Achieving the Bayes Error Rate in Synchronization and Block Models by SDP, Robustly

Yingjie Fei and Yudong Chen

***Abstract*— We study the statistical performance of semidefinite programming (SDP) relaxations for clustering under random graph models. Under the $\mathbb{Z}_2$ Synchronization model, Censored Block Model (CBM) and Stochastic Block Model (SBM), we show that SDP achieves an error rate of the form $\exp[-(1-o(1))\bar{n}I^*]$. Here $\bar{n}$ is an appropriate multiple of the number of nodes and I^* is an information-theoretic measure of the signal-to-noise ratio. We provide matching lower bounds on the Bayes error for each model and therefore demonstrate that the SDP approach is Bayes optimal. As a corollary, our results imply that SDP achieves the optimal exact recovery threshold under each model. Furthermore, we show that SDP is robust: the above bound remains valid under semirandom versions of the models in which the observed graph is modified by a monotone adversary. Our proof is based on a novel primal-dual analysis of SDP under a unified framework for all three models, and the analysis shows that SDP tightly approximates a joint majority voting procedure.**

***Index Terms*— Stochastic Block Model (SBM), Censored Block Model (CBM), synchronization, semidefinite programming (SDP), convex relaxation, Bayes error rates, semi-random robustness.**

I. Introduction

Clustering and community detection in graphs is an important problem lying at the intersection of computer science, optimization, statistics and information theory. Random graph models provide a venue for studying the average-case behavior of these problems. In these models, noisy pairwise observations are generated randomly according to the unknown clustering structure of the nodes. In its basic form, such a model involves n nodes divided into two clusters, which can be represented by a vector $\boldsymbol{\sigma}^* \in \{\pm 1\}^n$. For each pair of nodes i and j, one observes a number $A_{ij} \in \mathbb{R}$ generated independently based on the sign of $\sigma_i^*\sigma_j^*$, that is, whether the two nodes are in the same cluster or not. Given one realization of the random graph $\mathbf{A} = (A_{ij}) \in \mathbb{R}^{n\times n}$, the goal is to estimate the vector $\boldsymbol{\sigma}^*$, or equivalently, the matrix $\mathbf{Y}^* := (\sigma_i^*\sigma_j^*) \in \{\pm 1\}^{n\times n}$. Among the most popular random graph models are the $\mathbb{Z}_2$ Synchronization (Z2) model, Censored Block Model (CBM) and Stochastic Block Model (SBM), where A_{ij} follows the Gaussian, censored ± 1 and Bernoulli distributions, respectively (see Section III for the details). We consider these three models in this paper.

Clustering is a challenging problem involving discrete and hence non-convex optimization. semidefinite programming (SDP) relaxations have emerged as an efficient and robust approach to this problem, and recent work has witnessed the advances in establishing rigorous performance guarantees for SDP (see Section II for a review of this literature). Such guarantees are typically stated in terms of a signal-to-noise ratio (SNR) measure I^* that depends on the specific random model (see Equation (5)). In terms of controlling the *estimation error* of SDP, the best and most general results to date are given in the line of work in [2], [3], which proves that the optimal SDP solution $\widehat{\mathbf{Y}}$ satisfies the bound

$$\mathtt{err}(\widehat{\boldsymbol{\sigma}}^{\text{sdp}}, \boldsymbol{\sigma}^*) \lesssim \frac{1}{n^2}\|\widehat{\mathbf{Y}} - \mathbf{Y}^*\|_1 \lesssim \exp\left[-\frac{nI^*}{C}\right], \tag{1}$$

where $C > 0$ is a large constant, $\|\cdot\|_1$ denotes the entrywise ℓ_1 norm, and $\mathtt{err}(\widehat{\boldsymbol{\sigma}}^{\text{sdp}}, \boldsymbol{\sigma}^*)$ denotes the fraction of nodes mis-clustered by an estimate $\widehat{\boldsymbol{\sigma}}^{\text{sdp}} \in \{\pm 1\}^n$, extracted from $\widehat{\mathbf{Y}}$, of the ground-truth cluster labels $\boldsymbol{\sigma}^*$. The above result is, however, unsatisfactory due to the presence of a large multiplicative constant C in the exponent, rendering the bound fundamentally sub-optimal. In particular, the interesting regime for proving an error bound is when $nI^* \le 2\log n$, as otherwise SDP is already known to attain zero error. With a large C in the exponent, the result in (1) provides a rather loose, sometimes even uninformative,[1] bound in this regime. Moreover, this sub-optimality is intrinsic to the proof techniques used and cannot be avoided simply by more careful calculations.

In this paper, we establish a strictly tighter, and essentially optimal, error bound on SDP. Let $\bar{n} = n$ for Z2 and CBM, and $\bar{n} = \frac{n}{2}$ for SBM.

Theorem 1 (Informal): As $n \to \infty$, with probability tending to one, the optimal solution $\widehat{\mathbf{Y}}$ of the SDP relaxation satisfies

$$\frac{1}{n^2}\|\widehat{\mathbf{Y}} - \mathbf{Y}^*\|_1 \le \exp\Big[-\big(1-o(1)\big)\bar{n}I^*\Big], \tag{2}$$

Moreover, the explicit label estimate $\widehat{\boldsymbol{\sigma}}^{\text{sdp}}$ computed by taking entrywise signs of the top eigenvector of $\widehat{\mathbf{Y}}$ satisfies

$$\mathtt{err}(\widehat{\boldsymbol{\sigma}}^{\text{sdp}}, \boldsymbol{\sigma}^*) \le \exp\Big[-\big(1-o(1)\big)\bar{n}I^*\Big]. \tag{3}$$

Manuscript received April 23, 2019; revised October 19, 2019; accepted December 2, 2019. Date of publication January 13, 2020; date of current version May 20, 2020. This work was supported in part by the National Science Foundation under Grant 1657420 and Grant CCF-1704828 and in part by the School of Operations Research and Information Engineering, Cornell University. This article was presented in part at the 2019 Conference on Learning Theory (COLT).

The authors are with the School of Operations Research and Information Engineering, Cornell University, Ithaca, NY 14850 USA (e-mail: yf275@cornell.edu; yudong.chen@cornell.edu).

Communicated by M. Davenport, Associate Editor for Signal Processing.
Digital Object Identifier 10.1109/TIT.2020.2966438

[1]Note that $\frac{1}{n^2}\|\widehat{\mathbf{Y}} - \mathbf{Y}^*\|_1$ is trivially upper bounded by 2 since $\widehat{\mathbf{Y}}, \mathbf{Y}^* \in [-1,1]^{n\times n}$.

In all three models, the error exponent I^* is a form of Renyi divergence. See Theorem 4 for the precise statement of our results as well as an explicit, non-asymptotic estimate of the $o(1)$ term. One should compare this result with the following minimax lower bound, which shows that any estimator $\widehat{\boldsymbol{\sigma}}$ must incur an error

$$\mathtt{err}(\widehat{\boldsymbol{\sigma}}, \boldsymbol{\sigma}^*) \geq \exp\Big[-\big(1+o(1)\big)\bar{n}I^*\Big]. \tag{4}$$

as the latter represents the best achievable *Bayes risk* of the problem. For SBM, this bound is established in [4]; for Z2 and CBM, the above lower bound is new and formally established in Theorem 3. In view of the above upper and lower bound, we see that SDP achieves the optimal Bayes error under all three models.

a) Optimality as a surprise: The result above has come as unexpected to us, as it shows that relaxing the original discrete clustering problem via SDP incurs essentially no loss in terms of statistical accuracy. As we discuss in Section I-A and further elaborate in Section V, we prove this result by showing, via a novel primal-dual analysis, that SDP tightly approximates a majority voting procedure, and this procedure leads to the optimal error exponent I^*. Interestingly, our analysis is not tethered to the optimality of $\widehat{\mathbf{Y}}$ to the SDP; rather, it only relies on the fact that $\widehat{\mathbf{Y}}$ is feasible and no worse in objective value than $\mathbf{Y}^*$, and thus the bounds (2) and (3) in fact hold for any matrix $\mathbf{Y}$ with these two properties. This kind of leeway in the analysis makes the bounds robust, as we elaborate next.

b) Robustness: We show that the bounds in Theorem 1 continue to hold under the so-called *monotone semirandom model* [5], where an adversary is allowed to make arbitrary changes to the graph in a way that apparently strengthens connections within each cluster and weakens connections between clusters. While this model seemingly makes the clustering problem easier, they in fact foil, provably, many existing algorithms, particularly those that over-exploit the specific structures of standard SBM in order to achieve tight recovery guarantees [5], [6]. In contrast, our results show that SDP relaxations enjoy a robustness property that is possessed by few other algorithms. Importantly, this generalization can be achieved with little extra effort from our main result (see Theorem 7 and its proof).

c) Exact recovery: As another illustration of the strength of Theorem 1, we note that it implies sharp condition for SDP to recover $\boldsymbol{\sigma}^*$ *exactly*. In particular, when $\bar{n}I^* > (1+\delta)\log n$ for any positive constant δ, the bound (3) ensures that $\mathtt{err}(\widehat{\boldsymbol{\sigma}}^{\mathrm{sdp}}, \boldsymbol{\sigma}^*) < \frac{1}{n}$ and hence $\mathtt{err}(\widehat{\boldsymbol{\sigma}}^{\mathrm{sdp}}, \boldsymbol{\sigma}^*) = 0$. Moreover, the lower bound (4) shows that exact recovery is information-theoretically impossible when $\bar{n}I^* < \log n$. In the literature, establishing such tight exact recovery thresholds often involves specialized and sophisticated arguments, and has been the milestones in the remarkable recent development on community detection (see Section II for a discussion of related work). We recover these results, for all three models, as a corollary of our main theorem by plugging in the corresponding expressions of I^* and $\bar{n}$. In fact, the non-asymptotic version of Theorem 1 guarantees exact recovery via SDP with an explicit second-order term $\delta = O\big(1/\sqrt{\log n}\big)$, which is a refinement of existing results.

A. Primal-Dual Analysis

Key to the establishment of our results is a novel analysis that exploits both primal and dual characterizations of the SDP. To set the context, we note that the sub-optimal bound (1) in [3], [7] is established by utilizing the primal optimality of the SDP solution $\widehat{\mathbf{Y}}$. Their arguments, however, are too crude to provide a tight estimate of the multiplicative constant C in the exponent. On the other hand, work on exact recovery for SDP typically makes use of a dual analysis [8], [9]; in particular, the optimality of $\mathbf{Y}^*$ is certified by showing the existence of a corresponding dual optimal solution, often explicitly in the form of a diagonal matrix $\mathbf{D}$ with $D_{ii} = \sigma_i^* \sum_j A_{ij}\sigma_j^*$. As this "dual certificate" $\mathbf{D}$ is tied to (and constructed using) $\mathbf{Y}^*$, such a certification approach would only succeed when the SDP indeed admits $\mathbf{Y}^*$ as an optimal solution.

Here we are concerned with the setting where the optimal solution $\widehat{\mathbf{Y}}$ is different from $\mathbf{Y}^*$, and our goal is to bound their difference. As it is a priori unknown what $\widehat{\mathbf{Y}}$ should look like, we do not know which matrix to certify or how to construct its associated dual solution, rendering the above dual certification argument inapplicable. Instead, we make use of the fact that $\widehat{\mathbf{Y}}$ is feasible to the SDP and has a better primal objective value than $\mathbf{Y}^*$, that is, $\widehat{\mathbf{Y}}$ lies in the sublevel set defined by $\mathbf{Y}^*$ and the constraints of the SDP. We then characterize the diameter of this sublevel set by using, perhaps surprisingly, the dual certificate $\mathbf{D}$ of $\mathbf{Y}^*$. Our analysis is thus fundamentally different from the dual certification analysis in existing work, which only applies when the sublevel set consists of a single element $\mathbf{Y}^*$. At the same time, we make use of $\mathbf{D}$ in a crucial way to achieve an exponential improvement over previous primal analysis.

Note that our analysis, and hence our error bounds as well, actually apply to *every element* of this sublevel set, not just the optimal solution $\widehat{\mathbf{Y}}$. As can be seen in our proof, this flexibility plays an important role in establishing the aforementioned robustness results under semirandom and heterogeneous SBMs. On the other hand, however, with this level of generality we probably should not expect the second-order $o(1)$ term in our bounds to be optimal.

Finally, we emphasize that our results for Z2, CBM and SBM are proved under a unified framework. The main proof steps are deterministic and hold for the three models at once; only certain probabilistic arguments are model-specific. In Section V we outline this proof framework, and provide intuitions on the majority voting mechanism that drives the error rate $e^{-\bar{n}I^*}$. We believe that this unified framework may be broadly useful in studying SDP relaxations for other discrete problems under average-case/probabilistic settings.

B. Paper Organization

In Section II, we review related work on Z2, CBM and SBM. In Section III, we formally introduce the models and the SDP relaxation approach. In Section IV, we present our main results, with a discussion on their consequences and comparison with existing work. We outline the main steps of the proofs and discuss the intuitions in Section V, with the complete proofs deferred to the appendix. The paper is concluded in Section VI with a discussion on future directions.

II. Related Work

There is a large array of recent results on community detection and graph clustering, in particular, under Z2, CBM and SBM. The readers are referred to the surveys [10]–[12] for comprehensive reviews. Without trying to enumerate this body of work, here we restrict attention to those that study sharp performance bounds, with a particular focus on work on the SDP relaxation approach. A more detailed, quantitative comparison with our results is provided in Section IV after our main theorems.

To begin, we note that existing work has considered several recovery criteria for an estimator $\widehat{\boldsymbol{\sigma}}$ of $\boldsymbol{\sigma}^*$: *weak recovery* means $\widehat{\boldsymbol{\sigma}}$ is better than random guess, that is, $\texttt{err}(\widehat{\boldsymbol{\sigma}}, \boldsymbol{\sigma}^*) < \frac{1}{2}$; *partial recovery* means $\texttt{err}(\widehat{\boldsymbol{\sigma}}, \boldsymbol{\sigma}^*) \leq \delta$ for a given $\delta \in (0, \frac{1}{2})$; *exact recovery* means $\texttt{err}(\widehat{\boldsymbol{\sigma}}, \boldsymbol{\sigma}^*) = 0$ [10].

A. $\mathbb{Z}_2$ Synchronization and Censored Block Model

The Z2 model, being a simplified version of the angular/phase synchronization problem, is studied in [13], which argues that exact recovery is possible if and only if $I^* > \frac{\log n}{n}$. The work in [9] and [14] shows this optimal exact recovery threshold is achieved by SDP and a spectral algorithm, respectively. The work in [15], [16] considers low-rank matrix estimation under a spiked Wigner model—of which Z2 is a special case—and identifies the weak recovery threshold.

CBM is considered in [17], [18], which identifies sufficient and necessary conditions for exact recovery. They also show that SDP achieves a sub-optimal exact recovery threshold, which is further improved to be optimal in [9], [18]. CBM is a special case of the so-called Labelled SBM, whose weak recovery threshold is studied in [19], [20]. Achieving tight partial/weak recovery guarantees in CBM is challenging due to the sparsity of the observations. A sub-optimal partial recovery error bound can be achieved by a spectral algorithm with trimming [21]. The work of [22] studies a sophisticated spectral algorithm based on the non-backtracking operator or Bethe Hessian, and shows that it achieves the optimal weak recovery threshold.

For both Z2 and CBM, we establish for the first time that SDP has the optimal error rates for partial recovery. Our results also imply, as an immediate corollary, that SDP achieves the optimal exact recovery threshold as well as a sub-optimal weak recovery threshold.

B. Stochastic Block Model

SBM is arguably the most studied out of these three models. Most related to us is a line of work that characterizes minimax optimal error rates for partial recovery. For the binary symmetric SBM, the work [4] establishes the aforementioned minimax lower bound (4). They also provide an exponential-time algorithm that achieves a matching upper bound (up to an $o(1)$ factor in the exponent). Much research effort focuses on developing computationally feasible algorithms, and identifying the minimax rates in more general settings [23]–[29]. The monograph [30] provides a review on recent work on this front. We note that this line of work does not consider the SDP relaxation approach nor deliver robustness guarantees as we do. Nevertheless, we will compare our results with theirs after stating our main theorems.

For exact recovery under binary symmetric SBM with $p, q \asymp \frac{\log n}{n}$ (see Model 3 in Section III-B for the definitions of p and q), the work in [31], [32] establishes the sufficient and necessary condition $(\sqrt{p} - \sqrt{q})^2 > \frac{2 \log n}{n}$. Follow-up work develops efficient algorithms for exact recovery and considers extensions to more general SBMs; see, e.g., [14], [33]–[36]. As mentioned, our results imply sharp bounds for exact recovery.

Weak recovery under the binary symmetric SBM is most relevant in the sparse regime $p, q \asymp \frac{1}{n}$. Work of [20], [37], [38] establishes that the necessary and sufficient condition of weak recovery is $\frac{n(p-q)^2}{p+q} > 2$. Subsequent work proves similar phase transitions and shows that various algorithms achieve weak recovery above the optimal threshold for the SBM with $k \geq 2$ and possibly unbalanced clusters; see, e.g., [39]–[46]. As discussed later, our results also imply weak recovery guarantees with a sub-optimal constant.

C. Optimality and Robustness of SDP

For SBM, SDP has been proven to succeed in exact and weak recovery above the corresponding optimal thresholds (sometimes under additional assumptions). In particular, see [8], [9], [47] for exact recovery, and [48] for weak recovery. Prior to our work, SDP was not known to achieve the optimal error rate between the exact and weak recovery regimes. Sub-optimal polynomial rates are first proved in [2], later improved to exponential in [3], and further generalized in [7], [49].

Robustness has been recognized as a distinct feature of the SDP approach as compared to other more specialized algorithms for SBMs. Work in this direction has established robustness of SDP against random erasures [18], [50], atypical node degrees [2] and adversarial corruptions [8], [48], [51], [52]. The work in [6] investigates the relationship between statistical optimality and robustness under monotone semirandom models; we revisit this result in more details later.

Preliminary version of this work has appeared in [1], in which error upper bounds for SBM are presented. The present work proves that SDP achieves the Bayes error rates for a more general class of models including Z2, CBM and SBM, as well as their semirandom versions, under a unified framework. This work also establishes matching minimax lower bounds for Z2 and CBM.

III. Problem Set-Up

In this section, we formally define the models and introduce the SDP relaxation approach.

A. Notations

Vectors and matrices are denoted by bold letters. For a vector $\mathbf{u}$, u_i and $u(i)$ both denote its i-th entry. For a matrix $\mathbf{M}$, we let M_{ij} denote its (i,j)-th entry, $\mathrm{Tr}(\mathbf{M})$ its trace, and

$\|\mathbf{M}\|_1 := \sum_{i,j} |M_{ij}|$ its entry-wise ℓ_1 norm. We write $\mathbf{M} \succeq 0$ if $\mathbf{M}$ is symmetric positive semidefinite. The trace inner product between two matrices is $\langle \mathbf{M}, \mathbf{G} \rangle := \mathrm{Tr}(\mathbf{M}^\top \mathbf{G}) = \sum_{i,j} M_{ij} G_{ij}$. Denote by $\mathbf{I}$ and $\mathbf{J}$ the $n \times n$ identity matrix and all-one matrix, respectively, and denote by $\mathbf{1}$ the all-one column vector of length n.

