

Spectral Planting and the Hardness of Refuting Cuts, Colorability, and Communities in Random Graphs

Afonso S. Bandeira

Department of Mathematics, ETH Zurich

BANDEIRA@MATH.ETHZ.CH

Jess Banks

Department of Mathematics, UC Berkeley

JESS.M.BANKS@BERKELEY.EDU

Dmitriy Kunisky

Department of Mathematics, Courant Institute of Mathematical Sciences, NYU

KUNISKY@CIMS.NYU.EDU

Cristopher Moore

Santa Fe Institute

MOORE@SANTAFE.EDU

Alexander S. Wein

Department of Mathematics, Courant Institute of Mathematical Sciences, NYU

AWEIN@CIMS.NYU.EDU

Editors: Mikhail Belkin and Samory Kpotufe

Abstract

We study the problem of efficiently *refuting* the k -colorability of a graph, or equivalently, *certifying* a lower bound on its chromatic number. We give formal evidence of average-case computational hardness for this problem in sparse random regular graphs, suggesting that there is no polynomial-time algorithm that improves upon a classical spectral algorithm. Our evidence takes the form of a *computationally-quiet planting*: we construct a distribution of d -regular graphs that has significantly smaller chromatic number than a typical regular graph drawn uniformly at random, while providing evidence that these two distributions are indistinguishable by a large class of algorithms. We generalize our results to the more general problem of certifying an upper bound on the maximum k -cut.

This quiet planting is achieved by minimizing the effect of the planted structure (e.g. colorings or cuts) on the graph spectrum. Specifically, the planted structure corresponds exactly to eigenvectors of the adjacency matrix. This avoids the pushout effect of random matrix theory, and delays the point at which the planting becomes visible in the spectrum or local statistics. To illustrate this further, we give similar results for a Gaussian analogue of this problem: a quiet version of the spiked model, where we plant an eigenspace rather than adding a generic low-rank perturbation.

Our evidence for computational hardness of distinguishing two distributions is based on three different heuristics: stability of belief propagation, the local statistics hierarchy, and the low-degree likelihood ratio. Of independent interest, our results include general-purpose bounds on the low-degree likelihood ratio for multi-spiked matrix models, and an improved low-degree analysis of the stochastic block model.

Keywords: Refutation, quiet planting, chromatic number, random graphs, spectral methods

1. Introduction

Assuming the widely believed $P \neq NP$ hypothesis, many combinatorial problems in graphs are known to be computationally hard. Prominent examples from graph theory and network science include finding large cliques or independent sets, clustering or maximizing cuts, and finding vertex colorings or computing chromatic numbers. Fortunately, the worst-case computational difficulty of many of these problems appears to not be predictive of their feasibility in typical graphs, motivating the study of forms of *average-case complexity* for many of these problems. Many remarkable examples exist dating back to at least the work of Karp and others in the mid 70s (Karp, 1976),¹ they include the problem of vertex colorings (Grimmett and McDiarmid, 1975) in random graphs, the related problem of finding the largest independent set² (or, equivalently, clique) and many others (Karp, 1976). In both of these problems, for certain natural distributions of random graphs, a multiplicative gap of 2 is identified between the typical optimal solution and the solution found by the best known polynomial-time algorithm (Karp, 1986), and improving over this has since been a standing open problem. Motivated by this question, Kucera (1995) and Alon et al. (1998) studied random graph models with planted structures (either a large independent set or clique, or a coloring with an unusually small number of colors) and investigate when such structures are easy to detect. Foreshadowing to what follows, we point out that the existence of a planted structure that cannot be detected efficiently implies that it is impossible to efficiently refute the existence of such a structure in the underlying unplanted model. In this paper we will focus on the problem of computing the chromatic number, and the related problem of understanding the size of the largest k -cut. The random graph model throughout is the uniform distribution over d -regular graphs on n nodes.

Refuting colorability. For an integer $k \geq 1$, a graph $G = (V, E)$ is k -colorable if there exists an assignment $\sigma: V \rightarrow [k]$ of “colors” to the vertices such that $\sigma(i) \neq \sigma(j)$ for every edge $(i, j) \in E$. The *chromatic number* $\chi(G)$ of G is defined as the minimum value of k for which G is k -colorable.

A random d -regular graph G on n vertices (we will write random variables in bold-face font throughout the paper) is known to have a typical chromatic number $\chi(G) \sim \frac{1}{2} \frac{d}{\log d}$ (see, e.g., Coja-Oghlan et al. (2016) and references therein) in the double limit $n \rightarrow \infty$ followed by $d \rightarrow \infty$, where $f \sim g$ denotes $f/g \rightarrow 1$. The problem we will study is that of algorithmically *refuting* the k -colorability of a graph, which we define in Section 2.2. Informally speaking, an algorithm that refutes k -colorability provides an efficiently-verifiable *proof* that G is not k -colorable. As a simple example, an algorithm exhibiting a $(k + 1)$ -clique refutes k -colorability. More generally, one may encode a k -coloring as a collection of boolean variables satisfying certain logical relations depending on G , and refute coloring by deriving a contradiction from those axioms.

We will provide evidence that the refutation problem is computationally hard when G is a uniformly random d -regular graph. The proof strategy is to construct a different distribution over d -regular graphs whose typical chromatic number is $\chi(G) \sim \frac{1}{2} \sqrt{d}$, and to argue that this distribution is computationally hard to distinguish from a uniformly random d -regular graph (whose chromatic number is instead $\chi(G) \sim \frac{1}{2} \frac{d}{\log d}$)—we think of this new distribution as having a *computationally-quiet planting*³ of a coloring with few colors. We will see below that the value $\frac{1}{2} \sqrt{d}$ coincides with a simple spectral bound on $\chi(G)$ for G uniformly random, and so our result is a tight computational

1. See also Karp’s lecture on the occasion of his Turing Award (Karp, 1986).

2. Recall that the set of nodes of the same color in a vertex coloring is an independent set.

3. This terminology is inspired by the notion of *quiet planting* from prior work (Krzakala and Zdeborová, 2009; Zdeborová and Krzakala, 2011), although our notion is somewhat different.

lower bound on refutation algorithms. The formal evidence for computational hardness used in this paper is threefold: we provide consistent pieces of evidence based on (i) the Kesten–Stigum threshold for the belief propagation message-passing algorithm, (ii) the Local Statistics semidefinite programming hierarchy, and (iii) the low-degree likelihood ratio for an analogous Gaussian model.

The spectral refutation. Instead of just refuting k -colorability, we will consider the more general task of refuting the existence of large k -cuts in a graph. We define the fractional size of the largest such cut as

$$\text{MC}_k(G) := \max_{\sigma: V \rightarrow [k]} \frac{|\{(u, v) \in E : \sigma(u) \neq \sigma(v)\}|}{|E|} \in (0, 1]. \quad (1)$$

Intuitively, $\text{MC}_k(G)$ describes how close G is to being k -colorable, as the cut counts the fraction of polychromatic edges under the coloring σ . $\text{MC}_k(G)$ is non-decreasing in k , and for any k , G is k -colorable if and only if $\text{MC}_k(G) = 1$; therefore, the chromatic number is given by

$$\chi(G) = \min\{k : \text{MC}_k(G) = 1\} = 1 + \max\{k : \text{MC}_k(G) < 1\}. \quad (2)$$

Accordingly, upper bounds on MC_k away from the maximum value of 1 yield lower bounds on the chromatic number. This task is often called *certifying* a bound on the optimization problem MC_k . The relations (2) show how certifying such a bound in turn refutes colorability. In fact, Hoffman’s early work (Hoffman, 1970) proposed a technique for any d -regular graph G that gives, when rephrased in our notation, a refutation of colorability by bounding MC_k via the minimum eigenvalue $\lambda_{\min}(A_G)$ of the adjacency matrix A_G of G . Namely, Hoffman showed that for any d -regular graph G ,

$$\text{MC}_k(G) \leq \frac{k-1}{k} \left(1 + \frac{-\lambda_{\min}(A_G)}{d} \right). \quad (3)$$

(Note that $\lambda_{\min}(A_G) \leq 0$ since $\text{Tr}(A_G) = 0$.) A short proof of (3) will be given in Section 2.2. We note that $\frac{k-1}{k}$ is the expected value of the objective of (1) when $\sigma: V \rightarrow [k]$ is chosen uniformly at random, so expressions like the right-hand side of (3) should be viewed as expressing a factor of “gain” over this value. For $\mathbf{G}$ a uniformly *random* d -regular graph on n vertices, whose law we denote $\mathcal{G}_{n,d}$, a theorem due to Friedman (2003) states that $\lambda_{\min}(\mathbf{G}) = -2\sqrt{d-1} + o(1)$ with high probability⁴ for any fixed $d \geq 3$. This implies that when $\mathbf{G} \sim \mathcal{G}_{n,d}$, Hoffman’s spectral approach with high probability certifies the upper bound

$$\text{MC}_k(\mathbf{G}) \leq \frac{k-1}{k} \left(1 + \frac{2\sqrt{d-1}}{d} + o(1) \right). \quad (4)$$

This in turn translates to a lower bound on the chromatic number of $\chi(\mathbf{G}) \geq (1 - o_d(1))\frac{1}{2}\sqrt{d}$, where o_d pertains to the double limit $n \rightarrow \infty$ followed by $d \rightarrow \infty$ (see Section 2.1).

Equipped with this direct analysis of a simple technique, the natural question arises: can any polynomial-time algorithm produce a bound that, like Hoffman’s, is valid for any graph G , but is typically *tighter* for $\mathbf{G} \sim \mathcal{G}_{n,d}$? Most prior work on this question has studied bounds provided by a semidefinite program computing the Lovász ϑ function, which, as shown by Banks et al. (2017), is equivalent to the degree-2 sum-of-squares relaxation of $\text{MC}_k(G)$. The work of Coja-Oghlan (2003)

4. Here and throughout, $o(1)$ pertains to the limit $n \rightarrow \infty$ with d fixed. We say that an event occurs *with high probability* (w.h.p.) if it has probability $1 - o(1)$.

and a later, more precise analysis by Banks et al. (2017) showed that, with high probability when $\mathbf{G} \sim \mathcal{G}_{n,d}$, this relaxation certifies a bound still no better than

$$\text{MC}_k(\mathbf{G}) \leq \frac{k-1}{k} \left(1 + \frac{2\sqrt{d-1}}{d+2\sqrt{d-1}} - o(1) \right) = \frac{k-1}{k} \left(1 + \frac{2\sqrt{d-1}}{d} - O_d\left(\frac{1}{d}\right) \right). \quad (5)$$

Note that (5) matches the spectral bound (4) in the double limit $n \rightarrow \infty$ followed by $d \rightarrow \infty$. No polynomial-time certifier is known to asymptotically improve upon this bound. It appears plausible, then, that the spectral bound (4) is an optimal efficiently-computable certificate on $\text{MC}_k(\mathbf{G})$ for $\mathbf{G} \sim \mathcal{G}_{n,d}$. In this work, we will argue that this is indeed the case.

Detecting a planted k -cut. The following line of reasoning will be central to this work: to prove computational hardness of a refutation problem, it is sufficient to construct a *computationally-quiet planted distribution* (e.g. Bandeira et al. (2020)). To illustrate this, suppose our goal is to show computational hardness of refuting (w.h.p.) k -colorability of a graph drawn from $\mathcal{G}_{n,d}$. Suppose we are able to construct a *planted* distribution $\mathcal{P}$ over d -regular graphs such that (i) a typical graph drawn from $\mathcal{P}$ is k -colorable, and (ii) $\mathcal{P}$ is *computationally quiet* in the sense that no polynomial-time algorithm can distinguish (w.h.p.) between a sample from $\mathcal{P}$ and a sample from $\mathcal{G}_{n,d}$. It then follows that no polynomial-time algorithm can refute k -colorability in $\mathcal{G}_{n,d}$, because if such a refutation algorithm were to exist, it must succeed w.h.p. on $\mathcal{G}_{n,d}$ and must fail w.h.p. on $\mathcal{P}$, thus providing a solution to the distinguishing problem. More generally, to show hardness of certifying an upper bound on the maximum k -cut, we need a planted distribution for which there exists a large k -cut.

As discussed in Banks et al. (2017), a natural planted distribution with a large k -cut is a d -regular variant of the popular *stochastic block model (SBM)*. To sample a graph from this d -regular distribution, which we will denote $\mathcal{G}_{n,d,k,\eta}^{\text{sbm}}$, first sample a balanced labelling $\sigma : [n] \rightarrow [k]$ uniformly at random (*balanced* means that $|\sigma^{-1}(i)| = n/k$ for every $i \in [k]$), and then choose $\mathbf{G}$ uniformly among d -regular graphs conditional on the event⁵

$$\frac{|\{(u,v) \in E : \sigma(u) \neq \sigma(v)\}|}{|E|} = \frac{k-1}{k}(1-\eta). \quad (6)$$

Note that this ensures $\text{MC}_k(\mathbf{G}) \geq \frac{k-1}{k}(1-\eta)$, since the planted partition σ witnesses a k -cut with that fraction of bichromatic edges. We will mostly be concerned with the *disassortative* regime of this model where $\eta \in [-\frac{1}{k-1}, 0]$, and the planted k -cut σ is larger than a typical one; for instance, if $\eta = -\frac{1}{k-1}$ the planted cut includes every edge, giving the well-studied planted coloring model.

The relevance of this distribution to certifying bounds on MC_k for $\mathbf{G} \sim \mathcal{G}_{n,d}$ is as follows: if an algorithm can with high probability over $\mathcal{G}_{n,d}$ certify that $\text{MC}_k(\mathbf{G}) < \frac{k-1}{k}(1-\eta)$, then (as discussed above) it is simple to build another *testing* (or *detection*) algorithm that distinguishes between $\mathbf{G} \sim \mathcal{G}_{n,d}$ and $\mathbf{G} \sim \mathcal{G}_{n,d,k,\eta}^{\text{sbm}}$ with high probability. The advantage of taking this point of view is that there is a rich literature, originating in heuristic methods from statistical physics, that has provided a great deal of evidence that polynomial-time testing between $\mathcal{G}_{n,d}$ and $\mathcal{G}_{n,d,k,\eta}^{\text{sbm}}$ is impossible below the *Kesten–Stigum threshold*, i.e., when

$$d < d_{\text{KS}}^{\text{sbm}} = d_{\text{KS}}^{\text{sbm}}(\eta) := \frac{1}{\eta^2} + 1. \quad (7)$$

5. For now, we will not consider the integrality conditions on the parameters (n, d, k, η) that this implies, although these types of considerations will be important later.

We will refer to this claim as the *SBM conjecture*. Polynomial-time algorithms are known to succeed when $d > d_{\text{KS}}^{\text{sbm}}$ (Massoulié, 2014; Mossel et al., 2018; Abbe and Sandon, 2016; Banks et al., 2019). While proving such a conjecture seems to be beyond the reach of current techniques (even under an assumption such as $\mathbb{P} \neq \text{NP}$), various forms of concrete evidence have been given (either for the d -regular SBM and other variants). These include results on stability of belief propagation (Decelle et al., 2011a,b), the local statistics hierarchy (Banks et al., 2019), and the low-degree likelihood ratio (Hopkins and Steurer, 2017; Hopkins, 2018). We will discuss all of these methods further in Section 1.1.

Rearranging (7) we find that, conditional on the SBM conjecture, the above argument implies that when $G \sim \mathcal{G}_{n,d}$, no polynomial-time algorithm can certify with high probability a bound stronger than

$$\text{MC}_k(G) \leq \frac{k-1}{k} \left(1 + \frac{1}{\sqrt{d-1}} - o(1) \right). \quad (8)$$

Comparing this to (4) and (5), we see a discrepancy between the best known certification algorithms and the above hardness result. For large d , this discrepancy amounts to a factor of 2 in the “gain” term. This begs the question of whether better certification algorithms exist, or whether the hardness result can be improved. We will see below that it is the latter.

A quieter planting. Our main contribution is to show an improved hardness result by using a “better” planted distribution: namely, a more rigid version of the SBM in which, for each $i, j \in [k]$, every vertex with label i has a fixed number of neighbors with label j . This is in contrast to the d -regular SBM in (6) above, where we condition on the *total* number of bichromatic edges.

Definition 1 (Equitable stochastic block model) *The equitable stochastic block model (eSBM), denoted $\mathcal{G}_{n,d,k,\eta}^{\text{eq}}$, is the probability distribution over d -regular graphs on n vertices sampled as follows: first, choose a uniformly random balanced partition $\sigma : V \rightarrow [k]$. Then, letting*

$$M = \eta I_k + \frac{1-\eta}{k} J_k$$

(where I_k and J_k are the $k \times k$ identity and all-ones matrices, respectively), for each $i \in [k]$ place a random $dM_{i,i}$ -regular graph on the color class $\sigma^{-1}(i)$, and for each $i < j \in [k]$ place a random bipartite $dM_{i,j}$ -regular graph between $\sigma^{-1}(i)$ and $\sigma^{-1}(j)$. This model is only defined when $k|n$, dM is a nonnegative integer matrix, and $dM_{i,i}n/k$ is even⁶ for all i .

As in $\mathcal{G}_{n,d,k,\eta}^{\text{sbm}}$, the planted cut σ has fractional size $\frac{k-1}{k}(1-\eta)$, and we will again restrict to the disassortative case $\eta \in [-\frac{1}{k-1}, 0]$. This model is discussed in Brito et al. (2016), where it is called the “regular block model.” We instead follow Newman and Martin (2014); Barucca (2017) in using the term “equitable,” so as to differentiate it from the standard d -regular variant of the SBM discussed in the previous section.

Using some of the same methods that provide concrete evidence for the SBM conjecture—namely stability of belief propagation and the local statistics hierarchy—we will show that the equitable SBM appears to exhibit a different computational threshold from the standard SBM. Specifically, we conjecture that no polynomial-time algorithm can distinguish (w.h.p.) between $\mathcal{G}_{n,d}$ and $\mathcal{G}_{n,d,k,\eta}^{\text{eq}}$ when

$$d < d_{\text{KS}}^{\text{eq}} = d_{\text{KS}}^{\text{eq}}(\eta) := \frac{2}{\eta^2} \left(1 + \sqrt{1-\eta^2} \right). \quad (9)$$

6. This last condition ensures that it is possible to place a $dM_{i,i}$ -regular graph on n/k vertices.

We will state a formal version of this “eSBM conjecture” later (Conjecture 8), which actually pertains to a slightly “noisy” version of the equitable block model. We remark that when k is large, η must be close to zero (in the disassortative case), and so we have approximately $d_{\text{KS}}^{\text{eq}} \approx 4d_{\text{KS}}^{\text{sbm}}$.

Repeating our earlier argument for hardness of certification, with the eSBM in place of the SBM, yields the following corresponding result: conditional on the eSBM conjecture, no polynomial-time algorithm can certify a better bound than

$$\text{MC}_k(\mathbf{G}) \leq \frac{k-1}{k} \left(1 + \frac{2\sqrt{d-1}}{d} - o(1) \right) \quad (10)$$

when $\mathbf{G} \sim \mathcal{G}_{n,d}$, which matches the spectral bound (4). However, we have ignored an important caveat here: the equitable block model only exists when the parameters (n, d, k, η) satisfy certain integrality conditions. For this reason, our actual lower bound (Theorem 9) is sometimes weaker than (10) would suggest; see Section A.2 for discussion.

When $d \gg k^2$, the integrality conditions are negligible and we obtain a tight lower bound, essentially matching (10). Another setting where we obtain tight results is for the problem of refuting colorability (or more accurately, near-colorability) in the double limit $n \rightarrow \infty$ followed by $d \rightarrow \infty$, which is discussed in Remarks 10 and 11. This corresponds to the choice $\eta = -\frac{1}{k-1}$. Here, when $\mathbf{G} \sim \mathcal{G}_{n,d}$, the following results hold asymptotically: the true value of $\chi(\mathbf{G})$ is $\frac{1}{2} \frac{d}{\log d}$, the spectral approach certifies a lower bound of $\frac{1}{2} \sqrt{d}$ on $\chi(\mathbf{G})$, the basic SBM planting implies hardness of certifying a lower bound better than $\sqrt{d}$, and the improved eSBM planting implies hardness of certifying a lower bound better than $\frac{1}{2} \sqrt{d}$ (which is tight, matching the spectral bound).

Theorem 2 (Informal) *Assuming Conjecture 8 on the computational hardness of hypothesis testing for the eSBM, there exists no polynomial-time algorithm certifying a bound on $\text{MC}_k(\mathbf{G})$ asymptotically tighter than the spectral bound (4) for $\mathbf{G} \sim \mathcal{G}_{n,d}$ with $d \gg k^2$. Furthermore, for any $\epsilon > 0$ there exists no polynomial-time algorithm refuting k -colorings with at most $\epsilon|E|$ monochromatic edges for $k \leq \frac{1}{2}(1 - o_d(1))\sqrt{d}$.*

See Theorem 9 and Remark 10 for a formal and more complete version of this statement, where the case $d \lesssim k^2$ is also addressed.

Why is the equitable SBM quieter? Here we give some intuition for why the equitable model is a good quiet planting. For the sake of illustration, it helps to consider the simple rank-1 Wigner spiked matrix model: $\mathbf{Y} = \eta \mathbf{v} \mathbf{v}^\top + \mathbf{W}$ where $\eta > 0$ (the signal-to-noise ratio), $\|\mathbf{v}\| = 1$ (the planted “signal,” drawn from some prior), and $\mathbf{W}$ (the “noise”) is a GOE matrix, i.e., a symmetric matrix with $\mathcal{N}(0, 1/n)$ entries (see Definition 14). For large n , the eigenvalues of $\mathbf{W}$ follow the semicircle law and are contained in the interval $[-2, 2]$. For $1 < \eta < 2$, a surprising “pushout” effect occurs: although the planted signal $\mathbf{v}$ has quadratic form $\mathbf{v}^\top \mathbf{Y} \mathbf{v} \approx \eta < 2$, its presence causes there to exist some other unit vector $\mathbf{u}$ achieving $\mathbf{u}^\top \mathbf{Y} \mathbf{u} \approx \eta + 1/\eta > 2$ (Baik et al., 2005; Féral and Péché, 2007; Capitaine et al., 2009); as a result, simply checking the largest eigenvalue allows one to distinguish $\mathbf{Y}$ from $\mathbf{W}$. Even though the signal $\mathbf{v}$ is “small”, the vector $\mathbf{u}$ (which is the leading eigenvector of $\mathbf{Y}$) is able to achieve a “large” quadratic form by correlating nontrivially with both the signal $\mathbf{v}$ and the noise $\mathbf{W}$. The main result of Bandeira et al. (2020) can be interpreted as giving a “quieter” way to plant the signal “orthogonal to the noise” with no pushout effect, i.e., $\mathbf{v}^\top \mathbf{Y} \mathbf{v} \approx 2$ and the maximum eigenvalue of $\mathbf{Y}$ is ≈ 2 .

It turns out that a similar effect is at play in the SBM. Here the minimum eigenvalue of the adjacency matrix A of a random d -regular graph converges to $\lambda_{\min} := -2\sqrt{d-1}$. As will be made clear in Section 2.2, the SBM is in some sense analogous to the spiked Wigner model with multiple planted vectors, namely k “coloring vectors” v_i for $i \in [k]$, that encode the planted labelling σ as follows: $(v_i)_u = c(k \cdot \mathbf{1}[\sigma(u) = i] - 1)$, where c is chosen so that each $\|v_i\| = 1$. Planting a k -cut of value $\frac{k-1}{k}(1 + |\eta|)$ via either the SBM or eSBM has the effect that all planted coloring vectors achieve a small quadratic form: $v_i^\top A v_i \approx -d|\eta|$. In the SBM there is a pushout effect similar to the spiked Wigner model, whereby an eigenvalue less than $\lambda_{\min}$ can be created even when $d|\eta| < |\lambda_{\min}|$ (see [Nadakuditi and Newman \(2012\)](#) for results when the degree grows slowly with n). The eSBM, however, has the property that each coloring vector v_i is an eigenvector. In particular, the subspace spanned by $\{v_i\}_{i \in [k]}$ is orthogonal to the other eigenvectors, which are thus unaffected by the planted structure. As a result, there is no pushout effect in the eSBM, allowing for a larger k -cut to be planted without disrupting the minimum eigenvalue.

An alternative viewpoint is that the standard plantings (the spiked Wigner model or SBM) pick a random solution (e.g., a cut) and then condition on that *particular* solution having the desired value. In contrast, the quieter plantings are perhaps more similar to conditioning on the event “there exists a solution having the desired value.”

Remark 3 *The fact that the coloring vectors v_i are eigenvectors of the eSBM can actually be exploited to give a polynomial-time algorithm for distinguishing $\mathcal{G}_{n,d}$ from $\mathcal{G}_{n,d,k,\eta}^{\text{eq}}$ for any settings of the parameters. See, for example, [Barucca \(2017\)](#) for some discussion of such algorithms. For this reason, it is crucial that our eSBM conjecture (Conjecture 8) adds a small amount of noise to the graph in order to “defeat” these types of algorithms. We discuss this issue further in Section 1.1.4.*

The case $k = 2$: large cuts in $\mathcal{G}_{n,d}$. Here, we briefly discuss the specific case $k = 2$, which is better understood in the existing literature. In this case, $\text{MC}_2(G)$ is merely the fraction of edges crossing the largest cut of G , and thus up to this scaling is the solution to the well-known *max-cut* problem. Equivalently, letting A_G be the adjacency matrix of G , we have $\text{MC}_2(G) = \frac{1}{2}(1 + \frac{1}{2|E|}\Gamma_2(A_G))$ where

$$\Gamma_2(A) := - \min_{x \in \{\pm 1\}^n} x^\top A x \geq 0.$$

When $G \sim \mathcal{G}_{n,d}$, the behavior of $\Gamma_2(A_G)$ turns out to be deeply connected to its Gaussian analogue $\Gamma_2(W)$ where W is a GOE matrix. The quantity $\Gamma_2(W)$ has been studied in statistical physics, being the ground state energy of the *Sherrington-Kirkpatrick model* of spin glasses ([Sherrington and Kirkpatrick, 1975](#)). The deep but non-rigorous analysis of [Parisi \(1979\)](#) proposed an asymptotic value $\frac{1}{n} \mathbb{E} \Gamma_2(W) \rightarrow 2P_* \approx 1.526$, which was later proven rigorously in a sequence of mathematical works ([Guerra, 2003](#); [Talagrand, 2006](#); [Panchenko, 2011, 2013](#)). By relating the graph model to the Gaussian model, [Dembo et al. \(2017\)](#) gave an asymptotic formula for the size of the largest cut in a random regular graph with large degree:

$$\lim_{n \rightarrow \infty} \mathbb{E}_{G \sim \mathcal{G}_{n,d}} \text{MC}_2(G) = \frac{1}{2} \left(1 + \frac{2P_*}{\sqrt{d}} + o_d \left(\frac{1}{\sqrt{d}} \right) \right). \quad (11)$$

For the Gaussian setting, it was shown in [Bandeira et al. \(2020\)](#) using a quiet planting approach that (conditional on a certain complexity assumption based on the low-degree likelihood ratio) the best possible upper bound on $\frac{1}{n} \Gamma_2(W)$ that can be certified in polynomial time is 2, in contrast with the

true value $2P_*$; the optimal bound is given by a simple spectral certificate involving the maximum eigenvalue of $\mathbf{W}$. Given this, we might expect that the best efficiently-certifiable bound on $\text{MC}_2(\mathbf{G})$ for $\mathbf{G} \sim \mathcal{G}_{n,d}$ is given by replacing $2P_*$ in (11) with 2. Indeed, Montanari and Sen (2016) and Mohanty et al. (2019) showed respectively that the degree-2 and degree-4 sum-of-squares relaxations can certify a bound no better than

$$\text{MC}_2(\mathbf{G}) \leq \frac{1}{2} \left(1 + \frac{2}{\sqrt{d}} + o_d \left(\frac{1}{\sqrt{d}} \right) \right). \quad (12)$$

Our results extend the picture emerging from this literature in two important ways. We show (see Theorem 9 and discussion in Section A.2) that conditional on the eSBM conjecture, (12) is in fact the optimal bound on $\text{MC}_2(\mathbf{G})$ certifiable in polynomial time. Thus, if the eSBM conjecture holds then no constant-degree sum-of-squares relaxation can improve upon the known results for degree-2 and degree-4. Furthermore, we justify this bound with an explicit quiet planting of a large cut in a random regular graph.

The Gaussian k -cut model. Above, we have seen an intimate connection between $\Gamma_2(A_{\mathbf{G}})$ where $A_{\mathbf{G}}$ is the adjacency matrix of $\mathbf{G} \sim \mathcal{G}_{n,d}$, and its Gaussian counterpart $\Gamma_2(\mathbf{W})$ where $\mathbf{W} \sim \text{GOE}(n)$. More generally, $\text{MC}_k(\mathbf{G})$ can be written in terms of a certain quantity $\Gamma_k(A_{\mathbf{G}})$ defined in Section 2.2, which has a natural Gaussian counterpart $\Gamma_k(\mathbf{W})$.

In fact, explicit formulas similar to (11) are also known that relate the asymptotic values (in the double limit $n \rightarrow \infty$ followed by $d \rightarrow \infty$) of $\Gamma_k(A_{\mathbf{G}})$ and $\Gamma_k(\mathbf{W})$ even for $k > 2$ (Sen, 2018; Jagannath et al., 2018). The broadly applicable techniques used in these results—the Lindeberg exchange method and related probabilistic interpolation arguments—suggest that there is a general and fundamental relationship between the graph model in the large-degree limit and the Gaussian model. Yet, it is not clear whether this implies any relation between the respective thresholds for efficient certification.

