the following, we will call ℓ simply the overlap between $\boldsymbol{\sigma}$ and σ. The notation of Z is no coincidence here, since the partition function of the planted model for the group testing problem indeed simply counts the number of alternative configurations consistent with the test results. With Γ being the σ-algebra generated by the random test degrees $(\Gamma_a)_{a \in F_m}$ we find

$$\mathbb{E}\left[Z_{k,\ell}(\mathbb{G}, \hat{\boldsymbol{\sigma}}) \mid \Gamma\right] \leq O\left((\Delta k)^{3/2}\right) \cdot \binom{k}{\ell}\binom{n-k}{k-\ell} \prod_{i=1}^{m} \left(1 - 2\left(1 - k/n\right)^{\Gamma_i} + 2\left(1 - 2k/n + \ell/n\right)^{\Gamma_i}\right)$$

The proof of this expression requires a careful application of the balls-into-bins principle due to the regularities of the underlying graph model. Here, let us briefly explain the combinatorial meaning. The two binomial coefficients count the number of all possible configurations that have overlap ℓ with the true configuration. The latter expression describes the probability that such a configuration σ yields the same test results as $\boldsymbol{\sigma}$. Thus, for each test it calculates the joint probability of both tests being negative and both being positive. Note that with the definition of ℓ, we have ℓ individuals that are infected under $\boldsymbol{\sigma}$ and σ, $k - \ell$ individuals that are infected under $\boldsymbol{\sigma}$, but healthy under σ and vice versa and $n - 2k + \ell$ individuals that are healthy under both configurations. Neglecting the intricate regularities of the model, the joint probability of a negative test under both $\boldsymbol{\sigma}$ and σ is thus given by $(1 - 2k/n + \ell/n)^{\Gamma_i}$. The joint probability of a positive test follows from the inclusion-exclusion principle. To be precise, we will *not* have a positive test under $\boldsymbol{\sigma}$ and σ if there are only individuals that are healthy under $\boldsymbol{\sigma}$ and infected under σ or healthy under both $\boldsymbol{\sigma}$ and σ. The same goes for individuals that are infected under $\boldsymbol{\sigma}$ and healthy under σ. This insight yields a joint probability for a positive test under $\boldsymbol{\sigma}$ and σ as $1 -$

$2(1-k/n)^{\Gamma_i} + (1-2k/n+\ell/n)^{\Gamma_i}$. Putting these findings together yields the above expression. Standard simplifications using Stirling's formula then show that $Z_{k,\ell}(\mathbb{G},\boldsymbol{\sigma}) = 0$ w.h.p. for all $\ell < (1-1/\log n)k$ as long as $c > 1/\log 2$.

We proceed with large overlaps. As before, we will calculate $Z_{k,\ell}(\mathbb{G},\hat{\boldsymbol{\sigma}})$. However, this time we will consider the local changes introduced by flipping infected individuals under $\boldsymbol{\sigma}$ to healthy under σ. A crucial insight is that in the model laid out above w.h.p. every infected individual shows up in at least $\delta\Delta$ of its tests as the only infected individual for some constant $\delta > 0$ as long as $c > \theta/((1-\theta)\log^2 2)$. Let us denote this event by $\mathscr{R}$ and recall $\boldsymbol{m}_0$ from (3.2). We readily find

$$\mathbb{E}\left[Z_{k,\ell}(\mathbb{G},\hat{\boldsymbol{\sigma}}) \mid \Gamma, \mathscr{R}, \boldsymbol{m}_0\right] \leq O\left((\Delta k)^{3/2}\right)\binom{k}{\ell}\binom{n-k}{k-\ell} \cdot\left(1-\left(1-\frac{k-\ell}{n-k}\right)^{\max_a \Gamma_a}\right)^{\delta\Delta(k-\ell)}\left(\frac{n-2k+\ell}{n-k}\right)^{(1+n^{-\Omega(1)})\boldsymbol{m}_0 \min_a \Gamma_a}.$$

Again, it turns out that $\mathbb{E}\left[Z_{k,\ell}(\mathbb{G},\hat{\boldsymbol{\sigma}}) \mid \Gamma, \mathscr{R}, \boldsymbol{m}_0\right] = 0$ w.h.p. for $(1-1/\log n)k \leq \ell < k$ for any constant $c > 0$. In combination with the bound for the event $\mathscr{R}$ and the bound for ruling out small-overlap configurations yields the first statement of Theorem 3.1.

The careful reader will notice that the first statement of Theorem 3.1 does not make any reference to knowing the specific value of k. Indeed, we show that if there does not exist a second configuration of Hamming weight k consistent with the test results there will also not be a second configuration of Hamming weight less than k that yields the same test results as $\hat{\boldsymbol{\sigma}}$ w.h.p. Thus, finding $\boldsymbol{\sigma}$ for $m > m_{\text{inf}}$ boils down to finding the configuration with the lowest Hamming weight that is still consistent with the test results.

The second statement follows from two separate arguments. First, the universal counting bound ensures $m > m_{\text{ad}}$ also for the regular non-adaptive test design. The crucial next step relates to the individual types we described above. To be precise, we find that for $c < \theta/((1-\theta)\log^2 2)$, we have $V_0^+, V_1^+ = n^{\Omega(1)}$ with high probability. This result entails that we can easily construct many configurations satisfying the test results by simply flipping healthy individuals in V_0^+ under $\boldsymbol{\sigma}$ to infected and flipping an equal number of infected individuals in V_1^+ to healthy. Such a configuration clearly satisfies the test results and the Nishimori property guarantees that there is no hidden information on which configuration is the correct one. Since we can easily construct $n^{\Omega(1)}$ many alternative configurations that yield the same test results as $\boldsymbol{\sigma}$, we do not have any means to detect the true configuration $\boldsymbol{\sigma}$ with non-vanishing probability.

Having established an information-theoretic phase transition for the regular test design, we are interested in how close efficient algorithms can get to this bound. In earlier work, [45] already showed that the DD algorithm will succeed to recover $\boldsymbol{\sigma}$ if $m > m_{\text{alg}}$. But a converse result was still missing. At the same time, [5] had proposed an extension of the DD algorithm called SCOMP that was widely believed to improve upon the algorithmic bound of DD. SCOMP performs the same first two steps as DD. Being left with the sets V_0^+ and $V_1 \setminus V_1^{--}$, a greedy vertex cover algorithm is employed by iteratively classifying the remaining individual with the largest degree (breaking ties uniformly at random) as infected, removing it and all adjacent tests from the design. The following result establishes that the algorithmic upper bound for DD is tight and that, surprisingly, SCOMP fails at exactly the same bound.

Theorem 3.2 (Theorem 1.2 in [17])**.** *Suppose that* $0 < \theta < 1$ *and* $\varepsilon > 0$ *and let*

$$m_{alg} = m_{alg}(n,\theta) = \max\left\{\frac{1}{\log^2 2}, \frac{\theta}{(1-\theta)\log^2 2}\right\} k\log(n/k).$$

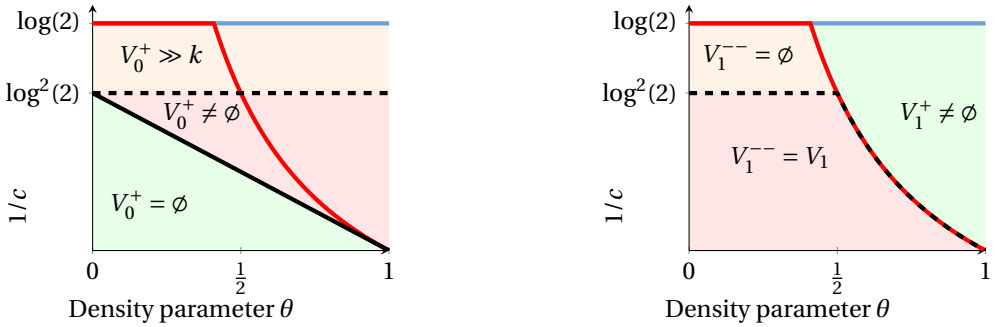

FIGURE 3. Size of different types of healthy (left figure) and infected individuals (right figure) with respect to the number of tests

If $m < (1-\varepsilon)m_{alg}(n,\theta)$*, then given* $\mathbb{G}, \hat{\boldsymbol{\sigma}}$ *w.h.p. both* SCOMP *and* DD *fail to output* $\boldsymbol{\sigma}$.

To discuss the proof idea of Theorem 3.2, we will focus on the SCOMP algorithm since if SCOMP fails in some regime for m, so does DD. The key stepping stone towards the proof of Theorem 3.2 is to determine the size of individual types V_0^+ and V_1^{--}. W.h.p. we have

$$\left|V_0^+\right| = \left(1 + n^{-\Omega(1)}\right) n2^{-\Delta}. \tag{3.5}$$

Moreover, it turns out that for $m < m_{\text{alg}}$ we have w.h.p.

$$V_1^{--} = \emptyset.$$

Even though not necessary for the proof of the theorem, let us visualise the size of the sets V_0^+, V_1^{--} and V_1^+ for various values of c and θ in Figure 3. It becomes immediately clear from this figure why DD succeeds for $m > m_{\text{alg}}$ since all infected individuals will show up in at least one test with only individuals from V_0^- and thus are easily identified as definitely infected by the second step of DD. Let us now consider the regime $m < m_{\text{alg}}$. In this case, the second step of DD which underlies SCOMP proceeds without identifying any infected individuals as definitively infected, since no such individual shows up in at least one test with only individuals from V_0^-. Thus, DD terminates with leaving the sets V_0^+ and V_1 unclassified. Now, SCOMP selects the individual out of these sets with the largest degree. However, by the former argument, no infected individual was classified and thus, all positive tests are still intact. Therefore, all individuals both from V_0^+ and V_1 will have degree Δ. Furthermore, (3.5) informs us that for $c < 1/\log^2 2$, we have $V_0^+ = kn^{\Omega(1)}$ with high probability. So breaking ties arbitrarily, SCOMP will select a healthy individual already in the first step of the greedy vertex cover algorithm w.h.p. and thus be unable to recover $\boldsymbol{\sigma}$. With some minor technical additions, the theorem follows from this reasoning.

Theorems 3.1 and 3.2 leave three interesting questions open that we will tackle in the following work. First, does there exist a non-adaptive test design that is superior to the regular model we employed and that allows inference of $\boldsymbol{\sigma}$ for $m_{\text{ad}} < m < m_{\text{inf}}$? Second, does there exist an efficient algorithm (and potentially an alternative test design) that also succeeds for $m_{\text{inf}} < m < m_{\text{alg}}$? Put differently, is the regime $m_{\text{inf}} < m < m_{\text{alg}}$ computationally hard? Third, does there exist an efficient two-stage algorithm that succeeds for $m > m_{\text{ad}}$?

Our next result sheds light on the first question showing that there exists an adaptivity gap in binary group testing.

Theorem 3.3 (Theorem 1.1 in [21])**.** *For any* $0 < \theta < 1, \varepsilon > 0$ *there exists* $n_0 = n_0(\theta,\varepsilon)$ *such that for all* $n > n_0$*, all test designs* G *with* $m \leq (1-\varepsilon)m_{inf}$ *tests and for every function* $\mathscr{A}_G : \{0,1\}^m \to$

$\{0,1\}^n$ *we have*

$$\mathbb{P}\left[\mathscr{A}_G(\hat{\boldsymbol{\sigma}}) = \boldsymbol{\sigma}\right] < \varepsilon.$$

The function $\mathscr{A}_G$ can be understood as any algorithm trying to infer $\boldsymbol{\sigma}$ from $\hat{\boldsymbol{\sigma}}$ and G with the algorithm not necessarily being efficient. What Theorem 3.3 tells us is that for $m < m_{\text{inf}}$ no non-adaptive test design and algorithm exist that can recover $\boldsymbol{\sigma}$ w.h.p. The proof of Theorem 3.3 is technically quite involved, so let us focus on the highlights here. The proof proceeds in two steps. First, we show that for $\theta = 1-\delta$ for some small $\delta > 0$ and $m < (1-\varepsilon)m_{\text{inf}}$ and any test design, we have

$$V_0^+, V_1^+ = n^{\Omega(1)}.$$

This result entails that we can easily construct $n^{\Omega(1)}$ alternative configurations consistent with the test results. By the Nishimori property, we know that there is no hidden information in the ground truth $\boldsymbol{\sigma}$ that would make it distinguishable from these alternative configurations. So our best pick is a uniform sample from all satisfying configurations leading to an error w.h.p. Second, we show that if we cannot solve the group testing problem for $\theta = 1-\delta$, we cannot solve it for smaller θ. In conjunction with the universal information-theoretic lower bound m_{ad}, we will thus yield the theorem. Let us look at each of the two steps in more detail.

For the first step, we aim to characterise individuals in V_0^+ and V_1^+ since the joint presence of such individuals allows us to construct alternative configurations that are indistinguishable from the ground truth. To this end, the concept of a disguised individual is central. An individual is said to be disguised if it features at least one (other) infected individual in all of its neighbouring tests. Clearly, the challenge in proving a general statement as Theorem 3.3 is to evince for *every* possible graph that there exists a large number of disguised infected and healthy individuals. To do so, we start off with two simple tricks. First, in order to avoid stochastic dependencies, we employ a slightly different approach to constructing the ground truth, i.e. we let each individual be infected with a certain probability independently of each other. Our choice of this probability ensures that if we have a large number of disguised individuals under this ground truth, it will also be true for our original model. Second, we rule out any tests of size $n^{1-\theta}\log n$ since all such tests are positive w.h.p. so it is pointless to carry them out in the first place. An important observation at this point is that almost all individuals in such a graph have a relatively moderate degree.

In what follows, we construct a sequence of graphs where we iteratively remove the individual with the largest probability of being disguised as well as its neighbourhood up to depth four to establish independence between the probability of being disguised between these iteratively removed individuals. Now, the key is to get a handle on the probability of being disguised. Observe that an individual is only disguised if it is disguised in every test it is assigned to. With $\mathscr{D}(x)$ denoting the event that x is disguised and $\mathscr{D}(x,a)$ denoting the event that x is disguised in test a, we have

$$\mathbb{P}\left[\mathscr{D}(x)\right] = \mathbb{P}\left[\cap_{a\in\partial x}\mathscr{D}(x,a)\right].$$

Since $\mathscr{D}(x,a)$ is increasing with respect to $\boldsymbol{\sigma}$ (and also the ground truth where individuals infected independently of each other), we can employ the FKG inequality to find

$$\mathbb{P}\left[\mathscr{D}(x)\right] \geq \prod_{a\in\partial x}\mathbb{P}\left[\mathscr{D}(x,a)\right].$$

With some calculations, we find that this lower bound can be readily evaluated by finding the minimum of the function

$$z \in (0,\infty) \to z\log\left(1-(1-p)^{z-1}\right)$$

where p denotes the prior probability of being infected. Solving this optimization problem, we obtain a probability for an individual to be disguised. From this result, some calculations in which we carefully control the error terms suffice to show that for $\theta = 1-\delta$ for some small $\delta > 0$ and $m < (1-\varepsilon)m_{\inf}$ we have $V_0^+, V_1^+ = n^{\Omega(1)}$.

For the second step, we will extend this result from $\theta = 1-\delta$ to the entire range of θ. To this end, we start off with a test design featuring n individuals and $k = \lceil n^\theta \rceil$ for $\theta = 1-\delta$. Next, for any smaller $\theta' < \theta$ we increase the number of individuals to n' in such a way that

$$k = \lceil n^\theta \rceil = \lceil n'^{\theta'} \rceil.$$

It turns out that if a test design exists for $\theta' < \theta$ that succeeds for $m < m_{\inf}(n,\theta)$, there also exists a test design for density θ' that outperforms $m_{\inf}(n,\theta)$. Since we have demonstrated that no such test design exists for $\theta = 1-\delta$ for δ arbitrarily close to 0, it follows that it also does not exist for $m < m_{\inf}(n,\theta')$ for any $\theta' < \theta$. The theorem follows from these two steps.

Having established the information-theoretic lower bound for non-adaptive designs, the interesting remaining question is whether there exists a test design and inference algorithm that jointly allow inference all the way down to this lower bound. The next result provides a positive answer to this question.

Theorem 3.4. *For any $0 < \theta < 1$, $\varepsilon > 0$, there is $n_0 = n_0(\theta,\varepsilon)$ such that for every $n > n_0$ there exist a randomised test design $\mathbb{G}$ comprising $m \le (1+\varepsilon)m_{inf}$ tests and a polynomial time algorithm* SPIV *that given $\mathbb{G}$ and the test results $\hat{\boldsymbol{\sigma}}$ outputs $\boldsymbol{\sigma}$ w.h.p.*

To understand the intuition behind the test design and the algorithm, let us go back one step and consider the simple DD algorithm in the regime $m_{\text{alg}} < m < m_{\inf}$ again. By Figure 3 we know that in this regime $V_1^{--}, V_1^+ = \emptyset$ and $V_0^+ = kn^{\Omega(1)}$ w.h.p. Thus, DD performs equally well as COMP by only classifying definitively healthy individuals from the set V_0^-, i.e. individuals that are included in at least one test. However, in this regime there does not exist any individual from V_1^{--} w.h.p. Thus, after performing the DD algorithm we are left with the sets V_1 and V_0^+. The key question is how we can set those two sets apart from each other. Fundamentally, there is clearly a difference between the two. To be precise, every individual from V_0^+ - we will call these individuals disguised healthy hereafter - must feature at least one infected individual in each test since otherwise this test would be negative. Conversely, the event $\mathscr{R}$ from above informs us that for $m < m_{\inf}$ every infected individual shows up in at least a constant fraction of its tests as the only infected individual. Thus, if we had the pleasant situation that we have already identified all but one individuals and need to determine whether this remaining individual is disguised healthy or infected, our job would be straightforward: simply consider each test and check whether it features at least one infected individual. If it does, the individual is disguised healthy; otherwise, it is infected. Unfortunately, we do not know the infection status of all other individuals or at least most other individuals that would facilitate this inference. Even worse, in the absence of any information on the ground truth (other than identifying definitively healthy individuals) we are not aware of any approach that could distinguish V_0^+ and V_1 in the plain random regular design. Thus, we appear to be in a situation that is predisposed for the application of spatial coupling that we have described in Section 2.6. What spatial coupling does is to iteratively identify individuals using information from previously classified individuals. It turns out that this head start towards the ground truth that we have at our disposal for previous compartments is precisely the right amount of information that allows us to infer $\boldsymbol{\sigma}$ w.h.p. for $m < m_{\inf}$.

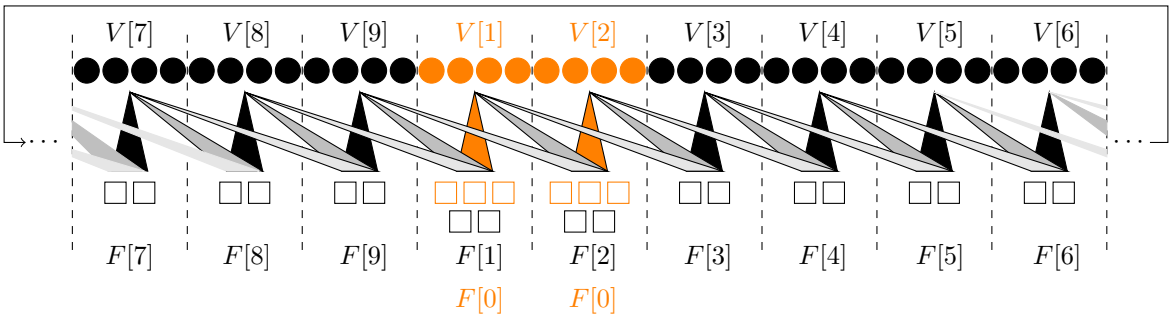

FIGURE 4. Visualization of spatially coupled test design. The figure is adopted from [21].

The crucial idea behind spatial coupling is to carefully exercise our discretion of selecting the test design and enforce a certain geometric structure upon the random regular design from before. Take two parameters s, ℓ such that $s \ll \ell \ll k$ [3]. We now divide the individuals into ℓ compartments which we arrange along a ring structure (see Figure 4). The same goes for the tests. We denote the set of individuals in compartment $0, 1, \dots, \ell$ by $V[0], V[1], \dots, V[\ell]$ and the set of tests accordingly by $F[0], F[1], \dots, F[\ell]$. Now, any individual $x \in V[i]$ joins Δ/s tests in each of its neighbouring compartments $F[i], F[i+1], \dots, F[i+s-1]$ uniformly at random without replacement. This assignment of individuals to tests gives rise to the test design in Figure 4. Finally, the first s compartments will act as our starting point of inference and thus be called *seed* compartments. For these compartments, we will conduct more tests such that a simple algorithm such as DD can infer the infection status of individuals in the seed compartments w.h.p. Let us denote this set of additional tests by $F[0]$.

So how does this test design help us to set the sets V_0^+ and V_1 apart? To this end, we developed a combinatorial algorithm called SPIV that proceeds in three steps. While the first and third are relatively straightforward, the key idea lies in the second step.

In the first step, we classify the individuals in the seed compartments. Since $m_{\text{alg}} = \Theta(m_{\text{inf}})$ and the number of individuals in the seed compartments is $sn/\ell = o(n)$, we immediately see that $o(m_{\text{inf}})$ tests suffice to infer the infection status of the individuals in the seed compartments w.h.p. using a simple algorithm such as DD.

With the seed compartments identified we move on to the next unidentified compartment $s+1$ on the right. We next aim to classify the individuals $V[s+1]$ that are included in test compartments $F[s+1], \dots, F[2s]$. Note that individual V_0^- can be readily found as healthy, so our focus is on the sets V_0^+ and V_1. Let us first consider test compartment $F[s+1]$. This compartment features a few individuals from $V[s+1]$, but the vast majority comes from the seed compartments we already classified. It should become clear now how spatial coupling can help us in distinguishing V_0^+ and V_1. To be precise, what we do is to consider the adjacent tests of an individual $x \in V[s+1]$ and count the number of *unexplained* tests. Clearly, every adjacent test is positive since otherwise, this individual would not be in V_0^+ or V_1. We call such a positive test unexplained if it does not contain an individual from previous compartments that we have identified as infected. There are two possibilities for such an unexplained test. Either the infected individual rendering this test positive has not been identified yet or the individual under consideration is itself infected. Thus, the number of unexplained tests in adjacent tests provides a way to distinguish V_0^+ and V_1 since individuals from V_0^+ should have meaningfully fewer unexplained tests than individuals from V_1.

[3] Note that the proof requires specific choices of s and ℓ, which we, however, neglect here for purposes of a high-level summary.

The key question is whether this approach suffices to discriminate between V_0^+ and V_1 all the way down to m_{inf}. If we count all unexplained tests equally, it does not. However, if we take into account that an unexplained test in an earlier test compartment such as $F[s+1]$ is a much stronger indication that we are dealing with an infected individual than an unexplained test in later test compartments where we have classified fewer individuals yet, we can close the gap to the information-theoretic lower bound by weighting the tests from different compartments using

$$w_j = \log \frac{(1-\delta)2^{j/s-1}(2-2^{j/s})}{(1-(1-\delta)2^{j/s-1})(2^{j/s}-1} \qquad \big(j \in [s]\big)$$

where $\delta = 2/s^2$. With $V_1[i] = V_1 \cap V[i]$ and $V_{0+}[i] = V_0^+ \cap V[i]$ and $\boldsymbol{W}_x$ denoting the weighted sum of unexplained tests for an individual x we find

$$\sum_{s\le i<\ell} \sum_{x\in V_1[i]} \mathbf{1}\left\{\boldsymbol{W}_x < (1-\delta/4)\frac{\Delta}{s}\sum_{j=1}^{s-1} 2^{j/s-1} w_j\right\} \le k\exp\left(-\frac{\Omega(\log n)}{(\log\log n)^4}\right) \tag{3.6}$$

$$\sum_{s\le i<\ell} \sum_{x\in V_{0+}[i]} \mathbf{1}\left\{\boldsymbol{W}_x > (1-\delta)\frac{\Delta}{s}\sum_{j=1}^{s-1} 2^{j/s-1} w_j\right\} \le k^{1-\Omega(1)} \tag{3.7}$$

whenever $(1+\varepsilon)m_{\text{ad}} \le m = O\big(n^\theta \log n\big)$. The proof of (3.6) is based on a careful large deviation analysis using Lagrangian optimization and the Chernoff bound for hypergeometric distributions. By (3.6) we find that we will make at most $kn^{-\Omega(1)}$ mistakes in each compartment when classifying individuals according to the weighted neighbourhood sum using a suitable threshold. Fortunately, it turns out that (3.6) not only holds when the individuals in preceding compartments have been perfectly classified such as for the first non-seed compartment, but even if the misclassifications in previous compartments are of order $kn^{-\Omega(1)}$. Thus, after iteratively thresholding the weighted sum of unexplained tests we are left with an estimate of the ground truth $\boldsymbol{\sigma}$ where we have made at most $kn^{-\Omega(1)}$ mistakes. With this estimate, we can move to a third clean-up step.

To tidy up things and eventually get to an estimate of the ground truth that is correct w.h.p. we perform a clean-up step in which we iteratively loop through the ring, count the number of unexplained tests in the neighbourhood of an individual and (re-)classify an individual as healthy if fewer than $\log^{1/4} n$ tests contain an infected individual and are thus unexplained. It turns out that as long as

$$m > (1+\varepsilon)\frac{\theta}{(1-\theta)\log^2 2} k\log(n/k) \tag{3.8}$$

$\log n$ rounds suffice until we have recovered the ground truth $\boldsymbol{\sigma}$ w.h.p. Combining (3.6) and (3.8) yields the theorem and thus evinces that the SPIV algorithm applied on a spatially coupled test design attains the information-theoretic lower bound for non-adaptive group testing.

In terms of our initial questions, these results leave us with one open end, namely the question of whether there exists an efficient two-stage algorithm that attains the universal information-theoretic lower bound m_{ad}. It turns out that there is a natural adaptation of the SPIV algorithm to adaptive settings that gives us the desired result.

Theorem 3.5 (Theorem 1.3 in [21])**.** *For any $0<\theta<1$, $\varepsilon>0$ there is $n_0 = n_0(\theta,\varepsilon)$ such that for every $n > n_0$ there exist a two-stage test design with no more than $(1+\varepsilon)m_{ad}$ tests in total and a polynomial time inference algorithm that outputs $\boldsymbol{\sigma}$ with high probability.*

This algorithm which we will denote by ASPIV is based on the crucial observation that $(1+\varepsilon)m_{\text{ad}}$ tests suffice such that we misclassify $kn^{-\Omega(1)}$ in the first stage using a spatially

coupled test design and the first two steps of SPIV - see (3.6) for reference. However, in contrast to SPIV we do not perform a clean-up step but enter a second group testing stage. In this stage, all individuals that were deemed infected in the first stage are tested individually requiring $k + kn^{-\Omega(1)} = o(m_{\text{ad}})$ tests. Moreover, we construct a simple random regular test design for the individuals deemed healthy in the first stage and apply the DD algorithm. Since we have $kn^{-\Omega(1)}$ infected individuals in this test design we require $\Theta(kn^{-\Omega}\log(n/k)) = o(m_{\text{ad}})$ tests to recover the correct infection status for all such individuals w.h.p. Combining these findings, it is clear that the two-stage ASPIV algorithm can recover $\boldsymbol{\sigma}$ with $(1+\varepsilon)m_{\text{ad}}$ tests for an arbitrarily small constant $\varepsilon > 0$.

These results completely resolve open questions that have been studied for almost two decades. In particular, we demonstrated that group testing undergoes a plain impossible-easy transition with no computationally hard regime and that there indeed exists an adaptivity gap for certain values of θ that disappears as soon as we allow two stages of group testing. With the phase diagram for noiseless sublinear group testing completely resolved, we next relax the assumption of perfect test results and consider a general $p-q$-noise model where every truly negative test is flipped to positive with probability p and every truly positive test is flipped to negative with probability q. We will consider a noisy variant of the COMP and DD algorithm.

The following result will first exhibit performance guarantees for a noisy version of the COMP algorithm. In the noiseless setting, a single negative test sufficed to definitely declare every member of such a test as healthy. For the noisy setting, things are not as straightforward since a negative test might also result from a truly positive test that was flipped to negative. Therefore, we will define a threshold $\alpha\Delta$ for the number of displayed negative tests such that we classify an individual as healthy if it exhibits more than $\alpha\Delta$ displayed negative tests in its neighbourhood and infected otherwise. In the following, let $H(\cdot)$ and $D_{\text{KL}}(\cdot\|\cdot)$ denote the entropy and Kullback-Leibler divergence, respectively. We find the following performance guarantees for this noisy version of COMP.

Theorem 3.6 (Theorem 2.1 in [39])**.** *Let $p,q \geq 0$, $p+q<1, d \in (0,\infty), \alpha \in (q, e^{-d}(1-p) + \left(1-e^{-d}\right)q)$. Suppose that $0<\theta<1$ and let*

$$m_{COMP} = m_{COMP}(n,\theta,p,q) = \min_{\alpha,d} \max\{b_1(\alpha,d), b_2(\alpha,d)\}\, k\log(n/k)$$

$$\text{where} \qquad b_1(\alpha,d) = \frac{\theta}{1-\theta}\frac{1}{dD_{\text{KL}}\left(\alpha\|q\right)}$$

$$\text{and} \qquad b_2(\alpha,d) = \frac{1}{1-\theta}\frac{1}{dD_{\text{KL}}\left(\alpha\|e^{-d}(1-p)+\left(1-e^{-d}\right)q\right)}$$

If $m \geq (1+\varepsilon)m_{COMP}$ for some $\varepsilon > 0$, noisy COMP *will recover $\boldsymbol{\sigma}$ w.h.p. given test design $\mathbb{G}$ and test results $\hat{\boldsymbol{\sigma}}$.*

Similarly, the noisy DD algorithm will in addition only classify an individual as infected in the second step if it shows up in at least $\beta\Delta$ displayed positive tests as the only unclassified individual giving rise to the following result.

Theorem 3.7 (Theorem 2.1 in [39])**.** *Let $p,q \geq 0$, $p+q<1, d \in (0,\infty), \alpha \in (q, e^{-d}(1-p) + \left(1-e^{-d}\right)q)$ and $\beta \in (0, e^{-d}(1-q))$ and define $w = e^{-d}p + (1-e^{-d})(1-q)$. Suppose that $0<$*

$\theta < 1$ *and let*

$$m_{DD} = m_{DD}(n,\theta,p,q) = \min_{\alpha,\beta,d} \max\left\{c_1(\alpha,d), c_2(\alpha,d), c_3(\beta,d), c_4(\alpha,\beta,d)\right\} k\log(n/k)$$

$$\textit{where} \qquad c_1(\alpha,d) = \frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\alpha\|q\right)}$$

$$\textit{and} \qquad c_2(\alpha,d) = \frac{1}{dD_{\mathrm{KL}}\left(\alpha\|1-w\right)}$$

$$\textit{and} \qquad c_3(\beta,d) = \frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\beta\|(1-q)e^{-d}\right)}$$

$$\textit{and} \qquad c_4(\alpha,\beta,d) = \max_{1-\alpha\le z\le 1}\left\{\frac{1}{1-\theta}\frac{1}{d\left(D_{\mathrm{KL}}\left(z\|w\right) + \mathbf{1}\left\{\beta > \frac{ze^{-d}p}{w}\right\} zD_{\mathrm{KL}}\left(\frac{\beta}{z}\|\frac{e^{-d}p}{w}\right)\right)}\right\}$$

If $m \ge (1+\varepsilon)m_{DD}$ for some $\varepsilon > 0$, then noisy DD *will recover $\boldsymbol{\sigma}$ w.h.p. given test design $\mathbb{G}$ and test results $\hat{\boldsymbol{\sigma}}$.*

The proofs of both theorems are based on similar techniques that we will jointly discuss hereafter. The starting point is the derivation of concentration bounds on the different types of tests (displayed positive, but truly negative etc.). Since tests are flipped independently of each other, a straightforward application of the Chernoff bound does the job. In the next step, we derive the distribution of displayed negative tests for infected and healthy individuals for COMP and the first step of DD as well as the distribution of displayed positive tests in which an individual shows up as the only yet unclassified individual for the second step of DD. Unsurprisingly given the model, these distributions turn out to be hypergeometric. What directly springs to mind is that the expectation of the distributions are well separated between healthy and infected individuals since a healthy individual shows up in more displayed negative tests in expectation. Similarly, an infected individual will show up in more displayed positive tests in expectation as the only yet unclassified individual after we have identified healthy individuals. Therefore, we define thresholds $\alpha\Delta$ and $\beta\Delta$ and perform a large-deviations analysis using the Chernoff bound for hypergeometric distributions on the distributions we derived before. In conjunction with a union bound, we yield the bounds stated in Theorem 3.6 and 3.7.

While the bounds in Theorem 3.6 and 3.7 appear unwieldy at first glance, the optimizations can be efficiently solved for every value of θ, p and q up to arbitrary precision. The generality of these statements has the significant benefit that many common noise models can be derived from the theorems as special cases. The specific noise models we present in our work are the Z-channel ($p = 0, q > 0$), the reverse Z-channel ($p > 0, q = 0$) and the binary symmetric channel ($p = q > 0$). As a corollary, we recover the bounds for the noiseless case. The interested reader is referred to Corollaries 2.5-2.12 in [39].

Given the algorithmic achievability results, it remains to be seen how close the noisy variant of COMP and DD get to the information-theoretic lower bound. A bound for the latter is our next result.

Theorem 3.8 (Theorem 2.3 in [39])**.** *Let $p, q \ge 0$, $p+q < 1$ and $\epsilon > 0$, write $H(\cdot)$ for the binary entropy in nats (logarithms taken to base e) and $\phi = \phi(p,q) = (H(p) - H(q))/(1-p-q)$. If we define*

$$m_{COUNT} = \left(\frac{1}{D_{\mathrm{KL}}\left(q\|1/(1+e^{\phi})\right)}\right) k\log(n/k),$$

then for $m \le (1-\epsilon)m_{COUNT}$ no algorithm can recover $\boldsymbol{\sigma}$ w.h.p. for any matrix design.

The proof of the theorem is based on deriving the capacity of the $p-q$-noise channel. To this end, we maximise the mutual information between the input signal (the true test result, denoted by X) and the output signal (the displayed test result denoted by Y) with respect to the probability of a truly negative test. The mutual information can be formulated using the standard equality

$$I(X,Y) = H(Y) - H(Y \mid X)$$

which can be readily evaluated to find that the optimal probability of a truly negative test is $(T^* - q)/(1-p-q)$ where $T^* = 1/(1+e^{\phi})$ with ϕ defined as above.

While Theorem 3.8 provides an explicit information-theoretic lower bound, we remain agnostic as to whether this bound is actually information-theoretically achievable. Moreover, in the light of the results for the noiseless setting it is not surprising that - even though we improve on prior research results - the noisy variant of the DD algorithm does not achieve the information-theoretic lower bound and there is likely a more sophisticated test design and algorithm that outperforms DD.

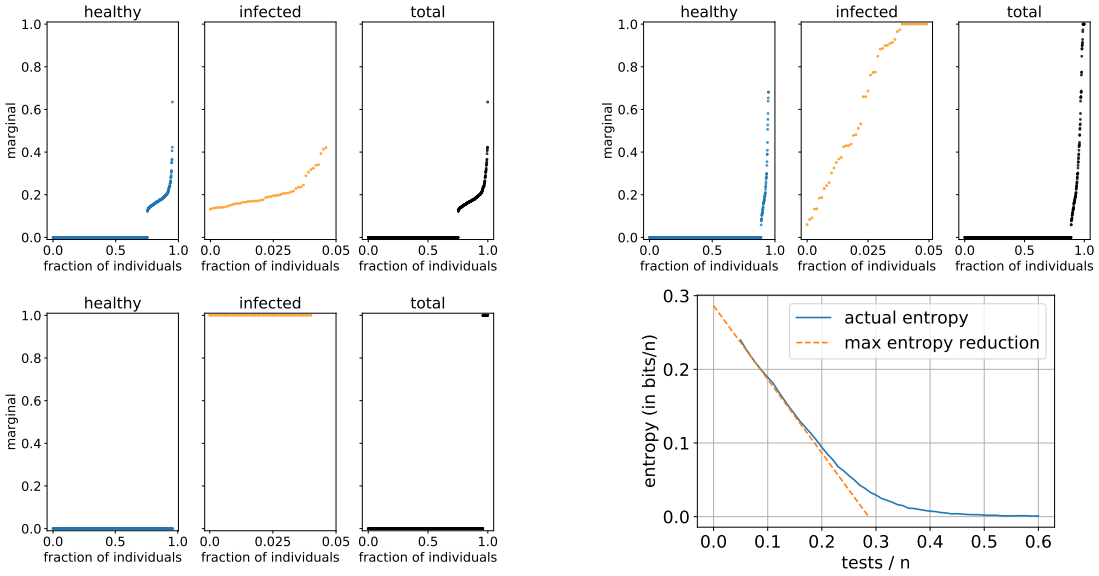

FIGURE 5. Illustration of the posterior distribution from running belief propagation on a random biregular test design for with 0.15 (top left), 0.25 (top right) and 0.6 (bottom left) tests/n and remaining entropy (bottom right) for $\lambda = 0.05$ and the noiseless setting. The figure is adopted from [23].

Considering the asymptotic results from above, a tempting next question is whether the presented algorithms hold merit for practical instances of group testing where the number of individuals is moderately small. In a follow-up empirical work, we tackle this question. In the light of our asymptotic algorithmic results and the success of message-passing algorithms in related problems, it seems natural to run the belief propagation algorithm on a test design where we fix both the degree of individuals and tests. This deliberation is the starting point for our empirical work. Figure 5 shows the posterior distribution of running belief propagation on an instance with 1000 individuals and a prior infection probability of 5% for varying number of tests in the noiseless setting. The graphic evinces that as we increase the number of tests, the posterior distribution becomes more polarised [4] and for

[4]We say that the marginal of an individual from running belief propagation is polarised if it is either 1 or 0, i.e. if the posterior infection probability is either 1 or 0.

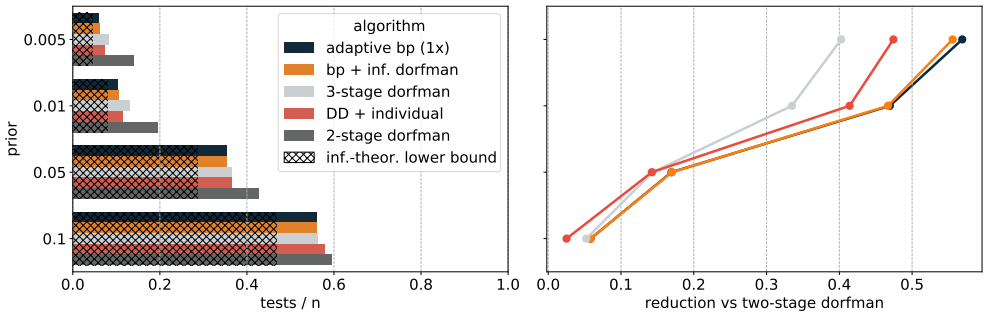

FIGURE 6. Number of tests required (left) and savings potential versus two-stage Dorfman (right) in noiseless setting. The figure is adopted from [23].

sufficiently large tests even allows perfect discrimination between healthy and infected individuals. But is such a one-stage approach optimal? If it was, we would expect the entropy of the factor graph messages to be reduced by one bit per test until the posterior distribution is completely polarised. However, it turns out this information-theoretically optimal reduction only holds up to a number of tests much below the point of polarised posterior distributions. Thus, we considered two- and three-stage algorithms that take advantage of the posterior distribution obtained by running belief propagation. To be precise, we analysed three such algorithms in addition to the traditional two- and three-stage Dorfman procedure. The first algorithm consists of running belief propagation and subsequently testing individuals with non-polarised posterior individually. It should be noted that for the noiseless case, this algorithm boils down to running the well-known DD algorithm presented above with individual follow-up testing for healthy individuals not in V_0^- and infected individuals not in V_1^{--}. For the second algorithm, rather than performing individual testing on non-polarised individuals in the second stage, we employ a procedure called informative Dorfman that is reminiscent of the conventional two-stage Dorfman approach but takes into account the posterior distribution from belief propagation. Third, we split the non-polarised individuals into two groups - one consisting of individuals with small posteriors on which we perform another round of belief propagation on a regular graph design and one consisting of individuals with larger posteriors for which we employ the informative Dorfman procedure. Here, the difference between practical group testing and asymptotic settings becomes most pronounced. If we had $n \to \infty$, each (arbitrarily small) interval of posteriors would feature an unbounded number of individuals and we could perform a more advanced and likely more efficient second-stage algorithm. However, since n is moderately small and thus also the number of individuals with non-polarised marginals, we resort to heuristics that exploit the posterior distribution but are likely not information-theoretically optimal.

Figure 6 shows the number of tests needed for each of these algorithms to succeed in the noiseless case for varying prior infection probabilities. It turns out that particularly for small priors, running belief propagation in the first stage leads to meaningful reductions in the number of tests compared to both two- and three stage Dorfman as well as the conventional DD scheme and comes close to the information-theoretic lower bound.

When it comes to the noisy setting, we find that the more complicated constructions of the second and third algorithm hold their merit. To be precise, with a similar number of tests as belief propagation plus individual testing or the traditional Dorfman procedure, the third algorithm achieves a significantly lower false positive and false negative rate. This algorithm

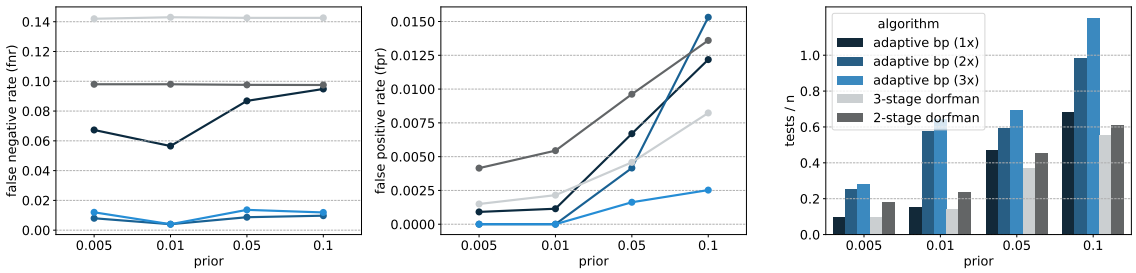

FIGURE 7. Number of tests (left), false positive rate (mid) and false negative rate (right) in noisy setting. The figure is adopted from [23].

also lends itself well if we desire to minimise the number of false positives and false negatives rather than minimising the number of tests with reasonable false positives and negatives. Since most false positives and false negatives originate from the informative Dorfman procedure performed on individuals with large marginals, we can simply perform this step twice or three times. Figure 7 evinces that the false positive and false negative rate can be drastically reduced if we are willing to accept moderately more tests. Thus, the bottom line of the empirical work is that belief propagation allows us to meaningfully enhance the reliability of group testing.

3.4. **Outlook.** While our work on group testing was able to completely resolve the phase diagram for the sublinear noiseless regime, some interesting questions remain. First, it would be interesting to see whether the notion of spatial coupling was truly needed to solve the group testing problem all the way down to the information-theoretic bound or whether an alternative algorithm might also succeed on a plain random regular graph. A suitable candidate might be the belief propagation algorithm. A closer look at the `SPIV` algorithm actually reveals that the crucial second step is indeed the first iteration of belief propagation on a spatially coupled test design. So can belief propagation (possibly with many iterations) also succeed on a plain random regular design? Initial heuristic arguments indicate that without the head start towards the ground truth that spatial coupling provides belief propagation gets trapped in a trivial fix point where the sets V_0^+ and V_1 remain indistinguishable. Further work is warranted to either verify or refute this conjecture.

Moreover, the noisy case for sublinear group testing remains wide open. It is well imaginable that some form of spatially coupled test design with a noisy variant of `SPIV` or belief propagation solves the group testing problem down to the information-theoretic lower bound. For the special case of the binary symmetric channel, we have some indication that it might indeed be the case, but further work is needed to generalise this finding to the $p-q$-noise model.

Finally, and maybe most importantly, the linear regime for group testing where k/n is a small constant remains wide open from a mathematical perspective. For some applications, this regime appears to be more suitable than the sublinear regime. We know from [6] that no non-adaptive algorithm can succeed in the linear regime with high probability. But what about two-stage algorithms or even non-adaptive algorithms where we allow some mistakes? Unfortunately, many of the amenable mathematical properties of the sublinear regime cease to hold in the linear regime, complicating the analysis. A promising route might be to analyse belief propagation on a random regular version for the linear regime, but this task is far from trivial. In a note [40] we already analysed the offspring distribution for the linear regime, which in conjunction with the population dynamics algorithms allows

us to simulate the posterior distributions of the marginal distribution resulting from running belief propagation on a random regular graph. But this starting point is nowhere near a complete analysis of belief propagation and significant work is left. Given the numerous applications of group testing, it seems like an important route for future research.

4. Quantitative group testing

4.1. **Setting & notation.** Let us next consider a variant of the binary group testing problem, namely *quantitative* group testing. In the literature, this problem has alternatively been labelled as the coin weighing problem or as a special case of the pooled data problem. As for binary group testing, we have a large population n out of which k individuals suffer from a rare disease. We will again consider the sublinear regime where $k \sim n^\theta$ for some $\theta \in (0,1)$. Again, we have a testing scheme at our disposal by which we can test groups of individuals. However, in contrast to binary group testing where each test only returns positive or negative, we are provided with the exact number of infected individuals in a test. Clearly, the incremental information that we are provided better facilitates inference than binary group testing and we should achieve lower information-theoretic and algorithmic bounds. Quantitative group testing has a long tradition in mathematics and information theory. Some notable contributions come from Erdős and Rényi [31], Djackov [27], Shapiro [67] and Soderberg and Shapiro [69]. Applications range from DNA screening [66] over identifying genetic carriers in a population [12] to machine learning [53].

Before we get to our results, let us get the notation and model straight. As for the binary group testing problem, quantitative group testing can be understood as an inference problem on random factor graphs. We have a set $V_n = \{x_1, \dots, x_n\}$ of individuals on the one side and a set $F_m = \{a_1, \dots, a_m\}$ of tests on the other side. In contrast to binary group testing, we employ a model where each test selects $\Gamma = n/2$ individuals uniformly at random with replacement. The resulting fluctuating variable degrees will be denoted by $(\Delta_x)_{x \in V_n}$. Thus, our resulting graph $\mathbb{G}$ is much denser than the relatively sparse graph we used for binary group testing. The counting bound in the next section will exhibit why this choice is likely optimal. While individuals will be assigned to $m/2$ tests on average, standard arguments reveal that the expected number of distinct tests Δ^* for an individual is much smaller, i.e. $\big(1 - \exp(-1/2)\big)\, m$. The inference problem starts off as usual with a uniformly sampled ground truth $\boldsymbol{\sigma} \in \{0,1\}^{V_n}$ of Hamming weight k encoding the infection status of the individuals. The ground truth together with the graph $\mathbb{G}$ which is generated independently of the ground truth gives rise to the test results $(\hat{\boldsymbol{\sigma}}_a)_{a \in F_m}$ according to the rule

$$\hat{\boldsymbol{\sigma}}_a = \sum_{x \in \partial a} \boldsymbol{\sigma}_x \qquad (a \in F_m),$$

Note that ∂a denotes the neighbourhood of test a including potential multi-edges. Put differently, an individual might contribute to the test result of test a more than once. If we refer to the distinct neighbourhood of test a or individual x we will write $\partial^* a$ and $\partial^* x$, respectively. The inference task now comes down to the question of whether given $\hat{\boldsymbol{\sigma}}$ and $\mathbb{G}$ we can completely infer the ground truth $\boldsymbol{\sigma}$ w.h.p.? Again, we have the information-theoretic and algorithmic perspective to this question. Let us start by revisiting some prior research.

4.2. **Prior research.** A similar counting bound argument as for binary group testing provides a natural information-theoretic lower bound. Since each test result can in principle exhibit $k+1$ different values, we need to ensure that the number of possible test results $(k+1)^m$ exceeds the number of possible configurations $\binom{n}{k}$. Using Stirling's formula, we

yield

$$m \geq k\frac{\log(n/k)}{\log(k)} \tag{4.1}$$

However, heuristically speaking it is not feasible that a test result exhibits all $k+1$ possible results. Instead, for each test we should see test results fluctuating by the standard deviation around the expectation. This restriction leads to an additional factor of 2 in (4.1). Indeed, Djackov [27] proved that all test designs require at least

$$m_{\text{inf}} = 2k\frac{\log(n/k)}{\log(k)}$$

tests. Using separating matrices, Grebinski and Kucherov [41] provide a (non-constructive) non-adaptive design with an exponential-time algorithm which achieves reliable recovery using $(2+\varepsilon)m_{\text{inf}}$ tests. For the linear regime $k = \Theta(n)$, [4] and [64] evinced the lower bound m_{inf} is information-theoretically achievable. Thus, for the sublinear regime a (constant) gap remains that we close in our work.

When it comes to efficient non-adaptive algorithms, [3] applied the approximate message passing algorithm to quantitative group testing in the linear regime. Approximate message passing can be understood as a computationally efficient version of the belief propagation algorithm for dense graphs. However, the shortcuts taken to make the algorithm more efficient also make the algorithm less analysable from a rigorous standpoint. In any regards, heuristic arguments used by [3] evince that approximate message passing succeeds for quantitative group testing only with $\Theta(\log(n)m_{\text{inf}})$ tests. The state of the art for the sublinear regime is similar. Donoho and Tanner [28] and Foucart and Rauhut [33] presented two efficient algorithms based on ℓ_1-minimisation and the basis pursuit algorithm that get by with

$$2k\log(n/k) \qquad \text{and} \qquad \frac{2}{1-\theta}k\log(n/k)$$

tests, thus also being off from the information-theoretic bound by a multiplicative factor of $\log(n)$. Recently, Karimi et al. [46, 47] presented two algorithms based on graph codes for relatively sparse test designs requiring approximately

$$1.72k\log(n/k) \qquad \text{and} \qquad 1.515k\log(n/k)$$

These results are astonishing in the sense that despite the wealth of additional information provided by quantitative group testing tests, no efficient algorithm is known that can get to a better order than efficient algorithms for *binary* group testing. Indeed, most known algorithm even play in the same league in terms of constants as the algorithms we analysed for binary group testing. These results raise the question of whether there might exist a computationally hard regime in non-adaptive quantitative group testing. In our work, we study a natural greedy neighbourhood algorithm. However, we also only find that it needs $\Theta(k\log(n/k)$ tests to succeed.

4.3. **Results.** Let us start with the non-adaptive information-theoretic upper bound that matches m_{inf} for the sublinear regime.

Theorem 4.1 (Theorem 1.1 in [38])**.** *Suppose that* $0 < \theta < 1$, $k = n^{\theta}$ *and* $\varepsilon > 0$ *and let*

$$m_{inf} = m_{inf}(n,\theta) = 2k\frac{\log(n/k)}{\log(k)}.$$

If $m > (1+\varepsilon)m_{inf}(n,\theta)$*, there exists an (exponential time) algorithm that given* $\mathbb{G}$ *and* $\hat{\boldsymbol{\sigma}}$ *outputs* $\boldsymbol{\sigma}$ *w.h.p.*

The proof of Theorem 4.1 is based on the same technique we used in the first work on binary group testing. To be precise, we calculate the planted first moment and show that for $m > (1+\varepsilon)m_{\mathrm{inf}}$ there does not exist a second configuration that yields consistent test results to the ground truth. Our proof is again split into two parts - one argument for configurations with small overlaps and one for large overlaps.

We start with small overlaps, i.e. $0 \leq \ell < k - (1-\exp(-1/2))\log k$. Recall the definition of $S_k(\mathbb{G}, \hat{\boldsymbol{\sigma}})$ and $Z_{k,\ell}$ from (3.3) and (3.4) with the overlap parameter ℓ. The crucial step towards establishing the absence of alternative configurations for small overlaps lies in the derivation of the following expression.

$$\mathbb{E}[Z_{k,\ell}(\mathbb{G}, \hat{\boldsymbol{\sigma}})] \leq \binom{k}{\ell}\binom{n-k}{k-\ell} \prod_{a \in F_m} \sum_{j=1}^{\hat{\boldsymbol{\sigma}}_a} \binom{\Gamma}{j, j, \Gamma - 2j} \left((1-\ell/k)\frac{k}{n}\right)^{2j} \left(1 - 2(1-\ell/k)\frac{k}{n}\right)^{\Gamma - 2j} \tag{4.2}$$

We shed light on its combinatorial meaning. The binomial coefficients simply count the number of configurations that have overlap ℓ with $\boldsymbol{\sigma}$. The second expression provides an upper bound on the probability that such an alternative configuration yields the same test results. To this end, the sum ranging over the test result considers the probability that we yield the same test result if we flip j infected individuals under $\boldsymbol{\sigma}$ to healthy in the alternative configurations and vice versa. While the binomial coefficient captures the number of such possibilities, the following terms account for the probability that we observe j individuals infected under $\hat{\boldsymbol{\sigma}}$, but healthy under the alternative configuration and vice versa, as well as the probability that $\Gamma - 2j$ individuals remain unflipped. Using standard results on one-dimensional random walks, the expression (4.2) can be readily reformulated to obtain

$$\mathbb{E}[Z_{k,\ell}(\mathbb{G}, \hat{\boldsymbol{\sigma}})] \leq (1+O(1))\binom{k}{\ell}\binom{n-k}{k-\ell}\left(\frac{1}{\sqrt{2\pi}}\mathbb{E}\left[\frac{1}{\sqrt{\boldsymbol{X}}}\right]\right)^m$$

where $\boldsymbol{X}$ is a random variable with distribution $\mathrm{Bin}_{\geq 1}(\Gamma, 2(1-\ell/k)k/n)$. Next, we find that we can move the expectation inside the square root since the Jensen gap vanishes. Some further manipulations finally evince that for sufficiently small ℓ and $m \geq (1+\varepsilon)m_{\mathrm{inf}}$ we have

$$\log\left(\mathbb{E}[Z_{k,\ell}(\mathbb{G}, \hat{\boldsymbol{\sigma}})]\right)/n < 0.$$

Ruling out alternative configurations with a large overlap with $\boldsymbol{\sigma}$ relies on the application of the classical coupon collector argument. While again some technical care is needed, the general strategy is as follows. In any alternative configuration, there must be at least one infected individual under $\boldsymbol{\sigma}$ that was flipped to healthy - otherwise, we would simply consider $\boldsymbol{\sigma}$. Now using standard coupon collector arguments, we can show that in order to compensate for the effect on test results of flipping this one individual at least $(1-\exp(-1/2))\log k$ formerly healthy individuals need to be flipped from healthy to infected. While these additional changes likely lead to further inconsistencies of the test results, the argument suffices to rule out configurations with a large overlap. Combining the large and small overlap argument readily yields the desired statement of the theorem.

Having established a sharp information-theoretic bound, we can move on to efficient algorithms. In our work, we consider a simple greedy algorithm called *Maximum Neighbourhood* (MN) and analyse how it fares in comparison to other more complicated efficient algorithms. The algorithm is based on calculating the adjusted neighbourhood sum for individuals. To be precise, let

$$\Psi_x = \sum_{a \in \partial^* x} \hat{\boldsymbol{\sigma}}_a$$

be the sum of all test results from the distinct tests that individual $x \in V_n$ is assigned to. We will refer to Ψ_x as the neighbourhood sum of x. Also, recall the definition of the number

of distinct tests Δ_x^* for an individual $x \in V_n$. Since each infected individual contributes Δ_x to its neighbourhood sum, the expected sum should be larger for infected than for healthy individuals. Indeed, if we calculate $\Psi_x - \Delta_x^* k/2$ for each individual $x \in V_n$ and classify the largest k individuals as infected, we can recover $\boldsymbol{\sigma}$ with sufficient tests.

Theorem 4.2 (Theorem 1.2 in [38])**.** *Suppose that* $0 < \theta < 1$, $k = n^\theta$ *and* $\varepsilon > 0$. *Further, let*

$$m_{MN}(n,\theta) = 4\left(1 - \frac{1}{\sqrt{e}}\right)\frac{1+\sqrt{\theta}}{1-\sqrt{\theta}}k\log(n/k).$$

If $m > (1+\varepsilon)m_{MN}$, *then the* MN *algorithm outputs* $\boldsymbol{\sigma}$ *w.h.p. given* $\mathbb{G}, \hat{\boldsymbol{\sigma}}, k$.

The proof of the theorem is based on a careful large deviation analysis of the distributions of the neighbourhood sums for infected and healthy individuals. The first item on the agenda is to derive the distributions of the neighbourhood sum if we neglect the impact of individual x in these tests. Let us denote this reduced sum by Φ_x. It turns out that the distribution of this random variable Φ_x is almost independent of the infection status of individual x. We find

$$\Phi_x \sim \mathrm{Bin}\left(\Delta_x^*\Gamma - \Delta_x, \frac{k - \mathbf{1}\{x \in V_1\}}{n-1}\right).$$

Thus, we have $\Psi_x = \Phi_x + \mathbf{1}\{x \in V_1\}\Delta_x$. Clearly, the expectation of the observable neighbourhood sum Ψ_x depends on the infection status of individual x and is thus a natural candidate for thresholding infected and healthy individuals. However, it turns out that the standard deviation of the (unadjusted) variable Ψ_x is too large to attain $m_{MN}(n,\theta)$ due to the fluctuations in Δ_x^*. To account for these fluctuations, we instead consider the variable $\Psi_x - \Delta^* k/2$, i.e. the neighbourhood sum adjusted for the conditional expectation given the number of distinct tests Δ_x^*. With this modification we can define a threshold $\alpha m/2$ and employ a standard Chernoff bound for the binomial distribution to show that when $m > (1+\varepsilon)m_{MN}$ we have

$$\begin{aligned}\Psi_x &\geq \Delta^* k/2 + \alpha m/2 \qquad \forall x \in V_1\\ \Psi_x &< \Delta^* k/2 + \alpha m/2 \qquad \forall x \in V_0.\end{aligned}$$

Above m_{MN}, we can thus recover $\boldsymbol{\sigma}$ w.h.p. Admittedly, the MN algorithm is a simple greedy algorithm and we should not have reasonably hoped for it to attain the information-theoretic lower bound we identified above. Nevertheless, it is interesting to see that despite its simplicity it is on par with significantly more complicated quantitative group testing algorithms for low sparsity regimes.

4.4. **Outlook.** The quantitative group testing problem gives rise to similar open questions as the binary group testing problem, namely considering noisy variants of the problem or adaptive algorithms. However, the central open question pertains to the gap between the information-theoretic bound and the best known algorithmic bounds. For the sparse graph algorithms developed by Karimi et al. [46, 47], it is not surprising that their bound does not coincide with the information-theoretic lower bound since the sparsity of the graph is not information-theoretically optimal. However, the fact that even a sophisticated algorithm such as approximate message passing fails to improve upon the test order $\Theta(k\log(n/k))$ raises the question of whether we face a computationally hard regime for m being larger than m_{inf}, but $o(k\log(n/k))$. In some follow-up work, we applied the notion of spatial coupling with a slightly modified version of the MN algorithm to the quantitative group testing problem. But other than an improvement in the constant of m_{MN}, we did not see major improvements from this approach. This finding is insofar interesting, as the MN algorithm

is reminiscent of a one-iteration message-passing algorithm. To be precise, it is the first iteration that approximate message passing performs when applied to the group testing problem. Clearly, the large number of update iterations that approximate message passing performs might paint a completely different picture, but the fact that approximate message passing without spatial coupling and our first-iteration spatially coupled message-passing algorithm fail in the same order provide some indication that even approximate message passing applied to a spatially coupled test design does not lead to improvements beyond $\Theta(k\log(n/k))$ tests. Thus, it is of fundamental interest to understand whether quantitative group testing indeed exhibits a computationally hard regime.

5. Ising antiferromagnet and max cut

5.1. **Setting & notation.** Our next two works take us to Ising antiferromagnet which is a cornerstone model in combinatorics and statistical physics. Take a graph G with V_n vertices each of which has one of two possible spins ± 1. The interactions of the vertices are encoded by a vertex set E. For any spin configuration $\sigma \in \{\pm 1\}^{V_n}$, let us define the Hamiltonian

$$\mathscr{H}_G(\sigma) = \sum_{(v,w)\in E} \frac{1+\sigma_v\sigma_w}{2}.$$

In addition, we introduce a real parameter $\beta > 0$ - the inverse temperature in physics jargon. Together with β, the Hamiltonian gives rise to the well-known Boltzmann distribution on a configuration $\sigma \in \{\pm 1\}^{V_n}$ defined by

$$\mu_{G,\beta}(\sigma) = \frac{\exp\left(-\beta\mathscr{H}_G(\sigma)\right)}{Z_{G,\beta}} \qquad \text{where} \qquad Z_{G,\beta} = \sum_{\tau\in\{\pm 1\}^{V_n}} \exp\left(-\beta\mathscr{H}_G(\tau)\right). \tag{5.1}$$

The normalising term $Z_{G,\beta}$ is our well-known partition function which will turn out to be of central importance. The distribution $\mu_{G,\beta}$ favours configurations with few monochromatic edges, i.e. edges between vertices of the same spin. This model is known as the Ising antiferromagnet. In the following, we will study the Ising antiferromagnet on random regular graphs $\mathbb{G} = \mathbb{G}(n,d)$. There are a number of interesting research questions that the Ising antiferromagnet raises in its own right. To be precise, it is easy to see that the distribution $\mu_{G,\beta}$ gives rise to short-range interactions between vertices in the sense that under $\mu_{G,\beta}$ the spins of two close vertices are correlated. The strength of these correlations is determined by the choice of β with large values of β yielding configurations with relatively few monochromatic edges. A key question pertains to the degree of correlation between two distant vertices in the graph. According to physics predictions, there should exist a threshold value such that for smaller values of β we should observe a rapid decay of correlations and thus no long-range correlations between vertices. This regime is known as the replica symmetric phase. The defining characteristic of this phase is that w.h.p. two independent samples $\boldsymbol{\sigma}_1, \boldsymbol{\sigma}_2$ from the Boltzmann distribution $\mu_{\mathbb{G},\beta}$ exhibit an almost flat overlap in the sense that $|\boldsymbol{\sigma}_1 \cdot \boldsymbol{\sigma}_2| = o(n)$. Conversely, for larger values of β correlations should persist between distant vertices giving rise to what physicists call replica-symmetry breaking. This threshold for the emergence of long-range correlations is predicted to be at the combinatorially meaningful Kesten-Stigum bound

$$\beta_{\mathrm{KS}} = \log\left(\frac{\sqrt{d-1}+1}{\sqrt{d-1}-1}\right).$$

The Ising antiferromagnet is also intimately related to the MaxCut problem - a prominent example of an inference problem. Given a graph G, the MaxCut problem asks for a

partition of the vertex set into two classes such that the maximum number of edges are between vertices of different classes. This problem can be readily formulated as an inference problem on random factor graphs. The vertices in the graph G form the set of variable nodes in this factor graph. Similarly, for each edge in the original graph we create a factor node in the factor graph and connect it to the respective variable nodes. Given the resulting factor graph, the goal is to infer a configuration $\boldsymbol{\sigma} \in \{\pm 1\}^{V_n}$ such that the maximum number of factor nodes are connected to variables nodes of different spins. Here, we are interested in the information-theoretic thresholds of the problems rather than efficient algorithms. As for other inference problems, we find that this problem is tightly connected to the partition function of the Boltzmann distribution (5.1). To be precise, for a random d-regular graph $\mathbb{G}$ on n vertices we find

$$\frac{2\mathbb{E}\left[\textsc{MaxCut}(\mathbb{G})\right]}{dn} = 1 - \frac{2\mathbb{E}\left[\min_{\sigma\in\{\pm1\}^{V_n}} \mathscr{H}_{\mathbb{G}}(\sigma)\right]}{dn} \leq 1 + \frac{2}{\beta dn}\mathbb{E}\left[\log Z_{\mathbb{G},\beta}\right]. \tag{5.2}$$

While the above inequality holds for any $\beta > 0$, it seems natural to take $\beta \to \infty$ to derive a tighter bound on the maximum cut size. Indeed, note that for small values of β, the Boltzmann distribution from (5.1) is close to a uniform distribution over all configurations $\sigma \in \{\pm 1\}$ irrespective of the number of monochromatic edges. However, when $\beta \to \infty$ the distribution becomes a point mass on the maximum cut configuration. Therefore, in order to bound the maximum cut size we need to get a handle on $\lim_{\beta\to\infty}\lim_{n\to\infty}\mathbb{E}\left[\log Z_{\mathbb{G},\beta}\right]/(n\beta)$. To this end, we will employ the interpolation method from mathematical physics already touched upon in the introduction.

5.2. **Prior research.** The Ising model has a long tradition in statistical physics. Invented by Lenz in 1920 [52] to explain magnetism, it is suggested to be one of the simplest models where replica-symmetry breaking occurs. Mézard and Parisi [54, 55] were the first to investigate replica-symmetry breaking in the Ising model using the non-rigorous cavity method. For the Ising ferromagnet, where vertices of identical spin are attracted to each other, these predictions were verified by Dembo and Montanari [24] by analysing the belief propagation recurrences on a random tree.

Since the disassortative stochastic block model with two communities that we touched upon in the introduction is nothing but the planted version of the Ising antiferromagnet, the result by Mossel, Neeman and Sly [60] shows the existence and threshold for replica-symmetry breaking in the Ising antiferromagnet for the Erdős-Rényi model. Moreover, a result by [15] evinces replica-symmetry breaking to occur for the Ising antiferromagnet on random regular graphs. However, this result only comes as an infinite-dimensional variational problem. The contribution of our work is to provide an explicit formula for the replica symmetry breaking phase transition. Moreover, we will derive the limiting distribution of $Z_{\mathbb{G},\beta}$ in the replica symmetric phase.

When it comes to the MaxCut problem on random graphs, several attempts have been made to derive lower and upper bounds. The derivation of upper bounds has thus far been mostly based on the classical first moment method, while the analysis of greedy algorithms yielded lower bounds. Table 1 provides the best known rigorous bounds for random regular graphs.

Prior techniques and algorithms mostly relied on local arguments. Thus, it is not surprising that up to now lower and upper bounds lie apart by quite a margin. The first non-local and more advanced approach to resolve the bounds for the MaxCut problem was taken by Dembo, Montanari and Sen [25] using the interpolation method and an inherent connection between Erdős-Rényi graphs and the Sherrington-Kirkpatrick model. Their results pertain to the case that $d \to \infty$, but do not carry any information about finite d regimes and

d	3	4	5	6	7	8	9	10
best previous upper bound	0.9320	0.8900	0.8539	0.8260	0.8038	0.7855	0.7701	0.7570
new upper bound	0.9241	0.8683	0.8350	0.8049	0.7851	0.7659	0.7523	0.7388
best lower bound	0.9067	0.8333	0.7989	0.7775	0.7571	0.7404	0.7263	0.7144
expected cut size at β_{KS}	0.8536	0.7887	0.7500	0.7236	0.7041	0.6890	0.6768	0.6667

TABLE 1. Bounds on the fraction of edges in a maximum cut of $\mathbb{G}(n,d)$. The table is adopted from [18].

random *regular* graphs. For that, we need to look at works by Panchenko [63] and Coja-Oghlan and Perkins [15] from which bounds on the MAXCUT problem come as the solution to an infinite-dimensional optimisation problem. Getting any specific bounds out of these results is far from trivial, if at all possible. At the same time, Zdeborová and Boettcher [72] put forward a beautiful conjecture on the maximum cut and min bisection size on random regular graphs based on non-rigorous statistical physics techniques. In our work, we will derive explicit upper bounds on the maximum cut size on random regular graphs that precisely match the prediction by Zdeborová and Boettcher [72].

5.3. **Results.** As the first key result, we pinpoint the replica-symmetry breaking phase transition of the Ising antiferromagnet on random regular graphs. To this end, let

$$\Phi_d : \beta \in (0,\infty) \to \lim_{n\to\infty} \mathbb{E}\left[\log Z_{\mathbb{G},\beta}\right]/n.$$

Theorem 5.1 (Theorem 1.1 in [18])**.** *For any $d \geq 3$ let*

$$\beta_{KS} = \beta_{KS}(d) = \log\left(\frac{\sqrt{d-1}+1}{\sqrt{d-1}-1}\right). \tag{5.3}$$

(i) If $\beta < \beta_{KS}(d)$, then

$$\Phi_d(\beta) = \log 2 + \frac{d}{2}\log\frac{1+e^{-\beta}}{2}.$$

(ii) If $\beta > \beta_{KS}(d)$, then

$$\Phi_d(\beta) < \log 2 + \frac{d}{2}\log\frac{1+e^{-\beta}}{2}.$$

The proof of the first statement is based on the method of moments. In its classical application, we would derive expressions for $\mathbb{E}\left[Z_{\mathbb{G},\beta}\right]$ and $\mathbb{E}\left[Z^2_{\mathbb{G},\beta}\right]$ and show that $\mathbb{E}\left[Z^2_{\mathbb{G},\beta}\right] = O\left(\mathbb{E}\left[Z_{\mathbb{G},\beta}\right]^2\right)$ for $\beta < \beta_{\mathrm{KS}}$. Then, standard concentration arguments imply the first statement of the theorem. However, in the case of the Ising antiferromagnet things turn out be mildly more complicated and we have to resort to different methods. To be precise, we first exhibit an event $\mathcal{O}$ as

$$\mathcal{O} = \left\{\mathbb{E}\left[\left|\boldsymbol{\sigma}_{\mathbb{G}}\cdot\boldsymbol{\sigma}'_{\mathbb{G}}\right| \mid \mathbb{G}\right] < \varepsilon_n n\right\}$$

for a sequence $\varepsilon_n = o(1)$ and two independent samples $\boldsymbol{\sigma}_{\mathbb{G}}, \boldsymbol{\sigma}'_{\mathbb{G}}$ from the Boltzmann distribution. Put differently, $\mathcal{O}$ is the event that two typical samples from the Boltzmann distribution are almost orthogonal. For the proof of the first statement, we will apply the methods of moments to the partition function restricted on event $\mathcal{O}$. To this end, we will first show that for $\beta < \beta_{\mathrm{KS}}$ we have $\mathbb{E}\left[Z_{\mathbb{G},\beta}\mathbf{1}\{\mathcal{O}\}\right] = \Theta\left(\mathbb{E}\left[Z_{\mathbb{G},\beta}\right]\right)$. Second, we will calculate $\mathbb{E}\left[Z^2_{\mathbb{G},\beta}\mathbf{1}\{\mathcal{O}\}\right]$ which turns out to be a more amenable opponent than $\mathbb{E}\left[Z^2_{\mathbb{G},\beta}\right]$. The proof of the first component is based on the insight that if $\mathbb{P}[\mathbb{G}^* \in \mathcal{O}] \sim 1$, then indeed $\mathbb{E}\left[Z_{\mathbb{G},\beta}\mathbf{1}\{\mathcal{O}\}\right] = \Theta\left(\mathbb{E}\left[Z_{\mathbb{G},\beta}\right]\right)$ where

$\mathbb{G}^*$ denotes the planted model of the Ising antiferromagnet. With this result at hand, we couple the planted model $\mathbb{G}^*$ with a broadcasting process on the infinite $(d-1)$-ary tree. The key idea is that such a $(d-1)$-ary tree suitably models the local neighbourhood of a sparse planted model. At the root of the tree, we select a spin uniformly at random. Then, we construct the tree downwards by sampling a spin for the children of a given parent as follows. With probability $e^{-\beta}/(1+e^{-\beta})$ we select the same spin as for the parent, while with probability $1/(1+e^{-\beta})$ we select a different spin. Let $\mathscr{F}_\ell$ denote the σ-algebra generated by the spins of vertices at distance larger than ℓ from the root. For $\beta < \beta_{\mathrm{KS}}$, it turns out that the tree exhibits a rapid correlation decay and quickly 'forgets' the spin of the root - a phenomenon known as non-reconstruction. Formally, with $\boldsymbol{\tau}_{v_0}$ denoting the spin at the root of the tree we have

$$\lim_{\ell\to\infty} \mathbb{E}\left[\mathbb{P}\left[\boldsymbol{\tau}_{v_0} = 1 \mid \mathscr{F}_\ell\right] - \frac{1}{2}\right] = 0. \tag{5.4}$$

With (5.4) at hand, we readily find that

$$\mathbb{E}\left[Z_{\mathbb{G},\beta}\mathbf{1}\left\{\mathscr{O}\right\}\right] = \Theta(\mathbb{E}\left[Z_{\mathbb{G},\beta}\right]) = \Theta\left(2^n\left(\frac{1+e^{-\beta}}{2}\right)^{dn/2}\right)$$

For the second moment, we have to consider

$$\mathbb{E}\left[Z^2_{\mathbb{G},\beta}\mathbf{1}\left\{\mathscr{O}\right\}\right] = \sum_{\sigma,\sigma'\in\{\pm1\}^{V_n}} \mathbf{1}\left\{|\sigma\mathbf{1}|, \left|\sigma'\mathbf{1}\right|, \left|\sigma,\sigma'\right| \le \delta n\right\}\mathbb{E}\left[\exp\left(-\beta\mathscr{H}_{\mathbb{G}}(\sigma) - \beta\mathscr{H}_{\mathbb{G}}(\sigma')\right)\right].$$

The event $\mathscr{O}$ thus allows us to focus our attention to spin configurations σ,σ' that have an almost flat overlap. Starting off with combinatorial arguments and bounding the second moment of random regular graphs by the second moment of the Erdős-Rényi graph we find

$$\mathbb{E}\left[Z^2_{\mathbb{G},\beta}\right] \le \exp\left(n \max_{\alpha\in[-1,1]} f_d(\alpha,\beta)\right) \qquad \text{where}$$

$$f_d(\alpha,\beta) = (1-d)\log(2) + H((1+\alpha)/2) + \frac{d}{2}\log\left(\left(1+e^{-\beta}\right)^2 + \alpha^2\left(1-e^{-\beta}\right)^2\right).$$

Since we can calculate $Z^2_{\mathbb{G},\beta}$ conditioned on event $\mathscr{O}$, we finally obtain

$$\mathbb{E}\left[Z^2_{\mathbb{G},\beta}\mathbf{1}\left\{\mathscr{O}\right\}\right] \le \exp\left(nf_d(0,\beta) + o(1)\right) = \mathbb{E}\left[Z_{\mathbb{G},\beta}\right]^2\exp(o(n)).$$

For the second part of the theorem, we strike a chord to the free energy in the planted model and leverage a result from [22]. To state the result, let $\mathscr{P}_*([-1,1])$ be the space of all probability measures on the interval $[-1,1]$ with mean 0. Moreover, for a probability measure π, let $(\boldsymbol{\mu}_{\pi,i})_{i\ge1}$ be a family of independent samples from π. Finally, we let $\Lambda(x) = x\log x$. The key towards the proof lies in the so-called Bethe free energy defined as

$$\mathscr{B}_{\mathrm{Ising}}(\pi,\beta,d) = \mathbb{E}\left[\frac{\Lambda\left(\sum_{\sigma\in\pm1}\prod_{i=1}^d 1-(1-e^{-\beta})(1+\sigma\boldsymbol{\mu}_{\pi,i})/2\right)}{2^{1-d}(1+e^{-\beta})^d} - \frac{d\Lambda\left(1-(1-e^{-\beta})(1+\boldsymbol{\mu}_{\pi,1}\boldsymbol{\mu}_{\pi,2})/2\right)}{1+e^{-\beta}}\right] \tag{5.5}$$

From [22] we know that the second statement of our theorem holds if

$$\sup_{\pi\in\mathscr{P}_*([-1,1])} \mathscr{B}_{\mathrm{Ising}}(\pi,\beta,d) > \lim_{n\to\infty}\log\left(\mathbb{E}\left[Z_{\mathbb{G},\beta}\right]\right)/n. \tag{5.6}$$

In order to establish (5.6), it suffices to show the inequality for any probability measure $\pi\in\mathscr{P}_*([-1,1])$. To this end, we consider a slightly perturbed probability distribution and perform a Taylor expansion of $\mathscr{B}_{\mathrm{Ising}}(\pi,\beta,d)$ around a zero perturbation to the fourth order. While the first order precisely matches the first moment, the second and third order turn out

to be zero. But for the fourth order, we eventually find that as long as $\beta > \beta_{\mathrm{KS}}$ we indeed have a positive value thereby establishing the second statement of the theorem and pinpointing the replica-symmetry breaking phase transition.

As an immediate corollary of Theorem 5.1 and the fact that the disassortative stochastic block model with two communities is the planted Ising antiferromagnet, we can describe the Kullback-Leibler divergence between the planted and the null model.

Theorem 5.2 (Theorem 1.4 in [18]). *For any $d \geq 3$ the following are true.*

(i) If $\beta < \beta_{KS}(d)$, then $\lim_{n\to\infty} D_{\mathrm{KL}}(\mathbb{G}^* \| \mathbb{G})/n = 0$ and $\lim_{n\to\infty} D_{\mathrm{KL}}(\mathbb{G} \| \mathbb{G}^*)/n = 0$.

(ii) If $\beta > \beta_{KS}(d)$, then $\lim_{n\to\infty} D_{\mathrm{KL}}(\mathbb{G}^* \| \mathbb{G})/n > 0$ and $\lim_{n\to\infty} D_{\mathrm{KL}}(\mathbb{G} \| \mathbb{G}^*)/n > 0$.

Arguably, the approximation of the partition function in Theorem 5.1 is relatively crude with an error term of $\exp(o(n))$ for $\beta < \beta_{\mathrm{KS}}$. As the following result evinces, we can do better than that and derive the limiting distribution of $Z_{\mathbb{G},\beta}$ in the replica symmetric phase.

Theorem 5.3 (Theorem 1.1 in [32]). *Assume that $0 < \beta < \beta_{KS}$ and $d \geq 3$. Let $(\Lambda_i)_i$ be a sequence of independent Poisson variables with $\mathbb{E}[\Lambda_i] = \lambda_i$ where $\lambda_i = \frac{(d-1)^i}{2i}$. Further, let $\delta_i = \left(\frac{e^{-\beta}-1}{e^{-\beta}+1}\right)^i$. Then as $n \to \infty$ we have*

$$\log\left(Z_{\mathbb{G},\beta}\right) - \frac{1}{2}\log\left(\frac{1+e^{\beta}}{2+de^{\beta}-d}\right) - n\left(\left(1-\frac{d}{2}\right)\log(2) + \frac{d}{2}\log\left(1+e^{-\beta}\right)\right)$$
$$\xrightarrow{d} \log(W) \quad \text{where} \quad W := \exp\left(-\lambda_1\delta_1 - \lambda_2\delta_2\right)\prod_{i=3}^{\infty}(1+\delta_i)^{\Lambda_i}\exp\left(-\lambda_i\delta_i\right).$$

The infinite product defining W converges a.s. and in L^2.

To derive the limiting distribution for $Z_{\mathbb{G},\beta}$ in the replica symmetric regime, we need to identify the sources of fluctuations in $Z_{\mathbb{G},\beta}$. One obvious such source are the number of short cycles in $\mathbb{G}$. Indeed, it turns out that once we condition on the number of short cycles, the variance of $Z_{\mathbb{G},\beta}$ vanishes. More formally, let $C_i(G)$ be the number of short cycles of length i in a graph G and $\mathscr{F}_\ell$ the σ-algebra generated by the random variables $C_i(\mathbb{G})$ for $i \leq \ell$. By standard decomposition of the variance we have

$$\mathbb{E}\left[Z_{\mathbb{G},\beta}^2\right] - \mathbb{E}\left[Z_{\mathbb{G},\beta}\right]^2 = \mathbb{E}\left[\mathbb{E}\left[Z_{\mathbb{G},\beta} \mid \mathscr{F}_\ell\right]^2 - \mathbb{E}\left[Z_{\mathbb{G},\beta}\right]^2\right] + \mathbb{E}\left[\mathbb{E}\left[Z_{\mathbb{G},\beta}^2 \mid \mathscr{F}_\ell\right] - \mathbb{E}\left[Z_{\mathbb{G},\beta} \mid \mathscr{F}_\ell\right]^2\right].$$

It turns out that when $\beta < \beta_{\mathrm{KS}}$ the second term accounting for the conditional variance given the number of short cycles vanishes. Thus, the entire variance of $Z_{\mathbb{G},\beta}$ is due to fluctuations in the number of short cycles in $\mathbb{G}$. Our proof of the limiting distribution relies on a combination of the methods of moments and small subgraph conditioning enriched by similar spatial mixing arguments as employed in our work to pinpoint the replica-symmetry breaking phase transition in the Ising model. To be precise, we leverage a result by Janson [44] that allows us to eventually state the limiting distribution of $Z_{\mathbb{G},\beta}$. In order to apply this result, a few steps are needed. As a first step, we need to get a handle on the distribution of short cycles in random regular graphs $\mathbb{G}$ and the planted model $\mathbb{G}^*$, i.e. the disassortative stochastic block model. While the former is a well-established result, the latter constitutes a major contribution of our work. Let

$$\delta_i = \left(\frac{e^{-\beta}-1}{e^{-\beta}+1}\right)^i \qquad \text{and} \qquad \lambda_i = \frac{(d-1)^i}{2i}$$

and consider a sequence of independent Poisson random variables for $i \geq 3$ defined by

$$\Lambda_i \sim \mathrm{Po}(\lambda_i) \qquad \text{and} \qquad \Xi_i \sim \mathrm{Po}(\lambda_i(1+\delta_i)).$$

While it is well known that jointly for all i the number of short cycles of length i in $\mathbb{G}$ converges in distribution to Λ_i, we show that the number of short cycles of length i in the planted model $\mathbb{G}^*$ converges in distribution to Ξ_i.

As a second step, we need to sharpen our pencil and derive a sharper approximation of the first and second moment of $Z_{\mathbb{G},\beta}$. As before, instead of considering $Z_{\mathbb{G},\beta}$ directly we will analyse $Z_{\mathbb{G},\beta}\mathbf{1}\{\mathscr{O}\}$. As before, we use the fact that for $\beta < \beta_{\mathrm{KS}}$ we have $\mathbb{E}\left[Z_{\mathbb{G},\beta}\mathbf{1}\{\mathscr{O}\}\right] = (1 + o(1))\mathbb{E}\left[Z_{\mathbb{G},\beta}\right]$. With some rather complicated and tedious calculations, we find for $\beta < \beta_{\mathrm{KS}}$

$$\begin{aligned}\mathbb{E}\left[Z_{\mathbb{G},\beta}\right] &= \exp\left(-\lambda_1\delta_1 - \lambda_2\delta_2 + O(1/n)\right)\sqrt{\left(1+e^{\beta}\right)/\left(2+de^{\beta}-d\right)} \\ &\quad\cdot\exp\left(n\left((1-d/2)\log(2) + d\log\left(1+e^{-\beta}\right)/2\right)\right) \\ \mathbb{E}\left[Z_{\mathbb{G},\beta}^2\mathscr{O}\right] &= \exp\left(\lambda_1+\lambda_2 - \frac{4\lambda_1}{\left(1+e^{\beta}\right)^2} - \frac{4\lambda_2\left(1+e^{2\beta}\right)^2}{\left(1+e^{\beta}\right)^4} + O\left(\frac{1}{n}\right)\right) \\ &\quad\cdot\frac{\left(1+e^{\beta}\right)^2\exp\left(n\left((2-d)\log(2) + d\log\left(1+e^{-\beta}\right)\right)\right)}{\left(de^{\beta}-d+2\right)\sqrt{2e^{2\beta}+2de^{\beta}-de^{2\beta}-d+2}}.\end{aligned}$$

With this tighter expression of the first and second moment at hand, we can establish that

$$\frac{\mathbb{E}\left[Z_{\mathbb{G},\beta}^2\{\mathscr{O}\}\right]}{\mathbb{E}\left[Z_{\mathbb{G},\beta}\{\mathscr{O}\}\right]^2} = (1+o(1))\exp\left(\sum_{i\geq 3}\lambda_i\delta_i^2\right)$$

which suffices to apply the result by Janson [44] and obtain the limiting distribution of $Z_{\mathbb{G},\beta}\mathbf{1}\{\mathscr{O}\}$. A rather straightforward application of Markov's inequality then allows us to transfer this result to $Z_{\mathbb{G},\beta}$ and thus establish the theorem.

Having established the replica-symmetry breaking phase transition at the combinatorially meaningful Kesten-Stigum bound and having derived the limiting distribution of $Z_{\mathbb{G},\beta}$ in the replica symmetric phase, we turn our attention to the inference problem of MaxCut. Our following result establishes an upper bound on the maximum cut size of random regular graphs. To this end, let $\mathscr{M}$ be a right stochastic block matrix of size $(d+1)\times(d+1)$

$$\mathscr{M} = \begin{bmatrix} 0 & 1 & 0 & \cdots & \cdots & \cdots & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} & \ddots & & & \vdots \\ 0 & \frac{1}{2} & 0 & \frac{1}{2} & \ddots & & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ \vdots & & \ddots & \ddots & \ddots & \ddots & 0 \\ \vdots & & & \ddots & \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & \cdots & \cdots & \cdots & 0 & 1 & 0 \end{bmatrix}. \tag{5.7}$$

Moreover, we define

$$F_d(\alpha, z) = -\frac{\log\left(\zeta\mathscr{A}^d\xi\right)}{\log z} + \frac{d\log\left(1-2\alpha^2+2\alpha^2 z\right)}{2\log z}, \qquad \text{where} \tag{5.8}$$

$$\mathscr{A} = (1-2\alpha)\mathrm{id} + 2\alpha\sqrt{z}\mathscr{M}, \tag{5.9}$$

$$\zeta = \begin{bmatrix}1, & 0, & 0, & \cdots\end{bmatrix} \in \mathbb{R}^{1\times(d+1)}, \tag{5.10}$$

$$\xi = \begin{bmatrix}1, & z^{-1/2}, & z^{-1}, & z^{-3/2}, & \cdots\end{bmatrix}^T \in \mathbb{R}^{(d+1)\times 1}. \tag{5.11}$$

Theorem 5.4 (Theorem 1.2 in [17]). *For any $d \geq 3$ we have*

$$\lim_{\beta\to\infty} \beta^{-1}\Phi_d(\beta) \leq \inf_{\substack{0<\alpha\leq 1/2\\ 0<z<1}} F_d(\alpha, z).$$

As a corollary to this theorem, we obtain the following bound on the max cut problem that matches precisely the conjectured values by Zdeborová and Boettcher [72].

Corollary 5.5 (Corollary 1.3 in [18]). *Let* $\textsc{MaxCut}(\mathbb{G})$ *be the number of edges cut by a maximum cut of* $\mathbb{G}$. *Then, w.h.p.,*

$$\textsc{MaxCut}(\mathbb{G}) \leq \frac{dn}{2} \inf_{\substack{0<\alpha<1/2\\ 0<z<1}} \left(1 + \frac{2}{d} F_d(\alpha, z)\right) + o(n).$$

As described above, the key towards deriving an upper bound on the max cut problem is to upper bound

$$\lim_{\beta\to\infty} \lim_{n\to\infty} \mathbb{E}\left[\log Z_{\mathbb{G},\beta}\right]/(n\beta).$$

One of the techniques in our toolbox is the so-called interpolation method. The basic idea of the interpolation method is to construct a sequence of graphs parametrised over $t \in [0,1]$. Figure 8 provides an illustration. At $t = 1$, the graph coincides with our original factor graph. At time $t = 0$, we have decomposed the graph into independent variables nodes (white) and a set of negative factor nodes (red). Each variable nodes is still connected to d factor nodes (blue), but instead of the variables interacting with each other, they only interact through a central node (orange) that models the interaction between variable nodes in the original graph. As we move from the original graph $\mathbb{G}$ at $t = 1$ to the decomposed graph at $t = 0$ we essentially remove the factor nodes connecting the variable nodes and replace them with independent factor nodes only connected to the central orange node. Since in effect we are adding twice the number of factor nodes that the original graph features, we obtain a set of negative factor nodes. The crucial insight behind the interpolation method is that the partition function only increases as we decrease t. Thus, the partition function of the original graph $\mathbb{G}$ is upper bounded by the partition function of the decomposed graph. In our work, we do not perform the interpolation method in full but leverage results on the Potts model which constitutes a generalisation of the Ising model to more spins. Thus, we get an upper bound on $\lim_{n\to\infty} \mathbb{E}\left[\log Z_{\mathbb{G},\beta}\right]/n$ in terms of an expression over any probability distribution on $[-1,1]$. While the general bound is cumbersome and does not lend itself well for explicit bounds, we are free to choose the probability distribution in order to get an explicit, yet ideally tight bound. To this end, we follow physics intuition and consider a candidate distribution that should recover the bound by Zdeborová and Boettcher [72]. While the expression pertaining to the set of negative factor nodes in the interpolation method can be readily evaluated under this distribution, matters are more complicated for the independent variable nodes. Fortunately, we find a connection between the candidate distribution and a certain random walk that allows us to derive the expression in Theorem 5.4. While the expression looks unwieldy at first glance, it can be numerically evaluated and we indeed find that the upper bound on the maximum cut size matches the conjecture by [72].

5.4. **Outlook.** While our results provide a meaningful step to understand the Ising antiferromagnet on random regular graphs, they leave a few questions open. First, the upper bound on the maximum cut size must not need to be tight. Indeed, in the light of the result of Coja-Oghlan & Perkins [15] we should not expect it to be tight and it would be interesting to see whether we can derive tighter lower and upper bounds. Second, having derived the limiting distribution of $Z_{\mathbb{G},\beta}$ in the replica symmetric phase, one might wonder if something similar is

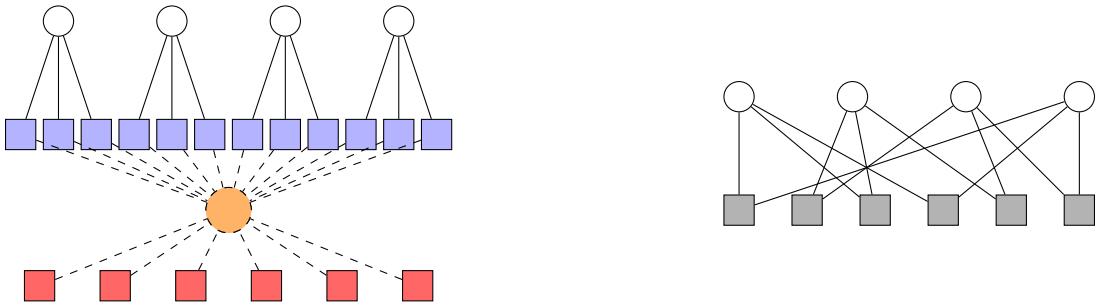

FIGURE 8. Illustration of the interpolation method adopted from [17] showing the decomposed factor graph at $t = 0$ and the original factor graph at $t = 1$.

also possible for $\beta > \beta_{\mathrm{KS}}$. We should not be too optimistic in this regard, since for $\beta > \beta_{\mathrm{KS}}$ we should expect long-range correlations to emerge in $\mathbb{G}$ giving rise to far more intricate fluctuations in $Z_{\mathbb{G},\beta}$. These fluctuations likely result in a far more complicated distribution of $Z_{\mathbb{G},\beta}$. Deriving this distribution seems out of reach with today's techniques. However, what seems to be in reach is vindicating the absence and presence of long-range correlations in $\mathbb{G}$. While it should be a stone's throw from Theorem 5.3 to the absence of long-range correlations in $\mathbb{G}$, establishing the presence of long-range correlations in $\mathbb{G}$ above the Kesten-Stigum bound is a far more challenging, yet important undertaking.

6. Zusammenfassung

6.1. **Einleitung.** Viele Fragestellungen der Mathematik und Informatik können als Inferenzprobleme auf zufälligen Faktorgraphen formuliert werden. Nehmen wir das Beispiel einer hoch-dimensionalen linearen Regression der Form

$$y = X\beta + w$$

mit $y \in \mathbb{R}^m$, $X \in \mathbb{R}^{mxn}$ mit unabhängigen, identisch verteilten (u.i.v.) $N(0,1)$ Einträgen, einem Parametervektor $\beta \in \{0,1\}^n$ und einem Rauschen $w \in \mathbb{R}^m$ mit u.i.v. $N(0,\sigma^2)$ Einträgen [62]. Das Ziel in diesem Inferenzproblem ist es, aus m verrauschten Beobachtungen (der Vektor y) und der ursprünglichen Eingabe X den Parametervektor β zu ermitteln. Dieses Problem kann man sich als einen bipartiten Graphen vorstellen, bei der wir auf der einen Seite n Variablenknoten haben, die die Einträge im Parametervektor β repräsentieren und auf der anderen Seite m Faktorknoten, die die verrauschten Beobachtungen darstellen. Die Interaktion der Parameter β, um die Beobachtungen y zu erhalten, sind in der Matrix X abgebildet, die somit die Graphenstruktur des Faktorgraphen enkodiert.

Hoch-dimensionale lineare Regression is nur eines von vielen Problemen, die man sich als Inferenzproblem auf (zufälligen) Faktorgraphen vorstellen kann [57, 73]. Andere prominente Probleme, die an dieser Stelle genannt werden sollten, sind Graph Clustering [1, 43], Hauptkomponentenanalyse [11] oder Bedingungserfüllungsprobleme [13, 26]. Dieses Gebiet ist ein aktives Forschungsfeld und in den vergangenen Jahren sind verschiedene Techniken entwickelt worden, um Inferenzprobleme auf (zufälligen) Faktorgraphen informationstheoretisch und algorithmisch zu analysieren. Die meisten dieser Probleme lassen sich mit der Analogie des Lehrer-Schüler-Modells verdeutlichen [73]. Nehmen Sie an, dass ein Lehrer eine Anfangsbelegung aus einer bestimmten, meist bekannten Verteilung zufällig auswählt. Im Fall der Regression wäre dies der Parametervektor β. Im Anschluss erzeugt er einen zufälligen Faktorgraphen auf der Grundlage dieser Anfangsbelegung. Dieser Faktorgraph enthält neben den Kanten zwischen Variablen und Faktoren auch Informationen auf den Faktoren, die sich aus der Anfangsbelegung ableiten. Im Falle der linearen Regression ist die Matrix X unabhängig von dem Parametervektor β, aber die verrauschten Beobachtungen y, die auf den Faktorknoten gespeichert werden, enthalten Informationen über β. Nun übergibt der Lehrer den Faktorgraphen mit den auf den Faktorknoten gespeicherten Informationen an einen Schüler, allerdings ohne die Anfangsbelegung β zu verraten. Die Aufgabe des Schülers ist es nun, aus der Graphenstruktur und den auf den Faktoren gespeicherten Informationen die Anfangsbelegung zu rekonstruieren. Natürlich hängt die Möglichkeit, die Anfangsbelegung zu rekonstruieren, vom Signal-Rausch-Verhältnis ab, das dem Problem zugrunde liegt. Für die linear Regression ist das Signal die Anzahl der Messungen y, das Rauschen hingegen die Varianz σ^2 im Vektor w. Die zentrale Frage ist, ab welchem Verhältnis von Signal zu Rauschen eine Rekonstruktion der Anfangsbelegung möglich ist. Es gibt zwei Perspektiven auf diese Frage. Die informationshteoretische Perspektive interessiert sich für das minimale Verhältnis von Signal zu Rauschen, sodass der Graph mit dem auf ihm gespeicherten Informationen genug Informationen enthält, um Rückschlüsse auf die Anfangsbelegung zu ziehen - ungeachtet der Rechenleistung, die dafür notwendig wäre. Die algorithmische Perspektive hingegen beschäftigt sich mit der Frage, ab welchem Verhältnis effiziente Algorithmen existieren, die das Problem lösen. Man mag vermuten, dass die Fähigkeit der Rekonstruktion kontinuierlich mit einem steigenden Verhältnis von Signal zu Rauschen zunimmt. Tatsächlich zeigen viele Inferenzprobleme sogenannte Phasenübergänge, bei denen sogar eine teilweise Approximation der Anfangsbelegung unterhalb

des Überganges nicht möglich ist und bei denen ab dem Übergang eine volle Rekonstruktion erreicht werden kann [73]. Viele Inferenzprobleme scheinen dabei einen *unmöglich-schwer-einfach* Übergang zu durchlaufen. Konkret lässt sich für viele Probleme zeigen, dass unter einer Schwelle m_{inf} Inferenz nicht möglich ist - ungeachtet der Rechenleistung, die man bereit ist auf das Problem zu verwenden. Der Graph zusammen mit den auf ihm gespeicherten Informationen enthält einfach nicht genug Informationen über die Anfangsbelegung. Auf der anderen Seite gibt es eine Schwelle m_{alg}, über der effiziente Algorithmen bekannt sind, die das Problem lösen. Oft stimmen m_{inf} und m_{alg} nicht überein. Der Grund dafür mag natürlich darin liegen, dass ein effizienter Algorithmus einfach noch nicht bekannt ist, der oberhalb der informationstheoretischen, aber unter der bekannten algorithmischen Schranke funktioniert. Allerdings mehren sich in den letzten Jahren die Anzeichen, dass wir es häufig mit einem *schweren* Regime zu tun haben, in dem das Inferenzproblem zwar prinzipiell lösbar ist, aber kein effizienter Algorithmus existiert, der es in diesem Bereich tatsächlich in Polynomialzeit löst [73].

In dieser Dissertation betrachten wir drei klassische Inferenzprobleme und untersuchen deren informationstheoretische und algorithmische Schwellen. Für das erste Problem des *binären Group Testing* zeigen wir in einer Reihe von Arbeiten, dass kein schweres Regime vorliegt und das Problem einen simplen *unmöglich-einfach*-Übergang durchläuft und lösen dabei teils 20 Jahre offene Fragestellungen. Im *quantitativen Group Testing* leiten wir die ausstehende obere informationstheoretische Schranke her und untersuchen einen gängigen Algorithmus, der allerdings deutlich über der informationstheoretischen Schwelle versagt. Vor dem Hintergrund verschiedener anderer Algorithmen, die in der gleichen Größenordnung wie unserer spielen, stellt sich bei diesem Problem die Frage, ob es ein *schweres* Regime gibt. Schlussendlich untersuchen wir den Ising Antiferromagneten - ein klassisches Modell der statistischen Physik. Zum einen zeigen wir für dieses Modell auf regulären zufälligen Graphen den *replica-symmetry breaking (RSB)*-Phasenübergang auf. Zum anderen nutzen wir die inherente Verbindung zum MAXCUT-Problem, einem gängigen Inferenzproblem der Kombinatorik, um eine schärfere obere Schranke für dasselbe herzuleiten und damit eine lang bestehende Vermutung der Physik zu bestätigen. Im folgenden werden wir diese drei Probleme sowie unsere Kernergebnisse beleuchten.

6.2. **Binäres *Group Testing*.** Binäres *Group Testing* ist ein Musterbeispiel für ein statistisches Inferenzproblem auf zufälligen Faktorengraphen. In dem Problem betrachten wir eine große Menge von n Individuen, von denen eine kleine Gruppe k an einer seltenen Krankheit leidet. Wer an der Krankheit leidet wird in einer Belegung $\boldsymbol{\sigma} \in \{0,1\}^n$ mit Hamming-Gewicht k enkodiert. Im Einklang mit der Literatur nehmen wir an, dass $k \sim n^\theta$ skaliert für eine Konstante $\theta \in (0,1)$. Statt jede Person einzeln auf die Krankheit zu testen steht uns ein Testschema zur Verfügung, in dem wir Gruppen von Individuen testen können. Ein Testergebnis ist genau dann positiv, wenn mindestens eine Person in der Gruppe infiziert ist, ansonsten erhalten wir ein negatives Ergebnis. Die wesentliche Erkenntnis hinter *Group Testing* ist, dass wir uns durch geschickte Gruppenbildung viele Tests gegenüber individuellem Testen sparen können. Betrachten wir dazu eine zweistufige Variante, die auf Dorfman [30] zurückgeht, der *Group Testing* erstmalig im zweiten Weltkrieg anwandte, um Soldaten auf Syphilis zu testen. Statt einzeln zu testen schloss Dorfman Gruppen von Individuen zusammen und führte einen Gruppentest durch. Wenn das Ergebnis desselben positiv war, war ein Test vergeudet und jede Person wurde im Anschluss einzeln getestet. Wenn hingegen das Ergebnis negativ war, konnte man sich sicher sein, dass keine infizierte Person in der Gruppe war. Somit genügte ein Test gegenüber einzelnen Tests für jede Person in der Gruppe. Auch wenn Dorfmans *Group Testing*-Variante einen Fortschritt im Vergleich

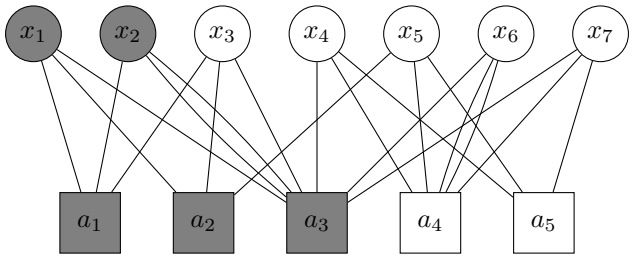

FIGURE 9. Darstellung eines zufälligen regulären Testdesign für nicht-adaptives *Group Testing*, übernommen aus [17]

zu individuellem Testen darstellt, ist es weit entfernt davon, optimal zu sein. Aus dieser Initialzündung entwickelte sich seitdem ein breiter Forschungsstrang, der sich zum Ziel gemacht hat, die informationstheoretischen und algorithmischen Aspekte des *Group Testing*-Problems zu beleuchten. Gerade in den vergangenen zwanzig Jahren gewann das sublineare Regime, das wir hier mit $k \sim n^\theta$ betrachten, an Relevanz.

Zwei Varianten des binären *Group Testing*s seien an dieser Stelle unterschieden. So gibt es adaptive Testdesigns und Algorithmen wie den oben erwähnten zweistufigen Algorithmus von Dorfman, bei dem das Design der zweiten Runde von den Ergebnissen der ersten Runde abhängt. Aus Automatisierungs- und Zeitgründen hat sich die Forschung der vergangenen Jahre jedoch auch für den nicht-adaptiven Fall interessiert, bei dem der Infektionsstatus aller Individuen in einer Stufe ermittelt wird. Zu diesem Zweck eignet sich besonders ein zufälliges reguläres Design, bei dem jedes Individuen nicht nur in einen Gruppentest, sondern eine fixe Anzahl von Tests Δ aufgenommen wird. Dieses Design kann man sich wie in Graphik 9 vorstellen.

Für beide Fälle - den adaptiven und nicht-adaptiven - ergibt ein klassisches Zählargument eine universelle informationstheoretische untere Schranke

$$m_{\text{ad}} = \frac{1}{\log 2} k \log(n/k) \qquad \left(k \sim n^\theta\right)$$

Für das nicht-adaptive reguläre Testdesign, das wir oben kurz vorgestellt haben, ergibt sich eine untere informationstheoretische Schranke der Größenordnung

$$m_{\text{inf}} = \max\left\{\frac{1}{\log 2}, \frac{\theta}{(1-\theta)\log^2 2}\right\} k \log(n/k).$$

In Bezug auf effiziente nicht-adaptive Algorithmen haben sich zwei Algorithmen namens `COMP` and `DD` hervorgetan. Deren Vorgehensweise ist schnell erklärt. Der `COMP`-Algorithmus klassifiziert alle Individuen, die in mindestens einem negativen Test vorkommen, als gesund und alle anderen als krank. Der `DD` Algorithmus geht einen Schritt weiter. Nach der Klassifizierung der Individuen in negativen Tests als gesund sucht er nach positiven Tests, die nunmehr nur noch ein nicht klassifiziertes Individuum enthalten, das nun notwendigerweise infiziert sein muss. Alle Individuen, die nicht in diese Kategorie fallen, werden als gesund eingestuft. Vor unseren Arbeiten am *Group Testing* lag die algorithmische obere Schranke für den `DD` Algorithmus bei

$$m_{\text{alg}} = \max\left\{\frac{1}{\log^2 2}, \frac{\theta}{(1-\theta)\log^2 2}\right\} k \log(n/k).$$

In Bezug auf effiziente adaptive Algorihtmen entwarf Scarlett [65] einen dreistufigen Algorithmus, der die universelle informationstheoretische Schranke m_{ad} erreicht. Diese vorherigen Ergebnisse sind in dem Phasendiagram in Graphik 10 visualisiert.

In unserer ersten Arbeit zum binären *Group Testing*

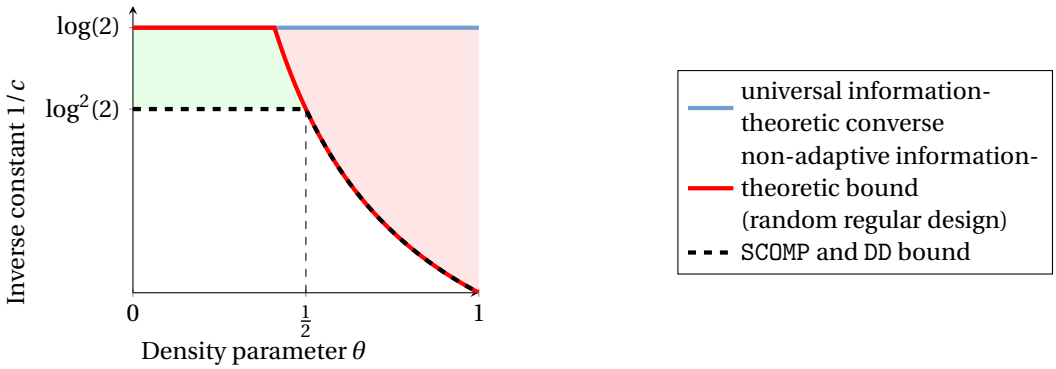

FIGURE 10. Das Phasendiagram für binäres *Group Testing*. Im roten Bereich ist Inferenz informationstheoretisch für das zufällige reguläre Design nicht möglich. Vor unseren Arbeiten war es unklar, ob es ein bessers nicht-adaptives Testdesign gäbe, das in diesem Bereich Inferenz erlauben würde. Der grüne Bereich markiert das Regime, in dem Inferenz informationstheoretisch möglich ist, aber für das kein effizienter Algorithmus bekannt war.

Information-theoretic and algorithmic threshold for group testing

von Amin Coja-Oghlan, Oliver Gebhard, Max Hahn-Klimroth und Philipp Loick konnten wir die voherigen Ergebnisse um die folgendenden Erkenntnisse erweitern.

- Für $m > (1+\varepsilon)m_{\text{inf}}$ ist es in der Tat informationstheoretisch möglich, den Infektionsstatus von jedem Individuum mit einem regulären nicht-adaptiven Testdeisgn mit hoher Wahrscheinlichkeit [5] zu bestimmen. Damit stellen wir eine passende obere Schranke zu der bestehenden unteren informationstheoretischen Schranke auf
- Für $m < (1-\varepsilon)m_{\text{alg}}$ scheitern sowohl der DD als auch ein leicht ausgefeilterer Algorithmus names SCOMP, den Infektionsstatus jedes Individuums mit hoher Wahrscheinlichkeit zu bestimmen

Der Beweis des ersten Ergebnisses basiert auf einer Anwendung von Techniken aus Bedingungserfüllungsprobleme auf das *Group Testing* Problem. Konkret berechnen wir die erwartete Anzahl von alternativen Belegungen $\sigma \in \{0,1\}^n$, die zu den gleichen Testergebnissen führen, und zeigen, dass für $m > (1+\varepsilon)m_{\text{inf}}$ mit hoher Wahrscheinlichkeit keine solche Belegung existiert. Für das zweite Ergebnis zeigen wir, dass für $m < (1-\varepsilon)m_{\text{alg}}$ kein infiziertes Individuum im zweiten Schritt von DD identifiziert wird.

Während diese erste Arbeit bestehende Ergebnisse erweitert, stellt unsere zweite Arbeit

Optimal group testing

von Amin Coja-Oghlan, Oliver Gebhard, Max Hahn-Klimroth und Philipp Loick den großen Wurf für das binäre *Group Testing* dar. Konkret zeigen wir folgende drei Ergebnisse und lösen damit die wesentlichen offenen Punkte des Phasendiagramms.

- Es existiert kein nicht-adaptives Testdesign, sodass das *Group Testing*-Problem für $m < (1-\varepsilon)m_{\text{inf}}$ gelöst werden kann. Die vorherigen Ergebnisse bezogen sich ausschließlich auf den spezifischen Fall des zufälligen regulären Graphens als nicht-adaptives Testdesign. Unser Beitrag ist somit, dass es keine Graphkonstruktion gibt, die besser abschneidet als das zufällige reguläre Testdesign. Unser Ergebnis impliziert damit, dass es im *Group Testing* eine Adaptivitätslücke gibt
- Für $m > (1+\varepsilon)m_{\text{inf}}$ existiert ein nicht-adaptives Testdesign und ein effizienter Algorithmus, der das *Group Testing*-Problem löst. Dieses Ergebnis beantwortet eine der

[5] Eine Folge von Ereignissen $\mathscr{E}_n$ tritt mit hoher Wahrscheinlichkeit ein, wenn $\lim_{n\to\infty}\mathbb{P}[\mathscr{E}_n] = 1$

wesentlichen offenen Fragen im binären *Group Testing*, nämlich die Frage, ob es im nicht-adaptiven Fall ein *schweres* Regime gibt, in dem das *Group Testing*-Problem zwar informationstheoretisch, aber nicht algorithmisch lösbar ist. Der Kerngedanke hinter unserem Testdesign ist die Anwendung der Idee des *Spatial Coupling* aus der Codierungstheorie. Wir verbinden dieses Testdesign mit einem neuen kombinatorischen Algorithmus, der es uns erlaubt, den Infektionsstatus aller Individuen mit hoher Wahrscheinlichkeit zu bestimmen

- Für $m > (1+\varepsilon)m_{\text{ad}}$ gibt es einen zweistufiges Testdesign und einen effizienten Algorithmus, der das *Group Testing*-Problem löst. Zu diesem Zweck verbinden wir Teilergebnisse des vorherigen Resultates und einer geschickt konstruierten zweiten Stufe des *Group Testing*s. Dieses Result zeigt, dass bereits ein zweistufiges (statt dreistufiges) Design genügt, um das *Group Testing*-Problem bis zur universellen informationstheoretischen unteren Schranke zu lösen

Der Beweis der ersten Aussage beruht im wesentlichen auf einer Anwendung der FKG-Ungleichung, um die Wahrscheinlichkeit abzuschätzen, das in einem beliebigen Testdesign ein Individuum nur in Tests mit mindestens einem (weiteren) infizierten Individuum auftritt. Wie oben angedeutet, nutzen wir für den Beweis der zweiten Aussage die Idee des Spatial Coupling, die es uns erlaubt, anhand einer Messgröße die infizierten und gesunden Individuen zu unterscheiden. Der eigentliche Beweis der Aussage fusst auf einer präzisen Abweichungsanalyse der Verteilung dieser Messgröße für infizierte und gesunde Individuen gepaart mit einer Lagrange-Optimierung. Auf der Grundlage dieses Ergebnisses ist es ein leichtes, ein entsprechendes zweistufiges Design zu erstellen, das mit einem geeigneten Algorithmus das *Group Testing*-Problem bis zur universellen informationstheoretischen Schranke löst. Somit zeigen wir, dass es im binären *Group Testing* eine informationstheoretische Adaptivitätslücke gibt, aber sowohl für den adaptiven und nicht-adaptiven Fall effiziente Algorithmen existieren, die bis zur informationstheoretischen unteren Schranke erfoglreich den Infektionsstatus rekonstruieren können.

In einer Erweiterungsarbeit

Improved bounds for noisy group testing with constant tests per item

von Oliver Gebhard, Oliver Johnson, Philipp Loick und Maurice Rolvien betrachten wir den verrauschten Fall, in dem jeder tatsächlich negativer Test mit einer konstanten Wahrscheinlichkeit p als positiv angezeigt wird und ein tatsächlich positiver Test mit Wahrscheinlichkeit q als negativ angezeigt wird. Für dieses Setting analysieren wir eine verrauschte Variante des bekannten `COMP` und `DD` Algorithmus. Im Gegensatz zum nicht verrauschten Fall, in dem ein einzelner negativer Test genügt, um ein Individuum als gesund einzustufen, ist die Situation im verrauschten Fall komplizierter, sodass mehrere negative Tests notwendig sind, um relativ sicher darauf zu schließen, dass ein Individuum gesund ist. Zu diesem Zweck definieren wir eine Schwelle $\alpha\Delta$ für die Anzahl von negativen Tests in der Nachbarschaft eines Individuums, über der ein Individuum als gesund und unter der ein Individuum als infiziert eingestuft wird. Ähnlich verfahren wir mit der zweiten Stufe der `DD` Algorithmus mit einer korrespondierenden Schwelle $\beta\Delta$. Unsere Ergebnisse leiten obere algorithmische Schranken für diese verrauschte Variante des `COMP` und `DD` Algorithmus wie folgt her. Für

$p, q > 0$ und $\alpha \in (q, e^{-d}(1-p) + \left(1-e^{-d}\right) q)$ definieren wir

$$m_{\mathrm{COMP}} = m_{\mathrm{COMP}}(n, \theta, p, q) = \min_{\alpha, d} \max\{b_1(\alpha, d), b_2(\alpha, d)\}\, k \log(n/k)$$

$$\text{mit} \qquad b_1(\alpha, d) = \frac{\theta}{1-\theta} \frac{1}{d D_{\mathrm{KL}}\left(\alpha \| q\right)}$$

$$\text{und} \qquad b_2(\alpha, d) = \frac{1}{1-\theta} \frac{1}{d D_{\mathrm{KL}}\left(\alpha \| e^{-d}(1-p) + \left(1-e^{-d}\right) q\right)}$$

wobei $D_{\mathrm{KL}}\left(\mu \| \nu\right)$ the Kullback-Leibler-Divergenz zwischen μ und ν ist. Als erstes Ergebnis zeigen wir, dass für $m > (1+\varepsilon) m_{\mathrm{COMP}}$ diese verrauschte Version des `COMP` Algorithmus zuverlässig den Infektionsstatus jedes Individuums bestimmen kann. Als zweites Ergebnis leiten wir eine entsprechende Schranke für den `DD` Algorithmus her. Zu diesem Zweck definieren wir

$$m_{\mathrm{DD}} = m_{\mathrm{DD}}(n, \theta, p, q) = \min_{\alpha, \beta, d} \max\left\{c_1(\alpha, d), c_2(\alpha, d), c_3(\beta, d), c_4(\alpha, \beta, d)\right\} k \log(n/k)$$

$$\text{mit} \qquad c_1(\alpha, d) = \frac{\theta}{1-\theta} \frac{1}{d D_{\mathrm{KL}}\left(\alpha \| q\right)}$$

$$\text{und} \qquad c_2(\alpha, d) = \frac{1}{d D_{\mathrm{KL}}\left(\alpha \| 1-w\right)}$$

$$\text{und} \qquad c_3(\beta, d) = \frac{\theta}{1-\theta} \frac{1}{d D_{\mathrm{KL}}\left(\beta \| (1-q) e^{-d}\right)}$$

$$\text{und} \qquad c_4(\alpha, \beta, d) = \max_{1-\alpha \le z \le 1} \left\{ \frac{1}{1-\theta} \frac{1}{d\left(D_{\mathrm{KL}}\left(z \| w\right) + \mathbf{1}\left\{\beta > \frac{z e^{-d} p}{w}\right\} z D_{\mathrm{KL}}\left(\frac{\beta}{z} \| \frac{e^{-d} p}{w}\right)\right)} \right\}.$$

Es stellt sich heraus, dass für $m > (1+\varepsilon) m_{\mathrm{DD}}$ die verrauschte Variante des `DD` Algorithmus zuverlässig den Infektionsstatus jedes Individuums mit hoher Wahrscheinlich ermittelt. Der Beweis der beiden Ergebnisse erfordert die Herleitung der Verteilung der negativ angezeigten Tests für infizierte und gesunde Individuen sowie die Verteilung der positiven Tests, die die Bedingung des zweiten Schrittes von `DD` erfüllen. Nachfolgend nutzen wir die Chernoff-Schranke für die hypergeometrische Verteilung, um die obigen Ergebnisse herzuleiten. Auch wenn die Ergebnisse auf den ersten Blick unhandlich erscheinen, lassen sich für jedes p, q und θ die entsprechende Schranke beliebig genau berechnen.

Die obigen Ergebnisse beziehen sich ausschließlich auf den asymptotischen Fall $n \to \infty$. In der empirischen Arbeit

Efficient and accurate group testing via Belief Propagation: an empirical study

von Amin Coja-Oghlan, Max Hahn-Klimroth, Philipp Loick und Manuel Penschuck betrachten wir praktische Instanzen des *Group Testing* mit 100 bis 10000 Individuen und wenden den Belief Propagation-Algorithmus auf ein reguläres Graphendesign an. Auf dieser Grundlage entwickeln wir effiziente zwei- und dreistufige Algorithmen, die die sich ergebende Posterior-Verteilung der ersten Stufe ausnutzen. Zwar können wir mit diesen Algorithmen nur mäßige Verbesserungen in der Anzahl der Tests im nicht verrauschten Fall erzielen. Allerdings schneiden unsere neuen Algorithmen deutlich besser im verrauschten Fall ab, was die Anzahl der falsche klassifizierten Individuen bei vergleichbaren Tests angeht. Auch können unsere Algorithmen natürlich erweitert werden, wenn man darauf bedacht ist, die Anzahl der falsch klassifizierten Individuen bedeutend zu reduzieren und dafür eine überschaubare Anzahl an zusätzlichen Tests in Kauf nimmt.

6.3. **Quantitatives *Group Testing*.** Als zweites Inferenzproblem betrachten wir quantitatives *Group Testing*, das in der Literatur alternativ auch unter dem Namen Münzwägung und *Pooled Data*-Problem untersucht wird. Der Aufbau des Problems unterscheidet sich zum *binären Group Testing* nur insofern, als dass jedes Testergebnis die (genaue) Anzahl der infizierten Individuuen in einem Test angibt. Wie vorher betrachten wir den sublinearen Fall einer Gruppe von n Individuen, von den $k \sim n^\theta$ für $\theta \in (0,1)$ an einer seltenen Krankheit erkrankt sind. Wieder stellen sich die Fragen der informationstheoretischen und algorithmischen Schranke. Ein analoges, aber ausgefeilteres Zählargument zu oben zeigt, dass mindestens

$$m_{\text{inf}} = 2k\frac{\log(n/k)}{\log(k)} \tag{6.1}$$

Tests informationstheoretisch notwendig sind, um das quantitative *Group Testing*-Problem zu lösen. Interessanterweise liegt die algorithmische Schranke für alle bekannten effizienten Algorithmen in der Größenordnung $\Theta(k\log(n/k))$ und somit einen Faktor $\log(n)$ von der unteren Schranke entfernt. In unserer Arbeit

Quantitative group testing in the sublinear regime

von Oliver Gebhard, Max Hahn-Klimroth, Dominik Kaaser und Philipp Loick tragen wir folgende Ergebnisse zum Verständnis des quantitativen *Group Testing*s bei.

- Die informationstheoretische untere Schranke aus (6.1) is scharf, d.h. für $m > (1+\varepsilon)m_{\text{inf}}$ ist es informationstheoretisch möglich, die Belegung $\boldsymbol{\sigma}$ aus dem Graphen und den Testergebnissen zu rekonstruieren
- Wir untersuchen einen effizienten Algorithmus namens *Maximum Neighbourhood*, der Individuen entsprechend der (angepassten) Summe ihrer Testergebnisse klassifiziert. Zu diesem Zweck definieren wir

  $$m_{MN}(n,\theta) = 4\left(1-\frac{1}{\sqrt{e}}\right)\frac{1+\sqrt{\theta}}{1-\sqrt{\theta}}k\log(n/k).$$

  und zeigen, dass für $m > (1+\varepsilon)m_{MN}$ unser Algorithmus den tatsächlichen Infektionsstatus jedes Individuums mit hoher Wahrscheinlichkeit rekonstruiert.

Der Beweis der ersten Aussage basiert auf der gleichen Methode, die wir schon für das binäre *Group Testing* verwendet haben. Konkret zeigen wir mittels einer Berechnung des ersten Moments und des klassischen Couponsammler-Arguments, dass es für $m > (1+\varepsilon)m_{\text{inf}}$ keine alternative Belegung zur tatsächlichen Belegung gibt, die die gleichen Testergebnisse erzeugt. Der Beweis des effizienten Algorithmus basiert auf der Herleitung der Verteilung der (angepassten) Nachbarschaftssumme für infizierte und gesunde Individuen. Dabei ist die Kernidee, dass infizierte Individuen zu den Testergebnissen in ihrer Umgebung jeweils 1 beisteuern. Somit ist die Erwartung der Nachbarschaftssumme höher für infizierte als für gesunde Individuen. Mit den entsprechenden Verteilungen, der Chernoff-Schranke für die Binomialverteilung und einer austarierten Klassifikationsschwelle sind wir in der Lage zu zeigen, dass wir $\boldsymbol{\sigma}$ rekonstruieren können, wenn wir die k größten Individuen als infiziert einstufen. Obwohl dieses Ergebnis in einer Liga mit bereits bekannten, deutlich komplizierteren Algorithmen spielt, ist es natürlich insofern unzufriedenstellend, als dass wir weiterhin um einen Faktor $\log(n)$ von der informationstheoretischen Schranke entfernt sind. Die Frage liegt daher nahe, ob das quantitative *Group Testing*-Problem anders als sein binäres Pendant tatsächlich einen *unmöglich-schwer-einfach* Übergang durchläuft, der für viele andere Inferenzprobleme vermutet wird.

6.4. **Ising Antiferromagnet und Max Cut.** Als drittes Problem betrachten wir den Ising Antiferromagneten und das eng verwandte Max Cut-Problem auf zufälligen regulären Graphen. Der Ising Antiferromagnet ist ein klassisches Model der statistischen Physik und lässt sich wie folgt beschreiben. Wir betrachten einen Graphen G mit einer Knotenmenge V_n und einer Kantenmenge E. Jeder Knoten kann einen von zwei möglichen Zuständen ± 1 haben. Für eine Zustandsbelegung $\sigma \in \{\pm 1\}^{V_n}$ können wir den sogenannten *Hamiltonian* $\mathscr{H}_G(\sigma)$ definieren als

$$\mathscr{H}_G(\sigma) = \sum_{(v,w)\in E} \frac{1+\sigma_v\sigma_w}{2}.$$

Zusammen mit einem reellen Parameter $\beta > 0$, lässt sich eine Wahrscheinlichkeitsverteilung auf Zustandsbelegungen $\sigma \in \{\pm 1\}^{V_n}$ entsprechend

$$\mu_{G,\beta}(\sigma) = \frac{\exp\left(-\beta\mathscr{H}_G(\sigma)\right)}{Z_{G,\beta}} \qquad \text{mit} \qquad Z_{G,\beta} = \sum_{\tau\in\{\pm 1\}^{V_n}} \exp\left(-\beta\mathscr{H}_G(\tau)\right).$$

definieren. Die Verteilung $\mu_{G,\beta}$ wird als Boltzmann-Verteilung bezeichnet und $Z_{G,\beta}$ als die Zustandssumme. Aus dieser Formulierung wird klar, dass $\mu_{G,\beta}$ Belegungen bevorzugt, bei denen wenige Kanten zwischen Knoten des gleichen Zustandes verlaufen. Dieses Modell ist als der Ising Antiferromagnet bekannt und wirft eine Reihe von interessanten Fragen auf, die wir in den beiden Arbeiten

The Ising antiferromagnet and max cut on random regular graphs

von Amin Coja-Oghlan, Philipp Loick, Balazs Mezei und Gregory Sorkin sowie

The Ising antiferromagnet in the replica symmetric phase

von Christian Fabian und Philipp Loick für den zufälligen d-regulären Graph $\mathbb{G} = \mathbb{G}(n,d)$ beleuchten. Eine Kernfrage des Ising Antiferromagneten betrifft die Korrelation zwischen den Zuständen zweier weit entfernter Knoten bei einem Sample aus der Boltzmann-Verteilung. Vor dem Hintergrund der obigen Erklärung ist es offenkundig, dass die Zustände von nah beieinander liegenden Knoten miteinander korrelieren. Der Grad dieser Korrelationen wird von dem Parameter β gesteuert. Die Frage ist, ob diese Korrelationen auch für zwei Knoten bestehen, die weit voneinander entfernt liegen. Es wird vermutet, dass es einen bestimmten Wert von β gibt, bis zu dem es einen rapiden Korrelationsabfall gibt, sodass die Zustände von zwei weit entfernten Knoten nicht miteinander korrelieren. Dieses Regime wird als replica-symmetrische Phase bezeichnet. Ab diesem Punkt jedoch sollen solche Korrelationen über weite Strecken auftreten - ein Phänomen, das sich *replica-symmetry breaking* (RSB) nennt. Die Schwelle für β wird an der kombinatorisch bedeutsamen Kesten-Stigum-Schranke

$$\beta_{\mathrm{KS}} = \log\left(\frac{\sqrt{d-1}+1}{\sqrt{d-1}-1}\right)$$

vermutet. In unserer Arbeit weisen wir den RSB-Phasenübergang an der Kesten-Stigum-Schranke nach. Der Zustandssumme $Z_{\mathbb{G},\beta}$ kommt dabei eine zentrale Rolle zu. Wir definieren die freie Energie

$$\Phi_d : \beta \in (0,\infty) \to \lim_{n\to\infty} \mathbb{E}\left[\log Z_{\mathbb{G},\beta}\right]/n.$$

Unser erstes Resultat liest sich wie folgt

- Wenn $\beta < \beta_{\mathrm{KS}}$, gilt $\Phi_d(\beta) = \log 2 + \frac{d}{2}\log\frac{1+e^{-\beta}}{2}$
- Wenn $\beta > \beta_{\mathrm{KS}}$, gilt $\Phi_d(\beta) < \log 2 + \frac{d}{2}\log\frac{1+e^{-\beta}}{2}$

Der Beweis der ersten Aussage basiert auf einer Berechnung des ersten und zweiten Moments des Zustandssumme. Wir erweitern diesen klassischen Ansatz um *Spatial Mixing*-Argumente, die den Bogen zu einem Verzweigungsprozess auf einem $(d-1)$-regulären Baum schlagen und uns erlauben die Aussage bis zur Kesten-Stigum-Schranke zu formulieren. Für die zweite Aussage nutzen wir den Zusammenhang zwischen dem Ising Antiferromagneten und dem disassortativen Stochastischem Block Modell. Für letzteres haben wurde in [22] ein Ausdruck für die freie Energie hergeleitet, der ein kompliziertes Optimierungsproblem über Wahrscheinlichkeitsverteilung ist. Wir nutzen diesen Ausdruck, um zu zeigen, dass für $\beta > \beta_{\mathrm{KS}}$ tatsächlich $\Phi_d(\beta) < \lim_{n\to\infty} \mathbb{E}\left[Z_{\mathbb{G},\beta}\right]/n$ gilt.

Durch den Zusammenhang zwischen the Ising Antiferromagneten und dem disassortativen Stochastischen Block Modell impliziert dieses erste Ergebnis auch unmittelbar, dass

- wenn $\beta < \beta_{\mathrm{KS}}(d)$, $\lim_{n\to\infty} D_{\mathrm{KL}}(\mathbb{G}^* \| \mathbb{G})/n = 0$ und $\lim_{n\to\infty} D_{\mathrm{KL}}(\mathbb{G} \| \mathbb{G}^*)/n = 0$ und
- wenn $\beta > \beta_{\mathrm{KS}}(d)$, $\lim_{n\to\infty} D_{\mathrm{KL}}(\mathbb{G}^* \| \mathbb{G})/n > 0$ and $\lim_{n\to\infty} D_{\mathrm{KL}}(\mathbb{G} \| \mathbb{G}^*)/n > 0$.

Zugegebenermaßen ist die Approximation der Zustandssumme in unserem ersten Ergebnis relativ krude mit einem Fehlerterm der Ordnung $\exp(o(n))$. Wenn wir genauer rechnen, stellt sich heraus, dass wir die Grenzverteilung der Zustandssumme im replica-symmetrischen Bereich, also für $\beta < \beta_{\mathrm{KS}}$, bestimmen können. Zu diesem Zweck definieren wird mit $(\Lambda_i)_i$ eine Folge unabhängiger Poisson-Variablen mit $\mathbb{E}[\Lambda_i] = \lambda_i$ für $\lambda_i = \frac{(d-1)^i}{2i}$. Weiterhin sei $\delta_i = \left(\frac{e^{-\beta}-1}{e^{-\beta}+1}\right)^i$. Dann gilt für $n \to \infty$

$$\begin{aligned}&\log\left(Z_{\mathbb{G},\beta}\right) - \frac{1}{2}\log\left(\frac{1+e^{\beta}}{2+de^{\beta}-d}\right) - n\left(\left(1-\frac{d}{2}\right)\log(2) + \frac{d}{2}\log\left(1+e^{-\beta}\right)\right)\\ &\quad\xrightarrow{d} \log(W) \quad \text{where} \quad W := \exp\left(-\lambda_1\delta_1 - \lambda_2\delta_2\right)\prod_{i=3}^{\infty}(1+\delta_i)^{\Lambda_i}\exp\left(-\lambda_i\delta_i\right).\end{aligned}$$

Damit zeigt sich, dass die Varianz in der Zustandssumme im replica-symmetrischen Bereich sich gänzlich den Fluktuationen in der Anzahl von kurzen Kreisen im Graphen $\mathbb{G}$ zuschreiben lässt. Für den Beweis verwenden wir ein Result von Janson [44], das eine Reihe von Bedingungen formuliert, um die Grenzverteilung einer Zufallsvariablen herzuleiten. Zu diesen Bedingungen gehört die Herleitung der Verteilung von kurzen Kreisen im d-regulären Graphen sowie im regulären disassortativen Stochastischen Block Modell. Während ersteres wohl-bekannt ist, ist die Herleitung zweiteres ein wesentlicher Beitrag unserer Arbeit. Ferner müssen wir genauere Ausdrücke für das erste und zweite Moment der Zustandssumme erarbeiten als für die vorherigen Ergebnisse notwendig. Um diese Genauigkeit zu erreichen, verwenden wir analoge *Spatial Mixing*-Argumente zu oben. Die Rückführung der Fluktuationen in der Zustandssumme auf die kurzen Kreise wie hier für den Ising Antiferromagneten durchgeführt ist unter dem Namen *Small Subgraph Conditioning* bekannt.

Neben dem disassortativen Stochastischen Block Modell gibt es auch einen direkten Zusammenhang zwischen dem Ising Antiferromagneten und dem MaxCut-Problem - einem klassischen Inferenzproblem der Kombinatorik. Gegegeben ein Graph G besteht das Ziel darin, die Knotenmenge in zwei Gruppen zu partitionieren, sodass möglichst viele Kanten zwischen den Gruppen verlaufen. Es stellt sich heraus, dass die Größe des MaxCut direkt mit der Zustandssumme zusammenhängt.

$$\frac{2\mathbb{E}\left[\textsc{MaxCut}(\mathbb{G})\right]}{dn} = 1 - \frac{2\mathbb{E}\left[\min_{\sigma\in\{\pm1\}^{V_n}} \mathscr{H}_{\mathbb{G}}(\sigma)\right]}{dn} \leq 1 + \frac{2}{\beta dn}\mathbb{E}\left[\log Z_{\mathbb{G},\beta}\right].$$

Wenn wir den Erwartungswert vom Logarithmus der Zustandssumme - die freie Energie im Physik-Jargon - also nach oben abschätzen können, erhalten wir gleichzeitig eine obere

Schranke an die Größe des MAXCUT. Wir erhalten diese Abschätzung, in dem wir die Interpolationsmethode der mathematischen Physik auf den zufälligen regulären Graphen anwenden. Wir befinden uns dabei in der glücklichen Position, dass die obere Schranke durch eine Verbindung zu einem bestimmten Random Walk explizit gemacht werden kann. So können wir eine explizite obere Schranke des MAXCUT herleiten und damit eine lange bestehende Vermutung von Zdeborová and Boettcher [72] bestätigen.