$\mathrm{Ber}(\mu)$ denotes the Bernoulli distribution with mean $\mu \in [0, 1]$. For a positive integer i, let $[i] := \{1, 2, \ldots, i\}$. For a real number x, $\lceil x \rceil$ denotes its ceiling and $\lfloor x \rfloor$ denotes its floor. $\mathbb{I}\{\cdot\}$ is the indicator function. For two non-negative sequences $\{a_n\}$ and $\{b_n\}$, we write $a_n = O(b_n)$, $b_n = \Omega(a_n)$ or $a_n \lesssim b_n$ if there exists a universal constant $C > 0$ such that $a_n \leq C b_n$ for all n. We write $a_n \asymp b_n$ if both $a_n = O(b_n)$ and $a_n = \Omega(b_n)$ hold. Asymptotic statements are with respect to the regime $n \to \infty$, in which case we write $a_n = o(b_n)$ and $b_n = \omega(a_n)$ if $\lim_{n\to\infty} a_n/b_n = 0$.

B. Models

In this section, we formally describe four models for generating the observed matrix $\mathbf{A}$ from the unknown ground-truth label vector $\boldsymbol{\sigma}^* \in \{\pm 1\}^n$.

In Z2 [9], each A_{ij} is generated by adding Gaussian noise to $\sigma_i^* \sigma_j^*$. Therefore, the matrix $\mathbf{A}$ contains noisy observations of the true relative signs between each pair of nodes.

Model 1 (*$\mathbb{Z}_2$ Synchronization):* The observed matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ is symmetric with its entries $\{A_{ij}, i \leq j\}$ generated independently by

$$A_{ij} \sim N(\sigma_i^* \sigma_j^*, \tau^2),$$

where $\tau > 0$ is allowed to scale with n.[2]

In CBM, each A_{ij} is generated by flipping $\sigma_i^* \sigma_j^*$ with probability ϵ and then erasing it with probability $1 - \alpha$. One may interpret $\mathbf{A}$ as the edge-censored version of a noisy signed network [17].

Model 2 (*Censored Block Model):* The observed matrix $\mathbf{A} \in \{0, \pm 1\}^{n \times n}$ is symmetric with its entries $\{A_{ij}, i \leq j\}$ generated independently by

$$A_{ij} = \begin{cases} \sigma_i^* \sigma_j^* & \text{with probability (w.p.) } \alpha(1-\epsilon), \\ -\sigma_i^* \sigma_j^* & \text{w.p. } \alpha\epsilon, \\ 0 & \text{w.p. } 1 - \alpha, \end{cases}$$

where $\alpha \in (0, 1]$ is allowed to scale with n, and $\epsilon \in (0, \frac{1}{2})$ is a constant.

In SBM, each A_{ij} is a Bernoulli random variable, whose mean is higher if $\sigma_i^* \sigma_j^* = 1$. Therefore, $\mathbf{A}$ is the adjacency matrix of a random graph in which nodes in the same cluster are more likely to be connected than those in different clusters [53].

Model 3 (*Binary symmetric SBM):* Suppose that the ground-truth $\boldsymbol{\sigma}^* \in \{\pm 1\}^n$ satisfies $\langle \boldsymbol{\sigma}^*, \mathbf{1} \rangle = 0$. The observed matrix $\mathbf{A} \in \{0, 1\}^{n \times n}$ is symmetric with its entries $\{A_{ij}, i \leq j\}$ generated independently by

$$A_{ij} \sim \begin{cases} \mathrm{Ber}(p) & \text{if } \sigma_i^* \sigma_j^* = 1, \\ \mathrm{Ber}(q) & \text{if } \sigma_i^* \sigma_j^* = -1, \end{cases}$$

where $0 < q < p < 1$ are allowed to scale with n.

In both Model 1 (Z2) and Model 2 (CBM), there can be any number of ± 1's in the ground-truth label vector $\boldsymbol{\sigma}^* \in \{\pm 1\}^n$. In Model 3 (binary symmetric SBM), the cluster labels $\boldsymbol{\sigma}^*$ are assumed to contain the same number of 1's and -1's, so the two clusters have equal size. Despite their simple forms, the above models have been of central importance in studying fundamental limits of clustering problems [8], [9], [17], [20], [31], [32], [37], [38], [48].

For the purpose of studying the robustness properties of SDP relaxation, we consider a semirandom generalization of the binary symmetric SBM. In this model, a so-called monotone adversary, upon observing the random adjacency matrix $\mathbf{A}$ generated from SBM and the ground-truth clustering $\boldsymbol{\sigma}^*$, modifies $\mathbf{A}$ by *arbitrarily* adding edges between nodes of the same cluster and deleting edges between nodes of different clusters.

Model 4 (*Semirandom SBM):* A monotone adversary observes $\mathbf{A}$ and $\boldsymbol{\sigma}^*$ from Model 3, picks an arbitrary set of pairs of nodes $\mathcal{L} \subset \{(i, j) \in [n] \times [n] : i < j\}$, and outputs a symmetric matrix $\mathbf{A}^{\mathrm{SR}} \in \{0, 1\}^{n \times n}$ such that for each $i < j$,

$$A_{ij}^{\mathrm{SR}} = \begin{cases} 1 & \text{if } (i,j) \in \mathcal{L},\ \sigma_i^* \sigma_j^* = 1, \\ 0 & \text{if } (i,j) \in \mathcal{L},\ \sigma_i^* \sigma_j^* = -1, \\ A_{ij}, & \text{if } (i,j) \notin \mathcal{L}. \end{cases}$$

Note that the set $\mathcal{L}$ is allowed to depend on the realization of $\mathbf{A}$.

Semirandom models have a long history with many variants [54]. Model 4 above has been considered in [5], [6] for SBM. While seemingly revealing more information about the underlying cluster structure, the semirandom model in fact destroys many local structures of the basic SBM, thus frustrating many algorithms that over-exploit such structures. In contrast, SDP is robust against the monotone adversary under Model 4, as we shall see in Section IV-C below.

Remark 2: One may define semirandom versions of Z2 and CBM in an analogous fashion as above; that is, the adversary may choose a set $\mathcal{L}$ and positive numbers $\{c_{ij}, i < j\}$, and then change A_{ij} and A_{ji} to $A_{ij} + c_{ij} \sigma_i^* \sigma_j^*$ for each $(i, j) \in \mathcal{L}$. It can be shown that SDP achieves the same performance guarantees in these semirandom settings of Z2 and CBM as in the original models. For conciseness we omit such details and only focus on the semirandom extension of SBM.

For each model discussed above, we define a measure of the signal-to-noise ratio (SNR):

$$I^* := \begin{cases} (2\tau^2)^{-1}, & \text{for Model 1}, \\ \left(\sqrt{\alpha(1-\epsilon)} - \sqrt{\alpha\epsilon}\right)^2, & \text{for Model 2}, \\ -2\log\big[\sqrt{pq} & \text{for Models 3 and 4.} \\ \quad + \sqrt{(1-p)(1-q)}\big], & \end{cases} \tag{5}$$

[2] In this and the next three models, we assume that the diagonal entries of $\mathbf{A}$ are random, which is inconsequential: these entries are independent of the ground-truth $\boldsymbol{\sigma}^*$, and they have no effect on the solutions of the SDP relaxations (7) or (8) due to the diagonal constraints therein.

In each case, I^* is a form of Renyi divergence of order $\frac{1}{2}$ [55] between the distributions of A_{ij} and $A_{ij'}$ with $\sigma^*_{ij} = -\sigma^*_{ij'} = 1$. In particular, for Z2, I^* is half of the Renyi divergence (or equivalently, the Kullback–Leibler divergence) between $N(1,\tau^2)$ and $N(-1,\tau^2)$. For CBM, we have $I^* \approx -\log(1-I^*)$ with the latter being half of the Renyi divergence between two random variables H and $-H$, where H has probability mass function $\alpha(1-\epsilon)\cdot\delta_1+\alpha\epsilon\cdot\delta_{-1}+(1-\alpha)\cdot\delta_0$ and δ_a denotes the Dirac delta function centered at a.[3] In SBM, I^* is the Renyi divergence between $\mathrm{Ber}(p)$ and $\mathrm{Ber}(q)$. These divergences, and their first-order approximations (discussed in Appendix A), are commonly used as SNR measures in previous work on these models (e.g., [4], [9], [17]).

Finally, we define the following distance measure between two vectors of cluster labels $\boldsymbol{\sigma}, \boldsymbol{\sigma'} \in \{\pm 1\}^n$:

$$\mathtt{err}(\boldsymbol{\sigma}, \boldsymbol{\sigma'}) := \min_{g\in\{\pm1\}} \frac{1}{n} \sum_{i\in[n]} \mathbb{I}\{g\sigma_i \neq \sigma'_i\}.$$

In words, $\mathtt{err}(\boldsymbol{\sigma}, \boldsymbol{\sigma'})$ is the fraction of nodes that are assigned a different label under $\boldsymbol{\sigma}$ and $\boldsymbol{\sigma'}$, modulo a global flipping of signs. With $\boldsymbol{\sigma}^*$ being the true labels, $\mathtt{err}(\widehat{\boldsymbol{\sigma}}, \boldsymbol{\sigma}^*)$ measures the relative error of the estimator $\widehat{\boldsymbol{\sigma}}$.

C. SDP Relaxation

The SDP formulations we consider can be derived as the convex relaxation of the MLE of $\boldsymbol{\sigma}^*$. Under Models 1 or 2, the MLE $\widehat{\boldsymbol{\sigma}}^{\text{mle}}$ is given by the solution of the discrete and non-convex optimization problem

$$\max_{\boldsymbol{\sigma}\in\{\pm1\}^n} \left\langle \mathbf{A}, \boldsymbol{\sigma}\boldsymbol{\sigma}^\top \right\rangle. \tag{6}$$

The MLE under Model 3 includes the extra constraint $\langle \boldsymbol{\sigma}, \mathbf{1}\rangle = 0$ due to the balanced-cluster assumption. Derivation of the MLE in this form is now standard; see for example [9], [13] for Z2, [50] for CBM, and [12] for SBM. Now define the lifted variable $\mathbf{Y} = \boldsymbol{\sigma}\boldsymbol{\sigma}^\top$, and observe that $\mathbf{Y}$ satisfies $\mathbf{Y} \succeq 0$, $Y_{ii} = (\sigma_i)^2 = 1$ for $i \in [n]$. Dropping the constraints that $\mathbf{Y}$ has rank one and binary entries, we obtain the following SDP relaxation of the MLE (6) for Models 1 or 2:

$$\begin{aligned} \widehat{\mathbf{Y}} = \arg\max_{\mathbf{Y}\in\mathbb{R}^{n\times n}} \ & \langle \mathbf{A}, \mathbf{Y}\rangle \\ \text{s.t. } & \mathbf{Y} \succeq 0, \\ & Y_{ii} = 1, \quad \forall i \in [n]. \end{aligned} \tag{7}$$

For Model 3, using the same reasoning and in addition replacing the $\langle \boldsymbol{\sigma}, \mathbf{1}\rangle = 0$ constraint by $\langle \mathbf{Y}, \mathbf{J}\rangle = \left\langle \boldsymbol{\sigma}\boldsymbol{\sigma}^\top, \mathbf{1}\mathbf{1}^\top \right\rangle = \langle \boldsymbol{\sigma}, \mathbf{1}\rangle^2 = 0$, we arrive at the relaxation:

$$\begin{aligned} \widehat{\mathbf{Y}} = \arg\max_{\mathbf{Y}\in\mathbb{R}^{n\times n}} \ & \langle \mathbf{A}, \mathbf{Y}\rangle \\ \text{s.t. } & \mathbf{Y} \succeq 0, \\ & Y_{ii} = 1, \ \forall i \in [n], \\ & \langle \mathbf{Y}, \mathbf{J}\rangle = 0. \end{aligned} \tag{8}$$

We also use this SDP for the semirandom Model 4.

[3] In fact, in this case I^* is the squared Hellinger distance between H and $-H$.

The optimization problems (7) and (8) are standard SDPs solvable in polynomial time. We remark that neither SDP requires knowing the parameters of the data generating processes (that is, τ^2, α, ϵ, p and q in Models 1–3).[4] The SDP (7) was considered in [9], [13] and [50] for studying the exact recovery threshold in Z2 and CBM, respectively, and the SDP (8) was considered in [8] for exact recovery under the binary symmetric SBM. These formulations can be further traced back to the work of [5] on SDP relaxation for MIN BISECTION.

We consider the SDP solution $\widehat{\mathbf{Y}}$ as an estimate of the ground-truth matrix $\mathbf{Y}^* := \boldsymbol{\sigma}^*(\boldsymbol{\sigma}^*)^\top$, and seek to characterize the accuracy of $\widehat{\mathbf{Y}}$ in terms of the ℓ_1 error $\|\widehat{\mathbf{Y}} - \mathbf{Y}^*\|_1$. Note that $\widehat{\mathbf{Y}}$ is not necessarily a rank-one matrix of the form $\widehat{\mathbf{Y}} = \boldsymbol{\sigma}\boldsymbol{\sigma}^\top$. To extract from $\widehat{\mathbf{Y}}$ a vector of binary estimates of cluster labels, we take the signs of the entries of the top eigenvector of $\widehat{\mathbf{Y}}$ (where the sign of 0 is 1, an arbitrary choice). Letting $\widehat{\boldsymbol{\sigma}}^{\text{sdp}} \in \{\pm1\}^n$ be the vector obtained in this way, we study the error of $\widehat{\boldsymbol{\sigma}}^{\text{sdp}}$ as an estimate of the ground-truth label vector $\boldsymbol{\sigma}^*$, as measured by $\mathtt{err}(\widehat{\boldsymbol{\sigma}}^{\text{sdp}}, \boldsymbol{\sigma}^*)$.

IV. Main Results

We present our main results in this section. Henceforth, let $\bar{n} := n$ in Models 1 (Z2) and 2 (CBM), and $\bar{n} := \frac{n}{2}$ in Models 3 and 4 (SBM and its semirandom version). To see why this definition of $\bar{n}$ is natural, we note that in Z2 and CBM, the cluster sizes (i.e., numbers of 1's and -1's in $\boldsymbol{\sigma}^*$) do not affect the hardness of the problem as the distribution is symmetric, and hence $\bar{n}$ is simply the number of nodes; in binary symmetric SBMs, recovery is most difficult when the clusters have equal size[5]—which is the setting we consider—and accordingly $\bar{n}$ is the cluster size.

A. Minimax Lower Bounds

Let $\ell_1(\boldsymbol{\sigma})$ denote the number of 1's in $\boldsymbol{\sigma}$. To state the lower bounds, we consider the following parameter space:

$$\Theta(n) := \begin{cases} \{\pm1\}^n, & \text{for Models 1 and 2,} \\ \{\boldsymbol{\sigma} \in \{\pm1\}^n : & \text{for Model 3,} \\ \quad \ell_1(\boldsymbol{\sigma}) \in \left[\frac{n}{2\beta}, \frac{n\beta}{2}\right]\}, & \end{cases} \tag{9}$$

where β is any number larger than $1+C/n$ with $C > 0$ being a large enough numerical constant. For Z2 and CBM, $\Theta(n)$ is the set of all possible cluster label vectors. For SBM, $\Theta(n)$ consists of label vectors with (roughly) equal-sized clusters; here we allow for a slight fluctuation in the cluster sizes in SBM following [4].[6]

The following theorem gives the minimax lower bound for each model.

Theorem 3 (Lower bound): For any constant $c_0 \in (0,1)$, the following holds for Model 1, Model 2, and Model 3 with

[4] The SDP (8) for SBM does require the knowledge of two equal-size clusters.

[5] Otherwise one could recover the large cluster first.

[6] This assumption is not essential but makes the proof therein somewhat simpler.

$0 < q < p < 1 - c_0$ (for Model 2 we additionally assume $I^* = o(1)$). If $nI^* \to \infty$ as $n \to \infty$, then we have

$$\inf_{\widehat{\boldsymbol{\sigma}}} \sup_{\boldsymbol{\sigma} \in \Theta(n)} \mathbb{E}_{\boldsymbol{\sigma}} \operatorname{\texttt{err}}(\widehat{\boldsymbol{\sigma}}, \boldsymbol{\sigma}) \geq \exp\left[-\big(1 + o(1)\big)\bar{n}I^*\right],$$

where $\mathbb{E}_{\boldsymbol{\sigma}}$ denotes expectation under the distribution of $\mathbf{A}$ with $\boldsymbol{\sigma}$ being the ground truth, and the infimum is taken over all estimators of the ground truth (i.e., measurable functions of $\mathbf{A}$).

For Models 1 and 2, the proof is given in Appendix B. For Model 3, the above result is part of [4, Theorem 2.1].

B. Upper Bounds on the SDP Errors

We next provide our main results on the error rate of the SDP relaxations (7) and (8). Define the following sublevel set (or superlevel set to be precise):

$$\mathcal{Y}(\mathbf{A}) := \Big\{ \mathbf{Y} \in \mathbb{R}^{n \times n} : \langle \mathbf{A}, \mathbf{Y} \rangle \geq \langle \mathbf{A}, \mathbf{Y}^* \rangle, \ \mathbf{Y} \text{ is feasible to the SDP} \Big\}, \tag{10}$$

where feasibility is with respect to the program (7) for Model 1 or 2, and to the program (8) for Model 3 or 4. In words, $\mathcal{Y}(\mathbf{A})$ is the set of feasible SDP solutions that attain an objective value no worse than the ground-truth $\mathbf{Y}^*$. Note that $\mathcal{Y}(\mathbf{A})$ is non-empty as it contains $\mathbf{Y}^*$. As mentioned, our upper bounds in fact hold for any solution in $\mathcal{Y}(\mathbf{A})$. With a slight abuse of notation, in the sequel we use $\widehat{\mathbf{Y}}$ to denote an arbitrary matrix in $\mathcal{Y}(\mathbf{A})$; accordingly, we let $\widehat{\boldsymbol{\sigma}}^{\text{sdp}}$ denote the corresponding vector of labels extracted from this $\widehat{\mathbf{Y}}$.

Our main theorem is a non-asymptotic bound on the error rates of the SDP relaxations.