To clarify this matter, we also give results showing that the problem of certifying upper bounds on $\Gamma_k(\mathbf{W})$ under the Gaussian model exhibits similar behavior to the graph model (see Section A.3): no polynomial-time certifier can improve over the basic spectral bound. The proof is again based on quiet planting, and can be seen as an extension of the results of Bandeira et al. (2020) which handled the $k = 2$ case. Our results rely on a complexity assumption concerning the low-degree likelihood ratio, which we discuss further in Section 1.1.3. As a by-product, we develop a general framework for bounding the low-degree likelihood ratio of certain Gaussian models, which may be of independent interest. Specifically, we conduct the low-degree analysis of a broad class of multi-spiked matrix models (both Wigner and Wishart), and give an improved low-degree analysis of the stochastic block model which suggests that fully-exponential time is needed below the Kesten–Stigum threshold (see Section A.3.2).

1.1. Heuristics for Average-Case Computational Hardness

Our results on hardness of certification rely on unproven conjectures about average-case hardness, such as the eSBM conjecture. Proving these types of conjectures seems to be beyond the reach of current techniques (even when assuming standard complexity conjectures such as $\text{P} \neq \text{NP}$), as illustrated by the fact that no such proof is known for the famous *planted clique* problem. However, a myriad of heuristic techniques have emerged for predicting hardness of average-case problems by proving lower bounds against certain classes of algorithms. Taken together, these methods create a

fairly coherent theory of computational complexity for a large class of high-dimensional Bayesian inference problems. In this section we describe the three methods that will be used in this work: belief propagation, the local statistics hierarchy, and the low-degree likelihood ratio. We remark that these are not the only such methods, some others being average-case reductions (Berthet and Rigollet, 2013; Brennan et al., 2018), sum-of-squares lower bounds (Barak et al., 2019), statistical query lower bounds (Feldman et al., 2017), the overlap gap property (Gamarnik and Zadik, 2017), and analysis of non-convex loss landscapes (Auffinger et al., 2013). Finally, in Section 1.1.4 we discuss the XOR-SAT problem—an important counterexample for many of the above heuristics—and the related issues of “robustness” that arise in our eSBM conjecture.

1.1.1. BELIEF PROPAGATION AND THE KESTEN–STIGUM THRESHOLD

The sharp computational phase transition known as the *Kesten–Stigum (KS) threshold* in the stochastic block model, was first predicted by Decelle et al. (2011a,b) using non-rigorous ideas inspired by statistical physics. The idea is to consider the *belief propagation (BP)* algorithm, an iterative method that attempts to recover the planted community structure by keeping track of “beliefs” about each node’s community label and updating these in a locally-Bayesian-optimal way. BP has an “uninformative” fixed point, which is a natural starting point for the algorithm where the beliefs reflect no knowledge of the communities. It was shown in Decelle et al. (2011a,b) that if the signal-to-noise ratio (SNR) lies above the KS threshold then the uninformative fixed point is unstable, suggesting that BP should leave it and find a community assignment that correlates with the truth; and if the SNR lies below the KS threshold then the uninformative fixed point is stable, meaning that BP will remain there and fail to find a nontrivial solution. It was later proven that indeed it is possible to nontrivially recover the communities in polynomial time when above the KS threshold (Massoulié, 2014; Mossel et al., 2018; Abbe and Sandon, 2016), whereas no such algorithm is known below the KS threshold. More generally, similar computational thresholds have been predicted in various models (e.g., Lesieur et al. (2015b,a)) by examining the stability of BP or its simplified variant, *approximate message passing (AMP)* (Donoho et al., 2009). For many high-dimensional inference problems, it is known that BP and AMP achieve optimal information-theoretic performance (e.g. Deshpande et al. (2017)); and when they don’t, it is often conjectured that they achieve the best possible performance among efficient algorithms. Thus, stability of BP provides concrete evidence for computational hardness. (Strictly speaking, stability of BP suggests hardness of recovering a partition that nontrivially correlates with the communities, whereas our eSBM conjecture pertains to the detection problem. However, we expect that the detection and recovery thresholds for the eSBM are the same, as is believed to be the case for the ordinary SBM. Also, the local statistics hierarchy discussed below gives direct evidence for hardness of detection.)

1.1.2. THE LOCAL STATISTICS HIERARCHY

Introduced by Banks, Mohanty, and Raghavendra (Banks et al., 2019) and building off of the work of Hopkins and Steurer (2017), the Local Statistics hierarchy is a family of increasingly powerful semidefinite programming algorithms for solving Bayesian hypothesis testing problems. By analogy with the Sum-of-Squares algorithm, we take this hierarchy as a proxy for hardness: the higher we need to go to perform the hypothesis test, the harder it is.

Consider a generic inference scheme where we are to distinguish between a null model $\mathbb{Q}$ which outputs unstructured data $\mathbf{G} \in \mathbb{R}^m$, and a planted model $\mathbb{P}$ generating structured data $\mathbf{G} \in \mathbb{R}^m$

according to some random and hidden signal $\mathbf{x} \in \mathbb{R}^n$. Letting $x = (x_1, \dots, x_n)$ be a set of variables, we may regard the conditional expectation $\mathbb{E}_{\mathbf{x} \sim \mathbb{P}}[p(\mathbf{x})|\mathbf{G}]$ as a random linear functional from $\mathbb{R}[x] \mapsto \mathbb{R}$ that is positive in a certain sense: $\mathbb{E}_{\mathbf{x} \sim \mathbb{P}}[p(\mathbf{x})^2|\mathbf{G}] \geq 0$ for every polynomial p .

The Local Statistics hierarchy is parameterized by two integers (D_x, D_G) . Given as input some $G_0 \in \mathbb{R}^m$, it attempts to find a linear functional that approximates this conditional expectation $\mathbb{E}[p(x)|G]$ in the planted model. In particular, borrowing terminology from Sum-of-Squares programming, we search for a “pseudoexpectation” functional $\tilde{\mathbb{E}}$ that assigns a real number to every polynomial of degree at most D_x in $\mathbb{R}[x]$, with the constraints that (i) $\tilde{\mathbb{E}}p(x)^2 \geq 0$, and (ii) $\tilde{\mathbb{E}}p(x) \approx \mathbb{E}_{(\mathbf{x}, \mathbf{G}) \sim \mathbb{P}} p(\mathbf{x})$ for every polynomial $p(x) \in \mathbb{R}[x]$ whose coefficients are of degree at most D_G in the input G_0 . It is well-known that this may be written as a SDP on matrices of size $O(n^{D_x})$ with $O(m^{D_G})$ affine constraints.

In many cases, this paper included, this problem is solvable with high probability when the input G_0 is sampled from the planted model $\mathbb{P}$. For instance, the evaluation map $p(x) \mapsto p(\mathbf{x})$ is a feasible solution provided that the $\mathbb{P}$ is sufficiently concentrated and these polynomials do not fluctuate too much about their expectations. On the other hand, by taking D_x and D_G sufficiently large it becomes infeasible when the input is drawn from a different distribution. The D_x and D_G necessary measure the hardness of the hypothesis testing problem.

1.1.3. THE LOW-DEGREE LIKELIHOOD RATIO

As was first discovered in a series of works in the sum-of-squares literature (Barak et al., 2019; Hopkins and Steurer, 2017; Hopkins et al., 2017; Hopkins, 2018), analyzing the *low-degree likelihood ratio* gives predictions of computational hardness that match widely-believed conjectures for many hypothesis testing problems (as corroborated by the various other heuristics mentioned above). In essence, this method takes low-degree polynomials as a proxy for all polynomial-time algorithms, and analyzes whether any low-degree polynomial can distinguish two distributions $\mathbb{Q}$ and $\mathbb{P}$ (with the same interpretations as above) as $n \rightarrow \infty$.

The key to making this analysis tractable is to choose the correct soft notion of “successfully distinguishing.” This is done by considering the maximization

$$\begin{aligned} & \text{maximize} && \frac{\mathbb{E}_{\mathbf{x} \sim \mathbb{P}} p(\mathbf{x})}{(\mathbb{E}_{\mathbf{x} \sim \mathbb{Q}} p(\mathbf{x})^2)^{1/2}} \\ & \text{such that} && p \neq 0 \text{ is a polynomial of degree } \leq D. \end{aligned} \tag{13}$$

Thus we try to maximize $p(\mathbf{x})$ ’s expectation under $\mathbb{P}$, while keeping its typical size under $\mathbb{Q}$ modest.

We note that, if we did not restrict p to be a low-degree polynomial, then the optimal p would equal the classical *likelihood ratio*, $L := \frac{d\mathbb{P}}{d\mathbb{Q}}$. In absence of computational constraints, thresholding L gives an optimal test between $\mathbb{P}$ and $\mathbb{Q}$ in the sense of minimizing error probabilities, as shown in the classical Neyman-Pearson lemma (Neyman and Pearson, 1933). Moreover, the value of the above problem would be the norm $\|L\|$ of L in $L^2(\mathbb{Q})$. If that norm is bounded as $n \rightarrow \infty$, then $\mathbb{P}$ and $\mathbb{Q}$ cannot be distinguished w.h.p. by *any* test, by an application of Le Cam’s second moment method for contiguity (see Le Cam and Yang (2012); Kunisky et al. (2019) for further exposition).

With the further constraint to low-degree polynomials, the result is similar: the optimal p is the aforementioned low-degree likelihood ratio, the orthogonal projection of L to the subspace of degree- D polynomials in $L^2(\mathbb{Q})$, which we denote $L^{\leq D}$. The value of the problem is its norm, $\|L^{\leq D}\|$. We again consider whether, as $n \rightarrow \infty$, for $D = D(n)$ slowly growing with n , this norm diverges or remains bounded. If it diverges, we expect that low-degree polynomials can distinguish

$\mathbb{P}$ from $\mathbb{Q}$ w.h.p. (at an intuitive level, the algorithm we have in mind is thresholding $L^{\leq D}$); if it remains bounded, we conclude that low-degree polynomials cannot distinguish $\mathbb{P}$ from $\mathbb{Q}$ w.h.p. in this particular sense, and therefore expect that no polynomial-time algorithm can do so either.

The precise scaling of $D(n)$ relates to the efficiency of algorithms that the heuristic pertains to—higher degree polynomials describe more time-consuming computations. However, while constant-degree polynomials may be evaluated in polynomial time, some other polynomial-time computations require slightly higher degree polynomials to express. A crucial example is approximating the spectral norm of a matrix with dimensions polynomial in n , which requires a polynomial of degree $\Theta(\log n)$ to approximate accurately. Taking this into account, a low-degree lower bound with $D(n) \gg \log n$ is taken as evidence that no polynomial-time test exists. Similarly, we view a low-degree lower bound with $D(n) \gg n^\delta$ for some $\delta \in (0, 1)$ as suggesting a lower bound against tests with runtime $O(\exp(n^\delta))$. See [Kunisky et al. \(2019\)](#) for further discussion.

1.1.4. A CAVEAT: XOR-SAT

In the planted 3-XOR-SAT problem, we are given m clauses of the form $x_i x_j x_k = b$ for some choice of $i, j, k \in [n]$ and $b \in \{\pm 1\}$. The goal is to distinguish the case where the clauses are completely random from the case where there is a planted assignment of $\{\pm 1\}$ values to the variables $x_1, \dots, x_n$ such that all clauses are satisfied. This is a notable counterexample for many heuristics for average-case hardness, including sum-of-squares and all three methods mentioned above (see e.g., Lecture 3.2 of [Barak and Steurer \(2016\)](#) or Chapter 18 of [Mezard and Montanari \(2009\)](#)). These heuristics predict that the distinguishing task is possible in polynomial time only when $m \gtrsim n^{3/2}$, whereas in reality the problem is much easier: Gaussian elimination can be used to decide with certainty whether or not there is a satisfying assignment. However, if we change the planted distribution so that only a $1 - \varepsilon$ fraction of the clauses are satisfied, Gaussian elimination breaks down and the best known algorithms indeed require $m \gtrsim n^{3/2}$. In this sense, the above heuristics seem to predict the threshold for “robust” algorithms. We expect that a similar phenomenon is at play in the eSBM model: while there exist algorithms to solve the problem by exploiting brittle algebraic structure in the eigenvectors (see Remark 3), we conjecture (Conjecture 8) that the threshold predicted by the heuristics is the correct computational threshold for a *noisy* version of the problem.

2. Main Results

In this section we state our main results for the graph model, with additional results on the Gaussian model deferred to Section A.

2.1. Notation

We use standard asymptotic notation such as $o(\cdot)$ and $O(\cdot)$; unless stated otherwise, this always pertains to the limit $n \rightarrow \infty$ with other parameters (such as d, k, η) held fixed. We use e.g., $o_d(\cdot)$ or $O_d(\cdot)$ when considering the double limit $n \rightarrow \infty$ followed by $d \rightarrow \infty$; for example, $f(n, d) = o_d(g(n, d))$ means that for any $\varepsilon > 0$ there exists $d_0 > 0$ such that for all $d \geq d_0$ there exists $n_0 > 0$ such that for all $n \geq n_0$ we have $|f(n, d)/g(n, d)| \leq \varepsilon$. An event occurs *with high probability* (w.h.p.) if it has probability $1 - o(1)$. We write $f \sim g$ to mean $f/g \rightarrow 1$.

Throughout, all graphs are assumed to be *simple*—i.e., without self-loops and multiple edges—unless stated otherwise. A d -regular graph has degree d at every vertex. $\mathcal{G}_{n,d}$ denotes the uniform distribution over d -regular n -vertex graphs.

We will use I_k denote the $k \times k$ identity matrix and J_k for the $k \times k$ all-ones matrix; $e_1, e_2, \dots$ are the standard unit basis vectors, and $\mathbb{1}$ will denote the all-ones vector. For matrices, $\|\cdot\|$ denotes the operator (spectral) norm and $\|\cdot\|_F$ denotes the Frobenius norm. We use $\mathbf{1}[A]$ for the $\{0, 1\}$ -valued indicator of an event A . We write $[k] = \{1, 2, \dots, k\}$ and $\mathbb{N} = \{0, 1, 2, \dots\}$.

2.2. Setup and Definitions

We consider a general framework that captures the problem of certifying bounds on max- k -cut in a random graph, as well as a Gaussian variant of this problem.

Definition 4 For a labeling $\sigma : [n] \rightarrow [k]$, the associated partition matrix $P = P^{(\sigma)} \in \mathbb{R}^{n \times n}$ is given by

$$P_{i,j} = \begin{cases} 1 & \sigma(i) = \sigma(j), \\ -1/(k-1) & \sigma(i) \neq \sigma(j). \end{cases}$$

A basic fact is that $P \succeq 0$ and $\text{Rank}(P) = k - 1$. This can be seen by realizing P as the Gram matrix of a certain collection of vectors: assign each $u \in [n]$ one of the k unit vectors in $\mathbb{R}^{k-1}$ pointing to the corners of a simplex, according to its label $\sigma(u) \in [k]$.

Let Π be the set of all partition matrices. For a given matrix $A \in \mathbb{R}^{n \times n}$ and a given k , we will be interested the problem of algorithmically certifying an upper bound on the value

$$\Gamma_k(A) := \max_{P \in \Pi} \langle P, -A \rangle. \quad (14)$$

If $A_G \in \{0, 1\}^{n \times n}$ is the adjacency matrix of a graph $G = (V, E)$ (defined with $(A_G)_{i,i} = 0$), then $\langle P^{(\sigma)}, -A_G \rangle = 2(k \cdot m_\sigma(G) - |E|)/(k-1)$, where $m_\sigma(G)$ is the number of monochromatic edges of G under the labeling σ . As a result,

$$\text{MC}_k(G) = \frac{k-1}{k} \left(1 + \frac{\Gamma_k(A_G)}{2|E|} \right). \quad (15)$$

Thus, an upper bound on $\Gamma_k(A_G)$ translates to an upper bound on $\text{MC}_k(G)$, which in turn can be used to refute k -colorability. We now formally define the certification task.

Definition 5 Let $\mathbb{Q} = \mathbb{Q}_n$ be a sequence of distributions $\mathbb{R}^{n \times n}$ and let $\mathcal{A} = \mathcal{A}_n : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}$ be a sequence of algorithms. We say that $\mathcal{A}$ certifies the upper bound B on Γ_k over $\mathcal{A} \sim \mathbb{Q}$ if both of the following hold:

- (i) for every $A \in \mathbb{R}^{n \times n}$, $\mathcal{A}(A) \geq \Gamma_k(A)$, and
- (ii) if $A \sim \mathbb{Q}_n$ then $\mathcal{A}_n(A) \leq B$ with probability $1 - o(1)$.

Crucially, $\mathcal{A}(A)$ must *always* be a valid upper bound, even if A is atypical under $\mathbb{Q}$. In exponential time it is possible to compute $\Gamma_k(A)$ exactly and thus achieve perfect certification; we are interested instead in polynomial-time certification procedures.

There is a simple spectral approach (Hoffman, 1970) to certifying bounds on $\Gamma_k(A)$. Let $\lambda_{\min} = \lambda_{\min}(A)$ denote the minimum eigenvalue of A . For any $P \in \Pi$ we have $P \succeq 0$ and $A - \lambda_{\min}I \succeq 0$, so

$$0 \leq \langle P, A - \lambda_{\min}I \rangle = \langle P, A \rangle - \lambda_{\min}n$$

and $-\lambda_{\min}n$ is an efficiently-computable upper bound on $\Gamma_k(A)$. Our main results give evidence that it is computationally hard to improve upon this spectral bound when A is drawn from certain distributions $\mathbb{Q}$: (i) random d -regular graphs, and (ii) Gaussian matrices.

2.3. Random Regular Graphs

We will be concerned with the following task of distinguishing two distributions, also called *hypothesis testing* or *detection*.

Definition 6 Let $\mathbb{P}_n$ and $\mathbb{Q}_n$ be probability measures on the same space Ω_n . We say that an algorithm $t_n : \Omega_n \rightarrow \{P, Q\}$ distinguishes $\mathbb{P}_n$ and $\mathbb{Q}_n$ with high probability, or (equivalently) achieves strong detection (between $\mathbb{P}_n$ and $\mathbb{Q}_n$) if

$$\mathbb{P}_n[t_n(\mathbf{x}) = Q] + \mathbb{Q}_n[t_n(\mathbf{x}) = P] = o(1)$$

as $n \rightarrow \infty$.

Our results for random regular graphs are conditional on a conjecture regarding computational hardness of detection in a noisy variant of the equitable stochastic block model (eSBM). The extra noise is crucial, as discussed in Remark 3 and Section 1.1.4. Specifically, the noise takes the form of “rewiring” a small constant fraction of the edges as follows.

Definition 7 (Noise Operator) If $G = (V, E)$ is a d -regular n -vertex graph and $\delta > 0$, let $T_\delta(G)$ denote the random d -regular graph obtained from G by making $\lfloor \delta n \rfloor$ ‘swaps.’ That is, repeatedly choose a pair of distinct edges $(i, j), (k, \ell)$ uniformly at random⁷ conditioned on the following events: $i \neq k, j \neq \ell, (i, k) \notin E$, and $(j, \ell) \notin E$. Remove edges (i, j) and (k, ℓ) , and add edges (i, k) and (j, ℓ) .

Recall the definition of the eSBM (Definition 1) and the associated threshold $d_{\text{KS}}^{\text{eq}}(\eta)$ defined in (9).

Conjecture 8 (eSBM Conjecture) Let $\tilde{\mathcal{G}}_{n,d,k,\eta,\delta}^{\text{eq}}$ denote the distribution over d -regular n -vertex graphs given by $T_\delta(\mathbf{G})$ where $\mathbf{G} \sim \mathcal{G}_{n,d,k,\eta}^{\text{eq}}$. Suppose $\delta > 0$, $\eta \in [-\frac{1}{k-1}, 1]$, $k \geq 2$, and $d < d_{\text{KS}}^{\text{eq}}(\eta)$ are all fixed. Also suppose $k \mid (1 - \eta)d$ so that the equitable model is defined for an infinite sequence of values of n . Then, there exists no polynomial-time algorithm that with high probability distinguishes $\tilde{\mathcal{G}}_{n,d,k,\eta,\delta}^{\text{eq}}$ from $\mathcal{G}_{n,d}$ as $n \rightarrow \infty$ (in the sense of Definition 6).

The above conjecture implies the following result on hardness of certifying bounds on max- k -cut.

Theorem 9 Assume the eSBM conjecture (Conjecture 8) holds. Fix $d > 0$ and $k \geq 2$ and let $\eta \in [-\frac{1}{k-1}, 0]$ be such that $|\eta| < \frac{2\sqrt{d-1}}{d}$ and $k \mid (1 - \eta)d$. Then for any $\epsilon > 0$, no polynomial-time algorithm can certify, in the sense of Definition 5, the upper bound $nd(|\eta| - \epsilon)$ on Γ_k over $\mathcal{G}_{n,d}$. Equivalently, no polynomial-time algorithm can certify for $\mathbf{G} \sim \mathcal{G}_{n,d}$ the bound

$$\text{MC}_k(\mathbf{G}) \leq \frac{k-1}{k} (1 + |\eta| - \epsilon). \quad (16)$$

7. Here, the ordering of the tuples matters, so e.g., (i, j) and (j, i) should be chosen with equal probability.

The proof idea is described in the Introduction; the full details are given in Section B.1. The effects of the integrality condition $k \mid (1 - \eta)d$ are discussed in Section A.2; this condition can be ignored when $d \gg k^2$.

Remark 10 (Near-Coloring) *An important special case where Theorem 9 is tight is $\eta = -\frac{1}{k-1}$, corresponding to refutation of near colorability. Here the integrality condition $k \mid (1 - \eta)d$ reduces to $(k - 1) \mid d$. Thus for any d, k satisfying $(k - 1) \mid d$ and $k > 1 + \frac{d}{2\sqrt{d-1}}$ we have that for any $\epsilon > 0$, no polynomial-time algorithm can certify $\text{MC}_k(\mathbf{G}) \leq 1 - \epsilon$. There is an infinite sequence of (d, k) values with $k \sim \frac{1}{2}\sqrt{d}$ for which the conditions on d, k are satisfied, namely $d = 4(k - 1)(k - 2)$ for all $k > 2$. Thus, our result is asymptotically tight, matching the spectral algorithm in the double limit $n \rightarrow \infty$ followed by $d \rightarrow \infty$.*

Remark 11 (Exact Coloring) *As stated, Theorem 9 only shows hardness of refuting a near-coloring, as opposed to an exact coloring. While we expect a similar hardness result to hold for exact coloring, this does not follow from the eSBM conjecture because the noise operator $\mathbf{T}_\delta$ prevents us from planting an exact coloring. Hardness of refuting exact coloring would follow from a variant of the eSBM conjecture where the noise operator only makes swaps that do not affect the value of the planted cut.*

Our next two results give concrete evidence for the eSBM conjecture using two different heuristics (discussed in Section 1.1): stability of belief propagation, and the local statistics hierarchy. Both of these methods predict $d_{\text{KS}}^{\text{eq}}(\eta)$ as the computational threshold for the eSBM.

Theorem 12 (Informal; Kesten–Stigum threshold of eSBM) *For $k \geq 4$, the Kesten–Stigum threshold of the eSBM with parameters (n, d, k, η) , defined as the smallest number $d_{\text{KS}}^{\text{eq}}$ so that, for all $d > d_{\text{KS}}^{\text{eq}}$, the “uninformative” fixed point of the belief propagation iteration is unstable, is given by*

$$d_{\text{KS}}^{\text{eq}} = d_{\text{KS}}^{\text{eq}}(\eta) := \frac{2}{\eta^2} \left(1 + \sqrt{1 - \eta^2} \right).$$

Further details, as well as the full analysis leading to the above result, can be found in Section C.

Theorem 13 (Local Statistics analysis of eSBM) *If $d > d_{\text{KS}}^{\text{eq}}(\eta)$, then there exist D sufficiently large and $\delta > 0$ so that the degree- $(2, D)$ Local Statistics algorithm with error tolerance δ can distinguish $\mathcal{G}_{n,d}$ and $\mathcal{G}_{n,d,k,\eta}^{\text{eq}}$ with high probability. If $d \leq d_{\text{KS}}^{\text{eq}}(\eta)$, no such D and δ exist.*

Further details on the Local Statistics hierarchy, and the proof of Theorem 13, can be found in Section D.

Acknowledgments

Some of this work was done while ASB was with the Department of Mathematics at the Courant Institute of Mathematical Sciences, and the Center for Data Science, at New York University; and partially supported by NSF grants DMS-1712730 and DMS-1719545, and by a grant from the Sloan Foundation. JB was supported by the NSF Graduate Research Fellowship Program under Grant DGE-1752814. DK was partially supported by NSF grants DMS-1712730 and DMS-1719545. CM was partially supported by NSF grant IIS-1838251. ASW was partially supported by NSF grant DMS-1712730 and by the Simons Collaboration on Algorithms and Geometry.

References

- Emmanuel Abbe and Colin Sandon. Achieving the KS threshold in the general stochastic block model with linearized acyclic belief propagation. In *Advances in Neural Information Processing Systems*, pages 1334–1342, 2016.
- N. Alon, M. Krivelevich, and B. Sudakov. Finding a large hidden clique in a random graph. *Random Structures & Algorithms*, 13:457–466, 1998.
- Antonio Auffinger, Gérard Ben Arous, and Jiří Černý. Random matrices and complexity of spin glasses. *Communications on Pure and Applied Mathematics*, 66(2):165–201, 2013.
- Jinho Baik and Jack W Silverstein. Eigenvalues of large sample covariance matrices of spiked population models. *Journal of multivariate analysis*, 97(6):1382–1408, 2006.
- Jinho Baik, Gérard Ben Arous, and Sandrine Péché. Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *The Annals of Probability*, 33(5):1643–1697, 2005.
- Afonso S Bandeira, Dmitriy Kunisky, and Alexander S Wein. Computational hardness of certifying bounds on constrained pca problems. In *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.
- Jess Banks, Robert Kleinberg, and Cristopher Moore. The Lovász theta function for random regular graphs and community detection in the hard regime. *arXiv preprint arXiv:1705.01194*, 2017.
- Jess Banks, Sidhanth Mohanty, and Prasad Raghavendra. Local statistics, semidefinite programming, and community detection. *arXiv preprint arXiv:1911.01960*, 2019.
- Boaz Barak and David Steurer. Proofs, beliefs, and algorithms through the lens of sum-of-squares. *Course notes: <http://www.sumofsquares.org/public/index.html>*, 2016.
- Boaz Barak, Samuel Hopkins, Jonathan Kelner, Pravesh K Kothari, Ankur Moitra, and Aaron Potechin. A nearly tight sum-of-squares lower bound for the planted clique problem. *SIAM Journal on Computing*, 48(2):687–735, 2019.
- Paolo Barucca. Spectral partitioning in equitable graphs. *Phys. Rev. E*, 95:062310, 2017.
- Quentin Berthet and Philippe Rigollet. Complexity theoretic lower bounds for sparse principal component detection. In *Conference on Learning Theory*, pages 1046–1066, 2013.
- Charles Bordenave and Benoît Collins. Eigenvalues of random lifts and polynomials of random permutation matrices. *Annals of Mathematics*, 190(3):811–875, 2019.
- Matthew Brennan, Guy Bresler, and Wasim Huleihel. Reducibility and computational lower bounds for problems with planted sparse structure. In *Conference On Learning Theory*, pages 48–166, 2018.
- Gerandy Brito, Ioana Dumitriu, Shirshendu Ganguly, Christopher Hoffman, and Linh V Tran. Recovery and rigidity in a regular stochastic block model. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 1589–1601. Society for Industrial and Applied Mathematics, 2016.

- Mireille Capitaine, Catherine Donati-Martin, and Delphine Féral. The largest eigenvalues of finite rank deformation of large wigner matrices: convergence and nonuniversality of the fluctuations. *The Annals of Probability*, 37(1):1–47, 2009.
- Amin Coja-Oghlan. The Lovász number of random graphs. In *Approximation, Randomization, and Combinatorial Optimization.. Algorithms and Techniques*, pages 228–239. Springer, 2003.
- Amin Coja-Oghlan, Charilaos Efthymiou, and Samuel Hetterich. On the chromatic number of random regular graphs. *Journal of Combinatorial Theory, Series B*, 116:367–439, 2016.
- Aurelien Decelle, Florent Krzakala, Cristopher Moore, and Lenka Zdeborová. Inference and phase transitions in the detection of modules in sparse networks. *Physical Review Letters*, 107(6):065701, 2011a.
- Aurelien Decelle, Florent Krzakala, Cristopher Moore, and Lenka Zdeborová. Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. *Physical Review E*, 84(6):066106, 2011b.
- Amir Dembo, Andrea Montanari, and Subhabrata Sen. Extremal cuts of sparse random graphs. *The Annals of Probability*, 45(2):1190–1217, 2017.
- Yash Deshpande, Emmanuel Abbe, and Andrea Montanari. Asymptotic mutual information for the balanced binary stochastic block model. *Information and Inference: A Journal of the IMA*, 6(2):125–170, 2017.
- David L Donoho, Arian Maleki, and Andrea Montanari. Message-passing algorithms for compressed sensing. *Proceedings of the National Academy of Sciences*, 106(45):18914–18919, 2009.
- Josep Díaz, Norman Do, Maria Serna, and Nicholas Wormald. Bounds on the max and min bisection of random cubic and random 4-regular graphs. *Theoretical Computer Science*, 307:531–547, 10 2003.
- Vitaly Feldman, Elena Grigorescu, Lev Reyzin, Santosh S Vempala, and Ying Xiao. Statistical algorithms and a lower bound for detecting planted cliques. *Journal of the ACM (JACM)*, 64(2):1–37, 2017.
- Delphine Féral and Sandrine Péché. The largest eigenvalue of rank one deformation of large wigner matrices. *Communications in mathematical physics*, 272(1):185–228, 2007.
- Joel Friedman. A proof of Alon’s second eigenvalue conjecture. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 720–724, 2003.
- David Gamarnik and Ilias Zadik. High-dimensional regression with binary coefficients. Estimating squared error and a phase transition. *arXiv preprint arXiv:1701.04455*, 2017.
- G. R. Grimmett and C. J. H. McDiarmid. On colouring random graphs. *Math. Proc. Camb. Phil. Soc.*, 77:313–324, 1975.
- Francesco Guerra. Broken replica symmetry bounds in the mean field spin glass model. *Communications in mathematical physics*, 233(1):1–12, 2003.