6.5. **Diskussion.** Die erörterten Probleme stellen klassische Probleme der Mathematik und Informatik dar, die man als Inferenzprobleme auf zufälligen Faktorgraphen auffassen kann. Die Liste solcher Probleme lässt sich dabei beliebig fortsetzen. Auch wenn wir für das binäre *Group Testing* nachweisen konnten, dass die informationstheoretische und algorithmische Schranke zusammenfallen, bleibt für das quantitative *Group Testing* weiterhin ein Faktor $\log(n)$ zwischen beiden Schranken bestehen. Gleiches gilt für eine Reihe anderer Probleme, bei denen informationstheoretische Schranken und die best bekannte algorithmische Schranke um Konstanten entfernt sind oder sogar in verschiedenen Ordnungen liegen. Die Ergründung, ob es sich in diesen Fällen um rechnerisch schwere Regime handelt oder ob doch effiziente Algorithmen existieren, stellt eine herausfordernde, aber wichtige Richtung für zukünftige Forschung dar.

REFERENCES

[1] E. Abbe: Community detection and stochastic block models: recent developments. The Journal of Machine Learning Research **18** (2017) 6446–6531.

[2] M. Aizenman, R. Sims, S. Starr: An extended variational principle for the SK spin-glass model. Phys. Rev. B **68** (2003) 214403.

[3] A. Alaoui, A. Ramdas, F. Krzakala, L. Zdeborová, M. Jordan: Decoding From Pooled Data: Phase Transitions of Message Passing. IEEE Transactions on Information Theory **65** (2019) 572–585.

[4] A. Alaoui, A. Ramdas, F. Krzakala, L. Zdeborová, M. Jordan: Decoding from pooled data: Sharp information-theoretic bounds. SIAM Journal on Mathematics of Data Science **1** (2019) 161–188.

[5] M. Aldridge, L. Baldassini, O. Johnson: Group testing algorithms: bounds and simulations. IEEE Transactions on Information Theory **60** (2014) 3671–3687.

[6] M. Aldridge: Individual testing is optimal for nonadaptive group testing in the linear regime. IEEE Transactions on Information Theory **65** (2019) 2058–2061.

[7] M. Aldridge, O. Johnson, J. Scarlett: Group testing: an information theory perspective. Foundations and Trends in Communications and Information Theory (2019).

[8] J. Barbier, C. Chan, N. Macris: Mutual information for the stochastic block model by the adaptive interpolation method. Proc. IEEE International Symposium on Information Theory (2019) 405—409.

[9] M. Bayati, A. Montanari: The dynamics of message passing on dense graphs, with applications to compressed sensing. IEEE Transactions on Information Theory **57** (2011) 764–785.

[10] T. Berger, V. Levenshtein: Asymptotic efficiency of two-stage disjunctive testing. IEEE Transactions on Information Theory, **48** (2002) 1741–1749.

[11] M. Brennan, G. Bresler: Optimal average-case reductions to sparse pca: From weak assumptions to strong hardness. arXiv preprint arXiv:1902.07380 (2019).

[12] C. Cao, C. Li, X. Sun: Quantitative group testing-based overlapping pool sequencing to identify rare variant carriers. BMC bioinformatics **15** (2014) 1–14.

[13] A. Coja-Oghlan, K. Panagiotou: Going after the k-SAT threshold. In Proc. 45th STOC (2013) 705–714.

[14] A. Coja-Oghlan, F. Krzakala, W. Perkins, L. Zdeborová: Information-theoretic thresholds from the cavity method. Advances in Mathematics **333** (2018) 694–795.

[15] A. Coja-Oghlan, W. Perkins: Spin systems on Bethe lattices. Communications in Mathematical Physics **372** (2019) 441–523.

[16] A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimroth, P. Loick: Information-theoretic and algorithmic thresholds for group testing. Proc. 46th ICALP (2019) #43.

[17] A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimroth, P. Loick: Information-theoretic and algorithmic thresholds for group testing. IEEE Transactions on Information Theory **66** (2020) 7911–7928.

[18] A. Coja-Oghlan, P. Loick, B. Mezei, G. Sorkin: The Ising antiferromagnet and max cut on random regular graphs. arXiv preprint arXiv:2009.10483 (2020).
[19] A. Coja-Oghlan, N. Müller, J. Ravelomanana: Belief Propagation on the random k-SAT model. arXiv preprint arXiv:2011.02303 (2020).
[20] A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimroth, P. Loick. Optimal group testing. Proceedings of Machine Learning Research (COLT) (2020) 1374–1388.
[21] A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimroth, P. Loick: Optimal group testing. Combinatorics, Probability and Computing (2021) 1–38.
[22] A. Coja-Oghlan, M. Hahn-Klimroth, P. Loick, N. Müller, K. Panagiotou, M. Pasch: Inference and mutual information on random factor graphs. Proc. of 38th International Symposium on Theoretical Aspects of Computer Science (2021).
[23] A. Coja-Oghlan, M. Hahn-Klimroth, P. Loick, M. Penschuck: Efficient and accurate group testing via Belief Propagation: an empirical study. arXiv preprint arXiv:2105.07882 (2021).
[24] A. Dembo, A. Montanari: Ising models on locally tree-like graphs. Annals of Applied Probability **20** (2010) 565–592.
[25] A. Dembo, A. Montanari, S. Sen: Extremal cuts of sparse random graphs. Annals of Probability **45** (2017) 1190–1217.
[26] J. Ding, A. Sly, N. Sun: Proof of the satisfiability conjecture for large k. Proc. 47th STOC (2015) 59–68.
[27] A. Djackov: On a search model of false coins. Topics in Information Theory. Hungarian Acad. Sci **16** (1975) 163–170.
[28] D. Donoho, J. Tanner: Thresholds for the recovery of sparse solutions via l1 minimization. In 2006 40th Annual Conference on Information Sciences and Systems IEEE (2006) 202–206.
[29] D. Donoho, A. Javanmard , A. Montanari: Information-theoretically optimal compressed sensing via spatial coupling and approximate message passing. IEEE Transactions on Information Theory **59** (2013) 7434–7464.
[30] R. Dorfman: The detection of defective members of large populations. Annals of Mathematical Statistics **14** (1943) 436–440.
[31] P. Erdős, A. Rényi: On two problems of information theory. Magyar Tud. Akad. Mat. Kutató Int. Közl **8** (1963) 229–243.
[32] C. Fabian, P. Loick: The Ising antiferromagnet in the replica symmetric phase. arXiv preprint arXiv:2103.09775 (2021).
[33] S. Foucart, H. Rauhut: An invitation to compressive sensing. In A mathematical introduction to compressive sensing. Birkhäuser, New York, NY (2013) 1–39.
[34] S. Franz, M. Leone: Replica bounds for optimization problems and diluted spin systems. Journal of Statistical Physics **111** (2003) 535–564.
[35] A. Giurgiu, N. Macris, R. Urbanke: Spatial coupling as a proof technique and three applications. IEEE Transactions on Information Theory **62** (2016) 5281—5295.
[36] M. Jerrum: Large cliques elude the Metropolis process. Random Structures & Algorithms, **3** (1992) 347–359.
[37] D. Gamarnik, I. Zadik: Sparse high-dimensional linear regression. algorithmic barriers and a local search algorithm. arXiv preprint arXiv:1711.04952 (2017).
[38] O. Gebhard, M. Hahn-Klimroth, D. Kaaser, P. Loick: Quantitative group testing in the sublinear regime. arXiv preprint arXiv:1905.01458 (2019).
[39] O. Gebhard, O. Johnson, P. Loick, M. Rolvien: Improved bounds for noisy group testing with constant tests per item. arXiv preprint arXiv:2007.01376 (2020).
[40] O. Gebhard, P. Loick: Note on the offspring distribution for group testing in the linear regime. arXiv preprint arXiv:2103.13039 (2021).
[41] V. Grebinski, G. Kucherov: Optimal reconstruction of graphs under the additive model. Algorithmica **28** (2000) 104–124.
[42] E. Guerra: Broken replica symmetry bounds in the mean field spin glass model. Communications in Mathematical Physics **233** (2003) 1—12.
[43] P. Holland, K. Laskey, S. Leinhardt: Stochastic blockmodels: First steps. Social networks, **5** (1983) 109-137.
[44] S. Janson: Random regular graphs: asymptotic distributions and contiguity. Combinatorics, Probability and Computing **4** (1995) 369–405.
[45] O. Johnson, M. Aldridge, J. Scarlett: Performance of group testing algorithms with near-constant tests per item. IEEE Transactions on Information Theory **65** (2018) 707–723.

[46] E. Karimi, F. Kazemi, A. Heidarzadeh, K. Narayanan, A. Sprintson: Sparse graph codes for non-adaptive quantitative group testing. In 2019 IEEE Information Theory Workshop (ITW) IEEE (2019) 1–5.
[47] E. Karimi, F. Kazemi, A. Heidarzadeh, K. Narayanan, A. Sprintson: Non-adaptive quantitative group testing using irregular sparse graph codes. In 2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton) IEEE (2019) 608–614.
[48] S. Kudekar, H. Pfister: The effect of spatial coupling oncompressive sensing. Proc. 48th Allerton (2010) 347–353.
[49] S. Kudekar, T. Richardson, R. Urbanke: Threshold saturation via spatial coupling: why convolutional LDPC ensembles perform so wellover the BEC. IEEE Transaction on Information Theory **57** (2011) 803—834.
[50] S. Kudekar, T. Richardson, R. Urbanke: Spatially coupled ensembles universally achieve capacity under belief propagation. IEEE Transaction on Information Theory **59** (2013) 7761—7813.
[51] H. Kwang-Ming, D. Ding-Zhu: Pooling designs and nonadaptive group testing: important tools for DNA sequencing. World Scientific (2006).
[52] W. Lenz: Beiträge zum Verständnis der magnetischen Eigenschaften in festen Körpern. Physikalische Zeitschrift **21** (1920) 613–615.
[53] J. Martins, R. Santos, R. Sousa: Testing the maximum by the mean in quantitative group tests. In New Advances in Statistical Modeling and Applications. Springer, Cham (2014) 55–63.
[54] M. Mézard, G. Parisi: The Bethe lattice spin glass revisited. Eur. Phys. J. B **20** (2001) 217—233.
[55] M. Mézard, G. Parisi: The cavity method at zero temperature. Journal of Statistical Physics **111** (2003) 1–34.
[56] M. Mézard, M. Tarzia, C. Toninelli: Group testing with random pools: phase transitions and optimal strategy. Journal of Statistical Physics **131** (2008) 783–801.
[57] M. Mézard, A. Montanari: Information, physics and computation. Oxford University Press 2009.
[58] M. Mézard, C. Toninelli: Group testing with random pools: optimal two-stage algorithms. IEEE Transactions on Information Theory **57** (2011) 1736–1745.
[59] R. Mourad, Z. Dawy, F. Morcos: Designing pooling systems for noisy high-throughput protein-protein interaction experiments using Boolean compressed sensing. IEEE/ACM Transactions on Computational Biology and Bioinformatics **10** (2013) 1478–1490.
[60] E. Mossel, J. Neeman, A. Sly: Reconstruction and estimation in the planted partition model. Probability Theory and Related Fields **162** (2015) 431–461.
[61] H. Ngo, D. Du: A survey on combinatorial group testing algorithms with applications to DNA library screening. Discrete Mathematical Problems with Medical Applications **7** (2000) 171–182.
[62] G. Reeves, J. Xu, I. Zadik: The all-or-nothing phenomenon in sparse linear regression. Proceedings of Machine Learning Research (COLT) (2019) 2652–2663.
[63] D. Panchenko: Spin glass models from the point of view of spin distributions. Annals of Probability **41** (2013) 1315–1361.
[64] J. Scarlett, V. Cevher: Phase Transitions in the Pooled Data Problem. Advances in Neural Information Processing Systems (2017) 376–384.
[65] J. Scarlett: An efficient algorithm for capacity-approaching noisy adaptive group testing. Proc. IEEE International Symposium on Information Theory (2019) 2679–2683.
[66] P. Sham, J. Bader, I. Craig, M. O'Donovan, M. Owen: DNA pooling: a tool for large-scale association studies. Nature Reviews Genetics **3** (2002) 862–871.
[67] H. Shapiro: Problem E 1399. The American Mathematical Monthly **67** (1960) 82.
[68] N. Shental, S. Levy, V. Wuvshet, S. Skorniakov, B. Shalem, A. Ottolenghi, M. Goldhirsh et al: Efficient high-throughput SARS-CoV-2 testing to detect asymptomatic carriers. Science advances **6** (2020) eabc5961.
[69] S. Söderberg, H. Shapiro: A combinatory detection problem. The American Mathematical Monthly **70** (1963) 1066–1070.
[70] M. Talagrand: The Parisi formula. Annals of Mathematics (2006) 221—263.
[71] N. Thierry-Mieg: A new pooling strategy for high-throughput screening: the shifted transversal design. BMC Bioinformatics **7** (2006) 28.
[72] L. Zdeborová, S. Boettcher: A conjecture on the maximum cut and bisection width in random regular graphs. Journal of Statistical Mechanics: Theory and Experiment (2010).
[73] L. Zdeborová, F. Krzakala: Statistical physics of inference: Thresholds and algorithms. Advances in Physics **65** (2016), 453–552.

APPENDIX A. CONTAINED PUBLICATIONS AND THE AUTHOR'S CONTRIBUTIONS

All subsequent publications are included in their current arXiv version. In the following, I will detail the contributions of the author of this dissertation, Philipp Loick, abbreviated by PL. For the sake of brevity, all other authors are also refered to by their initials.

A.1. **Information-theoretic and algorithmic thresholds for group testing.** This work by A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimorth and P. Loick appeared in the journal *IEEE Transactions on Information Theory* [17]. An extended abstract was published in the *Proceedings of the 46th ICALP* [16].

The basic idea of the paper originated during the master thesis of OG (supervised by MHK and PL). The main contribution of the author is the idea and formalisation of Theorem 1.1 jointly with MHK. The main idea of Theorem 1.2 came from the author. The constituting lemmas were split between MHK and PL with PL working particularly on the proof that the number of individuals in V_1^{--} vanishes below m_{alg} and that many healthy individuals have maximum degree.

A.2. **Optimal *Group Testing*.** This work by A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimorth and P. Loick appeared in the journal *Combinatorics, Probability and Computing* [21]. An extended abstract was published in the *Proceedings of the 33rd Conference on Learning Theory (COLT)* [20].

As described above, Theorem 1.1 is based on an adaptation of an argument by Aldridge put forward in earlier work [6]. The idea of this adaptation came from MHK and PL. The formal derivation is due to ACO, MHK and PL. Theorem 1.2 constitutes the joint work of all authors. The particular contribution by PL was the idea behind the combinatorial algorithm and the clean-up step after receiving input from ACO on the notion of spatial coupling. The idea for the adaptive variant of the `SPIV` algorithm is due to PL. The formalisation was done jointly by MHK and PL.

A.3. **Improved bounds for noisy group testing with constant tests per item.** This work by O. Gebhard, O. Johnson, P. Loick and M. Rolvien is under review at the journal *IEEE Transactions on Information Theory*.

The idea behind Theorem 2.1 and 2.2 was developed jointly by OJ, PL and OG. The proof of the theorems including groundwork and the special cases of the Z channel, reverse Z channel and Binary Symmetric Channel was done by PL. The information-theoretic lower bound (Theorem 2.3) was contributed by OJ. MR contributed the simulations. Proposition 2.13 was proved by PL, Proposition 2.15 is due to OG.

A.4. **Efficient and accurate group testing via Belief Propagation: an empirical study.** This work by A. Coja-Oghlan, M. Hahn-Klimroth, P. Loick and M. Penschuck is under review at the journal *SIAM Journal on Mathematics of Data Science*.

The ideas behind the algorithms were jointly developed by ACO, MHK and PL. The source code for generating graph designs and running different algorithms was written by ACO. PL performed the simulations and created the graphics describing the simulation outcome.

A.5. **Quantitative group testing in the sublinear regime.** This work by O. Gebhard, M. Hahn-Klimroth, D. Kaaser and P. Loick is under review at the *International Symposium on Mathematical Foundations of Computer Science (MFCS)*.

The combinatorial ideas underlying the proof of the information-theoretic upper bound and the idea for the `MN` algorithm are due to MHK, PL and DK with support by Amin Coja-Oghlan. These ideas were formalised jointly by OG, PL and MHK. The specific contributions of PL are the derivation of the distribution of the neighbourhood sum and the optimisation

of the threshold for the MN algorithm as well as deriving the information-theoretic bound from an expression simplified by MHK.

A.6. **The Ising antiferromagnet and max cut on random regular graphs.** This work by A. Coja-Oghlan, P. Loick, B. Mezei and G. Sorkin is under review at the *SIAM Journal on Discrete Mathematics*.

Theorem 1.1 is due to ACO and PL with the main ideas coming from ACO. For the first statement, ACO contributed the coupling with a branching process on the tree and PL worked out the first and second moment including the optimisation over the distribution $\mu(\alpha)$. For the second statement, while the idea came from ACO, PL performed the Taylor expansion to the fourth order showing that the Bethe functional is larger than the first moment precisely starting at the Kesten-Stigum bound. Theorem 1.2 is due to ACO. Theorem 1.3 is the joint work of all authors.

A.7. **The Ising antiferromagnet in the replica symmetric phase.** An extended abstract of this work by C. Fabian and P. Loick is accepted for publication in the *Proceedings of the 2021 European Conference on Combinatorics, Graph Theory and Applications (EUROCOMB)*.

The proof idea for the distribution of short cycles in the planted model $\mathbb{G}^*$ is due to PL. It was technically worked out by CF with support by PL. The precise calculation of the first moment is due to CF. The calculation of the second moment in terms of an optimisation problem is mainly due to CF with some input by PL. The idea for solving the optimisation problem in terms of μ and α is mainly due to PL with CF having worked out the technical details.

APPENDIX B. INFORMATION-THEORETIC AND ALGORITHMIC THRESHOLDS FOR GROUP TESTING

INFORMATION-THEORETIC AND ALGORITHMIC THRESHOLDS FOR GROUP TESTING

AMIN COJA-OGHLAN, OLIVER GEBHARD, MAX HAHN-KLIMROTH, PHILIPP LOICK

ABSTRACT. In the group testing problem we aim to identify a small number of infected individuals within a large population. We avail ourselves to a procedure that can test a group of multiple individuals, with the test result coming out positive iff at least one individual in the group is infected. With all tests conducted in parallel, what is the least number of tests required to identify the status of all individuals? In a recent test design [Aldridge et al. 2016] the individuals are assigned to test groups randomly with replacement, with every individual joining an almost equal number of groups. We pinpoint the sharp threshold for the number of tests required in this randomised design so that it is information-theoretically possible to infer the infection status of every individual. Moreover, we analyse two efficient inference algorithms. These results settle conjectures from [Aldridge et al. 2014, Johnson et al. 2019].

1. INTRODUCTION

1.1. **Background and motivation.** The group testing problem goes back to the work of Dorfman from the 1940s [24]. Among a large population a few individuals are infected with a rare disease. The objective is to identify the infected individuals effectively. At our disposal we have a testing procedure capable of not merely testing one individual, but several. The test result will be positive if at least one individual in the test group is infected, and negative otherwise; all tests are conducted in parallel. We are at liberty to assign a single individual to several test groups. The aim is to devise a test design that identifies the status of every single individual correctly while requiring as small a number of tests as possible. A recently proposed test design allocates the individuals to tests randomly [10, 12, 13, 30, 33]. To be precise, given integers $n, m, \Delta > 0$ we create a random bipartite multi-graph by choosing independently for each of the n vertices $x_1, \ldots, x_n$ 'at the top' Δ neighbours among the m vertices $a_1, \ldots, a_m$ 'at the bottom' uniformly at random with replacement. The vertices $x_1, \ldots, x_n$ represent the individuals, the $a_1, \ldots, a_m$ represent the test groups and an individual joins a test group iff the corresponding vertices are adjacent (see Figure 1). The wisdom behind this construction is that the expansion properties of the random bipartite graph precipitate virtuous correlations, facilitating inference. Given n and (an estimate of) the number k of infected individuals, what is the least m for which, with a suitable choice of Δ, the status of every individual can be inferred correctly from the test results with high probability?Like in many other inference problems the answer comes in two instalments. First, we might ask for what m it is *information-theoretically* possible to detect the infected individuals. In other words, regardless of computational resources, do the test results contain enough information in principle to identify the infection status of every individual? Second, for what m does this problem admit *efficient algorithms*? The first main result of this paper resolves the information-theoretic question completely. Specifically, Aldridge, Johnson and Scarlett [13] obtained a function $m_{\mathrm{inf}} = m_{\mathrm{inf}}(n,k)$ such that for any fixed $\varepsilon > 0$ the inference problem is information-theoretically infeasible if $m < (1-\varepsilon)m_{\mathrm{inf}}$. They conjectured that this bound is tight, i.e., that for $m > (1+\varepsilon)m_{\mathrm{inf}}(n,k)$ there is an (exponential) algorithm that correctly identifies the infected individuals with high probability. We prove this conjecture. Furthermore, concerning the algorithmic question, Johnson, Aldridge and Scarlett [30] obtained a function $m_{\mathrm{alg}} = m_{\mathrm{alg}}(n,k)$ that exceeds m_{inf} by a constant factor for small k such that for $m > (1+\varepsilon)m_{\mathrm{alg}}$ certain efficient algorithms successfully identify the infected individuals with high probability. They conjectured that SCOMP, their most sophisticated algorithm, actually succeeds for smaller values of m. We refute this conjecture and show that SCOMP asymptotically fails to outperform a much simpler algorithm called DD. A technical novelty of the present work is that we investigate the group testing problem from a new perspective. While most prior contributions rely either on elementary calculations and/or information-theoretic arguments [12, 13, 30, 39], here we bring to bear techniques from the theory of random constraint satisfaction problems [5, 32].

Supported by DFG CO 646/3 and Stiftung Polytechnische Gesellschaft. An extended abstract of this work appeared in the 2019 ICALP proceedings. A revised version is to appear in IEEE Transactions on Information Theory (Copyright (c) 2017 IEEE DOI: 10.1109/TIT.2020.3023377).

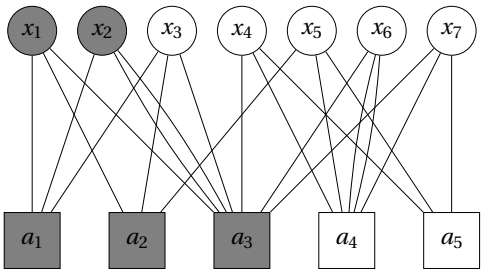

FIGURE 1. The graph illustrates a small example of a group testing instance, with the individuals $x_1,\dots,x_7$ at the top and the tests $a_1,\dots,a_5$ at the bottom. Infected individuals and positive tests are coloured in grey.

Indeed, group testing can be viewed naturally as a constraint satisfaction problem: the tests provide the constraints and the task is to find all possible ways of assigning a status ('infected' or 'not infected') to the n individuals in a way consistent with the given test results. Since the allocation of individuals to tests is random, this question is similar in nature to, e.g., the random k-SAT problem that asks for a Boolean assignment that satisfies a random collection of clauses [4, 6, 20, 23]. It also puts the group testing problem in the same framework as the considerable body of recent work on other inference problems on random graphs such as the stochastic block model (e.g., [1, 18, 22, 35, 37, 43]) or decoding from pooled data [7, 8].

We proceed to state the main results of the paper precisely, followed by a detailed discussion of the prior literature on group testing. The proofs of the information-theoretic and algorithmic bounds follow in 3, Section 4, and 5. The technical details can be found in the appendix.

1.2. **The information-theoretic threshold.** Throughout the paper we labour under the assumptions commonly made in the context of group testing; we will revisit their merit in Section 1.4. Specifically, we assume that the number k of infected individuals satisfies $k \sim n^\theta$ for a fixed $0 < \theta < 1$ [1]. Moreover, let $\boldsymbol{\sigma} \in \{0,1\}^{\{x_1,\dots,x_n\}}$ be a vector of Hamming weight k chosen uniformly at random. The (one-)entries of $\boldsymbol{\sigma}$ indicate which of the n individuals are infected. Moreover, let $\boldsymbol{G} = \boldsymbol{G}(n,m,\Delta)$ signify the aforementioned random bipartite graph with multi-edges. Then $\boldsymbol{\sigma}$ induces a vector $\hat{\boldsymbol{\sigma}} \in \{0,1\}^{\{a_1,\dots,a_m\}}$ that indicates which of the m tests come out positive. To be precise, $\hat{\boldsymbol{\sigma}}_i = 1$ iff test a_i is adjacent to an individual x_j with $\boldsymbol{\sigma}_{x_j} = 1$. For what m is it possible to recover $\boldsymbol{\sigma}$ from $\boldsymbol{G}, \hat{\boldsymbol{\sigma}}$? (Throughout the paper all logarithms are base e.)

Theorem 1.1. *Suppose that $0 < \theta < 1$, $k \sim n^\theta$ and $\varepsilon > 0$ and let*

$$m_{\inf} = m_{\inf}(n,\theta) = \frac{k\log(n/k)}{\min\left\{1, \frac{1-\theta}{\theta}\log 2\right\}\log 2}.$$

(i) If $m > (1+\varepsilon)m_{\inf}(n,\theta)$, then there exists an algorithm that given $\boldsymbol{G}, \hat{\boldsymbol{\sigma}}$ outputs $\boldsymbol{\sigma}$ with high probability.
(ii) If $m < (1-\varepsilon)m_{\inf}(n,\theta)$, then there does not exist any algorithm that given $\boldsymbol{G}, \hat{\boldsymbol{\sigma}}, k$ outputs $\boldsymbol{\sigma}$ with a non-vanishing probability.

Since for $\theta \le \log(2)/(1+\log(2))$ the first part of Theorem 1.1 readily follows from a folklore argument [25], the interesting regime is $\theta > \log(2)/(1+\log(2)) \approx 0.41$. The negative part of Theorem 1.1 strengthens a result from [13], who showed that for $m < (1-\varepsilon)m_{\inf}$ any inference algorithm has a strictly positive error probability. By comparison, Theorem 1.1 shows that any algorithm fails with *high* probability.

But the main contribution of Theorem 1.1 is the first, positive statement. While the problem was solved for $\theta < 1/3$ for a different test design [39, 40] and the case $\theta > 1/2$ is easy because a plain greedy algorithm succeeds [30], the case $1/3 < \theta < 1/2$ proved more challenging. Only heuristic arguments predicting the result of Theorem 1.1 have been put forward for this regime so far [33]. Indeed, Aldridge et al. [12] conjectured that in this case inferring $\boldsymbol{\sigma}$ from $\boldsymbol{G}, \hat{\boldsymbol{\sigma}}$ is equivalent to solving a hypergraph minimum vertex cover problem. The proof of Theorem 1.1 vindicates this conjecture. Specifically, the vertex set of the hypergraph comprises all 'potentially infected' individuals, i.e., those that do not appear in any negative test. The hyperedges are the neighbourhoods

[1] While we write that $k \sim n^\theta$ for the sake of brevity, our results immediately extend to the case $k \sim Cn^\theta$ for some constant C.

∂a_i of the positive tests a_i in $\boldsymbol{G}$. Exhaustive search solves this vertex cover problem in time $\exp(O(n^\theta \log n))$. But how about efficient algorithms for general θ?

1.3. **Efficient algorithms for group testing.** Several polynomial time group testing algorithms have been proposed. A very simple greedy strategy called DD (for 'definitive defectives') first labels all individuals that are members of negative test groups as uninfected. Subsequently it checks for positive tests in which all individuals but one have been identified as uninfected in the first step. Clearly, the single as yet unlabelled individual in such a test group must be infected. Up to this point all decisions made by DD are correct. But in the final step DD marks all as yet unclassified individuals as uninfected, possibly causing false negatives. In fact, the output of DD may be inconsistent with the test results as possibly some positive tests may fail to include an individual classified as 'infected'. While an achievability result is known for the DD algorithm, a corollary of the work in this paper is a matching converse.

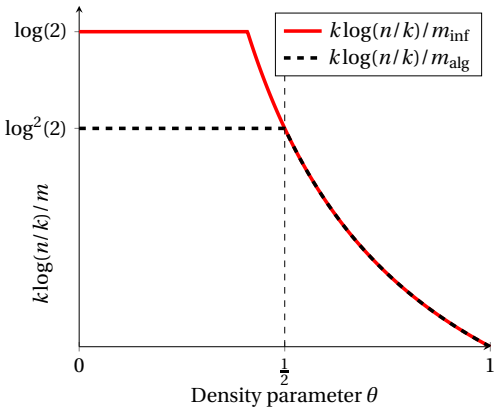

FIGURE 2. The red line shows the information theoretic threshold $m_{\inf}$, the dashed black line signifies the bound m_{alg} which is achieved by the both the SCOMP and the DD algorithm.

The more sophisticated SCOMP algorithm is roughly equivalent to the well-known greedy algorithm for the hypergraph vertex cover problem applied to the hypergraph from the previous paragraph. Specifically, in its first step SCOMP proceeds just like DD, classifying all individuals that occur in negative tests as uninfected. Then SCOMP identifies as infected all unmarked individuals that appear in at least one test whose other participants are already known to be uninfected. Subsequently the algorithm keeps picking an individual that appears in the largest number of as yet 'unexplained' (viz. uncovered) positive tests and marks that individual as infected, with ties broken randomly, until every positive test contains an individual classified as infected. Clearly, SCOMP may produce false positives as well as false negatives. But at least the output is consistent with the test results. Algorithm 1 summarises the procedure of SCOMP.

Input: $\boldsymbol{G}$, $\hat{\boldsymbol{\sigma}}$, k
Output: estimate of $\boldsymbol{\sigma}$
1 Classify all individuals in negative tests as healthy & remove such individuals and tests from $\boldsymbol{G}$;
2 Classify all individuals that appear in at least one positive test as the only yet unclassified individuals as infected & remove such individuals and tests from $\boldsymbol{G}$;
3 **while** *there exists at least one test in* $\boldsymbol{G}$ **do**
4 Classify the individual appearing in the largest number of remaining tests as infected & remove this individual and all adjacent tests from $\boldsymbol{G}$
5 Classify all remaining individuals as healthy;

Algorithm 1: Description of the SCOMP algorithm

Analysing SCOMP has been prominently posed as an open problem in the group testing literature [9, 12, 30]. Indeed, Aldridge et al. [12] opined that "the complicated sequential nature of SCOMP makes it difficult to analyse mathematically". On the positive side, [12] proved that SCOMP succeeds in recovering $\boldsymbol{\sigma}$ correctly given $(\boldsymbol{G}, \hat{\boldsymbol{\sigma}})$ if $m > (1+\varepsilon) m_{\text{alg}}(n,\theta)$ w.h.p.[2], where

$$m_{\text{alg}} = m_{\text{alg}}(n,\theta) = \frac{k \log(n/k)}{\min\left\{1, \frac{1-\theta}{\theta}\right\} \log^2 2}. \tag{1}$$

[2]W.h.p.refers to a probability of $1 - o(1)$ as $n \to \infty$.

However, the algorithm succeeds for a trivial reason; namely, for $m > (1+\varepsilon)m_{\mathrm{alg}}$ even DD suffices to recover $\boldsymbol{\sigma}$ w.h.p. Yet based on experimental evidence [12, 30] conjectured that SCOMP strictly outperforms DD. The following theorem refutes this conjecture.

Theorem 1.2. *Suppose that* $0 < \theta < 1$ *and* $\varepsilon > 0$. *If* $m < (1-\varepsilon)m_{\mathrm{alg}}(n,\theta)$, *then given* $\boldsymbol{G}, \hat{\boldsymbol{\sigma}}$ *w.h.p. both* SCOMP *and* DD *fail to output* $\boldsymbol{\sigma}$.

For $\theta < 1/2$ the information-theoretic bound provided by Theorem 1.1 and the algorithmic bound m_{alg} supplied by Theorem 1.2 remain a modest constant factor apart; see Figure 2. Whether there exists an efficient algorithm for group testing that can close the gap to the information-theoretic bound has long been an open research question. A recent result by Coja-Oghlan et al. [19] shows that such a polynomial-time algorithm indeed exists. The proposed algorithm which is inspired by the notion of spatial coupling from coding theory is able to recover $\boldsymbol{\sigma}$ whenever $m > (1+\varepsilon)m_{\mathrm{inf}}$. Moreover, the authors prove that below the information-theoretic threshold from Theorem 1.1 no non-adaptive algorithm can succeed under any test design (not only the random regular test design considered here) thereby establishing the presence of an adaptivity gap in the group testing problem. An exciting avenue for future research is to investigate the merits of the results and techniques of this paper and [19, 28] for the noisy variant of group testing.

1.4. **Discussion and related work.** Dorfman's original group testing scheme, intended to test the American army for syphilis, was *adaptive*. In a first round of tests each soldier would be allocated to precisely one test group. If the test result came out negative, none of the soldiers in the group were infected. In a second round the soldiers whose group was tested positively would be tested individually. Of course, Dorfman's scheme was not information-theoretically optimal. A first-order optimal adaptive scheme that involves several test stages, with the tests conducted in the present stage governed by the results from the previous stages, is known [15, 25]. In the adaptive scenario the information-theoretic threshold works out to be

$$m_{\mathrm{inf}}^{\mathrm{adapt}}(n,\theta) = \frac{k\log(n/k)}{\log 2}.$$

The lower bound, i.e., that no adaptive design gets by with $(1-\varepsilon)m_{\mathrm{inf}}^{\mathrm{adapt}}(n,\theta)$ tests, follows from a very simple information-theoretic consideration. Namely, with a total of m tests at our disposal there are merely 2^m possible test outcomes, and we need this number to exceed the count $\binom{n}{k}$ of possible vectors $\boldsymbol{\sigma}$, i.e., [14].

More recently there has been a great deal of interest in non-adaptive group testing, where the infection status of each individual is to be determined after just one round of tests [14, 17, 27, 33]. This is the version of the problem that we deal with in the present paper. An important advantage of the non-adaptive scenario is that tests, which may be time-consuming, can be conducted in parallel. Indeed, some of today's most popular applications of group testing are non-adaptive such as DNA screening [17, 31, 38] or protein interaction experiments [36, 42] in computational molecular biology. The randomised test design that we deal with here is the best currently known non-adaptive design (in terms of the number of tests required).

The most interesting regime for the group testing problem is when the number k of infected individuals scales as a power n^θ of the entire population. Mathematically this is because in the linear regime $k = \Omega(n)$ the optimal strategy is to perform n individual tests [11] in order to achieve a vanishing error probability. Similarly, the case of constant k has been solved for some time [41]. Thus, for k linear in n and k constant the theory is already well established. But the sublinear case is also of practical relevance, as witnessed by Heap's law in epidemiology [16] or biological applications [27].

Apart from the randomised test design $\boldsymbol{G}$ where each individual chooses precisely Δ tests (with replacement), the so-called Bernoulli design assigns each individual to every test with a certain probability independently. A considerable amount of attention has been devoted to this model, and its information-theoretic threshold as well as the thresholds for various algorithms have been determined [9, 10, 12, 39]. However, the Bernoulli test design, while easier to analyse, for $\theta > 1/3$ is provably inferior to the test design $\boldsymbol{G}$ that we study here. This is because in the Bernoulli design there are likely quite a few individuals that participate in far fewer tests than expected due to degree fluctuations. We note that our proofs can easily be adapted to reprove the known results for the Bernoulli design. In fact, many technical parts of the proofs become significantly easier and shorter, since we can assume independence between tests, whereas for the constant-column design under consideration here gives rise to subtle dependencies between the tests. A significant portion of the tests is devoted to getting a handle o these dependencies.

1.5. **Notation.** Throughout the paper $\boldsymbol{G} = \boldsymbol{G}(n, m, \Delta)$ denotes the random bipartite graph that describes which individuals take part in which test groups, the vector $\boldsymbol{\sigma} \in \{0,1\}^{\{x_1,\dots,x_n\}}$ encodes which individuals are infected, and $\hat{\boldsymbol{\sigma}} \in \{0,1\}^{\{a_1,\dots,a_m\}}$ indicates the test results. Clearly, $\boldsymbol{G}$ is independent of $\boldsymbol{\sigma}$. Moreover, $k \sim n^\theta$ signifies the number of infected individuals. Additionally, we write

$$V = V_n = \{x_1, \dots, x_n\}, \qquad V_0 = \{x_i \in V : \boldsymbol{\sigma}_{x_i} = 0\} \qquad \text{and} \qquad V_1 = V \setminus V_0$$

for the set of all individuals, the set of uninfected and infected individuals, respectively. For an individual $x \in V$ we write ∂x for the multi-set of tests a_i adjacent to x with $|\partial x| = \Delta$. Analogously, for a test a_i we denote by ∂a_i the multi-set of individuals that take part in the test and $\Gamma_i = |\partial a_i|$. These are multi-sets since individuals are assigned to tests uniformly at random with replacement and therefore $\boldsymbol{G}$ features multi-edges w.h.p.. Let Γ be the vector $(\Gamma_i)_{i\in[m]}$. Furthermore, all asymptotic notation refers to the limit $n \to \infty$. Thus, $o(1)$ denotes a term that vanishes in the limit of large n, while $\omega(1)$ stands for a function that diverges to ∞ as $n \to \infty$. We also let $c, d > 0$ denote reals such that

$$m = ck\log(n/k) \qquad\qquad \Delta = d\log(n/k).$$

Later, we will prove that $c, d = \Theta(1)$ as $n \to \infty$ is optimal for inference. Finally, let $\Gamma_{\min} = \min_{i\in[m]} \Gamma_i$, $\Gamma_{\max} = \max_{i\in[m]} \Gamma_i$. The following sections will outline the proofs of the information-theoretic bounds and the analysis of the SCOMP algorithm and feature the important proofs. The technical details are left to the appendix

2. Getting started

The very first item on the agenda is to get a handle on the posterior distribution of $\boldsymbol{\sigma}$ given $\boldsymbol{G}$ and $\hat{\boldsymbol{\sigma}}$. To this end, let $S_k(\boldsymbol{G}, \hat{\boldsymbol{\sigma}})$ be the set of all vectors $\sigma \in \{0,1\}^V$ of Hamming weight k such that

$$\hat{\boldsymbol{\sigma}}_{a_i} = \mathbf{1}\{\exists x \in \partial a_i : \sigma_x = 1\} \qquad\qquad \text{for all } i \in [m].$$

In words, $S_k(\boldsymbol{G}, \hat{\boldsymbol{\sigma}})$ contains the set of all vectors σ with k ones that label the individuals infected/uninfected in a way consistent with the test results, i.e. that are "satisfying sets" [12, 14]. Let $Z_k(\boldsymbol{G}, \hat{\boldsymbol{\sigma}}) = |S_k(\boldsymbol{G}, \hat{\boldsymbol{\sigma}})|$. The following proposition shows that the posterior of $\boldsymbol{\sigma}$ given $\boldsymbol{G}, \hat{\boldsymbol{\sigma}}$ is uniform on $S_k(\boldsymbol{G}, \hat{\boldsymbol{\sigma}})$.

Proposition 2.1 ([10])**.** *For all $\tau \in \{0,1\}^{\{x_1,\dots,x_n\}}$ we have $\mathbb{P}[\boldsymbol{\sigma} = \tau \mid \boldsymbol{G}, \hat{\boldsymbol{\sigma}}] = \dfrac{\mathbf{1}\{\tau \in S_k(\boldsymbol{G}, \hat{\boldsymbol{\sigma}})\}}{Z_k(\boldsymbol{G}, \hat{\boldsymbol{\sigma}})}$.*

Adopting the jargon of the recent literature on inference problems on random graphs, we refer to Proposition 2.1 as the *Nishimori identity* [18, 43]. The proposition shows that apart from the actual test results, there is no further 'hidden information' about $\boldsymbol{\sigma}$ encoded in $\boldsymbol{G}, \hat{\boldsymbol{\sigma}}$. In particular, the information-theoretically optimal inference algorithm just outputs a uniform sample from $S_k(\boldsymbol{G}, \hat{\boldsymbol{\sigma}})$. In effect, we obtain the following.

Corollary 2.2. *(1) If $Z_k(\boldsymbol{G}, \hat{\boldsymbol{\sigma}}) = \omega(1)$ w.h.p., then for any algorithm $\mathscr{A}$ we have*

$$\mathbb{P}[\mathscr{A}(\boldsymbol{G}, \hat{\boldsymbol{\sigma}}, k) = \boldsymbol{\sigma}] = o(1).$$

(2) If $Z_k(\boldsymbol{G}, \hat{\boldsymbol{\sigma}}) = 1$ w.h.p., then there is an algorithm $\mathscr{A}$ such that

$$\mathbb{P}[\mathscr{A}(\boldsymbol{G}, \hat{\boldsymbol{\sigma}}, k) = \boldsymbol{\sigma}] = 1 - o(1).$$

Both the positive and the negative part of Corollary 2.2 assume that the precise number k of infected individuals is known to the algorithm. This assumption makes the negative part stronger, but weakens the positive part. Yet we will see in due course how in the positive scenario the assumption that k be known can be removed.

For the information-theoretic bound, the proof hinges on analysing the number of individuals that can be flipped without affecting the test results. We encounter two kinds of such individuals. The first kind consists of healthy individuals that only appear in positive tests and which we will denote by V_0^+. In symbols,

$$V_0^+ = \{x_i \in V_0 : \forall a \in \partial x_i \exists y \in \partial a : \boldsymbol{\sigma}_y = 1\}. \tag{2}$$

Similarly, let V_1^+ be the set of all infected individuals x_i such that every test in which x_i occurs features another infected individual; in symbols,

$$V_1^+ = \{x_i \in V_1 : \forall a \in \partial x_i \exists y \in \partial a \setminus \{x_i\} : \boldsymbol{\sigma}_y = 1\}.$$

5

We think of the individuals in V_0^+ as the 'potential false positives'. Indeed, if for any $x_i \in V_0^+$ we obtain $\boldsymbol{\sigma}'$ from $\boldsymbol{\sigma}$ by setting x_i to one, then $\boldsymbol{\sigma}'$ will render the same test results as $\boldsymbol{\sigma}$. Similarly, the individuals in V_1^+ are potential false negatives. For completeness, we also define V_0^- and V_1^- as

$$V_0^- = V_0 \setminus V_0^+ \qquad \text{and} \qquad V_1^- = V_1 \setminus V_1^+ \tag{3}$$

In the following, let us get a handle on the size of sets V_0^+ and V_1^+. Specifically, we prove the following five statements.

Proposition 2.3. *Let $c, d = \Theta(1)$. Then, the following statements hold w.h.p.*

(1) $\left|V_0^+\right| = (1 + n^{-\Omega(1)}) n \left(1 - \exp(-d/c)\right)^{\Delta}$.
(2) If $k(1-\exp(-d/c))^{\Delta} \geq n^{\Omega(1)}$, *then* $\left|V_1^+\right| = n^{\Omega(1)}$.
(3) If $k(1-\exp(-d/c))^{\Delta} = o(1)$, *then* $\left|V_1^+\right| = o(1)$.
(4) If $c < \frac{\theta}{1-\theta}\frac{1}{\log^2 2}$, *then* $\left|V_1^+\right|, \left|V_0^+\right| = n^{\Omega(1)}$.
(5) If $c > \frac{\theta}{1-\theta}\frac{1}{\log^2 2}$, *then* $\left|V_1^+\right| = o(1)$.

The proof of Proposition 2.3, while not fundamentally difficult, requires a bit of care because we are dealing with a random bipartite multi-graph whose (test-)degrees scale as a power of n. In effect, the diameter of the bipartite graph is quite small and the neighbourhoods of different tests may have a sizeable intersection. The technical workout follows in Section B.6. In the next step, let us get a handle on the size of the test degrees.

Lemma 2.4. *With probability at least $1 - o(n^{-2})$ we have*

$$\Delta n/m - \sqrt{\Delta n/m}\log n \leq \Gamma_{\min} \leq \Gamma_{\max} \leq \Delta n/m + \sqrt{\Delta n/m}\log n.$$

The proof of this and the subsequent elementary lemmas are included in Section B. Next, we calculate the number of positive and negative tests. Let $\boldsymbol{m}_1$ be the number of positive tests and let $\boldsymbol{m}_0$ be the number of negative tests. Clearly $\boldsymbol{m}_0 + \boldsymbol{m}_1 = m$.

Lemma 2.5. *With probability at least $1 - o(n^{-2})$ we have*

$$\boldsymbol{m}_0 = \exp(-d/c)m + O(\sqrt{m}\log^2 n).$$

Finally, we justify that setting $c, d = \Theta(1)$ as $n \to \infty$ is optimal for inference. The fact that $c = \Theta(1)$ immediately follows from the information-theoretic counting bound, i.e., [14].

Lemma 2.6. *(1) If $\Delta = o(\log(n/k))$ and $m = \Theta(k\log(n/k))$, then $Z_k(\boldsymbol{G}, \hat{\boldsymbol{\sigma}}) = \omega(1)$ w.h.p.*
(2) If $\Delta = \omega(\log(n/k))$ and $m = \Theta(k\log(n/k))$, then $Z_k(\boldsymbol{G}, \hat{\boldsymbol{\sigma}}) = \omega(1)$ w.h.p.

3. The information-theoretic upper bound

We proceed to discuss the proof of Theorem 1.1. The proof of the first, positive statement and of the second, negative statement hinge on two separate arguments. We begin with the proof of the information-theoretic upper bound which is the principal achievement of the present work. The proof rests upon techniques that have come to play an important role in the theory of random constraint satisfaction problems. Specifically, we need to show that $Z_k(\boldsymbol{G}, \hat{\boldsymbol{\sigma}}) = 1$ w.h.p., i.e., that $\boldsymbol{\sigma}$ is the only assignment compatible with the test results w.h.p. We establish this result by combining two separate arguments. First, we use a moment calculation to show that w.h.p. there are no other solutions that have a small 'overlap' with $\boldsymbol{\sigma}$. Then we use an expansion argument to show that w.h.p. there are no alternative solutions with a big overlap. Both these arguments are variants of the arguments that have been used to study the solution space geometry of random constraint satisfaction problems such as random k-SAT or random k-XORSAT [3, 4, 26], as well as the freezing thresholds of random constraint satisfaction problems [2, 34]. Yet to our knowledge these methods have thus far not been applied to the group testing problem. In this section we choose $\Delta = \lceil \frac{m}{k}\log 2 \rceil$ which maximises the entropy of the test results. Formally, we define

$$Z_{k,\ell}(\boldsymbol{G}, \hat{\boldsymbol{\sigma}}) = |\{\sigma \in S_k(\boldsymbol{G}, \hat{\boldsymbol{\sigma}}) : \langle \boldsymbol{\sigma}, \sigma \rangle = \ell\}|$$

as the number of assignments $\sigma \in S_k(\boldsymbol{G}, \hat{\boldsymbol{\sigma}})$ different from the true configuration $\boldsymbol{\sigma}$ whose *overlap*

$$\langle \boldsymbol{\sigma}, \sigma \rangle = \sum_{i=1}^{n} \mathbf{1}\{\sigma_{x_i} = \boldsymbol{\sigma}_{x_i} = 1\}$$

with $\boldsymbol{\sigma}$ is equal to ℓ. The following two propositions rule out assignments with a small and a big overlap, respectively. In either case we choose $\Delta = \lceil \frac{m}{k} \log 2 \rceil$ to take its optimal value.

Proposition 3.1. *Let $\varepsilon > 0$ and $0 < \theta < 1$ and assume that $m > (1+\varepsilon) m_{\inf}(k, \theta)$. W.h.p. we have $Z_{k,\ell}(\boldsymbol{G}, \hat{\boldsymbol{\sigma}}) = 0$ for all $\ell < (1 - 1/\log n)k$.*

Proof. For $i \in [m]$ let Γ_i be the degree of a_i in $\boldsymbol{G}$, i.e., the number of edges incident with a_i; this number may exceed the number of different individuals that participate in test a_i as $\boldsymbol{G}$ may feature multi-edges. Let Γ be the σ-algebra generated by the random variables $(\Gamma_i)_{i \in [m]}$. Whenever we condition on Γ, we assume that the bounds from Lemma 2.4 and 2.5 hold. Given Γ we can generate $\boldsymbol{G}$ from the well-known *pairing model* [29]. Specifically, we create a set $\{x_i\} \times [\Delta]$ of Δ clones of each individual as well as sets $\{a_i\} \times [\Gamma_i]$ of clones of the tests. Then we draw a perfect matching of the complete bipartite graph on the vertex sets $\bigcup_{i=1}^{n} \{x_i\} \times [\Delta], \bigcup_{i=1}^{m} \{a_i\} \times [\Gamma_i]$ uniformly at random. For each matching edge linking a clone of x_i with a clone of a_j we insert an i-j-edge. The resulting bipartite random multi-graph has the same distribution as $\boldsymbol{G}$ given Γ. As an application of this observation we obtain for every integer $0 \le \ell < k$

$$\mathbb{E}[Z_{k,\ell}(\boldsymbol{G}, \hat{\boldsymbol{\sigma}}) \mid \Gamma] \le O\big((\Delta k)^{3/2}\big) \cdot \binom{k}{\ell}\binom{n-k}{k-\ell} \prod_{i=1}^{m} \big(1 - 2(1-k/n)^{\Gamma_i} + 2\,(1 - 2k/n + \ell/n)^{\Gamma_i}\big) \tag{4}$$

To see why (4) holds we use the linearity of expectation. The product of the two binomial coefficients simply accounts for the number of assignments σ that have overlap ℓ with $\boldsymbol{\sigma}$. Hence, with $\mathscr{S}$ the event that one specific $\sigma \in \{0,1\}^V$ that has overlap ℓ with $\boldsymbol{\sigma}$ belongs to $S_{k,\ell}(\boldsymbol{G}, \hat{\boldsymbol{\sigma}})$, we need to show that

$$\mathbb{P}[\mathscr{S} \mid \Gamma] \le O\big((\Delta k)^{3/2}\big) \prod_{i=1}^{m} 1 - 2(1-k/n)^{\Gamma_i} + 2\,(1 - 2k/n + \ell/n)^{\Gamma_i}. \tag{5}$$

By symmetry we may assume that $\boldsymbol{\sigma}_{x_i} = \mathbf{1}\{i \le k\}$ and that $\sigma_{x_i} = \mathbf{1}\{i \le \ell\} + \mathbf{1}\{k < i \le 2k - \ell\}$.

To establish (5) we harness the pairing model. Namely, given Γ we can think of each test a_i as a bin of capacity Γ_i. Moreover, we think of each clone (x_i, h), $h \in [\Delta]$, of an individual as a ball. The ball is labelled $(\boldsymbol{\sigma}_{x_i}, \sigma_{x_i}) \in \{0,1\}^2$. The random matching that creates $\boldsymbol{G}$ effectively tosses the Δn balls randomly into the bins. Hence, for $i \in [m]$ and for $j \in [\Gamma_i]$ let us write $\boldsymbol{A}_{i,j} = (\boldsymbol{A}_{i,j,1}, \boldsymbol{A}_{i,j,2}) \in \{0,1\}^2$ for the label of the jth ball that ends up in bin number i. Then we are left to calculate the probability that for every test a_i either $\boldsymbol{A}_{i,j,1} = \boldsymbol{A}_{i,j,2} = 0$ for every $j \in [\Gamma_i]$ or there is at least one pair $(j,k) \in [\Gamma_i]^2$ such that $\boldsymbol{A}_{i,j,1} = \boldsymbol{A}_{i,k,2} = 1$

$$\mathbb{P}[\mathscr{S} \mid \Gamma] = \mathbb{P}\left[\forall i \in [m] : \max_{j \in [\Gamma_i]} \boldsymbol{A}_{i,j,1} = \max_{j \in [\Gamma_i]} \boldsymbol{A}_{i,j,2} \,\middle|\, \Gamma\right], \tag{6}$$

To calculate this probability we borrow a trick from the analysis of the random k-SAT model [20]. Namely, we consider a new set of $\{0,1\}^2$-valued random variables $\boldsymbol{A}'_{i,j} = (\boldsymbol{A}'_{i,j,1}, \boldsymbol{A}'_{i,j,2})$ such that $(\boldsymbol{A}'_{i,j})_{i \in [m], j \in [\Gamma_i]}$ are mutually independent and such that

$$\mathbb{P}\left[\boldsymbol{A}'_{i,j} = (1,1)\right] = \ell/n, \qquad \mathbb{P}\left[\boldsymbol{A}'_{i,j} = (0,1)\right] = \mathbb{P}\left[\boldsymbol{A}'_{i,j} = (1,0)\right] = (k-\ell)/n,$$
$$\mathbb{P}\left[\boldsymbol{A}'_{i,j} = (0,0)\right] = (n - 2k + \ell)/n$$

for all i, j. Due to their independence, these multinomially distributed random variables are much easier to handle than $\boldsymbol{A}_{i,j}$. It will turn out, that given a (not too unlikely) event, it suffices to analyse these independent variables instead of $\boldsymbol{A}_{i,j}$. Now, let $\mathscr{T}$ be the event that

$$\sum_{i=1}^{m}\sum_{j=1}^{\Gamma_i} \mathbf{1}\left\{\boldsymbol{A}'_{i,j} = (1,1)\right\} = \ell\Delta, \qquad \sum_{i=1}^{m}\sum_{j=1}^{\Gamma_i} \mathbf{1}\left\{\boldsymbol{A}'_{i,j} = (0,0)\right\} = (n - 2k + \ell)\Delta, \tag{7}$$

$$\sum_{i=1}^{m}\sum_{j=1}^{\Gamma_i} \mathbf{1}\left\{\boldsymbol{A}'_{i,j} = (1,0)\right\} = \sum_{i=1}^{m}\sum_{j=1}^{\Gamma_i} \mathbf{1}\left\{\boldsymbol{A}'_{i,j} = (0,1)\right\} = (k-\ell)\Delta, \tag{8}$$

i..e, that all of the sums on the l.h.s. are *precisely* equal to their expected values. Then $\boldsymbol{A}' = (\boldsymbol{A}'_{i,j})_{i,j}$ given $\mathscr{T}$ is distributed precisely as $\boldsymbol{A} = (\boldsymbol{A}_{i,j})_{i,j}$. Hence, (6) yields

$$\mathbb{P}[\mathscr{S} \mid \Gamma] = \mathbb{P}\left[\forall i \in [m] : \max_{j \in [\Gamma_i]} \boldsymbol{A}'_{i,j,1} = \max_{j \in [\Gamma_i]} \boldsymbol{A}'_{i,j,2} \mid \Gamma, \mathscr{T}\right]. \tag{9}$$

Thus, let

$$\mathscr{A} = \left\{ \forall i \in [m] : \max_{j \in [\Gamma_i]} \boldsymbol{A}'_{i,j,1} = \max_{j \in [\Gamma_i]} \boldsymbol{A}'_{i,j,2} \right\}.$$

The grand idea is now to calculate the probability $\mathbb{P}[\mathscr{A} \mid \Gamma]$. Subsequently, we employ Bayes' Theorem to derive a bound for the conditional probability $\mathbb{P}[\mathscr{A} \mid \mathscr{T}, \Gamma]$ for which we know by the above application of the balls-into-bins principle

$$\mathbb{P}[\mathscr{S} \mid \Gamma] = \mathbb{P}[\mathscr{A} \mid \mathscr{T}, \Gamma].$$

Because the $(\boldsymbol{A}'_{i,j})_{i,j}$ are mutually independent, we can easily compute the unconditional probability $\mathbb{P}[\mathscr{A} \mid \Gamma]$: by inclusion/exclusion,

$$\mathbb{P}[\mathscr{A} \mid \Gamma] = \prod_{i=1}^{m} \left(1 - 2(1-k/n)^{\Gamma_i} + 2(1-2k/n+\ell/n)^{\Gamma_i}\right) \tag{10}$$

(the probability that $\max_{j \in [\Gamma_i]} \boldsymbol{A}'_{i,j,1} = \max_{j \in [\Gamma_i]} \boldsymbol{A}'_{i,j,2} = 1$, i.e., both tests positive, equals one minus the probability that $\max_{j \in [\Gamma_i]} \boldsymbol{A}'_{i,j,1} = 0$ minus the probability that $\max_{j \in [\Gamma_i]} \boldsymbol{A}'_{i,j,2} = 0$ plus the probability that $\max_{j \in [\Gamma_i]} \boldsymbol{A}'_{i,j,1} = \max_{j \in [\Gamma_i]} \boldsymbol{A}'_{i,j,2} = 0$; then add the probability that $\max_{j \in [\Gamma_i]} \boldsymbol{A}'_{i,j,1} = \max_{j \in [\Gamma_i]} \boldsymbol{A}'_{i,j,2} = 0$, i.e., both tests negative).

Finally, to deal with the conditioning we use Bayes' rule:

$$\mathbb{P}[\mathscr{A} \mid \mathscr{T}, \Gamma] = \frac{\mathbb{P}[\mathscr{A} \mid \Gamma]\,\mathbb{P}[\mathscr{T} \mid \mathscr{A}, \Gamma]}{\mathbb{P}[\mathscr{T} \mid \Gamma]}. \tag{11}$$

Since the $(\boldsymbol{A}'_{i,j})_{i,j}$ are independent, Stirling's formula yields

$$\mathbb{P}[\mathscr{T} \mid \Gamma] = \Omega\left((\Delta k)^{-3/2}\right).$$

A short justification can be found in Section B.1. Moreover, by definition we have $\mathbb{P}[\mathscr{T} \mid \mathscr{A}, \Gamma] \leq 1$. Hence, (5) follows from (9)–(11). To complete the proof of the proposition, we claim that

$$\sum_{0 \leq \ell \leq \lceil (1-1/\log n)k \rceil} O\left((\Delta k)^{3/2}\right) \binom{k}{\ell} \binom{n-k}{k-\ell} \prod_{i=1}^{m} \left(1 - 2(1-k/n)^{\Gamma_i} + 2(1-2k/n+\ell/n)^{\Gamma_i}\right) = o(1). \tag{12}$$

To prove Equation (12), let $\alpha = \ell/k$. Using Lemma 2.4 and recalling $m = ck\log(n/k)$ and $\Delta = d\log(n/k)$, we find

$$\begin{aligned}
\mathbb{E}[Z_{k,l}(\boldsymbol{G},\boldsymbol{\sigma})] &\leq O\left((\Delta k)^{3/2}\right) \binom{k}{(1-\alpha)k} \binom{n-k}{(1-\alpha)k} \prod_{i=1}^{m} \left(1 - 2\left(1-\frac{k}{n}\right)^{\Gamma_i} + 2\left(1-\frac{2k}{n}+\frac{\alpha k}{n}\right)^{\Gamma_i}\right) \\
&\leq O\left((\Delta k)^{3/2}\right) \left(\frac{e}{(1-\alpha)} \frac{en}{(1-\alpha)k}\right)^{(1-\alpha)k} \left(1 - 2\left(1-\frac{k}{n}\right)^{\Gamma_{\max}} + 2\left(1-\frac{2k}{n}+\frac{\alpha k}{n}\right)^{\Gamma_{\min}}\right)^{m} \\
&\leq O\left((\Delta k)^{3/2}\right) \left(\frac{e}{(1-\alpha)} \frac{en}{(1-\alpha)k}\right)^{(1-\alpha)k} \left(1 - 2\left(1-\frac{k}{n}\right)^{\frac{n\log 2}{k}\left(1+n^{-\Omega(1)}\right)}\right. \\
&\qquad \left. + 2\left(1-\frac{2k}{n}+\frac{\alpha k}{n}\right)^{\frac{n\log 2}{k}\left(1+n^{-\Omega(1)}\right)}\right)^{m} \qquad (13) \\
&\leq O\left((\Delta k)^{3/2}\right) \left(\frac{e}{(1-\alpha)} \frac{en}{(1-\alpha)k}\right)^{(1-\alpha)k} \left(1 - \left(1-2^{-(1-\alpha)}\right)\exp\left(n^{-\Omega(1)}\right)\right)^{m} \\
&= O\left((\Delta k)^{3/2}\right) \left(\frac{e}{(1-\alpha)} \frac{en}{(1-\alpha)k} (k/n)^{c\log(2)+n^{-\Omega(1)}} (1+o(1))\right)^{(1-\alpha)k} \\
&= O\left((\Delta k)^{3/2}\right) \left(\frac{e^2 (k/n)^{c\log(2)-1+n^{-\Omega(1)}}}{(1-\alpha)^2}\right)^{(1-\alpha)k}. \qquad (14)
\end{aligned}$$

By the definition of $m > (1+\varepsilon)m_{\inf}$ and $\ell < \lceil k(1-\log^{-1} n) \rceil$, we have

$$c\log 2 = 1+\varepsilon \qquad \text{and} \qquad (1-\alpha)^2 \geq 1/\log^2 n \tag{15}$$

Moreover, as $\ell < \lceil k(1-\log^{-1} n) \rceil$ we have $(1-\alpha)k = \omega(1)$. Thus (15) implies that (14) tends to zero with $n \to \infty$. Therefore, the proposition follows from Equations (14), (15) and Markov's inequality.

□

The argument from Proposition 3.1 does not extend to large overlaps (close to k) because the expression on the r.h.s. of (4) gets too large. In other words, merely computing the expected number of solutions with a given overlap does not do the trick. This 'lottery phenomenon' is ubiquitous in random constraint satisfaction problems: for big overlap values rare solution-rich instances drive up the expected number of solutions [4, 5]. Fortunately, we can find a remedy.

Proposition 3.2. *Let $\varepsilon > 0$ and $0 < \theta < 1$ and assume that $m > (1+\varepsilon)m_{\inf}(k,\theta)$. W.h.p. we have $Z_{k,\ell}(\boldsymbol{G},\hat{\boldsymbol{\sigma}}) = 0$ for all $(1-1/\log n)k \le \ell < k$.*

In order to cope with this issue we take another leaf out of the random CSP literature [2, 34]. Namely, we show that the solution $\boldsymbol{\sigma}$ is locally rigid. That is, the expansion properties of the random bipartite graph $\boldsymbol{G}$ preclude the existence of other solutions that have a big overlap with $\boldsymbol{\sigma}$. The following lemma holds the key to this effect.

Lemma 3.3. *For any $\varepsilon > 0$ there exists $\delta = \delta(\varepsilon) > 0$ such that for all $m > (1+\varepsilon)m_{\inf}$ the following is true. Let $\mathscr{R}$ be the event that for every x_i with $\boldsymbol{\sigma}_{x_i} = 1$ there are at least $\delta\Delta$ tests $a \in \partial x_i$ such that $\partial a \setminus \{x_i\} \subseteq V_0$. Then $\mathbb{P}[\mathscr{R}] = 1 - o(1)$.*

Proof. Let $(\boldsymbol{X}_i)_{i\in[m]}$ be a sequence of independent $\mathrm{Bin}(\Gamma_i, k/n)$-variables as in Section 2. Also let $W = \sum_{i=1}^{m} \mathbf{1}\{\boldsymbol{Y}_i = 1\}$ as in Section 2. Proceeding along the lines of the proof of Lemma 2.3 (see (35) in Section B.6), we obtain

$$\mathbb{P}\left[W = \left(1 + n^{-\Omega(1)}\right) k\Delta/2 \mid \Gamma\right] = 1 - o(n^{-7}). \tag{16}$$

Let T be the number of infected individuals which only show up less than $\delta\Delta$ of their tests as the only infected individual, i.e.

$$T = \left| x \in V_1 : \sum_{a \in \partial x} \mathbf{1}\{\partial a \setminus \{x\} \subseteq V_0\} < \delta\Delta \right|.$$

Moreover, let $\boldsymbol{H} = \boldsymbol{H}(N, K, n')$ be a hypergeometric random variable with parameters $N = k\Delta$ (total eligible assignments for infected individuals), $K = W$ (tests with only one infected individual) and $n' = \Delta$ (number of tests per individuals). Then the union bound over k infected individuals yields

$$\mathbb{E}[T \mid \Gamma, W] \le k\mathbb{P}[\boldsymbol{H} < \delta\Delta]. \tag{17}$$

Further, the Chernoff bound for the hypergeometric distribution implies

$$\mathbb{P}[\boldsymbol{H} < \delta\Delta] \le \exp(-\Delta D_{\mathrm{KL}}(\delta \| W/(k\Delta))) \tag{18}$$

Recall $\Delta = d\log(n/k)$. Since $D_{\mathrm{KL}}(\delta\|1/2 + o(1)) = \delta\log\delta + (1-\delta)\log(1-\delta) + \log 2 + o(1)$ and $\delta\log\delta + (1-\delta)\log(1-\delta) \nearrow 0$ as $\delta \to 0$ and $c > \frac{\theta}{(1-\theta)\log^2 2}$, we can choose $\delta > 0$ small enough so that

$$\Delta(\delta\log\delta + (1-\delta)\log(1-\delta) + \log 2 + o(1)) > \log k \tag{19}$$

Finally, the assertion follows from (16)–(19). □

Hence, w.h.p. any infected individual appears in plenty of tests where all the other individuals are uninfected. This property causes $\boldsymbol{\sigma}$ to be locally rigid. To see why, consider the repercussions of just changing the status of a single individual x_i from infected to uninfected. Because given $\mathscr{R}$ the individual x_i appears as the only infected individual in at least $\delta\Delta$ tests, in order to maintain the same tests results we will also need to flip at least one individual in each of these tests from 'uninfected' to 'infected'. Since tests typically have relatively few individuals in common, the necessary number of flips from 0 to 1 will be $\Omega(\Delta) = \Omega(\log n)$. But then in order to keep the total number of infected individuals constant k, we will need to perform another $\Omega(\Delta)$ flips from 1 to 0. Yet given $\mathscr{R}$ each of these 'second generation' individuals that we flip from infected to uninfected is itself the only infected individual in many tests. Thus, the single flip that we started from triggers a veritable avalanche of flips, which will stop only after the overlap has dropped significantly. The next lemma formalises this intuition. The lemma shows that while the unconditional expectation of $Z_{k,\ell}(\boldsymbol{G},\hat{\boldsymbol{\sigma}})$ is 'too big', the conditional expectation of $Z_{k,\ell}(\boldsymbol{G},\hat{\boldsymbol{\sigma}})$ given $\mathscr{R}$ (as defined in Lemma 3.3) is much smaller. Let $\boldsymbol{m}_0 = \boldsymbol{m}_0(\boldsymbol{G},\hat{\boldsymbol{\sigma}})$ be the total number of negative tests.

Lemma 3.4. *Suppose that $(1-1/\log n)k \le \ell < k$ and let $\Gamma_{\min} = \min_{i\in[m]}\Gamma_i$, $\Gamma_{\max} = \max_{i\in[m]}\Gamma_i$. Then*

$$\mathbb{E}[Z_{k,\ell}(\boldsymbol{G},\hat{\boldsymbol{\sigma}}) \mid \Gamma, \mathscr{R}, \boldsymbol{m}_0] \le O\left((\Delta k)^{3/2}\right)\binom{k}{\ell}\binom{n-k}{k-\ell}\left(1 - \left(1 - \frac{k-\ell}{n-k}\right)^{\Gamma_{\max}}\right)^{\delta\Delta(k-\ell)}\left(\frac{n-2k+\ell}{n-k}\right)^{\left(1+n^{-\Omega(1)}\right)\Gamma_{\min}\boldsymbol{m}_0}. \tag{20}$$

The proof of Lemma 3.4 is somehow subtle as we need to get a handle on the dependencies in $\boldsymbol{G}$ and is included in Section C.1. To convey the intuition behind the expression in Lemma 3.4, the term $\binom{k}{\ell}\binom{n-k}{k-\ell}$ accounts for the number of assignments $\tau \in \{0,1\}^V$ of Hamming weight k whose overlap with $\boldsymbol{\sigma}$ is equal to ℓ. The terms thereafter capture the probability that such an assignment τ exhibits the same test results as the true configuration $\boldsymbol{\sigma}$. The first term provides a necessary condition for a positive test under $\boldsymbol{\sigma}$ to stay positive under τ. By Lemma 3.3, we know that every infected individual shows up in at least $\delta\Delta$ tests as the only infected individual. Now, there are $k-\ell$ infected under $\boldsymbol{\sigma}$, but healthy under τ. For any of these $\delta\Delta(k-\ell)$ tests, we need to have at least one individual that is healthy under $\boldsymbol{\sigma}$, but infected under τ included in this test. Next, we need to ensure that any negative test under $\boldsymbol{\sigma}$ stay negative under τ. To this end, every individual included in a negative test under $\boldsymbol{\sigma}$ of which we have at least $\Gamma_{\min}\boldsymbol{m}_0$ must be healthy under τ. The second term captures this probability.

Proof of Proposition 3.2. In order to establish the proposition it suffices to show that there is $\varepsilon' \le (1-1/\log(n))k$ such that

$$\sum_{\varepsilon' \le \ell \le k} \mathbb{E}[Z_{k,\ell}(\boldsymbol{G},\hat{\boldsymbol{\sigma}})|\Gamma,\mathscr{R},\boldsymbol{m}_0] = o(1). \tag{21}$$

Starting from the expression in Lemma 3.4, setting $\alpha = \ell/k$ and recalling $m = ck\log(n/k)$ and $\Delta = d\log(n/k)$, we obtain

$$\begin{aligned}
&\mathbb{E}[Z_{k,\ell}(\boldsymbol{G},\hat{\boldsymbol{\sigma}})|\Gamma,\mathscr{R},\boldsymbol{m}_0] \\
&\le O\big((\Delta k)^{3/2}\big)\binom{k}{k-\ell}\binom{n-k}{k-\ell}\left(\frac{n-2k+\ell}{n-k}\right)^{(1+n^{-\Omega(1)})\Gamma_{\min}\boldsymbol{m}_0}\left(1-\left(1-\frac{k-\ell}{n-k}\right)^{\Gamma_{\max}}\right)^{\delta\Delta(k-\ell)} \\
&\le O\big((\Delta k)^{3/2}\big)\left(\frac{e}{1-\alpha}\right)^{(1-\alpha)k}\left(\frac{e(n-k)}{(1-\alpha)k}\right)^{(1-\alpha)k}\left(1-\frac{(1-\alpha)k}{n-k}\right)^{\frac{mn\log 2}{2k}\left(1+n^{-\Omega(1)}\right)}\left(1-2^{-(1-\alpha)\left(1+n^{-\Omega(1)}\right)}\right)^{\delta\Delta(1-\alpha)k} \qquad (22)\\
&\le O\big((\Delta k)^{3/2}\big)\left(\frac{e^2 n}{(1-\alpha)^2 k}\right)^{(1-\alpha)k}\exp\left((1-\alpha)k\frac{c\log 2}{2}\left(1+n^{-\Omega(1)}\right)\log(k/n)\right) \\
&\qquad\qquad \cdot\exp\left(-c\delta\log(2)\log\left(1-2^{-(1-\alpha)\left(1+n^{-\Omega(1)}\right)}\right)\log(k/n)(1-\alpha)k\right) \\
&\le O\big((\Delta k)^{3/2}\big)\left(\frac{e^2 n}{(1-\alpha)^2 k}\exp\left(\log(k/n)\left(1+n^{-\Omega(1)}\right)\left(\frac{c\log 2}{2}-c\delta\log(2)\log\left(1-2^{-(1-\alpha)\left(1+n^{-\Omega(1)}\right)}\right)\right)\right)\right)^{(1-\alpha)k}. \qquad (23)
\end{aligned}$$

As long as $1-\alpha = o(1)$, we find

$$(k/n)^{-\log\left(1-2^{-(1-\alpha)}\right)}(1-\alpha)^{-2} \to 0 \qquad \text{as } n\to\infty.$$

Moreover, $(1-\alpha)k \ge 1$. Thus, the expression (23) is of order

$$O\big((\Delta k)^{3/2}\big)(k/n)^{\omega(1)} = n^{-\omega(1)}. \tag{24}$$

Since (24) holds for any constant $c>0$ and any value of α s.t. $1-\alpha = o(1)$, it also holds for $\alpha \ge 1-1/\log n$. Consequently (21) is established w.h.p. □

Propositions 3.1 and 3.2 readily imply that $Z_k(\boldsymbol{G},\hat{\boldsymbol{\sigma}}) = 1$ w.h.p. if $m > (1+\varepsilon)m_{\inf}(k,\theta)$. Hence, Corollary 2.2 shows that there exists an inference algorithm that given $\boldsymbol{G},\hat{\boldsymbol{\sigma}}$ *and* k outputs $\boldsymbol{\sigma}$ w.h.p. Up to now, the algorithm relies on exactly knowing the number of infected individuals k, which in practice could be rather difficult to learn. Fortunately, this assumption can be removed. Namely, the following proposition shows that w.h.p. there is no assignment σ that is compatible with the test results and that has Hamming weight less than k.

Proposition 3.5. *Let $\varepsilon>0$ and $0<\theta<1$ and assume that $m > (1+\varepsilon)m_{\inf}(k,\theta)$. W.h.p. we have $\sum_{k'<k} Z_{k'}(\boldsymbol{G},\hat{\boldsymbol{\sigma}}) = 0$.*

Proof. To get started, suppose that $0<\theta<1$ and $c<\log^{-2}2$. We claim that for any value of $d>0$, $\left|V_0^+\right| \ge k\log n$ w.h.p.. Indeed, from Proposition 2.3(1), we know that

$$\left|V_0^+\right| = \left(1+n^{-\Omega(1)}\right)n\left(1-\exp(-d/c)\right)^{\Delta}.$$

Recalling $\Delta = d\log(n/k)$, the expression takes the minimum at $d = c\log 2$. It follows that

$$\left|V_0^+\right| \ge \left(1+n^{-\Omega(1)}\right)n(k/n)^{c\log^2 2}.$$

If $c = (1-\varepsilon)\log^{-2} 2$ for $\epsilon > 0$, then

$$\left|V_0^+\right| \geq \left(1+n^{-\Omega(1)}\right) n(k/n)^{1-\epsilon} = \left(1+n^{-\Omega(1)}\right) kn^{(1-\theta)\epsilon} \geq k\log n \qquad w.h.p. \tag{25}$$

Now, the following two statements establish that if there does not exist a second satisfying set of Hamming weight k, there does also not exist a satisfying set with smaller Hamming weight w.h.p..

First, we claim that if $m > (1+\epsilon)m_{\inf}(k,\theta)$, w.h.p. there does not exist a satisfying configuration with Hamming weight smaller than the correct configuration, where the set of infected individuals is not a subset of the true set of infected individuals. To see why, suppose there existed a satisfying configuration with a smaller Hamming weight, whose infected individuals are not a subset of the true infected individuals. By (25), we know that $\left|V_0^+\right| \gg k$ for $m < (1-\epsilon)m_{\text{alg}}$ w.h.p. Therefore, we could construct a satisfying configuration of identical Hamming weight as the true configuration by flipping individuals in V_0^+ from healthy to infected. Observe that by the definition of V_0^+, flipping individuals in V_0^+ does not change the test result. Therefore, we would be left with a second satisfying configuration of identical Hamming weight as the true configuration, a contradiction to Propositions 3.1 and 3.2.

Second, we argue that if $m > (1+\epsilon)m_{\inf}(k,\theta)$, w.h.p. there does not exist a satisfying configuration with Hamming weight smaller than the correct configuration, where the set of infected individuals is a subset of the true set of infected individuals. Suppose there existed a satisfying configuration with a smaller Hamming weight, whose infected individuals are a subset of the true infected individuals. Then, the true configuration would need to contain individuals in V_1^+, which can be flipped from infected to healthy without affecting the test result. However, Proposition 2.3(5) shows that for $m > (1+\epsilon)m_{\inf}$, $V_1^+ = \emptyset$ w.h.p. □

As an immediate consequence of Proposition 3.5 we conclude that for $m > (1+\varepsilon)m_{\inf}(k,\theta)$ the problem of inferring $\boldsymbol{\sigma}$ boils down to a minimum vertex cover problem, as previously conjectured by Aldridge, Baldassini and Johnson [12]. Namely, let $\mathscr{P}$ be the set of all positive tests, i.e., all tests a_i, $i \in [m]$, with $\hat{\boldsymbol{\sigma}}_{a_i} = 1$. Moreover, let V^+ be the set of all variables $x_i \in V$ such that $\partial x_i \subseteq \mathscr{P}$; in words, x_i takes part in positive tests only. We set up a hypergraph $\boldsymbol{H}$ with vertex set V^+ and hyperedges $\partial a_i \cap V^+$, $a_i \in \mathscr{P}$. Clearly, the set of all individuals x_i with $\boldsymbol{\sigma}_{x_i} = 1$ provides a valid vertex cover of $\boldsymbol{H}$ (as any positive test must feature an infected individual). Conversely, Propositions 3.1 and 3.2 show that w.h.p. this is the unique vertex cover of size k, and Proposition 3.5 shows that there is no strictly smaller vertex cover w.h.p. Therefore, w.h.p. we can infer $\boldsymbol{\sigma}$ even without prior knowledge of k by way of solving this minimum vertex cover instance.

4. The information-theoretic lower bound

We proceed with the negative statement that w.h.p. $\boldsymbol{\sigma}$ cannot be inferred if $m < (1-\varepsilon)m_{\inf}$. In light of Corollary 2.2 in order to prove the first part of Theorem 1.1 we need to show that the number $Z_k(\boldsymbol{G},\hat{\boldsymbol{\sigma}})$ of assignments consistent with the test results $\hat{\boldsymbol{\sigma}}$ is unbounded w.h.p. The proof of this fact is based on a very simple idea: we just identify a moderately large number of individuals whose infection status could be flipped without affecting the test results. The following lemma yields a bound on m below which the number of such potential false positives ($\left|V_0^+\right|$) and negatives ($\left|V_1^+\right|$) abound.

Proposition 4.1. *Let $\varepsilon > 0$ and $0 < \theta < 1$ and assume that*

$$m < \frac{(1-\varepsilon)\theta}{(1-\theta)\log^2 2} n^\theta (1-\theta)\log n.$$

Then for any choice of Δ we have $|V_0^+|, |V_1^+| = n^{\Omega(1)}$ w.h.p.

Proof. Thanks to Lemma 2.6 we may assume that $\Delta = d(\log(n/k))$, for a constant d as this choice minimizes the number of individuals in V_1^+. Then Proposition 2.3(4) guarantees that for every such constant as long as $c < \frac{\theta}{1-\theta}\frac{1}{\log^2 2}$, there are $n^{\Omega(1)}$ individuals in both V_1^+ and V_0^+, which yields to Proposition 4.1. □

As an immediate application we obtain the following information-theoretic lower bound.

Corollary 4.2. *Let $\varepsilon > 0$ and $0 < \theta < 1$ and assume that*

$$m < \frac{(1-\varepsilon)\theta}{(1-\theta)\log^2 2} n^\theta (1-\theta)\log n. \tag{26}$$

Then $Z_k(\boldsymbol{G},\hat{\boldsymbol{\sigma}}) = \omega(1)$ w.h.p.

Proof. We need to exhibit alternative vectors $\boldsymbol{\sigma}' \in \{0,1\}^V$ with Hamming weight k that render the same test results as $\boldsymbol{\sigma}$. Thus, pick any $x_i \in V_0^+$ and any $x_j \in V_1^+$ and obtain $\boldsymbol{\sigma}'$ from $\boldsymbol{\sigma}$ by setting $\boldsymbol{\sigma}'_{x_i} = 1$ and $\boldsymbol{\sigma}'_{x_j} = 0$. By construction, $\boldsymbol{\sigma}'$ has Hamming weight k and renders the same test results. Hence, Proposition 4.1 shows that $Z_k(\boldsymbol{G}, \hat{\boldsymbol{\sigma}}) \geq |V_0^+ \times V_1^+| = \Omega(n^{2\theta}) \gg 1$ w.h.p. □

The bound (26) matches m_{inf} for $\theta \gtrapprox 0.41$. A simpler, purely information-theoretic argument covers the remaining θ.

Proposition 4.3. *Let $\varepsilon > 0$, $0 < \theta < 1$. If $m < \frac{1-\varepsilon}{\log 2} n^\theta (1-\theta) \log n$, then $Z_k(\boldsymbol{G}, \hat{\boldsymbol{\sigma}}) = \omega(1)$ w.h.p.*

Proof. This Lemma follows from the classical information-theoretic lower bound for the group testing problem. Namely, m tests allow for 2^m possible test results. Hence, if

$$m < \frac{(1-\varepsilon)}{\log 2} n^\theta (1-\theta) \log n,$$

then the number of possible test results is far smaller than the number of vectors $\boldsymbol{\sigma} \in \{0,1\}^V$ with Hamming weight k. Therefore, w.h.p. there exists an unbounded number of vectors of Hamming weight k that render the same test results as $\boldsymbol{\sigma}$. □

We thus conclude that for all $0 < \theta < 1$, w.h.p. $Z_k(\boldsymbol{G}, \hat{\boldsymbol{\sigma}}) = \omega(1)$ if $m < (1-\varepsilon) m_{\mathrm{inf}}$. Therefore, the desired information-theoretic lower bound follows from Corollary 2.2.

5. The SCOMP algorithm

For $\theta \geq 1/2$ we have $m_{\mathrm{alg}} = m_{\mathrm{inf}}$ and thus Theorem 1.1 implies that SCOMP as described in Section 1.3 w.h.p. fails to infer $\boldsymbol{\sigma}$ for $m < (1-\varepsilon) m_{\mathrm{alg}}$. Therefore, we are left to establish Theorem 1.2 for $\theta < 1/2$, in which case

$$m_{\mathrm{alg}} = \frac{k \log(n/k)}{\log^2 2}. \tag{27}$$

The proof of Theorem 1.2 for $\theta < 1/2$ hinges on two propositions. First we show that below m_{alg}, the set V_1^{--} of infected individuals that the second step of SCOMP identifies correctly is empty. Formally, with V_0^- from (3), let

$$V_1^{--} = \left\{ x \in V_1 : \exists a \in \partial x : \partial a \setminus \{x\} \subseteq V_0^- \right\}.$$

Proposition 5.1. *Suppose that $0 < \theta < 1/2$ and $\varepsilon > 0$. If $m < (1-\varepsilon) m_{\mathrm{alg}}$, then for all $\Delta > 0$ we have $V_1^{--}(\boldsymbol{G}, \hat{\boldsymbol{\sigma}}^*) = \emptyset$ w.h.p.*

The proofs of Propositions 5.1 and 5.2 are based on moment calculations that turn out to be mildly subtle due to the potentially very large degrees of the underlying graph $\boldsymbol{G}$. The technical workout in included in Section D.1 and D.2.

With the second step of SCOMP failing to 'explain' (viz. cover) any positive tests, the greedy vertex cover algorithm takes over. This algorithm is applied to the hypergraph whose vertices are the as yet unclassified individuals and whose edges are the neighbourhoods of the positive tests. Our second lemma shows that the set $V^{+,\Delta}$ of potentially false positive individuals $x \in V_0^+$ that participate in the maximum number Δ of different tests is far greater than the actual number k of infected individuals. Formally, let

$$V_0^{+,\Delta} = \left\{ x \in V_0^+ : |\partial x| = \Delta \right\}.$$

Proposition 5.2. *Suppose that $0 < \theta < 1/2$ and $\varepsilon > 0$. If $m < (1-\varepsilon) m_{\mathrm{alg}}$, then for $\Delta = d \log(n/k)$ for all constant d we have $\left| V_0^{+,\Delta} \right| \geq k \log n$ w.h.p.*

We complete the proof of Theorem 1.2 as follows.

Proof of Theorem 1.2. The first step of SCOMP (correctly) marks all individuals that appear in negative tests as healthy. Moreover, Proposition 5.1 implies that the second step of SCOMP is void w.h.p., because there is no single infected individual that appears in a test whose other individuals have already been identified as healthy by the first step. Consequently, SCOMP simply applies the greedy vertex cover algorithm. Now, thanks to Proposition 5.2 it suffices to prove that SCOMP will fail w.h.p. if $\left| V_0^{+,\Delta} \right| = \omega(k)$. Because they belong to positive tests only, all the individuals of $V_0^{+,\Delta}$ are present in the vertex cover instance that SCOMP attempts to solve. Moreover, in the hypergraph no vertex

has degree greater than Δ, because the degrees of $x_1, \ldots, x_n$ in $\boldsymbol{G}$ are equal to Δ. (Some of the hypergraph degrees may be strictly smaller than Δ because $\boldsymbol{G}$ is a multi-graph.) Therefore, since $|V_0^{+,\Delta}| \geq k \log n$ while the actual set of infected individuals only has size k, w.h.p. the individual classified as infected by the very first step of the greedy set cover algorithm belongs to V_0^+. Hence, this individual is not actually infected, i.e., SCOMP errs w.h.p. □

Since the success probability of the SCOMP algorithm is at least as high as of the DD algorithm, we can prove the conjecture of [30] regarding the upper bound of the DD algorithm.

Corollary 5.3. *If $m < (1-\varepsilon) m_{\text{alg}}$, the* DD *algorithm will fail to retrieve the correct set of infected individuals w.h.p..*

Acknowledgment. We thank Arya Mazumdar for bringing the group testing problem to our attention.

A. NOTATION

Notation	Definition & Properties	Description
n		population size
k	$k \sim n^\theta$ for $\theta \in (0,1)$	number of infected individuals
m	$m = ck\log(n/k)$	number of tests
$x_1, \dots, x_n$		variable nodes
$V = V_n$	$\{x_i, \dots, x_n\}$	set of all individuals
$a_1, \dots, a_n$		factor nodes
$F = F_m$	$\{a_i, \dots, a_m\}$	set of all tests
Δ	$\Delta = d\log(n/k)$	tests per individual, variable node degree
$\Gamma_1, \dots, \Gamma_m$	$\left(\sum_{i=1}^m \Gamma_i\right)/m = dn/(ck)$	individuals per test, factor node degree
Γ	$(\Gamma_i)_{i\in[m]}$	σ-algebra generated by the random variables $(\Gamma_i)_{i\in[m]}$
$\boldsymbol{\sigma} \in \{0,1\}^V$	$\sum_{i=1}^n \boldsymbol{\sigma}_i = k$	n-dimensional vector of Hamming weight k indicating the individuals' infection status
$\boldsymbol{G} = \boldsymbol{G}(n,m,\Delta)$		random bipartite graph on n variable nodes, m factor nodes and variable degree Δ
$\partial x_i = \partial_{\boldsymbol{G}} x_i$ for $i \in [n]$	$\partial x_i \subseteq F, \lvert\partial x_i\rvert = \Delta$	set of tests that individual x_i participates in under $\boldsymbol{G}$
$\partial a_i = \partial_{\boldsymbol{G}} a_i$ for $i \in [m]$	$\partial a_i \subseteq V, \lvert\partial a_i\rvert = \Gamma_i$	set of individuals in test a_i under $\boldsymbol{G}$
$\hat{\boldsymbol{\sigma}} \in \{0,1\}^F$	$\hat{\boldsymbol{\sigma}}_i = \mathbf{1}\{\exists x \in \partial a_i : \boldsymbol{\sigma}_x = 1\}$	m-dimensional vector indicating the test outcomes
$\boldsymbol{m}_1, \boldsymbol{m}_0$	$\boldsymbol{m}_1 = \lvert\{a \in F : \hat{\boldsymbol{\sigma}}_a = 1\}\rvert, \boldsymbol{m}_0 = m - \boldsymbol{m}_1$	number of positive and negative tests
V_0	$V_0 = \{x \in V : \boldsymbol{\sigma}_x = 0\}$	set of healthy individuals
V_1	$V_1 = V \setminus V_0, \lvert V\rvert = k$	set of infected individuals
V_0^+	$\{x \in V_0 : \forall a \in \partial x : \hat{\boldsymbol{\sigma}}_a = 1\}$	set of healthy individuals only included in positive tests
V_0^-	$V_0^- = V_0 \setminus V_0^+$	set of healthy individuals included in at least one negative test
V_1^+	$\left\{x \in V_1 : \forall a \in \partial x : \exists y \in \partial a \setminus \{x\} : \boldsymbol{\sigma}_y = 1\right\}$	set of infected individuals that have another infected individual in all their tests
V_1^{--}	$\left\{x \in V_1 : \exists a \in \partial x : \partial a \setminus \{x\} \subseteq V_0^-\right\}$	Set of infected individuals that occur in at least one test with only healthy individuals
$\Gamma_{\min}, \Gamma_{\max}$	$\Gamma_{\min} = \min_{i\in[m]} \Gamma_i, \Gamma_{\max} = \max_{i\in[m]} \Gamma_i$	minimum and maximum test degree
$S_k(\boldsymbol{G}, \hat{\boldsymbol{\sigma}})$	$S_k(\boldsymbol{G}, \hat{\boldsymbol{\sigma}}) = \left\{\sigma \in \{0,1\}^V : \forall a_i \in [m] : \hat{\boldsymbol{\sigma}}_{a_i} = \mathbf{1}\{\exists x \in \partial a_i : \sigma_x = 1\}\right\}$	set of configurations consistent with the test results under $\boldsymbol{G}$
$Z_k(\boldsymbol{G}, \hat{\boldsymbol{\sigma}})$	$Z_k(\boldsymbol{G}, \hat{\boldsymbol{\sigma}}) = \lvert S_k(\boldsymbol{G}, \hat{\boldsymbol{\sigma}})\rvert$	number of configurations consistent with the test results
$Z_{k,\ell}(\boldsymbol{G}, \hat{\boldsymbol{\sigma}})$	$Z_{k,\ell}(\boldsymbol{G}, \hat{\boldsymbol{\sigma}}) = \lvert\{\sigma \in S_k(\boldsymbol{G}, \hat{\boldsymbol{\sigma}}) : \langle \boldsymbol{\sigma}, \sigma\rangle = \ell\}\rvert$	number of configuration consistent with the test results and with overlap ℓ with $\boldsymbol{\sigma}$
$\boldsymbol{Y}_i$ for $i \in [m]$	$\boldsymbol{Y}_i = \lvert\{x \in \partial a_i : \boldsymbol{\sigma}_x = 1\}\rvert$	number of edges that connect test a_i with an infected individual
$\boldsymbol{X}_i$ for $i \in [m]$	$\boldsymbol{X}_i \sim \mathrm{Bin}(\Gamma_i, k/n)$	binomially-distributed random variable with parameters Γ_i and k/n

W, W'	$W = \sum_{i=1}^{m} \mathbf{1}\{\boldsymbol{Y}_i = 1\}, W' = \sum_{i=1}^{m} \mathbf{1}\{\boldsymbol{X}_i = 1\}$	W is the number of tests containing a single infected individual, W' is a random variable depending on $(\boldsymbol{X}_i)_{i\in[m]}$
U	$U = \lvert\{x \in V_1 : \forall a_i \in \partial x : \boldsymbol{Y}_i > 1\}\rvert$	number of infected individuals not adjacent not any test with precisely one infected individual
T	$\left\lvert\left\{x \in V_1 : \sum_{a\in\partial x} \mathbf{1}\{\partial a \setminus \{x\} \subseteq V_0\} < \delta\Delta\right\}\right\rvert$	number of infected individuals who appear in less than $\delta\Delta$ tests as the only infected individual for some constant $\delta > 0$
R	$R = \lvert\{x \in V_1 : \exists a_i \in \partial x :$ $\boldsymbol{Y}_i > 1$ and $\partial a \setminus \{x\} \subseteq V_0\}\rvert$	number of infected individual adjacent to some test multiple times with no other infected individual besides themselves
$\boldsymbol{A}'_{i,j}, \boldsymbol{A}'_{i,j,k}$		auxiliary random variables, defined in proof of Proposition 3.1
$\mathscr{A}$	$\mathscr{A} = \{\forall i \in [m] :$ $\max_{j\in[\Gamma_i]} \boldsymbol{A}'_{i,j,1} = \max_{j\in[\Gamma_i]} \boldsymbol{A}'_{i,j,2}\}$	event that every test under the balls-and-bins experiment features the same test result
$\mathscr{E}$	$\mathscr{E} = \{\sum_{i\in[m]} \boldsymbol{X}_i = k\Delta\}$	event that the sum of $\boldsymbol{X}_i$ is exactly $k\Delta$
$\mathscr{M}$		set of all indices $i \in [m]$ for which there exists precisely one $g_i \in [\Gamma_i]$ such that $\boldsymbol{A}'_{i,g_i,1} = 1$
$\mathscr{N}$		set of indices $i \in [m]$ such that $\max_{j\in[\Gamma_i]} \boldsymbol{A}'_{i,j,1} = 0$
$\mathscr{R}$	$\mathscr{R} = \{\forall x \in V_1 : \lvert\{a \in \partial x : \partial a \setminus \{x\} \subseteq V_0\}\rvert \geq \delta\Delta\}$	event that for every $x \in V_1$ there are at least $\delta\Delta$ tests $a \in \partial x$ for some $\delta > 0$ such that $\partial a \setminus \{x\} \subseteq V_0$.
$\mathscr{S}$		event that one specific σ that has overlap ℓ with $\boldsymbol{\sigma}$ belongs to $S_k(\boldsymbol{G}, \hat{\boldsymbol{\sigma}})$
$\mathscr{T}$		event that sum of independent random variable is equal to specific value, defined in (7)
$\mathscr{V}$	$\mathscr{V} = \{\boldsymbol{m}_1 = \frac{m}{2}(1 + o(1))\}$	event that around half of the tests are positive
$\mathscr{W}$	$\mathscr{W} =$ $\{\lvert V_0^+\rvert = (1 + o(1))(n-k)(1 - \exp(-d/c))^{\Delta}\}$	event that the size of V_0^+ is concentrated around its mean
$o(1), \omega(1)$		$o(1)$ [$\omega(1)$] denotes a term that vanishes [diverges] in the limit of large n
w.h.p.		probability of $1 - o(1)$ as $n \to \infty$

The following sections contain the proofs of the lemmas omitted so far.

B. Preliminaries

B.1. **Preliminaries.** We start by stating the Chernoff bound as applied in this paper.

Lemma B.1 (Chernoff bound, [29] (Section 2.1))**.** *Let $\boldsymbol{X} \sim \mathrm{Bin}(n, p)$ be a binomially-distributed random variable with $\lambda = \mathbb{E}[\boldsymbol{X}]$. Further, let*

$$\varphi : (-1, \infty) \to \mathbb{R}_{\geq 0}, x \mapsto (1 + x)\log(1 + x) - x$$

Then for some $t \geq 0$,

$$\mathbb{P}\left(|\boldsymbol{X} - \lambda| \geq t\right) \leq \exp\left(-\lambda\varphi\left(t/\lambda\right) - (n - \lambda)\varphi\left(-t/(n - \lambda)\right)\right)$$

As an application, we readily find

$$\mathbb{P}\left(|\boldsymbol{X}-\lambda| \geq \sqrt{n}\log n\right) \leq n^{-\omega(1)}$$

Next, we justify that the Stirling approximation of Section 3 is accurate. Namely, let $\boldsymbol{A}'_{i,j} = (\boldsymbol{A}'_{i,j,1}, \boldsymbol{A}'_{i,j,2})$ be $\{0,1\}^2$-valued random variables such that $(\boldsymbol{A}'_{i,j})_{i\in[m],j\in[\Gamma_i]}$ are mutually independent and such that

$$\mathbb{P}\left[\boldsymbol{A}'_{i,j}=(1,1)\right] = \ell/n, \qquad \mathbb{P}\left[\boldsymbol{A}'_{i,j}=(0,1)\right] = \mathbb{P}\left[\boldsymbol{A}'_{i,j}=(1,0)\right] = (k-\ell)/n,$$
$$\mathbb{P}\left[\boldsymbol{A}'_{i,j}=(0,0)\right] = (n-2k+\ell)/n$$

for all i,j. As before, we denote by $\mathcal{T}$ the event that

$$\sum_{i=1}^{m}\sum_{j=1}^{\Gamma_i}\mathbf{1}\left\{\boldsymbol{A}'_{i,j}=(1,1)\right\} = \ell\Delta, \qquad \sum_{i=1}^{m}\sum_{j=1}^{\Gamma_i}\mathbf{1}\left\{\boldsymbol{A}'_{i,j}=(0,0)\right\} = (n-2k+\ell)\Delta,$$
$$\sum_{i=1}^{m}\sum_{j=1}^{\Gamma_i}\mathbf{1}\left\{\boldsymbol{A}'_{i,j}=(1,0)\right\} = \sum_{i=1}^{m}\sum_{j=1}^{\Gamma_i}\mathbf{1}\left\{\boldsymbol{A}'_{i,j}=(0,1)\right\} = (k-\ell)\Delta,$$

i..e, that all of the sums on the l.h.s. are *precisely* equal to their expected values. Since the $(\boldsymbol{A}'_{i,j})_{i,j}$ are independent, Stirling's formula yields

$$\mathbb{P}\left[\mathcal{T}\right] = \Omega\left((\Delta k)\right)^{-3/2}. \tag{28}$$

This can be seen as follows. For the sake of brevity, define

$$p_{00} = (n-2k+\ell)/n, \quad p_{11} = \ell/n, \quad \text{and} \quad p_{10} = p_{01} = (k-\ell)/n.$$

As $\boldsymbol{A}'_{i,j}$ is a family of independent multinomial variables

$$\boldsymbol{A}'_{i,j} \sim \mathrm{Mult}\left(1,(p_{11},p_{00},p_{10},p_{01})\right),$$

we find

$$\boldsymbol{X} \sim \sum_{i=1}^{m}\sum_{j=1}^{\Gamma_i}\boldsymbol{A}'_{i,j} \sim \mathrm{Mult}\left(n\Delta,(p_{11},p_{00},p_{10},p_{01})\right).$$

Hence, the probability of event $\mathcal{T}$ occurring is the probability, that $\boldsymbol{X}$ hits its expectation. Thus, using the very basic approximation $n! = \Theta\left(\sqrt{n}\right)(n/\mathrm{e})^n$ we find

$$\begin{aligned}
\mathbb{P}\left(\mathcal{T}\right) &= \frac{(n\Delta)!(\ell/n)^{\ell\Delta}((n-2k+\ell)/n)^{(n-2k+\ell)\Delta}((k-\ell)/n)^{2(k-\ell)\Delta}}{(\ell\Delta)!((n-2k+\ell)\Delta)!((k-\ell)\Delta)!((k-\ell)\Delta)!} \\
&= \Theta\left(\frac{\sqrt{n\Delta}}{\sqrt{\ell(n-2k\ell)(k-\ell)^2\Delta^4}}\right)\left(\frac{(n\Delta)^n(\ell/n)^\ell((n-2k+\ell)/n)^{n-2k+\ell}((k-\ell)/n)^{2(k-\ell)}}{\ell^\ell(n-2k+\ell)^{n-2k+\ell}(k-\ell)^{2(k-\ell)}}\right)^\Delta \\
&= (1+O(1/n))\Theta\left(\frac{\sqrt{n}}{\sqrt{n}\sqrt{\ell k^2-2\ell^2k+\ell^3-k\left(2\ell k^2/n-2k\ell^2/n\right)+\ell^4/n}}\right) \\
&= \Omega\left(\sqrt{\Delta^{-3}(\ell k^2+\ell^2k+\ell^3)^{-1}}\right) = \Omega\left((\Delta k)^{-3/2}\right),
\end{aligned} \tag{29}$$

where (29) follows immediately from $\ell \leq k = o(n)$ and directly implies (28). In due course we apply similar calculations often, some calculations involve conditional probabilities. These conditions are only restricting Γ_i to take specific (common) values and clearly the above argument is totally invariant under different values of Γ_i, as long as $\sum_i^m \Gamma_i = n\Delta$.

B.2. **Getting started.** In the next step, recall that neighbourhoods of different tests in the random multi-graph seizably intersect. To cope with the ensuing correlations, we introduce a new family of random variables that, as we will see, are closely related to the statistics of the appearances of infected/uninfected individuals in the various tests. Specifically, recalling that Γ_i signifies the degree of test a_i and that $\sum_{i=1}^m \Gamma_i = n\Delta$, let $(\boldsymbol{X}_i)_{i\in[m]}$ be a sequence of independent $\mathrm{Bin}(\Gamma_i, k/n)$-variables. Moreover, let

$$\mathscr{E} = \left\{ \sum_{i\in[m]} \boldsymbol{X}_i = k\Delta \right\}.$$

Because the $\boldsymbol{X}_i$ are mutually independent, Stirling's formula shows that

$$\mathbb{P}[\mathscr{E}] = \Omega(1/\sqrt{\Delta k}), \tag{30}$$

which follows along the lines of Section B.1. Additionally, let $\boldsymbol{Y}_i$ be the number of edges that connect test a_i with an infected individual. (Since $\boldsymbol{G}$ is a multi-graph, it is possible that an infected individual contributes more than one to $\boldsymbol{Y}_i$.) Further, let Γ be the σ-algebra generated by the random variables $(\Gamma_i)_{i\in[m]}$. Whenever we condition on Γ, we assume that the bounds from Lemma 2.4 and 2.5 hold.

Lemma B.2. *Given Γ, the vectors $(\boldsymbol{Y}_1,\ldots,\boldsymbol{Y}_m)$ and $(\boldsymbol{X}_1,\ldots,\boldsymbol{X}_m)$ given $\mathscr{E}$ are identically distributed.*

Proof. For any integer sequence $(y_i)_{i\in[m]}$ with $y_i \geq 0$ and $\sum_{i\in[m]} y_i = k\Delta$ we have

$$\mathbb{P}\left[\forall i \in [m] : \boldsymbol{Y}_i = y_i \mid \Gamma\right] = \frac{\binom{k\Delta}{y_1,\ldots,y_m}\binom{(n-k)\Delta}{\Gamma_1-y_1,\ldots,\Gamma_m-y_m}}{\binom{n\Delta}{\Gamma_1,\ldots,\Gamma_m}} = \frac{\prod_{i=1}^m \frac{\Gamma_i!}{y_i!(\Gamma_i-y_i)!}}{\frac{(n\Delta)!}{(k\Delta)!((n-k)\Delta)!}} = \binom{n\Delta}{k\Delta}^{-1}\prod_{i=1}^m \binom{\Gamma_i}{y_i}.$$

Hence, for any sequences $(y_i), (y_i')$ we obtain

$$\frac{\mathbb{P}\left[\forall i \in [m] : \boldsymbol{Y}_i = y_i \mid \Gamma\right]}{\mathbb{P}\left[\forall i \in [m] : \boldsymbol{Y}_i = y_i' \mid \Gamma\right]} = \prod_{i=1}^m \frac{\binom{\Gamma_i}{y_i}}{\binom{\Gamma_i}{y_i'}} = \frac{\mathbb{P}\left[\forall i \in [m] : \boldsymbol{X}_i = y_i \mid \Gamma, \mathscr{E}\right]}{\mathbb{P}\left[\forall i \in [m] : \boldsymbol{X}_i = y_i' \mid \Gamma, \mathscr{E}\right]},$$

as claimed. □

B.3. **Proof of Lemma 2.4.** Since each variable draws a sequence of Δ tests uniformly at random, for every $i \in [m]$ the degree Γ_i has distribution $\mathrm{Bin}(n\Delta, 1/m)$. Therefore, the assertion follows from the Chernoff bound.

B.4. **Proof of Lemma 2.5.** Let $\boldsymbol{m}_0' = \sum_{i=1}^m \mathbf{1}\{\boldsymbol{X}_i = 0\}$. Then $\mathbb{E}[\boldsymbol{m}_0'] = \sum_{i=1}^m \mathbb{P}[\mathrm{Bin}(\Gamma_i, k/n)) = 0] = \sum_{i=1}^m (1-k/n)^{\Gamma_i}$. Hence, Lemma 2.4 shows that with probability $1-o(n^{-2})$,

$$\mathbb{E}[\boldsymbol{m}_0' \mid \Gamma] \geq m(1-k/n)^{\Gamma_{\max}} = m\exp\left((\Delta n/m + O(\sqrt{\Delta n/m}\log n))\log(1-k/n)\right) \tag{31}$$

$$= m\left(\exp(-d/c) + O(\sqrt{k/n}\log n)\right), \tag{32}$$

$$\mathbb{E}[\boldsymbol{m}_0' \mid \Gamma] \leq m(1-k/n)^{\Gamma_{\min}} = m\left(\exp(-d/c) + O(\sqrt{k/n}\log n)\right). \tag{33}$$

Because the $\boldsymbol{X}_i$ are mutually independent, $\boldsymbol{m}_0'$ is a binomial variable. Therefore, the Chernoff bound (e.g. Lemma B.1) shows that

$$\mathbb{P}\left[\left|\boldsymbol{m}_0' - \mathbb{E}[\boldsymbol{m}_0' \mid \Gamma]\right| > \sqrt{m}\log n \mid \Gamma\right] = o(n^{-10}). \tag{34}$$

Finally, the assertion follows from (30), (31)–(34) and Lemma B.2.

B.5. **Proof of Lemma 2.6.** The expected degree of a test a_i equals $\Delta n/m$. Therefore, if $\Delta = o(\log(n/k))$, then by Lemma 2.5, $\boldsymbol{m}_1 = o(m)$ w.h.p. To exploit this fact, call $\sigma \in \{0,1\}^V$ of Hamming weight k *bad* for $\boldsymbol{G}$ if given $\boldsymbol{\sigma} = \sigma$ we indeed have $\boldsymbol{m}_1 = o(m)$. Let $B(\boldsymbol{G})$ be the set of all such bad σ. Then w.h.p. $\boldsymbol{G}$ has the property that $|B(\boldsymbol{G})| \sim \binom{n}{k}$, i.e. asymptotically most configurations will have few positive tests. Now, condition on the event that $|B(\boldsymbol{G})| \sim \binom{n}{k}$ and let $\mathscr{B}$ be the set of all subsets of $[m]$ of size $o(m)$. Further, let $f_{\boldsymbol{G}} : B(\boldsymbol{G}) \to \mathscr{B}$ map $\sigma \in \{0,1\}^V$ to the corresponding set of positive tests. Finally, let $B'(\boldsymbol{G})$ be the set of all $\sigma \in B(\boldsymbol{G})$ such that $|f_{\boldsymbol{G}}^{-1}(f_{\boldsymbol{G}}(\sigma))| < n$, i.e. the set of all configurations for which there are less than n other configurations rendering the same test results. Then

$$|B'(\boldsymbol{G})| \leq n|\mathscr{B}| \leq n\binom{m}{o(m)} = \exp(o(m)) = o\left(\binom{n}{k}\right).$$

Consequently, w.h.p. over the choice of $\boldsymbol{G}$ and $\boldsymbol{\sigma}$ we have $Z_k(\boldsymbol{G},\hat{\boldsymbol{\sigma}}) \geq n$. The same argument applies for $\log(n/k) = o(\Delta)$ with the term 'positive test' replaced by 'negative test'.

B.6. **Proof of Proposition 2.3.** We start by proving part (1) using a straightforward second-moment calculation. Recall $\Delta = d\log(n/k)$ and $m = ck\log(n/k)$. Lemma 2.4 and Lemma 2.5 show that with probability at least $1-o(n^{-2})$ the total degree of the negative tests comes to

$$\sum_{i=1}^{m} \mathbf{1}\{\partial a_i \subseteq V_0\}\,\Gamma_i = \Delta n\exp(-d/c) + O\Big(\sqrt{m}\log^2(n)\Delta n/m + m\sqrt{\Delta n/m}\log n\Big)$$
$$= \Delta n\exp(-d/c) + O\Big(\Big(\sqrt{nk} + n/\sqrt{k}\Big)\log^3 n\Big) = \Delta n\big(\exp(-d/c) + n^{-\Omega(1)}\big).$$

Consequently, with probability at least $1-o(n^{-2})$ the total number of edges between V_0 and the set of positive tests is $\Delta n\big(1-\exp(-d/c)+n^{-\Omega(1)}\big)$. Moreover, the total number of edges between V_0 and all tests comes down to $\Delta(n-k)$. Given these events and since each individual is assigned to tests uniformly at random with replacement, the probability that a given $x \in V_0$ belongs to V_0^+ comes out as

$$\binom{\Delta n\big(1-\exp(-d/c)+n^{-\Omega(1)}\big)}{\Delta}\binom{\Delta(n-k)}{\Delta}^{-1} = \big(1+n^{-\Omega(1)}\big)\big(1-\exp(-d/c)\big)^{\Delta}.$$

Next, we estimate the probability that $x, x' \in V_0$ both belong to V_0^+:

$$\binom{\Delta n\big(1-\exp(-d/c)+n^{-\Omega(1)}\big)}{2\Delta}\binom{\Delta(n-k)}{2\Delta}^{-1} = \big(1+n^{-\Omega(1)}\big)\big(1-\exp(-d/c)\big)^{2\Delta},$$

Hence, $\mathbb{E}[|V_0^+|^2 \mid \Gamma] - \mathbb{E}[|V_0^+| \mid \Gamma]^2 = O(n^{2-\Omega(1)})$. Therefore, the assertion follows from Chebyshev's inequality.

Proceeding with part (2), let the number of tests containing a single infected individual be

$$W = \sum_{i=1}^{m} \mathbf{1}\{\boldsymbol{Y}_i = 1\}, \qquad W' = \sum_{i=1}^{m} \mathbf{1}\{\boldsymbol{X}_i = 1\}.$$

Then Lemma 2.4 shows that w.h.p.

$$\mathbb{E}[W'] = \sum_{i=1}^{m} \frac{\Gamma_i k}{n}(1-k/n)^{\Gamma_i - 1} \leq \frac{\Gamma_{\max} km}{n}(1-k/n)^{\Gamma_{\min}-1}$$
$$= \big(1+n^{-\Omega(1)}\big)\,k\Delta(1-k/n)^{\Delta n/m} = \big(1+n^{-\Omega(1)}\big)\,k\Delta\exp(-d/c)$$

Analogously,

$$\mathbb{E}[W'] \geq \frac{\Gamma_{\min} km}{n}(1-k/n)^{\Gamma_{\max}} = \big(1+n^{-\Omega(1)}\big)\,k\Delta\exp(-d/c).$$

Hence, because W' is a binomial random variable, the Chernoff bound (e.g. Lemma B.1) shows that

$$\mathbb{P}\big[W' = \big(1+n^{-\Omega(1)}\big)\,k\Delta\exp(-d/c) \mid \Gamma\big] = 1-o(n^{-9}).$$

Therefore, (30) yields

$$\mathbb{P}\big[W = \big(1+n^{-\Omega(1)}\big)\,k\Delta\exp(-d/c) \mid \Gamma\big] = 1-o(n^{-7}). \tag{35}$$

Now, let U be the number of $x \in V_1$ that are not adjacent to any test with precisely one positive individual. An individual $x \in V_1$ counts towards U, if out of all possible assignment $k\Delta$, it is only assigned to those tests where it is not the only infected individual (there are a total of $k\Delta - W$ such assignments). Using the notation $n^{\underline{k}} = n(n-1)\ldots(n-k+1)$ and recalling $\Delta = \Theta(\log n)$, the bound on W yields

$$\mathbb{E}[U \mid \Gamma, W] = k\binom{k\Delta - W}{\Delta}\binom{k\Delta}{\Delta}^{-1} = k\frac{(k\Delta-W)^{\underline{\Delta}}}{(k\Delta)^{\underline{\Delta}}} = \big(1+n^{-\Omega(1)}\big)\,k\left(\frac{k\Delta-W}{k\Delta}\right)^{\Delta}$$
$$= \big(1+n^{-\Omega(1)}\big)\,k(1-W/k\Delta)^{\Delta} = \big(1+n^{-\Omega(1)}\big)\,k(1-\exp(-d/c))^{\Delta}.$$

By a similar token we obtain

$$\mathbb{E}[U^2 \mid \Gamma, W] = k^2\binom{k\Delta - W}{2\Delta}\binom{k\Delta}{2\Delta}^{-1} = \big(1+n^{-\Omega(1)}\big)\,\mathbb{E}[U \mid \Gamma, W]^2.$$

Therefore, Chebyshev's inequality shows that w.h.p.

$$U = \left(1 + n^{-\Omega(1)}\right) k(1-\exp(-d/c))^{\Delta}. \tag{36}$$

To complete the proof we need to compare U and $\left|V_1^+\right|$. Clearly, $U \geq \left|V_1^+\right|$. But the inequality may be strict because U includes positive individuals that appear twice in the same test. To be precise, an individual might be assigned to one test twice as the only infected individual. Such an individual should not be in V_1^+, but it shows up in U. Indeed, letting R be the number of such individuals, we obtain $\left|V_1^+\right| \geq U - R$. Hence, we are left to estimate R. To this end, we observe that the probability that an individual appears in a specific test twice is upper-bounded by $(\Delta/m)^2$. Recall $m = ck\log(n/k)$ and $\Delta = d\log(n/k)$. Consequently, taking the union bound over all tests and infected individuals we yield

$$\mathbb{E}[R \mid \Gamma] \leq km\left(\frac{\Delta}{m}\right)^2 = O(\log n).$$

Since by assumption the r.h.s. of (36) is $n^{\Omega(1)}$, we conclude that $\left|V_1^+\right| \geq U - R = n^{\Omega(1)}$ w.h.p., as claimed.

Next, we consider (3). Define U as in the proof of Proposition 2.3(2). Then we know that $U \geq \left|V_1^+\right|$. Hence, if $k(1-\exp(-d/c))^{\Delta} = o(1)$ then $\left|V_1^+\right| = o(1)$ due to (36).

For part (4), we observe for a given c that $\min_d (1-\exp(-d/c))^{\Delta}$ is attained at $d = c\log 2$. To see this, consider the function $f(d) = (1-\exp(-d/c))^{\Delta} = n^{(1-\theta)d\log(1-\exp(-d/c))}$ and observe that the minimum of $f(d)$ coincides with the minimum of $g(d) = d\log(1-\exp(-d/c))$. Letting $x = d/c$, the derivatives read as

$$\begin{aligned} g(x) &= cx\log(1-\exp(-x)) \\ g'(x) &= c\left(\log(1-\exp(-x)) + \frac{x\exp(-x)}{1-\exp(-x)}\right) \\ g''(x) &= c\left(-\frac{(x-2)\exp(x)+2}{(\exp(x)-1)^2}\right) \end{aligned}$$

For $d > 0$, the unique maximum is attained at $x = \log 2$ and accordingly, $d = c\log 2$. Furthermore, it is the case that $k(1-\exp(-\log 2))^{c\log 2\log(n/k)} \geq n^{\Omega(1)}$ and therefore by Proposition 2.3(2), $\left|V_1^+\right| = n^{\Omega(1)}$. By a similar token by Proposition 2.3(1), $\left|V_0^+\right| = n^{\Omega(1)}$.

Finally, for part (5), setting $d = c\log 2$, we see that $k(1-\exp(-\log 2))^{c\log 2\log(n/k)} = o(1)$ and therefore by Proposition 2.3(3), $\left|V_1^+\right| = o(1)$.

C. The information-theoretic upper bound

C.1. **Proof of Lemma 3.4.** The term $\binom{k}{\ell}\binom{n-k}{k-\ell}$ accounts for the number of assignments $\sigma \in \{0,1\}^V$ of Hamming weight k whose overlap with $\boldsymbol{\sigma}$ is equal to ℓ. Hence, with $\mathscr{S}$ being the event that one specific $\sigma \in \{0,1\}^V$ that has overlap ℓ with $\boldsymbol{\sigma}$ belongs to $S_{k,\ell}(\boldsymbol{G}, \hat{\boldsymbol{\sigma}})$, we need to show that

$$\mathbb{P}\left[\mathscr{S} \mid \Gamma, \mathscr{R}, \boldsymbol{m}_0\right] \leq O\left((\Delta k)^{3/2}\right) \cdot \left(1 - \left(1 - \frac{k-\ell}{n-k}\right)^{\Gamma_{\max}}\right)^{\delta\Delta(k-\ell)} \left(\frac{n-2k+\ell}{n-k}\right)^{\Gamma_{\min}\boldsymbol{m}_0} \tag{37}$$

Due to symmetry we may assume that $\boldsymbol{\sigma}_{x_i} = \mathbf{1}\{i \leq k\}$ and that $\sigma_{x_i} = \mathbf{1}\{i \leq \ell\} + \mathbf{1}\{k < i \leq 2k-\ell\}$.

Proceeding as in the proof of Proposition 3.1, we think of each test a_i as a bin of capacity Γ_i and of each clone (x_i, h), $h \in [\Delta]$, of an individual as a ball labelled $(\boldsymbol{\sigma}_{x_i}, \sigma_{x_i}) \in \{0,1\}^2$. We toss the Δn balls randomly into the bins. For $i \in [m]$ and for $j \in [\Gamma_i]$ we let $\boldsymbol{A}_{i,j} = (\boldsymbol{A}_{i,j,1}, \boldsymbol{A}_{i,j,2}) \in \{0,1\}^2$ be the label of the jth ball that ends up in bin number i. To cope with this experiment we introduce a new set $\{0,1\}^2$-valued random variables $\boldsymbol{A}'_{i,j} = (\boldsymbol{A}'_{i,j,1}, \boldsymbol{A}'_{i,j,2})$ such that $(\boldsymbol{A}'_{i,j})_{i\in[m], j\in[\Gamma_i]}$ are mutually independent and

$$\begin{aligned} &\mathbb{P}\left[\boldsymbol{A}'_{i,j} = (1,1)\right] = \ell/n, && \mathbb{P}\left[\boldsymbol{A}'_{i,j} = (0,1)\right] = \mathbb{P}\left[\boldsymbol{A}'_{i,j} = (1,0)\right] = (k-\ell)/n, \\ &\mathbb{P}\left[\boldsymbol{A}'_{i,j} = (0,0)\right] = (n-2k+\ell)/n \end{aligned}$$

for all i, j. With $\mathscr{T}$ being the event that

$$\sum_{i=1}^{m}\sum_{j=1}^{\Gamma_i}\mathbf{1}\left\{\boldsymbol{A}'_{i,j}=(1,1)\right\}=\ell\Delta,\qquad \sum_{i=1}^{m}\sum_{j=1}^{\Gamma_i}\mathbf{1}\left\{\boldsymbol{A}'_{i,j}=(0,0)\right\}=(n-2k+\ell)\Delta, \tag{38}$$

$$\sum_{i=1}^{m}\sum_{j=1}^{\Gamma_i}\mathbf{1}\left\{\boldsymbol{A}'_{i,j}=(1,0)\right\}=\sum_{i=1}^{m}\sum_{j=1}^{\Gamma_i}\mathbf{1}\left\{\boldsymbol{A}'_{i,j}=(0,1)\right\}=(k-\ell)\Delta, \tag{39}$$

the vector $\boldsymbol{A}'=(\boldsymbol{A}'_{i,j})_{i,j}$ given $\mathscr{T}$ is distributed as $\boldsymbol{A}=(\boldsymbol{A}_{i,j})_{i,j}$ given Γ. Moreover, with similar arguments as in Section B.1, Stirling's formula yields

$$\mathbb{P}[\mathscr{T}]=\Omega((\Delta k)^{-3/2}). \tag{40}$$

Let $\mathscr{N}$ be the set of indices $i\in[m]$ such that $\max_{j\in[\Gamma_i]}\boldsymbol{A}'_{i,j,1}=0$. Moreover, let $\mathscr{M}$ be the set of all indices $i\in[m]$ for which there exists precisely one $g_i\in[\Gamma_i]$ such that $\boldsymbol{A}'_{i,g_i,1}=1$ and such that for this index we have $\boldsymbol{A}'_{i,g_i,2}=0$. Further, let

$$\mathscr{S}'=\left\{\forall i\in\mathscr{N}:\max_{j\in[\Gamma_i]}\boldsymbol{A}'_{i,j,2}=0\right\},\qquad\qquad \mathscr{S}''=\left\{\forall i\in\mathscr{M}:\max_{j\in[\Gamma_i]}\boldsymbol{A}'_{i,j,2}=1\right\}.$$

Then

$$\mathscr{A}=\left\{\forall i\in[m]:\max_{j\in[k]}\boldsymbol{A}'_{i,j,1}=\max_{j\in[k]}\boldsymbol{A}'_{i,j,2}\right\}\subseteq\mathscr{S}'\cap\mathscr{S}''.$$

Furthermore, given $\mathscr{N},\mathscr{M}$ the events $\mathscr{S}',\mathscr{S}''$ are independent and

$$\mathbb{P}\left[\mathscr{S}'\mid\mathscr{N}\right]=\prod_{i\in\mathscr{N}}\left(\frac{n-2k+\ell}{n-k}\right)^{\Gamma_i}\le\left(\frac{n-2k+\ell}{n-k}\right)^{\Gamma_{\min}|\mathscr{N}|},$$

$$\mathbb{P}\left[\mathscr{S}''\mid\mathscr{M}\right]=\prod_{i\in\mathscr{M}}\left(1-\left(1-\frac{k-\ell}{n-k}\right)^{\Gamma_i-1}\right)\le\left(1-\left(1-\frac{k-\ell}{n-k}\right)^{\Gamma_{\max}}\right)^{|\mathscr{M}|}.$$

For an intuitive explanation of the above expressions, please refer to the section immediately following the statement of the Lemma 3.4. Given $|\mathscr{N}|\ge\left(1-n^{-\Omega(1)}\right)\boldsymbol{m}_0$ and $|\mathscr{M}|\ge\delta\Delta(k-\ell)$, we obtain

$$\mathbb{P}\left[\mathscr{A}\mid|\mathscr{N}|\ge\left(1-n^{-\Omega(1)}\right)\boldsymbol{m}_0,|\mathscr{M}|\ge\delta\Delta(k-\ell)\right]\le\left(\frac{n-2k+\ell}{n-k}\right)^{\Gamma_{\min}\boldsymbol{m}_0}\left(1-\left(1-\frac{k-\ell}{n-k}\right)^{\Gamma_{\max}}\right)^{\delta\Delta(k-\ell)}. \tag{41}$$

Moreover, we find by 3.3, the concentration of $|\mathscr{N}|$ and the fact that $\mathbb{E}[|\mathscr{N}|]=\mathbb{E}[\boldsymbol{m}_0]=m/2$

$$\mathbb{P}\left(|\mathscr{N}|\ge\left(1-n^{-\Omega(1)}\right)\boldsymbol{m}_0,|\mathscr{M}|\ge\delta\Delta(k-\ell)\right)=1-o(1)$$

and thus

$$\mathbb{P}\left[\mathscr{T}\,||\mathscr{N}|\ge\left(1-n^{-\Omega(1)}\right)\boldsymbol{m}_0,|\mathscr{M}|\ge\delta\Delta(k-\ell)\right]=\Omega((\Delta k)^{-3/2}).$$

Combining (40)–(41) and using the trivial bound

$$\mathbb{P}\left[\mathscr{T}\mid\mathscr{S},\mathscr{S}',|\mathscr{N}|\ge\left(1-n^{-\Omega(1)}\right)\boldsymbol{m}_0,|\mathscr{M}|\ge\delta\Delta(k-\ell)\right]\le 1, \tag{42}$$

we obtain by Bayes Theorem

$$\mathbb{P}\left[\mathscr{A}\mid\mathscr{T},|\mathscr{N}|\ge\left(1-n^{-\Omega(1)}\right)\boldsymbol{m}_0,|\mathscr{M}|\ge\delta\Delta(k-\ell)\right]\le O\left((\Delta k)^{3/2}\right)\left(\frac{n-2k+\ell}{n-k}\right)^{\left(1-n^{-\Omega(1)}\right)\Gamma_{\min}\boldsymbol{m}_0}\left(1-\left(1-\frac{k-\ell}{n-k}\right)^{\Gamma_{\max}}\right)^{\delta\Delta(k-\ell)}. \tag{43}$$

Because $\boldsymbol{A}'=(\boldsymbol{A}'_{i,j})_{i,j}$ given $\mathscr{T}$ is distributed as $\boldsymbol{A}=(\boldsymbol{A}_{i,j})_{i,j}$ given Γ, (37) follows from (43).

D. THE SCOMP ALGORITHM

D.1. **Proof of Proposition 5.1.** The proof of Proposition 5.1 proceeds in three steps. First, we show that $\left|V_0^+\right|$ is concentrated around its expectation. $\mathscr{W}$ denotes the corresponding event. Second, we need to get a handle on the subtle dependencies in $\boldsymbol{G}$. To this end, we introduce a set of independent multinomial random variables indexed over the tests. Whereas $\boldsymbol{Y}_1^i, \boldsymbol{Y}_{0+}^i, \boldsymbol{Y}_{0-}^i$ denotes the number of infected, potentially false positive and definitively healthy individuals in test a_i, respectively, the triple $(\boldsymbol{X}_1^i, \boldsymbol{X}_{0+}^i, \boldsymbol{X}_{0-}^i)$ denote the corresponding multinomial random variable. We will show that conditioned on the sum of $\boldsymbol{X}_1^i, \boldsymbol{X}_{0+}^i, \boldsymbol{X}_{0-}^i$ hitting the total number of individuals of the three types, $(\boldsymbol{X}_1^i, \boldsymbol{X}_{0+}^i, \boldsymbol{X}_{0-}^i)$ is distributed like $\boldsymbol{Y}_1^i, \boldsymbol{Y}_{0+}^i, \boldsymbol{Y}_{0-}^i$. The technical workout is delicate, but is based on standard results from balls-into-bins experiments. Third, we show that for $m < (1-\varepsilon)m_{\text{alg}}$, the number of tests W for which $\boldsymbol{X}_1^i = 1$ and $\boldsymbol{X}_{0+}^i = 0$ decays exponentially in n, which implies that $V_1^{--} = \emptyset$ w.h.p.

Proof. Lemma 2.6 implies that the optimal choice for the variable degree is $\Delta = d\log(n/k)$ for a constant d. Let $\boldsymbol{m}_1$ be the amount of positive tests and, w.l.o.g. assume that $a_1 \ldots a_{\boldsymbol{m}_1}$ are the positive tests and define

$$\mathscr{W} = \left\{|V_0^+| = (1+o(1))(n-k)(1-\exp(-d/c))^{\Delta}\right\}.$$

as the event that the number of 'potential false positives' $\left|V_0^+\right|$ is highly concentrated around its mean. Then by Proposition 2.3(1), we find

$$\mathbb{P}[\mathscr{W}] \geq 1 - o(1) \tag{44}$$

Similarly as before, we introduce a family of independent random variables corresponding to the tests.