Theorem 4 (Upper Bound): For any constants $c_0, c_1 \in (0,1)$, there exist constants $C_{I^*}, C_e, C_e' > 0$ such that the following holds for Model 1, Model 2, and Model 3 with $0 < c_0 p \leq q < p \leq 1 - c_1$. If $nI^* \geq C_{I^*}$, then with probability at least $1 - 10\exp\left(-\sqrt{\log n}\right)$ and for any $\widehat{\mathbf{Y}} \in \mathcal{Y}(\mathbf{A})$, we have

$$\frac{1}{n}\|\widehat{\mathbf{Y}} - \mathbf{Y}^*\|_1 \leq \left\lfloor n \exp\left[-\left(1 - C_e\sqrt{\frac{1}{nI^*}}\right)\bar{n}I^*\right]\right\rfloor,$$
$$\operatorname{\texttt{err}}(\widehat{\boldsymbol{\sigma}}^{\text{sdp}}, \boldsymbol{\sigma}^*) \leq \exp\left[-\left(1 - C_e'\sqrt{\frac{1}{nI^*}}\right)\bar{n}I^*\right].$$

The proof is given in Appendices C and E. Note the floor operation in the first inequality above; consequently, we have $\|\widehat{\mathbf{Y}} - \mathbf{Y}^*\|_1 = 0$ whenever the exponent is strictly less than $-\log n$. We later explore the implication of this fact for exact recovery.

Remark 5: The assumption $c_0 p \leq q$ for Model 3 is common in the literature on minimax rates [23], [24], [28], [29]. It stipulates that p and q are on the same order (but their difference can be vanishingly small). This is the regime where the clustering problem is hard, and it is the regime we focus on. The assumption arises from a technical step in our proof, and it is currently not clear to us whether this assumption is necessary. We would like to point out that in Section 5.1 of [23], a weaker minimax upper bound is obtained with this assumption dropped.

Letting $n \to \infty$ in Theorem 4, we immediately obtain the following asymptotic result.

Corollary 6 (Upper bound, asymptotic): For any constants $c_0, c_1 \in (0,1)$, the following holds for Model 1, Model 2, and Model 3 with $0 < c_0 p \leq q < p \leq 1 - c_1$. If $nI^* \to \infty$, then with probability $1 - o(1)$ and for any $\widehat{\mathbf{Y}} \in \mathcal{Y}(\mathbf{A})$, we have

$$\frac{1}{n}\|\widehat{\mathbf{Y}} - \mathbf{Y}^*\|_1 \leq \left\lfloor n \exp\left[-\big(1 - o(1)\big)\bar{n}I^*\right]\right\rfloor,$$
$$\operatorname{\texttt{err}}(\widehat{\boldsymbol{\sigma}}^{\text{sdp}}, \boldsymbol{\sigma}^*) \leq \exp\left[-\big(1 - o(1)\big)\bar{n}I^*\right].$$

Comparing the upper bound in Corollary 6 with the minimax lower bound in Theorem 3, we see that the SDP achieves the optimal error rate, up to a second-order $o(1)$ term in the exponent.[7] Moreover, Theorem 4 provides an explicit, non-asymptotic upper bound for the $o(1)$ term in the exponent. This bound, taking the form of $O(1/\sqrt{nI^*})$, yields second-order characterization of various recovery thresholds and is strong enough to provide non-trivial guarantees in the sparse graph regime—these points are discussed in Section IV-D to follow. We do not expect this $O(1/\sqrt{nI^*})$ bound to be information-theoretic optimal, for reasons discussed in Section I-A.

As a passing note, the above error upper bounds also apply to the MLE solution $\widehat{\boldsymbol{\sigma}}^{\text{mle}}$, since the optimality of $\widehat{\boldsymbol{\sigma}}^{\text{mle}}$ to the program (6) implies that $(\widehat{\boldsymbol{\sigma}}^{\text{mle}})(\widehat{\boldsymbol{\sigma}}^{\text{mle}})^\top \in \mathcal{Y}(\mathbf{A})$. In fact, our proof of the upper bounds involves showing that the SDP solutions closely approximate the MLE; we elaborate on this point in Section V.

C. Robustness Under Semirandom Models

Our next result shows that the error rate of the SDP is unaffected by passing to the semirandom model. Recall the definition in Equation (10), so $\mathcal{Y}(\mathbf{A}^{\text{SR}})$ is the sublevel set of the SDP (8) with $\mathbf{A}^{\text{SR}}$ as the input.

Theorem 7 (Semirandom SBM): Suppose that $\mathbf{A}^{\text{SR}}$ is generated according to Model 4. The conclusions of Theorem 4 and Corollary 6 continue to hold for the program (8) with $\mathcal{Y}(\mathbf{A})$ replaced by $\mathcal{Y}(\mathbf{A}^{\text{SR}})$.

Proof: This theorem admits a short proof, thanks to the validity of Theorem 4 for any $\widehat{\mathbf{Y}} \in \mathcal{Y}(\mathbf{A})$. Recall that $\mathbf{A}^{\text{SR}}$ is obtained by monotonically modifying the matrix $\mathbf{A}$ generated from Model 3 (binary symmetric SBM). Let $\widehat{\mathbf{Y}}$ be an arbitrary element of $\mathcal{Y}(\mathbf{A}^{\text{SR}})$. By definition of $\mathbf{A}^{\text{SR}}$ and the feasibility of $\widehat{\mathbf{Y}}$, we have for all $i, j \in [n]$

$$\begin{cases} A_{ij}^{\text{SR}} \geq A_{ij}, \ \widehat{Y}_{ij} - Y_{ij}^* \leq 0, & \text{if } Y_{ij}^* = 1, \\ A_{ij}^{\text{SR}} \leq A_{ij}, \ \widehat{Y}_{ij} - Y_{ij}^* \geq 0, & \text{if } Y_{ij}^* = -1. \end{cases}$$

The fact $\widehat{\mathbf{Y}}$ has objective value no worse than $\mathbf{Y}^*$ under $\mathbf{A}^{\text{SR}}$, together with the above inequalities, implies that $0 \leq \langle \mathbf{A}^{\text{SR}}, \widehat{\mathbf{Y}} - \mathbf{Y}^* \rangle \leq \langle \mathbf{A}, \widehat{\mathbf{Y}} - \mathbf{Y}^* \rangle$. This further implies that $\widehat{\mathbf{Y}} \in \mathcal{Y}(\mathbf{A})$. Therefore, invoking Theorem 4 gives the desired result. ■

[7] We note that Theorem 3 bounds the error in expectation and holds for a parameter space containing $\boldsymbol{\sigma}^*$ with slightly unequal-sized clusters. The results in Theorem 4 and Corollary 6 are high-probability bounds. Extending these upper bounds to the setting of slightly unequal-sized clusters is possible, albeit tedious; we leave this to future work.

Remark 8: As mentioned in Remark 2, one may define semirandom versions of Models 1 (Z2) and 2 (CBM). It is easy to see that the proof above applies to these models without change. Therefore, the SDP approach is also robust under semirandom Z2 and CBM.

As an immediate consequence of Theorem 7, we obtain error bounds for a generalization of the standard SBM (Model 3) with heterogeneous edge probabilities, where

$$A_{ij} \overset{\text{i.i.d.}}{\sim} \begin{cases} \mathrm{Ber}(p_{ij}) & \text{if } \sigma_i^* \sigma_j^* = 1, \\ \mathrm{Ber}(q_{ij}) & \text{if } \sigma_i^* \sigma_j^* = -1, \end{cases}$$

with $0 \leq q_{ij} \leq q < p \leq p_{ij} \leq 1$.

Corollary 9 (*Heterogeneous SBM):* Under the above generalization of Model 3, the conclusions of Theorem 4 and Corollary 6 continue to hold for the program (8).

Proof: The corollary follows from the same coupling argument as in [3, Appendix V], which shows that the Heterogeneous SBM can be reduced to the semirandom Model 4. ■

The results above show that SDP is insensitive to monotone modification and heterogeneous probabilities. We emphasize that such robustness is by no means automatic. With non-uniformity in the probabilities, the likelihood function no longer has a known, rigid form, a property heavily utilized in many algorithms. The monotone adversary can similarly alter the graph structure by creating hotspots and short cycles. Even worse, the adversary is allowed to make changes *after* observing the realized graph,[8] thus producing unspecified dependency among all edges in the observed data and leading to major obstacles for existing analysis of iterative algorithms.

We would like to mention that the work in [6] shows that the semirandom model makes weak recovery strictly harder. While not contradicting their results technically, the fact that our error bounds remain unaffected under this model does demand a closer look. We note that our bounds are optimal only up to a second-order term in the exponent and consequently do not attain the optimal weak recovery limit. Also, our robustness results on error rates are tied to a specific form of SDP analysis (using the sublevel set $\mathcal{Y}(\mathbf{A})$). In comparison, for exact recovery SDP is robust *by design* to the semirandom model, as is well recognized in past work [5], [8], [56].

D. Consequences

Theorem 4 and Corollary 6 imply sharp sufficient conditions for several types of recovery:

- **Exact recovery:** Whenever $\bar{n}I^* \geq (1+\delta)\log n$ for any constant $\delta > 0$, we have $\|\widehat{\mathbf{Y}} - \mathbf{Y}^*\|_1 = 0$ by Corollary 6 (note the floor operation therein) and hence SDP achieves exact recovery by itself without any rounding/post-processing steps.
- **Second-order refinement:** Using the non-asymptotic Theorem 4, we can obtain the following refinement of the above result: exact recovery provided that $\frac{\bar{n}I^*}{\log n} \geq 1 + \frac{C_1}{\sqrt{\log n}}$ for some constant $C_1 > 0$.
- **Weak recovery:** When $\bar{n}I^* \geq C$ for a sufficiently large constant C, Theorem 4 ensures that $\mathtt{err}(\widehat{\boldsymbol{\sigma}}^{\text{sdp}}, \boldsymbol{\sigma}^*) < \frac{1}{2}$ and hence SDP achieves weak recovery.
- **Sparse regime:** Theorem 4 ensures that SDP achieves an arbitrarily small constant error when nI^* is a sufficiently large but finite constant. This corresponds to the sparse graph regime with constant expected degrees, namely $\alpha, p, q = \Theta(1/n)$ in CBM and SBM. Many results on minimax rates require nI^*, and hence the degrees, to diverge (e.g., [23], [28]).

Moreover, these conditions remain sufficient under the semirandom model. Below we specialize the above results to each of the three models.

1) $\mathbb{Z}_2$ Synchronization: Recall that $I^* := \frac{1}{2\tau^2}$ and $\bar{n} = n$ under Model 1. Consequently, SDP achieves exact recovery if $\tau^2 \leq \frac{n}{2\log n + C\sqrt{\log n}}$. This is a refinement of the best existing threshold $\tau^2 \leq \frac{n}{(2+\delta)\log n}$ in [9], [14].

We also have weak recovery by SDP if $\tau^2 \leq \frac{n}{C}$, which matches, up to constants, the optimal threshold $\tau^2 < n$ established in [15], [16].

2) Censored Block Model: Recall that $I^* := \alpha\big(\sqrt{1-\epsilon} - \sqrt{\epsilon}\big)^2$ and $\bar{n} = n$ under Model 2. Consequently, SDP achieves exact recovery if $\frac{n}{\log n} I^* \geq 1 + \frac{C}{\sqrt{\log n}}$. This result is a second-order improvement over the threshold $\frac{n}{\log n} I^* \geq 1 + \delta$ for SDP established in the work [50]. The same work also proves that exact recovery is impossible if $\frac{n}{\log n} I^* < 1 - \delta$.

Noting that $I^* \asymp \alpha(1-2\epsilon)^2$ (cf. Fact 10(b) in Appendix A), we also have weak recovery by SDP if $n\alpha(1-2\epsilon)^2 \geq C$, which matches, up to constants, the optimal threshold $n\alpha(1-2\epsilon)^2 > 1$ proved in [19], [20], [22].

3) Stochastic Block Model: Recall that $I^* := -2\log\big[\sqrt{pq} + \sqrt{(1-p)(1-q)}\big]$ and $\bar{n} = \frac{n}{2}$ under Model 3, and note the equivalence $I^* = (1+o(1))\big(\sqrt{p} - \sqrt{q}\big)^2$ valid for $0 < q \asymp p = o(1)$. Consequently, SDP achieves exact recovery if $n(\sqrt{p} - \sqrt{q})^2 \geq (2+\delta)\log n$, recovering the result established in [8], [9].

We also have the following refinement: exact recovery provided that $\frac{nI^*}{\log n} \geq 2 + \frac{C_1}{\sqrt{\log n}}$. This result is comparable to the sufficient condition $\frac{n(\sqrt{p}-\sqrt{q})^2}{\log n} \geq 2 + \frac{C}{\sqrt{\log n}} + \omega\left(\frac{1}{\log n}\right)$ for SDP established in [8], whereas the necessary and sufficient condition for the optimal estimator (MLE) is $\frac{n(\sqrt{p}-\sqrt{q})^2}{\log n} \geq 2 - \frac{\log\log n}{\log n} + \omega\left(\frac{1}{\log n}\right)$ [31], [32].

Finally, noting that $I^* \asymp (p-q)^2/p$ (cf. Fact 11(b) in Appendix A), we have weak recovery by SDP if $n(p-q)^2/p \geq C$. This condition matches, up to constants, the so-called Kesten-Stigum (KS) threshold $n(p-q)^2/(p+q) > 2$, which is optimal [37]–[39], [57].

E. Comparison With Existing Results

In this section we focus on partial recovery under the binary symmetric SBM (Model 3), and compare with the existing work that derives sharp error rate bounds achievable by polynomial-time algorithms. To be clear, the algorithms considered in this line of work are very different from ours. In particular, most existing results require a good enough

[8]We therefore strengthen the robustness results in the previous work [3], which does not allow such adaptivity.

initial estimate of the true clusters. Obtaining such an initial solution (typically using spectral clustering) is itself a non-trivial task.

Using neighbor voting and variational inference algorithms, the work in [23], [28] obtains an error bound of the same form as our Corollary 6, though they do not provide non-asymptotic results as in our Theorem 4. The work in [26] considers a spectral algorithm and proves the error bound $\mathtt{err}(\widehat{\boldsymbol{\sigma}},\boldsymbol{\sigma}^*) \le \exp\left[-(1-\delta)(\sqrt{p}-\sqrt{q})^2\cdot n/2\right]$ for any constant $\delta > 0$ if $np \to \infty$. Recalling $I^* = (1+o(1))(\sqrt{p}-\sqrt{q})^2$, we find that our Corollary 6 is better as we allow the δ term to vanish. The recent work in [14] uses a novel perturbation analysis to show that a very simple spectral algorithm achieves the error bound in Corollary 6 under the assumption $\frac{n}{\log n}(\sqrt{p}-\sqrt{q})^2 \ge \delta'$ for any constant $\delta' > 0$; their assumption excludes the sparse regime with $p, q = o\left(\frac{\log n}{n}\right)$ and is stronger than our assumption $nI^* \to \infty$ in Corollary 6. Compared to the above works, another strength of our results is that we provide an explicit bound for the second-order term in the exponent; we know of few error rate results (with the exception discussed below) that offer this level of accuracy.

Concurrently to our work, the paper [29] establishes a tight non-asymptotic error bound for an EM-type algorithm. Translated to our notation, their bound takes the form

$$\mathtt{err}(\widehat{\boldsymbol{\sigma}},\boldsymbol{\sigma}^*) \le \exp\left[-\left(1+\frac{2}{nI^*}\log(np)\right)\frac{nI^*}{2}\right],$$

which is valid under the assumption $nI^* \gtrsim \sqrt{np} \gtrsim 1$. Their assumption is order-wise more restrictive than that in our Theorem 4, but their error bound has a better second-order term in the exponent. We do note that their algorithm is fairly technical: it requires data partition and the leave-one-out tricks to ensure independence, degree truncation to regularize spectral clustering, and blackbox solvers for K-means and matching problems. In comparison, the SDP approach is much simpler conceptually.

Finally, we emphasize that we also provide robustness guarantees under the monotone semirandom model and non-uniform edge probabilities. In comparison, it is unclear if comparable robustness results can be established for the algorithms above, as these algorithms and their analyses make substantial use of the properties of the standard SBM, particularly the complete independence among edges and the specific form of the likelihood function.

V. Proof Outline

In this section we outline our proofs of the lower and upper bounds. In the process we provide insights on how the error rate $e^{-\bar{n}I^*}$ arises and why SDP achieves it.

A. Proof Outline of Theorem 3

The intuition behind the lower bound is relatively easy to describe. To illustrate the idea, take as an example the Z2 model, where $A_{ij} \overset{\text{i.i.d.}}{\sim} N(1,\tau^2)$, and assume that $\sigma_i^* = 1, \forall i$. It is not hard to see that the error fraction $\mathtt{err}(\widehat{\boldsymbol{\sigma}},\boldsymbol{\sigma}^*)$ for any estimator $\widehat{\boldsymbol{\sigma}}$ is lower bounded by the probability of recovering the label σ_1^* for the first node given true labels of the other nodes; see the Lemma 12 for the precise argument. For this one-node problem, the optimal Bayes estimate of σ_1^* is given by the sign of the majority vote $\sum_{j=1}^n A_{1j}$:

$$\widehat{\sigma}_1 = \underset{\sigma\in\{\pm1\}}{\arg\max}\left\{\sigma\cdot\sum\nolimits_{j=1}^n A_{1j}\right\} = \operatorname{sign}\left(\sum\nolimits_{j=1}^n A_{1j}\right); \tag{11}$$

see Lemma 13. It follows that the error probability of recovering $\sigma_1^* = 1$ is

$$\begin{aligned}\mathbb{P}\Big\{\widehat{\sigma}_1 \ne 1\Big\} &= \mathbb{P}\left\{\sum\nolimits_{j=1}^n A_{1j} < 0\right\}\\ &= \exp\Big[-(1+o(1))\cdot nI(0)\Big],\end{aligned}$$

where the last step can be justified in general by the large deviation theory, with $I(x) := \sup_{t>0}[tx - \log \mathbb{E}e^{t(-A_{11})}]$ being the rate function (see, e.g., Cramer's Theorem [58, Theorem 2.2.3]). In our setting, a direct calculation suffices, as is done in Lemma 14. The error exponent

$$I(0) = -\inf_{t>0}\left\{\log\mathbb{E}e^{t(-A_{11})}\right\} = I^* \tag{12}$$

is precisely our SNR measure, a quantity we will encounter again in proving the upper bound.

The above intuition remains valid for CBM and SBM, though the specific forms of the majority voting procedure and the rate I^* vary. The complete proof is given in Appendix B.

B. Proof Outline of Theorem 4

To prove the upper bound for SDP, we proceed in three steps:

Step 1: As mentioned in Section I, we construct a diagonal matrix $\mathbf{D}$ with $D_{ii} = \sigma_i^*\sum_j A_{ij}\sigma_j^*$, which takes the same form as the "dual certificate" used in previous work. The construction of $\mathbf{D}$ allows us to establish the *basic inequality*:

$$0 \le \left\langle -\mathbf{D}, \mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}})\right\rangle + \left\langle \mathbf{A}-\mathbb{E}\mathbf{A}, \mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}})\right\rangle,$$

for any $\widehat{\mathbf{Y}} \in \mathcal{Y}(\mathbf{A})$; see the proof of Lemma 19 for the details of this critical step. Here $\mathcal{P}_{T^\perp}$ is an appropriate projection operator that satisfies $\operatorname{Tr}\left[\mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}})\right] = \frac{1}{n}\|\widehat{\mathbf{Y}}-\mathbf{Y}^*\|_1$, thus exposing the ℓ_1 error of $\widehat{\mathbf{Y}}$ that we seek to control.