- Alan J Hoffman. On eigenvalues and colorings of graphs. *Graph Theory and its Applications*, 1970.
- Samuel Hopkins. *Statistical Inference and the Sum of Squares Method*. PhD thesis, Cornell University, 2018.
- Samuel B Hopkins and David Steurer. Bayesian estimation from few samples: community detection and related problems. *arXiv preprint arXiv:1710.00264*, 2017.
- Samuel B Hopkins, Pravesh K Kothari, Aaron Potechin, Prasad Raghavendra, Tselil Schramm, and David Steurer. The power of sum-of-squares for detecting hidden structures. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 720–731. IEEE, 2017.
- Aukosh Jagannath, Justin Ko, and Subhabrata Sen. Max κ -cut and the inhomogeneous potts spin glass. *Annals of Applied Probability*, 28(3):1536–1572, 2018.
- Richard M. Karp. The probabilistic analysis of some combinatorial search algorithms. 1976.
- Richard M. Karp. Combinatorics, complexity, and randomness. 1986.
- Florent Krzakala and Lenka Zdeborová. Hiding quiet solutions in random constraint satisfaction problems. *Physical review letters*, 102(23):238701, 2009.
- Ludek Kucera. Expected complexity of graph partitioning problems. *Discrete Applied Mathematics*, 57:193–212, 1995.
- Dmitriy Kunisky, Alexander S Wein, and Afonso S Bandeira. Notes on computational hardness of hypothesis testing: Predictions using the low-degree likelihood ratio. *arXiv preprint arXiv:1907.11636*, 2019.
- Monique Laurent. Sums of squares, moment matrices and optimization over polynomials. In *Emerging applications of algebraic geometry*, pages 157–270. Springer, 2009.
- Lucien Le Cam and Grace Lo Yang. *Asymptotics in statistics: some basic concepts*. Springer Science & Business Media, 2012.
- Thibault Lesieur, Florent Krzakala, and Lenka Zdeborová. MMSE of probabilistic low-rank matrix estimation: Universality with respect to the output channel. In *2015 53rd Annual Allerton Conference on Communication, Control, and Computing*, pages 680–687. IEEE, 2015a.
- Thibault Lesieur, Florent Krzakala, and Lenka Zdeborová. Phase transitions in sparse PCA. In *2015 IEEE International Symposium on Information Theory (ISIT)*, pages 1635–1639. IEEE, 2015b.
- Laurent Massoulié. Community detection thresholds and the weak ramanujan property. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 694–703, 2014.
- Brendan D McKay. The expected eigenvalue distribution of a large regular graph. *Linear Algebra and its Applications*, 40:203–216, 1981.
- Marc Mezard and Andrea Montanari. *Information, physics, and computation*. Oxford University Press, 2009.

- Sidhanth Mohanty, Prasad Raghavendra, and Jeff Xu. Lifting sum-of-squares lower bounds: Degree-2 to degree-4. *arXiv preprint arXiv:1911.01411*, 2019.
- Andrea Montanari and Subhabrata Sen. Semidefinite programs on sparse random graphs and their application to community detection. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 814–827. ACM, 2016.
- Elchanan Mossel, Joe Neeman, and Allan Sly. A proof of the block model threshold conjecture. *Combinatorica*, 38(3):665–708, 2018.
- Raj Rao Nadakuditi and M. E. J. Newman. Graph spectra and the detectability of community structure in networks. *Phys. Rev. Lett.*, 108:188701, May 2012.
- M. E. J. Newman and Travis Martin. Equitable random graphs. *Phys. Rev. E*, 90:052824, 2014.
- Jerzy Neyman and Egon Sharpe Pearson. IX. on the problem of the most efficient tests of statistical hypotheses. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 231(694-706):289–337, 1933.
- Dmitry Panchenko. The Parisi ultrametricity conjecture. *arXiv preprint arXiv:1112.1003*, 2011.
- Dmitry Panchenko. *The Sherrington-Kirkpatrick model*. Springer Science & Business Media, 2013.
- Giorgio Parisi. Infinite number of order parameters for spin-glasses. *Physical Review Letters*, 43(23):1754, 1979.
- Steven Roman. *The umbral calculus*. Springer, 2005.
- Subhabrata Sen. Optimization on sparse random hypergraphs and spin glasses. *Random Structures & Algorithms*, 53(3):504–536, 2018.
- David Sherrington and Scott Kirkpatrick. Solvable model of a spin-glass. *Physical review letters*, 35(26):1792, 1975.
- Sasha Sodin. Random matrices, nonbacktracking walks, and orthogonal polynomials. *Journal of Mathematical Physics*, 48(12):123503, 2007.
- Patrick Solé. Spectra of regular graphs and hypergraphs and orthogonal polynomials. *European Journal of Combinatorics*, 17(5):461–477, 1996.
- Gabor Szeg. *Orthogonal polynomials*, volume 23. American Mathematical Soc., 1939.
- Michel Talagrand. The Parisi formula. *Annals of mathematics*, pages 221–263, 2006.
- Lenka Zdeborová and Stefan Boettcher. A conjecture on the maximum cut and bisection width in random regular graphs. *Journal of Statistical Mechanics: Theory and Experiment*, 2010(02):P02020, 2010.
- Lenka Zdeborová and Florent Krzakala. Quiet planting in the locked constraint satisfaction problems. *SIAM Journal on Discrete Mathematics*, 25(2):750–770, 2011.

Appendix A. Additional Results

A.1. Outline of Paper

In Section A.2 we give some additional discussion regarding the integrality condition in Theorem 9. In Section A.3 we state additional results for a Gaussian variant of the k -cut problem; these results are analogous to those for the graph model in Section 2. One by-product of these results are a number of new results on the low-degree likelihood ratio for spiked models and the stochastic block model (see Section A.3.2). The rest of the paper is devoted to proofs. In Section B we give formal reductions from the detection problem to the certification problem, for both the graph and Gaussian models. Sections C and D contain evidence for the eSBM conjecture based on belief propagation and the local statistics hierarchy, respectively. Sections E,F,G contain results on the low-degree likelihood ratio.

A.2. Discussion of Integrality Condition

The integrality constraints in the eSBM (Definition 1) and in Theorem 9 may seem somewhat mystifying. In this section we will try to shed some light on the limits of our method for quiet planting, and how they compare to the power of spectral refutation. The integrality condition $k \mid (1 - \eta)d$ from Theorem 9 is required so that the equitable block model exists for an infinite sequence of values of n . This condition can be written as $\eta d \equiv d \pmod{k}$, which means η is constrained to lie in a certain grid of spacing $\frac{k}{d}$ within the interval $[-\frac{1}{k-1}, 0]$. In order to extract the strongest possible hardness result from Theorem 9, the goal is to pick η as to maximize $|\eta|$ subject to the integrality condition and $|\eta| < \frac{2\sqrt{d-1}}{d}$. If $d \gg k^2$ then we have $\frac{k}{d} \ll \frac{2\sqrt{d-1}}{d}$ and so the grid allows η to be chosen extremely close to $\frac{2\sqrt{d-1}}{d}$; this means the effects of the integrality condition are negligible and we essentially obtain the ideal hardness result (10) which matches the spectral algorithm. In particular, our results are tight in the regime $d \rightarrow \infty$ with k fixed.

As discussed in Remark 10, we also get tight results in the particular case $\eta = -\frac{1}{k-1}$ corresponding to near-coloring, because the integrality condition simplifies in a helpful way.

On the other hand, for small degree the integrality conditions can impose significant limitations. A particularly severe example is that of $k = 2$ and $d = 3$, which corresponds to Max-Cut in a random 3-regular graph. The spectral bound corresponds to

$$\text{MC}_2(\mathbf{G}) \leq \frac{1}{2} \left(1 + \frac{2\sqrt{3-1}}{3} \right) \approx 0.97. \quad (17)$$

On the other hand, the integrality condition from Theorem 9 requires that $2 \mid 3(1 - \eta)$, in addition to the condition $|\eta| < \frac{2\sqrt{3-1}}{3}$. Since $3 \left(1 + \frac{2\sqrt{3-1}}{3} \right) \approx 5.83$ the largest $|\eta|$ admissible is such that that $3(1 - \eta) = 4$, corresponding to $\eta = \frac{1}{3}$. This means that Theorem 9 only addresses certification below

$$\text{MC}_2(\mathbf{G}) \leq \frac{1}{2} \left(1 + \frac{1}{3} \right) = \frac{2}{3} \approx 0.67.$$

This is unfortunate since random 3-regular graphs have $\text{MC}_2(\mathbf{G}) > 0.88$ with high probability (Díaz et al., 2003), and calculations from statistical physics (Zdeborová and Boettcher, 2010) suggest the even larger value $\text{MC}_2(\mathbf{G}) \approx 0.92$. That is, using the equitable block model to plant a 2-cut in a 3-regular graph, by giving each vertex one neighbor in its own group and two neighbors in the other group, plants a cut that is smaller than naturally arising cuts in random 3-regular graphs.

In general, aside from the special case of near-coloring (Remark 10), our results are most compelling when k is small and d is large, so that the integrality condition in Theorem 9 does not create a significant gap from the limiting threshold $|\eta| = \frac{2\sqrt{d-1}}{d}$. We believe the investigation of quiet plantings in small-degree graphs to be an interesting direction of future research.

A.3. The Gaussian k -Cut Model

A.3.1. HARDNESS OF CERTIFICATION

In this section we discuss a Gaussian analogue of the coloring problem, namely the problem of certifying upper bounds on $\Gamma_k(\mathbf{W})$ where $\mathbf{W}$ is a GOE matrix defined as follows.

Definition 14 *The Gaussian orthogonal ensemble is the following distribution $\text{GOE}(n)$ over random matrices: $\mathbf{W} \sim \text{GOE}(n)$ is symmetric ($\mathbf{W}_{u,v} = \mathbf{W}_{v,u}$) with diagonal entries $\mathbf{W}_{u,u} \sim \mathcal{N}(0, 2/n)$ and off-diagonal entries $\mathbf{W}_{u,v} \sim \mathcal{N}(0, 1/n)$, where the values $\{\mathbf{W}_{u,v} : u \leq v\}$ are independent.*

It is well known that (as $n \rightarrow \infty$) the eigenvalues of $\mathbf{W} \sim \text{GOE}(n)$ follow the Wigner semicircle law supported on $[-2, 2]$, and in particular, $\lambda_{\min}(\mathbf{W}) \rightarrow -2$ almost surely. Thus, the spectral approach (see Section 2.2) certifies the upper bound $\Gamma_k(\mathbf{W}) \leq (2 + o(1))n$. Our main results give rigorous evidence in support of the following conjecture, which states that improving upon this spectral bound requires fully exponential time.

Conjecture 15 *For any constants $k \geq 2$, $\epsilon > 0$ and $\delta > 0$, there is no algorithm of runtime $\exp(O(n^{1-\delta}))$ that certifies the upper bound $(2 - \epsilon)n$ on $\Gamma_k(\mathbf{W})$ over $\mathbf{W} \sim \text{GOE}(n)$.*

Our evidence for this conjecture can be seen as a generalization of [Bandeira et al. \(2020\)](#), which handles the $k = 2$ case. In Section B.2 we give a reduction from a certain hypothesis testing problem to the certification problem in question. In analogy to our results on coloring, this reduction can be seen as constructing a planted distribution that directly plants an eigenspace with no pushout effect; this planting is computationally quiet, conditional on hardness of the testing problem. The testing problem is a particular instance of the spiked Wishart model with a rank- $(k - 1)$ negative spike. In Section A.3.2 below, we state results analyzing the low-degree likelihood ratio (see Section 1.1.3) for this model; these results suggest that fully exponential time is required (in the appropriate parameter regime).

As a by-product of our analysis, we give new bounds on the low-degree likelihood ratio for a wide class of multi-spiked matrix models (both Wigner and Wishart), which may be of independent interest. These results also extend to certain binary-values analogues of these problems, including (a variant of) the stochastic block model. These results are also discussed in Section A.3.2 below.

A.3.2. LOW-DEGREE HARDNESS FOR GENERAL SPIKED MODELS

We consider general variants of the spiked Wigner and Wishart models, defined as follows.

Definition 16 *Let $\mathcal{X} = (\mathcal{X}_n)$ be a probability measure over $\mathbb{R}_{\text{sym}}^{n \times n}$, the symmetric $n \times n$ matrices. The general spiked Wigner model with spike prior $\mathcal{X}$ and signal-to-noise ratio $\lambda \in \mathbb{R}$ is specified by the following null and planted distributions over $\mathbb{R}_{\text{sym}}^{n \times n}$.*

- Under $\mathbb{Q}$, draw $\mathbf{Y} \sim \text{GOE}(n)$.
- Under $\mathbb{P}$, let $\mathbf{Y} = \lambda \mathbf{X} + \mathbf{W}$ where $\mathbf{X} \sim \mathcal{X}$ and $\mathbf{W} \sim \text{GOE}(n)$, independently.

Definition 17 Let $\mathcal{X} = (\mathcal{X}_n)$ be a probability measure over $\mathbb{R}_{\text{sym}}^{n \times n}$. The general spiked Wishart model with spike prior $\mathcal{X}$, signal-to-noise ratio $\beta > -1$, and number of samples $N \in \mathbb{N}$, is specified by the following null and planted distributions over $(\mathbf{y}_1, \dots, \mathbf{y}_N) \in (\mathbb{R}^n)^N$.

- Under $\mathbb{Q}$, draw $\mathbf{y}_u \sim \mathcal{N}(0, I_n)$ independently for $u \in [N]$.
- Under $\mathbb{P}$, first draw $\widetilde{\mathbf{X}} \sim \mathcal{X}$ and define

$$\mathbf{X} = \begin{cases} \widetilde{\mathbf{X}} & \text{if } \beta \widetilde{\mathbf{X}} \succ -I_n, \\ 0 & \text{else.} \end{cases} \quad (18)$$

Then draw $\mathbf{y}_u \sim \mathcal{N}(0, I_n + \beta \mathbf{X})$ independently for $u \in [N]$.

The purpose of (18) is to ensure that $I_n + \beta \mathbf{X}$ is a valid covariance matrix. We will consider priors for which the first case of (18) occurs with high probability. Specifically, we focus on the following class of priors that are PSD with constant rank.

Definition 18 Fix an integer $k \geq 1$. Let π be a probability measure supported on a bounded subset of $\mathbb{R}^k$, satisfying $\mathbb{E}[\pi] = 0$ and $\|\text{Cov}(\pi)\| = 1$. Let $\mathcal{X}(\pi)$ denote the spike prior which outputs $\mathbf{X} = \frac{1}{n} \mathbf{U} \mathbf{U}^\top$, where $\mathbf{U}$ is $n \times k$ with each row distributed according to π . (We do not allow π to depend on n .)

It is well known in random matrix theory that polynomial-time detection in the Wigner model is possible when $|\lambda| > 1$, by thresholding the maximum (or minimum if $\lambda < 0$) eigenvalue of $\mathbf{Y}$ (Féral and Péché, 2007; Capitaine et al., 2009). Similarly, poly-time detection in the Wishart model is possible when $\beta^2 > n/N$, by thresholding the maximum or minimum eigenvalue of the sample covariance matrix $\mathbf{Y} = \frac{1}{N} \sum_i \mathbf{y}_i \mathbf{y}_i^\top$ (Baik et al., 2005; Baik and Silverstein, 2006). In the general setting above, we have the following bounds on the norm of the low-degree likelihood ratio $\|L^{\leq D}\|$ (see Section 1.1.3), which suggest that fully exponential time is required to solve the detection problem below this spectral threshold. Our results are consistent with the computational thresholds predicted by Lesieur et al. (2015a), where it was shown that the approximate message passing algorithm fails below the spectral threshold. The proofs are deferred to Section E.

Theorem 19 Fix constants $k \geq 1$ and $\lambda \in \mathbb{R}$. Fix π satisfying the requirements in Definition 18. Consider the general spiked Wigner model with spike prior $\mathcal{X}(\pi)$. If $|\lambda| < 1$ then $\|L^{\leq D}\| = O(1)$ for any $D = o(n/\log n)$.

Theorem 20 Fix constants $k \geq 1$, $\beta > -1$, and $\gamma > 0$. Fix π satisfying the requirements in Definition 18. Consider the general spiked Wishart model with any $N = N_n$ satisfying $n/N \rightarrow \gamma$ as $n \rightarrow \infty$, and with spike prior $\mathcal{X}(\pi)$. If $\beta^2 < \gamma$ then $\|L^{\leq D}\| = O(1)$ for any $D = o(n/\log n)$.

Remark 21 In the setting of Theorem 20, the first case of (18) holds with high probability because $\beta > -1$ and $\|\widetilde{\mathbf{X}}\| \rightarrow \|\text{Cov}(\pi)\| = 1$ (in probability). To see this, write $\|\widetilde{\mathbf{X}}\| = \|\frac{1}{n} \mathbf{U} \mathbf{U}^\top\| = \|\frac{1}{n} \mathbf{U}^\top \mathbf{U}\|$; being an average of n i.i.d. $k \times k$ matrices, $\frac{1}{n} \mathbf{U}^\top \mathbf{U}$ converges in probability to its expectation, which is $\text{Cov}(\pi)$.

We also extend our framework to binary-valued problems. In Proposition 66 we give a general result analyzing the low degree likelihood ratio in binary-valued problems via a comparison to the analogous Gaussian-valued problem. As an application, we study the following variant of the stochastic block model (SBM).

Definition 22 *The stochastic block model with parameters $k \geq 2$, $d > 0$, $\eta \in [-1/(k-1), 1]$ (constants not depending on n) is specified by the following null and planted distributions over n -vertex graphs.*

- Under $\mathbb{Q}$, for every $u < v$ the edge (u, v) occurs independently with probability d/n .
- Under $\mathbb{P}$, each vertex is independently assigned a community label drawn uniformly from $[k]$. Conditioned on these labels, edges occur independently. If u, v belong to the same community then edge (u, v) occurs with probability $(1 + (k-1)\eta)d/n$; otherwise (u, v) occurs with probability $(1 - \eta)d/n$.

Here k is the number of communities, d is the average degree, and η is a signal-to-noise ratio: the planted k -cut cuts a fraction $\frac{k-1}{k}(1-\eta)$ of the edges on average. Known polynomial-time algorithms only succeed at distinguishing $\mathbb{P}$ from $\mathbb{Q}$ above the so-called Kesten–Stigum (KS) threshold, i.e., when $d\eta^2 > 1$ (Massoulié, 2014; Mossel et al., 2018; Abbe and Sandon, 2016). We prove the following in Appendix G.

Theorem 23 *Consider the stochastic block model as in Definition 22 with parameters k, d, η fixed. If $d\eta^2 < 1$ then $\|L^{\leq D}\| = O(1)$ for any $D = o(n/\log n)$.*

Prior work (Hopkins and Steurer, 2017; Hopkins, 2018) has already given a low-degree analysis of this variant of the SBM, showing that the problem is low-degree-hard below the KS bound. Our result offers two advantages: (i) the proof is streamlined, following easily from our general-purpose machinery, without the need for direct combinatorial calculations, and (ii) we bound $\|L^{\leq D}\|$ for D all the way up to $o(n/\log n)$ instead of only $n^{0.01}$. As discussed in Section 1.1.3, item (ii) constitutes evidence that distinguishing $\mathbb{P}$ from $\mathbb{Q}$ in the SBM requires fully exponential time $\exp(n^{1-o(1)})$ below the KS bound.

Appendix B. Reduction from Detection to Certification

In this section we give formal proofs, for both the graph and Gaussian models, that hardness of a particular detection problem implies hardness of certification.

B.1. The Graph Model

We now give the proof of Theorem 9, which shows that hardness of detection in the noisy eSBM model implies hardness of certifying bounds on max- k -cut.

Proof [Proof of Theorem 9] Assume for the sake of contradiction that some algorithm $\mathcal{A}$ certifies the upper bound $\text{MC}_k(\mathbf{G}) \leq \frac{k-1}{k}(1 + |\eta| - \varepsilon)$ when $\mathbf{G} \sim \mathcal{G}_{n,d}$. We will use this to distinguish between $\tilde{\mathcal{G}}_{n,d,k,\eta,\delta}^{\text{eq}}$ and $\mathcal{G}_{n,d}$ for $\delta = \frac{d\varepsilon(k-1)}{5k}$; this contradicts Conjecture 8 because the assumption $|\eta| < \frac{2\sqrt{d-1}}{d}$ implies $d < d_{\text{KS}}^{\text{eq}}(\eta)$. Our detection algorithm takes as input a graph G and outputs $\mathbb{Q}$ (“null”) if $\mathcal{A}(G) \leq \frac{k-1}{k}(1 + |\eta| - \varepsilon)$ and $\mathbb{P}$ (“planted”) otherwise. If $\mathbf{G} \sim \mathcal{G}_{n,d}$ then $\mathcal{A}(\mathbf{G}) \leq$

$\frac{k-1}{k}(1 + |\eta| - \varepsilon)$ with high probability by assumption, and so the distinguisher outputs Q. Now consider the case $\mathbf{G} \sim \tilde{\mathcal{G}}_{n,d,k,\eta,\delta}^{\text{eq}}$. A graph drawn from $\mathcal{G}_{n,d,k,\eta}^{\text{eq}}$ has a planted k -cut of fractional size $\frac{k-1}{k}(1 + |\eta|)$, and the noise operator $\mathbf{T}_\delta$ can remove at most $2\delta n$ edges from this cut. Thus,

$$\mathcal{A}(\mathbf{G}) \geq \text{MC}_k(\mathbf{G}) \geq \frac{k-1}{k}(1 + \eta) - \frac{2\delta n}{|E|}$$

where

$$\frac{2\delta n}{|E|} = \frac{4\delta}{d} = \frac{4}{d} \cdot \frac{d\varepsilon(k-1)}{5k} < \varepsilon \cdot \frac{k-1}{k}.$$

This implies $\mathcal{A}(\mathbf{G}) > \frac{k-1}{k}(1 + |\eta| - \varepsilon)$ and so the distinguisher outputs P. ■

B.2. The Gaussian Model

Definition 24 Let π_k be the distribution over $\mathbb{R}^k$ given by $\sqrt{k}e_i - \mathbb{1}/\sqrt{k}$ where $i \sim [k]$ uniformly at random. Let $\mathcal{X}_k = \mathcal{X}(\pi_k)$ be the associated spike prior, as defined in Definition 18.

Here, $e_1, e_2, \dots$ denote the standard unit basis vectors and $\mathbb{1}$ denotes the all-ones vector.

Theorem 25 Suppose there exist constants $k \geq 2$ and $\epsilon > 0$ such that there is a time- $t(n)$ algorithm to certify the upper bound $(2 - \epsilon)n$ on $\Gamma_k(\mathbf{W})$ over $\mathbf{W} \sim \text{GOE}(n)$. Then for some constants $\beta \in (-1, 0)$ and $\gamma > 1$, there is a time- $(t(n) + \text{poly}(n))$ algorithm achieving strong detection in the general spiked Wishart model with spike prior $\mathcal{X}_k$.

Note that the above parameters satisfy $\beta^2 < \gamma$, which is in the “hard” regime of the Wishart model. Thus, Theorem 25 together with the low-degree-hardness of the Wishart model in that regime (Theorem 20) constitute rigorous evidence for Conjecture 15.

Proof Let $\mathcal{A}$ be the purported certification algorithm. We will use this to solve detection in the Wishart model, given Wishart samples $\mathbf{y}_1, \dots, \mathbf{y}_N$. Sample $\tilde{\mathbf{W}} \sim \text{GOE}(n)$ and let $\lambda_1 \leq \dots \leq \lambda_n$ denote its (random) eigenvalues. Sample a uniformly random orthonormal basis $\mathbf{v}_{n-N+1}, \dots, \mathbf{v}_n$ for $\mathbf{V} := \text{span}\{\mathbf{y}_1, \dots, \mathbf{y}_N\}$ and a uniformly random orthonormal basis $\mathbf{v}_1, \dots, \mathbf{v}_{n-N}$ for the orthogonal complement $\mathbf{V}^\perp$. Let $\mathbf{W} = \sum_{i=1}^n \lambda_i \mathbf{v}_i \mathbf{v}_i^\top$. Our Wishart detection algorithm is as follows: if $\mathcal{A}(\mathbf{W}) \leq (2 - \epsilon)n$ then output Q; otherwise, output P.

We now prove that this achieves strong detection. If the Wishart samples were drawn from $\mathbb{Q}$ then $\mathbf{V}$ is a uniformly random N -dimensional subspace and so $\mathbf{W} \sim \text{GOE}(n)$. This means $\mathcal{A}(\mathbf{W}) \leq (2 - \epsilon)n$ with high probability by assumption, and so our algorithm correctly outputs Q. It remains to show that if the Wishart samples were drawn from $\mathbb{P}$, then $\Gamma_k(\mathbf{W}) > (2 - \epsilon)n$ with high probability, and so the algorithm is forced to output P.

Suppose the Wishart samples were drawn from $\mathbb{P}$ with planted matrix $\mathbf{X}$. With high probability we are in the first case of (18), i.e., $\mathbf{X} = \frac{1}{n} \mathbf{U} \mathbf{U}^\top$ where each row of $\mathbf{U}$ is drawn independently from π_k . Note that $\frac{1}{k-1} \mathbf{U} \mathbf{U}^\top$ is a partition matrix, and so

$$\Gamma_k(\mathbf{W}) \geq -\frac{1}{k-1} \langle \mathbf{U} \mathbf{U}^\top, \mathbf{W} \rangle.$$

We can bound

$$\begin{aligned}
 \langle \mathbf{U}\mathbf{U}^\top, \mathbf{W} \rangle &= \left\langle \mathbf{U}\mathbf{U}^\top, \sum_{i=1}^n \lambda_i \mathbf{v}_i \mathbf{v}_i^\top \right\rangle \\
 &\leq \left\langle \mathbf{U}\mathbf{U}^\top, \lambda_{n-N} \sum_{i=1}^{n-N} \mathbf{v}_i \mathbf{v}_i^\top + \lambda_n \sum_{i=n-N+1}^n \mathbf{v}_i \mathbf{v}_i^\top \right\rangle \\
 &= \left\langle \mathbf{U}\mathbf{U}^\top, \lambda_{n-N} \left(I_n - \sum_{i=n-N+1}^n \mathbf{v}_i \mathbf{v}_i^\top \right) + \lambda_n \sum_{i=n-N+1}^n \mathbf{v}_i \mathbf{v}_i^\top \right\rangle \\
 &= \left\langle \mathbf{U}\mathbf{U}^\top, \lambda_{n-N} I_n + (\lambda_n - \lambda_{n-N}) \sum_{i=n-N+1}^n \mathbf{v}_i \mathbf{v}_i^\top \right\rangle
 \end{aligned}$$

where we have used the fact that $\sum_{i=1}^n \mathbf{v}_i \mathbf{v}_i^\top = I_n$ since $\{\mathbf{v}_i\}_{i \in [n]}$ is an orthonormal basis. We will bound the pieces of this expression separately.

First we bound $\langle \mathbf{U}\mathbf{U}^\top, I_n \rangle$. The nonzero eigenvalues of $\mathbf{U}\mathbf{U}^\top$ are the same as the nonzero eigenvalues of $\mathbf{U}^\top \mathbf{U}$. Since $\mathbf{U}^\top \mathbf{U}$ is the sum of n i.i.d. $k \times k$ matrices, $\frac{1}{n} \mathbf{U}^\top \mathbf{U}$ converges in probability to its expectation, which is $\text{Cov}(\pi) = I_k - J_k/k$ where J_k is the $k \times k$ all-ones matrix. This means $\langle \mathbf{U}\mathbf{U}^\top, I_n \rangle = \text{Tr}(\mathbf{U}\mathbf{U}^\top) = n \text{Tr}(\frac{1}{n} \mathbf{U}^\top \mathbf{U}) = (1 + o(1))n(k-1)$ with high probability.