Let $\boldsymbol{Y}_1^1, \ldots, \boldsymbol{Y}_1^{\boldsymbol{m}_1}$ be the number of ones in the tests corresponding to $a_1, \ldots, a_{\boldsymbol{m}_1}$ respectively. Let $\boldsymbol{Y}_{0+}^1, \ldots, \boldsymbol{Y}_{0+}^{\boldsymbol{m}_1}$ count the V_0^+ occurrences in $a_1, \ldots, a_{\boldsymbol{m}_1}$. Let $\boldsymbol{Y}_{0-}^1, \ldots, \boldsymbol{Y}_{0-}^{\boldsymbol{m}_1}$ count the V_0^- occurrences in $a_1, \ldots, a_{\boldsymbol{m}_1}$. By definition we find $\boldsymbol{Y}_{0-}^i = \Gamma_i - \boldsymbol{Y}_{0+}^i - \boldsymbol{Y}_1^i$. We introduce auxiliary variables $\boldsymbol{X}_1^1, \ldots, \boldsymbol{X}_1^{\boldsymbol{m}_1}$, $\boldsymbol{X}_{0+}^1, \ldots, \boldsymbol{X}_{0+}^{\boldsymbol{m}_1}, \boldsymbol{X}_{0-}^1, \ldots, \boldsymbol{X}_{0-}^{\boldsymbol{m}_1}$ such that $(\boldsymbol{X}_1^i, \boldsymbol{X}_{0+}^i, \boldsymbol{X}_{0-}^i)$ have distribution

$$\text{Mult}_{\geq(1,0,0)}(\Gamma_i, p, q, 1-p-q),$$

a multinomial distribution conditioned on the first variable being at least one. The triples $\left((X_1^i, X_{0^+}^i, X_{0-}^i)\right)_{i\in\boldsymbol{m}_1}$ are mutually independent. We seek a choice of p satisfying the equation

$$p := \frac{k\Delta}{\sum_{i=1}^{\boldsymbol{m}_1} \frac{\Gamma_i}{1-(1-p)^{\Gamma_i}}} \qquad \text{and} \qquad q := \frac{\left|V_0^+\right|\Delta}{\sum_{i=1}^{\boldsymbol{m}_1} \frac{\Gamma_i}{1-(1-p)^{\Gamma_i}}}.$$

and will show following equation (48) that such a choice exists. Define

$$\mathscr{E} = \left\{\sum_{i=1}^{\boldsymbol{m}_1} \boldsymbol{X}_1^i = k\Delta, \sum_{i=1}^{\boldsymbol{m}_1} \boldsymbol{X}_{0+}^i = |V_0^+|\Delta\right\}.$$

Along the lines of Section B.1 , Stirling's formula implies

$$\mathbb{P}[\mathscr{E}] = \Omega(1/n). \tag{45}$$

Moreover, $(\boldsymbol{Y}_1^1, \boldsymbol{Y}_{0+}^1, \boldsymbol{Y}_{0-}^1, \ldots, \boldsymbol{Y}_1^{\boldsymbol{m}_1}, \boldsymbol{Y}_{0+}^{\boldsymbol{m}_1}, \boldsymbol{Y}_{0-}^{\boldsymbol{m}_1})$ and $(\boldsymbol{X}_1^1, \boldsymbol{X}_{0+}^1, \boldsymbol{X}_{0-}^1, \ldots, \boldsymbol{X}_1^{\boldsymbol{m}_1}, \boldsymbol{X}_{0+}^{\boldsymbol{m}_1}, \boldsymbol{X}_{0-}^{\boldsymbol{m}_1})$ given $\mathscr{E}$ are identically distributed. This can be seen as follows:

$$\begin{aligned}
&\mathbb{P}\left[\forall i \in [\boldsymbol{m}_1] : (\boldsymbol{Y}_1^i, \boldsymbol{Y}_{0+}^i, \boldsymbol{Y}_{0-}^i) = (y_i, y_i', y_i'') \mid \Gamma, |V_0^+|, \boldsymbol{m}_1\right] \\
&= \frac{\binom{k\Delta}{y_1 \ldots y_{\boldsymbol{m}_1}}\binom{|V_0^+|\Delta}{y_1' \ldots y_{\boldsymbol{m}_1}'}\binom{\sum_{i=1}^{\boldsymbol{m}_1}\Gamma_i - (k+|V_0^+|)\Delta}{\Gamma_1 - y_1 - y_1', \ldots, \Gamma_{\boldsymbol{m}_1} - y_{\boldsymbol{m}_1} - y_{\boldsymbol{m}_1}'}}{\binom{\sum_{i=1}^{\boldsymbol{m}_1}\Gamma_i}{\Gamma_1, \ldots, \Gamma_{\boldsymbol{m}_1}}} \mathbf{1}\{\forall i \in [\boldsymbol{m}_1] : y_i'' = \Gamma_i - y_i - y_i'\} \\
&= \binom{\sum_{i=1}^{\boldsymbol{m}_1}\Gamma_i}{k\Delta, |V_0^+|\Delta, \sum_{i=1}^{\boldsymbol{m}_1}\Gamma_i - \left(k+|V_0^+|\right)\Delta} \prod_{i=1}^{\boldsymbol{m}_1} \binom{\Gamma_i}{y_i, y_i', \Gamma - y_i - y_i'} \mathbf{1}\{\forall i \in [\boldsymbol{m}_1] : y_i'' = \Gamma_i - y_i - y_i'\}.
\end{aligned}$$

Thus, given $y_i'' = \Gamma_i - y_i - y_i'$ and $\tilde{y}_i'' = \Gamma_i - \tilde{y}_i - \tilde{y}_i'$ for all $i \in [\boldsymbol{m}_1]$, we find

$$\frac{\mathbb{P}\left[\forall i \in [\boldsymbol{m}_1] : (\boldsymbol{Y}_1^1, \boldsymbol{Y}_{0+}^1, \boldsymbol{Y}_{0-}^1) = (y_i, y_i', y'') \mid \Gamma, |V_0|^+, \boldsymbol{m}_1\right]}{\mathbb{P}\left[\forall i \in [\boldsymbol{m}_1] : (\boldsymbol{Y}_1^1, \boldsymbol{Y}_{0+}^1, \boldsymbol{Y}_{0-}^1) = (\tilde{y}_i, \tilde{y}_i', \tilde{y}_i'') \mid \Gamma, |V_0|^+, \boldsymbol{m}_1\right]} = \prod_{i=1}^{\boldsymbol{m}_1} \frac{\binom{\Gamma_i}{y_i, y_i', \Gamma - y_i - y_i'}}{\binom{\Gamma_i}{\tilde{y}_i, \tilde{y}_i', \Gamma - \tilde{y}_i - \tilde{y}_i'}}. \tag{46}$$

Given $x_i'' = \Gamma_i - x_i - x_i'$, we find:

$$\begin{aligned}
&\mathbb{P}\left[\forall i \in [\boldsymbol{m}_1] : (\boldsymbol{X}_1^1, \boldsymbol{X}_{0+}^1, \boldsymbol{X}_{0-}^1) = (x_i, x_i', x_i'') \mid \mathscr{E}, \Gamma, |V_0|^+, \boldsymbol{m}_1\right] \\
&= \prod_{i=1}^{\boldsymbol{m}_1} \binom{\Gamma_i}{x_i, x_i', x_i''} p^{x_i} q^{x_i'} (1-p-q)^{x_i''} \frac{1}{1-(1-p)^{\Gamma_i}} \\
&= p^{k\Delta} q^{|V_0^+|\Delta} (1-p-q)^{\sum_{i=1}^{\boldsymbol{m}_1} \Gamma_i - \Delta(k+|V_0^+|)} \prod_{i=1}^{\boldsymbol{m}_1} \frac{1}{1-(1-p)^{\Gamma_i}} \binom{\Gamma_i}{x_i, x_i', x_i''}
\end{aligned}$$

where the last equality follows from the fact that we conditioned on $\mathscr{E}$. Since the first terms are independent of x_i, x_i', x_i'', we find

$$\frac{\mathbb{P}\left[\forall i \in [\boldsymbol{m}_1] : (\boldsymbol{X}_1^i, \boldsymbol{X}_{0+}^i, \boldsymbol{X}_{0-}^i) = (x_i, x_i', x_i'') \mid \mathscr{E}, \Gamma, |V_0|^+, \boldsymbol{m}_1\right]}{\mathbb{P}\left[\forall i \in [\boldsymbol{m}_1] : (\boldsymbol{X}_1^i, \boldsymbol{X}_{0+}^i, \boldsymbol{X}_{0-}^i) = (\tilde{x}_i, \tilde{x}_i', \tilde{x}_i'') \mid \mathscr{E}, \Gamma, |V_0|^+, \boldsymbol{m}_1\right]} = \prod_{i=1}^{\boldsymbol{m}_1} \frac{\binom{\Gamma_i}{x_i, x_i', \Gamma - x_i - x_i'}}{\binom{\Gamma_i}{\tilde{x}_i, \tilde{x}_i', \Gamma - \tilde{x}_i - \tilde{x}_i'}}.$$

Therefore, given $\Gamma_i = x_i + x_i' + x_i'' = \tilde{x}_i + \tilde{x}_i' + \tilde{x}_i''$, we have by comparison with (46),

$$\begin{aligned}
&\frac{\mathbb{P}\left[\forall i \in [\boldsymbol{m}_1] : (\boldsymbol{X}_1^i, \boldsymbol{X}_{0+}^i, \boldsymbol{X}_{0-}^i) = (x_i, x_i', x_i'') \mid \mathscr{E}, \Gamma, |V_0|^+, \boldsymbol{m}_1\right]}{\mathbb{P}\left[\forall i \in [\boldsymbol{m}_1] : (\boldsymbol{X}_1^i, \boldsymbol{X}_{0+}^i, \boldsymbol{X}_{0-}^i) = (\tilde{x}_i, \tilde{x}_i', \tilde{x}_i'') \mid \mathscr{E}, \Gamma, |V_0|^+, \boldsymbol{m}_1\right]} \\
=&\frac{\mathbb{P}\left[\forall i \in [\boldsymbol{m}_1] : (\boldsymbol{Y}_1^i, \boldsymbol{Y}_{0+}^i, \boldsymbol{Y}_{0-}^i) = (x_i, x_i', x_i'') \mid \Gamma, \boldsymbol{m}_1\right]}{\mathbb{P}\left[\forall i \in [\boldsymbol{m}_1] : (\boldsymbol{Y}_1^i, \boldsymbol{Y}_{0+}^i, \boldsymbol{Y}_{0-}^i) = (\tilde{x}_i, \tilde{x}_i', \tilde{x}_i'') \mid \Gamma, \boldsymbol{m}_1\right]},
\end{aligned}$$

which yields the claim. Let

$$W = \sum_{i=1}^{\boldsymbol{m}_1} \mathbf{1}\left\{\boldsymbol{X}_1^i + \boldsymbol{X}_{0+}^i = 1\right\}.$$

be the number of positive tests that contain exactly one infected individual and no healthy individuals in V_0^+. Note that this split is the only possibility for the test to be positive. Then

$$\mathbb{E}[W|\Gamma, \mathscr{E}, \left|V_0^+\right|, \boldsymbol{m}_1] = \sum_{i=1}^{\boldsymbol{m}_1} \mathbb{P}[\boldsymbol{X}_1^i = 1, \boldsymbol{X}_{0+}^i = 0, \boldsymbol{X}_{0-}^i = \Gamma_i - 1] = \sum_{i=1}^{\boldsymbol{m}_1} \frac{\Gamma_i p(1-p-q)^{\Gamma_i - 1}}{1-(1-p)^{\Gamma_i}}.$$

By Lemma 2.5 we readily find for any choice of $c, d = \Theta(1)$ that

$$\sum_{i=1}^{\boldsymbol{m}_1} \frac{\Gamma_i p(1-p-q)^{\Gamma_i - 1}}{1-(1-p)^{\Gamma_i}} = (1+o(1)) \sum_{i=1}^{m} \Gamma_i p(1-p-q)^{\Gamma_i - 1} \tag{47}$$

Hence,

$$m\Gamma_{\min} p(1-p-q)^{\Gamma_{\max}} \le \mathbb{E}[W \mid \Gamma, \mathscr{E}, |V_0^+|, \boldsymbol{m}_1] \le m\Gamma_{\max} p(1-p-q)^{\Gamma_{\min}-1}.$$

Moreover, since W is a binomial random variable, the Chernoff bound (e.g. Lemma B.1) shows that

$$\mathbb{P}\left[\left|W - \mathbb{E}[W \mid \Gamma, \mathscr{E}, |V_0^+|]\right| > \sqrt{m}\log n\right] \le O(n^{-2}).$$

Further, Lemma 2.4 yields approximations for $\Gamma_{\min}$ and $\Gamma_{\max}$. Now assume that $c < \log^{-2} 2$. Using a similar reformulation as in (47), we find that $p = (1+o(1))k/n$. Thus, we have

$$\begin{aligned}
&\mathbb{E}[W|\Gamma, \mathscr{E}, \mathscr{W}] \\
&= (1+o(1)) m \frac{dn}{ck} \frac{k}{n} \exp\left((1+o(1)) \frac{dn}{ck} \log\left((1-k/n)\left(1+n^{-\Omega(1)}\right)(1-(1-\exp(-d/c)))^{\Delta}\right)\right) \\
&= (1+o(1)) m \exp(-d/c) \frac{d}{c} \left(1-(k/n)^{-d\log(1-\exp(-d/c))}\right)^{dn/(ck)}
\end{aligned} \tag{48}$$

As Lemma 2.6 shows, the optimal value of d is a constant. For a fixed c the same d that maximizes $-d/c\log(1-\exp(-d/c))$ in (48), also maximizes $\mathbb{E}[W|\Gamma, \mathscr{E}, \left|V_0^+\right|]$. This maximum is attained at $d = c\log 2$. Consequently $p = o(q)$ and

$$q \sim \left(\frac{k}{n}\right)^{c\log^2 2}.$$

Hence,

$$\mathbb{E}[W \mid \Gamma, \mathscr{E}, \mathscr{W}] \sim \frac{k\Delta}{2} \exp\left(-(\log 2)\left(\frac{n}{k}\right)^{1-c\log^2 2}\right) = \exp(-n^{\Omega(1)}).$$

As before, we find $\mathbb{E}[W] \to 0$ w.h.p. since $\mathbb{P}(\mathscr{W}) = 1 - o(1)$ and $\mathbb{P}(\mathscr{E}) = \Omega(1/\sqrt{\Delta k})$ and Markov's inequality leads to $V_1^{--} = \emptyset$. Proposition 5.1 follows. □

D.2. **Proof of Proposition 5.2.** By Lemma 2.3, we have $\left|V_0^+\right| \geq k \log n$ for $m < (1-\varepsilon) m_{\text{alg}}$. To prove Proposition 5.2, we need to show that for such m, we also have $\left|V_0^{+,\Delta}\right| \geq k \log n$. We proceed in two steps. First, we show that every individual $x \in V$ is assigned to at least $\Delta - O(1)$ distinct tests. Second, we show that a constant fraction of individuals $x \in V_0^+$ are assigned to exactly Δ tests establishing Proposition 5.2.

Proof. Let $d^\star(x)$ be the number of distinct neighbors of a vertex x. We claim that w.h.p. the following statements are true.

$$\min_{x \in V} d^\star(x) \geq \Delta - 2/\theta^2.$$

The probability that a given $x \in V$ appears $\ell \geq 2$ times in the same test is upper-bounded by

$$\binom{\Delta}{\ell} m^{1-\ell} \leq \frac{m}{\ell!}\left(\frac{d}{ck}\right)^\ell = \frac{ck\log(n/k)}{\ell!}\left(\frac{d}{ck}\right)^\ell \leq \frac{c(d/c)^\ell}{\ell!} n^{(1-\ell)\theta + o(1)} = o(1/n),$$

provided that $\ell > 1 + 1/\theta$. Moreover, the probability that x appears in one test twice is upper-bounded by $\Delta\dot{\Delta}/m$. Thus, the probability that x appears in at least ℓ tests at least twice is upper-bounded by

$$\sum_{i=\ell}^{\lfloor \Delta/2 \rfloor} \left(\frac{\Delta^2}{m}\right)^i = (1+o(1))\left(\frac{\Delta^2}{m}\right)^\ell \leq (1+o(1))\left(\frac{O(\log^2 n)}{ck\log(n/k)}\right)^\ell = n^{-\theta\ell + o(1)} = o(1/n),$$

provided that $\ell > 1/\theta$ and since $m = ck\log(n/k)$ and $\Delta = d\log(n/k)$. The bound follows.

By Lemma 2.3, we know that for $m < (1-\varepsilon) m_{\text{alg}}$, $\left|V_0^+\right| \geq k \log n$ w.h.p.. Since the SCOMP algorithm in its third stage selects the individual with the highest number of adjacent unexplained tests, we are left to show that also $\left|V_0^{+,\Delta}\right| \geq k \log n$, which implies that w.h.p. we erroneously classify a healthy individual as infected. The prior bounds ensure that each individual is in at least $\Delta - O(1)$ tests. The question remains which fraction of individuals in V_0^+ are in $V_0^{+,\Delta}$. In principle, it could be the case that most potentially false positive individuals of V_0^+ appear in less than Δ different tests. Indeed, it is more likely for such an individual in V_0^+ to be in fewer than Δ different tests since each additional test increases the probability for such an individual to be assigned to a negative test. However, we claim that a constant fraction of all potentially false positive individuals in V_0^+ will have degree Δ, thus be in $V_0^{+,\Delta}$. To see this, let p be the maximum proportion of $\left|V_0^{+,\Delta-i}\right|$ and $\left|V_0^{+,\Delta-i+1}\right|$ for $i \in [2/\theta^2]$, i.e.

$$p = \max_{i \in [2/\theta^2]} \frac{\left|V_0^{+,\Delta-i}\right|}{\left|V_0^{+,\Delta-i+1}\right|}$$

By conditioning on a test degree sequence $\Gamma_1, \ldots, \Gamma_m$, we find

$$p \geq (1 - (1 - (k/n))^{\Gamma_{\min}}) = \Theta(1),$$

as long as $c, d = \Theta(1)$, which by Lemma 2.6 we can safely assume. Since each individual in V_0^+ is in at least $\Delta - O(1)$ different tests and the probability of being in any number of different tests $\Delta, \Delta-1\ldots$ is constant, a constant fraction of individuals in V_0^+ will be in exactly Δ tests. Since $\left|V_0^+\right| = \Omega(k\log n)$, the claim follows. □

REFERENCES

[1] E. Abbe: Community detection and stochastic block models: recent developments. The Journal of Machine Learning Research **18** (2017) 6446–6531.

[2] D. Achlioptas, A. Coja-Oghlan: Algorithmic barriers from phase transitions. Proc. 49th FOCS (2008) 793–802.

[3] D. Achlioptas, A. Coja-Oghlan, F. Ricci-Tersenghi: On the solution space geometry of random formulas. Random Structures and Algorithms **38** (2011) 251–268.

[4] D. Achlioptas, C. Moore: Random k-SAT: two moments suffice to cross a sharp threshold. SIAM Journal on Computing **36** (2006) 740–762.

[5] D. Achlioptas, A. Naor, and Y. Peres: Rigorous location of phase transitions in hard optimization problems. Nature **435** (2005) 759–764.

[6] D. Achlioptas, Y. Peres: The threshold for random k-SAT is $2^k \log 2 - O(k)$. Journal of the AMS **17** (2004) 947–973.

[7] A. Alaoui, A. Ramdas, F. Krzakala, L. Zdeborová, M. Jordan: Decoding from pooled data: Sharp information-theoretic bounds. SIAM Journal on Mathematics of Data Science **1** (2019) 161–188.

[8] A. Alaoui, A. Ramdas, F. Krzakala, L. Zdeborová, M. Jordan: Decoding from pooled data: Phase transitions of message passing. IEEE Transactions on Information Theory **65** (2019) 572–585.

[9] M. Aldridge: On the optimality of some group testing algorithms. IEEE International Symposium on Information Theory (2017).

[10] M. Aldridge: The capacity of Bernoulli nonadaptive group testing. IEEE Transactions on Information Theory **63** (2017) 7142–7148.

[11] M. Aldridge: Individual testing is optimal for nonadaptive group testing in the linear regime. IEEE Transactions on Information Theory **65** (2019) 2058–2061.

[12] M. Aldridge, L. Baldassini, O. Johnson: Group testing algorithms: bounds and simulations. IEEE Transactions on Information Theory **60** (2014) 3671–3687.

[13] M. Aldridge, O. Johnson, J. Scarlett: Improved group testing rates with constant column weight designs. IEEE International Symposium on Information Theory (2016).

[14] M. Aldridge, O. Johnson, J. Scarlett: Group testing: an information theory perspective. arXiv preprint arXiv:1902.06002 (2019).

[15] A. Alleman: An efficient algorithm for combinatorial group testing. H. Aydinian, F. Cicalese, C. Deppe (eds) Information Theory, Combinatorics, and Search Theory. Lecture Notes in Computer Science. Springer, Berlin, Heidelberg **7777** (2013), 569–596.

[16] R. Benz, S. Swamidass, P. Baldi: Discovery of power-laws in chemical space. Journal of Chemical Information and Modeling **48** (2008) 1138–1151.

[17] H. Chen, F. Hwang: A survey on nonadaptive group testing algorithms through the angle of decoding. Journal of Combinatorial Optimization **15** (2008) 49–59.

[18] A. Coja-Oghlan, F. Krzakala, W. Perkins , L. Zdeborová: Information-theoretic thresholds from the cavity method. Advances in Mathematics **333** (2018) 694–795.

[19] A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimroth, P. Loick: Optimal non-adaptive group testing. (2019) arXiv: 1911.02287.

[20] A. Coja-Oghlan, K. Panagiotou: The asymptotic k-SAT threshold. Advances in Mathematics **288** (2016) 985–1068.

[21] B. Davis, D. McDonald: An elementary proof of the local central limit theorem. Journal of Theoretical Probability **8** (1995) 693–702.

[22] A. Decelle, F. Krzakala, C. Moore, L. Zdeborová: Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. Phys. Rev. E **84** (2011) 066106.

[23] J. Ding, A. Sly, N. Sun: Proof of the satisfiability conjecture for large k. Proc. 47th STOC (2015) 59–68.

[24] R. Dorfman: The detection of defective members of large populations. Annals of Mathematical Statistics **14** (1943) 436–440.

[25] D. Du, F. Hwang: Combinatorial group testing and its applications. World Scientific (1993).

[26] O. Dubois, J. Mandler: The 3-XORSAT Threshold. Proc. 43rd FOCS (2002) 769–778.

[27] A. Emad, O. Milenkovic: Poisson group testing: a probabilistic model for nonadaptive streaming Boolean compressed sensing. Proc. ICASSP (2014) 3335–3339.

[28] M. Hahn-Klimroth, P. Loick: Optimal adaptive group testing. (2019) arXiv:1911.06647.

[29] S. Janson, T. Luczak, A. Ruciński: Random Graphs, Wiley 2000.

[30] O. Johnson, M. Aldridge, J. Scarlett: Performance of group testing algorithms with near-constant tests per item. IEEE Transactions on Information Theory **65** (2019) 707–723.

[31] H. Kwang-Ming, D. Ding-Zhu: Pooling designs and nonadaptive group testing: important tools for DNA sequencing. World Scientific (2006)

[32] M. Mézard, A. Montanari: Information, physics and computation. Oxford University Press 2009.

[33] M. Mézard, M. Tarzia, C. Toninelli: Group Testing with Random Pools: Phase Transitions and Optimal Strategy. Journal of Statistical Physics **131** (2008) 783–801.

[34] M. Molloy: The freezing threshold for k-colourings of a random graph. Proc. 43rd STOC (2012) 921–930.

[35] C. Moore: The computer science and physics of community detection: landscapes, phase transitions, and hardness. Bulletin of the EATCS **121** (2017).

[36] R. Mourad, Z. Dawy, F. Morcos: Designing pooling systems for noisy high-throughput protein-protein interaction experiments using Boolean compressed sensing. IEEE/ACM Transactions on Computational Biology and Bioinformatics **10** (2013) 1478–1490.

[37] E. Mossel, J. Neeman, A. Sly: Reconstruction and estimation in the planted partition model. Probability Theory and Related Fields (2014) 1–31.

[38] H. Ngo, D. Du: A survey on combinatorial group testing algorithms with applications to DNA library screening. Discrete Mathematical Problems with Medical Applications **7** (2000) 171–182.

[39] J. Scarlett, V. Cevher: Phase transitions in group testing. Proc. 27th SODA (2016) 40–53.

[40] J. Scarlett, V. Cevher: Limits on support recovery with probabilistic models: an information-theoretic framework. IEEE Transactions on Information Theory **63** (2017) 593-620.

[41] A. Sebo: On two random search problems. Journal of Statistical Planning and Inference **11** (1985) 23–31.

[42] N. Thierry-Mieg: A new pooling strategy for high-throughput screening: the shifted transversal design. BMC Bioinformatics **7** (2006) 28
[43] L. Zdeborová, F. Krzakala: Statistical physics of inference: thresholds and algorithms. Advances in Physics **65** (2016) 453–552.

AMIN COJA-OGHLAN, acoghlan@math.uni-frankfurt.de, GOETHE UNIVERSITY, MATHEMATICS INSTITUTE, 10 ROBERT MAYER ST, FRANKFURT 60325, GERMANY.

OLIVER GEBHARD, gebhard@math.uni-frankfurt.de, GOETHE UNIVERSITY, MATHEMATICS INSTITUTE, 10 ROBERT MAYER ST, FRANKFURT 60325, GERMANY.

MAX HAHN-KLIMROTH, hahnklim@math.uni-frankfurt.de, GOETHE UNIVERSITY, MATHEMATICS INSTITUTE, 10 ROBERT MAYER ST, FRANKFURT 60325, GERMANY.

PHILIPP LOICK, loick@math.uni-frankfurt.de, GOETHE UNIVERSITY, MATHEMATICS INSTITUTE, 10 ROBERT MAYER ST, FRANKFURT 60325, GERMANY.

APPENDIX C. OPTIMAL GROUP TESTING

OPTIMAL GROUP TESTING

AMIN COJA-OGHLAN, OLIVER GEBHARD, MAX HAHN-KLIMROTH, PHILIPP LOICK

ABSTRACT. In the group testing problem the aim is to identify a small set of $k \sim n^\theta$ infected individuals out of a population size n, $0 < \theta < 1$. We avail ourselves of a test procedure capable of testing groups of individuals, with the test returning a positive result iff at least one individual in the group is infected. The aim is to devise a test design with as few tests as possible so that the set of infected individuals can be identified correctly with high probability. We establish an explicit sharp information-theoretic/algorithmic phase transition m_{inf} for non-adaptive group testing, where all tests are conducted in parallel. Thus, with more than m_{inf} tests the infected individuals can be identified in polynomial time w.h.p., while learning the set of infected individuals is information-theoretically impossible with fewer tests. In addition, we develop an optimal adaptive scheme where the tests are conducted in two stages. *MSc: 05C80, 60B20, 68P30*

1. INTRODUCTION

1.1. **Background and motivation.** Various intriguing combinatorial problems come as inference tasks where we are to learn a hidden ground truth by means of indirect queries. The goal is to get by with as small a number of queries as possible. The ultimate solution to such a problem should consist of a positive algorithmic result showing that a certain number of queries suffice to learn the ground truth efficiently, complemented by a matching information-theoretic lower bound showing that with fewer queries the problem is insoluble, regardless of computational resources.

Group testing is a prime example of such an inference problem [6]. The objective is to identify within a large population of size n a subset of k individuals infected with a rare disease. We presume that the number of infected individuals scales as a power $k = \lceil n^\theta \rceil$ of the population size with an exponent $\theta \in (0,1)$, a parametrisation suited to modelling the pivotal early stages of an epidemic [36]. Indeed, since early on in an epidemic test kits might be in short supply, it is vital to get the most diagnostic power out the least number of tests. To this end we assume that the test gear is capable of not merely testing a single individual but an entire group. The test comes back positive if any one individual in the group is infected and negative otherwise. While in *non-adaptive* group testing all tests are conducted in parallel, in *adaptive* group testing test are conducted in several stages. In either case we are free to allocate individuals to test groups as we please. Randomisation is allowed. What is the least number of tests required so that the set of infected individuals can be inferred from the test results with high probability? Furthermore, in adaptive group testing, what is the smallest depth of test stages required?

Closing the considerable gaps that the best prior bounds left, the main results of this paper furnish matching algorithmic and information-theoretic bounds for both adaptive and non-adaptive group testing. Specifically, the best prior information-theoretic lower bound derives from the following folklore observation. Suppose that we conduct m tests that each return either 'positive' or 'negative'. Then to correctly identify the set of infected individuals we need the total number 2^m of conceivable test results to asymptotically exceed the number $\binom{n}{k}$ of possible sets of infected individuals. Hence, $2^m \geq (1+o(1))\binom{n}{k}$. Thus, Stirling's formula yields the lower bound

$$m_{\mathrm{ad}} = \frac{1-\theta}{\ln 2} n^\theta \ln n, \tag{1.1}$$

which applies to both adaptive and non-adaptive testing. On the positive side, a randomised non-adaptive test design with

$$m_{\mathrm{DD}} \sim \frac{\max\{\theta, 1-\theta\}}{\ln^2 2} n^\theta \ln n \tag{1.2}$$

Supported by DFG CO 646/3 and Stiftung Polytechnische Gesellschaft. An extended abstract version of this work has been submitted to the COLT 2020 conference.

tests exists from which a greedy algorithm called DD correctly infers the set of infected individuals w.h.p. [22]. Clearly, $m_{\mathrm{ad}} < m_{\mathtt{DD}}$ for all infection densities θ and $m_{\mathtt{DD}}/m_{\mathrm{ad}} \to \infty$ as $\theta \to 1$. In addition, there is an efficient adaptive three-stage group testing scheme that asymptotically matches the lower bound m_{ad} [33].

We proceed to state the main results of the paper. First, improving both the information-theoretic and the algorithmic bounds, we present optimal results for non-adaptive group testing. Subsequently we show how the non-adaptive result can be harnessed to perform adaptive group testing with the least possible number $(1+o(1))m_{\mathrm{ad}}$ of tests in only two stages.

1.2. **Non-adaptive group testing.** A *non-adaptive test design* is a bipartite graph $G = (V \cup F, E)$ with one vertex class $V = V_n = \{x_1, \ldots, x_n\}$ representing individuals and the other class $F = F_m = \{a_1, \ldots, a_m\}$ representing tests. For a vertex v of G denote by $\partial v = \partial_G v$ the set of neighbours of v. Thus, an individual x_j takes part in a test a_i iff $x_j \in \partial a_i$. Since we can shuffle the individuals randomly, we may safely assume that the vector $\boldsymbol{\sigma} \in \{0,1\}^V$ whose 1-entries mark the infected individuals is a uniformly random vector of Hamming weight k. Furthermore, the test results induced by $\boldsymbol{\sigma}$ read

$$\hat{\boldsymbol{\sigma}}_{a_i} = \hat{\boldsymbol{\sigma}}_{G,a_i} = \max_{x \in \partial a_i} \boldsymbol{\sigma}_x.$$

Hence, given $\hat{\boldsymbol{\sigma}} = \hat{\boldsymbol{\sigma}}_G = (\hat{\boldsymbol{\sigma}}_{G,a})_{a \in F}$ and G we aim to infer $\boldsymbol{\sigma}$. Thus, we can represent an inference procedure by a function $\mathscr{A}_G : \{0,1\}^m \to \{0,1\}^n$. The following theorem improves the lower bound on the number of tests required for successful inference. Let

$$m_{\mathrm{inf}} = m_{\mathrm{inf}}(n,\theta) = \max\left\{\frac{\theta}{\ln^2 2}, \frac{1-\theta}{\ln 2}\right\} n^\theta \ln n. \tag{1.3}$$

Theorem 1.1. *For any $0 < \theta < 1$, $\varepsilon > 0$ there exists $n_0 = n_0(\theta,\varepsilon)$ such that for all $n > n_0$, all test designs G with $m \le (1-\varepsilon)m_{\mathrm{inf}}$ tests and for every function $\mathscr{A}_G : \{0,1\}^m \to \{0,1\}^n$ we have*

$$\mathbb{P}\left[\mathscr{A}_G(\hat{\boldsymbol{\sigma}}_G) = \boldsymbol{\sigma}\right] < \varepsilon. \tag{1.4}$$

Theorem 1.1 rules out both deterministic and randomised test designs and inference procedures because (1.4) holds uniformly for all G and all $\mathscr{A}_G$. Thus, no test design, randomised or not, with fewer than m_{inf} tests allows to infer the set of infected individuals with a non-vanishing probability. Since m_{inf} matches $m_{\mathtt{DD}}$ from (1.2) for $\theta \ge 1/2$, Theorem 1.1 shows that the positive result from [22] is optimal in this regime. The following theorem closes the remaining gap by furnishing an optimal positive result for all θ.

Theorem 1.2. *For any $0 < \theta < 1$, $\varepsilon > 0$ there is $n_0 = n_0(\theta,\varepsilon)$ such that for every $n > n_0$ there exist a randomised test design $\boldsymbol{G}$ comprising $m \le (1+\varepsilon)m_{\mathrm{inf}}$ tests and a polynomial time algorithm* SPIV *that given $\boldsymbol{G}$ and the test results $\hat{\boldsymbol{\sigma}}_{\boldsymbol{G}}$ outputs $\boldsymbol{\sigma}$ w.h.p.*

An obvious candidate for an optimal test design appears be a plain random bipartite graph. In fact, prior to the present work the best known test design consisted of a uniformly random bipartite graph where all vertices in V_n have the same degree Δ. In other words, every individual independently joins Δ random test groups. Applied to this random Δ-out test design the DD algorithm correctly recovers the set of infected individuals in polynomial time provided that the number of tests exceeds $m_{\mathtt{DD}}$ from (1.2). However, $m_{\mathtt{DD}}$ strictly exceeds m_{inf} for $\theta < 1/2$. While the random Δ-out test design with $(1+o(1))m_{\mathrm{inf}}$ tests is known to admit an exponential time algorithm that successfully infers the set of infected individuals w.h.p. [11], we do not know of a polynomial time that solves this inference problem. Instead, to facilitate the new efficient inference algorithm SPIV the test design for Theorem 1.2 relies on a blend of a geometric and a random construction that is inspired by recent advances in coding theory known as spatially coupled low-density parity check codes [18, 26].

Finally, for

$$\theta \le \frac{\ln 2}{1+\ln 2} \approx 0.41 \tag{1.5}$$

the number m_{inf} of tests required by Theorem 1.2 matches the folklore lower bound m_{ad} from (1.2) that applies to both adaptive and non-adaptive group testing. Hence, in this regime adaptivity confers no advantage. By contrast, for $\theta > \ln(2)/(1+\ln 2)$ the adaptive bound m_{ad} is strictly smaller than m_{inf}. Consequently, in this regime at least two test stages are necessary to match the lower bound. Indeed, the next theorem shows that two stages suffice.

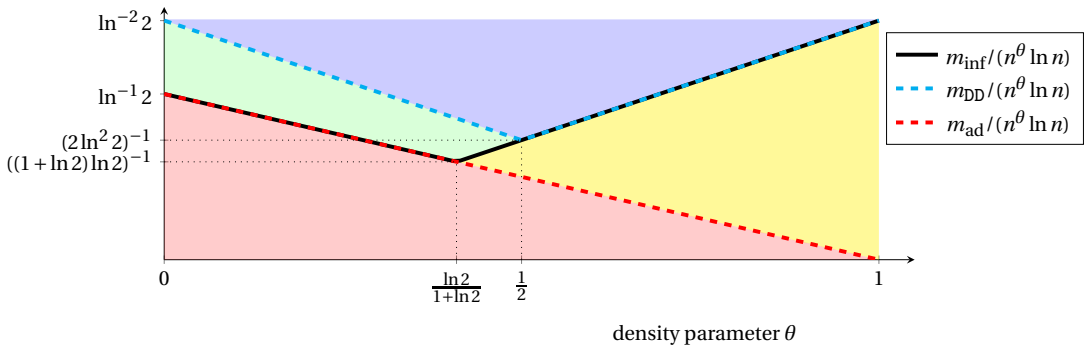

FIGURE 1. The phase transitions in group testing. The best previously known algorithm DD succeeds in the blue but not in the green region. The new algorithm SPIV succeeds in both the blue and the green region. The black line indicates the non-adaptive information-theoretic threshold $m_{\inf}$, below which non-adaptive group testing is impossible. In the red area even (multi-stage) adaptive inference is impossible. Finally, the two-stage adaptive group testing algorithm from Theorem 1.3 succeeds in the yellow region.

1.3. **Adaptive group testing.** A *two-stage test design* consists of a bipartite graph $G = (V, F)$ along with a second bipartite graph $G' = G'(G, \hat{\boldsymbol{\sigma}}_G) = (V', F')$ with $V' \subset V$ that may depend on the tests results $\hat{\boldsymbol{\sigma}}_G$ of the first test design G. Hence, the task is to learn $\boldsymbol{\sigma}$ correctly w.h.p. from $G, \hat{\boldsymbol{\sigma}}_G, G'$ and the test results $\hat{\boldsymbol{\sigma}}_{G'}$ from the second stage while minimising the total number $|F| + |F'|$ of tests. The following theorem shows that a two-stage test design and an efficient inference algorithm exist that meet the multi-stage adaptive lower bound (1.1).

Theorem 1.3. *For any $0 < \theta < 1$, $\varepsilon > 0$ there is $n_0 = n_0(\theta, \varepsilon)$ such that for every $n > n_0$ there exist a two-stage test design with no more than $(1+\varepsilon)m_{\mathrm{ad}}$ tests in total and a polynomial time inference algorithm that outputs $\boldsymbol{\sigma}$ with high probability.*

Theorem 1.3 improves over [33] by reducing the number of stages from three to two, thus potentially significantly reducing the overall time required to complete the test procedure [10, 28]. The proof of Theorem 1.3 combines the test design and efficient algorithm from Theorem 1.2 with ideas from [32].

The question of whether an 'adaptivity gap' exists for group testing, i.e., if the number of tests can be reduced by allowing multiple stages, has been raised prominently [6]. Theorems 1.1–1.3 answer this question comprehensively. While for $\theta \leq \ln(2)/(1+\ln(2)) \approx 0.41$ adaptivity confers no advantage, Theorem 1.1 shows that for $\theta > \ln(2)/(1+\ln(2))$ there is a widening gap between m_{ad} and the number $m_{\inf}$ of tests required by the optimal non-adaptive test design. Further, Theorem 1.3 demonstrates that this gap can be closed by allowing merely two stages. Figure 1 illustrates the thresholds from Theorems 1.1–1.3.

1.4. **Discussion.** The group testing problem was first raised in 1943, when Dorfman [15] proposed a two-stage adaptive test design to test the US Army for syphilis: in a first stage disjoint groups of equal size are tested. All members of negative test groups are definitely uninfected. Then, in the second stage the members of positive test groups get tested individually. Of course, this test design is far from optimal, but Dorfman's contribution triggered attempts at devising improved test schemes.

At first combinatorial group testing, where the aim is to construct a test design that is guaranteed to succeed on *all* vectors $\boldsymbol{\sigma}$, attracted significant attention. This version of the problem was studied, among others, by Erdős and Rényi [17], D'yachkov and Rykov [16] and Kautz and Singleton [23]. Hwang [20] was the first to propose an adaptive test design that asymptotically meets the information-theoretic lower bound m_{ad} from (1.1) for all $\theta \in [0,1]$. However, this test design requires an unbounded number of stages. Conversely, D'yachkov and Rykov [16] showed that m_{ad} tests do not suffice for non-adaptive group testing. Indeed, $m \geq \min\{\Omega(k^2), n\}$ tests are required non-adaptively, making individual testing optimal for $\theta > 1/2$. For an excellent survey of combinatorial group testing see [6].

Since the early 2000s attention has shifted to probabilistic group testing, which we study here as well. Thus, instead of asking for test designs and algorithms that are guaranteed to work for *all* $\boldsymbol{\sigma}$, we are content with recovering $\boldsymbol{\sigma}$ with high probability. Berger and Levenshtein [8] presented a two-stage probabilistic group testing design and

algorithm requiring

$$m_{\mathrm{BL,ad}} \sim 4n^{\theta}\ln n$$

tests in expectation. Their test design, known as the Bernoulli design, is based on a random bipartite graph where each individual joins every test independently with a carefully chosen edge probability. For a fixed θ the number $m_{\mathrm{BL,ad}}$ of tests is within a bounded factor of the information-theoretic lower bound m_{ad} from (1.1), although the gap $m_{\mathrm{ad}}/m_{\mathrm{BL,ad}}$ diverges as $\theta \to 1$. Unsurprisingly, the work of Berger and Levenshtein spurred efforts at closing the gap. Mézard, Tarzia and Toninelli proposed a different two-stage test design whose first stage consists of a random bipartite graph called the constant weight design [29]. Here each individual independently joins an equal number of random tests. For their two-stage design they obtained an inference algorithm that gets by with about

$$m_{\mathrm{MTT,ad}} \sim \frac{1-\theta}{\ln^2 2}n^{\theta}\ln n. \tag{1.6}$$

tests, a factor of $1/\ln 2$ above the elementary bound m_{ad}. Conversely, Mézard, Tarzia and Toninelli showed by means of the FKG inequality and positive correlation arguments that two-stage test algorithms from a certain restricted class cannot beat the bound (1.6). Furthermore, Aldridge, Johnson and Scarlett analysed non-adaptive test designs and inference algorithms [4, 22]. For the Bernoulli test design their best efficient algorithm DD requires

$$m_{\mathrm{DD,Be}} \sim \mathrm{e}\cdot\max\{\theta, 1-\theta\}\, n^{\theta}\ln n.$$

tests. For the constant weight design they obtained the bound m_{DD} from (1.2). In addition, in a previous article [11] we showed that on the constant weight design an exponential time algorithm correctly identifies the set of infected individuals w.h.p. if the number of tests exceeds m_{inf} from (1.3). Furthermore, Scarlett [33] discovered the aforementioned three-stage test design and polynomial time algorithm that matches the universal lower bound m_{ad} from (1.1). Finally, concerning lower bounds, in the case of a linear number $k = \Theta(n)$ infected individuals Aldridge [5] showed via arguments similar to [29] that individual testing is optimal in the non-adaptive case, while Ungar [35] proved that individual testing is optimal even adaptively once $k \geq (3-\sqrt{5})n/2$.

A further variant of group testing is known as the quantitative group testing or the coin weighing problem. In this problem tests are assumed to not merely indicate the presence of at least one infected individual but to return the number of infected individuals. Thus, the tests are significantly more powerful. For quantitative group testing with k infected individuals Alaoui, Ramdas, Krzakala, Zdeborová and Jordan [3] presented a test design with

$$m_{\mathrm{QGT}} \sim 2\left(1+\frac{(n-k)\ln(1-k/n)}{k\ln(k/n)}\right)\frac{k\ln(n/k)}{\ln(k)} \qquad \text{for} \qquad k=\Theta(n)$$

tests from which the set of infected individuals can be inferred in exponential time; the paper actually deals with the slightly more general pooled data problem. However, no efficient algorithm is known to come within a constant factor of m_{QGT}. Indeed, the best efficient algorithm, due to the same authors [2], requires $\Omega(k\ln(n/k))$ tests.

More broadly, the idea of harnessing random graphs to tackle inference problems has been gaining momentum. One important success has been the development of capacity achieving linear codes called spatially coupled low-density parity check ('LDPC') codes [26, 27]. The Tanner graphs of these codes, which represent their check matrices, consist of a linear sequence of sparse random bipartite graphs with one class of vertices corresponding to the bits of the codeword and the other class corresponding to the parity checks. The bits and the checks are divided equitably into a number of compartments, which are arranged along a line. Each bit of the codeword takes part in random checks in a small number of preceding and subsequent compartments of checks along the line. This combination of a spatial arrangement and randomness facilitates efficient decoding by means of the Belief Propagation message passing algorithm. Furthermore, the general design idea of combining a linear spatial structure with a random graph has been extended to other inference problems. Perhaps the most prominent example is compressed sensing, i.e., solving an underdetermined linear system subject to a sparsity constraint [13, 14, 24, 25], where a variant of Belief Propagation called Approximate Message Passing matches an information-theoretic lower bound from [37].

While in some inference problems such as LDPC decoding or compressed sensing the number of queries required to enable an efficient inference algorithm matches the information-theoretic lower bound, in many other problems gaps remain. A prominent example is the stochastic block model [1, 12, 30], an extreme case of which is the notorious planted clique problem [7]. For both these models the existence of a genuine computationally

intractable phase where the problem can be solved in exponential but not in polynomial time appears to be an intriguing possibility. Further examples include code division multiple access [34, 38], quantitative group testing [2], sparse principal component analysis [9] and sparse high-dimensional regression [31]. The problem of solving the group testing inference problem on the test design from [22] could be added to the list. Indeed, while an exponential time algorithm (that reduces the problem to minimum hypergraph vertex cover) infers the set of infected individuals w.h.p. with only $(1+\varepsilon)m_{\inf}$ tests, the best known polynomial algorithm requires $(1+\varepsilon)m_{\mathtt{DD}}$ tests.

Instead of developing a better algorithm for the test design from [22], here we exercise the discretion of constructing a different test design that the group testing problem affords. The new design is tailored to enable an efficient algorithm SPIV for Theorem 1.2 that gets by with $(1+\varepsilon)m_{\inf}$ tests. While prior applications of the idea of spatial coupling such as coding and compressed sensing required sophisticated message passing algorithms [18, 26, 27], the SPIV algorithm is purely combinatorial and extremely transparent. The main step of the algorithm merely computes a weighted sum to discriminate between infected individuals and 'disguised' healthy individuals. Furthermore, the analysis of the algorithm is based on a technically subtle but conceptually clean large deviations analysis. This technique of blending combinatorial ideas and large deviations methods with spatial coupling promises to be an exciting route for future research. Applications might include noisy versions of group testing, quantitative group testing or the coin weighing problem [2]. Beyond these immediate extensions, it would be most interesting to see if the SPIV strategy extends to other inference problems for sparse data.

1.5. **Organisation.** After collecting some preliminaries and introducing notation in Section 2, we prove Theorem 1.1 in Section 3. Section 4 then deals with the test design and the inference algorithm for Theorem 1.2. Finally, in Section 5 we prove Theorem 1.3.

2. Preliminaries

As we saw in Section 1.2 a non-adaptive test design can be represented by a bipartite graph $G = (V \cup F, E)$ with one vertex class V representing the individuals and the other class F representing the tests. We refer to the number $|V|$ of individuals as the *order* of the test design and to the number $|F|$ of tests as its *size*. For a vertex v of G we denote by $\partial_G v$ the set of neighbours. Where G is apparent from the notation we just write ∂v. Furthermore, for an integer $k \le |V|$ we denote by $\boldsymbol{\sigma}_{G,k} = (\boldsymbol{\sigma}_{G,k,x})_{x\in V} \in \{0,1\}^V$ a random vector of Hamming weight k. Additionally, we let

$$\hat{\boldsymbol{\sigma}}_{G,k} = (\hat{\boldsymbol{\sigma}}_{G,k,a})_{a\in F} \in \{0,1\}^F \qquad \text{with} \qquad \hat{\boldsymbol{\sigma}}_{G,k,a} = \max_{x\in\partial_G a} \boldsymbol{\sigma}_{G,k,x} \tag{2.1}$$

be the associated vector of test results. Where G and/or k are apparent from the context, we drop them from the notation. More generally, for a given vector $\tau \in \{0,1\}^V$ we introduce a vector $\hat{\tau}_G = (\hat{\tau}_{G,a})_{a\in F}$ by letting $\hat{\tau}_{G,a} = \max_{x\in\partial_G a}\tau_x$, just as in (2.1). Furthermore, for a given $\tau \in \{0,1\}^V$ we let

$$V_0(G,\tau) = \{x \in V : \tau_x = 0\}, \quad V_1(G,\tau) = \{x \in V : \tau_x = 1\}, \quad F_0(G,\tau) = \left\{a \in F : \hat{\tau}_{G,a} = 0\right\}, \quad F_1(G,\tau) = \left\{a \in F : \hat{\tau}_{G,a} = 1\right\}.$$

The *Kullback-Leibler divergence* of $p, q \in (0,1)$ is denoted by

$$D_{\mathrm{KL}}\left(q \| p\right) = q\ln\left(\frac{q}{p}\right) + (1-q)\ln\left(\frac{1-q}{1-p}\right).$$

We will occasionally apply the following Chernoff bound.

Lemma 2.1 ([21])**.** *Let $\boldsymbol{X}$ be a binomial random variable with parameters N, p. Then*

$$\mathbb{P}\left[X \ge qN\right] \le \exp\left(-ND_{\mathrm{KL}}\left(q\|p\right)\right) \quad \text{for } p < q < 1, \tag{2.2}$$

$$\mathbb{P}\left[X \le qN\right] \le \exp\left(-ND_{\mathrm{KL}}\left(q\|p\right)\right) \quad \text{for } 0 < q < p. \tag{2.3}$$

In addition, we recall that the *hypergeometric distribution* $\mathrm{Hyp}(L,M,N)$ is defined by

$$\mathbb{P}\left[\mathrm{Hyp}(L,M,N) = k\right] = \binom{M}{k}\binom{L-M}{N-k}\binom{L}{N}^{-1}. \qquad (k \in \{0,1,\ldots,M\wedge N\}).$$

Hence, out of a total of L items of which M are special we draw N items without replacement and count the number of special items in the draw. The mean of the hypergeometric distribution equals MN/L. It is well known that the Chernoff bound extends to the hypergeometric distribution.

Lemma 2.2 ([19])**.** *For a hypergeometric variable $\boldsymbol{X} \sim \mathrm{Hyp}(L,M,N)$ the bounds* (2.2)–(2.3) *hold with $p = M/L$.*

Throughout the paper we use asymptotic notation $o(\cdot), \omega(\cdot), O(\cdot), \Omega(\cdot), \Theta(\cdot)$ to refer to limit $n \to \infty$. It is understood that the constants hidden in, e.g., a $O(\cdot)$-term may depend on the density parameter θ or other parameters.

3. The information theoretic lower bound

In this section we prove Theorem 1.1. The proof combines techniques based on the FKG inequality and positive correlation that were developed in [6, 29] with new combinatorial ideas. Throughout this section we fix a number $\theta \in (0,1)$ and we let $k = \lceil n^\theta \rceil$.

3.1. **Outline.** The starting point is a simple and well known observation. Namely, for a test design $G = G_{n,m} = (V_n, F_m)$ and a vector $\tau \in \{0,1\}^{F_m}$ of test results let

$$\mathscr{S}_k(G,\tau) = \left\{\sigma \in \{0,1\}^{V_n} : \sum_{x \in V_n} \sigma_x = k,\ \hat{\sigma}_G = \tau\right\}$$

be the set of all possible vectors σ of Hamming weight k that give rise to the test results τ. Further, let $Z_k(G,\tau) = |\mathscr{S}_k(G,\tau)|$ be the number of such vectors σ. Also recall that $\boldsymbol{\sigma} = \boldsymbol{\sigma}_{G,k} \in \{0,1\}^{V_n}$ is a random vector of Hamming weight k and that $\hat{\boldsymbol{\sigma}} = \hat{\boldsymbol{\sigma}}_{G,k}$ comprises the test results that $\boldsymbol{\sigma}$ renders under the test design G. We observe that the posterior of $\boldsymbol{\sigma}$ given $\hat{\boldsymbol{\sigma}}$ is the uniform distribution on $\mathscr{S}_k(G,\hat{\boldsymbol{\sigma}})$.

Fact 3.1. *For any $G, \sigma \in \{0,1\}^{V_n}$ we have $\mathbb{P}[\boldsymbol{\sigma} = \sigma \mid \hat{\boldsymbol{\sigma}}] = \mathbf{1}\{\sigma \in \mathscr{S}_k(G,\hat{\boldsymbol{\sigma}})\}/Z_k(G,\hat{\boldsymbol{\sigma}})$.*

As an immediate consequence of Fact 3.1, the success probability of any inference scheme $\mathscr{A}_G : \{0,1\}^{F_m} \to \{0,1\}^{V_n}$ is bounded by $1/Z_k(G,\hat{\boldsymbol{\sigma}})$. Indeed, an optimal inference algorithm is to simply return a uniform sample from $\mathscr{S}_k(G,\hat{\boldsymbol{\sigma}})$.

Fact 3.2. *For any test design G and for any $\mathscr{A}_G : \{0,1\}^{F_m} \to \{0,1\}^{V_n}$ we have $\mathbb{P}[\mathscr{A}_G(\hat{\boldsymbol{\sigma}}) = \boldsymbol{\sigma} \mid \hat{\boldsymbol{\sigma}}] \le 1/Z_k(G,\hat{\boldsymbol{\sigma}})$.*

Hence, in order to prove Theorem 1.1 we just need to show that $Z_k(G,\hat{\boldsymbol{\sigma}})$ is large for any test design G with $m < (1-\varepsilon)m_{\inf}$ tests. In other words, we need to show that w.h.p. there are many vectors $\sigma \in \mathscr{S}_k(G,\hat{\boldsymbol{\sigma}})$ that give rise to the test results $\hat{\boldsymbol{\sigma}}$.

We obtain these σ by making diligent local changes to $\boldsymbol{\sigma}$. More precisely, we identify two sets $V_{0+} = V_{0+}(G,\boldsymbol{\sigma})$, $V_{1+} = V_{1+}(G,\boldsymbol{\sigma})$ of individuals whose infection status can be flipped without altering the test results. Specifically, following [5] we call an individual $x \in V_n$ *disguised* if every test $a \in \partial_G x$ contains another individual $y \in \partial_G a \setminus \{x\}$ with $\boldsymbol{\sigma}_y = 1$. Let $V_+ = V_+(G,\boldsymbol{\sigma})$ be the set of all disguised individuals. Moreover, let

$$V_{0+} = V_{0+}(G,\boldsymbol{\sigma}) = \{x \in V_+ : \boldsymbol{\sigma}_x = 0\}, \qquad V_{1+} = V_{1+}(G,\boldsymbol{\sigma}) = \{x \in V_+ : \boldsymbol{\sigma}_x = 1\}. \tag{3.1}$$

Hence, V_{0+} is the set of all healthy disguised individuals while V_{1+} contains all infected disguised individuals.

Fact 3.3. *We have $Z_k(G,\hat{\boldsymbol{\sigma}}) \ge |V_{0+}(G,\boldsymbol{\sigma})| \cdot |V_{1+}(G,\boldsymbol{\sigma})|$.*

Proof. For a pair $(x,y) \in V_{0+}(G,\boldsymbol{\sigma}) \times V_{1+}(G,\boldsymbol{\sigma})$ obtain $\boldsymbol{\tau}$ from $\boldsymbol{\sigma}$ by letting $\boldsymbol{\tau}_x = 1, \boldsymbol{\tau}_y = 0$ and $\boldsymbol{\tau}_z = \boldsymbol{\sigma}_z$ for all $z \neq x, y$. Then $\boldsymbol{\tau}$ has Hamming weight k and $\hat{\boldsymbol{\tau}}_G = \hat{\boldsymbol{\sigma}}$. Thus, $\boldsymbol{\tau} \in \mathscr{S}_k(G,\hat{\boldsymbol{\sigma}})$. □

Hence, an obvious proof strategy for Theorem 1.1 is to exhibit a large number of disguised individuals. A similar strategy has been pursued in the proof of the conditional lower bound of Mézard, Tarzia and Toninelli [29] and the proof of Aldridge's lower bound for the linear case $k = \Theta(n)$ [5]. Both [5, 29] exhibit disguised individuals via positive correlation and the FKG inequality. However, we do not see how to stretch such arguments to obtain the desired lower bound for all $\theta \in (0,1)$. Yet for θ *extremely* close to one it is possible to combine the positive correlation argument with new combinatorial ideas to obtain the following.

Proposition 3.4. *For any $\varepsilon > 0$ there exists $\theta_0 = \theta_0(\varepsilon) < 1$ such that for every $\theta \in (\theta_0, 1)$ there exists $n_0 = n_0(\theta,\varepsilon)$ such that for all $n > n_0$ and all test designs $G = G_{n,m}$ with $m \le (1-\varepsilon)m_{\inf}$ we have*

$$\mathbb{P}[|V_{0+}(G,\boldsymbol{\sigma})| \wedge |V_{1+}(G,\boldsymbol{\sigma})| \ge \ln n] > 1 - \varepsilon.$$

The proof of Proposition 3.4 can be found in Section 3.2.

The second step towards Theorem 1.1 is a reduction from larger to smaller values of θ. Suppose we wish to apply a test scheme designed for an infection density $\theta \in (0,1)$ to a larger infection density $\theta' \in (\theta, 1)$. Then we could dilute the larger infection density by adding a large number of healthy 'dummy' individuals. A careful analysis of this dilution process yields the following result. Due to the elementary lower bound (1.1) we need not worry about $\theta \le \ln(2)/(1+\ln 2)$.

Proposition 3.5. *For any* $\ln(2)/(1+\ln(2)) < \theta < \theta' < 1$, $t > 0$ *there exists* $n_0 = n_0(\theta,\theta',t) > 0$ *such that for every* $n > n_0$ *and for every test design* G *of order* n *there exist an integer* n' *such that*

$$k = \lceil n^{\theta} \rceil = \lceil n'^{\theta'} \rceil$$

and a test design G' *of order* n' *with the same number of tests as* G *such that the following is true. Let* $\boldsymbol{\tau} \in \{0,1\}^{V_{n'}}$ *be a random vector of Hamming weight* k *and let* $\hat{\boldsymbol{\tau}}_a = \max_{x \in \partial_{G'} a} \boldsymbol{\tau}_x$ *comprise the tests results of* G'*. Then*

$$\mathbb{P}\left[Z_k(G,\hat{\boldsymbol{\sigma}}) \le t\right] \le \mathbb{P}\left[Z_k(G',\hat{\boldsymbol{\tau}}) \le t\right].$$

Hence, if a test design exists for $\theta < \theta'$ that beats $m_{\inf}(n,\theta)$, then there is a test design for infection density θ' that beats $m_{\inf}(n',\theta')$. We prove Proposition 3.4 in Section 3.2. Theorem 1.1 is an easy consequence of Propositions 3.4 and 3.5.

Proof of Theorem 1.1. For $\theta \le \ln(2)/(1+\ln(2))$ the assertion follows from the elementary lower bound (1.1). Hence, fix $\varepsilon > 0$ and assume for contradiction that some $\theta \in (\ln(2)/(1+\ln(2)),1)$ for infinitely many n admits a test design G of order n and size $m \le (1-\varepsilon)m_{\inf}(n,\theta)$ such that $\mathbb{P}\left[Z_k(G,\hat{\boldsymbol{\sigma}}_G) \le t\right] \ge \varepsilon$. Then Proposition 3.5 shows that for $\theta' > \theta$ arbitrarily close to one for an integer n' with $k = \lceil n'^{\theta'} \rceil$ a test design $G' = G_{n',m}$ exists such that

$$\mathbb{P}\left[Z_k(G',\hat{\boldsymbol{\tau}}) \le 1/\varepsilon\right] \ge \varepsilon. \tag{3.2}$$

Furthermore, (1.3) shows that for large n,

$$m_{\inf}(n',\theta') = \frac{\theta'}{\ln^2 2} n'^{\theta'} \ln n' = \frac{\theta + o(1)}{\ln^2 2} n^{\theta} \ln n = (1+o(1)) m_{\inf}(n,\theta).$$

Hence, the number m of tests of G' satisfies $m \le (1-\varepsilon+o(1))m_{\inf}(n',\theta')$. Thus, (3.2) contradicts Fact 3.3 and Proposition 3.4. □

3.2. **Proof of Proposition 3.4.** Given a small $\varepsilon > 0$ we choose $\theta_0 = \theta_0(\varepsilon) \in (0,1)$ sufficiently close to one and fix $\theta \in (\theta_0,1)$. Additionally, pick $\xi = \xi(\varepsilon,\theta) \in (0,1)$ such that

$$2(1-\theta) < \xi < \theta\varepsilon. \tag{3.3}$$

We fix ε,θ,ξ throughout this section.

To avoid the (mild) stochastic dependencies that result from the total number of infected individuals being fixed, instead of $\boldsymbol{\sigma}$ we will consider a vector $\boldsymbol{\chi} \in \{0,1\}^{V_n}$ whose entries are stochastically independent. Specifically, every entry of $\boldsymbol{\chi}$ equals one with probability

$$p = \frac{k - \sqrt{k}\ln n}{n}$$

independently. Let $\hat{\boldsymbol{\chi}}_G \in \{0,1\}^{F_m}$ be the corresponding vector of test results. The following lemma shows that it suffices to estimate $|V_{0+}(G,\boldsymbol{\chi})|, |V_{1+}(G,\boldsymbol{\chi})|$. Let G denote an arbitrary test design with individuals $V_n = \{x_1,\dots,x_n\}$ and tests $F_m = \{a_1,\dots,a_m\}$.

Lemma 3.6. *There is* $n_0 = n_0(\theta,\varepsilon)$ *such that for all* $n > n_0$ *and for all* G *with* $m \le m_{\inf}$ *the following is true:*

$$\textit{if } \mathbb{P}\left[|V_{0+}(G,\boldsymbol{\chi})| \wedge |V_{1+}(G,\boldsymbol{\chi})| \ge 2\ln n\right] > 1-\varepsilon/4, \textit{ then } \mathbb{P}\left[|V_{0+}(G,\boldsymbol{\sigma})| \wedge |V_{1+}(G,\boldsymbol{\sigma})| \ge \ln n\right] > 1-\varepsilon.$$

Proof. Let $\mathscr{X} = \{k - 2\sqrt{k}\ln n \le \sum_{x \in V_n} \boldsymbol{\chi}_x \le k\}$. The Chernoff bound shows for large enough n,

$$\mathbb{P}[\mathscr{X}] > 1-\eta/4. \tag{3.4}$$

Further, given $\mathscr{X}$ we can couple $\boldsymbol{\chi},\boldsymbol{\sigma}$ such that the latter is obtained by turning $k - \sum_{x \in V_n} \boldsymbol{\chi}_x$ random zero entries of the former into ones. Since turning zero entries into ones can only increase the number of disguised individuals, on $\mathscr{X}$ we have

$$V_{1+}(G,\boldsymbol{\sigma}) \ge V_{1+}(G,\boldsymbol{\chi}). \tag{3.5}$$

Of course, it is possible that $|V_{0+}(G,\boldsymbol{\sigma})| < |V_{0+}(G,\boldsymbol{\chi})|$. But since on $\mathscr{X}$ the two vectors $\boldsymbol{\sigma},\boldsymbol{\chi}$ differ in no more than $2\sqrt{k}\ln n$ entries, we obtain the bound

$$\mathbb{E}\left[|V_{0+}(G,\boldsymbol{\chi})| - |V_{0+}(G,\boldsymbol{\sigma})| \,|\, \mathscr{X}\right] \le \frac{2\sqrt{k}\ln n}{n-k} |V_{0+}(G,\boldsymbol{\chi})| < n^{-1/3} |V_{0+}(G,\boldsymbol{\chi})|,$$

provided n is sufficiently large. Hence, Markov's inequality shows that for large enough n,

$$\mathbb{P}\left[|V_{0+}(G,\boldsymbol{\chi})| - |V_{0+}(G,\boldsymbol{\sigma})| > |V_{0+}(G,\boldsymbol{\chi})|/2 \mid \mathscr{X}\right] < \varepsilon/4. \tag{3.6}$$

Combining (3.4), (3.5) and (3.6) completes the proof. □

As a next step we show that there is no point in having very big tests a that contain more than, say, $\Gamma = \Gamma(n,\theta) = n^{1-\theta}\ln n$ individuals. This is because anyway all such tests are positive w.h.p., so there is little point in actually conducting them. Indeed, the following lemma shows that w.h.p. all tests of very high degree contain at least two infected individuals.

Lemma 3.7. *There exists $n_0 = n_0(\theta,\varepsilon) > 0$ such that for all $n > n_0$ and all test designs G with $m \le m_{\inf}$ tests,*

$$\mathbb{P}\left[\exists a \in F_m : |\partial_G a| > \Gamma \wedge |\partial_G a \cap V_1(G,\boldsymbol{\chi})| \le 1\right] < \varepsilon/8.$$

Proof. Consider a test a of degree $\gamma = |\partial_G a| \ge \Gamma$. Because in $\boldsymbol{\chi}$ each of the γ individuals that take part in a is infected with probability p independently, we have

$$\mathbb{P}\left[|\partial_G a \cap V_1(G,\boldsymbol{\sigma})| \le 1\right] = \mathbb{P}\left[\mathrm{Bin}(\gamma,p) \le 1\right] = (1-p)^\gamma + \gamma p(1-p)^{\gamma-1} \le (1+\gamma p/(1-p))\exp(-\gamma p) = n^{o(1)-1}. \tag{3.7}$$

Since $m \le m_{\inf} = O(n^\theta)$ for a fixed $\theta < 1$, the assertion follows from (3.7) and the union bound. □

Let G^* be test design obtained from $G = G_{n,m}$ by deleting all tests of degree larger than Γ. If indeed every test of degree at least Γ contains at least two infected individuals, then $V_{0+}(G^*,\boldsymbol{\chi}) = V_{0+}(G,\boldsymbol{\chi})$ and $V_{1+}(G^*,\boldsymbol{\chi}) = V_{1+}(G,\boldsymbol{\chi})$. Hence, Lemma 3.7 shows that it suffices to bound $|V_{0+}(G^*,\boldsymbol{\chi})|, |V_{1+}(G^*,\boldsymbol{\chi})|$. To this end we observe that G^* contains few individuals of very high degree.

Lemma 3.8. *There is $n_0 = n_0(\theta,\varepsilon) > 0$ such that for all $n > n_0$ and all test designs G with $m \le m_{\inf}$ we have*

$$\left|\left\{x \in V_n : |\partial_{G^*} x| > \ln^3 n\right\}\right| \le \frac{n\ln\ln n}{\ln n}.$$

Proof. Since $\max_{a\in F_m}|\partial_{G^*} a| \le \Gamma = n^{1-\theta}\ln n$, double counting yields

$$\sum_{x\in V_n}|\partial_{G^*} x| = \sum_{a\in F_m}|\partial_{G^*} a| \le m_{\inf}\Gamma = O(n\ln^2 n).$$

Consequently, there are no more than $O(n/\ln n)$ individuals $x \in V_n$ with $|\partial_{G^*} x| > \ln^3 n$. □

Further, obtain $G^{(0)}$ from G^* by deleting all individuals of degree greater than $\ln^3 n$ (but keeping all tests). Then the degrees of $G^{(0)}$ satisfy

$$\max_{a\in F(G^{(0)})}|\partial_{G^{(0)}} a| \le \Gamma, \qquad \max_{x\in V(G^{(0)})}|\partial_{G^{(0)}} x| \le \ln^3 n. \tag{3.8}$$

Let $\boldsymbol{\chi}^{(0)} = (\boldsymbol{\chi}_x)_{x\in V(G^{(0)})}$ signify the restriction of $\boldsymbol{\chi}$ to the individuals that remain in $G^{(0)}$.

With these preparations in place we are ready to commence the main step of the proof of Proposition 3.4. Given a test design G with $m \le (1-\varepsilon)m_{\inf}$ we are going to construct a sequence $y_1, y_2, \ldots, y_N$, $N = \lceil n^{1-\xi}\rceil$, of individuals of $G^{(0)}$ such that each y_i individually has a moderately high probability of being disguised. Of course, to conclude that in the end a large number of disguised y_i actually materialise, we need to cope with stochastic dependencies. To this end we will pick individuals y_i that have pairwise distance at least five in $G^{(0)}$. The degree bounds (3.8) guarantee a sufficient supply of such far apart individuals.

To be precise, starting from $G^{(0)}$ we construct a sequence of test designs $G^{(1)}, G^{(2)}, \ldots, G^{(N)}$ inductively as follows. For each $i \ge 1$ select a variable $y_{i-1} \in V(G^{(i-1)})$ whose probability of being disguised is maximum; ties are broken arbitrarily. In formulas,

$$\mathbb{P}\left[y_{i-1} \in V_+(G^{(i-1)},\boldsymbol{\chi}^{(i-1)})\right] = \max_{y\in V(G^{(i-1)})}\mathbb{P}\left[y \in V_+(G^{(i-1)},\boldsymbol{\chi}^{(i-1)})\right],$$

where, of course, $\boldsymbol{\chi}^{(i-1)}$ is the only random object. Then obtain $G^{(i)}$ from $G^{(i-1)}$ by removing y_{i-1} along with all vertices (i.e., tests or individuals) at distance at most four from y_{i-1}. Moreover, let $\boldsymbol{\chi}^{(i)}$ denote the restriction $(\boldsymbol{\chi}_x)_{x\in V(G^{(i)})}$ of $\boldsymbol{\chi}$ to $G^{(i)}$. The following lemma estimates the probability of y_i being disguised. Let $m^* = |F(G^*)|$ be the total number of tests of G of degree at most Γ.

Lemma 3.9. *There exists $n_0 = n_0(\varepsilon,\theta,\xi)$ such that for all $n > n_0$ and all G with $m \le (1-\varepsilon)m_{\inf}$ we have*

$$\min_{1\le i\le N} \mathbb{P}\left[y_i \in V_+(G^{(i)})\right] \ge \exp\left(-\frac{m\ln^2 2}{n^\theta} - 1\right).$$

The proof of Lemma 3.9 requires three intermediate steps. First, we need a lower bound on number of individuals in $G^{(i)}$. Recall that $N = \lceil n^{1-\xi}\rceil$.

Claim 3.10. *We have* $\min_{0\le i\le N} |V(G^{(i)})| \ge n - N\Gamma^2\ln^6 n$.

Proof. Since throughout the construction of the $G^{(i)}$ we only delete vertices, the degree bound (3.8) implies

$$\max_{a\in F(G^{(i)})} |\partial_{G^{(i)}} a| \le \Gamma = n^{1-\theta}\ln n, \qquad \max_{x\in V(G^{(i)})} |\partial_{G^{(i)}} x| \le \ln^3 n \qquad \text{for all } i \le N. \tag{3.9}$$

We now proceed by induction on i. For $i = 0$ there is nothing to show. Going from i to $i+1 \le N$, we notice that because all individuals $x \in V(G^{(i)}) \setminus V(G^{(i+1)})$ have distance at most four from y_{i+1}, (3.9) ensures that

$$|V(G^{(i)}) \setminus V(G^{(i+1)})| \le \Gamma^2\ln^6 n. \tag{3.10}$$

Iterating (3.10), we obtain $|V(G^{(0)}) \setminus V(G^{(i+1)})| \le (i+1)\Gamma^2\ln^6 n$, whence $|V(G^{(i+1)})| \ge n - (i+1)\Gamma^2\ln^6 n$. □

The following claim resembles the proof of [5, Theorem 1] (where the case $k = \Omega(n)$ is considered).

Claim 3.11. *Let $\mathscr{D}^{(i)}(x) = \{x \in V_+(G^{(i)})\}$ and let*

$$L^{(i)} = \frac{1}{|V(G^{(i)})|}\sum_{x\in V(G^{(i)})} \ln\mathbb{P}\left[\mathscr{D}^{(i)}(x)\right]. \tag{3.11}$$

Then

$$L^{(i)} \ge \frac{|F(G^{(i)})|}{|V(G^{(i)})|}\min_{a\in F(G^{(i)})} \left|\partial_{G^{(i)}} a\right| \ln\left(1-(1-p)^{|\partial_{G^{(i)}} a|-1}\right). \tag{3.12}$$

Proof. For an individual $x \in V(G^{(i)})$ and a test $a \in \partial_{G^{(i)}} x$ let $\mathscr{D}^{(i)}(x,a)$ be the event that there is another individual $z \in \partial_{G^{(i)}} a \setminus \{x\}$ such that $\boldsymbol{\chi}_z = 1$. Then for every $x \in V(G^{(i)})$ we have

$$\mathbb{P}\left[\mathscr{D}^{(i)}(x)\right] = \mathbb{P}\left[\bigcap_{a\in\partial_{G^{(i)}} x} \mathscr{D}^{(i)}(x,a)\right]. \tag{3.13}$$

Furthermore, the events $\mathscr{D}^{(i)}(x,a)$ are increasing with respect to $\boldsymbol{\chi}$. Therefore, (3.13) and the FKG inequality imply

$$\mathbb{P}\left[\mathscr{D}^{(i)}(x)\right] \ge \prod_{a\in\partial_{G^{(i)}} x} \mathbb{P}\left[\mathscr{D}^{(i)}(x,a)\right]. \tag{3.14}$$

Moreover, because each entry of $\boldsymbol{\chi}$ is one with probability p independently, we obtain

$$\mathbb{P}\left[\mathscr{D}^{(i)}(x,a)\right] = 1-(1-p)^{|\partial_{G^{(i)}} a|-1} \tag{3.15}$$

Finally, combining (3.13)–(3.15), we obtain

$$\begin{aligned}|V(G^{(i)})|L^{(i)} &\ge \sum_{x\in V(G^{(i)})}\sum_{a\in F(G^{(i)})} \mathbf{1}\left\{a\in\partial_{G^{(i)}} x\right\}\ln\left(1-(1-p)^{|\partial_{G^{(i)}} a|-1}\right)\\ &= \sum_{a\in F(G^{(i)})} \left|\partial_{G^{(i)}} a\right|\ln\left(1-(1-p)^{|\partial_{G^{(i)}} a|-1}\right) \ge |F(G^{(i)})|\min_{a\in F(G^{(i)})}\left|\partial_{G^{(i)}} a\right|\ln\left(1-(1-p)^{|\partial_{G^{(i)}} a|-1}\right),\end{aligned}$$

as claimed. □

As a final preparation for the proof of Lemma 3.9 we need the following estimate.

Claim 3.12. *The function $z \in (0,\infty) \mapsto z\ln(1-(1-p)^{z-1})$ attains its minimum at $z = \left(1+O(n^{-\Omega(1)})\right)\ln(2)/p$.*

Proof. We consider three separate cases.

Case 1: $z = o(1/p)$: we obtain

$$\begin{aligned}z\ln\left(1-(1-p)^{z-1}\right) &= z\ln\left(1-\exp\left(-pz+O(p^2z)\right)\right) = z\ln\left(1-\left(1-pz+O(p^2z^2)\right)\right)\\ &= \frac{z}{\ln}(zp+O(zp)^2) = o(1/p).\end{aligned} \tag{3.16}$$

Case 2: $z = \omega(1/p)$**:** we find

$$z\ln\bigl(1-(1-p)^{z-1}\bigr) = z\ln\bigl(1-\exp\bigl(-pz+O(p^2z)\bigr)\bigr) = -z\bigl(\exp(-pz)+O\bigl(\exp(-2pz)\bigr)\bigr)$$
$$= -\frac{1}{p}pz\bigl(\exp\bigl(-pz\bigr)+\exp\bigl(-2pz\bigr)\bigr) = o(1/p). \tag{3.17}$$

Case 3: $z = \Theta(1/p)$**:** letting $d = zp$, we obtain

$$z\ln\bigl(1-(1-p)^{z-1}\bigr) = \frac{d}{p}\ln\bigl(1-\exp\bigl(-d+O(p)\bigr)\bigr) = \frac{d}{p}\ln\bigl(1-\exp\bigl(-d\bigr)\bigr)+O(1). \tag{3.18}$$

Since the strictly convex function $d \in (0,\infty) \mapsto d\ln(1-\exp(-d))$ attains its minimum at $d = \ln 2$, (3.18) dominates (3.16) and (3.17). Thus, the minimiser reads $z = \ln(2)/p + O(p^{-1/2})$. □

Proof of Lemma 3.9. Combining Claims 3.11 and 3.12, we see that for all test designs G with $m \le (1-\varepsilon)m_{\inf}$ and for all $i \le N$,

$$L^{(i)} \ge -\bigl(1+O(n^{-\Omega(1)})\bigr)\frac{|F(G^{(i)})|\ln^2 2}{|V(G^{(i)})|p} \ge -\bigl(1+O(n^{-\Omega(1)})\bigr)\frac{m\ln^2 2}{|V(G^{(i)})|p}.$$

Hence, Claim 3.10, (3.3) and the choice $p = (k+\sqrt{k}\ln n)/n$ imply that for all $i \le N$,

$$L^{(i)} \ge -\bigl(1+O(n^{-\Omega(1)})\bigr)\frac{m\ln^2 2}{(n-N\Delta^2\ln^6 n)p} \ge -\bigl(1+O(n^{-\Omega(1)})\bigr)\frac{m\ln^2 2}{n^\theta}. \tag{3.19}$$

Further, combining the definition (3.11) of $L^{(i)}$ with (3.19), we conclude that for every $i \le N$ there exists an individual $y_i \in V(G^{(i)})$ such that

$$\mathbb{P}\Bigl[y_i \in V_+(G^{(i)})\Bigr] = \mathbb{P}\Bigl[\mathscr{D}^{(i)}(y_i)\Bigr] \ge \exp\Bigl(L^{(i)}\Bigr) \ge \exp\biggl(-\bigl(1+O(n^{-\Omega(1)})\bigr)\frac{m\ln^2(2)}{n^\theta}\biggr),$$

which implies the assertion. □

Lemma 3.9 implies the following bound on $|V_{0+}(G^*,\boldsymbol{\chi})|$, $|V_{1+}(G^*,\boldsymbol{\chi})|$.

Corollary 3.13. *There exists $n_0 = n_0(\varepsilon,\theta,\xi)$ such that for all $n > n_0$ and all $G = G_{n,m}$ with $m \le (1-\varepsilon)m_{\inf}$ we have*

$$\mathbb{P}\bigl[\bigl|V_{0+}(G^*,\boldsymbol{\chi})\bigr| \wedge \bigl|V_{1+}(G^*,\boldsymbol{\chi})\bigr| < \ln^4 n\bigr] < \varepsilon/8.$$

Proof. We observe that $V_+(G^{(i)},\boldsymbol{\chi}) \subset V_+(G^*,\boldsymbol{\chi})$ for all $i \le N$ because by construction for any individual $x \in V(G^{(i)})$ every test $a \in \partial_{G^*}x$ of G^* that x belongs to is still present in $G^{(i)}$. Consequently, we obtain the bound

$$\mathbb{P}\bigl[x \in V_+(G^*)\bigr] \ge \mathbb{P}\Bigl[x \in V(G^{(i)})\Bigr] \qquad \text{for all } i \in [N],\ x \in V(G^*). \tag{3.20}$$

Combining (3.20) with Lemma 3.9 we obtain

$$\mathbb{P}\Bigl[y^{(i)} \in V_+(G^*)\Bigr] \ge \exp\Bigl(-\ln^2(2)n^{-\theta}m-1\Bigr) \ge \exp\Bigl(-(1-\varepsilon)\ln^2(2)n^{-\theta}m_{\inf}-1\Bigr) \qquad \text{for all } i \in [N].$$

Hence, recalling the definition of $m_{\inf}$ from (1.3), we obtain

$$\mathbb{P}\Bigl[y^{(i)} \in V_+(G^*)\Bigr] \ge \exp\bigl(-(1-\varepsilon)\theta\ln(n)-1\bigr) = n^{(\varepsilon-1)\theta}/\mathrm{e}. \qquad \text{for all } i \in [N]. \tag{3.21}$$

Since the entry $\boldsymbol{\chi}_{y^{(i)}}$ is independent of the event $\{y^{(i)} \in V_+(G^*)\}$, the definitions (3.1) of $V_{0+}(G^*,\boldsymbol{\chi})$ and $V_{1+}(G^*,\boldsymbol{\chi})$ and (3.21) yield

$$\mathbb{P}\Bigl[y^{(i)} \in V_{0+}(G^*,\boldsymbol{\chi})\Bigr] \ge (1-p)\cdot\frac{n^{(\varepsilon-1)\theta}}{\mathrm{e}} \ge \frac{n^{\varepsilon\theta-1}}{3}, \quad \mathbb{P}\Bigl[y^{(i)} \in V_{1+}(G^*,\boldsymbol{\chi})\Bigr] \ge p\cdot\frac{n^{(\varepsilon-1)\theta}}{\mathrm{e}} \ge \frac{n^{\varepsilon\theta-1}}{3} \quad \text{for all } i \in [N],$$

provided n is sufficiently large. Therefore, recalling $N = \lceil n^{1-\xi}\rceil$ we obtain for large enough n,

$$\mathbb{E}\bigl|\{y^{(1)},\ldots,y^{(N)}\} \cap V_{0+}(G^*,\boldsymbol{\chi})\bigr| \ge n^{\varepsilon\theta-\xi}/3, \qquad \mathbb{E}\bigl|\{y^{(1)},\ldots,y^{(N)}\} \cap V_{1+}(G^*,\boldsymbol{\chi})\bigr| \ge n^{\varepsilon\theta-\xi}/3. \tag{3.22}$$

Further, because the pairwise distances of $y^{(1)},\ldots,y^{(N)}$ in G^* exceed four, the events $\{y^{(i)} \in V_{0+}(G^*,\boldsymbol{\chi})\}_{i\le N}$ are mutually independent. So are the events $\{y^{(i)} \in V_{1+}(G^*,\boldsymbol{\chi})\}_{i\le N}$. Finally, since (3.3) ensures that $\varepsilon\theta-\xi>0$, (3.22) and the Chernoff bound yield

$$\mathbb{P}\left[\left|\{y^{(1)},\ldots,y^{(N)}\}\cap V_{0+}(G^*,\boldsymbol{\chi})\right|\le\ln^2 n\right]\le\mathbb{P}\left[\mathrm{Bin}(N,n^{\varepsilon\theta-1}/3)\le\ln^2 n\right]\le\exp(-n^{\Omega(1)}),$$
$$\mathbb{P}\left[\left|\{y^{(1)},\ldots,y^{(N)}\}\cap V_{1+}(G^*,\boldsymbol{\chi})\right|\le\ln^2 n\right]\le\mathbb{P}\left[\mathrm{Bin}(N,n^{\varepsilon\theta-1}/3)\le\ln^2 n\right]\le\exp(-n^{\Omega(1)}),$$

whence the assertion is immediate. □

Proof of Proposition 3.4. Suppose that $n>n_0(\varepsilon,\theta,\xi)$ is large enough and let $G=G_{n,m}$ be a test design with $m\le(1-\varepsilon)m_{\inf}$ tests. If for every test $a\in F_m$ of degree $|\partial_G a|>\Gamma$ we have $|\partial_G a\cap V_1(G,\boldsymbol{\chi})|\ge 2$, then $V_{0+}(G,\boldsymbol{\chi})=V_{0+}(G^*,\boldsymbol{\chi})$ and $V_{1+}(G,\boldsymbol{\chi})=V_{1+}(G^*,\boldsymbol{\chi})$. Therefore, the assertion is an immediate consequence of Lemma 3.6, Lemma 3.7 and Corollary 3.13. □

3.3. **Proof of Proposition 3.5.** Given $\varepsilon>0$ and $\ln(2)/(1+\ln(2))\le\theta<\theta'<1$ we choose a large enough $n_0=n_0(\varepsilon,\theta,\theta')$ and assume that $n>n_0$. Furthermore, let G be a test design with $m\le(1-\varepsilon)m_{\inf}(n,\theta)$ for the purpose of identifying $k=\lceil n^\theta\rceil$ infected individuals. Starting from the test design G infection for density θ we are going to construct a random test design $\boldsymbol{G}'$ for infection density θ' with the same number m of tests as G. The following lemma fixes the order of $\boldsymbol{G}'$.

Lemma 3.14. *There exists an integer $n^{\theta/\theta'}/2\le n'\le 2n^{\theta/\theta'}\wedge n$ such that $k'=\lceil n'^{\theta'}\rceil=k$.*

Proof. Let $n''=\lceil n^{\theta/\theta'}/2\rceil$. Then $(4n'')^{\theta'}>k$ but $n''^{\theta'}<k$ because the function $z\in(1,\infty)\mapsto z^{\theta'}$ has derivative less than one. For the same reason for any integer $n''<N<4n''$ we have $(N+1)^{\theta'}-N^{\theta'}\le 1$ and thus

$$\lceil(N+1)^{\theta'}\rceil-\lceil N^{\theta'}\rceil\le 1.$$

Consequently, there exists an integer $n'\in(n'',4n'')$ such that $\lceil n'^{\theta'}\rceil=k$. □

Given the test design G with individuals $V_n=\{x_1,\ldots,x_n\}$ and tests $F_m=\{a_1,\ldots,a_m\}$ we now construct the test design $\boldsymbol{G}'$ as follows. Choose a subset $V(\boldsymbol{G}')\subset V_n$ of n' individuals uniformly at random. Then $\boldsymbol{G}'$ is the subgraph that G induces on $V(\boldsymbol{G}')\cup F_m$. Thus, $\boldsymbol{G}'$ has the same tests as G but we simply leave out from every test the individuals that do not belong to the random subset $V(\boldsymbol{G}')$. Let $\boldsymbol{\tau}\in\{0,1\}^{V(\boldsymbol{G}')}$ be a random vector of Hamming weight k and let $\hat{\boldsymbol{\tau}}\in\{0,1\}^{F_m}$ be the induced vector of tests results

$$\hat{\boldsymbol{\tau}}_a=\max_{x\in\partial_{\boldsymbol{G}'}a}\boldsymbol{\tau}_x\qquad\qquad(a\in F_m).$$

Lemma 3.15. *For any integer $t>0$ we have $\mathbb{P}[Z_k(G,\hat{\boldsymbol{\sigma}})\ge t]\ge\mathbb{P}\left[Z_k(\boldsymbol{G}',\hat{\boldsymbol{\tau}})\ge t\right]$.*

Proof. The choice of n' ensures that $k'=\lceil n'^{\theta'}\rceil=k$. Therefore, the random sets $\{x\in V:\boldsymbol{\sigma}_x=1\}$ and $\{x\in V(\boldsymbol{G}'):\boldsymbol{\tau}_x=1\}$ are identically distributed. Indeed, we obtain the latter by first choosing the random subset $V(\boldsymbol{G}')$ of V_n and then choosing a random subset of $V(\boldsymbol{G}')$ size k. Clearly, this two-step procedure is equivalent to just choosing a random subset of size k out of V_n. Hence, we can couple $\boldsymbol{\sigma},\boldsymbol{\tau}$ such that the sets $\{x\in V:\boldsymbol{\sigma}_x=1\}$, $\{x\in V:\boldsymbol{\tau}_x=1\}$ are identical. Then the construction of $\boldsymbol{G}'$ ensures that the vectors $\hat{\boldsymbol{\sigma}},\hat{\boldsymbol{\tau}}$ coincide as well.

Now consider a vector $\sigma'\in\mathscr{S}_k(\boldsymbol{G}',\hat{\boldsymbol{\tau}})$ that explains the test results. Extend σ' to a vector $\sigma\in\{0,1\}^{V_n}$ by setting $\sigma_x=0$ for all $x\in V_n\setminus V(\boldsymbol{G}')$. Then $\sigma\in\mathscr{S}_k(\boldsymbol{G},\hat{\boldsymbol{\sigma}})$. Hence, $Z_k(G,\hat{\boldsymbol{\sigma}})\ge Z_k(\boldsymbol{G}',\hat{\boldsymbol{\tau}})$. □

Proof of Proposition 3.5. Lemma 3.15 shows that for any $t>0$,

$$\mathbb{P}[Z_k(G,\hat{\boldsymbol{\sigma}})\ge t]\ge\mathbb{P}\left[Z_k(\boldsymbol{G}',\hat{\boldsymbol{\tau}})\ge t\right]=\mathbb{E}\left[\mathbb{P}\left[Z_k(\boldsymbol{G}',\hat{\boldsymbol{\tau}})\ge t\mid\boldsymbol{G}'\right]\right].$$

Consequently, there exists an outcome G' of $\boldsymbol{G}'$ such that $\mathbb{P}[Z_k(G,\hat{\boldsymbol{\sigma}})\ge t]\ge\mathbb{P}\left[Z_k(G',\hat{\boldsymbol{\tau}})\ge t\right]$. □

4. The non-adaptive group testing algorithm SPIV

In this section we describe the new test design and the associated inference algorithm SPIV for Theorem 1.2. Throughout we fix $\theta\in(0,1)$ and $\varepsilon>0$ and we tacitly assume that $n>n_0(\varepsilon,\theta)$ is large enough for the various estimates to hold.

4.1. **The random bipartite graph and the DD algorithm.** To motivate the new test design we begin with a brief discussion of the plain random design used in prior work and the best previously known inference algorithm DD [11, 22]. At first glance a promising candidate test design appears to be a random bipartite graph with one vertex class $V_n = \{x_1, \dots, x_n\}$ representing individuals and the other class $F_m = \{a_1, \dots, a_m\}$ representing tests. Indeed, two slightly different random graph models have been proposed [6]. First, in the *Bernoulli model* each V_n–F_m-edge is present with a certain probability (the same for every pair) independently of all others. However, due to the relatively heavy lower tail of the degrees of the individuals, this test design turns out to be inferior to a second model where the degrees of the individuals are fixed. Specifically, in the Δ-*out model* every individual independently joins an equal number of Δ tests drawn uniformly at random without replacement [29].

Clearly, in order to extract the maximum amount of information Δ should be chosen so as to maximise the entropy of the vector of test results. Specifically, since the average test degree equals $\Delta n/m$ and a total of k individuals are infected, the average number of infected individuals per test comes to $\Delta k/m$. Indeed, since $k \sim n^\theta$ for a fixed $\theta < 1$, the number of infected individuals in test a_i can be well approximated by a Poisson variable. Therefore, setting

$$\Delta \sim \frac{m}{k} \ln 2 \tag{4.1}$$

ensures that about half the tests are positive w.h.p.

With respect to the performance of the Δ-out model, [11, Theorem 1.1] implies together with Theorem 1.1 that this simple construction is information-theoretically optimal. Indeed, $m = (1+\varepsilon+o(1))m_{\inf}$ test suffice so that an exponential time algorithm correctly infers the set of infected individuals. Specifically, the algorithm solves a minimum hypergraph vertex cover problem with the individuals as the vertex set and the positive test groups as the hyperedges. For $m = (1+\varepsilon+o(1))m_{\inf}$ the unique optimal solution is precisely the correct set of infected individuals w.h.p. While the worst case NP-hardness of hypergraph vertex cover does not, of course, preclude the existence of an algorithm that is efficient on random hypergraphs, despite considerable efforts no such algorithm has been found. In fact, as we saw in Section 1.4 for a good number of broadly similar inference and optimisation problems on random graphs no efficient information-theoretically optimal algorithms are known.

But for m exceeding the threshold $m_{\mathtt{DD}}$ from (1.2) an efficient greedy algorithm DD correctly recovers $\boldsymbol{\sigma}$ w.h.p. The algorithm proceeds in three steps.

- **DD1:** declare every individual that appears in a negative test uninfected and subsequently remove all negative tests and all individuals that they contain.
- **DD2:** for every remaining (positive) test of degree one declare the individual that appears in the test infected.
- **DD3:** declare all other individuals as uninfected.

The decisions made by the first two steps **DD1**–**DD2** are clearly correct but **DD3** might produce false negatives. Prior to the present work DD was the best known polynomial time group testing algorithm. While DD correctly identifies the set of infected individuals w.h.p. if $m > (1+\varepsilon)m_{\mathtt{DD}}$ [22], the algorithm fails if $m < (1-\varepsilon)m_{\mathtt{DD}}$ w.h.p. [11].

4.2. **Spatial coupling.** The new efficient algorithm SPIV for Theorem 1.2 that gets by with the optimal number $(1+\varepsilon+o(1))m_{\inf}$ of tests comes with a tailor-made test design that, inspired by spatially coupled codes [18, 26, 27], combines randomisation with a superimposed geometric structure. Specifically, we divide both the individuals and the tests into

$$\ell = \lceil \ln^{1/2} n \rceil \tag{4.2}$$

compartments of equal size. The compartments are arranged along a ring and each individual joins an equal number of random tests in the

$$s = \lceil \ln \ln n \rceil = o(\ell) \tag{4.3}$$

topologically subsequent compartments. Additionally, to get the algorithm started we equip the first s compartments with extra tests so that they can be easily diagnosed via the DD algorithm. Then, having diagnosed the initial compartments correctly, SPIV will work its way along the ring, diagnosing one compartment after the other.

To implement this idea precisely we partition the set $V = V_n = \{x_1, \dots, x_n\}$ of individuals into pairwise disjoint subsets $V[1], \dots, V[\ell]$ of sizes $|V[j]| \in \{\lfloor n/\ell \rfloor, \lceil n/\ell \rceil\}$. With each compartment $V[i]$ of individuals we associate a compartment $F[i]$ of tests of size $|F[i]| = m/\ell$ for an integer m that is divisible by ℓ. Additionally, we introduce

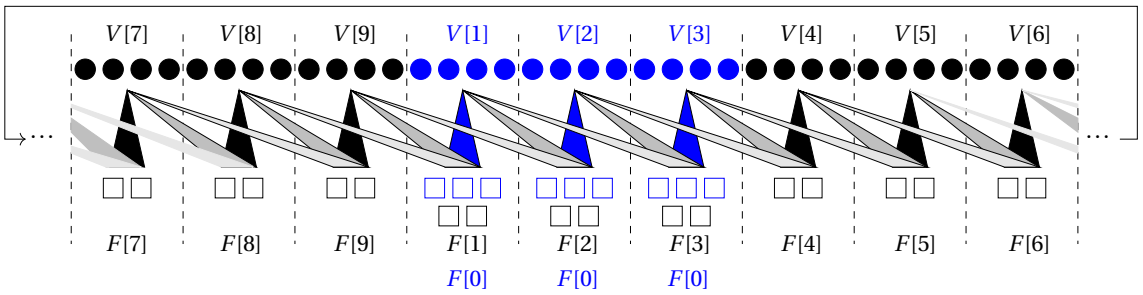

FIGURE 2. The spatially coupled test design with $n = 36, \ell = 9, s = 3$. The individuals in the seed groups $V[1] \cup \cdots \cup V[s]$ (blue) are equipped with additional test $F[0]$ (blue rectangles). The black rectangles represent the tests $F[1] \cup \cdots \cup F[\ell]$.

a set $F[0]$ of $10\lceil (ks/\ell)\ln n\rceil$ extra tests to facilitate the greedy algorithm for diagnosing the first s compartments. Thus, the total number of tests comes to

$$|F[0]| + \sum_{i=1}^{\ell} |F[i]| = (1 + O(s/\ell))m = (1 + o(1))m. \tag{4.4}$$

Finally, for notational convenience we define $V[\ell + i] = V[i]$ and $F[\ell + i] = F[i]$ for $i = 1, \ldots, s$.

The test groups are composed as follows: let

$$k = \lceil n^{\theta} \rceil \qquad \text{and let} \qquad \Delta = \frac{m \ln 2}{k} + O(s) \tag{4.5}$$

be an integer divisible by s; cf. (4.1). Then we construct a random bipartite graph as follows.

SC1: for $i = 1, \ldots, \ell$ and $j = 1, \ldots, s$ every individual $x \in V[i]$ joins Δ/s tests from $F[i + j - 1]$ chosen uniformly at random without replacement. The choices are mutually independent for all individuals x and all j.

SC2: additionally, each individual from $V[1] \cup \cdots \cup V[s]$ independently joins $\lceil 10 \ln(2) \ln n \rceil$ random tests from $F[0]$, drawn uniformly without replacement.

Thus, **SC1** provides that the individuals in compartment $V[i]$ take part in the next s compartments $F[i], \ldots, F[i + s - 1]$ of tests along the ring. Furthermore, **SC2** supplies the tests required by the DD algorithm to diagnose the first s compartments. Figure 2 provides an illustration of the resulting random test design,

From here on the test design produced by **SC1–SC2** is denoted by $\boldsymbol{G}$. Furthermore $\boldsymbol{\sigma} \in \{0,1\}^V$ denotes a uniformly random vector of Hamming weight k, drawn independently of $\boldsymbol{G}$, and $\hat{\boldsymbol{\sigma}} = (\hat{\boldsymbol{\sigma}}_a)_{a \in F[0] \cup \cdots \cup F[\ell]}$ signifies the vector of test results

$$\hat{\boldsymbol{\sigma}}_a = \max_{x \in \partial a} \boldsymbol{\sigma}_x.$$

In addition, let $V_1 = \{x \in V : \boldsymbol{\sigma}_x = 1\}$ be the set of infected individuals and let $V_0 = V \setminus V_1$ be the set of healthy individuals. Moreover, let $F = F[0] \cup F[1] \cup \cdots \cup F[\ell]$ be the set of all tests, let $F_1 = \{a \in F : \hat{\boldsymbol{\sigma}}_a = 1\}$ be the set of all positive tests and let $F_0 = F \setminus F_1$ be the set of all negative tests. Finally, let

$$V_0[i] = V[i] \cap V_0, \qquad V_1[i] = V[i] \cap V_1, \qquad F_0[i] = F[i] \cap F_0, \qquad F_1[i] = F[i] \cap F_1.$$

The following proposition summarises a few basic properties of the test design $\boldsymbol{G}$.

Proposition 4.1. *If $m = \Theta(n^{\theta} \ln n)$ then $\boldsymbol{G}$ enjoys the following properties with probability $1 - o(n^{-2})$.*

(i) The infected individual counts in the various compartments satisfy

$$\frac{k}{\ell} - \sqrt{\frac{k}{\ell}} \ln n \le \min_{i \in [\ell]} |V_1[i]| \le \max_{i \in [\ell]} |V_1[i]| \le \frac{k}{\ell} + \sqrt{\frac{k}{\ell}} \ln n.$$

(ii) For all $i \in [\ell]$ and all $j \in [s]$ the test degrees satisfy

$$\frac{\Delta n}{ms} - \sqrt{\frac{\Delta n}{ms}} \ln n \le \min_{a \in F[i+j-1]} |V[i] \cap \partial a| \le \max_{a \in F[i+j-1]} |V[i] \cap \partial a| \le \frac{\Delta n}{ms} + \sqrt{\frac{\Delta n}{ms}} \ln n.$$

(iii) For all $i \in [\ell]$ the number of negative tests in compartment $F[i]$ satisfies

$$\frac{m}{2\ell} - \sqrt{m}\ln^3 n \le |F_0[i]| \le \frac{m}{2\ell} + \sqrt{m}\ln^3 n.$$

We prove Proposition 4.1 in Section 4.4. Finally, as a preparation for things to come we point out that for any specific individual $x \in V[i]$ and any particular test $a \in F[i+j]$, $j = 0,\dots,s-1$, we have

$$\mathbb{P}[x \in \partial a] = 1 - \mathbb{P}[x \notin \partial a] = 1 - \binom{|F[i+j]|-1}{\Delta/s}\binom{|F[i+j]|}{\Delta/s}^{-1} = \frac{\Delta\ell}{ms} + O\left(\left(\frac{\Delta\ell}{ms}\right)^2\right). \tag{4.6}$$

4.3. **The Spatial Inference Vertex Cover ('SPIV') algorithm.** The SPIV algorithm for Theorem 1.2 proceeds in three phases. The plan of attack is for the algorithm to work its way along the ring, diagnosing one compartment after the other aided by what has been learned about the preceding compartments. Of course, we need to start somewhere. Hence, in its first phase SPIV diagnoses the seed compartments $V[1],\dots,V[s]$.

4.3.1. *Phase 1: the seed.* Specifically, the first phase of SPIV applies the DD greedy algorithm from Section 4.1 to the subgraph of $\boldsymbol{G}$ induced on the individuals $V[1]\cup\ldots\cup V[s]$ and the tests $F[0]$. Throughout the vector $\tau \in \{0,1\}^V$ signifies the algorithm's current estimate of the ground truth $\boldsymbol{\sigma}$.

Input: $\boldsymbol{G}$, $\hat{\boldsymbol{\sigma}}$
Output: an estimate of $\boldsymbol{\sigma}$
1 Let $(\tau_x)_{x\in V[1]\cup\cdots\cup V[s]} \in \{0,1\}^{V[1]\cup\cdots\cup V[s]}$ be the result of applying DD to the tests $F[0]$;
2 Set $\tau_x = 0$ for all individuals $x \in V \setminus (V[1]\cup\cdots\cup V[s])$;

Algorithm 1: SPIV, phase 1

The following proposition, whose proof can be found in Section 4.5, summarises the analysis of phase 1.

Proposition 4.2. *W.h.p. the output of* DD *satisfies $\tau_x = \boldsymbol{\sigma}_x$ for all $x \in V[1]\cup\cdots\cup V[s]$.*

4.3.2. *Phase 2: enter the ring.* This is the main phase of the algorithm. Thanks to Proposition 4.2 we may assume that the seed has been diagnosed correctly. Now, the programme is to diagnose one compartment after the other, based on what the algorithm learned previously. Hence, assume that we managed to diagnose compartments $V[1],\dots,V[i]$ correctly. How do we proceed to compartment $V[i+1]$?

For a start, we can safely mark as uninfected all individuals in $V[i+1]$ that appear in a negative test. But a simple calculation reveals that this will still leave us with many more than k undiagnosed individuals w.h.p. To be precise, consider the set of uninfected disguised individuals

$$V_{0+}[i+1] = \{x \in V_0[i+1] : \hat{\boldsymbol{\sigma}}_a = 1 \text{ for all } a \in \partial x\},$$

i.e., uninfected individuals that fail to appear in a negative test. In Section 4.6 we prove the following.

Lemma 4.3. *Suppose that $(1+\varepsilon)m_{\mathrm{ad}} \le m = O(n^\theta \ln n)$. Then w.h.p. for all $s \le i < \ell$ we have*

$$|V_{0+}[i+1]| = \left(1 + O\left(n^{-\Omega(1)}\right)\right)\frac{n}{\ell 2^\Delta}.$$

Hence, by the definition (4.5) of Δ for m close to $m_{\inf}$ the set $V_{0+}[i+1]$ has size $k^{1+\Omega(1)} \gg k$ w.h.p.

Thus, the challenge is to discriminate between $V_{0+}[i+1]$ and the set $V_1[i+1]$ of actual infected individuals in compartment $i+1$. The key observation is that we can tell these sets apart by counting currently 'unexplained' positive tests. To be precise, for an individual $x \in V[i+1]$ and $1 \le j \le s$ let $\boldsymbol{W}_{x,j}$ be the number of tests in compartment $F[i+j]$ that contain x but that do not contain an infected individual from the preceding compartments $V[1]\cup\cdots\cup V[i]$. In formulas,

$$\boldsymbol{W}_{x,j} = \left|\left\{a \in \partial x \cap F[i+j] : \partial a \cap (V_1[1]\cup\cdots\cup V_1[i]) = \emptyset\right\}\right|. \tag{4.7}$$

Crucially, the following back-of-the-envelope calculation shows that the mean of this random variable depends on whether x is infected or healthy but disguised.

Infected individuals ($x \in V_1[i+1]$)**:** consider a test $a \in \partial x \cap F[i+j]$, $j = 1,\dots,s$. Because the individuals join tests independently, conditioning on x being infected does not skew the distribution of the individuals from the $s-j$ prior compartments $V[i+j-s+1],\dots,V[i]$ that appear in a. Furthermore, we chose Δ so that for each of these compartments $V[h]$ the expected number of infected individuals that join a has mean $(\ln 2)/s$. Indeed, due to independence it is not difficult to see that $|V_1[h] \cap \partial a|$ is approximately a Poisson variable. Consequently,

$$\mathbb{P}\left[(V_1[i+j-s+1] \cup \dots \cup V_1[i]) \cap \partial a = \emptyset\right] \sim 2^{-(s-j)/s}. \tag{4.8}$$

Hence, because x appears in Δ/s tests $a \in F[i+j]$, the linearity of expectation yields

$$\mathbb{E}\left[\boldsymbol{W}_{x,j} \mid x \in V_1[i+1]\right] \sim 2^{j/s-1}\frac{\Delta}{s}. \tag{4.9}$$

Disguised healthy individuals ($x \in V_{0+}[i+1]$)**:** similarly as above, for any individual $x \in V[i+1]$ and any $a \in \partial x \cap F[i+j]$ the *unconditional* number of infected individuals in a is asymptotically $\mathrm{Po}(\ln 2)$. But given $x \in V_{0+}[i+1]$ we know that a is positive. Thus, $\partial a \setminus \{x\}$ contains at least one infected individual. In effect, the number of positives in a approximately turns into a conditional Poisson $\mathrm{Po}_{\geq 1}(\ln 2)$. Consequently, for test a not to include any infected individual from one of the known compartments $V[h]$, $h = i+j-s+1,\dots,i$, every infected individual in test a must stem from the j yet undiagnosed compartments. Summing up the conditional Poisson and recalling that x appears in Δ/s tests $a \in F[j]$, we thus obtain

$$\mathbb{E}\left[\boldsymbol{W}_{x,j} \mid x \in V_{0+}[i+1]\right] \sim \frac{\Delta}{s}\sum_{t\geq 1}\mathbb{P}\left[\mathrm{Po}_{\geq 1}(\ln 2) = t\right](j/s)^t = (2^{j/s}-1)\frac{\Delta}{s}. \tag{4.10}$$

A first idea to tell $V_{0+}[i+1]$ and $V_1[i+1]$ apart might thus be to simply calculate

$$\boldsymbol{W}_x = \sum_{j=1}^{s-1}\boldsymbol{W}_{x,j} \qquad\qquad (x \in V[i+1]). \tag{4.11}$$

Indeed, (4.9) and (4.10) yield

$$\mathbb{E}\left[\boldsymbol{W}_x \mid x \in V_1[i+1]\right] \sim \frac{\Delta}{2\ln 2} = 0.721\dots\Delta \qquad \text{whereas} \qquad \mathbb{E}\left[\boldsymbol{W}_x \mid x \in V_{0+}[i+1]\right] \sim \frac{\Delta(1-\ln 2)}{\ln 2} = 0.442\dots\Delta.$$

But unfortunately a careful large deviations analysis reveals that $\boldsymbol{W}_x$ is not sufficiently concentrated. More precisely, even for $m = (1+\varepsilon+o(1))m_{\mathrm{inf}}$ there are as many as $k^{1+\Omega(1)}$ 'outliers' $x \in V_{0+}[i+1]$ whose $\boldsymbol{W}_x$ grows as large as the mean $\Delta/(2\ln 2)$ of actual infected individuals w.h.p.

At second thought the plain sum (4.11) does seem to leave something on the table. While $\boldsymbol{W}_x$ counts all as yet unexplained positive tests equally, not all of these tests reveal the same amount of information. In fact, we should really be paying more attention to 'early' unexplained tests $a \in F[i+1]$ than to 'late' ones $b \in F[i+s]$. For we already diagnosed $s-1$ out of the s compartments of individuals that a draws on, whereas only one of the s compartments that contribute to b has already been diagnosed. Thus, the unexplained test a is a much stronger indication that x might be infected. Consequently, it seems promising to replace $\boldsymbol{W}_x$ by a weighted sum

$$\boldsymbol{W}_x^\star = \sum_{j=1}^{s-1} w_j \boldsymbol{W}_{x,j} \tag{4.12}$$

with $w_1,\dots,w_{s-1} \geq 0$ chosen so as to gauge the amount of information carried by the different compartments.

To find the optimal weights $w_1,\dots,w_{s-1}$ we need to investigate the rate function of $\boldsymbol{W}_x^\star$ given $x \in V_{0+}[i+1]$. More specifically, we should minimise the probability that $\boldsymbol{W}_x^\star$ given $x \in V_{0+}[i+1]$ grows as large as the mean of $\boldsymbol{W}_x^\star$ given $x \in V_1[i+1]$, which we read off (4.9) easily:

$$\mathbb{E}\left[\boldsymbol{W}_x^\star \mid x \in V_1[i+1]\right] \sim \frac{\Delta}{s}\sum_{j=1}^{s-1} 2^{j/s-1} w_j. \tag{4.13}$$

A careful large deviations analysis followed by a Lagrangian optimisation leads to the optimal choice

$$w_j = \ln\frac{(1-2\zeta)2^{j/s-1}(2-2^{j/s})}{(1-(1-2\zeta)2^{j/s-1})(2^{j/s}-1)} \qquad \text{where} \qquad \zeta = 1/s^2. \tag{4.14}$$

The following two lemmas show that with these weights the scores $\boldsymbol{W}^\star_x$ discriminate well between the potential false positives and the infected individuals. More precisely, thresholding $\boldsymbol{W}^\star_x$ we end up misclassifying no more than $o(k)$ individuals x w.h.p.

Lemma 4.4. *Suppose that* $(1+\varepsilon)m_{\mathrm{ad}} \le m = O(n^\theta \ln n)$. *W.h.p. we have*

$$\sum_{s\le i<\ell}\sum_{x\in V_1[i]} \mathbf{1}\left\{\boldsymbol{W}^\star_x < (1-\zeta/2)\frac{\Delta}{s}\sum_{j=1}^{s-1} 2^{j/s-1}w_j\right\} \le k\exp\left(-\frac{\Omega(\ln n)}{(\ln\ln n)^4}\right). \tag{4.15}$$

Lemma 4.5. *Suppose that* $(1+\varepsilon)m_{\mathrm{ad}} \le m = O(n^\theta \ln n)$. *W.h.p. we have*

$$\sum_{s\le i<\ell}\sum_{x\in V_{0+}[i]} \mathbf{1}\left\{\boldsymbol{W}^\star_x > (1-2\zeta)\frac{\Delta}{s}\sum_{j=1}^{s-1} 2^{j/s-1}w_j\right\} \le k^{1-\Omega(1)}. \tag{4.16}$$

We prove these two lemmas in Sections 4.7 and 4.8.

Lemmas 4.4–4.5 leave us with only one loose end. Namely, calculating the scores $\boldsymbol{W}^\star_x$ requires knowledge of the correct infection status $\boldsymbol{\sigma}_x$ of *all* the individuals $x \in V[1]\cup\cdots\cup V[i]$ from the previous compartments. But since the r.h.s. expressions in (4.15) and (4.16) are non-zero, it is unrealistic to assume that the algorithm's estimates τ_x will consistently match the ground truth $\boldsymbol{\sigma}_x$ beyond the seed compartments. Hoping that the algorithm's estimate will not stray too far, we thus have to make do with the approximate scores

$$W^\star_x(\tau) = \sum_{j=1}^{s-1} w_j W_{x,j}(\tau), \qquad \text{where} \qquad W_{x,j}(\tau) = \left|\left\{a\in\partial x\cap F[i+j-1] : \max_{y\in\partial a\cap(V[1]\cup\cdots V[i])}\tau_y = 0\right\}\right|. \tag{4.17}$$

Hence, phase 2 of SPIV reads as follows.

```
for i = s, ..., ℓ − 1 do
    for x ∈ V[i + 1] do
        if ∃a ∈ ∂x : σ̂_a = 0 then
            τ_x = 0 // classify as uninfected
        else if W*_x(τ) < (1 − ζ)(Δ/s) Σ_{j=1}^{s−1} 2^{j/s−1} w_j then
            τ_x = 0 // tentatively classify as uninfected
        else
            τ_x = 1 // tentatively classify as infected
```

Algorithm 2: SPIV, phase 2.

Since phase 2 of SPIV uses the approximations from (4.17), there seems to be a risk of errors amplifying as we move along. Fortunately, it turns out that errors proliferate only moderately and the second phase of SPIV will misclassify only $o(k)$ individuals. The following proposition summarises the analysis of phase 2.

Proposition 4.6. *Suppose that* $(1+\varepsilon)m_{\mathrm{ad}} \le m = O(k\ln n)$. *W.h.p. the assignment* τ *obtained after steps 1–10 satisfies*

$$\sum_{x\in V}\mathbf{1}\{\tau_x \ne \boldsymbol{\sigma}_x\} \le k\exp\left(-\frac{\ln n}{(\ln\ln n)^6}\right).$$

The proof of Proposition 4.6 can be found in Section 4.9.

4.3.3. *Phase 3: cleaning up.* The final phase of the algorithm rectifies the errors incurred during phase 2. The combinatorial insight that makes this possible is that for $m \ge (1+\varepsilon)m_{\mathrm{inf}}$ every infected individual has at least $\Omega(\Delta)$ positive tests to itself w.h.p. Thus, these tests do not feature a second infected individual. Phase 3 of the algorithm exploits this observation by simply thresholding the number S_x of tests where there is no other infected individual besides potentially x. Thanks to the expansion properties of the graph $\boldsymbol{G}$, each iteration of the thresholding procedure reduces the number of misclassified individuals by at least a factor of three. In effect, after $\ln n$ iterations all individuals will be classified correctly w.h.p. Of course, due to Proposition 4.2 we do not need to reconsider the seed $V[1]\cup\cdots\cup V[s]$.

Let $\tau^{(1)} = \tau$;
for $i = 1, \dots, \lceil \ln n \rceil$ **do**
  For all $x \in V[s+1] \cup \dots \cup V[\ell]$ calculate
    $S_x(\tau^{(i)}) = \sum_{a \in \partial x : \hat{\boldsymbol{\sigma}}_a = 1} \mathbf{1}\left\{\forall y \in \partial a \setminus \{x\} : \tau_y^{(i)} = 0\right\}$;
  Let $\tau_x^{(i+1)} = \begin{cases} \tau_x^{(i)} & \text{if } x \in V[1] \cup \dots \cup V[s], \\ \mathbf{1}\left\{S_x\left(\tau^{(i)}\right) > \ln^{1/4} n\right\} & \text{otherwise} \end{cases}$;
return $\tau^{(\lceil \ln n \rceil)}$

Algorithm 3: `SPIV`, phase 3.

Proposition 4.7. *Suppose that* $(1+\varepsilon)m_{\inf} \le m = O(n^\theta \ln n)$. *W.h.p. for all* $1 \le i \le \lceil \ln n \rceil$ *we have*

$$\sum_{x \in V} \mathbf{1}\{\tau_x^{(i+1)} \neq \boldsymbol{\sigma}_x\} \le \frac{1}{3} \sum_{x \in V} \mathbf{1}\{\tau_x^{(i)} \neq \boldsymbol{\sigma}_x\}.$$

We prove Proposition 4.7 in Section 4.10.

Proof of Theorem 1.2. The theorem is an immediate consequence of Propositions 4.2, 4.6 and 4.7. □

4.4. **Proof of Proposition 4.1.** The number $|V_1[i]|$ of infected individuals in compartment $V[i]$ has distribution $\mathrm{Hyp}(n, k, |V[i]|)$. Since $||V[i]| - n/\ell| \le 1$, (i) is an immediate consequence of the Chernoff bound from Lemma 2.2.

With respect to (ii), we recall from (4.6) that $\mathbb{P}[x \in \partial a] = \frac{\Delta\ell}{ms}(1 + O(\frac{\Delta\ell}{ms}))$. Hence, because the various individuals $x \in V[i]$ join tests independently, the number $|V[i] \cap \partial a|$ of test participants from $V[i]$ has distribution

$$|V[i] \cap \partial a| \sim \mathrm{Bin}(|V[i]|, \Delta\ell/(ms) + O((\Delta\ell/ms)^2)).$$

Since $|V[i]| = n/\ell + O(1)$, assertion (ii) follows from (4.5) and the Chernoff bound from Lemma 2.1.

Coming to (iii), due to part (i) we may condition on $\mathscr{E} = \{\forall i \in [\ell] : |V_1[i]| = k/\ell + O(\sqrt{k/\ell}\ln n)\}$. Hence, with h ranging over the s compartments whose individuals join tests in $F[i]$, (4.6) implies that for every test $a \in F[i]$ the number of infected individuals $|V_1 \cap \partial a|$ is distributed as a sum of independent binomial variables

$$|V_1 \cap \partial a| \sim \sum_h \boldsymbol{X}_h \qquad \text{with} \qquad \boldsymbol{X}_h \sim \mathrm{Bin}\left(V_1[h], \frac{\Delta\ell}{ms} + O\left(\left(\frac{\Delta\ell}{ms}\right)^2\right)\right).$$

Consequently, (4.5) ensures that the event $V_1 \cap \partial a = \emptyset$ has conditional probability

$$\begin{aligned}\mathbb{P}[V_1 \cap \partial a = \emptyset \mid \mathscr{E}] &= \prod_h \mathbb{P}[\boldsymbol{X}_h = 0 \mid \mathscr{E}] = \exp\left[s\left(\frac{k}{\ell} + O\left(\sqrt{\frac{k}{\ell}}\ln n\right)\right)\ln\left(1 - \frac{\Delta\ell}{ms} + O\left(\left(\frac{\Delta\ell}{ms}\right)^2\right)\right)\right] \\ &= \exp\left[-\frac{sk}{\ell}\cdot\frac{\Delta\ell}{ms} + O\left(\sqrt{\frac{k}{\ell}}\cdot\frac{\Delta\ell}{m}\right) + O\left(\frac{sk}{\ell}\cdot\left(\frac{\Delta\ell}{ms}\right)^2\right)\right] = \frac{1}{2} + O(\sqrt{\ell/k}).\end{aligned}$$

Therefore, we obtain the estimate

$$\mathbb{E}[|F_0[i]| \mid \mathscr{E}] = \frac{m}{2\ell} + O(\sqrt{m}\ln n). \tag{4.18}$$

Finally, changing the set of tests that a specific infected individual $x \in V_1[h]$ joins shifts $|F_0[i]|$ by at most Δ (while tinkering with uninfected ones does not change $|F_0[i]|$ at all). Therefore, the Azuma–Hoeffding inequality yields

$$\mathbb{P}[||F_0[i]| - \mathbb{E}[|F_0[i]| \mid \mathscr{E}]| \ge t \mid \mathscr{E}] \le 2\exp\left(-\frac{t^2}{2k\Delta^2}\right) \qquad \text{for any } t > 0. \tag{4.19}$$

Thus, (iii) follows from (4.5), (4.18) and (4.19) with $t = \sqrt{m}\ln^3 n$.

4.5. **Proof of Proposition 4.2.** Let $D = \lceil 10\ln(2)\ln n\rceil$ and recall that $|F[0]| = \lceil 10ks\ln(n)/\ell\rceil$. Since by **SC2** every individual from $\in V[1]\cup\cdots\cup V[s]$ joins D random tests from $F[0]$, in analogy to (4.6) for every $x \in V[1]\cup\cdots\cup V[s]$ and every test $a \in F[0]$ we obtain

$$\mathbb{P}[x \in \partial a] = 1 - \mathbb{P}[x \notin \partial a] = 1 - \binom{|F[0]|-1}{D}\binom{|F[0]|}{D}^{-1} = \frac{D}{|F[0]|}\left(1 + O\left(\frac{D}{|F[0]|}\right)\right) = \frac{\ell\ln 2}{ks}\left(1 + O(n^{-\Omega(1)})\right). \tag{4.20}$$

Let $F_1[0]$ be the set of tests $a \in F[0]$ with $\hat{\boldsymbol{\sigma}}_a = 1$.

Lemma 4.8. *W.h.p. the number of positive tests $a \in F[0]$ satisfies $|F_1[0]| = |F[0]|(\frac{1}{2} + O(n^{-\Omega(1)}))$.*

Proof. By Proposition 4.1 we may condition on the event $\mathscr{E}$ that $|V_1[1]\cup\cdots\cup V_1[s]| = \frac{ks}{\ell}(1+O(n^{-\Omega(1)}))$. Hence, (4.20) implies that given $\mathscr{E}$ the expected number of infected individuals in a test $a \in F[0]$ comes to

$$\mathbb{E}[|\partial a \cap V_1| \mid \mathscr{E}] = \ln 2 + O(n^{-\Omega(1)}). \tag{4.21}$$

Moreover, since individuals join tests independently, $|\partial a \cap V_1|$ is a binomial random variable. Hence, (4.21) implies $\mathbb{P}[\partial a \cap V_1 = \emptyset \mid \mathscr{E}] = \frac{1}{2} + O(n^{-\Omega(1)})$. Consequently, since $\mathbb{P}[\mathscr{E}] = 1 - o(n^{-2})$ by Proposition 4.1,

$$\mathbb{E}|F_1 \cap F[0]| = \mathbb{E}|F_1[0]| = \frac{|F[0]|}{2}(1 + O(n^{-\Omega(1)})). \tag{4.22}$$

Finally, changing the set ∂x of neighbours of an infected individual can shift $|F_1[0]|$ by at most D. Therefore, the Azuma–Hoeffding inequality implies that

$$\mathbb{P}[||F_1[0]| - \mathbb{E}|F_1[0]|| > t] \le 2\exp\left(-\frac{t^2}{2D^2k}\right) \qquad \text{for any } t > 0. \tag{4.23}$$

Since $D = O(\ln n)$, combining (4.22) and (4.23) and setting, say, $t = k^{2/3}$ completes the proof. □

As an application of Lemma 4.8 we show that w.h.p. every seed individual x appears in a test $a \in F[0]$ whose other individuals are all healthy.

Corollary 4.9. *W.h.p. every individual $x \in V[1]\cup\cdots\cup V[s]$ appears in a test $a \in F[0]\cap\partial x$ such that $\partial a \setminus \{x\} \subset V_0$.*

Proof. We expose the random bipartite graph induced on $V[1]\cup\cdots\cup V[s]$ and $F[0]$ in two rounds. In the first round we expose $\boldsymbol{\sigma}$ and all neighbourhoods $(\partial y)_{y\in(V[1]\cup\cdots\cup V[s])\setminus\{x\}}$. In the second round we expose ∂x. Let $\boldsymbol{X}$ be the number of negative tests $a \in F[0]$ after the first round. Since x has degree $D = O(\ln n)$, Lemma 4.8 implies that $\boldsymbol{X} = |F[0]|(\frac{1}{2} + O(n^{-\Omega(1)}))$ w.h.p. Furthermore, given $\boldsymbol{X}$ the number of tests $a \in \partial x$ all of whose other individuals are uninfected has distribution $\mathrm{Hyp}(|F[0]|, \boldsymbol{X}, D)$. Hence,

$$\mathbb{P}[\forall a \in \partial x : V_1 \cap \partial a \setminus \{x\} \neq \emptyset \mid \boldsymbol{X}] = \binom{|F[0]|-\boldsymbol{X}}{D}\binom{|F[0]|}{D}^{-1} \le \exp(-D\boldsymbol{X}/|F[0]|). \tag{4.24}$$

Assuming $\boldsymbol{X}/|F[0]| = \frac{1}{2} + O(n^{-\Omega(1)})$ and recalling that $D = \lceil 10\ln(2)\ln n\rceil$, we obtain $\exp(-D\boldsymbol{X}/|F[0]|) = o(1/n)$. Thus, the assertion follows from (4.24) and the union bound. □

Proof of Proposition 4.2. Due to Corollary 4.9 we may assume that for every $x \in V[1]\cup\cdots\cup V[s]$ there is a test $a_x \in F[0]$ such that $\partial a_x \setminus \{x\} \subset V_0$. Hence, recalling the `DD` algorithm from Section 4.1, we see that the first step **DD1** will correctly identify all healthy individuals $x \in V_0[1]\cup\cdots\cup V_0[s]$. Moreover, the second step **DD2** will correctly classify all remaining individuals $V_1[1]\cup\cdots\cup V_1[s]$ as infected, and the last step **DD3** will be void. □

4.6. **Proof of Lemma 4.3.** Let $\mathscr{E}$ be the event that properties (i) and (iii) from Proposition 4.1 hold; then $\mathbb{P}[\mathscr{E}] = 1 - o(n^{-2})$. Moreover, let $\mathfrak{E}$ be the σ-algebra generated by $\boldsymbol{\sigma}$ and the neighbourhoods $(\partial x)_{x\in V_1}$. Then the event $\mathscr{E}$ is $\mathfrak{E}$-measurable while the neighbourhoods $(\partial x)_{x\in V_0}$ of the healthy individuals are independent of $\mathfrak{E}$. Recalling from **SC1** that the individuals $x \in V_0[i]$ choose Δ/s random tests in each of the compartments $F[i+j]$, $0 \le j \le s-1$ independently and remembering that $x \in V_{0+}[i]$ iff none of these tests is negative, on $\mathscr{E}$ we obtain

$$\begin{aligned}\mathbb{P}[x \in V_{0+}[i] \mid \mathfrak{E}] &= \binom{m/(2\ell) + O(\sqrt{m}\ln^3 n)}{\Delta/s}^s \binom{m/\ell}{\Delta/s}^{-s} = \left(\frac{1 + O(m^{-1/2}\ell\ln^3 n)}{2}\right)^{\Delta} \\ &= 2^{-\Delta} + O(m^{-1/2}\Delta\ell\ln^3 n) = 2^{-\Delta}(1 + O(n^{-\theta/2}\ln^4 n)) \qquad \text{[due to (4.2) and (4.5)]}.\end{aligned} \tag{4.25}$$

Because all $x \in V_0[i]$ choose their neighbourhoods independently, (4.25) implies that the conditional random variable $|V_{0+}[i]|$ given $\mathfrak{E}$ has distribution $\mathrm{Bin}(|V_0[i]|, 2^{-\Delta}(1+O(n^{-\Omega(1)})))$. Therefore, since on $\mathscr{E}$ we have $|V_0[i]| = |V[i]| + O(n^\theta) = n/\ell + O(n^\theta)$, the assertion follows from the Chernoff bound from Lemma 2.1.

4.7. **Proof of Lemma 4.4.** The aim is to estimate the weighted sum $\boldsymbol{W}^\star_x$ for infected individuals $x \in V[i+1]$ with $s \le i < \ell$. These individuals join tests in the s compartments $F[i+j]$, $j \in [s]$. Conversely, for each such j the tests $a \in F[i+j]$ recruit their individuals from the compartments $V[i+j-s+1], \dots, V[i+j]$. Thus, the compartments preceding $V[i+1]$ that the tests in $F[i+j]$ draw upon are $V[h]$ with $i+j-s < h \le i$. We begin by investigating the set $\mathscr{W}_{i,j}$ of tests $a \in F[i+j]$ without an infected individual from these compartments, i.e.,

$$\mathscr{W}_{i,j} = \big\{a \in F[i+j] : (V_1[1] \cup \dots \cup V_1[i]) \cap \partial a = \emptyset\big\} = \left\{a \in F[i+j] : \bigcup_{i+j-s+1<h\le i} V_1[h] \cap \partial a = \emptyset\right\}.$$

Claim 4.10. *With probability* $1-o(n^{-2})$ *for all* $s \le i < \ell$, $j \in [s]$ *we have* $|\mathscr{W}_{i,j}| = 2^{-(s-j)/s}\frac{m}{\ell}(1+O(n^{-\Omega(1)}))$.

Proof. We may condition on the event $\mathscr{E}$ that (i) from Proposition 4.1 occurs. To compute the mean of $|\mathscr{W}_{i,j}|$ fix a test $a \in F[i+j]$ and an index $i+j-s < h \le i$. Then (4.6) shows that the probability that a fixed individual $x \in V[h]$ joins a equals $\mathbb{P}[x \in \partial a] = \frac{\Delta\ell}{ms}(1+O(\frac{\Delta\ell}{ms}))$. Hence, the choices (4.2) and (4.5) of Δ and ℓ and the assumption $m = \Theta(k\ln n)$ ensure that

$$\begin{aligned}\mathbb{E}\big[\big|(V_1[i+j-s+1] \cup \dots \cup V_1[i]) \cap \partial a\big| \mid \mathscr{E}\big] &= (s-j)\left(\frac{\Delta\ell}{ms}\cdot\frac{k}{\ell} + O\left(\frac{\Delta^2 k}{m^2 s^2}\right) + O\left(\frac{\Delta\ell\sqrt{k}\ln n}{ms}\right)\right)\\ &= \frac{s-j}{s}\ln 2 + O(n^{-\Omega(1)}).\end{aligned} \tag{4.26}$$

Since by **SC1** the events $\{x \in \partial a\}_x$ are independent, $|V_1[h] \cap \partial a|$ is a binomial random variable for every h and all these random variables $(|V_1[h] \cap \partial a|)_h$ are mutually independent. Therefore, (4.26) implies that

$$\mathbb{P}\big[(V_1[i+j-s+1] \cup \dots V_1[i]) \cap \partial a = \emptyset \mid \mathscr{E}\big] = 2^{-(s-j)/s} + O(n^{-\Omega(1)}). \tag{4.27}$$

Hence,

$$\mathbb{E}\big[|\mathscr{W}_{i,j}| \mid \mathscr{E}\big] = \sum_{a \in F[i+j]} \mathbb{P}\big[(V_1[i+j-s+1] \cup \dots \cup V_1[i]) \cap \partial a = \emptyset \mid \mathscr{E}\big] = \frac{m}{\ell}2^{-(s-j)/s}(1+O(n^{-\Omega(1)})). \tag{4.28}$$

Finally, changing the neighbourhood ∂x of one infected individual $x \in V_1$ can alter $|\mathscr{W}_{i,j}|$ by at most Δ. Therefore, the Azuma–Hoeffing inequality shows that for any $t > 0$,

$$\mathbb{P}\big[\big||\mathscr{W}_{i,j}| - \mathbb{E}[|\mathscr{W}_{i,j}| \mid \mathscr{E}]\big| > t \mid \mathscr{E}\big] \le 2\exp\left(-\frac{t^2}{2k\Delta^2}\right). \tag{4.29}$$

Combining (4.28) and (4.29), applied with $t = \sqrt{m}\ln^2 n$, and taking a union bound on i, j completes the proof. □

As a next step we use Claim 4.10 to estimate the as yet unexplained tests counts $\boldsymbol{W}_{x,j}$ from (4.7).

Claim 4.11. *For all* $s \le i < \ell$, $x \in V_1[i+1]$ *and* $j \in [s]$ *we have*

$$\mathbb{P}\left[\boldsymbol{W}_{x,j} < (1-\varepsilon/2)2^{j/s-1}\Delta/s\right] \le \exp\left(-\frac{\Omega(\ln n)}{(\ln\ln n)^4}\right).$$

Proof. Fix a pair of indices i, j and an individual $x \in V_1[i+1]$. We also condition on the event $\mathscr{E}$ that (i) from Proposition 4.1 occurs. Additionally, thanks to Claim 4.10 we may condition on the event

$$\mathscr{E}' = \left\{|\mathscr{W}_{i,j}| = 2^{-(s-j)/s}\frac{m}{\ell}(1+O(n^{-\Omega(1)}))\right\}.$$

Further, let $\mathfrak{E}$ be the σ-algebra generated by $\boldsymbol{\sigma}$ and by the neighbourhoods $(\partial y)_{y \in V[1]\cup\dots\cup V[i]}$. Recall from **SC1** that x simply joins Δ/s random tests in compartment $F[i+j]$, independently of all other individuals, and remember from (4.7) that $\boldsymbol{W}_{x,j}$ counts tests $a \in \mathscr{W}_{i+j} \cap \partial x$. Therefore, since the events $\mathscr{E}, \mathscr{E}'$ and the random variable $|\mathscr{W}_{i,j}|$ are $\mathfrak{E}$-measurable while ∂x is independent of $\mathfrak{E}$, given $\mathfrak{E}$ the random variable $\boldsymbol{W}_{x,j}$ has a hypergeometric distribution $\mathrm{Hyp}(m/\ell, |\mathscr{W}_{i,j}|, \Delta/s)$. Thus, the assertion follows from the hypergeometric Chernoff bound from Lemma 2.2 and the choice (4.14) of ζ. □

Proof of Lemma 4.4. Since $\boldsymbol{W}^\star_x = \sum_{j=1}^s w_j \boldsymbol{W}_{x,j}$, the lemma is an immediate consequence of Markov's inequality and Claim 4.11. □

4.8. **Proof of Lemma 4.5.** We need to derive the rate functions of the random variable $\boldsymbol{W}_{x,j}$ that count as yet unexplained tests for $x \in V_{0+}[i+1]$. To this end we first investigate the set of positive tests in compartment $i+j$ that do not contain any infected individuals from the first i compartments. In symbols,

$$\mathscr{P}_{i+1,j} = \{a \in F_1[i+j] : \partial a \cap (V_1[1] \cup \cdots \cup V_1[i]) = \emptyset\} \qquad (s \le i < \ell,\ j \in [s]).$$

Claim 4.12. *W.h.p. for all $s \le i < \ell$, $j \in [s]$ we have $|\mathscr{P}_{i+1,j}| = \big(1+O(n^{-\Omega(1)})\big)\big(2^{j/s}-1\big)\frac{m}{2\ell}$.*

Proof. We may condition on the event $\mathscr{E}$ that (i) from Proposition 4.1 occurs. As a first step we calculate the probability that $(V_1[i+1] \cup \cdots \cup V_1[i+j]) \cap \partial a \neq \emptyset$ for a specific test $a \in F[i+j]$. To this end we follow the steps of the proof of Claim 4.10. Since by (4.6) a specific individual $x \in V[h]$, $i < h \le i+j$, joins a with probability $\mathbb{P}[x \in \partial a] = (\Delta\ell/(ms))(1+O(\Delta\ell/(ms)))$ and since given $\mathscr{E}$ each compartment $V[h]$ contains $k/\ell + O(\sqrt{k/\ell}\ln n)$ infected individuals, we obtain, in perfect analogy to (4.26),

$$\mathbb{E}\big[\big|(V_1[i+1] \cup \cdots \cup V_1[i+j]) \cap \partial a\big| \mid \mathscr{E}\big] = \frac{j}{s}\ln 2 + O(n^{-\Omega(1)}). \tag{4.30}$$

Since the individuals $x \in V[i+1] \cup \cdots \cup V[i+j]$ join tests independently, (4.30) implies that

$$\mathbb{P}\big[(V_1[i+1] \cup \cdots \cup V_1[i+j]) \cap \partial a \neq \emptyset \mid \mathscr{E}\big] = 1 - 2^{-j/s} + O(n^{-\Omega(1)}). \tag{4.31}$$

Furthermore, we already verified in (4.27) that

$$\mathbb{P}\big[(V_1[i+j-s+1] \cup \cdots V_1[i]) \cap \partial a = \emptyset \mid \mathscr{E}\big] = 2^{-(s-j)/s} + O(n^{-\Omega(1)}). \tag{4.32}$$

Because the choices for the compartments $V[i+j-s+1] \cup \cdots \cup V[i+j]$ from which a draws its individuals are mutually independent, we can combine (4.31) with (4.32) to obtain

$$\mathbb{P}\left[\bigcup_{i+j-s<h\le i} V_1[h] \cap \partial a = \emptyset \neq \bigcup_{i<h\le i+j} V_1[h] \cap \partial a \mid \mathscr{E}\right] = \frac{2^{j/s}-1}{2} + O(n^{-\Omega(1)}). \tag{4.33}$$

Further, (4.33) implies

$$\mathbb{E}\big[|\mathscr{P}_{i+1,j}| \mid \mathscr{E}\big] = \mathbb{E}\left[\left|\left\{a \in F_1[i+j] : \bigcup_{h\le i} V_1[h] \cap \partial a = \emptyset \neq \bigcup_{i<h} V_1[h] \cap \partial a\right\}\right| \mid \mathscr{E}\right] = (2^{j/s}-1)\frac{m}{2\ell}\big(1+O(n^{-\Omega(1)})\big). \tag{4.34}$$

Finally, altering the neighbourhood ∂x of any infected individual can shift $|\mathscr{P}_{i+1,j}|$ by at most Δ. Therefore, the Azuma–Hoeffding inequality implies that

$$\mathbb{P}\big[\big||\mathscr{P}_{i+1,j}| - \mathbb{E}[|\mathscr{P}_{i+1,j}| \mid \mathscr{E}]\big| > t \mid \mathscr{E}\big] \le 2\exp\left(-\frac{t^2}{2k\Delta^2}\right). \tag{4.35}$$

Thus, the assertion follows from (4.5), (4.34) and (4.35) by setting $t = \sqrt{m}\ln^2 n$. □

Thanks to Proposition 4.1 (iii) and Lemma 4.12 in the following we may condition on the event

$$\mathscr{U} = \left\{\forall s < i \le \ell, j \in [s] : |F_1[i+j]| = \big(1+O(n^{-\Omega(1)})\big)\frac{m}{2\ell} \wedge |\mathscr{P}_{i+1,j}| = \big(1+O(n^{-\Omega(1)})\big)\big(2^{j/s}-1\big)\frac{m}{2\ell}\right\}. \tag{4.36}$$

As a next step we will determine the conditional distribution of $\boldsymbol{W}_{x,j}$ for $x \in V_{0+}[i+1]$ given $\mathscr{U}$.

Claim 4.13. *Let $s < i \le \ell$ and $j \in [s]$. Given $\mathscr{U}$ for every $x \in V_{0+}[i+1]$ we have*

$$\boldsymbol{W}_{x,j} \sim \mathrm{Hyp}\left(\big(1+O(n^{-\Omega(1)})\big)\frac{m}{2\ell}, \big(1+O(n^{-\Omega(1)})\big)\big(2^{j/s}-1\big)\frac{m}{2\ell}, \frac{\Delta}{s}\right). \tag{4.37}$$

Proof. By **SC1** each individual $x \in V_{0+}[i+1]$ joins Δ/s positive test from $F[i+j]$, drawn uniformly without replacement. Moreover, by (4.7) given $x \in V_{0+}[i+1]$ the random variable $\boldsymbol{W}_{x,j}$ counts the number of tests $a \in \mathscr{P}_{i+1,j} \cap \partial x$. Therefore, $\boldsymbol{W}_{x,j} \sim \mathrm{Hyp}(|F_1[i+j]|, |\mathscr{P}_{i+1,j}|, \Delta/s)$. Hence, given $\mathscr{U}$ we obtain (4.37). □

The estimate (4.37) enables us to bound the probability that $\boldsymbol{W}^\star_x$ gets 'too large'.