Step 2: We proceed by showing that the second term $S_2 := \langle \mathbf{A}-\mathbb{E}\mathbf{A}, \mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}})\rangle$ in the basic inequality is negligible compared to the first term $\langle -\mathbf{D}, \mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}})\rangle$; see Proposition 20 for a quantitative version of this claim, whose proof involves certain trimming argument in the case of CBM and SBM. Dropping S_2 from the basic inequality hence yields

$$0 \le \langle -\mathbf{D}, \mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}})\rangle = \sum\nolimits_{i=1}^n(-D_{ii})b_i, \tag{13}$$

where $b_i := \big(\mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}})\big)_{ii}$ satisfies $\sum_{i\in[n]} b_i = m := \frac{1}{n}\|\widehat{\mathbf{Y}}-\mathbf{Y}^*\|_1$.

Step 3: The b_i's take fractional values in general, but must be bounded in $[0,4]$ (cf. 18). We use this fact to upper bound the RHS of (13) by its worst-case value, hence obtaining

$$0 \le \sum\nolimits_{i=1}^n(-D_{ii})b_i \lesssim \max_{\substack{\mathcal{M}\subseteq[n]\\|\mathcal{M}|=m}}\sum\nolimits_{i\in\mathcal{M}}(-D_{ii}). \tag{14}$$

This argument is reminiscent of the "order statistics" analysis in [3], though here we provide a more fine-grained bound; see Appendix C-D for details. The rest of this step, done in Lemma 23, establishes a probabilistic bound for the RHS of (14) and ultimately gives rise to the error exponent $-I^*$. To illustrate the idea, we again consider Z2 with $\sigma_i^* \equiv 1$, in which case $D_{ii} = \sum_{j\in[n]} A_{ij}$. For a fixed set $\mathcal{M}$ with $|\mathcal{M}| = m$, the RHS of (14) can be controlled using the Chernoff bound:

$$\begin{aligned}&\mathbb{P}\left\{\sum\nolimits_{i\in\mathcal{M}}\left(\sum\nolimits_{j=1}^{n}(-A_{ij})\right)>0\right\}\\&\le\inf_{t>0}\mathbb{E}\exp\left[t\sum\nolimits_{i\in\mathcal{M}}\left(\sum\nolimits_{j=1}^{n}(-A_{ij})\right)\right]\\&=\left(\inf_{t>0}\mathbb{E}\exp(-tA_{11})\right)^{nm}=e^{-nmI^*},\end{aligned}$$

where the last two steps follow from independence and the expression (12) for I^*. In the first equality above, we ignore the dependence between symmetric entries of A for now and note that a more careful calculation is able to deal with the dependency (see Appendix D-G). By a union bound over all $\binom{n}{m} \approx \left(\frac{n}{m}\right)^m$ such $\mathcal{M}$'s, we obtain

$$\begin{aligned}&\mathbb{P}\left\{\max_{\substack{\mathcal{M}\subseteq[n]\\|\mathcal{M}|=m}}\sum\nolimits_{i\in\mathcal{M}}\left(\sum\nolimits_{j=1}^{n}(-A_{ij})\right)>0\right\}\\&\le\left(\frac{n}{m}\right)^m\cdot e^{-nmI^*}=\left(e^{\log(n/m)-nI^*}\right)^m.\end{aligned}$$

If $\log(n/m)-nI^* < 0$, then RHS above is $\ll 1$ and thus with high probability the negation of (14) holds, a contradiction. We therefore must have $\log(n/m) - nI^* \ge 0$, which implies the desired error bound $\frac{m}{n} \le e^{-nI^*}$. The second-order term in the exponent comes from a more accurate calculation for Steps 2 and 3.

The previous arguments are closely connected to our proof for the lower bound outlined above. Note that the MLE of the entire vector $\boldsymbol{\sigma}^* = \mathbf{1}$ is given by the "joint majority voting" procedure

$$\widehat{\boldsymbol{\sigma}}^{\mathrm{mle}} = \arg\max_{\boldsymbol{\sigma}\in\{\pm1\}^n}\left\{\sum\nolimits_{i=1}^{n}\sigma_i\cdot\left(\sum\nolimits_{j=1}^{n}A_{ij}\right)\right\};$$

one should compare this equation with the "single-node majority voting" in (11). The maximality of the above $\widehat{\boldsymbol{\sigma}}^{\mathrm{mle}}$ over $\boldsymbol{\sigma}^* = \mathbf{1}$, as well as the fact that $1-\widehat{\sigma}_i^{\mathrm{mle}} \in \{0, 2\}$, implies that

$$\begin{aligned}0&\le\sum\nolimits_{i=1}^{n}\left(1-\widehat{\sigma}_i^{\mathrm{mle}}\right)\cdot\left(\sum\nolimits_{j=1}^{n}(-A_{ij})\right)\\&\lesssim\max_{\substack{\mathcal{M}\subseteq[n]\\|\mathcal{M}|=m}}\sum_{i\in\mathcal{M}}(-D_{ii}),\end{aligned}$$

if we set $m = \frac{1}{2}\sum_{i=1}^{n}(1-\widehat{\sigma}_i^{\mathrm{mle}}) = n\cdot\mathtt{err}(\widehat{\boldsymbol{\sigma}}^{\mathrm{mle}}, \boldsymbol{\sigma}^*)$. Note that this inequality is the same as (14), so following the arguments above shows that the MLE satisfies the same error bound $\frac{m}{n} \le e^{-nI^*}$. We therefore see that the SDP solution closely approximates the MLE in the above precise sense, and both of them achieve the Bayes rate. The form of the rate e^{-nI^*} arises from a majority voting mechanism, in the proofs for both lower and upper bounds.

Again, the above intuition remains valid for CBM and SBM, though the calculation of the rate I^* varies. The details of the proof are given in Appendices C and E.

VI. Discussion

In this paper, we analyze the error rates of the SDP relaxation approach for clustering under several random graph models, namely Z2, CBM and the binary symmetric SBM, via a unified framework. We show that SDP achieves an exponentially-decaying error with a sharp exponent, matching the minimax lower bound for all three models. We also show that these results continue to hold under monotone semirandom models, demonstrating the robustness of SDP.

Immediate future directions include extensions to problems with multiple and unbalanced clusters, as well as to closely related models such as weighted SBM. It is also of interest to see if better estimates of the second order term can be obtained, and if there is a fundamental tradeoff between statistical optimality and robustness. More broadly, it would be interesting to explore the applications of the techniques in this paper in analyzing SDP relaxations for other discrete problems.

Appendix A
Preliminaries

In this section we record several notations and facts that are useful for subsequent proofs.

We first define a random variable H that encapsulates the distributions of the three models:

- For Model 1 (Z2), let $H \sim N(1, \tau^2)$.
- For Model 2 (CBM), let H have probability mass function $\alpha(1-\epsilon)\cdot\delta_1 + \alpha\epsilon\cdot\delta_{-1} + (1-\alpha)\cdot\delta_0$, where δ_a denotes the Dirac delta function centered at a.
- For Model 3 (SBM), let $H = Y - Z$, where $Y \sim \mathrm{Ber}(p)$, $Z \sim \mathrm{Ber}(q)$, and Y, Z are independent.

It can be seen that under Model 1 or 2, we have $A_{ij} \sim \sigma_i^*\sigma_j^* H$ (here $\sim$ means equality in distribution); under Model 3, we have $A_{ii'} - A_{ij} \sim H$ if $\sigma_i^* = \sigma_{i'}^* = -\sigma_j^*$.

Let t^* be the minimizer of the moment generating function $t \mapsto \mathbb{E}e^{-tH}$, which has the explicit expression

$$t^* := \begin{cases}\frac{1}{\tau^2}, & \text{for Model 1},\\ \frac{1}{2}\log\frac{1-\epsilon}{\epsilon}, & \text{for Model 2},\\ \frac{1}{2}\log\frac{p(1-q)}{q(1-p)}, & \text{for Model 3}.\end{cases} \tag{15}$$

Note that $t^* > 0$. We later verify that $\mathbb{E}e^{-tH} \approx e^{-I^*}$ for all three models (see Facts 35, 36 and 37). Also define the quantity

$$\lambda^* := \begin{cases}0 & \text{for Model 1 and 2},\\ \frac{1}{2t^*}\log\frac{1-q}{1-p}, & \text{for Model 3},\end{cases} \tag{16}$$

which plays a role only in Model 3 (SBM).

Finally, for Model 2, we let $p := \alpha(1-\epsilon)$ and $q := \alpha\epsilon$; this notation is chosen to bring out the similarity between Models 2 and 3.

We record several simple estimates for the above quantities t^* and λ^* as well as the SNR measure I^* defined in (5). The proofs are given in Appendices A-A and A-B to follow.

Fact 10: Under Model 2 with the notation $p := \alpha(1-\epsilon)$ and $q := \alpha\epsilon$, if $0 < q \le p \le 1$, then

(a) $t^* \le \frac{1-\epsilon}{2\epsilon} \cdot \frac{p-q}{p}$.

(b) $I^* \in \left[\frac{(p-q)^2}{4p}, \frac{(p-q)^2}{p}\right]$.

Fact 11: Under Model 3, if $0 < q < p < 1$, then the following hold.

(a) $\lambda^* \in (q, p)$.

(b) If in addition $p \le 1-c$ for some constant $c \in (0,1)$, then $I^* \asymp \frac{(p-q)^2}{p}$.

(c) If in addition $p \le 1-c$ and $q \ge c_0 p$ for some constants $c \in (0,1)$, $c_0 \in (0,1)$, then $t^* \lesssim \frac{p-q}{p}$.

A. Proof of Fact 10

Recall the shorthands $p := \alpha(1-\epsilon)$ and $q := \alpha\epsilon$ introduced for Model 2. For part (a) of the fact, by definition of t^* in Equation (15), we have

$$t^* = \frac{1}{2}\log\left(1 + \frac{p-q}{q}\right) \overset{(i)}{\le} \frac{p-q}{2q} \overset{(ii)}{=} \frac{1-\epsilon}{2\epsilon} \cdot \frac{p-q}{p},$$

where step (i) holds since the fact that $1 + x \le e^x$ for $x \in \mathbb{R}$ implies $\log(1+x) \le x$ for $x > -1$, and step (ii) holds by the fact that $q = \frac{\epsilon}{1-\epsilon}p$.

For part (b), recalling the definition of I^* in Equation (5), we have

$$I^* = \frac{\left(\sqrt{p} - \sqrt{q}\right)^2\left(\sqrt{p} + \sqrt{q}\right)^2}{\left(\sqrt{p} + \sqrt{q}\right)^2} = \frac{(p-q)^2}{p + q + 2\sqrt{pq}}.$$

Some algebra shows that $I^* \le \frac{(p-q)^2}{p}$ and $I^* \ge \frac{(p-q)^2}{p+p+2p} = \frac{(p-q)^2}{4p}$.

B. Proof of Fact 11

For part (a), recalling the definition of λ^* in Equation (16), we obtain by direct calculation the identity

$$p - \lambda^* = \left[\log\frac{p(1-q)}{q(1-p)}\right]^{-1}\left[p\log\frac{p}{q} + (1-p)\log\frac{1-p}{1-q}\right].$$

The quantity inside the second bracket on the RHS is positive, as it is the KL divergence between $\mathrm{Ber}(p)$ and $\mathrm{Ber}(q)$ with $p \ne q$. We also have $\log\frac{p(1-q)}{q(1-p)} > 0$ since $0 < q < p < 1$. It follows that $p - \lambda^* > 0$ as claimed. A similar argument shows that $\lambda^* - q > 0$.

Part (b) is a partial result of [4, Lemma B.1].

For part (c), recall the definition $t^* := \frac{1}{2}\left(\log\frac{p}{q} + \log\frac{1-q}{1-p}\right)$ in Eq. (15). We consider two cases. If $\frac{p}{q} \ge \frac{1-q}{1-p}$, we have

$$t^* \le \log\frac{p}{q} \overset{(i)}{\le} \frac{p}{q} - 1 \overset{(ii)}{\le} \frac{p-q}{c_0 p},$$

where step (i) holds since $\log(x) \le x - 1, \forall x > 0$, and step (ii) holds by assumption $c_0 p \le q$. If $\frac{p}{q} \le \frac{1-q}{1-p}$, we have

$$t^* \le \log\frac{1-q}{1-p} \le \frac{1-q}{1-p} - 1 \overset{(i)}{\le} \frac{p-q}{c} \le \frac{p-q}{cp},$$

where step (i) holds by the assumption that $p \le 1-c$. In both cases, we have $t^* \lesssim \frac{p-q}{p}$ as claimed.

Appendix B
Proof of Theorem 3

In this section we prove Theorem 3 under Models 1 and 2, following a similar strategy as in the proof of [4, Theorem 1.1]. We make use of the definitions and facts given in Appendix A.

For simplicity, in the sequel we write $\Theta \equiv \Theta(n) := \{\pm 1\}^n$. Let ϕ be the uniform prior over all the elements in Θ. Define the global Bayesian risk

$$B_\phi\big(\Theta, \widehat{\boldsymbol{\sigma}}\big) := \frac{1}{|\Theta|}\sum_{\boldsymbol{\sigma}\in\Theta}\mathbb{E}_{\boldsymbol{\sigma}}\,\mathtt{err}(\widehat{\boldsymbol{\sigma}}, \boldsymbol{\sigma}),$$

and the local Bayesian risk for the first node

$$B_\phi\big(\Theta, \widehat{\sigma}(1)\big) := \frac{1}{|\Theta|}\sum_{\boldsymbol{\sigma}\in\Theta}\mathbb{E}_{\boldsymbol{\sigma}}\,\mathtt{err}\left(\widehat{\sigma}(1), \sigma(1)\right).$$

In the above, the quantity $\mathtt{err}(\widehat{\sigma}(1), \sigma(1))$ denotes the loss on the first node, defined as

$$\mathtt{err}\left(\widehat{\sigma}(1), \sigma(1)\right) := \frac{1}{|\mathcal{S}_{\boldsymbol{\sigma}}(\widehat{\boldsymbol{\sigma}})|}\sum_{\boldsymbol{\sigma}'\in\mathcal{S}_{\boldsymbol{\sigma}}(\widehat{\boldsymbol{\sigma}})}\mathbb{I}\left\{\sigma'(1) \ne \sigma(1)\right\},$$

where

$$\mathcal{S}_{\boldsymbol{\sigma}}(\widehat{\boldsymbol{\sigma}}) := \Big\{ g\widehat{\boldsymbol{\sigma}} : g \in \{\pm 1\}, \\ \frac{1}{n}\sum_{i\in[n]}\mathbb{I}\{g\widehat{\sigma}_i \ne \sigma_i\} = \mathtt{err}(\widehat{\boldsymbol{\sigma}}, \boldsymbol{\sigma})\Big\}.$$

The following lemma shows that these risks are equal.

Lemma 12: Under Models 1 and 2, we have

$$\inf_{\widehat{\boldsymbol{\sigma}}} B_\phi\big(\Theta, \widehat{\boldsymbol{\sigma}}\big) = \inf_{\widehat{\boldsymbol{\sigma}}} B_\phi\big(\Theta, \widehat{\sigma}(1)\big).$$

Proof: The lemma essentially follows from the symmetry/exchangeability property of Models 1 (Z2) and 2 (CBM). Rigorous proof of this intuitive result is however quite technical, as the definition of clustering error involves a global sign flipping. Fortunately, most of the work has been done in [4]. In particular, note that the parameter space Θ is closed under permutation in the sense that for any label vector $\boldsymbol{\sigma} \in \Theta$ and any permutation π on $[n]$, the new label vector $\boldsymbol{\sigma}'$ defined by $\sigma'(i) := \sigma(\pi^{-1}(i))$ also belongs to Θ. It can also be seen that both Models 1 (Z2) and 2 (CBM) are homogeneous, i.e., the distribution of each A_{ij} is uniquely determined by the sign of $\sigma_i^*\sigma_j^*$. Consequently, for Model 2 (CBM) this lemma immediately follows from Lemma 2.1 in [4], as its proof applies without change. For Model 1 (Z2) in which the distribution of $\mathbf{A}$ is continuous, we note that the proof of Lemma 2.1 in [4] continues to hold when summations therein are replaced by appropriate integrations. ■

With the above lemma, it suffices to lower bound the local Bayes risk. This task can be further reduced to computing the tail probability of a certain sum of independent copies of the random variable H defined in Appendix A. This is done in the following lemma.

Lemma 13: Let ϕ be the uniform prior over all elements in Θ. Under Models 1 and 2, we have

$$B_\phi\big(\Theta, \widehat{\sigma}(1)\big) \geq \mathbb{P}\left(\sum_{i\in[n-1]} Z_i \geq 0\right),$$

where $\{Z_i\}$ are i.i.d. copies of $-H$.

This lemma is analogous to Lemma 5.1 in [4]. We provide the proof in Appendix B-A.

Finally, the lemma below provides an explicit lower bound of the above tail probability in terms of the SNR measure I^*.

Lemma 14: Let $\{Z_i\}$ be i.i.d. copies of $-H$. For Model 1 and Model 2 with $I^* = o(1)$, if $nI^* \to \infty$, then there exists $\xi = o(1)$ such that

$$\mathbb{P}\left(\frac{1}{n-1}\sum_{i\in[n-1]} Z_i \geq 0\right) \geq \exp\left[-\left(1+\xi\right)(n-1)I^*\right].$$

This lemma is analogous to Lemma 5.2 in [4]. We provide the proof in Appendix B-B for Model 1 (Z2) and in Appendix B-C for Model 2 (CBM).

We are now ready to prove Theorem 3. Note that

$$\inf_{\widehat{\boldsymbol{\sigma}}} \sup_{\boldsymbol{\sigma}\in\Theta} \mathbb{E}_{\boldsymbol{\sigma}} \, \mathtt{err}(\widehat{\boldsymbol{\sigma}}, \boldsymbol{\sigma}) \geq \inf_{\widehat{\boldsymbol{\sigma}}} B_\phi(\Theta, \widehat{\boldsymbol{\sigma}}),$$

since the Bayes risk lower bounds the minimax risk. To complete the proof, we continue the above inequality by successively invoking Lemmas 12, 13 and 14.