Next we bound $\langle \mathbf{U}\mathbf{U}^\top, \sum_{i=n-N+1}^n \mathbf{v}_i \mathbf{v}_i^\top \rangle$. Recall that $\{\mathbf{v}_i\}_{i=n-N+1}^n$ is an orthonormal basis for $\text{span}\{\mathbf{y}_1, \dots, \mathbf{y}_N\}$, and so $\sum_{i=n-N+1}^n \mathbf{v}_i \mathbf{v}_i^\top \preceq \frac{1}{\mu} \mathbf{Y}$ where $\mathbf{Y} = \frac{1}{N} \sum_{i=1}^N \mathbf{y}_i \mathbf{y}_i^\top$ and μ is the smallest nonzero eigenvalue of $\mathbf{Y}$. Since $\mathbf{Y}$ is a spiked covariance matrix, Theorem 1.2 of [Baik and Silverstein \(2006\)](#) gives $\mu \rightarrow (\sqrt{\gamma} - 1)^2 > 0$ in probability. Therefore,

$$\left\langle \mathbf{U}\mathbf{U}^\top, \sum_{i=n-N+1}^n \mathbf{v}_i \mathbf{v}_i^\top \right\rangle \leq \left\langle \mathbf{U}\mathbf{U}^\top, \frac{1}{\mu} \mathbf{Y} \right\rangle = \frac{1}{\mu N} \sum_{i=1}^N \|\mathbf{U}^\top \mathbf{y}_i\|^2.$$

For fixed $\mathbf{U}$ (and therefore fixed $\mathbf{X}$), note that $\mathbf{U}^\top \mathbf{y}_i$ follows a multivariate Gaussian distribution with mean zero and covariance

$$\mathbb{E}[\mathbf{U}^\top \mathbf{y}_i \mathbf{y}_i^\top \mathbf{U}] = \mathbf{U}^\top \mathbb{E}[\mathbf{y}_i \mathbf{y}_i^\top] \mathbf{U} = \mathbf{U}^\top (I + \beta \mathbf{X}) \mathbf{U} = \mathbf{U}^\top \mathbf{U} + \frac{\beta}{n} \mathbf{U}^\top \mathbf{U} \mathbf{U}^\top \mathbf{U}.$$

Recalling that $\frac{1}{n} \mathbf{U}^\top \mathbf{U} \rightarrow I_k - J_k/k$, we have $\frac{1}{n} (\mathbf{U}^\top \mathbf{U} + \frac{\beta}{n} \mathbf{U}^\top \mathbf{U} \mathbf{U}^\top \mathbf{U}) \rightarrow (1 + \beta)(I_k - J_k/k)$ in probability. Thus, $\frac{1}{nN} \sum_{i=1}^N \|\mathbf{U}^\top \mathbf{y}_i\|^2$ converges in probability to $(1 + \beta) \text{Tr}(I_k - J_k/k) = (1 + \beta)(k-1)$, and we can conclude that

$$\langle \mathbf{U}\mathbf{U}^\top, \sum_{i=n-N+1}^n \mathbf{v}_i \mathbf{v}_i^\top \rangle \leq (1 + o(1)) \frac{n}{\mu} (1 + \beta)(k-1) = (1 + o(1))n(\sqrt{\gamma} - 1)^2 (1 + \beta)(k-1)$$

with high probability.

The eigenvalues of $\widetilde{\mathbf{W}}$ converge to the Wigner semicircle law on $[-2, 2]$, and so we have $\lambda_n - \lambda_{n-N} \leq 4 + o(1)$ with high probability. Also, by taking $\gamma > 1$ close enough to 1, we can ensure $\lambda_{n-N} \leq -2 + \epsilon/2$ with high probability.

Putting it all together, we now have

$$\begin{aligned}
\Gamma_k(\mathbf{W}) &\geq -\frac{1}{k-1} \langle \mathbf{U}\mathbf{U}^\top, \mathbf{W} \rangle \\
&\geq -\frac{1}{k-1} \left\langle \mathbf{U}\mathbf{U}^\top, \lambda_{n-N} \mathbf{I}_n + (\boldsymbol{\lambda}_n - \boldsymbol{\lambda}_{n-N}) \sum_{i=n-N+1}^n \mathbf{v}_i \mathbf{v}_i^\top \right\rangle \\
&\geq -\frac{1}{k-1} [(-2 + \epsilon/2)(1 + o(1))n(k-1) + (4 + o(1))n(\sqrt{\gamma} - 1)^{-2}(1 + \beta)(k-1)] \\
&> (2 - \epsilon)n
\end{aligned}$$

for sufficiently large n , provided we choose $\beta > -1$ sufficiently close to -1 . This completes the proof. $\blacksquare$

Appendix C. Belief Propagation and the Kesten–Stigum Transition

In this section, we carry out a stability analysis of belief propagation (as discussed in Section 1.1.1) and derive the result presented in Theorem 12. The analysis resembles that of the original work of Decelle et al. (2011a,b) that predicted the Kesten–Stigum threshold in the ordinary stochastic block model, but the equitability constraints create additional technical complexity in our setting. We start with the special case of the equitable coloring model in Section C.1, and generalize to the equitable block model in Section C.2. Throughout this section, it turns out to be convenient to parametrize the equitable model in a different way (defined below) than used in the Introduction.

C.1. The Equitable Coloring Model

The equitable coloring model is obtained by setting $\eta = -\frac{1}{k-1}$ in Definition 1. For brevity let's define

$$c := \frac{1 - \eta}{k} d, \quad (19)$$

so that in the planted coloring, each vertex has exactly c neighbors of every other color. From the point of view of the vertices, this is a complicated constraint affecting a star of $d + 1$ vertices. As a result, a factor graph with a variable node for each vertex, and a constraint node corresponding to each vertex and its neighbors, is not even locally treelike. Instead, we define a variable for each edge, giving a pair of colors. The constraint then demands that the d edges incident to each vertex agree on its color, and that the colors of their other endpoints are equitable.

This lets us define a message-passing algorithm. Regarding each edge (u, v) of the graph $\mathbf{G} \sim \mathcal{G}_{n,d,k,(k-1)-1}^{\text{eq}}$ as a pair of directed edges $u \rightarrow v, v \rightarrow u$, each directed edge $u \rightarrow v$ sends a message $\mu_{u \rightarrow v}^{r,s}$ to vertex v consisting of the estimated probabilities $\mu_{r,s}^{u \rightarrow v}$ that u and v are color r and s respectively, for each $r, s \in [k]$ with $r \neq s$. Vertex v then sends out messages $\mu_{v \rightarrow w}^{r,s}$ to the directed edges (v, w) which are computed as follows:

1. For each of v 's neighbors u other than w , choose a pair of colors (r_u, s_u) independently from the distribution $\mu_{u \rightarrow v}^{r,s} = (\mu_{r,s}^{u \rightarrow v})$.
2. Condition on the event that the s_u are identical for all u . Call this color s .

3. Condition on the event that all but one of the colors other than s appear c times in the list (r_u) , and that one color appears $c - 1$ times. Call this color t .
4. Then $\mu^{v \rightarrow w} = (\mu_{st}^{v \rightarrow w})$ is the resulting conditional distribution of the pair (s, t) , i.e., the probability that v and w are color s and t respectively.

Formally we can write

$$\mu_{st}^{v \rightarrow w} = \frac{\psi_{st}^{v \rightarrow w}}{z^{v \rightarrow w}} \quad (20)$$

where

$$\psi_{st}^{v \rightarrow w} = \sum_{(r_u : u \in \partial v \setminus w) \in [k]^{d-1}} \left(\prod_u \mu_{r_u, s}^{u \rightarrow v} \right) \left(\prod_{q \in [k], q \neq s} \mathbf{1} \left[|\{u : r_u = q\}| = \begin{cases} c & q \neq s, t \\ c-1 & q = t \end{cases} \right] \right) \quad (21)$$

$$z^{v \rightarrow w} = \sum_{r, s \in [k] : r \neq s} \psi_{rs}^{v \rightarrow w}, \quad (22)$$

where the product over q in (21) enforces condition 3 above.

Clearly the uniform messages $\mu_{rs}^{u \rightarrow v} = 1/(k(k-1))$ are a fixed point of this algorithm. We want to study its stability to small perturbations, and in particular the matrix of partial derivatives

$$Y_{rs, s't} = \frac{\partial \mu_{s't}^{v \rightarrow w}}{\partial \mu_{rs}^{u \rightarrow v}}. \quad (23)$$

To compute this matrix, suppose that we perturb the incoming message $\mu^{u \rightarrow v}$ for a pair (r, s) with $r \neq s$,

$$\mu_{r's'}^{u \rightarrow v} = \frac{1}{k(k-1)} (1 + \varepsilon \delta_{rr'} \delta_{ss'})$$

where δ is the Kronecker delta, $\delta_{rr'} = 1$ if $r = r'$ and 0 otherwise. This perturbation does not respect the normalization $\sum_{rs} \mu_{rs}^{u \rightarrow v} = 1$, but this will simply show up as Y having zero row and column sums since normalization projects perturbations to the subspace perpendicular to the uniform vector.

We then have several cases. If $s' \neq s$, for all $t \neq s'$ then $\psi_{s't}^{v \rightarrow w}$ is unchanged from its value at the uniform fixed point, namely

$$\psi_{s't}^{v \rightarrow w} = \left(\frac{1}{k(k-1)} \right)^{d-1} \binom{d-1}{c-1, c, \dots, c} = \left(\frac{1}{k(k-1)} \right)^{d-1} \frac{c(d-1)!}{c!^{k-1}} := \psi. \quad (24)$$

For $s' = s$ and $t = r$, we have

$$\begin{aligned} \psi_{st}^{v \rightarrow w} &= \left(\frac{1}{k(k-1)} \right)^{d-1} \left((1 + \varepsilon) \binom{d-2}{c-2, c, \dots, c} + (k-2) \binom{d-2}{c-1, c-1, c, \dots, c} \right) \\ &= \psi \left((1 + \varepsilon) \frac{c-1}{d-1} + (k-2) \frac{c}{d-1} \right) \\ &= \psi \left(1 + \varepsilon \frac{c-1}{d-1} \right), \quad (t = r) \end{aligned} \quad (25)$$

where the two terms in the first line come from $r_u = r$ and $r_u \neq r, s$ respectively. Finally, for $s' = s$ and $t \neq r$, we have

$$\begin{aligned}\psi_{st}^{v \rightarrow w} &= \left(\frac{1}{k(k-1)} \right)^{d-1} \left((1+\varepsilon) \binom{d-2}{c-1, c-1, c, \dots, c} + \binom{d-2}{c-2, c, \dots, c} + (k-3) \binom{d-2}{c-1, c-1, c, \dots, c} \right) \\ &= \psi \left((k-2+\varepsilon) \frac{c}{d-1} + \frac{c-1}{d-1} \right) \\ &= \psi \left(1 + \varepsilon \frac{c}{d-1} \right), \quad (t \neq r)\end{aligned}\tag{26}$$

where the three terms in the first line come from $r_u = r$, $r_u = t$, and $r_u \neq r, s, t$ respectively.

While this level of bookkeeping is comforting, both (25) and (26) are simply $\psi(1 + \varepsilon P)$ where P is the fraction of $(d-1)$ -tuples that contribute to (21) such that $r_u = r$. We can write (24), (25), and (26) as

$$\psi_{s't}^{v \rightarrow w} = \psi \left(1 + \varepsilon \delta_{ss'} \frac{c - \delta_{rt}}{d-1} \right).\tag{27}$$

Summing over all distinct s', t gives the normalization factor

$$\begin{aligned}z^{v \rightarrow w} &= k(k-1)\psi + \varepsilon\psi \left(\frac{c-1}{d-1} + (k-2)\frac{c}{d-1} \right) \\ &= k(k-1)\psi + \varepsilon\psi \\ &= k(k-1)\psi \left(1 + \frac{\varepsilon}{k(k-1)} \right),\end{aligned}\tag{28}$$

with the multiplicative factor $1 + \varepsilon P$ where $P = 1/(k(k-1))$ is now the probability that a random edge $u \rightarrow v$ has colors r and s on its endpoints.

Combining (28) with (27), and (20) gives

$$\mu_{s't}^{v \rightarrow w} = \frac{1}{k(k-1)} \left(1 + \varepsilon \left(-\frac{\delta_{ss'}\delta_{rt}}{d-1} + \frac{c\delta_{ss'}}{d-1} - \frac{1}{k(k-1)} \right) + O(\varepsilon^2) \right).\tag{29}$$

Canceling the factor $1/(k(k-1))$ gives the matrix of partial derivatives (23),

$$Y_{rs,s't} = -\frac{\delta_{ss'}\delta_{rt}}{d-1} + \frac{c\delta_{ss'}}{d-1} - \frac{1}{k(k-1)}.\tag{30}$$

Using (30) and (19) the reader can check that the rows and columns of Y sum to zero, as alluded to above:

$$\forall r, s : \sum_{\substack{s', t: \\ s' \neq t}} Y_{rs,s't} = 0, \quad \forall s', t : \sum_{\substack{r, s: \\ r \neq s}} Y_{rs,s't} = 0.\tag{31}$$

To diagonalize Y , it is useful to treat the $k(k-1)$ -dimensional space $\mathcal{U}$ spanned by ordered pairs (r, s) with $r \neq s$ as the space of $k \times k$ matrices $U = (U_{rs})$ with zeroes on the diagonal. Then we can interpret the three terms in (30) as follows:

- The term $\delta_{ss'}\delta_{rt}$ is the transpose operator, sending U to $U^\top$.

- The term $\delta_{ss'}$ sends U to $(UJ)^\top$ where J is the all-1s matrix.
- The term $-1/k(k-1)$ subtracts the mean entry of U from each entry of $Y(U)$.
- Finally, we set all the diagonal elements of $Y(U)$ to zero.

Thus we can rewrite (30) as

$$Y(U) = \Pi \left[-\frac{1}{d-1} U^\top + \frac{c}{d-1} (UJ)^\top - \frac{1}{k(k-1)} JUJ \right] \quad (32)$$

where Π is the projection operator that sets the diagonal entries of a matrix to zero.

To diagonalize Y , recall that if two linear operators commute, they share the same eigenvectors. Clearly Y commutes with relabelings of the colors, i.e., with the S_k -action that conjugates U with a permutation matrix. This action preserves the following subspaces of matrices:

- The symmetric matrices with zero diagonal
- The antisymmetric matrices
- The matrices whose row (resp. column) sums are zero
- The matrices whose rows (resp. columns) are uniform, other than being zero on the diagonal

... and their intersections. More abstractly, $\mathcal{U}$ is the $k(k-1)$ -dimensional combinatorial representation where S_k acts on distinct ordered pairs (r, s) by sending (r, s) to $(\pi(r), \pi(s))$. We can find the eigenvectors and eigenvalues of Y by decomposing $\mathcal{U}$ into a direct sum of irreducible representations of S_k . This decomposition includes one copy of the trivial representation $\rho_{(k)} = I$, and one copy each of $\rho_{(k-2,1,1)}$ and $\rho_{(k-2,2)}$. (To avoid some case-checking we assume that $k \geq 4$. In particular, if $k = 3$ then $\rho_{(k-2,2)}$ disappears.) By Schur's Lemma, when restricted to each of these irreducible subspaces Y is a scalar matrix with a single eigenvalue. These are as follows:

- The trivial representation is spanned by the matrix with 1s everywhere off the diagonal. By (31) this has eigenvalue zero.
- The copy of $\rho_{(k-2,1,1)}$ consists of antisymmetric matrices with zero row and column sums. These are annihilated by v and are eigenvectors of the transpose with eigenvalue -1 . Thus they are eigenvectors of Y with eigenvalue $+1/(d-1)$. This eigenspace has dimension $(k-1)(k-2)/2$.
- The copy of $\rho_{(k-2,2)}$ consists of symmetric matrices with zero row and column sums and zeroes on the diagonal. These are annihilated by v and are eigenvectors of the transpose with eigenvalue $+1$. Thus they are eigenvectors of Y with eigenvalue $-1/(d-1)$. This eigenspace has dimension $(k-1)(k-2)/2 - 1 = k(k-3)/2$.

However, we are not done. In addition to these multiplicity-free irreducible representations, $\mathcal{U}$ includes two copies of the “standard” representation $\rho_{(k-1,1)}$, one each in the symmetric and antisymmetric subspace. Each one has dimension $k-1$, and together they span an isotypic subspace of dimension $2(k-1)$. By Schur's lemma, when restricted to this subspace, Y is the tensor product of the identity with a 2×2 matrix, giving it two additional eigenvalues. One of these will turn out to be the dominant one and will control where the Kesten–Stigum transition occurs.

This isotypic subspace is spanned by matrices of the form

$$U_{ij} = \begin{cases} 0 & i = j \\ \alpha & i = 1, j \neq 1 \\ \beta & j = 1, i \neq 1 \\ \gamma & i \neq 1, j \neq 1, i \neq j \end{cases} \quad \text{where } \gamma = -\frac{\alpha + \beta}{k - 2}, \quad (33)$$

and their images under conjugation by permutation matrices, i.e., where the “special” row and column ranges from 1 to k . That is,

$$U = \begin{pmatrix} 0 & \alpha & \alpha & \cdots & \alpha \\ \beta & 0 & \gamma & \cdots & \gamma \\ \beta & \gamma & 0 & & \vdots \\ \vdots & \vdots & & \ddots & \gamma \\ \beta & \gamma & \cdots & \gamma & 0 \end{pmatrix}$$

where γ is set so that U ’s entries sum to zero. The reader can check that these matrices are orthogonal to the other irreducible subspaces with respect to the trace inner product $\langle U, U' \rangle = \text{Tr } U^\top U'$.

Using (32), we find that $Y(U)$ is also of this form but with entries α' and β' , where

$$\begin{pmatrix} \alpha' \\ \beta' \end{pmatrix} = m \cdot \begin{pmatrix} \alpha \\ \beta \end{pmatrix} \quad \text{where } m = \begin{pmatrix} 0 & 1 \\ \frac{-1}{d-1} & \frac{-c}{d-1} \end{pmatrix}. \quad (34)$$

Thus Y on this isotypic subspace is $m \otimes I$ where I is the $(k - 1)$ -dimensional identity. The corresponding eigenvalues of Y are those of m , namely the roots κ of

$$(d - 1)\kappa^2 + c\kappa + 1 = 0,$$

which are

$$\kappa_{\pm} = \frac{-c \pm \sqrt{c^2 - 4(d - 1)}}{2(d - 1)}.$$

When $c^2 < 4(d - 1)$ the discriminant is negative, so these eigenvalues are complex and lie on the unit circle $|\kappa_{\pm}| = 1/\sqrt{d - 1}$. But when $c^2 > 4(d - 1)$ they are real, and $\kappa_+ > 1/\sqrt{d - 1}$.

The full Jacobian of belief propagation is the tensor product of this local matrix Y with the non-backtracking matrix B . It is a consequence of Theorem 2 in [Bordenave and Collins \(2019\)](#) that, with high probability over G sampled from the equitable SBM with $d < d_{\text{KS}}^{\text{eq}}$, the spectrum of the non-backtracking matrix consists of a “trivial” eigenvalue d whose left and right eigenvectors is the all-ones vector, and remaining eigenvalues with modulus at most $\sqrt{d - 1} + o_n(1)$ in the complex plane. Since perturbations along the uniform eigenvector of B would violate the balance of colors, we are left with these remaining eigenvalues. Multiplying κ_+ by $\sqrt{d - 1} + o_n(1)$ tells us that the Kesten–Stigum transition, where the largest eigenvalue of the Jacobian exceeds 1 in absolute value, occurs when $c^2 \approx 4(d - 1)$, with $\approx$ hiding $o_n(1)$ terms. Given (19) this is

$$c \approx 2 \left((k - 1) + \sqrt{k(k - 2)} \right)$$

or

$$d \approx 2(k-1)^2 \left(1 + \sqrt{\frac{k(k-2)}{(k-1)^2}} \right) \approx \frac{2}{\lambda^2} \left(1 + \sqrt{1 - \lambda^2} \right)$$

where $\lambda = -1/(k-1)$. Solving this condition for λ in terms of d gives

$$|\lambda| = \frac{2\sqrt{d-1}}{d} + o_n(1). \quad (35)$$

C.2. Generalizing to the Equitable Block Model

Referring again to Definition 1, let a and b be $M_{i,j}d$ where $i = j$ and $i \neq j$ respectively:

$$a := \frac{1 + (k-1)\eta}{k} d \quad \text{and} \quad b := \frac{1 - \eta}{k} d$$

so that

$$d = a + b(k-1). \quad (36)$$

Then a legal labeling of the equitable block model on a d -regular graph is a k -coloring where each vertex has exactly a neighbors of its own color, and exactly b neighbors of each of the $k-1$ other colors. The case of equitable k -colorings corresponds to $a = 0$ and $b = c$.

We can generalize the message-passing algorithm of the previous section as follows. Each directed edge (u, v) again sends a message $\mu^{u \rightarrow v}$ to vertex v consisting of the estimated probability $\mu_{rs}^{u \rightarrow v}$ that u and v are color r and s respectively, for each $r, s \in [k]$ with $r \neq s$. Vertex v then sends out messages $\mu^{v \rightarrow w}$ to the directed edges (v, w) which are computed as follows:

1. Give pairs (s, t) the prior distribution

$$P(s, t) = \frac{1}{kd} \begin{cases} a & (s = t) \\ b & (s \neq t). \end{cases} \quad (37)$$

2. Multiplicatively reweight this distribution by the probability that, if for each of v 's neighbors u other than w we choose a pair of colors (r_u, s_u) independently from the distribution $\mu^{u \rightarrow v} = (\mu_{rs}^{u \rightarrow v})$, then:
 - $s_u = s$ for all u , and
 - (if $s = t$) s appears $a - 1$ times in the list (r_u) , and every color other than s appears b times,
 - (if $s \neq t$) s appears a times in the list (r_u) , t appears $b - 1$ times, and all other colors appear b times each.

3. Then $\mu^{v \rightarrow w} = (\mu_{st}^{v \rightarrow w})$ is the resulting posterior distribution of (s, t) .

Generalizing (21) and (22), we can write

$$\mu_{st}^{v \rightarrow w} = \frac{\psi_{st}^{v \rightarrow w}}{z^{v \rightarrow w}} \quad (38)$$

where

$$\psi_{st}^{v \rightarrow w} = \sum_{(r_u : u \in \partial v \setminus k) \in [k]^{d-1}} \left(\prod_u \mu_{r_u, s}^{u \rightarrow v} \right) \begin{cases} a \prod_{q \in [k]} \mathbf{1} \left[|\{u : r_u = q\}| = \begin{cases} a-1 & q = s \\ b & q \neq s \end{cases} \right] & (s = t) \\ b \prod_{q \in [k]} \mathbf{1} \left[|\{u : r_u = q\}| = \begin{cases} a & q = s \\ b-1 & q = t \\ b & q \neq s, t \end{cases} \right] & (s \neq t) \end{cases} \quad (39)$$

$$z^{v \rightarrow w} = \sum_{r, s \in [k] : r \neq s} \psi_{rs}^{v \rightarrow w}, \quad (40)$$

where the product over q in (39) enforces the second and third parts of condition 2 above.

Note that the factors of a and b in (39), which did not appear in the coloring case, come from the prior distribution (37). The reader can check that the prior distribution on (s, t) is now a fixed point of this algorithm,

$$\mu_{st}^{u \rightarrow v} = \frac{1}{kd} \begin{cases} a & (s = t) \\ b & (s \neq t). \end{cases}$$

At this fixed point, (39) gives

$$\psi_{st}^{v \rightarrow w} = \begin{cases} a\psi & (s = t) \\ b\psi & (s \neq t) \end{cases} \quad (41)$$

where

$$\psi := \frac{a^{a-1} b^{b(k-1)}}{(kd)^{d-1}} \binom{d-1}{a-1, b, \dots, b} = \frac{a^a b^{b(k-1)-1}}{(kd)^{d-1}} \binom{d-1}{a, b-1, b, \dots, b} = \frac{a^a b^{b(k-1)} (d-1)!}{(kd)^{d-1} a! b!^{k-1}},$$

and where as always $0^0 = 0! = 1$.

We again consider perturbing the incoming message $\mu^{u \rightarrow v}$ for a pair (r, s) ,

$$\mu_{r's'}^{u \rightarrow v} = (1 + \varepsilon \delta_{rr'} \delta_{ss'}) \frac{1}{kd} \begin{cases} a & (r = s) \\ b & (r \neq s). \end{cases} \quad (42)$$

As before, if $s' \neq s$ then $\psi_{s't}^{v \rightarrow w}$ is unchanged for all t , and is still given by (41). For the other cases where $s' = s$, let us first assume that $r \neq s$. There are now three cases: $t = s$ (now that some neighbors have the same color), $t = r$, and $t \notin \{r, s\}$. We have

$$\psi_{st}^{v \rightarrow w} = \begin{cases} a\psi \left(1 + \varepsilon \frac{b}{d-1}\right) & t = s \\ b\psi \left(1 + \varepsilon \frac{b-1}{d-1}\right) & t = r \\ b\psi \left(1 + \varepsilon \frac{b}{d-1}\right) & t \notin \{r, s\} \end{cases} \quad (r \neq s) \quad (43)$$

where in each case we multiply by $1 + \varepsilon P$ where P is the fraction of $(d - 1)$ -tuples that contribute to (39) where $r_u = r$. Summing over all s', t gives

$$\begin{aligned} z^{v \rightarrow w} &= kd\psi + \varepsilon b\psi \left(\frac{b-1}{d-1} + (k-2)\frac{b}{d-1} + \frac{a}{d-1} \right) \\ &= kd\psi + \varepsilon b\psi \\ &= kd\psi \left(1 + \varepsilon \frac{b}{kd} \right) \quad (r \neq s). \end{aligned} \quad (44)$$

This multiplicative factor is $1 + \varepsilon P$ where $P = b/(kd)$ is the prior probability that a random edge $u \rightarrow v$ has colors r, s on its endpoints where $r \neq s$.

Analogously, if $r = s = s'$ we have

$$\psi_{st}^{v \rightarrow w} = \begin{cases} a\psi \left(1 + \varepsilon \frac{a-1}{d-1} \right) & t = s \\ b\psi \left(1 + \varepsilon \frac{a}{d-1} \right) & t \neq s \end{cases} \quad (r = s), \quad (45)$$

and

$$\begin{aligned} z^{v \rightarrow w} &= kd\psi + \varepsilon a\psi \left(\frac{a-1}{d-1} + (k-1)\frac{b}{d-1} \right) \\ &= kd\psi + \varepsilon a\psi \\ &= kd\psi \left(1 + \varepsilon \frac{a}{kd} \right) \quad (r = s), \end{aligned} \quad (46)$$

where the multiplicative factors are again $1 + \varepsilon P$ where P is the fraction of $(d - 1)$ -tuples contributing to (39), or in (46) the prior probability $P = a/(kd)$ of an edge having colors $r = s$ on its endpoints.

Putting all this together generalizes (29) to

$$\mu_{s't}^{v \rightarrow w} = \frac{1}{kd} \begin{cases} a \left(1 + \varepsilon \left(\delta_{ss'} \frac{a-1}{d-1} - \frac{a}{kd} \right) + O(\varepsilon^2) \right) & [s' = t, r = s] \\ a \left(1 + \varepsilon \left(\delta_{ss'} \frac{b}{d-1} - \frac{b}{kd} \right) + O(\varepsilon^2) \right) & [s' = t, r \neq s] \\ b \left(1 + \varepsilon \left(\delta_{ss'} \frac{a}{d-1} - \frac{a}{kd} \right) + O(\varepsilon^2) \right) & [s' \neq t, r = s] \\ b \left(1 + \varepsilon \left(\delta_{ss'} \frac{b-\delta_{rt}}{d-1} - \frac{b}{kd} \right) + O(\varepsilon^2) \right) & [s' \neq t, r \neq s], \end{cases} \quad (47)$$

the fourth case of which coincides with (29) when $a = 0$ and $b = c$. Comparing with (42) and accounting for factors of a and b gives the matrix of partial derivatives,

$$Y_{rs,s't} = -\frac{\delta_{ss'}\delta_{rt}}{d-1} + \left(\frac{\delta_{ss'}}{d-1} - \frac{1}{kd} \right) \begin{cases} a & (s' = t) \\ b & (s' \neq t) \end{cases}. \quad (48)$$

The reader can check the normalization conditions: the rows of Y sum to zero, so that the uniform vector is a right eigenvector of eigenvalue zero, but the columns are orthogonal to the prior distribution (37) so that it is a left eigenvector of eigenvalue zero. Thus (31) becomes

$$\forall r, s : \sum_{s', t} Y_{rs,s't} = 0, \quad \forall s', t : \sum_{r, s} Y_{rs,s't} \begin{cases} a & (r = s) \\ b & (r \neq s) \end{cases} = 0. \quad (49)$$

At the risk of multiplying entities without necessity, we can also write Y in the style of (32). If we think of Y 's action by right multiplication on k^2 -dimensional vectors as a linear operator on k -dimensional matrices $U = (U_{s't})$, then

$$Y(U) = -\frac{1}{d-1}U^\top + \frac{1}{d-1}(\Upsilon(U)J)^\top - \frac{1}{kd}J\Upsilon(U)J, \quad (50)$$

where J is again the all-1s matrix and Υ is a linear operator on matrices that reweights diagonal and off-diagonal elements by a and b respectively,

$$\Upsilon(U)_{s't} = U_{s't} \begin{cases} a & (s' = t) \\ b & (s' \neq t) \end{cases}. \quad (51)$$

As in the coloring case, we use representation theory to diagonalize Y . The k -dimensional matrices form a k^2 -dimensional representation of S_k where permutation matrices act by conjugation. Since this representation sends pairs of colors (r, s) to $(\pi(r), \pi(s))$, this is the tensor product of the natural permutation representation with itself. The permutation representation is a direct sum of the trivial representation (spanned by the uniform vector) with the $(k-1)$ -dimensional standard representation (spanned by vectors that sum to zero). Taking its tensor square and decomposing gives the representations described above, as well as the subspace spanned by diagonal matrices, giving one additional copy each of the trivial representation and the standard representation.

Thus in total we have a two-dimensional trivial subspace, one copy each of $\rho_{(k-2,1,1)}$ and $\rho_{(k-2,2)}$ with dimension $(k-1)(k-2)/2$ and $k(k-3)/2$ respectively, and a $3(k-1)$ -dimensional subspace consisting of three copies of the standard representation $\rho_{(k-1,1)}$. We go through each of these subspaces, focusing on Y 's right eigenvectors.

First, the trivial subspace is spanned by the identity matrix $I = (\delta_{s't})$ and the all-1s matrix J . As in (49) J is a right eigenvector with eigenvalue zero, so $Y(J) = 0$. Observing (50), we have $I^\top = I$, $\Upsilon(I) = aI$, $IJ = J$, and $JIJ = kJ$. This gives

$$Y(I) = -\frac{1}{d-1}I + \left(\frac{a}{d-1} - \frac{a}{d}\right)J = -\frac{1}{d-1}I + \frac{a}{d(d-1)}J.$$

Thus in this two-dimensional subspace Y acts as the matrix

$$\frac{1}{d-1} \begin{pmatrix} -1 & 0 \\ a/d & 0 \end{pmatrix}$$

giving the eigenvalues $-1/(d-1)$ and 0.