Claim 4.14. *Let*

$$\mathscr{M} = \min \frac{1}{s}\sum_{j=1}^{s-1} \mathbf{1}\left\{z_j \geq 2^{j/s}-1\right\} D_{\mathrm{KL}}\left(z_j \| 2^{j/s}-1\right)$$
$$s.t. \quad \sum_{j=1}^{s-1}\left(z_j - (1-2\zeta)2^{j/s-1}\right) w_j = 0, \qquad z_1,\dots,z_{s-1} \in [0,1].$$

Then for all $s \leq i < \ell$ and all $x \in V[i+1]$ we have

$$\mathbb{P}\left[\boldsymbol{W}^\star_x > (1-2\zeta)\frac{\Delta}{s}\sum_{j=1}^{s-1} 2^{j/s-1} w_j \mid \mathscr{U}, x \in V_{0+}[i+1]\right] \leq \exp(-(1+o(1))\mathscr{M}\Delta).$$

Proof. Let $s \leq i < \ell$ and $x \in V_{0+}[i+1]$. Step **SC1** of the construction of $\boldsymbol{G}$ ensures that the random variables $(\boldsymbol{W}_{x,j})_{j\in[s]}$ are independent because the tests in the various compartments $F[i+j]$, $j \in [s]$, that x joins are drawn independently. Therefore, the definition (4.12) of $\boldsymbol{W}^\star_x$ and Lemma 4.13 yield

$$\begin{aligned}\mathbb{P}\left[\boldsymbol{W}^\star_x > (1-2\zeta)\frac{\Delta}{s}\sum_{j=1}^{s-1} 2^{j/s-1} w_j \mid \mathscr{U}, x \in V_{0+}[i+1]\right] &= \mathbb{P}\left[\sum_{j=1}^{s-1} w_j \boldsymbol{W}_{x,j} \geq \frac{1-2\zeta}{s}\sum_{j=1}^{s-1} 2^{j/s-1} w_j \mid \mathscr{U}, x \in V_{0+}[i+1]\right] \\ &\leq \sum_{y_1,\dots,y_s=0}^{\Delta} \mathbf{1}\left\{\sum_{j=1}^{s-1} w_j y_j \geq \frac{1-2\zeta}{s}\sum_{j=1}^{s-1} 2^{j/s-1} w_j\right\} \prod_{j=1}^{s-1} \mathbb{P}[\boldsymbol{W}_{x,j} \geq y_j \mid \mathscr{U}, x \in V_{0+}[i+1]]. \end{aligned} \tag{4.38}$$

Further, let

$$\mathscr{Z} = \left\{(z_1,\dots,z_{s-1}) \in [0,1]^{s-1} : \sum_{j=1}^{s-1}\left(z_j - (1-2\zeta)2^{j/s-1}\right) w_j = 0\right\}.$$

Substituting $y_j = \Delta z_j / s$ in (4.38) and bounding the total number of summands by $(\Delta+1)^s$, we obtain

$$\mathbb{P}\left[\boldsymbol{W}^\star_x > (1-2\zeta)\frac{\Delta}{s}\sum_{j=1}^{s-1} 2^{j/s-1} w_j \mid \mathscr{U}, x \in V_{0+}[i+1]\right] \leq (\Delta+1)^s \max_{(z_1,\dots,z_s)\in\mathscr{Z}} \prod_{j=1}^{s-1} \mathbb{P}[\boldsymbol{W}_{x,j} \geq \Delta z_j/s \mid \mathscr{U}, x \in V_{0+}[i+1]]. \tag{4.39}$$

Moreover, Claim 4.13 and the Chernoff bound from Lemma 2.2 yield

$$\mathbb{P}[\boldsymbol{W}_{x,j} \geq \Delta z_j/s \mid \mathscr{U}, x \in V_{0+}[i+1]] \leq \exp\left(-\mathbf{1}\left\{z_j \geq p_j\right\}\frac{\Delta}{s} D_{\mathrm{KL}}\left(z_j \| p_j\right)\right) \qquad \text{where } p_j = 2^{j/s}-1+O(n^{-\Omega(1)}).$$

Consequently, since (4.5) and the assumption $m = \Theta(k\ln n)$ ensure that $\Delta = \Theta(\ln n)$, we obtain

$$\mathbb{P}[\boldsymbol{W}_{x,j} \geq \Delta z_j/s \mid \mathscr{U}, x \in V_{0+}[i+1]] \leq \exp\left(-\mathbf{1}\left\{z_j \geq 2^{j/s}-1\right\}\frac{\Delta}{s} D_{\mathrm{KL}}\left(z_j \| 2^{j/s}-1\right) + O(n^{-\Omega(1)})\right). \tag{4.40}$$

Finally, the assertion follows from (4.39) and (4.40). □

As a next step we solve the optimisation problem $\mathscr{M}$ from Claim 4.14.

Claim 4.15. *We have $\mathscr{M} = 1 - \ln 2 + O(\ln(s)/s)$.*

Proof. Fixing an auxiliary parameter $\delta \geq 0$ we set up the Lagrangian

$$\mathscr{L}_\delta(z_1,\dots,z_s,\lambda) = \sum_{j=1}^{s-1}\left(\mathbf{1}\left\{z_j \geq 2^{j/s}-1\right\} + \delta\mathbf{1}\left\{z_j < 2^{j/s}-1\right\}\right) D_{\mathrm{KL}}\left(z_j \| 2^{j/s}-1\right) + \frac{\lambda}{s}\sum_{j=1}^{s-1} w_j\left(z_j - (1-2\zeta)2^{j/s-1}\right).$$

The partial derivatives come out as

$$\frac{\partial\mathscr{L}_\delta}{\partial\lambda} = -\frac{1}{s}\sum_{j=1}^{s-1}((1-2\zeta)2^{j/s-1} - z_j) w_j, \quad \frac{\partial\mathscr{L}_\delta}{\partial z_j} = -\lambda w_j + \left(\mathbf{1}\left\{z_j \geq 2^{j/s}-1\right\} + \delta\mathbf{1}\left\{z_j < 2^{j/s}-1\right\}\right)\ln\frac{z_j(2-2^{j/s})}{(1-z_j)(2^{j/s}-1)}.$$

Set $z^*_j = (1-2\zeta)2^{j/s-1}$ and $\lambda^* = 1$. Then clearly

$$\frac{\partial\mathscr{L}_\delta}{\partial\lambda}\Big|_{\lambda^*, z^*_1,\dots,z^*_{s-1}} = 0. \tag{4.41}$$

Moreover, the choice (4.14) of ζ guarantees that $z_j^* \geq 2^{j/s}-1$. Hence, by the choice (4.14) of the weights w_j,

$$\frac{\partial \mathscr{L}_\delta}{\partial z_j}\Big|_{\lambda^*,z_1^*,\dots,z_{s-1}^*} = 0. \tag{4.42}$$

Since $\mathscr{L}_\delta(y_1,\dots,y_s,\lambda)$ is strictly convex in $z_1,\dots,z_s$ for every $\delta>0$, (4.41)–(4.42) imply that $\lambda^*, z_1^*,\dots,z_{s-1}^*$ is a global minimiser. Furthermore, since this is true for any $\delta>0$ and since $z_j^* \geq 2^{j/s}-1$, we conclude that $(z_1^*,\dots,z_{s-1}^*)$ is an optimal solution to the minimisation problem $\mathscr{M}$. Hence,

$$\mathscr{M} = \frac{1}{s}\sum_{j=1}^{s-1} D_{\mathrm{KL}}\Big(z_j^* \| 2^{j/s}-1\Big) = \frac{1}{s}\sum_{j=1}^{s-1} D_{\mathrm{KL}}\Big((1-2\zeta)2^{j/s-1}\|2^{j/s}-1\Big). \tag{4.43}$$

Since

$$\frac{\partial}{\partial\alpha} D_{\mathrm{KL}}\big((1-2\alpha)2^{z-1}\|2^z-1\big) = 2^z\big[-z\ln(2)+\ln(1-2^{z-1}+\alpha 2^z)-\ln(1-2^{z-1})-\ln(1-2\alpha)+\ln(2^z-1)\big],$$

we obtain $\frac{\partial}{\partial\alpha} D_{\mathrm{KL}}\big((1-2\alpha)2^{z-1}\|2^z-1\big) = O(\ln s)$ for all $z=1/s,\dots,(s-1)/s$ and $\alpha\in[0,2\zeta]$. Combining this bound with (4.43), we arrive at the estimate

$$\mathscr{M} = O(\zeta\ln s) + \frac{1}{s}\sum_{j=1}^{s-1} D_{\mathrm{KL}}\Big(2^{j/s-1}\|2^{j/s}-1\Big). \tag{4.44}$$

Additionally, the function $f: z\in[0,1]\mapsto D_{\mathrm{KL}}\big(2^{z-1}\|2^z-1\big)$ is strictly decreasing and convex. Indeed,

$$f'(z) = \frac{2^{z-1}\ln 2}{2^z-1}\left((2^z-1)\ln\left(\frac{2^z}{2^z-1}\right)-1\right), \qquad f''(z) = \big(2^{z-1}\ln^2 2\big)\left(\ln\left(\frac{2^z}{2^z-1}\right)+\frac{2-2^z}{(2^z-1)^2}\right).$$

The first derivative is negative because $2^{z-1}/(2^z-1)>0$ while $(2^z-1)\ln(2^z/(2^z-1))<1$ for all $z\in(0,1)$. Moreover, since evidently $f''(z)>0$ for all $z\in(0,1)$, we obtain convexity. Further, l'Hôpital's rule yields

$$D_{\mathrm{KL}}\big(2^{1/s-1}\|2^{1/s}-1\big) = O(\ln s).$$

As a consequence, we can approximate the sum (4.44) by an integral and obtain

$$\begin{aligned}\mathscr{M} &= O(\ln(s)/s) + \int_0^1 D_{\mathrm{KL}}\big(2^{z-1}\|2^z-1\big)\,\mathrm{d}z\\ &= O(\ln(s)/s) + \frac{2(1-z)\ln^2(2)+2^z\ln 2^z+(1-2^z)\ln(2^z-1)}{2\ln 2}\Big|_{z=0}^{z=1} = 1-\ln(2)+O(\ln(s)/s),\end{aligned}$$

as claimed. □

Proof of Lemma 4.5. Fix $s\leq i<\ell$ and let $\boldsymbol{X}_i$ be the number of $x\in V_{0+}[i]$ such that $\boldsymbol{W}_x^\star > (1-2\zeta)\frac{\Delta}{s}\sum_{j=1}^{s-1}2^{j/s-1}w_j$. Also recall that Proposition 4.1 (iii) and Claim 4.12 imply that $\mathbb{P}[\mathscr{U}] = 1-o(1)$. Combining Lemma 4.3 with Claims 4.14 and 4.15, we conclude that

$$\mathbb{E}[\boldsymbol{X}_i\,|\,\mathscr{U}] \leq \big(1+O\big(n^{-\Omega(1)}\big)\big)2^{-\Delta}n\exp(-(1-\ln(2)+o(1))\Delta) = \exp\left(\ln n-(1+o(1))\Delta\right). \tag{4.45}$$

Recalling the definition (4.5) of Δ and using the assumption that $m\geq(1+\varepsilon)m_{\mathrm{ad}}$ for a fixed $\varepsilon>0$, we obtain $\Delta\geq(1-\theta+\Omega(1))\ln n$. Combining this estimate with (4.45), we find

$$\mathbb{E}[\boldsymbol{X}_i\,|\,\mathscr{U}] \leq n^{\theta-\Omega(1)}. \tag{4.46}$$

Finally, the assertion follows from (4.46) and Markov's inequality. □

4.9. **Proof of Proposition 4.6.** The following lemma establishes an expansion property of $\boldsymbol{G}$. Specifically, if T is a small set of individuals, then there are few individuals x that share many tests with another individual from T.

Lemma 4.16. *Suppose that $m=\Theta(n^\theta\ln n)$. W.h.p. for any set $T\subset V$ of size at most $\exp(-\ln^{7/8}n)k$ we have*

$$\left|\left\{x\in V:\sum_{a\in\partial x\setminus F[0]}\mathbf{1}\{T\cap\partial a\setminus\{x\}\neq\emptyset\}\geq\ln^{1/4}n\right\}\right| \leq \frac{|T|}{3}.$$

Proof. Fix a set $T \subset V$ of size $t = |T| \le \exp(-\ln^{7/8} n)k$, a set $R \subset V$ of size $r = \lceil t/3 \rceil$ and let $\gamma = \lceil \ln^{1/4} n \rceil$. Furthermore, let $U \subset F[1] \cup \cdots \cup F[\ell]$ be a set of tests of size $\gamma r \le u \le \Delta t$. Additionally, let $\mathscr{E}(R,T,U)$ be the event that every test $a \in U$ contains two individuals from $R \cup T$. Then

$$\mathbb{P}\left[R \subset \left\{x \in V : \sum_{a \in \partial x \setminus F[0]} \mathbf{1}\{T \cap \partial a \setminus \{x\} \neq \emptyset\} \ge \gamma\right\}\right] \le \mathbb{P}\left[\mathscr{E}(R,T,U)\right]. \tag{4.47}$$

Hence, it suffices to estimate $\mathbb{P}[\mathscr{E}(R,T,U)]$.

Given a test $a \in U$ there are at most $\binom{r+t}{2}$ way to choose two individuals $x_a, x_a' \in R \cup T$. Moreover, (4.6) shows that the probability of the event $\{x_a, x_a' \in \partial a\}$ is bounded by $(1+o(1))(\Delta\ell/(ms))^2$. Therefore,

$$\mathbb{P}[\mathscr{E}(R,T,U)] \le \left[\binom{r+t}{2}\left(\frac{(1+o(1))\Delta\ell}{ms}\right)^2\right]^u.$$

Consequently, the event $\mathscr{E}(t,u)$ that there exist sets R,T,U of sizes $|R| = r = \lceil t/3 \rceil, |T| = t, |U| = u$ such that $\mathscr{E}(R,T,U)$ occurs has probability

$$\mathbb{P}[\mathscr{E}(t,u)] \le \binom{n}{r}\binom{n}{t}\binom{m}{u}\left[\binom{r+t}{2}\left(\frac{(1+o(1))\Delta\ell}{ms}\right)^2\right]^u.$$

Hence, the bounds $\gamma t/3 \le \gamma r \le u \le \Delta t$ yield

$$\begin{aligned}\mathbb{P}[\mathscr{E}(t,u)] &\le \binom{n}{t}^2\binom{m}{u}\left[\binom{2t}{2}\left(\frac{(1+o(1))\Delta\ell}{ms}\right)^2\right]^u \le \left(\frac{\mathrm{e}n}{t}\right)^{2t}\left(\frac{2\mathrm{e}\Delta^2\ell^2t^2}{ms^2u}\right)^u \\ &\le \left[\left(\frac{\mathrm{e}n}{t}\right)^{3/\gamma}\frac{6\mathrm{e}\Delta^2\ell^2 t}{\gamma ms^2}\right]^u \le \left[\left(\frac{\mathrm{e}n}{t}\right)^{3/\gamma}\cdot\frac{t\ln^4 n}{m}\right]^u \qquad \text{[due to (4.2), (4.5)].}\end{aligned}$$

Further, since $\gamma = \Omega(\ln^{1/4} n)$ and $m = \Omega(k\ln n)$ while $t \le \exp(-\ln^{7/8} n)k$, we obtain $\mathbb{P}[\mathscr{E}(t,u)] \le \exp(-u\sqrt{\ln n})$. Thus,

$$\sum_{\substack{1 \le t \le k^{1-\alpha} \\ \gamma t/3 \le u \le \Delta t}} \mathbb{P}[\mathscr{E}(t,u)] \le \sum_{1 \le u \le \Delta t} u\exp(-u\sqrt{\ln n}) = o(1). \tag{4.48}$$

Finally, the assertion follows from (4.47) and (4.48). □

Proof of Proposition 4.6. With τ the result of steps 1–10 of SPIV let $\mathscr{M}[i] = \{x \in V[i] : \tau_x \neq \boldsymbol{\sigma}_x\}$ be the set of misclassified individuals in compartment $V[i]$. Proposition 4.2 shows that w.h.p. $\mathscr{M}[i] = \emptyset$ for all $i \le s$. Further, we claim that for every $s \le i < \ell$ and any individual $x \in \mathscr{M}[i+1]$ one of the following three statements is true.

- **M1:** $x \in V_1[i+1]$ and $\boldsymbol{W}_x^\star < (1-\zeta/2)\frac{\Delta}{s}\sum_{j=1}^{s-1} 2^{j/s-1} w_j$,
- **M2:** $x \in V_{0+}[i+1]$ and $\boldsymbol{W}_x^\star > (1-2\zeta)\frac{\Delta}{s}\sum_{j=1}^{s-1} 2^{j/s-1} w_j$, or
- **M3:** $x \in V[i+1]$ and $\sum_{a \in \partial x} \mathbf{1}\{\partial a \cap (\mathscr{M}[1] \cup \cdots \cup \mathscr{M}[i]) \neq \emptyset\} \ge \ln^{1/4} n$.

To see this, assume that $x \in \mathscr{M}[i+1]$ while **M3** does not hold. Then comparing (4.7) and (4.17) we obtain

$$\left|W_{x,j}(\tau) - \boldsymbol{W}_{x,j}\right| \le \ln^{1/4} n \qquad \text{for all } 1 \le j < s. \tag{4.49}$$

Moreover, the definition (4.14) of the weights, the choice (4.3) of s, and the choices (4.14) of ζ and the weights w_j ensure that $0 \le w_j \le O(s) = O(\ln\ln n)$. This bound implies together with the definition (4.12) of the scores $\boldsymbol{W}_x^\star$ and (4.49) that

$$|\boldsymbol{W}_x^\star - W_x^\star(\tau)| = o(\zeta\Delta). \tag{4.50}$$

Thus, combining (4.50) with the definition of τ_x in Steps 5–10 of SPIV, we conclude that either **M1** or **M2** occurs.

Finally, to bound $\mathscr{M}[i+1]$ let $\mathscr{M}_1[i+1]$, $\mathscr{M}_2[i+1]$, $\mathscr{M}_3[i+1]$ be the sets of individuals $x \in V[i+1]$ for which **M1**, **M2** or **M3** occurs, respectively. Then Lemmas 4.4 and 4.5 imply that w.h.p.

$$|\mathscr{M}_1[i+1]|, |\mathscr{M}_2[i+1]| \le k\exp\left(-\frac{\ln n}{(\ln\ln n)^5}\right).$$

Furthermore, Lemma 4.16 shows that $|\mathscr{M}_3[i+1]| \le \sum_{h=1}^{i} |\mathscr{M}[h]|$ w.h.p. Hence, we obtain the relation

$$|\mathscr{M}[i+1]| \le k \exp\left(-\frac{\ln n}{(\ln\ln n)^5}\right) + \sum_{h=1}^{i} |\mathscr{M}[h]|. \tag{4.51}$$

Because (4.2) ensures that the total number of compartments is $\ell = O(\ln^{1/2} n)$, the bound (4.51) implies that $|\mathscr{M}[i+1]| \le O(\ell^2 k \exp(-(\ln n)/(\ln\ln n)^5)$ for all $i \in [\ell]$ w.h.p. Summing on i completes the proof. □

4.10. **Proof of Proposition 4.7.** For an infected individual $x \in V$ let

$$\boldsymbol{S}_x[j] = \left|\left\{a \in F[j] \cap \partial x : V_1 \cap \partial a = \{x\}\right\}\right| \qquad \text{and} \qquad \boldsymbol{S}_x = \sum_{j=1}^{\ell} \boldsymbol{S}_x[j].$$

Thus, $\boldsymbol{S}_x[j]$ is the number of positive sets $a \in F[j]$ that x has to itself, i.e., tests that do not contain a second infected individual, and $\boldsymbol{S}_x$ is the total number of such tests.

Lemma 4.17. *Assume that $m \ge (1+\varepsilon)m_{\inf}$. W.h.p. we have $\min_{x\in V_1} \boldsymbol{S}_x \ge \sqrt{\Delta}$.*

Proof. Due to Proposition 4.1 we may condition on the event

$$\mathscr{N} = \left\{\forall i \in [\ell] : \frac{m}{2\ell} - \sqrt{m}\ln n \le |F_0[i]| \le \frac{m}{2\ell} + \sqrt{m}\ln n\right\}.$$

We claim that given $\mathscr{N}$ for each $x \in V_1[i]$, $i \in [\ell]$, the random variable $\boldsymbol{S}_x$ has distribution

$$\boldsymbol{S}_x[i+j-1] \sim \mathrm{Hyp}\left(\frac{m}{\ell}, \frac{m}{2\ell} + O(\sqrt{m}\ln n), \frac{\Delta}{s}\right). \tag{4.52}$$

To see this, consider the set $F_x[i+j-1] = \left\{a \in F[i+j-1] : \partial a \cap V_1 \setminus \{x\} = \emptyset\right\}$ of all tests in compartment $F[i+j-1]$ without an infected individual besides possibly x. Since x joins $\Delta/s = O(\ln n)$ tests in $F[i+j-1]$, given $\mathscr{N}$ we have

$$\left|F_{0,x}[i+j]\right| = |F_0[i+j]| + O(\ln n) = \frac{m}{2\ell} + O(\sqrt{m}\ln n). \tag{4.53}$$

Furthermore, consider the experiment of first constructing the test design $\boldsymbol{G}$ and then re-sampling the set ∂x of neighbours of x; i.e., independently of $\boldsymbol{G}$ we have x join Δ/s random tests in each compartment $F[i+j]$. Then the resulting test design $\boldsymbol{G}'$ has the same distribution as $\boldsymbol{G}$ and hence the random variable $\boldsymbol{S}'_x[i+j-1]$ that counts tests $a \in F[i+j-1] \cap \partial x$ that do not contain another infected individual has the same distribution as $\boldsymbol{S}_x[i+j-1]$. Moreover, the conditional distribution of $\boldsymbol{S}'_x[i+j-1]$ given $\boldsymbol{G}$ reads

$$\boldsymbol{S}'_x[i+j-1] \sim \mathrm{Hyp}\left(\frac{m}{\ell}, |F_{0,x}[i+j-1]|, \frac{\Delta}{s}\right). \tag{4.54}$$

Combining (4.53) and (4.54), we obtain (4.52).

To complete the proof we combine (4.52) with Lemma 2.2, which implies that

$$\mathbb{P}\left[\boldsymbol{S}_x[i+j-1] \le \sqrt{\Delta} \mid x \in V_1\right] \le \exp\left(-\frac{\Delta}{s} D_{\mathrm{KL}}\left((1+o(1))s/\sqrt{\Delta} \| 1/2 + o(1)\right)\right) = \exp\left(-(1+o(1))\frac{\Delta \ln 2}{s}\right). \tag{4.55}$$

Since **SC1** ensures that the random variables $(\boldsymbol{S}_x[i+j-1])_{j\in[s]}$ are mutually independent, (4.55) yields

$$\mathbb{P}\left[\boldsymbol{S}_x \le \sqrt{\Delta} \mid x \in V_1\right] \le 2^{-(1+o(1))\Delta}. \tag{4.56}$$

Finally, the assumption $m \ge (1+\varepsilon)m_{\inf}$ for a fixed $\varepsilon > 0$ and the choice (4.5) of Δ ensure that $2^{-(1+o(1))\Delta} = o(1/k)$. Thus, the assertion follows from (4.56) by taking a union bound on $x \in V_1$. □

Proof of Proposition 4.7. For $j = 1 \ldots \lceil \ln n \rceil$, let

$$\mathscr{M}_j = \left\{x \in V : \tau_x^{(j)} \ne \boldsymbol{\sigma}_x\right\}$$

contain all individuals that remain misclassified at the j-th iteration of the clean-up step. Proposition 4.6 shows that w.h.p.

$$|\mathscr{M}_1| \le k \exp\left(-\frac{\ln n}{(\ln\ln n)^6}\right). \tag{4.57}$$

Furthermore, in light of Lemma 4.17 we may condition on the event $\mathscr{A} = \{\min_{x\in V_1} \boldsymbol{S}_x \ge \sqrt{\Delta}\}$.

We now claim that given $\mathscr{A}$ for every $j \geq 1$

$$\mathscr{M}_{j+1} \subset \left\{ x \in V : \sum_{a \in \partial x \setminus F[0]} \left| \partial a \cap \mathscr{M}_j \setminus \{x\} \right| \geq \lceil \ln^{1/4} n \rceil \right\}. \tag{4.58}$$

To see this, suppose that $x \in \mathscr{M}_{j+1}$ and recall that the assumption $m \geq m_{\mathrm{inf}}$ and (4.5) ensure that $\Delta = \Omega(\ln n)$. Also recall that SPIV's Step 15 thresholds the number

$$S_x(\tau^{(j)}) = \sum_{a \in \partial x : \hat{\boldsymbol{\sigma}}_a = 1} \mathbf{1}\left\{ \forall y \in \partial a \setminus \{x\} : \tau_y^{(j)} = 0 \right\}$$

of positive tests containing x whose other individuals are deemed uninfected. There are two cases to consider.

- **Case 1:** $x \in V_0$**:** in this case every positive tests $a \in \partial x$ contains an individual that is actually infected. Hence, if $\tau_y^{(j)} = 0$ for all $y \in \partial a \setminus \{x\}$, then $\partial a \cap \mathscr{M}_j \setminus \{x\} \neq \emptyset$. Consequently, since Step 15 of SPIV applies the threshold of $S_x(\tau^{(j)}) \geq \ln^{1/4} n$, there are at least $\ln^{1/4} n$ tests $a \in \partial x$ such that $\partial a \cap \mathscr{M}_j \setminus \{x\} \neq \emptyset$.
- **Case 2:** $x \in V_1$**:** given $\mathscr{A}$ every infected x participates in at least $\boldsymbol{S}_x \geq \sqrt{\Delta} = \Omega(\ln^{1/2} n)$ tests that do not actually contain another infected individual. Hence, if $S_x(\tau^{(j)}) \leq \ln^{1/4} n$, then at least $\sqrt{\Delta} - \ln^{1/4} n \geq \ln^{1/4} n$ tests $a \in \partial x$ contain an individual from $\mathscr{M}_j \setminus \{x\}$.

Thus, we obtain (4.58). Finally, (4.57), (4.58) and Lemma 4.16 show that w.h.p. $|\mathscr{M}_{j+1}| \leq |\mathscr{M}_j|/3$ for all $j \geq 1$. Consequently, $\mathscr{M}_{\lceil \ln n \rceil} = \emptyset$ w.h.p. □

5. Optimal adaptive group testing

In this final section we show how the test design $\boldsymbol{G}$ from Section 4 can be extended into an optimal two-stage adaptive design. The key observation is that Proposition 4.6, which summarises the analysis of the first two phases of SPIV (i.e., steps 1–10) only requires $m \geq (1+\varepsilon)m_{\mathrm{ad}}$ tests. In other words, the excess number $(1+\varepsilon)(m_{\mathrm{inf}} - m_{\mathrm{ad}})$ of tests required for non-adaptive group testing is necessary only to facilitate the clean-up step, namely phase 3 of SPIV.

Replacing phase 3 of SPIV by a second test stage, we obtain an optimal adaptive test design. To this end we follow Scarlett [32], who observed that a single-stage group testing scheme that correctly diagnoses all but $o(k)$ individuals with $(1+o(1))m_{\mathrm{ad}}$ tests could be turned into a two-stage design that diagnoses all individuals correctly w.h.p. with $(1+o(1))m_{\mathrm{ad}}$ tests in total. (Of course, at the time no such optimal single-stage test design and algorithm were known.) The second test stage works as follows. Let τ denote the outcome of phases 1 and 2 of SPIV applied to $\boldsymbol{G}$ with $m = (1+\varepsilon)m_{\mathrm{ad}}$.

- **T1:** Test every individual from the set $V_1(\tau) = \{x \in V : \tau_x = 1\}$ of individuals that SPIV diagnosed as infected separately.
- **T2:** To the individuals $V_0(\tau) = \{x \in V : \tau_x = 0\}$ apply the random d-out design and the DD-algorithm from Section 4.1 with a total of $m = k$ tests and $d = \lceil 10 \ln n \rceil$.

Let $\tau' \in \{0,1\}^V$ be the result of **T1**–**T2**.

Proposition 5.1. *W.h.p. we have $\tau'_x = \boldsymbol{\sigma}_x$ for all $x \in V$.*

As a matter of course **T1** renders correct results, i.e., for all individuals $x \in V_1(\tau)$ we have $\tau'_x = \boldsymbol{\sigma}_x$. Further, to analyse **T2** we use a similar argument as in the analysis of the first phase of SPIV in Section 4.5; we include the analysis for the sake of completeness. We begin by investigating the number of negative tests. Let $\boldsymbol{G}'$ denote the test design set up by **T2**, let $F' = \{b_1, \ldots, b_k\}$ denote its set of tests and let $\hat{\boldsymbol{\sigma}}_{b_1}, \ldots, \hat{\boldsymbol{\sigma}}_{b_k}$ signify the corresponding test results. Further, let $F'_0 = \{b \in F' : \hat{\boldsymbol{\sigma}}_b = 0\}$ and $F'_1 = \{b \in F' : \hat{\boldsymbol{\sigma}}_b = 1\}$ be the set of negative and positive tests, respectively.

Lemma 5.2. *W.h.p. we have $|F'_1| \leq \frac{k}{2}$.*

Proof. Proposition 4.6 implies that w.h.p.

$$|V_0(\tau) \cap V_1| \leq \sum_{x \in V} \mathbf{1}\{\tau_x \neq \boldsymbol{\sigma}_x\} \leq k \exp\left(-\frac{\ln n}{(\ln \ln n)^6} \right). \tag{5.1}$$

Moreover, since every individual $x \in V_0(\tau)$ joins d random tests, for any specific test $b \in F'$ we have

$$\mathbb{P}\left[x \in \partial_{\boldsymbol{G}'} b\right] = 1 - \mathbb{P}\left[x \notin \partial_{\boldsymbol{G}'} b\right] = 1 - \binom{k-1}{d}\binom{k}{d}^{-1} = \frac{d}{k}(1 + O(n^{-\Omega(1)})).$$

Hence, for every test $b \in F'$,

$$\mathbb{E}\left[|\partial b \cap V_1| \,\middle|\, |V_0(\tau) \cap V_1| \le k \exp\left(-\frac{\ln n}{(\ln\ln n)^6}\right)\right] = O(1/\ln n).$$

Consequently,

$$\mathbb{E}\left[|F_1'| \,|\, |V_0(\tau) \cap V_1| \le k/\ln n\right] = O(k/\ln n). \tag{5.2}$$

Finally, combining (5.1) and (5.2) and applying Markov's inequality, we conclude that $|F_1'| \le \frac{k}{2}$ w.h.p. □

Corollary 5.3. *W.h.p. for every $x \in V_0(\tau)$ there is a test $b \in F'$ such that $\partial b \setminus \{x\} \subset V_0$.*

Proof. We construct the random graph $\boldsymbol{G}'$ in two rounds. In the first round we first expose the neighbourhoods $(\partial_{\boldsymbol{G}'} y)_{y \in V_0(\tau)\setminus\{x\}}$. Lemma 5.2 implies that after the first round the number $\boldsymbol{X}$ of tests that do not contain an infected individual $y \in V_0(\tau) \cap V_1$ exceeds $k/2$ w.h.p. In the second round we expose $\partial_{\boldsymbol{G}'} x$. Because $\partial_{\boldsymbol{G}'} x$ is chosen independently of the neighbourhoods $(\partial_{\boldsymbol{G}'} y)_{y \in V_0(\tau)\setminus\{x\}}$, the number of tests $b \in \partial_{\boldsymbol{G}'} x$ that do not contain an infected individual $y \in V_0(\tau) \cap V_1$ has distribution $\mathrm{Hyp}(k, \boldsymbol{X}, d)$. Therefore, since $d \ge 10 \ln n$ we obtain

$$\mathbb{P}\left[\forall b \in \partial x : V_1 \cap \partial b \setminus \{x\} \ne \emptyset \mid \boldsymbol{X} \le k/2\right] \le \mathbb{P}\left[\mathrm{Hyp}(k, k/2, d) = 0\right] \le 2^{-d} = o(1/n). \tag{5.3}$$

Finally, the assertion follows (5.3) and the union bound. □

Proof of Proposition 5.1. Corollary 5.3 shows that we may assume that for every $x \in V_0(\tau)$ there is a test $b_x \in F'$ with $\partial b_x \setminus \{x\} \subset V_0$. As a consequence, upon executing the first step **DD1** of the `DD` algorithm, **T2** will correctly diagnose all individuals $x \in V_0(\tau) \cap V_0$. Therefore, if $x \in V_0(\tau) \cap V_1$, then **DD2** will correctly identify x as infected because all other individuals $y \in \partial b_x$ were already identified as healthy by **DD1**. Thus, $\tau'_x = \boldsymbol{\sigma}_x$ for all $x \in V$. □

Proof of Theorem 1.3. Proposition 5.1 already establishes that the output of the two-stage adaptive test is correct w.h.p. Hence, to complete the proof we just observe that the total number of tests comes to $(1+\varepsilon)m_{\mathrm{ad}}$ for the first stage plus $|V_1(\tau)| + k$ for the second stage. Furthermore, Proposition 4.6 implies that w.h.p.

$$|V_1(\tau)| \le |V_1| + \sum_{x \in V} \mathbf{1}\{\tau_x \ne \boldsymbol{\sigma}_x\} \le k\left(1 + \exp\left(-\frac{\ln n}{(\ln\ln n)^6}\right)\right) = (1 + o(1))k.$$

Thus, the second stage conducts $O(k) = o(m_{\mathrm{ad}})$ tests. □

Acknowledgment. We thank Arya Mazumdar for bringing the group testing problem to our attention.

References

[1] E. Abbe: Community detection and stochastic block models: recent developments. The Journal of Machine Learning Research **18** (2017) 6446–6531.

[2] A. Alaoui, A. Ramdas, F. Krzakala, L. Zdeborová, M. Jordan: Decoding from pooled data: phase transitions of message passing. IEEE Transactions on Information Theory **65** (2019) 572–585.

[3] A. Alaoui, A. Ramdas, F. Krzakala, L. Zdeborová, M. Jordan: Decoding from pooled data: Sharp information-theoretic bounds. SIAM Journal on Mathematics of Data Science **1** (2019) 161–188.

[4] M. Aldridge, L. Baldassini, O. Johnson: Group testing algorithms: bounds and simulations. IEEE Transactions on Information Theory **60** (2014) 3671–3687.

[5] M. Aldridge: Individual testing is optimal for nonadaptive group testing in the linear regime. IEEE Transactions on Information Theory **65** (2019) 2058–2061.

[6] M. Aldridge, O. Johnson, J. Scarlett: Group testing: an information theory perspective. Foundations and Trends in Communications and Information Theory (2019).

[7] N. Alon, M. Krivelevich, B. Sudakov: Finding a large hidden clique in a random graph. Proc. 9th SODA (1998) 594–598.

[8] T. Berger, V. Levenshtein: Asymptotic efficiency of two-stage disjunctive testing. IEEE Transactions on Information Theory, **48** (2002) 1741–1749.

[9] M. Brennan, G. Bresler: Optimal average-case reductions to sparse pca: From weak assumptions to strong hardness. arXiv:1902.07380.

[10] H. Chen, F. Hwang: A survey on nonadaptive group testing algorithms through the angle of decoding. Journal of Combinatorial Optimization **15** (2008) 49–59.

[11] A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimroth, P. Loick: Information-theoretic and algorithmic thresholds for group testing. Proc. 46th ICALP (2019) #43.

[12] A. Decelle, F. Krzakala, C. Moore, L. Zdeborová: Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. Phys. Rev. E **84** (2011) 066106.
[13] D. Donoho: Compressed sensing. IEEE Transactions on Information Theory **52** (2006) 1289–1306.
[14] D. Donoho, A. Javanmard , A. Montanari: Information-theoretically optimal compressed sensing via spatial coupling and approximate message passing. IEEE Transactions on Information Theory **59** (2013) 7434–7464.
[15] R. Dorfman: The detection of defective members of large populations. Annals of Mathematical Statistics **14** (1943) 436–440.
[16] A. D'yachkov, V. Rykov: Bounds on the length of disjunctive codes. Problemy Peredachi Informatsii **18** (1982) 166–171.
[17] P. Erdős, A. Rényi: On Two Problems of Information Theory. Magyar Tud. Akad. Mat. Kutató Int. Közl **8** (1963) 229–243.
[18] A. Felstrom, K. Zigangirov: Time-varying periodic convolutional codes with low-density parity-check matrix. IEEE Transactions on Information Theory **45** (1999) 2181–2191.
[19] W. Hoeffding: Probability inequalities for sums of bounded random variables. In N. Fisher, P. Sen (eds.): The collected works of Wassily Hoeffding. Springer Series in Statistics (Perspectives in Statistics). Springer, New York, NY (1994) 409–426.
[20] F. Hwang: A method for detecting all defective members in a population by group testing. Journal of the American Statistical Association **67** (1972) 605–608.
[21] S. Janson, T. Luczak, A. Rucinski: Random Graphs. John Wiley & Sons (2011).
[22] O. Johnson, M. Aldridge, J. Scarlett: Performance of group testing algorithms with near-constant tests per item. IEEE Transactions on Information Theory **65** (2018) 707–723.
[23] W. Kautz, R. Singleton: Nonrandom binary superimposed codes. IEEE Transactions on Information Theory **10** (1964), 363–377.
[24] F. Krzakala, M. Mézard, F. Sausset, Y. Sun, L. Zdeborová: Statistical-physics-based reconstruction in compressed sensing. Physical Review X **2** (2012) 021005.
[25] S. Kudekar, H. Pfister: The effect of spatial coupling on compressive sensing. Proc. 48th Allerton (2010) 347–353.
[26] S. Kudekar, T. Richardson, R. Urbanke: Threshold saturation via spatial coupling: why convolutional LDPC ensembles perform so well over the BEC. IEEE Transaction on Information Theory **57** (2011) 803–834.
[27] S. Kudekar, T. Richardson, R. Urbanke: Spatially coupled ensembles universally achieve capacity under belief propagation. IEEE Transaction on Information Theory **59** (2013) 7761–7813.
[28] H. Kwang-Ming, D. Ding-Zhu: Pooling designs and nonadaptive group testing: important tools for DNA sequencing. World Scientific (2006).
[29] M. Mézard, M. Tarzia, C. Toninelli: Group testing with random pools: phase transitions and optimal strategy. Journal of Statistical Physics **131** (2008) 783–801.
[30] C. Moore: The computer science and physics of community detection: landscapes, phase transitions, and hardness. Bulletin of the EATCS **121** (2017).
[31] G. Reeves, H. Pfister (2019). Understanding phase transitions via mutual information and MMSE. arXiv:1907.02095.
[32] J. Scarlett: Noisy adaptive group testing: Bounds and algorithms. IEEE Transactions on Information Theory **65** (2018) 3646–3661.
[33] J. Scarlett: An efficient algorithm for capacity-approaching noisy adaptive group testing. Proc. IEEE International Symposium on Information Theory (2019) 2679–2683.
[34] K. Takeuchi, T. Tanaka, T. Kawabata: Improvement of BP-based CDMA multiuser detection by spatial coupling. Proc. IEEE International Symposium on Information Theory Proceedings (2011) 1489–1493.
[35] P. Ungar: The cutoff point for group testing. Communications on Pure and Applied Mathematics **13** (1960) 49–54.
[36] L. Wang, X. Li, Y. Zhang, K. Zhang: Evolution of scaling emergence in large-scale spatial epidemic spreading. PLoS ONE **6** (2011).
[37] Y. Wu, S. Verdu, Renyi information dimension: fundamental limits of almost lossless analog compression. IEEE Transactions on Information Theory **56** (2010) 3721–3748.
[38] L. Zdeborová, F. Krzakala: Statistical physics of inference: thresholds and algorithms. Advances in Physics **65** (2016) 453–552.

AMIN COJA-OGHLAN, acoghlan@math.uni-frankfurt.de, GOETHE UNIVERSITY, MATHEMATICS INSTITUTE, 10 ROBERT MAYER ST, FRANKFURT 60325, GERMANY.

OLIVER GEBHARD, gebhard@math.uni-frankfurt.de, GOETHE UNIVERSITY, MATHEMATICS INSTITUTE, 10 ROBERT MAYER ST, FRANKFURT 60325, GERMANY.

MAX HAHN-KLIMROTH, hahnklim@math.uni-frankfurt.de, GOETHE UNIVERSITY, MATHEMATICS INSTITUTE, 10 ROBERT MAYER ST, FRANKFURT 60325, GERMANY.

PHILIPP LOICK, loick@math.uni-frankfurt.de, GOETHE UNIVERSITY, MATHEMATICS INSTITUTE, 10 ROBERT MAYER ST, FRANKFURT 60325, GERMANY.

APPENDIX D. IMPROVED BOUNDS FOR NOISY GROUP TESTING WITH CONSTANT TESTS PER ITEM

IMPROVED BOUNDS FOR NOISY GROUP TESTING WITH CONSTANT TESTS PER ITEM

OLIVER GEBHARD, OLIVER JOHNSON, PHILIPP LOICK, MAURICE ROLVIEN

{Gebhard, Loick, Rolvien}@math.uni-frankfurt.de, *Goethe University, Mathematics Institute, 10 Robert Mayer St, Frankfurt 60325, Germany.*

O.Johnson@bristol.ac.uk, *University of Bristol, School of Mathematics, Woodland Road, Bristol, BS8 1UG, United Kingdom*

ABSTRACT. The group testing problem is concerned with identifying a small set of infected individuals in a large population. At our disposal is a testing procedure that allows us to test several individuals together. In an idealized setting, a test is positive if and only if at least one infected individual is included and negative otherwise. Significant progress was made in recent years towards understanding the information-theoretic and algorithmic properties in this noiseless setting. In this paper, we consider a noisy variant of group testing where test results are flipped with certain probability, including the realistic scenario where sensitivity and specificity can take arbitrary values. Using a test design where each individual is assigned to a fixed number of tests, we derive explicit algorithmic bounds for two commonly considered inference algorithms and thereby improve on results by Scarlett & Cevher (SODA 2016) and Scarlett & Johnson (2020) and providing the strongest performance guarantees currently proved for efficient algorithms in these noisy group testing models.

Oliver Gebhard and Philipp Loick are supported by DFG CO 646/3.

1. INTRODUCTION

1.1. **Motivation and background.** Suppose we have a large collection of n people, a small number k of whom are infected by some disease, and where only $m \ll n$ tests are available.

In a landmark paper [15] from 1943, Dorfman introduced the idea of group testing. The basic idea is as follows: rather than screen one person using one test, we could mix samples from individuals in one pool, and use a single test for this whole pool. The task is to recover the infection status of all individuals using the pooled test results.

Dorfman's original work was motivated by a biological application, namely identifying individuals with syphilis. Subsequently, group testing has found a number of related applications, including detection of HIV [51], DNA sequencing [30, 37] and protein interaction experiments [35, 49]. More recently, it has been recognised as an essential tool to moderate pandemic spread [12], where identifiying infected individuals fast and at a low cost is indispensable [33]. In particular, group testing has been identified as a testing scheme for the detection of COVID-19 [16, 19].

From a mathematical perspective, group testing is a prime example of an inference problem where one wants to learn a ground truth from (possibly noisy) measurements [1, 2, 9, 21, 22, 28, 42]. Over the last decade, it has regained popularity and today is a field of active research. Results on its information-theoretic and algorithmic properties were recently presented by Scarlett et al. at SODA'16, ISIT'16, ISIT'19 [44, 46, 45], and Baldassini et al. at ISIT'13 [8] and Coja-Oghlan et al. at ICALP'19, COLT'20 [13, 14]. In this paper, we provide improved upper bounds on the number of tests that guarantee successful inference for the noisy variant of group testing.

1.2. **Related Work.** In the simplest version of group testing, we suppose that a test is positive if and only if the pool contains at least one infected individual. We refer to this as the noiseless case. In this setting, each negative test guarantees that every member of the corresponding pool is not infected, so they can be removed from further consideration. However, a positive test only tells us that at least one item in the test is defective (but not which one), and so requires further investigation.

Dorfman's original work [15] proposed a simple adaptive strategy where a small pool of individuals is tested, and where each positive test is followed up by testing every individual in the corresponding pool individually. Since then it has been an important problem to find the optimal way to recover the population's infection status in the noiseless case. A simple counting argument (see for example [7, Section 1.4]) shows that to ensure recovery with zero error probability, since every possible defective set must give different test outcomes, the following must hold in the noiseless setting:

$$2^m \geq \binom{n}{k} \qquad \Rightarrow \qquad m \geq m^0_{\inf} := \frac{1}{\log 2} k \log(n/k) \tag{1.1}$$

Hwang [24] provided an algorithm based on repeated binary search, which is essentially optimal in terms of the number of tests required in that it requires $m^0_{\inf} + O(k)$ tests, but may require many stages of testing. As described for example in pandemic plans developed by the EU, US and WHO [18, 38, 39], and in COVID-specific work [36], adaptive strategies may not be suitable for pandemic prevention. For example, if a test takes one day to prepare and for the results to be known, then each stage will require an extra day to perform, meaning that adaptive group testing information can be received too late to be useful.

Hence the need to perform large-scale testing to identify infected individuals fast relative to the doubling time [12, 33, 36] can make adaptive group testing unsuitable to prevent an infectious disease from spreading. Furthermore the preservation of uncharted viruses in a large scale may be challenging due to structural and chemical differences [20]. Due to its automation potential and the fact that tests can be completed in parallel (for example by the use of 96-well PCR plates [17]), the main application of group testing such as DNA screening [11, 30, 37], HIV testing [51] and protein interaction analysis [35, 49] are non-adaptive where all tests are specified upfront and performed in parallel.

The question of whether non-adaptive algorithms (or even adaptive algorithms with a limited number of stages) can attain the bound (1.1) remained open until recently. [4, 14] showed that the answer depends on the prevalence of the disease, for example on the value of $\theta \in (0,1)$ in a parameterisation where the number of infected individuals $k \sim n^\theta$. Non-adaptive testing schemes can be represented through a binary $(m \times n)$-matrix that represents which individual participates in which test. Significant research was dedicated to see which design attains the optimal performance. Since deterministic designs were shown to not attain the optimal order [7], research focused on randomized designs. Initial research focused on the case where the matrix entries are iid [3, 5, 45]. Later work considered a constant column design where each individual is assigned to a (near-)constant number of tests [6, 14, 13, 26]. Indeed [14] showed that such a design is information-theoretically optimal in the *noiseless* setting and it is to be expected that this remains true for the noisy case.

To recover the ground truth from the test results and the pooling scheme, this paper focuses on two non-adaptive algorithms, COMP and DD, which are relatively simple to perform and interpret in the noiseless case. We describe them in more detail below, but in brief COMP [10] simply builds a list of all the individuals who ever appear in a negative test and are hence certainly healthy, and assumes that the other individuals are infected. DD [5] uses COMP as a first stage and builds on it by looking for individuals who appear in a positive test that only otherwise contains individuals known to be healthy.

While the noiseless case provides an interesting mathematical abstraction, it is clear that it may not be realistic in practice [40]. In medical applications the two occurring types of noise in a testing procedure are related to sensitivity (positive correct) and specificity (negative correct), and in that language we cannot assume the gold standard of tests with unit specificity and sensitivity. Thus, research attention in recent years has shifted towards the noisy version of group testing [10, 43, 44, 45, 47, 48]. On the one hand, *adaptive* noisy case was considered in [43, 44]. On the other hand [10, 27, 29, 34, 45, 47, 48] looked at the *non-adaptive* noise case from different angles (for instance linear programming, belief propagation, Bernoulli-pooling, Markov-Chain Monte Carlo).

In this paper we focus on the COMP and DD algorithms, since it is possible to deduce explicit performance guarantees for them. The deductions made by the original COMP and DD algorithms are designed for the noiseless case and do not hold in general. However, recent work of Scarlett and Johnson [48] has shown that noisy versions of these algorithms can perform well under certain noise models using Bernoulli i.i.d. test designs, particularly focusing on Z channel and reverse Z channel noise.

As common medical tests have different values for sensitivity and specificity [32] the analysis of a generalized noise model beyond the Z and reverse Z channel is warranted. For example, while group testing strategies appear to be useful to identify individuals infected with COVID-19 (see for example [16, 19]), testing for the presence of the SARS-CoV-19 virus is not perfect [52], and so we need to understand the effect of both false positive and false negative errors in this context, with non-identical error probabilities. For this reason, we consider a general $p-q$ noise model in this paper. Under this model, a truly negative test is flipped with probability p to display a positive test result, while a truly positive test is flipped to negative with probability q (Figure 1). Its formulation is sufficiently general to accommodate the recovery of the noiseless results ($p = q = 0$), Z channel ($p = 0$), reverse Z channel ($q = 0$) and the Binary Symmetric Channel ($p = q$). However, our results include the case of non-zero p and q without having to make the somewhat artificial assumption that false negative and false positive errors are equally likely.

1.3. **Contribution.** This paper provides a simultaneous extension of [13] and [26, 48], by analysing noisy versions of COMP and DD under more general noise models for constant-column weight designs. We provide explicit bounds on the performance of these algorithms in a generalized noise model. For all typical noise channels (Z, reverse Z and BSC) we compare the constant-column and Bernoulli design and find for all such instances that the former meaningfully outperforms the latter thereby improving on results from [26] and providing the strongest performance guarantees currently proved for efficient algorithms in noisy group testing. As group testing offers an essential tool for pandemic prevention [33] and as the

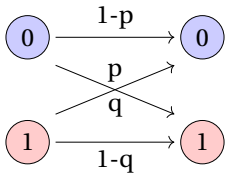

FIGURE 1. The p-q-noise model: the result of each standard noiseless group test is transmitted independently through the given noisy communication channel.

the accuracy of medical testing is limited [32, 40] this paper provides the natural next step in the group testing literature.

1.4. **Test design and notation.** To formalise our notation, we write n for the number of individuals in the population, $\boldsymbol{\sigma}$ for a binary vector representing the infection status of each individual, k (the Hamming weight of $\boldsymbol{\sigma}$) for the number of infected individuals and m for the number of tests performed. We assume that k is known for purposes of matrix design, though in practice (see [7, Remark 2.3]) it is generally enough to know k up to a constant factor to design a matrix with good properties. In this paper, in line with other work such as [5], we consider a scaling $k \sim n^\theta$ for some fixed $\theta \in (0,1)$, referred to in [7, Remark 1.1] as the sparse regime. We believe a similar analysis should be possible in the very sparse regime ($k = O(1)$) and linear regime ($k \sim \beta n$ for a fixed β). In addition to the interesting phase transitions observed using this scaling, this sparse regime is particularly relevant as it is the parametrisation to model the early state of a pandemic [50].

Let us next introduce the test design. With $V = (x_i)_{i\in[n]}$ [1] denoting the set of n individuals and $F = (a_i)_{i\in[m]}$ the set of m tests, the test design can be envisioned as a bipartite factor graph with n variable nodes "on the left" and m factor nodes "on the right". We draw a configuration $\boldsymbol{\sigma} \in \{0,1\}^V$, encoding the infection status of each individual, uniformly at random from vectors of Hamming weight k. The set of healthy individuals will be denoted by V_0 and the set of infected individuals by V_1. In symbols,

$$V_0 = \{x \in V : \boldsymbol{\sigma}_x = 0\} \qquad \text{and} \qquad V_1 = V \setminus V_0 = \{x \in V : \boldsymbol{\sigma}_x = 1\}$$

The lower bound from (1.1) suggests that in the noisy group testing setting it is natural to compare the performance of algorithms and matrix designs in terms of the prefactor of $k\log(n/k)$ in the number of tests required. To be precise, we carry out m tests, and each item is assigned to exactly Δ tests chosen uniformly at random without replacement. We parameterise m and Δ as

$$m = ck\log(n/k) \qquad \text{and} \qquad \Delta = cd\log(n/k) \tag{1.2}$$

for some suitably chosen constants $c, d \geq 0$.

Let ∂x denote the set of tests that individual x appears in and ∂a the set of individuals assigned to test a. The resulting (non-constant) collection of test degrees will be denoted by the vector $\Gamma = (\Gamma_a)_{a\in[m]}$. Further, let

$$\Gamma_{\min} = \min_{a\in[m]} \Gamma_a \qquad \text{and} \qquad \Gamma_{\max} = \max_{a\in[m]} \Gamma_a. \tag{1.3}$$

Throughout, $\boldsymbol{G} = \boldsymbol{G}(n,m,\Delta)$ describes the random bipartite factor graph from this construction.

Now consider the outcome of the tests. Recall from above that a standard noiseless group test a gives a positive result if and only if there is at least one defective item contained in the pool, or equivalently if $\sum_{x\in\partial a} \boldsymbol{\sigma}(x) > 0$. Even in the noisy case, this sum is a useful object to consider. Writing $\mathbf{1}$ for the indicator

[1] $[n]$ will be used as an abbreviated notation for the set $\{1,\dots,n\}$.

function, we define

$$\boldsymbol{\sigma}^*(a) = \mathbf{1}\left\{\sum_{x\in\partial a}\boldsymbol{\sigma}(x) > 0\right\} \tag{1.4}$$

to be the outcome we would observe in the noiseless case using the test matrix corresponding to $\boldsymbol{G}$. We will say that test a is *truly positive* if $\boldsymbol{\sigma}^*(a) = 1$ and truly negative otherwise.

However, we do not observe the values of $\boldsymbol{\sigma}^*(a)$ directly, but rather see what we will refer to as the *displayed* test outcomes $\hat{\boldsymbol{\sigma}}(a)$ – the outcomes of sending the true outcomes $\boldsymbol{\sigma}^*(a)$ independently through the $p-q$ channel of Figure 1. Since in this model a truly positive test remains positive with probability $1-q$ and a truly negative test is displayed as positive with probability p we can write

$$\hat{\boldsymbol{\sigma}}(a) = \mathbf{1}\left\{\mathrm{Be}(p)=1\right\}\left(1-\boldsymbol{\sigma}^*(a)\right) + \mathbf{1}\left\{\mathrm{Be}(1-q)=1\right\}\boldsymbol{\sigma}^*(a) \tag{1.5}$$

where $\mathrm{Be}(r)$ denotes a Bernoulli random variable with parameter r. For models with binary outputs, this is the most general channel satisfying the noisy defective channel property of [7, Definition 3.3], though more general models are possible under the only defects matter property [7, Definition 3.2], where the probability of a test being positive depends on the number of contained infected individuals.

Note that if $p+q>1$, we can preprocess the outputs from (1.5) by flipping them, i.e. setting $\widetilde{p} = 1-p$ and $\widetilde{q} = 1-q$, where $\widetilde{p}+\widetilde{q} < 1$. Hence without loss of generality we will assume throughout that $p+q<1$. In the case $p+q=1$, the test outcomes are independent of the inputs, and we cannot hope to find the infected individuals – see Theorem 2.3.

With $\boldsymbol{m}_0$ being the number of truly negative tests, let $\boldsymbol{m}_0^f$ be the number of truly negative tests that are flipped to display a positive test result and $\boldsymbol{m}_0^u$ be the number of truly negative tests that are unflipped. Similarly, define $\boldsymbol{m}_1$ as the number of truly positive tests, of which $\boldsymbol{m}_1^f$ are flipped to a negative test result and of which $\boldsymbol{m}_1^u$ are unflipped. For reference, for $t\in\{0,1\}$ we write

$$\boldsymbol{m}_t = \left|\left\{a : \boldsymbol{\sigma}^*(a) = t\right\}\right|$$
$$\boldsymbol{m}_t^f = \left|\left\{a : \boldsymbol{\sigma}^*(a) = t, \hat{\boldsymbol{\sigma}}(a) \neq t\right\}\right| \quad \text{and} \quad \boldsymbol{m}_t^u = \left|\left\{a : \boldsymbol{\sigma}^*(a) = t, \hat{\boldsymbol{\sigma}}(a) = t\right\}\right|$$

Throughout the paper, we use the standard Landau notation $o(\cdot), O(\cdot), \Theta(\cdot), \Omega(\cdot), \omega(\cdot)$ and define $0\log 0 = 0$. In order to quantify the performance of our algorithms, for any $0 < r \neq s < 1$, we write

$$D_{\mathrm{KL}}(r\|s) := r\log\left(\frac{r}{s}\right) + (1-r)\log\left(\frac{1-r}{1-s}\right),$$

for the relative entropy of a Bernoulli random variable with parameter r to a Bernoulli random variable with parameter s, commonly referred to as the Kullback–Leibler divergence. Here and throughout the paper we use log to denote the natural logarithm. For r or s equal to 0 or 1 we define the value of $D_{\mathrm{KL}}(\cdot\|\cdot)$ (possibly infinite) on grounds of continuity, so for example $D_{\mathrm{KL}}(0\|s) = -\log(1-s)$.

2. Main results

With the test design and notation in place, we are now in a position to state our main results. Theorems 2.1, 2.2 and 2.3 are the centerpiece of this paper featuring improved bounds for the noisy group testing problem for the general $p-q$ model. We follow up with a discussion of the combinatorics underlying both algorithms. Subsequently, we show how the bounds simplify when we consider the special cases of the Z, the reverse Z and Binary Symmetric Channel. Finally, we derive sufficient conditions under which DD provably outperforms the COMP algorithm and compare the bounds of our constant-column design against the Bernoulli design employed in prior literature.

2.1. **Bounds for Noisy Group Testing.** We will consider two well-known algorithms from the noiseless setting to identify infected individuals in this paper. First, we study a noisy variant of the COMP algorithm, originally introduced in [10].

1 Declare every individual that appears in $\alpha\Delta$ or more displayed negative tests as healthy.
2 Declare all remaining individuals as infected.

Algorithm 1: The noisy COMP algorithm

Note that for $\alpha = 1/\Delta$ we recover the standard COMP algorithm where an individual is classified as healthy if it appears in at least one displayed negative test which constitutes a sufficient condition in the noiseless case. We now state the first main result of this paper.

Theorem 2.1 (Noisy COMP)**.** *Let $p, q \geq 0$, $p+q<1$, $d \in (0,\infty)$, $\alpha \in (q, e^{-d}(1-p)+\left(1-e^{-d}\right)q)$. Suppose that $0<\theta<1$ and let*

$$m_{COMP} = m_{COMP}(n,\theta,p,q) = \min_{\alpha,d} \max\left\{b_1(\alpha,d), b_2(\beta,d)\right\} k\log(n/k)$$

$$\text{where} \qquad b_1(\alpha,d) = \frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\alpha\|q\right)}$$

$$\text{and} \qquad b_2(\alpha,d) = \frac{1}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\alpha\|e^{-d}(1-p)+\left(1-e^{-d}\right)q\right)}$$

If $m \geq (1+\varepsilon)m_{COMP}$ for some $\varepsilon>0$, noisy COMP will recover $\boldsymbol{\sigma}$ w.h.p. given test design $\boldsymbol{G}$ and test results $\hat{\boldsymbol{\sigma}}$.

The noisy variant of the DD algorithm of [5] was introduced in [48] and reads as follows:

1 Declare every individual that appears in $\alpha\Delta$ or more displayed negative tests as healthy and remove such individual from every assigned test.
2 Declare every yet unclassified individual who is now the only unclassified individual in $\beta\Delta$ or more displayed positive tests as infected.
3 Declare all remaining individuals as healthy.

Algorithm 2: The noisy DD algorithm [48]

This reduces to the noiseless version of DD introduced in [5] by taking $\alpha=\beta=1/\Delta$. We now state the second main result of the paper.

Theorem 2.2 (Noisy DD)**.** *Let $p, q \geq 0$, $p+q<1$, $d \in (0,\infty)$, $\alpha \in (q, e^{-d}(1-p)+\left(1-e^{-d}\right)q)$ and $\beta \in (0, e^{-d}(1-q))$ and define $w = e^{-d}p+(1-e^{-d})(1-q)$. Suppose that $0<\theta<1$ and let*

$$m_{DD} = m_{DD}(n,\theta,p,q) = \min_{\alpha,\beta,d} \max\left\{c_1(\alpha,d), c_2(\alpha,d), c_3(\beta,d), c_4(\alpha,\beta,d)\right\} k\log(n/k)$$

$$\text{where} \qquad c_1(\alpha,d) = \frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\alpha\|q\right)}$$

$$\text{and} \qquad c_2(\alpha,d) = \frac{1}{dD_{\mathrm{KL}}\left(\alpha\|1-w\right)}$$

$$\text{and} \qquad c_3(\beta,d) = \frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\beta\|(1-q)e^{-d}\right)}$$

$$\text{and} \qquad c_4(\alpha,\beta,d) = \max_{1-\alpha\leq z\leq 1}\left\{\frac{1}{1-\theta}\frac{1}{d\left(D_{\mathrm{KL}}\left(z\|w\right)+\mathbf{1}\left\{\beta>\frac{ze^{-d}p}{w}\right\}zD_{\mathrm{KL}}\left(\frac{\beta}{z}\|\frac{e^{-d}p}{w}\right)\right)}\right\}$$

If $m \geq (1+\varepsilon)m_{DD}$ for some $\varepsilon>0$, then noisy DD will recover $\boldsymbol{\sigma}$ w.h.p. given test design $\boldsymbol{G}$ and test results $\hat{\boldsymbol{\sigma}}$.

While the bounds appear cumbersome at first glance due to the numerous optimizations, the optimizations are of finite dimensions and for every specific value of p and q can be efficiently solved to arbitrary precision yielding explicit values for m_{COMP} and m_{DD}. For illustration purposes, we will calculate those bounds for several values of p, q and θ. Motivated by (1.1), we can describe the bounds in terms of rate, in a Shannon-theoretic sense. That is, we define the rate (bits learned per test) of an algorithm in this setting to be

$$R := \frac{\log\binom{n}{k}}{m\log 2} \sim \frac{k\log(n/k)}{m\log 2}.$$

(Recall that we take logarithms to base e throughout this paper). For example the fact that Theorems 2.1 and 2.2 show that noisy COMP and DD respectively can succeed w.h.p. with $m \geq (1+\epsilon)ck\log(n/k)$ tests for some c is equivalent to the fact that $R = 1/(c\log 2)$ is an achievable rate in a Shannon-theoretic sense.

We now give a counterpart to these two theorems by stating a universal converse for the $p-q$ channel below, improving on the universal counting bound from (1.1). The starting observation (see [7, Theorem 3.1]) is that no group testing algorithm can succeed w.h.p. with rate greater than C_{Chan} – the Shannon capacity of the corresponding noisy communication channel. Thus, we cannot hope to succeed w.h.p. with $m < (1-\epsilon)ck\log(n/k)$ tests where $c = 1/(C_{\mathrm{Chan}}\log 2)$. Hence as a direct consequence of the value of the channel capacity of the $p-q$ channel given in Lemma F.1 below, we deduce the following theorem.

Theorem 2.3. *Let $p, q \geq 0$, $p+q<1$ and $\epsilon > 0$, write $h(\cdot)$ for the binary entropy in nats (logarithms taken to base e) and $\phi = \phi(p,q) = (h(p)-h(q))/(1-p-q)$. If we define*

$$m_{COUNT} = \left(\frac{1}{D_{\mathrm{KL}}\left(q \| 1/(1+e^{\phi})\right)}\right) k\log(n/k),$$

then for $m \leq (1-\epsilon)m_{COUNT}$ no algorithm can recover $\boldsymbol{\sigma}$ w.h.p. for any matrix design.

Remark 2.4. *Note that the derivation of this result in Lemma 2.3 suggests a choice of density for the matrix:*

$$d = d^*_{\mathrm{ch}} = \log(1-p-q) - \log\left(\frac{1}{1+e^{\phi}} - q\right).$$

While this is not optimal, it may be regarded as a sensible heuristic that provides good rates for a range of p and q values.

2.2. **The combinatorics of the noisy group testing algorithms.** In the following, we outline the combinatorial structures that Algorithm 1 and 2 take advantage of.

2.2.1. *The noisy COMP algorithm.* To get started, let us shed light on the combinatorics of noisy COMP (Algorithm 1). For the *noiseless* case, the COMP algorithm classifies each individual that appears in at least one negative test as healthy and all other individuals as infected, since the participation in a negative test is a sufficient condition for the individual to be healthy.

For the noisy case, the situation is not as straightforward, since an infected individual might appear in *displayed* negative tests that were flipped when sent through the noisy channel. Thus, a single negative test is not definitive evidence that an individual is healthy. Yet, we can use the number of negative tests to tell the infected individuals apart from the healthy individuals.

Clearly, noisy COMP (Algorithm 1) using a threshold $\alpha\Delta$ succeeds if no healthy individual appears in less than $\alpha\Delta$ displayed negative tests and no infected individual appears in more than $\alpha\Delta$ displayed negative tests. To this end, we define

$$\boldsymbol{N}_x = |\{a \in \partial x : \hat{\boldsymbol{\sigma}}(a) = 0\}| \tag{2.1}$$

for the number of displayed negative tests that item x appears in. In terms of Figure 2, the algorithm determines the infection status by counting the number of tests of type I.

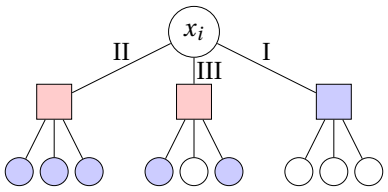

FIGURE 2. Rectangles represent tests (displayed positive in red, displayed negative in blue). Blue circles represent individuals that have been classified as healthy in the first step of DD (or by COMP). White circles represent individuals that are yet unclassified. On the one hand (Type II and III) this can happen before the first round of DD (or by COMP).On the other hand (Type I) it is the case before the algorithms start

2.2.2. *The noisy* DD *algorithm.* As in the prior section, let us first consider the *noiseless* DD algorithm. The first step is identical to COMP classifying all individuals that are contained in at least one negative test as healthy. In a second step, the algorithm checks each individual to see if they are contained in a positive test as the only yet unclassified individual and thus must be infected.

Again, the situation is more intricate when we add noise, since neither a single negative test gives us confidence that an individual is healthy nor does a positive test where the individual is the single yet unclassified individual inform us that this individual must be infected. Instead we count and compare the number of such tests. The first step of the noisy DD algorithm is identical to noisy COMP, but we are not required to identify all healthy individuals in the first step. Thus, after the first step, we are left with all infected individuals V_1 and a set of yet unclassified healthy individuals which we will denote by $V_{0,\mathrm{PD}}$. These are healthy individuals who did not appear in sufficiently many displayed negative tests to be declared healthy with confidence in the first step. In symbols, for some $\alpha \in (0,1)$

$$V_{0,\mathrm{PD}} = \{x \in V_0 : \boldsymbol{N}_x < \alpha\Delta\}$$

To tell V_1 and $V_{0,\mathrm{PD}}$ apart, we consider the number of displayed positive tests $\boldsymbol{P}_x$ where the individual x appears on its own after removing the definitely healthy individuals $V_0 \setminus V_{0,\mathrm{PD}}$ from the first step, i.e.

$$\boldsymbol{P}_x = \left|\left\{a \in \partial x : \hat{\boldsymbol{\sigma}}(a) = 1 \text{ and } \partial a \setminus \{x\} \subset V_0 \setminus V_{0,\mathrm{PD}}\right\}\right| \tag{2.2}$$

Referring to Figure 2, the second step of the algorithm is based on counting tests of type II. Tests of type III contain another yet unclassified individual from $V_{0,PD} \cup V_1$. The noisy DD algorithm takes advantage of the fact that it is less likely for an individual $x \in V_{0,\mathrm{PD}}$ to appear as the only yet unclassified individual in a displayed positive test than it is for an individual in $x \in V_1$. For $x \in V_{0,\mathrm{PD}}$ such a test would be truly negative and would have been flipped (which occurs with probability p) to display a positive test result. Conversely, an individual $x \in V_1$ renders any of its tests truly positive and thus the only requirement is that the test otherwise contains only definitely healthy individuals and is not flipped (which occurs with probability $1-q$). For this reason, we will see that the distribution of $\boldsymbol{P}_x$ differs between $x \in V_1$ and $x \in V_{0,\mathrm{PD}}$, and the difference $(1-q)-p > 0$ helps determine the size of this difference.

2.3. **Applying the results to standard channels.** With Theorem 2.1 and Theorem 2.2 we derived achievable rates for the generalized *p-q-model* (see Figure 1). prior research considered the Z channel where $p = 0$ and $q > 0$, the Reverse Z channel where $p > 0$ and $q = 0$ and the Binary Symmetric Channel with $p = q > 0$. These channels are the common models in coding theory [41], but are also often considered in medical applications [31, 32] concerned with taking sensitivity ($q > 0$), specificity ($p > 0$) or both ($p > 0$ and $q > 0$) into account. In the following section we will demonstrate how performance guarantees on these channels can directly be obtained from our main theorems.

2.3.1. *Recovery of the noiseless model.* First, we show the noiseless bounds can be simply recovered by setting $p = q = 0$. In the noiseless setting, it is optimal to set both α and β to $1/\Delta$. To see why, observe that in the absence of noise a single negative test is sufficient evidence that an individual is healthy. Conversely, a single positive test where the individual only appears with definitely healthy individuals implies that particular individual must surely be infected. As shown in [13] the optimal parameter choice for d in the constant-column design in the noiseless setting is $\log(2)$. Applying these values to Theorem 2.1 we recover the noiseless bound for COMP.

Corollary 2.5 (COMP in the noiseless setting)**.** *Let $p = q = 0$, $0 < \theta < 1$ and $\varepsilon > 0$. Further, let*

$$m_{\texttt{COMP},noiseless} = \frac{1}{(1-\theta)\log^2 2} k\log(n/k).$$

If $m > (1+\varepsilon)\, m_{\texttt{COMP},noiseless}$, COMP will recover $\boldsymbol{\sigma}$ w.h.p.given $\boldsymbol{G}, \hat{\boldsymbol{\sigma}}$.

We also recover the noiseless bounds for the DD algorithm as stated in [26].

Corollary 2.6 (DD in the noiseless setting)**.** *Let $p = q = 0, 0 < \theta < 1$ and $\varepsilon > 0$. Further, let*

$$m_{\texttt{DD},noiseless} = \max\left\{1, \frac{\theta}{1-\theta}\right\} \frac{1}{\log^2 2} k\log(n/k).$$

If $m > (1+\varepsilon)\, m_{\texttt{DD},noiseless}$, DD will recover $\boldsymbol{\sigma}$ w.h.p.given $\boldsymbol{G}, \hat{\boldsymbol{\sigma}}$.

2.3.2. *The Z channel.* In the Z channel, we have $p = 0$ and $q > 0$, i.e. no truly negative test displays a positive test result. Thus, we set $\beta = 1/\Delta$ and remain agnostic about α and d. The bounds for COMP and DD thus read.

Corollary 2.7 (Noisy COMP for the Z channel)**.** *Let $p = 0, 0 < q < 1, 0 < \theta < 1$ and $\varepsilon > 0$. Further, let*

$$m_{\texttt{COMP},Z} = \min_{\alpha,d} \max\{b_1(\alpha,d), b_2(\alpha,d)\}\, k\log(n/k)$$
$$\text{with} \quad b_1(\alpha,d) = \frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\alpha \| q\right)} \quad \text{and} \quad b_2(\alpha,d) = \frac{1}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\alpha \| e^{-d} + \left(1-e^{-d}\right) q\right)}.$$

If $m > (1+\varepsilon)\, m_{\texttt{COMP},Z}$, noisy COMP will recover $\boldsymbol{\sigma}$ w.h.p.given $\boldsymbol{G}, \hat{\boldsymbol{\sigma}}$.

Corollary 2.8 (Noisy DD for the Z channel)**.** *Let $p = 0, 0 < q < 1, 0 < \theta < 1$ and $\varepsilon > 0$. Further, let*

$$m_{\texttt{DD},Z} = \min_{\alpha,d} \max\{c_1(\alpha,d), c_2(\alpha,d), c_3(d)\}\, k\log(n/k)$$
$$\text{with} \quad c_1(\alpha,d) = \frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\alpha \| q\right)} \quad \text{and} \quad c_2(\alpha,d) = \frac{1}{dD_{\mathrm{KL}}\left(\alpha \| e^{-d} + \left(1-e^{-d}\right) q\right)}$$
$$\text{and} \quad c_3(d) = \frac{\theta}{1-\theta}\frac{1}{-d\log\left(1-e^{-d}(1-q)\right)}.$$

If $m > (1+\varepsilon)\, m_{\texttt{DD},Z}$, noisy DD will recover $\boldsymbol{\sigma}$ w.h.p.given $\boldsymbol{G}, \hat{\boldsymbol{\sigma}}$.

Proof. The bounds c_1 and c_2 follow directly from Theorem 2.2 by setting $p = 0$. For c_3 we use the fact that $D_{\mathrm{KL}}\left(1/\Delta \| e^{-d}(1-q)\right) = -\log\left(1-e^{-d}(1-q)\right) + o(1)$. An immediate consequence of $p = 0$ is $c_4 = 0$. □

An illustration of the bounds from Corollary 2.7 and 2.8 for sample values of q is shown in Figure 5.

2.3.3. *Reverse Z channel.* In the reverse Z channel, we have $q = 0$ and $p > 0$, i.e. no truly positive test displays a negative test result. thus, we set $\alpha = 1/\Delta$ and remain agnostic about β and d. The bounds for the noisy COMP and DD thus read as follows.

Corollary 2.9 (Noisy COMP for the Reverse Z channel)**.** *Let* $0 < p < 1, q = 0, 0 < \theta < 1$ *and* $\varepsilon > 0$*. Further, let*

$$m_{\mathtt{COMP},rev\,Z} = \frac{1}{1-\theta} \min_d \left\{ \frac{1}{-d\log\left(1-e^{-d}(1-p)\right)} \right\} k\log(n/k).$$

If $m > (1+\varepsilon)\, m_{\mathtt{COMP},rev\,Z}$*, noisy* COMP *will recover* $\boldsymbol{\sigma}$ *w.h.p.given* $\boldsymbol{G}, \hat{\boldsymbol{\sigma}}$*.*

Proof. The corollary follows from Theorem 2.1 and the fact that $D_{\mathrm{KL}}(1/\Delta \| 0)$ diverges and $D_{\mathrm{KL}}\left(1/\Delta \| e^{-d}(1-p)\right) = -\log\left(1-e^{-d}(1-p)\right)$. □

Note that the optimal d arising from Corollary 2.9 cannot be expressed in terms of standard functions.

Corollary 2.10 (Noisy DD in the Reverse Z channel)**.** *Let* $0 < p < 1, q = 0, 0 < \theta < 1$ *and* $\varepsilon > 0$*. Further, let*

$$m_{\mathtt{DD},rev\,Z} = \min_{\beta,d} \max\left\{c_2(d), c_3(\beta,d), c_4(\beta,d)\right\} k\log(n/k)$$

$$\text{with } c_2(d) = \frac{1}{-d\log\left(1-e^{-d}(1-p)\right)} \quad \text{and} \quad c_3(\beta,d) = \frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\beta \| e^{-d}\right)}$$

$$\text{and } c_4(\beta,d) = \frac{1}{1-\theta}\frac{1}{d\left(-\log\left(1-e^{-d}(1-p)\right) + D_{\mathrm{KL}}\left(\beta \| \frac{e^{-d}p}{e^{-d}p+(1-e^{-d})}\right)\right)}$$

If $m > (1+\varepsilon)\, m_{\mathtt{DD},rev\,Z}$*, noisy* DD *will recover* $\boldsymbol{\sigma}$ *w.h.p.given* $\boldsymbol{G}, \hat{\boldsymbol{\sigma}}$*.*

Proof. The bounds $c_1 = 0, c_2, c_3$ follow from Theorem 2.2 and the same manipulations as above. For c_4, note that z needs to take the value 1 since $1-\alpha = 1-1/\Delta$, whence the simplification follows immediately. □

An illustration of the bounds of Corollary 2.9 and 2.10 for sample values of p is shown in Figure 6.

2.3.4. *Binary Symmetric Channel.* In the Binary Symmetric Channel (BSC), we set $p = q > 0$. Even though information-theoretic arguments would suggest setting $d = \log 2$, we formulate the expression below with general d. We also keep the threshold parameters α and β. The bounds for the noisy DD and COMP only simplify slightly.

Corollary 2.11 (Noisy COMP in the Binary Symmetric Channel)**.** *Let* $0 < p = q < 1/2, 0 < \theta < 1$ *and* $\varepsilon > 0$*. Further, let*

$$m_{\mathtt{COMP},BSC} = \min_{\alpha,d} \max\{b_1(\alpha,d), b_2(\alpha,d)\}\, k\log(n/k)$$

$$\text{with} \quad b_1(\alpha,d) = \frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\alpha \| p\right)} \quad \text{and} \quad b_2(\alpha,d) = \frac{1}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\alpha \| e^{-d}+p-2e^{-d}p\right)}.$$

If $m > (1+\varepsilon)\, m_{\mathtt{COMP},BSC}$*, noisy* COMP *will recover* $\boldsymbol{\sigma}$ *w.h.p.given* $\boldsymbol{G}, \hat{\boldsymbol{\sigma}}$*.*

Corollary 2.12 (Noisy DD in the Binary Symmetric Channel)**.** *Let* $0 < p = q < 1/2, 0 < \theta < 1$ *and* $\varepsilon > 0$ *and define* $\nu = 1 - e^{-d} - p + 2e^{-d}p$*. Further, let*

$$m_{\mathtt{DD},BSC} = \min_{\alpha,\beta,d} \max\left\{c_1(\alpha,d), c_2(\alpha,d), c_3(\beta,d), c_4(\alpha,\beta,d)\right\} k\log(n/k)$$

$$\text{with} \quad c_1(\alpha,d) = \frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\alpha \| p\right)} \quad \text{and} \quad c_2(\alpha,d) = \frac{1}{dD_{\mathrm{KL}}\left(\alpha \| e^{-d}+p-2e^{-d}p\right)}$$

$$\text{and} \quad c_3(\beta,d) = \frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\beta \| (1-p)e^{-d}\right)}$$

$$\text{and} \quad c_4(\alpha,\beta,d) = \max_{1-\alpha \le z \le 1} \left\{ \frac{1}{1-\theta}\frac{1}{d\left(D_{\mathrm{KL}}(z\|\nu) + \mathbf{1}\left\{\beta > \frac{ze^{-d}p}{\nu}\right\} zD_{\mathrm{KL}}\left(\frac{\beta}{z} \| \frac{e^{-d}p}{\nu}\right)\right)} \right\}.$$

If $m > (1+\varepsilon)\, m_{\mathtt{DD},BSC}$*, noisy* DD *will recover* $\boldsymbol{\sigma}$ *w.h.p.given* $\boldsymbol{G}, \hat{\boldsymbol{\sigma}}$*.*

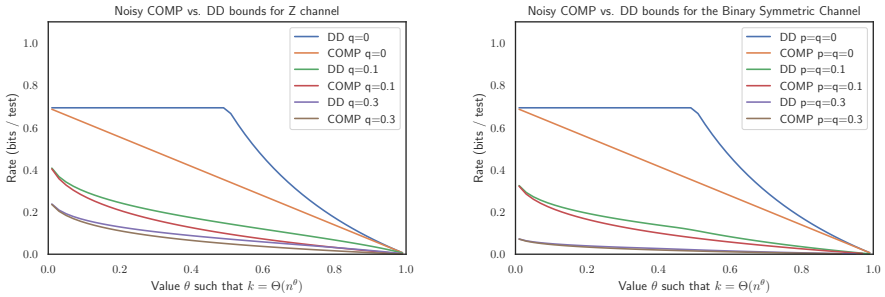

FIGURE 3. Comparison of the bound for noisy DD and noisy COMP in the Z-channel and the Binary Symmetric Channel for different noise level.

An illustration of the bounds of Corollary 2.11 and 2.12 is shown in Figure 7.

2.4. **Comparison of noisy COMP and DD.** An obvious next question is to find conditions under which the noisy DD algorithm outperforms noisy COMP. For the noiseless setting, it can be easily shown that DD provably outperforms COMP for all $\theta \in (0,1)$. For the noisy case, matters are slightly more complicated.

Recall that noisy COMP classifies all individuals appearing in less than $\alpha\Delta$ displayed negative tests as infected while noisy DD additionally requires such individuals to appear in more than $\beta\Delta$ displayed positive tests as the only yet unclassified individual. Thus, it might well be that an infected individual is classified correctly by noisy COMP, while it is missed by the noisy DD algorithm.

That being said, our simulations indicate that noisy DD generally outperforms noisy COMP, but for the reason mentioned above we can only prove that noisy DD outperforms noisy COMP for the reverse Z channel while remaining agnostic about the Z channel and the Binary Symmetric Channel, as the next proposition evinces.

Proposition 2.13. *For all $p, q \geq 0$ with $p+q<1$ there exists a $d^* \in (0,\infty)$ such that $m_{COMP} \geq m_{DD}$ as long as $e^{-d^*} p \geq q$.*

In terms of the common noise channels Proposition 2.13 gives the following corollary.

Corollary 2.14. *In the reverse Z channel, $m_{COMP} \geq m_{DD}$.*

Our simulations suggest that this superior performance of noisy DD holds as well for the Z channel and Binary Symmetric Channel. Please refer to Figure 3 for an illustration.

2.5. **Relation to Bernoulli testing.** [26] derived sufficient bounds for noisy group testing and a Bernoulli test design where each individual joins every test with some fixed probability. Thus, the variable degrees fluctuate and we end up with some individuals assigned only to few tests. In contrast, we work under a model in this paper where each individual joins an equal number of tests Δ chosen uniformly at random without replacement. For the noiseless case, it is by now clear that the constant-column design better facilitates inference than the Bernoulli test design [13, 26]. We find that the same holds true for the noisy variant of the COMP algorithm. Let us denote by $m_{\mathrm{COMP}}^{\mathrm{Ber}}$ the number of tests required for the noisy COMP to succeed under a Bernoulli test design.

Proposition 2.15. *For all $p+q<1$, we have*

$$m_{COMP}^{Ber} \geq m_{COMP}$$

We see the same effect for the noisy variant of the DD algorithm for all simulations, but for technical reasons only prove it for the Z channel.

Proposition 2.16. *For the Z channel where $p = 0$ and $0 < q < 1$, we have*

$$m_{DD}^{Ber} > m_{DD}$$

For an illustration on the magnitude of the difference, we refer to Figure 4 and Figure 8.

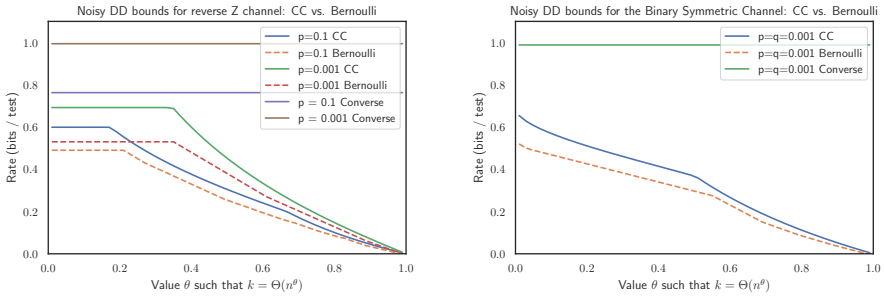

FIGURE 4. Comparison of DD bounds under a Bernoulli test design ([48]) and constant column test design (present paper) for the reverse Z and Binary Symmetric Channel

APPENDIX

The core of the technical sections is the proof of Theorems 2.1 and Theorem 2.2. Some groundwork with standard concentration bounds and group testing properties can be found in Section A. We continue with the proof of Theorems 2.1 and 2.2 in Sections B and C, respectively. The structure of the proofs follows a similar logic. First, we derive the distributions for the number of displayed positive and negative tests for infected and healthy individuals. Second, we threshold these distributions using sharp Chernoff concentration bounds to deduce the bounds stated in Theorem 2.1 and Theorem 2.2. Thereafter, we proceed to the proof of Proposition 2.13 in Section D, while the proofs of Propositions 2.15 and 2.16 follow in Section E. We conclude with the proof of the converse result from Theorem 2.3 in Section F.

APPENDIX A. GROUNDWORK

For starters, let us recall the Chernoff bound for binomial and hypergeometric distributions.

Lemma A.1 (Chernoff bound for the binomial distribution [25])**.** *Let $p < q < r \in (0,1)$ and $\boldsymbol{X} \sim \mathrm{Bin}(n,q)$ be a binomially distributed random variable. Then*

$$\mathbb{P}\left(\boldsymbol{X} \le \lceil pn \rceil\right) = \exp\left(-\left(1+n^{-\Omega(1)}\right) nD_{\mathrm{KL}}\left(p \| q\right)\right)$$
$$\mathbb{P}\left(\boldsymbol{X} \ge \lceil rn \rceil\right) = \exp\left(-\left(1+n^{-\Omega(1)}\right) nD_{\mathrm{KL}}\left(r \| q\right)\right)$$

Lemma A.2 (Chernoff bound for the hypergeometric distribution [23])**.** *Let $p < q < r \in (0,1)$ and $\boldsymbol{Y} \sim H(N,Q,n)$ be a hypergeometrically distributed random variable. Further, let $q = Q/N$. Then*

$$\mathbb{P}\left(H(N,Q,n) \le \lceil pn \rceil\right) = \exp\left(-\left(1+n^{-\Omega(1)}\right) nD_{\mathrm{KL}}\left(p \| q\right)\right)$$
$$\mathbb{P}\left(H(N,Q,n) \ge \lceil rn \rceil\right) = \exp\left(-\left(1+n^{-\Omega(1)}\right) nD_{\mathrm{KL}}\left(r \| q\right)\right)$$

The next lemma provides that the test degrees, as defined in (1.3) above, are tightly concentrated. Recall from (1.2) that the number of tests $m = ck\log(n/k)$ and each item appears in $\Delta = cd\log(n/k)$ tests.

Lemma A.3. *With probability $1-o(n^{-2})$ we have*

$$dn/k - \sqrt{dn/k}\log n \le \Gamma_{\min} \le \Gamma_{\max} \le dn/k + \sqrt{dn/k}\log n$$

Proof. The probability that an individual x is assigned to test a is given by

$$\mathbb{P}(x \in \partial a) = 1 - \mathbb{P}(x \notin \partial a) = 1 - \binom{m-1}{\Delta}\binom{m}{\Delta}^{-1} = \Delta/m = d/k \tag{A.1}$$

Since each individual is assigned to tests independently, the total number of individuals in a given test follows the binomial distribution $\mathrm{Bin}(n, d/k)$. The assertion now follows from the Chernoff bound for binomial distributions (Lemma A.1). □

Next, we show that the number of truly negative tests $\boldsymbol{m}_0$ (and thus the number of truly positive tests $\boldsymbol{m}_1$) are tightly concentrated.

Lemma A.4. *With probability $1-o(n^{-2})$ we have $\boldsymbol{m}_0 = e^{-d}m + O(\sqrt{m}\log^3 n)$.*

Proof. Recall from (A.1) that

$$\mathbb{P}(x \in \partial a) = d/k$$

Since infected individuals are assigned to tests mutually independently, we find for a test a that

$$\mathbb{P}(V_1 \cap \partial a = \emptyset) = \mathbb{P}(\mathrm{Bin}(k, d/k) = 0) = (1-d/k)^k = \left(1+n^{-\Omega(1)}\right)e^{-d}.$$

Consequently, $\mathbb{E}[\boldsymbol{m}_0] = \left(1 + n^{-\Omega(1)}\right) e^{-d} m$. Finally, changing the set of tests for a specific infected individual shifts the total number of negative tests by at most Δ. Therefore, the Azuma-Hoeffding inequality yields

$$\mathbb{P}\left(|\boldsymbol{m}_0 - \mathbb{E}[\boldsymbol{m}_0]| \geq t\right) \leq 2\exp\left(-\frac{t^2}{4k\Delta^2}\right).$$

The lemma follows from setting $t = \sqrt{m}\log^3 n$. □

With the concentration of $\boldsymbol{m}_0$ and $\boldsymbol{m}_1$ at hand, we readily obtain estimates for $\boldsymbol{m}_0^f, \boldsymbol{m}_0^u, \boldsymbol{m}_1^f$ and $\boldsymbol{m}_1^u$.

Corollary A.5. *With probability* $1 - o(n^{-2})$ *we have*

(i) $\boldsymbol{m}_0^f = e^{-d} p m + O\left(\sqrt{m}\log^4 n\right)$
(ii) $\boldsymbol{m}_0^u = e^{-d}(1-p) m + O\left(\sqrt{m}\log^4 n\right)$
(iii) $\boldsymbol{m}_1^f = (1 - e^{-d}) q m + O\left(\sqrt{m}\log^4 n\right)$
(iv) $\boldsymbol{m}_1^u = (1 - e^{-d})(1-q) m + O\left(\sqrt{m}\log^4 n\right)$

Proof. Since each test is flipped with probability p and q independently, the claims follow from Lemma A.4 and the Chernoff bound for the binomial distribution (Lemma A.1). □

In the following, let $\mathscr{E}$ be the event that the bounds from Lemma A.4 and A.5 hold.

Appendix B. Proof of COMP bound, Theorem 2.1

Recall from (2.1) that we write $\boldsymbol{N}_x$ for the number of displayed negative tests that item x appears in (as illustrated by the right branch of Fig. 2). The proof of Theorem 2.1 is based on two pillars. First, Lemmas B.1 and B.2 provide the distribution of $\boldsymbol{N}_x$ for healthy and infected individuals, respectively. We will see that these distributions differ according to the infection status of the individual. Second, we will derive a suitable threshold $\alpha\Delta$ via Lemma B.3 and B.4 to tell healthy and infected individuals apart w.h.p. We start by analysing individuals in the infected set V_1. Throughout the section, we assume $\alpha \in (q, e^{-d}(1-p) + \left(1 - e^{-d}\right) q)$.

Lemma B.1. *Given* $x \in V_1$*, its number of displayed negative tests* $\boldsymbol{N}_x$ *is distributed as* $\mathrm{Bin}(\Delta, q)$.

Proof. Any test containing an infected individual is truly positive because of the presence of the infected individual. Since an infected individual is assigned to Δ different tests and each such test is flipped with probability q independently, the lemma follows immediately. □

Next, we consider the distribution for healthy individuals. Recall that $\mathscr{E}$ denotes the event that the bounds from Lemma A.4 and Corollary A.5 hold.

Lemma B.2. *Given* $x \in V_0$ *and* $\mathscr{E}$*,* $\boldsymbol{N}_x$ *is distributed as* $H\left(m, m\left(e^{-d}(1-p) + \left(1 - e^{-d}\right) q + n^{-\Omega(1)}\right), \Delta\right)$.

Proof. Since x is healthy, the outcome of all the tests remains the same if it is removed from consideration (if we perform group testing with $n-1$ items and the corresponding reduced matrix).

Thus, given $\mathscr{E}$, we find that with x removed the $\boldsymbol{m}_0^f, \boldsymbol{m}_0^u, \boldsymbol{m}_1^f, \boldsymbol{m}_1^u$ still satisfy the bounds from Corollary A.5. As a result the number of displayed negative tests (which consist of unflipped truly negative tests and flipped truly positive tests) is given by

$$\boldsymbol{m}_0^u + \boldsymbol{m}_1^f = \left(e^{-d}(1-p) + (1 - e^{-d}) q\right) m + O\left(\sqrt{m}\log^4 n\right) \tag{B.1}$$

Now, adding x back into consideration: $x \in V_0$ chooses Δ tests without replacement independently of this. Hence the number of displayed negative tests it appears in $\boldsymbol{N}_x$ is distributed as $H(m, \boldsymbol{m}_0^u + \boldsymbol{m}_1^f, \Delta)$ and the lemma follows. □

Moving to the second pillar of the proof, we need to demonstrate that no infected individual is assigned to more than $\alpha\Delta$ displayed negative tests as shown by the following lemma.

Lemma B.3. *If $c > (1+\eta)\frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}(\alpha\|q)}$ for some small $\eta > 0$, $\boldsymbol{N}_x < \alpha\Delta$ for all $x \in V_1$ w.h.p.*

Proof. We have to ensure that $\mathbb{P}(\exists x \in V_1 : \boldsymbol{N}_x \geq \alpha\Delta) = o(1)$. By Lemma B.1 and the union bound, we thus need to have

$$o(1) = k \cdot \mathbb{P}\left(\boldsymbol{N}_x \geq \alpha\Delta : x \in V_1\right) = k \cdot \mathbb{P}\left(\mathrm{Bin}(\Delta, q) \geq \alpha\Delta\right) = k \cdot \exp\left(-\left(1+\Delta^{-\Omega(1)}\right)\Delta D_{\mathrm{KL}}\left(\alpha\|q\right)\right),$$

by the Chernoff bound for the binomial distribution (Lemma A.1). Since $k \sim n^\theta$ and $\Delta = cd(1-\theta)\log n$ this implies

$$\theta - cd(1-\theta)D_{\mathrm{KL}}\left(\alpha\|q\right) < 0$$

The lemma follows from rearranging terms. □

We proceed to show that no healthy individual is assigned to less than $\alpha\Delta$ displayed negative tests.

Lemma B.4. *If $c > (1+\eta)\frac{1}{1-\theta}\frac{1}{dD_{\mathrm{KL}}(\alpha\|e^{-d}(1-p)+(1-e^{-d})q)}$ for some small $\eta > 0$, $\boldsymbol{N}_x > \alpha\Delta$ for all $x \in V_0$ w.h.p.*

Proof. We need to ensure that $\mathbb{P}(\exists x \in V_0 : \boldsymbol{N}_x < \alpha\Delta) = o(1)$. Since $\mathscr{E}$ occurs w.h.p. by Lemma A.4 and Corollary A.5, we need to have by Lemma B.2 and the union bound that

$$(n-k) \cdot \mathbb{P}\left(\boldsymbol{N}_x \leq \alpha\Delta | x \in V_0, \mathscr{E}\right) \leq n \cdot \mathbb{P}\left(H\left(m, m\left(e^{-d}(1-p) + \left(1-e^{-d}\right)q + n^{-\Omega(1)}\right), \Delta\right) \leq \alpha\Delta\right) = o(1). \tag{B.2}$$

Together with the Chernoff bound for the hypergeometric distribution (Lemma A.2) this implies

$$1 - cd(1-\theta)D_{\mathrm{KL}}\left(\alpha\|(1-pe^{-d}+(1-e^{-d})q\right) < 0$$

in a similar way to the proof of Lemma B.3. The lemma follows from rearranging terms. □

Proof of Theorem 2.1. The theorem is now an immediate consequence of Lemma B.3 and B.4 which guarantee that w.h.p. classifying individuals according to the threshold $\alpha\Delta$ for negative displayed tests recovers $\boldsymbol{\sigma}$, and the fact that the choice of α and d is at our disposal. □

Appendix C. Proof of DD bound, Theorem 2.2

The proof of Theorem 2.2 follows a similar two-step approach as the proof of Theorem 2.1 by first finding the distribution of $\boldsymbol{P}_x$ (the number of displayed positive tests where individual x appears on its own after removing the definitely healthy individuals $V_0 \setminus V_{0,\mathrm{PD}}$, illustrated by the left branch of Fig. 2). We then threshold the distributions for healthy and infected individuals. To get started, we revise the second bound from Theorem 2.1 to allow $kn^{-\Omega(1)}$ healthy individuals to not be classified yet after the first step of DD. Throughout the section, we assume $\alpha \in (q, e^{-d}(1-p)+\left(1-e^{-d}\right)q)$ and $\beta \in (0, e^{-d}(1-q))$.

Lemma C.1. *If*

$$c > (1+\eta)\frac{1}{dD_{\mathrm{KL}}\left(\alpha\|e^{-d}(1-p)+\left(1-e^{-d}\right)q\right)}$$

for some small $\eta > 0$, we have $\left|V_{0,PD}\right| = kn^{-\Omega(1)}$ w.h.p.

Proof. The lemma follows immediately by replacing the r.h.s. of (B.2) with $kn^{-\delta}$ for some small $\delta = \delta(\eta)$, rearranging terms and applying Markov's inequality. □

For the next lemmas, we need an auxiliary notation denoting the number of tests $\boldsymbol{m}_{0,\mathrm{nd}}$ that only contain individuals from $V_0 \setminus V_{0,\mathrm{PD}}$. In symbols,

$$\boldsymbol{m}_{0,\mathrm{nd}} = \left|\left\{a \in F : \partial a \subset V_0 \setminus V_{0,\mathrm{PD}}\right\}\right|.$$

Lemma C.2. *If*

$$c > (1+\eta)\frac{1}{dD_{\mathrm{KL}}\left(\alpha \| e^{-d}(1-p)+\left(1-e^{-d}\right)q\right)}$$

for some small $\eta > 0$, we have $\boldsymbol{m}_{0,nd} = \left(1-n^{-\Omega(1)}\right)e^{-d}m$ with probability $1-o(n^{-2})$.

Proof. As in the proof of Lemma B.2 above, we consider the graph in two rounds: first we consider the tests containing infected individuals. Since each healthy individual $x \in V_0$ does not impact the number of positive and negative tests, we know by Lemma A.4 that with probability $1-o(n^{-2})$ we have $\boldsymbol{m}_0 = e^{-d}m + O\left(\sqrt{m}\log^4 n\right)$ after the first round.

Now consider some particular negative test a. The probability that a healthy individual x is assigned to this test is d/k by (A.1). By Lemma C.1, we know that $\left|V_{0,\mathrm{PD}}\right| = kn^{-\Omega(1)}$. Since each such individual is assigned to tests mutually independently, we find for the truly negative test a that

$$\mathbb{P}\left(V_{0,\mathrm{PD}} \cap \partial a = \emptyset\right) = \mathbb{P}\left(\mathrm{Bin}\left(\left|V_{0,\mathrm{PD}}\right|, d/k\right) = 0\right) = (1-d/k)^{kn^{-\Omega(1)}} = 1-n^{-\Omega(1)}$$

We therefore have $\mathbb{E}\left[\boldsymbol{m}_{0,\mathrm{nd}}\right] = \left(1-n^{-\Omega(1)}\right)e^{-d}m$. Finally, changing the set of tests for a specific individual $x \in V_1 \cup V_{0,\mathrm{PD}}$ shifts $\boldsymbol{m}_{0,\mathrm{nd}}$ by at most Δ. The lemma follows by a similar application of the Azuma-Hoeffding inequality as used in the proof of Lemma A.4. □

Let $\mathscr{F}$ be the event that $\boldsymbol{m}_{0,\mathrm{nd}} = \left(1-n^{-\Omega(1)}\right)e^{-d}m$ indeed. By Lemma C.2, $\mathbb{P}(\mathscr{F}) = 1-o(n^{-2})$ if

$$c > (1+\eta)\frac{1}{dD_{\mathrm{KL}}\left(\alpha \| e^{-d}(1-p)+\left(1-e^{-d}\right)q\right)}$$

for some small $\eta > 0$. With Lemma C.2 at hand, we are in a position to describe the distribution of $\boldsymbol{P}_x$ for healthy and infected individuals. Let us start with infected individuals.

Lemma C.3. *Given an infected individual $x \in V_1$ and assuming $\mathscr{F}$ holds, $\boldsymbol{P}_x$ is distributed as $H\left(m, m\left(e^{-d}(1-q)+n^{-\Omega(1)}\right), \Delta\right)$.*

Proof. Consider an infected individual $x \in V_1$. As before, if we remove x from tests, this will change $\boldsymbol{m}_{0,\mathrm{nd}}$ by at most Δ.

Thus, by Lemma C.2 the number of tests that x is assigned to that contain neither infected individuals nor individuals from $V_{0,\mathrm{PD}}$ is distributed as $H\left(m, m\left(e^{-d}+n^{-\Omega(1)}\right), \Delta\right)$ given $\mathscr{F}$. Since each test featuring x will truly be positive and will be displayed positive with probability $1-q$ independently, the lemma follows immediately. □

To describe the distribution of $\boldsymbol{P}_x$ for healthy individuals, let us introduce the random variable $\boldsymbol{P}_x(P)$, which is $\boldsymbol{P}_x$ conditioned on the individual appearing in P displayed positive tests, as follows:

$$\mathbb{P}\left(\boldsymbol{P}_x(P) = t\right) = \mathbb{P}\left(\boldsymbol{P}_x = t | \boldsymbol{N}_x = \Delta - P\right)$$

Then, we find for healthy individuals the following conditional distribution.

Lemma C.4. *Given $x \in V_0$ and $\mathscr{F}$, $\boldsymbol{P}_x(P)$ is distributed as*

$$H\left(m\left(e^{-d}p+(1-e^{-d})(1-q)+n^{-\Omega(1)}\right), m\left(e^{-d}p+n^{-\Omega(1)}\right), P\right).$$

Proof. We proceed with the same exposition as in the proof of Lemma C.3. Since individual $x \in V_0$ is assigned to exactly P displayed positive, $\boldsymbol{P}_x(P)$ is distributed as $H\left(\boldsymbol{m}_0^f + \boldsymbol{m}_1^u, \boldsymbol{m}_{0,\mathrm{nd}}, P\right)$. The lemma follows from Corollary A.5 and Lemma C.2. □

Having derived the distributions for $\boldsymbol{P}_x$ for $x \in V_1$ and $\boldsymbol{P}_x(P)$ for $x \in V_0$ we can now determine a threshold $\beta\Delta$ of displayed positive tests where the individual appears only with individuals from the set $V_0 \setminus V_{0,\mathrm{PD}}$ such that we can tell V_1 and $V_{0,\mathrm{PD}}$ apart and thus recover $\boldsymbol{\sigma}$. Let us start with infected individuals.

Lemma C.5. *As long as*

$$c > (1+\eta)\max\left\{\frac{1}{dD_{\mathrm{KL}}\left(\alpha\|e^{-d}(1-p)+\left(1-e^{-d}\right)q\right)}, \frac{\theta}{1-\theta}\frac{1}{dD_{\mathrm{KL}}\left(\beta\|(1-q)e^{-d}\right)}\right\}$$

for some small $\eta > 0$, we have $\boldsymbol{P}_x > \beta\Delta$ for all $x \in V_1$ w.h.p.

Proof. We need to ensure that $\mathbb{P}(\exists x \in V_1 : \boldsymbol{P}_x < \beta\Delta) = o(1)$. For the bound on c from the lemma, we know that $\mathscr{F}$ occurs w.h.p. by Lemma C.2. In combination with Lemma C.3 and the union bound we need to ensure

$$k \cdot \mathbb{P}\left(\boldsymbol{P}_x \le \beta\Delta | x \in V_1, \mathscr{F}\right) = k \cdot \mathbb{P}\left(H\left(m, m\left(e^{-d}(1-q)+n^{-\Omega(1)}\right), \Delta\right) \le \beta\Delta\right) = o(1) \tag{C.1}$$

Using the Chernoff bound for the hypergeometric distribution (Lemma A.2), (C.1) holds if

$$\theta - cd(1-\theta)D_{\mathrm{KL}}\left(\beta\|(1-q)e^{-d}\right) < 0 \tag{C.2}$$

The lemma follows from rearranging terms in (C.2).

□

We proceed with the set of individuals $V_{0,\mathrm{PD}}$.

Lemma C.6. *As long as*

$$c > (1+\eta)\max\left\{\frac{1}{dD_{\mathrm{KL}}\left(\alpha\|e^{-d}(1-p)+\left(1-e^{-d}\right)q\right)}, \max_{1-\alpha\le z\le 1}\left\{\frac{1}{1-\theta}\frac{1}{d\left(D_{\mathrm{KL}}\left(z\|e^{-d}p+(1-e^{-d})(1-q)\right)+zD_{\mathrm{KL}}\left(\frac{\beta}{z}\|\frac{e^{-d}p}{e^{-d}p+(1-e^{-d})(1-q)}\right)\right)}\right\}\right\}$$

for some small $\eta > 0$, we have $\boldsymbol{P}_x < \beta\Delta$ for all $x \in V_{0,PD}$ w.h.p.

Proof. We need to ensure that $\mathbb{P}(\exists x \in V_{0,PD} : \boldsymbol{P}_x > \beta\Delta) = o(1)$. For the bound on c from the lemma, we know that $\mathscr{F}$ occurs w.h.p. by Lemma C.2. Moreover, $\mathscr{E}$ occurs w.h.p. by Lemma A.4 and Corollary A.5. We write $w = e^{-d}p + \left(1-e^{-d}(1-q)\right)$ for brevity. Combining this fact with Lemma B.2 and C.4 we need to ensure

$$\begin{aligned}(n-k)\sum_{P=(1-\alpha)\Delta}^{\Delta}\mathbb{P}\left(\boldsymbol{N}_x = \Delta - P | x \in V_0, \mathscr{E}\right)\mathbb{P}\left(\boldsymbol{P}_x(P) \ge \beta\Delta | x \in V_0, \mathscr{F}\right) \\ = \left(1-n^{-\Omega(1)}\right)n\sum_{P=(1-\alpha)\Delta}^{\Delta}\mathbb{P}\left(H\left(m, m\left(w+n^{-\Omega(1)}\right), \Delta\right) = P\right)\end{aligned} \tag{C.3}$$

$$\mathbb{P}\left(H\left(m\left(w+n^{-\Omega(1)}\right), m\left(e^{-d}p+n^{-\Omega(1)}\right), P\right) \ge \beta\Delta\right) = o(1) \tag{C.4}$$

By the Chernoff bound for the hypergeometric distribution (Lemma A.2) and setting $z = P/\Delta$, we reformulate the left-hand-side of (C.4) to

$$\begin{aligned}n\sum_{P=(1-\alpha)\Delta}^{\Delta}\exp\left(-(1+o(1))\Delta\left(D_{\mathrm{KL}}(z\|w) + \mathbf{1}\left\{\beta > \frac{ze^{-d}p}{w}\right\}zD_{\mathrm{KL}}\left(\frac{\beta}{z}\|\frac{e^{-d}p}{w}\right)\right)\right) \\ = \left(1+n^{-\Omega(1)}\right)n\max_{1-\alpha\le z\le 1}\left\{\exp\left(-(1+o(1))\Delta\left(D_{\mathrm{KL}}(z\|w) + \mathbf{1}\left\{\beta > \frac{ze^{-d}p}{w}\right\}zD_{\mathrm{KL}}\left(\frac{\beta}{z}\|\frac{e^{-d}p}{w}\right)\right)\right)\right\}\end{aligned}$$

where the second equality follows since the sum consists of $\Theta(\Delta) = \Theta(\log n)$ many summands. Since $\mathbb{P}(\mathscr{F}) = 1 - n^{-\Omega(1)}$ for our choice of c by Lemma C.2 rearranging terms readily yields that the expression in (C.3) is indeed of order $o(1)$. □

Proof of Theorem 2.2. The theorem is now immediate from Lemma B.3, C.1, C.5 and C.6 and the fact that the choice of α, β and d is at our disposal. □

APPENDIX D. COMPARISON OF THE NOISY DD AND COMP BOUNDS

The following section is intended to prove sufficient conditions under which the DD algorithm is guaranteed to outperform COMP. However, these conditions are not necessary and DD might (and for all performed simulations does) outperform COMP for even wider settings.

Proof of Proposition 2.13. In order to prove the proposition, we need to find conditions under which

$$\min_{\alpha,d} \max\{b_1(\alpha,d), b_2(\alpha,d)\} \geq \min_{\alpha,\beta,d} \max\left\{c_1(\alpha,d), c_2(\alpha,d), c_3(\beta,d), c_4(\alpha,\beta,d)\right\}$$

We write α^* and d^* for the values that minimise the maximum of the two terms at the LHS, at which point we know that $b_1(\alpha^*,d^*) = b_2(\alpha^*,d^*)$. Then it is sufficient to show that there exists β^* such that

$$b_1(\alpha^*,d^*) = b_2(\alpha^*,d^*) \geq \max\left\{c_1(\alpha^*,d^*), c_2(\alpha^*,d^*), c_3(\beta^*,d^*), c_4(\alpha^*,\beta^*,d^*)\right\}$$

By inspection for any α and d $b_1(\alpha,d) = c_1(\alpha,d)$ and $b_2(\alpha,d) \geq c_2(\alpha,d)$ since $\theta \in (0,1)$.

Next, we will show that $b_2(\alpha,d) \geq c_4(\alpha,\beta,d)$ for any α,β in the respective bounds and $d \in (0,\infty)$. Writing $w = e^{-d}p + (1-e^{-d})(1-q)$, and recalling that by assumption that $\alpha \leq 1-w$ (or $w \leq 1-\alpha$) we readily find that

$$D_{\mathrm{KL}}(\alpha\|1-w) = \min_{1-\alpha\leq z\leq 1}(D_{\mathrm{KL}}(z\|w)) \leq \min_{1-\alpha\leq z\leq 1}\left(D_{\mathrm{KL}}(z\|w) + z\mathbf{1}\left\{\beta > \frac{ze^{-d}p}{w}\right\}D_{\mathrm{KL}}\left(\frac{\beta}{z}\|\frac{e^{-d}p}{w}\right)\right) \tag{D.1}$$

where the first equality follows since $D_{\mathrm{KL}}(\alpha\|1-w) = D_{\mathrm{KL}}(1-\alpha\|w)$ and $D_{\mathrm{KL}}(z\|w) > D_{\mathrm{KL}}(1-a\|w)$ for any $z > 1-\alpha$. The bound follows. Note that (D.1) indeed holds for any choice of α,β and d in the respective bounds stated in the theorem.

Finally, we need to demonstrate that $c_3(\beta^*,d^*) \leq b_2(\alpha^*,d^*)$. Since β is not an optimisation parameter in $b_2(\alpha^*,d^*)$ and the bound in (D.1) holds for any value of β, we can simply set it to the value that minimizes $c_3(\beta^*,d^*)$ which is $\beta = 1/\Delta$ and for which we find

$$c_3(\beta^*,d^*) = \frac{\theta}{1-\theta}\frac{1}{d^*\log\left(1-e^{-d^*}(1-q)\right.}.$$

Thus, to obtain the desired inequality we need to ensure that for the optimal choice α^* from COMP

$$\theta D_{\mathrm{KL}}\left(\alpha^*\|e^{-d^*}(1-p) + \left(1-e^{-d^*}\right)q\right) \leq -\log\left(1-e^{-d^*}(1-q)\right)$$

Using the bound

$$\theta D_{\mathrm{KL}}\left(\alpha\|e^{-d}(1-p) + \left(1-e^{-d}\right)q\right) \leq -\theta\log\left(1-\left(e^{-d}(1-p) + \left(1-e^{-d}\right)q\right)\right) \leq -\log\left(1-\left(e^{-d}(1-p) + \left(1-e^{-d}\right)q\right)\right)$$

which is obtained by setting $\alpha = 1/\Delta$, we find that $c_3(\beta^*,d^*) \leq b_2(\alpha^*,d^*)$ if

$$-\log\left(1-e^{-d^*}(1-q)\right) \geq -\log\left(1-e^{-d^*}(1-p) + \left(1-e^{-d^*}\right)q\right) \Leftrightarrow e^{-d^*}p \geq q$$

□

As mentioned before, due to bounding $b_2(\alpha^*,d^*)$ the result is not sharp. However, one immediate consequence of Proposition 2.13 is that DD is guaranteed to outperform COMP for the reverse Z channel.

APPENDIX E. RELATION TO BERNOULLI TESTING

In the noiseless case [26] shows that the constant column weight design (where each individual joins exactly Δ different tests) requires fewer tests to recover $\boldsymbol{\sigma}$ than the Bernoulli design (where each individual is included in each test with a certain probability independently). In this section we show that in the noisy case, the COMP algorithm requires fewer tests for the constant column weight design than for the Bernoulli design, and derive sufficient conditions under which the same is true for the noisy DD algorithm.

To get started, let us state the relevant bounds for the Bernoulli design. [48] derived these bounds for the Z channel, reverse Z channel and Binary Symmetric Channel. Building on this work, let us extend these bounds for the general $p-q$-model. The test design and notation is identical to the constant column design employed so far with the key difference that individuals are not assigned to Δ tests uniformly at random without replacement, but that each individual is included in each test with probability $\Delta/m = d/k$ independently. Our first observation is the size of $\boldsymbol{m}_0, \boldsymbol{m}_0^f, \boldsymbol{m}_0^u, \boldsymbol{m}_1^f$ and $\boldsymbol{m}_1^u$ carry over without further ado.

Lemma E.1. *The bounds from Lemma A.4 and Corollary A.5 hold for the Bernoulli test design.*

Proof. The crucial observation is that (A.1) now becomes $\mathbb{P}(x \in \partial a) = \Delta/m$ for any individual $x \in V$ and test $a \in F$ where we avoid any dependencies between tests that we encountered before. The rest of the proof follows exactly the proof of Lemma A.4 and Corollary A.5. □

Proposition E.2 (Noisy COMP under Bernoulli)**.** *Let $p,q \geq 0$, $p+q<1$, $d \in (0,\infty)$, $\alpha \in (q, e^{-d}(1-p) + (1-e^{-d})q)$. Suppose that $0<\theta<1$ and $\varepsilon>0$ and let*

$$m_{COMP}^{Ber} = m_{COMP}^{Ber}(n,\theta,p,q) = \min_{\alpha,d} \max\{b_1(\alpha,d), b_2(\alpha,d)\}\, k\log(n/k)$$

$$\text{where} \qquad b_1(\alpha,d) = \frac{\theta}{1-\theta}\frac{1}{kD_{\mathrm{KL}}\left(\alpha d/k \| qd/k\right)}$$

$$\text{and} \qquad b_2(\alpha,d) = \frac{1}{1-\theta}\frac{1}{kD_{\mathrm{KL}}\left(\alpha d/k \| (e^{-d}(1-p)+(1-e^{-d})q)d/k\right)}$$

If $m > (1+\varepsilon)m_{COMP}^{Ber}$, COMP *will recover $\boldsymbol{\sigma}$ under the Bernoulli test design w.h.p. given $\boldsymbol{G}, \hat{\boldsymbol{\sigma}}$.*

Proof. Using the same two-round exposition of the graph as in prior proofs and again denoting by $\boldsymbol{N}_x$ the number of displayed negative tests for an individual x, we readily find

$$\boldsymbol{N}_x \sim \mathrm{Bin}\left(m, qd/k\right) \qquad \text{for } x \in V_1$$

$$\boldsymbol{N}_x \sim \mathrm{Bin}\left(\boldsymbol{m}_0^u + \boldsymbol{m}_1^f, d/k\right) \qquad \text{for } x \in V_0$$

Using the union bound, we thus have

$$k \cdot \mathbb{P}(\boldsymbol{N}_x > \alpha\Delta | x \in V_1) = o(1) \Leftrightarrow c > b_1(\alpha,d) \tag{E.1}$$

$$(n-k) \cdot \mathbb{P}(\boldsymbol{N}_x < \alpha\Delta | x \in V_0) = o(1) \Leftrightarrow c > b_2(\alpha,d) \tag{E.2}$$

closing the proof of the proposition. □

Along the same lines, we obtain the bounds of the DD algorithm under the Bernoulli design.

Proposition E.3 (Noisy DD under Bernoulli)**.** *Let $p,q \geq 0$, $p+q<1$, $d \in (0,\infty)$, $\alpha \in (q, e^{-d}(1-p) + (1-e^{-d})q)$ and $\beta \in (e^{-d}p, e^{-d}(1-q))$. Suppose that $0<\theta<1, \zeta \in (0,\theta)$ and $\varepsilon>0$ and let*

$$m_{DD}^{Ber} = m_{DD}^{Ber}(n,\theta,p,q) = \min_{\alpha,\beta,d} \max\left\{c_1(\alpha,d), c_2(\alpha,d), c_3(\beta,d), c_4(\beta,d)\right\} k\log(n/k)$$

$$\text{where} \qquad c_1(\alpha,d) = \frac{\theta}{1-\theta}\frac{1}{kD_{\mathrm{KL}}\left(\alpha d/k \| qd/k\right)}$$

$$\text{and} \qquad c_2(\alpha,d) = \frac{1-\zeta}{1-\theta}\frac{1}{kD_{\mathrm{KL}}\left(\alpha d/k \| (e^{-d}(1-p)+(1-e^{-d})q)d/k\right)}$$

$$\text{and} \qquad c_3(\beta,d) = \frac{\theta}{1-\theta}\frac{1}{k\cdot D_{\mathrm{KL}}\left(\beta d/k \| e^{-d}(1-q)d/k\right)}$$

$$\text{and} \qquad c_4(\beta,d) = \frac{\zeta}{1-\theta}\frac{1}{k\cdot D_{\mathrm{KL}}\left(\beta d/k \| e^{-d}pd/k\right)}$$

If $m > (1+\varepsilon)m_{DD}^{Ber}$, DD *will recover* $\boldsymbol{\sigma}$ *under the Bernoulli test design w.h.p. given* $\boldsymbol{G}, \hat{\boldsymbol{\sigma}}$.

Proof. The bounds for $c_1(\alpha, d)$ and $c_2(\alpha, d)$ follow as in the proof of Proposition E.2 by replacing the right-hand side of (E.2) with n^ζ for some $\zeta \in (0, \theta)$. Next, we note that the bound for $\boldsymbol{m}_{0,\mathrm{nd}}$ of Lemma C.2 still holds as long as $\zeta \in (0, \theta)$. Using the same two-round exposition of the graph as in prior proofs and denoting by $\boldsymbol{P}_x$ the number of displayed positive tests for an individual x such that the remaining neighbourhood of the test is a subset of $V_0 \setminus V_{0,\mathrm{PD}}$, we readily find

$$k \cdot \mathbb{P}\left(\boldsymbol{P}_x < \beta\Delta | x \in V_1\right) = o(1) \Leftrightarrow c > c_3(\beta, d)$$
$$(n-k) \cdot \mathbb{P}\left(\boldsymbol{P}_x > \beta\Delta | x \in V_0\right) = o(1) \Leftrightarrow c > c_4(\beta, d)$$

concluding the proof of the proposition. □

To compare the bounds of the Bernoulli and constant-column test design we employ the following handy observation.

Lemma E.4. *Let* $0 < x, y < 1$ *and* $d > 0$ *be constants independent of* k. *As* $k \to \infty$

$$k D_{\mathrm{KL}}\left(\frac{xd}{k} \| \frac{yd}{k}\right) = d\left(D_{\mathrm{KL}}\left(x \| y\right) + v(x, y)\right) + o(1/k)$$

with

$$v(x, y) = y - x + (1-x)\log\left(\frac{1-y}{1-x}\right) \leq 0 \tag{E.3}$$

Proof. Applying the definition of the Kullback-Leibler divergence and Taylor expanding the logarithm we obtain

$$\begin{aligned} k \cdot D_{\mathrm{KL}}\left(\frac{xd}{k} \| \frac{yd}{k}\right) &= xd \cdot \log\left(\frac{x}{y}\right) + (k - xd)\left(\log\left(1 - \frac{xd}{k}\right) - \log\left(1 - \frac{yd}{k}\right)\right) \\ &= xd \cdot \log\left(\frac{x}{y}\right) + (k - xd)\left(-\frac{xd}{k} + \frac{yd}{k} + o\left(\frac{1}{k^2}\right)\right) \\ &= d\left(x \cdot \log\left(\frac{x}{y}\right) - x + y\right) + o(1/k) \\ &= d\left(D_{\mathrm{KL}}\left(x \| y\right) + y - x - (1-x)\log\left(\frac{1-x}{1-y}\right)\right) + o(1/k). \end{aligned}$$

We can bound $v(x, y)$ from above by writing the final term as $(1-x)\log\left(1 + \frac{x-y}{1-x}\right) \leq (1-x)\frac{x-y}{1-x} = x - y$, using the standard linearisation of the logarithm. □

We are now in a position to prove Proposition 2.15 and 2.16.

Proof of Proposition 2.15. The lemma follows by comparing the bounds from Theorem 2.1 and Proposition E.2 and applying Lemma E.4. □

Proof of Proposition 2.16. As evident from Corollary 2.8, the fourth bound $c_4(\alpha, \beta, d)$ vanishes under the Z channel. Now comparing the bounds from Theorem 2.2 and Proposition E.3, observing that $(1-\zeta)/(1-\theta) > 1$ for $\zeta < \theta$ and applying Lemma E.4 immediately implies the lemma. □

Appendix F. Converse bound

We can give some sense of the sharpness of these results by considering the $p-q$ communication channel. That is, we write X for the channel input and Y for the output of a noisy channel with error probabilities given exactly by Figure 1. Recall that [7, Theorem 3.1] shows that the capacity of a particular noisy group testing problem is bounded above by the Shannon capacity of the corresponding channel. For completeness we derive the capacity and optimal signalling strategy of the $p-q$ channel in terms of $h(\cdot)$, the binary entropy in nats (logarithms taken to base e):

Lemma F.1. *If $p+q<1$ the Shannon capacity of the $p-q$ channel of Figure 1 measured in nats is*

$$C_{Chan} = D_{\mathrm{KL}}\left(q\|\frac{1}{1+e^{\phi}}\right) = D_{\mathrm{KL}}\left(p\|\frac{1}{1+e^{-\phi}}\right), \tag{F.1}$$

where $\phi = (h(p)-h(q))/(1-p-q)$. This is achieved by taking

$$\mathbb{P}(X=0) = \frac{1}{1-p-q}\left(\frac{1}{1+e^{\phi}} - q\right). \tag{F.2}$$

Proof. Write $\mathbb{P}(X=0)=\gamma$ and $\mathbb{P}(Y=0)=T(\gamma) := (1-p)\gamma + q(1-\gamma)$. Then since the mutual information

$$I(X;Y) = h(Y) - h(Y|X) = h\big(T(\gamma)\big) - \big(\gamma h(p) + (1-\gamma)h(q)\big), \tag{F.3}$$

we can find the optimal T by solving

$$0 = \frac{\partial}{\partial\gamma} I(X;Y) = (1-p-q)\log\left(\frac{1-T(\gamma)}{T(\gamma)}\right) - \big(h(p)-h(q)\big),$$

which implies that the optimal $T^* = 1/(1+e^{\phi})$. We can solve for this for $\gamma^* = (T^*-q)/(1-p-q)$ to find the expression above. As $\frac{\partial}{\partial^2\gamma} I(X;Y) < 0$ it is indeed a maximum. Substituting this in (F.3) we obtain that the capacity is given by

$$\begin{aligned} h(T^*) - \big(\gamma^* h(p) + (1-\gamma^*)h(q)\big) &= h\left(\frac{1}{1+e^{\phi}}\right) - \big((T^*-q)\phi + h(q)\big) \\ &= \log(1+e^{\phi}) - \phi(1-q) - h(q) \\ &= D_{\mathrm{KL}}\big(q\|1/(1+e^{\phi})\big) \end{aligned} \tag{F.4}$$

as claimed in the first expression in (F.1) above. We can see that the second expression in (F.1) matches the first by writing the corresponding expression as $D_{\mathrm{KL}}\big(1-p\|1/(1+e^{\phi})\big) = \log(1+e^{\phi}) - \phi p - h(p)$, which is equal to (F.4) by the definition of ϕ. □

Note that this result suggests a choice of density for the matrix: since each test is negative with probability e^{-d}, equating this with (F.2) suggests that we take

$$d = d^*_{\mathrm{ch}} = \log(1-p-q) - \log\left(\frac{1}{1+e^{\phi}} - q\right).$$

This is unlikely to be optimal in a group testing sense, since we make different inferences from positive and negative tests, but gives a closed form expression that may perform well in practice. For the noiseless and BSC case observe that $\phi = 0$, and we obtain $d^*_{\mathrm{ch}} = \log 2$.

Appendix G. Illustration of bounds for Z, reverse Z channel and the BSC

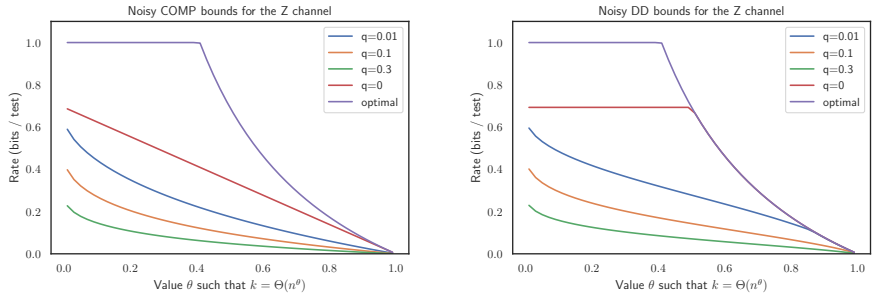

FIGURE 5. Illustration of achievability bounds for noisy COMP and DD under the Z channel. The *optimal* curve refers to the information-theoretic non-adaptive lower bound in the *noiseless* setting

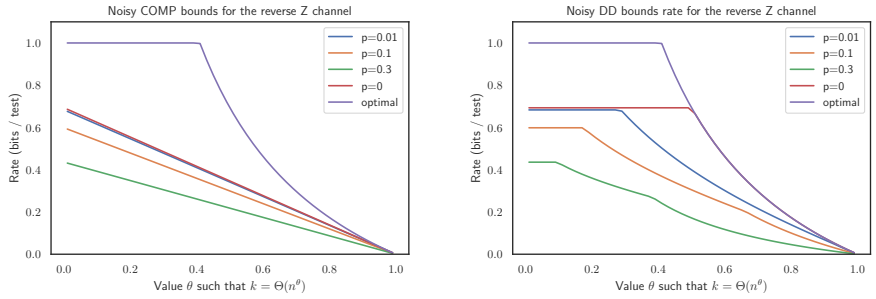

FIGURE 6. Illustration of achievability bounds for noisy COMP and DD under the reverse Z channel. The *optimal* curve refers to the information-theoretic non-adaptive lower bound in the *noiseless* setting

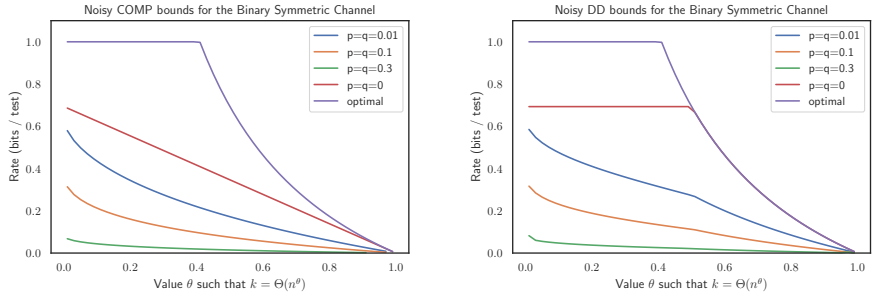

FIGURE 7. Illustration of achievability bounds for noisy COMP and DD under the Binary Symmetric Channel. The *optimal* curve refers to the information-theoretic non-adaptive lower bound in the *noiseless* setting

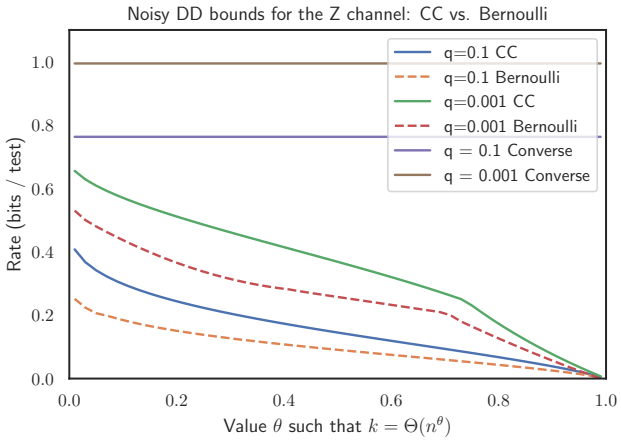

FIGURE 8. Comparison of the noisy DD rates under Bernoulli pooling ([48]) with the DD bounds and converse with constant-column design as provided in the paper at hand within the Z-Channel

REFERENCES

[1] D. Achlioptas, P. Beame, and M. Molloy (2004): Exponential bounds for dpll below the satisfiability threshold. *Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA04)*, 132–133

[2] D. Achlioptas and F. Iliopoulos (2016): Focused stochastic local search and the lovasz local lemma. *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA16)*, 2024–2038

[3] M. Aldridge (2017): The capacity of Bernoulli non-adaptive group testing. *IEEE Transactions on Information Theory*, 63:7142–7148

[4] M. Aldridge (2019): Individual testing is optimal for non-adaptive group testing in the linear regime. *IEEE Transactions on Information Theory*, 65:2058–2061

[5] M. Aldridge, L. Baldassini, and O. Johnson (2014): Group testing algorithms: bounds and simulations. *IEEE Transactions on Information Theory*, 60:3671–3687

[6] M. Aldridge, O. Johnson, and J. Scarlett (2016): Improved group testing rates with constant column weight designs. *Proceedings of 2016 IEEE International Symposium on Information Theory (ISIT16)*, 1381–1385

[7] M. Aldridge, O. Johnson, and J. Scarlett (2019): Group testing: an information theory perspective. *Foundations and Trends in Communications and Information Theory*, 15(3–4):196–392

[8] L. Baldassini, O. Johnson, and M. Aldridge (2013): The capacity of adaptive group testing. *Proceedings of 2013 IEEE International Symposium on Information Theory (ISIT13)*, 1:2676–2680

[9] C. Canonne, A. De, and R. Servedio (2020): Learning from satisfying assignments under continuous distributions. *Proceedings of the 31st Annual ACM-SIAM Symposium on Discrete Algorithms(SODA20)*, 82–101

[10] C. Chan, P. Che, S. Jaggi, and V. Saligrama (2011): Non-adaptive probabilistic group testing with noisy measurements: near-optimal bounds with efficient algorithms. *Proceedings of 49th Annual Allerton Conference on Communication, Control,and Computing*, 1: 1832–1839

[11] H. Chen and F. Hwang (2008): A survey on non-adaptive group testing algorithms through the angle of decoding. *Journal of Combinatorial Optimization*, 15:49–59

[12] I. Cheong (2020): The experience of South Korea with COVID-19. *Mitigating the COVID Economic Crisis: Act Fast and Do Whatever It Takes (CEPR Press)*, 113–120

[13] A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimroth, and P. Loick (2019): Information-theoretic and algorithmic thresholds for group testing. *Proceedings of 46th International Colloquium on Automata, Languages, and Programming (ICALP19)*, 132(43):1–14

[14] A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimroth, and P. Loick (2020): Optimal group testing. *Proceedings of 33rd Conference on Learning Theory (COLT20)*

[15] R. Dorfman(1943): The detection of defective members of large populations. *Annals of Mathematical Statistics*, 14:436–440

[16] S. Ciesek E. Seifried (2020): Pool testing of SARS-CoV-02 samples increases worldwide test capacities many times over. *https://www.bionity.com/en/news/1165636/pool-testing-of-sars-cov-02-samples-increases-worldwide-test-capacities-many-times-over.html*, last accessed on 2020-04

[17] Y. Erlich, A. Gilbert, H. Ngo, A. Rudra, N. Thierry-Mieg, M. Wootters, D. Zielinski, and O. Zuk(2015): Biological screens from linearcodes: theory and tools. *bioRxiv, page 035352*

[18] European Centre for Disease Prevention and Control (2009): Surveillance and studies in a pandemic in Europe. *https://www.ecdc.europa.eu/en/publications-data/surveillance-and-studies-pandemic-europe* (last: 06/30/2020)

[19] Y. Gefen, M. Szwarcwort-Cohen and R. Kishony (2020): Pooling method for accelerated testing of COVID-19. *https://www.technion.ac.il/en/2020/03/pooling-method-for-accelerated-testing-of-covid-19/* (last:06/30/20)

[20] E. Gould (1999) Methods for long-term virus preservation. *Mol Biotechnol*, 13:57–66

[21] A. Harrow and A. Wei (2020): Adaptive quantum simulated annealing for Bayesian inference and estimating partition functions. *Proceedings of the 31st Annual ACM-SIAM Symposium on Discrete Algorithms(SODA20)*, 193–212, 2020

[22] J. Hartline, A. Johnson, D. Nekipelov, and Z. Wang(2020): Inference from auction prices. *Proceedings of the 31st Annual ACM-SIAM Symposium on Discrete Algorithms(SODA20)*, 2472–2491

[23] W. Hoeffding (1963): Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58:301:13–30

[24] F. Hwang (1972): A method for detecting all defective members in a population by group testing. *Journal of the American Statistical Association*, 67:605–608

[25] S. Janson, T. Luczak, and A. Rucinski (2011): Random graphs *John Wiley and Sons*

[26] O. Johnson, M. Aldridge, and J. Scarlett (2018): Performance of group testing algorithms with near-constant tests per item. *IEEE Transactions on Information Theory*, 65:707–723

[27] O. Johnson and D. Sejdinovic (2010): Note on noisy group testing: Asymptotic bounds and belief propagation reconstruction. *Proceedings of 48th Allerton Conference on Communication, Control, and Computing*

[28] G. Kamath and C. Tzamos (2019): Anaconda: A non-adaptive conditional sampling algorithm for distribution testing. *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms(SODA19)*, 679–693

[29] E. Knill, A. Schliep, and D. Torney (1996): Interpretation of pooling experiments using the markov chain monte carlo method. *Journal of Computational Biology*, 3:395–406,
[30] H. Kwang-Ming and D. Ding-Zhu (2006): Pooling designs and nonadaptive group testing: important tools for dna sequencing. *World Scientific*
[31] A. Lalkhen (2008): Clinical tests: sensitivity and specificity. *Continuing Education in Anaesthesia Critical Care and Pain*, 8
[32] S. Long, C. Prober, and M. Fischer (2018): Principles and practice of pediatric infectious diseases. *Principles and practice of pediatric infectious diseases*,Elsevier
[33] N. Madhav, B. Oppenheim, M. Gallivan, P. Mulembakani, E. Rubin, and N. Wolfe (2017): Pandemics: Risks, impacts and mitiga-tion. *The World Bank:Disease control priorities*, 9:315–345
[34] D. M. Malioutov and M. Malyutov (2012): Boolean compressed sensing: Lp relaxation for group testing. *Proceedings of IEEE Inter-national Conference on Acoustics, Speech and Signal Processing*
[35] R. Mourad, Z. Dawy, and F. Morcos (2013): Designing pooling systems for noisy high-throughput protein-protein interaction experiments using boolean compressed sensing. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*,10:1478âĂŞ1490.
[36] L. Mutesa, P. Ndishimye, Y. Butera, J. Souopgui, A. Uwineza, R. Rutayisire, E. Musoni, N. Rujeni, T. Nyatanyi, E. Ntagwabira,M. Semakula, C. Musanabaganwa, D. Nyamwasa, M. Ndashimye, E. Ujeneza, I. Mwikarago, C. Muvunyi, J. Mazarati, S. Nsanzimana, N. Turok, and W. Ndifon (2020): A strategy for finding people infected with SARS-CoV-2: optimizing pooled test-ing at low prevalence *arxiv preprint: 2004.14934*
[37] H. Ngo and D. Du. (2000): A survey on combinatorial group testing algorithms with applications to dna library screening. *Discrete Mathematical Problems with Medical Applications*, 7:171–182.
[38] U.S. Department of Health and Human Services (2017): Pandemic influenza *plan.https://www.cdc.gov/flu/pandemic-resources/pdf/pandemic-influenza-implementation.pdf (last access:06/30/20)*, 2017
[39] World Health Origanisation (2009): Global surveillance during an influenza pandemic. *www.who.int/csr/resources/publications/swineflu/surveillance/en/(last access 06/30/20)*
[40] M. Plebani (2015): Diagnostic errors and laboratory medicine – causes and strategies. *Electronic Journal of the International Federation of Clinical Chemistry and Laboratory Medicine*, 26:7–14
[41] T. Richardson and R. Urbanke (2007): Modern coding theory. *Cambridge University Press*
[42] A. Sankararaman and F. Baccelli (2018): Community detection on Euclidean random graphs. *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms(SODA'18)*, 2181–2200
[43] J. Scarlett (2018): Noisy adaptive group testing: Bounds and algorithms. *IEEE Transactions on Information Theory*, 65:3646–3661.
[44] J. Scarlett (2019): An efficient algorithm for capacity-approaching noisy adaptive group testing. *Proceedings of 2019 IEEE Interna-tional Symposium on Information Theory (ISIT19)*, pages 2679–2683, 2019.
[45] J. Scarlett and V. Cevher (2016): Converse bounds for noisy group testing with arbitrary measurement matrices. *Proceedings of2016 IEEE International Symposium on Information Theory (ISIT16)*, pages 2868–2872.
[46] J. Scarlett and V. Cevher (2016): Phase transitions in group testing. *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms(SODA16)*, 1:40–53.
[47] J. Scarlett and V. Cevher (2017): Near-optimal noisy group testing via separate decoding of items. *IEEE Journal of Selected Topics in Signal Processing*, 2017.
[48] J. Scarlett and O. Johnson (2020): . Noisy non-adaptive group testing: A (near-)definite defectives approach. *IEEE Transactions on Information Theory*, 66(6):3775-3797
[49] N. Thierry-Mieg (2006): A new pooling strategy for high-throughput screening: the shifted transversal design. *BMC Bioinformatics*, 7:28, 2006
[50] L. Wang, X. Li, Y. Zhang, and K. Zhang (2011): Evolution of scaling emergence in large-scale spatial epidemic spreading. *PublicLibrary of Science ONE* 6, 2011.
[51] L. Wein and S. Zenios. Pooled testing for HIV screening (1996) : Capturing the dilution effect. *Operations Research*,44:543–569,
[52] S. Woloshin, N. Patel, and A. Kesselheim (2020) : False negative tests for SARS–CoV–2 infection:challenges and implications *New England Journal of Medicine*

APPENDIX E. EFFICIENT AND ACCURATE GROUP TESTING VIA BELIEF PROPAGATION: AN EMPIRICAL STUDY

EFFICIENT AND ACCURATE GROUP TESTING VIA BELIEF PROPAGATION: AN EMPIRICAL STUDY

AMIN COJA-OGHLAN, MAX HAHN-KLIMROTH, PHILIPP LOICK, MANUEL PENSCHUCK

ABSTRACT. The group testing problem asks for efficient pooling schemes and algorithms that allow to screen moderately large numbers of samples for rare infections. The goal is to accurately identify the infected samples while conducting the least possible number of tests. Exploring the use of techniques centred around the Belief Propagation message passing algorithm, we suggest a new test design that significantly increases the accuracy of the results. The new design comes with Belief Propagation as an efficient inference algorithm. Aiming for results on practical rather than asymptotic problem sizes, we conduct an experimental study. *MSc: 05C80, 60B20, 68P30*

1. INTRODUCTION

1.1. **The group testing problem.** In science generally and in applied science particularly there is much to be said for simplicity. But occasionally a modest degree of sophistication carries extraordinary rewards. Group testing is a case in point. Every single day medical labs around the globe screen moderately large numbers of samples for rare pathogens. The vast majority of samples, anywhere between 90% and 99.9%, are actually uninfected [7, 25, 28, 32, 35, 36, 37, 38, 39]. Labs therefore test pools of samples rather than individual samples. The *group testing problem* asks for pooling strategies that minimise the total number of tests required while maximising the accuracy of the results. The latter is crucial because test results are generally not perfectly accurate.