A. Proof of Lemma 13

Recall that $B_\phi(\Theta, \widehat{\sigma}(1))$ is defined as

$$B_\phi\big(\Theta, \widehat{\sigma}(1)\big) \coloneqq \frac{1}{|\Theta|}\sum_{\boldsymbol{\sigma}\in\Theta} \mathbb{E}_{\boldsymbol{\sigma}} \, \mathtt{err}\left(\sigma(1), \widehat{\sigma}(1)\right).$$

For each $\boldsymbol{\sigma}_0 \in \Theta$, we generate a new assignment $\boldsymbol{\sigma}[\boldsymbol{\sigma}_0]$ based on $\boldsymbol{\sigma}_0$ by setting $\sigma[\sigma_0](1) \coloneqq -\sigma_0(1)$ and $\sigma[\sigma_0](i) \coloneqq \sigma_0(i)$ for all $i \in [n]\setminus\{1\}$. It can be seen that $\boldsymbol{\sigma}[\boldsymbol{\sigma}_0] \in \Theta$ and the Hamming distance between $\boldsymbol{\sigma}_0$ and $\boldsymbol{\sigma}[\boldsymbol{\sigma}_0]$ is 1. In addition, for any $\boldsymbol{\sigma}_1, \boldsymbol{\sigma}_2 \in \Theta$ with $\boldsymbol{\sigma}_1 \neq \boldsymbol{\sigma}_2$, we have $\boldsymbol{\sigma}[\boldsymbol{\sigma}_1] \neq \boldsymbol{\sigma}[\boldsymbol{\sigma}_2]$. This bijection implies that $\{\boldsymbol{\sigma}[\boldsymbol{\sigma}_0] : \boldsymbol{\sigma}_0 \in \Theta\} = \Theta$. Consequently, continuing from the last displayed equation we obtain

$$\begin{aligned} B_\phi(\Theta, \widehat{\sigma}(1)) = \frac{1}{|\Theta|}\sum_{\boldsymbol{\sigma}_0\in\Theta}\frac{1}{2}\Big(&\mathbb{E}_{\boldsymbol{\sigma}_0} \, \mathtt{err}\left(\widehat{\sigma}(1), \sigma_0(1)\right) \\ &+ \mathbb{E}_{\boldsymbol{\sigma}[\boldsymbol{\sigma}_0]} \, \mathtt{err}\left(\widehat{\sigma}(1), \sigma[\sigma_0](1)\right)\Big), \end{aligned}$$

whence

$$\begin{aligned} \inf_{\widehat{\boldsymbol{\sigma}}} B_\phi(\Theta, \widehat{\sigma}(1)) \geq \frac{1}{|\Theta|}\sum_{\boldsymbol{\sigma}_0\in\Theta}\frac{1}{2}\inf_{\widehat{\boldsymbol{\sigma}}}\Big(&\mathbb{E}_{\boldsymbol{\sigma}_0} \, \mathtt{err}\left(\widehat{\sigma}(1), \sigma_0(1)\right) \\ &+ \mathbb{E}_{\boldsymbol{\sigma}[\boldsymbol{\sigma}_0]} \, \mathtt{err}\left(\widehat{\sigma}(1), \sigma[\sigma_0](1)\right)\Big). \quad (17) \end{aligned}$$

We proceed to compute the infimum above for a given $\boldsymbol{\sigma}_0 \in \Theta$. Let $\tilde{\boldsymbol{\sigma}}$ be the Bayes estimator that attains the infimum. Since $\boldsymbol{\sigma}_0$ and $\boldsymbol{\sigma}[\boldsymbol{\sigma}_0]$ only differ at the first node, we must have $\tilde{\sigma}(i) = \sigma[\sigma_0](i) = \sigma_0(i)$ for all $i \in [n]\setminus\{1\}$, and either $\tilde{\sigma}(1) = \sigma_0(1)$ or $\tilde{\sigma}(1) = \sigma[\sigma_0](1)$. Now the problem is reduced to a test between two distributions $\mathbb{P}_{\boldsymbol{\sigma}_0}$ and $\mathbb{P}_{\boldsymbol{\sigma}[\boldsymbol{\sigma}_0]}$. Since the prior ϕ is uniform, $\tilde{\sigma}(1)$ is given by the likelihood ratio test $\frac{\mathbb{P}_{\boldsymbol{\sigma}_0}(\mathbf{A})}{\mathbb{P}_{\boldsymbol{\sigma}[\boldsymbol{\sigma}_0]}(\mathbf{A})} \gtrless 1$. The following lemma gives the explicit form of this test. Here we let $J_0 \coloneqq \{u \in [n]\setminus\{1\} : \sigma_0(u) = \sigma_0(1)\}$ and $J_1 \coloneqq \{u \in [n] : \sigma_0(u) = \sigma[\sigma_0](1)\}$.[9]

Lemma 15: Let $\boldsymbol{\sigma}_0$ and $\boldsymbol{\sigma}[\boldsymbol{\sigma}_0]$ be defined as above. Under Models 1 and 2, we have that

$$\tilde{\sigma}(1) = \begin{cases} \sigma_0(1), & \text{if } \sum_{u\in J_0} A_{1u} \geq \sum_{u\in J_1} A_{1u}, \\ \sigma[\sigma_0](1), & \text{otherwise.} \end{cases}$$

The lemma follows from a routine calculation of the likelihood. We give proof in Appendix B-A.1 for Model 1 (Z2) and in Appendix B-A.2 for Model 2 (CBM).

From Lemma 15 we have

$$\begin{aligned} &\mathbb{E}_{\boldsymbol{\sigma}_0} \, \mathtt{err}\left(\tilde{\sigma}(1), \sigma_0(1)\right) \\ &\quad = \mathbb{P}_{\boldsymbol{\sigma}_0}\Big(\sum_{u\in J_0} A_{1u} < \sum_{u\in J_1} A_{1u}\Big), \qquad \text{and} \\ &\mathbb{E}_{\boldsymbol{\sigma}[\boldsymbol{\sigma}_0]} \, \mathtt{err}\left(\tilde{\sigma}(1), \sigma[\sigma_0](1)\right) \\ &\quad = \mathbb{P}_{\boldsymbol{\sigma}[\boldsymbol{\sigma}_0]}\Big(\sum_{u\in J_0} A_{1u} \geq \sum_{u\in J_1} A_{1u}\Big). \end{aligned}$$

Recalling the distribution of $\{A_{1u}\}$, the definition of H and that $Z_i \overset{\text{i.i.d.}}{\sim} -H$, we see that both probabilities above equal $\mathbb{P}\big(\sum_{i\in[n-1]} Z_i \geq 0\big)$. Combining with the bound (17), we obtain the desired inequality $\inf_{\widehat{\boldsymbol{\sigma}}} B_\phi(\Theta, \widehat{\sigma}(1)) \geq \mathbb{P}\big(\sum_{i\in[n-1]} Z_i \geq 0\big)$.

1) Proof of Lemma 15 for Model 1 (Z2): Since $\boldsymbol{\sigma}_0$ and $\boldsymbol{\sigma}[\boldsymbol{\sigma}_0]$ only differ at the first node, the likelihood ratio $\frac{\mathbb{P}_{\boldsymbol{\sigma}_0}(\mathbf{A})}{\mathbb{P}_{\boldsymbol{\sigma}[\boldsymbol{\sigma}_0]}(\mathbf{A})}$ only depends on the first row and column of $\mathbf{A}$. In particular, recalling that $\{A_{1u}\}$ are Gaussian under Model 1, we have

$$\begin{aligned} &\frac{\mathbb{P}_{\boldsymbol{\sigma}_0}(\mathbf{A})}{\mathbb{P}_{\boldsymbol{\sigma}[\boldsymbol{\sigma}_0]}(\mathbf{A})} \\ &= \frac{\prod_{u\in J_0}\exp\left[-\left(A_{1u}-1\right)^2/(2\tau^2)\right]}{\prod_{u\in J_0}\exp\left[-\left(A_{1u}+1\right)^2/(2\tau^2)\right]} \\ &\quad \times \frac{\prod_{u\in J_1}\exp\left[-\left(A_{1u}+1\right)^2/(2\tau^2)\right]}{\prod_{u\in J_1}\exp\left[-\left(A_{1u}-1\right)^2/(2\tau^2)\right]} \\ &= \frac{\exp\left[-\left(\sum_{u\in J_0}\left(A_{1u}-1\right)^2+\sum_{u\in J_1}\left(A_{1u}+1\right)^2\right)/(2\tau^2)\right]}{\exp\left[-\left(\sum_{u\in J_0}\left(A_{1u}+1\right)^2+\sum_{u\in J_1}\left(A_{1u}-1\right)^2\right)/(2\tau^2)\right]}. \end{aligned}$$

Some algebra shows that

$$\frac{\mathbb{P}_{\boldsymbol{\sigma}_0}(\mathbf{A})}{\mathbb{P}_{\boldsymbol{\sigma}[\boldsymbol{\sigma}_0]}(\mathbf{A})} \gtrless 1 \iff \sum_{u\in J_0} A_{1u} \gtrless \sum_{u\in J_1} A_{1u}.$$

The result follows from the fact that the likelihood ratio test is Bayes-optimal for binary hypotheses under a uniform prior.

[9] In Lemma 15 and its proof, we adopt the convention that $\sum_{u\in J} f(u) = 0$ and $\prod_{u\in J} f(u) = 1$ if $J = \emptyset$.

2) Proof of Lemma 15 for Model 2 (CBM): Similarly to the previous section, we recall the distribution of $\{A_{1u}\}$ under Model 2 to obtain

$$\frac{\mathbb{P}_{\boldsymbol{\sigma}_0}(\mathbf{A})}{\mathbb{P}_{\boldsymbol{\sigma}[\boldsymbol{\sigma}_0]}(\mathbf{A})} = \frac{\prod_{u\in J_0}[\alpha(1-\epsilon)]^{\mathbb{I}\{A_{1u}=1\}}[\alpha\epsilon]^{\mathbb{I}\{A_{1u}=-1\}}}{\prod_{u\in J_0}[\alpha(1-\epsilon)]^{\mathbb{I}\{A_{1u}=-1\}}[\alpha\epsilon]^{\mathbb{I}\{A_{1u}=1\}}} \times \frac{\prod_{u\in J_1}[\alpha(1-\epsilon)]^{\mathbb{I}\{A_{1u}=-1\}}[\alpha\epsilon]^{\mathbb{I}\{A_{1u}=1\}}}{\prod_{u\in J_1}[\alpha(1-\epsilon)]^{\mathbb{I}\{A_{1u}=1\}}[\alpha\epsilon]^{\mathbb{I}\{A_{1u}=-1\}}}.$$

Since $\epsilon \in (0, \frac{1}{2})$, some algebra shows that

$$\frac{\mathbb{P}_{\boldsymbol{\sigma}_0}(\mathbf{A})}{\mathbb{P}_{\boldsymbol{\sigma}[\boldsymbol{\sigma}_0]}(\mathbf{A})} \gtrless 1 \iff \sum_{u\in J_0} A_{1u} \gtrless \sum_{u\in J_1} A_{1u}.$$

The result follows from the fact that the likelihood ratio test is Bayes-optimal for binary hypotheses under a uniform prior.

B. Proof of Lemma 14 for Model 1 (Z2)

Let $n' := n-1$, $p(z)$ be the pdf of Z_1, and $M(t)$ be the moment generating function of Z_1. Since $Z_1 \sim -H \sim N(-1,\tau^2)$, we can compute $M(t) = \exp(-t + \frac{1}{2}t^2\tau^2)$. Recalling $t^* = \frac{1}{\tau^2}$ as defined in Equation (15) and the definition of I^* in Equation (5), we obtain

$$M(t^*) = \exp\left(-\frac{1}{2\tau^2}\right) = \exp(-I^*).$$

Let $\delta := (2nI^*)^{-\frac{1}{4}}$, $S_{n'} := \sum_{i\in[n']} Z_i$ and $S_{n'}(\mathbf{z}) := \sum_{i\in[n']} z_i$. We have

$$\begin{aligned}
&\mathbb{P}(S_{n'} \geq 0)\\
&\geq \int_{\{\mathbf{z}:S_{n'}(\mathbf{z})\in[0,n'\delta]\}} \prod_{i\in[n']} p(z_i)\mathrm{d}\mathbf{z}\\
&\geq \frac{(M(t^*))^{n'}}{\exp(n't^*\delta)} \int_{\{\mathbf{z}:S_{n'}(\mathbf{z})\in[0,n'\delta]\}} \prod_{i\in[n']} \frac{\exp(t^*z_i)\,p(z_i)}{M(t^*)}\mathrm{d}\mathbf{z},
\end{aligned}$$

where the last step holds since $\exp(n't^*\delta) \geq \exp\left(t^*\sum_{i\in[n']} z_i\right) = \prod_{i\in[n']}\exp(t^*z_i)$ given that $\sum_{i\in[n']} z_i = S_{n'}(\mathbf{z}) \leq n'\delta$. Let $q(w) := \frac{\exp(t^*w)p(w)}{M(t^*)}$ and we have

$$\mathbb{P}(S_{n'} \geq 0) \geq \exp(-n'I^*)\exp(-n't^*\delta) \times \int_{\{\mathbf{z}:S_{n'}(\mathbf{z})\in[0,n'\delta]\}} \prod_{i\in[n']} q(z_i)\mathrm{d}\mathbf{z}.$$

Note that $q(w)$ is a pdf since $\int_w q(w)\mathrm{d}w = 1$ and $q(w) \geq 0$ for any w. Let $W_1, W_2, \ldots, W_{n'}$ be i.i.d. random variables with pdf $q(w)$. We have

$$\begin{aligned}
&\mathbb{P}(S_{n'} \geq 0)\\
&\geq \exp(-n't^*\delta)\,\mathbb{P}\left(\frac{1}{n'}\sum_{i\in[n']} W_i \in [0,\delta]\right)\exp(-n'I^*)\\
&=: Q_1Q_2\exp(-n'I^*). \qquad (18)
\end{aligned}$$

1) Controlling Q_1: It can be seen that $t^*\delta = 2I^*\cdot(2nI^*)^{-\frac{1}{4}}$ by the definitions of t^* and δ. Therefore, for some constant $C' > 0$ we have

$$Q_1 \geq \exp\left[-C'\left(\frac{1}{nI^*}\right)^{\frac{1}{4}} n'I^*\right].$$

2) Controlling Q_2: Recall that $p(w)$ is the pdf for $N(-1,\tau^2)$. A closer look at $q(w)$ yields

$$\begin{aligned}
q(w) &= \exp\left(\frac{w}{\tau^2}\right)\exp\left(\frac{1}{2\tau^2}\right)\frac{1}{\sqrt{2\pi\tau^2}}\exp\left(-\frac{(w+1)^2}{2\tau^2}\right)\\
&= \frac{1}{\sqrt{2\pi\tau^2}}\exp\left(-\frac{w^2}{2\tau^2}\right).
\end{aligned}$$

Therefore, $q(w)$ is the pdf for $N(0,\tau^2)$. Define $V := \mathrm{Var}\left(\frac{1}{n'}\sum_{i\in[n']} W_i\right)$ and we have $V = \frac{1}{n'}\mathrm{Var}(W_1) = \frac{\tau^2}{n'} = \frac{1}{2n'I^*}$.

Recall that $\delta := (2nI^*)^{-\frac{1}{4}}$ and $n' := n-1$. Using Chebyshev's inequality we have

$$\mathbb{P}\left(\left|\frac{1}{n'}\sum_{i\in[n']} W_i\right| > \delta\right) \leq \frac{V}{\delta^2} \leq \frac{C}{\sqrt{nI^*}},$$

for some constant $C > 0$, where the second step holds by $n \asymp n'$. Therefore, there exist some constants $C_{I^*} > 0$ and $c' \in (0,1)$ depending only on C_{I^*} such that

$$\begin{aligned}
Q_2 &= \mathbb{P}\left(\frac{1}{n'}\sum_{i\in[n']} W_i \in [0,\delta]\right)\\
&= \frac{1}{2}\left(1 - \mathbb{P}\left(\left|\frac{1}{n'}\sum_{i\in[n']} W_i\right| > \delta\right)\right) \geq c'
\end{aligned}$$

when $nI^* \geq C_{I^*}$ (implied by our assumption $nI^* \to \infty$).

3) Putting Together: Returning to Equation (18), we have

$$\begin{aligned}
&\mathbb{P}(S_{n'} \geq 0)\\
&\geq c'\cdot\exp\left[-C'\left(\frac{1}{nI^*}\right)^{\frac{1}{4}} n'I^*\right]\cdot\exp[-n'I^*]\\
&= \exp\left[-\left(1 + C'\left(\frac{1}{nI^*}\right)^{\frac{1}{4}} + \frac{1}{n'I^*}\log\frac{1}{c'}\right)n'I^*\right].
\end{aligned}$$

The desired inequality follows by taking $\xi := C'\left(\frac{1}{nI^*}\right)^{\frac{1}{4}} + \frac{1}{n'I^*}\log\frac{1}{c'}$ and noting that $\xi = o(1)$ under our assumption $nI^* \to \infty$.

C. Proof of Lemma 14 for Model 2 (CBM)

Let $n' := n-1$, $p(z)$ be the probability mass function of Z_1, and $M(t)$ be the moment generating function of Z_1. Recall that $p := \alpha(1-\epsilon)$ and $q := \alpha\epsilon$. Since $Z_1 \sim -H$, we can compute $M(t) = (1-\alpha) + pe^{-t} + qe^{t}$. Noting that $\alpha = p + q$ and recalling $t^* = \frac{1}{2}\log\frac{p}{q}$ as defined in Equation (15), we obtain

$$M(t^*) = (1-\alpha) + 2\sqrt{pq} = 1 - I^*.$$

Let $\delta := V^{\frac{1}{4}} I^{*\frac{1}{2}} (t^*)^{-\frac{1}{2}}$, $S_{n'} := \sum_{i\in[n']} Z_i$ and $S_{n'}(\mathbf{z}) := \sum_{i\in[n']} z_i$. We have

$$\begin{aligned}&\mathbb{P}(S_{n'} \geq 0)\\ &\geq \sum_{\{\mathbf{z}: S_{n'}(\mathbf{z})\in[0,n'\delta]\}} \prod_{i\in[n']} p(z_i)\\ &\geq \frac{(M(t^*))^{n'}}{\exp(n't^*\delta)} \sum_{\{\mathbf{z}: S_{n'}(\mathbf{z})\in[0,n'\delta]\}} \prod_{i\in[n']} \frac{\exp(t^* z_i)\, p(z_i)}{M(t^*)},\end{aligned}$$

where the last step holds since $\exp(n't^*\delta) \geq \exp\left(t^* \sum_{i\in[n']} z_i\right) = \prod_{i\in[n']} \exp(t^* z_i)$ given that $\sum_{i\in[n']} z_i = S_{n'}(\mathbf{z}) \leq n'\delta$. Let $q(w) := \frac{\exp(t^* w) p(w)}{M(t^*)}$ for $w \in \{-1, 0, 1\}$ and we have

$$\begin{aligned}\mathbb{P}(S_{n'} \geq 0) &\geq \exp(n' \log(1 - I^*)) \exp(-n't^*\delta)\\ &\quad\times \sum_{\{\mathbf{z}: S_{n'}(\mathbf{z})\in[0,n'\delta]\}} \prod_{i\in[n']} q(z_i).\end{aligned}$$

Noting that $q(w)$ is a pmf, we let $W_1, W_2, \ldots, W_{n'}$ be i.i.d. random variables with pmf $q(w)$. We have

$$\begin{aligned}\mathbb{P}(S_{n'} \geq 0) &\geq \exp(-n't^*\delta)\, \mathbb{P}\left(\frac{1}{n'} \sum_{i\in[n']} W_i \in [0, \delta]\right)\\ &\quad\times \exp(n' \log(1 - I^*))\\ &=: Q_1 Q_2 Q_3. \qquad (19)\end{aligned}$$

1) Controlling Q_2: A closer look at $q(w)$ yields

$$q(w) = \begin{cases} \frac{1}{M(t^*)}\sqrt{pq}, & \text{if } w = 1 \text{ or } w = -1,\\ \frac{1}{M(t^*)}(1-\alpha), & \text{if } w = 0,\end{cases}$$

whence $\operatorname{Var}(W_1) = \frac{2}{M(t^*)}\sqrt{pq}$. Define $V := \operatorname{Var}\left(\frac{1}{n'} \sum_{i\in[n']} W_i\right)$ and we have $V = \frac{1}{n'} \operatorname{Var}(W_1) = \frac{2\sqrt{pq}}{n' M(t^*)}$. We need the following estimates.