Next, as before the copy of $\rho_{(k-2,1,1)}$ consists of antisymmetric matrices U with zero row and column sums. For these matrices we have $\Upsilon(U) = bU$ and $UJ = 0$, while $U^\top = -U$. Thus they are again eigenvectors of Y with eigenvalue $+1/(d-1)$.

The copy of $\rho_{(k-2,2)}$ consists of symmetric matrices U with zero row and column sums and zeroes on the diagonal. Now we have $\Upsilon(U) = bU$, $UJ = 0$, and $U^\top = U$, and (50) again makes them eigenvectors with eigenvalue $-1/(d-1)$.

This leaves the three copies of the standard representation. This isotypic subspace is spanned by matrices like those in (33) but with nonzero diagonal entries, namely

$$U_{ij} = \begin{cases} \delta & i = j = 1 \\ \zeta & i = j \neq 1 \\ \alpha & i = 1, j \neq 1 \\ \beta & j = 1, i \neq 1 \\ \gamma & i \neq 1, j \neq 1, i \neq j \end{cases} \quad \text{where } \gamma = -\frac{\alpha + \beta}{k-2} \quad \text{and} \quad \zeta = -\frac{\delta}{k-1}, \quad (52)$$

and their images under conjugation by permutation matrices, i.e., where the “special” row and column ranges from 1 to k . That is,

$$U = \begin{pmatrix} \delta & \alpha & \alpha & \cdots & \alpha \\ \beta & \zeta & \gamma & \cdots & \gamma \\ \beta & \gamma & \zeta & & \vdots \\ \vdots & \vdots & & \ddots & \gamma \\ \beta & \gamma & \cdots & \gamma & \zeta \end{pmatrix}$$

where γ and ζ are set so that U 's entries sum to zero and U has zero trace. In particular, U is orthogonal to both J and I , and hence to the trivial subspace. The reader can confirm that it is orthogonal to $\rho_{(k-2,1,1)}$ and $\rho_{(k-2,2)}$ as well.

Using (50) and a little work, we find that $Y(U)$ is also of this form but with entries α', β', δ' , where

$$\begin{pmatrix} \alpha' \\ \beta' \\ \delta' \end{pmatrix} = m \cdot \begin{pmatrix} \alpha \\ \beta \\ \delta \end{pmatrix} \quad \text{where} \quad m = \frac{1}{d-1} \begin{pmatrix} -b & -1 & \frac{-a}{k-1} \\ b(k-1)-1 & 0 & a \\ b(k-1) & 0 & a-1 \end{pmatrix}. \quad (53)$$

Thus Y on this isotypic subspace is $m \otimes I$ where I is the $(k-1)$ -dimensional identity. The corresponding eigenvalues of Y are those of m , which are namely the roots κ of

$$-\frac{1}{d-1} \quad \text{and} \quad \kappa_{\pm} = \frac{a-b \pm \sqrt{(a-b)^2 - 4(d-1)}}{2(d-1)}.$$

Analogous with the coloring case, if $(a-b)^2 < 4(d-1)$ these eigenvalues are complex and lie on the unit circle $|\kappa_{\pm}| = 1/\sqrt{d-1}$. But when $(a-b)^2 > 4(d-1)$, they are real, and $\kappa_+ > 1/\sqrt{d-1}$.

We again multiply κ_+ by the modulus of the largest non-trivial eigenvalue of the non-backtracking matrix, $\sqrt{d-1} + o_n(1)$, to obtain the dominant eigenvalue of the Jacobian of belief propagation. The Kesten–Stigum transition occurs when this eigenvalue exceeds the unit circle, or when

$$(a-b)^2 = 4(d-1) + o_n(1).$$

Since in the equitable block model we have

$$\lambda = \frac{a-b}{d},$$

this again occurs at

$$|\lambda| = \frac{2\sqrt{d-1}}{d} + o_n(1).$$

Appendix D. Local Statistics

Throughout this section, we will for the sake of brevity write $\mathbb{Q} = (\mathbb{Q}_n)$ for the uniform distribution $\mathcal{G}_{n,d}$ on d -regular graphs, and $\mathbb{P} = (\mathbb{P}_n)$ for the equitable stochastic block model $\mathcal{G}_{n,k,d,\eta}^{\text{eq}}$ from Definition 1.

As in the preceding text, we are most concerned with the behavior of the null and planted models when the number of vertices is very large, and we will write *with high probability (w.h.p.)* to describe a sequence of events that hold with probability $1 - o_n(1)$ in $\mathbb{P}_n$ or $\mathbb{Q}_n$ as $n \rightarrow \infty$, with other parameters (d, k, η) held fixed. The constant in the $o_n(1)$ may depend on these other parameters, and we will not make any attempt to quantify its rate, leaving us free to take union bounds over constantly many events.

In this section we study a family of semidefinite programming algorithms for the $\mathbb{P}$ vs. $\mathbb{Q}$ distinguishing problem. Like Sum of Squares, the *Local Statistics algorithm* is phrased in the language of polynomials. Let us define a set of variables $x = \{x_{u,i}\}$ indexed by vertices $u \in [n]$ and group labels $i \in [k]$, and $G = \{G_{u,v}\}$ indexed by pairs of distinct vertices. We think of the planted model as outputting a random evaluation of these variables, namely a pair $(\mathbf{x}, \mathbf{G})$, where $\mathbf{x} \in \mathbb{R}^{n \times k}$ encodes the hidden community structure—with $x_{u,i} = 1$ if $\sigma(u) = i$ and zero otherwise—and $\mathbf{G} \in \mathbb{R}^{\binom{[n]}{2}}$ is the Boolean vector indicating which edges are present in the graph. This allows us to regard polynomials $p \in \mathbb{R}[x, G]$ as statistics of the planted distribution $\mathbb{P}$, and we will in particular focus on the quantities $\mathbb{E}_{(\mathbf{x}, \mathbf{G}) \sim \mathbb{P}} p(\mathbf{x}, \mathbf{G})$.

The planted model outputs random variables $\mathbf{x}$ and $\mathbf{G}$ with a particular combinatorial structure: each variable is $\{0, 1\}$ -valued, and each vertex has exactly one label. This can be encoded in a set of polynomial constraints:

$$\begin{aligned} G_{u,v}^2 - G_{u,v} &= 0 & \forall (u, v) \in \binom{[n]}{2} \\ x_{u,i}^2 - x_{u,i} &= 0 & \forall u \in [n], i \in [k] \\ \sum_{i \in [k]} x_{u,i} - 1 &= 0 & \forall u \in [n]. \end{aligned}$$

Calling $\mathcal{I}_k$ the ideal of $\mathbb{R}[x, G]$ generated by the polynomials on the left hand side of the equations above, then for any $p \in \mathcal{I}_k$, $p(\mathbf{x}, \mathbf{G}) = 0$. Moreover, the planted distribution has a pleasant symmetry property: the symmetric group S_n acts naturally and simultaneously on the variables x and G , with a permutation ξ acting as $x_{u,i} \mapsto x_{\xi(u),i}$ and $G_{u,v} \mapsto G_{\xi(u),\xi(v)}$, and for any polynomial $p \in \mathbb{R}[x, G]$, the expectation $\mathbb{E}_{(\mathbf{x}, \mathbf{G}) \sim \mathbb{P}} p(\mathbf{x}, \mathbf{G})$ is constant on the orbits of this action. The Local Statistics algorithm, given as input a graph G_0 , endeavors to find a “pseudoexpectation” that mimics the conditional expectation $\mathbb{E}[\cdot | G_0]$ on polynomials $p(x, G_0)$ of sufficiently low degree.

Definition 26 (Local Statistics Algorithm with Informal Moment Constraints) *The degree- (D_x, D_G) Local Statistics algorithm is the following SDP: given an input graph G_0 , find $\tilde{\mathbb{E}} : \mathbb{R}[x]_{\leq D_x} \rightarrow \mathbb{R}$ s.t.*

1. (Positivity) $\tilde{\mathbb{E}}p(x)^2 \geq 0$ whenever $\deg p^2 \leq D_x$
2. (Hard Constraints) $\tilde{\mathbb{E}}p(x, G_0) = 0$ for every $p \in \mathcal{I}_k$

3. (Moment Constraints) $\tilde{\mathbb{E}}p(x, G_0) \approx \mathbb{E}_{(\mathbf{x}, \mathbf{G}) \sim \mathbb{P}} p(\mathbf{x}, \mathbf{G})$ whenever $\deg_G p(x, G) \leq D_G$, $\deg_x p(x, G) \leq D_x$, and p is fixed under the S_n action.

See, e.g., the survey [Laurent \(2009\)](#) for detailed discussion of how optimization problems of this form can be solved as SDPs.

We use the Local Statistics SDP for the $\mathbb{P}$ vs. $\mathbb{Q}$ hypothesis testing problem as follows: given $\mathbf{G}$ sampled from one of these two distributions, we run Local Statistics, outputting $\mathbb{P}$ if the SDP is feasible, and $\mathbb{Q}$ otherwise. The symbol $\approx$ in the moment constraints indicates that we will permit some additive error; this is necessary so that, when $\mathbf{G} \sim \mathbb{P}$, the SDP is with high probability satisfiable, by setting $\tilde{\mathbb{E}}p(x, G_0) = \mathbb{E}p(\mathbf{x}, \mathbf{G})$. In fact, we will instantiate these moment constraints only on the elements of a certain combinatorially meaningful basis, and we will allow different additive error for different basis elements. In so doing we automatically satisfy positivity and the hard constraints, and the additive slack allows for fluctuations of the $p(\mathbf{x}, \mathbf{G})$ around their expectations. When we make this precise below, we will write this additive slack in terms of an ‘error tolerance’ $\delta > 0$.

Theorem 27 *If $(d\eta)^2 > 4(d-1)$, then there exists a constant (in n) $\delta > 0$ so that the degree $(2, 2)$ Local Statistics algorithm with error tolerance δ can distinguish $\mathbb{P}$ and $\mathbb{Q}$. If $(d\eta)^2 \leq 4(d-1)$, then there do not exist constants D and δ for which the degree $(2, D)$ Local Statistics algorithm with tolerance δ can do so.*

The proof of Theorem 27 closely follows [Banks et al. \(2019\)](#). As in that work, we will first study a simpler SDP hierarchy for the $\mathbb{P}$ vs. $\mathbb{Q}$ hypothesis testing problem, and then show that its feasibility is equivalent to that of degree $(2, D)$ local statistics SDP. We begin with some standard facts about non-backtracking walks, which will be a central tool in our analysis.

D.1. Non-Backtracking Walks

Let A_G be the adjacency matrix for a d -regular graph G (which may have self-loops and multi-edges). A length- s non-backtracking walk on G is an alternating sequence of vertices and edges $v_1, e_1, v_2, e_2, \dots, v_s$ without terms of the form v, e, w, e, v . The matrices $A_G^{(s)}$ whose u, v entries count the number of such walks between vertices u and v are given by

$$\begin{aligned} A_G^{(0)} &= 1 \\ A_G^{(1)} &= A_G \\ A_G^{(2)} &= A_G^2 - d \\ A_G^{(s+1)} &= A A_G^{(s)} - (d-1) A_G^{(s-1)} \quad s \geq 2. \end{aligned}$$

In particular, $A_G^{(s)} = q_s(A_G)$, for a sequence of monic univariate polynomials $q_s \in \mathbb{R}[z]$, with $\deg q_s = s$, which are known to be orthogonal with respect to the Kesten-McKay measure

$$d\mu_{\text{KM}}(z) := \frac{d}{2\pi} \frac{\sqrt{4(d-1) - z^2}}{d^2 - z^2} \mathbf{1}_{\left[|z| < 2\sqrt{d-1}\right]} dz$$

on the interval $(-2\sqrt{d-1}, 2\sqrt{d-1})$. This fact has appeared innumerable times in the literature, dating back at least to [McKay \(1981\)](#). Thus the polynomials q_s are a basis for the Hilbert space of

square integrable functions on this interval, equipped with the inner product

$$\langle f, g \rangle_{\text{KM}} := \int f(z)g(z)d\mu_{\text{KM}}(z),$$

and associated norm

$$\|f\|_{\text{KM}}^2 := \langle f, f \rangle_{\text{KM}} = \int f(z)^2 d\mu_{\text{KM}}(z),$$

and in particular for any polynomial $f \in \mathbb{R}[z]$, we have the orthogonal decomposition

$$f = \sum_{s \geq 0} \frac{\langle f, q_s \rangle_{\text{KM}}}{\|q_s\|_{\text{KM}}^2} q_s.$$

We record for later use that

$$\|q_s\|_{\text{KM}}^2 = q_s(d) = \begin{cases} 1 & s = 1 \\ d(d-1)^{s-1} & s > 1 \end{cases};$$

this is equal to the number of vertices at depth s in a rooted d -regular tree, or equivalently n^{-1} times the total number of length- s non-backtracking walks in a d -regular graph on n vertices. For a derivation of this and other related facts, the reader may refer to [Solé \(1996\)](#) or [Sodin \(2007\)](#), but should beware of differing normalization conventions.

We will also require some standard and generic properties sequences of univariate polynomials orthogonal with respect to a measure on an interval of $\mathbb{R}$ ([Szeg, 1939](#), Theorems 3.3.1, 6.6.1, and 3.4.1-2): each q_s has s roots in the interval $(-2\sqrt{d-1}, 2\sqrt{d-1})$, the union of these roots over all $s \geq 0$ are dense in this interval, and we have a nonnegative quadrature rule

Lemma 28 (Quadrature) *For each s , call $r_1 < r_2 < \dots < r_s$ the roots of q_s . There exist weights $w_1, \dots, w_s \geq 0$ with the property that*

$$\langle f, q_s \rangle_{\text{KM}} = \sum_{i \in [s]} f(r_i) q_s(r_i) w_i$$

for every polynomial f of degree at most $2s - 1$.

D.2. The Symmetric Path Statistics SDP

In this section we study a simplified version of the Local Statistics SDP, which will ultimately be key to our analysis of the full Local Statistics SDP. Let $(\mathbf{x}, \mathbf{G}) \sim \mathbb{P}$, thinking of $\mathbf{x}$ as a collection of k vectors $\mathbf{x}_1, \dots, \mathbf{x}_k \in \{0, 1\}^n$. One can check that the partition matrix for the planted labelling, in the sense of Definition 4, is

$$\mathbf{P} := \frac{k}{k-1} \left(\sum_{i \in [k]} \mathbf{x}_i \mathbf{x}_i^\top - \frac{1}{k} J_k \right)$$

Since deterministically $\mathbf{x}_1 + \dots + \mathbf{x}_k = \mathbf{1}$, we as well have

$$= \frac{k}{k-1} \left(\sum_{i \in [k]} \mathbf{x}_i \mathbf{x}_i^\top - \frac{1}{k} \sum_{i,j} \mathbf{x}_i \mathbf{x}_j^\top \right)$$

As we observed in the Introduction, $\mathbf{P}$ is PSD with ones on the diagonal, and $(k/n)\mathbf{P}$ is the orthogonal projector onto the $(k-1)$ -dimensional subspace spanned by $\mathbf{x}_1, \dots, \mathbf{x}_k$ and orthogonal to the vector of all-ones.

We will be particularly interested in the inner products $\langle \mathbf{P}, A_{\mathbf{G}}^{(s)} \rangle$, which count non-backtracking walks on $\mathbf{G}$, with weight 1 if the endpoints share a group label, and $-(k-1)^{-1}$ if they do not. The following is a consequence of Lemma 39 in the sequel.

Lemma 29 *For every $s \geq 1$ and increasing, nonnegative function $\Delta(n)$,*

$$\mathbb{P} \left[\left| \langle \mathbf{P}, A_{\mathbf{G}}^{(s)} \rangle - q_s(d\eta)n \right| > \Delta(n) \right] = O \left(\frac{n}{\Delta(n)^2} \right).$$

The *Symmetric Path Statistics SDP*, given as input a graph G_0 , attempts to find a “pseudo-partition matrix,” i.e. PSD matrix with ones on the diagonal, and whose inner products with the matrices $A_{G_0}^{(s)}$ are equal to $q_s(d\eta)n$, at least up to the fluctuations in Lemma 29.

Definition 30 (Symmetric Path Statistics) *The level- D Symmetric Path Statistics Algorithm with error tolerance $\delta > 0$, on input a d -regular graph G_0 on n vertices, is the following SDP: find $\tilde{\mathbf{P}} \succeq 0$ so that*

1. $\tilde{P}_{u,u} = 1$ for every $u \in [n]$
2. $|\langle \tilde{\mathbf{P}}, J_n \rangle| \leq \delta n^2$
3. $|\langle \tilde{\mathbf{P}}, A_{G_0}^{(s)} \rangle - q_s(d\eta)n| \leq \delta n$ for every $s \in [D]$.

Theorem 31 *If $(d\eta)^2 > 4(d-1)$, then for every $D \geq 2$ there exists an error tolerance $\delta > 0$ at which the level D Symmetric Path Statistics SDP can w.h.p. distinguish $\mathbb{P}$ and $\mathbb{Q}$. If $(d\eta)^2 \leq 4(d-1)$, then no such D and δ exist.*

Proof By Lemma 8.5, this SDP is with high probability feasible on input $\mathbf{G} \sim \mathbb{P}$. Our proof will therefore show that when η^2 is sufficiently large, the SDP for some constant D is *infeasible* on input $\mathbf{G} \sim \mathbb{Q}$, whereas for η^2 sufficiently small, it is feasible for every constant D .

First, fix $D \geq 2$ and assume $(d\eta)^2 > 4(d-1)$. We will show that there exists $\delta > 0$ so that with high probability the level- D Symmetric Path Statistics is infeasible on input $\mathbf{G} \sim \mathbb{Q}$. Our strategy will be to find a polynomial f with the property that $f(A_{\mathbf{G}}) \succeq 0$ with high probability, but we can deduce $\langle \mathbf{P}, f(A_{\mathbf{G}}) \rangle < 0$ from the affine constraints in Definition 30.

Let f be a degree D polynomial which is strictly positive on the closed interval $[-2\sqrt{d-1}, 2\sqrt{d-1}]$ and satisfies $f(d\eta) < 0$; our assumption on η ensures that this is possible, for instance by setting $f(z) = 2(d-1) + \frac{1}{2}(d\eta)^2 - z^2$. From our preliminaries on non-backtracking walks and the polynomials q_s we know

$$f = \sum_{s=0}^D \frac{\langle f, q_s \rangle_{\text{KM}}}{\|q_s\|_{\text{KM}}^2} q_s.$$

When $\mathbf{G} \sim \mathbb{Q}$, $A_{\mathbf{G}}$ has an eigenvalue at d whose eigenvector is the all-ones vector and by Friedman’s Theorem (Friedman, 2003) its remaining eigenvalues with high probability have absolute value at

most $2\sqrt{d-1}+o(1)$. Our assumptions on f therefore imply $f(A_G) - f(d)J/n = f(A_G - dJ/n) \succeq 0$ with high probability.

On the other hand, if $\tilde{\mathbf{P}} \succeq 0$ is a feasible solution for the degree- D Symmetric Path Statistics SDP on input $\mathbf{G} \sim \mathbb{Q}_n$, then

$$\begin{aligned} 0 &\leq \langle \tilde{\mathbf{P}}, f(A_G) - f(d)J/n \rangle \\ &= \sum_{s=0}^D \frac{\langle f, q_s \rangle_{\text{KM}}}{\|q_s\|_{\text{KM}}^2} \langle \tilde{\mathbf{P}}, q_s(A_G) \rangle + \delta |f(d)|n \\ &\leq \sum_{s=0}^D \frac{\langle f, q_s \rangle_{\text{KM}}}{\|q_s\|_{\text{KM}}^2} q_s(d\eta)n + \delta (\|f\|_{\text{KM}}^2 + |f(d)|)n \\ &= (f(d\eta) + \delta (\|f\|_{\text{KM}}^2 + |f(d)|))n < 0, \end{aligned}$$

if we set δ sufficiently small.

We can now turn to the case $(d\eta)^2 \leq 4(d-1)$, seeking to prove that on input $\mathbf{G} \sim \mathbb{Q}$, with high probability *every* level of the Symmetric Path Statistics hierarchy is feasible, for every error tolerance δ . We will use the following lemma, which may be proved by adapting the proof of (Banks et al., 2019, Proposition 4.8). The proof proceeds by setting $\tilde{\mathbf{P}}$ equal to a mild modification of the matrix $g(A_G) - g(d)J_n/n$.⁸

Lemma 32 *Assume there exists a constant-degree polynomial $g \in \mathbb{R}[z]$ that is strictly positive on $[-2\sqrt{d-1}, 2\sqrt{d-1}]$ and satisfies*

$$|\langle g, q_s \rangle_{\text{KM}} - q_s(d\eta)| \leq \delta$$

for every $s = 1, \dots, D$. Then the level- D Symmetric Path Statistics SDP with error tolerance δ is w.h.p. feasible on input $\mathbf{G} \sim \mathbb{Q}$.

Lemma 8.9 in hand, we need only to construct such a polynomial. Assume that $(d\eta)^2 \leq 4(d-1)$, so that $d\eta \in [-2\sqrt{d-1}, 2\sqrt{d-1}]$, the interval in which lie the roots of every polynomial q_s . Let $P \gg D$, and write $r_1 < r_2 < \dots < r_P$ for the roots of q_P . Let $I \subset [P]$ contain the indices of the $(D+1)/2$ roots of q_P closest to $d\eta$, and set

$$g_\eta = \frac{1}{\zeta} \prod_{i \notin I} (z - r_i)^2,$$

where ζ is a normalizing factor to ensure that $\langle g_\eta, 1 \rangle_{\text{KM}} = 1$.

This polynomial is certainly nonnegative, and its degree is $2P - D - 1$. From Lemma 8.4, then, there exist $w_1, \dots, w_P \geq 0$ so that for any $s = 0, \dots, D$

$$\langle g_\eta, q_s \rangle_{\text{KM}} = \sum_{i \in I} w_i g_\eta(r_i) q_s(r_i).$$

8. The referenced result in Banks et al. (2019) was proved for an SDP which shared constraints (1) and (2) from Definition 30, but in which constraint (3) read $\langle P, A_{G_0}^{(s)} \rangle = \lambda^s \|q_s\|_{\text{KM}}^2 n$. The proof may be adapted simply by adopting the hypotheses below, changing every instance of the aforementioned constraints, and taking some care with the δ slack.

In particular, setting $s = 0$ and recalling the definition of ζ , we have

$$1 = \langle g_\eta, 1 \rangle_{\text{KM}} = \sum_{i \in I} w_i g_\eta(r_i).$$

This means that for any $s = 0, \dots, D$, the inner product $\langle g_\eta, q_s \rangle$ is a weighted average of q_s evaluated at the $(D+1)/2$ roots of q_P closest to the point $d\eta$. Since the roots of the q polynomials are dense in $(-2\sqrt{d-1}, 2\sqrt{d-1})$, and $d\eta$ is in the closure of this interval, for each D and $\delta > 0$, there exists a constant P for which $|\langle g_\eta, q_s \rangle_{\text{KM}} - q_s(d\eta)| < \delta$ for every $s = 0, \dots, D$. $\blacksquare$

D.3. Partially Labelled Subgraphs

We are now ready to study the full Local Statistics algorithm. To start, we will need to develop a basis for the symmetric polynomials appearing as affine moment matching constraints in Definition 26. Because $\mathbb{E}p(\mathbf{x}, \mathbf{G}) = 0$ for any $p \in \mathcal{I}_k$, a constraint shared by the pseudoexpectation, it would suffice to study the subspace of $\mathbb{R}[x, G]/\mathcal{I}_k$ fixed under the S_n action inherited from $\mathbb{R}[x, G]$. However, it is computationally favorable to work in a slightly larger vector space instead.

Definition 33 *Let us write $\mathbb{S}[x, G]_{\leq D_x, D_G} \subset \mathbb{R}[x, G]$ for the vector space of polynomials that (1) satisfy $\deg_x \leq D_x$ and $\deg_G \leq D_G$, (2) are symmetric with respect to the S_n action, (3) are multilinear in G and x , and (4) for which at most one of $x_{u,1}, \dots, x_{u,k}$ appears in each monomial, for every $u \in V$.*

Every polynomial appearing as a moment constraint in the level- (D_x, D_G) Local Statistics algorithm belongs to $\mathbb{S}[x, G]_{\leq D_x, D_G}$. We now give a combinatorially structured basis for this vector space, similar to the ‘shapes’ of Barak et al. (2019).

Definition 34 *A partially labelled graph (H, S, τ) consists of a graph H , a subset $S \subset V(H)$, and a map $\tau : S \rightarrow [k]$; we say that a graph is fully labelled if $S = V(H)$, and in this case write (H, τ) for short. A homomorphism from (H, S, τ) into a fully labelled graph (G, σ) is a map $\phi : V(H) \rightarrow V(G)$ that takes edges to edges and agrees on labels; an occurrence of (H, S, τ) in (G, σ) is an injective homomorphism. For each partially labelled graph (H, S, τ) there is an associated polynomial in $\mathbb{S}[x, G]$,*

$$p_{(H, S, \tau)}(x, G) = \sum_{\phi: V(H) \hookrightarrow [n]} \prod_{(u, v) \in E(H)} G_{\phi(u), \phi(v)} \prod_{u \in S} x_{\phi(u), \tau(u)}.$$

Each point in the zero locus of $\mathcal{I}_k$ may be identified with a fully labelled graph (G, σ) . Evaluated at such a point, this polynomial counts the number of occurrences of (H, S, τ) in (G, σ) . Finally, $\deg_x p_{(H, S, \tau)} = |S|$ and $\deg_G p_{(H, S, \tau)} = |E(H)|$.

Lemma 35 *The polynomials $p_{(H, S, \tau)}$ with $|E(H)| \leq D_G$ and $|S| \leq D_x$ are a vector space basis for $\mathbb{S}[x, G]_{\leq D_x, D_G}$*

Proof Let $s : \mathbb{R}[x, G] \rightarrow \mathbb{S}[x, G]$ be the map that sends a polynomial p to the sum over its S_n orbit. The vector space $\mathbb{S}[x, G]_{\leq D_x, D_G}$ is spanned by the images under s of the multilinear monomials with x -degree D_x and G -degree D_G in $\mathbb{R}[x]$ in which at most one of $x_{u,1}, \dots, x_{u,k}$ appears

for each $u \in [n]$. From each such monomial $m(x, G)$ one can extract a partially labelled graph (H_m, S_m, τ_m) , where $|S_m| \leq D_x$, $|E(H_m)| \leq D_G$, and H_m is a subgraph of the complete graph: $E(H_m)$ is the union of all pairs (u, v) appearing as an index of a G variable, S_m is the union of all u occurring as an index of an x variable, $V(H_m)$ is the union of S and the endpoints of every edge, and $\tau(u) = i$ if the variable $x_{u,i}$ appears. Because S_n acts transitively on $[n]$, the orbit of $m(x)$ corresponds to every possible injection $V(H_m) \hookrightarrow [n]$, and thus

$$s(m(x, G)) = p_{(H_m, S_m, \tau_m)}(x, G).$$

This shows that the $p_{(H, S, \tau)}$ span. To see that they are independent, observe that each monomial appears as a term in exactly one $p_{(H, S, \tau)}$. ■

In view of this lemma, the moment constraints in Definition 26 are equivalent to the requirement that $\mathbb{E}p_{(H, S, \tau)}(x, G_0) \approx \mathbb{E}p_{(H, S, \tau)}(\mathbf{x}, \mathbf{G})$ for every (H, S, τ) with at most D_G edges and D_x distinguished vertices, up to isomorphism. In order to instantiate and analyze the Local Statistics algorithm, we now need to compute these expectations, and bound the fluctuations around them.

D.4. Local Statistics in the Planted Model

Instead of working directly in $\mathbb{P}$, we will as usual work in the *configuration model* $\widehat{\mathbb{P}}$, a distribution on multigraphs with two key properties: (i) with probability bounded away from zero as $n \rightarrow \infty$, $\widehat{\mathbf{G}} \sim \widehat{\mathbb{P}}$ is simple, and (ii) the conditional distribution of $\widehat{\mathbb{P}}$ on this event is equal to $\mathbb{P}$. In this model, conditional on a balanced partition σ , we adorn each vertex in every group i with $dM_{i,j}$ ‘half-edges’ labelled $i \rightarrow j$, and then for each $i, j \in [k]$ randomly match the $i \rightarrow j$ half-edges with the $j \rightarrow i$ half-edges. When $k = 1$, this is the usual configuration model on d -regular graphs, and when $k = 2$ and $M_{i,i} = 0$, it gives bipartite regular graphs. Thus many results we prove for $\mathbb{P}$ will apply to $\mathbb{Q}$ as well.

Claim 36 Write $\mathcal{G}_n$ and $\widehat{\mathcal{G}}_n$ for the sets of all d -regular, n -vertex graphs and multigraphs, respectively. If $\widehat{\mathcal{E}}_n \subset \widehat{\mathcal{G}}_n$ is a sequence of events holding w.h.p. in $\widehat{\mathbb{P}}_n$, then $\widehat{\mathcal{E}}_n \cap \mathcal{G}_n$ holds w.h.p. in $\mathbb{P}_n$.

Proof Since $\widehat{\mathbb{P}}_n[\mathcal{G}_n]$ is bounded away from zero and $\widehat{\mathbb{P}}_n[\widehat{\mathcal{E}}_n] = 1 - o_n(1)$, we have

$$\mathbb{P}_n[\widehat{\mathcal{E}}_n \cap \mathcal{G}_n] = \frac{\widehat{\mathbb{P}}_n[\widehat{\mathcal{E}}_n \cap \mathcal{G}_n]}{\widehat{\mathbb{P}}_n[\mathcal{G}_n]} \geq \frac{\widehat{\mathbb{P}}_n[\mathcal{G}_n] - o_n(1)}{\widehat{\mathbb{P}}_n[\mathcal{G}_n]} = 1 - o_n(1). \quad \blacksquare$$

As a warm-up, let us recall some standard calculations of subgraph probabilities in these two simpler situations.