Coming up with practical solutions to this problem turns out to be challenging precisely because the total number of samples in a real-world scenario is moderate–say, in the hundreds or thousands. To elaborate, on the one hand the group testing problem has inspired a body of beautiful mathematical work that deals with the asymptotical scenario where the number of samples grows to infinity [4, 10, 11]. However, such asymptotical results do not directly bear on practical problem sizes. Besides, the theoretical test designs tend to suffer other drawbacks such as asking for excessively large test pools or subdivisions of individual samples into very many sub-samples [4, 10, 11]. On the other hand, practical problem sizes far exceed the limits up to which an exhaustive search for an optimal test design seems remotely feasible. As a consequence, the pooling schemes in practical use remain the self-same extremely simple ones suggested in the 1940s [7, 25, 28, 32, 35, 36, 37, 38, 39].

The aim of this paper is to investigate better test designs for practical problem sizes. The focus is on improving the *accuracy* of the results, i.e., avoiding false positive and/or negative diagnoses while keeping the number of tests as small as possible. Indeed, the thrust of this paper is that the idea of group testing, originally invented to reduce the number of tests, actually excels at improving the accuracy of the results. This may seem surprising at first glance because one might deem individual testing optimal in terms of accuracy. It is not. Group testing does better in much the same way as error-correcting codes gain power from encoding entire blocks of data simultaneously.

Given the moderate number of samples in real-world scenarios, an empirical approach is the only feasible way to obtain practically meaningful results. Thus, taking on board the intuition from theoretical work on group testing as well as recent ideas from information theory and statistical physics, we conduct an extensive experimental study. The main finding is that a novel test scheme called *adaptive Belief Propagation* greatly improves the accuracy of the overall results while keeping the number of tests conducted low. Furthermore, the new test design requires only relatively small test pools and only assigns each sample to a small number of tests. Finally, the design comes with an efficient, easy-to-implement algorithm to infer the status of the individual samples from the test results, namely the Belief Propagation message passing algorithm.

We proceed to discuss the mathematical model of tests and samples that we work with. Subsequently, we present the results of adaptive Belief Propagation by comparison to other test schemes. These test schemes, which are partly incorporated into adaptive Belief Propagation, are discussed in detail in Section 2. In Section 3 we then present the new test design and the corresponding inference algorithm. Section 4 details the theoretical and

Amin Coja-Oghlan, Max Hahn-Klimroth and Philipp Loick are supported by DFG CO 646/3 and DFG CO 646/5. Manuel Penschuck is supported by ME 2088/5-1.

heuristical considerations that underpin adaptive Belief Propagation. Finally, in Section 5 we discuss the potential impact of the new results and future directions for both empirical and theoretical work.

1.2. **The model.** We work with a simple but standard model of group testing that allows for test results to not be entirely accurate [4]. Let $x_1,\dots,x_n$ represent the samples submitted for testing. We assume that with a prior probability of $\lambda \in [0,1]$ any one sample is infected is known. The true infection status of each sample is indicated by $\boldsymbol{\sigma}(x_j) \in \{0,1\}$, with 1 representing 'infected'. The $\boldsymbol{\sigma}(x_j)$ are assumed to be independent Bernoulli variables with mean λ. We refer to the vector $\boldsymbol{\sigma} = (\boldsymbol{\sigma}(x_j))_{j=1,\dots,n}$ as the *ground truth*. Let $\boldsymbol{k} = \sum_{j=1}^{n} \mathbf{1}\{\boldsymbol{\sigma}(x_j) = 1\}$ signify the actual number of infected samples.

The way how test pools are formed is represented by a bipartite graph. To be precise, a *test design* is a bipartite graph G with one class $\mathscr{X} = \{x_1,\dots,x_n\}$ of vertices representing the n samples and the other class $\mathscr{A} = \{a_1,\dots,a_m\}$ representing the test pools. An edge between x_j and a_i indicates that x_j is included in test pool a_i. For each x_j we let $\partial x_j = \partial_G x_j$ be the set of test pools that include x_j. Similarly, for each test pool a_i we write ∂a_i for the set of individual samples x_j included in that pool.

Each test a_i reports a positive or negative result $\hat{\boldsymbol{\sigma}}(a_i) \in \{0,1\}$. Ideally a test a_i should come back positive iff at least one sample $x_j \in \partial a_i$ is actually infected. But the test results need not be completely accurate. We therefore posit two parameters p, called the *specificity*, and q, the *sensitivity*, both between 0 and 1, such that the tests return results

$$\hat{\boldsymbol{\sigma}}(a_i) = \begin{cases} 0 & \text{with probability } p \\ 1 & \text{with probability } 1-p \end{cases} \qquad \text{if } \boldsymbol{\sigma}(x_j) = 0 \text{ for all } x_j \in \partial a_i \tag{1.1}$$

$$\hat{\boldsymbol{\sigma}}(a_i) = \begin{cases} 0 & \text{with probability } 1-q \\ 1 & \text{with probability } q \end{cases} \qquad \text{if } \boldsymbol{\sigma}(x_j) = 1 \text{ for some } x_j \in \partial a_i. \tag{1.2}$$

The random outcomes in (1.1)–(1.2) are mutually independent given $\boldsymbol{\sigma}$. Let $\hat{\boldsymbol{\sigma}} = (\hat{\boldsymbol{\sigma}}(a_i))_{i=1,\dots,m}$ encompass the test results.

Generally the ground truth $\boldsymbol{\sigma}$ cannot be inferred with perfect accuracy form the test results $\hat{\boldsymbol{\sigma}}$ of a single "one-shot" test design (unless $p = q = 1$ and we test every x_j separately) [1]. Indeed, under the noise model (1.1)–(1.2) the posterior of the ground truth given the test results reads[1]

$$\mu_G(\sigma) = \mathbb{P}\left[\boldsymbol{\sigma} = \sigma \mid \hat{\boldsymbol{\sigma}}\right] \propto \prod_{i=1}^{n} \lambda^{\sigma(x_i)}(1-\lambda)^{1-\sigma(x_i)} \prod_{i=1}^{m} \psi_{\hat{\boldsymbol{\sigma}}(a_i)}\left((\sigma(y))_{y\in\partial a_i}\right) \qquad (\sigma = (\sigma(x_i))_{i=1,\dots,n} \in \{0,1\}^n) \tag{1.3}$$

$$\text{where} \quad \psi_0(\sigma_1,\dots,\sigma_\ell) = p^{1-\bigvee_{i=1}^{\ell}\sigma_i}(1-q)^{\bigvee_{i=1}^{\ell}\sigma_i}, \qquad \psi_1(\sigma_1,\dots,\sigma_\ell) = (1-p)^{1-\bigvee_{i=1}^{\ell}\sigma_i} q^{\bigvee_{i=1}^{\ell}\sigma_i}. \tag{1.4}$$

Hence, the information-theoretically optimal inference algorithm just draws a random sample from the distribution μ_G. In effect, the accuracy with which the ground truth can potentially be recovered is governed by the entropy of the posterior μ_G: the smaller the entropy the better the results. Furthermore, depending on the specific design G there may or may not exists an *efficient* algorithm for sampling from μ_G.

To deal with these challenges, in *adaptive group testing* tests are not deployed in a single stage like in (1.1)–(1.2) but in several stages. To be precise, an *ℓ-stage test design* is an increasing sequence $G^{(0)}, G^{(1)},\dots,G^{(\ell)}$ of test designs such that $G^{(i+1)}$ is obtained from $G^{(i)}$ by adding further tests. How many tests are added and which samples they contain depends on results from the previous stages. The results of the new tests are assumed to be distributed independently according to (1.1)–(1.2). The aim, of course, is to diligently add tests so as to maximally reduce the entropy of the posterior.

In summary, the group testing problem poses the following challenges.

(i) To come up with an adaptive test design that allows to infer the true infection status $\boldsymbol{\sigma}(x_j)$ of as many x_j as possible while conducting as small a number of tests as possible.
(ii) To devise an *efficient* algorithm that actually infers the $\boldsymbol{\sigma}(x_j)$ from the observed $\hat{\boldsymbol{\sigma}}(a_i)$ with reasonable computational effort.
(iii) To facilitate practical adoption we need to avoid high degrees because very large test pools may be infeasible, as may be dividing an individual sample into very many pools.
(iv) To ensure a timely reporting of test outcomes we should aim for a small number of test stages, or at least ensure that most samples can be diagnosed after the first or second stage.

[1] In (1.3) the $\propto$-symbol hides the normalisation required to turn μ_G into a probability distribution.

1.3. **Results.** To meet these objectives we devise a new test scheme called adaptive Belief Propagation. We investigate its performance empirically for the following parameter choices.

- The results in this section refer to $n = 1000$ samples. In Section 4.3 we discuss that the performance is similar on instances with $n = 100$ and slightly better with $n = 10000$.
- We study prior infection probabilities $\lambda = 0.5\%,\ 1\%,\ 5\%,\ 10\%$.
- Three different specificity/sensitivity scenarios are investigated:
  (a) perfectly reliable tests, i.e. $p = q = 1$,
  (b) moderately high values $p = 0.99$ and $q = 0.98$ which reflects, among others, the reliability of certain Covid-19 tests [7, 8, 31, 38, 41] and
  (c) a noisy scenario with $p = q = 0.95$.
- Each experiment is run 100 times independently for each parameter combination.

The experiments show that adaptive Belief Propagation improves the accuracy of the results by an order of magnitude by comparison to known test designs while keeping the number of tests at a reasonable level. Let us begin with the high-noise scenario $p = q = 0.95$, where we reap the greatest gains. We propose three different test designs adaptive Belief Propagation 1, adaptive Belief Propagation 2 and adaptive Belief Propagation 3. The first strikes a balance between accuracy of results and the number of tests, while the latter emphasises accuracy. In the following, let the *false positive rate (fpr)* be the number of healthy samples falsely classified as infected over all healthy samples. Correspondingly, let the *false negative rate (fnr)* be the number of infected samples falsely classified as healthy over all infected samples. Figure 1 displays the results of adaptive Belief Propagation 1 in comparison to several previously known approaches. These include the *2-stage Dorfman* and the 3-stage Dorfman designs, which are widely used in practice, as well as *Belief Propagation* followed by individual testing advocated in the theoretical literature [2]. A further scheme that we included is *informative Dorfman*, a 2-stage design proposed in [29]. We will discuss these approaches in some detail in Section 2. Figure 1 shows that with about the same number of tests as 2-stage Dorfman, adaptive Belief Propagation achieves up to 78% reduction in the number of false positives and an up to 42% reduction in the number of false negatives. The gains are particularly high for small priors.

Nevertheless, the absolute value of the false positive and false negative rate of all test designs in Figure 1, particularly for large priors, may still be unacceptably high for many real-world applications. Here our other two designs adaptive Belief Propagation 2 and adaptive Belief Propagation 3 come to the rescue. As Figure 2 shows, these designs, particularly adaptive Belief Propagation 3, dramatically reduce the number of false positives and negatives. Of course, these improvements come at the expense of a larger number of tests. But for priors $\lambda \leq 0.05$ the number of extra tests is moderate, and for the largest prior $\lambda = 0.1$ adaptive Belief Propagation 2 and adaptive Belief Propagation 3 require not many more tests than individual testing while being the only designs that deliver decent accuracy.

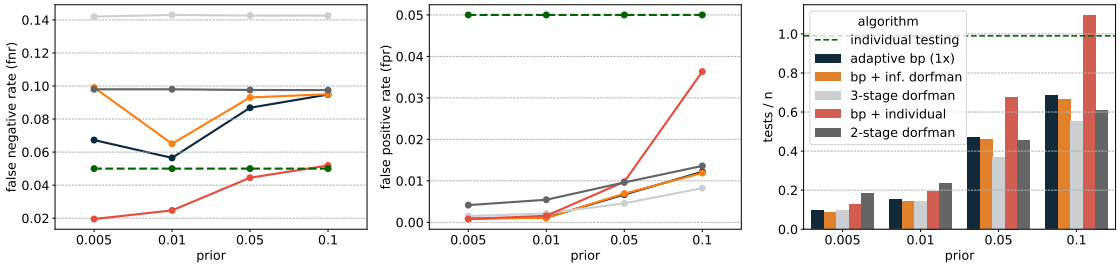

FIGURE 1. Simulation results for high noise scenario (sensitivity and specificity of 95%)

Matters turns out similar in the case of moderately high sensitivity and specificity $p = 0.99$, $q = 0.98$. Figure 3 displays the results. In comparison to the classical two- and three-stage Dorfman scheme, adaptive Belief Propagation requires at most 11% more tests for high priors of $\lambda = 0.1$ - for small priors even fewer tests. The benefit is

[2]Note that with perfectly reliable tests, this approach is equivalent to the so-called *definite defectives (DD)* algorithm followed by individual testing.

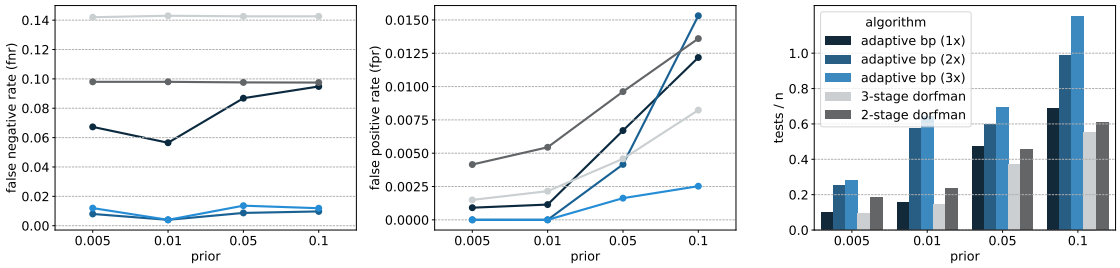

FIGURE 2. Simulation results of reliability-enhanced adaptive Belief Propagation for high noise scenario (sensitivity and specificity of 95%)

that adaptive Belief Propagation boosts accuracy compared to all the previously known designs, particularly so for low priors. We point out that the gains vis-a-vis *informative Dorfman* for moderately high priors are modest. The key benefit in adaptive Belief Propagation however, lies in its versatility to meaningfully enhance the accuracy at the expense of somewhat more tests as shown in Figure 4. A similar extension of *informative Dorfman* would yield a similar accuracy but require significantly more tests than adaptive Belief Propagation.

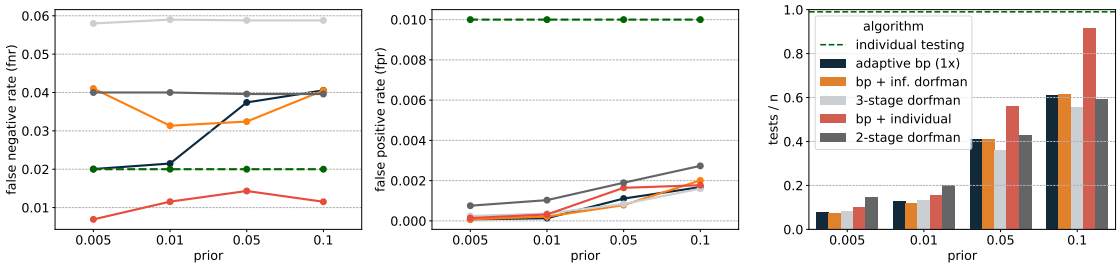

FIGURE 3. Simulation results for sensitivity for moderate noise scenario (sensitivity of 99%, specificity of 98%)

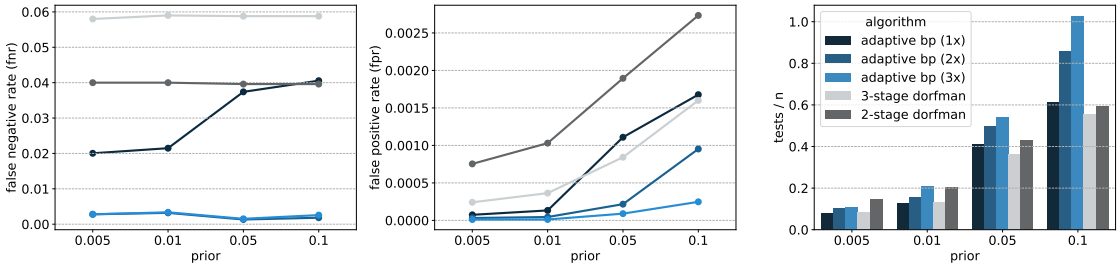

FIGURE 4. Simulation results of reliability-enhanced adaptive Belief Propagation for moderate noise scenario (sensitivity of 99%, specificity of 98%)

Even with perfectly reliable tests, the conventional *definite defectives (DD)* approach in the literature can be improved upon by adaptive Belief Propagation or the *informative Dorfman* approach. Both schemes are able to reduce the number of tests compared to the former by up to 18% and comes within 19% to 32% of the information-theoretic lower bound. The gains vis-a-vis two-stage Dorfman with up to 57% and individual testing with up to 94% are even more pronounced. We do not need to consider the accuracy in the noiseless case since all test designs recover the entire ground truth by construction.

4

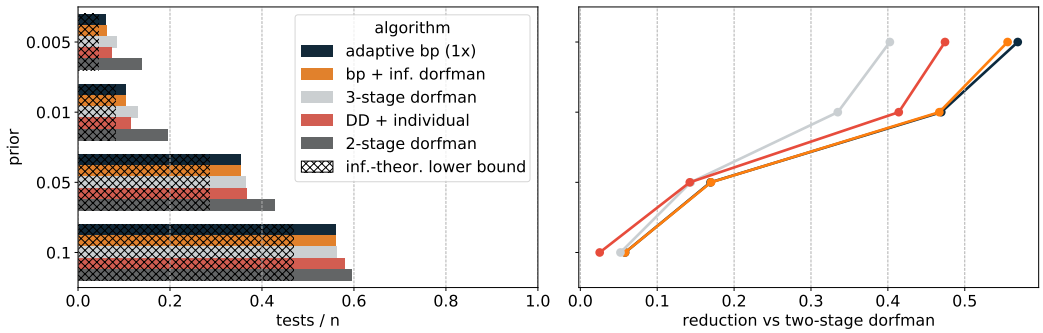

FIGURE 5. Simulation results for the noiseless setting. The left plot displays the numbers of tests required by the different designs; the black hatched area represents a plausible information-theoretic lower bound for the number of tests. The right plot shows the reduction achieved by comparison to the 2-stage Dorfman procedure, a classical and widely used test design.

All examined algorithms require reasonable pool sizes and splits of the individual sample that are in line with common pooling procedures [22, 24, 29]. The maximum pool size is between 8 and 170 depending on noise level and prior, while the splits of the individual sample range between 3 and 19. It should be noted that the proposed algorithms and test designs can readily be adjusted to accommodate smaller pool sizes or individual sample splits — at the expense of somewhat more tests.

Organisation. In Section 2, we will discuss designs and algorithms that are in practical use or have been studied in the mathematical literature on group testing. Subsequently, we present the details behind our novel test design named adaptive Belief Propagation in Section 3. In Section 4 we relate adaptive Belief Propagation to the theoretical work on group testing and asymptotic considerations.

2. DESIGNS AND ALGORITHMS

We discuss the various previously studied test designs and inference algorithms. In Section 3 we will then see how we extend and modify these known constructions in order to obtain the new adaptive Belief Propagation design.

2.1. **Individual testing.** The most straightforward test strategy, of course, is to conduct $m = n$ individual tests for each of the n samples. At first glance, individual testing may appear to be the gold standard in terms of accuracy. Naturally, in the case $p = q = 1$, individual testing will register the status of each sample correctly. However, realistic values for p and q range between 0.95 and 0.99 [7, 8, 31, 38, 41]. If p, q are less than one, then individual testing will produce numbers of false positives/negatives distributed as $\mathrm{Bin}(n - \boldsymbol{k}, 1 - p)$ and $\mathrm{Bin}(\boldsymbol{k}, 1 - q)$, respectively.

The accuracy of the results could obviously be boosted by conducting two or three individual tests per sample. Indeed, if we test each x_j twice and report x_j as infected only if both tests come back positive, then we could reduce the expected number of false positives to $(n - \boldsymbol{k})(1 - p)^2$. But we would now expect a slightly larger number of $2\boldsymbol{k}(1 - q)$ false negatives. To reduce the number of false positives and negatives simultaneously we could test each x_j thrice and report the majority of the three test results.

2.2. **Dorfman and grid designs.** The test designs that currently appear to be most widely used in practice date back to the 1940s. Indeed, the idea of group testing was first brought up by Dorfman in 1943 [18]. He suggested a two-stage test procedure. In the first stage, every sample gets placed in precisely one pool. All pools are the same size, which depends on the prior λ only. Pools with a positive test result get tested separately in the second stage. An illustration is provided in Figure 6. Depending on the prior, this scheme can significantly reduce the number of tests required. For example, with $\lambda = 0.05$ this scheme uses pools of size five and the expected overall number of tests conducted in both stages comes to about $0.426n$. At the same time, Dorfman's two-stage procedure reduces the number of false positives because a sample is ultimately reported as positive only if both the tests are positive. But for the same reason, the expected number of false negatives increases. For instance, with $n = 10^4$ and $k = \lambda n = 500$ as above, we expect 18.2 false positives and 9.95 false negatives.

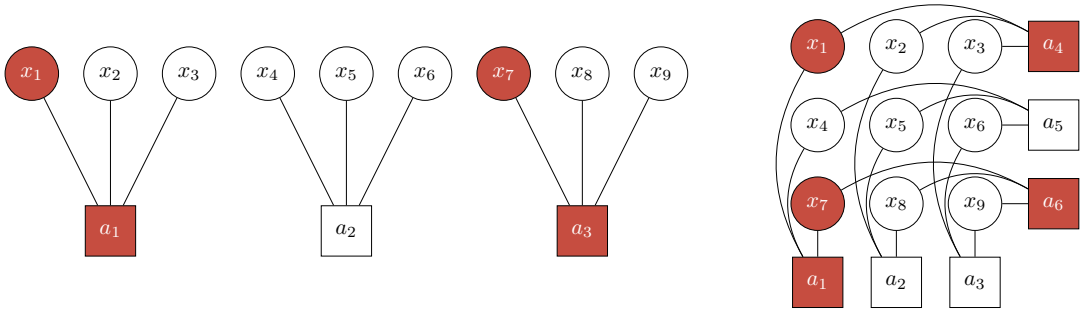

FIGURE 6. Illustration of first stage of the first stage of the Dorfman scheme (left) and grid test designs (right)

A natural extension of the Dorfman procedure employs three stages. In the first stage, relatively large pools are formed. The second stage then splits the positive ones into smaller sub-pools and the third stage resorts to individual testing. In effect, as with the two-stage procedure, the expected number of false positives drops while the expected number of false negatives increases. For $n = 10^4$ and $\lambda = 0.05$ as above the expected numbers of false positives/negatives work out to be 11.76 and 14.8, respectively.

Grid designs are a variation on the Dorfman scheme. They partition the set of all individual samples into equal-sized subsets. For instance, if $\lambda = 0.05$ the size would be 16. Each of these subsets is mapped onto a 4×4 grid. Its rows and columns constitute the pools for the first stage. Thus, each sample lands in two first-stage pools. Depending on the results, further tests are conducted in a second stage; see Figure 6 for an illustration. Grid designs significantly reduce the number of false negatives by comparison to individual testing while increasing the number of false positives. However, the number of tests required exceeds that of the two-stage Dorfman procedure.

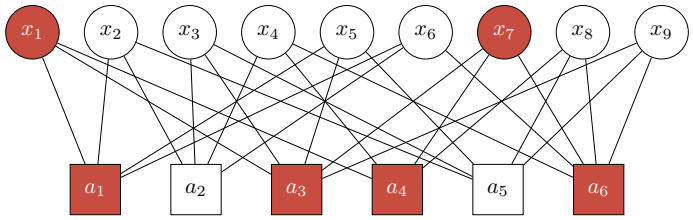

FIGURE 7. Illustration of a random biregular test design with $\Delta = 3$ and $\Gamma = 4$

2.3. **Probabilistic constructions.** More sophisticated test designs have been proposed in the mathematical theory of group testing. The best current, and in certain asymptotic settings provably optimal, test designs harness randomisation [4, 11]. For instance, in the *random biregular test design* every test pool has an equal size Γ and every individual sample joins an equal number Δ of pools; see Figure 7 for an illustration. In other words, the test design $G = G_{n,m}(\Gamma,\Delta)$ is chosen uniformly at random from the set of all (Δ,Γ)-regular bipartite graphs [21].[3] In order to extract the maximum amount of information about the ground truth, the parameters Γ,Δ should be chosen so as to maximise the conditional entropy of the vector $\hat{\boldsymbol{\sigma}}$ of test results. Hence, Γ,Δ should be chosen so that about half the tests will be positive:[4]

$$\Delta = \frac{m\log(2)}{n\lambda} \qquad \Gamma = \frac{\log(2)}{\lambda}. \tag{2.1}$$

Why does such a randomised construction seem promising? Intuitively the randomness of the test design reduces dependencies between the different test results $\hat{\boldsymbol{\sigma}}(a_i)$ to a minimum. Thus, with Δ,Γ chosen as above and

[3]To be precise, G is typically drawn from the pairing model of graphs with given degrees. In this model, it is rare but possible that the same individual joins a test pool twice. In practice, such double occurrence could, of course, be reduced to single occurrences.

[4]Of course, due to rounding issues we cannot ensure that the expected number of positive tests is precisely equal to $m/2$.

for a number m of tests up to a certain threshold, we can hope to squeeze as much as one bit's worth of information from each test. Similar randomised constructions have proved powerful in coding theory and compressed sensing as well [16, 15, 26, 34].

While in the designs that we discussed previously obvious inference algorithms suggested themselves, in the case of the random biregular design matters are not quite so straightforward. In the case $p = q = 1$ maximum a posteriori inference boils down to a minimum hypergraph vertex cover problem [10]. However, this problem is NP-hard and even on random instances no efficient algorithm is known.

A blunt but efficient algorithm that has been analysed in the case $p = q = 1$ goes by the name of *definite defectives* ('DD') [3]. The algorithm classifies as infected every sample that is included in positive test pools only and that appears in at least one positive test pool where all other samples appear in a negative test. All other samples are classified as uninfected. In symbols,

$$\boldsymbol{\sigma}_{\mathtt{DD}}(x_j) = \bigwedge_{a \in \partial x_j} \hat{\boldsymbol{\sigma}}(a) \wedge \bigvee_{a \in \partial x_j} \bigwedge_{\substack{y \in \partial a \\ y \neq x_j}} \bigvee_{b \in \partial y} (1 - \hat{\boldsymbol{\sigma}}(b)).$$

For $p = q = 1$ this algorithm will never produce false positives but may render false negatives. Several similarly-flavoured algorithms have been analysed mathematically. Aldridge analysed an adaptive test design whose different stages employ random biregular test designs with suitably chosen degrees [2]. This adaptive test design carried out over an unbounded number of stages achieves rates in excess of 0.95 bits per tests. However, the large number of stages might render the scheme impractical.

2.4. **Glauber dynamics.** The DD algorithm merely extracts binary information about each sample. For a more fine-grained picture we would need to get a better handle on the posterior distribution (1.3) of the random test design. An immediate idea is to use a Markov Chain Monte Carlo algorithm to approximate the marginals of the posterior. Specifically, the Glauber dynamics starts at a random initial configuration $\sigma^{(0)} = (\sigma^{(0)}(x_i))_{i=1,\ldots,n}$ drawn from the prior. Thus, the individual $\sigma^{(0)}(x_i)$ are independent $\mathrm{Be}(\lambda)$ variables. Glauber then proceeds to generate a random sequence $(\sigma^{(t)})_{t=0,\ldots,T}$ of configurations by updating the status of a random sample at each time step according to (1.3); see [27] for a detailed derivation of the Glauber update rule. The hope is that for moderate T the empirical means of the sequence approximate the actual posteriors well, i.e.,

$$\mu_G(\{\boldsymbol{\sigma}(x_j) = s\}) \approx \frac{1}{T}\sum_{i=0}^{T} \mathbf{1}\left\{\sigma^{(t)}(x_j) = s\right\} \qquad (j = 1, \ldots, n;\ s = 0, 1). \tag{2.2}$$

At this point, no mathematical analysis of Glauber exists. Furthermore, an empirical assessment of (2.2) is difficult because even for moderate values of n we cannot hope to compute the marginals of the posterior (1.3) exactly by exhaustive enumeration. Nonetheless, an experimental study of Glauber has been conducted in [14].

2.5. **Informative Dorfman.** Even if we assume that Glauber (or some other algorithm) approximates the posterior marginals well, how could we use this information in the second stage? A simple idea is to revisit the original Dorfman design. Hence, equipped with the posterior marginals from the first round, we could set up test pools such that each sample gets placed in precisely one pool. But now we could try to take the posteriors from the first stage into consideration in compiling the pools. Finally, just like in the original Dorfman scheme one could test the samples in each pool that returns a positive result separately. This procedure goes by the name of *informative Dorfman* [29].

How exactly do we take advantage of the marginals to set up the pools? A natural idea is to sort the samples increasingly according to their marginals and pool them in this order. A simple optimisation given the sequence of marginals then yields the optimal sequence of pool sizes. The pools containing samples with small marginals are relatively large, while samples with marginals above 0.3 get tested individually. A combination of Glauber and informative Dorfman has been studied empirically in [14]. The key finding was that for a given number of tests this procedure worked decently well but was still outperformed by quite a margin by more complicated multi-stage test designs and algorithms. In our study, we find that the marginals obtained by running Belief Propagation closely resemble the empirical marginals of Glauber and thus consistently use Belief Propagation in our analyses.

3. Adaptive Belief Propagation

In this section we discuss the new design and inference algorithm. The first stage employs the random biregular test design from Section 2.3. Given the results of the first stage, in the second and third stage we use a blend

of the random biregular design and informative Dorfman. For the inference algorithm we seize upon the Belief Propagation message passing paradigm [33]. Since Belief Propagation and the mathematical theory behind this algorithm inform the entire construction, that is where we start.

3.1. **Belief Propagation.** In recent years the Belief Propagation message passing paradigm has been applied in combination with randomised constructions with stunning success. Prominent examples include coding theory and other signal processing tasks such as compressed sensing [16, 26, 34]. The development of the Belief Propagation technique in conjunction with randomised constructions has been inspired by deep ideas from the statistical mechanics of disordered systems [30]. More recently, a substantial body of mathematical research has been devoted to Belief Propagation; e.g., [5, 12, 19, 40]. Although most of the theoretical work from both the physics and maths communities is intrinsically asymptotical, we let these ideas guide our quest for a practical group testing design.

Belief Propagation is a generic message passing technique for approximating the marginals of Boltzmann distributions on factor graphs. The posterior distribution (1.3) turns out to be a specimen of such a Boltzmann distribution. The basic intuition behind Belief Propagation, which has been substantiated mathematically to a good extent, is that under certain assumptions the posterior distribution admits a succinct representation in terms of *messages* [12, 13, 30, 42]. These assumptions are provably met in many Bayes-optimal inference problems on random factor graphs, at least asymptotically as the problem size tends to infinity [6, 9]. The group testing problem as modelled in Section 1.2 is an example of such a Bayes-optimal inference problem.

At first glance the posterior distribution (1.3) appears to be quite a difficult object of study. For instance, if we were to estimate the entropy of this distribution we might have to inspect all 2^n possible vectors $\sigma \in \{0,1\}^n$. But according to the Belief Propagation paradigm we can get a handle on the posterior distribution in terms of messages associated with the edges of the test design $G = G_{n,m}(\Gamma, \Delta)$. Formally, the *message space* of $\mathscr{M}(G)$ consists of vectors

$$(\mu_{x_j \to a_i}(s), \mu_{a_i \to x_j}(s))_{j=1,\ldots,n;\, i=1,\ldots,m;\, x_j \in \partial a_i;\, s \in \{0,1\}}.$$

The idea is that there are two messages $\mu_{x_j \to a_i}(\cdot)$, $\mu_{a_i \to x_j}(\cdot)$ associated with every edge of G, one directed from the sample x_j to the test a_i and one in the opposite direction. The messages themselves are probability distributions on $\{0,1\}$. Thus,

$$\mu_{x_j \to a_i}(0), \mu_{x_j \to a_i}(1) \in [0,1] \quad \text{and} \quad \mu_{x_j \to a_i}(0) + \mu_{x_j \to a_i}(1) = 1,$$

and similarly for $\mu_{a_i \to x_j}(\cdot)$.

Roughly speaking, $\mu_{a_i \to x_j}(\cdot)$ is meant to represent the impact that a_i has on x_j in the absence of all other tests $b \in \partial x_j$. Moreover, $\mu_{x_j \to a_i}(\cdot)$ represents the status of x_j in the absence of test a_i. More formally, we define the *standard message* $\mu_{G, x_j \to a_i}(s)$ as the posterior probability that $\boldsymbol{\sigma}(x_j) = s$ given the test design $G - a_i$ obtained from G by omitting test a_i and given the test results $(\hat{\boldsymbol{\sigma}}(a_h))_{h \neq i}$. In light of (1.3) we can write this probability out explicitly as

$$\mu_{G, x_j \to a_i}(s) \propto \sum_{\substack{\sigma \in \{0,1\}^{\mathscr{X}} \\ \sigma(x_j) = s}} \prod_{i=1}^{n} \lambda^{\sigma(x_i)} (1-\lambda)^{1-\sigma(x_i)} \prod_{i=1}^{m} \psi_{\hat{\boldsymbol{\sigma}}(a_i)}\big((\sigma_y)_{y \in \partial a_i}\big) \quad (j = 1, \ldots, n;\, i = 1, \ldots, m;\, x_j \in \partial a_i;\, s \in \{0,1\}),$$

with the $\propto$ -sign hiding the normalisation to ensure that $\mu_{G, x_j \to a_i}(0) + \mu_{G, x_j \to a_i}(1) = 1$. Similarly, the standard message $\mu_{G, a_i \to x_j}(s)$ is defined as the posterior probability that $\boldsymbol{\sigma}(x_j) = s$ given the test design $G - (\partial x_j \setminus \{a_i\})$ obtained by removing all tests that x_j takes part in except for a_i and given the test results $\hat{\boldsymbol{\sigma}}(a_h)$ of all tests $a_h \notin \partial x_j \setminus \{a_i\}$.

Conceived wisdom, vindicated mathematically for a broad family of inference problems, predicts that asymptotically these messages satisfy the following *Belief Propagation equations* [6, 9, 12, 42]:

$$\mu_{G, x \to a}(s) \propto \lambda^s (1-\lambda)^{1-s} \prod_{b \in \partial x \setminus \{a\}} \mu_{G, b \to x}(s), \tag{3.1}$$

$$\mu_{G, a \to x}(0) \propto 1 - q + (p + q - 1) \prod_{y \in \partial a \setminus x} \mu_{G, y \to a}(0), \qquad \mu_{G, a \to x}(1) \propto 1 - q \qquad \text{if } \hat{\boldsymbol{\sigma}}(a) = 0, \tag{3.2}$$

$$\mu_{G, a \to x}(0) \propto q + (1 - p - q) \prod_{y \in \partial a \setminus x} \mu_{G, y \to a}(0), \qquad \mu_{G, a \to x}(1) \propto q \qquad \text{if } \hat{\boldsymbol{\sigma}}(a) = 1. \tag{3.3}$$

These equations express the notion that the random biregular design $G_{n,m}(\Gamma,\Delta)$ minimises dependencies between the test results. Furthermore, we expect that the marginals of the posterior distribution can be well approximated in terms of the messages:

$$\mu_G(\{\boldsymbol{\sigma}(x_i)=s\}) \propto \lambda^s(1-\lambda)^{1-s}\prod_{b\in\partial x_i}\mu_{G,b\to x_i}(s). \tag{3.4}$$

Apart from the marginals, asymptotic results also suggest that the entropy of the posterior distribution can be approximated in terms of the messages [9, 12, 30]. This approximation comes in terms of a functional called the *Bethe free energy*, defined as

$$\mathscr{B}_G = \sum_{x\in\mathscr{X}}\mathscr{B}_{G,x} + \sum_{a\in\mathscr{A}}\mathscr{B}_{G,a} - \sum_{x\in\mathscr{X},a\in\partial x}\mathscr{B}_{G,x,a} \qquad \text{with} \tag{3.5}$$

$$\mathscr{B}_{G,x} = \log\sum_{s\in\{0,1\}}\prod_{a\in\partial x}\mu_{G,a\to x}(s) \tag{3.6}$$

$$\mathscr{B}_{G,a} = \begin{cases}\log\bigl(1-q+(p+q-1)\prod_{x\in\partial a}\mu_{G,x\to a}(0)\bigr) & \text{if } \hat{\boldsymbol{\sigma}}(a)=0\\ \log\bigl(q+(1-p-q)\prod_{x\in\partial a}\mu_{G,x\to a}(0)\bigr) & \text{if } \hat{\boldsymbol{\sigma}}(a)=1\end{cases} \tag{3.7}$$

$$\mathscr{B}_{G,x,a} = \log\sum_{s\in\{0,1\}}\mu_{G,x\to a}(s)\mu_{G,a\to x}(s). \tag{3.8}$$

The resulting approximation of the entropy reads

$$\begin{aligned}\mathscr{H}_G = \mathscr{B}_G - n\log\lambda + \sum_{i=1}^{n}\mu_G(\{\boldsymbol{\sigma}_x=0\})\log\frac{\lambda}{1-\lambda} \\ - \sum_{\substack{i=1\\ \hat{\boldsymbol{\sigma}}(a_i)=0}}^{m}\frac{p\prod_{x\in\partial a_i}\mu_{G,x\to a_i}(0)}{1-q+(p+q-1)\prod_{x\in\partial a_i}\mu_{G,x\to a_i}(0)}\log p + \frac{(1-q)(1-\prod_{x\in\partial a_i}\mu_{G,x\to a_i}(0))}{1-q+(p+q-1)\prod_{x\in\partial a_i}\mu_{G,x\to a_i}(0)}\log(1-q) \\ - \sum_{\substack{i=1\\ \hat{\boldsymbol{\sigma}}(a_i)=1}}^{m}\frac{(1-p)\prod_{x\in\partial a_i}\mu_{G,x\to a_i}(0)}{q+(1-p-q)\prod_{x\in\partial a_i}\mu_{G,x\to a_i}(0)}\log(1-p) + \frac{q(1-\prod_{x\in\partial a_i}\mu_{G,x\to a_i}(0))}{q+(1-p-q)\prod_{x\in\partial a_i}\mu_{G,x\to a_i}(0)}\log q.\end{aligned} \tag{3.9}$$

Hence, in order to estimate the marginals and the entropy of the posterior we need to calculate the Belief Propagation messages. A natural idea is to perform a fixed point iteration using the Belief Propagation equations (3.1)–(3.3). Of course, the equations (3.1)–(3.3) may possess several solutions; they usually do [9, 42]. Whether or not the fixed point iteration homes in on the correct solution then depends on the initialisation.

While there is no generic recipe for choosing an appropriate initialisation $\mu^{(0)}\in\mathscr{M}(G)$, two choices suggest themselves. First, we could initialise the messages according to the prior λ, i.e.,

$$\mu^{(0)}_{x_j\to a_i}(s) = \lambda^s(1-\lambda)^{1-s}. \tag{3.10}$$

Second, we could initialise the messages in accordance with the ground truth, i.e.,

$$\mu^{(0)}_{x_j\to a_i}(s) = \boldsymbol{\sigma}(x_j). \tag{3.11}$$

The latter is not practically useful for the obvious reason. But the analogy with other applications of Belief Propagation for inference problems suggests that *if* the fixed point iteration converges to the same solution to (3.1)–(3.3) from the two initialisations (3.10) and (3.11), then this solution actually is a good approximation to the correct messages. Furthermore, whether or not (3.10) and (3.11) yield the same solution we can try out experimentally.

One last but crucial point remains to be clarified: how precisely do we perform the fixed point iteration? The textbook method is to perform message updates in parallel. This means that, starting from the initialisation $(\mu^{(0)}_{x_j\to a_i})_{i,j}$, we compute all test-to-sample approximations $\mu^{(0)}_{a_i\to x_j}$ via (3.2)–(3.3). Then we use these together with (3.1) to compute the next approximation $(\mu^{(1)}_{x_j\to a_i}(\cdot))_{i,j}$ to all sample-to-test messages, and so forth.

This parallel updates mechanism was tried out experimentally in [14]. However, this method does not converge. Instead, the messages oscillate between odd and even rounds as shown in Figure 8. Similar oscillations emerge in other applications of Belief Propagation. They may result from an instability of the empirical mean of the messages. To elaborate, if in some particular iteration t the deviation from the prior

$$\sum_{j=1}^{n}\sum_{i=1}^{m}\mathbf{1}\{a_i\in\partial x_j\}(\mu^{(t)}_{x_j\to a_i}(1)-\lambda) \tag{3.12}$$

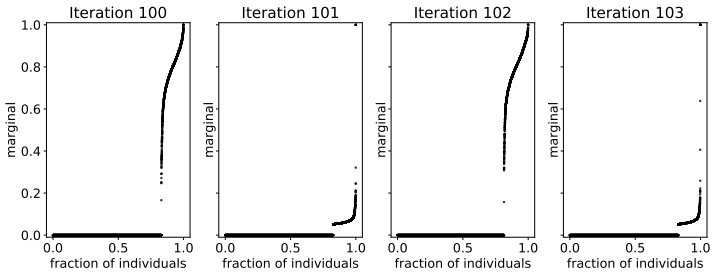

FIGURE 8. Illustration of the oscillatory behavior of Belief Propagation when performing parallel updates for $\lambda = 0.05$ and 0.2 tests/n for the noiseless setting

is positive, then we should expect a negative deviation in the next round. This is because due to (3.12) in the next iteration many tests will receive a relatively large indication from that one of their samples may be infected. The test will therefore send out "less urgent" messages to the other samples. Conversely, if (3.12) is negative, then in iteration $t+1$ we expect to see a positive deviation. Due to the analytic nature of the update rules (3.1)–(3.3) these oscillations do not dampen down but actually blow up, leading to oscillations between odd and even rounds. This observation led the authors of [14] to turn to the computationally more intensive Glauber dynamics algorithm.

But actually oscillations of this type have been observed in other problems as well and several ideas for tackling the problem are on the market. Perhaps the most organic solution, and the method to which we resort, is to update the messages in a randomised fashion rather than in parallel. Hence, starting from the initialisation $(\mu^{(0)}_{x_j \to a_i}(\cdot))_{i,j}$, we apply (3.2)–(3.3) once to initialise the test-to-sample messages $\mu_{a_i \to x_j}(\cdot)$ as well. Then at each time $t \geq 1$ we choose an edge $a_i x_j$ of G randomly and also flip a fair coin. If the coin comes up heads we update the message $\mu^{(t)}_{x_j \to a_i}(\cdot)$ according to (3.1). Otherwise we update $\mu^{(t)}_{a_i \to x_j}(\cdot)$ according to (3.2)–(3.3). The random choices break the cycle of oscillations. We stop the fixed point iteration after a fixed number T of steps. The precise choice of T is guided by experiments but of course T should be chosen large enough so that every message will likely get updated several times. We note that this update scheme does not impede practical matters from using our algorithm in a laboratory setting since it purely pertains to the computations behind the scene and does not impact how samples are split and combined.

Beyond relying on asymptotic ideas and comparing the messages that result from the two aforementioned initialisations we take two additional steps to corroborate the results of Belief Propagation. First, we compared the marginals obtained by Belief Propagation with the empirical marginals of Glauber dynamics on a number of samples. They match. Second, we compared the marginals obtained via Belief Propagation on moderately sized biregular test designs with the marginal distributions obtained via *population dynamics*, a heuristic intended to approximate the limiting distribution of the marginals as $n \to \infty$ [30]. Once again the Belief Propagation results align very well. Figure 9 displays the typical outcome of the Belief Propagation for different numbers of tests along with the estimate (3.9) of the remaining entropy.

What conclusions are to be drawn regarding a promising test design? We see three different scenarios.

- For small numbers m of tests we can extract some information from the negative tests. For instance, in the case $p = 1$ of perfect specificity we can rest assured that any sample included in a negative test is indeed uninfected. But beyond the direct effect of the negative tests the marginals do not align particularly well with the ground truth.
- The second scenario concerns intermediate values of m. Here Belief Propagation gains information from both positive and negative test results. As a consequence, the marginals start to align better with the ground truth.
- Finally, once m gets quite large the ground truth leaves a clear imprint on the test results. In this scenario we can recover the ground truth with good accuracy, albeit at the expense of investing many tests.

In light of what we learned on Belief Propagation, we now move on to describe the new adaptive Belief Propagation test design.

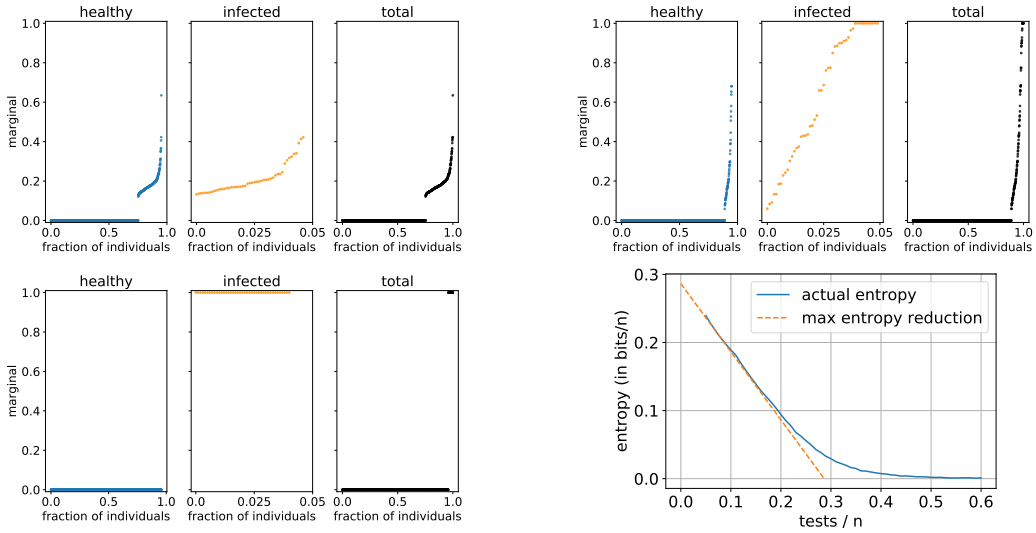

FIGURE 9. Illustration of the posterior distribution from running Belief Propagation on a random biregular test design for with 0.15 (top left), 0.25 (top right) and 0.6 (bottom left) tests/n and remaining entropy (bottom right) for $\lambda = 0.05$ and the noiseless setting.

3.2. **The first stage.** As the first stage we use the random biregular design $G = G_{n,m}(\Delta,\Gamma)$ with the optimal choice of Δ,Γ from (2.1) subject to rounding. Thus, the only free parameter is the total number m of tests conducted in the first stage. Its choice is informed by Belief Propagation.

Specifically, we choose the largest number m of tests up to which each test yields the optimal entropy reduction of $\ln 2$. Practically, this means that we choose m to match the point at which the entropy plot for the corresponding parameter values flattens. The fourth graphic in Figure 9 shows the approximation of the entropy as a function of the number of tests for our illustrative case of $n = 1000$ and $\lambda = 0.05$ in the noiseless setting. For other priors and noise levels, the story turns out to be analogous.

3.3. **The second and third stage.** Given the approximation of the marginals from the first stage, how should we proceed? As we saw in Section 2.3 and 2.5, two ideas for the subsequent stages proposed in the literature include individual testing of all samples whose marginals are not entirely polarised after the first round and informative Dorfman. The former strategy, known as Definite Defectives, seems wasteful as it completely disregards any non-trivial information about the marginals resulting from the Belief Propagation computation. The latter suffers from the same problem as the original Dorfman scheme, namely a potentially fairly large number of false positives and negatives.

To remedy these issues, we propose a new design that blends the random biregular design with the informative Dorfman scheme from Section 2.5. For a start we threshold marginals obtained from the first stage at 0.1% and 99.9%. Thus, we report samples with Belief Propagation marginals less than 0.1% as healthy and those with marginals beyond 99.9% as infected right after the first stage. The remaining samples are split into two groups, one comprising samples with marginals below 12.4% and one with marginals above. Let us refer to these as the *low risk* and the *high risk* groups, respectively. The choice of 12.4% marks precisely the threshold beyond which the expressions (2.1) suggest that any sample should be placed in one test only. Figure 10 provides an illustration.

For the low risk group we set up another random biregular test design on which we run Belief Propagation once again. The posterior of the first stage now acts as the prior of the second stage. A range of different tests tested for the biregular test design depending on the prior, the posteriors of the first stage and noise level and the final recommended test numbers obtained via optimisation over this range. The resulting marginals are again thresholded at 0.1% and 99.9%. Those samples whose marginals fall in between are subsequently retested individually with their classification being solely determined by the outcome of the individual test.

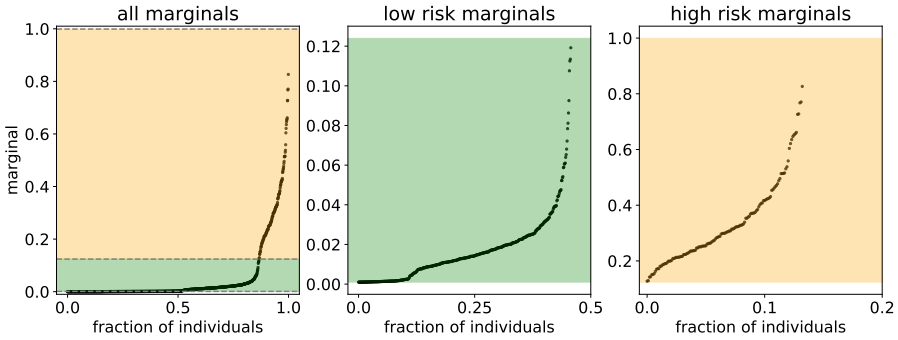

FIGURE 10. Illustration of split between low and high risk marginals for $\lambda = 0.05$ in the high noise setting with $m/n = 0.25$.

To be more precise, let $\mathscr{X}'$ be the samples in the low risk group, let $n' = |\mathscr{X}'|$ and let m' be the number of tests dedicated to this group. Thanks to the Belief Propagation results from the first stage we can (approximately) calculate the average marginal

$$\lambda' = \frac{1}{n'} \sum_{x \in \mathscr{X}'} \mu_G(\{\boldsymbol{\sigma}(x) = 1\}.$$

Mimicing (2.1) we then choose the degrees

$$\Delta' = \frac{m' \log(2)}{n' \lambda'} \qquad \Gamma' = \frac{\log(2)}{\lambda'} \tag{3.13}$$

subject to rounding and set up a random biregular test design $G' = G_{n',m'}(\Delta', \Gamma')$ on $\mathscr{X}'$. Furthermore, we modify the Belief Propagation equations on this random biregular design to accommodate the marginals computed in the first stage. Hence, instead of using the universal prior λ' for all the samples, we substitute the separate marginals computed in the first stage:

$$\mu_{G',x \to a}(s) \propto \mu_G(\{\hat{\boldsymbol{\sigma}}(x) = 1\})^s (1 - \mu_G(\{\hat{\boldsymbol{\sigma}}(x))^{1-s} \prod_{b \in \partial x \setminus \{a\}} \mu_{b \to x}(s). \tag{3.14}$$

The test-to-sample equations remain the same as in (3.2)–(3.3).

For the high risk group we set up an informative Dorfman design G'' as described in Section 2.5. If such a pooled test turns out to be negative, we classify all samples in this pool as healthy. Otherwise, we conduct individual tests and classify samples solely based on this individual test result.

3.4. **Enhanced accuracy.** The construction that we described up to this point is the one labelled adaptive Belief Propagation 1 in Section 1.3. Enhanced constructions adaptive Belief Propagation 2 and adaptive Belief Propagation 3 further reduce the number of false positives and negatives, at the expense of increasing the number of tests. Indeed, the adaptive Belief Propagation 1 construction facilitates such enhancements explicitly. This is because almost all false positives and negatives actually originate from the informative Dorfman procedure in the second stage, while neither the thresholding nor the second-stage random biregular design tend to produce a notable number of mistakes. Therefore, in adaptive Belief Propagation 2 and adaptive Belief Propagation 3 we simply perform the informative Dorfman procedure twice or thrice independently in parallel. Thus, in adaptive Belief Propagation 2 and adaptive Belief Propagation 3 we double or triple the number of tests required for the informative Dorfman bit of the construction, *but only for that bit.*

If we perform informative Dorfman twice (adaptive Belief Propagation 2), we need to choose whether to reduce false negatives or false positives. Accordingly, we classify a sample as healthy (infected) if both Dorfman procedures classify it as healthy (infected). In adaptive Belief Propagation 3 we get to avoid both false positives and false negatives. To this end we classify according to the majority vote of the three informative Dorfman schemes. Table 3.4 illustrates the number of tests to be performed in the first and second stage depending on the prior and noise level.

		noiseless		moderate noise		high noise	
algorithm	prior	m1/n	c	m1/n	c	m1/n	c
Belief Propagation + individual testing	0.5%	0.05	n/a	0.09	n/a	0.11	n/a
	1%	0.08	n/a	0.12	n/a	0.16	n/a
	5%	0.23	n/a	0.37	n/a	0.45	n/a
	10%	0.3	n/a	0.7	n/a	0.34	n/a
Belief Propagation + informative Dorfman	0.5%	0.045	n/a	0.05	n/a	0.045	n/a
	1%	0.075	n/a	0.075	n/a	0.1	n/a
	5%	0.28	n/a	0.24	n/a	0.16	n/a
	10%	0.125	n/a	0.1	n/a	0.1	n/a
adaptive Belief Propagation (1x)	0.5%	0.035	1.0	0.05	2.0	0.05	2.0
	1%	0.075	1.0	0.085	2.0	0.1	2.0
	5%	0.28	1.0	0.18	2.0	0.16	2.0
	10%	0.125	0.25	0.15	4.0	0.1	2.0
adaptive Belief Propagation (2x)	0.5%	n/a	n/a	0.075	8.0	0.02	8.0
	1%	n/a	n/a	0.12	8.0	0.03	8.0
	5%	n/a	n/a	0.4	2.0	0.36	2.0
	10%	n/a	n/a	0.5	2.0	0.325	2.0
adaptive Belief Propagation (3x)	0.5%	n/a	n/a	0.075	8.0	0.02	8.0
	1%	n/a	n/a	0.085	8.0	0.03	8.0
	5%	n/a	n/a	0.4	2.0	0.4	2.0
	10%	n/a	n/a	0.55	2.0	0.5	2.0

TABLE 1. Number of tests for the first and second stage found via optimization for various algorithms, priors and noise levels. The number of tests in the second stage in terms of the stated parameter c can be obtained as $c\lambda' n' \log(n')$ with λ' and n' defined as the average marginal and size of the low risk group, respectively.

4. ASYMPTOTIC CONSIDERATIONS

Clearly, adaptive Belief Propagation relies on heuristics and is not asymptotically optimal. This begs the question of how we would adapt the design and algorithm if we decide to live unburdened by practical considerations and consider the case $n \to \infty$?

4.1. **Variations on adaptive Belief Propagation.** The optimal drop in entropy seen in Figure 9 shows that running Belief Propagation on a random biregular test design in the first stage seems like a good idea. The discrete partition into three groups in the second stage, however, gives something away. Indeed, in the asymptotic regime infinitesimal intervals of posterior marginals contain an unbounded number of samples.[5] Thus, it seems information-theoretically optimal to construct a random biregular design for every single small marginal interval and repeat this procedure over a few stages. However, such an approach does not seem practical since for moderate n each random biregular design would only contain very few samples.

A simpler alternative that we considered is to still include all samples in one single second-stage test design, in which we choose the number of tests in which each sample takes part according to the posterior marginal from the first stage. Specifically, we chose these numbers so that in expectation half the tests should be positive. However, this design turned out to be unstable for small values of n because of random fluctuations.

[5]Of course, depending on the prior and the noise setting the distribution of the posterior marginals need not be supported on the entire unit interval.

4.2. **Plain Belief Propagation.** Thus far we disregarded what might seem at first glance the most straightforward scheme: just run Belief Propagation on a random biregular design and then simply threshold the marginals at, say, 50%. An obvious plus of this approach is that it requires one stage only. Indeed, when we simulated this scheme for large group testing instances such as $n = 10000$, this approach turned out to work extremely well. Particularly for small priors such as 0.5% and 1%, the plain Belief Propagation plus thresholding design is on par or even outperforms adaptive Belief Propagation in terms of both efficiency and reliability. However, for smaller values of n plain Belief Propagation plus threshold turns out to be extremely vulnerable to fluctuations of the number $\boldsymbol{k}$ of infected samples. This is because such fluctuations might cause the fraction of positive tests to significantly deviate from half.

4.3. **Scale effects.** All the simulation results were presented for group testing instances with $n = 1000$. However, we might be interested in smaller instances of say $n = 100$ or larger ones such as $n = 10000$. We performed extensive simulations in these directions and found that our results, particularly the power of the adaptive Belief Propagation scheme carry over to those group testing sizes as well, subject to rounding issues and few samples in a second stage for small instances necessitating slightly more tests.

4.4. **Population dynamics.** In the light of these scaling results for different instance sizes, let us spare a few more lines on the *population dynamics* already touched upon above. As mentioned above, this heuristic allows us to get a glimpse of the marginal distribution resulting from running Belief Propagation as $n \to \infty$ [30]. To this end, we require as input the distribution of infected and healthy samples in the local neighbourhood of a sample which is provided in [23]. Subsequently, we iteratively sample the local neighbourhood for infected and healthy samples and perform one-step Belief Propagation updates to model the marginal distribution of those samples whose marginal is not completely polarised. The resulting distribution which is shown in Figure 11 for a prior of 5% and the noiseless setting for illustration purposes closely resembles the marginal distribution that we observe from running Belief Propagation in our simulation in the first stage. As a side product, we obtain the proportion of polarised healthy and infected samples which lines up nicely with our simulation results. It should be noted that the population dynamics heuristic is nowhere near a complete analysis of Belief Propagation on random biregular graphs. Given the gains in efficiency and reliability that we observe in this empirical work for moderately-sized instances, a formal analysis of Belief Propagation seems to be an important next step in group testing research.

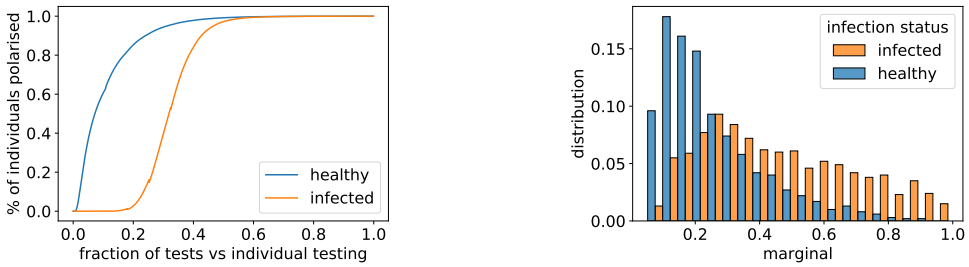

FIGURE 11. Illustration of the asymptotic fraction of samples with polarised marginals and the posterior distribution for non-polarised samples obtained by running population dynamics on the offspring distribution by [23] for $\lambda = 0.05$ in the noiseless setting

5. DISCUSSION

Group testing is a powerful method to efficiently and accurately detect infected samples. Since the mathematical work on group testing deals with the asymptotic $n \to \infty$ scenario, practical adoption of methods proposed in this literature has been limited. Instead practitioners tend to apply very simple test designs dating back to the 1940s. In this paper we therefore conducted an experimental study that shows how a mildly more sophisticated test design can significantly improve the accuracy of the overall test results by comparison to classical methods without asking for many more tests. The new test design comes with an efficient, easy-to-run and easy-to-implement algorithm that determines the status of each sample from the test results. Since the new design employs randomisation, its adoption is probably feasible only in a practical setting that employs a degree of automation in preparing test

pools. But on the plus side the new adaptive Belief Propagation design keeps the pool sizes and the number of pools that each sample has to be placed in fairly low.

Apart from the group testing model studied in the present paper, there are several other, more complicated models. For example, in *quantitative group testing* each test returns the *number* of infected samples rather than a binary positive or negative result. Further variants include the pooled data problem, the generalised coin weighing problem or the compressed sensing problem [17, 20].

What are the loose ends of the present work? On the one hand, it seems worthwhile to consider alternative noise models. A candidate might be one where the specificity decreases in the test size. Both the fixed noise model considered in this work and this diluted model have value from a practical perspective and it would be interesting to see whether our results carry over. On the other hand, the success of Belief Propagation in practical group testing leaves us wondering whether it is guaranteed to converge to a fixed point reminiscent of the ground truth. Hence, a mathematical analysis of Belief Propagation remains as an outstanding open problem.

ACKNOWLEDGEMENT

The authors thank Oliver Johnson for his helpful comments on group testing algorithms.

REFERENCES

[1] M. ALDRIDGE, *Individual testing is optimal for nonadaptive group testing in the linear regime*, IEEE Transactions on Information Theory, 65 (2019), pp. 2058–2061, https://doi.org/10.1109/TIT.2018.2873136.

[2] M. ALDRIDGE, *Conservative two-stage group testing*, 2020, https://arxiv.org/abs/2005.06617.

[3] M. ALDRIDGE, L. BALDASSINI, AND O. JOHNSON, *Group testing algorithms: Bounds and simulations*, IEEE Transactions on Information Theory, 60 (2014), pp. 3671–3687, https://doi.org/10.1109/TIT.2014.2314472.

[4] M. ALDRIDGE, O. JOHNSON, AND J. SCARLETT, *Group Testing: An Information Theory Perspective*, Foundations and Trends in Communications and Information Theory, 2019.

[5] V. BAPST AND A. COJA-OGHLAN, *Harnessing the bethe free energy*, Random Structures & Algorithms, 49 (2016), pp. 694–741, https://doi.org/10.1002/rsa.20692.

[6] J. BARBIER AND D. PANCHENKO, *Strong replica symmetry in high-dimensional optimal bayesian inference*, 2020, https://arxiv.org/abs/2005.03115.

[7] V. BRAULT, B. MALLEIN, AND J.-F. RUPPRECHT, *Group testing as a strategy for COVID-19 epidemiological monitoring and community surveillance*, PLOS Computational Biology, 17 (2021), p. e1008726, https://doi.org/10.1371/journal.pcbi.1008726.

[8] A. N. COHEN, B. KESSEL, AND M. G. MILGROOM, *Diagnosing SARS-CoV-2 infection: the danger of over-reliance on positive test results*, Cold Spring Harbor Laboratory, (2020), https://doi.org/10.1101/2020.04.26.20080911.

[9] A. COJA-OGHLAN, C. EFTHYMIOU, N. JAAFARI, M. KANG, AND T. KAPETANOPOULOS, *Charting the replica symmetric phase*, Communications in Mathematical Physics, 359 (2018), pp. 603–698, https://doi.org/10.1007/s00220-018-3096-x.

[10] A. COJA-OGHLAN, O. GEBHARD, M. HAHN-KLIMROTH, AND P. LOICK, *Information-theoretic and algorithmic thresholds for group testing*, IEEE Transactions on Information Theory, 66 (2020), pp. 7911–7928, https://doi.org/10.1109/TIT.2020.3023377.

[11] A. COJA-OGHLAN, O. GEBHARD, M. HAHN-KLIMROTH, AND P. LOICK, *Optimal group testing*, Combinatorics, Probability and Computing, (2021), pp. 1–38, https://doi.org/10.1017/s096354832100002x.

[12] A. COJA-OGHLAN AND W. PERKINS, *Belief propagation on replica symmetric random factor graph models*, Annales de l'Institut Henri Poincaré D, 5 (2018), pp. 211–249, https://doi.org/10.4171/aihpd/53.

[13] A. COJA-OGHLAN AND W. PERKINS, *Bethe states of random factor graphs*, Communications in Mathematical Physics, 366 (2019), pp. 173–201, https://doi.org/10.1007/s00220-019-03387-7.

[14] M. CUTURI, O. TEBOUL, Q. BERTHET, A. DOUCET, AND J.-P. VERT, *Noisy adaptive group testing using bayesian sequential experimental design*, 2020, https://arxiv.org/abs/2004.12508.

[15] D. DONOHO, *Compressed sensing*, IEEE Transactions on Information Theory, 52 (2006), pp. 1289–1306, https://doi.org/10.1109/TIT.2006.871582.

[16] D. L. DONOHO, A. JAVANMARD, AND A. MONTANARI, *Information-theoretically optimal compressed sensing via spatial coupling and approximate message passing*, IEEE Transactions on Information Theory, 59 (2013), pp. 7434–7464, https://doi.org/10.1109/TIT.2013.2274513.

[17] D. L. DONOHO, A. MALEKI, AND A. MONTANARI, *Message-passing algorithms for compressed sensing*, Proceedings of the National Academy of Sciences, 106 (2009), pp. 18914–18919, https://doi.org/10.1073/pnas.0909892106.

[18] R. DORFMAN, *The detection of defective members of large populations*, The Annals of Mathematical Statistics, 14 (1943), pp. 436–440, https://doi.org/10.1214/aoms/1177731363.

[19] C. EFTHYMIOU, T. P. HAYES, D. ŠTEFANKOVIČ, E. VIGODA, AND Y. YIN, *Convergence of MCMC and loopy BP in the tree uniqueness region for the hard-core model*, SIAM Journal on Computing, 48 (2019), pp. 581–643, https://doi.org/10.1137/17m1127144.

[20] A. EL ALAOUI, A. RAMDAS, F. KRZAKALA, L. ZDEBOROVÁ, AND M. I. JORDAN, *Decoding from pooled data: Phase transitions of message passing*, IEEE Transactions on Information Theory, 65 (2019), pp. 572–585, https://doi.org/10.1109/TIT.2018.2855698.

[21] A. FRIEZE AND M. KARONSKI, *Introduction to Random Graphs*, Cambridge University Press, 2015, https://doi.org/10.1017/cbo9781316339831.

[22] L. P. Garrison, J. B. Babigumira, A. Masaquel, B. C. Wang, D. Lalla, and M. Brammer, *The lifetime economic burden of inaccurate HER2 testing: Estimating the costs of false-positive and false-negative HER2 test results in US patients with early-stage breast cancer*, Value in Health, 18 (2015), pp. 541–546, https://doi.org/10.1016/j.jval.2015.01.012.
[23] O. Gebhard and P. Loick, *Note on the offspring distribution for group testing in the linear regime*, 2021, https://arxiv.org/abs/2103.13039.
[24] E. Joly and B. Mallein, *Group testing and pcr: a tale of charge value*, 2020, https://arxiv.org/abs/2012.09096.
[25] S. Kleinman, D. Strong, G. Tegtmeier, P. Holland, J. Gorlin, C. Cousins, R. Chiacchierini, and L. Pietrelli, *Hepatitis b virus (HBV) DNA screening of blood donations in minipools with the COBAS AmpliScreen HBV test*, Transfusion, 45 (2005), pp. 1247–1257, https://doi.org/10.1111/j.1537-2995.2005.00198.x.
[26] F. Krzakala, M. Mézard, F. Sausset, Y. F. Sun, and L. Zdeborová, *Statistical-physics-based reconstruction in compressed sensing*, Phys. Rev. X, 2 (2012), p. 021005, https://doi.org/10.1103/PhysRevX.2.021005.
[27] D. A. Levin and Y. Peres, *Markov Chains and Mixing Times (Second Edition)*, American Mathematical Society, 2017.
[28] S. Mallapaty, *The mathematical strategy that could transform coronavirus testing*, Nature, 583 (2020), pp. 504–505, https://doi.org/10.1038/d41586-020-02053-6.
[29] C. S. McMahan, J. M. Tebbs, and C. R. Bilder, *Informative dorfman screening*, Biometrics, 68 (2011), pp. 287–296, https://doi.org/10.1111/j.1541-0420.2011.01644.x.
[30] M. Mézard and A. Montanari, *Information, Physics, and Computation*, Oxford University Press, Inc., 2009, https://doi.org/10.5555/1592967.
[31] M. Müller, P. M. Derlet, C. Mudry, and G. Aeppli, *Testing of asymptomatic individuals for fast feedback-control of COVID-19 pandemic*, Physical Biology, 17 (2020), p. 065007, https://doi.org/10.1088/1478-3975/aba6d0.
[32] Y. Ohhashi, A. Pai, H. Halait, and R. Ziermann, *Analytical and clinical performance evaluation of the cobas TaqScreen MPX test for use on the cobas s 201 system*, Journal of Virological Methods, 165 (2010), pp. 246–253, https://doi.org/10.1016/j.jviromet.2010.02.004.
[33] J. Pearl, *Probabilistic Reasoning in Intelligent Systems*, Elsevier, 1988, https://doi.org/10.1016/c2009-0-27609-4.
[34] T. Richardson and R. Urbanke, *Modern Coding Theory*, Cambridge University Press, 2008, https://doi.org/10.1017/cbo9780511791338.
[35] N. Shental, S. Levy, V. Wuvshet, S. Skorniakov, B. Shalem, A. Ottolenghi, Y. Greenshpan, R. Steinberg, A. Edri, R. Gillis, M. Goldhirsh, K. Moscovici, S. Sachren, L. M. Friedman, L. Nesher, Y. Shemer-Avni, A. Porgador, and T. Hertz, *Efficient high-throughput SARS-CoV-2 testing to detect asymptomatic carriers*, Science Advances, 6 (2020), p. eabc5961, https://doi.org/10.1126/sciadv.abc5961.
[36] M. Sherlock, N. M. Zetola, and J. D. Klausner, *Routine detection of acute HIV infection through RNA pooling: Survey of current practice in the united states*, Sexually Transmitted Diseases, 34 (2007), pp. 314–316, https://doi.org/10.1097/01.olq.0000263262.00273.9c.
[37] J. M. Tebbs, C. S. McMahan, and C. R. Bilder, *Two-stage hierarchical group testing for multiple infections with application to the infertility prevention project*, Biometrics, 69 (2013), pp. 1064–1073, https://doi.org/10.1111/biom.12080.
[38] L. N. Theagarajan, *Group testing for covid-19: How to stop worrying and test more*, 2020, https://arxiv.org/abs/2004.06306.
[39] G. U. van Zyl, W. Preiser, S. Potschka, A. T. Lundershausen, R. Haubrich, and D. Smith, *Pooling strategies to reduce the cost of HIV-1 RNA load monitoring in a resource-limited setting*, Clinical Infectious Diseases, 52 (2010), pp. 264–270, https://doi.org/10.1093/cid/ciq084.
[40] P. O. Vontobel, *Counting in graph covers: A combinatorial characterization of the bethe entropy function*, IEEE Transactions on Information Theory, 59 (2013), pp. 6018–6048, https://doi.org/10.1109/TIT.2013.2264715.
[41] J. Watson, P. F. Whiting, and J. E. Brush, *Interpreting a covid-19 test result*, BMJ, (2020), p. m1808, https://doi.org/10.1136/bmj.m1808.
[42] L. Zdeborová and F. Krzakala, *Statistical physics of inference: thresholds and algorithms*, Advances in Physics, 65 (2016), pp. 453–552, https://doi.org/10.1080/00018732.2016.1211393.

Appendix A. Sample splits and test degree

The algorithms required the following number of maximum test degree and the following maximum and average split of samples. The algorithms can be readily adjusted to work with smaller test degrees or sample splits at the expense of slightly more tests.

		noiseless			moderate noise			high noise		
algorithm	prior	$\Gamma_{\max}$	$\Delta_{\max}$	Δ_{avg}	$\Gamma_{\max}$	$\Delta_{\max}$	Δ_{avg}	$\Gamma_{\max}$	$\Delta_{\max}$	Δ_{avg}
Belief Propagation + individual testing	0.5%	140	8	7.0	134	13	12.0	137	16	15.0
	1%	75	7	6.0	67	9	8.0	69	12	11.0
	5%	14	4	3.1	14	6	5.2	14	7	6.2
	10%	7	3	2.3	8	6	5.2	6	3	2.8
Belief Propagation + informative Dorfman	0.5%	134	8	6.0	140	9	7.1	134	8	6.2
	1%	67	7	5.1	67	7	5.2	70	9	7.2
	5%	15	6	4.1	20	5	3.5	13	4	2.9
	10%	8	3	1.8	14	3	2.3	10	3	2.3
adaptive Belief Propagation	0.5%	143	8	5.2	140	11	7.6	140	13	8.0
	1%	67	8	5.1	71	12	6.7	70	13	8.0
	5%	15	7	4.1	147	12	6.2	66	12	6.5
	10%	8	3	1.8	172	19	10.3	50	10	4.9

APPENDIX B. DISTRIBUTION BETWEEN STAGES

Based on the number of tests in the first and second stage, the following table shows the fraction of samples identified in each round. It evinces that despite a total of three stages needed for adaptive Belief Propagation the majority of samples are identified already in the first and second stage, depending on the prior and noise level.

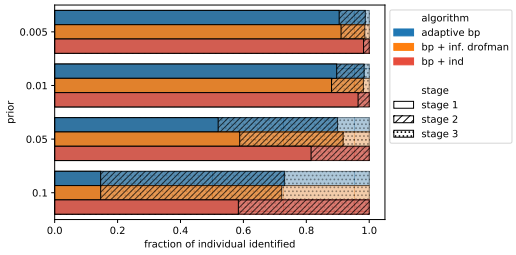

FIGURE 12. Fraction of samples identified in each stage by Belief Propagation followed by individual testing, Belief Propagation followed by informative Dorfman and adaptive Belief Propagation

AMIN COJA-OGHLAN, acoghlan@math.uni-frankfurt.de, GOETHE UNIVERSITY, MATHEMATICS INSTITUTE, 10 ROBERT MAYER ST, FRANKFURT 60325, GERMANY.

MAX HAHN-KLIMROTH, hahnklim@math.uni-frankfurt.de, GOETHE UNIVERSITY, MATHEMATICS INSTITUTE, 10 ROBERT MAYER ST, FRANKFURT 60325, GERMANY.

PHILIPP LOICK, loick@math.uni-frankfurt.de, GOETHE UNIVERSITY, MATHEMATICS INSTITUTE, 10 ROBERT MAYER ST, FRANKFURT 60325, GERMANY.

MANUEL PENSCHUCK, mpenschuck@ae.cs.uni-frankfurt.de, GOETHE UNIVERSITY, INSTITUTE FOR COMPUTER SCIENCE, 11–15 ROBERT MAYER ST, FRANKFURT 60325, GERMANY.

APPENDIX F. QUANTITATIVE GROUP TESTING IN THE SUBLINEAR REGIME

Quantitative Group Testing in the Sublinear Regime: Information-Theoretic and Algorithmic Bounds

Oliver Gebhard
Goethe University, Frankfurt, Germany

Max Hahn-Klimroth
Goethe University, Frankfurt, Germany

Dominik Kaaser
Universität Hamburg, Germany

Philipp Loick
Goethe University, Frankfurt, Germany

Abstract

The quantitative group testing (QGT) problem deals with efficiently identifying a small number of infected individuals among a large population. To this end, we can test groups of individuals where each test returns the total number of infected individuals in the tested group. For the regime where the number of infected individuals is sublinear in the population size we derive a sharp information-theoretic threshold for the minimum number of tests required to identify the infected individuals with high probability. Such a threshold was so far only known for the case where the infected individuals form a constant fraction of the population (Alaoui et al. 2014, Scarlett & Cevher 2017). Moreover, we propose and analyze a simple and efficient greedy reconstruction algorithm that matches the performance guarantees of much more involved constructions (Karimi et al. 2019).

2012 ACM Subject Classification Mathematics of computing → Information theory ; Mathematics of computing → Probabilistic inference problems

Keywords and phrases Information Theory, Quantitative Group Testing, Coin Weighing, Phase Transitions

Funding *Oliver Gebhard*: DFG CO 646/4
Max Hahn-Klimroth: Stiftung Polytechnische Gesellschaft and DFG FOR 2975
Philipp Loick: DFG CO 646/4

Acknowledgements The authors thank Uriel Feige for various detailed comments which improved the quality of the paper significantly. Furthermore, the authors thank Petra Berenbrink and Amin Coja-Oghlan for helpful discussions and important hints.

1 Introduction

We consider the problem of *quantitative group testing (QGT)* which is defined as follows. Suppose that k individuals out of a population of size n suffer from a rare disease. Then the goal is to identify the infected individuals with as few tests as possible. To this end, we are equipped with a testing procedure which allows us to test groups of individuals. In the *quantitative* variant we consider, each test outputs the exact number of infected individuals in the tested group (see Figure 1).

Group testing has its roots in the work of Dorfman [11], Erdős and Rényi [12], Djackov [9], and Shapiro [25]. More recently, QGT has gained a lot of interest in the literature [1, 5, 13, 18, 23], with applications in a multitude of disciplines such as DNA screening [24], identifying genetic carriers [6] and machine learning [20].

In this paper we study the *sublinear regime* $k \ll n$, where the number of infected individuals scales sublinearly in the population size. Furthermore, we restrict ourselves to

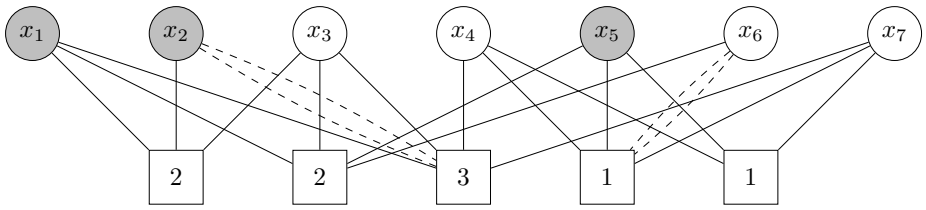

Figure 1 A small example of a QGT instance with the individuals $x_1, \ldots, x_7$ at the top and the tests $a_1, \ldots, a_5$ at the bottom. Infected individuals are colored in gray and test results are given in the test nodes. The dashed lines highlight the appearance of multi-edges.

non-adaptive test-designs, where all tests have to be conducted in parallel in a single step. This setting has recently gained a lot of attention in the closely related *binary* group testing problem [2, 7], in which the tests only output whether at least one individual is infected in the tested group. In this setting, we are interested in two types of *phase-transitions* that commonly arise in the analysis of inference problems:

- First, what is the minimum number of tests m_{inf} which allows us to infer the infected individuals from the test results?
- Second, how many tests m_{alg} are required such that an efficient algorithm can compute the infected individuals from given test results?

We will refer to the first phase-transition as the *information-theoretic threshold* m_{inf} and to the second phase-transition as the *algorithmic threshold* m_{alg}.

1.1 Related Work

Information-Theoretic Aspects. A simple information-theoretic lower bound can be obtained by a folklore counting argument: each test outputs a result in $\{0, \ldots, k\}$, thus a test design with m tests can produce at most $(k+1)^m$ different outcomes. This number must be larger than $\binom{n}{k}$ in order to distinguish all possible configurations of k infected individuals. Applying standard asymptotic bounds, we obtain

$$m^{\text{QGT}}_{\text{count}} \geq \frac{\ln \frac{n}{k}}{\ln k} k. \tag{1}$$

On the positive side, Bshouty [5] prove that inference of the set of infected individuals is efficiently possible with $(2+\varepsilon)m^{\text{QGT}}_{\text{count}}$ tests by using an adaptive procedure that runs in multiple rounds. If we restrict the analysis to non-adaptive designs, inference of the infected individuals in one round of testing is not possible by any design which uses less than

$$m^{\text{QGT}}_{\text{non-ada}} = 2\frac{\ln \frac{n}{k}}{\ln k} k \tag{2}$$

tests [9]. Grebinski and Kucherov [16] provide a non-adaptive design with an exponential-time decoding algorithm which guarantees inference with $(2+\varepsilon)m^{\text{QGT}}_{\text{non-ada}}$ tests using *separating matrices*. So far, these results hold independently of k. If we restrict ourselves to the linear regime where $k = \Theta(n)$, much stricter results are already known: Alaoui et al. [1] and Scarlett and Cevher [23] show that there is an exponential time construction that achieves inference with $(1+\varepsilon)m^{\text{QGT}}_{\text{non-ada}}$ tests.

Algorithmic Aspects. Bshouty [5] presents an efficient adaptive algorithm that succeeds at $m^{\text{QGT}}_{\text{non-ada}}$. However, for non-adaptive schemes, there are significant gaps of $\Theta(\ln n)$ between

the information-theoretic lower bound and the currently best known efficient algorithms [1, 10, 13, 14, 19, 21]. For instance, Alaoui et al. [1] present an *Approximate Message Passing* algorithm for the linear regime of QGT. Donoho and Tanner [10] present a decoding strategy based on ℓ_1-minimization, and Foucart and Rauhut [14] introduce the *Basis Pursuit*-algorithm which both solve the quantitative group testing problem with

$$2k\ln\frac{n}{k} \quad \text{and} \quad 2k\ln n \sim \frac{2}{1-\theta}k\ln\frac{n}{k}$$

tests, respectively, in the sublinear regime $k \ll n$. (Note that these algorithms solve the more general Compressed Sensing problem. There are various improvements over the Basis Pursuit algorithm known (e.g., the Orthogonal Matching Pursuit [21] and its improved version for discrete signals [26]) but as Wang and Yin [28] discuss, they do not perform asymptotically better in the quantitative group testing setting. The most recent algorithms explicitly designed for QGT in the sublinear regime are due to Karimi et al. [19]. The authors provide two algorithms based on graph codes that require at least

$$1.72k\ln\frac{n}{k} \quad \text{and} \quad 1.515k\ln\frac{n}{k}$$

tests, respectively [18, 19], again leaving a multiplicative $\ln n$-gap to the lower bound?. Furthermore, after the first version of this paper appeared on arXiv, Feige and Lellouche [13] introduced a relaxation of the QGT problem called *Subset Select problem*. They prove that, under mild assumptions, an algorithm succeeding at this relaxation can be easily turned into an algorithm for QGT without significantly increasing the required number of tests.

1.2 Our Contributions

We study the QGT problem under the random regular model $\boldsymbol{G}$ which is known to be information-theoretically optimal in the linear regime of QGT and in similar inference problems [7]. More precisely, we let $\boldsymbol{G} = (V \cup F, E)$ be a random bipartite multi-graph with factor nodes $F = \{a_1, \dots, a_m\}$ representing the tests, variable nodes $V = \{x_1, \dots, x_n\}$ representing the vertices and edges E indicating how often a specific individual takes part in a given test. Each test $a_i \in F$ chooses $\Gamma = n/2$ individuals uniformly at random with replacement. Hence, the number of tests per individual (with multi-edges counted multiple times) is a binomially distributed random variable with mean $\Delta = m/2$. Furthermore, the expected number of *distinct* tests to which an individual belongs is $\Delta^\star = (1 - \exp(-1/2))m$. Finally, we denote by $\boldsymbol{\sigma} \in \{0,1\}^n$ the infection status of the individuals and by $\boldsymbol{y} = \boldsymbol{y}(\boldsymbol{G}, \boldsymbol{\sigma})$ the vector of test results.

Information-Theoretic Results. We prove that in the sublinear regime where $k = n^\theta$ for some $\theta \in (0,1)$ it is possible to reconstruct $\boldsymbol{\sigma}$ from $(\boldsymbol{G}, \boldsymbol{y})$ with high probability if $(1+\varepsilon)m^{\mathrm{QGT}}_{\text{non-ada}}$ tests are conducted. More precisely, we show that there is, with high probability, no second configuration $\tau \in \{0,1\}^n$ leading to the same test results. Note that in an independent work, Feige and Lellouche [13] submitted an alternative proof of Theorem 1 based on the analysis of random matrices to arXiv.

▶ **Theorem 1.** *Suppose that* $0 < \theta < 1$, $k = n^\theta$, *and* $\varepsilon > 0$ *and let*

$$m_{\inf} = m_{\inf}(n,\theta) = 2\frac{k\ln(n/k)}{\ln k} = 2\frac{1-\theta}{\theta}k.$$

If $m > (1+\varepsilon)m_{\inf}(n,\theta)$, $\boldsymbol{\sigma}$ *can be computed from* $\boldsymbol{G}$ *and* $\boldsymbol{y}$ *w.h.p.*

Algorithmic Results. Beside this novel information-theoretic result, we also present a simple greedy algorithm called *Maximum Neighborhood (MN) Algorithm*. It follows a thresholding approach that is much simpler than the known algorithms by Karimi et al. [18, 19], which are technically highly challenging. A formal definition of the MN-Algorithm is given in Algorithm 1.

Algorithm 1 The Maximum Neighborhood Algorithm, a greedy algorithm for QGT

Input: $\boldsymbol{G}$, $\boldsymbol{y}$, k
Output: Estimation $\tilde{\boldsymbol{\sigma}}$ for $\boldsymbol{\sigma}$.
1 For every x_i for $i \in [n]$ calculate $\Psi_i = \sum_{j \in \partial^\star x_i} \boldsymbol{y}_j$;
2 Sort the individuals i in decreasing order by $\Psi_i - \boldsymbol{\Delta}_i^\star \frac{k}{2}$;
3 Declare the first k individuals as infected, declare the other individuals as uninfected;

On an intuitive level, the MN-Algorithm works as follows. First, we sum up the test results in the neighborhood of each individual, counting multi-edges only once. The sum is then centralized by its expected value. Finally, we declare those individuals with a large score as infected and the remaining individuals as uninfected. In our second main theorem, we analyze how many tests are required for the MN-Algorithm to recover the correct $\boldsymbol{\sigma}$ w.h.p.

▶ **Theorem 2.** *Suppose that* $0 < \theta < 1$, $k = n^\theta$, *and* $\varepsilon > 0$ *and let*

$$m_{MN}(n, \theta) = 4\left(1 - \frac{1}{\sqrt{\mathrm{e}}}\right)\frac{1+\sqrt{\theta}}{1-\sqrt{\theta}} k \ln(n/k).$$

If $m > (1+\varepsilon) m_{MN}(n, \theta)$, *then Algorithm 1 outputs* $\boldsymbol{\sigma}$ *w.h.p. on input* $\boldsymbol{G}$, $\boldsymbol{y}$, *and* k.

1.3 Discussion

Our results extend the results of Alaoui et al. [1] from the linear regime to the sublinear regime. For $\theta \to 1$, our threshold of Theorem 2 turns out to converge towards the threshold of [1]. The study of the sublinear regime is inspired by studies of the compressed sensing problem with a sparse underlying signal [3]. In the special case of QGT, it was initiated by [19]. It turns out that this regime is indeed interesting in real-world applications, as prominent examples show. For instance, Heaps law of epidemiology [4] models the early spread of pandemics in that way. It is not surprising that also other variants of group testing including binary group testing have been analyzed recently in the sublinear regime. By now, a vast body of related literature exists (see, e.g., the survey by Aldridge et al. [2]). Interestingly, for the binary (presumably more difficult) variant an efficient algorithm is known which requires $m_{GT} \sim \ln^{-1}(2) k \ln \frac{n}{k}$ tests for $\theta \leq \ln 2/(1+\ln 2) \approx 0.409$ [7]. Thus, dropping most of the available information and using this binary algorithm outperforms not only the simple greedy approach discussed in this paper, but also the quite involved algorithms by Karimi et al. [18, 19] if θ is small enough.

As in state-of-the art designs for binary group testing [2, 7], we allow individuals to participate in the same test multiple times. While this seems counterintuitive in the first place, it does not affect practicability of the proposed design. Finally, note that all known algorithms do not achieve the order of the information-theoretic bound. An exciting avenue for future research is to investigate whether other algorithms can be order-optimal or even achieve the information-theoretic bound. However, it might also be the case that QGT exhibits a similar impossible-hard-easy transition presumed for other statistical inference problems, where the best known efficient algorithms do not attain the information-theoretic bounds [17].

2 Analysis

Model and Notation. We consider a set of n individuals out of which k are infected, where $k = n^\theta$ for some $\theta \in (0,1)$. We use $\boldsymbol{G} = \boldsymbol{G}(n, m, \Delta)$ to denote the random bipartite multi-graph that models the test design, where m denotes the total number of tests and $\boldsymbol{\Delta} = \{\boldsymbol{\Delta}_1, \dots, \boldsymbol{\Delta}_n\}$ describes the number of tests in which each individual participates. Observe that $\boldsymbol{\Delta}_i \sim \mathrm{Bin}(mn/2, 1/n)$. Similarly, we let $\boldsymbol{\Delta}^\star = \{\boldsymbol{\Delta}^\star_1, \dots, \boldsymbol{\Delta}^\star_n\}$ denote the number of *distinct* tests with expected value $\mathbb{E}[\boldsymbol{\Delta}^\star_i] = (1 - \exp(-1/2))m$. In the following analysis, all asymptotic notation refers to the limit $n \to \infty$.

The vector $\boldsymbol{\sigma} \in \{0,1\}^n$ encodes which individuals are infected by assigning the value 1 to infected individuals and the value 0 to uninfected individuals. The vector $\boldsymbol{y} \in \{0, \dots, \Gamma\}^m$ denotes the test results. When we refer to any other configuration than $\boldsymbol{\sigma}$, we simply write σ for the configuration and $y = y(\boldsymbol{G}, \sigma)$ for the corresponding test result vector. Additionally, we write $V = \{x_1, \dots, x_n\}$ for the set of all individuals and $V_0 = \{x_i \in V : \boldsymbol{\sigma}_{x_i} = 0\}$ and $V_1 = V \setminus V_0$ for the set of uninfected and infected individuals, respectively. For an individual $x_i \in V$ we write ∂x_i for the multiset of tests a_j adjacent to x_i. Similarly, we write $\partial^\star x_i$ for the set of *distinct* tests a_j adjacent to x_i. Analogously, for a test a_i we write ∂a_i for the multiset of individuals that take part in test a_i.

Recall that in our model every test has size exactly $\Gamma = n/2$, and individuals are assigned uniformly at random with replacement. In the presence of multi-edges, one individual may appear more than once in ∂a_i. If an infected individual x_i participates in a test a_j more than once, it increases $\boldsymbol{y}_j$ multiple times. For each $x_i \in V$, we let Ψ_i be the sum of test results for *distinct* tests adjacent to x_i. That is, even if an individual appears more than once in a test and thus contributes to the test result multiple times, this test contributes to Ψ_i only once. The infection status of x_i has a significant impact on this sum, increasing it by $\boldsymbol{\Delta}_i$, if individual x_i is infected. To account for this effect in our analysis, we introduce a second variable Φ_i that sums the adjacent test results and excludes the impact of the status of individual x_i. Formally, for any configuration $\sigma \in \{0,1\}^n$ we define

$$\Psi_i(\sigma) = \sum_{j \in \partial^\star x_i} y_{a_j} \quad \text{and} \quad \Phi_i(\sigma) = \Psi_i(\sigma) - \mathbf{1}\{\sigma(i) = 1\}\Delta_i$$

and let $\Psi = (\Psi_1, \dots, \Psi_n)$ and $\Phi = (\Phi_1, \dots, \Phi_n)$. When we consider a specific instance $(\boldsymbol{G}, \boldsymbol{y})$, we will write $\boldsymbol{\Psi}_i = \Psi_i(\boldsymbol{\sigma})$ and $\boldsymbol{\Phi}_i = \Phi_i(\boldsymbol{\sigma})$ for the sake of brevity. Notably, while $\boldsymbol{\Psi}_i$ is known to the observer or an algorithm instantly from the test results, $\boldsymbol{\Phi}_i$ is not, since the individuals' infection status is unknown.

We let $c(n) > 0$ denote a positive function from the natural numbers to $\mathbb{R}^+$ such that

$$m = c(n)k\frac{\ln(n/k)}{\ln k}.$$

While we will assume that $c(n) = \Theta(1)$ for the information-theoretic bound, we will see that the algorithmic bound requires $c(n)$ to scale as $\Theta(\ln n)$.

Define $\mathcal{R}$ as the event that, for all $i \in [n]$, we have

$$\boldsymbol{\Delta}_i = \frac{m}{2} + O\left(\sqrt{m}\ln n\right) \quad \text{and} \quad \boldsymbol{\Delta^\star}_i = (1 - \exp(-1/2))\, m + O\left(\sqrt{m}\ln n\right).$$

Due to standard concentration results, $\mathcal{R}$ is a high probability event. A proof of the statement can be found in Appendix A.2.

▶ **Lemma 3.** *For the random experiment leading to graph $\boldsymbol{G}$, we find $\mathbb{P}(\mathcal{R}) = 1 - o(1)$.*

As Theorems 1 and 2 only contain w.h.p.-assertions, we can safely condition on $\mathcal{R}$.

2.1 The Teacher-Student Model

As in many related inference problems [8] the teacher-student model provides the fundamental means towards analyzing information-theoretic questions in random constraint satisfaction problems and will also be employed in the present paper. Specifically, the challenge in random constraint satisfaction problems lies in deriving probability distributions that are dependent on a variety of random variables and per-se hard to express. However, deriving probability distributions conditioned on certain high-probability events is feasible. For an excellent introduction and mathematical justification of the model, we refer to the reader to [8]. The setup is the following: a teacher aims to convey some *ground truth* to the student. Rather than directly providing the *ground truth* to the student, the teacher generates observable data from the *ground truth* via some statistical model and passes both information to the student. The student now aims to infer the ground truth from the observed data and the model.

In terms of this paper, we see $\boldsymbol{\sigma}$ as the generated ground truth. Its distribution is inherited from all vectors in $\{0,1\}^n$ of Hamming weight k. The observable data $\boldsymbol{y}$, together with the used pooling scheme $\boldsymbol{G}$ are passed to the student in order to infer $\boldsymbol{\sigma}$. In the following, we analyze the chances of the student to infer the ground truth from the observable data. First, we derive the model distribution from the provided pair $(\boldsymbol{G}, \boldsymbol{y})$. Afterwards, we use the gained knowledge to analyze the chances of the student to recover the ground-truth by estimating the number of solutions that fulfill the necessary prerequisites. As our goal is to recover $\boldsymbol{\sigma}$ w.h.p., we condition on the event that the underlying bipartite multi-graph behaves almost *as expected*. We exploit the knowledge about the pooling scheme to derive high-probability events which we can condition on. Eventually, our analysis outputs the information whether there is a unique or multiple solutions the student might guess.

2.2 Information-Theoretic Achievability

In order to proof Theorem 1 we count alternative configurations yielding the same test results as the true configuration. This approach rests on techniques that are regularly employed for random constraint satisfaction problems [8]. To this end, let $S_k(\boldsymbol{G}, \boldsymbol{y})$ be the set of all vectors $\sigma \in \{0,1\}^n$ of Hamming weight k such that

$$\boldsymbol{y}_{a_i} = |\{x_j \in \partial a_i : \sigma(x_j) = 1\}| \quad \text{for all } i \in [m].$$

In words, $S_k(\boldsymbol{G}, \boldsymbol{y})$ contains the set of all vectors σ with k ones that label the individuals infected and uninfected in a way consistent with the test results. Let $Z_k(\boldsymbol{G}, \boldsymbol{y}) = |S_k(\boldsymbol{G}, \boldsymbol{y})|$. We need to prove that $Z_k(\boldsymbol{G}, \boldsymbol{y}) = 1$ w.h.p. as soon as the number of tests m exceeds $m^{\mathrm{QGT}}_{\text{non-ada}}$. In this case, we can find $\boldsymbol{\sigma}$ via exhaustive search. It turns out that it is much more convenient to study $Z_{k,\ell}(\boldsymbol{G}, \boldsymbol{y})$, the number of alternative configurations that are consistent with the test results and have an *overlap* of ℓ with $\boldsymbol{\sigma}$. The overlap signifies the number of infected individuals under $\boldsymbol{\sigma}$ that are also infected under the alternative configuration. Formally, we define

$$Z_{k,\ell}(\boldsymbol{G}, \boldsymbol{y}) = |\{\sigma \in S_k(\boldsymbol{G}, \boldsymbol{y}) : \sigma \neq \boldsymbol{\sigma}, \langle \boldsymbol{\sigma}, \sigma \rangle = \ell\}| \,.$$

Thus it suffices to prove that, for $m > m_{\mathrm{inf}}$, w.h.p., $\sum_{\ell=0}^{k-1} Z_{k,\ell}(\boldsymbol{G}, \boldsymbol{y}) = 0$. To this end, two separate arguments are needed. First, we show via a first moment argument that no second satisfying configuration can exist with a small overlap with $\boldsymbol{\sigma}$. Second, we employ the classical coupon collector argument to show that a second satisfying configuration cannot exist for large overlaps, i.e., one individual flipped from uninfected under $\boldsymbol{\sigma}$ to infected under

an alternative configuration initiates a cascade of other changes in infection status to correct for this initial change. The following two propositions rule out configurations with a small and large overlap, respectively.

▶ **Proposition 4.** *Let $\varepsilon > 0$, $0 < \theta < 1$ and assume that $m > (1+\varepsilon)m_{\inf}(n,k,\theta)$. W.h.p., we have*

$$\sum_{\ell=0}^{k(1-\exp(-1/2))} Z_{k,\ell}(\boldsymbol{G},\boldsymbol{y}) = 0.$$

By Markov's inequality it clearly suffices to show that $\mathbb{E}[Z_{k,\ell}(\boldsymbol{G},\boldsymbol{y})] \to 0$ fast enough for all $0 \leq \ell < k - (1-\exp(-1/2))\ln k$ if $m \geq (1+\varepsilon)m_{\inf}$ for some $\varepsilon > 0$. A rigorous proof and the proofs of the following lemmas can be found in Appendix B. It turns out that $\mathbb{E}[Z_{k,\ell}(\boldsymbol{G},\boldsymbol{y})]$ satisfies

$$\mathbb{E}[Z_{k,\ell}(\boldsymbol{G},\boldsymbol{y})] \leq \binom{k}{\ell}\binom{n-k}{k-\ell}\prod_{i=1}^{m}\sum_{j=1}^{\boldsymbol{y}_{a_i}}\binom{\Gamma}{j,j,\Gamma-2j}\left((1-\ell/k)\frac{k}{n}\right)^{2j}\left(1-2(1-\ell/k)\frac{k}{n}\right)^{\Gamma-2j}.$$

The combinatorial meaning is immediate. The binomial coefficients count the number of configurations of overlap ℓ with $\boldsymbol{\sigma}$. The subsequent term measures the probability that a specific configuration σ yields the same test result vector as $\boldsymbol{\sigma}$. To this end, we divide individuals into three categories. The first contains those individuals exhibiting the same status under σ and $\boldsymbol{\sigma}$, while the second and third category feature those individuals that are infected under σ and uninfected under $\boldsymbol{\sigma}$ and vice versa. The probability for an individual to be in the second or third category is $(1-\ell/k)k/n$ each, while the probability in the first category is $1-2(1-\ell/k)k/n$. The key observation is that a test result is the same between σ and $\boldsymbol{\sigma}$, if the number of individuals in the second category is identical to the number in the third category. We compute the sum over the number of individuals which are flipped. Since the probability term allows for an individual included in a test multiple times to be both infected and uninfected, the expression considers cases that do not occur in the model and therefore is an upper bound to $\mathbb{E}[Z_{k,\ell}(\boldsymbol{G},\boldsymbol{y})]$. Simplifying the term and conditioning on $\mathcal{R}$ yields the first lemma. Let $\mathrm{Bin}_{\geq i}(n,p)$ be the binomial distribution with parameters n and p conditioned on being at least i.

▶ **Lemma 5.** *For every $0 \leq \ell \leq k - (1-\exp(-1/2))\ln k$ and a random variable $\boldsymbol{X} \sim \mathrm{Bin}_{\geq 1}(\Gamma, 2(1-\ell/k)k/n)$, we have*

$$\mathbb{E}[Z_{k,\ell}] \leq (1+O(1))\mathbb{E}[Z_{k,\ell}(\boldsymbol{G},\boldsymbol{y}) \mid \mathcal{R}] \leq (1+O(1))\binom{k}{\ell}\binom{n-k}{k-\ell}\left(\frac{1}{\sqrt{2\pi}}\mathbb{E}\left[\frac{1}{\sqrt{\boldsymbol{X}}}\right]\right)^{m}.$$

Using standard asymptotics, we are able to simplify this expression.

▶ **Lemma 6.** *For every $0 \leq \ell \leq k - (1-\exp(-1/2))\ln k$ and $n \to \infty$, we have*

$$\frac{1}{n}\ln\left(\mathbb{E}[Z_{k,\ell}(\boldsymbol{G},\boldsymbol{y}) \mid \mathcal{R}]\right)$$
$$\leq (1+o(1))\left(\frac{k}{n}H\left(\frac{\ell}{k}\right) + \left(1-\frac{k}{n}\right)H\left(\frac{k-\ell}{n-k}\right) - \frac{ck/n\ln(n/k)}{2\ln k}\ln\left(2\pi\left(1-\frac{\ell}{k}\right)k\right)\right).$$

The key is to choose c such that $Z_{k,\ell}(\boldsymbol{G},\boldsymbol{y}) \to 0$ for every $\ell \leq k - (1-\exp(-1/2))\ln k$ and $n \to \infty$. We find that, asymptotically, $\ln \mathbb{E}[Z_{k,\ell}(\boldsymbol{G},\boldsymbol{y})]/n$ takes its maximum at $\ell = \Theta\left(k^2/n\right)$. Therefore, the r.h.s. of (6) becomes negative, if and only if the number of tests m, parametrized by c, is larger than $m_{\inf}(\theta,k)$. This is formalized in the following lemma, concluding the proof of Proposition 4.

▶ **Lemma 7.** *For every* $0 \le \ell \le k - (1 - \exp(-1/2)) \ln k$, $0 < \theta < 1$ *and* $\varepsilon > 0$ *the following holds. If* $m \ge (1+\varepsilon) m_{\inf}(\theta, k)$ *then*

$$\frac{1}{n} \ln \mathbb{E}[Z_{k,\ell}(\boldsymbol{G}, \boldsymbol{y}) \mid \mathcal{R}] < 0.$$

Proof of Proposition 4. The proposition is a direct consequence of Lemmas 5–7 and Markov's inequality. ◀

While we could already establish that there are no feasible configurations that have a small overlap with the true configuration $\boldsymbol{\sigma}$, we still need to ensure that there are no feasible configurations that are close to $\boldsymbol{\sigma}$. Indeed, we can exclude configurations with a large overlap with the next proposition.

▶ **Proposition 8.** *Let* $\varepsilon > 0$ *and* $0 < \theta \le 1$ *and assume that* $m > (1+\varepsilon) m_{\inf}(k, \theta)$. *Given* $\mathcal{R}$ *we have* $Z_{k,\ell}(\boldsymbol{G}, \boldsymbol{y}) = 0$ *for all* $k - (1 - \exp(-1/2)) \ln k < \ell < k$ *w.h.p.*

Appendix B.4 is devoted to prove Proposition 8 rigorously. The proof, while fundamentally easy as it follows the classical coupon collector argument, needs some technical attention. If we consider a configuration σ different from $\boldsymbol{\sigma}$ with the same Hamming weight k, at least one individual that is infected under $\boldsymbol{\sigma}$, is labeled uninfected under an alternative configuration σ. Given $\mathcal{R}$, this individual participates in at least $\boldsymbol{\Delta}_i^\star > m/4$ different tests, whose results all change by at least -1, depending on how often the individual participates. To compensate for these changes, we need to find individuals $x_1 \dots x_l$ that are uninfected under $\boldsymbol{\sigma}$ and infected under σ and such that their joint neighborhood is a super-set of the changed tests. The balls-into-bins experiment [22] shows that the size of this super-set is of order at least

$$(1 - \exp(-1/2)) \ln m \ge (1 - \exp(-1/2)) \ln k.$$

Now Theorem 1 follows directly.

Proof of Theorem 1. The theorem is a direct consequence of Propositions 4 and 8. ◀

2.3 Greedy Algorithm

We now sketch how to prove Theorem 2. Subsequently, we present simulations that analyze the performance of the algorithm empirically.

Performance Guarantees. Recall that Ψ_i is the sum over all test results in the neighborhood of individual x_i (multi-edges counted only once) and $\boldsymbol{\Delta}_i^\star$ is the (random) number of disjoint tests x_i belongs to. Finally, let $\mathcal{E}_j$ be the $\sigma-$algebra generated by the edges connected with x_j. As already discussed, we find that

$$\boldsymbol{\Delta}_i^\star = (1 + o(1)) \left(1 - \exp\left(-1/2\right)\right) m$$

with high probability. Intuitively, an infected individual x_i increases the value of $\boldsymbol{\Psi}_i$ by $\boldsymbol{\Delta}_i = (1 + o(1))m/2$ while this is not true for uninfected individuals. Furthermore, by construction of the random graph, we find that the second neighborhood of x_i contains $\mathrm{Bin}\left(\Gamma \boldsymbol{\Delta}_i^\star, k/n\right)$ infected individuals, thus we expect

$$\mathbb{E}\left[\Psi_i - \boldsymbol{\Delta}_i^\star \frac{k}{2} \middle| \mathcal{E}_i\right] = \mathbf{1}\{\boldsymbol{\sigma}(i) = 1\} \boldsymbol{\Delta}_i,$$

thus the *scores* $\Psi_i - \boldsymbol{\Delta}_i^\star \frac{k}{2}$ differ between infected and uninfected individuals. The whole proof of the algorithmic performance boils down to identify a threshold value $T(\alpha)$ such that, if sufficiently many tests are conducted, all scores of uninfected individuals are below $T(\alpha)$ while the scores of all infected individuals exceed this threshold w.h.p. If we conduct $m = dk \ln \frac{n}{k}$ tests, we find by a standard application of a Chernoff bound and a union bound over all infected and, respectively, uninfected individuals that $T(\alpha)$ is a valid threshold whenever

$$\frac{-(1-\theta)\alpha^2 d}{4\left(1-\exp\left(-1/2\right)\right)\left(1+o(1)\right)} + \theta < 0 \quad \text{and} \quad \frac{-(1-\theta)(1-\alpha)^2 d}{4\left(1-\exp\left(-1/2\right)\right)\left(1+o(1)\right)} + 1 < 0. \quad (3)$$

Optimizing this expression with respect to α and plugging d into $m = dk\ln(n/k)$ yields, for any $\varepsilon > 0$, the sufficient condition

$$m \geq (4+\varepsilon)(1+o(1))\left(1-\exp\left(-1/2\right)\right)\frac{1+\sqrt{\theta}}{1-\sqrt{\theta}} k \ln(n/k).$$

A formal derivation of those statements can be found in Appendix C.

Empirical Analysis. In this section we present simulation results for the MN-Algorithm (Algorithm 1). In Figure 2 we plot the number of tests required to reconstruct $\boldsymbol{\sigma}$ for $n \in [10^2, 10^6]$ and different values of θ. The dotted lines show our theoretical bounds. Note that the discontinuities in the theoretical bound stem from rounding the number of infected individuals to the closest integer. We remark our simulation results align well with the theoretical predictions for larger values of n. For smaller values of n, our theoretical results are too pessimistic: the lower-order term hidden in the $o(1)$ in Equation (3) scales as $\Theta\left(\frac{\sqrt{\ln n}}{k}\right)$, and while this expression decreases polynomially fast in n, it is far from vanishing for small values of n and θ.

In Figures 3 and 4 we analyze the success probability for exact reconstruction of $\boldsymbol{\sigma}$ and the number of correctly identified infected individuals. For different numbers of tests we conducted 100 independent simulation runs for population sizes of $n = 10^3$ and $n = 10^4$ and different values of θ. The dashed lines show the phase transition predicted by Theorem 2. The data in Figure 4 indicate that all but a small fraction of infected individuals are correctly detected, even if the exact reconstruction of $\boldsymbol{\sigma}$ is still quite unlikely according to Figure 3. Overall, the implementation hints at the practical usability of the MN-Algorithm, even for small population sizes.

▶ Remark 9. The formal proof of the algorithmic bound directly gives an insight about the convergence speed and thus about the expected performance of the MN-Algorithm for finite n: we can compute that the MN-Algorithm requires an additional multiplicative factor of at least

$$\left(1 + \frac{\sqrt{2}\ln n}{\sqrt{4\left(1-\exp\left(-1/2\right)\right)mk}}\right)$$

tests in addition to the asymptotic analysis for $n \to \infty$. This explains the (slight) deviation of the theoretical and the empirical results for small values of n. See the proof of Corollary 19 in Appendix C for the rigorous analysis.

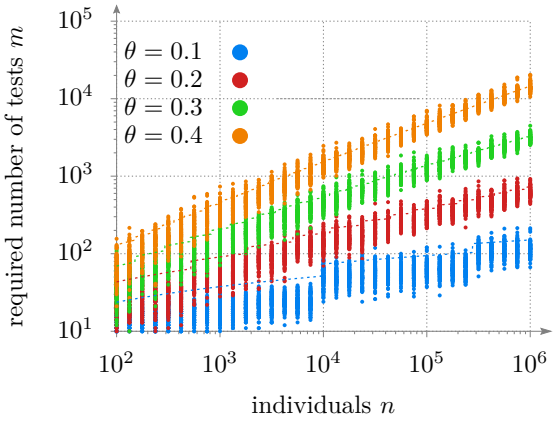

Figure 2 The required number of tests until $\boldsymbol{\sigma}$ can be exactly reconstructed for different population sizes and θ regimes. For each value of n, 100 simulations were carried out independently.

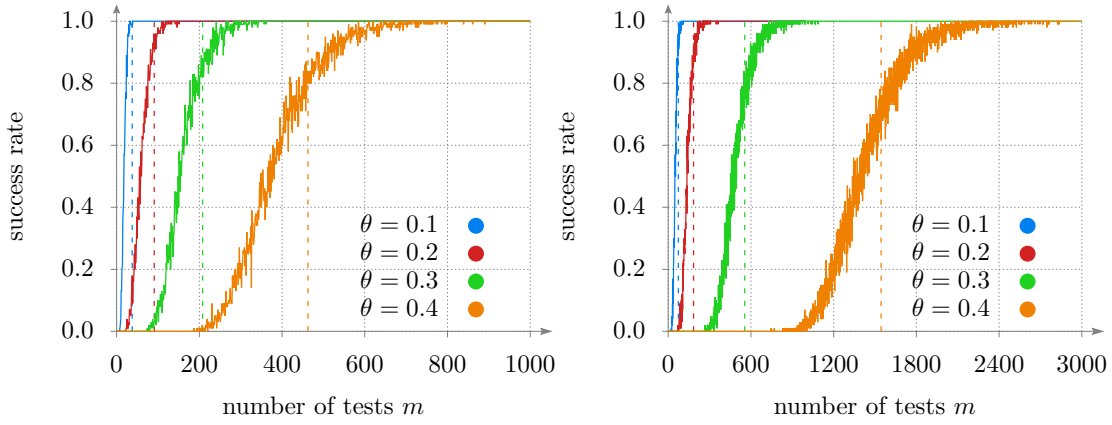

(a) $n = 1000$ **(b)** $n = 10000$

Figure 3 The plot shows the rate of successful recovery of $\boldsymbol{\sigma}$ among 100 independent simulation runs over the number of tests m for different values of θ and $n = 10^3$ (left) and $n = 10^4$ (right).

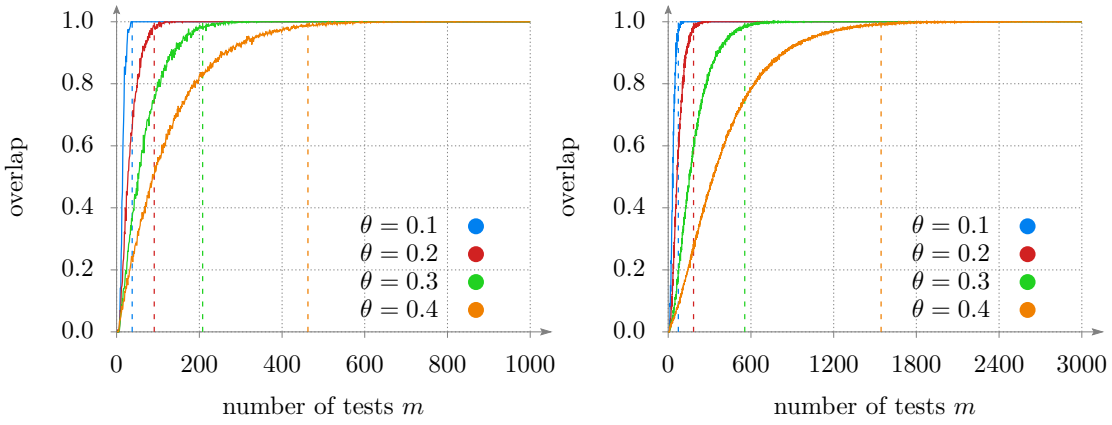

(a) $n = 1000$ **(b)** $n = 10000$

Figure 4 The plots show the *overlap* – the fraction of correctly classified infected individuals – among 100 independent simulation runs over the numbers of tests m for different values of θ and $n = 10^3$ (left) and $n = 10^4$ (right).

References

1 A. Alaoui, A. Ramdas, F. Krzakala, L. Zdeborová, and M. I. Jordan. Decoding from pooled data: Phase transitions of message passing. *IEEE Trans. Information Theory*, 65(1):572–585, 2019.

2 M. Aldridge, O. Johnson, and J. Scarlett. Group testing: An information theory perspective. *Foundations and Trends in Communications and Information Theory*, 15(3–4):196–392, 2019.

3 Y. Arjoune, N. Kaabouch, H. El Ghazi, and A. Tamtaoui. Compressive sensing: Performance comparison of sparse recovery algorithms. *IEEE 7th Annual Computing and Communication Workshop and Conference (CCWC'17)*, 2017.

4 R. W. Benz, S. J. Swamidass, and P. Baldi. Discovery of power-laws in chemical space. *Journal of Chemical Information and Modeling*, 48(6):1138–1151, 2008.

5 N. H. Bshouty. Optimal algorithms for the coin weighing problem with a spring scale. *The 22nd Conference on Learning Theory (COLT'09)*, 2009.

6 C. C. Cao, C. Li, and X. Sun. Quantitative group testing-based overlapping pool sequencing to identify rare variant carriers. *BMC Bioinformatics*, 15:195, 2014.

7 A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimroth, and P. Loick. Optimal group testing. *Combinatorics, Probability and Computing*, page 1–38, 2021.

8 A. Coja-Oghlan, F. Krzakala, W. Perkins, and L. Zdeborová. Information-theoretic thresholds from the cavity method. *Advances in Mathematics*, 333:694–795, 2018.

9 A. G. Djackov. On a search model of false coins. In *Topics in Information Theory. Hungarian Acad. Sci*, volume 16, pages 163–170, 1975.

10 D.L. Donoho and J. Tanner. Thresholds for the recovery of sparse solutions via l1 minimization. In *40th Annual Conference on Information Sciences and Systems*, pages 202–206, 2006.

11 R. Dorfman. The detection of defective members of large populations. *The Annals of Mathematical Statistics*, 14(4):436–440, 1943.

12 P. Erdős and A. Rényi. On two problems of information theory. *Magyar Tud. Akad. Mat. Kutató Int. Közl*, 8:229–243, 1963.

13 U. Feige and A. Lellouche. Quantitative group testing and the rank of random matrices, 2020. `arXiv:2006.09074`.

14 S. Foucart and H. Rauhut. *An Invitation to Compressive Sensing*, pages 1–39. Springer New York, New York, NY, 2013.

15 X. Gao, M. Sitharam, and A. E. Roitberg. Bounds on the jensen gap, and implications for mean-concentrated distributions. *The Australian Journal of Mathematical Analysis and Applications*, 16(16):1–16, 2019.

16 V. Grebinski and G. Kucherov. Optimal reconstruction of graphs under the additive model. *Algorithmica*, 28(1):104–124, 2000.

17 M. Jerrum. Large cliques elude the metropolis process. *Random Struct. Algorithms*, 3(4):347–360, 1992.

18 E. Karimi, F. Kazemi, A. Heidarzadeh, K. R. Narayanan, and A. Sprintson. Non-adaptive quantitative group testing using irregular sparse graph codes. *2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton) IEEE*, pages 608–614, 2019.

19 E. Karimi, F. Kazemi, A. Heidarzadeh, K. R. Narayanan, and A. Sprintson. Sparse graph codes for non-adaptive quantitative group testing. In *2019 IEEE Information Theory Workshop (ITW)*, pages 1–5, 2019.

20 J. P. Martins, R. Santos, and R. Sousa. Testing the maximum by the mean in quantitative group tests. In *New Advances in Statistical Modeling and Applications*, pages 55–63. Springer, 2014.

21 Y. C. Pati, R. Rezaiifar, and P. S. Krishnaprasad. Orthogonal matching pursuit: recursive function approximation with applications to wavelet decomposition. In *Proceedings of 27th Asilomar Conference on Signals, Systems and Computers*, pages 40–44 vol.1, 1993.

22 M. Raab and A. Steger. "balls into bins" — A simple and tight analysis. In *Randomization and Approximation Techniques in Computer Science (RANDOM'98)*, pages 159–170, 1998.

23 J. Scarlett and V. Cevher. Phase transitions in the pooled data problem. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 376–384, 2017.

24 P. Sham, J. S. Bader, I. Craig, M. O'Donovan, and M. Owen. Dna pooling: a tool for large-scale association studies. *Nature Reviews Genetics*, 3:862–871, 2002.

25 H. S. Shapiro. Problem e 1399. *Amer. Math. Monthly*, 67:82, 1960.

26 S. Sparrer and R. F. H. Fischer. Soft-feedback omp for the recovery of discrete-valued sparse signals. In *2015 23rd European Signal Processing Conference (EUSIPCO)*, pages 1461–1465, 2015.

27 J. Spencer. *Asymptopia*. American Mathematical Society, 2014.

28 Y. Wang and W. Yin. Sparse signal reconstruction via iterative support detection. *SIAM Journal on Imaging Sciences*, 3(3):462–491, January 2010.

A Groundwork

A.1 Preliminaries

In this section we present some standard results on concentration bounds and comparing distributions. Afterwards, we present some technical lemmas from the theory of concentration inequalities of the binomial distribution and approximating results for random walks that are used throughout the proof section.

We begin, for the convenience of the reader, with the basic Chernoff bound in the form which we will employ due to [27].

▶ **Lemma 10.** *Let $X \sim \mathrm{Bin}(n,p), \varepsilon > 0$ and $\delta \in (0,1)$. Then we find that*

$$\mathbb{P}\left[X > (1+\delta)np\right] \leq \exp\left(-np\delta^2/(2+\delta)\right) \quad \text{and} \quad \mathbb{P}\left[X < (1-\delta)np\right] \leq \exp\left(-np\delta^2/2\right).$$

The following lemmas are results on the asymptotic behavior of random walks. A random walk $\boldsymbol{R}$ can be described by its transition probabilities $R(x,y)$. The simple random walk on $\mathbb{Z}$ has the transition probabilities $R(x,x+1) = R(x,x-1) = 1/2$.

▶ **Lemma 11** (Section 1.5 of [27])**.** *The probability that a one-dimensional simple random walk with $2j$ steps will end at its original position is asymptotically given by $(\pi j)^{-1/2}(1+O(j^{-1}))$.*

▶ **Lemma 12.** *The following asymptotic equivalence holds for every $0 < p = p(n) < 1$ s.t. $np \to \infty$.*

$$\sum_{j=1}^{n/2} \binom{n}{2j} p^{2j}(1-p)^{n-2j} j^{-1/2} = 2^{-1/2} \sum_{j=1}^{n} \binom{n}{j} p^j (1-p)^{n-j} j^{-1/2} + O((np)^{-1}) \tag{4}$$

Proof. Let $\boldsymbol{X} \sim \mathrm{Bin}_{\geq 1}(n,p)$ and define $a_j = \mathbb{P}\left(\boldsymbol{X} = j\right)/\sqrt{j/2}$ for $j = 1 \dots n$. Then

$$a_{j+1}/a_j = (p/(1-p))\left(j/(j+1)^3\right)^{1/2}(n-j)$$

is larger than 1 up to $j^\star \in \{\lfloor (n+1)p \rfloor, \lfloor (n+1)p - 1 \rfloor\}$, depending on n being even or odd, and strictly less than 1 for $j = j^\star + 1, ..., n$. Furthermore, $a_j = o(1)$ for every j. Define j' as the largest even integer s.t. $j' \leq j^\star$. Then

$$\begin{aligned}\sum_{j=1}^{n/2} a_{2j} &\geq \frac{1}{2}\left(\sum_{j=1}^{j'/2} a_{2j} + a_{2j-1} + \sum_{j=j'/2+1}^{n/2-1} a_{2j} + a_{2j+1}\right) \\ &= \left(\frac{1}{2}\sum_{j=1}^{n} a_j\right) + \frac{1}{2}\left(a_{j'+1} + a_n + a_{n-1}\cdot \mathbf{1}(n \text{ odd })\right) = \left(\frac{1}{2}\sum_{j=1}^{n} a_j\right) + O((np)^{-2}),\end{aligned} \tag{5}$$

and similarly

$$\sum_{j=1}^{n/2} a_{2j} \leq \frac{1}{2}\left(\sum_{j=1}^{j'/2} a_{2j} + a_{2j+1} + \sum_{j=j'/2+1}^{n/2-1} a_{2j-1} + a_{2j}\right) \leq \left(\frac{1}{2}\sum_{j=1}^{n} a_j\right) + O((np)^{-2}). \tag{6}$$

Equations (5) and (6) jointly imply (4). ◀

Let μ be a distribution, f a real-valued function and $\boldsymbol{X} \sim \mu$. Then the *Jensen gap* $\mathcal{J}(f, \mu)$ is defined as

$$\mathcal{J}(f, \mu) = \left|\mathbb{E}\left[f(\boldsymbol{X})\right] - f\left(\mathbb{E}\left[\boldsymbol{X}\right]\right)\right|.$$

A well known upper bound on the Jensen gap for functions $f : I \to \mathbb{R}$ s.t. $|f(x) - f(\mathbb{E}[\boldsymbol{X}])| \leq M\,|x - \mathbb{E}[\boldsymbol{X}]|$ for all $x \in I$ (see equation (1.1) of [15]) is given by

$$\mathcal{J}(f, \mu) \leq M\mathbb{E}\left[|\boldsymbol{X} - \mathbb{E}[\boldsymbol{X}]|\right]. \tag{7}$$

An immediate consequence is the following corollary.

▶ **Corollary 13.** *Let $\boldsymbol{X} \sim \mathrm{Bin}_{x\geq 1}(n, p)$ be a binomial random variable conditioned on being at least 1, s.t. $\lim_{n\to\infty} np = \infty$. Then, as $n \to \infty$, the following holds.*

$$\mathbb{E}\left[\boldsymbol{X}^{-1/2}\right] = \left(1 + o(n^{-1})\right)\mathbb{E}[\boldsymbol{X}]^{-1/2} \quad \text{and} \quad \mathbb{E}\left[\boldsymbol{X}^{-1}\right] = \left(1 + o(n^{-1})\right)\mathbb{E}[\boldsymbol{X}]^{-1} \tag{8}$$

Proof. Let $\boldsymbol{X}' \sim \mathrm{Bin}(n, p)$. As $\mathbb{P}\left(\boldsymbol{X} = 0\right) = (1-p)^n = o(1)$, we find $\left\|\boldsymbol{X} - \boldsymbol{X}'\right\|_{\mathrm{TV}} = o(1)$. Then Lemma 10 and (7) imply

$$\left|\mathbb{E}\left[\boldsymbol{X}^{-1/2}\right] - \mathbb{E}[\boldsymbol{X}]^{-1/2}\right| \leq \mathbb{E}\left[|\boldsymbol{X} - \mathbb{E}[\boldsymbol{X}]|\right] = o(n^{-1}) \quad \text{and} \tag{9}$$

$$\left|\mathbb{E}\left[\boldsymbol{X}^{-1}\right] - \mathbb{E}[\boldsymbol{X}]^{-1}\right| \leq \mathbb{E}\left[|\boldsymbol{X} - \mathbb{E}[\boldsymbol{X}]|\right] = o(n^{-1}). \tag{10}$$

The corollary follows directly from Equation (10). ◀

Finally, we require the following result for the balls-into-bins experiment.

▶ **Lemma 14.** *Suppose that B balls are thrown uniformly at random into D bins. Then, with probability $\left(1 - \left(1 - \frac{1}{D}\right)^B\right)^D$, we find no empty bins.*

A.2 Properties of the Pooling Scheme

For the analysis of QGT, the underlying pooling scheme can be seen as a bipartite factor graph $\boldsymbol{G}$. The structure is induced by degree sequences $\boldsymbol{\Delta} = (\boldsymbol{\Delta}_i)_{i\in[n]}$ and Γ and the chosen test design is randomized. Nevertheless we can apply standard techniques to gain insight into the form of the underlying graph. In order to prove Lemma 3, we show each of the two statements.

▶ **Lemma 15.** *With probability $1 - o\left(n^{-2}\right)$ we have for all $i \in [n]$*

$$\boldsymbol{\Delta}_i = \frac{m}{2} + O\left(\sqrt{m}\ln n\right).$$

Proof. From the construction of $\boldsymbol{G}$, it follows that $\boldsymbol{\Delta}_i$ is distributed as $\mathrm{Bin}(mn/2, 1/n)$. Then Lemma 10 implies for an individual $x_i \in V$

$$\mathbb{P}\left(\boldsymbol{\Delta}_i > m/2 + O\left(\sqrt{m}\ln^2 n\right)\right) = n^{-\omega(1)}$$

Taking the union bound over all n individuals implies the lemma. ◀

▶ **Lemma 16.** *With probability $1 - o\left(n^{-2}\right)$ we have for all $i \in [n]$*

$$\boldsymbol{\Delta}_i^\star = \left(1 - 1/\sqrt{\mathrm{e}}\right)m + O\left(\sqrt{m}\ln n\right).$$

Proof. The probability that an individual x_i is assigned to a specific test a_j is given by

$$p = 1 - \left(1 - \frac{1}{n}\right)^{\Gamma} = \left(1 + n^{-\Omega(1)}\right)\left(1 - 1/\sqrt{\mathrm{e}}\right).$$

Since tests select their participating individuals independently of each other, we observe that $\boldsymbol{\Delta}_i^\star \sim \mathrm{Bin}\,(m, p)$. Thus, a standard application of the Chernoff bound and the union bound over all n individuals implies the lemma. ◀

Those results suffice in order to show that $\mathcal{R}$ as given through Lemma 3 is indeed a high probability event.