Lemma 16: If $\epsilon \in (0, \frac{1}{2})$ is a constant and $0 < q < p \leq 1 - c$ for some constant $c \in (0, 1)$, then there exist constants $C, C_1 > 0$ such that

$$V \leq \frac{4p}{cn}, \qquad \frac{V(t^*)^2}{I^{*2}} \leq \frac{C}{nI^*}, \qquad t^*\sqrt{V} \leq C_1\sqrt{\frac{I^*}{n}}.$$

Proof: By Fact 10(b), we have $I^* \leq \frac{(p-q)^2}{p} \leq p \leq 1 - c$ and therefore $M(t^*) \geq c$. This implies

$$V \leq \frac{4\sqrt{pq}}{nM(t^*)} \leq \frac{4p}{cn},$$

where the first step holds by $n' = n - 1 \geq \frac{1}{2} n$ for $n \geq 2$. Furthermore, there exist some constants $C, C', C'' > 0$ such that

$$\frac{V(t^*)^2}{I^{*2}} \leq C'' \frac{\frac{4p}{cn}\left(\frac{p-q}{p}\right)^2}{\left(\frac{(p-q)^2}{p}\right)^2} = C' \frac{p}{n(p-q)^2} \leq \frac{C}{nI^*},$$

where the first step holds by Facts 10(a) and 10(b), and the last step holds by Fact 10(b). Finally, we have

$$t^*\sqrt{V} \leq C_0 \frac{p-q}{p}\sqrt{\frac{p}{n}} \leq C_1\sqrt{\frac{I^*}{n}}$$

for some constants $C_0, C_1 > 0$, where the first step holds by Fact 10(a) and the last step holds by Fact 10(b). ■

We return to controlling Q_2. Recalling that $\delta := V^{\frac{1}{4}} I^{*\frac{1}{2}} (t^*)^{-\frac{1}{2}}$ and using Chebyshev's inequality, we have

$$\mathbb{P}\left(\left|\frac{1}{n'} \sum_{i\in[n']} W_i\right| > \delta\right) \leq \frac{V}{\delta^2} \leq \frac{C}{\sqrt{nI^*}},$$

for some constant $C > 0$, where the second step holds since $\frac{V}{\delta^2} = \frac{t^*\sqrt{V}}{I^*} \lesssim \frac{1}{\sqrt{nI^*}}$ by Lemma 16. Therefore, there exist some constants $C_{I^*} > 0$ and $c' \in (0, 1)$ that only depends on C_{I^*} such that

$$\begin{aligned}Q_2 &= \mathbb{P}\left(\frac{1}{n'} \sum_{i\in[n']} W_i \in [0, \delta]\right)\\ &= \frac{1}{2}\left(1 - \mathbb{P}\left(\left|\frac{1}{n'} \sum_{i\in[n']} W_i\right| > \delta\right)\right) \geq c'\end{aligned}$$

when $nI^* \geq C_{I^*}$ (implied by our assumption $nI^* \to \infty$).

2) Controlling Q_1: The last two inequalities of Lemma 16 implies that $t^*\delta = \sqrt{I^* t^* \sqrt{V}} \lesssim \sqrt{I^* \sqrt{\frac{I^*}{n}}} \lesssim I^* \left(\frac{1}{nI^*}\right)^{\frac{1}{4}}$. Therefore, for some constant $C' > 0$ we have

$$Q_1 \geq \exp\left[-C'\left(\frac{1}{nI^*}\right)^{\frac{1}{4}} n'I^*\right].$$

3) Controlling Q_3: By Taylor's theorem, we have $\log(1 - I^*) = -I^* - \frac{I^{*2}}{2(1-u)^2}$ for some $u \in [0, I^*]$. This implies that when $I^* \leq 1 - c_0$ for some constant $c_0 \in (0, 1)$, we have $\log(1 - I^*) \geq -I^* - C_2 I^{*2}$ for some constant $C_2 \in (0, 1)$ that only depends on c_0. It follows that

$$Q_3 \geq \exp\left[-(1 + C_2 I^*) n'I^*\right].$$

4) Putting Together: Returning to Equation (19), we obtain

$$\begin{aligned}\mathbb{P}(S_{n'} \geq 0) &\geq c' \exp\left[-(1 + C_2 I^*) n'I^*\right]\\ &\quad\times \exp\left[-C'\left(\frac{1}{nI^*}\right)^{\frac{1}{4}} n'I^*\right]\\ &= \exp\left[-\left(1 + C_2 I^* + C'\left(\frac{1}{nI^*}\right)^{\frac{1}{4}}\right.\right.\\ &\quad\left.\left. + \frac{1}{n'I^*}\log\frac{1}{c'}\right) n'I^*\right].\end{aligned}$$

The desired inequality follows by taking $\xi := C_2 I^* + C'\left(\frac{1}{nI^*}\right)^{\frac{1}{4}} + \frac{1}{n'I^*}\log\frac{1}{c'}$ and noting that $\xi = o(1)$ under our assumptions $I^* = o(1)$ and $nI^* \to \infty$.

Appendix C
Proof of the First Inequality in Theorem 4

Here we prove the first inequality in Theorem 4. The proof of the second inequality is given in Appendix E.

A. Preliminaries

Recall the definitions given in Appendix A. We introduce some additional notations. For a matrix $\mathbf{M}$, we let $\|\mathbf{M}\|_F := \sqrt{\sum_{i,j} M_{ij}^2}$ denote its Frobenius norm, $\|\mathbf{M}\|_{\text{op}}$ its spectral norm (the maximum singular value), and $\|\mathbf{M}\|_\infty := \max_{i,j} |M_{ij}|$ its entrywise ℓ_∞ norm. With another matrix $\mathbf{G}$ of the same shape as $\mathbf{M}$, we use $\mathbf{M} \geq \mathbf{G}$ to mean that $M_{ij} \geq G_{ij}$ for all i, j.

Let $\mathbf{U} := \frac{1}{\sqrt{n}}\boldsymbol{\sigma}^*$; it can be seen that $\mathbf{U}\mathbf{U}^\top = \frac{1}{n}\mathbf{Y}^*$ and in particular $\mathbf{U}$ is a singular vector of $\mathbf{Y}^*$. Define the projections $\mathcal{P}_T(\mathbf{M}) := \mathbf{U}\mathbf{U}^\top\mathbf{M} + \mathbf{M}\mathbf{U}\mathbf{U}^\top - \mathbf{U}\mathbf{U}^\top\mathbf{M}\mathbf{U}\mathbf{U}^\top$ and $\mathcal{P}_{T^\perp}(\mathbf{M}) := \mathbf{M} - \mathcal{P}_T(\mathbf{M})$ for any $\mathbf{M} \in \mathbb{R}^{n\times n}$. Recall that $\mathbf{A}$ is the observed matrix from Model 1, 2 or 3. For any $\widehat{\mathbf{Y}}$ in the sublevel set $\mathcal{Y}(\mathbf{A})$, we introduce the shorthand $\gamma := \|\widehat{\mathbf{Y}} - \mathbf{Y}^*\|_1$ for the ℓ_1 error we aim to bound. Define the shifted adjacency matrix $\mathbf{A}^0 := \mathbf{A} - \lambda^*\mathbf{J}$, where λ^* is defined in Equation (16), and the centered adjacency matrix (or noise matrix) $\mathbf{W} := \mathbf{A} - \mathbb{E}\mathbf{A} = \mathbf{A}^0 - \mathbb{E}\mathbf{A}^0$. Crucial to our analysis is a "dual certificate" $\mathbf{D}$, which is an $n \times n$ diagonal matrix with diagonal matrices

$$D_{ii} := \sum_{j\in[n]} A_{ij}^0 Y_{ij}^* = \sigma_i^* \sum_{j\in[n]} A_{ij}^0 \sigma_j^*, \qquad \text{for each } i \in [n].$$

See [9] for how $\mathbf{D}$ arises as a candidate solution to the dual program of the SDPs (7) and (8), though we do not rely on the explicit form of the dual program in the subsequent proof.

Let us record some facts about any feasible solution $\mathbf{Y}$ to program (7) or (8).

Fact 17: For any $\mathbf{Y}$ feasible to program (7) or (8), we have

$$\begin{cases} Y_{ij} - Y_{ij}^* = 0, & \text{if } i = j, \\ Y_{ij} - Y_{ij}^* \leq 0, & \text{if } Y_{ij}^* = 1, \\ Y_{ij} - Y_{ij}^* \geq 0, & \text{if } Y_{ij}^* = -1. \end{cases}$$

Proof: Since $Y_{ii} = 1$ for $i \in [n]$ and $\mathbf{Y} \succeq 0$, we have $\|\mathbf{Y}\|_\infty = 1$. The result follows from the fact that $\mathbf{Y}^* \in \{\pm 1\}^{n\times n}$. ■

Fact 18: For any $\mathbf{Y}$ feasible to program (7) or (8), the following hold.

(a) $\mathcal{P}_{T^\perp}(\mathbf{Y}) \succeq 0$ and $\operatorname{Tr}[\mathcal{P}_{T^\perp}(\mathbf{Y})] = \frac{\gamma}{n}$.
(b) $\|\mathcal{P}_{T^\perp}(\mathbf{Y})\|_\infty \leq 4$.

Proof: For part (a), note that $\mathcal{P}_{T^\perp}(\mathbf{Y}) = (\mathbf{I} - \mathbf{U}\mathbf{U}^\top)\mathbf{Y}(\mathbf{I} - \mathbf{U}\mathbf{U}^\top)$, which is positive semidefinite since $\mathbf{Y} \succeq 0$ by feasibility to program (7) or (8). We also have

$$\begin{aligned} \operatorname{Tr}[\mathcal{P}_{T^\perp}(\mathbf{Y})] &\overset{(i)}{=} \operatorname{Tr}\left[\left(\mathbf{I} - \mathbf{U}\mathbf{U}^\top\right)(\mathbf{Y} - \mathbf{Y}^*)\right] \\ &\overset{(ii)}{=} \operatorname{Tr}\left[\left(-\mathbf{U}\mathbf{U}^\top\right)(\mathbf{Y} - \mathbf{Y}^*)\right] \\ &= \frac{1}{n}\operatorname{Tr}\left[(-\mathbf{Y}^*)(\mathbf{Y} - \mathbf{Y}^*)\right] = \frac{\gamma}{n}, \end{aligned}$$

where step (i) holds since trace is invariant under cyclic permutations and the matrix $\mathbf{I} - \mathbf{U}\mathbf{U}^\top$ is idempotent, and step (ii) holds since $Y_{ii}^* - Y_{ii} = 0$ for $i \in [n]$.

For part (b), the definition of $\mathcal{P}_{T^\perp}(\cdot)$ and direct calculation give

$$\begin{aligned} \|\mathcal{P}_{T^\perp}(\mathbf{Y})\|_\infty \leq \|\mathbf{Y}\|_\infty + \|\mathbf{U}\mathbf{U}^\top\mathbf{Y}\|_\infty + \|\mathbf{Y}\mathbf{U}\mathbf{U}^\top\|_\infty \\ + \|\mathbf{U}\mathbf{U}^\top\mathbf{Y}\mathbf{U}\mathbf{U}^\top\|_\infty \leq 4, \end{aligned}$$

where the last step holds because for all (i, j), $(\mathbf{U}\mathbf{U}^\top)_{ij} = \frac{1}{n}Y_{ij}^* \in [-\frac{1}{n}, \frac{1}{n}]$ and $Y_{ij} \in [-1, 1]$. ■

We now proceed to the proof of the first inequality in Theorem 4. Following the strategy outlined in Section V, we perform the proof in three steps.

B. Step 1: Basic Inequality

As our first step, we establish the following critical basic inequality.

Lemma 19: Any $\widehat{\mathbf{Y}} \in \mathcal{Y}(\mathbf{A})$ satisfies the inequality

$$0 \leq \left\langle -\mathbf{D}, \mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}}) \right\rangle + \left\langle \mathbf{W}, \mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}}) \right\rangle.$$

We prove this lemma in Appendix D-A.

With the basic inequality, we can reduce the problem of bounding the error $\frac{1}{n}\gamma = \frac{1}{n}\|\widehat{\mathbf{Y}} - \mathbf{Y}^*\|_1$ to that of studying the two random matrices $\mathbf{D}$ (the dual certificate) and $\mathbf{W}$ (the noise matrix). In particular, recall that the matrix $\mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}})$ satisfies $\operatorname{Tr}\left[\mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}})\right] = \frac{1}{n}\gamma$ and the other properties in Fact 18. The rest of the proof relies only on these properties of $\mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}})$, and it suffices to study how matrices with such properties interact with $\mathbf{D}$ and $\mathbf{W}$.

Henceforth we use $S_1 := \left\langle -\mathbf{D}, \mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}}) \right\rangle$ and $S_2 := \left\langle \mathbf{W}, \mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}}) \right\rangle$ to denote the two terms on the RHS of the basic inequality. As the next two steps of the proof, we first control S_2, and then derive the desired exponential error rate by analyzing the sum $S_1 + S_2$.

C. Step 2: Controlling S_2

The proposition below provides a bound on S_2.

Proposition 20: Under the conditions of Theorem 4, with probability at least $1 - \frac{4}{\sqrt{n}}$, at least one of the following inequalities holds:

$$\begin{aligned} \frac{\gamma}{n} &\leq \left\lfloor n \exp\left[-\left(1 - C_e\sqrt{\frac{1}{nI^*}}\right)\bar{n}I^*\right]\right\rfloor, \\ S_2 &\leq C_{S_2}\frac{\gamma}{n}\sqrt{\frac{1}{nI^*}}\frac{\bar{n}I^*}{t^*}, \end{aligned}$$

where $C_e > 0$ and $C_{S_2} > 0$ are numeric constants.
The proof of this proposition is model-dependent, and is given in Appendices D-B, D-C and D-D for Models 1, 2 and 3, respectively.

While technical in its form, Proposition 20 has a simple interpretation: either the desired exponential error bound already holds, or S_2 is bounded by a small quantity that eventually dictates the second order term in the error exponent. To proceed, we may assume that the second bound in Proposition 20 holds. Plugging this bound into the basic inequality in Lemma 19, we obtain that

$$0 \leq \underbrace{\left\langle -\mathbf{D}, \mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}}) \right\rangle}_{S_1} + C_{S_2}\frac{\gamma}{n} \cdot \frac{1}{t^*}\sqrt{\frac{1}{nI^*}}\bar{n}I^*. \tag{20}$$

D. Step 3: Analyzing $S_1 + S_2$

If $\gamma = 0$ then we are done, so we assume $\gamma > 0$ in the sequel. To show that γ decays exponentially in I^*, we need a simple pilot bound on γ that is polynomial in I^*.

Lemma 21: Suppose that $\tau > 0$ for Model 1, $n\alpha \geq 1$ for Model 2, and $np \geq 1$ for Model 3. There exists a constant $C_{\text{pilot}} > 0$ such that with probability at least $1 - 2(e/2)^{-2n}$,

$$\gamma \leq C_{\text{pilot}} \sqrt{\frac{n^3}{I^*}}.$$

The proof is model-dependent, and is given in Appendices D-E.1, D-E.2 and D-E.3 for Models 1, 2 and 3, respectively. We verify that the premise of Lemma 21 is satisfied: for Model 1 this is clear; for Models 2 and 3, we have $p \geq \frac{(p-q)^2}{p} \gtrsim I^*$ thanks to Facts 10(b) and 11(b) (recall $p := \alpha(1-\epsilon)$ in Model 2), and $nI^* \geq C_{I^*}$ for some large enough $C_{I^*} > 0$ under the premise of Theorem 4.

Now recall that the positive semidefinite matrix $\mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}})$ has non-negative diagonal entries and satisfies $\operatorname{Tr}\left[\mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}})\right] = \frac{\gamma}{n}$ (Fact 18(a)). Define the (non-negative) numbers

$$\begin{aligned} b_i &:= \left(\mathcal{P}_{T^\perp}\left(\widehat{\mathbf{Y}}\right)\right)_{ii}, \quad b_{\max} = 4, \\ \beta &:= \frac{1}{b_{\max}} \sum_{i\in[n]} b_i = \frac{\gamma}{b_{\max} n} \end{aligned}$$

and the random variables

$$X_i := -D_{ii}, \quad i \in [n], \tag{21}$$

which is the i-th diagonal entry of the dual certificate $\mathbf{D}$ defined in Appendix C-A. With the above notations, we have

$$\begin{aligned} S_1 &:= \left\langle -\mathbf{D}, \mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}}) \right\rangle \\ &= \sum_{i\in[n]} X_i b_i \\ &= b_{\max} \sum_{i\in[n]} X_i \left(\frac{b_i}{b_{\max}}\right), \end{aligned} \tag{22}$$

where $\frac{b_i}{b_{\max}} \in [0,1]$ by Fact 18(b).

To proceed, we shall employ a technique that is reminiscent of the "order statistics argument" in [3], [49]. Let $X_{(1)} \geq X_{(2)} \geq \cdots \geq X_{(n)}$ be the order statistics of $\{X_i\}$. Let C be a constant to be chosen later. For ease of presentation, we define shorthands $\eta := C\sqrt{\frac{1}{nI^*}}$,

$$\begin{aligned} \varrho\,(m,c) &:= \frac{1}{t^*}\left[(1+\eta)\log\left(\frac{ne}{m}\right) - (1-c\eta)\bar{n}I^*\right], \quad \text{and} \\ \vartheta\,(c) &:= \exp\left[-(1-c\eta)\bar{n}I^*\right]. \end{aligned}$$

Below we consider two cases: $\beta > 1$ and $0 < \beta \leq 1$, where we recall that $\beta := \sum_{i\in[n]}\left(\frac{b_i}{b_{\max}}\right) = \frac{\gamma}{b_{\max} n}$.