Lemma 37 Let $\widehat{\mathbf{G}}$ be a multi-graph produced by the d -regular configuration model on n vertices. If H is a simple graph, the probability that a fixed injection $\phi : V(H) \hookrightarrow V(\widehat{\mathbf{G}})$ is a homomorphism is

$$\prod_{v \in V(H)} \frac{d!}{(d - \deg(v))!} \cdot \frac{(dn - 2|E(H)| - 1)!!}{(dn - 1)!!} = \prod_{v \in V(H)} \frac{d!}{(d - \deg(v))!} \cdot (dn)^{-|E(H)|} + O\left(n^{-|E(H)|-1}\right).$$

Similarly, let $(\widehat{\mathbf{G}}, \sigma)$ be generated by the d -biregular configuration model on $2n$ vertices, $\sigma : V(\widehat{\mathbf{G}}) \rightarrow [2]$ its left/right labelling. If (H, τ) is a simple bipartite graph with left/right labelling τ , then the probability that a fixed injection $\phi : V(H) \hookrightarrow V(\widehat{\mathbf{G}})$ agreeing on labels is a homomorphism is

$$\prod_{v \in V(H)} \frac{d!}{(d - \deg(v))!} \cdot \frac{(dn - |E(H)|)!}{(dn)!} = \prod_{v \in V(H)} \frac{d!}{(d - \deg(v))!} \cdot (dn)^{-|E(H)|} + O\left(n^{-|E(H)|-1}\right).$$

Proof [Sketch] In the case of the d -regular configuration model, once an injective ϕ has been chosen, there are $\frac{d!}{(d - \deg(v))!}$ ways to choose $\deg(v)$ stubs from the d available at each vertex v to be matched with the appropriate stubs at each intended neighbor of v , and then $(dn - 2|E(H)| - 1)!!$ ways to match the remaining stubs. On the other hand, the total number of multi-graphs possible to output is $(dn - 1)!!$. The calculation is analogous in the bipartite regular case. $\blacksquare$

We will use Lemma 37 to compute local statistics in the planted model, but first we will need a bit more notation. Let (H, S, τ) be a partially labelled graph, and $\widehat{\tau}$ an extension of τ , i.e. $\widehat{\tau} : V(H) \rightarrow [k]$, and $\widehat{\tau}|_S = \tau$; for each $i, j \in [k]$ and $u \in V_i(H)$, write $\deg_j(u)$ for the number of neighbors that u has in group j according to $\widehat{\tau}$. Let us define

$$(dM)_{\widehat{\tau}}^{(H, S)} := \sum_{\widehat{\tau} : \widehat{\tau}|_S = \tau} \frac{\prod_v \prod_{j \in [k]} dM_{\widehat{\tau}(v), j} (dM_{\widehat{\tau}(v), j} - 1) \cdots (dM_{\widehat{\tau}(v), j} - \deg_{i \rightarrow j}(v) + 1)}{\prod_{(u, v) \in E(H)} dM_{\widehat{\tau}(u), \widehat{\tau}(v)}} \quad (54)$$

if $dM_{i, j} \geq \deg_j(v)$ for every $i, j \in [k]$ and $v \in \widehat{\tau}^{-1}(i)$, and zero otherwise. Note that this operation is multiplicative on disjoint unions:

$$(dM)_{\tau_1 \sqcup \tau_2}^{(H_1 \sqcup H_2, S_1 \sqcup S_2)} = (dM)_{\tau_1}^{(H_1, S_1)} (dM)_{\tau_2}^{(H_2, S_2)}.$$

Finally, let us write $\chi(H) := |V(H)| - |E(H)|$ and $\text{cc}(H)$ for the number of connected components.

Since we are aiming to prove high probability statements regarding $p_{(H, S, \tau)}(\mathbf{x}, \mathbf{G})$ for $(\mathbf{x}, \mathbf{G}) \sim \mathbb{P}$ by studying the configuration model, we need to extend the quantities $p_{(H, S, \tau)}(\mathbf{x}, G)$ to the case when G is a multigraph with self-loops. For convenience, we will define an *occurrence* of (H, S, τ) in a fully labelled, loopy multigraph as an occurrence of (H, S, τ) in the simple graph obtained by deleting all self-loops and merging all multiedges between each pair of vertices. Our key lemma computes the expected number of occurrences in the configuration model.

Lemma 38 *Let (H, S, τ) be a partially labelled graph on $O(1)$ edges, and $(\mathbf{x}, \widehat{\mathbf{G}})$ be drawn from the configuration model $\mathbb{P}$. Then*

$$\mathbb{E} p_{(H, S, \tau)}(\mathbf{x}, \widehat{\mathbf{G}}) = (n/k)^{\chi(H)} (dM)_{\widehat{\tau}}^{(H, S)} + O\left(n^{\chi(H)-1}\right).$$

Proof Let $V(\widehat{\mathbf{G}}) = [n]$, fix a labelling $\sigma : [n] \rightarrow [k]$, and let $\widehat{\mathbf{G}}$ be drawn from the configuration model. Fix an extension $\widehat{\tau}$ of τ . If $\phi : V(H) \hookrightarrow V(\widehat{\mathbf{G}})$ is an injection that agrees on labels, applying

Lemma 8.9 to the multigraphs on each set of vertices $\sigma^{-1}(i)$ and between each pair of sets $\sigma^{-1}(i)$ and $\sigma^{-1}(j)$,

$$\begin{aligned} \mathbb{P}[\phi \text{ is an occurrence}] &= (n/k)^{|E(H)|} \cdot \prod_{i \leq j} (dM_{i,j})^{-|E_{i,j}(H, \hat{\tau})|} \cdot \prod_{v \in V(H, \hat{\tau})} \prod_{i,j} \frac{(dM_{i,j})!}{(dM_{i,j} - \deg_j(v))!} \\ &\quad + O\left(n^{-|E(H)|-1}\right), \end{aligned}$$

and there are $(n/k)^{V(H)} + O(n^{|V(H)|-1})$ injective choices for ϕ that agree on labels.

Finally, writing $\Phi(H, S, \tau)$ for the total number of occurrences of (H, S, τ) in $\hat{\mathbf{G}}$, and $\Phi(H, \hat{\tau})$ for the number of occurrences of the fully labelled graph $(H, \hat{\tau})$,

$$\begin{aligned} \mathbb{E}\Phi(H, S, \tau) &= \mathbb{E} \sum_{\hat{\tau}: \hat{\tau}|_S = \tau} \Phi(H, \hat{\tau}) \\ &= (n/k)^{\chi(H)} (dM)_\tau^{(H, S)} + O\left(n^{|V(H)|-|E(H)|-1}\right). \end{aligned}$$

■

This lemma has some immediate consequences. First, it tells us that occurrences of partially labelled forests are sharply concentrated.

Lemma 39 *Let $(H, S, \tau) = \bigsqcup_{t \in [\text{cc}(H)]} (H_t, S_t, \tau_t)$ be a partially labelled graph with $O(1)$ vertices, $\text{cc}(H)$ connected components H_t , and no cycles. Then for any function $f(n) > 0$,*

$$\mathbb{P} \left[\left| p_{(H, S, \tau)}(\mathbf{x}, \mathbf{G}) - (n/k)^{\text{cc}(H)} \prod_{t \in [\text{cc}(H)]} (dM)_{\tau_t}^{(H_t, S_t)} \right| > f(n) \right] = O\left(\frac{n^{2\text{cc}(H)-1}}{f(n)^2}\right).$$

Proof The expectation $\mathbb{E}[p_{(H, S, \tau)}(\mathbf{x}, \mathbf{G})^2]$ is a sum over all pairs of injective, label-consistent maps ϕ_1, ϕ_2 , of the probability that both maps are occurrences. The image of two disjoint copies of H under these two injective maps is a graph H' that is either $H \sqcup H$, or is obtained by identifying some pairs of vertices whose τ labels agree—each pair with one vertex from each copy of H . We can promote H' to a partially labelled graph by taking the induced partial labelling τ' from τ . Thus, let us think of the pair ϕ_1, ϕ_2 as a single injective map $\varphi : V(H') \hookrightarrow V(\hat{\mathbf{G}})$ that agrees with τ' . Thus

$$\begin{aligned} \mathbb{E}[p_{(H, S, \tau)}(\mathbf{x}, \mathbf{G})^2] &= \sum_{H'} \mathbb{E} p_{(H', S', \tau')}(\mathbf{x}, \mathbf{G}) \\ &= \sum_{H'} \left((n/k)^{\chi(H')} (dM)_{\tau'}^{(H', S')} + O\left(n^{\chi(H')-1}\right) \right). \end{aligned}$$

When $H' = H \sqcup H$, from our observation above

$$(dM)_{\tau'}^{(H', S')} = ((dM)_\tau^{(H, S)})^2 = \left(\prod_{t \in [\text{cc}(H)]} ((dM)_{\tau_t}^{(H_t, S_t)}) \right)^2.$$

As H has no cycles, every other H' satisfies $\chi(H') < 2\chi(H)$. Thus since $\text{cc}(H) = \chi(H)$, the assertion is true in the configuration model by Chebyshev, and transfers immediately to the planted model. For good measure, an application of the triangle inequality shows as well that

$$\mathbb{P} \left[\left| p_{(H,S,\tau)}(\mathbf{x}, \mathbf{G}) - \prod_{t \in [\text{cc}(H)]} p_{(H_t, S_t, \tau_t)}(\mathbf{x}, \mathbf{G}) \right| > f(n) \right] = O \left(\frac{n^{2\text{cc}(H)-1}}{f(n)^2} \right).$$

■

Thus with high probability, the counts of partially labelled forests enjoy concentration of $\pm o(n^{\text{cc}(H)})$. On the other hand, an immediate application of Markov in $\tilde{\mathbb{P}}$ tells us that there are very few occurrences of partially labelled graphs with cycles.

Lemma 40 *Let (H, S, τ) be a partially labelled graph with $O(1)$ edges and at least one cycle. Then for any function $f(n) > 0$,*

$$\mathbb{P} [p_{(H,S,\tau)}(\mathbf{x}, \mathbf{G}) > f(n)] = O \left(\frac{n^{\text{cc}(H)-1}}{f(n)} \right).$$

In particular, we will need later on the fact that there are very few vertices within constant distance of a constant length cycle. The proof is once again Markov, combined with a union bound.

Lemma 41 *Let $\mathbf{G} \sim \mathbb{P}$ or $\mathbb{Q}$. Fix constants L and C , and call a vertex bad if it is at most L steps a way from a cycle of length at most C . Then w.h.p. there are fewer than $f(n)$ bad vertices, for any increasing function $f(n)$.*

We can now restate the local statistics algorithm more precisely.

Definition 42 (Local Statistics Algorithm with Formal Moment Constraints) *The degree- (D_x, D_G) Local Statistics algorithm with error tolerance δ is the following SDP. Given an input graph G_0 , find $\tilde{\mathbb{E}} : \mathbb{R}[x]_{\leq D_x} \rightarrow \mathbb{R}$ s.t.*

1. (Positivity) $\tilde{\mathbb{E}}p(x)^2 \geq 0$ whenever $\deg p^2 \leq D_x$
2. (Hard Constraints) $\tilde{\mathbb{E}}p(x, G_0) = 0$ for every $p \in \mathcal{I}_k$
3. (Moment Constraints) For every (H, S, τ) with at most D_x distinguished vertices, D_G edges, and ℓ connected components,

$$\tilde{\mathbb{E}}p_{(H,S,\tau)}(x, G_0) = (dM)_\tau^{(H,S)}(n/k)^{\chi(H)} \pm \delta n^{\text{cc}(H)}.$$

The $n^{\chi(H)}$ vs. $n^{\text{cc}(H)}$ scaling may seem ad hoc, but as promised above we have arranged things so that the the SDP is w.h.p. feasible when its input is drawn from the planted model.

Lemma 43 *Fix D_x, D_G constant, and $\delta > 0$. with high probability, the degree- (D_x, D_G) Local Statistics Algorithm with error tolerance δ is feasible on input $\mathbf{G} \sim \mathbb{P}$.*

Proof Let $(\mathbf{x}, \mathbf{G}) \sim \mathbb{P}$, and for each $p(x, G) \in \mathbb{R}[x, G]_{\leq D_x, D_G}$ define

$$\tilde{\mathbb{E}}p(\mathbf{x}, \mathbf{G}) = p(\mathbf{x}, \mathbf{G}).$$

This satisfies positivity, as $\tilde{\mathbb{E}}p(\mathbf{x}, \mathbf{G})^2 = p(\mathbf{x}, \mathbf{G})^2 \geq 0$, and obeys the hard constraints because $(\mathbf{x}, \mathbf{G})$ lies in the zero locus of $\mathcal{I}_k$. Finally, let (H, S, τ) be a partially labelled graph. If H has a cycle, then by Corollary 8.15, w.h.p.

$$\left| p_{(H, S, \tau)}(\mathbf{x}, \mathbf{G}) - (n/k)^{\chi(H)} (dM)_\tau^{(H, S)} \right| \leq p_{(H, S, \tau)}(\mathbf{x}, \mathbf{G}) + O(n^{\chi(H)}) \leq \delta n^{\mathbf{cc}(H)}.$$

On the other hand, if H has no cycles, then $\chi(H) = \mathbf{cc}(H)$ and w.h.p.

$$\left| p_{(H, S, \tau)}(\mathbf{x}, \mathbf{G}) - (n/k)^{\chi(H)} (dM)_\tau^{(H, S)} \right| \leq \delta n^{\mathbf{cc}(H)}$$

by Proposition 8.14. There are only constantly many partially labelled subgraphs with at most D_x vertices and D_G edges, so a union bound finishes the proof. $\blacksquare$

Finally, we end this subsection with the proof of Lemma 29, which concerned the affine constraints in the Symmetric Path Statistics SDP

Proof of Lemma 29 Recall the partition matrix

$$\mathbf{P} = \frac{k}{k-1} \left(\sum_{i \in [k]} \mathbf{x}_i \mathbf{x}_i^\top - \frac{1}{k} \sum_{i, j \in [k]} \mathbf{x}_i \mathbf{x}_j^\top \right)$$

from Section 8.2, where each $\mathbf{x}_i \in \{0, 1\}^n$ is the indicator vector for membership in the i th group. We are interested in $\langle \mathbf{P}, A_{\mathbf{G}}^{(s)} \rangle$.

Let $(P_s, \{0, s\}, \{i, j\})$ denote a path of length s with distinguished endpoints labelled i and j , and write its vertices as $V(P_s) = \{0, 1, \dots, s\}$. From Lemma 39, w.h.p. for $(\mathbf{x}, \mathbf{G}) \sim \mathbb{P}$,

$$p_{(P_s, \{0, s\}, \{i, j\})}(\mathbf{x}, \mathbf{G}) = \frac{1}{k} (dM)_{i, j}^{(P_s, \{0, s\})} n \pm o(n).$$

Expanding the right hand side,

$$\begin{aligned} (dM)_{i, j}^{(P_s, \{0, s\})} &= \sum_{\hat{\tau}: \hat{\tau}|_S = \tau} \frac{\prod_{v \in V(P_s)} \prod_{j' \in [k]} dM_{\hat{\tau}(v), j'} (dM_{\hat{\tau}(v), j'} - 1) \cdots (dM_{\hat{\tau}(v), j'} - \deg_j(v))}{\prod_{(u, v) \in E(P_s)} dM_{\hat{\tau}(u), \hat{\tau}(v)}} \\ &= \sum_{\hat{\tau}: \hat{\tau}|_S = \tau} dM_{i, \hat{\tau}(1)} \prod_{t=1, \dots, s-1} (dM_{\hat{\tau}(t), \hat{\tau}(t+1)} - \{\hat{\tau}(t+1) = \hat{\tau}(t-1)\}) \\ &= q_s(dM)_{i, j}. \end{aligned}$$

The expression in the second to last line counts the number of non-backtracking walks of length s between vertices i and j on the multi-graph whose adjacency matrix is dM ; from Section D.1 these may be enumerated using the polynomial q_s applied to dM .

Now, let us define $A_{\mathbf{G}}^{(s)}$ as the $n \times n$ matrix whose entries count *self-avoiding* (as opposed to non-backtracking) walks on $\mathbf{G}$. By definition

$$p_{(P_s, \{0, s\}, \{i, j\})}(\mathbf{x}, \mathbf{G}) = \sum_{u, v} \left(A_{\mathbf{G}}^{(s)} \right)_{u, v} \cdot x_{u, i} x_{v, j}.$$

Thus, with high probability

$$\begin{aligned}
 \langle \mathbf{P}, A_{\mathbf{G}}^{(s)} \rangle &= \frac{k}{k-1} \left(\sum_i p_{(P_s, \{0,s\}, \{i,i\})}(\mathbf{x}, \mathbf{G}) - \frac{1}{k} \sum_{i,j} p_{(P_s, \{0,s\}, \{i,j\})}(\mathbf{x}, \mathbf{G}) \right) \\
 &= \frac{1}{k-1} \langle q_s(dM), I - J/k \rangle n \pm o(n) \\
 &= \frac{1}{k-1} \left(\text{Tr } q_s(dM) - q_s(d) \right) n + o(n) \\
 &= q_s(d\lambda) n + o(n).
 \end{aligned}$$

The last equality follows from the fact that the spectrum of dM consists of an eigenvalue d with multiplicity one, and an eigenvalue $d\lambda$ with multiplicity $k-1$.

We need finally to ensure that the same inner product constraint holds for the matrices $A_{\mathbf{G}}^{(s)}$. It is an easy consequence of Lemma 8.16 that w.h.p. for $\mathbf{G} \sim \mathbb{Q}$, the matrices $A_{\mathbf{G}}^{(s)}$ and $A_{\mathbf{G}}^{(s)}$ disagree in at most $o(n)$ rows. Thus, since the L^1 norm of every row is bounded by a constant (by degree-regularity), w.h.p. $\|A_{\mathbf{G}}^{(s)} - A_{\mathbf{G}}^{(s)}\|_F^2 = o(n)$. Since X is PSD with ones on the diagonal, every off-diagonal element has magnitude at most one—thus

$$\langle X, A_{\mathbf{G}}^{(s)} \rangle = \langle X, A_{\mathbf{G}}^{(s)} \rangle + o(n),$$

and we are done. ■

D.5. Proof of Theorem 27: Upper Bound

We will show that if $(d\lambda)^2 > 4(d-1)$, then the degree $(2, D)$ Local Statistics algorithm can distinguish $\mathbb{P}$ and $\mathbb{Q}$ for every $D \geq 2$. Specifically, we will show that for any such D , with high probability over input $\mathbf{G} \sim \mathbb{Q}$ there exists a δ at which the SDP is *infeasible*. Our goal, here and in the proof of the lower bound, will be to reduce to our characterization of the Symmetric Path Statistics SDP in Theorem 8.7.

Assume that $\tilde{\mathbb{E}}$ is a feasible pseudoexpectation for the degree $(2, D)$ Local Statistics SDP with tolerance $\delta > 0$, on input $\mathbf{G} \sim \mathbb{Q}$, and consider the matrix X with entries

$$\tilde{P}_{u,v} = \frac{k}{k-1} \left(\sum_{i \in [k]} \tilde{\mathbb{E}}(x_{u,i} - 1/k)(x_{v,i} - 1/k) \right) = \frac{k}{k-1} \left(\sum_{i \in [k]} \tilde{\mathbb{E}}x_{u,i}x_{v,i} - \frac{1}{k} \sum_{i,j \in [k]} \tilde{\mathbb{E}}x_{u,i}x_{v,j} \right).$$

We will show that $\tilde{\mathbf{P}}$ is a feasible solution to the level- D Symmetric Path Statistics SDP with the input $\mathbf{G}$, at some tolerance $\delta' = c\delta$ —thus by Theorem 8.7, when $(d\lambda)^2 > 4(d-1)$ and δ is sufficiently small, we will have a contradiction. We observe first that X is PSD with ones on the diagonal—these facts follow immediately from the hard constraints in Definition 8.17, which tell us that $\tilde{\mathbb{E}}x_{u,i}^2 = \tilde{\mathbb{E}}x_{u,i}$ for every u and i , and $\tilde{\mathbb{E}} \sum_{i \in [k]} x_{u,i} = 1$ for every u .

We turn now to the moment constraints, with the goal of showing

$$\begin{aligned}
 \langle \tilde{\mathbf{P}}, A_{\mathbf{G}}^{(s)} \rangle &= q_s(d\lambda)n \pm \delta' n & \forall s \in [D] \\
 \langle \tilde{\mathbf{P}}, J \rangle &= 0 \pm \delta' n^2.
 \end{aligned}$$

Rehashing our calculations from the proof of Lemma 29, we find that w.h.p.

$$\begin{aligned}
\langle \mathbf{P}, A_{\mathbf{G}}^{(s)} \rangle &= \langle \tilde{\mathbf{P}}, A_{\mathbf{G}}^{(s)} \rangle + o(n) \\
&= \tilde{\mathbb{E}} \frac{k}{k-1} \left(\sum_i p_{(P_s, \{0, s\}, \{i, i\})}(x, \mathbf{G}) - \frac{1}{k} \sum_{i, j} p_{(P_s, \{0, s\}, \{i, j\})}(x, \mathbf{G}) \right) + o(n) \\
&= \frac{1}{k-1} \langle q_s(dM), I - J/k \rangle n \pm \frac{k}{k-1} \cdot 2k\delta n + o(n) \\
&= q_s(d\lambda)n \pm \frac{k^3}{k-1} \delta n + o(n).
\end{aligned}$$

To verify that X has the correct inner product against the all-ones matrix, consider two partially labelled subgraphs: a single vertex labelled $i \in [k]$, and two disjoint vertices both labelled $i \in [k]$. The sum of their corresponding polynomials is $\sum_{u, v} x_{u, i} x_{v, i}$, and identically in the planted model

$$\sum_{u, v} x_{u, i} x_{v, i} = (n/k)^2.$$

Our pseudoexpectation is required to match this up to an additive $\delta(n + n^2)$, the n and n^2 terms respectively coming from the additive slack in the one vs. two vertex graphs. Thus

$$\langle \tilde{\mathbf{P}}, J \rangle = \frac{k}{k-1} (n^2/k \pm \delta(n^2 + n) - n^2/k) = 0 \pm \frac{2k}{k-1} \delta n^2.$$

D.6. Proof of Theorem 27: Lower Bound

Assume that $(d\lambda)^2 \leq 4(d-1)$; we need to explicitly construct a degree- $(2, D)$ pseudoexpectation that is with high probability feasible for $\mathbf{G} \sim \mathbb{Q}$. Our tactic will be to show that such an operator can be constructed from a feasible solution to the Symmetric Path Statistics SDP guaranteed us by Theorem 8.7.

Before building the degree- $(2, D)$ pseudoexpectation asserted to exist in the theorem statement, we will first prove a series of structural lemmas showing that it suffices to check only a subset of the moment constraints of a Local Statistics pseudoexpectation.

First, we show that the moment constraints regarding the pseudo-expected counts of partially labelled graphs containing cycles are satisfied more or less for free.

Lemma 44 *Let $\mathbf{G} \sim \mathbb{Q}$, and $\tilde{\mathbb{E}} = \tilde{\mathbb{E}}(\mathbf{G})$ be a degree (D_x, D_G) pseudoexpectation, perhaps dependent on $\mathbf{G}$, that satisfies positivity and the hard constraints. For every error tolerance δ , w.h.p. $\tilde{\mathbb{E}}$ satisfies the moment constraints for all partially labelled subgraphs containing a cycle.*

Proof It is a routine sum-of-squares calculation that for any monomial $\mu(x)$,

$$(\tilde{\mathbb{E}}\mu(x))^2 \leq \tilde{\mathbb{E}}\mu(x)^2 = \tilde{\mathbb{E}}\mu(x),$$

meaning that $|\tilde{\mathbb{E}}\mu(x)| \leq 1$. Thus for any (H, S, τ) ,

$$|\tilde{\mathbb{E}}p_{(H, S, \tau)}(x, \mathbf{G})| = \left| \sum_{\phi: V(H) \hookrightarrow [n]} \prod_{(\alpha, \beta) \in E(H)} \mathbf{G}_{\phi(\alpha), \phi(\beta)} \tilde{\mathbb{E}} \prod_{\alpha \in S} x_{\phi(\alpha), \tau(\alpha)} \right| \leq \left| \sum_{\phi: V(H)} \prod_{(\alpha, \beta) \in E(H)} \mathbf{G}_{\phi(\alpha), \phi(\beta)} \right|,$$

and the right hand side is simply the number of occurrences of the unlabelled graph H in G . From Proposition 8.15, if H has $\text{cc}(H)$ connected components and at least one cycle, (i) w.h.p. this quantity is smaller than $\delta' n^{\text{cc}(H)}$ for every $\delta' > 0$, and (ii) $\chi(H) < \text{cc}(H)$. Thus trivially, if we set $\delta' < \delta$, w.h.p.

$$|\tilde{\mathbb{E}}p_{(H,S,\tau)}(x, \mathbf{G}) - (dM)_\tau^{(H,S)}| \leq \delta' n^{\text{cc}(H)} + (dM)_\tau^{(H,S)} n^{\chi(H)} \leq \delta n^{\text{cc}(H)}.$$

■

It therefore suffices to check only the moment constraints for partially labelled forests. In fact, only a subset of these are important.

Definition 45 *Let (H, S, τ) be a partially labelled tree. The pruning of (H, S, τ) is the unique partially labelled subtree in which every leaf belongs to S . The pruning of an unlabelled graph is the empty graph, and the pruning of a forest is defined tree-by-tree. We say that a partially labelled forest is pruned if it is equal to its pruning.*

Lemma 46 *Let (H, S, τ) be a partially labelled forest with maximal degree d , $(\tilde{H}, S, \tau)$ its pruning, and write $\deg$ and $\widetilde{\deg}$ for the vertex degrees in H and $\tilde{H}$ respectively. If X is a symmetric nonnegative integer matrix with row and column sums equal to d , then*

$$\frac{X_\tau^{(H,S)}}{X_\tau^{(\tilde{H},S)}} = \prod_{v \in V(H)} \prod_{q=\widetilde{\deg}(v)}^{\deg(v)-1} (d - q).$$

Proof Let's reason combinatorially. Any such matrix X can be thought of as the adjacency matrix for a d -regular multigraph with self-loops; let's fix X and call this graph Γ . Since Γ 's vertex set is $[k]$, we will think of it as a fully labelled graph. By multiplicativity on disjoint unions, we may freely assume that (H, S, τ) is a tree. Let's choose a root $r \in V(\tilde{H}) \subset V(H)$; having done so, $E(H)$ is in bijection with $V(H) \setminus r$ (and similarly for $E(\tilde{H})$ and $V(\tilde{H})$). Let's write $p(v)$ for the unique parent of every vertex. We can thus write

$$\begin{aligned} X_\tau^{(H,S)} &= \sum_{\tilde{\tau}: \tilde{\tau}|_S = \tau} \prod_{j \in [k]} X_{\tilde{\tau}(r),j} (X_{\tilde{\tau}(r),j} - 1) \cdots (X_{\tilde{\tau}(r),j} - \deg_j(r) + 1) \\ &\quad \times \prod_{v \in V(H) \setminus r} \frac{\prod_{j \in [k]} X_{\tilde{\tau}(v),j} \cdots (X_{\tilde{\tau}(v),j} - \deg_j(v) + 1)}{X_{\tilde{\tau}(v),\tilde{\tau}(p(v))}}. \end{aligned}$$

Thinking of each $\tilde{\tau}$ as a map $V(H) \rightarrow V(\Gamma)$, the summand above gives the number of ways to map $\eta : E(H) \rightarrow E(\Gamma)$ with the following constraints: (1) each edge (u, v) must be mapped to one of the $X_{\tau(u),\tau(v)}$ edges between $\tau(u)$ and $\tau(v)$, and (2) no two edges in $E(H)$ with the same endpoint may be mapped to the same edge in Γ . We'll call the pair $(\tilde{\tau}, \eta)$ a *locally injective occurrence* of (H, S, τ) in the fully labelled graph Γ . Thus the expression $X_\tau^{(H,S)}$ gives the number of such occurrences.

The same argument applies to the pruning $(\tilde{H}, S, \tau)$. Now, the graph (H, S, τ) consists pruning $(\tilde{H}, S, \tau)$, plus some trees hanging off the edges. For each locally injective occurrence of $(\tilde{H}, S, \tau)$, there are

$$\prod_{v \in V(H)} \prod_{q=\widetilde{\deg}(v)}^{\deg(v)-1} (d - q)$$

ways to extend it to a locally injective occurrence of (H, S, τ) , since Γ is d -regular. $\blacksquare$

Lemma 47 *Let $\mathbf{G} \sim \mathbb{Q}$, (H, S, τ) be a partially labelled forest, and $(\tilde{H}, S, \tau)$ its pruning. Then w.h.p.*

$$\left\| p_{(H, S, \tau)}(x, \mathbf{G}) - n^{\text{cc}(H) - \text{cc}(\tilde{H})} \frac{(dM)_\tau^{(H, S)}}{(dM)_\tau^{(\tilde{H}, S)}} p_{(\tilde{H}, S, \tau)}(x, \mathbf{G}) \right\|_1 = o(n^{\text{cc}(H)}),$$

where $\|\cdot\|_1$ is the coefficient-wise L^1 norm.