Proof of Lemma 3. The lemma follows from Lemmas 15 and 16. ◀

B Proof of Theorem 1

B.1 Proof of Lemma 5

The product of the two binomial coefficients simply accounts for the number of configurations σ that have overlap l with $\boldsymbol{\sigma}$. Hence, with $\mathcal{S}$ denoting the event that one specific $\sigma \in \{0,1\}^V$ that has overlap ℓ with $\boldsymbol{\sigma}$ belongs to $S_{k,\ell}(\boldsymbol{G}, \boldsymbol{y})$, it suffices to show for $\boldsymbol{X} \sim \mathrm{Bin}_{\geq 1}(\Gamma, 2(1-\ell/k)k/n)$ that

$$\mathbb{P}[\mathcal{S} \mid \mathcal{R}] \leq (1 + O(1)) \left(\frac{1}{\sqrt{2\pi}} \mathbb{E}\left[\frac{1}{\sqrt{\boldsymbol{X}}}\right]\right)^m$$

By the pooling scheme, the size of each test is fixed to $\Gamma = n/2$ with individuals chosen uniformly at random with replacement. Observe that all tests are independent of each other. Therefore, we need to determine the probability that for a specific σ and a specific test a_i the test result is consistent with the test result under $\boldsymbol{\sigma}$, i.e., $\boldsymbol{y}_i = y_i$. Given the overlap ℓ, we know for a uniformly at random drawn σ that $\mathbb{P}\left[\boldsymbol{\sigma}_i = \sigma_i = 1\right] = \ell/n$, $\mathbb{P}\left[\boldsymbol{\sigma}_i = \sigma_i = 0\right] = (n - 2k + \ell)/n$ and finally $\mathbb{P}\left[\boldsymbol{\sigma}_i \neq \sigma_i\right] = (k - \ell)/n$ holds for all individuals x_i. We get

$$\begin{aligned}
\mathbb{P}[\mathcal{S} \mid \mathcal{R}] &\leq \prod_{i=1}^{m} \sum_{j=1}^{\boldsymbol{y}_i} \binom{\Gamma}{j, j, \Gamma - 2j} \left((1-\ell/k)\frac{k}{n}\right)^{2j} \left(1 - 2(1-\ell/k)\frac{k}{n}\right)^{\Gamma-2j} \\
&\leq \left(\sum_{j=1}^{\Gamma/2} \binom{\Gamma}{2j} \left(2(1-\ell/k)\frac{k}{n}\right)^{2j} \left(1 - 2(1-\ell/k)\frac{k}{n}\right)^{\Gamma-2j} \binom{2j}{j} 2^{-2j}\right)^m \qquad (11)
\end{aligned}$$

The last two components of (11) describe the probability that a one-dimensional simple random walk will return to its original position after $2j$ steps, which is by Lemma 11 equal to $(1+O(j^{-1}))/\sqrt{\pi j}$. The term before describes the probability that a $\mathrm{Bin}_{\geq 1}(\Gamma, 2(1-\ell/k)k/n))$ random variable $\boldsymbol{X}$ takes the value $2j$. For $\ell \leq k - (1 - \exp(-1/2)) \ln k$ the expectation of $\boldsymbol{X}$ given $\boldsymbol{G}$ is at least of order $\ln k$ such that the asymptotic description of the random walk return probability is feasible. Note that if ℓ gets closer to k, the expectation of $\boldsymbol{X}$ gets finite, s.t. the random walk approximation is not feasible anymore. Therefore, using Lemma 12, we

can, as long as $\Gamma(2(1-\ell/k)k/n) = \Omega(\ln n)$, simplify (11) in the large-system limit to

$$\begin{aligned}\mathbb{P}[\mathcal{S} \mid \mathcal{R}] &\leq (1+O(1)) \left(\sum_{j=1}^{\Gamma/2} \binom{\Gamma}{2j} \left(2(1-\ell/k)\frac{k}{n}\right)^{2j} \left(1-2(1-\ell/k)\frac{k}{n}\right)^{\Gamma-2j} \frac{1}{\sqrt{\pi j}} \right)^m \\ &= (1+O(1)) \left(\frac{1}{2} \sum_{j=1}^{\Gamma} \binom{\Gamma}{j} \left(2(1-\ell/k)\frac{k}{n}\right)^{j} \left(1-2(1-\ell/k)\frac{k}{n}\right)^{\Gamma-j} \frac{1}{\sqrt{\pi j/2}} \right)^m \\ &= (1+O(1)) \left(\frac{1}{\sqrt{2\pi}} \mathbb{E}\left[\frac{1}{\sqrt{\boldsymbol{X}}}\right] \right)^m\end{aligned}$$

which implies Lemma 5. ◀

B.2 Proof of Lemma 6

Let $\boldsymbol{X} \sim \mathrm{Bin}(\Gamma, 2(1-\ell/k)k/n))$. Then Lemma 5 and Corollary 13 imply

$$\mathbb{E}[Z_{k,\ell}(\boldsymbol{G}, \boldsymbol{y}) \mid \mathcal{R}] \leq (1+O(1)) \binom{k}{\ell} \binom{n-k}{k-\ell} (2\pi \mathbb{E}[\boldsymbol{X}])^{-m/2}$$

In the following we use the well known fact [27] that as $n \to \infty$ we have

$$n^{-1} \ln \binom{n}{np} \to H(p),$$

where the expression $H(a)$ with $a \in [0,1]$ denotes the entropy of a $\mathrm{Be}(a)$-variable such that $H(a) = -a\ln(a) - (1-a)\ln(1-a)$. Correspondingly by taking the $\ln(\cdot)$ on the r.h.s. and scaling with $1/n$, we find that

$$\begin{aligned}&\frac{1}{n} \ln\left(\mathbb{E}\left[Z_{k,\ell}(\boldsymbol{G}, \boldsymbol{y}) \mid \mathcal{R}\right]\right) \qquad (12)\\ &= \frac{1}{n}\left(\ln(1+O(1)) + \ln\left(\binom{k}{\ell}\right) + \ln\left(\binom{n-k}{k-\ell}\right) - \frac{m}{2}\ln(2\pi\mathbb{E}[\boldsymbol{X}]) \right) \\ &\leq (1+o(1)) \left(\frac{k}{n} H\left(\frac{\ell}{k}\right) + \left(1-\frac{k}{n}\right) H\left(\frac{k-\ell}{n-k}\right) - \frac{m}{2n} \ln\left(4\pi\Gamma\left(1-\frac{\ell}{k}\right)\frac{k}{n}\right) \right) \\ &= (1+o(1)) \left(\frac{k}{n} H\left(\frac{\ell}{k}\right) + \left(1-\frac{k}{n}\right) H\left(\frac{k-\ell}{n-k}\right) - \frac{ck/n\ln(n/k)}{2\ln k} \ln\left(2\pi k\left(1-\ell/k\right)\right) \right) \qquad (13)\end{aligned}$$

Lemma 6 follows. ◀

B.3 Proof of Lemma 7

Recall that for $\gamma = 1 - \exp(-1/2)$ and c being a constant we have

$$m = ck \ln\left(\frac{n}{k}\right) / \ln k = c\frac{1-\theta}{\theta} k \quad \text{and} \quad 0 \leq \ell \leq k - \gamma \ln k. \qquad (14)$$

Then define $f_{n,k} : [0, k-\gamma\ln k] \to \mathbb{R}$ as

$$\ell \mapsto \left(\frac{k}{n} H\left(\frac{\ell}{k}\right) + \left(1-\frac{k}{n}\right) H\left(\frac{k-\ell}{n-k}\right) - \frac{ck/n\ln(n/k)}{2\ln k} \ln(2\pi(1-\ell/k)k) \right) \qquad (15)$$

and suppose, as usually, $0 \ln 0 = 0$. By Lemma 6 we find $n^{-1} \ln \left(\mathbb{E}\left[Z_{k,\ell}(\boldsymbol{G}, \boldsymbol{y}) \mid \boldsymbol{G} \right] \right) \leq (1 + o(1)) f_{n,k}(\ell)$. Expanding the entropy yields

$$f_{n,k}(\ell) = \frac{1}{n} \Bigg(- \ell \ln \left(\frac{\ell}{k} \right) - (k - \ell) \ln \left(1 - \frac{\ell}{k} \right) - (k - \ell) \ln \left(\frac{k - \ell}{n - k} \right) - \\ (n - 2k + \ell) \ln \left(1 - \frac{k - \ell}{n - k} \right) + \frac{ck \ln(k/n)}{2 \ln k} \ln \left(2\pi k \left(1 - \frac{\ell}{k} \right) \right) \Bigg).$$

Therefore,

$$f'_{n,k}(\ell) = \frac{1}{n} \left(- \ln \left(\frac{\ell}{k} \right) + \ln \left(1 - \frac{\ell}{k} \right) + \ln \left(\frac{k - \ell}{n - k} \right) - \ln \left(1 - \frac{k - \ell}{n - k} \right) - \frac{ck \ln(k/n)}{2(k - \ell) \ln k} \right),$$

$$f''_{n,k}(\ell) = \frac{1}{n} \left(- \frac{1}{\ell} - \frac{2}{k - \ell} - \frac{1}{n - 2k + \ell} - \frac{ck \ln(k/n)}{2 \ln k (k - \ell)^2} \right).$$

As long as $\ell = o(k)$ we find that

$$n f''_{n,k}(\ell) = -\frac{1}{\ell} - \frac{1}{n - 2k + \ell} - \frac{1}{k - \ell} \left(2 - \frac{ck(1 - \theta)}{2\theta(k - \ell)} \right) < 0 \tag{16}$$

as $\left| \frac{1}{k - \ell} \left(2 - \frac{ck(1-\theta)}{2\theta(k-\ell)} \right) \right| \ll \frac{1}{\ell}$. In this case, Equation (16) shows that $f'_{n,k}$ is monotonously decreasing in ℓ for large enough n. Furthermore, $f'_{n,k}$ is continuous on $(0, k - \gamma \ln k]$. Let $\tilde{c} > 0$ be an arbitrary constant, then we have

$$\begin{aligned} & n f'_{n,k} \left(\tilde{c} \frac{k^2}{n} \right) \\ & \quad = - \ln \left(\tilde{c} \frac{k}{n} \right) + \ln \left(1 - \tilde{c} \frac{k}{n} \right) + \ln \left(\frac{k}{n} \right) + \ln \left(\frac{1 - \tilde{c}k/n}{1 - k/n} \right) && (17) \\ & \qquad - \ln \left(1 - \frac{k(1 - \tilde{c}k/n)}{n(1 - k/n)} \right) + \frac{c(1 - \theta)}{2\theta(1 - \tilde{c}k/n)} && (18) \\ & \quad = - \ln (\tilde{c}) + \frac{c(1 - \theta)}{\theta} + o(1) \end{aligned}$$

Thus, Equation (18) implies that there are constants $0 < \tilde{c}_1 < \tilde{c}_2 < \infty$ such that

$$n f'_{n,k} \left(\tilde{c}_1 \frac{k^2}{n} \right) > 0 \quad \text{and} \quad n f' \left(\tilde{c}_2 \frac{k^2}{n} \right) < 0.$$

By the intermediate value theorem we conclude that there is $\hat{c} \in [\tilde{c}_1, \tilde{c}_2]$ such that $\hat{c} \frac{k^2}{n}$ is the unique zero of $f'_{n,k}$ for all $\ell = o(k)$ and, respectively, the unique maximizer of $f_{n,k}$ in this regime. Finally, by putting this value into Equation (15) we find that the highest order terms satisfy

$$n f_{n,k} \left(\hat{c} \frac{k^2}{n} \right) < 0 \iff n H\left(k/n \right) - \frac{ck \ln(n/k)}{2} < 0 \iff c > -2 \frac{H(k/n)}{k/n \ln(k/n)} = 2 + o(1). \tag{19}$$

We are left to show that, for large enough n, we have $f_{n,k}(\ell) < 0$ for any $k - \gamma \ln k \geq \ell = \Theta(k)$. It is immediate from the definition that, in this case, we find

$$n f_{n,k}(\ell) = -\frac{c(1 - \theta)}{2\theta} k \ln(k) + O(k) \tag{20}$$

which is negative.

Therefore, the assertion of the lemma follows from Equations (14), (19), and (20) ◀

B.4 Proof of Proposition 8

Assume, $\sigma \in \{0,1\}^n$ is a second configuration consistent with the test results vector $\boldsymbol{y}$. By definition, there is at least one infected individual $x_j \in V$ s.t. $\sigma_j = 0$.

By Lemma 3 the size of $\partial^\star x_j$ is at least $\boldsymbol{\Delta}_i^\star \geq (1-\exp(-1/2))\, m - O\left(\sqrt{m}\ln n\right)$ and for any test $a_l \in \partial x_j$ we find $|y_l(\sigma) - y_l(\boldsymbol{\sigma})| \geq 1$. In order to guarantee $y(\boldsymbol{\sigma}) = y(\sigma)$, it is certainly necessary (admittedly not likely sufficient) to identify individuals $x_1, ..., x_h$ s.t. $\sigma_i = 1 - \boldsymbol{\sigma}_i$ for $i = 1 \dots h$ s.t. $\partial\{x_1, ..., x_h\} \supseteq \partial x_j$.

By construction of $\boldsymbol{G}$, the amount of tests in $\partial^\star x_j$ that do not contain any of the individuals $x_1, ..., x_h$, i.e., $\boldsymbol{H} = |\{a \in \partial^\star x_j : \{x_1, \dots, x_h\} \cap \partial a = \emptyset\}|$, can be coupled with the distribution of the number of empty bins in a balls-into-bins experiment, described as follows. Given $\boldsymbol{G}$, throw $b = \sum_{i=1}^h \deg(x_i)$ balls into $\deg(x_i) \geq (1-\exp(-1/2))\, m - O\left(\sqrt{m}\ln n\right)$ bins. Denote by $\boldsymbol{H'}$ the number of empty bins. Since for any x_i the $\deg(x_i)$ edges are not only distributed over the $(1-o(1))\,(1-\exp(-1/2))\, m$ factor nodes in ∂x_j but over all m factor nodes in $\boldsymbol{G}$, we find

$$\mathbb{P}\left[\boldsymbol{H} = 0 \mid \mathcal{R}\right] \leq \mathbb{P}\left[\boldsymbol{H'} = 0 \mid \mathcal{R}\right]. \tag{21}$$

Since given $\mathcal{R}$, $b = (1+o(1))hm/2$ by Lemma 14, the r.h.s. of (21) becomes by using $h = L\ln m$ and letting $\gamma = (1-\exp(-1/2))$

$$\mathbb{P}\left[\boldsymbol{H'} = 0 \mid \mathcal{R}\right] \leq \left(1 - \left(1 - \frac{1}{\gamma m}\right)^{hm/2}\right)^{\gamma m} = (1+o(1))\exp\left(-\gamma m^{1-L/(2\gamma)}\right).$$

This value is $n^{-\omega(1)}$ whenever $L < 2\gamma$, thus if

$$h < 2\gamma \ln(m) \sim 2\gamma\left(\ln k + \ln\ln k\right).$$

We conclude that if the Hamming distance of $\boldsymbol{\sigma}$ and σ is at least one, it is w.h.p. at least $2\gamma\left(\ln k + \ln\ln k\right)$ with probability $1 - n^{-\omega(1)}$. Thus, a union bound over all k infected individuals implies the proposition. ◀

C Proof of Theorem 2

Recall that $\boldsymbol{\Delta}_j^\star$ is the number of distinct tests individual x_j is connected to and $\boldsymbol{\Delta}_j$ signifies the total number of tests in which x_j is contained (multi-edges counted multiple times. Furthermore, let $\mathcal{E}_j$ be the $\sigma-$algebra generated by the edges incident to x_j. Furthermore, let $\boldsymbol{A}_{ij} \in \mathbb{N}_0$ denote how often individual x_j appears in test a_j.

▶ **Corollary 17.** *Let $1 \leq j \leq n$. Given $\mathcal{E}_j$ the random variable*

$$\boldsymbol{S}_j = \boldsymbol{\psi}_j - \boldsymbol{\Delta}_j = \sum_{i=1}^m \mathbf{1}\left\{\boldsymbol{A}_{ij} > 0\right\}\left(\boldsymbol{y}_j - \boldsymbol{A}_{ij}\right)$$

has distribution

$$\mathrm{Bin}\left(\boldsymbol{\Delta}_j^\star\Gamma - \boldsymbol{\Delta}_j, \frac{k - \mathbf{1}\{\boldsymbol{\sigma}(j) = 1\}}{n-1}\right).$$

Proof. This is immediate from the model definition. There are $\Gamma\boldsymbol{\Delta}_j^\star - \boldsymbol{\Delta}_j$ half-edges connected to tests in the neighborhood of x_j which are connected to different individuals than x_j. Each of those half-edges is connected to one of $k - \mathbf{1}\left\{\boldsymbol{\sigma}(j) = 1\right\}$ infected individuals independently out of the $n-1$ remaining individuals. ◀

For the sake of brevity let $\gamma = 1 - \exp(-1/2)$. Given the event $\mathcal{R}$ which guarantees concentration properties of the underlying graph we find, with high probability,

$$\mathbb{E}[\boldsymbol{S}_j \mid \mathcal{E}_j, \mathcal{R}] = (1 \pm \delta)\frac{\gamma km}{2} \quad \text{where} \quad \delta := \frac{\sqrt{2}\ln n}{\sqrt{\gamma mk}} = o(1). \tag{22}$$

The Chernoff bound allows us to bound the tails of $\boldsymbol{S}_j$ as follows.

▶ **Lemma 18.** *Let $\alpha \in (0,1)$ be a constant and $m = dk\ln\frac{n}{k}$. Given $\mathcal{R}$ we find*

$$\mathbb{P}(|\boldsymbol{S}_j - \mathbb{E}[\boldsymbol{S}_j \mid \mathcal{E}_j, \mathcal{R}]| \geq (1-\alpha)m/2 \mid \mathcal{E}_j, \mathcal{R}) \leq \exp\left(-(1+o(1))\frac{(1-\alpha)^2 d}{4\gamma(1+o(1))}\ln\frac{n}{k}\right).$$

Proof. The Chernoff bound (Lemma 10) directly implies

$$\begin{aligned}
\mathbb{P}(|\boldsymbol{S}_j - \mathbb{E}[\boldsymbol{S}_j \mid \mathcal{E}_j] \mid \mathcal{E}_j, \mathcal{R}| \geq (1-\alpha)m/2) &\leq \exp\left(-(1+o(1))\frac{(1-\alpha)^2 m}{8\mathbb{E}[\boldsymbol{S}_j \mid \mathcal{E}_j, \mathcal{R}]}\right) \\
&= \exp\left(-(1+o(1))\frac{(1-\alpha)^2 m}{4\gamma k(1+o(1))}\right) \\
&= \exp\left(-(1+o(1))\frac{(1-\alpha)^2 dk\ln(n/k)}{4\gamma k(1+o(1))}\right) \\
&= \exp\left(-(1+o(1))\frac{(1-\alpha)^2 d}{4\gamma(1+o(1))}\ln(n/k)\right),
\end{aligned}$$

as claimed. ◀

Next we show that, depending on a suitable choice for a threshold, the scores of infected and uninfected individuals are indeed well separated.

▶ **Corollary 19.** *Let $\varepsilon > 0$ be an arbitrary constant.*
If $m \geq (4+\varepsilon)(1-\exp(-1/2))\frac{1+\sqrt{\theta}}{1-\sqrt{\theta}}k\ln\frac{n}{k}$ there is an $\alpha \in (0,1)$ such that, w.h.p., we have

$$\begin{aligned}
\boldsymbol{S}_j + \boldsymbol{\Delta}_j &\geq \mathbb{E}[\boldsymbol{S}_j \mid \mathcal{E}_j] + (1-\alpha)m/2 && \text{for all } x_j \text{ s.t. } \boldsymbol{\sigma}(j) = 1, \\
\boldsymbol{S}_j &< \mathbb{E}[\boldsymbol{S}_j \mid \mathcal{E}_j] + (1-\alpha)m/2 && \text{for all } x_j \text{ s.t. } \boldsymbol{\sigma}(j) = 0.
\end{aligned}$$

Proof. Let x_j be an infected individual. We may condition on the event $\mathcal{R}$, therefore, we suppose that $\boldsymbol{\Delta}_j = m/2 + O(\sqrt{m}\ln n)$. Then Lemma 18 ensures

$$\begin{aligned}
\mathbb{P}(\boldsymbol{S}_j + \boldsymbol{\Delta}_j \leq \mathbb{E}[\boldsymbol{S}_j \mid \mathcal{E}_j] + (1-\alpha)m/2 \mid \mathcal{E}_j, \mathcal{R}) &\leq \exp\left(-\alpha^2 d/(4\gamma(1+o(1)))\ln\frac{n}{k}\right) \\
&= \exp\left(\frac{(\theta-1)\alpha^2 d}{4\gamma(1+o(1))}\ln n\right).
\end{aligned}$$

Hence, the union bound shows that the first inequality holds for all infected individuals x_j w.h.p. if

$$\frac{(\theta-1)\alpha^2 d}{4\gamma(1+o(1))} + \theta < 0. \tag{23}$$

Analogously, the second inequality holds for all uninfected individuals w.h.p. if

$$\begin{aligned}
\mathbb{P}[\boldsymbol{S}_j \geq \mathbb{E}[\boldsymbol{S}_j \mid \mathcal{E}_j] + (1-\alpha)m/2 \mid \mathcal{E}_j] &\leq \exp\left(\left((1-\alpha)^2 d/(4\gamma(1+o(1)))\right)\ln\frac{n}{k}\right) \\
&= \exp\left(\frac{(\theta-1)(1-\alpha)^2 d}{4\gamma(1+o(1))}\ln n\right).
\end{aligned}$$

Thus, the union bound shows that the second inequality holds w.h.p. if

$$\frac{(\theta-1)(1-\alpha)^2 d}{4\gamma(1+o(1))}+1<0. \tag{24}$$

Thus, as one sufficient condition is monotonously increasing in α while the other is monotonously decreasing, the optimal choice of α is the one that makes the two terms (23) and (24) equal:

$$\frac{(\theta-1)\alpha^2 d}{4\gamma(1+o(1))}+\theta=\frac{(\theta-1)(1-\alpha)^2 d}{4\gamma(1+o(1))}+1,$$

which boils down to

$$\alpha=\frac{d-4\gamma(1+o(1))}{2d}.$$

By putting this solution for α into (23) we find that (23) equals

$$\frac{(\theta-1)(d-4(\gamma+o(1))^2}{16\gamma d+o(1)}+\theta.$$

Hence, it suffices to find $d=d(\theta)$ such that

$$\frac{(\theta-1)(d-4\gamma+o(1))^2}{16\gamma d+o(1)}+\theta=0.$$

There are two solutions, the greater of which works out to be

$$d=4\gamma\cdot\frac{1+\sqrt{\theta}}{1-\sqrt{\theta}}+o(1).$$

Hence, for d exceeding this value the two desired inequalities are valid w.h.p. ◀

C.1 Proof of Theorem 2

Theorem 2 follows directly from Corollary 19 and the definition $m=dk\ln\frac{n}{k}$.

APPENDIX G. THE ISING ANTIFERROMAGNET AND MAX CUT ON RANDOM REGULAR GRAPHS

THE ISING ANTIFERROMAGNET AND MAX CUT ON RANDOM REGULAR GRAPHS

AMIN COJA-OGHLAN, PHILIPP LOICK, BALÁZS F. MEZEI, GREGORY B. SORKIN

ABSTRACT. The Ising antiferromagnet is an important statistical physics model with close connections to the MAX CUT problem. Combining spatial mixing arguments with the method of moments and the interpolation method, we pinpoint the replica symmetry breaking phase transition predicted by physicists. Additionally, we rigorously establish upper bounds on the MAX CUT of random regular graphs predicted by Zdeborová and Boettcher [Journal of Statistical Mechanics 2010]. As an application we prove that the information-theoretic threshold of the disassortative stochastic block model on random regular graphs coincides with the Kesten-Stigum bound. *MSc: 05C80.*

1. INTRODUCTION

1.1. **Motivation.** The Ising model is to statistical physics what the k-SAT problem is to computer science or the Ramsey problem to combinatorics: it serves as a benchmark for new techniques to prove their mettle. Devised by Lenz in the 1920s to explain magnetism, the Ising model can be defined on an arbitrary graph G. Think of the vertices of G as iron atoms that each carry one of two possible magnetic spins, ± 1. With the topology of interactions defined by the edges of G, the Hamiltonian $\mathscr{H}_G$ (the 'energy' function) maps a spin configuration $\sigma \in \{\pm 1\}^V$ to the number of edges of G that link two vertices with the same spin, i.e.,

$$\mathscr{H}_G(\sigma) = \sum_{\{v,w\}\in E} \frac{1+\sigma_v\sigma_w}{2}. \tag{1.1}$$

Together with a real parameter β the Hamiltonian induces a probability distribution $\mu_{G,\beta}$ on the set of spin configurations via

$$\mu_{G,\beta}(\sigma) = \frac{\exp(-\beta\mathscr{H}_G(\sigma))}{Z_{G,\beta}} \qquad (\sigma \in \{\pm 1\}^V) \qquad \text{where} \qquad Z_{G,\beta} = \sum_{\tau\in\{\pm 1\}^V} \exp(-\beta\mathscr{H}_G(\tau)). \tag{1.2}$$

This probability measure is called the Boltzmann distribution. The normalising term $Z_{G,\beta}$ is known as the partition function. If $\beta > 0$, then $\mu_{G,\beta}$ favours spin configurations σ with a small number of edges joining vertices with the same spin; this case is known as the antiferromagnetic Ising model. By contrast, in the ferromagnetic case $\beta < 0$ configurations with many aligned spins receive a boost.

Both variants of the Ising model are of keen interest in physics and the literature on each, rigorous as well as non-rigorous, is vast [30, 37]. But the antiferromagnetic Ising model appears to be more challenging. According to physics lore this is because its Boltzmann distribution is prone to a complicated type of long-range correlation known as 'replica symmetry breaking'. Another way to see the challenge is that from the partition function we could solve the NP-complete problem MAX CUT: as β increases the mass of the Boltzmann distribution shifts to spin configurations with more edges joining vertices with opposite spins. Ultimately the measure concentrates on the maximum cuts of the graph G, and it is well known (and easy to check) that

$$\mathrm{MAXCUT}(G) = \frac{dn}{2} + \lim_{\beta\to\infty} \frac{\partial}{\partial\beta} \log Z_{G,\beta}. \tag{1.3}$$

We study the Ising antiferromagnet on the random d-regular graph $\mathbb{G} = \mathbb{G}(n,d)$. From a statistical physics perspective, this example has been suggested as one of the simplest models where replica symmetry breaking is expected to occur. Fond of lattice-like geometries, physicists favour the random regular graph, which converges to the d-regular tree in the Benjamini-Schramm topology, over the Erdős-Rényi model. In particular, regularity greatly simplifies the physics 'cavity equations' that Zdeborová and Boettcher [63] employed to put forward a beautiful, well-known conjecture about MAX CUT on random regular graphs. From a combinatorics perspective, the random regular graph provides a neat but notoriously challenging model for MAX CUT, both structurally (determining the fraction of edges it should be possible to cut, asymptotically almost surely) and algorithmically (finding algorithms that give large cuts in such graphs). The problem has received a great deal of attention in the combinatorics community, e.g. [18, 19, 26, 27, 38, 39, 61]. Additionally, the Ising model is intimately related to the regular version of the disassortative stochastic block model [16], a prominent case study in Bayesian inference.

Amin Coja-Oghlan and Philipp Loick are supported by DFG CO 646/3.

1.2. **Our contributions and paper outline.** Our first contribution is to identify the precise value of β where the replica symmetry breaking phase transition occurs; see Theorem 1.1. A common approach to problems of this type would be the trick of bounding the second moment of random regular graphs by that of the Erdős-Rényi as applied in [1, 13]. Since this approach fails in our case, we instead turn to harnessing spatial mixing arguments to establish the phase transition. As a ramification of the replica symmetry breaking phase transition, our second contribution is to derive the information-theoretic threshold of the disassortative regular stochastic block model; see Theorem 1.4. Our third contribution is to establish rigorously the upper bound on the MAX CUT of the random regular graph predicted by Zdeborová and Boettcher; see Corollary 1.3. Specifically, their prediction is based on the so-called '1-step replica symmetry breaking' formalism from physics. Using the interpolation method it is easy to obtain a rigorous upper bound that comes as a variational problem. However, this variational problem appears rather unwieldy at first glance. But by expressing the variational problem as a certain random walk that we can analyze, we obtain an elegant explicit expression whose numerical results match those of Zdeborová and Boettcher.

In the remainder of Section 1 we state the main results of the paper precisely. In Section 2 we outline the proof strategy, and in Section 3 we discuss the advances over earlier work. Details of the proofs of Theorem 1.1, Theorem 1.2, and Corollary 1.3 are given respectively in Sections 4, 5, and 6.

1.3. **Replica symmetry breaking.** The key quantity associated with the Ising model on $\mathbb{G}$ is the partition function $Z_{\mathbb{G},\beta}$. This is because various combinatorially meaningful observables derive from the partition function via differentiation; see, for example, (1.3). Because $Z_{\mathbb{G},\beta}$ scales exponentially in n, it is common to consider the normalised logarithm $n^{-1}\log Z_{\mathbb{G},\beta}$, known as the free energy. Routine arguments show that this random variable concentrates about its mean. Hence, we are led to investigate the function

$$\Phi_d : \beta \in (0,\infty) \mapsto \lim_{n\to\infty} \frac{1}{n}\mathbb{E}\left[\log Z_{\mathbb{G},\beta}\right]; \tag{1.4}$$

the limit is known to exist for all $d \geq 3, \beta > 0$ [8]. In particular, for a fixed $d \geq 3$ the singularities β of Φ_d (the points at which Φ_d cannot be expanded to an absolutely convergent power series) are called the *phase transitions* of the Ising model. Hence, from a mathematical physics point of view computing Φ_d and pinpointing the phase transitions is the key challenge associated with the model.

Jensen's inequality immediately yields the inequality

$$\Phi_d(\beta) \leq \lim_{n\to\infty} \frac{1}{n}\log\mathbb{E}\left[Z_{\mathbb{G},\beta}\right] = \log 2 + \frac{d}{2}\log\frac{1+\mathrm{e}^{-\beta}}{2}, \tag{1.5}$$

where the equality is taken from the (easy) calculation of $\mathbb{E}[Z_{\mathbb{G},\beta}]$ as (2.10) in Lemma 2.1. A tempting first guess might be that (1.5) is generally tight. Combinatorially this would indicate that the Boltzmann distribution $\mu_{\mathbb{G},\beta}$ is free from long-range correlations. To see this, consider the experiment of removing a single random edge $\boldsymbol{e} = \{v, w\}$ from $\mathbb{G}$. Because short cycles are scarce, w.h.p. the vertices v, w have distance $\Omega(\log n)$ in $\mathbb{G}-\boldsymbol{e}$. Hence, in the absence of long-range correlations, in a sample $\boldsymbol{\sigma}$ from the Boltzmann distribution of $\mathbb{G}-\boldsymbol{e}$, the spins $\boldsymbol{\sigma}_v, \boldsymbol{\sigma}_w$ should be asymptotically independent, i.e., $\mathbb{P}[\boldsymbol{\sigma}_v = \boldsymbol{\sigma}_w \mid \mathbb{G}, \boldsymbol{e}] = 1/2 + o(1)$. Therefore, adding $\boldsymbol{e}$ back in should change the partition function by

$$\log\frac{Z_{\mathbb{G},\beta}}{Z_{\mathbb{G}-\boldsymbol{e},\beta}} = \log\Big(1-(1-\mathrm{e}^{-\beta})\mathbb{P}[\boldsymbol{\sigma}_v = \boldsymbol{\sigma}_w \mid \mathbb{G}, \boldsymbol{e}]\Big) \sim \log\frac{1+\mathrm{e}^{-\beta}}{2}.$$

Removing a random edge $dn/2$ times until all the edges are gone and observing that the partition function of the empty graph equals 2^n, we would thus obtain equality in (1.5). However, the following theorem shows that (1.5) is tight only for β up to an explicit threshold β^*.

Theorem 1.1. *For any $d \geq 3$ let*

$$\beta^*(d) = \log\left(\frac{\sqrt{d-1}+1}{\sqrt{d-1}-1}\right). \tag{1.6}$$

(i) If $\beta < \beta^*(d)$, then

$$\Phi_d(\beta) = \log 2 + \frac{d}{2}\log\frac{1+\mathrm{e}^{-\beta}}{2}.$$

(ii) If $\beta > \beta^*(d)$, then

$$\Phi_d(\beta) < \log 2 + \frac{d}{2}\log\frac{1+\mathrm{e}^{-\beta}}{2}.$$

Because the function $\beta \mapsto \log 2 + \frac{d}{2}\log\frac{1+e^{-\beta}}{2}$ is analytic, Theorem 1.1 implies that $\Phi_d(\beta)$ is non-analytic at the point $\beta = \beta^*$. Hence, there occurs a phase transition at β^* that separates a regime where $Z_{\mathbb{G},\beta}$ concentrates about its mean from a regime where the mean is driven up by rare events. In physics jargon this phase transition is called the *replica symmetry breaking* transition. The value β^* has a special combinatorial meaning: it is the reconstruction threshold for a broadcasting process first studied by Kesten and Stigum [40], and is thus known as the 'Kesten-Stigum bound'. Thus, Theorem 1.1 shows that the replica symmetry breaking phase transition in the Ising antiferromagnet on $\mathbb{G}$ occurs precisely at the Kesten-Stigum bound.

1.4. **Bounding** MAX CUT**.** Theorem 1.1 does not provide a simple expression for $\Phi_d(\beta)$ for $\beta > \beta^*$. Indeed, such a simple expression may not exist. This is because according to physics predictions the value $\Phi_d(\beta)$ for $\beta > \beta^*$ results from a complicated variational problem over an infinite-dimensional space of probability measures that meticulously characterises the long-range correlations of the Boltzmann distribution [15].

Yet in the limit $\beta \to \infty$ it is possible to derive an explicit upper bound on the value of $\Phi_d(\beta)$. To state this bound consider the following right stochastic band matrix $\mathscr{M}$ of size $(d+1)\times(d+1)$:

$$\mathscr{M} = \begin{bmatrix} 0 & 1 & 0 & \cdots & \cdots & \cdots & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} & \ddots & & & \vdots \\ 0 & \frac{1}{2} & 0 & \frac{1}{2} & \ddots & & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ \vdots & & \ddots & \ddots & \ddots & \ddots & 0 \\ \vdots & & & \ddots & \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & \cdots & \cdots & \cdots & 0 & 1 & 0 \end{bmatrix}. \tag{1.7}$$

Moreover, let

$$F_d(\alpha, z) = -\frac{\log\left(\zeta \mathscr{A}^d \xi\right)}{\log z} + \frac{d\log\left(1-2\alpha^2+2\alpha^2 z\right)}{2\log z}, \qquad \text{where} \tag{1.8}$$

$$\mathscr{A} = (1-2\alpha)\mathrm{id} + 2\alpha\sqrt{z}\mathscr{M}, \tag{1.9}$$

$$\zeta = \begin{bmatrix} 1, & 0, & 0, & \cdots \end{bmatrix} \in \mathbb{R}^{1\times(d+1)}, \tag{1.10}$$

$$\xi = \begin{bmatrix} 1, & z^{-1/2}, & z^{-1}, & z^{-3/2}, & \cdots \end{bmatrix}^T \in \mathbb{R}^{(d+1)\times 1}. \tag{1.11}$$

Theorem 1.2. *For any $d \geq 3$ we have*

$$\lim_{\beta\to\infty} \beta^{-1}\Phi_d(\beta) \leq \inf_{\substack{0<\alpha\leq 1/2 \\ 0<z<1}} F_d(\alpha, z).$$

Since (1.3) shows that the MAX CUT problem is tied to $\Phi_d(\beta)$ for large β, we can use Theorem 1.2 to derive upper bounds on the maximum cut size of the random regular graph.

Corollary 1.3. *Let* MAXCUT($\mathbb{G}$) *be the number of edges cut by a maximum cut of $\mathbb{G}$. Then, w.h.p.,*

$$\mathrm{MAXCUT}(\mathbb{G}) \leq \frac{dn}{2} \inf_{\substack{0<\alpha<1/2 \\ 0<z<1}} \left(1 + \frac{2}{d} F_d(\alpha, z)\right) + o(n).$$

Zdeborová and Boettcher [63] conjectured that the expected maximum cut size in a random regular graph is upper bounded by the solution to the one-step replica-symmetry breaking equations and provided numerical estimates of the resulting cut size. Corollary 1.3 matches their numbers.

Table 1 displays the upper bounds from Corollary 1.3 for $d = 3, \ldots, 10$. For comparison the table also contains the previous best rigorous upper bounds we are aware of, and the best rigorous lower bounds. Upper bounds appear to have received little attention. For $d > 3$ the upper bounds shown come from straightforward application of the first moment method, counting cuts of the given size; this can be done either by standard counting arguments or using [13, Corollary 2.8], in either case followed by a small numerical computation. For $d = 3$ better upper bounds come from the first moment method but restricting to cuts satisfying some local maximality conditions; the bound shown is from [61]. The lower bounds result from analyses of algorithms, and are from [31] via [20] for $d = 3$, [26] for $d = 4$ and [27] for $d > 4$.

The article of Zdeborová and Boettcher contains a second, more prominent conjecture that ties together the MIN BISECTION and MAX CUT problems on random regular graphs, namely that the two cases result w.h.p. in asymptotically equal numbers of edges 'dissatisfied' (respectively, cut and not cut). Unfortunately the methods of the present work do not appear to shed light on this question.

d	3	4	5	6	7	8	9	10
best previous upper bound	0.9320	0.8900	0.8539	0.8260	0.8038	0.7855	0.7701	0.7570
Corollary 1.3 upper bound	0.9241	0.8683	0.8350	0.8049	0.7851	0.7659	0.7523	0.7388
best lower bound	0.9067	0.8333	0.7989	0.7775	0.7571	0.7404	0.7263	0.7144
expected cut size at β^*	0.8536	0.7887	0.7500	0.7236	0.7041	0.6890	0.6768	0.6667
expected cut size at Gibbs uniqueness	0.7500	0.6667	0.6250	0.6000	0.5833	0.5714	0.5625	0.5556

TABLE 1. Bounds on the fraction of edges in a maximum cut of $\mathbb{G}(n,d)$.

However, the work does shed light on a different question of interest: as an application of Theorem 1.1 we can calculate the information-theoretic threshold of the disassortative stochastic block model.

1.5. **The stochastic block model.** Over the past decade the stochastic block model has become a prominent benchmark for Bayesian inference as well as graph clustering. The impressive literature on the model is surveyed in [2, 49]. Like the Ising model, the stochastic block model comes in two variants. In the assortative version edges are more likely join vertices with the same spin while in the disassortative model edges are more likely to occur between vertices with opposite spins. Thus, the disassortative variant resembles the Ising antiferromagnet.

Formally the d-regular disassortative stochastic block model is defined by way of the following experiment. Let $V_n = \{v_1, \ldots, v_n\}$ be a set of n vertices. In a first step we draw a spin assignment $\boldsymbol{\sigma}^* \in \{\pm 1\}^{V_n}$ uniformly at random. Subsequently we draw a d-regular graph $\mathbb{G}^* = \mathbb{G}^*(\boldsymbol{\sigma}^*)$ from the distribution

$$\mathbb{P}\left[\mathbb{G}^* = G \mid \boldsymbol{\sigma}^*\right] \propto \exp(-\beta \mathscr{H}_G(\boldsymbol{\sigma}^*)). \tag{1.12}$$

Thus, the probability that a given d-regular graph G comes up is proportional to the Boltzmann weight $\exp(-\beta\mathscr{H}_G(\boldsymbol{\sigma}^*))$ of the 'ground truth' $\boldsymbol{\sigma}^*$.

The obvious question is whether the bias introduced by (1.12) has a discernible impact on the distribution of the graph. In other words, is it possible to tell $\mathbb{G}^*$ apart from the 'null model' $\mathbb{G}$? To formalise this we use the Kullback-Leibler divergence of $\mathbb{G}^*$ from $\mathbb{G}$,

$$D_{\mathrm{KL}}\left(\mathbb{G}^* \| \mathbb{G}\right) = \sum_G \mathbb{P}\left[\mathbb{G}^* = G\right] \log \frac{\mathbb{P}\left[\mathbb{G}^* = G\right]}{\mathbb{P}\left[\mathbb{G} = G\right]}.$$

The Kullback-Leibler divergence is an information-theoretic potential that gauges the difference between random objects. Specifically, if $D_{\mathrm{KL}}(\mathbb{G}^* \| \mathbb{G}) = o(n)$ then extensive observables such as the maximum cut value or the logarithm of the partition function in the two random graph models are asymptotically equal [45]. By contrast, if $D_{\mathrm{KL}}(\mathbb{G}^* \| \mathbb{G}) = \Omega(n)$, then one can tell the two random graph models apart by calculating the partition function [17]. In particular, in the latter case there exists a (not necessarily efficient) algorithm $\mathtt{A}$ that given a graph G outputs $\mathtt{A}(G) \in \{0,1\}$ such that

$$\lim_{n\to\infty} \mathbb{P}\left[\mathtt{A}(\mathbb{G}) = 0\right] = \lim_{n\to\infty} \mathbb{P}\left[\mathtt{A}(\mathbb{G}^*) = 1\right] = 1. \tag{1.13}$$

Hence, $\mathtt{A}$, the essence of which is calculating the partition function, distinguishes the stochastic block model from the null model with high probability.

Theorem 1.4. *For any $d \geq 3$ the following are true.*

(i) If $\beta < \beta^*(d)$, then $\lim_{n\to\infty} D_{\mathrm{KL}}(\mathbb{G}^* \| \mathbb{G})/n = 0$ and $\lim_{n\to\infty} D_{\mathrm{KL}}(\mathbb{G} \| \mathbb{G}^*)/n = 0$.
(ii) If $\beta > \beta^*(d)$, then $\lim_{n\to\infty} D_{\mathrm{KL}}(\mathbb{G}^* \| \mathbb{G})/n > 0$ and $\lim_{n\to\infty} D_{\mathrm{KL}}(\mathbb{G} \| \mathbb{G}^*)/n > 0$.

By definition, then, β^* is the information-theoretic threshold of the stochastic block model.

2. TECHNIQUES

This section contains a survey of the proofs of the main results and the techniques they are based on. We begin with the proof of the first part of Theorem 1.1, which combines moment computations with a spatial mixing argument. To motivate this combination we first discuss the Erdős-Rényi case, in which a straightforward moment calculation does the trick. Subsequently we discuss the proof of the second part of Theorem 1.1, which relies on the connection between the Ising model and the stochastic block model. This connection also shows how Theorem 1.4 follows from Theorem 1.1. The final subsection then deals with the the proof of Theorem 1.2, based on the interpolation method.

2.1. **The second moment method.** To get started, we will compute the typical value of the Ising partition function using the method of moments for the Erdős-Rényi model. To this end, we reproduce the calculation by Mossel, Neeman and Sly [50] for the Erdős-Rényi model $\mathbb{G}_{\mathrm{ER}}$ where $m = dn/2$ edges are drawn uniformly at random. (We skip their supplementing of the second moment method with small subgraph conditioning for increased precision.) We will show why this does not directly extend to the random regular model.

For the first moment we simply obtain

$$\mathbb{E}\left[Z_{\mathbb{G}_{\mathrm{ER}},\beta}\right] = \sum_{\sigma\in\{\pm1\}^{V_n}} \mathbb{E}\left[\exp(-\beta\mathscr{H}_{\mathbb{G}_{\mathrm{ER}}}(\sigma))\right] = \sum_{\sigma\in\{\pm1\}^{V_n}} \left(1-\frac{1-\mathrm{e}^{-\beta}}{n^2}\sum_{i,j=1}^{n}\mathbf{1}\{\sigma_{v_i}=\sigma_{v_j}\}\right)^{m+o(n)}. \tag{2.1}$$

The second equality holds because the edges of $\mathbb{G}_{\mathrm{ER}}$ are asymptotically independent. A moment's reflection reveals that the expression in the braces is maximised by σ such that $\sum_i \sigma_{v_i} = o(n)$. Combinatorially this means that σ corresponds to an approximately balanced cut. Since there are $2^{n+o(n)}$ such σ, (2.1) yields

$$\mathbb{E}\left[Z_{\mathbb{G}_{\mathrm{ER}},\beta}\right] = 2^{n+o(n)}\left(\frac{1+\mathrm{e}^{-\beta}}{2}\right)^{dn/2} = \exp\left(n\left(\left(1-\frac{d}{2}\right)\log(2)+\frac{d}{2}\log(1+\mathrm{e}^{-\beta})+o(1)\right)\right). \tag{2.2}$$

Calculating the second moment is similarly straightforward. Indeed, we obtain

$$\begin{aligned}\mathbb{E}\left[Z^2_{\mathbb{G}_{\mathrm{ER}},\beta}\right] &= \sum_{\sigma,\sigma'\in\{\pm1\}^{V_n}} \mathbb{E}\left[\exp(-\beta\mathscr{H}_{\mathbb{G}_{\mathrm{ER}}}(\sigma)-\beta\mathscr{H}_{\mathbb{G}_{\mathrm{ER}}}(\sigma'))\right] \qquad (2.3)\\ &= \sum_{\sigma,\sigma'}\left(1-\frac{1}{n^2}\sum_{i,j=1}^{n}\left(1-\mathrm{e}^{-\beta}\right)\left(\mathbf{1}\{\sigma_{v_i}=\sigma_{v_j}\}+\mathbf{1}\{\sigma'_{v_i}=\sigma'_{v_j}\}\right)-\left(1-\mathrm{e}^{-\beta}\right)^2\mathbf{1}\{\sigma_{v_i}=\sigma_{v_j}\wedge\sigma'_{v_i}=\sigma'_{v_j}\}\right)^{m+o(n)}.\end{aligned}$$

As in the first moment calculation it is easy to see that asymptotically balanced σ,σ' dominate. Moreover, rearranging the sum according to the inner product $a = \sigma\cdot\sigma'$, we obtain

$$\mathbb{E}\left[Z^2_{\mathbb{G}_{\mathrm{ER}},\beta}\right] = \sum_{a=-n}^{n}\binom{n}{(n-a)/4,(n-a)/4,(n+a)/4,(n+a)/4}\left(\frac{(1+\mathrm{e}^{-\beta})^2}{4}+\left(\frac{a}{n}\right)^2\frac{(1-\mathrm{e}^{-\beta})^2}{4}\right)^{m+o(n)}. \tag{2.4}$$

Introducing $\alpha = a/n$ and the entropy function $H(p) = -p\log p-(1-p)\log(1-p)$ for $0<p<1$, we can apply Stirling's formula to simplify (2.4) to

$$\mathbb{E}\left[Z^2_{\mathbb{G}_{\mathrm{ER}},\beta}\right] = \exp\left(n\max_{-1<\alpha<1} f_d(\alpha,\beta)+o(n)\right), \qquad \text{where} \tag{2.5}$$

$$f_d(\alpha,\beta) = (1-d)\log(2)+H((1+\alpha)/2)+\frac{d}{2}\log\left((1+\mathrm{e}^{-\beta})^2+\alpha^2(1-\mathrm{e}^{-\beta})^2\right). \tag{2.6}$$

Substituting $\alpha = 0$ into (2.6) yields

$$f_d(0,\beta) = (2-d)\log(2)+d\log\left(1+\mathrm{e}^{-\beta}\right) \tag{2.7}$$

which is twice the exponent from (2.2). Hence, if $f_d(\alpha,\beta)$ attains its maximum at $\alpha = 0$, then (2.1) and (2.5) show that $\mathbb{E}[Z^2_{\mathbb{G}_{\mathrm{ER}},\beta}]/\mathbb{E}[Z_{\mathbb{G}_{\mathrm{ER}},\beta}]^2 = \exp(o(n))$. Routine concentration arguments therefore apply and show that $Z_{\mathbb{G}_{\mathrm{ER}},\beta}$ concentrates about its expectation. In particular, we obtain

$$\lim_{n\to\infty}\frac{1}{n}\mathbb{E}\left[\log Z_{\mathbb{G}_{\mathrm{ER}},\beta}\right] = \lim_{n\to\infty}\frac{1}{n}\log\mathbb{E}\left[Z_{\mathbb{G}_{\mathrm{ER}},\beta}\right] = \log 2+\frac{d}{2}\log\frac{1+\mathrm{e}^{-\beta}}{2} \qquad \text{if } \max_{-1<\alpha<1} f_d(\alpha,\beta) = f_d(0,\beta). \tag{2.8}$$

By contrast, if the maximum in (2.5) is attained at $\alpha\neq 0$, then the second moment exceeds the square of the first moment exponentially. Hence, the moment method succeeds iff $f_d(\alpha,\beta)$ attains its maximum at $\alpha = 0$.

Whether or not this is the case depends on the value of β. Specifically, for a given $d\geq 3$ the function $f_d(\alpha,\beta)$ is maximised at $\alpha = 0$, and (2.8) is satisfied, if

$$\beta\leq\beta^\dagger(d) = \log\frac{\sqrt{d}+1}{\sqrt{d}-1};$$

as mentioned above, this discovery belongs to Mossel, Neeman and Sly [50]. Note that $\beta^*(d) = \beta^\dagger(d-1) > \beta^\dagger(d)$. Conversely, results from [5, 50] imply that

$$\lim_{n\to\infty}\frac{1}{n}\mathbb{E}\left[\log Z_{\mathbb{G}_{\mathrm{ER}},\beta}\right] < \log 2+\frac{d}{2}\log\frac{1+\mathrm{e}^{-\beta}}{2} \qquad \text{for } \beta>\beta^\dagger(d). \tag{2.9}$$

Hence, $\beta^\dagger(d)$ marks the replica symmetry breaking threshold for the Ising model on the Erdős-Rényi graph.

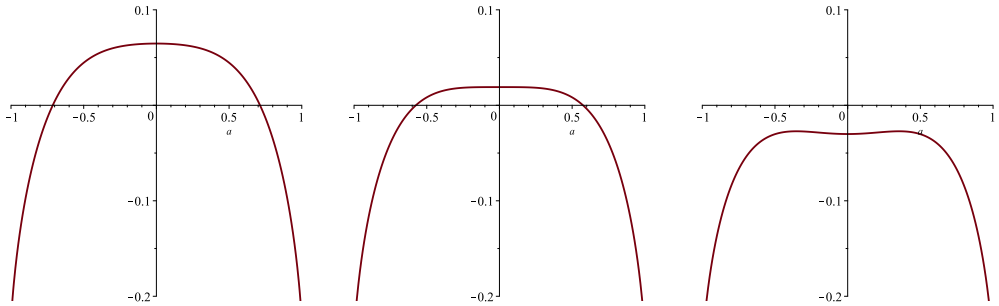

FIGURE 1. The function $f_d(\alpha,\beta)$ for $d=3$ and $\beta=1.25$ (left), $\beta=1.32$ (middle) and $\beta=1.40$ (right). For $d=3$, we have $\beta^\dagger\approx 1.32$ and $\beta^*\approx 1.76$.

To what extent do these considerations carry over to the random regular graph $\mathbb{G}(n,d)$? The following lemma shows that the first moment for random regular graphs is about the same as in the Erdős-Rényi case; the calculations, given in [50], are similar to those above.

Lemma 2.1. *For any $d\geq 3,\beta>0$ we have*

$$\mathbb{E}[Z_{\mathbb{G},\beta}]=\Theta\left(2^n\left(\frac{1+\mathrm{e}^{\beta}}{2}\right)^{dn/2}\right). \tag{2.10}$$

For the second moment, the expression from (2.5) for the Erdős-Rényi model carries over to the random regular graph and yields the upper bound

$$\mathbb{E}\left[Z^2_{\mathbb{G}(n,d),\beta}\right]\leq\exp\left(n\max_{-1<\alpha<1}f_d(\alpha,\beta)+o(n)\right). \tag{2.11}$$

The fact that the bound extends to random regular graphs may not appear entirely immediate; the analytic explanation derives from the convexity of the Kullback-Leibler divergence [1, 13]. A similar trick has been applied with some success to various random regular graph problems, notably graph colouring [1]. Unfortunately, in our case the second moment trick only yields the desired solution for $\beta<\beta^\dagger(d)$, while Theorem 1.1 requires it for all $\beta<\beta^*(d)$. Indeed, for $\beta>\beta^\dagger(d)$ the trick fails in a rather spectacular way: once β crosses above $\beta^\dagger(d)$ the value $\alpha=0$ turns from a global maximum of the function $f_d(\alpha,\beta)$ into a local minimum! Figure 1 provides an illustration. Thus, we have to turn to other means to establish the first part of Theorem 1.1, which we explore next.

2.2. **Broadcasting and non-reconstruction.** Our approach towards proving the first part of Theorem 1.1 relies on combining spatial mixing arguments with the method of moments. To be precise, we will exhibit an event $\mathcal{O}$ such that for all $\beta<\beta^*(d)$,

$$\mathbb{E}\left[Z_{\mathbb{G}(n,d),\beta}\mathbf{1}\{\mathcal{O}\}\right]=\Theta(\mathbb{E}\left[Z_{\mathbb{G}(n,d),\beta}\right])=\Theta\left(2^n\left(\frac{1+\mathrm{e}^{-\beta}}{2}\right)^{dn/2}\right),\quad \mathbb{E}\left[Z^2_{\mathbb{G}(n,d),\beta}\mathbf{1}\{\mathcal{O}\}\right]=4^{n+o(n)}\left(\frac{1+\mathrm{e}^{-\beta}}{2}\right)^{dn}. \tag{2.12}$$

Together with routine concentration arguments (2.12) will imply the first part of Theorem 1.1.

To elaborate, the event $\mathcal{O}$ concerns the relative location of two typical samples from the Boltzmann distribution. Hence, for a graph G let $\boldsymbol{\sigma}_G,\boldsymbol{\sigma}'_G$ denote two independent samples from $\mu_{G,\beta}$. Then for a sequence $\varepsilon_n=o(1)$ that tends to zero slowly enough (and that we will specify precisely in due course) we let

$$\mathcal{O}=\left\{\mathbb{E}\left[|\boldsymbol{\sigma}_{\mathbb{G}}\cdot\boldsymbol{\sigma}'_{\mathbb{G}}|\,|\,\mathbb{G}\right]<\varepsilon_n n\right\}. \tag{2.13}$$

Thus, $\mathcal{O}$ is the event that two typical samples from the Boltzmann distribution are nearly orthogonal. Since the combinatorial interpretation of α in (2.11) is to pinpoint the value of the inner product of spin configurations that renders the largest contribution, one might reasonably hope that conditioning on $\mathcal{O}$ will eliminate the need for taking values $\alpha\neq 0$ into consideration. Indeed, the proof of the second part of (2.12) will be relatively straightforward. Unfortunately, it turns out that the same cannot quite be said of the proof of the first part.

Proposition 2.2. *The event $\mathcal{O}$ from* (2.13) *satisfies* (2.12) *for all $d\geq 3$, $\beta\leq\beta^*(d)$.*

The proof of Proposition 2.2 uses two tools: the stochastic block model $\mathbb{G}^*$ from (1.12) and the analysis of a broadcasting process on the infinite d-regular tree from [10]. Specifically, the stochastic block model will help to derive the first part of (2.12). Indeed, the definition (1.12) suggests that the probability that $\mathbb{G}^*=G$ for

a given graph G should be roughly proportional to the partition function $Z_{G,\beta}$ (see e.g. [17]). This is because G has a chance proportional to $\sum_{\sigma\in\{\pm1\}^{V_n}}\left(\exp(-\beta\mathscr{H}_G(\sigma))/\sum_G\exp(-\beta\mathscr{H}_G(\sigma))\right)$, and if the denominators were the same over all σ, this would be proportional to $Z_{G,\beta}=\sum_\sigma\exp(-\beta\mathscr{H}_G(\sigma))$. By symmetry each denominator depends only on the magnetisation of σ (the sum of its entries), and summands with magnetisation near 0 are far more frequent, so it is reasonable to hope that they dominate the sum. Capitalising on this intuition, the following lemma shows that we can make use of $\mathbb{G}^*$ to establish the first part of (2.12).

Lemma 2.3. *Let $d\geq 3,\beta>0$. If $\mathbb{P}[\mathbb{G}^*\in\mathscr{O}]\sim 1$, then $\mathbb{E}\left[Z_{\mathbb{G},\beta}\mathbf{1}\{\mathscr{O}\}\right]=\Theta\left(\mathbb{E}\left[Z_{\mathbb{G},\beta}\right]\right)$.*

To show that $\mathbb{P}[\mathbb{G}^*\in\mathscr{O}]\sim 1$ we will couple the planted model $\mathbb{G}^*$ with a broadcasting process on the infinite $(d-1)$-ary tree $\mathbb{T}_{d-1}$. Let u_0 signify the (degree-d) root of $\mathbb{T}_{d-1}$. Proceeding down the tree, the broadcasting process constructs an assignment $\boldsymbol{\tau}\in\{\pm1\}^{V(\mathbb{T}_{d-1})}$ as follows. Initially we choose $\boldsymbol{\tau}_{u_0}\in\{\pm1\}$ uniformly at random. Subsequently, having defined $\boldsymbol{\tau}_u$ for all u at distance at most ℓ from u_0 already, we define the value $\boldsymbol{\tau}_w$ of a child w of such a vertex u by letting

$$\mathbb{P}[\boldsymbol{\tau}_w=\boldsymbol{\tau}_u\mid\boldsymbol{\tau}_u]=\mathrm{e}^{-\beta}/(1+\mathrm{e}^{-\beta}). \tag{2.14}$$

In words, w retains the spin of its parent with probability $\mathrm{e}^{-\beta}/(1+\mathrm{e}^{-\beta})$, and is assigned the opposite spin with the remaining probability $1/(1+\mathrm{e}^{-\beta})$. Let $\mathscr{T}_\ell$ denote the σ-algebra generated by the spins $\boldsymbol{\tau}_u$ of all vertices u at distance greater than ℓ from u_0. The following result shows that the spin τ_{v_0} decorrelates from $\mathscr{T}_\ell$ in the limit of large ℓ if $\beta<\beta^*(d)$. In other words, the broadcasting process 'forgets' the spin of the root, a property known as non-reconstruction [10].

Lemma 2.4 ([10])**.** *Let $d\geq 3$ and $\beta<\beta^*(d)$. Then*

$$\lim_{\ell\to\infty}\mathbb{E}\left|\mathbb{P}\left[\boldsymbol{\tau}_{v_0}=1\mid\mathscr{T}_\ell\right]-\frac{1}{2}\right|=0. \tag{2.15}$$

As an aside, $\beta^*(d)$ actually is the sharp non-reconstruction threshold [36]. Thus, (2.15) ceases to hold for $\beta>\beta^*(d)$.

Equipped with Lemma 2.4 the proof of the condition $\mathbb{E}\left[Z_{\mathbb{G},\beta}\mathbf{1}\{\mathscr{O}\}\right]\sim\mathbb{E}\left[Z_{\mathbb{G},\beta}\right]$ proceeds as follows. For a typical vertex of $\mathbb{G}^*$, say v_1, we couple the spins that the planted configuration $\boldsymbol{\sigma}^*$ assigns to vertices in the ℓ-ball around v_1 with the broadcasting process. This coupling is based on the fact that the random regular graph $\mathbb{G}^*$ converges to the d-regular tree in the Benjamini-Schramm topology. Then we re-sample the spins inside the ℓ-ball given the spins assigned to all the vertices at distance greater than ℓ from v_1 according to the Boltzmann distribution $\mu_{\mathbb{G}^*,\beta}$. Let $\boldsymbol{\sigma}^{**}$ denote the resulting spin configuration. Lemma 2.4 will enable us to conclude that the re-sampled spin $\boldsymbol{\sigma}^{**}_{v_1}$ is asymptotically independent from the original spin $\boldsymbol{\sigma}^*_{v_1}$. Finally, we will show that both $\boldsymbol{\sigma}^*$ and $\boldsymbol{\sigma}^{**}$ are distributed approximately as two samples from the Boltzmann distribution $\mu_{\mathbb{G}^*,\beta}$, thereby deriving the following.

Lemma 2.5. *Let $d\geq 3$ and $\beta<\beta^*(d)$. Then $\mathbb{P}[\mathbb{G}^*\in\mathscr{O}]\sim 1$.*

As shown in Section 4, Proposition 2.2 will be an easy consequence of Lemma 2.3 and Lemma 2.5, proved respectively in Sections 4.2 and 4.3. Moreover, the first part of Theorem 1.1 follows from Proposition 2.2 and a few lines of calculations; this is show just below.

The above argument highlights the difference between the Erdős-Rényi graph and the random regular graph and the reason why we have the strict inequality $\beta^\dagger(d)<\beta^*(d)$ for all $d\geq 3$. Indeed, in the d-regular tree, to which the random regular graph converges locally, every vertex has $d-1$ children. By contrast, the Erdős-Rényi graph of average degree d converges locally to a Galton-Watson tree with offspring distribution $\mathrm{Po}(d)$. Hence, the average number of children has mean d rather than $d-1$. The effect is that the broadcasting process on the Galton-Watson tree is able to remember the spin of the root for smaller values of β than in the regular case. Therefore, it is natural to expect that on the Erdős-Rényi graph long-range correlations emerge for smaller β.

Proof of Theorem 1.1, part 1. First, we argue by Azuma's inequality that for any $d\geq 3$ and $\beta>0$, $\log Z_{\mathbb{G},\beta}$ is concentrated about $\mathbb{E}[\log Z_{\mathbb{G},\beta}]$. As is standard, construct $\mathbb{G}$ using $dn/2$ independent random variables X_i each giving the matching of the next point in the configuration model. Compare this to a uniform random reference matching. If X_i matches a point A to B but the reference matching matched A to $C\neq B$ and B to D, update the reference by matching A to B and C to D. The reference copy has two edges added and two deleted, and with X_i uniformly random the reference remains uniformly random, so X_i changes the expectation of $\log Z_{\mathbb{G},\beta}$ conditioned on $X_1,\ldots,X_i$ by at most 2β. Azuma's inequality yields

$$\mathbb{P}\left[\left|\log Z_{\mathbb{G},\beta}-\mathbb{E}[\log Z_{\mathbb{G},\beta}]\right|>t\right]\leq 2\exp\left(-\frac{t^2}{4\beta^2dn}\right)\qquad(t>0). \tag{2.16}$$

For $\beta < \beta^*(d)$, Proposition 2.2 gives $\mathbb{E}[Z_{\mathbb{G},\beta}] = \Theta(\mathbb{E}[Z_{\mathbb{G},\beta}\mathbf{1}\{\mathscr{O}\}])$ while trivially $\mathbb{E}[Z^2_{\mathbb{G},\beta}] \le \mathbb{E}[Z^2_{\mathbb{G},\beta}\mathbf{1}\{\mathscr{O}\}]$, so from the Paley-Zygmund inequality,

$$\mathbb{P}\left[Z_{\mathbb{G},\beta} \ge \frac{1}{2}\mathbb{E}[Z_{\mathbb{G},\beta}]\right] \ge \frac{1}{4}\frac{\mathbb{E}[Z_{\mathbb{G},\beta}]^2}{\mathbb{E}[Z^2_{\mathbb{G},\beta}]} = \Omega(1)\frac{\mathbb{E}[Z_{\mathbb{G},\beta}\mathbf{1}\{\mathscr{O}\}]^2}{\mathbb{E}[Z^2_{\mathbb{G},\beta}\mathbf{1}\{\mathscr{O}\}]} = \Omega(\exp(-o(n))), \tag{2.17}$$

the last inequality again from Proposition 2.2. Thus, $\mathbb{P}\left[\log Z_{\mathbb{G},\beta} \ge \log\mathbb{E}[Z_{\mathbb{G},\beta}] - 2\right] = \Omega(\exp(-o(n)))$. For any $\varepsilon > 0$, were there arbitrarily large n for which $\log\mathbb{E}[Z_{\mathbb{G},\beta}] > \mathbb{E}[\log Z_{\mathbb{G},\beta}] + \varepsilon n$ this would contradict the result from Azuma. But by Jensen's inequality $\log\mathbb{E}[Z_{\mathbb{G},\beta}] \ge \mathbb{E}[\log Z_{\mathbb{G},\beta}]$, so $\mathbb{E}[\log Z_{\mathbb{G},\beta}] = \log\mathbb{E}[Z_{\mathbb{G},\beta}] + o(n)$. Taking the value of $\log\mathbb{E}[Z_{\mathbb{G},\beta}]$ from Lemma 2.1 (or Proposition 2.2) establishes the first part of Theorem 1.1.

□

2.3. **The Bethe free energy.** In the next step we prove the the second statement of Theorem 1.1. As discussed earlier, the probability that a given graph G comes up as the result $\mathbb{G}^*$ of the stochastic block model is (nearly) proportional to the partition function $Z_{G,\beta}$. Therefore, if the partition function $Z_{\mathbb{G},\beta}$ is tightly concentrated about its mean $\mathbb{E}[Z_{\mathbb{G},\beta}]$, then we might expect that the distribution of $\mathbb{G}^*$ and of the plain random d-regular graph are 'close'. By contrast, if $Z_{\mathbb{G},\beta}$ is not concentrated but prone to a lottery phenomenon where a few unlikely outcomes render a disproportionate contribution to $\mathbb{E}[Z_{\mathbb{G},\beta}]$, then we should expect that this discrepancy is exacerbated upon passing to size-biased model $\mathbb{G}^*$ as outliers receive an extra boost. The following lemma formalises this intuition. It replaces the vague 'concentration' phrasing with asymptotic equality of $\mathbb{E}\left[\log Z_{\mathbb{G},\beta}\right]$ and $\log\mathbb{E}\left[Z_{\mathbb{G},\beta}\right]$, and the equivalent for $\mathbb{G}^*$, with $\log\mathbb{E}\left[Z_{\mathbb{G},\beta}\right]$ known from (2.10); this equality certainly follows from sufficient concentration, while without concentration the equality would be an odd coincidence.

Lemma 2.6 ([16, Lemma 4.4])**.** *Let $d \ge 3$ and $\beta > 0$. We have $\Phi_d(\beta) = \log 2 + \frac{d}{2}\log\frac{1+\mathrm{e}^{-\beta}}{2}$ if and only if*

$$\lim_{n\to\infty}\frac{1}{n}\mathbb{E}\left[\log Z_{\mathbb{G}^*,\beta}\right] = \log 2 + \frac{d}{2}\log\frac{1+\mathrm{e}^{-\beta}}{2}. \tag{2.18}$$

By the first part of Theorem 1.1 the lemma's hypothesis holds for $\beta < \beta^*(d)$. To prove the second part of Theorem 1.1 we will show that the conclusion (and thus the hypothesis) is violated if $\beta > \beta^*(d)$. As a stepping stone we use a variational formula for $\mathbb{E}[\log Z_{\mathbb{G}^*,\beta}]$ from [16]. Let $\mathscr{P}_*([-1,1])$ be the space of all probability measures on the interval $[-1,1]$ with mean zero. Moreover, for a given such probability measure π let $(\boldsymbol{\mu}_{\pi,i})_{i\ge1}$ be a family of independent samples from π and let $\Lambda(x) = x\log x$. The expression

$$\mathscr{B}_{\mathrm{Ising}}(\pi,\beta,d) = \mathbb{E}\left[\frac{\Lambda\left(\sum_{\sigma\in\{\pm1\}}\prod_{i=1}^{d}1-(1-e^{-\beta})(1+\sigma\boldsymbol{\mu}_{\pi,i})/2\right)}{2^{1-d}(1+\mathrm{e}^{-\beta})^d} - \frac{d\Lambda\left(1-(1-\mathrm{e}^{-\beta})(1+\boldsymbol{\mu}_{\pi,1}\boldsymbol{\mu}_{\pi,2})/2\right)}{1+\mathrm{e}^{-\beta}}\right] \tag{2.19}$$

is called the Bethe free energy.

Lemma 2.7 ([16, Theorem 2.3])**.** *For any $\beta > 0$ and any $d \ge 3$, we have*

$$\lim_{n\to\infty}\frac{1}{n}\mathbb{E}[\log Z_{\mathbb{G}^*,\beta}] = \sup_{\pi\in\mathscr{P}_*([-1,1])}\mathscr{B}_{Ising}(\pi,\beta,d).$$

Combining Lemmas 2.6 and 2.7, we see that the second part of Theorem 1.1 boils down to showing that

$$\sup_{\pi\in\mathscr{P}_*(\{\pm1\})}\mathscr{B}_{\mathrm{Ising}}(\pi,\beta,d) > \log 2 + \frac{d}{2}\log\frac{1+\mathrm{e}^{-\beta}}{2} \qquad \text{for } \beta > \beta^*(d). \tag{2.20}$$

Luckily, the variational formula (2.20) asks to take a supremum over distributions π. Therefore, it suffices to point to a specific distribution π such that $\mathscr{B}_{\mathrm{Ising}}(\pi,\beta,d)$ exceeds the first moment bound. Specifically, for a small $\varepsilon > 0$ let us introduce

$$\pi^*_\varepsilon = \frac{1}{2}(\delta_{2\varepsilon} + \delta_{-2\varepsilon}) \tag{2.21}$$

where δ_z is the point mass on z. It is easy to see that for $\varepsilon = 0$ we precisely obtain $\mathscr{B}_{\mathrm{Ising}}(\pi^*_0,\beta,d) = \log(2) + d\log((1+\mathrm{e}^{-\beta})/2)/2$. The following proposition shows that for $\beta > \beta^*(d)$ small $\varepsilon > 0$ yield a slightly but strictly larger value.

Proposition 2.8. *For any $\beta > \beta^*(d)$ there exists $\varepsilon > 0$ such that*

$$\mathscr{B}_{Ising}(\pi^*_\varepsilon,\beta,d) > \log 2 + \frac{d}{2}\log\frac{1+\mathrm{e}^{-\beta}}{2}.$$

Proof of Theorem 1.1, part 2. This follows directly from Lemma 2.7 and Proposition 2.8. □

Proof of Theorem 1.4. The theorem follows from Theorem 1.1, Theorem 17.1 in [16] and the fact that the disassortative stochastic block model with two communities is the planted Ising antiferromagnet. □

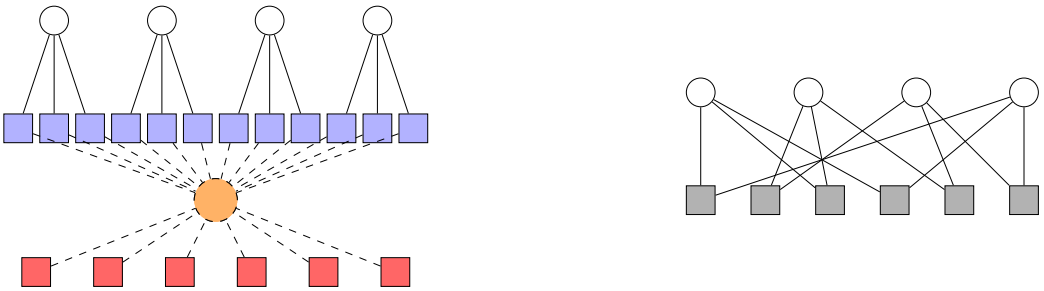

FIGURE 2. The factor graphs $\mathbb{G}_0$ and the original factor graph $\mathbb{G}_1$ with $n = 4, d = 3$.

2.4. **The interpolation method.** With Theorem 1.1 in place, let us consider the case that $\beta \to \infty$ which eventually allows us to derive improved upper bounds on the expected maximum cut size of random regular graphs. Unfortunately, the stochastic dependencies between vertex spins make it difficult to get a handle on a simple expression like we had for $\beta < \beta^*(d)$ where we simply obtained the first moment bound. Apart from the obvious short-range dependencies that, for example, induce adjacent vertices to prefer opposite spins, we expect long-range dependencies to occur above the Kesten-Stigum bound. Thus, for $\beta > \beta^*(d)$ the spin of a vertex impacts those of distant vertices.

The '1-step replica symmetry breaking ansatz' from physics attempts to describe these long-range dependencies by means of an additional hidden variable [42, 46]. The basic hypothesis is that for $\beta > \beta^*(d)$ the phase space, i.e., the set $\{\pm 1\}^{V_n}$ of all possible spin configurations, decomposes into a number $\boldsymbol{S}_1, \ldots, \boldsymbol{S}_\ell$ of 'pure states' w.h.p. Mathematically $\boldsymbol{S}_1, \ldots, \boldsymbol{S}_\ell$ are pairwise disjoint subsets of $\{\pm 1\}^{V_n}$ such that $\mu_{\mathbb{G},\beta}(\boldsymbol{S}_1) + \cdots + \mu_{\mathbb{G}.\beta}(\boldsymbol{S}_\ell) \sim 1$. Thus, the sets cover nearly the entire support of the Boltzmann distribution. Furthermore, once we condition on a pure state $\boldsymbol{S}_h$, long-range effects disappear; formally,

$$\mathbb{E}\left|\sum_{h=1}^{\ell} \mu_{\mathbb{G},\beta}(\boldsymbol{S}_h)\left(\mathbb{P}\left[\boldsymbol{\sigma}_{\mathbb{G},v_1} = s, \boldsymbol{\sigma}_{\mathbb{G},v_2} = s' \mid \mathbb{G}, \boldsymbol{\sigma}_{\mathbb{G}} \in \boldsymbol{S}_h\right] - \mathbb{P}\left[\boldsymbol{\sigma}_{\mathbb{G},v_1} = s \mid \mathbb{G}, \boldsymbol{\sigma}_{\mathbb{G}} \in \boldsymbol{S}_h\right]\mathbb{P}\left[\boldsymbol{\sigma}_{\mathbb{G},v_2} = s' \mid \mathbb{G}, \boldsymbol{\sigma}_{\mathbb{G}} \in \boldsymbol{S}_h\right]\right)\right| = o(1).$$

Hence, long-range correlations of the unconditional measure $\mu_{\mathbb{G},\beta}$ arise because the spin $\boldsymbol{\sigma}_{\mathbb{G},v_1}$ of a vertex hints at the pure state $\boldsymbol{S}_h$ to which $\boldsymbol{\sigma}_{\mathbb{G}}$ belongs, which in turn skews our expectations as to the other spins. The existence of such a pure state decomposition has been established rigorously [15].

How does this picture help to estimate the partition function $Z_{\mathbb{G},\beta}$? The basic idea behind the interpolation method is to set up an synthetic model of a spin system that exhibits precisely the long-range dependencies predicted by the 1-step rsb ansatz, and no others. Mathematically this model is represented by a factor graph (or a Markov random field); see the left panel of Figure 2. The factor graph contains variable nodes (the white circles) that represent the vertices of our graph. Each of these variable nodes is connected to an external field (the blue box) that is meant to represent the impact of the short-range dependencies imposed by one of the incident edges of the corresponding vertex of $\mathbb{G}$. But instead of the complicated direct interactions between the vertices through actual edges as in the original random graph $\mathbb{G}$, the variable nodes only interact with each other through the yellow node. This node represents the hidden variable postulated by the 1-step rsb ansatz, i.e., the index of the pure state. Finally, the red boxes are 'negative edges'. They are necessary because the variable nodes do not interact directly. In effect, the number of blue nodes is twice the number of edges of the actual graph $\mathbb{G}$, and thus we have to compensate for the impact of $dn/2$ spare blue boxes.

The cunning idea behind the interpolation method is to build a family of factor graph models parametrised by time $t \in [0,1]$. The interpolation scheme starts from the artificial factor graph model at time $t = 0$. At each intermediate time step $t \in (0,1)$ the model blends the synthetic $t = 0$ case and the actual Ising model on $\mathbb{G}$. Ultimately at time $t = 1$ all the synthetic ingredients (the blue and red boxes) disappeared and we are left with just the Ising antiferromagnet on $\mathbb{G}$. Remarkably, it is possible to prove that the partition function decreases monotonically in terms of t. As a consequence, the partition function of the synthetic model upper bounds that of the Ising model on $\mathbb{G}$. Fortunately we do not need to carry out the interpolation method in full. The result that we need follows from a more general version of the interpolation bound derived in [60].

To state the resulting upper bound precisely, fix any probability measure $\mathfrak{r}$ on $[-1,1]$. Let $(\boldsymbol{r}_i)_{i \geq [d]}$ be a family of independent random variables with distribution $\mathfrak{r}$; thus, $\boldsymbol{r}_i \in [-1,1]$ for all i. Further, define

$$\boldsymbol{\rho}_i(\sigma) = \frac{1 + \sigma \boldsymbol{r}_i}{2} \qquad (\sigma = \pm 1).$$

The idea is that $\boldsymbol{\rho}_1,\dots,\boldsymbol{\rho}_d$ represent the short-range influences that the neighbours of some vertex, say v_1, exercise on the spin of that vertex within a single pure state. More specifically, think of $\boldsymbol{\rho}_i(s)$ as the probability that the i-th neighbour of v_1 would take spin $s \in \{\pm 1\}$ if we removed v_1 from the random graph. The following lemma is an immediate consequence of [60, Theorem E.5].

Lemma 2.9. *Let $d \geq 3, \beta > 0$. Then for any $y > 0$ and any $\mathfrak{r} \in \mathscr{P}([-1,1])$ we have $\Phi_d(\beta) \leq \phi_{\beta,y}(\mathfrak{r})$, where*

$$\phi_{\beta,y}(\mathfrak{r}) = \frac{1}{y}\log\mathbb{E}[X_1^y] - \frac{d}{2y}\log\mathbb{E}[X_2^y], \tag{2.22}$$
$$X_1 = \sum_{\tau\in\{\pm1\}}\prod_{h=1}^{d} 1-(1-\mathrm{e}^{-\beta})\boldsymbol{\rho}_h(\tau), \qquad X_2 = 1-(1-\mathrm{e}^{-\beta})\sum_{\tau\in\{\pm1\}}\boldsymbol{\rho}_1(\tau)\boldsymbol{\rho}_2(\tau).$$

Clearly, (2.22) is not exactly what we had in mind when aiming for an explicit expression of the upper bound for $\Phi_d(\beta)$. However, a key feature of Lemma 2.9 is that the inequality holds for *any* $y, \mathfrak{r}$. We are thus free to choose these parameters so that we obtain a reasonable expression and, hopefully, at the same time a good upper bound.

Following physics intuition [46, 47] we define the measure $\mathfrak{r}$ as follows. Let $\delta_x \in \mathscr{P}([-1,1])$ be the atom on $x \in [-1,1]$. Then for $\alpha \in [0,1/2]$ we let

$$\mathfrak{r}_\alpha = \alpha\delta_{-1} + (1-2\alpha)\delta_0 + \alpha\delta_1 \in \mathscr{P}([-1,1]). \tag{2.23}$$

Intuitively, we 'freeze' a spin to $+1$ or -1 with probability α. Otherwise, if the spin does not freeze we leave it unbiased, i.e., it takes either spin ± 1 with equal probability. The following proposition shows that for the distribution $\mathfrak{r}$ from (2.23) the function $\phi_{\beta,y}(\mathfrak{r})$ boils down to a manageable expression.

Proof of Theorem 1.2. The theorem is an immediate consequence of Lemma 2.9 and Proposition 2.10. □

What does the bound look like in the trivial case $\alpha = 0$? For any $y > 0$ we obtain

$$\Phi_d(\beta) \leq \phi_{\beta,y}(\mathfrak{r}_0) \leq \frac{1}{y}\left(\log\mathbb{E}[X_1^y] - \frac{d}{2}\log\mathbb{E}[X_2^y]\right) = \frac{1}{y}\log\left(2\left(\frac{1+\mathrm{e}^{-\beta}}{2}\right)^d\right)^y - \frac{d}{2y}\log\left(\frac{1+\mathrm{e}^{-\beta}}{2}\right)^y \tag{2.24}$$
$$= \log 2 + \frac{d}{2}\log\frac{1+\mathrm{e}^{-\beta}}{2}. \tag{2.25}$$

Hence, we simply recover the first moment bound (1.5). However, for large β the strictly positive α render a better bound. The following proposition simplifies the expression from Theorem 1.2 for large β. Recall F from (1.8).

Proposition 2.10. *Let $d \geq 3, \beta > 0, 0 < z < 1, 0 < \alpha < 1/2$. Then with $y = y(\beta) = -\log(z)/\beta$ we have*

$$\lim_{\beta\to\infty}\frac{1}{\beta y}\left(\log\mathbb{E}\left[X_1^y\right] - \frac{d}{2}\log\mathbb{E}[X_2^y]\right) = F_d(\alpha,z).$$

Proof of Corollary 1.3. For any d-regular graph G on n vertices and any $\beta > 0$, we have

$$\frac{2}{dn}\textsc{MaxCut}(G) = 1 - \frac{2}{dn}\min_{\sigma\in\{\pm1\}^n}\mathscr{H}_G(\sigma) \leq 1 + \frac{2}{\beta dn}\log Z_\beta(G).$$

Thus by Theorem 1.2, we obtain

$$\limsup_{n\to\infty}\frac{2}{dn}\mathbb{E}[\textsc{MaxCut}(\mathbb{G})] \leq 1 + \frac{2}{d}\lim_{\beta\to\infty}\Phi_d(\beta)/\beta \tag{2.26}$$
$$\leq 1 + \frac{2}{d}\inf_{\substack{0<\alpha\leq 1/2\\ 0<z<1}} F_d(\alpha,z). \tag{2.27}$$

The corollary follows. □

2.5. **Organisation.** Let us outline how the remainder of the paper is organized. Section 4 is devoted to proving Proposition 2.2, while Section 5 contains the proof of Proposition 2.8. Jointly, the two sections provide the missing pieces for the proof of Theorem 1.1. Finally, in Section 6 we prove the two key Lemmas 6.1 and 6.2 needed for the proof of Corollary 1.3.

2.6. **Notation.** We will denote a random d-regular graph on n vertices by $\mathbb{G}(n,d)$. When the context is clear, we will simply write $\mathbb{G} = \mathbb{G}(n,d)$. We tacitly assume that dn is even. Throughout the paper we will use standard Landau notation with the usual symbols $o(\cdot), O(\cdot), \Theta(\cdot), \omega(\cdot)$ and $\Omega(\cdot)$. These symbols refer to the limit $n \to \infty$ by default, but may refer to other limits where specified.

For a subset $I \subset \mathbb{R}$ we denote by $\mathscr{P}(I)$ the set of all Borel probability measures on I. Moreover, for a finite set $\Omega \neq \emptyset$ let $\mathscr{P}(\Omega)$ be the set of all probability distributions on Ω. We recall that the entropy $H(\mu)$ of such a probability distribution $\mu \in \mathscr{P}(\Omega)$ is defined as

$$H(\mu) = -\sum_{\omega \in \Omega} \mu(\omega) \log \mu(\omega).$$

We will also need the Kullback-Leibler divergence of $\mu, \nu \in \mathscr{P}(\Omega)$, defined as

$$D_{\mathrm{KL}}(\mu \| \nu) = \sum_{\omega \in \Omega} \mu(\omega) \log \frac{\mu(\omega)}{\nu(\omega)} \in [0, \infty],$$

with the conventions $0 \log 0 = 0$, $0 \log \frac{0}{0} = 0$ and $-\log 0 = \infty$.

3. Discussion

In this section we relate the contributions of the present paper to prior work. We begin with the statistical physics perspective.

3.1. **Replica symmetry breaking.** The Ising model, proposed by Lenz in 1920 [43], has become a cornerstone of statistical physics generally [30, 37]. Moreover, the Ising model on random graphs in particular has proved a testbed for the investigation of the idea of replica symmetry breaking that was proposed by Mézard and Parisi on the basis of the non-rigorous 'cavity method' [46, 47]. The corroboration of the cavity method's predictions for the ferromagnetic Ising model on Erdős-Rényi graphs by Dembo and Montanari [22] was a first success, although replica symmetry breaking does not occur in this model. The proof was based on the analysis of the Belief Propagation recurrences on random trees. These techniques have subsequently been extended to the Potts model, a generalisation of the Ising model with more than two possible spin values [23, 24].

The antiferromagnetic version of the Potts and Ising models is closely related to the stochastic block model. The results of Mossel, Neeman and Sly [52] on the block model with two communities therefore imply the existence and location of a replica symmetry breaking phase transition in the Ising antiferromagnet on the Erdős-Rényi graph. Thus, as we saw above a new contribution of the present paper is the extension to random regular graphs. Moreover, results from Coja-Oghlan, Krzakala, Perkins and Zdeborová [14] imply the existence and location of a replica symmetry breaking phase transition for the Potts model on Erdős-Rényi graphs. The recent work of Coja-Oghlan, Hahn-Klimroth, Loick, Müller, Panagiotou and Pasch [16] extend these results to graphs with given degree sequences. However, the results from [16] determine the location of the replica symmetry breaking phase transition only implicitly as the solution to an infinite-dimensional variational problem. Thus, the contribution of Theorem 1.1 is the explicit analytic formula for the phase transition $\beta^*(d)$, which matches the combinatorially meaningful Kesten-Stigum bound [40].

Apart from the Potts and Ising models, replica symmetry breaking phase transitions have been pinpointed in several other models. Examples include random (hyper)graph colouring, several other random constraint satisfaction problems and further models from mathematical physics, such as the Viana-Bray spin glass model [34]. But usually the formula for the phase transition comes in as a complicated variational problem. Indeed, the question whether the replica symmetry breaking transition equals the explicit Kesten-Stigum threshold has been linked to the order of the phase transition [59], a question that merits further rigorous attention.

3.2. **The** Max Cut **problem.** The semidefinite programming based Max Cut algorithm of Goemans and Williamson [32] has been one of the most important contributions to algorithms research. The algorithm achieves an approximation ratio of $\min_{0 \le \theta \le \pi} \frac{2}{\pi} \frac{\theta}{1-\cos(\theta)} \approx 0.878$ on graphs with non-negative edge weights. On regular graphs better approximation ratios can be achieved (also via semidefinite programming) [28]. The question whether the Goemans-Williamson approximation ratio is optimal has sparked an important line of research. Håstad [35] derived from the PCP theorem that no approximation better than 0.941 can be attained unless P=NP. Moreover, Koth [41] showed that the unique games conjecture implies the optimality of the Goemans-Williamson approximation ratio; see Barak [6] for a discussion.

Given the great interest in Max Cut generally, it is hardly surprising that the problem has been studied intensively on random graphs, too. In the classical combinatorics literature upper bounds have typically been based on the first moment method, while greedy algorithms were employed to derive lower bounds [9, 12, 18, 19, 26, 27, 38, 39]. A semidefinite programming approach was taken in [48] on Erdős-Rényi graphs. Table 1 summarises the best explicit prior bounds for random regular graphs. Naturally, arguments based

on the method of moments or greedy algorithms suffer from the shortcoming of being inherently local, i.e., confined to short-range interactions. In effect, they remain oblivious to the long-range interactions that, according to physics prediction, shape the MAX CUT problem on random graphs. Therefore, it is unsurprising that these techniques only carry so far.

The first complex model where the long-range interactions predicted by the theory of replica symmetry breaking were well understood is the Sherrington-Kirkpatrick spin glass. The model can be viewed as a weighted MAX CUT problem on a complete graph. Specifically, the weight of the edge between vertices v, w is a Gaussian $\boldsymbol{J}_{v,w}$. The random variables $(\boldsymbol{J}_{v,w})_{1\le v<w\le n}$ are mutually independent. Hence, the model is described by the random Hamiltonian

$$\mathscr{H}_{\mathrm{SK}}(\sigma) = -\frac{1}{\sqrt{n}}\sum_{1\le i<j\le n}\boldsymbol{J}_{ij}\sigma_i\sigma_j \qquad (\sigma\in\{\pm1\}^n),$$

which induces a partition function and a Boltzmann distribution as in (1.2). Clearly, the 'ground state energy' $\min_\sigma \mathscr{H}_{\mathrm{SK}}(\sigma)$ corresponds to the maximum cut weight. Parisi's seminal work [58] predicted formulas for the free energy and the ground state energy of the Sherrington-Kirkpatrick model. After several decades Talagrand established the 'Parisi formula' rigorously [62]. An important ingredient to this work was the interpolation method, which Guerra had proposed [33]. Panchenko developed a different argument [56], which also led to a proof of Parisi's ultrametricity conjecture [55].

Franz and Leone [29] extended the interpolation method to sparse random graphs; see also [54]. The version of the interpolation method quoted in Lemma 2.9 is an adaptation to random regular graphs. Furthermore, Dembo, Montanari and Sen [25] used interpolation techniques to strike a chord between the sparse Erdős-Rényi graph and the Sherrington-Kirkpatrick model. Specifically, they proved that

$$\lim_{d\to\infty}\lim_{n\to\infty}\frac{2}{\sqrt{d}n}\left[\textsc{MaxCut}(\mathbb{G}(n,d/n))-\frac{dn}{4}\right]=p_\star\approx0.7632, \tag{3.1}$$

where $p_\star$ derives from the ground state energy of the Sherrington-Kirkpatrick model. Conceptually it seems natural to expect that the Sherrington-Kirkpatrick model occurs as the limit of sparse random graphs as the average degree gets large, basically due to central limit theorem-like effects. Yet this result says nothing about any finite d and, indeed, sparse random graphs for fixed finite values of d appear to exhibit a more diverse and potentially even more intricate behavior. As a result, we are only just beginning to understand the genuine behaviour of sparse models in the replica symmetry breaking phase; see, e.g., [7].

Finally, Panchenko [57] obtained a variational formula for the free energy of the Ising antiferromagnet on Erdős-Rényi graphs. The formula involves an optimisation over exchangeable distributions on $\{\pm1\}^{\mathbb{N}\times\mathbb{N}}$ subject to certain invariance conditions. Coja-Oghlan and Perkins [15] extended this result to random regular graphs, also pointing out that a corresponding variational formula can be derived for the MAX CUT of $\mathbb{G}(n,d)$ for any fixed d. However, the formula is not explicit, and it appears difficult (to put it mildly) to extract any numerical estimates. Thus, the contribution of Corollary 1.3 is that we obtain a (relatively) simple explicit formula that incorporates at least the first level of the physicists' replica symmetry breaking formalism.

3.3. **The stochastic block model.** The Ising antiferromagnet is intimately related to the stochastic block model which has gained significant attention in recent years [2]. The model provides a benchmark for both Bayesian inference and graph clustering, the basic idea being to create a random graph with a community structure. In the simplest version the vertex set is partitioned into q communities and edges between vertices in the same community are either more likely (assortative) or less (disassortative). The question is for what discrepancy of edge densities it is possible to at least partially recover the community structure or, less ambitiously, to at least discriminate a random graph drawn from the block model from a null model. The modern study of the stochastic block model originated with conjectures that Decelle, Krzakala, Moore and Zdeborová [21] derived via the cavity method. Specifically, they predicted a phase diagram that splits the model parameters into regions where recovering the community structure is information-theoretically and/or algorithmically feasible.

Mathematically the most complete picture exists for graphs with independent edges in the case of $q=2$ communities. In this case the information-theoretic and algorithmic thresholds were established in a series of papers by Mossel, Neeman and Sly [50, 51, 52, 53] and Massoulie [44]. For $q>2$ communities algorithms that match the conjecture from [21] have been proposed by Abbe and Sandon [3] and Bordenave, Lelarge and Massoulie [11]. As explained above, the contribution of Theorem 1.4 is to show that for the disassortative regular case with two communities the information-theoretic threshold equals the explicit Kesten-Stigum bound $\beta^*(d)$. Finally, as an interesting direction for future research we point to the question of developing an efficient algorithm that (partially) recovers the community structure $\boldsymbol{\sigma}^*$ for $\beta>\beta^*(d)$.

4. Proof of Proposition 2.2

The proof of Proposition 2.2 requires several steps. First we perform some preparatory calculations; in particular, we compute the first moment of the partition function of a random multi-graph drawn from the pairing model. Subsequently we establish a relationship between the stochastic block model $\mathbb{G}^*$ and the 'null model' $\mathbb{G}$. Then we construct a coupling of the spin configuration around a typical vertex of $\mathbb{G}^*$ with the broadcasting process from Lemma 2.4 to estimate the probability that $\mathbb{G}^* \in \mathscr{O}$. Finally, we perform a truncated moment computation to obtain (2.12).

4.1. **The pairing model.** In order to calculate the first moment, as well as for some of the manoeuvres to follow, it will be convenient to replace the simple random d-regular graph $\mathbb{G}$ by a random graph chosen from the pairing model. Hence, think of the elements of $V_n \times [d]$ as vertex clones. Moreover, let $\mathbf{\Gamma}$ be a random perfect matching of the complete graph on $V_n \times [d]$. Finally, let $\mathbf{G}$ be the d-regular multigraph on V_n obtained by contracting the clones $V_n \times [d]$. With $\mathscr{S}$ the set of all simple graphs, it is well known that

$$\mathbb{P}[\mathbf{G} \in \mathscr{S}] = \Omega(1) \qquad \text{and} \qquad \mathbb{P}[\mathbb{G} \in \mathscr{E}] = \mathbb{P}[\mathbf{G} \in \mathscr{E} \mid \mathscr{S}] \qquad \text{for any event } \mathscr{E}. \tag{4.1}$$

In order to compute the first moment $\mathbb{E}[Z_{\mathbb{G},\beta}]$ we will compute $\mathbb{E}[Z_{\mathbf{G},\beta}]$ and then investigate the impact of conditioning on $\mathscr{S}$.

To calculate $\mathbb{E}[Z_{\mathbf{G},\beta}]$ we proceed as follows. For $\sigma \in \{\pm 1\}^{V_n}$ let $\rho(\sigma) = (\rho_1(\sigma), \rho_{-1}(\sigma))$ be the distribution on ± 1 defined by

$$\rho_1(\sigma) = \frac{1}{n}\sum_{i=1}^{n} \mathbf{1}\{\sigma_{v_i} = 1\}, \qquad \rho_{-1}(\sigma) = \frac{1}{n}\sum_{i=1}^{n} \mathbf{1}\{\sigma_{v_i} = -1\}.$$

Thanks to the linearity of expectation we can write the first moment as

$$\mathbb{E}[Z_{\mathbf{G},\beta}] = \sum_{\sigma \in \{\pm 1\}^{V_n}} \mathbb{E}[\psi_{\mathbf{G},\beta}(\sigma)].$$

Naturally, $\psi_{\mathbf{G},\beta}(\sigma)$ depends on the number of edges that join vertices with the same spin. Hence, to calculate $\mathbb{E}[\psi_{\mathbf{G},\beta}(\sigma)]$ we need to know the number of graphs with a given number of such edges.

The following lemma solves this problem. Let $\mathscr{M}(\sigma)$ be the set of all probability distributions

$$\mu_{11} + \mu_{1-1} = \rho_1, \qquad \mu_{-1-1} + \mu_{1-1} = \rho_{-1}, \qquad \mu_{1-1} = \mu_{-11} \tag{4.2}$$

and such that $\mu_{11}dn, \mu_{-1-1}dn$ are even integers and $\mu_{1-1}dn$ is an integer. Moreover, let $\mathscr{G}(\sigma,\mu)$ be the event that $\mathbf{G}$ has $\mu_{11}dn/2$ edges that join vertices v, w with $\sigma_v, \sigma_w = 1$. Then due to regularity there are $\mu_{-1-1}dn/2$ edges joining vertices that both carry a -1 spin and $\mu_{1-1}dn$ edges that connect vertices with opposite spins.

Lemma 4.1. *For $\sigma \in \{\pm 1\}^{V_n}$ and $\mu \in \mathscr{M}(\sigma)$ we have*

$$\mathbb{P}\left[\mathbf{G} \in \mathscr{G}(\sigma,\mu)\right] = \binom{dn\rho_1(\sigma)}{dn\mu_{11}}\binom{dn\rho_{-1}(\sigma)}{dn\mu_{-1-1}}\frac{(dn\mu_{11}-1)!!(dn\mu_{-1-1}-1)!!(dn\mu_{1-1})!}{(dn-1)!!}.$$

Proof. The denominator $(dn-1)!!$ simply counts the total number of possible perfect matchings $\mathbf{\Gamma}$. Moreover, the two binomial coefficients account for the number of ways of selecting clones of vertices with spin ± 1 to constitute edges of the four possible types. Finally, the numerator equals the number of possible ways to match the clones up according to these designated types. □

As it stands the formula from Lemma 4.1 does not yet lend itself to asymptotical calculations. But Stirling's formula yields the following approximation.

Corollary 4.2. *For $\sigma \in \{\pm 1\}^{V_n}$ and $\mu \in \mathscr{M}(\sigma)$ we have* $\mathbb{P}\left[\mathbf{G} \in \mathscr{G}(\sigma,\mu)\right] = \exp\left(-\frac{dn}{2}D_{\mathrm{KL}}\left(\mu \| \rho \otimes \rho\right) + O(\log n)\right)$.

Corollary 4.2 follows from a more general lemma about partitions of random regular graphs from [13]. But since we will encounter similar calculations again in due course and because the proof is quite short, we include it here. We need Stirling's formula

$$k! = \sqrt{2\pi k}\left(\frac{k}{\mathrm{e}}\right)^k \exp(O(1/k)) \tag{4.3}$$

and the elementary formula

$$(2k-1)!! = \frac{(2k)!}{k!2^k}. \tag{4.4}$$

Proof of Corollary 4.2. Applying (4.4), we obtain

$$(dn\mu_{11}-1)!! = \frac{(dn\mu_{11})!}{2^{dn\mu_{11}/2}(dn\mu_{11}/2)!} \quad (dn\mu_{-1-1}-1)!! = \frac{(dn\mu_{-1-1})!}{2^{dn\mu_{-1-1}/2}(dn\mu_{-1-1}/2)!}, \quad (dn-1)!! = \frac{(dn)!}{2^{dn/2}(dn/2)!}.$$

Hence,

$$\begin{aligned}\frac{(dn\mu_{11}-1)!!(dn\mu_{-1-1}-1)!!(dn\mu_{1-1})!}{(dn-1)!!} &= 2^{dn(1-\mu_{11}-\mu_{-1-1})/2}\binom{dn\mu_{1-1}}{dn\mu_{1-1}/2}^{-1}\binom{dn/2}{dn\mu/2}\binom{dn}{dn\mu}^{-1}\\ &= 2^{dn\mu_{1-1}}\binom{dn\mu_{1-1}}{dn\mu_{1-1}/2}^{-1}\binom{dn/2}{dn\mu/2}\binom{dn}{dn\mu}^{-1}. \end{aligned} \tag{4.5}$$

Thus, Stirling's formula (4.3) gives

$$\frac{(dn\mu_{11}-1)!!(dn\mu_{-1-1}-1)!!(dn\mu_{1-1})!}{(dn-1)!!} = \exp\bigl(-dnH(\mu)/2+O(\log n)\bigr). \tag{4.6}$$

Further, combining (4.2) and (4.3), we obtain

$$\binom{dn\rho_1(\sigma)}{dn\mu_{11}}\binom{dn\rho_{-1}(\sigma)}{dn\mu_{-1-1}} = \exp\bigl(dn(H(\mu)-H(\rho(\sigma))+O(\log n)\bigr). \tag{4.7}$$

Finally, combining Lemma 4.1 with (4.6) and (4.7), we obtain

$$\begin{aligned}\mathbb{P}\left[\mathbf{G}\in\mathscr{G}(\sigma,\mu)\right] &= \exp\bigl(dn(H(\mu)-2H(\rho(\sigma)))/2+O(\log n)\bigr) = \exp\bigl(dn(H(\mu)-H(\rho(\sigma)\otimes\rho(\sigma)))/2+O(\log n)\bigr)\\ &= \exp\bigl(-dnD_{\mathrm{KL}}\bigl(\mu\|\rho(\sigma)\otimes\rho(\sigma)\bigr)/2+O(\log n)\bigr),\end{aligned}$$

as claimed. □

Let $\mathscr{M}_n = \bigcup_{\sigma\in\{\pm1\}^{V_n}}\mathscr{M}(\sigma)$ be the set of all conceivable distributions μ. Moreover, for $\mu\in\mathscr{M}_n$ set $\rho_1(\mu) = \mu_{11}+\mu_{1-1}$ and $\rho_{-1}(\mu) = 1-\rho_1(\mu)$. Additionally, let $\mu^* = \mu^*_\beta$ be the distribution

$$\mu^*_{11} = \mu^*_{-1-1} = \frac{1}{2(1+\mathrm{e}^\beta)}, \qquad \mu^*_{1-1} = \mu^*_{-11} = \frac{\mathrm{e}^\beta}{2(1+\mathrm{e}^\beta)}. \tag{4.8}$$

Furthermore, let $\mathscr{M}^*_n$ be the set of all $\mu\in\mathscr{M}_n$ such that $d_{\mathrm{TV}}(\mu,\mu^*) < n^{-0.49}$. Finally, let $\mathscr{G}(\mu)$ be the set of all pairs (G,σ) such that $\sigma\in\{\pm1\}^{V_n}$ satisfies $\rho_1(\sigma) = \rho_1(\mu)$ and $G\in\mathscr{G}(\sigma,\mu)$. The following lemma supplies the promised formula for the first moment of $Z_{\mathbf{G},\beta}$.

Lemma 4.3. *For all $d\geq3, \beta>0$ we have*

$$\mathbb{E}[Z_{\mathbf{G},\beta}] = (1+\exp(-n^{\Omega(1)}))\sum_{\mu\in\mathscr{M}^*_n}|\mathscr{G}(\mu)|\exp\left(-\frac{dn}{2}\bigl(\mu_{11}+\mu_{-1-1}\bigr)\right) = \Theta\left(2^n\left(\frac{1+\mathrm{e}^{-\beta}}{2}\right)^{dn/2}\right). \tag{4.9}$$

Proof. For a given $\mu\in\mathscr{M}_n$ the total number of $\sigma\in\{\pm1\}^{V_n}$ with $\mu\in\mathscr{M}(\sigma)$ equals $\binom{n}{\rho_1(\mu)n}$. Therefore, Corollary 4.2 and (4.3) yield

$$\begin{aligned}\mathbb{E}[Z_{\mathbf{G},\beta}] &= \sum_{\mu\in\mathscr{M}_n}|\mathscr{G}(\mu)|\exp\left(-\frac{dn}{2}\bigl(\mu_{11}+\mu_{-1-1}\bigr)\right)\\ &= \sum_{\mu\in\mathscr{M}_n}\binom{n}{\rho_1(\sigma)n}\exp\left(-\frac{dn}{2}\left[D_{\mathrm{KL}}\bigl(\mu\|\rho(\mu)\otimes\rho(\mu)\bigr)+\beta\bigl(\mu_{11}+\mu_{-1-1}\bigr)\right]+O(\log n)\right) &(4.10)\\ &= \max_{\mu\in\mathscr{M}_n}\exp\left(n\left[H(\rho(\mu))-\frac{d}{2}D_{\mathrm{KL}}\bigl(\mu\|\rho(\mu)\otimes\rho(\mu)\bigr)-\frac{d\beta}{2}\bigl(\mu_{11}+\mu_{-1-1}\bigr)\right]+O(\log n)\right). &(4.11)\end{aligned}$$

Due to the linear relations (4.2) we can view the expression inside the square brackets, i.e.,

$$\varphi_{d,\beta}(\mu) = H(\rho(\mu))-\frac{d}{2}D_{\mathrm{KL}}\bigl(\mu\|\rho(\mu)\otimes\rho(\mu)\bigr)-\frac{d\beta}{2}\bigl(\mu_{11}+\mu_{-1-1}\bigr), \tag{4.12}$$

as a function of the two variables μ_{11} and μ_{-1-1}. The function is strictly concave because the entropy function is strictly concave and the Kullback-Leibler divergence is convex. Hence, the unique stationary point of $\varphi_{d,\beta}$ is its maximiser. Since the derivatives of $\varphi_{d,\beta}$ work out to be

$$\frac{\partial\varphi_{d,\beta}}{\partial\mu_{11}} = \frac{d-1}{2}\log\frac{\rho_1(\mu)}{\rho_{-1}(\mu)}+\frac{d}{2}\log\frac{\mu_{1-1}}{\mu_{11}}-\frac{d\beta}{2}, \qquad \frac{\partial\varphi_{d,\beta}}{\partial\mu_{-1-1}} = \frac{1-d}{2}\log\frac{\rho_1(\mu)}{\rho_{-1}(\mu)}+\frac{d}{2}\log\frac{\mu_{1-1}}{\mu_{-1-1}}-\frac{d\beta}{2},$$

the stationary point occurs at μ^*. Substituting the solution (4.8) into (4.11), we obtain

$$\mathbb{E}[Z_{\mathbf{G},\beta}] = \exp\left(n\left[\log 2 + \frac{d}{2}\log\frac{1+\mathrm{e}^{-\beta}}{2}\right] + O(\log n)\right). \tag{4.13}$$

as well as the first equality sign in (4.9). To obtain the second part of (4.9) we take another look at Lemma 4.1, which shows together with Stirling's formula that there exists $c = c(d,\beta)$ such that

$$|\mathscr{G}(\mu)|\frac{\exp\left(-\frac{\beta dn}{2}(\mu_{11}+\mu_{-1-1})\right)}{(dn-1)!!} = \frac{c}{n}\exp\bigl(-n\varphi_{d,\beta}(\mu)\bigr) \qquad \text{uniformly for all } \mu \in \mathscr{M}_n^*. \tag{4.14}$$

Since the function $\varphi_{d,\beta}$ is strictly concave, (4.14) shows together with the first part of (4.9) and the Laplace method that

$$\mathbb{E}[Z_{\mathbf{G},\beta}] = \Theta(\exp(-n\varphi_{d,\beta}(\mu^*))) = \Theta\left(2^n\left(\frac{1+\mathrm{e}^{-\beta}}{2}\right)^{dn/2}\right),$$

which completes the proof. □

Having calculated $\mathbb{E}[Z_{\mathbf{G},\beta}]$ sufficiently accurately, we proceed to extend this formula to the simple random graph $\mathbb{G}$ and to the truncated first moment $\mathbb{E}[Z_{\mathbb{G},\beta}\mathbf{1}\{\mathscr{O}\}]$. Fortunately we can kill these two birds with one stone.

4.2. **The truncated first moment.** We need to calculate truncated first moments of the form $\mathbb{E}[Z_{\mathbf{G},\beta}\mathbf{1}_{\mathscr{A}}]$ for some event $\mathscr{A}$. To this end we define a pairing model variant of the stochastic block model. In analogy to (1.12) we draw $\boldsymbol{\sigma}^* \in \{\pm1\}^{V_n}$ uniformly at random. Further, given $\boldsymbol{\sigma}^*$ for any possible outcome G of $\mathbf{G}$ we let

$$\mathbb{P}\left[\mathbf{G}^* = G \mid \boldsymbol{\sigma}^*\right] \propto \exp(-\beta\mathscr{H}_G(\boldsymbol{\sigma}^*)). \tag{4.15}$$

The following lemma will enable us to reduce the task of computing $\mathbb{E}[Z_{\mathbf{G},\beta}\mathbf{1}_{\mathscr{A}}]$ for an event $\mathscr{A}$ to estimating the probability of $\mathbf{G}^* \in \mathscr{A}$. Similar lemmas have been known for other random problems since the work of Achlioptas and Coja-Oghlan [4].

Lemma 4.4. *Let $d \geq 3, \beta > 0$ and let $\mathscr{E}$ be a set of graph/spin configuration pairs. Then*

$$\frac{1}{\mathbb{E}[Z_{\mathbf{G},\beta}]}\sum_{\sigma\in\{\pm1\}^{V_n}}\mathbb{E}[\psi_{\mathbf{G},\beta}(\sigma)\mathbf{1}\{(\mathbf{G},\sigma)\in\mathscr{E}\}] = \Theta(\mathbb{P}[(\mathbf{G}^*,\boldsymbol{\sigma}^*)\in\mathscr{E}]) + o(1).$$

Proof. The definition (4.15) of $\mathbf{G}^*$ ensures that

$$\sum_{\sigma\in\{\pm1\}^{V_n}}\mathbb{E}[\psi_{\mathbf{G},\beta}(\sigma)\mathbf{1}\{(\mathbf{G},\sigma)\in\mathscr{E}\}] = \sum_{\sigma\in\{\pm1\}^{V_n}}\mathbb{P}[(\mathbf{G}^*,\boldsymbol{\sigma}^*)\in\mathscr{E}\mid\boldsymbol{\sigma}^*=\sigma]\mathbb{E}[\psi_{\mathbf{G},\beta}(\sigma)]. \tag{4.16}$$

We now split the above sums up into three parts: for a small $\varepsilon > 0$ pick $C > 0$ large and let

$$S = \left\{\sigma\in\{\pm1\}^{V_n} : |\rho_1(\sigma) - 1/2| \leq Cn^{-1/2}\right\}, \quad S' = \left\{\sigma\in\{\pm1\}^{V_n}\setminus S : |\rho_1(\sigma)-1/2| \leq n^{-0.49}\right\}, \quad S'' = \{\pm1\}^{V_n}\setminus(S\cup S').$$

Then Lemma 4.3 implies that

$$\sum_{\sigma\in S''}\mathbb{E}[\psi_{\mathbf{G},\beta}(\sigma)] = o(\mathbb{E}[Z_{\mathbf{G},\beta}]). \tag{4.17}$$

In fact, (4.14) implies that for large enough C,

$$\sum_{\sigma\in S'}\mathbb{E}[\psi_{\mathbf{G},\beta}(\sigma)] \leq \varepsilon\mathbb{E}[Z_{\mathbf{G},\beta}]. \tag{4.18}$$

In addition, (4.14) implies together with the fact that μ^* is the unique stationary point of the concave function $\varphi_{d,\beta}$ that

$$\mathbb{E}[\psi_{\mathbf{G},\beta}(\sigma)] = O(2^{-n}\mathbb{E}[Z_{\mathbf{G},\beta}]) \qquad \text{uniformly for all } \sigma\in\{\pm1\}^{V_n}, \tag{4.19}$$

$$\mathbb{E}[\psi_{\mathbf{G},\beta}(\sigma)] = \Theta(2^{-n}\mathbb{E}[Z_{\mathbf{G},\beta}]) \qquad \text{uniformly for all } \sigma\in S. \tag{4.20}$$

Furthermore, because $\boldsymbol{\sigma}^*$ is uniformly random we can choose C so large that

$$\sum_{\sigma\in\{\pm1\}^{V_n}}\mathbb{P}[(\mathbf{G}^*,\boldsymbol{\sigma}^*)\in\mathscr{E}\mid\boldsymbol{\sigma}^*=\sigma] \leq \varepsilon + \sum_{\sigma\in S}\mathbb{P}[(\mathbf{G}^*,\boldsymbol{\sigma}^*)\in\mathscr{E}\mid\boldsymbol{\sigma}^*=\sigma]. \tag{4.21}$$

Combining (4.16)–(4.21) and taking $\varepsilon\to0$ slowly, we obtain the assertion. □

As an immediate consequence of Lemma 4.4 we obtain the following.

Corollary 4.5. *For all $d \geq 3, \beta > 0$ and for any event $\mathscr{A}$ the following two statements are true.*

(i) If $\mathbb{P}[\mathbf{G}^*\in\mathscr{A}] = \Omega(1)$, then $\mathbb{E}[Z_{\mathbf{G},\beta}\mathbf{1}_{\mathscr{A}}] = \Theta(\mathbb{E}[Z_{\mathbf{G},\beta}])$.

(ii) We have $\mathbb{P}[\mathbf{G}^*\in\mathscr{A}] = 1-o(1)$ iff $\mathbb{E}[Z_{\mathbf{G},\beta}\mathbf{1}_{\mathscr{A}}] \sim \mathbb{E}[Z_{\mathbf{G},\beta}]$.

As an application of Corollary 4.5 we will compute $\mathbb{E}[Z_{\mathbb{G},\beta}] = \mathbb{E}[Z_{\mathbf{G},\beta}\mathbf{1}\mathscr{S}]$. To this end we need bound the probability of the event $\mathbf{G}^* \in \mathscr{S}$ away from zero.

Lemma 4.6. *For all $d \geq 3, \beta > 0$ we have $\mathbb{P}[\mathbf{G}^* \in \mathscr{S}] = \Omega(1)$.*

Proof. Following the well known proof that $\mathbb{P}[\mathbf{G} \in \mathscr{S}] = \Omega(1)$, we will use the method of moments. Thus, fix $\mu \in \mathscr{M}_n^*$ and $\sigma \in \{\pm 1\}^{V_n}$ with $\rho_1(\sigma) = \rho_1(\mu)$. Let X be the number of self-loops of $\mathbf{G}$ and let Y be the number of double-edges. We will show that for any fixed integers $k, \ell \geq 1$,

$$\mathbb{E}\left[\prod_{j=1}^{k}(X-j+1)\prod_{j=1}^{\ell}(Y-j+1)\right] \sim \kappa^k\lambda^\ell \qquad \text{with } \kappa = \frac{d-1}{\mathrm{e}^\beta+1},\ \lambda = \frac{(d-1)^2(1+\mathrm{e}^{2\beta})}{2(1+\mathrm{e}^\beta)^2}. \tag{4.22}$$

Clearly (4.22) implies that $\mathbb{P}[\mathbf{G}^* \in \mathscr{S}] = \mathbb{P}[X = Y = 0] \sim \exp(-\kappa - \lambda) = \Omega(1)$.

To verify (4.22) we start by computing the means of X, Y. To be more precise, let X_1 be the number of self-loops at vertices v_i with $\sigma_{v_i} = 1$. In order to construct a self-loop we need to pick a vertex and two of its clones and calculate the probability that these clones get matched. Thus, the number of choices equals $\binom{d}{2}\rho_1(\sigma)n$. Therefore, Lemma 4.1 and (4.8) yield

$$\mathbb{E}[X_1 \mid \mathscr{G}(\sigma,\mu)] = \frac{\binom{d}{2}\rho_1(\sigma)n\binom{dn\rho_1(\sigma)-2}{dn\mu_{11}-2}\binom{dn\rho_{-1}(\sigma)}{dn\mu_{-1-1}}(dn\mu_{11}-3)!!(dn\mu_{-1-1}-1)!!(dn\mu_{1-1})!}{\binom{dn\rho_1(\sigma)}{dn\mu_{11}}\binom{dn\rho_{-1}(\sigma)}{dn\mu_{-1-1}}(dn\mu_{11}-1)!!(dn\mu_{-1-1}-1)!!(dn\mu_{1-1})!} \sim \frac{\kappa}{2}. \tag{4.23}$$

Because (4.8) ensures that $\rho_1(\sigma) \sim 1/2$, (4.23) implies that

$$\mathbb{E}[X \mid \mathscr{G}(\sigma,\mu)] \sim \kappa. \tag{4.24}$$

Similar considerations yield the mean of Y. Specifically, we decompose Y into Y_{11}, Y_{-1-1} and Y_{1-1}, which, respectively, count double-edges among vertices assigned spin 1, among vertices with spin -1, and between vertices with different spins. To work out Y_{11} we need to select two vertices with spin 1, two clones of each and a perfect matching. Thus, the number of choices comes to $2\binom{\rho_1(\sigma)n}{2}\binom{d}{2}^2$. Hence, Lemma 4.1 and (4.8) yield

$$\begin{aligned}\mathbb{E}[Y_{11} \mid \mathscr{G}(\sigma,\mu)] &= \frac{2\binom{\rho_1(\sigma)n}{2}\binom{d}{2}^2\binom{dn\rho_1(\sigma)-4}{dn\mu_{11}-4}\binom{dn\rho_{-1}(\sigma)}{dn\mu_{-1-1}}(dn\mu_{11}-5)!!(dn\mu_{-1-1}-1)!!(dn\mu_{1-1})!}{\binom{dn\rho_1(\sigma)}{dn\mu_{11}}\binom{dn\rho_{-1}(\sigma)}{dn\mu_{-1-1}}(dn\mu_{11}-1)!!(dn\mu_{-1-1}-1)!!(dn\mu_{1-1})!}\\ &\sim \frac{(d-1)^2\mu_{11}^2}{4\rho_1(\sigma)^2} \sim \frac{(d-1)^2}{4(\mathrm{e}^\beta+1)^2}.\end{aligned} \tag{4.25}$$

The same calculation applies to the mean of Y_{-1-1}. Moreover, analogously we obtain

$$\begin{aligned}\mathbb{E}[Y_{1-1} \mid \mathscr{G}(\sigma,\mu)] &= \frac{2\rho_1(\sigma)\rho_{-1}(\sigma)n^2\binom{d}{2}^2\binom{dn\rho_1(\sigma)-2}{dn\mu_{11}}\binom{dn\rho_{-1}(\sigma)-2}{dn\mu_{-1-1}}(dn\mu_{11}-1)!!(dn\mu_{-1-1}-1)!!(dn\mu_{1-1}-2)!}{\binom{dn\rho_1(\sigma)}{dn\mu_{11}}\binom{dn\rho_{-1}(\sigma)}{dn\mu_{-1-1}}(dn\mu_{11}-1)!!(dn\mu_{-1-1}-1)!!(dn\mu_{1-1})!}\\ &\sim \frac{(d-1)^2\mu_{1-1}^2}{2\rho_1\rho_{-1}} = \frac{(d-1)^2\mathrm{e}^{2\beta}}{2(1+\mathrm{e}^\beta)^2}.\end{aligned} \tag{4.26}$$

Combining (4.25) and (4.26), we obtain

$$\mathbb{E}[Y \mid \mathscr{G}(\sigma,\mu)] \sim \lambda. \tag{4.27}$$

The calculations that we performed towards (4.24) and (4.27) easily extend to a proof of (4.22). Indeed, instead of just accounting for the choice of placing a single double-edge or loop, we need to place fixed numbers k, ℓ. Since k, ℓ remain bounded as $n \to \infty$, the probability that any choices overlap is $O(1/n)$. Therefore, the joint factorial moment of X, Y works out to be $\kappa^k\lambda^\ell$, which is (4.22). □

Proof of Lemma 2.1. The lemma follows from Lemma 4.3, Corollary 4.5 and Lemma 4.6. □

Proof of Lemma 2.3. This is an immediate consequence of Corollary 4.5 and Lemma 4.6. □

4.3. **Coupling with the broadcasting process.** In this section we are going to establish a coupling of the local structure of $\mathbf{G}^*$ around a given vertex v_i with the broadcasting process from Lemma 2.4. Specifically, we are going to prove the following statement.

Lemma 4.7. *For any $d \geq 3, \beta > 0$ there exists $\varepsilon_n = o(1)$ such that the event $\mathscr{O}$ from (2.13) satisfies $\mathbb{E}[Z_{\mathbf{G},\beta}\mathbf{1}\{\mathscr{O}\}] \sim \mathbb{E}[Z_{\mathbf{G},\beta}]$.*

We begin the proof of of Lemma 4.7 by showing that the bounded-depth neighbourhoods in $\mathbf{G}^*$ are typically acyclic.

Lemma 4.8. *Let $d \geq 3$, $\beta > 0$. Moreover, for an integer $\ell \geq 1$ let C_ℓ be the number of cycles of length ℓ in $\mathbf{G}^*$. Then for any fixed integer L we have $\sum_{\ell \leq L} C_\ell = O(\log n)$ w.h.p.*

Proof. By Lemma 4.3 and Corollary 4.5 we may condition on the event $\mathcal{G}(\mu)$ for some $\mu \in \mathcal{M}_n^*$ and on the event $|\rho_1(\boldsymbol{\sigma}^*)| \sim 1/2$. A cycle of length ℓ passes through (not necessarily distinct) vertices $\boldsymbol{u} = (u_1, \ldots, u_\ell)$. For each step of the cycle we select a clone i_t where the cycle enters and one $j_t \neq i_t$ where it leaves. Set $\boldsymbol{i} = (\boldsymbol{i}_1, \ldots, \boldsymbol{i}_t)$ and $\boldsymbol{j} = (\boldsymbol{j}_1, \ldots, \boldsymbol{j}_t)$. However, we overcounted by a factor of 2ℓ (for the direction and the choice of the starting point). Given these choices let e_{11} be the number of edges of the cycle that connect two vertices of spin 1 under $\boldsymbol{\sigma}^*$ and define e_{-1-1} similarly. Moreover, let e_{1-1} be the number of cycle edges that join vertices with different spins. Following Lemma 4.1 we estimate the probability of the event $\mathscr{C}(\boldsymbol{u}, \boldsymbol{i}, \boldsymbol{j})$ that the specified cycle actually appears in $\mathbf{G}^*$ by

$$\begin{aligned}
\mathbb{P}\left[\mathscr{C}(\boldsymbol{u}, \boldsymbol{i}, \boldsymbol{j}) \mid \mathcal{G}(\mu), \boldsymbol{\sigma}^*\right] &\sim \binom{dn\rho_1(\sigma) - 2e_{11} - e_{1-1}}{dn\mu_{11} - 2e_{11}} \binom{dn\rho_{-1}(\sigma) - 2e_{-1-1} - e_{1-1}}{dn\mu_{-1-1} - 2e_{-1-1}} \binom{dn\rho_1(\sigma)}{dn\mu_{11}}^{-1} \binom{dn\rho_{-1}(\sigma)}{dn\mu_{-1-1}}^{-1} \\
&\quad \cdot \frac{(dn\mu_{11} - 2e_{11} - 1)!!(dn\mu_{-1-1} - 2e_{-1-1} - 1)!!(dn\mu_{1-1} - e_{1-1})!}{(dn\mu_{11} - 1)!!(dn\mu_{-1-1} - 1)!!dn\mu_{1-1}!} \\
&\sim (dn)^{-\ell} \left(\frac{\mu_{11}}{\rho_1(\boldsymbol{\sigma}^*)^2}\right)^{e_{11}} \left(\frac{\mu_{-1-1}}{\rho_{-1}(\boldsymbol{\sigma}^*)^2}\right)^{e_{-1-1}} \left(\frac{\mu_{1-1}}{\rho_1(\boldsymbol{\sigma}^*)\rho_{-1}(\boldsymbol{\sigma}^*)}\right)^{e_{1-1}} \\
&\sim \left(\frac{2}{dn(\mathrm{e}^\beta + 1)}\right)^\ell \mathrm{e}^{\beta e_{1-1}} \qquad \text{[due to (4.8)]}. \qquad (4.28)
\end{aligned}$$

Since the total number of choices for $\boldsymbol{u}, \boldsymbol{i}, \boldsymbol{j}$ is bounded by $n^\ell \binom{d}{2}^\ell$, (4.28) implies that $\mathbb{E}[C_\ell \mid \mathcal{G}(\mu), \boldsymbol{\sigma}^*] = O(1)$. Therefore, the assertion follows from Markov's inequality. □

For a vertex v of $\mathbf{G}^*$ and an integer $\ell \geq 0$ let $\boldsymbol{\sigma}^*_{v,\ell}$ be the spin configuration that $\boldsymbol{\sigma}^*$ induces on the vertices at distance at most ℓ from v. Furthermore, let $\boldsymbol{\tau}_\ell, \boldsymbol{\tau}'_\ell$ be two independent copies of the spin configuration that the broadcasting process from Section 2.2 induces on the vertices of the infinite d-regular tree $\mathbb{T}_d$ at distance at most ℓ from its root.

Lemma 4.9. *For any $d \geq 3, \beta > 0, \ell \geq 0$ the spin configurations $\boldsymbol{\sigma}^*_{v_1,\ell}$ and $\boldsymbol{\tau}_\ell$ have total variation distance $o(1)$.*

Proof. Thanks to Lemma 4.3 and Corollary 4.5 we may condition on the event $\mathcal{G}(\mu)$ for a $\mu \in \mathcal{M}_n^*$ and on $|\rho_1(\boldsymbol{\sigma}^*)| \sim 1/2$. Moreover, due to Lemma 4.8 we may confine ourselves to the case that the depth-ℓ neighbourhood of v_1 is acyclic. Let T be a possible outcome of the depth-ℓ neighbourhood of v_1 under these assumptions. Moreover, let $e_{11}, e_{-1-1}, e_{1-1}$ be the numbers of edges of T that join vertices both assigned spin 1 under $\boldsymbol{\sigma}^*$, or both assigned spin -1, or assigned different spins, respectively. Further, set $e = e_{11} + e_{1-1} + e_{-1-1}$. Finally, let $\mathscr{E}(T)$ be the event that T occurs in $\mathbf{G}^*$. Then Lemma 4.1 and (4.8) show that

$$\begin{aligned}
\mathbb{P}\left[\mathscr{E}(T) \mid \mathcal{G}(\mu), \boldsymbol{\sigma}^*\right] &= \binom{dn\rho_1(\boldsymbol{\sigma}^*) - 2e_{11} - e_{1-1}}{dn\mu_{11} - 2e_{11}} \binom{dn\rho_1(\boldsymbol{\sigma}^*)}{dn\mu_{11}}^{-1} \\
&\quad \cdot \binom{dn\rho_{-1}(\boldsymbol{\sigma}^*) - 2e_{-1-1} - e_{1-1}}{dn\mu_{-1-1} - 2e_{-1-1}} \binom{dn\rho_{-1}(\boldsymbol{\sigma}^*)}{dn\mu_{-1-1}}^{-1} \\
&\quad \cdot \frac{(dn\mu_{11} - 2e_{11} - 1)!!(dn\mu_{-1-1} - 2e_{-1-1} - 1)!!(dn\mu_{1-1} - e_{1-1})!}{(dn\mu_{11} - 1)!!(dn\mu_{-1-1} - 1)!!(dn\mu_{1-1})!} \\
&\sim \left(\frac{2}{dn}\right)^{2e} \left(\frac{\mathrm{e}^\beta}{1 + \mathrm{e}^\beta}\right)^{e_{1-1}} \left(\frac{1}{1 + \mathrm{e}^\beta}\right)^{e_{11} + e_{-1-1}}. \qquad (4.29)
\end{aligned}$$

Hence, (4.29) shows that the probability of observing a given spin assignment $\boldsymbol{\sigma}^*_{v_1,\ell}$ depends only on the number of edges joining vertices with the same spin, and that this dependence is precisely the same as in the case (2.14) of the broadcasting process. Thus, $\boldsymbol{\sigma}^*_{v_1,\ell}$ and $\boldsymbol{\tau}_\ell$ have total variation distance $o(1)$. □

For a graph G, a vertex v of G and an integer $\ell > 0$ let $\partial^\ell(G, v)$ be the set of vertices at distance precisely ℓ from v. Further, for a spin configuration $\chi \in \{\pm 1\}^{V(G)}$ let

$$\mu_{G,\beta,v,\ell}(s \mid \chi) = \frac{\sum_{\sigma \in \{\pm 1\}^{V(G)}} \mathbf{1}\{\sigma_v = s, \forall u \in \partial^\ell(G, v) : \sigma_u = \chi_u\} \exp(-\beta \mathscr{H}_G(\sigma))}{\sum_{\sigma \in \{\pm 1\}^{V(G)}} \mathbf{1}\{\forall u \in \partial^\ell(G, v) : \sigma_u = \chi_u\} \exp(-\beta \mathscr{H}_G(\sigma))} \qquad (s = \pm 1);$$

in words, this is the conditional Boltzmann marginal of v given the 'boundary condition' χ at the vertices at distance precisely ℓ from v.

Corollary 4.10. *For any $d \geq 3, \beta > 0, \varepsilon > 0$ there exists $\ell > 0$ such that $\mathbb{E}\sum_{i=1}^{n} \left|\mu_{\mathbf{G}^*,\beta,v_i,\ell}(1|\boldsymbol{\sigma}^*) - \frac{1}{2}\right| < \varepsilon n$.*

Proof. Because the random pair $(\mathbf{G}^*, \boldsymbol{\sigma}^*)$ is invariant under vertex permutations, we have

$$\mathbb{E}\sum_{i=1}^{n}\left|\mu_{\mathbf{G}^*,\beta,v_i,\ell}(1|\boldsymbol{\sigma}^*) - \frac{1}{2}\right| = n\mathbb{E}\left|\mu_{\mathbf{G}^*,\beta,v_i,\ell}(1|\boldsymbol{\sigma}^*) - \frac{1}{2}\right|$$

and Lemma 2.4 and Lemma 4.9 show that the r.h.s. gets small in the limit of large ℓ. □

Proof of Lemma 4.7. We apply Corollary 4.10 to a function $\varepsilon'_n = o(1)$ that tends to zero slowly. Specifically, let $X(\mathbf{G}^*, \boldsymbol{\sigma}^*) = \sum_{i=1}^{n} \mathbf{1}\{|\mu_{\mathbf{G}^*,\beta,v_i,\ell}(1|\boldsymbol{\sigma}^*) - \frac{1}{2}| > \varepsilon'\}$. Corollary 4.10 implies together with Markov's inequality that $\mathbb{P}[X(\mathbf{G}^*, \boldsymbol{\sigma}^*) > \varepsilon'' n] \leq \varepsilon''$ for a suitable $1 \ll \ell = o(\log n)$, provided that $\varepsilon', \varepsilon'' = o(1)$ tend to zero slowly enough. Hence, Lemma 4.4 shows that

$$\mathbb{E}[Z_{\mathbb{G},\beta}\mathbb{P}[X(\mathbb{G}, \boldsymbol{\sigma}_{\mathbb{G}}) \leq \varepsilon'' n \mid \mathbb{G}]] \sim \mathbb{E}[Z_{\mathbb{G},\beta}]. \tag{4.30}$$

We claim that (4.30) implies that there exists $n^{-1/4} \ll \delta = o(1)$ such that

$$\mathbb{E}[Z_{\mathbb{G},\beta}\mathbf{1}\{\mathbb{P}[X(\mathbb{G}, \boldsymbol{\sigma}_{\mathbb{G}}) > \delta n \mid \mathbb{G}] < \delta\}] \sim \mathbb{E}[Z_{\mathbb{G},\beta}]. \tag{4.31}$$

Indeed, (4.30) shows that for a suitable δ,

$$\mathbb{E}[Z_{\mathbb{G},\beta}\mathbf{1}\{\mathbb{P}[X(\mathbb{G}, \boldsymbol{\sigma}_{\mathbb{G}}) > \delta n \mid \mathbb{G}] \geq \delta\}] \leq \delta^{-1}\mathbb{E}[Z_{\mathbb{G},\beta}\mathbb{P}[X(\mathbb{G}, \boldsymbol{\sigma}_{\mathbb{G}}) > \varepsilon'' n \mid \mathbb{G}]] = o(\mathbb{E}[Z_{\mathbb{G},\beta}]).$$

Due to (4.31) it suffices to prove that there exists $\varepsilon = o(1)$ such that for any d-regular graph G of sufficiently large order n the following is true:

$$\text{if } \mathbb{P}[X(G, \boldsymbol{\sigma}_G) > \delta n] < \delta \text{ then } \mathbb{E}\left|\boldsymbol{\sigma}_G \cdot \boldsymbol{\sigma}'_G\right| \leq \varepsilon n. \tag{4.32}$$

Indeed, since $\mathbb{E}|\boldsymbol{\sigma}_G \cdot \boldsymbol{\sigma}'_G| = \mathbb{E}[\mathbb{E}[|\boldsymbol{\sigma}_G \cdot \boldsymbol{\sigma}'_G| \mid \boldsymbol{\sigma}'_G]]$, we may condition on $\boldsymbol{\sigma}'_G$. Hence, let V'_1 contain all vertices $v \in V(G)$ such that $\boldsymbol{\sigma}'_{G,v} = 1$ and let $V'_{-1} = V(G) \setminus V'_1$. Further, let

$$Y_s = \sum_{v \in V'_s} \mathbf{1}\{\boldsymbol{\sigma}_{G,v} = 1\} \qquad (s = \pm 1).$$

Then to establish (4.32) we just need to prove that

$$|Y_s - |V'_s|/2| = o(n) \qquad \text{w.h.p. for } s = \pm 1. \tag{4.33}$$

To deduce (4.33) fix $s = \pm 1$. If $|V'_s| < \delta^{1/3} n$, say, then (4.33) is immediate. Hence, we may assume that $|V'_s| \geq \delta^{1/3} n$. Draw a vertex $\boldsymbol{v} \in V'_s$ uniformly at random, independently of $\boldsymbol{\sigma}_G$. Then the assumption $\mathbb{P}[X(G, \boldsymbol{\sigma}_G) \leq \delta n \mid G]$ implies that

$$\mathbb{P}\left[|\mu_{G,\beta,\boldsymbol{v},\ell}(1|\boldsymbol{\sigma}_G) - 1/2| > \varepsilon' \mid \boldsymbol{\sigma}'_G\right] < \delta^{2/3}. \tag{4.34}$$

Now, consider a spin configuration $\boldsymbol{\sigma}''_G$ drawn from the Boltzmann distribution given the event

$$\mathscr{A}(\boldsymbol{v}, \boldsymbol{\sigma}_G) = \left\{\sigma \in \{\pm 1\}^{V(G)} : \sigma_w = \boldsymbol{\sigma}_{G,w} \text{ for all } w \text{ at distance } \ell \text{ or more from } \boldsymbol{v}\right\}.$$

In other words, $\boldsymbol{\sigma}''_G$ is obtained by re-sampling the spins of the vertices at distance less than ℓ from $\boldsymbol{v}$ from the Boltzmann distribution with the boundary condition that $\boldsymbol{\sigma}_G$ induces on the vertices at distance precisely ℓ from $\boldsymbol{v}$. Since $\boldsymbol{\sigma}_G$ is a sample from $\mu_{G,\beta}$, so is $\boldsymbol{\sigma}''_G$. Moreover, $\boldsymbol{\sigma}''_G$ is independent of $\boldsymbol{\sigma}'_G$. Therefore, (4.34) yields

$$\mathbb{E}[Y_s \mid \boldsymbol{\sigma}'_G] = |V'_s|\mathbb{P}\left[\boldsymbol{\sigma}''_{G,\boldsymbol{v}} = 1 \mid \boldsymbol{\sigma}'_G\right] = |V'_s|\mathbb{E}[\mu_{G,\beta,\boldsymbol{v},\ell}(1|\boldsymbol{\sigma}_G) \mid G, \boldsymbol{\sigma}'_G] \sim |V'_s|/2. \tag{4.35}$$

To complete the proof we apply similarly reasoning to estimate $\mathbb{E}[Y_s^2 \mid G, \boldsymbol{\sigma}'_G]$. Specifically, let $\boldsymbol{v}'$ be a second vertex drawn uniformly from V'_s, independently of $\boldsymbol{v}$. Then in analogy to (4.34) we obtain

$$\mathbb{P}\left[|\mu_{G,\beta,\boldsymbol{v},\ell}(1|\boldsymbol{\sigma}_G) - 1/2| > \varepsilon' \vee |\mu_{G,\beta,\boldsymbol{v}',\ell}(1|\boldsymbol{\sigma}_G) - 1/2| > \varepsilon' \mid \boldsymbol{\sigma}'_G\right] < 2\delta^{2/3}. \tag{4.36}$$

Further, draw $\boldsymbol{\sigma}'''_G$ from the Boltzmann distribution given

$$\mathscr{A}(\boldsymbol{v}, \boldsymbol{v}', \boldsymbol{\sigma}_G) = \left\{\sigma \in \{\pm 1\}^{V(G)} : \sigma_w = \boldsymbol{\sigma}_{G,w} \text{ for all } w \text{ at distance } \ell \text{ or more from both } \boldsymbol{v}, \boldsymbol{v}'\right\}.$$

Since G is d-regular, ℓ is fixed and $|V'_s| \geq \delta^{1/3} n \geq \sqrt{n}$, the vertices $\boldsymbol{v}, \boldsymbol{v}'$ are distance more than 2ℓ apart w.h.p. In this case the spins $\boldsymbol{\sigma}'''_{\boldsymbol{v}}, \boldsymbol{\sigma}'''_{\boldsymbol{v}'}$ are conditionally independent given $\boldsymbol{\sigma}'_G$. Consequently, (4.36) implies that

$$\begin{aligned}\mathbb{E}[Y_s^2 \mid G, \boldsymbol{\sigma}'_G] &= |V'_s|^2\mathbb{P}\left[\boldsymbol{\sigma}'''_{G,\boldsymbol{v}} = 1, \boldsymbol{\sigma}'''_{G,\boldsymbol{v}'} = 1 \mid G, \boldsymbol{\sigma}'_G\right] \\ &= |V'_s|^2\mathbb{E}[\mu_{G,\beta,\boldsymbol{v},\ell}(1|\boldsymbol{\sigma}_G) \mid G, \boldsymbol{\sigma}'_G]\mathbb{E}[\mu_{G,\beta,\boldsymbol{v}',\ell}(1|\boldsymbol{\sigma}_G) \mid G, \boldsymbol{\sigma}'_G] + o(2^2) = |V'_s|^2/4 + o(n^2).\end{aligned} \tag{4.37}$$

Finally, combining (4.35), (4.37) and Chebyshev's inequality, we obtain (4.33), completing the proof. □

Proof of Lemma 2.5. The lemma follows directly from Corollary 4.5, Lemma 4.6 and Lemma 4.7. □

4.4. **The truncated second moment.** The aim in this section is to show the following.