Case 1 ($\beta > 1$): In this case, the expression (22) implies that

$$S_1 \leq b_{\max}\left[\sum_{i\in[\lfloor\beta\rfloor]} X_{(i)} + (\beta - \lfloor\beta\rfloor)\, X_{(\lceil\beta\rceil)}\right].$$

Combining with Equation (20), we obtain that

$$\begin{aligned} 0 &\leq S_1 + C_{S_2} b_{\max}\beta\frac{1}{t^*}\sqrt{\frac{1}{nI^*}}\bar{n}I^* \\ &\leq b_{\max}\Bigg[\sum_{i\in[\lfloor\beta\rfloor]}\left(X_{(i)} + C_{S_2}\frac{1}{t^*}\sqrt{\frac{1}{nI^*}}\bar{n}I^*\right) \\ &\qquad + (\beta - \lfloor\beta\rfloor)\left(X_{(\lceil\beta\rceil)} + C_{S_2}\frac{1}{t^*}\sqrt{\frac{1}{nI^*}}\bar{n}I^*\right)\Bigg]. \end{aligned}$$

When β is not an integer, the residual term above involving $(\beta - \lfloor\beta\rfloor)$ is cumbersome. Fortunately, the following simple lemma (proved in Appendix D-F) allows us to take the integer part of β.

Lemma 22: Suppose that $\beta \in [1,n]$, and $\phi_1 \geq \phi_2 \geq \ldots \geq \phi_n$ are n fixed numbers. Define $V(u) := \sum_{i\in[\lfloor u\rfloor]}\phi_i + (u - \lfloor u\rfloor)\phi_{\lceil u\rceil}$. If $0 \leq V(\beta)$, then $0 \leq V(\beta_0)$ for any $\beta_0 \in [1,\beta]$.

Letting $\beta_0 := \lfloor\beta\rfloor$ and invoking Lemma 22, we deduce from the last displayed equation that

$$\begin{aligned} 0 &\leq b_{\max}\sum_{i\in[\beta_0]}\left(X_{(i)} + C_{S_2}\frac{1}{t^*}\sqrt{\frac{1}{nI^*}}\bar{n}I^*\right) \\ &= b_{\max}\sum_{i\in[\beta_0]} X_{(i)} + b_{\max}\beta_0 C_{S_2}\frac{1}{t^*}\sqrt{\frac{1}{nI^*}}\bar{n}I^* \\ &\leq b_{\max}\cdot\max_{\substack{\mathcal{M}\subset[n]\\|\mathcal{M}|=\beta_0}}\left\{\sum_{i\in\mathcal{M}} X_i\right\} + b_{\max}\beta_0 C_{S_2}\frac{1}{t^*}\sqrt{\frac{1}{nI^*}}\bar{n}I^*. \end{aligned}$$

The following lemma, proved in Appendix D-G, provides a tight bound on the first sum above.

Lemma 23: Recall the definition of $\{X_i\}$ in Equation (21). Let C be any constant satisfying $C \geq 2\sqrt{2}$. Let $\eta := C\sqrt{\frac{1}{nI^*}}$ and M be any positive number satisfying $1 \leq M \leq \max\left\{1, C\sqrt{\frac{n}{I^*}}\right\}$. If $nI^* \geq C_{I^*}$ for some constant $C_{I^*} \geq 4$, then we have

$$\begin{aligned} &\max_{\substack{\mathcal{M}\subset[n]\\|\mathcal{M}|=m}}\left\{\sum_{i\in\mathcal{M}} X_i\right\} \\ &\quad \leq \frac{1}{t^*}\left((1+\eta)m\log\left(\frac{ne}{m}\right) - (1-2\eta)m\bar{n}I^*\right), \\ &\quad \forall m = 1,2,\ldots,\lfloor M\rfloor \end{aligned}$$

with probability at least $1 - 3\exp\left(-\sqrt{\log n}\right)$.

Set $C = C'_{\text{pilot}} := \frac{C_{\text{pilot}}}{b_{\max}}$ and $\eta = C\sqrt{\frac{1}{nI^*}}$. Note that Lemma 21 ensures that $\beta_0 \leq \beta \leq C'_{\text{pilot}}\sqrt{\frac{n}{I^*}}$ with high probability.[10] Therefore, applying Lemma 23 with $M = C'_{\text{pilot}}\sqrt{\frac{n}{I^*}}$ and the above C, we obtain that with high probability,

$$\begin{aligned} 0 &\leq \beta_0\cdot\varrho\,(\beta_0, 2) + C_{S_2}\frac{1}{t^*}\beta_0\sqrt{\frac{1}{nI^*}}\bar{n}I^* \\ &= \beta_0\cdot\varrho\,(\beta_0, 2) + \beta_0\frac{1}{t^*}\frac{C_{S_2}}{C'_{\text{pilot}}}\eta\bar{n}I^* \\ &= \beta_0\cdot\varrho\left(\beta_0, 2 + \frac{C_{S_2}}{C'_{\text{pilot}}}\right), \end{aligned}$$

[10]We assume that $C'_{\text{pilot}}\sqrt{\frac{n}{I^*}} \geq 1$ here; if this is not true, we can skip to the proof under case $0 < \beta \leq 1$ that is presented later.

which implies $0 \le \varrho\left(\beta_0, 2+\frac{C_{S_2}}{C'_{\text{pilot}}}\right)$. Rearranging this inequality using the definition of ϱ, we obtain

$$\beta_0 \le en \cdot \exp\left[-\frac{1-(2+C_{S_2}/C'_{\text{pilot}})\eta}{1+\eta}\bar{n}I^*\right].$$

To simplify the last RHS, we use the following elementary lemma.

Lemma 24: If $0 \le \eta \le \frac{1}{C_0+1}$ for some $C_0 > 0$, then $\frac{1-C_0\eta}{1+\eta} \ge 1-(C_0+1)\eta \ge 0$.

Proof: We have $\frac{1-C_0\eta}{1+\eta} = 1-\frac{(1+C_0)\eta}{1+\eta} \overset{(a)}{\ge} 1-(1+C_0)\eta \overset{(b)}{\ge} 0$, where step (a) holds since $\eta \ge 0$ and step (b) holds since $\eta \le \frac{1}{C_0+1}$. ■

The premise of Theorem 4, i.e., $nI^* \ge C_{I^*}$ for C_{I^*} sufficiently large, implies that $\eta \le \frac{1}{C_0+1}$ for $C_0 := 2+\frac{C_{S_2}}{C'_{\text{pilot}}}$. Applying the above lemma gives $\frac{1-C_0\eta}{1+\eta} \ge 1-(C_0+1)\eta$. Combining with the last displayed equation, we obtain that

$$\beta_0 \le en \cdot \vartheta\Big(1+C_0\Big) \implies \beta_0 \le \Big\lfloor en \cdot \vartheta\Big(1+C_0\Big)\Big\rfloor$$

since β_0 is an integer. Because $\frac{\gamma}{n} = b_{\max}\cdot\beta \le b_{\max}\cdot 2\beta_0$, it follows that

$$\frac{\gamma}{n} \le 2b_{\max}\Big\lfloor en \cdot \vartheta\Big(1+C_0\Big)\Big\rfloor \le \Big\lfloor 2b_{\max}en \cdot \vartheta\Big(1+C_0\Big)\Big\rfloor,$$

where in the last step we use the fact that $c\lfloor x\rfloor \le \lfloor cx\rfloor$ for any real number $x \ge 0$ and integer $c > 0$.[11] As long as $nI^* \ge 1$, we have $\frac{1}{nI^*} \le \sqrt{\frac{1}{nI^*}}$ and hence

$$\begin{aligned}\frac{\gamma}{n} &\le \Big\lfloor n \cdot \exp\Big(\log(2b_{\max}e)\Big)\cdot\vartheta\Big(1+C_0\Big)\Big\rfloor \\ &\le \left\lfloor n \cdot \vartheta\left(1+C_0+\frac{1}{C'_{\text{pilot}}}\log(2b_{\max}e)\right)\right\rfloor.\end{aligned}$$

Recalling the definition of ϑ, we see that we have proved the first inequality in Theorem 4.

Case 2 ($0 < \beta \le 1$): In this case, continuing from the expression (22), we have

$$S_1 \le b_{\max}\beta \cdot X_{(1)} \le b_{\max}\beta\cdot\varrho(1,2),$$

where in the last step we apply Lemma 23 with $m = M = 1$, $C = 2\sqrt{2}$ and $\eta = C\sqrt{\frac{1}{nI^*}}$. Combining with Equation (20), we obtain that

$$\begin{aligned}0 &\le b_{\max}\beta\cdot\varrho(1,2) + C_{S_2}\frac{\gamma}{n}\frac{1}{t^*}\sqrt{\frac{1}{nI^*}}\bar{n}I^* \\ &= b_{\max}\beta\cdot\varrho(1,2) + b_{\max}\beta\frac{1}{t^*}\frac{C_{S_2}}{2\sqrt{2}}\eta\bar{n}I^* \\ &= b_{\max}\beta\cdot\varrho\left(1, 2+\frac{C_{S_2}}{2\sqrt{2}}\right),\end{aligned}$$

which implies $\varrho\left(1, 2+\frac{C_{S_2}}{2\sqrt{2}}\right) \ge 0$. Rearranging this inequality using the definition of ϱ, we obtain

$$1 \le en\cdot\exp\left[-\frac{1-(2+C_{S_2}/2\sqrt{2})\eta}{1+\eta}\bar{n}I^*\right].$$

[11]Proof: we have $\lfloor cx\rfloor = \lfloor c\lfloor x\rfloor + (cx - c\lfloor x\rfloor)\rfloor \ge \lfloor c\lfloor x\rfloor\rfloor = c\lfloor x\rfloor$ by noting that $cx - c\lfloor x\rfloor \ge 0$ and $c\lfloor x\rfloor$ is an integer.

Applying Lemma 24 with $C_0 = 2+\frac{C_{S_2}}{2\sqrt{2}}$ gives $\frac{1-C_0\eta}{1+\eta} \ge 1-(C_0+1)\eta$. Combining with the last displayed equation, we obtain

$$1 \le en\cdot\vartheta(C_0+1) \implies 1 \le \lfloor en\cdot\vartheta(C_0+1)\rfloor.$$

But we have $\frac{\gamma}{b_{\max}n} = \beta \le 1$ by the case assumption. It follows that

$$\frac{\gamma}{n} \le b_{\max}\lfloor en\cdot\vartheta(C_0+1)\rfloor \le \lfloor b_{\max}en\cdot\vartheta(C_0+1)\rfloor,$$

where in the last step we use the fact that $c\lfloor x\rfloor \le \lfloor cx\rfloor$ for any real $x \ge 0$ and integer $c > 0$. As long as $nI^* \ge 1$, we have $\frac{1}{nI^*} \le \sqrt{\frac{1}{nI^*}}$ and hence

$$\begin{aligned}\frac{\gamma}{n} &\le \Big\lfloor n\cdot\exp\Big(\log(b_{\max}e)\Big)\cdot\vartheta\Big(1+C_0\Big)\Big\rfloor \\ &\le \left\lfloor n\cdot\vartheta\left(1+C_0+\frac{1}{2\sqrt{2}}\log(b_{\max}e)\right)\right\rfloor.\end{aligned}$$

Recalling the definition of ϑ, we see that we have proved the first inequality in Theorem 4.

Appendix D
Proofs of Technical Lemmas in Appendix C

A. Proof of Lemma 19

Recall the matrices $\mathbf{W}$ and $\mathbf{D}$ defined in Appendix C-A. Let $\mathbf{d} := (D_{11},\ldots,D_{nn})^\top \in \mathbb{R}^n$ be the vector of diagonal entries of $\mathbf{D}$. Note that $\mathbf{D}\boldsymbol{\sigma}^* = \boldsymbol{\sigma}^*\circ\mathbf{d} = \boldsymbol{\sigma}^*\circ\left(\boldsymbol{\sigma}^*\circ(\mathbf{A}^0\boldsymbol{\sigma}^*)\right) = \mathbf{A}^0\boldsymbol{\sigma}^*$, where $\circ$ denotes element-wise product. Therefore, we have the identity

$$\mathbf{D}\mathbf{Y}^* = \mathbf{D}\boldsymbol{\sigma}^*(\boldsymbol{\sigma}^*)^\top = \mathbf{A}^0\boldsymbol{\sigma}^*(\boldsymbol{\sigma}^*)^\top = \mathbf{A}^0\mathbf{Y}^*.$$

To prove the basic inequality, let us fix an arbitrary $\widehat{\mathbf{Y}} \in \mathcal{Y}(\mathbf{A})$ and observe that $0 \le \left\langle \mathbf{A}, \widehat{\mathbf{Y}}-\mathbf{Y}^*\right\rangle$. On the other hand, we have

$$\left\langle \mathbf{A}, \widehat{\mathbf{Y}}-\mathbf{Y}^*\right\rangle = \left\langle \mathbf{A}^0-\mathbf{D}, \widehat{\mathbf{Y}}-\mathbf{Y}^*\right\rangle$$

thanks to the following facts: (i) $\widehat{\mathbf{Y}}-\mathbf{Y}^*$ has zero diagonal and $\mathbf{D}$ is a diagonal matrix; (ii) for Models 1 and 2 we have $\mathbf{A}^0 = \mathbf{A}$; (iii) for Model 3 we have $\mathbf{A}^0 = \mathbf{A}-\lambda^*\mathbf{J}$ but the program (8) used for this model ensures that $\left\langle \mathbf{J}, \widehat{\mathbf{Y}}\right\rangle = \langle \mathbf{J}, \mathbf{Y}^*\rangle = 0$. Using the equality $\mathbf{D}\mathbf{Y}^* = \mathbf{A}^0\mathbf{Y}^*$ proved above, we obtain that

$$\begin{aligned}&\left\langle \mathbf{A}, \widehat{\mathbf{Y}}-\mathbf{Y}^*\right\rangle \\ &= \left\langle \mathbf{A}^0-\mathbf{D}, \widehat{\mathbf{Y}}\right\rangle \\ &= \left\langle \mathbf{A}^0-\mathbf{D}, \mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}})\right\rangle + \left\langle \mathbf{A}^0-\mathbf{D}, \mathcal{P}_{T}(\widehat{\mathbf{Y}})\right\rangle.\end{aligned}$$

By definition of $\mathcal{P}_T$, we can write $\mathcal{P}_T(\widehat{\mathbf{Y}}) = \boldsymbol{\sigma}^*\mathbf{u}^\top + \mathbf{v}(\boldsymbol{\sigma}^*)^\top$ for some $\mathbf{u},\mathbf{v}\in\mathbb{R}^n$, hence

$$\left\langle \mathbf{A}^0-\mathbf{D}, \mathcal{P}_T(\widehat{\mathbf{Y}})\right\rangle = \left\langle \mathbf{A}^0-\mathbf{D}, \boldsymbol{\sigma}^*\mathbf{u}^\top + \mathbf{v}(\boldsymbol{\sigma}^*)^\top\right\rangle = 0$$

because $\mathbf{D}\boldsymbol{\sigma}^* = \mathbf{A}^0\boldsymbol{\sigma}^*$. It follows that

$$\begin{aligned}&\left\langle \mathbf{A}, \widehat{\mathbf{Y}} - \mathbf{Y}^* \right\rangle \\ &= \left\langle \mathbf{A}^0 - \mathbf{D}, \mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}}) \right\rangle \\ &= \left\langle (\mathbb{E}\mathbf{A} - \lambda^*\mathbf{J}) + \mathbf{W} - \mathbf{D}, \mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}}) \right\rangle .\end{aligned}$$

We shall prove later that

$$\left\langle \mathbb{E}\mathbf{A} - \lambda^*\mathbf{J}, \mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}}) \right\rangle = 0. \tag{23}$$

Taking this identity as given, we obtain $0 \le \left\langle \mathbf{A}, \widehat{\mathbf{Y}} - \mathbf{Y}^* \right\rangle \le \left\langle \mathbf{W} - \mathbf{D}, \mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}}) \right\rangle$, thereby completing the proof of Lemma 19.

Proof of inequality (23): Under Models 1 and 2, we have $\mathbb{E}\mathbf{A} = c\mathbf{Y}^*$ for some scalar c as well as $\lambda^* = 0$ as chosen in (16), hence

$$\left\langle \mathbb{E}\mathbf{A} - \lambda^*\mathbf{J}, \mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}}) \right\rangle = c \left\langle \mathcal{P}_{T^\perp}(\mathbf{Y}^*), \widehat{\mathbf{Y}} \right\rangle = 0$$

as desired. Under Model 3, we have the equalities

$$\begin{aligned}&\langle \mathbb{E}\mathbf{A} - \lambda^*\mathbf{J}, \mathcal{P}_{T^\perp}(\mathbf{Y}) \rangle \\ &\overset{(a)}{=} \left\langle \mathcal{P}_{T^\perp}\left(\frac{p-q}{2}\mathbf{Y}^* + \frac{p+q}{2}\mathbf{J} - \lambda^*\mathbf{J} \right), \widehat{\mathbf{Y}} \right\rangle \\ &\overset{(b)}{=} \left\langle \frac{p+q}{2}\mathbf{J} - \lambda^*\mathbf{J}, \widehat{\mathbf{Y}} \right\rangle \overset{(c)}{=} 0,\end{aligned}$$

where step (a) holds since $\mathbb{E}\mathbf{A} = \frac{p-q}{2}\mathbf{Y}^* + \frac{p+q}{2}\mathbf{J}$ and the projection $\mathcal{P}_{T^\perp}$ is self-adjoint, step (b) holds since $\mathcal{P}_{T^\perp}(\mathbf{Y}^*) = 0$ and $\mathcal{P}_{T^\perp}(\mathbf{J}) = \mathbf{J}$ (because $(\mathbf{I} - \mathbf{U}\mathbf{U}^\top)\mathbf{J} = (\mathbf{I} - n^{-1}\mathbf{Y}^*)\mathbf{J} = \mathbf{J}$), and step (c) holds since $\left\langle \mathbf{J}, \widehat{\mathbf{Y}} \right\rangle = 0$ by feasibility of the matrix $\widehat{\mathbf{Y}} \in \mathcal{Y}(\mathbf{A})$ to the program (8). ■

B. Proof of Proposition 20 for Model 1 (Z2)

We shall make use of the following matrix inequality:

$$\langle \mathbf{G}, \mathbf{M} \rangle \le \|\mathbf{G}\|_{\text{op}} \operatorname{Tr}(\mathbf{M}), \quad \forall \mathbf{G}, \forall \mathbf{M} \succeq 0, \tag{24}$$

which is due to the fact that for $\mathbf{M} \succeq 0$, $\operatorname{Tr}(\mathbf{M})$ equals the sum of its singular values. We also need the following spectral norm bound, which is from a direct application of [59, Theorem 2.11].