Proof Let (H, S, τ) be a partially labelled forest, $(\tilde{H}, S, \tau)$ its pruning. If $\tilde{\phi} : V(\tilde{H}) \hookrightarrow V(\mathbf{G})$ is an occurrence of $(\tilde{H}, S, \tau)$, then we call $\phi : V(H) \hookrightarrow V(\mathbf{G})$ an *extension* of $\tilde{\phi}$ if it is an occurrence of (H, S, τ) and agrees with $\tilde{\phi}$ on $V(\tilde{H})$. Let's write $\tilde{\Phi}$ for the set of occurrences of $(\tilde{H}, S, \tau)$ in $(\mathbf{G}, \sigma)$, and for each $\tilde{\phi} \in \tilde{\Phi}$, write $\Xi(\tilde{\phi})$ for its set of extensions. Thus, incorporating Lemma 8.20,

$$\begin{aligned} & \left\| p_{(H, S, \tau)}(x, \mathbf{G}) - n^{\text{cc}(H) - \text{cc}(\tilde{H})} \frac{(dM)_\tau^{(H, S)}}{(dM)_\tau^{(\tilde{H}, S)}} p_{(\tilde{H}, S, \tau)}(x, \mathbf{G}) \right\|_1 \\ & \leq \sum_{\tilde{\phi} \in \tilde{\Phi}} \left| |\Xi(\tilde{\phi})| - n^{\text{cc}(H) - \text{cc}(\tilde{H})} \prod_{v \in V(H)} \prod_{q=\widehat{\deg}(v)}^{\deg(v)-1} (d-q) \right|. \end{aligned}$$

By Proposition 8.15, with high probability there are $o(n^{\text{cc}(\tilde{H})})$ occurrences of $\tilde{H}$ whose $|E(H)|$ neighborhoods in $\tilde{\mathbf{G}}$ either intersect or contain a cycle, so we can safely restrict the right hand side above to the remaining ones. Let's fix such an occurrence and enumerate the possible extensions. First, for each connected component $\tilde{J}$ of $\tilde{H}$, and its corresponding component J of H , because $\mathbf{G}$ is d -regular and locally treelike in the neighborhood of $\phi(\tilde{J})$, there are exactly

$$\prod_{v \in V(J)} \prod_{q=\widehat{\deg}(v)}^{\deg(v)-1} (d-q)$$

ways to extend $\tilde{\phi}$ to the remainder of J . Having already chosen how to extend the occurrence on these connected components, call K the union of all connected components in H that have no distinguished vertex. We need to find an injective homomorphism from K into $\mathbf{G}$ that does not collide with $\tilde{\phi}(\tilde{H})$ or the portion of the extension that we have already constructed. Since $|V(H)| = O(1)$, there are

$$n^{\text{cc}(H) - \text{cc}(\tilde{H})} \prod_{v \in V(K)} \prod_{q=0}^{\deg(v)-1} (d-q) + O(n^{\text{cc}(H) - \text{cc}(\tilde{H}) - 1})$$

ways to do this. Since $|\tilde{\Phi}| = O(n^{\text{cc}(\tilde{H})})$,

$$\sum_{\tilde{\phi} \in \tilde{\Phi}} \left| |\Xi(\tilde{\phi})| - n^{\text{cc}(H) - \text{cc}(\tilde{H})} \prod_{v \in V(H)} \prod_{q=\widehat{\deg}(v)}^{\deg(v)-1} (d-q) \right| = O(n^{\text{cc}(\tilde{H})}) \cdot O(n^{\text{cc}(H) - \text{cc}(\tilde{H}) - 1}) = O(n^{\text{cc}(H) - 1})$$

as desired. $\blacksquare$

Lemma 48 *Let $\mathbf{G} \sim \mathbb{Q}$ and $\tilde{\mathbb{E}} = \tilde{\mathbb{E}}(\mathbf{G})$ be a degree- (D_x, D_G) pseudoexpectation, perhaps dependent on $\mathbf{G}$. If $\tilde{\mathbb{E}}$ w.h.p. satisfies the moment constraints for pruned partially labelled forests, w.h.p. it does so for every partially labelled forest.*

Proof This is a direct consequence of Lemma 8.21. Retaining (H, S, τ) and $(\tilde{H}, S, \tau)$ from the proof of that lemma, since the pseudoexpectation of any monomial has absolute value at most one,

$$\tilde{\mathbb{E}}_{p(H,S,\tau)}(x, \mathbf{G}) = \frac{(dM)_\tau^{(H,S)}}{(dM)_\tau^{(\tilde{H},S)}} p_{(\tilde{H},S,\tau)}(x, \mathbf{G}) \pm o(n^\ell) = (dM)_\tau^{(H,S)} \pm \delta n^\ell$$

for every $\delta > 0$. We can take a union bound over all finitely many (H, S, τ) . ■

We are finally ready to describe our own degree $(2, D)$ pseudoexpectation. Our key building block will be the feasible solution $\mathbf{P} \succeq 0$ to the degree- D Symmetric Path Statistics on input $\mathbf{G} \sim \mathbb{Q}$ SDP whose asymptotic almost sure existence is guaranteed us by Theorem 8.7. Recall that this matrix satisfies

1. $\mathbf{P}_{u,u} = 1$ for every $u \in [n]$
2. $\langle \mathbf{P}, J \rangle = 0$
3. $|\langle \mathbf{P}, A_{\mathbf{G}}^{(s)} \rangle - q_s(d\lambda)| \leq \delta n$ for every $s = 1, \dots, D$.

A degree-2 pseudoexpectation $\tilde{\mathbb{E}} : \mathbb{R}[x]_{\leq 2} \rightarrow \mathbb{R}$ may be expressed as a $(1 + nk) \times (1 + nk)$ block matrix

$$\begin{pmatrix} 1 & l^\top \\ l & Q \end{pmatrix} = \begin{pmatrix} 1 & l_1^\top & \cdots & l_n^\top \\ l_1 & Q_{1,1} & \cdots & Q_{1,n} \\ \vdots & \vdots & \ddots & \vdots \\ l_n & Q_{n,1} & \cdots & Q_{n,n} \end{pmatrix}$$

where $(l_u)_i = \tilde{\mathbb{E}}x_{u,i}$ and $(Q_{u,v})_{i,j} = \tilde{\mathbb{E}}x_{u,i}x_{v,j}$. Our construction will set $(l_u)_i = 1/k$ for every u and i , and

$$Q = \frac{1}{k} (J_{nk}/k + \mathbf{P} \otimes (I - J_k/k))$$

Let us first check the hard constraints. For positivity it suffices to observe that

$$Q - ll^\top = \frac{1}{k} \mathbf{P} \otimes (I - J_k/k) \succeq 0,$$

as $\mathbf{P}, (I - J_k/k) \succeq 0$. We also have

$$\tilde{\mathbb{E}}x_{u,i}^2 = (Q_{u,u})_{i,i} = \frac{1}{k} (1/k + (1 - 1/k)\mathbf{P}_{i,i}) = 1/k = \tilde{\mathbb{E}}x_{u,i}$$

since $\mathbf{P}_{u,u} = 1$ for every $u \in [n]$. It remains only to check that $\tilde{\mathbb{E}}(x_{u,1} + \cdots + x_{u,k})p(x) = \tilde{\mathbb{E}}p(x)$ for every $p(x)$ of degree one. For this it is sufficient to verify that

$$\tilde{\mathbb{E}}(x_{u,1} + \cdots + x_{u,k})x_{v,j} = \sum_i (Q_{u,v})_{i,j} = \frac{1}{k} \sum_i (1/k + \mathbf{P}_{u,v}(1 - J_k/k)_{i,j}) = 1/k = \tilde{\mathbb{E}}x_{v,j}.$$

Finally, we need to verify the moment constraints. By Lemmas 8.21-22, if we do so only on the minimal partially labelled forests with at most two distinguished vertices, then w.h.p. the remainder of the moment constraints are satisfied. A minimal partially labelled forest with one distinguished vertex is just a single vertex with a label $i \in [k]$; the associated polynomial in this case is just $x_{1,k} + \dots + x_{n,k}$, and its required pseudoexpectation is $n/k \pm \delta n$, since there are identically n/k vertices with each label in the planted model. Our pseudoexpectation assigns a value of

$$\tilde{\mathbb{E}} \sum_u x_{u,i} = \sum_u (l_u)_i = n/k$$

as desired.

A minimal partially labelled forest on two vertices is either a path of length $s \in [d]$ with endpoints labelled $i, j \in [k]$, or two isolated vertices labelled $i, j \in [k]$. In the former case, the pseudoexpectation is required to read $\frac{1}{k} q(dM)_{i,j} n \pm \delta n$. Recycling some calculations from Section 8.4, our pseudoexpectation on this polynomial reads

$$\begin{aligned} \langle Q_{i,j}, A_{\mathbf{G}}^{(s)} \rangle &= \langle Q_{i,j}, A_{\mathbf{G}}^{(s)} \rangle \pm o(n) \\ &= \frac{1}{k} \left(\langle J_n/k, A_{\mathbf{G}}^{(s)} \rangle + \langle \mathbf{P}, A_{\mathbf{G}}^{(s)} \rangle (I - J_k/k)_{i,j} \right) \\ &= \frac{1}{k} (q_s(d)(J/k)_{i,j} + q_s(d\lambda)(I - J_k/k)_{i,j}) n \pm (\delta/k)n \\ &= \frac{1}{k} q_s(dM)_{i,j} \pm (\delta/k)n. \end{aligned}$$

The last line follows since $dM = dJ_k/k + d\lambda(I - J_k/k)$ is the spectral decomposition of dM .

Finally, we verify the case of two disjoint vertices labelled $i, j \in [k]$. The polynomial here is $\sum_{u \neq v} x_{u,i} x_{v,j}$, and the pseudoexpectation is required to give a value of $(n/k)^2 \pm \delta n^2$. As needed, our pseudoexpectation gives

$$\langle Q_{i,j}, J_n \rangle = \frac{1}{k} (\langle J_n/k, J_n \rangle + \langle \mathbf{P}, J_n \rangle (I - J_k/k)_{i,j}) = (n/k)^2,$$

as $\langle \mathbf{P}, J_n \rangle = 0$.

Appendix E. Low-Degree Analysis of Spiked Models

In this section, we develop machinery for low-degree analysis of general spiked Wigner and Wishart models, culminating in the proofs of Theorems 19 and 20.

E.1. Preliminaries

Definition 49 A random vector $\mathbf{x}$ is ε -local c -subgaussian if for any fixed vector v with $\|v\| \leq \varepsilon$,

$$\mathbb{E} \exp(\langle v, \mathbf{x} \rangle) \leq \exp\left(\frac{c}{2} \|v\|^2\right).$$

It is straightforward to verify the following fact.

Fact 50 Suppose $\mathbf{x}$ is ε -local c -subgaussian. For a (non-random) scalar $\alpha \neq 0$, $\alpha\mathbf{x}$ is $\varepsilon/|\alpha|$ -local $c\alpha^2$ -subgaussian. Also, the sum of n independent copies of $\mathbf{x}$ is ε -local cn -subgaussian.

Proposition 51 *If a random vector $\mathbf{x}$ is ε -local c -subgaussian then it admits the following local Chernoff bound: for any $\|v\| = 1$ and $0 \leq t \leq \varepsilon c$,*

$$\Pr\{\langle v, \mathbf{x} \rangle \geq t\} \leq \exp\left(-\frac{t^2}{2c}\right).$$

Proof Apply the standard Chernoff bound argument: for any $\alpha > 0$,

$$\begin{aligned} \Pr\{\langle v, \mathbf{x} \rangle \geq t\} &= \Pr\{\exp(\alpha \langle v, \mathbf{x} \rangle) \geq \exp(\alpha t)\} \\ &\leq \mathbb{E}[\exp(\alpha \langle v, \mathbf{x} \rangle)] / \exp(\alpha t) \\ &\leq \exp(c\alpha^2/2 - \alpha t) \end{aligned} \quad \text{provided } \alpha \leq \varepsilon.$$

Set $\alpha = t/c$ to complete the proof. ■

Proposition 52 *If a random vector $\mathbf{x} \in \mathbb{R}^k$ is ε -local c -subgaussian then for any $\delta > 0$ and any $0 \leq t \leq \varepsilon c/(1 - \delta)$,*

$$\Pr\{\|\mathbf{x}\| \geq t\} \leq C(\delta, k) \exp\left(-\frac{1}{2c}(1 - \delta)^2 t^2\right)$$

where $C(\delta, k)$ is a constant depending only on δ and k .

Proof Let $\mathcal{N} \subseteq \mathbb{R}^k$ be a δ -net of the unit sphere in $\mathbb{R}^k$, in the sense that for any $v \in \mathbb{R}^k$,

$$\max_{u \in \mathcal{N}} \langle u, v \rangle \geq (1 - \delta)\|v\|$$

where $\|u\| = 1$ for all $u \in \mathcal{N}$. Let $C(\delta, k) = |\mathcal{N}|$. Using a union bound and the local Chernoff bound (Proposition 51), for all $0 \leq t \leq \varepsilon c/(1 - \delta)$,

$$\Pr\{\|\mathbf{x}\| \geq t\} \leq \Pr\left\{\max_{u \in \mathcal{N}} \langle u, \mathbf{x} \rangle \geq (1 - \delta)t\right\} \leq C(\delta, k) \exp\left(-\frac{1}{2c}(1 - \delta)^2 t^2\right).$$
■

Proposition 53 *Let $\delta > 0$. If a random vector $\mathbf{x} \in \mathbb{R}^k$ is ε -local c -subgaussian with $c < (1 - \delta)^2/2$ then*

$$\mathbb{E}[\mathbf{1}[\|\mathbf{x}\| \leq \varepsilon c/(1 - \delta)] \exp(\|\mathbf{x}\|^2)] \leq 1 + \frac{C(\delta, k)}{(1 - \delta)^2/(2c) - 1}$$

where $C(\delta, k)$ is a constant depending only on δ and k .

Proof Let $\Delta = \varepsilon c/(1 - \delta)$, and integrate the tail bound from Proposition 52:

$$\begin{aligned} \mathbb{E}[\mathbf{1}[\|\mathbf{x}\| \leq \Delta] \exp(\|\mathbf{x}\|^2)] &= \int_0^\infty \Pr\{\mathbf{1}[\|\mathbf{x}\| \leq \Delta] \exp(\|\mathbf{x}\|^2) \geq t\} dt \\ &\leq 1 + \int_1^{\exp(\Delta^2)} \Pr\{\exp(\|\mathbf{x}\|^2) \geq t\} dt \\ &= 1 + \int_1^{\exp(\Delta^2)} \Pr\{\|\mathbf{x}\| \geq \sqrt{\log t}\} dt \\ &\leq 1 + C(\delta, k) \int_1^\infty \exp\left(-\frac{1}{2c}(1 - \delta)^2 \log t\right) dt \\ &= 1 + \frac{C(\delta, k)}{(1 - \delta)^2/(2c) - 1}. \end{aligned}$$

■

E.2. The Wigner Model

Proof [Proof of Theorem 19] We start with a formula for $\|L^{\leq D}\|^2$ from Kunisky et al. (2019) (adapted slightly for the case of symmetric Gaussian noise):

$$\|L^{\leq D}\|^2 = \mathbb{E}_{\mathbf{X}, \mathbf{X}'} \exp^{\leq D} \left(\frac{\lambda^2 n}{2} \langle \mathbf{X}, \mathbf{X}' \rangle \right) \quad (55)$$

where $\mathbf{X}'$ is an independent copy of $\mathbf{X}$ and $\exp^{\leq D}(x) = \sum_{d=0}^D \frac{x^d}{d!}$ denotes the Taylor series truncation of $\exp$. Write $\|L^{\leq D}\|^2 = L_1 + L_2$ where L_1 is the *small deviations* term

$$L_1 := \mathbb{E}_{\mathbf{X}, \mathbf{X}'} \mathbf{1} [\langle \mathbf{X}, \mathbf{X}' \rangle \leq \Delta] \exp^{\leq D} \left(\frac{\lambda^2 n}{2} \langle \mathbf{X}, \mathbf{X}' \rangle \right),$$

and L_2 is the *large deviations* term

$$L_2 := \mathbb{E}_{\mathbf{X}, \mathbf{X}'} \mathbf{1} [\langle \mathbf{X}, \mathbf{X}' \rangle > \Delta] \exp^{\leq D} \left(\frac{\lambda^2 n}{2} \langle \mathbf{X}, \mathbf{X}' \rangle \right),$$

where $\Delta > 0$ is a small constant to be chosen later. Lemmas 54 and 55, proved in the following two subsections, show that for some choice of Δ , L_1 and L_2 are both $O(1)$, completing the proof. ■

E.2.1. SMALL DEVIATIONS

Lemma 54 *In the setting of Theorem 19, if $|\lambda| < 1$ then there exists $\Delta > 0$ such that $L_1 = O(1)$ for any D .*

Proof Note that

$$\langle \mathbf{X}, \mathbf{X}' \rangle = \frac{1}{n^2} \langle \mathbf{U} \mathbf{U}^\top, \mathbf{U}' (\mathbf{U}')^\top \rangle = \frac{1}{n^2} \|\mathbf{R}\|_F^2 \quad (56)$$

where $\mathbf{R} = \mathbf{U}^\top \mathbf{U}'$. In particular, $\langle \mathbf{X}, \mathbf{X}' \rangle \geq 0$. Since $\exp^{\leq D}(x) \leq \exp(x)$ for all $x \geq 0$,

$$L_1 \leq \mathbb{E} \mathbf{1} [\|\mathbf{R}\|_F^2 \leq \Delta n^2] \exp \left(\frac{\lambda^2}{2n} \|\mathbf{R}\|_F^2 \right). \quad (57)$$

We have $\mathbf{R} = \sum_{i=1}^n \mathbf{R}_i$ where the $\mathbf{R}_i$ are independent $k \times k$ matrices, each distributed as $\pi(\pi')^\top$. Since π has bounded support, the moment-generating function $M(T) = \mathbb{E} \exp(\langle T, \mathbf{R}_i \rangle)$ exists in a neighborhood of $T = 0$ (in fact, it exists everywhere) and thus, by the defining property of the MGF, has gradient $\nabla M(0) = \mathbb{E}[\mathbf{R}_i] = 0$ and Hessian $(\text{Hess } M)(0) = \text{Cov}(\mathbf{R}_i) = \text{Cov}(\pi)^{\otimes 2} \preceq I_{k^2}$. Thus for any $\eta > 0$ there exists $\varepsilon > 0$ such that

$$M(T) \leq \exp \left(\frac{1}{2} (1 + \eta) \|T\|_F^2 \right) \quad \text{for all } \|T\|_F \leq \varepsilon.$$

In other words, $\mathbf{R}_i$ is ε -local $(1 + \eta)$ -subgaussian. From Fact 50, this implies $\mathbf{R}$ is ε -local $(1 + \eta)n$ -subgaussian, and $\lambda \mathbf{R} / \sqrt{2n}$ is $\varepsilon \sqrt{2n} / \lambda$ -local $(1 + \eta) \lambda^2 / 2$ -subgaussian. Since $\lambda^2 < 1$, we can

choose $\delta > 0$ and $\eta > 0$ such that $(1 + \eta)\lambda^2 < (1 - \delta)^2$. Letting $\Delta = (\varepsilon(1 + \eta)/(1 - \delta))^2$ and using Proposition 53,

$$L_1 \leq \mathbb{E} \mathbf{1} \left[\|\lambda \mathbf{R}/\sqrt{2n}\|_{\text{F}} \leq \lambda\sqrt{\Delta n/2} \right] \exp \left(\|\lambda \mathbf{R}/\sqrt{2n}\|_{\text{F}}^2 \right) = O(1).$$

■

E.2.2. LARGE DEVIATIONS

Lemma 55 *In the setting of Theorem 19, for any constants $\lambda \in \mathbb{R}$ and $\Delta > 0$, and for any $D = o(n/\log n)$, we have $L_2 = o(1)$.*

Proof Recall from above that $\mathbf{R}$ is ε -local $(1 + \eta)n$ -subgaussian. By Proposition 52 (taking t to be n times a small constant),

$$\Pr\{\langle \mathbf{X}, \mathbf{X}' \rangle > \Delta\} \leq \Pr\{\|\mathbf{R}\|_{\text{F}} > \sqrt{\Delta n}\} = \exp(-\Omega(n)). \quad (58)$$

The boundedness of π guarantees $|\langle \mathbf{X}, \mathbf{X}' \rangle| \leq n^C$ for some constant $C > 0$, and so

$$L_2 \leq \exp(-\Omega(n)) \sum_{d=0}^D \left(\frac{\lambda^2 n}{2} n^C \right)^d \leq \exp(-\Omega(n)) (D + 1) n^{O(D)},$$

which is $o(1)$ provided $D = o(n/\log n)$. ■

E.3. The Wishart Model

Proof [Proof of Theorem 20] In Appendix F we derive a formula for $\|L^{\leq D}\|^2$ in the general spiked Wishart model. The version we will need here is summarized in Proposition 65. The formula takes the form

$$\|L^{\leq D}\|^2 = \mathbb{E}_{\mathbf{X}, \mathbf{X}'} \sum_{d=0}^D r_d(\beta \mathbf{X}, \beta \mathbf{X}') \quad (59)$$

for some polynomials r_d . As in the Wigner case, we write $\|L^{\leq D}\|^2 = L_1 + L_2$ where

$$L_1 := \mathbb{E}_{\mathbf{X}, \mathbf{X}'} \mathbf{1} [\langle \mathbf{X}, \mathbf{X}' \rangle \leq \Delta] \sum_{d=0}^D r_d(\beta \mathbf{X}, \beta \mathbf{X}')$$

and

$$L_2 := \mathbb{E}_{\mathbf{X}, \mathbf{X}'} \mathbf{1} [\langle \mathbf{X}, \mathbf{X}' \rangle > \Delta] \sum_{d=0}^D r_d(\beta \mathbf{X}, \beta \mathbf{X}')$$

for a small constant $\Delta > 0$ to be chosen later. Lemmas 58 and 59, proved in the following two subsections, show that for some choice of Δ , L_1 and L_2 are both $O(1)$, completing the proof. ■

E.3.1. SMALL DEVIATIONS

Before bounding the small deviations term in Lemma 58, we state two deterministic facts that will be useful in the proof.

Proposition 56 (Sylvester Determinant Identity) *For $n \times m$ matrices A and B ,*

$$\det(I_n - AA^\top BB^\top) = \det(I_m - A^\top BB^\top A).$$

Proof If A, B are square and A is nonsingular, this can be shown by taking determinants on both sides of the equation $(I - AA^\top BB^\top)A = A(I - A^\top BB^\top A)$. For the general case, pad A, B with zeros to make them square, and consider a sequence of nonsingular matrices converging to A . ■

Lemma 57 *For any $\eta > 0$ there exists $\varepsilon > 0$ such that for all $0 \leq t \leq \varepsilon$, we have $(1 - t)^{-1} \leq \exp((1 + \eta)t)$.*

Proof Letting $f(t) = (1 - t)^{-1}$ and $g(t) = \exp((1 + \eta)t)$, we have $f(0) = g(0) = 1$ and $f'(0) = 1 < 1 + \eta = g'(0)$. ■

We are now ready to bound the small deviations term.

Lemma 58 *In the setting of Theorem 20, if $\beta^2 < \gamma$ then there exists $\Delta > 0$ such that $L_1 = O(1)$ for any D .*

Proof Since $\mathbf{X} \succeq 0$, we have from Proposition 65(c,d) that $r_d(\beta\mathbf{X}, \beta\mathbf{X}') \geq 0$. Analogous to (56), we have either $\langle \mathbf{X}, \mathbf{X}' \rangle = 0$ or $\langle \mathbf{X}, \mathbf{X}' \rangle = \frac{1}{n^2} \|\mathbf{R}\|_F^2$. Recall from Definition 17 that, when drawing from $\mathbb{P}$, we first draw $\widetilde{\mathbf{X}} \sim \mathcal{X}$, and then threshold to form $\mathbf{X}$ as:

$$\mathbf{X} = \begin{cases} \widetilde{\mathbf{X}} & \text{if } \beta\widetilde{\mathbf{X}} \succ -I_n, \\ 0 & \text{else.} \end{cases} \quad (60)$$

We can upper bound L_1 by dropping the low-degree truncation and using Proposition 65(b):

$$\begin{aligned} L_1 &\leq \mathbb{E}_{\mathbf{X}, \mathbf{X}'} \mathbf{1}[\langle \mathbf{X}, \mathbf{X}' \rangle \leq \Delta] \sum_{d=0}^{\infty} r_d(\beta\mathbf{X}, \beta\mathbf{X}') \\ &= \mathbb{E}_{\mathbf{X}, \mathbf{X}'} \mathbf{1}[\langle \mathbf{X}, \mathbf{X}' \rangle \leq \Delta] \det(I_n - \beta^2 \mathbf{X} \mathbf{X}')^{-N/2} \\ &\leq 1 + \mathbb{E} \mathbf{1}[\|\mathbf{R}\|_F^2 \leq \Delta n^2] \det(I_n - \beta^2 \widetilde{\mathbf{X}} \widetilde{\mathbf{X}}')^{-N/2} \end{aligned}$$

where the +1 in the last line covers the second case of (60). Let $\{\lambda_i\}$ be the eigenvalues of $\mathbf{R}\mathbf{R}^\top$ (which are nonnegative). Let $\eta > 0$ (to be chosen later), take ε according to Lemma 57, and choose $\Delta \leq \varepsilon/\beta^2$. Note that $\|\mathbf{R}\|_F^2 = \sum_i \lambda_i$. Provided $\|\mathbf{R}\|_F^2 \leq \Delta n^2$, we have $\lambda_i \leq \Delta n^2$ for all i , i.e.,

$\beta^2 n^{-2} \lambda_i \leq \varepsilon$, and so

$$\begin{aligned}
 \det(I_n - \beta^2 \widetilde{\mathbf{X}} \widetilde{\mathbf{X}}')^{-N/2} &= \det(I_n - \beta^2 n^{-2} \mathbf{U} \mathbf{U}^\top \mathbf{U}' \mathbf{U}'^\top)^{-N/2} \\
 &= \det(I_n - \beta^2 n^{-2} \mathbf{U}^\top \mathbf{U}' \mathbf{U}'^\top \mathbf{U})^{-N/2} && \text{(using Proposition 56)} \\
 &= \det(I_n - \beta^2 n^{-2} \mathbf{R} \mathbf{R}^\top)^{-N/2} \\
 &= \prod_i (1 - \beta^2 n^{-2} \lambda_i)^{-N/2} \\
 &\leq \prod_i \exp((1 + \eta) \beta^2 n^{-2} \lambda_i)^{N/2} && \text{(using Lemma 57)} \\
 &= \exp\left((1 + \eta) \frac{\beta^2 N}{2n^2} \sum_i \lambda_i\right) \\
 &= \exp\left((1 + \eta) \frac{\beta^2 N}{2n^2} \|\mathbf{R}\|_{\text{F}}^2\right).
 \end{aligned}$$

We now have

$$L_1 \leq 1 + \mathbb{E} \mathbf{1} [\|\mathbf{R}\|_{\text{F}}^2 \leq \Delta n^2] \exp\left((1 + \eta) \frac{\beta^2 N}{2n^2} \|\mathbf{R}\|_{\text{F}}^2\right).$$

Comparing this to (57), we see that we have reduced to the case of Wigner small deviations. In place of λ^2 , we have $(1 + \eta) \frac{\beta^2 N}{n} \rightarrow (1 + \eta) \frac{\beta^2}{\gamma}$. Thus, provided $\beta^2 < \gamma$, we can choose $\eta > 0$ and $\Delta > 0$ small enough so that $L_1 = O(1)$. $\blacksquare$

E.3.2. LARGE DEVIATIONS

Lemma 59 *In the setting of Theorem 20, for any constants $\beta > -1$, $\gamma > 0$, and $\Delta > 0$, and for any $D = o(n/\log n)$, we have $L_2 = o(1)$.*

Proof Consider the case $\beta > 0$ so that $\beta \mathbf{X} \succeq 0$; the case $\beta < 0$ is handled similarly. We will use the formula for r_d given in Proposition 65(c):

$$r_d(\beta \mathbf{X}, \beta \mathbf{X}') = \sum_{\substack{d_1, \dots, d_N \in 2\mathbb{N} \\ \sum_{i=1}^N d_i = d}} \prod_{i=1}^N \frac{1}{d_i!} \mathbb{E}_{\substack{\mathbf{x} \sim \mathcal{N}(0, \beta \mathbf{X}) \\ \mathbf{x}' \sim \mathcal{N}(0, \beta \mathbf{X}')}} \langle \mathbf{x}, \mathbf{x}' \rangle^{d_i}.$$

Let $\{\lambda_i, \mathbf{v}_i\}$ be an eigendecomposition of $\mathbf{X}$, and let $\mathbf{X}_i = \lambda_i \mathbf{v}_i \mathbf{v}_i^\top$ so that $\mathbf{X} = \mathbf{X}_1 + \dots + \mathbf{X}_k$. Let $\mathbf{x}_i \sim \mathcal{N}(0, \beta \mathbf{X}_i)$ independently so that $\mathbf{x} := \sum_{i=1}^k \mathbf{x}_i \sim \mathcal{N}(0, \beta \mathbf{X})$. Similarly define $\mathbf{X}'_i, \mathbf{x}'_i, \mathbf{x}'$. For fixed $\mathbf{X}, \mathbf{X}', \{\mathbf{X}_i\}, \{\mathbf{X}'_i\}$,

$$\mathbb{E}[\langle \mathbf{x}, \mathbf{x}' \rangle^d] \leq \mathbb{E} \|\mathbf{x}\|^d \mathbb{E} \|\mathbf{x}'\|^d = \mathbb{E} \|\mathbf{x}\|^d \cdot \mathbb{E} \|\mathbf{x}'\|^d.$$

The boundedness of $\boldsymbol{\pi}$ guarantees $\sum_i \lambda_i^2 = \|\mathbf{X}\|_{\text{F}}^2 \leq n^C$ for some constant $C > 0$. Recall that $\mathbf{x}_i \sim \mathcal{N}(0, \beta \lambda_i \mathbf{v}_i \mathbf{v}_i^\top)$, so for $\mathbf{g}$ a standard Gaussian vector, $\mathbf{x}_i$ has the same law as $\sqrt{\beta \lambda_i} \cdot \langle \mathbf{v}_i, \mathbf{g} \rangle \mathbf{v}_i$. Since $\mathbf{v}_i$ is a unit vector, $\|\mathbf{x}_i\|$ then has the same law as $\sqrt{\beta \lambda_i} \cdot |\langle \mathbf{v}_i, \mathbf{g} \rangle|$, which in turn has the same law as $\sqrt{\beta \lambda_i} \cdot |\mathbf{g}_1|$. So, for d even, $\mathbb{E} \|\mathbf{x}_i\|^d = (\beta \lambda_i)^{d/2} (d-1)!! \leq \beta^{d/2} n^{Cd/4} d^d$. This means

$$\mathbb{E} \|\mathbf{x}\|^d \leq \mathbb{E} \left(\sum_{i=1}^k \|\mathbf{x}_i\| \right)^d \leq \mathbb{E} \left(k \max_i \|\mathbf{x}_i\| \right)^d \leq k^d \mathbb{E} \sum_{i=1}^k \|\mathbf{x}_i\|^d \leq k^{d+1} \beta^{d/2} n^{Cd/4} d^d.$$

Since $\mathbb{E}\|\mathbf{x}\|^d = 1$ when $d = 0$, we can rewrite this as $\mathbb{E}\|\mathbf{x}\|^d \leq k^{2d}\beta^{d/2}n^{Cd/4}d^d \leq (dn)^{O(d)}$. Thus for $d_1, \dots, d_N \in 2\mathbb{N}$ with $\sum_i d_i = d \leq D$,

$$\prod_{i=1}^N \mathbb{E}[\langle \mathbf{x}, \mathbf{x}' \rangle^{d_i}] \leq \prod_{i=1}^N (d_i n)^{O(d_i)} \leq (Dn)^{O(D)}.$$

Now we have

$$r_d(\mathbf{X}, \mathbf{X}') \leq N^D (Dn)^{O(D)}$$

and so

$$L_2 \leq \Pr\{\langle \mathbf{X}, \mathbf{X}' \rangle > \Delta\} \sum_{d=0}^D N^D (Dn)^{O(D)}.$$

Similarly to (58) we have $\Pr\{\langle \mathbf{X}, \mathbf{X}' \rangle > \Delta\} \leq \exp(-\Omega(n))$ and so $L_2 = o(1)$ provided $D = o(n/\log n)$. $\blacksquare$

Appendix F. Low-Degree Analysis of General Wishart Models

In this section we derive the formula (59) for $\|L^{\leq D}\|^2$ in the general spiked Wishart model.