Lemma 4.11. *For any $d \geq 3, \beta > 0$ there exists $\varepsilon_n = o(1)$ such that the event $\mathcal{O}$ from* (2.13) *satisfies*

$$\mathbb{E}\Big[Z^2_{\mathbb{G}(n,d),\beta}\mathbf{1}\{\mathcal{O}\}\Big] \leq \mathbb{E}\big[Z_{\mathbb{G}(n,d),\beta}\big]^2 \exp(o(n)).$$

Toward the proof of Lemma 4.11 we require the following observation.

Fact 4.12. *Suppose that $(\mu_n)_{n\geq 1}$ is a sequence of probability measures $\mu_n \in \mathcal{P}(\{\pm 1\}^n)$ such that*

$$\lim_{n\to\infty} \sum_{\sigma,\sigma'\in\{\pm 1\}^n} \frac{|\sigma\cdot\sigma'|}{n}\mu_n(\sigma)\mu_n(\sigma') = 0.$$

Then $\lim_{n\to\infty} n^{-1}\sum_{\sigma\in\{\pm 1\}^n}|\sigma\cdot\mathbf{1}|\mu_n(\sigma)\mu_n(\sigma') = 0$.

Corollary 4.13. *For any $d \geq 3, \beta > 0$ there exists $\delta = \delta_n = o(1)$ such that*

$$\mathbb{E}\Big[Z^2_{\mathbb{G}(n,d),\beta}\mathbf{1}\{\mathcal{O}\}\Big] \leq (1+o(1)) \sum_{\sigma,\sigma'\in\{\pm 1\}^{V_n}} \mathbf{1}\big\{|\sigma\cdot\mathbf{1}|, |\sigma'\cdot\mathbf{1}|, |\sigma\cdot\sigma'| \leq \delta n\big\}\mathbb{E}[\exp(-\beta\mathscr{H}_{\mathbb{G}}(\sigma) - \beta\mathscr{H}_{\mathbb{G}}(\sigma'))].$$

Proof. This is an immediate consequence of Fact 4.12 and the definition (2.13) of the event $\mathcal{O}$. □

Lemma 4.14. *For any $d \geq 3, \beta > 0$ we have*

$$\sum_{\sigma,\sigma'\in\{\pm 1\}^{V_n}} \mathbf{1}\big\{|\sigma\cdot\mathbf{1}|, |\sigma'\cdot\mathbf{1}|, |\sigma\cdot\sigma'| \leq \delta n\big\}\mathbb{E}[\exp(-\beta\mathscr{H}_{\mathbf{G}}(\sigma) - \beta\mathscr{H}_{\mathbf{G}}(\sigma'))] \leq \exp\big(nf_d(0,\beta) + O(\delta n)\big). \tag{4.38}$$

Proof. Given $\sigma,\sigma' \in \{\pm 1\}^{V_n}$ let $\rho = \rho(\sigma,\sigma') = (\rho_{s,t}(\sigma,\sigma))_{s,t\in\{\pm 1\}}$ and $\mu = \mu(\sigma,\sigma',\mathbf{G}) = (\mu_{r,s,t,u}(\sigma,\sigma',\mathbf{G}))_{r,s,t,u\in\{\pm 1\}}$ be the vectors with entries

$$\rho_{s,t}(\sigma,\sigma') = \frac{1}{n}\sum_{v=1}^{n}\mathbf{1}\{\sigma_v = s, \sigma'_w = t\} \qquad (s,t\in\{\pm 1\}),$$

$$\boldsymbol{\mu}_{r,s,t,u}(\sigma,\sigma') = \frac{2}{dn}\sum_{\{v,w\}\in E(\mathbf{G})}\mathbf{1}\{\sigma_v = r, \sigma'_v = s, \sigma_w = t, \sigma'_w = u\} \qquad (r,s,t,u\in\{\pm 1\}).$$

Thus, $\rho(\sigma,\sigma')$ is the empirical distribution of the spin combinations that σ,σ' assign to the vertices. Similarly, $\boldsymbol{\mu}$ comprises the statistics of the edges of $\mathbf{G}$ joining vertices with different spin combinations. Let

$$\Sigma^{\otimes} = \big\{\sigma,\sigma' \in \{\pm 1\}^{V_n} : |\sigma\cdot\mathbf{1}|, |\sigma'\cdot\mathbf{1}|, |\sigma\cdot\sigma'| \leq \delta n\big\}, \qquad \mathscr{R}^{\otimes} = \{\rho(\sigma,\sigma') : \sigma,\sigma' \in \Sigma^{\otimes}\}, \tag{4.39}$$

and let $\mathscr{M}^{\otimes}$ be the set of all possible outcomes of the random vector $\boldsymbol{\mu}(\sigma,\sigma')$ for any $\sigma,\sigma' \in \Sigma^{\otimes}$. For $\rho \in \mathscr{R}^{\otimes}, \mu \in \mathscr{M}^{\otimes}$ we use the shorthands $\rho_{++} = \rho_{+1,+1}$ and $\mu_{++++} = \mu_{+1,+1,+1,+1}$, and similarly for the other possible sign patterns. Further, let

$$\mathscr{H}(\mu) = 2\big(\mu_{++++} + \mu_{----} + \mu_{+-+-} + \mu_{-+-+}\big) + \tag{4.40}$$

$$\mu_{+++-} + \mu_{++-+} + \mu_{+-++} + \mu_{-+++} + \mu_{---+} + \mu_{--+-} + \mu_{-+--} + \mu_{+---}. \tag{4.41}$$

Finally, for $\mu \in \mathscr{M}^{\otimes}$ we define $\rho(\mu) \in \mathscr{R}^{\otimes}$ by

$$\rho_{ij}(\mu) = \sum_{k,l\in\{\pm 1\}}\mu_{ijkl} \qquad \text{for } i,j\in\{\pm 1\}.$$

We now claim that for any $\mu \in \mathscr{M}^{\otimes}$,

$$\sum_{\sigma,\sigma'\in\{\pm 1\}^{V_n}} \mathbb{P}\big[\boldsymbol{\mu}(\sigma,\sigma') = \mu\big] = \frac{\mathscr{X}_\mu\mathscr{Y}_\mu\mathscr{Z}_\mu}{(dn-1)!!} \qquad \text{where} \tag{4.42}$$

$$\mathscr{X}_\mu = \binom{n}{\rho_{++}(\mu)n, \rho_{+-}(\mu)n, \rho_{-+}(\mu)n, \rho_{--}(\mu)n},$$

$$\mathscr{Y}_\mu = \prod_{i,j\in\{\pm\}}\binom{dn\rho_{ij}(\mu)}{dn\mu_{ij++}, dn\mu_{ij+-}, dn\mu_{ij-+}, dn\mu_{ij--}},$$

$$\mathscr{Z}_\mu = (dn\mu_{++--})!\,(dn\mu_{-++-})!\prod_{i\in\{\pm\}}\big((dn\mu_{+i-i})!\,(dn\mu_{i+i-})!\big)\prod_{i,j\in\{\pm\}}(dn\mu_{ijij}-1)!!.$$

Indeed, the first factor $\mathscr{X}_\mu$ counts all pairs (σ,σ') with $\rho(\sigma,\sigma') = \rho(\mu)$. Moreover, $\mathscr{Y}_\mu$ accounts for the number of ways of selecting among clones of vertices with a given spin combination those that will be matched to vertices with another specific sign combination. Finally, $\mathscr{Z}_\mu$ equals the number of possible matchings of clones in accordance with their designations.

Combining (4.40) and (4.42), we obtain

$$\begin{aligned}
&\sum_{\sigma,\sigma'\in\{\pm1\}^{V_n}} \mathbf{1}\left\{|\sigma\cdot\mathbf{1}|,|\sigma'\cdot\mathbf{1}|,|\sigma\cdot\sigma'|\le\delta n\right\}\mathbb{E}[\exp(-\beta\mathscr{H}_{\mathbf{G}}(\sigma)-\beta\mathscr{H}_{\mathbf{G}}(\sigma'))]\\
&\quad=\sum_{\mu\in\mathscr{M}_n}\sum_{\sigma,\sigma'\in\{\pm1\}^{V_n}}\mathbb{P}\left[\boldsymbol{\mu}(\sigma,\sigma')=\mu\right]\exp(-\beta\mathscr{H}(\mu))\\
&\quad\le|\mathscr{M}_n|\max_{\mu\in\mathscr{M}_n}\sum_{\sigma,\sigma'\in\{\pm1\}^{V_n}}\mathbb{P}\left[\boldsymbol{\mu}(\sigma,\sigma')=\mu\right]\exp(-\beta\mathscr{H}(\mu))\\
&\quad\le\exp(O(\log n))\max_{\mu\in\mathscr{M}_n}\frac{\mathscr{X}_\mu\mathscr{Y}_\mu\mathscr{Z}_\mu}{(dn-1)!!}\exp(-\beta\mathscr{H}(\mu)),
\end{aligned}\tag{4.43}$$

because, naturally, $|\mathscr{M}_n|\le n^{16}$. Further, Stirling's formula (4.3) and the explicit formula (4.4) for the double factorial yield the approximations

$$\log\mathscr{X}_\rho=nH(\rho(\mu))+O\big(\log n\big),\quad \log\mathscr{Y}_{\rho,\mu}=dn(H(\mu)-H(\rho))+O\big(\log n\big),\quad \log\frac{\mathscr{Z}_{\rho,\mu}}{(dn-1)!!}=-\frac{dn}{2}H(\mu)+O(\log n).$$

Combining these formulas and recalling the the Kullback-Leibler divergence, we obtain

$$\max_{\mu\in\mathscr{M}^\otimes}\frac{\mathscr{X}_\mu\mathscr{Y}_\mu\mathscr{Z}_\mu}{(dn-1)!!}\mathrm{e}^{-\beta\mathscr{H}(\mu)}=\exp\left(n\max_{\mu\in\mathscr{M}^\otimes}\left\{H(\rho(\mu))-dD_{\mathrm{KL}}\left(\mu\|\rho(\mu)\otimes\rho(\mu)\right)/2-d\beta\mathscr{H}(\mu)/2\right\}+O(\log n)\right).\tag{4.44}$$

To simplify the last optimisation problem we reparametrise the exponent in terms of $\alpha\in[-1,1]$. Combinatorially the optimal choice of α will correspond to the overlap value $\sigma\cdot\sigma'/n$ that renders the largest contribution to the l.h.s. of (4.38). Hence, let $\mathscr{M}^\otimes(\alpha)$ be the set of all $\mu\in\mathscr{M}^\otimes$ such that

$$\rho_{++}(\mu)=\rho_{--}(\mu)=\frac{1+\alpha}{4},\qquad\qquad \rho_{+-}(\mu)=\rho_{-+}(\mu)=\frac{1-\alpha}{4}.\tag{4.45}$$

Then we claim that for any $\alpha\in(-1,1)$,

$$\max_{\mu\in\mathscr{M}^\otimes(\alpha)}H(\rho(\mu))-dD_{\mathrm{KL}}\left(\mu\|\rho(\mu)\otimes\rho(\mu)\right)/2-d\beta\mathscr{H}(\mu)/2\le f_d(\alpha,\beta).\tag{4.46}$$

To see this, we first notice that the constrained optimization problem on the l.h.s of 4.46 is upper bounded by the result of the unconstrained optimization problem, i.e.

$$\begin{aligned}
&\max_{\mu\in\mathscr{M}^\otimes(\alpha)}H(\rho(\mu))-dD_{\mathrm{KL}}\left(\mu\|\rho(\mu)\otimes\rho(\mu)\right)/2-d\beta\mathscr{H}(\mu)/2\\
&\qquad\le\max_{\mu\in M(\alpha)}H(\rho(\mu))-dD_{\mathrm{KL}}\left(\mu\|\rho(\mu)\otimes\rho(\mu)\right)/2-d\beta\mathscr{H}(\mu)/2
\end{aligned}\tag{4.47}$$

where $M(\alpha)$ is the set of all probability measures on $\mathscr{P}_\alpha(\{\pm1\}^4)$ parametrised by α. Moreover, we notice that the function $\mu\in M(\alpha)\mapsto H(\rho(\mu))-dD_{\mathrm{KL}}\left(\mu\|\rho(\mu)\otimes\rho(\mu)\right)/2-d\beta\mathscr{H}(\mu)/2$ is concave because $H(\rho(\mu))$ is constant on $M(\alpha)$, the Kullback-Leibler divergence is strictly convex and the function $\mathscr{H}(\mu)$ is linear. Hence, it suffices to find the (unique) zero of the derivative of $D_{\mathrm{KL}}\left(\mu\|\rho(\mu)\otimes\rho(\mu)\right)/2+\beta\mathscr{H}(\mu)/2$ subject to (4.45).

Letting $z_\alpha=(1+\mathrm{e}^{-2\beta})(1+\alpha^2)/4+\mathrm{e}^{-\beta}(1-\alpha^2)/2$, we claim that μ_α given by

$$\mu_{\alpha,++++}=\mu_{\alpha,----}=\frac{(1+\alpha)^2}{16z_\alpha}\mathrm{e}^{-2\beta},\qquad \mu_{\alpha,+-+-}=\mu_{\alpha,-+-+}=\frac{(1-\alpha)^2}{16z_\alpha}\mathrm{e}^{-2\beta},\tag{4.48}$$

$$\mu_{\alpha,+++-}=\mu_{\alpha,++-+}=\mu_{\alpha,+-++}=\mu_{\alpha,-+++}=\mu_{\alpha,---+}=\mu_{\alpha,--+-}=\mu_{\alpha,-+--}=\mu_{\alpha,+---}=\frac{1-\alpha^2}{16z_\alpha}\mathrm{e}^{-\beta},\tag{4.49}$$

$$\mu_{\alpha,++--}=\mu_{\alpha,--++}=\frac{(1+\alpha)^2}{16z_\alpha},\qquad \mu_{\alpha,+--+}=\mu_{\alpha,-++-}=\frac{(1-\alpha)^2}{16z_\alpha},\tag{4.50}$$

fits the bill. Indeed, we calculate

$$\mu_{\alpha,++++}\log\frac{\mu_{\alpha,++++}}{\rho_{\alpha,++}\rho_{\alpha,++}}=\mu_{\alpha,++++}\log\frac{\mathrm{e}^{-2\beta}}{z_\alpha}=-\mu_{\alpha,++++}(\log(z_\alpha)+2\beta).\tag{4.51}$$

The same formula holds with $++++$ replaced by any of the other sign patterns from (4.48), i.e., $----$, $+-+-$, $-+-+$. Moreover,

$$\mu_{\alpha,+++-}\log\frac{\mu_{\alpha,+++-}}{\rho_{\alpha,++}\rho_{\alpha,+-}}=\mu_{\alpha,+++-}\log\frac{\mathrm{e}^{-\beta}}{z_\alpha}=-\mu_{\alpha,+++-}(\log(z_\alpha)+\beta),\tag{4.52}$$

and similarly for the other seven patterns from (4.49). Further,

$$\mu_{\alpha,++--}\log\frac{\mu_{\alpha,++--}}{\rho_{\alpha,++}\rho_{\alpha,--}}=-\mu_{\alpha,++--}\log z_\alpha\tag{4.53}$$

and analogously for the other sign patterns from (4.50). Consequently, the derivatives work out to be

$$\frac{\partial}{\partial \mu_{\alpha,++++}}\left(D_{\mathrm{KL}}\left(\mu\|\rho(\mu)\otimes\rho(\mu)\right)+\beta\mathscr{H}(\mu)\right)\Big|_{\mu_\alpha}=1+\log\frac{\mu_{\alpha,++++}}{\rho_{++}(\mu_\alpha)^2}+2\beta=1-\log(z_\alpha), \tag{4.54}$$

$$\frac{\partial}{\partial \mu_{\alpha,+++-}}\left(D_{\mathrm{KL}}\left(\mu\|\rho(\mu)\otimes\rho(\mu)\right)+\beta\mathscr{H}(\mu)\right)\Big|_{\mu_\alpha}=1+\log\frac{\mu_{\alpha,+++-}}{\rho_{++}(\mu_\alpha)\rho_{+-}(\mu_\alpha)}+\beta=1-\log(z_\alpha), \tag{4.55}$$

$$\frac{\partial}{\partial \mu_{\alpha,++--}}\left(D_{\mathrm{KL}}\left(\mu\|\rho(\mu)\otimes\rho(\mu)\right)+\beta\mathscr{H}(\mu)\right)\Big|_{\mu_\alpha}=1+\log\frac{\mu_{\alpha,++--}}{\rho_{++}(\mu_\alpha)\rho_{--}(\mu_\alpha)}=1-\log(z_\alpha). \tag{4.56}$$

In each case the same calculation applies to the other sign patterns from the respective line (4.48)–(4.50). Since the right hand sides of (4.54)–(4.56) are identical, the constraint that μ belongs to the simplex $\mathscr{P}(\{\pm1\}^4)$ shows that μ_α is a stationary point and therefore the unique maximiser of $D_{\mathrm{KL}}\left(\mu\|\rho(\mu)\otimes\rho(\mu)\right)/2+\beta\mathscr{H}(\mu)/2$. Moreover, combining (4.51)–(4.53), we find

$$-D_{\mathrm{KL}}\left(\mu_\alpha\|\rho_\alpha\otimes\rho_\alpha\right)-\beta\mathscr{H}(\mu_\alpha)=\log(z_\alpha)=\log\left(1+\alpha^2\left(\frac{1-\mathrm{e}^{-\beta}}{1+\mathrm{e}^{-\beta}}\right)^2\right)+2\log\frac{1+\mathrm{e}^{-\beta}}{2}, \tag{4.57}$$

Thus, (4.46) follows from the definition (2.6) of $f_d(\alpha,\beta)$ and (4.57).

To complete the proof consider any $\mu\in M(\alpha)$. Then (4.39) ensures that there exists $\mu'\in M(0)$ such that $\|\mu-\mu'\|_2=O(\delta)$. Consequently, since the function $\mu\mapsto H(\rho(\mu))-dD_{\mathrm{KL}}\left(\mu\|\rho(\mu)\otimes\rho(\mu)\right)/2-d\beta\mathscr{H}(\mu)/2$ is differentiable, the bound (4.46) shows that

$$\max_{\mu\in M(\alpha)}H(\rho(\mu))-dD_{\mathrm{KL}}\left(\mu\|\rho(\mu)\otimes\rho(\mu)\right)/2-d\beta\mathscr{H}(\mu)/2\le f_d(0,\beta)+O(\delta). \tag{4.58}$$

Thus, the assertion follows from (4.43) and (4.58). □

Corollary 4.15. *For any $d\ge3,\beta>0$ we have*

$$\mathbb{E}\left[Z^2_{\mathbf{G}(n,d),\beta}\mathbf{1}\{\mathscr{O}\}\right]\le\mathbb{E}\left[Z_{\mathbf{G}(n,d),\beta}\right]^2\exp(o(n)).$$

Proof. This is an immediate consequence of Corollary 4.13 and Lemma 4.14. □

Proof of Lemma 4.11. The lemma follows from Lemma 4.3, Lemma 2.1 and Corollary 4.15. □

Proof of Proposition 2.2. The proposition follows from Lemma 2.1, Lemma 2.3, Lemma 2.5, Lemma 4.3, and Corollary 4.15. □

5. Proof of Proposition 2.8

We begin by deriving an explicit formula for $\mathscr{B}(\pi^*_\varepsilon,\beta)$. Recall that $\Lambda(x)=x\log x$.

Lemma 5.1. *Let $d\ge3,\beta>0$. Then for small enough $\varepsilon>0$ we have*

$$\mathscr{B}_{Ising}(\pi^*_\varepsilon,\beta,d)=\frac{\sum_{i=1}^d\binom{d}{i}2^{-d}\Lambda\left(\sum_{\sigma=\pm1}\left(1-\left(1-\mathrm{e}^{-\beta}\right)\left(\frac12+\sigma\varepsilon\right)\right)^i\left(1-\left(1-\mathrm{e}^{-\beta}\right)\left(\frac12-\sigma\varepsilon\right)\right)^{d-i}\right)}{2\left((1+\mathrm{e}^{-\beta})/2\right)^d}$$
$$-\frac{d\left(\Lambda\left(1-\left(1-\mathrm{e}^{-\beta}\right)\left(\frac12+2\varepsilon^2\right)\right)+\Lambda\left(1-\left(1-\mathrm{e}^{-\beta}\right)\left(\frac12-2\varepsilon^2\right)\right)\right)}{2(1+\mathrm{e}^{-\beta})}$$

Proof. The expression follows straight from plugging the distribution π^*_ε from (2.21) into the Bethe functional from (2.19). Let us shed light on its combinatorial meaning. The first term represents the 'weighted penalty factor' arising at a root vertex with d adjacent vertices. Since we polarise each of these adjacent vertices with probability 1/2 independently, the number of adjacent vertices polarised to ε and $-\varepsilon$ follows a binomial distribution. The term $\binom{d}{i}2^{-d}$ captures the corresponding probability while the term inside $\Lambda(\cdot)$ describes the resulting penalty factor over all adjacent vertices summed over the $+1$ and -1 spins at the root vertex. The second term represents the 'weighted penalty factor' between two vertices connected via an edge. Here, the first corresponds to the case that both vertices are polarised to the same spin, while the second summand picks up the penalty factor for polarisation towards different spins. □

Lemma 5.2. *Let $d\ge3,\beta>0$. Then as $\varepsilon\to0$,*

$$\mathscr{B}_{Ising}(\pi^*_\varepsilon,\beta,d)=\log2+\frac{d}{2}\log\frac{1+\mathrm{e}^{-\beta}}{2}+\frac{4\varepsilon^4d\mathrm{e}^{-2\beta}\left(\mathrm{e}^\beta-1\right)^2\left(\mathrm{e}^{2\beta}(d-2)-2d\mathrm{e}^\beta+d-2\right)}{(1+\mathrm{e}^\beta)^2(1+\mathrm{e}^{-\beta})^2}+O(\varepsilon^5).$$

Proof. Lemma 5.1 shows that the function $\varepsilon \mapsto \mathscr{B}_{\mathrm{Ising}}(\pi^*_\varepsilon, \beta, d)$ has five continuous derivatives in for small enough $\varepsilon > 0$. Hence, Taylor's formula yields

$$\begin{aligned}\mathscr{B}_{\mathrm{Ising}}(\pi^*_\varepsilon, \beta, d) =& \mathscr{B}_{\mathrm{Ising}}(\pi^*_0, \beta) + \varepsilon \frac{\partial}{\partial \varepsilon} \mathscr{B}_{\mathrm{Ising}}(\pi^*_\varepsilon, \beta, d)|_{\varepsilon=0} + \frac{\varepsilon^2}{2} \frac{\partial^2}{\partial \varepsilon^2} \mathscr{B}_{\mathrm{Ising}}(\pi^*_\varepsilon, \beta, d)|_{\varepsilon=0} \\ &+ \frac{\varepsilon^3}{6} \frac{\partial^3}{\partial \varepsilon^3} \mathscr{B}_{\mathrm{Ising}}(\pi^*_\varepsilon, \beta, d)|_{\varepsilon=0} + \frac{\varepsilon^4}{24} \frac{\partial^4}{\partial \varepsilon^4} \mathscr{B}_{\mathrm{Ising}}(\pi^*_\varepsilon, \beta, d)|_{\varepsilon=0} + O\left(\varepsilon^5\right). \end{aligned} \tag{5.1}$$

The formula for $\mathscr{B}_{\mathrm{Ising}}(\pi^*_\varepsilon, \beta, d)$ from Lemma 5.1 is complicated but explicit. Therefore, we can rely on a computer algebra system to calculate the first four derivatives of $\mathscr{B}_{\mathrm{Ising}}(\pi^*_\varepsilon, \beta, d)$ symbolically at $\varepsilon = 0$. The result of this calculation reads

$$\frac{\partial}{\partial \varepsilon} \mathscr{B}_{\mathrm{Ising}}(\pi^*, \beta, d)|_{\varepsilon=0} = \frac{\partial^2}{\partial \varepsilon^2} \mathscr{B}_{\mathrm{Ising}}(\pi^*, \beta, d)|_{\varepsilon=0} = \frac{\partial^3}{\partial \varepsilon^3} \mathscr{B}_{\mathrm{Ising}}(\pi^*, \beta, d)|_{\varepsilon=0} = 0, \tag{5.2}$$

$$\frac{\partial^4}{\partial \varepsilon^4} \mathscr{B}_{\mathrm{Ising}}(\pi^*, \beta, d)|_{\varepsilon=0} = \frac{96 d e^{-2\beta} \left(e^\beta - 1\right)^2 \left(e^{2\beta}(d-2) - 2de^\beta + d - 2\right)}{(1+e^\beta)^2 (1+e^{-\beta})^2} \tag{5.3}$$

Plugging (5.2)–(5.3) into (5.1) yields the desired formula. □

Proof of Proposition 2.8. Let $d \geq 3$. A few lines of algebra reveal that $\mathrm{e}^{2\beta}(d-2) - 2d\mathrm{e}^\beta + d - 2 > 0$ if $\beta > \beta^*(d)$. Therefore, Lemma 5.2 shows that for small enough $\varepsilon > 0$ and $\beta > \beta^*(d)$ we have $\mathscr{B}_{\mathrm{Ising}}(\pi^*_\varepsilon, \beta, d) > \log 2 + \frac{d}{2} \log \frac{1+\mathrm{e}^{-\beta}}{2}$, as claimed. □

6. Proof of Proposition 2.10

We treat X_1, X_2 from Lemma 2.9 separately. Let us begin with X_2, which is the easier case. In the following, fix any $d \geq 3$, $\alpha \in (0, 1/2)$ and $z \in (0,1)$ and set $y = y(\beta) = -\beta^{-1} \log z$.

Lemma 6.1. *We have* $\lim_{\beta\to\infty} \log \mathbb{E}[X_2^y] = \log\left(1 - 2\alpha^2 + 2\alpha^2 z\right)$.

Proof. Recalling the definition of $\mathfrak{r}_\alpha$ from (2.23) and writing the expectation out explicitly, we obtain

$$\mathbb{E}[X_2^y] = \mathbb{E}\left[\left(1 - (1-\mathrm{e}^{-\beta}) \sum_{\tau\in\{\pm1\}} \boldsymbol{\rho}_1(\tau)\boldsymbol{\rho}_2(\tau)\right)^y\right] = \sum_{r_1,r_2\in\{0,1,-1\}} \mathfrak{r}_\alpha(r_1)\mathfrak{r}_\alpha(r_2) \left[1 - \frac{(1-\mathrm{e}^{-\beta})(1+r_1r_2)}{2}\right]^y. \tag{6.1}$$

To evaluate this expression we consider the possible values of the product $r_1 r_2$.

Case 1: $r_1 r_2 = -1$**:** by the choice (2.23) of $\mathfrak{r}$ this event has probability $2\alpha^2$ and

$$1 - \frac{(1-\mathrm{e}^{-\beta})(1+r_1r_2)}{2} = 1. \tag{6.2}$$

Case 2: $r_1 r_2 = 1$**:** in this case, which occurs with probability $2\alpha^2$ as well, we obtain

$$1 - \frac{(1-\mathrm{e}^{-\beta})(1+r_1r_2)}{2} = \mathrm{e}^{-\beta}. \tag{6.3}$$

Case 3: $r_1 r_2 = 0$**:** naturally this event occurs with the remaining probability $1 - 4\alpha^2$ and

$$1 - \frac{(1-\mathrm{e}^{-\beta})(1+r_1r_2)}{2} = \frac{1+\mathrm{e}^{-\beta}}{2}. \tag{6.4}$$

Combining (6.1)–(6.4), we obtain

$$\mathbb{E}\left[\left(1 - (1-\mathrm{e}^{-\beta}) \sum_{\tau\in\{\pm1\}} \boldsymbol{\rho}_1(\tau)\boldsymbol{\rho}_2(\tau)\right)^y\right] = 2\alpha^2 + 2\alpha^2\mathrm{e}^{-\beta y} + (1-4\alpha^2)\left(\frac{1+\mathrm{e}^{-\beta}}{2}\right)^y. \tag{6.5}$$

Finally, since $z = \exp(-\beta y)$ and

$$\left(\frac{1+\mathrm{e}^{-\beta}}{2}\right)^y = \exp\left(-\frac{\log z}{\beta} \log(1+\mathrm{e}^{-b})\right) \to 1 \qquad \text{as } \beta \to \infty,$$

combining (6.1) and (6.5) shows that $\lim_{\beta\to\infty} \log \mathbb{E}[X_2^y] = \log(1 - 2\alpha^2 + 2\alpha^2 z)$, as desired. □

The computation of $\mathbb{E}\left[X_1^y\right]$ is a little more intricate. Combinatorially speaking, the basic idea is this. Consider the picture on the left of Figure 2. The expression

$$X_1 = \sum_{\tau\in\{\pm1\}} \prod_{h=1}^{d} 1 - (1-\mathrm{e}^{-\beta})\boldsymbol{\rho}_h(\tau)$$

represents the contribution to the partition function of a single white vertex along with its adjacent blue boxes. Each of these boxed represents an imaginary vertex, or a 'field' in physics jargon, that takes the spin τ with probability $\boldsymbol{\rho}_h(\tau)$. The spins of these imaginary vertices are mutually independent. Hence, the sum on τ in the definition of X_1 accounts for the two possible choices of spin for the white vertex. Thus, if we let $\boldsymbol{H}(\tau)$ be the number of imaginary vertices with spin τ, then the product $\prod_{h=1}^{d} 1-(1-\mathrm{e}^{-\beta})\boldsymbol{\rho}_h(\tau)$ equals the expected Boltzmann weight

$$\mathbb{E}[\exp(-\beta \boldsymbol{H}(\tau)) \mid \boldsymbol{\rho}_1, \ldots, \boldsymbol{\rho}_d].$$

Furthermore, the fields $\boldsymbol{\rho}_h(\tau)$ can be either 'soft', i.e., $\boldsymbol{\rho}_h(\tau) = 1/2$, or 'hard', meaning $\boldsymbol{\rho}_h(\tau) \in \{0,1\}$. As in the proof of Lemma 6.1 we will see that in the limit $\beta \to \infty$ and $y \to 0$ the soft fields are inconsequential. In effect, the computation of X_1 will come down to studying the random variable $\sum_{\tau \in \{\pm 1\}} \tau \mathbf{1}\{\boldsymbol{\rho}_h(\tau) = 1\}$, which gauges the relative strength of the hard fields. In other words, the calculation of $\mathbb{E}\left[X_1^y\right]$ comes down to analysing a random walk. Let us get down to the details.

Lemma 6.2. *We have* $\lim_{\beta\to\infty} \log \mathbb{E}\left[X_1^y\right] = \log\left(\zeta \mathscr{A}^d \xi\right)$.

Proof. Letting

$$\boldsymbol{R}_{-1} = \sum_{h=1}^{d} \mathbf{1}\left\{\boldsymbol{\rho}_h(1) = 0\right\}, \qquad \boldsymbol{R}_0 = \sum_{h=1}^{d} \mathbf{1}\left\{\boldsymbol{\rho}_h(1) = 1/2\right\}, \qquad \boldsymbol{R}_1 = \sum_{h=1}^{d} \mathbf{1}\left\{\boldsymbol{\rho}_h(1) = 1\right\}$$

we can write the random walk as $\sum_{\tau\in\{\pm1\}} \tau \mathbf{1}\{\boldsymbol{\rho}_h(\tau) = 1\} = \boldsymbol{R}_1 - \boldsymbol{R}_{-1}$. Hence,

$$\prod_{h=1}^{d} \left(1-(1-\mathrm{e}^{-\beta})\boldsymbol{\rho}_h(1)\right)^y = \exp(-\beta y \boldsymbol{R}_1)\left(\frac{1+\mathrm{e}^{-\beta}}{2}\right)^{y\boldsymbol{R}_0}, \tag{6.6}$$

$$\prod_{h=1}^{d} \left(1-(1-\mathrm{e}^{-\beta})\boldsymbol{\rho}_h(-1)\right)^y = \exp(-\beta y \boldsymbol{R}_{-1})\left(\frac{1+\mathrm{e}^{-\beta}}{2}\right)^{y\boldsymbol{R}_0}. \tag{6.7}$$

Since $y = -\beta^{-1}\log z$, for any non-negative integers $R_1, R_{-1}, R_0 \geq 0$ such that $R_1 + R_0 + R_{-1} = d$ we have

$$\lim_{\beta\to\infty} (\exp(-\beta y R_1) + \exp(-\beta y R_{-1}))\left(\frac{1+\mathrm{e}^{-\beta}}{2}\right)^{yR_0} = z^{R_1 \wedge R_{-1}}. \tag{6.8}$$

Thus, combining (6.6)–(6.8), we obtain

$$\lim_{\beta\to\infty} \log \mathbb{E}[X_1^y] = \log \mathbb{E}\left[z^{\boldsymbol{R}_1 \wedge \boldsymbol{R}_{-1}}\right]. \tag{6.9}$$

To calculate the mean on the r.h.s. consider a d-step symmetric random walk on $\{0,1,\ldots,d\}$ with a reflective barrier at 0. The walk starts at 0 and the available moves are $+1$, -1 or 0, with probabilities α, α and $1-2\alpha$ respectively. We couple this random walk with the probability space $(\boldsymbol{\rho}_1, \ldots, \boldsymbol{\rho}_d)$ such that $\boldsymbol{R}_1$ and $\boldsymbol{R}_{-1}$ count the ±1 moves of the random walk, respectively. Thus, $\boldsymbol{R}_0 = d - \boldsymbol{R}_1 - \boldsymbol{R}_{-1}$ equals the number of 0-moves and $|\boldsymbol{R}_1 - \boldsymbol{R}_{-1}|$ is the final position of the walk. To study this random walk we remember the matrix $\mathscr{M}$ from (1.7) and introduce

$$\mathfrak{A} = (1-2\alpha)\mathrm{id} + 2\alpha t.\mathscr{M},$$

where t is a formal variable that we introduce to track the walk's movements. Specifically, for any $i \in [d]$ the $(1,i)$-entry of the d-th power of $\mathfrak{A}$ works out to be

$$\mathfrak{A}_{1i}^d = \sum_{k=i-1}^{d} t^k \mathbb{P}\left[\boldsymbol{R}_1 + \boldsymbol{R}_{-1} = k \text{ and } |\boldsymbol{R}_1 - \boldsymbol{R}_{-1}| = i-1\right] \tag{6.10}$$

Finally, we introduce the vector

$$\mathfrak{x} = (1, t^{-1}, t^{-2}, t^{-3}, \cdots)^T \in \mathbb{R}^{(d+1)\times 1}.$$

Then recalling the definition of vector ζ from (1.10), we readily find

$$\zeta \mathfrak{A}^d \mathfrak{x} = \sum_{k=0}^{d} t^k \mathbb{P}\left[\boldsymbol{R}_1 + \boldsymbol{R}_{-1} - |\boldsymbol{R}_1 - \boldsymbol{R}_{-1}| = k\right] = \mathbb{E}\left[t^{\boldsymbol{R}_1 + \boldsymbol{R}_{-1} - |\boldsymbol{R}_1 - \boldsymbol{R}_{-1}|}\right]. \tag{6.11}$$

Let us shed light on the combinatorial meaning of (6.11). $\zeta\mathfrak{A}^d$ is a $(d+1)$-dimensional vector where the ith entry captures the probability of all random walks that end up at position $i-1$ and where the exponent of t measures the number of non-stationary steps performed to reach position $i-1$. Thus, using the definition from (6.10) we have

$$\zeta \mathfrak{A}^d = \left(\mathfrak{A}_{11}^d, \mathfrak{A}_{12}^d, \mathfrak{A}_{13}^d, \ldots\right).$$

The multiplication with vector $\mathfrak{x}$ then deducts the final position $|\boldsymbol{R}_1 - \boldsymbol{R}_{-1}|$ of the random walk from the exponent of t. In effect, the exponent now captures the total number of offsetting non-stationary steps of the random walk which is precisely twice the minimum of $\boldsymbol{R}_1$ and $\boldsymbol{R}_{-1}$. This relationship can be compactly written in the basic identity

$$\boldsymbol{R}_1 + \boldsymbol{R}_{-1} - |\boldsymbol{R}_1 - \boldsymbol{R}_{-1}| = 2\,(\boldsymbol{R}_1 \wedge \boldsymbol{R}_{-1})\,.$$

We are now in a position to relate (6.8) to (6.11) by writing

$$\lim_{\beta\to\infty} \log\mathbb{E}[X_1^y] = \log\mathbb{E}\left[z^{\boldsymbol{R}_1\wedge\boldsymbol{R}_{-1}}\right] = \mathbb{E}\left[\sqrt{z}^{\boldsymbol{R}_1+\boldsymbol{R}_{-1}-|\boldsymbol{R}_1-\boldsymbol{R}_{-1}|}\right] = \log\left(\zeta\mathfrak{A}^d\mathfrak{x}|_{t=\sqrt{z}}\right).$$

Since $\mathscr{A}$ from (1.9) and ξ from (1.11) were defined in terms of $\sqrt{z}$ rather than t we conclude that

$$\lim_{\beta\to\infty} \log\mathbb{E}\left[X_1^y\right] = \log\left(\zeta\mathfrak{A}^d\mathfrak{x}|_{t=\sqrt{z}}\right) = \log\left(\zeta\mathscr{A}^d\xi\right)$$

as claimed. □

Proof of Proposition 2.10. The proposition is an immediate consequence of Lemmas 6.1 and 6.2. □

REFERENCES

[1] D. Achlioptas, C. Moore: The chromatic number of random regular graphs. In Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques. Springer, Berlin, Heidelberg (2004) 219–228.

[2] E. Abbe: Community detection and stochastic block models: recent developments. The Journal of Machine Learning Research **18** (2017) 6446–6531.

[3] E. Abbe, C. Sandon: Proof of the achievability conjectures for the general stochastic block model. Communications on Pure and Applied Mathematics **71** (2018) 1334–1406.

[4] D. Achlioptas, A. Coja-Oghlan: Algorithmic barriers from phase transitions. Proc. 49th FOCS (2008) 793-802.

[5] P. Ayre, A. Coja-Oghlan, C. Greenhill: Lower bounds on the chromatic number of random graphs. arXiv preprint arXiv:1812.09691 (2018).

[6] B. Barak, D. Steurer: Sum-of-squares proofs and the quest toward optimal algorithms. Electronic Colloquium on Computational Complexity **21** (2014).

[7] Z. Bartha, N. Sun, Y. Zhang: Breaking of 1RSB in random MAX-NAE-SAT. Proc. 60th FOCS (2019) 1405–1416.

[8] M. Bayati, D. Gamarnik, P. Tetali. Combinatorial approach to the interpolation method and scaling limits in sparse random graphs. Annals of Probability **41** (2013) 4080–4115.

[9] A. Bertoni, P. Campadelli, R. Posenato: An upper bound for the maximum cut mean value. Proc. 23rd WG (1997) 78–84.

[10] P. Bleher, J. Ruiz, V. Zagrebnov: On the purity of the limiting Gibbs state for the Ising model on the Bethe lattice. Journal of Statistical Physics. **79** (1995) 473–482.

[11] C. Bordenave, M. Lelarge, L. Massoulié: Non-backtracking spectrum of random graphs: community detection and non-regular ramanujan graphs. Proc. 56th FOCS (2015) 1347–1357.

[12] A. Coja-Oghlan, C. Moore, V. Sanwalani: MAX k-CUT and approximating the chromatic number of random graphs. Random Structures & Algorithms **28** (2006) 289–322.

[13] A. Coja-Oghlan, C. Efthymiou, S. Hetterich: On the chromatic number of random regular graphs. Journal of Combinatorial Theory Series B **116** (2016) 367–439.

[14] A. Coja-Oghlan, F. Krzakala, W. Perkins, L. Zdeborová: Information-theoretic thresholds from the cavity method. Advances in Mathematics **333** (2018) 694–795

[15] A. Coja-Oghlan, W. Perkins: Spin systems on Bethe lattices. Communications in Mathematical Physics **372** (2019) 441–523.

[16] A. Coja-Oghlan, M. Hahn-Klimroth, P. Loick, N. Müller, K. Panagiotou, M. Pasch: Inference and mutual information on random factor graphs. arXiv preprint arXiv:2007.07494 (2020).

[17] A. Coja-Oghlan, F. Krzakala, W. Perkins, L. Zdeborová: Information-theoretic thresholds from the cavity method. Advances in Mathematics **333** (2018) 694–795.

[18] D. Coppersmith, D. Gamarnik, M. Hajiaghayi, G. Sorkin: Random MAX SAT, random MAX CUT, and their phase transitions. Random Structures & Algorithms, **24** (2004) 502–545.

[19] E. Csóka: Independent sets and cuts in large-girth regular graphs. arXiv preprint arXiv:1602.02747 (2016).

[20] E. Csóka, B. Gerencsér, V. Harangi, and B. Virág: Invariant Gaussian processes and independent sets on regular graphs of large girth. Random Structures & Algorithms, **47** (2015), 284-–303.

[21] A. Decelle, F. Krzakala, C. Moore, L. Zdeborová: Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. Physical Review E **84** (2011).

[22] A. Dembo, A. Montanari: Ising models on locally tree-like graphs. Annals of Applied Probability **20** (2010) 565–592.

[23] A. Dembo, A. Montanari, N. Sun: Factor models on locally tree-like graphs. Annals of Probability **41** (2013) 4162–4213.

[24] A. Dembo, A. Montanari, A. Sly, N. Sun: The replica symmetric solution for Potts models on d-regular graphs. Communications in Mathematical Physics **327** (2014) 551–575.

[25] A. Dembo, A. Montanari, S. Sen: Extremal cuts of sparse random graphs. Annals of Probability **45** (2017) 1190–1217.

[26] J. Diaz, N. Do, M. Serna, N. Wormald: Bounds on the max and min bisection of random cubic and random 4-regular graphs. Theoretical Computer Science **307** (2003) 531–547.

[27] J. Diaz, M. Serna, N. Wormald: Bounds on the bisection width for random d-regular graphs. Theoretical Computer Science **382** (2007) 120–130.

[28] U. Feige, M. Karpinski, M. Langberg: Improved approximation of Max-Cut on graphs of bounded degree. Journal of Algorithms **43** (2002) 201–219.

[29] S. Franz, M. Leone: Replica bounds for optimization problems and diluted spin systems. Journal of Statistical Physics **111** (2003) 535—564.
[30] S. Friedli, Y. Velenik: Statistical mechanics of lattice systems: a concrete mathematical introduction. Cambridge University Press (2017).
[31] D. Gamarnik, Q. Li: On the max-cut of sparse random graphs. Random Structures & Algorithms **52** (2018) 219–262.
[32] M. Goemans, D. Williamson: Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. Journal of the ACM **42** (1995) 1115–1145.
[33] F. Guerra: Broken replica symmetry bounds in the mean field spin glass model. Communications in Mathematical Physics **233** (2003) 1–12.
[34] F. Guerra, F. Toninelli: The high temperature region of the Viana-Bray diluted spin glass model. Journal of Statistical Physics **115** (2004) 531–555.
[35] J. Håstad: Some optimal inapproximability results. Journal of the ACM **48** (2001) 798–859.
[36] Y. Higuchi: Remarks on the limiting Gibbs states on a (d+ 1)-tree. Publications of the Research Institute for Mathematical Sciences, **13** (1977) 335–348.
[37] K. Huang (2009). Introduction to statistical physics. CRC press.
[38] V. Kalapala, C. Moore: MAX-CUT on sparse random graphs. University of New Mexico Technical Report TR-CS-2002-24 (2002).
[39] F. Kardos, D. Kral, J. Volec: Maximum edge-cuts in cubic graphs with large girth and in random cubic graphs. Random Structures & Algorithms **41** (2012) 506–520.
[40] H. Kesten, B. Stigum: Additional limit theorem for indecomposable multidimensional Galton-Watson processes. Ann. Math. Statist. **37** (1966) 1463–1481.
[41] S. Khot: On the power of unique 2-prover 1-round games. Proc. 34th STOC (2002) 767–775.
[42] F. Krzakala, A. Montanari, F. Ricci-Tersenghi, G. Semerjian, L. Zdeborová: Gibbs states and the set of solutions of random constraint satisfaction problems. Proc. National Academy of Sciences **104** (2007) 10318–10323.
[43] W. Lenz: Beiträge zum Verständnis der magnetischen Eigenschaften in festen Körpern. Physikalische Zeitschrift **21** (1920) 613–615.
[44] L. Massoulié: Community detection thresholds and the weak Ramanujan property. Proc. 46th STOC (2014) 694–703.
[45] M. Mézard, A. Montanari: Information, physics and computation. Oxford University Press 2009.
[46] M. Mézard, G. Parisi: The Bethe lattice spin glass revisited. Eur. Phys. J. B **20** (2001) 217—233.
[47] M. Mézard, G. Parisi: The cavity method at zero temperature. Journal of Statistical Physics **111** (2003) 1–34.
[48] A. Montanari, S. Sen: Semidefinite programs on sparse random graphs and their application to community detection. Proc. 48th STOC (2016) 814–827.
[49] C. Moore: The computer science and physics of community detection: Landscapes, phase transitions, and hardness. arXiv preprint arXiv:1702.00467 (2017).
[50] E. Mossel, J. Neeman, A. Sly: Reconstruction and estimation in the planted partition model. Probability Theory and Related Fields **162** (2015) 431–461.
[51] E. Mossel, J. Neeman, A. Sly: Consistency thresholds for binary symmetric block models. Proc. 47th STOC (2015).
[52] E. Mossel, J. Neeman, A. Sly: Belief propagation, robust reconstruction and optimal recovery of block models. The Annals of Applied Probability **26** (2016) 2211–2256.
[53] E. Mossel, J. Neeman, A. Sly: A proof of the block model threshold conjecture. Combinatorica **38** (2018) 665–708.
[54] D. Panchenko, M. Talagrand: Bounds for diluted mean-fields spin glass models. Probab. Theory Relat. Fields **130** (2004) 319–336.
[55] D. Panchenko: The Parisi ultrametricity conjecture. Annals of Mathematics (2013) 383–393.
[56] D. Panchenko: The Sherrington-Kirckpatrick model. Springer Science & Business Media (2013).
[57] D. Panchenko: Spin glass models from the point of view of spin distributions. Annals of Probability **41** (2013) 1315–1361.
[58] G. Parisi: A sequence of approximated solutions to the SK model for spin glasses. Journal of Physics A: Mathematical and General **13** (1980) L115.
[59] F. Ricci-Tersenghi, G. Semerjian, L. Zdeborová: Typology of phase transitions in Bayesian inference problems. Physical Review E **99** (2019) 042109.
[60] A. Sly, N. Sun, Y. Zhang: The number of solutions for random regular NAE-SAT. Proc. 57th FOCS (2016) 724–731.
[61] G. Sorkin: Extremal cuts in random cubic graphs. Manuscript (2019).
[62] M. Talagrand: The Parisi formula. Annals of Mathematics (2006) 221–263.
[63] L. Zdeborová, S. Boettcher: A conjecture on the maximum cut and bisection width in random regular graphs. Journal of Statistical Mechanics: Theory and Experiment (2010).

AMIN COJA-OGHLAN, acoghlan@math.uni-frankfurt.de, GOETHE UNIVERSITY, MATHEMATICS INSTITUTE, 10 ROBERT MAYER ST, FRANKFURT 60325, GERMANY.

PHILIPP LOICK, loick@math.uni-frankfurt.de, GOETHE UNIVERSITY, MATHEMATICS INSTITUTE, 10 ROBERT MAYER ST, FRANKFURT 60325, GERMANY.

BALÁZS F. MEZEI, balazs.mezei@cs.ox.ac.uk, OXFORD UNIVERSITY, WOLFSON BUILDING, PARKS ROAD, OXFORD OX1 3QD, UK

GREGORY B. SORKIN, g.b.sorkin@lse.ac.uk, LONDON SCHOOL OF ECONOMICS, DEPARTMENT OF MATHEMATICS, HOUGHTON ST, LONDON WC2A 2AE, UK

APPENDIX H. THE ISING ANTIFERROMAGNET IN THE REPLICA SYMMETRIC PHASE

THE ISING ANTIFERROMAGNET IN THE REPLICA SYMMETRIC PHASE

CHRISTIAN FABIAN, PHILIPP LOICK

ABSTRACT. Partition functions are an important research object in combinatorics and mathematical physics [Barvinok, 2016]. In this work, we consider the partition function of the Ising antiferromagnet on random regular graphs and characterize its limiting distribution in the replica symmetric phase up to the Kesten-Stigum bound. Our proof relies on a careful execution of the method of moments, spatial mixing arguments and small subgraph conditioning.

1. INTRODUCTION

1.1. **Motivation.** The Ising model, invented by Lenz in 1920 to explain magnetism, is a cornerstone in statistical physics. Consider any graph G with vertex set V and edge set E. Each vertex carries one of two possible spins ± 1 and the interactions between vertices are represented by E. For a spin configuration $\sigma \in \{\pm 1\}^V$ on G, we can consider the Hamiltonian $\mathscr{H}_G$

$$\mathscr{H}_G(\sigma) = \sum_{(v,w)\in E} \frac{1+\sigma_v\sigma_w}{2}.$$

Together with a real parameter $\beta > 0$ the Hamiltonian gives rise to a distribution on spin configurations defined by

$$\mu_{G,\beta}(\sigma) = \frac{\exp\left(-\beta\mathscr{H}_G(\sigma)\right)}{Z_{G,\beta}} \quad \left(\sigma \in \{\pm 1\}^V\right) \quad \text{where} \quad Z_{G,\beta} = \sum_{\tau\in\{\pm 1\}^V} \exp\left(-\beta\mathscr{H}_G(\tau)\right). \tag{1.1}$$

The probability measure $\mu_{G,\beta}$ is known as the Boltzmann distribution with the normalizing term $Z_{G,\beta}$ being the partition function. $\mu_{G,\beta}$ favors configurations with few edges between vertices of the same spin which is known as the antiferromagnetic Ising model. There is a corresponding formulation of (1.1) where edges between vertices of the same spin are preferred - the ferromagnetic Ising model. Both models are of great interest in combinatorics and physics and the literature on each is vast [7].

In this paper, we study the Ising antiferromagnet on the random d-regular graph $\mathbb{G} = \mathbb{G}(n,d)$. One might be tempted to think that the regularities of this graph model provide a more amenable study object than its well-known Erdős-Rényi counterpart with fluctuating vertex degrees. However, for the Ising model the reverse seems to be true. Indeed, the independence of edges in the Erdős-Rényi-model greatly facilitates deriving the distribution of short cycles in the planted model and simplifies the calculation of both the first and second moment.

Clearly, $\mu_{\mathbb{G},\beta}$ gives rise to correlations between spins of nearby vertices. The degree of such correlations is governed by the choice of β. A question which is of keen interest in combinatorics and statistical physics is whether such correlations persist for two uniformly sampled (and thus likely distant) vertices. According to physics predictions, for small values of β we should observe a rapid decay of correlation [10] and thus no *long-range correlations*. This regime is known as the *replica symmetric phase*. It is suggested that there exists a specific β which marks the onset of long-range correlations in $\mathbb{G}$. This value is conjectured to be at the combinatorially meaningful Kesten-Stigum bound [3]

$$\beta_{\mathrm{KS}} = \log\left(\frac{\sqrt{d-1}+1}{\sqrt{d-1}-1}\right).$$

The question of long-range correlations is tightly related to the partition function $Z_{\mathbb{G},\beta}$ from which also various combinatorially meaningful observables can be derived. The MAX CUT on random d-regular graphs is a case in point due to the well-known relation

$$\mathrm{MAXCUT}(G) = \frac{dn}{2} + \lim_{\beta\to\infty} \frac{\partial}{\partial\beta} \log Z_{G,\beta}.$$

for any graph G. Thus, it is of key interest to understand the behavior of $Z_{\mathbb{G},\beta}$.

The authors thank Amin Coja-Oghlan for helpful discussions and insights. The authors also thank Mark Sellke for helpful comments. Philipp Loick is supported by DFG CO 646/3.

1.2. **Result.** In recent work, [3] were able to pinpoint the replica symmetry breaking phase transition at the Kesten-Stigum bound, thus charting the replica symmetric phase for the Ising antiferromagnet on random d-regular graphs. The key feature of the replica-symmetric phase is that w.h.p. two independent samples $\boldsymbol{\sigma}_1, \boldsymbol{\sigma}_2$ from the Boltzmann distribution $\mu_{\mathbb{G},\beta}$ exhibit an almost flat overlap in the sense that $|\boldsymbol{\sigma}_1 \cdot \boldsymbol{\sigma}_2| = o(n)$. To be precise, [3] determined $Z_{\mathbb{G},\beta}$ up to an error term $\exp(o(n))$ for $\beta < \beta_{\mathrm{KS}}$. In this paper, we move beyond this crude approximation. By deriving the limiting distribution in the replica-symmetric phase, we show that $Z_{\mathbb{G},\beta}$ is tightly concentrated with bounded fluctuations which we can quantify and attribute to short cycles in $\mathbb{G}$.

Theorem 1.1. *Assume that $0 < \beta < \beta_{KS}$ and $d \geq 3$. Let $(\Lambda_i)_i$ be a sequence of independent Poisson variables with $\mathbb{E}[\Lambda_i] = \lambda_i$ where $\lambda_i = \frac{(d-1)^i}{2i}$. Then as $n \to \infty$ we have*

$$\log\left(Z_{\mathbb{G}(n,d),\beta}\right) - \frac{1}{2}\log\left(\frac{1+e^{\beta}}{2+de^{\beta}-d}\right) - n\left(\left(1-\frac{d}{2}\right)\log(2) + \frac{d}{2}\log\left(1+e^{-\beta}\right)\right) + \frac{d-1}{2}\frac{e^{-\beta}-1}{e^{-\beta}+1} + \frac{(d-1)^2}{4}\left(\frac{e^{-\beta}-1}{e^{-\beta}+1}\right)^2$$
$$\xrightarrow{d} \log(W) := \sum_{i=3}^{\infty} \Lambda_i \log\left(1+\left(\frac{e^{-\beta}-1}{e^{-\beta}+1}\right)^i\right) - \frac{(d-1)^i}{2i}\left(\frac{e^{-\beta}-1}{e^{-\beta}+1}\right)^i.$$

The infinite product defining W converges a.s. and in L^2.

Taking the expectation of this distribution readily recovers the first part of the result by [3]. The proof of Theorem 1.1 relies on the combination of the method of moments and small subgraph conditioning enriched in our case by spatial mixing arguments to make the calculation of the second moment tractable.

2. Techniques

2.1. **Notation.** Let $\mathbb{G} = \mathbb{G}(n,d)$ denote a random d-regular graph on n vertices. We consider sparse graphs with constant d as $n \to \infty$. Throughout the paper, we will employ standard Landau notation with the usual symbols $o(\cdot), O(\cdot), \Theta(\cdot), \omega(\cdot)$, and $\Omega(\cdot)$ to refer to the limit $n \to \infty$. We say that a sequence of events $(\mathscr{E}_n)_n$ holds *with high probability* (w.h.p.) if $\lim_{n\to\infty} \mathbb{P}[\mathscr{E}_n] = 1$. When the context is clear we might drop the index of the expectation. Moreover, we will use the proportional $\propto$ to hide necessary normalisations.

2.2. **Outline.** To get a handle on the distribution of $Z_{\mathbb{G},\beta}$ in the replica symmetric phase, we need to identify the sources of fluctuations of $Z_{\mathbb{G},\beta}$. One obvious source is the number of short cycles. Since $\mathbb{G}$ is sparse and random, standard arguments reveal that $\mathbb{G}$ contains only few short cycles. In the following, let $C_i(G)$ denote the number of short cycles of length i in a graph G and $\mathscr{F}_\ell$ the σ-algebra generated by the random variables $C_i(\mathbb{G})$ for $i \leq \ell$. A key quantity to consider is the variance of $Z_{\mathbb{G},\beta}$. By standard decomposition, we have

$$\mathbb{E}\left[Z_{\mathbb{G},\beta}^2\right] - \mathbb{E}\left[Z_{\mathbb{G},\beta}\right]^2 = \mathbb{E}\left[\mathbb{E}\left[Z_{\mathbb{G},\beta} \mid \mathscr{F}_\ell\right]^2 - \mathbb{E}\left[Z_{\mathbb{G},\beta}\right]^2\right] + \mathbb{E}\left[\mathbb{E}\left[Z_{\mathbb{G},\beta}^2 \mid \mathscr{F}_\ell\right] - \mathbb{E}\left[Z_{\mathbb{G},\beta} \mid \mathscr{F}_\ell\right]^2\right]$$

for any $\ell \geq 1$. Note that the first term of the r.h.s. describes the contribution to the variance by the fluctuations in the number of short cycles, while the second term accounts for the conditional variance given the number of short cycles. It turns out that as $\ell \to \infty$ after taking $n \to \infty$, the second summand vanishes. In other words, the entire variance of $Z_{\mathbb{G},\beta}$ is due to fluctuations in the number of short cycles.

To show this property formally, we leverage a result by [8] that stipulates conditions under which one is able to describe the limiting distribution of $Z_{\mathbb{G},\beta}$ (see Theorem 4.1 in the appendix). One ingredient is the distribution of short cycles in $\mathbb{G}$ and a planted model $\mathbb{G}^*$. In $\mathbb{G}^*$, we first select a spin configuration σ uniformly at random and subsequently sample a graph G with probability proportional to $\exp\left(-\beta\mathscr{H}_G(\sigma)\right)$. While the distribution of short cycles in $\mathbb{G}$ is well established, the distribution of short cycles in the planted model $\mathbb{G}^*$ is a key contribution of this paper. The second ingredient is a careful application of the method of moments. Unfortunately, standard results on the first and second moment on random regular graphs (see i.e. [3]), do not suffice in our case and we have to sharpen our pencils to yield an error term of order $O\left(\exp(1/n)\right)$. While the need for this lower error term prolongs calculations, it also poses some challenges that we resolve by a careful application of the Laplace's method as suggested by [5] and spatial mixing arguments.

2.3. **Short cycles.** To get started, let us write

$$\delta_i = \left(\frac{e^{-\beta}-1}{e^{-\beta}+1}\right)^i \qquad \text{and} \qquad \lambda_i = \frac{(d-1)^i}{2i}. \tag{2.1}$$

The first item on the agenda is to derive the distribution of short cycles in $\mathbb{G}$. This is a well-established result.

Fact 2.1 (Theorem 9.5 in [9])**.** *Let $\Lambda_i \sim \mathrm{Po}(\lambda_i)$ be a sequence of independent Poisson random variables for $i \geq 3$. Then jointly for all i we have $C_i(\mathbb{G}) \xrightarrow{d} \Lambda_i$ as $n \to \infty$.*

Deriving the distribution of short cycles in the planted model $\mathbb{G}^*$ informally introduced above requires some more work. Let us start with the definitions. Given $\sigma \in \{\pm 1\}^V$ and for any $\beta > 0$, let us define the distribution of $\mathbb{G}^*(\sigma)$ for any event $\mathscr{A}$ as

$$\mathbb{P}\left[\mathbb{G}^*(\sigma) \in \mathscr{A}\right] \propto \mathbb{E}\left[\exp\left(-\beta \mathscr{H}_{\mathbb{G}}(\sigma)\right) \mathbf{1}\left\{\mathbb{G} \in \mathscr{A}\right\}\right]. \tag{2.2}$$

This definition gives rise to the following experiment. First, draw a spin configuration $\boldsymbol{\sigma}^*$ uniformly at random among all configurations $\{\pm 1\}^V$. In the next step, draw $\mathbb{G}^* = \mathbb{G}^*(\boldsymbol{\sigma}^*)$ according to (2.2). Hereafter, $\mathbb{G}^*$ will be denoted the planted model.

Proposition 2.2. *Let*

$$\Xi_i \sim \mathrm{Po}\left(\lambda_i\left(1+\delta_i\right)\right)$$

be a sequence of independent Poisson random variables for $i \geq 3$. Then jointly for all i we have $C_i(\mathbb{G}^*) \xrightarrow{d} \Xi_i$ as $n \to \infty$.

Establishing the distribution of short cycles in $\mathbb{G}^*$ is one of the main contributions of this paper. To this end, we start off with similar arguments as used in [11], but need to diligently account for the subtle dependencies introduced by the regularities in $\mathbb{G}^*$.

Applying Fact 2.1 and Proposition 2.2 to Theorem 1 in [8] requires a slight detour via the Nishimori property. To this end, note that the random graph $\mathbb{G}$ induces a reweighted graph distribution $\hat{\mathbb{G}}$ which for any event $\mathscr{A}$ is defined by

$$\mathbb{P}\left[\hat{\mathbb{G}} \in \mathscr{A}\right] \propto \mathbb{E}\left[Z_{\mathbb{G},\beta}\mathbf{1}\left\{\mathbb{G} \in \mathscr{A}\right\}\right]. \tag{2.3}$$

Moreover, consider the distribution $\hat{\boldsymbol{\sigma}}$ on spin configurations defined by

$$\mathbb{P}\left[\hat{\boldsymbol{\sigma}} = \sigma\right] \propto \mathbb{E}\left[\exp\left(-\beta \mathscr{H}_{\mathbb{G}}(\sigma)\right)\right] \tag{2.4}$$

for any $\beta > 0$. $\hat{\mathbb{G}}, \mathbb{G}^*, \hat{\boldsymbol{\sigma}}, \boldsymbol{\sigma}^*$, and the Boltzmann distribution from (1.1) are connected via the well-known Nishimori property.

Fact 2.3 (Proposition 3.2 in [4])**.** *For any graph G and spin configuration $\sigma \in \{\pm 1\}^V$ we have*

$$\mathbb{P}\left[\hat{\mathbb{G}} = G\right]\mu_G(\sigma) = \mathbb{P}\left(\hat{\boldsymbol{\sigma}} = \sigma\right)\mathbb{P}\left(\mathbb{G}^* = G \,|\, \boldsymbol{\sigma}^* = \sigma\right).$$

2.4. **The first and second moment.** The second key ingredient towards the proof of Theorem 1.1 is the method of moments. As standard random regular graph results are too crude, we need a more precise calculation. Fortunately, with some patience and equipped with Laplace's method as stated in [5], the first moment is not too hard to find.

Proposition 2.4. *Assume that $0 < \beta < \beta_{KS}$ and $d \geq 3$. Then we have*

$$\mathbb{E}\left[Z_{\mathbb{G},\beta}\right] = \exp\left(-\lambda_1\delta_1 - \lambda_2\delta_2 + O\left(\frac{1}{n}\right)\right)\sqrt{\frac{1+e^{\beta}}{2+de^{\beta}-d}}\exp\left(n\left((1-d/2)\log(2) + d\log\left(1+e^{-\beta}\right)/2\right)\right)$$

The second moment is not as amenable. The key challenge for applying Laplace's method is to exhibit that the obvious choice of the optimum is indeed a global maximum. We resolve this issue by resorting to results on the broadcasting process on an infinite d-regular tree and the disassortative stochastic block model. This spatial mixing argument allows us to focus our attention on an area close to the anticipated optimum. To this end, let us exhibit an event $\mathscr{O}$ that is concerned with the location of two typical samples $\boldsymbol{\sigma}_{\mathbb{G}}, \boldsymbol{\sigma}'_{\mathbb{G}}$ from the Boltzmann distribution $\mu_{\mathbb{G},\beta}$, i.e.

$$\mathscr{O} = \left\{\mathbb{E}\left[|\boldsymbol{\sigma}_{\mathbb{G}} \cdot \boldsymbol{\sigma}'_{\mathbb{G}}| \,|\, \mathbb{G}\right] < \varepsilon_n n\right\} \tag{2.5}$$

for a sequence of $\varepsilon_n = o(1)$. Then we can leverage the following result from [3].

Lemma 2.5 (Lemma 4.7 in [3])**.** *For the event $\mathscr{O}$ defined in* (2.5) *we have for $d \geq 3, 0 < \beta < \beta_{KS}$*

$$\mathbb{E}\left[Z_{\mathbb{G},\beta}\mathbf{1}\left\{\mathscr{O}\right\}\right] = (1-o(1))\mathbb{E}\left[Z_{\mathbb{G},\beta}\right].$$

Conditioning on $\mathscr{O}$ greatly facilitates the calculation of the second moment.

Proposition 2.6. *For $0 < \beta < \beta_{KS}$ and $d \geq 3$ we have*

$$\mathbb{E}\left[Z^2_{\mathbb{G},\beta}\mathbf{1}\left\{\mathscr{O}\right\}\right] = \exp\left(\lambda_1 + \lambda_2 - \frac{4\lambda_1}{\left(1+e^{\beta}\right)^2} - \frac{4\lambda_2\left(1+e^{2\beta}\right)^2}{\left(1+e^{\beta}\right)^4} + O\left(\frac{1}{n}\right)\right)\frac{\left(1+e^{\beta}\right)^2\exp\left(n\left((2-d)\log(2) + d\log\left(1+e^{-\beta}\right)\right)\right)}{\left(de^{\beta}-d+2\right)\sqrt{2e^{2\beta}+2de^{\beta}-de^{2\beta}-d+2}}$$

2.5. **Proof of Theorem 1.1.** We apply Theorem 1 in [8] to the random variable $Z_{\mathbb{G},\beta}\mathbf{1}\{\mathscr{O}\}$. Condition (1) readily follows from Fact 2.1. For Condition (2) let us write

$$\mathscr{C}(G) = \{C_1(G) = c_1, \ldots, C_\ell(G) = c_\ell\}$$

for any graph G. By Lemma 2.5 considering $Z_{\mathbb{G},\beta}$ rather than $Z_{\mathbb{G},\beta}\mathbf{1}\{\mathscr{O}\}$ only introduces an error of order $1+o(1)$ in Condition (2). Using standard reformulations and the definition of $\hat{\mathbb{G}}$ from (2.3) we find

$$\frac{\mathbb{E}\left[Z_{\mathbb{G},\beta} \mid \mathscr{C}(\mathbb{G})\right]}{\mathbb{E}\left[Z_{\mathbb{G},\beta}\right]} = \frac{\mathbb{E}\left[Z_{\mathbb{G},\beta}\mathbf{1}\{\mathscr{C}(\mathbb{G})\}\right]}{\mathbb{P}\left[\mathscr{C}(\mathbb{G})\right]\mathbb{E}\left[Z_{\mathbb{G},\beta}\right]} = \frac{\mathbb{P}\left[\mathscr{C}(\hat{\mathbb{G}})\right]}{\mathbb{P}\left[\mathscr{C}(\mathbb{G})\right]} = \frac{\mathbb{E}_{\hat{\boldsymbol{\sigma}}}\left[\mathbb{P}\left[\mathscr{C}(\hat{\mathbb{G}}) \mid \hat{\boldsymbol{\sigma}}\right]\right]}{\mathbb{P}\left[\mathscr{C}(\mathbb{G})\right]}.$$

Since a typical sample σ from $\hat{\boldsymbol{\sigma}}$ has the property that $|\sigma\cdot\mathbf{1}| = O\left(n^{2/3}\right)$, i.e. is relatively balanced, the Nishimori property (Fact 2.3) implies

$$\mathbb{E}_{\hat{\boldsymbol{\sigma}}}\left[\mathbb{P}\left[\mathscr{C}(\hat{\mathbb{G}}) \mid \hat{\boldsymbol{\sigma}}\right]\right] \sim \mathbb{P}\left[\mathscr{C}(\mathbb{G}^*)\right].$$

Condition (2) now follows from Fact 2.1 and Proposition 2.2. For Condition (3) consider any $\beta = \beta_{\mathrm{KS}} - \varepsilon$ for some small $\varepsilon > 0$. Letting $\eta = \eta(\varepsilon) > 0$ a simple calculation reveals

$$\sum_{i\geq 1}\lambda_i\delta_i^2 \leq \sum_{i\geq 1}\lambda_i\left(\frac{e^{-\beta_{\mathrm{KS}}+\varepsilon}-1}{e^{-\beta_{\mathrm{KS}}+\varepsilon}+1}\right)^{2i} = \sum_{i\geq 1}\frac{(1-\eta)^i}{2i} < \infty$$

which also implies $\sum_{i\geq 3}\lambda_i\delta_i^2 < \infty$. Finally, by Lemma 2.5, Propositions 2.4 and 2.6 and the fact that for any $0 < x < 1$ $\log(1-x) = -\sum_{i\geq 1}x^i/i$ we find for $0 < \beta < \beta_{\mathrm{KS}}$ and $d \geq 3$

$$\begin{aligned}
&\frac{\mathbb{E}\left[Z_{\mathbb{G},\beta}^2\mathbf{1}\{\mathscr{O}\}\right]}{\mathbb{E}\left[Z_{\mathbb{G},\beta}\mathbf{1}\{\mathscr{O}\}\right]^2} = (1+o(1))\frac{\mathbb{E}\left[Z_{\mathbb{G},\beta}^2\mathbf{1}\{\mathscr{O}\}\right]}{\mathbb{E}\left[Z_{\mathbb{G},\beta}\right]^2}\\
&= (1+o(1))\frac{1+e^{\beta}}{\sqrt{2+2e^{2\beta}+2de^{\beta}-d-de^{2\beta}}}\exp\left(\lambda_1+\lambda_2-\frac{4\lambda_1}{\left(1+e^{\beta}\right)^2}-\frac{4\lambda_2\left(1+e^{2\beta}\right)^2}{\left(1+e^{\beta}\right)^4}+2\lambda_1\delta_1+2\lambda_2\delta_2\right)\\
&= (1+o(1))\left(1-(d-1)\left(\frac{e^{-\beta}-1}{e^{-\beta}+1}\right)\right)^{-1/2}\exp\left(-\lambda_1\delta_1^2-\lambda_2\delta_2^2\right) = (1+o(1))\exp\left(\sum_{i\geq 3}\lambda_i\delta_i^2\right)
\end{aligned}$$

establishing Condition (4) and thus the distribution of $Z_{\mathbb{G},\beta}\mathbf{1}\{\mathscr{O}\}$. Since $\mathbb{E}\left[Z_{\mathbb{G},\beta}(1-\mathbf{1}\{\mathscr{O}\})\right] = o\left(\mathbb{E}\left[Z_{\mathbb{G},\beta}\right]\right)$ by Lemma 2.5, Theorem 1.1 follows from Markov's inequality.

3. Discussion

Studying partition functions has a long tradition in combinatorics and mathematical physics. k-SAT, q-coloring or the stochastic block model are just some noteworthy examples where the partition function reveals fundamental and novel combinatorial insights. Due to its connection to the Max Cut problem and the disassortative stochastic block model, the Ising antiferromagnet fits nicely into this list. For random d-regular graphs, Coja-Oghlan et al. [3] pinpointed its replica symmetry breaking phase transition at the Kesten-Stigum bound. Using the method of moments and spatial mixing arguments, they they determine $Z_{\mathbb{G},\beta}$ up to $\exp(o(n))$. In this paper, we move beyond this approximation and derive the limiting distribution of $Z_{\mathbb{G},\beta}$ in the replica symmetric regime. We note that the distribution of $Z_{\mathbb{G},\beta}$ above the Kesten-Stigum bound is fundamentally different. A similar analysis for the Erdős-Rényi-model was carried out in [11].

Using the combination of the method of moments and small subgraph conditioning underlying our proof was initially pioneered by Robinson & Wormald [12] to prove that cubic graphs are w.h.p. Hamiltonian. Janson [8] subsequently showed that small subgraph conditioning can be used to obtain limiting distributions. This strategy was successfully applied, among others, to the stochastic block model [11] and the Viana-Bray model [6]. For other problems, the second moment appears to be too crude for the entire replica symmetric phase and enhanced techniques are needed [2]. In this work, we enrich the classical strategy of the method of moments and small subgraph conditioning by spatial mixing arguments to cover the entire replica symmetric phase.

An interesting remaining question is to throw a bridge between the properties of the partition function $Z_{\mathbb{G},\beta}$ and long-range correlations in $\mathbb{G}$. While it should be a small step from Theorem 1.1 to vindicate the absence of long-range correlations in the replica symmetric phase, proving the presence of long-range correlations above the Kesten-Stigum bound is a more challenging, yet important endeavour.

Before moving to the proofs of Propositions 2.2, 2.4 and 2.6, let us introduce some additional notation. With $\mathscr{P}$ denoting the set of all probability distributions on a finite set $\Omega \neq \emptyset$ and two probability measures $\mu, \nu \in \mathscr{P}(\Omega)$, let us introduce the entropy $H(\mu)$ and Kullback-Leibler divergence $D_{\mathrm{KL}}(\mu\|\nu)$

$$H(\mu) = -\sum_{\omega\in\Omega} \mu(\omega)\log\mu(\omega) \qquad \text{and} \qquad D_{\mathrm{KL}}(\mu\|\nu) = \sum_{\omega\in\Omega} \mu(\omega)\log\frac{\mu(\omega)}{\nu(\omega)} \in [0,\infty].$$

Note the convention $0\cdot\log\left(\frac{0}{0}\right) = 0$ and furthermore that if there exists some $\omega \in \Omega$ such that $\mu(\omega) > 0$ and $\nu(\omega) = 0$, this implies $D_{\mathrm{KL}}(\mu\|\nu) = \infty$. When we consider the product measure between two probability distribution μ and ν, we will use the notation $\mu \otimes \nu$.

Next, let us state a fundamental result by Janson [8] which stipulates conditions under which one is able to obtain the limiting distribution of the partition function.

Theorem 4.1 (Theorem 1 in [8])**.** *Let $\lambda_i > 0$ and $\delta_i \geq -1, i = 1,2,\ldots,$ be constants and suppose that for each n there are random variables $C_{in}, i = 1,2,\ldots,$ and Z_n (defined on the same probability space) such that X_{in} is non-negative integer valued and $\mathbb{E}[Z_n] \neq 0$ (at least of large n), and furthermore the following conditions are satisfied:*

(1) $C_{in} \xrightarrow{d} \Lambda_i$ as $n \to \infty$, jointly for all i where $\Lambda_i \sim \mathrm{Po}(\lambda_i)$ are independent Poisson random variables;
(2) For any finite sequence $c_1,\ldots,c_m$ of non-negative integers,

$$\frac{\mathbb{E}[Z_n \mid C_{1n} = c_1,\ldots,C_{mn} = c_m]}{\mathbb{E}[Z_n]} \to \prod_{i=1}^{m}(1+\delta_i)^{x_i}\exp(-\lambda_i\delta_i) \qquad \text{as} \quad n\to\infty;$$

(3) $\sum_i \lambda_i\delta_i^2 < \infty$;
(4) $\mathbb{E}\left[Z_n^2\right]/(\mathbb{E}[Z_n])^2 \to \exp\left(\sum_i \lambda_i\delta_i^2\right) \qquad \text{as} \quad n\to\infty.$

Then, we have

$$\frac{Z_n}{\mathbb{E}[Z_n]} \xrightarrow{d} W = \prod_{i\geq 1}(1+\delta_i)^{\Lambda_i}\exp(-\lambda_i\delta_i);$$

moreover, this and the convergence in (1) hold jointly. The infinite product defining W converges a.s. and in L^2, with $\mathbb{E}(W) = 1$ and $\mathbb{E}(W^2) = \exp\left(\sum_{i=1}^{\infty}\lambda_i\delta_i^2\right)$. Hence, the normalized variables $Y_n/\mathbb{E}(Y_n)$ are uniformly square integrable. Furthermore, the event $W > 0$ equals, up to a set of probability zero, the event that $Z_i > 0$ for some i with $\delta_i = -1$. In particular, $W > 0$ a.s. if and only if every $\delta_i > -1$.

A substantial part of this paper is devoted to determining the first and second moment of $Z_{\mathbb{G}}$. As we will see in due course, this task requires a special version of the well-known Laplace's method, which is usually formulated in terms of integrals. In contrast to that, the model considered here is discrete and therefore requires a variation of Laplace's method which is applicable to countable sums. Fortunately, [5] provides an adaptation that we can leverage here. Let us start by providing the result of interest:

Theorem 4.2 (Theorem 2.3 in [5])**.** *Suppose the following:*

(1) $\mathscr{L} \subset \mathbb{R}^N$ is a lattice with rank $r \leq N$.
(2) $V \subseteq \mathbb{R}^N$ is the r-dimensional subspace spanned by $\mathscr{L}$.
(3) $W = V + w$ is an affine subspace parallel to V, for some $w \in \mathbb{R}^N$.
(4) $K \subset \mathbb{R}^N$ is a compact convex set with non empty interior K°.
(5) $\phi : K \to \mathbb{R}$ is a continuous function and the restriction of ϕ to $K \cap W$ has a unique maximum at some point $x_0 \in K^\circ \cap W$.
(6) ϕ is twice continuously differentiable in a neighbourhood of x_0 and $H := D^2\phi(x_0)$ is its Hessian at x_0.
(7) $\psi : K_1 \to \mathbb{R}$ is a continuous function on some neighbourhood $K_1 \subseteq K$ of x_0 with $\psi(x_0) > 0$.
(8) For each positive integer n there is a vector $\ell_n \in \mathbb{R}^N$ with $\frac{\ell_n}{n} \in W$.
(9) For each positive integer n there is a positive real number b_n and a function $a_n : (\mathscr{L}+\ell_n)\cap nK \to \mathbb{R}$ such that, as $n \to \infty$,

$$a_n(\ell) = O\left(b_n e^{n\phi(\ell/n)+o(n)}\right), \qquad \ell \in (\mathscr{L}+\ell_n)\cap nK,$$

and

$$a_n(\ell) = b_n\left(\psi\left(\frac{\ell}{n}\right)+o(1)\right)e^{n\phi(\ell/n)}, \qquad \ell \in (\mathscr{L}+\ell_n)\cap nK_1,$$

uniformly for ℓ in the indicated sets.

Then, provided $\det(-H|_V) \neq 0$, *as* $n \to \infty$,

$$\sum_{\ell \in (\mathscr{L}+\ell_n)\cap nK} a_n(\ell) \sim \frac{(2\pi n)^{r/2}\,\psi(x_0)\,b_n e^{n\phi(x_0)}}{\det(\mathscr{L})\sqrt{\det(-H|_V)}}.$$

Theorem 4.2 is largely self-explanatory. The concept of lattices, however, is not obvious from the theorem itself. Therefore, we briefly revisit the idea of lattices and how they are connected to our model. In general, lattices are discrete subgroups of $\mathbb{R}^N$ where each lattice is isomorphic to $\mathbb{Z}^r$ for some $0 \le r \le N$. In this context, *discrete* simply means that the intersection of a lattice with an arbitrary, bounded set in $\mathbb{R}^N$ is finite. Furthermore, r is commonly called the rank of the respective lattice. This means that each lattice has a (not necessarily unique) basis consisting of the vectors $x_1, \dots, x_r$. The crucial characteristics of these basis vectors are on the one hand that they are independent. On the other hand, every element of the respective lattice has a unique representation of the form $\sum_{i=1}^r k_i \cdot x_i$ where $k_i \in \mathbb{Z}$ for all $i \in [r]$.
In applying Theorem 4.2 we are especially interested in understanding the determinant $\det(\mathscr{L})$ for a given lattice $\mathscr{L}$. Formally, $\det(\mathscr{L})$ is simply obtained by calculating the determinant of the matrix that consists of the basis vectors $x_1, \dots, x_r$ mentioned above. Intuitively, the determinant provides the r-dimensional volume of a unit cell of the lattice $\mathscr{L}$. Note that the term $(\det(\mathscr{L}))^{-1}$ in Theorem 4.2 is the key difference compared to more common versions of Laplace's method for integrals.

5. Short Cycles in the Regular Stochastic Block Model / Proof of Proposition 2.2

Let us start with a brief repetition of the Regular Stochastic Block Model (RSBM) which is the result of the following experiment. Given a vertex set $V_n = \{v_1, \dots, v_n\}$, we first sample a spin configuration uniformly at random. We denote this uniformly sampled configuration by $\boldsymbol{\sigma}^*$. Next, we draw a d-regular graph $\mathbb{G}^* = \mathbb{G}^*(\boldsymbol{\sigma}^*)$ from the distribution

$$\mathbb{P}\left[\mathbb{G}^* = G | \boldsymbol{\sigma}^* = \sigma\right] \propto \exp\left(-\beta \mathscr{H}_G(\sigma)\right).$$

For some graph d-regular G with n nodes and some spin configuration $\sigma \in \{\pm 1\}^n$ on the nodes of G we define

$$\mu_{++}(G,\sigma) := \frac{2}{dn} \sum_{(u,v)\in E} \mathbf{1}\{\sigma(v) = \sigma(u) = +1\}. \tag{5.1}$$

Since G has $\frac{dn}{2}$ edges in total, μ_{++} simply measures the fraction of edges that connect two positive vertices. Analogously, we define

$$\mu_{--}(G,\sigma) := \frac{2}{dn} \sum_{(u,v)\in E} \mathbf{1}\{\sigma(v) = \sigma(u) = -1\} \qquad \text{and} \tag{5.2}$$

$$\mu_{+-}(G,\sigma) = \mu_{-+}(G,\sigma) := \frac{1}{dn} \sum_{(u,v)\in E} \mathbf{1}\{\sigma(v) \neq \sigma(u)\}. \tag{5.3}$$

Due to the fact that our model is built on undirected edges, we just count all the edges connecting vertices with different spins and evenly 'split' them between μ_{+-} and μ_{-+}. In a similar way, we define

$$\rho_+(G,\sigma) := \frac{1}{n} \sum_{v\in V} \mathbf{1}\{\sigma(v) = +1\} \qquad \text{and} \qquad \rho_-(G,\sigma) := \frac{1}{n} \sum_{v\in V} \mathbf{1}\{\sigma(v) = -1\} \tag{5.4}$$

where ρ_+ and ρ_- depict the fractions of nodes that have been assigned a positive spin or a negative one, respectively. For notational convenience, we usually drop the reference to the graph G and the spin configuration σ. Accordingly, let $\mu' = \mu(\mathbb{G}^*, \boldsymbol{\sigma}^*)$ and $\mathscr{M}(\sigma)$ denote the set of all probability distributions fulfilling the obvious symmetry and marginalization conditions, i.e.

$$\mu_{++} + \mu_{+-} = \rho_+, \qquad \mu_{--} + \mu_{+-} = \rho_-, \qquad \mu_{+-} = \mu_{-+}$$

and where $\mu_{++}dn/2, \mu_{--}dn/2$ and $\mu_{+-}dn/2$ are integers. Further, we define a probability measure $\hat{\mu}$ with

$$\hat{\mu}_{++} = \hat{\mu}_{--} = \frac{e^{-\beta}}{2\left(1+e^{-\beta}\right)} \qquad \text{and} \qquad \hat{\mu}_{+-} = \hat{\mu}_{-+} = \frac{1}{2\left(1+e^{-\beta}\right)}.$$

To determine the distribution of short cycles in the RSBM, we start by considering the event

$$\mathscr{A}_\mu := \left\{\left\|\mu' - \mu\right\| = O\left(n^{-1/2}\log n\right)\right\}.$$

In the next lines, we establish that $\mathscr{A}_{\hat{\mu}}$ is a high probability event.

Lemma 5.1. *We have* $\mathbb{P}\left[\mathscr{A}_{\hat{\mu}}\right] = 1 - o(1)$.

Proof of Lemma 5.1. In the following we will write μ for $\mu(G,\sigma)$ when the reference to G and σ is obvious. For this proof, we leverage some results that are derived in detail in Section 6. More specifically, we consider equation (6.3), that is

$$\mathbb{E}\left[Z_{\mathbb{G}(n,d),\beta}\right] = \exp\left(O\left(\frac{1}{n}\right)\right) \cdot \sum_{(\rho_+,\mu_{++})\in\mathscr{Q}} \frac{1}{\pi n\sqrt{2\mu_{++}\mu_{--}\mu_{+-}d}} \exp\left(n\psi\left(\mu_{++},\rho_+\right)\right).$$

where

$$\psi\left(\mu_{++},\rho_+\right) := \mathrm{H}\left(\rho\right) - \frac{d}{2}\left(D_{\mathrm{KL}}(\mu||\rho\otimes\rho) + \beta\left(1+2\mu_{++}-2\rho_+\right)\right)$$

and $\mathscr{Q}$ is the set of all conceivable pairs $\left(\rho_+,\mu_{++}\right)$. Furthermore, from Lemma 6.5 we know that $\psi\left(\mu_{++},\rho_+\right)$ obtains it unique maximum on $\mathscr{Q}$ at $\left(\hat{\mu}_{++},\hat{\rho}_+\right) = \left(\frac{e^{-\beta}}{2(1+e^{-\beta})},\frac{1}{2}\right)$. The entries of the Hessian turn out to be

$$\begin{aligned}
\frac{\partial^2\psi}{\partial\mu_{++}^2}\left(\hat{\mu},\hat{\rho}\right) &= -2d\left(1+e^{-\beta}\right)^2 e^{\beta} = \Theta(1)\\
\frac{\partial^2\psi}{\partial\mu_{++}\partial\rho_+}\left(\hat{\mu},\hat{\rho}\right) &= 2d\left(1+e^{-\beta}\right)^2 e^{\beta} = \Theta(1)\\
\frac{\partial^2\psi}{\partial\rho_+^2}\left(\hat{\mu},\hat{\rho}\right) &= -4-2d\left(1+e^{-\beta}+2e^{\beta}\right) = \Theta(1).
\end{aligned}$$

Note that a detailed calculation of the Hessian can be found in Section 6. With all these results at hand, the two dimensional Taylor expansion of ψ at $\left(\hat{\mu},\hat{\rho}\right)$ turns out to be

$$\begin{aligned}
\psi\left(\mu_{++},\rho_+\right) &= \psi\left(\hat{\mu},\hat{\rho}\right) + \Theta\left(1\right)\left(\left(\rho_+-\hat{\rho}_+\right)^2 + \left(\mu_+-\hat{\mu}_+\right)^2 + \left(\rho_+-\hat{\rho}_+\right)\left(\mu_+-\hat{\mu}_+\right)\right) + O\left(\left\|\mu-\hat{\mu}\right\|^3\right)\\
&= \psi\left(\hat{\mu},\hat{\rho}\right) + \Theta\left(\left\|\mu-\hat{\mu}\right\|^2\right)
\end{aligned}$$

where we exploited that the higher order derivatives are bounded. Keeping this in mind, we obtain

$$\begin{aligned}
&\exp\left(O\left(\frac{1}{n}\right)\right) \cdot \sum_{(\rho_+,\mu_{++})\in\mathscr{Q}} \frac{1}{\pi n\sqrt{2\mu_{++}\mu_{--}\mu_{+-}d}} \exp\left(n\psi\left(\mu_{++},\rho_+\right)\right)\mathbf{1}\{1-\mathscr{A}_{\hat{\mu}}\}\\
&= \sum_{(\rho_+,\mu_{++})\in\mathscr{Q}} \exp\left(n\psi\left(\hat{\mu},\hat{\rho}\right) - \Omega\left(\log^2 n\right)\right)\mathbf{1}\{1-\mathscr{A}_{\hat{\mu}}\}\\
&= O\left(n^2\right)\exp\left(n\psi\left(\hat{\mu},\hat{\rho}\right) - \Omega\left(\log^2 n\right)\right) = O\left(n^{-\log n}\right)\exp\left(n\psi\left(\hat{\mu},\hat{\rho}\right)\right)
\end{aligned}$$

which in turn yields

$$\begin{aligned}
\mathbb{E}\left[Z_{\mathbb{G},\beta}\right] &= \exp\left(O\left(\frac{1}{n}\right)\right) \cdot \sum_{(\rho_+,\mu_{++})\in\mathscr{Q}} \frac{1}{\pi n\sqrt{2\mu_{++}\mu_{--}\mu_{+-}d}} \exp\left(n\psi\left(\mu_{++},\rho_+\right)\right)\mathbf{1}\{\mathscr{A}_{\hat{\mu}}\}\\
&\quad + \exp\left(O\left(\frac{1}{n}\right)\right) \cdot \sum_{(\rho_+,\mu_{++})\in\mathscr{Q}} \frac{1}{\pi n\sqrt{2\mu_{++}\mu_{--}\mu_{+-}d}} \exp\left(n\psi\left(\mu_{++},\rho_+\right)\right)\mathbf{1}\{1-\mathscr{A}_{\hat{\mu}}\}\\
&= (1+o(1)) \cdot \sum_{(\rho_+,\mu_{++})\in\mathscr{Q}} \frac{1}{\pi n\sqrt{2\mu_{++}\mu_{--}\mu_{+-}d}} \exp\left(n\psi\left(\mu_{++},\rho_+\right)\right)\mathbf{1}\{\mathscr{A}_{\hat{\mu}}\}\\
&= (1+o(1))\,\mathbb{E}\left[Z_{\mathbb{G},\beta}\mathbf{1}\{\mathscr{A}_{\hat{\mu}}\}\right].
\end{aligned}$$

Now, the proof is almost completed. Corollary 4.5 in [3] states that iff $\mathbb{E}\left[Z_{\mathbb{G},\beta}\right] = (1+o(1))\,\mathbb{E}\left[Z_{\mathbb{G},\beta}\mathbf{1}\{\mathscr{A}_{\hat{\mu}}\}\right]$ holds, we have $\mathbb{P}\left[\mathbb{G}^*\in\mathscr{A}_{\hat{\mu}}\right] = 1-o(1)$. This is just the desired statement. □

The following preliminary arguments combine ideas from [9] and [11] to derive the distribution of short cycles in $\mathbb{G}^*$. We apply the method of moments to derive expected values conditional on μ' being close to $\hat{\mu}$. Then, with Lemma 5.1, we draw conclusions for the unconditional expectation. Let $C_l(\mathbb{G}^*)$ be the number of cycles of length l in $\mathbb{G}^*$. Furthermore, let M denote the number of edges $e_1,\dots e_l$ that connect vertices with opposite spins. This construction immediately implies that M is an even number. Let us briefly recap the configuration model to construct a d-regular graph on n uniformly at random. To get started, we take d copies of each of the n nodes. Thus, we have dn nodes in total. In the next step, we choose a perfect matching uniformly at random. To obtain a graph with n nodes again, we merge the d copies of each node, providing a graph with $\frac{dn}{2}$ edges in total. Since this procedure does not rule out self-loops or double-edges, we condition on the event $\mathscr{S}$ that we obtain a simple graph. Note that standard results from the literature entail

that $\mathbb{P}[G \in \mathscr{S}] = \Omega(1)$. Similarly, conditional on $\mathscr{S}$, each of the admissable d-regular graphs is created with the same probability.
Now recall the probability to observe a specific graph in the regular stochastic block model

$$\mathbb{P}\left[\mathbb{G}^* = G|\sigma\right] \propto \exp\left(-\beta \mathscr{H}_G(\sigma)\right). \tag{5.5}$$

Clearly, the definition of $\mathbb{G}^*$ does not give rise to a uniform distribution over all admissable graphs. However, it is easy to see that (5.5) yields a uniform distribution over all graphs exhibiting a specific μ. This observation is central towards deriving the distribution of short cycles in $\mathbb{G}^*$.

Lemma 5.2. *Let*

$$\Xi_i \sim \mathrm{Po}(\lambda_i(1+\delta_i))$$

be a sequence of independent Poisson random variables for $i \geq 3$. Then jointly for all i we have $C_i(\mathbb{G}^*)|_{\hat{\mu}} \xrightarrow{d} \Xi_i$ as $n \to \infty$.

Proof. Let $p_{l,M}$ be the probability that any given set of l edges where l_{++} edges connect two positive vertices and l_{--} edges connect two negative edges results from the construction of $\mathbb{G}^*$ conditioned on some some fixed μ. We readily find

$$p_{l,M}(\mu) = \frac{\binom{dn\rho_+ - 2l_{++} - M}{dn\mu_{++} - 2l_{++}}(dn\mu_{++} - 2l_{++} - 1)!!}{\binom{dn\rho_+}{dn\mu_{++}}(dn\mu_{++}-1)!!} \cdot \frac{\binom{dn\rho_- - 2l_{--} - M}{dn\mu_{--} - 2l_{--}}(dn\mu_{--} - 2l_{--} - 1)!!(dn\mu_{+-} - M)!}{\binom{dn\rho_-}{dn\mu_{--}}(dn\mu_{--}-1)!!(dn\mu_{+-})!}.$$

Using the following well-known identity.

$$(2k-1)!! = \frac{(2k)!}{k!2^k} \tag{5.6}$$

we find

$$\begin{aligned}\frac{(dn\mu_{++} - 2l_{++} - 1)!!(dn\mu_{--} - 2l_{--} - 1)!!(dn\mu_{+-} - M)!}{(dn\mu_{++}-1)!!(dn\mu_{--}-1)!!(dn\mu_{+-})!} &= \frac{\frac{(dn\mu_{++}-2l_{++})!}{\left(\frac{dn}{2}\mu_{++}-l_{++}\right)!2^{\frac{dn}{2}\mu_{++}-l_{++}}} \cdot \frac{(dn\mu_{--}-2l_{--})!}{\left(\frac{dn}{2}\mu_{--}-l_{--}\right)!2^{\frac{dn}{2}\mu_{--}-l_{--}}}}{\frac{(dn\mu_{++})!}{\left(\frac{dn}{2}\mu_{++}\right)!2^{\frac{dn}{2}\mu_{++}}} \cdot \frac{(dn\mu_{--})!}{\left(\frac{dn}{2}\mu_{--}\right)!2^{\frac{dn}{2}\mu_{--}}} \cdot (dn\mu_{+-})_M} \\ &= 2^{l-M} \cdot \frac{\left(\frac{dn}{2}\mu_{++}\right)_{l_{++}}\left(\frac{dn}{2}\mu_{--}\right)_{l_{--}}}{(dn\mu_{++})_{2l_{++}}(dn\mu_{--})_{2l_{--}}(dn\mu_{+-})_M}. \end{aligned} \tag{5.7}$$

Moving on to the binomial coefficients and using Stirling's formula

$$k! = \sqrt{2\pi k}\left(\frac{k}{e}\right)^k \exp\left(O\left(\frac{1}{k}\right)\right) \tag{5.8}$$

we obtain

$$\begin{aligned}\frac{\binom{dn\rho_+ - 2l_{++} - M}{dn\mu_{++} - 2l_{++}}\binom{dn\rho_- - 2l_{--} - M}{dn\mu_{--} - 2l_{--}}}{\binom{dn\rho_+}{dn\mu_{++}}\binom{dn\rho_-}{dn\mu_{--}}} &= \frac{(dn\mu_{++})_{2l_{++}}(dn\mu_{--})_{2l_{--}}}{(dn\rho_+)_{2l_{++}+M}(dn\rho_-)_{2l_{--}+M}} \cdot \frac{(dn\rho_+ - dn\mu_{++})!(dn\rho_- - dn\mu_{--})!}{(dn\rho_+ - dn\mu_{++} - M)!(dn\rho_- - dn\mu_{--} - M)!} \\ &= \frac{(dn\mu_{++})_{2l_{++}}(dn\mu_{--})_{2l_{--}}(dn\rho_+ - dn\mu_{++})_M(dn\rho_- - dn\mu_{--})_M}{(dn\rho_+)_{2l_{++}+M}(dn\rho_-)_{2l_{--}+M}} \end{aligned} \tag{5.9}$$

Combining (5.7) and (5.9), we yield

$$p_{l,M}(\mu) = 2^{l-M} \cdot \frac{\left(\frac{dn}{2}\mu_{++}\right)_{l_{++}}\left(\frac{dn}{2}\mu_{--}\right)_{l_{--}}(dn\rho_+ - dn\mu_{++})_M(dn\rho_- - dn\mu_{--})_M}{(dn\rho_+)_{2l_{++}+M}(dn\rho_-)_{2l_{--}+M}(dn\mu_{+-})_M}.$$

In particular, we thus have for all $\hat{\mu}' \in \mathscr{A}_{\hat{\mu}}$

$$\begin{aligned} p_{l,M}(\hat{\mu}') &= 2^{l-M} \frac{\left(\frac{dn}{2}\frac{e^{-\beta}}{2(1+e^{-\beta})}\right)^{l_{++}}\left(\frac{dn}{2}\frac{e^{-\beta}}{2(1+e^{-\beta})}\right)^{l_{--}}\left(dn\frac{1}{2(1+e^{-\beta})}\right)^M\left(dn\frac{1}{2(1+e^{-\beta})}\right)^M}{\left(\frac{dn}{2}\right)^{2l_{++}+M}\left(\frac{dn}{2}\right)^{2l_{--}+M}\left(dn\frac{1}{2(1+e^{-\beta})}\right)^M}(1+o(1)) \\ &= \frac{2^{l-M}\left(\frac{e^{-\beta}}{2(1+e^{-\beta})}\right)^{l-M}\left(\frac{1}{1+e^{-\beta}}\right)^{2M}}{\left(\frac{dn}{2}\right)^{l_{++}}\left(\frac{dn}{2}\right)^{l_{--}}\left(dn\frac{1}{2(1+e^{-\beta})}\right)^M}(1+o(1)) = \left(\frac{2}{dn}\right)^l\left(\frac{e^{-\beta}}{(1+e^{-\beta})}\right)^{l-M}\left(\frac{1}{(1+e^{-\beta})}\right)^M(1+o(1)). \end{aligned}$$

We point out that $p_{l,M}(\hat{\mu}')$ can asymptotically be expressed without l_{++} and l_{--}. Next, we consider the number of possible cycles with length l and exactly M edges that connect vertices with opposite spins, subsequently denoted by $a_{l,M}(\mu)$. For starters, we have

$$2l \cdot a_{l,M}(\mu) = 2\binom{l}{M}(n\rho_+)_{l_+}(n\rho_-)_{l_-}(d(d-1))^l.$$

This implies for $\hat{\mu}' \in \mathscr{A}_{\hat{\mu}}$

$$a_{l,M}(\hat{\mu}') = \binom{l}{M}\frac{1}{l}n^l 2^{-l}(d(d-1))^l(1+o(1)).$$

Now, we are in a position to calculate the conditional expectation of the number of short cycles, that is

$$\begin{aligned}
\mathbb{E}\left[C_l(\mathbb{G}^*)|\mathscr{A}_{\hat{\mu}}\right] &= \sum_{i=0}^{l} p_{l,M=i}(\hat{\mu})\, a_{l,M=i}(\hat{\mu})(1+o(1)) \\
&\sim \sum_{i=0,i\text{ even}}^{l}\left(\frac{2}{dn}\right)^l\left(\frac{e^{-\beta}}{(1+e^{-\beta})}\right)^{l-M}\left(\frac{1}{(1+e^{-\beta})}\right)^M\binom{l}{i}\frac{1}{l}n^l 2^{-l}(d(d-1))^l \\
&= \frac{(d-1)^l}{l}\sum_{i=0,i\text{ even}}^{l}\binom{l}{i}\left(\frac{e^{-\beta}}{1+e^{-\beta}}\right)^{l-i}\left(\frac{1}{1+e^{-\beta}}\right)^i \\
&= \frac{(d-1)^l}{2l}\left(\left(\frac{e^{-\beta}}{1+e^{-\beta}}+\frac{1}{1+e^{-\beta}}\right)^l+\left(\frac{e^{-\beta}}{1+e^{-\beta}}-\frac{1}{1+e^{-\beta}}\right)^l\right) \\
&= \frac{(d-1)^l}{2l}\left(1+\left(\frac{e^{-\beta}-1}{1+e^{-\beta}}\right)^l\right) =: \lambda^*.
\end{aligned}$$

In order to establish Proposition 2.2 we next need to calculate the higher moments of the number of short cycles in $\mathbb{G}^*$. To this end, we consider $\mathbb{E}\left[C_l(\mathbb{G}^*)^2|\mathscr{A}_{\hat{\mu}}\right]$ which can be interpreted as the expected number of ordered pairs of cycles in $\mathbb{G}$. We introduce two new random variables, namely X' and X''. X' denotes the number of ordered cycle pairs that are vertex-disjoint whereas X'' counts the ordered cycle pairs that have at least one vertex in common. This immediately brings us to

$$\mathbb{E}\left[C_l(\mathbb{G}^*)^2|\mathscr{A}_{\hat{\mu}}\right] = \mathbb{E}\left[X'|\mathscr{A}_{\hat{\mu}}\right] + \mathbb{E}\left[X''|\mathscr{A}_{\hat{\mu}}\right].$$

Starting with X' and adopting a corresponding definition of $p'_{l,M}$ and $a'_{l,M}$ - just now referring to two vertex-disjoint cycles - an analogue calculation to the one above yields

$$p'_{l,M}(\hat{\mu}) \sim \left(\frac{2}{dn}\right)^{2l}\left(\frac{e^{-\beta}}{(1+e^{-\beta})}\right)^{2l-2M}\left(\frac{1}{(1+e^{-\beta})}\right)^{2M}$$

and

$$(2l)^2 \cdot a'_{l,M}(\mu) = 4\left(\binom{l}{M}\right)^2(n\rho_+)_{2l_+}(n\rho_-)_{2l_-}(d(d-1))^{2l}.$$

Therefore, we arrive at

$$\mathbb{E}\left[X'|\mathscr{A}_{\hat{\mu}}\right] \sim (\lambda^*)^2.$$

All that remains to do is to show that $\mathbb{E}\left[X''|\mathscr{A}_{\hat{\mu}}\right]$ is asymptotically dominated by $\mathbb{E}\left[X'|\mathscr{A}_{\hat{\mu}}\right]$. More precisely, we show that $\mathbb{E}\left[X''|\mathscr{A}_{\hat{\mu}}\right] = O\left(n^{-1}\right)$ where we adopt an argument from [9] to our case. Whenever we have two cycles of length l that have k vertices in common, the number of shared vertices will exceed the number of shared edges by at least one. Put differently, the number of shared edges is at most $k-1$. As a result of this insight we have

$$a_{l,M}(\hat{\mu}') = \Theta\left(n^{2l-k}\right) \qquad \text{and} \qquad p_{l,M}(\hat{\mu}') = O\left(n^{-2l+k-1}\right)$$

for any $k < l$ and $\hat{\mu}' \in \mathscr{A}_{\hat{\mu}}$. Summing up over all $k \in [l-1]$ yields the desired statement

$$\mathbb{E}\left[X''|\mathscr{A}_{\hat{\mu}}\right] = O\left(n^{-1}\right).$$

This same argumentation can be extended to arbitrary higher moments $\mathbb{E}\left[C_l(\mathbb{G}^*)^j|\mathscr{A}_{\hat{\mu}}\right]$ with $j \in \mathbb{N}$. Thus, the method of moments provides the desired statement. □

Proof of Proposition 2.2. The Proposition results from combining Lemmas 5.1 and 5.2. □

6. The First Moment/ Proof of Proposition 2.4

In this section, we first focus on the so-called pairing model $\mathbf{G} = \mathbf{G}(n,d)$. In pairing model, each of the n initial nodes is represented by d clones. Then, a perfect matching for these dn clones is chosen uniformly at random. Finally, the clones are merged back into their initial vertex, such that each node in the original vertex set has degree d. By design, this setup allows for loops and double edges. If the graph does not contain either of them, we call the graph simple. Furthermore, we denote the event that a graph is simple by $\mathscr{S}$. The following result (which we will prove first) can be leveraged for showing Proposition 2.4.

Proposition 6.1. *Assume that* $0 < \beta < \beta_{KS}$ *and* $d \geq 3$. *Then we have*

$$\mathbb{E}\left[Z_{\mathbf{G},\beta}\right] = \exp\left(O\left(\frac{1}{n}\right)\right)\sqrt{\frac{1+e^{\beta}}{2+de^{\beta}-d}}\exp\left(n\left((1-d/2)\log(2)+d\log\left(1+e^{-\beta}\right)/2\right)\right)$$

6.1. **Getting started.** Recall the definitions of $\mu(G,\sigma)$ and $\rho(G,\sigma)$ from (5.1)–(5.4). As a starting point for our first moment calculations, consider the following result due to [3] which encodes the combinatorial structure of the first moment of the partition function. Let $\mathscr{M}_n = \cup_{\sigma\in\{\pm1\}^{V_n}}\mathscr{M}(\sigma)$ be the set of all conceivable distributions μ.

Lemma 6.2 (Lemmas 4.1 and 4.3 in [3])**.** *We have*

$$\mathbb{E}\left[Z_{\mathbf{G},\beta}\right] = \sum_{\mu\in\mathscr{M}_n}\binom{n}{\rho_+ n}\frac{(dn\mu_{++}-1)!!\,(dn\mu_{--}-1)!!\,(dn\mu_{+-})!}{(dn-1)!!}\binom{dn\rho_+}{dn\mu_{++}}\binom{dn\rho_-}{dn\mu_{--}}\cdot\exp\left(-\beta\frac{dn}{2}\left(\mu_{++}+\mu_{--}\right)\right).$$

6.2. **Reformulation of the first moment.** Recall Stirling's formula (5.8) and the identity for the double factorial from (5.6). The next Lemma yields a simplified expression for the first moment which is obtained by applying (5.8) and (5.6) to the factorials and binomial coefficients in Lemma 6.2. The proof follows [3], but now explicitly accounting for smaller-order terms to yield an error term of order $O(\exp(1/n))$.

Lemma 6.3. *We have*

$$\mathbb{E}\left[Z_{\mathbf{G},\beta}\right] = \sum_{\mu\in\mathscr{M}_n}\frac{\exp\left(nH(\rho)-\frac{dn}{2}\left(D_{KL}(\mu||\rho\otimes\rho)+\beta\left(\mu_{++}+\mu_{--}\right)\right)+O\left(\frac{1}{n}\right)\right)}{\pi n\sqrt{2\mu_{++}\mu_{--}\mu_{+-}d}}.$$

Proof. Starting with Lemma 6.2 and considering the fraction of factorials first, we find

$$\begin{aligned}
\frac{(dn\mu_{++}-1)!!\,(dn\mu_{--}-1)!!\,(dn\mu_{+-})!}{(dn-1)!!} &= \frac{(dn\mu_{++})!\,(dn\mu_{--})!\,(dn\mu_{+-})!\left(\frac{dn}{2}\right)!\cdot 2^{\frac{dn}{2}}}{(dn)!\left(\frac{dn\mu_{++}}{2}\right)!\left(\frac{dn\mu_{--}}{2}\right)!\cdot 2^{\frac{dn}{2}(\mu_{++}+\mu_{--})}}\\
&= \exp\left(O\left(\frac{1}{n}\right)\right)\sqrt{2\pi\frac{dn\mu_{++}dn\mu_{--}dn\mu_{+-}\frac{dn}{2}}{dn\frac{dn\mu_{++}}{2}\frac{dn\mu_{--}}{2}}}\cdot 2^{\frac{dn}{2}(1-\mu_{++}-\mu_{--})}\cdot\left(\frac{dn}{e}\right)^{dn\left(\mu_{+-}-\frac{1}{2}+\frac{\mu_{++}}{2}+\frac{\mu_{--}}{2}\right)}\\
&\quad\cdot 2^{dn\left(\frac{\mu_{++}}{2}+\frac{\mu_{--}}{2}-\frac{1}{2}\right)}\cdot\mu_{++}^{dn\left(\mu_{++}-\frac{\mu_{++}}{2}\right)}\cdot\mu_{--}^{dn\left(\mu_{--}-\frac{\mu_{--}}{2}\right)}\cdot\mu_{+-}^{dn\mu_{+-}}\\
&= \exp\left(O\left(\frac{1}{n}\right)\right)2\sqrt{\pi dn\mu_{+-}}\cdot\mu_{++}^{dn\frac{\mu_{++}}{2}}\cdot\mu_{--}^{dn\frac{\mu_{--}}{2}}\cdot\mu_{+-}^{dn\mu_{+-}}\\
&= 2\exp\left(O\left(\frac{1}{n}\right)+\frac{1}{2}\log\left(\pi dn\mu_{+-}\right)+\underbrace{dn\frac{\mu_{++}}{2}\log\left(\mu_{++}\right)+dn\frac{\mu_{--}}{2}\log\left(\mu_{--}\right)+dn\mu_{+-}\log\left(\mu_{+-}\right)}_{=-\frac{dn}{2}\mathrm{H}(\mu)}\right)\\
&= \exp\left(-\frac{dn}{2}\mathrm{H}\left(\mu\right)+\frac{1}{2}\log\left(n\right)+\frac{1}{2}\log\left(4\pi d\mu_{+-}\right)+O\left(\frac{1}{n}\right)\right) \qquad (6.1)
\end{aligned}$$

where we used (5.6) for the first equality and Stirling's formula (5.8) for the second equality. Similarly, we rearrange the second term of interest:

$$\binom{dn\rho_+}{dn\mu_{++}}\binom{dn\rho_-}{dn\mu_{--}} = \frac{(dn\rho_+)!\,(dn\rho_-)!}{(dn\mu_{++})!\,(dn\mu_{--})!\,(dn(\rho_+-\mu_{++}))!\,(dn(\rho_--\mu_{--}))!} = \frac{(dn\rho_+)!\,(dn\rho_-)!}{(dn\mu_{++})!\,(dn\mu_{--})!\,((dn\mu_{+-})!)^2}.$$

Another application of (5.8) yields

$$\begin{aligned}
&\binom{dn\rho_+}{dn\mu_{++}}\binom{dn\rho_-}{dn\mu_{--}}\\
&= \exp\left(O\left(\frac{1}{n}\right)\right)\frac{1}{2\pi dn}\sqrt{\frac{\rho_+\rho_-}{\mu_{++}\mu_{--}\mu_{+-}^2}}\left(\frac{dn}{e}\right)^{dn(\rho_++\rho_--\mu_{++}-\mu_{--}-2\mu_{+-})}\\
&\qquad\cdot\rho_+^{dn\rho_+}\cdot\rho_-^{dn\rho_-}\cdot\mu_{++}^{-dn\mu_{++}}\cdot\mu_{--}^{-dn\mu_{--}}\cdot\mu_{+-}^{-2dn\mu_{+-}}\\
&= \exp\left(dn\rho_+\log(\rho_+)+dn\rho_-\log(\rho_-)-dn\mu_{++}\log(\mu_{++})-dn\mu_{--}\log(\mu_{--})-2dn\mu_{+-}\log(\mu_{+-})\right)\\
&\qquad\cdot\exp\left(-\log(n)+\frac{1}{2}\log\left(\frac{\rho_+\rho_-}{\mu_{++}\mu_{--}\mu_{+-}^2 4\pi^2d^2}\right)+O\left(\frac{1}{n}\right)\right)\\
&= \exp\left(dn\left(\mathrm{H}(\mu)-\mathrm{H}(\rho)\right)-\log(n)+\frac{1}{2}\log\left(\frac{\rho_+\rho_-}{\mu_{++}\mu_{--}\mu_{+-}^2 4\pi^2d^2}\right)+O\left(\frac{1}{n}\right)\right) \qquad (6.2)
\end{aligned}$$

Combining (6.1) and (6.2) and denoting by we have

$$\begin{aligned}
&\frac{(dn\mu_{++}-1)!!\,(dn\mu_{--}-1)!!\,(dn\mu_{+-})!}{(dn-1)!!}\binom{dn\rho_+}{dn\mu_{++}}\binom{dn\rho_-}{dn\mu_{--}}\\
&= \exp\left(\frac{dn}{2}\left(\mathrm{H}(\mu)-\mathrm{H}(\rho\otimes\rho)\right)-\frac{1}{2}\log(n)+\frac{1}{2}\log\left(\frac{\rho_+\rho_-}{\mu_{++}\mu_{--}\mu_{+-}\pi d}\right)+O\left(\frac{1}{n}\right)\right)\\
&= \exp\left(-\frac{dn}{2}D_{\mathrm{KL}}(\mu||\rho\otimes\rho)-\frac{1}{2}\log(n)+\frac{1}{2}\log\left(\frac{\rho_+\rho_-}{\mu_{++}\mu_{--}\mu_{+-}\pi d}\right)+O\left(\frac{1}{n}\right)\right)
\end{aligned}$$

As an immediate consequence, the first moment from Lemma 6.2 can be expressed as

$$\mathbb{E}\left[Z_{\mathbf{G},\beta}\right]=\sum_{\mu\in\mathscr{M}_n}\binom{n}{\rho_+n}\sqrt{\frac{\rho_+\rho_-}{\mu_{++}\mu_{--}\mu_{+-}\pi dn}}\exp\left(-\frac{dn}{2}\left(D_{\mathrm{KL}}(\mu||\rho\otimes\rho)+\beta\left(\mu_{++}+\mu_{--}\right)\right)+O\left(\frac{1}{n}\right)\right)$$

where $\mathscr{M}_n$ is again the set of all conceivable distributions μ. A short auxiliary calculation using Stirling's formula (5.8) yields

$$\begin{aligned}
\binom{n}{\rho_+n}&=\frac{n!}{(\rho_+n)!\,(\rho_-n)!}=\frac{1}{\sqrt{2\pi n\rho_+\rho_-}}\left(\frac{n}{e}\right)^n\left(\frac{\rho_+n}{e}\right)^{-\rho_+n}\left(\frac{\rho_-n}{e}\right)^{-\rho_-n}\exp\left(O\left(\frac{1}{n}\right)\right)\\
&=\frac{\rho_+^{-\rho_+n}\rho_-^{-\rho_-n}}{\sqrt{2\pi n\rho_+\rho_-}}\exp\left(O\left(\frac{1}{n}\right)\right)=\exp\left(n\mathrm{H}(\rho)-\frac{1}{2}\log(n)-\frac{1}{2}\log\left(2\pi\rho_+\rho_-\right)+O\left(\frac{1}{n}\right)\right)
\end{aligned}$$

which enables us to state

$$\mathbb{E}\left[Z_{\mathbf{G},\beta}\right]=\sum_{\mu\in\mathscr{M}_n}\frac{\exp\left(n\mathrm{H}(\rho)-\frac{dn}{2}\left(D_{\mathrm{KL}}(\mu||\rho\otimes\rho)+\beta\left(\mu_{++}+\mu_{--}\right)\right)+O\left(\frac{1}{n}\right)\right)}{\pi n\sqrt{2\mu_{++}\mu_{--}\mu_{+-}d}}$$

as claimed. □

Revisiting the setup of our model, we see that all values of μ and ρ are completely determined by the choice of μ_{++} and ρ_+. Exploiting the fact that ρ is a probability distribution, we have

$$\rho_-=1-\rho_+.$$

A similar argument can be made for μ. Since edges are by definition undirected in our setup, we have $\mu_{+-}=\mu_{-+}$. Keeping in mind that μ is also a probability measure, the missing weights of μ can be deduced from μ_{++} and ρ_+ by the equations

$$\begin{aligned}
\mu_{+-}&=\mu_{-+}=\rho_+-\mu_{++}\\
\mu_{--}&=1-2\left(\rho_+-\mu_{++}\right)-\mu_{++}=1+\mu_{++}-2\rho_+.
\end{aligned}$$

Substituting the above into Lemma 6.3 and some simplifications give us

$$\mathbb{E}\left[Z_{\mathbf{G},\beta}\right]=\exp\left(O\left(\frac{1}{n}\right)\right)\cdot\sum_{(\rho_+,\mu_{++})\in\mathscr{Q}}\frac{1}{\pi n\sqrt{2\mu_{++}\mu_{--}\mu_{+-}d}}\exp\left(n\psi\left(\mu_{++},\rho_+\right)\right). \qquad (6.3)$$

where

$$\psi\left(\mu_{++},\rho_+\right):=\mathrm{H}(\rho)-\frac{d}{2}\left(D_{\mathrm{KL}}(\mu||\rho\otimes\rho)+\beta\left(1+2\mu_{++}-2\rho_+\right)\right)$$

and $\mathcal{Q}$ is the set of all conceivable pairs (ρ_+, μ_{++}). The KL-divergence can also be expressed just in terms of μ_{++} and ρ_+, as the following calculation shows.

$$\begin{aligned}
D_{\mathrm{KL}}(\mu||\rho\otimes\rho) =& \mu_{++}\log\left(\frac{\mu_{++}}{\rho_+^2}\right) + \mu_{--}\log\left(\frac{\mu_{--}}{\rho_-^2}\right) + 2\mu_{+-}\log\left(\frac{\mu_{+-}}{\rho_+\rho_-}\right)\\
=& \mu_{++}\log\left(\frac{\mu_{++}}{\rho_+^2}\right) + (1+\mu_{++}-2\rho_+)\log\left(\frac{1+\mu_{++}-2\rho_+}{(1-\rho_+)^2}\right) + 2(\rho_+-\mu_{++})\log\left(\frac{\rho_+-\mu_{++}}{\rho_+(1-\rho_+)}\right)\\
=& \mu_{++}\log(\mu_{++}) + (1+\mu_{++}-2\rho_+)\log(1+\mu_{++}-2\rho_+)\\
& -2(1-\rho_+)\log(1-\rho_+) + 2(\rho_+-\mu_{++})\log(\rho_+-\mu_{++}) - 2\rho_+\log(\rho_+).
\end{aligned}$$

Having effectively reduced the number of involved variables, we now can move on to apply the Laplace method as stated in Theorem 2.3 in [5].

6.3. **Application of the Laplace method to the first moment.** Before we can apply the Laplace method to the expression for the first moment in (6.3), we need some preliminary work. To be precise, we need to determine the unique maximum $(\hat{\mu}_{++}, \hat{\rho}_+)$ of ψ on the set $\mathcal{Q}$ and evaluate the Hessian at this point. To this end, consider

$$\hat{\rho}_+ = \hat{\rho}_- = \frac{1}{2}, \tag{6.4}$$

i.e. balanced number of vertices with positive and negative spins. Moreover, let

$$\hat{\mu}_{++} = \hat{\mu}_{--} = \frac{e^{-\beta}}{2(1+e^{-\beta})} \qquad \text{and} \qquad \hat{\mu}_{+-} = \hat{\mu}_{-+} = \frac{1}{2(1+e^{-\beta})}. \tag{6.5}$$

We will see in due course in Lemma 6.5 that $(\hat{\mu}_{++}, \hat{\rho}_+)$ indeed constitutes the unique maximum of ψ. Let us first calculate partial derivatives and establish the Hessian of ψ at $(\hat{\mu}_{++}, \hat{\rho}_+)$.

Lemma 6.4 (Hessian for the first moment)**.** *We have*

$$\det\left(-Hes_\psi\left(\frac{1}{2}, \frac{e^{-\beta}}{2(1+e^{-\beta})}\right)\right) = 4d\left(1+e^{-\beta}\right)^2 e^{\beta}\left(2+d\left(e^{\beta}-1\right)\right).$$

Proof. Let us get started simple and state the partial derivatives of the Kullback-Leibler divergence from (6.3) with respect to μ_{++} and ρ_+.

$$\begin{aligned}
\frac{\partial D_{\mathrm{KL}}(\mu||\rho\otimes\rho)}{\partial\mu_{++}} &= \log(\mu_{++}) + \log(1+\mu_{++}-2\rho_+) - 2\log(\rho_+-\mu_{++})\\
\frac{\partial^2 D_{\mathrm{KL}}(\mu||\rho\otimes\rho)}{\partial\mu_{++}^2} &= \frac{1}{\mu_{++}} + \frac{1}{1+\mu_{++}-2\rho_+} + \frac{2}{\rho_+-\mu_{++}}\\
\frac{\partial D_{\mathrm{KL}}(\mu||\rho\otimes\rho)}{\partial\rho_+} &= -2\log(1+\mu_{++}-2\rho_+) + 2\log(1-\rho_+) + 2\log(\rho_+-\mu_{++}) - 2\log(\rho_+)\\
\frac{\partial^2 D_{\mathrm{KL}}(\mu||\rho\otimes\rho)}{\partial\rho_+^2} &= \frac{4}{1+\mu_{++}-2\rho_+} - \frac{2}{1-\rho_+} + \frac{2}{\rho_+-\mu_{++}} - \frac{2}{\rho_+}.
\end{aligned}$$

Furthermore, for the entropy we recall

$$\frac{\partial \mathrm{H}(\rho_+)}{\partial\rho_+} = \log(1-\rho_+) - \log(\rho_+) \qquad \text{and} \qquad \frac{\partial^2 \mathrm{H}(\rho_+)}{\partial\rho_+^2} = -\frac{1}{1-\rho_+} - \frac{1}{\rho_+}.$$

Keeping these auxiliary calculations in mind, the first derivatives of ψ turn out to be

$$\begin{aligned}
\frac{\partial\psi(\mu_{++},\rho_+)}{\partial\mu_{++}} &= -\frac{d}{2}\left(\log(\mu_{++}) + \log(1+\mu_{++}-2\rho_+) - 2\log(\rho_+-\mu_{++}) + 2\beta\right)\\
\frac{\partial\psi(\mu_{++},\rho_+)}{\partial\rho_+} &= \log(1-\rho_+) - \log(\rho_+) - d\left(-\log(1+\mu_{++}-2\rho_+) + \log(1-\rho_+) + \log(\rho_+-\mu_{++}) - \log(\rho_+) - \beta\right)
\end{aligned}$$

while the second derivatives of ψ are given by

$$\frac{\partial^2\psi(\mu_{++},\rho_+)}{\partial\mu_{++}^2} = -\frac{d}{2}\left(\frac{1}{\mu_{++}} + \frac{1}{1+\mu_{++}-2\rho_+} + \frac{2}{\rho_+-\mu_{++}}\right)$$
$$\frac{\partial^2\psi(\mu_{++},\rho_+)}{\partial\mu_{++}\partial\rho_+} = \frac{\partial^2\psi(\mu_{++},\rho_+)}{\partial\rho_+\partial\mu_{++}} = -\frac{d}{2}\left(-\frac{2}{1+\mu_{++}-2\rho_+} - 2\frac{1}{\rho_+-\mu_{++}}\right) = \frac{d(1-\rho_+)}{(1+\mu_{++}-2\rho_+)(\rho_+-\mu_{++})}$$
$$\frac{\partial^2\psi(\mu_{++},\rho_+)}{\partial\rho_+^2} = -\frac{1}{1-\rho_+} - \frac{1}{\rho_+} - d\left(\frac{2}{1+\mu_{++}-2\rho_+} - \frac{1}{1-\rho_+} + \frac{1}{\rho_+-\mu_{++}} - \frac{1}{\rho_+}\right).$$

With the above at hand, the entries of the Hessian turn out to be

$$\begin{aligned}\frac{\partial^2\psi}{\partial\mu_{++}^2}\left(\frac{1}{2},\frac{e^{-\beta}}{2(1+e^{-\beta})}\right) &= -\frac{d}{2}\left(\frac{1}{\frac{e^{-\beta}}{2(1+e^{-\beta})}} + \frac{1}{1+\frac{e^{-\beta}}{2(1+e^{-\beta})}-1} + \frac{2}{\frac{1}{2}-\frac{e^{-\beta}}{2(1+e^{-\beta})}}\right)\\ &= -2d\left(\frac{1+e^{-\beta}+e^{-\beta}+e^{-2\beta}}{e^{-\beta}}\right) = -2d\left(1+e^{-\beta}\right)^2 e^{\beta} < 0 \qquad (6.6)\end{aligned}$$

and

$$\begin{aligned}\frac{\partial^2\psi}{\partial\mu_{++}\partial\rho_+}\left(\frac{1}{2},\frac{e^{-\beta}}{2(1+e^{-\beta})}\right) &= \frac{\partial^2\psi}{\partial\rho_+\partial\mu_{++}}\left(\frac{1}{2},\frac{e^{-\beta}}{2(1+e^{-\beta})}\right) = \frac{\frac{d}{2}}{\frac{e^{-\beta}}{2(1+e^{-\beta})}\left(\frac{1}{2}-\frac{e^{-\beta}}{2(1+e^{-\beta})}\right)}\\ &= \frac{2d}{\frac{e^{-\beta}}{1+e^{-\beta}}\cdot\frac{1}{1+e^{-\beta}}} = 2d\left(1+e^{-\beta}\right)^2 e^{\beta}\end{aligned}$$

and

$$\begin{aligned}\frac{\partial^2\psi}{\partial\rho_+^2}\left(\frac{1}{2},\frac{e^{-\beta}}{2(1+e^{-\beta})}\right) &= -2-2-d\left(\frac{2}{\frac{e^{-\beta}}{2(1+e^{-\beta})}} - 2 + \frac{1}{\frac{1}{2(1+e^{-\beta})}} - 2\right)\\ &= -4-d\left(\frac{4(1+e^{-\beta})}{e^{-\beta}} - 2 + 2e^{-\beta}\right) = -4-2d\left(1+e^{-\beta}+2e^{\beta}\right).\end{aligned}$$

Combining the above, the determinant of the Hessian at $(\hat{\mu}_{++},\hat{\rho}_+)$ is given by

$$\begin{aligned}\det\left(-\mathrm{Hes}_{\psi}\left(\frac{1}{2},\frac{e^{-\beta}}{2(1+e^{-\beta})}\right)\right) &= 8d\left(1+e^{-\beta}\right)^2 e^{\beta} + 4d^2\left(1+e^{-\beta}+2e^{\beta}\right)\left(1+e^{-\beta}\right)^2 e^{\beta} - 4d^2\left(1+e^{-\beta}\right)^4 e^{2\beta}\\ &= 8d\left(1+e^{-\beta}\right)^2 e^{\beta} + 4d^2\left(1+e^{-\beta}\right)^2 e^{\beta}\underbrace{\left(1+e^{-\beta}+2e^{\beta}-e^{\beta}-2-e^{-\beta}\right)}_{=e^{\beta}-1}\\ &= 4d\left(1+e^{-\beta}\right)^2 e^{\beta}\left(2+d\left(e^{\beta}-1\right)\right) > 0. \qquad (6.7)\end{aligned}$$

closing the proof of the lemma. □

With the partial derivatives in place, we can proceed to establish that the unique maximum of ψ is indeed at $(\hat{\mu}_{++},\hat{\rho}_+)$.