F.1. Hermite Polynomial Facts

We first define and give the key facts that we will use about the Hermite polynomials, the orthogonal polynomials with respect to the standard Gaussian measure.

Definition 60 (Hermite polynomials) *The univariate Hermite polynomials are the sequence of polynomials $h_k(y) \in \mathbb{R}[y]$ for $k \in \mathbb{N}$, defined by the recursion*

$$h_0(y) = 1, \tag{61}$$

$$h_{k+1}(y) = yh_k(y) - h'_k(y). \tag{62}$$

The n -variate Hermite polynomials are the polynomials $H_\alpha(y) \in \mathbb{R}[y_1, \dots, y_n]$ indexed by $\alpha \in \mathbb{N}^n$ and $H_\alpha(y) = \prod_{i=1}^n h_{\alpha_i}(y_i)$. Finally, the normalized n -variate Hermite polynomials are $\widehat{H}_\alpha(y) = (\alpha!)^{-1/2} H_\alpha(y)$, where we abbreviate $\alpha! = \prod_{i=1}^n \alpha_i!$.

The main and defining property of the Hermite polynomials is their orthogonality, which we record below.

Proposition 61 (Orthogonality under Gaussian measure) *For any $\alpha, \beta \in \mathbb{N}^n$,*

$$\mathbb{E}_{\mathbf{y} \sim \mathcal{N}(0, I)} [\widehat{H}_\alpha(\mathbf{y}) \widehat{H}_\beta(\mathbf{y})] = \delta_{\alpha\beta}. \tag{63}$$

Beyond this, the key additional tool for our analysis is a generalization of the following fact from the ‘‘umbral calculus’’ of Hermite polynomials (a proof will be subsumed in our more general result below).

Proposition 62 (Mismatched Variance Formula) *Let $x > -1$. Then,*

$$\mathbb{E}_{\mathbf{y} \sim \mathcal{N}(0, 1+x)} h_k(\mathbf{y}) = \begin{cases} (k-1)!! \cdot x^{k/2} & k \text{ even} \\ 0 & k \text{ odd} \end{cases} \quad (64)$$

Note that the formula on the right-hand side is that for the moments of a Gaussian random variable with variance x , but we extend it to apply even for negative x , which is the “umbral” case of the result, admitting an interpretation in terms of a fictitious Gaussian of negative variance—even if $x \in (-1, 0)$, the right-hand side may be viewed formally as the value of “ $\mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, x)} \mathbf{g}^k$.” A thorough exposition of such analogies arising in combinatorics and the theory of orthogonal polynomials is given in [Roman \(2005\)](#).

In fact, the same holds even for multivariate Gaussians. The correct result in this case is given by imitating the formula for the moments of a multivariate Gaussian, via Wick’s (or Isserlis’) formula. While Proposition 62 is well-known in the literature on Hermite polynomials and the umbral calculus, we are not aware of previous appearances of the formula below.

Proposition 63 (Multivariate Mismatched Variance Formula) *Let $X \in \mathbb{R}_{\text{sym}}^{n \times n}$ with $X \succ -I_n$. For $\alpha \in \mathbb{N}^n$ viewed as a multiset of elements of $[n]$ (specifically, the multiset containing α_i copies of i , for each $i \in [n]$), let $\mathcal{P}(\alpha)$ be the set of pairings of elements of α , and for each $P \in \mathcal{P}(\alpha)$ write X^P denote the product of the entries of X located at paired indices from P . Then,*

$$\mathbb{E}_{\mathbf{y} \sim \mathcal{N}(0, I+X)} H_\alpha(\mathbf{y}) = \sum_{P \in \mathcal{P}(\alpha)} X^P. \quad (65)$$

Note that if $X \succeq 0$, then the right-hand side equals $\mathbb{E}_{\mathbf{x} \sim \mathcal{N}(0, X)} x^\alpha$ by Wick’s formula, but we again have an umbral extension to non-PSD matrices X .

Proof Define

$$\ell_\alpha := \mathbb{E}_{\mathbf{y} \sim \mathcal{N}(0, I+X)} H_\alpha(\mathbf{y}). \quad (66)$$

Let $e_i \in \mathbb{N}^n$ have i th coordinate equal to 1 and all other coordinates equal to zero, and write $0 \in \mathbb{N}^n$ for the vector with all coordinates equal to zero. Clearly $\ell_0 = 1$ and $\ell_{e_i} = 0$ for any $i \in [n]$. We

then proceed by induction and use Gaussian integration by parts:

$$\begin{aligned}
\ell_{\alpha+e_i} &= \mathbb{E}_{\mathbf{y} \sim \mathcal{N}(0, I+X)} h_{\alpha_i+1}(\mathbf{y}_i) \prod_{j \in [n] \setminus \{i\}} h_{\alpha_j}(\mathbf{y}_j) \\
&= \mathbb{E}_{\mathbf{y} \sim \mathcal{N}(0, I+X)} (\mathbf{y}_i h_{\alpha_i}(\mathbf{y}_i) - h'_{\alpha_i}(\mathbf{y}_i)) \prod_{j \in [n] \setminus \{i\}} h_{\alpha_j}(\mathbf{y}_j) && \text{(Definition 60)} \\
&= \mathbb{E}_{\mathbf{y} \sim \mathcal{N}(0, I+X)} \left[\sum_{k=1}^n (I+X)_{ik} \prod_{j \in [n]} h_{\alpha_j}^{(\delta_{jk})}(\mathbf{y}_j) - \prod_{j \in [n]} h_{\alpha_j}^{(\delta_{ij})}(\mathbf{y}_j) \right] && \text{(integration by parts)} \\
&= \sum_{\substack{k \in [n] \\ \alpha_k > 0}} X_{ik} \mathbb{E}_{\mathbf{y} \sim \mathcal{N}(0, I+X)} \prod_{j \in [n]} h_{\alpha_j}^{(\delta_{jk})}(\mathbf{y}_j) \\
&= \sum_{\substack{k \in [n] \\ \alpha_k > 0}} \alpha_k X_{ik} \mathbb{E}_{\mathbf{y} \sim \mathcal{N}(0, I+X)} \prod_{j \in [n]} h_{\alpha_j - \delta_{jk}}(\mathbf{y}_j) \\
&= \sum_{\substack{k \in [n] \\ \alpha_k > 0}} \alpha_k X_{ik} \ell_{\alpha - e_k} && \text{(inductive hypothesis)}
\end{aligned}$$

so ℓ_α satisfy the same recursion and initial condition as the sum-of-products formula on the right-hand side of (65). $\blacksquare$

F.2. Components of the LDLR

Let $\mathbb{Q}, \mathbb{P}$ be as in the general spiked Wishart model (Definition 17), and let L be the associated likelihood ratio. Throughout this section, we assume without loss of generality that (i) $\beta = 1$ (since β can be absorbed into X), and (ii) the prior $\mathcal{X}$ is supported on X for which $X \succ -I_n$. For $\alpha \in (\mathbb{N}^n)^N$, we denote

$$m_\alpha := \mathbb{E}_{\mathbf{Y} \sim \mathbb{Q}} H_\alpha(\mathbf{Y}) L(\mathbf{Y}), \quad (67)$$

$$\hat{m}_\alpha := \mathbb{E}_{\mathbf{Y} \sim \mathbb{Q}} \hat{H}_\alpha(\mathbf{Y}) L(\mathbf{Y}) = \frac{1}{\sqrt{\alpha!}} m_\alpha. \quad (68)$$

We may compute these numbers as follows. For any $\alpha \in (\mathbb{N}^n)^N$, we have, passing to an expectation under $\mathbb{P}$ rather than $\mathbb{Q}$ and then using Proposition 63,

$$\begin{aligned}
m_\alpha &= \mathbb{E}_{\mathbf{Y} \sim \mathbb{P}} H_\alpha(\mathbf{Y}) \\
&= \mathbb{E}_{\mathbf{X} \sim \mathcal{X}} \prod_{i=1}^N \mathbb{E}_{\mathbf{y} \sim \mathcal{N}(0, I+X)} H_{\alpha_i}(\mathbf{y}) \\
&= \mathbb{E}_{\mathbf{X} \sim \mathcal{X}} \prod_{i=1}^N \left(\sum_{P \in \mathcal{P}(\alpha_i)} X^P \right). && (69)
\end{aligned}$$

F.3. Taylor Expansion of the LDLR

Lemma 64 Suppose $\mathcal{X}$ is as in Definition 17. Denote by $\mathcal{T}^{\leq D}$ the operation of truncating the Taylor series of a function that is real-analytic in a neighborhood of zero to degree- D polynomials. For the sake of clarity, this will only ever apply to the variable named t . Then,

$$\|L^{\leq D}\|^2 = \mathbb{E}_{\mathbf{X}, \mathbf{X}' \sim \mathcal{X}} \mathcal{T}^{\leq D} \left[\det(I_n - t^2 \mathbf{X} \mathbf{X}')^{-N/2} \right] (1), \quad (70)$$

by which we mean evaluation of the function of t on the RHS at $t = 1$.

Proof We have

$$\begin{aligned} \|L^{\leq D}\|^2 &= \sum_{|\alpha| \leq D} \hat{m}_\alpha^2 \\ &= \mathbb{E}_{\mathbf{X}, \mathbf{X}' \sim \mathcal{X}} \sum_{|\alpha| \leq D} \prod_{i=1}^N \frac{1}{\alpha_i!} \left(\sum_{P \in \mathcal{P}(\alpha_i)} \mathbf{X}^P \right) \left(\sum_{P \in \mathcal{P}(\alpha_i)} (\mathbf{X}')^P \right) \\ &= \mathbb{E}_{\mathbf{X}, \mathbf{X}' \sim \mathcal{X}} \sum_{d=0}^D \sum_{|\alpha|=d} \prod_{i=1}^N \frac{1}{\alpha_i!} \left(\sum_{P \in \mathcal{P}(\alpha_i)} \mathbf{X}^P \right) \left(\sum_{P \in \mathcal{P}(\alpha_i)} (\mathbf{X}')^P \right). \end{aligned} \quad (71)$$

On the other hand, we may expand the right-hand side of (70) by repeatedly differentiating with respect to t to extract Taylor coefficients. In doing so we will repeatedly apply the chain rule, and since $\frac{d}{dt} \det(I - tA) = \text{Tr}(A)$ each derivative will be a rational function in (t, X, X') . Therefore, the Taylor coefficients are some rational functions $r_d(X, X')$ (depending on N) for $d \in \mathbb{N}$, and for these coefficients

$$\mathbb{E}_{\mathbf{X}, \mathbf{X}' \sim \mathcal{X}} \mathcal{T}^{\leq D} \left[\det(I_n - t^2 \mathbf{X} \mathbf{X}')^{-N/2} \right] (1) = \mathbb{E}_{\mathbf{X}, \mathbf{X}' \sim \mathcal{X}} \sum_{d=0}^D r_d(\mathbf{X}, \mathbf{X}'). \quad (72)$$

We will show that in fact termwise equality holds, *inside the expectations*, namely

$$r_d(X, X') = \sum_{|\alpha|=d} \prod_{i=1}^N \frac{1}{\alpha_i!} \left(\sum_{P \in \mathcal{P}(\alpha_i)} X^P \right) \left(\sum_{P \in \mathcal{P}(\alpha_i)} (X')^P \right) =: r'_d(X, X') \quad (73)$$

for each $d \in \mathbb{N}$ and for all (deterministic) $X, X' \in \mathbb{R}_{\text{sym}}^{n \times n}$. Since either side of (73) is a rational function of (X, X') , it suffices to show that this is true on a set of matrices of positive measure. We will show that it holds for all $X, X' \succeq 0$.

In this case, we have the convenient Gaussian interpretation of the expression for $r'_d(X, X')$ from Wick's formula:

$$\begin{aligned} r'_d(X, X') &= \sum_{|\alpha|=d} \prod_{i=1}^N \frac{1}{\alpha_i!} \left(\mathbb{E}_{\mathbf{x} \sim \mathcal{N}(0, X)} \mathbf{x}^{\alpha_i} \right) \left(\mathbb{E}_{\mathbf{x} \sim \mathcal{N}(0, X')} \mathbf{x}^{\alpha_i} \right) \\ &= \mathbb{E}_{\substack{\mathbf{x}_1, \dots, \mathbf{x}_N \sim \mathcal{N}(0, X) \\ \mathbf{x}'_1, \dots, \mathbf{x}'_N \sim \mathcal{N}(0, X')}} \sum_{|\alpha| \text{ even}} \prod_{i=1}^N \frac{1}{\alpha_i!} (\mathbf{x}_i)^{\alpha_i} (\mathbf{x}'_i)^{\alpha_i} \end{aligned} \quad (74)$$

and grouping by the values of $|\alpha_i|$,

$$\begin{aligned}
&= \mathbb{E}_{\substack{\mathbf{x}_1, \dots, \mathbf{x}_N \sim \mathcal{N}(0, X) \\ \mathbf{x}'_1, \dots, \mathbf{x}'_N \sim \mathcal{N}(0, X')}} \sum_{\substack{d_1, \dots, d_N \in 2\mathbb{N} \\ \sum_{i=1}^N d_i = d}} \frac{1}{\prod_{i=1}^N d_i!} \sum_{\substack{\alpha_1, \dots, \alpha_N \in \mathbb{N}^n \\ |\alpha_i| = d_i}} \prod_{i=1}^N \binom{d_i}{\alpha_i} \prod_{j=1}^n ((\mathbf{x}_i)_j (\mathbf{x}'_i)_j)^{\alpha_i(j)} \\
&= \mathbb{E}_{\substack{\mathbf{x}_1, \dots, \mathbf{x}_N \sim \mathcal{N}(0, X) \\ \mathbf{x}'_1, \dots, \mathbf{x}'_N \sim \mathcal{N}(0, X')}} \sum_{\substack{d_1, \dots, d_N \in 2\mathbb{N} \\ \sum_{i=1}^N d_i = d}} \prod_{i=1}^N \frac{\langle \mathbf{x}_i, \mathbf{x}'_i \rangle^{d_i}}{d_i!} \\
&= \sum_{\substack{d_1, \dots, d_N \in 2\mathbb{N} \\ \sum_{i=1}^N d_i = d}} \prod_{i=1}^N \frac{1}{d_i!} \mathbb{E}_{\substack{\mathbf{x} \sim \mathcal{N}(0, X) \\ \mathbf{x}' \sim \mathcal{N}(0, X')}} \langle \mathbf{x}, \mathbf{x}' \rangle^{d_i}. \tag{75}
\end{aligned}$$

From here, we introduce the moment-generating function of the inner overlap variables. Define

$$\phi_{X, X'}(t) := \mathbb{E}_{\substack{\mathbf{x} \sim \mathcal{N}(0, X) \\ \mathbf{x}' \sim \mathcal{N}(0, X')}} \exp(t \langle \mathbf{x}, \mathbf{x}' \rangle) = \sum_{d=0}^{\infty} \frac{t^d}{d!} \mathbb{E}_{\substack{\mathbf{x} \sim \mathcal{N}(0, X) \\ \mathbf{x}' \sim \mathcal{N}(0, X')}} \langle \mathbf{x}, \mathbf{x}' \rangle^d. \tag{76}$$

Then, $r'_d(X, X')$ is simply the coefficient of t^d in the Taylor series of $\phi_{X, X'}(t)^N$.

On the other hand, we may actually compute $\phi_{X, X'}(t)$:

$$\begin{aligned}
\phi_{X, X'}(t) &= \mathbb{E}_{\mathbf{g}, \mathbf{h} \sim \mathcal{N}(0, I_n)} \exp\left(t \mathbf{g}^\top \sqrt{X} \sqrt{X'} \mathbf{h}\right) \\
&= \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, I_{2n})} \exp\left(\mathbf{g}^\top \begin{bmatrix} 0 & \frac{t}{2} \sqrt{X} \sqrt{X'} \\ \frac{t}{2} \sqrt{X'} \sqrt{X} & 0 \end{bmatrix} \mathbf{g}\right)
\end{aligned}$$

which may be calculated as a moment generating function of the “matrix χ^2 ” variable $gg^\top$, giving

$$\begin{aligned}
&= \det\left(\begin{bmatrix} I_n & -t\sqrt{X'}\sqrt{X} \\ -t\sqrt{X}\sqrt{X'} & I_n \end{bmatrix}\right)^{-1/2} \\
&= \det\left(I_n - t^2 \sqrt{X} X' \sqrt{X}\right)^{-1/2}
\end{aligned}$$

and applying Proposition 56,

$$= \det(I_n - t^2 X X')^{-1/2}. \tag{77}$$

Thus $r'_d(X, X')$ is also the coefficient of t^d in the Taylor series of $\det(I_n - t^2 X X')^{-N/2}$, whereby $r'_d(X, X') = r_d(X, X')$, completing the proof. $\blacksquare$

Implicit in the above proof are the following facts which we have made use of in Section E.3.

Proposition 65 *Consider the general spiked Wishart model (Definition 17), and assume (without loss of generality) that $\beta = 1$ and $\mathcal{X}$ is supported on X for which $X \succ -I_n$. The following formulas hold:*

- (a) $\|L^{\leq D}\|^2 = \mathbb{E}_{\mathbf{X}, \mathbf{X}' \sim \mathcal{X}} \sum_{d=0}^D r_d(\mathbf{X}, \mathbf{X}')$ for some polynomials r_d ,
- (b) $\sum_{d=0}^{\infty} r_d(X, X') = \det(I_n - XX')^{-N/2}$,
- (c) if $X \succeq 0$ and $X' \succeq 0$ then $r_d(X, X') = \sum_{\substack{d_1, \dots, d_N \in 2\mathbb{N} \\ \sum_{i=1}^N d_i = d}} \prod_{i=1}^N \frac{1}{d_i!} \mathbb{E}_{\substack{\mathbf{x} \sim \mathcal{N}(0, X) \\ \mathbf{x}' \sim \mathcal{N}(0, X')}} \langle \mathbf{x}, \mathbf{x}' \rangle^{d_i}$, and
- (d) if $X \preceq 0$ and $X' \preceq 0$ then $r_d(X, X') = \sum_{\substack{d_1, \dots, d_N \in 2\mathbb{N} \\ \sum_{i=1}^N d_i = d}} \prod_{i=1}^N \frac{1}{d_i!} \mathbb{E}_{\substack{\mathbf{x} \sim \mathcal{N}(0, -X) \\ \mathbf{x}' \sim \mathcal{N}(0, -X')}} \langle \mathbf{x}, \mathbf{x}' \rangle^{d_i}$.

Appendix G. Exponential Low-Degree Hardness for SBM

The goal of this section is to prove Theorem 23. We first prove a general statement that reduces binary-valued models to the analogous Gaussian model.

G.1. Comparing Binary-Valued Models to Gaussian

Proposition 66 *Consider the following general binary-valued problem.*

- Under the null distribution $\mathbb{Q}$, we observe $\mathbf{Y} \in \mathbb{R}^N$ where $\mathbf{Y}_i$ are independent, satisfy $\mathbb{E}[\mathbf{Y}_i] = 0$ and $\mathbb{E}[\mathbf{Y}_i^2] = 1$, and each take two possible values: $\mathbf{Y}_i \in \{a_i, b_i\}$ with $a_i < b_i$.
- Under the planted distribution $\mathbb{P}$, a signal $\mathbf{X} \in \mathbb{R}^N$ is drawn from some prior, and then $\mathbf{Y}_i \in \{a_i, b_i\}$ are drawn independently (conditioned on $\mathbf{X}$) such that $\mathbb{E}[\mathbf{Y}_i | \mathbf{X}_i] = \mathbf{X}_i$. (This requires $\mathbf{X}_i \in [a_i, b_i]$.)

In the above setting,

$$\|L^{\leq D}\|^2 \leq \sum_{d=0}^D \frac{1}{d!} \mathbb{E}_{\mathbf{X}, \mathbf{X}'} \langle \mathbf{X}, \mathbf{X}' \rangle^d \quad (78)$$

where $\mathbf{X}'$ denotes an independent copy of $\mathbf{X}$.

The significance of (78) is that the right-hand side is the exact formula for $\|L^{\leq D}\|^2$ in the following additive Gaussian model (see Kunisky et al. (2019)): under $\mathbb{Q}$, $\mathbf{Y} \sim \mathcal{N}(0, I_N)$; and under $\mathbb{P}$, $\mathbf{Y} = \mathbf{X} + \mathbf{Z}$ with $\mathbf{Z} \sim \mathcal{N}(0, I_N)$ and $\mathbf{X}$ drawn from some prior. Thus, Proposition 66 can be interpreted as saying that a binary-valued problem is at least as hard as the corresponding Gaussian problem.

Proof [Proof of Proposition 66] The Fourier characters $\chi_S(\mathbf{Y}) = \prod_{i \in S} \mathbf{Y}_i$ for $S \subseteq [N]$ with $|S| \leq D$ are orthonormal with respect to $\mathbb{Q}$, in the sense that $\mathbb{E}_{\mathbf{Y} \sim \mathbb{Q}}[\chi_S(\mathbf{Y}) \chi_T(\mathbf{Y})] = \mathbf{1}[S = T]$. They also span the subspace of degree $\leq D$ polynomials, since for any $r \in \mathbb{N}$, any Y_i^r can be written as a degree-1 polynomial in Y_i . Thus, $\{\chi_S\}_{|S| \leq D}$ is an orthonormal basis for the degree $\leq D$ polynomials. It is a standard fact that this allows us to write $\|L^{\leq D}\|^2 = \sum_{|S| \leq D} (\mathbb{E}_{\mathbf{Y} \sim \mathbb{P}}[\chi_S(\mathbf{Y})])^2$

(see e.g., [Hopkins and Steurer \(2017\)](#); [Hopkins \(2018\)](#)). We will use S to denote a subset of $[N]$ and use α to denote an ordered multi-set of $[N]$, with $\chi_\alpha(Y) := \prod_{i \in \alpha} Y_i$. Now compute

$$\begin{aligned}
\|L^{\leq D}\|^2 &= \sum_{|S| \leq D} (\mathbb{E}_{\mathbf{Y} \sim \mathbb{P}} \chi_S(\mathbf{Y}))^2 \\
&= \sum_{|S| \leq D} (\mathbb{E}_{\mathbf{X}} \chi_S(\mathbf{X}))^2 \\
&\stackrel{(*)}{\leq} \sum_{d=0}^D \sum_{|\alpha|=d} \frac{1}{d!} (\mathbb{E}_{\mathbf{X}} \chi_\alpha(\mathbf{X}))^2 && \text{(see below)} \\
&= \sum_{d=0}^D \sum_{|\alpha|=d} \frac{1}{d!} \mathbb{E}_{\mathbf{X}, \mathbf{X}'} \chi_\alpha(\mathbf{X}) \chi_\alpha(\mathbf{X}') \\
&= \sum_{d=0}^D \frac{1}{d!} \mathbb{E}_{\mathbf{X}, \mathbf{X}'} \langle \mathbf{X}, \mathbf{X}' \rangle^d.
\end{aligned}$$

To see that inequality $(*)$ holds, note that it would be an *equality* if restricted to ordered multi-sets α that contain distinct elements. $\blacksquare$

G.2. Stochastic Block Model

We now specialize to the case of the stochastic block model (Definition 22), and use Proposition 66 to reduce to a certain spiked Wigner model.

Proof [Proof of Theorem 23] It will be convenient to consider a modification of the SBM that allows self-loops. Specifically, the edge (i, i) occurs with probability $\frac{d}{n}$ under $\mathbb{Q}$ and with probability $(1 + \frac{\eta}{\sqrt{2}}(k-1))\frac{d}{n}$ under $\mathbb{P}$. It is clear from the variational formula (13) for $\|L^{\leq D}\|$ that revealing this extra information can only increase $\|L^{\leq D}\|$.

In order to place the SBM in the setting of Proposition 66, take $N = n(n+1)/2$ with a variable $\mathbf{Y}_{i,j}$ for every $i \leq j$. Let $p = d/n$. In order to ensure $\mathbb{E}_{\mathbb{Q}}[\mathbf{Y}_{i,j}] = 0$ and $\mathbb{E}_{\mathbb{Q}}[\mathbf{Y}_{i,j}^2] = 1$, take $\mathbf{Y}_{i,j} = b := \sqrt{(1-p)/p}$ if edge (i, j) is present, and $\mathbf{Y}_{i,j} = a := -\sqrt{p/(1-p)}$ otherwise.

We now define $\mathbf{X}$ appropriately. Conditioned on the community structure, edge (i, j) occurs with probability $(1 + \Delta_{i,j})p$ where $\Delta_{i,i} = \eta(k-1)/\sqrt{2}$ and $\Delta_{i,j}$ (for $i < j$) is either $\eta(k-1)$ or $-\eta$ depending on whether i and j belong to the same community or not, respectively. Using the fact $pb + (1-p)a = 0$ (twice), this means we should take

$$\mathbf{X}_{i,j} = \mathbb{E}[\mathbf{Y}_{i,j} | \mathbf{X}_{i,j}] = (1 + \Delta_{i,j})pb + [1 - (1 + \Delta_{i,j})p]a = \Delta_{i,j}p(b-a) = -\Delta_{i,j}a = \Delta_{i,j}\sqrt{\frac{p}{1-p}}.$$

Let $\mathbf{U}$ be the $n \times k$ matrix whose i th row is $\sqrt{k}e_{\mathbf{k}_i} - \mathbb{1}/\sqrt{k}$ where $\mathbf{k}_i \in [k]$ is the community assignment of vertex i . One can check that $(\mathbf{U}\mathbf{U}^\top)_{i,j} = k\mathbb{1}[\mathbf{k}_i = \mathbf{k}_j] - 1$ and so $\mathbf{X}_{i,j} = \eta\sqrt{\frac{p}{1-p}}(\mathbf{U}\mathbf{U}^\top)_{i,j}$ for $i < j$, and $\mathbf{X}_{i,i} = \frac{\eta}{\sqrt{2}}\sqrt{\frac{p}{1-p}}(\mathbf{U}\mathbf{U}^\top)_{i,i}$. Therefore $\langle \mathbf{X}, \mathbf{X}' \rangle = \frac{\eta^2}{2} \frac{p}{1-p} \langle \mathbf{U}\mathbf{U}^\top, \mathbf{U}'(\mathbf{U}')^\top \rangle$. By Proposition 66,

$$\|L^{\leq D}\|^2 \leq \sum_{d=0}^D \frac{1}{d!} \mathbb{E} \left(\frac{\eta^2}{2} \frac{p}{1-p} \langle \mathbf{U}\mathbf{U}^\top, \mathbf{U}'(\mathbf{U}')^\top \rangle \right)^d.$$

Comparing this to (55) reveals that this is precisely the expression for $\|L^{\leq D}\|^2$ in the general spiked Wigner model with spike prior $\mathcal{X}_k$ (see Definition 24), except in place of λ^2 we have $\frac{\eta^2 d}{1-p} = (1 + o(1))\eta^2 d$. Therefore, appealing to the Wigner result (Theorem 19), we have $\|L^{\leq D}\| = O(1)$ provided $d\eta^2 < 1$. ■