Lemma 6.5 (Maximum for the First Moment Calculation)**.** *With the definitions of $\hat{\rho}_+$ and $\hat{\mu}_{++}$ from* (6.4) *and* (6.5) *we have*

$$\arg\max_{(\mu_{++},\rho_+)\in\mathcal{Q}} \psi(\mu_{++},\rho_+) = (\hat{\mu}_{++},\hat{\rho}_+)$$

Proof. As a starting point, we set the first derivatives equal to zero, resulting in

$$\frac{\partial\psi(\hat{\mu}_{++},\hat{\rho}_+)}{\partial\hat{\mu}_{++}} = -\frac{d}{2}\left(\log(\hat{\mu}_{++}) + \log(1+\hat{\mu}_{++}-2\hat{\rho}_+) - 2\log(\hat{\rho}_+-\hat{\mu}_{++}) + 2\beta\right) = 0$$

which is equivalent to

$$\begin{aligned}0 &= \log(\hat{\mu}_{++}) + \log(1+\hat{\mu}_{++}-2\hat{\rho}_+) - 2\log(\hat{\rho}_+-\hat{\mu}_{++}) + 2\beta\\ \Leftrightarrow 1 &= \frac{\hat{\mu}_{++}(1+\hat{\mu}_{++}-2\hat{\rho}_+)}{(\hat{\rho}_+-\hat{\mu}_{++})^2}e^{2\beta}\\ \Leftrightarrow 0 &= \hat{\mu}_{++}^2\left(1-e^{-2\beta}\right) + \hat{\mu}_{++}\left(1-2\hat{\rho}_+ + e^{-2\beta}2\hat{\rho}_+\right) - e^{-2\beta}\hat{\rho}_+^2.\end{aligned}$$

Then, the quadratic formula yields two candidates for the solution, namely

$$\begin{aligned}\hat{\mu}_{++,1/2} &= \frac{-1+2\hat{\rho}_+ - e^{-2\beta}2\hat{\rho}_+ \pm \sqrt{\left(1-2\hat{\rho}_+\left(1-e^{-2\beta}\right)\right)^2 + 4\left(1-e^{-2\beta}\right)e^{-2\beta}\hat{\rho}_+^2}}{2\left(1-e^{-2\beta}\right)} \\ &= \hat{\rho}_+ - \frac{1 \mp \sqrt{\left(1-2\hat{\rho}_+\right)^2 + 4\hat{\rho}_+ e^{-2\beta}\left(1-\hat{\rho}_+\right)}}{2\left(1-e^{-2\beta}\right)}.\end{aligned}$$

This result immediately poses the question of possible extrema. First we note that

$$\left(1-2\hat{\rho}_+\right)^2 + 4\hat{\rho}_+ e^{-2\beta}\left(1-\hat{\rho}_+\right) > 0$$

since both summands are positive. This in turn enables us to rule out $\hat{\mu}_{++,2} = \hat{\rho}_+ - \frac{1+\sqrt{\left(1-2\hat{\rho}_+\right)^2+4\hat{\rho}_+e^{-2\beta}\left(1-\hat{\rho}_+\right)}}{2\left(1-e^{-2\beta}\right)}$ as a solution since that would imply

$$\hat{\mu}_{-+} = \hat{\mu}_{+-} = \hat{\rho}_+ - \hat{\mu}_{++,2} = \frac{1+\sqrt{\left(1-2\hat{\rho}_+\right)^2 + 4\hat{\rho}_+ e^{-2\beta}\left(1-\hat{\rho}_+\right)}}{2\left(1-e^{-2\beta}\right)} > \frac{1}{2\left(1-e^{-2\beta}\right)} > \frac{1}{2}$$

which contradicts the fact that $\hat{\mu}$ is a probability measure. As a consequence, the only solution that is consistent with our model assumptions is

$$\hat{\mu}_{++} = \hat{\rho}_+ - \frac{1-\sqrt{\left(1-2\hat{\rho}_+\right)^2 + 4\hat{\rho}_+ e^{-2\beta}\left(1-\hat{\rho}_+\right)}}{2\left(1-e^{-2\beta}\right)} = \hat{\rho}_+ - \frac{1-\eta}{2\left(1-e^{-2\beta}\right)} \tag{6.8}$$

where

$$\eta := \sqrt{\left(1-2\hat{\rho}_+\right)^2 + 4\hat{\rho}_+ e^{-2\beta}\left(1-\hat{\rho}_+\right)} = \sqrt{1-4\hat{\rho}_+\left(1-e^{-2\beta}\right) + 4\hat{\rho}_+^2\left(1-e^{-2\beta}\right)}$$

In the next step, we plug (6.8) into the first derivative of ψ with respect to ρ_+

$$\begin{aligned}\frac{\partial\psi\left(\hat{\mu}_{++},\hat{\rho}_+\right)}{\partial\hat{\rho}_+} &= \log\left(1-\hat{\rho}_+\right) - \log\left(\hat{\rho}_+\right) \\ &\quad - d\left(-\log\left(1+\hat{\mu}_{++}-2\hat{\rho}_+\right) + \log\left(1-\hat{\rho}_+\right) + \log\left(\hat{\rho}_+-\hat{\mu}_{++}\right) - \log\left(\hat{\rho}_+\right) - \beta\right) = 0\end{aligned}$$

which yields

$$\begin{aligned}\xi\left(\hat{\rho}_+\right) &:= \log\left(1-\hat{\rho}_+\right) - \log\left(\hat{\rho}_+\right) \\ &\quad - d\left(-\log\left(1-\frac{1-\eta}{2\left(1-e^{-2\beta}\right)}-\hat{\rho}_+\right) + \log\left(1-\hat{\rho}_+\right) + \log\left(\frac{1-\eta}{2\left(1-e^{-2\beta}\right)}\right) - \log\left(\hat{\rho}_+\right) - \beta\right) \\ &= \log\left(1-\hat{\rho}_+\right) - \log\left(\hat{\rho}_+\right) - d\left(-\log\left(\frac{\left(1-\hat{\rho}_+\right)2\left(1-e^{-2\beta}\right)}{1-\eta}-1\right) + \log\left(1-\hat{\rho}_+\right) - \log\left(\hat{\rho}_+\right) - \beta\right) \\ &= 0\end{aligned}$$

Next, let us take a look at the derivative of ξ with respect to $\hat{\rho}_+$:

$$\begin{aligned}\frac{\partial\xi\left(\hat{\rho}_+\right)}{\partial\hat{\rho}_+} &= (1-d)\left(-\frac{1}{1-\hat{\rho}_+} - \frac{1}{\hat{\rho}_+}\right) + d\left(\frac{\frac{-(1-\eta)2\left(1-e^{-2\beta}\right)+\frac{2}{\eta}\left(2\hat{\rho}_+-1\right)\left(1-e^{-2\beta}\right)\left(1-\hat{\rho}_+\right)2\left(1-e^{-2\beta}\right)}{(1-\eta)^2}}{\frac{\left(1-\hat{\rho}_+\right)2\left(1-e^{-2\beta}\right)}{1-\eta}-1}\right) \\ &= \frac{d-1}{\left(1-\hat{\rho}_+\right)\hat{\rho}_+} + d\left(\frac{-\left(1-\eta\right)2\left(1-e^{-2\beta}\right)+\frac{4}{\eta}\left(2\hat{\rho}_+-1\right)\left(1-\hat{\rho}_+\right)\left(1-e^{-2\beta}\right)^2}{\left(1-\hat{\rho}_+\right)2\left(1-e^{-2\beta}\right)\left(1-\eta\right)-\left(1-\eta\right)^2}\right)\end{aligned}$$

where we made use of the simple fact

$$\frac{\partial\eta}{\partial\hat{\rho}_+} = \frac{1}{2\eta}\left(8\hat{\rho}_+-4\right)\left(1-e^{-2\beta}\right) = \frac{2}{\eta}\left(2\hat{\rho}_+-1\right)\left(1-e^{-2\beta}\right).$$

To simplify the first derivative, we focus on

$$\frac{-(1-\eta)\,2\left(1-e^{-2\beta}\right)+\frac{4}{\eta}\left(2\hat{\rho}_+-1\right)\left(1-\hat{\rho}_+\right)\left(1-e^{-2\beta}\right)^2}{\left(1-\hat{\rho}_+\right)2\left(1-e^{-2\beta}\right)(1-\eta)-(1-\eta)^2}=\frac{2\left(1-e^{-2\beta}\right)\cdot\left(-(1-\eta)+\frac{1-\eta^2}{\eta}+\frac{2}{\eta}\left(\hat{\rho}_+-1\right)\left(1-e^{-2\beta}\right)\right)}{(1-\eta)\left(1-2\hat{\rho}_+-2e^{-2\beta}+2e^{-2\beta}\hat{\rho}_++\eta\right)}$$

$$=\frac{2\left(1-e^{-2\beta}\right)}{(1-\eta)\eta}\cdot\frac{-\eta+2\hat{\rho}_+-1-2\hat{\rho}_+e^{-2\beta}+2e^{-2\beta}}{1-2\hat{\rho}_+-2e^{-2\beta}+2e^{-2\beta}\hat{\rho}_++\eta}=-\frac{2\left(1-e^{-2\beta}\right)}{(1-\eta)\eta}.$$

As a consequence, the derivative can be simplified to

$$\frac{\partial\xi\left(\hat{\rho}_+\right)}{\partial\hat{\rho}_+}=\frac{d-1}{\left(1-\rho_+\right)\rho_+}-2d\left(\frac{1-e^{-2\beta}}{(1-\eta)\eta}\right).$$

Before proceedings, we point out that

$$\left(1-\hat{\rho}_+\right)\hat{\rho}_+=\frac{1-\eta^2}{\left(1-e^{-2\beta}\right)4}$$

which brings us to

$$\begin{aligned}\frac{\partial\xi\left(\hat{\rho}_+\right)}{\partial\hat{\rho}_+}&=\frac{d-1}{\left(1-\hat{\rho}_+\right)\hat{\rho}_+}-2d\left(\frac{1-e^{-2\beta}}{(1-\eta)\eta}\right)=\frac{1-e^{-2\beta}}{1-\eta}\left(\frac{4d-4}{1+\eta}-\frac{2d}{\eta}\right)\\&=\frac{1-e^{-2\beta}}{\left(1-\eta^2\right)\eta}\left(4d\eta-4\eta-2d-2d\eta\right)=\frac{1-e^{-2\beta}}{\left(1-\eta^2\right)\eta}\left(2d\left(\eta-1\right)-4\eta-2d\right)<0\end{aligned}$$

where we implicitly assumed that $\hat{\rho}_+$ is conceivable which especially means that $0<\eta<1$ holds. $\frac{\partial\xi\left(\hat{\rho}_+\right)}{\partial\hat{\rho}_+}<0$ implies that if we can locate any root of $\xi\left(\hat{\rho}_+\right)$ it is automatically the unique one. Recalling our definition of ρ_+ and μ_{++} from (6.4) and (6.5), we conjecture that this root is located at $\hat{\rho}_+=\frac{1}{2}$. A short calculation indeed verifies

$$\begin{aligned}\xi\left(\frac{1}{2}\right)&=-d\left(-\log\left(\frac{\left(1-\frac{1}{2}\right)2\left(1-e^{-2\beta}\right)}{1-\eta}-1\right)+\log\left(1-\frac{1}{2}\right)-\log\left(\frac{1}{2}\right)-\beta\right)\\&=d\left(\log\left(\frac{1-e^{-2\beta}}{1-e^{-\beta}}-1\right)+\beta\right)=d\left(\log\left(e^{-\beta}\right)+\beta\right)=0\end{aligned}$$

where we used

$$\eta=\sqrt{\left(1-2\hat{\rho}_+\right)^2+4\hat{\rho}_+e^{-2\beta}\left(1-\hat{\rho}_+\right)}=\sqrt{\left(1-2\cdot\frac{1}{2}\right)^2+4\cdot\frac{1}{2}\cdot e^{-2\beta}\left(1-\frac{1}{2}\right)}=e^{-\beta}.$$

This immediately allows us to calculate the optimal $\hat{\mu}_{++}$ by plugging $\hat{\rho}_+=\frac{1}{2}$ into equation (6.8)

$$\hat{\mu}_{++}=\frac{1}{2}-\frac{1-e^{-\beta}}{2\left(1-e^{-2\beta}\right)}=\frac{1}{2}-\frac{1}{2\left(1+e^{-\beta}\right)}=\frac{e^{-\beta}}{2\left(1+e^{-\beta}\right)}.$$

The above establishes that $\left(\hat{\mu}_{++},\hat{\rho}_+\right)$ is the (only) extremum of ψ. Let us next show that it is indeed the global maximum (and not a minimum or stationary point). From the calculation of the Hessian (Lemma 6.4), we saw that the first leading principal minor is negative (see inequality (6.6)) and the second one is positive (see inequality (6.7)). Thus, ψ is strictly concave at $\left(\hat{\mu}_{++},\hat{\rho}_+\right)$ which makes it a local maximum. Due to the uniqueness of the extremum, $\left(\hat{\mu}_{++},\hat{\rho}_+\right)$ thereby also is the unique maximum. □

As an application of Lemma 6.5, we obtain the following corollary.

Corollary 6.6. *We have*

$$\max_{(\mu_{++},\rho_+)\in\mathcal{Q}}\psi\left(\mu_{++},\rho_+\right)=\psi\left(\hat{\mu}_{++},\hat{\rho}_+\right)=\left(1-\frac{d}{2}\right)\log(2)+\frac{d}{2}\log\left(1+e^{-\beta}\right)$$

Proof. We evaluate ψ at the optimal point $\left(\hat{\mu}_{++},\hat{\rho}_+\right)$. Starting with the entropy we have

$$\mathrm{H}\left(\hat{\rho}\right)=-\frac{1}{2}\log\left(\frac{1}{2}\right)-\frac{1}{2}\log\left(\frac{1}{2}\right)=\log(2).\tag{6.9}$$

Continuing with the Kullback-Leibler divergence, we find

$$D_{\mathrm{KL}}\left(\hat{\mu}||\hat{\rho}\otimes\hat{\rho}\right) = \frac{e^{-\beta}}{2\left(1+e^{-\beta}\right)}\log\left(\frac{e^{-\beta}}{2\left(1+e^{-\beta}\right)}\right) + \frac{e^{-\beta}}{2\left(1+e^{-\beta}\right)}\log\left(\frac{e^{-\beta}}{2\left(1+e^{-\beta}\right)}\right) \tag{6.10}$$

$$+2\left(\frac{1}{2}-\frac{e^{-\beta}}{2\left(1+e^{-\beta}\right)}\right)\log\left(\frac{1}{2}-\frac{e^{-\beta}}{2\left(1+e^{-\beta}\right)}\right)-2\log\left(\frac{1}{2}\right) \tag{6.11}$$

$$= \frac{e^{-\beta}}{\left(1+e^{-\beta}\right)}\log\left(\frac{e^{-\beta}}{2\left(1+e^{-\beta}\right)}\right)+2\log(2)+\frac{1}{\left(1+e^{-\beta}\right)}\log\left(\frac{1}{2\left(1+e^{-\beta}\right)}\right) \tag{6.12}$$

$$= \log(2)-\frac{\beta e^{-\beta}}{\left(1+e^{-\beta}\right)}-\log\left(1+e^{-\beta}\right). \tag{6.13}$$

Combining (6.9) and (6.10) we arrive at

$$\begin{aligned}\psi\left(\hat{\mu}_{++},\hat{\rho}_{+}\right) &= \mathrm{H}\left(\hat{\rho}\right)-\frac{d}{2}\left(D_{\mathrm{KL}}(\hat{\mu}||\hat{\rho}\otimes\hat{\rho})+\beta\left(1+2\hat{\mu}_{++}-2\hat{\rho}_{+}\right)\right)\\ &= \log(2)-\frac{d}{2}\left(\log(2)-\frac{\beta e^{-\beta}}{\left(1+e^{-\beta}\right)}-\log\left(1+e^{-\beta}\right)+2\beta\frac{e^{-\beta}}{2\left(1+e^{-\beta}\right)}\right)\\ &= \left(1-\frac{d}{2}\right)\log(2)+\frac{d}{2}\log\left(1+e^{-\beta}\right).\end{aligned}$$

as claimed. □

Proof of Proposition 6.1. With Lemmas 6.4 and 6.5 and Corollary 6.6 in place, all that is left for the application of Laplace's method from [5] is the determination of the appropriate lattice. Put differently, we are interested in the respective matrix A_{first} which consist of the basis elements of the lattice. For the first moment, the matrix can be constructed in a rather simple way. Since ρ_+ is of the form

$$\rho_{+}=\frac{1}{n}\sum_{v\in V}\mathbf{1}\{\sigma(v)=+1\}$$

the first entry $A_{\text{first},1,1}$ immediately turns out to be equal to one. Similarly, keeping in mind

$$\mu_{++}=\frac{2}{dn}\sum_{(u,v)\in E}\mathbf{1}\{\sigma(v)=\sigma(u)=+1\}$$

yields $A_{\text{first},2,2}=\frac{2}{d}$. Having constructed the matrix A_{first}, we are left to compute its determinant

$$\det\left(A_{\text{first}}\right)=\det\begin{pmatrix}1 & 0\\ 0 & \frac{2}{d}\end{pmatrix}=\frac{2}{d}.$$

Now, we can bring together all the findings of this section to obtain a precise statement of the first moment up to an error term of order $O(\exp(1/n))$. Applying the Laplace method, i.e. Theorem 2.3 in [5] to expression (6.3) yields

$$\begin{aligned}\mathbb{E}\left[Z_{\mathbf{G},\beta}\right] &= \exp\left(O\left(\frac{1}{n}\right)\right)\cdot\sum_{(\rho_+,\mu_{++})\in\mathcal{Q}}\frac{1}{\pi n\sqrt{2\mu_{++}\mu_{--}\mu_{+-}d}}\exp\left(n\psi\left(\mu_{++},\rho_{+}\right)\right)\\ &= \exp\left(O\left(\frac{1}{n}\right)\right)\frac{2\pi n\exp\left(n\psi\left(\hat{\mu}_{++},\hat{\rho}_{+}\right)\right)}{\det\left(A_{\text{first}}\right)\sqrt{\det\left(-\mathrm{Hes}_{\psi}\left(\hat{\mu}_{++},\hat{\rho}_{+}\right)\right)}\pi n\sqrt{2\hat{\mu}_{++}\hat{\mu}_{--}\hat{\mu}_{+-}d}}\\ &= \exp\left(O\left(\frac{1}{n}\right)\right)\frac{2\exp\left(n\left(\left(1-\frac{d}{2}\right)\log(2)+\frac{d}{2}\log\left(1+e^{-\beta}\right)\right)\right)}{\frac{2}{d}\sqrt{4d\left(1+e^{-\beta}\right)^{2}e^{\beta}\left(2+d\left(e^{\beta}-1\right)\right)}\sqrt{2\frac{e^{-2\beta}}{8\left(1+e^{-\beta}\right)^{3}}d}}\\ &= \exp\left(O\left(\frac{1}{n}\right)\right)\sqrt{\frac{1+e^{\beta}}{2+de^{\beta}-d}}\exp\left(n\left(\left(1-\frac{d}{2}\right)\log(2)+\frac{d}{2}\log\left(1+e^{-\beta}\right)\right)\right).\end{aligned}$$

as claimed. □

6.4. **The simple d-regular case.** Having established the first moment in the pairing model $\mathbf{G}$, we next adapt the result to the d-regular model $\mathbb{G}$ of interest. As we will see, a pairing variant $\mathbf{G}_1^*$ of the planted model will be a useful tool to do so. The pairing variant $\mathbf{G}_1^*$ is defined as follows. First, draw a spin assignment $\boldsymbol{\sigma}^* \in \{\pm 1\}^n$ uniformly at random. Then, draw a graph $\mathbf{G}_1^*$ according to the probability distribution

$$\mathbb{P}\left[\mathbf{G}_1^* = G|\boldsymbol{\sigma}^*\right] \propto \exp\left(-\beta \mathscr{H}_G\left(\boldsymbol{\sigma}^*\right)\right).$$

where G might contain self-loops and double-edges. In the following, we will call a graph G *simple* if it does not feature any such self-loops or double-edges. With this definition, we are able to prove Proposition 2.4.

Proof of Proposition 2.4. To get started, we note the asymptotic equality

$$\mathbb{E}\left[Z_{\mathbb{G},\beta}\right] \sim \frac{\mathbb{P}\left[\mathbf{G}_1^* \text{ is simple}\right]}{\mathbb{P}\left[\mathbf{G} \text{ is simple}\right]} \mathbb{E}\left[Z_{\mathbf{G},\beta}\right]. \tag{6.14}$$

Fortunately, both $\mathbb{P}\left[\mathbf{G}_1^* \text{ is simple}\right]$ and $\mathbb{P}\left[\mathbf{G} \text{ is simple}\right]$ can be readily found in the literature.

Fact 6.7 (Corollary 9.7 in [9])**.** *For $d \geq 3$, we have*

$$\mathbb{P}\left[\mathbf{G} \text{ is simple}\right] \sim \exp\left(-\frac{d-1}{2} - \frac{(d-1)^2}{4}\right).$$

Lemma 6.8 (Lemma 4.6 in [3])**.** *For $d \geq 0$ and $\beta > 0$ we have*

$$\mathbb{P}\left[\mathbf{G}_1^* \text{ is simple}\right] \sim \exp\left(-(d-1)\frac{1}{1+e^{\beta}} - (d-1)^2 \frac{1+e^{2\beta}}{2\left(1+e^{\beta}\right)^2}\right).$$

In combination with Proposition 6.1 and equation (6.14), Fact 6.7 and Lemma 6.8 yield the desired result. □

7. The Second Moment / Proof of Proposition 2.6

Similar to the first moment, we will first establish the following result for the paring model $\mathbf{G}$.

Proposition 7.1. *For $0 < \beta < \beta_{KS}$ and $d \geq 3$ we have*

$$\mathbb{E}\left[Z_{\mathbf{G},\beta}\mathbf{1}\left\{\mathscr{O}\right\}\right] = \exp\left(O\left(\frac{1}{n}\right)\right) \frac{\left(1+e^{\beta}\right)^2 \exp\left(n\left((2-d)\log(2) + d\log\left(1+e^{-\beta}\right)\right)\right)}{\left(de^{\beta}-d+2\right)\sqrt{2e^{2\beta}+2de^{\beta}-de^{2\beta}-d+2}}.$$

Once we have done so, we bridge the gap between $\mathbf{G}$ and $\mathbb{G}$.

7.1. **Getting started.** For the second moment calculation, we introduce a set of variables that is similar in meaning to the ones employed in the previous sections. Yet, the definitions become more complicated since for the second moment each node v in some graph G is assigned two spins σ_v and τ_v which can be either positive or negative. As before, we aim to measure the fractions of edges that connect two vertices with certain spin configurations. Since each node is equipped with two spins, there are 16 possible spin configurations for two connected vertices. Usually, we will denote such a configuration as $(\sigma_1, \tau_1, \sigma_2, \tau_2) \in \{\pm 1\}^4$ where σ_1 and τ_1 denote the spins assigned to the first node. Accordingly, σ_2 and τ_2 are the spins of the second node. With this notation of spin assignments in mind, we define

$$\mu_{r,s,t,u} := \frac{2}{dn} \sum_{(u,v)\in E} \mathbf{1}\left\{\sigma(u) = r, \sigma(v) = s, \tau(u) = t, \tau(v) = u\right\} \qquad \{r,s,t,u \in \pm 1\}.$$

with the shorthand notation $\mu_{++++} = \mu_{+1,+1,+1,+1}$ and so forth. Our choices of μ are constrained by the following relationship.

$$\mu_{(\sigma_1,\tau_1,\sigma_2,\tau_2)} = \mu_{(\sigma_2,\tau_2,\sigma_1,\tau_1)} \qquad \forall\, (\sigma_1,\tau_1,\sigma_2,\tau_2) \in \{\pm 1\}^4. \tag{7.1}$$

Note that $\mu_{++++}, \mu_{+-+-}, \mu_{-+-+}$, and μ_{----} get a special meaning: these four configurations satisfy both $\sigma(u) = \sigma(v)$ and $\tau(u) = \tau(v)$. Hence, they trivially fit condition (7.1). All of the remaining 12 μ's can be divided into pairs which are the same up to the order of the two vertices. Since the edges are undirected, for each of these pairs we simply count all the edges that could be assigned to either of the two components of μ. Then, to

ensure that the μ pairs satisfy (7.1), the count is equally split between the μ pair. Combining these thoughts yields

$$\mu_{++--} = \mu_{--++} := \frac{1}{dn} \sum_{(u,v)\in E} \mathbf{1}\{\sigma(u) = \tau(u) \neq \sigma(v) = \tau(v)\}$$
$$\mu_{+--+} = \mu_{-++-} := \frac{1}{dn} \sum_{(u,v)\in E} \mathbf{1}\{\sigma(u) = \tau(v) \neq \tau(u) = \sigma(v)\}$$
$$\mu_{+++-} = \mu_{+-++} := \frac{1}{dn} \sum_{(u,v)\in E} \mathbf{1}\{\sigma(u) = \sigma(v) = +1 \wedge \tau(u) \neq \tau(v)\}$$
$$\mu_{++-+} = \mu_{-+++} := \frac{1}{dn} \sum_{(u,v)\in E} \mathbf{1}\{\sigma(u) \neq \sigma(v) \wedge \tau(u) = \tau(v) = +1\}$$
$$\mu_{---+} = \mu_{-+--} := \frac{1}{dn} \sum_{(u,v)\in E} \mathbf{1}\{\sigma(u) = \sigma(v) = -1 \wedge \tau(u) \neq \tau(v)\}$$
$$\mu_{--+-} = \mu_{+---} := \frac{1}{dn} \sum_{(u,v)\in E} \mathbf{1}\{\sigma(u) \neq \sigma(v) \wedge \tau(u) = \tau(v) = -1\}.$$

Finally, we need expressions to indicate which fraction of vertices is assigned a certain spin configuration $(\sigma_1,\tau_1) \in \{\pm 1\}^2$. This is achieved rather easily by defining

$$\rho_{\sigma_1,\tau_1} := \frac{1}{n} \sum_{v\in V} \mathbf{1}\{\sigma(v) = \sigma_1, \tau(v) = \tau_1\}$$

for $(\sigma_1,\tau_1) \in \{\pm 1\}^2$. With the definitions in place, we can move on to calculating the second moment. As a starting point we choose an equation that was derived in detail in [3].

Lemma 7.2 ((4.42) in [3])**.** *We have*

$$\mathbb{E}\left[Z^2_{\mathbf{G},\beta}\right] = \sum_{\mu\in\mathscr{U}} \frac{\mathscr{X}_\mu \mathscr{Y}_\mu \mathscr{Z}_\mu}{(dn-1)!!} \cdot \exp\left(-\beta\frac{dn}{2}\left(\sum_{\sigma\in A_1}\mu(\sigma) + 2\sum_{\sigma\in A_2}\mu(\sigma)\right)\right) \tag{7.2}$$

where $\mathscr{U}$ is the set of conceivable distributions μ, σ is of the form $\sigma := (\sigma_1,\tau_1,\sigma_2,\tau_2) \in \{\pm 1\}^4$, A_1 is defined by

$$A_1 := \left\{x \in \{\pm 1\}^4 : \sigma_1 = \sigma_2 \text{ and } \tau_1 \neq \tau_2\right\} \cup \left\{x \in \{\pm 1\}^4 : \sigma_1 \neq \sigma_2 \text{ and } \tau_1 = \tau_2\right\},$$

A_2 is given by

$$A_2 := \left\{x \in \{\pm 1\}^4 : \sigma_1 = \sigma_2 \text{ and } \tau_1 = \tau_2\right\},$$

and

$$\mathscr{X}_\mu = \binom{n}{\rho_{++}n, \rho_{+-}n, \rho_{-+}n, \rho_{--}n},$$
$$\mathscr{Y}_\mu = \prod_{i,j\in\{\pm\}} \binom{dn\rho_{ij}}{dn\mu_{ij++}, dn\mu_{ij+-}, dn\mu_{ij-+}, dn\mu_{ij--}},$$
$$\mathscr{Z}_\mu = (dn\mu_{++--})!(dn\mu_{--++})! \prod_{k\in\{\pm\}} \left((dn\mu_{+k-k})!(dn\mu_{k+k-})!\right) \prod_{i,j\in\{\pm\}} (dn\mu_{ijij} - 1)!!.$$

7.2. **Reformulation of the second moment.** The next Lemma equips us with an useful reformulation of the second moment.

Lemma 7.3. *We have*

$$\mathbb{E}\left[Z^2_{\mathbf{G},\beta}\right] = \sum_{\mu\in\mathscr{U}} \frac{1}{8d^3\pi^{\frac{9}{2}} n^{\frac{9}{2}} \sqrt{\prod_{\sigma\in B}\mu_\sigma}} \exp\left(n\delta(\mu,\rho) + O\left(\frac{1}{n}\right)\right)$$

where $\mathscr{U}$ is the set of conceivable distributions μ, σ is of the form $\sigma := (\sigma_1,\tau_1,\sigma_2,\tau_2) \in \{\pm 1\}^4$, $\delta(\mu,\rho)$ is defined by

$$\delta(\mu,\rho) := H(\rho) - \frac{d}{2}\left(D_{KL}(\mu||\rho\otimes\rho) + \beta\sum_{\sigma\in A_1}\mu(\sigma) + 2\beta\sum_{\sigma\in A_2}\mu(\sigma)\right)$$

and

$$B := \{\pm\}^4 \setminus \{(++--),(-++-),(++-+),(+++-),(+---),(-+--)\}.$$

Proof. We start off the formulation of the second moment from Lemma 7.2. In the next lines, we will establish four asymptotic equalities[1]. Let us start with

$$(dn-1)!! = \frac{(dn)!}{\left(\frac{dn}{2}\right)!2^{\frac{dn}{2}}} = 2^{-\frac{dn}{2}}\cdot\sqrt{2}\left(\frac{dn}{e}\right)^{dn}\left(\frac{dn}{2e}\right)^{-\frac{dn}{2}}\exp\left(O\left(\frac{1}{n}\right)\right) = \exp\left(\frac{1}{2}\log(2)+\frac{dn}{2}\log(dn)-\frac{dn}{2}+O\left(\frac{1}{n}\right)\right).$$

Second, we take a closer look at

$$\begin{aligned}
\mathscr{X}_\mu &= \binom{n}{\rho_{++}n,\rho_{+-}n,\rho_{-+}n,\rho_{--}n} = \frac{n!}{(\rho_{++}n)!(\rho_{+-}n)!(\rho_{-+}n)!(\rho_{--}n)!}\\
&= (2\pi)^{-\frac{3}{2}}\left(n^3\rho_{++}\rho_{+-}\rho_{-+}\rho_{--}\right)^{-\frac{1}{2}}\left(\frac{n}{e}\right)^n\left(\frac{\rho_{++}n}{e}\right)^{-\rho_{++}n}\left(\frac{\rho_{+-}n}{e}\right)^{-\rho_{+-}n}\\
&\cdot\left(\frac{\rho_{-+}n}{e}\right)^{-\rho_{-+}n}\left(\frac{\rho_{--}n}{e}\right)^{-\rho_{--}n}\exp\left(O\left(\frac{1}{n}\right)\right)\\
&= (2\pi)^{-\frac{3}{2}}\left(n^3\rho_{++}\rho_{+-}\rho_{-+}\rho_{--}\right)^{-\frac{1}{2}}\underbrace{\rho_{++}^{-\rho_{++}n}\rho_{+-}^{-\rho_{+-}n}\rho_{-+}^{-\rho_{-+}n}\rho_{--}^{-\rho_{--}n}}_{=\exp(n\cdot\mathrm{H}(\rho))}\exp\left(O\left(\frac{1}{n}\right)\right)\\
&= \frac{1}{\sqrt{8\pi^3n^3\rho_{++}\rho_{+-}\rho_{-+}\rho_{--}}}\exp\left(n\cdot\mathrm{H}(\rho)+O\left(\frac{1}{n}\right)\right).
\end{aligned}$$

Moving on to the third term, we obtain

$$\begin{aligned}
\mathscr{Y}_\mu &= \prod_{i,j\in\{\pm\}}\binom{dn\rho_{ij}}{dn\mu_{ij++},dn\mu_{ij+-},dn\mu_{ij-+},dn\mu_{ij--}}\\
&= \prod_{i,j\in\{\pm\}}(2\pi dn)^{-\frac{3}{2}}\sqrt{\frac{\rho_{ij}}{\mu_{ij++}\mu_{ij+-}\mu_{ij-+}\mu_{ij--}}}\\
&\qquad\cdot\rho_{ij}^{dn\rho_{ij}}\mu_{ij++}^{-dn\mu_{ij++}}\mu_{ij+-}^{-dn\mu_{ij+-}}\mu_{ij-+}^{-dn\mu_{ij-+}}\mu_{ij--}^{-dn\mu_{ij--}}\exp\left(O\left(\frac{1}{n}\right)\right)\\
&= \frac{1}{(2\pi dn)^6}\left(\prod_{i,j\in\{\pm\}}\sqrt{\frac{\rho_{ij}}{\mu_{ij++}\mu_{ij+-}\mu_{ij-+}\mu_{ij--}}}\right)\exp\left(dn\left(\mathrm{H}(\mu)-\mathrm{H}(\rho)\right)+O\left(\frac{1}{n}\right)\right)
\end{aligned}$$

Last, we consider the fourth term

$$\begin{aligned}
\mathscr{Z}_\mu &= (dn\mu_{++--})!(dn\mu_{-++-})!\prod_{k\in\{\pm\}}\left((dn\mu_{+k-k})!(dn\mu_{k+k-})!\right)\prod_{i,j\in\{\pm\}}(dn\mu_{ijij}-1)!!\\
&= 2\pi dn\sqrt{\mu_{++--}\mu_{-++-}}\left(\frac{dn\mu_{++--}}{e}\right)^{dn\mu_{++--}}\left(\frac{dn\mu_{-++-}}{e}\right)^{dn\mu_{-++-}}\cdot\exp\left(O\left(\frac{1}{n}\right)\right)\\
&\qquad\cdot\prod_{k\in\{\pm\}}\left(2\pi dn\sqrt{\mu_{+k-k}\mu_{k+k-}}\left(\frac{dn\mu_{+k-k}}{e}\right)^{dn\mu_{+k-k}}\left(\frac{dn\mu_{k+k-}}{e}\right)^{dn\mu_{k+k-}}\right)\cdot\prod_{i,j\in\{\pm\}}\frac{(dn\mu_{ijij})!}{\left(\frac{dn\mu_{ijij}}{2}\right)!\cdot 2^{\frac{dn\mu_{ijij}}{2}}}
\end{aligned}$$

In order to proceed with the fourth equation, we keep in mind

$$\begin{aligned}
\frac{(dn\mu_{ijij})!}{\left(\frac{dn\mu_{ijij}}{2}\right)!\cdot 2^{\frac{dn\mu_{ijij}}{2}}} &= 2^{-\frac{dn\mu_{ijij}}{2}}\cdot\sqrt{2}\left(\frac{dn\mu_{ijij}}{e}\right)^{dn\mu_{ijij}}\left(\frac{dn\mu_{ijij}}{2e}\right)^{-\frac{dn\mu_{ijij}}{2}}\cdot\exp\left(O\left(\frac{1}{n}\right)\right)\\
&= \sqrt{2}\left(\frac{dn\mu_{ijij}}{e}\right)^{\frac{dn\mu_{ijij}}{2}}\cdot\exp\left(O\left(\frac{1}{n}\right)\right)
\end{aligned}$$

[1]The basic idea for the proof is the same as the one in [3]. The contribution of this paper is a more precise calculation that allows us to reduce the error term to order $\exp O\left(\frac{1}{n}\right)$.

which brings us back to

$$\begin{aligned}
\mathcal{I}_\mu &= (2\pi dn)^3 \sqrt{\mu_{++--}\mu_{-++-}} \left(\frac{dn\mu_{++--}}{e}\right)^{dn\mu_{++--}} \left(\frac{dn\mu_{-++-}}{e}\right)^{dn\mu_{-++-}} \cdot \exp\left(O\left(\frac{1}{n}\right)\right) \\
&\quad \cdot \prod_{k\in\{\pm\}} \left(\sqrt{\mu_{+k-k}\mu_{k+k-}} \left(\frac{dn\mu_{+k-k}}{e}\right)^{dn\mu_{+k-k}} \left(\frac{dn\mu_{k+k-}}{e}\right)^{dn\mu_{k+k-}}\right) \cdot \prod_{i,j\in\{\pm\}} \sqrt{2} \left(\frac{dn\mu_{ijij}}{e}\right)^{\frac{dn\mu_{ijij}}{2}} \\
&= 4\,(2\pi dn)^3 \sqrt{\mu_{++--}\mu_{-++-}} \left(\frac{dn\mu_{++--}}{e}\right)^{dn\mu_{++--}} \left(\frac{dn\mu_{-++-}}{e}\right)^{dn\mu_{-++-}} \cdot \exp\left(O\left(\frac{1}{n}\right)\right) \\
&\qquad \cdot \sqrt{\mu_{++-+}\mu_{+++-}} \left(\frac{dn\mu_{++-+}}{e}\right)^{dn\mu_{++-+}} \left(\frac{dn\mu_{+++-}}{e}\right)^{dn\mu_{+++-}} \sqrt{\mu_{+---}\mu_{-+--}} \\
&\qquad \cdot \left(\frac{dn\mu_{+---}}{e}\right)^{dn\mu_{+---}} \left(\frac{dn\mu_{-+--}}{e}\right)^{dn\mu_{-+--}} \left(\frac{dn\mu_{++++}}{e}\right)^{\frac{dn\mu_{++++}}{2}} \\
&\qquad \cdot \left(\frac{dn\mu_{+-+-}}{e}\right)^{\frac{dn\mu_{+-+-}}{2}} \left(\frac{dn\mu_{-+-+}}{e}\right)^{\frac{dn\mu_{-+-+}}{2}} \left(\frac{dn\mu_{----}}{e}\right)^{\frac{dn\mu_{----}}{2}}.
\end{aligned}$$

To simplify this rather complicated term further, we recall the symmetry of our model (see (7.1)). Applying this insight to our calculation yields

$$\mathcal{I}_\mu = 2^5\pi^3 d^3 n^3 \sqrt{\mu_{++--}\mu_{-++-}\mu_{++-+}\mu_{+++-}\mu_{+---}\mu_{-+--}} \cdot \exp\left(-\frac{dn}{2}\cdot \mathrm{H}(\mu) - \frac{dn}{2} + \frac{dn}{2}\log(dn) + O\left(\frac{1}{n}\right)\right).$$

Next, we combine these four results starting with

$$\frac{\mathcal{I}_\mu}{(dn-1)!!} = 2^{\frac{9}{2}}\pi^3 d^3 n^3 \sqrt{\mu_{++--}\mu_{-++-}\mu_{++-+}\mu_{+++-}\mu_{+---}\mu_{-+--}} \cdot \exp\left(-\frac{dn}{2}\cdot \mathrm{H}(\mu) + O\left(\frac{1}{n}\right)\right).$$

Finally, we arrive at

$$\begin{aligned}
\frac{\mathcal{X}_\mu\mathcal{Y}_\mu\mathcal{I}_\mu}{(dn-1)!!} &= \exp\left(\frac{dn}{2}\mathrm{H}(\mu) - n(d-1)\,\mathrm{H}(\rho) - \frac{9}{2}\log(n) - \frac{9}{2}\log(2\pi) - 3\log(d) + O\left(\frac{1}{n}\right)\right) \\
&\qquad \cdot \exp\left(-\frac{1}{2}\log\left(\rho_{++}\rho_{+-}\rho_{-+}\rho_{--}\right) + \frac{1}{2}\sum_{i,j\in\{\pm\}} \log\left(\frac{\rho_{ij}}{\mu_{ij++}\mu_{ij+-}\mu_{ij-+}\mu_{ij--}}\right)\right) \\
&\qquad \cdot \exp\left(\frac{1}{2}\log\left(\mu_{++--}\mu_{-++-}\mu_{++-+}\mu_{+++-}\mu_{+---}\mu_{-+--}\right) + \frac{3}{2}\log(2)\right) \\
&= \exp\left(\frac{dn}{2}\left(\mathrm{H}(\mu) - 2\mathrm{H}(\rho)\right) + n\mathrm{H}(\rho) - \frac{9}{2}\log(2\pi n) - 3\log(d) + O\left(\frac{1}{n}\right)\right) \\
&\qquad \cdot \exp\left(-\frac{1}{2}\log\left(\rho_{++}\rho_{+-}\rho_{-+}\rho_{--}\right) + \frac{1}{2}\log\left(\frac{\rho_{++}\rho_{+-}\rho_{-+}\rho_{--}}{\prod_{\sigma\in\{\pm\}^4}\mu_\sigma}\right)\right) \\
&\qquad \cdot \exp\left(\frac{1}{2}\log\left(\mu_{++--}\mu_{-++-}\mu_{++-+}\mu_{+++-}\mu_{+---}\mu_{-+--}\right) + \frac{3}{2}\log(2)\right) \\
&= \exp\left(-\frac{dn}{2}D_{\mathrm{KL}}(\mu||\rho\otimes\rho) + n\mathrm{H}(\rho) - \frac{9}{2}\log(\pi n) - 3\log(2d) - \frac{1}{2}\log\left(\prod_{\sigma\in\{\pm\}^4}\mu_\sigma\right)\right) \\
&\qquad \cdot \exp\left(\frac{1}{2}\log\left(\mu_{++--}\mu_{-++-}\mu_{++-+}\mu_{+++-}\mu_{+---}\mu_{-+--}\right) + O\left(\frac{1}{n}\right)\right)
\end{aligned}$$

To simplify this expression, we introduce the set

$$B = \{\pm\}^4 \setminus \{(++--),(-++-),(++-+),(+++-),(+---),(-+--)\}$$

to write

$$\frac{\mathcal{X}_\mu\mathcal{Y}_\mu\mathcal{I}_\mu}{(dn-1)!!} = \frac{1}{8d^3\pi^{\frac{9}{2}}n^{\frac{9}{2}}\sqrt{\prod_{\sigma\in B}\mu_\sigma}} \exp\left(n\mathrm{H}(\rho) - \frac{dn}{2}D_{\mathrm{KL}}(\mu||\rho\otimes\rho) + O\left(\frac{1}{n}\right)\right).$$

With this result , the second moment turns out to be

$$\begin{aligned}\mathbb{E}\left[Z_{\mathbf{G},\beta}^2\right] &= \sum_{\mu\in\mathscr{U}} \frac{\mathscr{X}_\mu \mathscr{Y}_\mu \mathscr{Z}_\mu}{(dn-1)!!} \cdot \exp\left(-\beta\frac{dn}{2}\left(\sum_{\sigma\in A_1}\mu(\sigma)+2\sum_{\sigma\in A_2}\mu(\sigma)\right)\right)\\ &= \sum_{\mu\in\mathscr{U}} \frac{1}{8d^3\pi^{\frac{9}{2}}n^{\frac{9}{2}}\sqrt{\prod_{\sigma\in B}\mu_\sigma}} \exp\left(n\delta\left(\mu,\rho\right)+O\left(\frac{1}{n}\right)\right)\end{aligned}$$

where $\mathscr{U}$ is the set of conceivable distributions μ, σ is of the form $\sigma := (\sigma_1,\tau_1,\sigma_2,\tau_2) \in \{\pm 1\}^4$, $\delta\left(\mu,\rho\right)$ is defined by

$$\delta\left(\mu,\rho\right) := \mathrm{H}\left(\rho\right) - \frac{d}{2}\left(D_{\mathrm{KL}}\left(\mu||\rho\otimes\rho\right)+\beta\sum_{\sigma\in A_1}\mu(\sigma)+2\beta\sum_{\sigma\in A_2}\mu(\sigma)\right)$$

A_1 is defined by

$$A_1 := \left\{x\in\{\pm1\}^4 : \sigma_1=\sigma_2 \text{ and } \tau_1\neq\tau_2\right\}\cup\left\{x\in\{\pm1\}^4 : \sigma_1\neq\sigma_2 \text{ and } \tau_1=\tau_2\right\},$$

and A_2 is given by

$$A_2 := \left\{x\in\{\pm1\}^4 : \sigma_1=\sigma_2 \text{ and } \tau_1=\tau_2\right\}.$$

□

Our ultimate goal is to apply the Laplace method. In order to keep things manageable, we will substitute certain variables using basic symmetry and composition arguments. First, we note that ρ can be simply obtained by calculating the marginals of μ, that is

$$\begin{aligned}\rho_{++} &= \mu_{++++}+\mu_{+++-}+\mu_{++-+}+\mu_{++--}\\ \rho_{+-} &= \mu_{+-++}+\mu_{+-+-}+\mu_{+--+}+\mu_{+---}\\ \rho_{-+} &= \mu_{-+++}+\mu_{-++-}+\mu_{-+-+}+\mu_{-+--}\\ \rho_{--} &= \mu_{--++}+\mu_{--+-}+\mu_{---+}+\mu_{----}.\end{aligned}$$

By construction we know that

$$\begin{aligned}&\mu_{+--+}=\mu_{-++-}, \quad \mu_{++--}=\mu_{--++}, \quad \mu_{+++-}=\mu_{+-++}\\ &\mu_{++-+}=\mu_{-+++}, \quad \mu_{+---}=\mu_{--+-}, \quad \mu_{-+--}=\mu_{---+}\\ &\mu_{----}=1-\sum_{\sigma\in\{\pm1\}^4,\sigma\neq(-1,-1,-1,-1)}\mu_\sigma\end{aligned}$$

Bringing these results together we are left with 9 variables, which we rename in the following order for notational convenience

$$\begin{aligned}&x_1 := \mu_{+--+}=\mu_{-++-}, \quad x_2 := \mu_{++--}=\mu_{--++}, \quad x_3 := \mu_{+++-}=\mu_{+-++},\\ &x_4 := \mu_{++-+}=\mu_{-+++}, \quad x_5 := \mu_{+---}=\mu_{--+-}, \quad x_6 := \mu_{-+--}=\mu_{---+}\\ &x_7 := \mu_{+-+-}, \quad x_8 := \mu_{-+-+}, \quad x_9 := \mu_{++++}\end{aligned}$$

which implies

$$\begin{aligned}\mu_{----} &= 1-\sum_{\sigma\in\{\pm1\}^4,\sigma\neq(-1,-1,-1,-1)}\mu_\sigma = 1-2x_1-2x_2-2x_3-2x_4-2x_5-2x_6-x_7-x_8-x_9\\ \rho_{++} &= x_9+x_3+x_4+x_2 = x_2+x_3+x_4+x_9\\ \rho_{+-} &= x_3+x_7+x_1+x_5 = x_1+x_3+x_5+x_7\\ \rho_{-+} &= x_4+x_1+x_8+x_6 = x_1+x_4+x_6+x_8\\ \rho_{--} &= x_2+x_5+x_6+\mu_{----} = 1-2x_1-x_2-2x_3-2x_4-x_5-x_6-x_7-x_8-x_9.\end{aligned}$$

In order to apply the Laplace method to the second moment, let us consider the function

$$\delta\left(\mu,\rho\right) = \mathrm{H}\left(\rho\right) - \frac{d}{2}\left(D_{\mathrm{KL}}\left(\mu||\rho\otimes\rho\right)+\beta\sum_{\sigma\in A_1}\mu(\sigma)+2\beta\sum_{\sigma\in A_2}\mu(\sigma)\right).$$

We continue by reformulating terms

$$\begin{aligned}
D_{\mathrm{KL}}\left(\mu||\rho\otimes\rho\right) &= 2x_1\log\left(\frac{x_1}{\rho_{+-}\rho_{-+}}\right)+2x_2\log\left(\frac{x_2}{\rho_{++}\rho_{--}}\right)+2x_3\log\left(\frac{x_3}{\rho_{++}\rho_{+-}}\right)+2x_4\log\left(\frac{x_4}{\rho_{++}\rho_{-+}}\right)\\
&\quad+2x_5\log\left(\frac{x_5}{\rho_{+-}\rho_{--}}\right)+2x_6\log\left(\frac{x_6}{\rho_{-+}\rho_{--}}\right)+x_7\log\left(\frac{x_7}{\rho_{+-}^2}\right)+x_8\log\left(\frac{x_8}{\rho_{-+}^2}\right)\\
&\quad+x_9\log\left(\frac{x_9}{\rho_{++}^2}\right)+\mu_{----}\log\left(\frac{\mu_{----}}{\rho_{--}^2}\right)\\
&= 2x_1\log\left(x_1\right)+2x_2\log\left(x_2\right)+2x_3\log\left(x_3\right)+2x_4\log\left(x_4\right)+2x_5\log\left(x_5\right)+2x_6\log\left(x_6\right)\\
&\quad+x_7\log\left(x_7\right)+x_8\log\left(x_8\right)+x_9\log\left(x_9\right)+\mu_{----}\log\left(\mu_{----}\right)-2\rho_{++}\log\left(\rho_{++}\right)\\
&\quad-2\rho_{+-}\log\left(\rho_{+-}\right)-2\rho_{-+}\log\left(\rho_{-+}\right)-2\rho_{--}\log\left(\rho_{--}\right).
\end{aligned}$$

As a consequence, we obtain

$$\begin{aligned}
\delta\left(\mu,\rho\right) &= \mathrm{H}\left(\rho\right)-\frac{d}{2}\left(D_{\mathrm{KL}}\left(\mu||\rho\otimes\rho\right)+\beta\sum_{\sigma\in A_1}\mu(\sigma)+2\beta\sum_{\sigma\in A_2}\mu(\sigma)\right)\\
&= (1-d)\,\mathrm{H}\left(\rho\right)+\frac{d}{2}\mathrm{H}\left(\mu\right)-d\beta\left(x_3+x_4+x_5+x_6+x_7+x_8+x_9+\mu_{----}\right)
\end{aligned}$$

7.3. **Application of the Laplace method to the second moment.** To apply the Laplace method we need to determine the maximum of $\delta\left(\mu,\rho\right)$. This is achieved with the following Lemma. Due to its technical and tedious nature, the proof of the lemma is outsourced to a separate section (see section 8).

Lemma 7.4. *For $0<\beta<\beta_{KS}$, we have*

$$\max_{\mu\in\mathcal{O}'}\delta\left(\mu,\rho\right)=\delta\left(\mu^*,\rho^*\right)=(2-d)\log(2)+d\log\left(1+e^{-\beta}\right)$$

where $\mathcal{O}'$ denotes the set of all μ that are conceivable under the assumption that the event $\mathcal{O}$ occurs. The unique maximum is obtained at

$$\begin{aligned}
\mu^*_{++++}=\mu^*_{----}=\mu^*_{+-+-}=\mu^*_{-+-+}&=\frac{e^{-2\beta}}{4\left(1+e^{-\beta}\right)^2}\\
\mu^*_{+--+}=\mu^*_{-++-}=\mu^*_{-+--}=\mu^*_{++--}&=\frac{1}{4\left(1+e^{-\beta}\right)^2}\\
\mu^*_{+++-}=\mu^*_{++-+}=\mu^*_{+-++}=\mu^*_{-+++}=\mu^*_{---+}=\mu^*_{--+-}=\mu^*_{-+--}=\mu^*_{+---}&=\frac{e^{-\beta}}{4\left(1+e^{-\beta}\right)^2}
\end{aligned}$$

which also implies

$$\rho^*_{++}=\rho^*_{+-}=\rho^*_{-+}=\rho^*_{--}=\frac{1}{4}.$$

Having determined the maximum, we next need to evaluate the Hessian at the optimal point. The derivation of the Hessian matrix and evaluation at the optimal point is not too difficult. Thus, we just state the result here and refer the interested reader to Section 9.

Lemma 7.5 (Hessian for the second moment)**.** *We have*

$$\det\left(-\mathrm{D}^2\delta\left(\mu^*,\rho^*\right)\right)=2^{17}d^6e^{-8\beta}\left(1+e^{\beta}\right)^{16}\left(de^{\beta}-d+2\right)^2\left(2e^{2\beta}+2de^{\beta}-de^{2\beta}-d+2\right).$$

Proof of Proposition 7.1. With Lemmas 7.4 and 7.5 in place, we still need to determine the lattice matrix and its determinant. Similar to the notation for the first moment, we let A_{second} denote the matrix consisting of the

elements in the basis of the lattice for the second moment. Recalling the following definitions

$$
\begin{aligned}
x_1 &= \mu_{+--+} = \mu_{-++-} = \frac{1}{dn} \sum_{(u,v)\in E} \mathbf{1}\{\sigma(u) = \tau(v) \neq \tau(u) = \sigma(v)\} \\
x_3 &= \mu_{+++-} = \mu_{+-++} = \frac{1}{dn} \sum_{(u,v)\in E} \mathbf{1}\{\sigma(u) = \sigma(v) = +1 \wedge \tau(u) \neq \tau(v)\} \\
x_4 &= \mu_{++-+} = \mu_{-+++} = \frac{1}{dn} \sum_{(u,v)\in E} \mathbf{1}\{\sigma(u) \neq \sigma(v) \wedge \tau(u) = \tau(v) = +1\} \\
x_7 &= \mu_{+-+-} = \frac{2}{dn} \sum_{(u,v)\in E} \mathbf{1}\{\sigma(u) = \sigma(v) = +1 \wedge \tau(u) = \tau(v) = -1\} \\
x_8 &= \mu_{-+-+} = \frac{2}{dn} \sum_{(u,v)\in E} \mathbf{1}\{\sigma(u) = \sigma(v) = -1 \wedge \tau(u) = \tau(v) = +1\} \\
x_9 &= \mu_{++++} = \frac{2}{dn} \sum_{(u,v)\in E} \mathbf{1}\{\sigma(u) = \sigma(v) = \tau(u) = \tau(v) = +1\}
\end{aligned}
$$

we immediately obtain the diagonal entries for the respective x's, i.e $\frac{1}{d}$ and $\frac{2}{d}$. From here on, things get more complicated. Since $\rho_{++}, \rho_{+-}, \rho_{-+}$, and ρ_{--} each count fractions of the set of nodes (which contains n nodes in total) their entries in the lattice matrix all have to be multiples of $\frac{1}{n}$. Furthermore, we recall the following binding conditions

$$
\begin{aligned}
\rho_{++} &= x_9 + x_3 + x_4 + x_2 = x_2 + x_3 + x_4 + x_9 \\
\rho_{+-} &= x_3 + x_7 + x_1 + x_5 = x_1 + x_3 + x_5 + x_7 \\
\rho_{-+} &= x_4 + x_1 + x_8 + x_6 = x_1 + x_4 + x_6 + x_8.
\end{aligned}
$$

Combining these two points, x_2, x_5, and x_6 each need to be chosen such that the sums consisting of four summands each add up to a number that is a multiple of $\frac{1}{n}$. Let us focus on x_2. Similar arguments apply to x_5 and x_6. For x_2, the above equation can be reformulated as

$$
x_2 = \rho_{++} - x_3 - x_4 - x_9 = \frac{b_2}{n} - \frac{b_3}{dn} - \frac{b_4}{dn} - \frac{2 \cdot b_9}{dn}
$$

where $b_i \in \mathbb{N}, i \in [9]$ are the scalars for the linear combination yielding the desired μ. From the reformulated equation we immediately obtain the matrix entries $A_{\text{second},2,2} = 1$, $A_{\text{second},2,3} = -1/d$, $A_{\text{second},2,4} = -1/d$, and $A_{\text{second},2,9} = -2/d$. Following through this procedure for x_5 and x_6, we obtain the remaining entries of A_{second} that are different from zero. This enables us to calculate the determinant of interest:

$$
\det(A_{\text{second}}) = \det \begin{pmatrix}
\frac{1}{d} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & -\frac{1}{d} & -\frac{1}{d} & 0 & 0 & 0 & 0 & -\frac{2}{d} \\
0 & 0 & \frac{1}{d} & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & \frac{1}{d} & 0 & 0 & 0 & 0 & 0 \\
-\frac{1}{d} & 0 & -\frac{1}{d} & 0 & 1 & 0 & -\frac{2}{d} & 0 & 0 \\
-\frac{1}{d} & 0 & 0 & -\frac{1}{d} & 0 & 1 & 0 & -\frac{2}{d} & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & \frac{2}{d} & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & \frac{2}{d} & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \frac{2}{d}
\end{pmatrix} = \frac{2^3}{d^6}. \tag{7.3}
$$

With these results in place, we can apply Theorem 4.1 to the expression for the second moment in Lemma 7.3 which yields

$$
\begin{aligned}
\mathbb{E}\left[Z^2_{\mathbf{G},\beta}\mathbf{1}\{\mathscr{O}\}\right] &= \exp\left(O\left(\frac{1}{n}\right)\right) \cdot \sum_{\mu\in\mathscr{U}} \frac{1}{8d^3\pi^{\frac{9}{2}} n^{\frac{9}{2}} \sqrt{\prod_{\sigma\in B}\mu_\sigma}} \exp\left(n\delta\left(\mu,\rho\right)\right) \\
&= \exp\left(O\left(\frac{1}{n}\right)\right) \cdot \frac{(2\pi n)^{\frac{9}{2}} \exp\left(n\delta\left(\mu^*,\rho^*\right)\right)}{8d^3\pi^{\frac{9}{2}} n^{\frac{9}{2}} \sqrt{\prod_{\sigma\in B}\mu^*_\sigma} \det(A_{\text{second}}) \sqrt{\det\left(-\mathrm{D}^2\delta\left(\mu^*,\rho^*\right)\right)}}.
\end{aligned}
$$

Using Lemmas 7.4 and 7.5 and the determinant of the lattice matrix from (7.3), we arrive at

$$\mathbb{E}\left[Z^2_{\mathbf{G},\beta}\mathbf{1}\{\mathcal{O}\}\right] = \exp\left(O\left(\frac{1}{n}\right)\right)\cdot\frac{\left(1+e^{-\beta}\right)^{10}\exp\left(n\left((2-d)\log(2)+d\log\left(1+e^{-\beta}\right)\right)\right)}{e^{-10\beta}\left(1+e^{\beta}\right)^8\left(de^{\beta}-d+2\right)\sqrt{\left(2e^{2\beta}+2de^{\beta}-de^{2\beta}-d+2\right)}}$$
$$= \exp\left(O\left(\frac{1}{n}\right)\right)\cdot\frac{\left(1+e^{\beta}\right)^{2}\exp\left(n\left((2-d)\log(2)+d\log\left(1+e^{-\beta}\right)\right)\right)}{\left(de^{\beta}-d+2\right)\sqrt{\left(2e^{2\beta}+2de^{\beta}-de^{2\beta}-d+2\right)}}.$$

□

7.4. **The simple d-regular case.** Having established the second moment in the pairing model $\mathbf{G}$, we still have to adapt the result to the d-regular model $\mathbb{G}$ of interest. As we will see, a pairing variant (not the same as for the first moment) $\mathbf{G}_2^*$ of the planted model will be a useful tool to do so. The pairing variant $\mathbf{G}_2^*$ is defined as follows. First, draw two spin assignments $\boldsymbol{\sigma}^*, \boldsymbol{\tau}^* \in \{\pm 1\}^n$ independently and uniformly at random. Then, draw a graph $\mathbf{G}_2^*$ according to the probability distribution

$$\mathbb{P}\left[\mathbf{G}_2^* = G|\boldsymbol{\sigma}^*, \boldsymbol{\tau}^*\right] \propto \exp\left(-\beta\mathscr{H}_G\left(\boldsymbol{\sigma}^*\right) - \beta\mathscr{H}_G\left(\boldsymbol{\tau}^*\right)\right).$$

where G might again feature self-loops and double-edges. With some effort, we obtain the next result.

Lemma 7.6. *For $d \geq 0$ and $\beta > 0$ we have*

$$\mathbb{P}\left[\mathbf{G}_2^* \text{ is simple}\right] \sim \exp\left(-(d-1)\frac{2}{\left(1+e^{\beta}\right)^2} - (d-1)^2\frac{\left(1+e^{2\beta}\right)^2}{\left(1+e^{\beta}\right)^4}\right).$$

Proof of Lemma 7.6. This proof is based on an idea in [3] (Lemma 4.6). First of all, we are interested in the number of self-loops X in $\mathbf{G}_2^*$ on the one hand, and the number of double edges Y on the other hand. For notational convenience, we let $\mathscr{G}(\sigma,\mu)$ be the event that the generated graph has $\frac{dn}{2}\mu_{++++}$ edges that connect two vertices that each have been assigned two positive spins; the same is assumed to hold for all entries of μ and the respective types of edges. With these definitions in place, we move on to the expectations of X and Y. Instead of calculating the two directly, we decompose the two to simplify the following calculations.

So let us start with the number of self-loops X. Basically, there are four different types of self-loops in our model, X_{++}, X_{+-}, X_{-+}, and X_{--}. The index in each of the four cases just refers to the spin pair assigned to the vertex of the self-loop. Then, the expectation of X_{++} can be formulated as

$$\mathbb{E}\left[X_{++}|\mathscr{G}(\sigma,\mu)\right] = \frac{\rho_{++}n\binom{d}{2}\binom{dn\rho_{++}-2}{dn\mu_{++++}-2}\left(dn\mu_{++++}-3\right)!!}{\binom{dn\rho_{++}}{dn\mu_{++++}}\left(dn\mu_{++++}-1\right)!!} = \frac{\rho_{++}n\frac{d!}{2(d-2)!}\frac{(dn\rho_{++}-2)!}{(dn\mu_{++++}-2)!(dn\rho_{++}-dn\mu_{++++})!}}{\frac{(dn\rho_{++})!}{(dn\mu_{++++})!(dn\rho_{++}-dn\mu_{++++})!}\left(dn\mu_{++++}-1\right)}$$
$$= \frac{n\mu_{++++}d(d-1)}{2\left(dn\rho_{++}-1\right)} \sim \frac{\mu_{++++}(d-1)}{2\rho_{++}}$$

where, in the first step, we already cancelled out the factors that appeared both in the numerator and denominator. By almost identical calculations, we obtain

$$\mathbb{E}\left[X_{+-}|\mathscr{G}(\sigma,\mu)\right] \sim \frac{\mu_{+-+-}(d-1)}{2\rho_{+-}}, \qquad \mathbb{E}\left[X_{-+}|\mathscr{G}(\sigma,\mu)\right] \sim \frac{\mu_{-+-+}(d-1)}{2\rho_{-+}},$$
$$\text{and} \qquad \mathbb{E}\left[X_{--}|\mathscr{G}(\sigma,\mu)\right] \sim \frac{\mu_{----}(d-1)}{2\rho_{--}}.$$

Bringing these four results together and plugging in the optimal point (μ^*,ρ^*), we arrive at

$$\mathbb{E}\left[X|\mathscr{G}(\sigma,\mu)\right] \sim (d-1)\frac{2}{\left(1+e^{\beta}\right)^2}. \tag{7.4}$$

With a similar argument, we determine the expectation of the number of double edges Y. More precisely, we decompose Y into the random variables $Y_{\sigma_1,\tau_1,\sigma_2,\tau_2}$ with $(\sigma_1,\tau_1,\sigma_2,\tau_2) \in \{\pm 1\}^4$. Each $Y_{\sigma_1,\tau_1,\sigma_2,\tau_2}$ is just the number of double edges between two vertices where the first vertex is assigned to the spin-pair (σ_1,τ_1) and the second to the pair (σ_2,τ_2). Let us start with the four spin configurations with $(\sigma_1,\tau_1) = (\sigma_2,\tau_2)$. In order to

keep the calculations simple, we focus on Y_{++++} and then extend the results to Y_{+-+-}, Y_{-+-+}, and Y_{----}.

$$\begin{aligned}\mathbb{E}\left[Y_{++++}|\mathcal{G}(\sigma,\mu)\right] &= \frac{2\binom{\rho_{++}n}{2}\binom{d}{2}^2\binom{dn\rho_{++}-4}{dn\mu_{++++}-4}(dn\mu_{++++}-5)!!}{\binom{dn\rho_{++}}{dn\mu_{++++}}(dn\mu_{++++}-1)!!} \\ &\sim \frac{(\rho_{++}n)^2\left(\frac{d(d-1)}{2}\right)^2(dn\mu_{++++})^4}{(\rho_{++}dn)^4(dn\mu_{++++})^2} = \frac{(d-1)^2}{4}\frac{\mu_{++++}^2}{\rho_{++}^2}\end{aligned}$$

where, in the first step, we already cancelled out the factors that occured both in the numerator and denominator. Following this line of thought, we can also state

$$\mathbb{E}\left[Y_{+-+-}|\mathcal{G}(\sigma,\mu)\right] \sim \frac{(d-1)^2}{4}\frac{\mu_{+-+-}^2}{\rho_{+-}^2}, \qquad \mathbb{E}\left[Y_{-+-+}|\mathcal{G}(\sigma,\mu)\right] \sim \frac{(d-1)^2}{4}\frac{\mu_{-+-+}^2}{\rho_{-+}^2},$$
$$\text{and} \qquad \mathbb{E}\left[Y_{----}|\mathcal{G}(\sigma,\mu)\right] \sim \frac{(d-1)^2}{4}\frac{\mu_{----}^2}{\rho_{--}^2}.$$

For the next calculation, we consider the sum of Y_{+---} and Y_{--+-}. Since the edges in our model are undirected, it is not suitable to make a distinction between the two.

$$\begin{aligned}\mathbb{E}\left[Y_{+---}+Y_{--+-}|\mathcal{G}(\sigma,\mu)\right] &= \frac{2\rho_{+-}\rho_{--}n^2\binom{d}{2}^2\binom{dn\rho_{+-}-2}{dn\mu_{+---}-2}\binom{dn\rho_{--}-2}{dn\mu_{+---}-2}(dn\mu_{+---}-2)!}{\binom{dn\rho_{+-}}{dn\mu_{+---}}\binom{dn\rho_{--}}{dn\mu_{+---}}(dn\mu_{+---})!} \\ &\sim \frac{2\rho_{+-}\rho_{--}n^2\left(\frac{d(d-1)}{2}\right)^2(dn\mu_{+---})^4}{(\rho_{+-}dn)^2(\rho_{--}dn)^2(dn\mu_{+---})^2} = \frac{(d-1)^2}{4}\frac{\mu_{+---}^2}{\rho_{+-}\rho_{--}}\end{aligned}$$

Here, we once again tacitly cancelled out the factors in the first expression that are included in both the numerator and denominator. The same approach can be iteratively applied to the remaining types of double edges, which eventually yields

$$\mathbb{E}\left[Y_{-+--}+Y_{---+}|\mathcal{G}(\sigma,\mu)\right] \sim \frac{(d-1)^2}{4}\frac{\mu_{-+--}^2}{\rho_{-+}\rho_{--}}, \qquad \mathbb{E}\left[Y_{+++-}+Y_{+-++}|\mathcal{G}(\sigma,\mu)\right] \sim \frac{(d-1)^2}{4}\frac{\mu_{+++-}^2}{\rho_{++}\rho_{+-}},$$
$$\mathbb{E}\left[Y_{-+++}+Y_{++-+}|\mathcal{G}(\sigma,\mu)\right] \sim \frac{(d-1)^2}{4}\frac{\mu_{++-+}^2}{\rho_{++}\rho_{-+}}, \qquad \mathbb{E}\left[Y_{++--}+Y_{--++}|\mathcal{G}(\sigma,\mu)\right] \sim \frac{(d-1)^2}{4}\frac{\mu_{++--}^2}{\rho_{++}\rho_{--}},$$
$$\text{and} \qquad \mathbb{E}\left[Y_{+--+}+Y_{-++-}|\mathcal{G}(\sigma,\mu)\right] \sim \frac{(d-1)^2}{4}\frac{\mu_{+--+}^2}{\rho_{+-}\rho_{-+}}.$$

Taking the sum of all these findings and plugging in the optimal point (μ^*,ρ^*), we finally obtain

$$\mathbb{E}\left[Y|\mathcal{G}(\sigma,\mu)\right] \sim (d-1)^2\frac{\left(1+e^{2\beta}\right)^2}{\left(1+e^{\beta}\right)^4}. \tag{7.5}$$

With the statements (7.4) and (7.5) in mind, we claim that for all $k,\ell \geq 1$

$$\mathbb{E}\left[\prod_{i=1}^{k}(X-i+1)\prod_{j=1}^{\ell}(Y-j+1)\right] \sim \left((d-1)\frac{2}{\left(1+e^{\beta}\right)^2}\right)^k\left((d-1)^2\frac{\left(1+e^{2\beta}\right)^2}{\left(1+e^{\beta}\right)^4}\right)^{\ell} \tag{7.6}$$

holds. This can be seen as follows. In (7.4) and (7.5), we placed just one loop or double edge, respectively. To obtain (7.6), we now have to place some fixed numbers k and ℓ of self-loops and double edges. Since n approaches infinity, the probability that any choices of self-loops and double-edges overlap is bounded by $O(1/n)$. Thus, the desired result can be leveraged from (7.4) and (7.5).

With (7.6) in place, we immediately see

$$\mathbb{P}\left[\mathbf{G}_2^* \in \mathscr{S}\right] = \mathbb{P}\left[X=Y=0\right] \sim \exp\left(-(d-1)\frac{2}{\left(1+e^{\beta}\right)^2}-(d-1)^2\frac{\left(1+e^{2\beta}\right)^2}{\left(1+e^{\beta}\right)^4}\right).$$

which concludes the proof. □

Now, we are equipped to prove Proposition 2.6.

Proof of Proposition 2.6. We again use the asymptotic equality

$$\mathbb{E}\left[Z^2_{\mathbb{G},\beta}\mathbf{1}\{\mathscr{O}\}\right] \sim \frac{\mathbb{P}\left[\mathbf{G}^*_2 \text{ is simple}\right]}{\mathbb{P}\left[\mathbf{G} \text{ is simple}\right]}\mathbb{E}\left[Z^2_{\mathbf{G},\beta}\mathbf{1}\{\mathscr{O}\}\right]. \tag{7.7}$$

Thus, the desired result is obtained by combining equation (7.7), Proposition 7.1, Lemma 7.6, and Fact 6.7. □

8. Second Moment Optimization / Proof of Lemma 7.4

In this section we solve the maximization problem

$$\max_{\mu\in\mathscr{O}'}\delta\left(\mu,\rho\right)$$

where

$$\begin{aligned}
\delta\left(\mu,\rho\right) &:= \mathrm{H}\left(\rho\right) - \frac{d}{2}\left(D_{\mathrm{KL}}\left(\mu||\rho\otimes\rho\right) + \beta\sum_{\sigma\in A_1}\mu(\sigma) + 2\beta\sum_{\sigma\in A_2}\mu(\sigma)\right) \qquad \text{and}\\
A_1 &:= \left\{x\in\{\pm1\}^4 : \sigma_1=\sigma_2 \text{ and } \tau_1\neq\tau_2\right\}\cup\left\{x\in\{\pm1\}^4 : \sigma_1\neq\sigma_2 \text{ and } \tau_1=\tau_2\right\} \qquad \text{and}\\
A_2 &:= \left\{x\in\{\pm1\}^4 : \sigma_1=\sigma_2 \text{ and } \tau_1=\tau_2\right\}.
\end{aligned}$$

with vectors of the form $\sigma=(\sigma_1,\tau_1,\sigma_2,\tau_2)\in\{\pm1\}^4$. Furthermore, $\mathscr{O}'$ denotes the set of all μ that are conceivable given that the event $\mathscr{O}$ from (2.5) holds. At this point, we exploit the spatial mixing argument. Keeping Lemma 2.5 in mind, we limit our attention to the event $\mathscr{O}$. This is a crucial step for the following calculations because it allows us to perform the reparametrization

$$\rho_\alpha(\sigma_1,\tau_1) := \frac{1+\alpha\cdot\mathbf{1}\{\sigma_1=\tau_1\}-\alpha\cdot\mathbf{1}\{\sigma_1\neq\tau_1\}}{4} \tag{8.1}$$

where $\alpha\in(-1,+1)$ and $(\sigma_1,\tau_1)\in\{\pm1\}^2$. Now, the proof strategy is as follows. First, we minimize δ with respect to μ. This will provide us with a solution of μ formulated in terms of ρ or α, respectively. In a second step, all that remains to do is to maximize the function δ with respect to α.

8.1. Minimization with respect to μ.

Instead of solving the optimization in one step, we start by considering

$$g(\mu,\rho_\alpha) := D_{\mathrm{KL}}(\mu||\rho_\alpha\otimes\rho_\alpha) + \beta\cdot\sum_{\sigma\in A_1}\mu(\sigma) + 2\cdot\beta\cdot\sum_{\sigma\in A_2}\mu(\sigma)$$

where A_1 is defined by

$$A_1 := \left\{x\in\{\pm1\}^4 : \sigma_1=\sigma_2 \text{ and } \tau_1\neq\tau_2\right\}\cup\left\{x\in\{\pm1\}^4 : \sigma_1\neq\sigma_2 \text{ and } \tau_1=\tau_2\right\}$$

and A_2 is given by

$$A_2 := \left\{x\in\{\pm1\}^4 : \sigma_1=\sigma_2 \text{ and } \tau_1=\tau_2\right\}.$$

Note that the entropy term $H(\rho)$ is independent of μ and thus not relevant for optimizing with respect to μ. The above formulation immediately brings us to the constrained minimization problem

$$\begin{aligned}
&\min_{\mu\in\mathscr{O}'} g(\mu,\rho_\alpha)\\
\text{s.t. } \forall\,(\sigma_1,\tau_1)\in\{\pm1\}^2 :\ &\sum_{(\sigma_2,\tau_2)\in\{\pm1\}^2}\mu(\sigma_1,\tau_1,\sigma_2,\tau_2) = \rho_\alpha(\sigma_1,\tau_1)\\
\forall\,(\sigma_2,\tau_2)\in\{\pm1\}^2 :\ &\sum_{(\sigma_1,\tau_1)\in\{\pm1\}^2}\mu(\sigma_1,\tau_1,\sigma_2,\tau_2) = \rho_\alpha(\sigma_2,\tau_2)
\end{aligned}$$

where $\mathscr{P}\left(\{\pm1\}^4\right)$ denotes the set of all probability distributions on $\{\pm1\}^4$. For ease of notation, we will drop the index α and just write ρ. As a first step, we point out that due to symmetry the optimal μ^* will have the following properties:

$$\begin{aligned}
\mu^*_{++++} = \mu^*_{----}, \qquad \mu^*_{++--} = \mu^*_{--++},\\
\mu^*_{+-+-} = \mu^*_{-+-+}, \qquad \mu^*_{+---} = \mu^*_{-+++},\\
\mu^*_{+++-} = \mu^*_{++-+} = \mu^*_{+-++} = \mu^*_{-+++} = \mu^*_{---+} = \mu^*_{--+-} = \mu^*_{-+--} = \mu^*_{+---}.
\end{aligned}$$

From the above reparametrization, we additionally emphasize that both

$$\rho_{++} = \rho_{--} \qquad \text{and} \qquad \rho_{+-} = \rho_{-+}$$

hold irrespective of the chosen α. This fact directly entails that setting up the Lagrangian function for our minimization problem will only require two distinct Lagrangian multipliers, namely λ_{++} and λ_{+-}. Put differently, we are going to consider the following Lagrangian function $\mathscr{L}$:

$$\begin{aligned}\mathscr{L}\left(\mu,\lambda_{++},\lambda_{+-}\right) := & g(\mu,\rho_\alpha) - \lambda_{++}\cdot\left(\sum_{(\sigma_1,\tau_1)\in\{(-1,-1),(+1,+1)\}}\left[\sum_{(\sigma_2,\tau_2)\in\{\pm1\}^2}\mu(\sigma_1,\tau_1,\sigma_2,\tau_2)\right]-\rho_\alpha(\sigma_1,\tau_1)\right)\\ & -\lambda_{++}\cdot\left(\sum_{(\sigma_2,\tau_2)\in\{(-1,-1),(+1,+1)\}}\left[\sum_{(\sigma_1,\tau_1)\in\{\pm1\}^2}\mu(\sigma_1,\tau_1,\sigma_2,\tau_2)\right]-\rho_\alpha(\sigma_2,\tau_2)\right)\\ & -\lambda_{+-}\cdot\left(\sum_{(\sigma_1,\tau_1)\in\{(-1,+1),(+1,-1)\}}\left[\sum_{(\sigma_2,\tau_2)\in\{\pm1\}^2}\mu(\sigma_1,\tau_1,\sigma_2,\tau_2)\right]-\rho_\alpha(\sigma_1,\tau_1)\right)\\ & -\lambda_{+-}\cdot\left(\sum_{(\sigma_2,\tau_2)\in\{(-1,+1),(+1,-1)\}}\left[\sum_{(\sigma_1,\tau_1)\in\{\pm1\}^2}\mu(\sigma_1,\tau_1,\sigma_2,\tau_2)\right]-\rho_\alpha(\sigma_2,\tau_2)\right)\end{aligned}$$

Keeping the symmetry in mind, it suffices to consider the following derivatives of the Lagrangian function

$$\begin{aligned}\frac{\partial\mathscr{L}\left(\mu,\lambda_{++},\lambda_{+-}\right)}{\partial\mu_{++++}} &= 1+\log\left(\frac{\mu_{++++}}{\rho_{++}^2}\right)+2\beta-2\lambda_{++}\\ \frac{\partial\mathscr{L}\left(\mu,\lambda_{++},\lambda_{+-}\right)}{\partial\mu_{++--}} &= 1+\log\left(\frac{\mu_{++--}}{\rho_{++}^2}\right)-2\lambda_{++}\\ \frac{\partial\mathscr{L}\left(\mu,\lambda_{++},\lambda_{+-}\right)}{\partial\mu_{+-+-}} &= 1+\log\left(\frac{\mu_{+-+-}}{\rho_{+-}^2}\right)+2\beta-2\lambda_{+-}\\ \frac{\partial\mathscr{L}\left(\mu,\lambda_{++},\lambda_{+-}\right)}{\partial\mu_{+--+}} &= 1+\log\left(\frac{\mu_{+--+}}{\rho_{+-}^2}\right)-2\lambda_{+-}\\ \frac{\partial\mathscr{L}\left(\mu,\lambda_{++},\lambda_{+-}\right)}{\partial\mu_{+++-}} &= 1+\log\left(\frac{\mu_{+++-}}{\rho_{++}\cdot\rho_{+-}}\right)+\beta-\lambda_{++}-\lambda_{+-}\end{aligned}$$

Setting these derivatives equal to zero, we instantly obtain

$$\mu^*_{++++} = \rho^2_{++}\exp\left(2\lambda_{++}-2\beta-1\right) = \rho^2_{++}x_1^2e^{-2\beta} \tag{8.2}$$

$$\mu^*_{++--} = \rho^2_{++}\exp\left(2\lambda_{++}-1\right) = \rho^2_{++}x_1^2 \tag{8.3}$$

$$\mu^*_{+-+-} = \rho^2_{+-}\exp\left(2\lambda_{+-}-2\beta-1\right) = \rho^2_{+-}x_2^2e^{-2\beta} \tag{8.4}$$

$$\mu^*_{+--+} = \rho^2_{+-}\exp\left(2\lambda_{+-}-1\right) = \rho^2_{+-}x_2^2 \tag{8.5}$$

$$\mu^*_{+++-} = \rho_{++}\rho_{+-}\exp\left(\lambda_{++}+\lambda_{+-}-\beta-1\right) = \rho_{++}\rho_{+-}x_1x_2e^{-\beta} \tag{8.6}$$

where the second equalities in each line represent a notational simplification by introducing $x_1 := e^{\lambda_{++}-\frac{1}{2}}$ and $x_2 := e^{\lambda_{+-}-\frac{1}{2}}$.

Lemma 8.1. *The above system of equations (8.2)-(8.6) and therefore the minimization problem of $g\left(\mu,\rho_\alpha\right)$ has the unique solution*

$$x_1 = 2\sqrt{\frac{\left(1+e^{-2\beta}\right)^2+\alpha\left(1-e^{-2\beta}\right)^2-2e^{-\beta}z}{(1+\alpha)^2\left(1+e^{-2\beta}\right)\left(1-e^{-2\beta}\right)^2}}.$$

and

$$x_2 = 2\sqrt{\frac{\left(1+e^{-2\beta}\right)^2-\alpha\left(1-e^{-2\beta}\right)^2\pm2e^{-\beta}z}{(1-\alpha)^2\left(1+e^{-2\beta}\right)\left(1-e^{-2\beta}\right)^2}}.$$

Proof. In order to solve this system of equations, we recall two of the initial constraints of our minimization problem, i.e.

$$\begin{aligned}\mu_{++++}+\mu_{+++-}+\mu_{++-+}+\mu_{++--} &= \rho_{++}\\ \mu_{+-++}+\mu_{+---}+\mu_{+--+}+\mu_{+-+-} &= \rho_{+-}.\end{aligned}$$

Plugging in the μ^* we derived above and once again keeping in mind the symmetry of the problem, the two constraints can be reformulated into

$$\rho_{++}x_1^2e^{-2\beta}+2\rho_{+-}x_1x_2e^{-\beta}+\rho_{++}x_1^2=1$$
$$\rho_{+-}x_2^2e^{-2\beta}+2\rho_{++}x_1x_2e^{-\beta}+\rho_{+-}x_2^2=1.$$

which in turn yields

$$x_1=\frac{1-x_2^2\rho_{+-}\left(1+e^{-2\beta}\right)}{2\rho_{++}e^{-\beta}x_2}.$$

Substituting x_1 into the first constraint, we arrive at

$$\frac{\left(1+e^{-2\beta}\right)\left(1-x_2^2\rho_{+-}\left(1+e^{-2\beta}\right)\right)^2}{4\rho_{++}e^{-2\beta}x_2^2}+\frac{\rho_{+-}}{\rho_{++}}\cdot\left(1-x_2^2\rho_{+-}\left(1+e^{-2\beta}\right)\right)=1$$

For notational convenience, we substitute $x:=x_2^2$ and $\kappa:=1+e^{-2\beta}$. As a consequence, the previous equation can be expressed as

$$x^2\left(\rho_{+-}^2\kappa^3-4\rho_{+-}^2\kappa e^{-2\beta}\right)+x\left(4\rho_{+-}e^{-2\beta}-2\rho_{+-}\kappa^2-4\rho_{++}e^{-2\beta}\right)+\kappa=0.$$

Now, we are able to apply the quadratic formula which yields

$$\begin{aligned}x&=\frac{-4\rho_{+-}e^{-2\beta}+2\rho_{+-}\kappa^2+4\rho_{++}e^{-2\beta}\pm\sqrt{\left(4\rho_{+-}e^{-2\beta}-2\rho_{+-}\kappa^2-4\rho_{++}e^{-2\beta}\right)^2-4\rho_{+-}^2\kappa^2\left(\kappa^2-4e^{-2\beta}\right)}}{2\left(\rho_{+-}^2\kappa^3-4\rho_{+-}^2\kappa e^{-2\beta}\right)}\\&=\frac{\frac{1}{2}(1-\alpha)\left(1+e^{-2\beta}\right)^2+(1+\alpha)e^{-2\beta}-(1-\alpha)e^{-2\beta}\pm\sqrt{\left(2\alpha e^{-2\beta}+\frac{1}{2}(1-\alpha)\kappa^2\right)^2-\frac{1}{4}(1-\alpha)^2\kappa^4+(1-\alpha)^2\kappa^2e^{-2\beta}}}{\frac{1}{8}(1-\alpha)^2\left(1+e^{-2\beta}\right)\left(\left(1+e^{-2\beta}\right)^2-4e^{-2\beta}\right)}\\&=4\frac{\left(1+e^{-2\beta}\right)^2-\alpha\left(1-e^{-2\beta}\right)^2\pm2\sqrt{4\alpha^2e^{-4\beta}+2\alpha e^{-2\beta}(1-\alpha)\left(1+e^{-2\beta}\right)^2+(1-\alpha)^2\left(1+e^{-2\beta}\right)^2e^{-2\beta}}}{(1-\alpha)^2\left(1+e^{-2\beta}\right)\left(1-e^{-2\beta}\right)^2}\end{aligned}$$

Focusing on the square root term, we note

$$\begin{aligned}&4\alpha^2e^{-4\beta}+2\alpha e^{-2\beta}(1-\alpha)\left(1+e^{-2\beta}\right)^2+(1-\alpha)^2\left(1+e^{-2\beta}\right)^2e^{-2\beta}\\&=4\alpha^2e^{-4\beta}-\alpha^2e^{-2\beta}\left(1+e^{-2\beta}\right)^2+\left(1+e^{-2\beta}\right)^2e^{-2\beta}=e^{-2\beta}\left(\left(1+e^{-2\beta}\right)^2-\alpha^2\left(1-e^{-2\beta}\right)^2\right)\end{aligned}$$

which leads us to

$$x=4\frac{\left(1+e^{-2\beta}\right)^2-\alpha\left(1-e^{-2\beta}\right)^2\pm2e^{-\beta}\sqrt{\left(1+e^{-2\beta}\right)^2-\alpha^2\left(1-e^{-2\beta}\right)^2}}{(1-\alpha)^2\left(1+e^{-2\beta}\right)\left(1-e^{-2\beta}\right)^2}=4\frac{\left(1+e^{-2\beta}\right)^2-\alpha\left(1-e^{-2\beta}\right)^2\pm2e^{-\beta}z}{(1-\alpha)^2\left(1+e^{-2\beta}\right)\left(1-e^{-2\beta}\right)^2}$$

where we have introduced $z:=\sqrt{\left(1+e^{-2\beta}\right)^2-\alpha^2\left(1-e^{-2\beta}\right)^2}$ for notational convenience. At this point, we recall that $x=x_2^2$ to arrive at

$$x_2=2\sqrt{\frac{\left(1+e^{-2\beta}\right)^2-\alpha\left(1-e^{-2\beta}\right)^2\pm2e^{-\beta}z}{(1-\alpha)^2\left(1+e^{-2\beta}\right)\left(1-e^{-2\beta}\right)^2}}$$

where we discard the negative square root since x_2 by definition is of the form $x_2=e^{\lambda_{+-}-\frac{1}{2}}$ and thereby always non-negative. This leaves us with two potential solutions for x_2 which only differ in the $\pm$ sign in the above equation. Leaving out the detailed calculation, it is easy to show that choosing + at the $\pm$ sign would result in a negative x_1. However, similar to x_2, $x_1=e^{\lambda_{++}-\frac{1}{2}}$ also cannot become negative by construction. As a consequence, the only remaining and suitable candidate for x_2 and thereby the solution is

$$x_2=2\sqrt{\frac{\left(1+e^{-2\beta}\right)^2-\alpha\left(1-e^{-2\beta}\right)^2-2e^{-\beta}z}{(1-\alpha)^2\left(1+e^{-2\beta}\right)\left(1-e^{-2\beta}\right)^2}}.$$

With this solution for x_2 we are now able to calculate the optimal x_1. More specifically, we recall the formula we have derived a few steps back

$$x_1 = \frac{1 - x_2^2 \rho_{+-} \left(1 + e^{-2\beta}\right)}{2\rho_{++} e^{-\beta} x_2}.$$

Plugging in the optimal x_2, we arrive at the expression

$$x_1 = \frac{1 - \frac{\left(1+e^{-2\beta}\right)^2 - \alpha\left(1-e^{-2\beta}\right)^2 - 2e^{-\beta}z}{(1-\alpha)\left(1-e^{-2\beta}\right)^2}}{(1+\alpha)\, e^{-\beta}\sqrt{\frac{\left(1+e^{-2\beta}\right)^2 - \alpha\left(1-e^{-2\beta}\right)^2 - 2e^{-\beta}z}{(1-\alpha)^2\left(1+e^{-2\beta}\right)\left(1-e^{-2\beta}\right)^2}}} = \frac{\left(-4e^{-\beta} + 2z\right)\sqrt{1+e^{-2\beta}}}{(1+\alpha)\left(1-e^{-2\beta}\right)\sqrt{\left(1+e^{-2\beta}\right)^2 - \alpha\left(1-e^{-2\beta}\right)^2 - 2e^{-\beta}z}}. \tag{8.7}$$

Next, we claim that

$$x_1 = 2\sqrt{\frac{\left(1+e^{-2\beta}\right)^2 + \alpha\left(1-e^{-2\beta}\right)^2 - 2e^{-\beta}z}{(1+\alpha)^2\left(1+e^{-2\beta}\right)\left(1-e^{-2\beta}\right)^2}}.$$

Indeed, we find starting at (8.7) that

$$x_1 = \frac{\left(-4e^{-\beta} + 2z\right)\sqrt{1+e^{-2\beta}}}{(1+\alpha)\left(1-e^{-2\beta}\right)\sqrt{\left(1+e^{-2\beta}\right)^2 - \alpha\left(1-e^{-2\beta}\right)^2 - 2e^{-\beta}z}}.$$

Thus, our claim is equivalent to

$$\left(z - 2e^{-\beta}\right)^2\left(1+e^{-2\beta}\right)^2 = \left(\left(1+e^{-2\beta}\right)^2 + \alpha\left(1-e^{-2\beta}\right)^2 - 2e^{-\beta}z\right)\left(\left(1+e^{-2\beta}\right)^2 - \alpha\left(1-e^{-2\beta}\right)^2 - 2e^{-\beta}z\right). \tag{8.8}$$

To see that (8.8) is indeed true, we execute the following auxiliary calculation:

$$\begin{aligned}
&\left(\left(1+e^{-2\beta}\right)^2 + \alpha\left(1-e^{-2\beta}\right)^2 - 2e^{-\beta}z\right)\left(\left(1+e^{-2\beta}\right)^2 - \alpha\left(1-e^{-2\beta}\right)^2 - 2e^{-\beta}z\right) \\
&= \left(1+e^{-2\beta}\right)^4 - 4e^{-\beta}z\left(1+e^{-2\beta}\right)^2 + 4e^{-2\beta}\left(\left(1+e^{-2\beta}\right)^2 - \alpha^2\left(1-e^{-2\beta}\right)^2\right) - \alpha^2\left(1-e^{-2\beta}\right)^4 \\
&= \left(1+e^{-2\beta}\right)^2\left(z^2 - 4e^{-\beta}z + 4e^{-2\beta}\right) = \left(1+e^{-2\beta}\right)^2\left(z - 2e^{-\beta}\right)^2.
\end{aligned}$$

Hence, we established our claim and thus know

$$x_1 = 2\sqrt{\frac{\left(1+e^{-2\beta}\right)^2 + \alpha\left(1-e^{-2\beta}\right)^2 - 2e^{-\beta}z}{(1+\alpha)^2\left(1+e^{-2\beta}\right)\left(1-e^{-2\beta}\right)^2}}.$$

□

Let us bring together our findings of this subsection. Due to the well-known fact that the Kullback-Leibler divergence is convex in its input parameters, we immediately see that the function $g(\mu, \rho_\alpha)$ is convex as well. As a consequence, the μ^* we have just calculated is indeed the minimum. Put differently, we are now able to state

$$\min_{\mu \in \mathscr{O}'} g(\mu, \rho_\alpha) = g(\mu^*, \rho_\alpha).$$

Although this statement is satisfactory, we would favor a more explicit expression. This is achieved by the following Lemma.

Lemma 8.2. *We have*

$$\begin{aligned}
g(\mu^*, \rho_\alpha) = {} & 2\log(2) - \log\left(\left(1+e^{-2\beta}\right)\left(1-e^{-2\beta}\right)^2\right) - (1+\alpha)\log(1+\alpha) \\
& - (1-\alpha)\log(1-\alpha) + \frac{1+\alpha}{2}\log\left(\left(1+e^{-2\beta}\right)^2 + \alpha\left(1-e^{-2\beta}\right)^2 - 2e^{-\beta}z\right) \\
& + \frac{1-\alpha}{2}\log\left(\left(1+e^{-2\beta}\right)^2 - \alpha\left(1-e^{-2\beta}\right)^2 - 2e^{-\beta}z\right).
\end{aligned}$$

Proof. In order to get to the desired expression, we take a closer look at the Kullback-Leibler divergence for the optimal μ^*

$$\begin{aligned}D_{\mathrm{KL}}(\mu^*||\rho_\alpha\otimes\rho_\alpha) &= 2\mu^*_{++++}\log\left(x_1^2e^{-2\beta}\right)+2\mu^*_{++--}\log\left(x_1^2\right)+2\mu^*_{+-+-}\log\left(x_2^2e^{-2\beta}\right)\\ &\quad+2\mu^*_{+--+}\log\left(x_1^2\right)+8\mu^*_{+++-}\log\left(x_1x_2e^{-\beta}\right)\\ &= -\beta\left(4\mu^*_{++++}+4\mu^*_{+-+-}+8\mu^*_{+++-}\right)+\log(x_1)\left(4\mu^*_{++++}+4\mu^*_{++--}+8\mu^*_{+++-}\right)\\ &\quad+\log(x_2)\left(4\mu^*_{+-+-}+4\mu^*_{+--+}+8\mu^*_{+++-}\right)\end{aligned}$$

Using the reformulation of $D_{\mathrm{KL}}(\mu^*||\rho_\alpha\otimes\rho_\alpha)$, $g(\mu^*,\rho_\alpha)$ can be formulated as

$$g(\mu^*,\rho_\alpha)=\log(x_1)\left(4\mu^*_{++++}+4\mu^*_{++--}+8\mu^*_{+++-}\right)+\log(x_2)\left(4\mu^*_{+-+-}+4\mu^*_{+--+}+8\mu^*_{+++-}\right).$$

This expression in turn is suitable for inserting x_1 and x_2 leading to

$$\begin{aligned}g(\mu^*,\rho_\alpha) &= 2\log(2)-\log\left(\left(1+e^{-2\beta}\right)\left(1-e^{-2\beta}\right)^2\right)\\ &\quad+\left[\log\left(\left(1+e^{-2\beta}\right)^2+\alpha\left(1-e^{-2\beta}\right)^2-2e^{-\beta}z\right)-2\log(1+\alpha)\right]\cdot\left(2\mu^*_{++++}+2\mu^*_{++--}+4\mu^*_{+++-}\right)\\ &\quad+\left[\log\left(\left(1+e^{-2\beta}\right)^2-\alpha\left(1-e^{-2\beta}\right)^2-2e^{-\beta}z\right)-2\log(1-\alpha)\right]\cdot\left(2\mu^*_{+-+-}+2\mu^*_{+--+}+4\mu^*_{+++-}\right).\end{aligned}$$

Since μ^* is a probability measure by definition, we can exploit the identity

$$2\mu^*_{+-+-}+2\mu^*_{+--+}+4\mu^*_{+++-}=1-2\mu^*_{++++}-2\mu^*_{++--}-4\mu^*_{+++-}$$

to rearrange $g(\mu^*,\rho_\alpha)$ as

$$\begin{aligned}g(\mu^*,\rho_\alpha) &= 2\log(2)-\log\left(\left(1+e^{-2\beta}\right)\left(1-e^{-2\beta}\right)^2\right)+\log\left(\left(1+e^{-2\beta}\right)^2-\alpha\left(1-e^{-2\beta}\right)^2-2e^{-\beta}z\right)-2\log(1-\alpha)\\ &\quad+\left[\log\left(\frac{\left(1+e^{-2\beta}\right)^2+\alpha\left(1-e^{-2\beta}\right)^2-2e^{-\beta}z}{\left(1+e^{-2\beta}\right)^2-\alpha\left(1-e^{-2\beta}\right)^2-2e^{-\beta}z}\right)-2\log\left(\frac{1+\alpha}{1-\alpha}\right)\right]\cdot\left(2\mu^*_{++++}+2\mu^*_{++--}+4\mu^*_{+++-}\right).\end{aligned}$$

To keep the terms relatively brief, we define

$$\begin{aligned}T_1 &:= 2\log(2)-\log\left(\left(1+e^{-2\beta}\right)\left(1-e^{-2\beta}\right)^2\right)+\log\left(\left(1+e^{-2\beta}\right)^2-\alpha\left(1-e^{-2\beta}\right)^2-2e^{-\beta}z\right)-2\log(1-\alpha)\\ T_2 &:= 2\mu^*_{++++}+2\mu^*_{++--}+4\mu^*_{+++-}\\ T_3 &:= \log\left(\frac{\left(1+e^{-2\beta}\right)^2+\alpha\left(1-e^{-2\beta}\right)^2-2e^{-\beta}z}{\left(1+e^{-2\beta}\right)^2-\alpha\left(1-e^{-2\beta}\right)^2-2e^{-\beta}z}\right)-2\log\left(\frac{1+\alpha}{1-\alpha}\right)\end{aligned}$$

which implies $g(\mu^*,\rho_\alpha)=T_1+T_2\cdot T_3$. In the next step, we will plug in μ^* in order to simplify T_2

$$\begin{aligned}T_2 &= 2\mu^*_{++++}+2\mu^*_{++--}+4\mu^*_{+++-}=2x_1^2\rho_{++}^2\left(1+e^{-2\beta}+2e^{-\beta}\frac{\rho_{+-}x_2}{\rho_{++}x_1}\right)\\ &= \frac{\left(1+e^{-2\beta}\right)^2+\alpha\left(1-e^{-2\beta}\right)^2-2e^{-\beta}z}{2\left(1-e^{-2\beta}\right)^2}\\ &\quad+e^{-\beta}\frac{\sqrt{\left(\left(1+e^{-2\beta}\right)^2+\alpha\left(1-e^{-2\beta}\right)^2-2e^{-\beta}z\right)\left(\left(1+e^{-2\beta}\right)^2-\alpha\left(1-e^{-2\beta}\right)^2-2e^{-\beta}z\right)}}{\left(1+e^{-2\beta}\right)\left(1-e^{-2\beta}\right)^2}\end{aligned}$$

Applying (8.8) to the term in the square root yields

$$\begin{aligned}T_2 &= \frac{\left(1+e^{-2\beta}\right)^2+\alpha\left(1-e^{-2\beta}\right)^2-2e^{-\beta}z}{2\left(1-e^{-2\beta}\right)^2}+e^{-\beta}\frac{\left(z-2e^{-\beta}\right)\left(1+e^{-2\beta}\right)}{\left(1+e^{-2\beta}\right)\left(1-e^{-2\beta}\right)^2}\\ &= \frac{\left(1+e^{-2\beta}\right)^2+\alpha\left(1-e^{-2\beta}\right)^2-4e^{-2\beta}}{2\left(1-e^{-2\beta}\right)^2}=\frac{1+\alpha}{2}.\end{aligned}$$

Coming back to $g(\mu^*, \rho_\alpha)$, we obtain the expression that Lemma 8.2 promised

$$g(\mu^*, \rho_\alpha) = 2\log(2) - \log\left(\left(1+e^{-2\beta}\right)\left(1-e^{-2\beta}\right)^2\right) - (1+\alpha)\log(1+\alpha) - (1-\alpha)\log(1-\alpha)$$
$$+ \frac{1+\alpha}{2}\log\left(\left(1+e^{-2\beta}\right)^2 + \alpha\left(1-e^{-2\beta}\right)^2 - 2e^{-\beta}z\right) + \frac{1-\alpha}{2}\log\left(\left(1+e^{-2\beta}\right)^2 - \alpha\left(1-e^{-2\beta}\right)^2 - 2e^{-\beta}z\right).$$

□

8.2. **Maximization with respect to α.** In this subsection we focus on the function

$$f_d(\alpha,\beta) := \log(2) + \mathrm{H}\left(\frac{1+\alpha}{2}\right) - \frac{d}{2}g(\mu^*, \rho_\alpha).$$

which results from plugging in the definition of ρ in terms of α from (8.1). More specifically, we are interested in solving the optimization

$$\max_{-1<\alpha<1} f_d(\alpha,\beta).$$

which will immediately yield the answer to our initial optimization problem over $\delta(\mu,\rho)$. Note that we tacitly exploit the results of both Lemma 8.1 and Lemma 8.2 to be able to state a function $f_d(\alpha,\beta)$ that only depends on d, α, and β. As a consequence, we have to prove the following statement.

Lemma 8.3. *Assume that* $0 < \beta < \beta_{KS}$. *Then we have*

$$\arg\max_{-1<\alpha<1} f_d(\alpha,\beta) = 0.$$

Proof. To solve the maximization with respect to α, we calculate the derivatives. Let us start with the simpler ones, namely the first and second derivative of the entropy with respect to α:

$$\frac{\partial \mathrm{H}\left(\frac{1+\alpha}{2}\right)}{\partial\alpha} = \frac{1}{2}\log(1-\alpha) - \frac{1}{2}\log(1+\alpha)$$

and

$$\frac{\partial^2 \mathrm{H}\left(\frac{1+\alpha}{2}\right)}{\partial\alpha^2} = \frac{1}{2}\left(\frac{-1}{1-\alpha} - \frac{1}{1+\alpha}\right) = -\frac{1}{1-\alpha^2}.$$

Before we continue with our main task, let us state a useful observation which will be helpful in the following calculations. Let

$$z = \sqrt{\left(1+e^{-2\beta}\right)^2 - \alpha^2\left(1-e^{-2\beta}\right)^2}$$

Then, we have

$$\frac{\partial z}{\partial\alpha} = \frac{-\alpha\left(1-e^{-2\beta}\right)^2}{\sqrt{\left(1+e^{-2\beta}\right)^2 - \alpha^2\left(1-e^{-2\beta}\right)^2}} = -\alpha\left(1-e^{-2\beta}\right)^2 z^{-1}.$$

Next, we determine the first two derivatives for $g(\mu^*, \rho_\alpha)$. Starting with the first derivative, we find

$$\frac{\partial g(\mu^*,\rho_\alpha)}{\partial\alpha} = -\frac{1+\alpha}{1+\alpha} - \log\left(\frac{1+\alpha}{2}\right) + \frac{1-\alpha}{1-\alpha} + \log\left(\frac{1-\alpha}{2}\right)$$
$$+ \frac{1}{2}\left[\log\left(\left(1+e^{-2\beta}\right)^2 + \alpha\left(1-e^{-2\beta}\right)^2 - 2e^{-\beta}z\right) - \log\left(\left(1+e^{-2\beta}\right)^2 - \alpha\left(1-e^{-2\beta}\right)^2 - 2e^{-\beta}z\right)\right]$$
$$+ \frac{1+\alpha}{2}\cdot\frac{\left(1-e^{-2\beta}\right)^2 + 2e^{-\beta}\alpha\left(1-e^{-2\beta}\right)^2 z^{-1}}{\left(1+e^{-2\beta}\right)^2 + \alpha\left(1-e^{-2\beta}\right)^2 - 2e^{-\beta}z} + \frac{1-\alpha}{2}\cdot\frac{-\left(1-e^{-2\beta}\right)^2 + 2e^{-\beta}\alpha\left(1-e^{-2\beta}\right)^2 z^{-1}}{\left(1+e^{-2\beta}\right)^2 - \alpha\left(1-e^{-2\beta}\right)^2 - 2e^{-\beta}z}$$

For the next simplification, we focus on the last two summands of the previously stated derivative, i.e.

$$\frac{1+\alpha}{2}\cdot\frac{\left(1-e^{-2\beta}\right)^2 + 2e^{-\beta}\alpha\left(1-e^{-2\beta}\right)^2 z^{-1}}{\left(1+e^{-2\beta}\right)^2 + \alpha\left(1-e^{-2\beta}\right)^2 - 2e^{-\beta}z} + \frac{1-\alpha}{2}\cdot\frac{-\left(1-e^{-2\beta}\right)^2 + 2e^{-\beta}\alpha\left(1-e^{-2\beta}\right)^2 z^{-1}}{\left(1+e^{-2\beta}\right)^2 - \alpha\left(1-e^{-2\beta}\right)^2 - 2e^{-\beta}z}$$
$$= \frac{\left(1-e^{-2\beta}\right)^2}{2z}\cdot\frac{(1+\alpha)\left(z+2e^{-\beta}\alpha\right)\left(\left(1+e^{-2\beta}\right)^2 - \alpha\left(1-e^{-2\beta}\right)^2 - 2e^{-\beta}z\right)}{\left(\left(1+e^{-2\beta}\right)^2 + \alpha\left(1-e^{-2\beta}\right)^2 - 2e^{-\beta}z\right)\left(\left(1+e^{-2\beta}\right)^2 - \alpha\left(1-e^{-2\beta}\right)^2 - 2e^{-\beta}z\right)}$$
$$+ \frac{\left(1-e^{-2\beta}\right)^2}{2z}\cdot\frac{(1-\alpha)\left(-z+2e^{-\beta}\alpha\right)\left(\left(1+e^{-2\beta}\right)^2 + \alpha\left(1-e^{-2\beta}\right)^2 - 2e^{-\beta}z\right)}{\left(\left(1+e^{-2\beta}\right)^2 + \alpha\left(1-e^{-2\beta}\right)^2 - 2e^{-\beta}z\right)\left(\left(1+e^{-2\beta}\right)^2 - \alpha\left(1-e^{-2\beta}\right)^2 - 2e^{-\beta}z\right)}$$

Again, we restrict our attention to one term, namely

$$\begin{aligned}
&(1+\alpha)\left(z+2e^{-\beta}\alpha\right)\left(\left(1+e^{-2\beta}\right)^2-\alpha\left(1-e^{-2\beta}\right)^2-2e^{-\beta}z\right)\\
&+(1-\alpha)\left(-z+2e^{-\beta}\alpha\right)\left(\left(1+e^{-2\beta}\right)^2+\alpha\left(1-e^{-2\beta}\right)^2-2e^{-\beta}z\right)\\
&=\left(\left(1+e^{-2\beta}\right)^2-2e^{-\beta}z\right)\left(4e^{-\beta}\alpha+2\alpha z\right)+\alpha\left(1-e^{-2\beta}\right)^2\left(-2z-4e^{-\beta}\alpha^2\right)\\
&=2\alpha z4e^{-2\beta}-8e^{-2\beta}\alpha z=0.
\end{aligned}$$

As a result, the first derivative can be reduced to

$$\frac{\partial g(\mu^*,\rho_\alpha)}{\partial\alpha}=-\log(1+\alpha)+\log(1-\alpha)+\frac{1}{2}\log\left(\frac{\left(1+e^{-2\beta}\right)^2+\alpha\left(1-e^{-2\beta}\right)^2-2e^{-\beta}z}{\left(1+e^{-2\beta}\right)^2-\alpha\left(1-e^{-2\beta}\right)^2-2e^{-\beta}z}\right).$$

Based on this result, we can instantly compute the second derivative

$$\frac{\partial^2 g(\mu^*,\rho_\alpha)}{\partial\alpha^2}=-\frac{2}{1-\alpha^2}+\frac{1}{2}\cdot\left[\frac{\left(1-e^{-2\beta}\right)^2+2e^{-\beta}\alpha\left(1-e^{-2\beta}\right)^2z^{-1}}{\left(1+e^{-2\beta}\right)^2+\alpha\left(1-e^{-2\beta}\right)^2-2e^{-\beta}z}-\frac{-\left(1-e^{-2\beta}\right)^2+2e^{-\beta}\alpha\left(1-e^{-2\beta}\right)^2z^{-1}}{\left(1+e^{-2\beta}\right)^2-\alpha\left(1-e^{-2\beta}\right)^2-2e^{-\beta}z}\right].$$

Once again, we apply (8.8) to get to

$$\begin{aligned}
\frac{\partial^2 g(\mu^*,\rho_\alpha)}{\partial\alpha^2}&=-\frac{2}{1-\alpha^2}+\left(1-e^{-2\beta}\right)^2\frac{2z\left(\left(1+e^{-2\beta}\right)^2-2e^{-\beta}z\right)-4e^{-\beta}\alpha^2\left(1-e^{-2\beta}\right)^2}{2z\left(z-2e^{-\beta}\right)^2\left(1+e^{-2\beta}\right)^2}\\
&=-\frac{2}{1-\alpha^2}+\left(1-e^{-2\beta}\right)^2\frac{\left(1+e^{-2\beta}\right)^2\left(z-2e^{-\beta}\right)}{z\left(z-2e^{-\beta}\right)^2\left(1+e^{-2\beta}\right)^2}=\frac{\left(1-e^{-2\beta}\right)^2}{z\left(z-2e^{-\beta}\right)}-\frac{2}{1-\alpha^2}.
\end{aligned}$$

Finally, combining the derivatives of the entropy and $g(\mu^*,\rho_a)$ we arrive at

$$\begin{aligned}
&\frac{\partial f_d(\alpha,\beta)}{\partial\alpha}=\frac{\partial \mathrm{H}\left(\frac{1+\alpha}{2}\right)}{\partial\alpha}-\frac{d}{2}\cdot\frac{\partial g(\mu^*,\rho_\alpha)}{\partial\alpha}\\
&=\frac{1}{2}\log(1-\alpha)-\frac{1}{2}\log(1+\alpha)\\
&\quad+\frac{d}{2}\left(\log(1+\alpha)-\log(1-\alpha)-\frac{1}{2}\log\left(\frac{\left(1+e^{-2\beta}\right)^2+\alpha\left(1-e^{-2\beta}\right)^2-2e^{-\beta}z}{\left(1+e^{-2\beta}\right)^2-\alpha\left(1-e^{-2\beta}\right)^2-2e^{-\beta}z}\right)\right)\\
&=\frac{d-1}{2}\log(1+\alpha)-\frac{d-1}{2}\log(1-\alpha)-\frac{d}{4}\log\left(\frac{\left(1+e^{-2\beta}\right)^2+\alpha\left(1-e^{-2\beta}\right)^2-2e^{-\beta}z}{\left(1+e^{-2\beta}\right)^2-\alpha\left(1-e^{-2\beta}\right)^2-2e^{-\beta}z}\right)
\end{aligned}$$

and

$$\begin{aligned}
\frac{\partial^2 f_d(\alpha,\beta)}{\partial\alpha^2}&=\frac{\partial^2 \mathrm{H}\left(\frac{1+\alpha}{2}\right)}{\partial\alpha^2}-\frac{d}{2}\cdot\frac{\partial^2 g(\mu^*,\rho_\alpha)}{\partial^2\alpha}=-\frac{1}{1-\alpha^2}-\frac{d}{2}\cdot\left(\frac{\left(1-e^{-2\beta}\right)^2}{z\left(z-2e^{-\beta}\right)}-\frac{2}{1-\alpha^2}\right)\\
&=d\cdot\frac{z^2-4ze^{-\beta}+\left(1+e^{-2\beta}\right)^2-\alpha^2\left(1-e^{-2\beta}\right)^2-\left(1-e^{-2\beta}\right)^2+\alpha^2\left(1-e^{-2\beta}\right)^2}{\left(1-\alpha^2\right)2z\left(z-2e^{-\beta}\right)}-\frac{1}{1-\alpha^2}\\
&=\frac{d-2}{2\left(1-\alpha^2\right)}-\frac{de^{-\beta}}{\left(1-\alpha^2\right)z}=\frac{d-2}{2\left(1-\alpha^2\right)}-\frac{d}{\left(1-\alpha^2\right)\sqrt{\left(e^{\beta}+e^{-\beta}\right)^2-\alpha^2\left(e^{\beta}-e^{-\beta}\right)^2}}\\
&=\frac{d-2}{2\left(1-\alpha^2\right)}-\frac{d}{\left(1-\alpha^2\right)\sqrt{\left(1-\alpha^2\right)\cdot\left(e^{2\beta}+e^{-2\beta}\right)+2+2\alpha^2}}.
\end{aligned}$$

Furthermore, we note that for every β we have for $\alpha=0$

$$\frac{\partial f_d}{\partial\alpha}(0,\beta)=0.$$

Now, to complete the maximization with respect to α, we claim is that the global maximum of $f_d(\alpha,\beta)$ is at $\alpha=0$ as long as $\beta<\beta_{\mathrm{KS}}$. We prove this claim in two steps. First, we show that $\frac{\partial^2 f_d}{\partial\alpha^2}(\alpha,\beta)$ is increasing in β. Subsequently, we establish that $\frac{\partial^2 f_d}{\partial\alpha^2}(\alpha,\beta^*)$ is smaller than zero for all $\alpha\in(-1,1)$. As a consequence, $\frac{\partial^2 f_d}{\partial\alpha^2}(\alpha,\beta)<0$ holds for all $\beta\in(0,\beta^*)$ and $\alpha\in(-1,1)$ and thereby implies that the maximum of $f_d(\alpha,\beta)$ is

attained at $\alpha = 0$ for $\beta < \beta^*$. The previously performed technical rearrangements are helpful for calculating the next derivative in a straightforward manner.

$$\begin{aligned}\frac{\partial}{\partial\beta}\left(\frac{\partial^2 f_d}{\partial\alpha^2}\right)(\alpha,\beta) &= \frac{d}{2\left(1-\alpha^2\right)}\frac{\left(1-\alpha^2\right)\cdot\left(2\beta e^{2\beta-1}-2\beta e^{-2\beta-1}\right)}{\left[\left(1-\alpha^2\right)\cdot\left(e^{2\beta}+e^{-2\beta}\right)+2+2\alpha^2\right]^{\frac{3}{2}}} \\ &= \underbrace{\frac{d}{2}2\beta e^{2\beta-1}\left(1-e^{-4\beta}\right)}_{>0}\underbrace{\left[\underbrace{\left(1-\alpha^2\right)}_{>0}\cdot\left(e^{2\beta}+e^{-2\beta}\right)+2+2\alpha^2\right]^{-\frac{3}{2}}}_{>0} > 0\end{aligned}$$

where we restrict our attention to $-1 < \alpha < 1$. All that remains to do is to plug in the Kesten-Stigum bound into the second derivative with respect to alpha which yields

$$\begin{aligned}\frac{\partial^2 f_d}{\partial\alpha^2}(\alpha,\beta^*) &= \frac{d-2}{2\left(1-\alpha^2\right)} - \frac{d}{1-\alpha^2}\cdot\left[\left(1-\alpha^2\right)\cdot\left(\frac{\left(\sqrt{d-1}+1\right)^2}{\left(\sqrt{d-1}-1\right)^2}+\frac{\left(\sqrt{d-1}-1\right)^2}{\left(\sqrt{d-1}+1\right)^2}\right)+2+2\alpha^2\right]^{-\frac{1}{2}} \\ &= \frac{d-2}{2\left(1-\alpha^2\right)} - \frac{d}{1-\alpha^2}\cdot\left[\left(1-\alpha^2\right)\cdot\left(\frac{\left(\sqrt{d-1}+1\right)^4+\left(\sqrt{d-1}-1\right)^4}{(d-1-1)^2}\right)+2+2\alpha^2\right]^{-\frac{1}{2}} \\ &= \frac{d-2}{2\left(1-\alpha^2\right)} - \frac{d}{1-\alpha^2}\cdot\left[\frac{4d^2+\alpha^2\left(16-16d\right)}{(d-2)^2}\right]^{-\frac{1}{2}} \\ &= \frac{d-2}{2\left(1-\alpha^2\right)}\cdot\underbrace{\left(1-\frac{d}{\sqrt{d^2-4\alpha^2\left(d-1\right)}}\right)}_{<0} < 0\end{aligned}$$

where we assume both $d > 2$ and $-1 < \alpha < 1$. This concludes the maximization problem. □

What remains is to bring all the findings of this section together.

Proof of Lemma 7.4. Substituting $\alpha = 0$ from Lemma 8.3 into the previous reformulations, we can state that $\delta(\mu,\rho)$ obtains its optimum at μ^* where

$$\begin{aligned}\mu^*_{++++} = \mu^*_{----} = \mu^*_{+-+-} = \mu^*_{-+-+} &= \frac{e^{-2\beta}}{4\left(1+e^{-\beta}\right)^2} \\ \mu^*_{+--+} = \mu^*_{--++} = \mu^*_{-++-} = \mu^*_{++--} &= \frac{1}{4\left(1+e^{-\beta}\right)^2} \\ \mu^*_{+++-} = \mu^*_{++-+} = \mu^*_{+-++} = \mu^*_{-+++} = \mu^*_{---+} = \mu^*_{--+-} = \mu^*_{-+--} = \mu^*_{+---} &= \frac{e^{-\beta}}{4\left(1+e^{-\beta}\right)^2}\end{aligned}$$

which also implies

$$\rho^*_{++} = \rho^*_{+-} = \rho^*_{-+} = \rho^*_{--} = \frac{1}{4}$$

and

$$\begin{aligned}&\delta(\mu^*,\rho^*) = \mathrm{H}(\rho^*) - \frac{d}{2}\left(D_{\mathrm{KL}}(\mu^*||\rho^*\otimes\rho^*) + \beta\sum_{\sigma\in A_1}\mu^*(\sigma) + 2\beta\sum_{\sigma\in A_2}\mu^*(\sigma)\right) \\ &= (2-2d)\log(2) + d\log\left(2\left(1+e^{-\beta}\right)\right) + \frac{d}{2}\left(-\frac{e^{-2\beta}}{\left(1+e^{-\beta}\right)^2}\log\left(e^{-2\beta}\right) - \frac{2e^{-\beta}}{\left(1+e^{-\beta}\right)^2}\log\left(e^{-\beta}\right)\right) - d\beta\frac{e^{-2\beta}+e^{-\beta}}{\left(1+e^{-\beta}\right)^2} \\ &= (2-d)\log(2) + d\log\left(1+e^{-\beta}\right) + d\beta\left(\frac{e^{-2\beta}}{\left(1+e^{-\beta}\right)^2}+\frac{e^{-\beta}}{\left(1+e^{-\beta}\right)^2}\right) - d\beta\frac{e^{-2\beta}+e^{-\beta}}{\left(1+e^{-\beta}\right)^2} \\ &= (2-d)\log(2) + d\log\left(1+e^{-\beta}\right).\end{aligned}$$

Lemma 7.4 readily follows. □

9. The Hessian for the second moment / Proof of Lemma 7.5

The proof of Lemma 7.5 boils down to tedious calculations of the first and second partial derivatives. As a starting point we reformulate $\delta(\mu,\rho)$ with the restricted number of variables.

$$\begin{aligned}\delta(\mu,\rho) &= \mathrm{H}(\rho) - \frac{d}{2}\left(D_{\mathrm{KL}}(\mu||\rho\otimes\rho) + \beta\sum_{\sigma\in A_1}\mu(\sigma) + 2\beta\sum_{\sigma\in A_2}\mu(\sigma)\right)\\ &= (1-d)\,\mathrm{H}(\rho) + \frac{d}{2}\mathrm{H}(\mu) - d\beta\left(x_3+x_4+x_5+x_6+x_7+x_8+x_9+\mu_{----}\right)\end{aligned}$$

Now, let us turn to the first derivatives of $\mathrm{H}(\rho)$

$$\begin{aligned}\frac{\partial\mathrm{H}(\rho)}{\partial x_1} &= -\log(\rho_{+-}) - 1 - \log(\rho_{-+}) - 1 + 2\log(\rho_{--}) + 2\\ \frac{\partial\mathrm{H}(\rho)}{\partial x_2} &= -\log(\rho_{++}) - 1 + \log(\rho_{--}) + 1\\ \frac{\partial\mathrm{H}(\rho)}{\partial x_3} &= -\log(\rho_{++}) - 1 - \log(\rho_{+-}) - 1 + 2\log(\rho_{--}) + 2\\ \frac{\partial\mathrm{H}(\rho)}{\partial x_4} &= -\log(\rho_{++}) - 1 - \log(\rho_{-+}) - 1 + 2\log(\rho_{--}) + 2\\ \frac{\partial\mathrm{H}(\rho)}{\partial x_5} &= -\log(\rho_{+-}) - 1 + \log(\rho_{--}) + 1\\ \frac{\partial\mathrm{H}(\rho)}{\partial x_6} &= -\log(\rho_{-+}) - 1 + \log(\rho_{--}) + 1\\ \frac{\partial\mathrm{H}(\rho)}{\partial x_7} &= -\log(\rho_{+-}) - 1 + \log(\rho_{--}) + 1\\ \frac{\partial\mathrm{H}(\rho)}{\partial x_8} &= -\log(\rho_{-+}) - 1 + \log(\rho_{--}) + 1\\ \frac{\partial\mathrm{H}(\rho)}{\partial x_9} &= -\log(\rho_{++}) - 1 + \log(\rho_{--}) + 1\end{aligned}$$

and the first derivatives of $\mathrm{H}(\mu)$

$$\begin{aligned}\frac{\partial\mathrm{H}(\mu)}{\partial x_1} &= -2\log(x_1) - 2 + 2\log(\mu_{----}) + 2\\ \frac{\partial\mathrm{H}(\mu)}{\partial x_2} &= -2\log(x_2) - 2 + 2\log(\mu_{----}) + 2\\ \frac{\partial\mathrm{H}(\mu)}{\partial x_3} &= -2\log(x_3) - 2 + 2\log(\mu_{----}) + 2\\ \frac{\partial\mathrm{H}(\mu)}{\partial x_4} &= -2\log(x_4) - 2 + 2\log(\mu_{----}) + 2\end{aligned}$$

$$\begin{aligned}\frac{\partial\mathrm{H}(\mu)}{\partial x_5} &= -2\log(x_5) - 2 + 2\log(\mu_{----}) + 2\\ \frac{\partial\mathrm{H}(\mu)}{\partial x_6} &= -2\log(x_6) - 2 + 2\log(\mu_{----}) + 2\\ \frac{\partial\mathrm{H}(\mu)}{\partial x_7} &= -\log(x_7) - 1 + \log(\mu_{----}) + 1\\ \frac{\partial\mathrm{H}(\mu)}{\partial x_8} &= -\log(x_8) - 1 + \log(\mu_{----}) + 1\\ \frac{\partial\mathrm{H}(\mu)}{\partial x_9} &= -\log(x_9) - 1 + \log(\mu_{----}) + 1.\end{aligned}$$

For the second derivatives we obtain

$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_1^2} = (1-d)\left(-\frac{1}{\rho_{+-}}-\frac{1}{\rho_{-+}}-\frac{4}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{2}{x_1}-\frac{4}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_2\partial x_1} = (1-d)\left(-\frac{2}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{4}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_3\partial x_1} = (1-d)\left(-\frac{1}{\rho_{+-}}-\frac{4}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{4}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_4\partial x_1} = (1-d)\left(-\frac{1}{\rho_{-+}}-\frac{4}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{4}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_5\partial x_1} = (1-d)\left(-\frac{1}{\rho_{+-}}-\frac{2}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{4}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_6\partial x_1} = (1-d)\left(-\frac{1}{\rho_{-+}}-\frac{2}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{4}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_7\partial x_1} = (1-d)\left(-\frac{1}{\rho_{+-}}-\frac{2}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{2}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_8\partial x_1} = (1-d)\left(-\frac{1}{\rho_{-+}}-\frac{2}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{2}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_9\partial x_1} = (1-d)\left(-\frac{2}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{2}{\mu_{----}}\right)$$

and

$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_2^2} = (1-d)\left(-\frac{1}{\rho_{++}}-\frac{1}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{2}{x_2}-\frac{4}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_3\partial x_2} = (1-d)\left(-\frac{1}{\rho_{++}}-\frac{2}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{4}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_4\partial x_2} = (1-d)\left(-\frac{1}{\rho_{++}}-\frac{2}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{4}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_5\partial x_2} = (1-d)\left(-\frac{1}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{4}{\mu_{----}}\right)$$

$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_6\partial x_2} = (1-d)\left(-\frac{1}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{4}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_7\partial x_2} = (1-d)\left(-\frac{1}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{2}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_8\partial x_2} = (1-d)\left(-\frac{1}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{2}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_9\partial x_2} = (1-d)\left(-\frac{1}{\rho_{++}}-\frac{1}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{2}{\mu_{----}}\right).$$

We continue with

$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_3^2} = (1-d)\left(-\frac{1}{\rho_{++}}-\frac{1}{\rho_{+-}}-\frac{4}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{2}{x_3}-\frac{4}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_4\partial x_3} = (1-d)\left(-\frac{1}{\rho_{++}}-\frac{4}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{4}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_5\partial x_3} = (1-d)\left(-\frac{1}{\rho_{+-}}-\frac{2}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{4}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_6\partial x_3} = (1-d)\left(-\frac{2}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{4}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_7\partial x_3} = (1-d)\left(-\frac{1}{\rho_{+-}}-\frac{2}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{2}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_8\partial x_3} = (1-d)\left(-\frac{2}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{2}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_9\partial x_3} = (1-d)\left(-\frac{1}{\rho_{++}}-\frac{2}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{2}{\mu_{----}}\right)$$

and

$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_4^2} = (1-d)\left(-\frac{1}{\rho_{++}}-\frac{1}{\rho_{-+}}-\frac{4}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{2}{x_4}-\frac{4}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_5\partial x_4} = (1-d)\left(-\frac{2}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{4}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_6\partial x_4} = (1-d)\left(-\frac{1}{\rho_{-+}}-\frac{2}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{4}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_7\partial x_4} = (1-d)\left(-\frac{2}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{2}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_8\partial x_4} = (1-d)\left(-\frac{1}{\rho_{-+}}-\frac{2}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{2}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_9\partial x_4} = (1-d)\left(-\frac{1}{\rho_{++}}-\frac{2}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{2}{\mu_{----}}\right)$$

and

$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_5^2} = (1-d)\left(-\frac{1}{\rho_{+-}}-\frac{1}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{2}{x_5}-\frac{4}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_6\partial x_5} = (1-d)\left(-\frac{1}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{4}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_7\partial x_5} = (1-d)\left(-\frac{1}{\rho_{+-}}-\frac{1}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{2}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_8\partial x_5} = (1-d)\left(-\frac{1}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{2}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_9\partial x_5} = (1-d)\left(-\frac{1}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{2}{\mu_{----}}\right)$$

and

$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_6^2} = (1-d)\left(-\frac{1}{\rho_{-+}}-\frac{1}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{2}{x_6}-\frac{4}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_7\partial x_6} = (1-d)\left(-\frac{1}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{2}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_8\partial x_6} = (1-d)\left(-\frac{1}{\rho_{-+}}-\frac{1}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{2}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_9\partial x_6} = (1-d)\left(-\frac{1}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{2}{\mu_{----}}\right)$$

and

$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_7^2} = (1-d)\left(-\frac{1}{\rho_{+-}}-\frac{1}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{1}{x_7}-\frac{1}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_8\partial x_7} = (1-d)\left(-\frac{1}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{1}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_9\partial x_7} = (1-d)\left(-\frac{1}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{1}{\mu_{----}}\right)$$

and

$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_8^2} = (1-d)\left(-\frac{1}{\rho_{-+}}-\frac{1}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{1}{x_8}-\frac{1}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_9\partial x_8} = (1-d)\left(-\frac{1}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{1}{\mu_{----}}\right)$$
$$\frac{\partial^2\delta(\mu,\rho)}{\partial x_9^2} = (1-d)\left(-\frac{1}{\rho_{++}}-\frac{1}{\rho_{--}}\right)+\frac{d}{2}\left(-\frac{1}{x_9}-\frac{1}{\mu_{----}}\right).$$

Recall the definition of μ^*

$$\mu^*_{++++} = \mu^*_{----} = \mu^*_{+-+-} = \mu^*_{-+-+} = \frac{e^{-2\beta}}{4\left(1+e^{-\beta}\right)^2}$$
$$\mu^*_{+--+} = \mu^*_{--++} = \mu^*_{-++-} = \mu^*_{++--} = \frac{1}{4\left(1+e^{-\beta}\right)^2}$$
$$\mu^*_{+++-} = \mu^*_{++-+} = \mu^*_{+-++} = \mu^*_{-+++} = \mu^*_{---+} = \mu^*_{--+-} = \mu^*_{-+--} = \mu^*_{+---} = \frac{e^{-\beta}}{4\left(1+e^{-\beta}\right)^2}$$

which implies

$$\rho^*_{++} = \rho^*_{+-} = \rho^*_{-+} = \rho^*_{--} = \frac{1}{4}$$

and

$$\begin{aligned}\delta(\mu^*,\rho^*) &= (1-d)\,\mathrm{H}(\rho^*)+\frac{d}{2}\mathrm{H}(\mu^*)-d\beta\left(x_3^*+x_4^*+x_5^*+x_6^*+x_7^*+x_8^*+x_9^*+\mu^*_{----}\right)\\ &= (2-d)\log(2)+d\log\left(1+e^{-\beta}\right).\end{aligned}$$

Evaluating the above derivatives at μ^*, ρ^* we obtain the Hessian at μ^*, ρ^*.

$$\mathrm{D}^2\delta\left(\mu^*,\rho^*\right) = 4(d-1)\begin{pmatrix} 6 & 2 & 5 & 5 & 3 & 3 & 3 & 3 & 2 \\ 2 & 2 & 3 & 3 & 1 & 1 & 1 & 1 & 2 \\ 5 & 3 & 6 & 5 & 3 & 2 & 3 & 2 & 3 \\ 5 & 3 & 5 & 6 & 2 & 3 & 2 & 3 & 3 \\ 3 & 1 & 3 & 2 & 2 & 1 & 2 & 1 & 1 \\ 3 & 1 & 2 & 3 & 1 & 2 & 1 & 2 & 1 \\ 3 & 1 & 3 & 2 & 2 & 1 & 2 & 1 & 1 \\ 3 & 1 & 2 & 3 & 1 & 2 & 1 & 2 & 1 \\ 2 & 2 & 3 & 3 & 1 & 1 & 1 & 1 & 2 \end{pmatrix}$$

$$-2d\frac{\left(1+e^{-\beta}\right)^2}{e^{-2\beta}}\begin{pmatrix} 4 & 4 & 4 & 4 & 4 & 4 & 2 & 2 & 2 \\ 4 & 4 & 4 & 4 & 4 & 4 & 2 & 2 & 2 \\ 4 & 4 & 4 & 4 & 4 & 4 & 2 & 2 & 2 \\ 4 & 4 & 4 & 4 & 4 & 4 & 2 & 2 & 2 \\ 4 & 4 & 4 & 4 & 4 & 4 & 2 & 2 & 2 \\ 4 & 4 & 4 & 4 & 4 & 4 & 2 & 2 & 2 \\ 2 & 2 & 2 & 2 & 2 & 2 & 1 & 1 & 1 \\ 2 & 2 & 2 & 2 & 2 & 2 & 1 & 1 & 1 \\ 2 & 2 & 2 & 2 & 2 & 2 & 1 & 1 & 1 \end{pmatrix}$$

$$-2d\left(1+e^{-\beta}\right)^2\begin{pmatrix} 2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 2e^{\beta} & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 2e^{\beta} & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 2e^{\beta} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 2e^{\beta} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & e^{2\beta} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & e^{2\beta} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & e^{2\beta} \end{pmatrix}$$

The lemma now follows from calculating the determinant of the preceding expression.

REFERENCES

[1] A. Barvinok: Combinatorics and complexity of partition functions. Switzerland: Springer **9** (2016).

[2] A. Coja-Oghlan, C. Efthymiou, N. Jaafari, M. Kang, T. Kapetanopoulos: Charting the replica symmetric phase. Communications in Mathematical Physics **359** (2018) 603–698.

[3] A. Coja-Oghlan, P. Loick, B. Mezei, G. Sorkin: The Ising antiferromagnet and max cut on random regular graphs. arXiv preprint arXiv:2009.10483 (2020).

[4] A. Coja-Oghlan, M. Hahn-Klimroth, P. Loick, N. Müller, K. Panagiotou, M. Pasch: Inference and mutual information on random factor graphs. Proc. of 37th International Symposium on Theoretical Aspects of Computer Science (2021).

[5] C. Greenhill, S. Janson, A. Ruciński: On the number of perfect matchings in random lifts. Combinatorics, Probability and Computing **19** (2010) 791–817.

[6] F. Guerra, F. Toninelli: The high temperature region of the Viana–Bray diluted spin glass model. Journal of statistical physics **115** (2004) 531–555.

[7] Huang, K. (2009). Introduction to statistical physics. CRC press.

[8] S. Janson: Random regular graphs: asymptotic distributions and contiguity. Combinatorics, Probability and Computing **4** (1995) 369–405.

[9] S. Janson, T. Luczak, A. Rucinski: Random graphs. John Wiley & Sons **45** (2011).

[10] M. Mézard, A. Montanari: Information, physics and computation. Oxford University Press 2009.

[11] E. Mossel, J. Neeman, A. Sly: Reconstruction and estimation in the planted partition model. Probability Theory and Related Fields **162** (2015) 431–461.

[12] R. Robinson, N. Wormald: Almost all cubic graphs are Hamiltonian. Random Structures & Algorithms **3** (1992) 117–125.

CHRISTIAN FABIAN, `cfabian@stud.uni-frankfurt.de`, GOETHE UNIVERSITY, MATHEMATICS INSTITUTE, 10 ROBERT MAYER ST, FRANKFURT 60325, GERMANY.

PHILIPP LOICK, `loick@math.uni-frankfurt.de`, GOETHE UNIVERSITY, MATHEMATICS INSTITUTE, 10 ROBERT MAYER ST, FRANKFURT 60325, GERMANY.

APPENDIX I. CURRICULUM VITAE

Curriculum Vitae – Philipp Loick

e-mail: loick@math.uni-frankfurt.de
phone: +49 163 3783702
date of birth: 03.11.1992
nationality: German

Education & Qualifications

Sep-18 – today **Goethe University** PhD Mathematics
Faculty: Mathematics, location: Frankfurt a.M., Germany, graduation: spring 2022
- Research project *Message-passing algorithms, information-theoretic thresholds and computational barriers* funded by the German Research Foundation
- Coursework in *Higher Stochastics* and *Probabilistic Combinatorics*, grade: 1.0

Sept-17 – Sep-18 **London School of Economics & Political Science** MSc Operations Research & Analytics
Faculty: Mathematics, location: London, UK
- GPA: 84/100 with distinction
- Coursework in *Algorithms & Computation, Combinatorial Optimization, Algorithmic Techniques for Data Mining, Statistics, Distributed Computing for Big Data*
- Master thesis on supervised learning algorithms in the context of customised pricing

Aug-10 – Jun-13 **Maastricht University** BSc International Business, major finance
Location: Maastricht, Netherlands
- GPA: 8.8/10.0 cum laude
- Participation in honors program with focus on economics

Aug-12 – Dec-12 **University of California, Berkeley** Undergraduate exchange student
Location: Berkeley, CA, USA
- GPA: 3.91/4.0
- Research assistant to Prof. Clayton Critcher's research on behavioral economics

Aug-03 – Jun-10 **Städtisches Gymnasium Marienschule** Highschool grade: 1.0
Location: Euskirchen, Germany

Work Experience

Feb-16 – Aug-17 **Relias Learning, part of Bertelsmann Education Group** Raleigh, NC, USA
Manager of Financial Planning & Analysis (FP&A)

May-14 – Jan-16 **Bertelsmann Education Group** New York City, NY, USA
Investment Manager (01/16 – 02/16), Financial Analyst (10/14 – 12/15), Financial Analyst Intern (05/14 – 09/14)

Dec-13 – May-14 **Bertelsmann SE & Co. KGaA** Gütersloh, Germany
Intern as part of gap-year rotation – Corporate Controlling & Strategy

Sep-13 – Nov-13 **Henkel AG & Co. KGaA** Düsseldorf, Germany
Intern as part of gap-year rotation – Office of the CFO

Additional information

Publications:

A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimroth, P. Loick: Information-theoretic and algorithmic thresholds for group testing.
-- *46th International Colloqium on Automata, Languages, and Programming (ICALP 2019), Leibniz International Proceedings in Informatics (LIPIcs)* (**132**) (2019), 43:1-43:14
-- *IEEE Transactions on Information Theory (2020) doi: 10.1109/TIT.2020.3023377*

O. Gebhard, M. Hahn-Klimroth, D. Kaaser, P. Loick: Quantitative group testing in the sublinear regime. *preprint arXiv 1905.01458* (2019)

A. Coja-Oghlan, O. Gebhard, M. Hahn-Klimroth, P. Loick: Optimal group testing.
-- *Proceedings of Machine Learning Research* (**125**) (2020) 1374—1388 (33rd COLT).
-- *Combinatorics, Probability and Computing, 1-38 (2021) doi:10.1017/S096354832100002X*

A. Coja-Oghlan, P. Loick, B. Mezei, G. Sorkin: The Ising antiferromagnet and max cut on random regular graphs. *preprint arXiv:2009.10483* (2020)

O. Gebhard, O. Johnson, P. Loick, M. Rolvien: Improved bounds for noisy group testing with constant tests per item. *preprint arXiv:2007.01376* (2020)

A. Coja-Oghlan, M. Hahn-Klimroth, P. Loick, N. Müller, K. Panagiotou, M. Pasch: Inference and mutual information on random factor graphs. *38th International Symposium on Theoretical Aspects of Computer Science (STACS 2021), Leibniz International Proceedings in Informatics* (2021), 24:1-24:15

C. Fabian, P. Loick: The Ising antiferromagnet in the replica symmetric phase. *preprint arXiv:2103.09775 (2021)*

O. Gebhard, P. Loick: Note on the offspring distribution for group testing in the linear regime. *preprint arXiv: 2103.13039 (2021)*

A. Coja-Oghlan, M. Hahn-Klimroth, P. Loick, M. Penschuck: Efficient and accurate group testing via Belief Propagation: an empirical study. *preprint arXiv: 2105.07882 (2021)*

Talks:

Maximum cut on random regular graphs. *19th International Conference on Random Structures and Algorithms,* ETH Zurich (2019)

Optimal group testing.
-- 33rd Annual Conference on Learning Theory, virtual (2020)
-- 5th Highlights of Algorithms Conference, virtual (2020)