Lemma 25: We have $\|\mathbf{W}\|_{\text{op}} \le (2 + \sqrt{2})\sqrt{\tau^2 n}$ with probability at least $1 - e^{-n/2}$.

Turning to bounding S_2, we have with probability at least $1 - e^{-n/2}$,

$$\begin{aligned}S_2 &:= \left\langle \mathbf{W}, \mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}}) \right\rangle \\ &\overset{(a)}{\le} \|\mathbf{W}\|_{\text{op}} \cdot \operatorname{Tr}\left[\mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}}) \right] \\ &\overset{(b)}{\le} 4\tau\sqrt{n} \cdot \frac{\gamma}{n} \\ &\overset{(c)}{=} 4\sqrt{2}\frac{\gamma}{n}\sqrt{\frac{1}{nI^*}\frac{nI^*}{t^*}},\end{aligned}$$

where step (a) follows from noting that $\mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}}) \succeq 0$ by Fact 18(a) and applying inequality (24), step (b) holds by Lemma 25 and Fact 18(a), and step (c) holds by definitions of I^* in Equation (5) and t^* in Equation (15). Setting $C_{S_2} = 4\sqrt{2}$ completes the proof of Proposition 20 for Z2.

C. Proof of Proposition 20 for Mode 2 (CBM)

Recall that $S_2 := \left\langle \mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}}), \mathbf{W} \right\rangle$. We control the right hand side by splitting it into two parts, one involving a trimmed version of $\mathbf{W}$ and the other the residual. This technique is similar to those in [3], [28], but here we use it for a different model.

a) Trimming: We consider an equivalent way of generating $\mathbf{A}$ under Mode 2 (CBM). Define a symmetric random matrix $\mathbf{G} \in \{\pm 1\}^{n\times n}$ such that $G_{ii} = 0$ for $i \in [n]$ and $\{G_{ij} : i < j\}$ are generated independently as

$$G_{ij} = \begin{cases} Y^*_{ij}, & \text{w.p. } 1-\epsilon, \\ -Y^*_{ij}, & \text{w.p. } \epsilon. \end{cases}$$

The observed matrix $\mathbf{A}$ from Model 2 can be equivalently generated by

$$A_{ij} = \begin{cases} G_{ij}, & \text{w.p. } \alpha, \\ 0 & \text{w.p. } 1-\alpha, \end{cases} \qquad \text{independently for } i < j,$$

with $A_{ii} = 0$ for $i \in [n]$.[12] We introduce a few additional notations. For a vector $\mathbf{v} \in \mathbb{R}^n$, we let $\|\mathbf{v}\|_0$ denote the number of nonzero entries in $\mathbf{v}$. For a matrix $\mathbf{M} \in \mathbb{R}^{n\times n}$, we let $\mathbf{M}^{\text{up}}$ be obtained from $\mathbf{M}$ by zeroing out its lower triangular entries, $\mathbf{M}_{i,:}$ and $\mathbf{M}_{:,j}$ be the i-th row and j-th column of $\mathbf{M}$ respectively, and we define the trimmed matrix $\widetilde{\mathbf{M}} := (M_{ij}\mathbb{I}\{\|\mathbf{M}_{i,:}\|_0 \le 40\alpha n, \|\mathbf{M}_{:,j}\|_0 \le 40\alpha n\})_{i,j\in[n]}$.

With the above notations, we note that $\mathbf{A}^{\text{up}}$ and $\mathbf{G}^{\text{up}}$ are both matrices with independent entries. We first record a series of lemmas that are useful for our proof to follow. The first lemma is a standard spectral norm bound for random matrices.

Lemma 26: Let $\mathbf{M} \in \mathbb{R}^{n\times n}$ be a random matrix whose entries M_{ij} are independent mean-zero random variables with $|M_{ij}| \le C'$ for some constant $C' \ge 0$. Then there exists a constant $C > 0$ such that with probability at least $1 - 2e^{-n}$, we have $\|\mathbf{M}\|_{\text{op}} \le C\sqrt{n}$.

Proof: Such a result is standard. For example, it follows as a corollary of Theorem 4.4.5 in [60] with $m = n$, $t = \sqrt{n}$ and $K \le C''$ for some constant C' since $\|\mathbf{M}\|_\infty \le C'$. ■

The next lemma controls spectral norm of the trimmed matrix $\widetilde{\mathbf{A}^{\text{up}}}$.

Lemma 27: For some absolute constant $C > 0$, we have

$$\mathbb{P}\left(\|\widetilde{\mathbf{A}^{\text{up}}} - \alpha\mathbf{G}^{\text{up}}\|_{\text{op}} \ge C\sqrt{\alpha n} \right) \le \frac{1}{n^3}.$$

Proof: Note that $\|\mathbf{G}^{\text{up}}\|_\infty \le 1$ surely. Applying [61, Lemma 3.2] with m and ϵ therein set to n and αn, we obtain $\mathbb{P}(\|\widetilde{\mathbf{A}^{\text{up}}} - \alpha\mathbf{G}^{\text{up}}\|_{\text{op}} \ge C\sqrt{\alpha n} \mid \mathbf{G}^{\text{up}}) \le \frac{1}{n^3}$ for any $\mathbf{G}^{\text{up}}$.[13] Integrating out $\mathbf{G}^{\text{up}}$ yields the result. ■

The next two lemmas are known results in the literature and control the "atypical" rows in a random matrix.

Lemma 28: Let $\mathbf{M}$ be an $n \times n$ random matrix whose entries are independent Bernoulli random variables. Let $\lambda \ge \max_{i,j} \mathbb{E}M_{ij}$. Define $\mathcal{S} := \{i \in [n] : \sum_j M_{ij} \ge 20n\lambda\}$

[12] As mentioned in Section III-B, it is inconsequential to change the diagonal entries of $\mathbf{A}$.

[13] [61, Lemma 3.2] involves trimming rows and columns that contain more than $2\alpha n$ nonzero entries. A closer inspection of their proof reveals that their bound still applies to our setting, albeit with a possibly larger constant C.

and $Z_i' := \sum_j M_{ij}\mathbb{I}\{i \in \mathcal{S}\}$. Then with probability at least $1 - e^{-5n\lambda}$, we have

$$\sum_i Z_i' \leq 20n^2\lambda e^{-5n\lambda}.$$

Proof: The proof exactly follows that of [28, Lemma C.5], with each Z_i therein replaced by Z_i' and p therein replaced by λ.[14] ■

Lemma 29: Let $\mathbf{M}$ be an $n \times n$ random matrix whose entries are independent Bernoulli random variables. Let $\lambda \geq \max_{i,j} \mathbb{E}M_{ij}$ and $\varepsilon \in (0, 1/2]$. Consider the rows of $\mathbf{M}$ with more than $21\lambda n + 2\ln\varepsilon^{-1}$ ones. Then with probability $1 - \exp(-\varepsilon n/2)$, those rows have at most εn ones altogether.

Proof: The proof exactly follows that of [62, Lemma 8.1].[15] ■

With the last two lemmas, we can establish the following result.

Lemma 30: Let $\mathbf{M} \in \{0,1\}^{n\times n}$ be a binary matrix with $M_{ii} = 0$ for all $i \in [n]$, and $\{M_{ij}\}_{i,j\in[n]}$ being independent Bernoulli random variables. Let $p' := \max_{ij} \mathbb{E}M_{ij}$. Define $\mathcal{T} := \{i \in [n] : \sum_j M_{ij} \geq 40np'\}$, $Z_i := \sum_j |M_{ij} - \mathbb{E}M_{ij}|\,\mathbb{I}\{i \in \mathcal{T}\}$ and $Z_i' := \sum_j M_{ij}\mathbb{I}\{i \in \mathcal{T}\}$. If $p' \geq \frac{C}{n}$ for a sufficiently large positive constant C, then with probability at least $1 - \frac{1}{\sqrt{n}}$, we have

$$\sum_i Z_i \leq 2\sum_i Z_i' \leq 40n^2p'e^{-5np'}.$$

Proof: Since $\sum_j M_{ij}\mathbb{I}\{i \in \mathcal{T}\} \geq 40np'\mathbb{I}\{i \in \mathcal{T}\}$ by definition of $\mathcal{T}$, we have

$$\begin{aligned}\sum_i Z_i &\leq \sum_i\sum_j M_{ij}\mathbb{I}\{i \in \mathcal{T}\} + np'\sum_i \mathbb{I}\{i \in \mathcal{T}\}\\ &\leq 2\sum_i\sum_j M_{ij}\mathbb{I}\{i \in \mathcal{T}\} = 2\sum_i Z_i',\end{aligned}$$

thereby proving the first inequality in the lemma. It remains to show that the event $B := \left\{\sum_i Z_i' > 20n^2p'e^{-5np'}\right\}$ holds with low probability. Let us consider the following two cases.

Case 1. First we consider the case where $np' \geq \frac{1}{10}\log n$. Applying Lemma 28 with λ therein set to $2p'$, we obtain that $\mathbb{P}\{B\} \leq e^{-10np'}$. Since $np' \geq \frac{1}{10}\log n$, this probability is at most $e^{-\log n} \leq \frac{1}{\sqrt{n}}$ as claimed.

Case 2. Next we consider the case where $C \leq np' < \frac{1}{10}\log n$. Set $\varepsilon = 20np'e^{-5np'}$. Note that $\varepsilon \in (0, 1/2]$ and $21np' + 2\log\varepsilon^{-1} \leq 40np'$ since $p' \geq \frac{C}{n}$. Applying Lemma 29 with the above ε and with λ therein set to p, we obtain that

$$\begin{aligned}\mathbb{P}\{B\} &\leq \exp\left(-10n^2p'e^{-5np'}\right)\\ &\leq \exp\left(-10Cne^{-\frac{1}{2}\log n}\right)\\ &= \exp\left(-10C\sqrt{n}\right) \leq \frac{1}{\sqrt{n}},\end{aligned}$$

where the second step holds since $np' < \frac{1}{10}\log n$.

[14]Inspecting their proof, we see that their bound holds without change for matrices with independent entries.

[15]Inspecting their proof, we see that their bound holds without change when the means of the Bernoulli are upper bounded by the same λ.

Putting together, we have shown that $\mathbb{P}\{B\} \leq \frac{1}{\sqrt{n}}$ for $p' \geq \frac{C}{n}$. The proof is completed. ■

We are now ready to bound S_2. Observe that

$$\begin{aligned}S_2 &= 2\left\langle \mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}}), \mathbf{W}^{\text{up}}\right\rangle\\ &= 2\Big\langle \mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}}), \left(\widetilde{\mathbf{A}^{\text{up}}} - \mathbb{E}\left[\mathbf{A}^{\text{up}} \mid \mathbf{G}\right]\right)\\ &\quad + \left(\mathbb{E}\left[\mathbf{A}^{\text{up}} \mid \mathbf{G}\right] - \mathbb{E}\mathbf{A}^{\text{up}}\right) + \left(\mathbf{A}^{\text{up}} - \widetilde{\mathbf{A}^{\text{up}}}\right)\Big\rangle\\ &= 2\Big\langle \mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}}), \left(\widetilde{\mathbf{A}^{\text{up}}} - \alpha\mathbf{G}^{\text{up}}\right) + \left(\alpha\mathbf{G}^{\text{up}} - \alpha\mathbb{E}\mathbf{G}^{\text{up}}\right)\\ &\quad + \left(\mathbf{A}^{\text{up}} - \widetilde{\mathbf{A}^{\text{up}}}\right)\Big\rangle\\ &\overset{(a)}{\leq} 2\,\mathrm{Tr}\left[\mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}})\right]\cdot\|\widetilde{\mathbf{A}^{\text{up}}} - \alpha\mathbf{G}^{\text{up}}\|_{\text{op}}\\ &\quad + 2\alpha\,\mathrm{Tr}\left[\mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}})\right]\cdot\|\mathbf{G}^{\text{up}} - \mathbb{E}\mathbf{G}^{\text{up}}\|_{\text{op}}\\ &\quad + 2\|\mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}})\|_\infty\|\mathbf{A}^{\text{up}} - \widetilde{\mathbf{A}^{\text{up}}}\|_1\\ &\overset{(b)}{\leq} 2\frac{\gamma}{n}\|\widetilde{\mathbf{A}^{\text{up}}} - \alpha\mathbf{G}^{\text{up}}\|_{\text{op}} + 2\alpha\frac{\gamma}{n}\|\mathbf{G}^{\text{up}} - \mathbb{E}\mathbf{G}^{\text{up}}\|_{\text{op}}\\ &\quad + 8\|\mathbf{A}^{\text{up}} - \widetilde{\mathbf{A}^{\text{up}}}\|_1,\end{aligned}$$

where step (a) follows noting that $\mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}}) \succeq 0$ (by Fact 18(a)) and applying inequality (24), and step (b) holds by Fact 18(a) and Fact 18(b).

We then apply Lemma 27 to bound $\|\widetilde{\mathbf{A}^{\text{up}}} - \alpha\mathbf{G}^{\text{up}}\|_{\text{op}}$, Lemma 26 to bound $\|\mathbf{G}^{\text{up}} - \mathbb{E}\mathbf{G}^{\text{up}}\|_{\text{op}}$, and Lemma 30 to $\mathbf{A}^{\text{up}}$ and $(\mathbf{A}^{\text{up}})^\top$ to bound $\|\mathbf{A}^{\text{up}} - \widetilde{\mathbf{A}^{\text{up}}}\|_1$ (we do so by setting $M_{ij} = |A_{ij}^{\text{up}}|$ for $i, j \in [n]$ and noting that $\{|A_{ij}^{\text{up}}|\}$ are independent Bernoulli random variables with means $u_{ij} \leq \alpha$). Note that the assumption $\alpha \geq \frac{C}{n}$ of Lemma 30 is satisfied by the assumption of this proposition that $nI^* \geq C_{I^*}$ for some large enough $C_{I^*} > 0$ (since Fact 10(b) implies $I^* \lesssim \frac{(p-q)^2}{p} \leq p \leq \alpha$). It follows that with probability at least $1 - \frac{4}{\sqrt{n}}$, there holds

$$S_2 \leq C_0\frac{\gamma}{n}\sqrt{n\alpha} + C_1n^2\alpha e^{-5n\alpha} =: C_0Q_1 + C_1Q_2$$

for some constants $C_0, C_1 > 0$. It remains to bound Q_1 and Q_2 above.

For Q_1, Fact 10(b) implies that $\sqrt{I^*} \leq C'\frac{p-q}{\sqrt{p}}$ for some constant $C' > 0$ and therefore

$$Q_1 = \gamma\frac{p-q}{p-q}\sqrt{\frac{p}{n}} \leq \frac{1}{C'}\gamma(p-q)\sqrt{\frac{1}{nI^*}}.$$

Bounding Q_2 involves some elementary manipulation.

b) Controlling Q_2: We record an elementary inequality.

Lemma 31: There exists a constant $C_{I^*} \geq 1$ such that if $nI^* \geq C_{I^*}$, then $pe^{-5pn} \leq (p-q)e^{-5nI^*/2}$.

Proof: Note that $pn \geq (p-q)n \geq \frac{(p-q)^2}{p}n \geq \frac{1}{C'}nI^* \geq \frac{1}{C'}C_{I^*}$ for some constant $C' > 0$ by Fact 10(b). As long as C_{I^*} is sufficiently large, we have $\frac{pn}{2} \leq e^{5pn/2}$. These inequalities imply that

$$\frac{p}{p-q} \leq \frac{pn}{2} \leq e^{5pn/2} \leq e^{5(2p-I^*)n/2}.$$

Multiplying both sides by $(p-q)e^{-5pn}$ yields the claimed inequality. ■

Equipped with the above bound, we are ready to bound Q_2. Let $\kappa := e^{-(1-\xi)nI^*}$ where $\xi := C_e\sqrt{\frac{1}{nI^*}}$ for some constant $C_e > 0$ such that $\lfloor n\kappa \rfloor > 0$. If $\frac{\gamma}{n} = 0$ or $\frac{\gamma}{n} \le \lfloor n\kappa \rfloor$, then the first inequality in Proposition 20 holds and we are done. It remains to consider the case $\frac{\gamma}{n} > \lfloor n\kappa \rfloor > 0$. We have that $\lfloor n\kappa \rfloor$ is a positive integer and $\gamma > n\lfloor n\kappa \rfloor \ge \frac{1}{2}n^2\kappa$. Hence,

$$\begin{aligned} Q_2 &= \alpha n^2 e^{-5\alpha n} \\ &\overset{(a)}{\le} 2pn^2 e^{-5pn} \\ &\overset{(b)}{\le} 2(p-q)n^2 e^{-5nI^*/2} \\ &\le 2(p-q)e^{-nI^*}\cdot n^2 e^{-nI^*} \\ &\overset{(c)}{\le} 2(p-q)e^{-nI^*}\cdot n^2\kappa \\ &\le 4(p-q)e^{-nI^*}\cdot\gamma, \end{aligned}$$

where step (a) holds since $\alpha = \frac{1}{1-\epsilon}p$ and $\epsilon \in [0, \frac{1}{2}]$ imply $\alpha \in [p, 2p]$, step (b) holds by Lemma 31, and step (c) holds by definition of κ. Choosing $C_{I^*} > 0$ large enough so that $e^{-nI^*/2} \le \sqrt{\frac{1}{nI^*}}$, we have $Q_2 \le 4\gamma(p-q)\sqrt{\frac{1}{nI^*}}$.

c) Putting together: Combining the above bounds for Q_1 and Q_2, we obtain that

$$S_2 \le C_2\gamma(p-q)\sqrt{\frac{1}{nI^*}} = C_2\frac{\gamma}{n}\sqrt{\frac{1}{nI^*}}n(p-q)$$

for some constant $C_2 > 0$. Under the assumption $0 < q \le p \le 1$, we have $p - q \le C'\frac{I^*}{t^*}\cdot\frac{1-\epsilon}{\epsilon}$ for some constant $C' > 0$ by Facts 10(a) and 10(b). This completes the proof of Proposition 20 for CBM.

D. Proof of Proposition 20 for Model 3 (SBM)

Similarly to the proof for Model 2, we control the RHS of $S_2 = \left\langle \mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}}), \mathbf{W}\right\rangle$ by splitting it into two parts, one involving a trimmed version of $\mathbf{W}$ and the other the residual. This technique is similar to those in [3], [28], but here we provide somewhat tighter bounds.

a) Trimming: We record a technical lemma concerning a trimmed Bernoulli matrix

Lemma 32: Suppose $\mathbf{M} \in \mathbb{R}^{n\times n}$ is a random matrix with zero on the diagonal and independent entries $\{M_{ij}\}$ with the following distribution: $M_{ij} = 1 - p_{ij}$ with probability p_{ij}, and $M_{ij} = -p_{ij}$ with probability $1 - p_{ij}$. Let $p' := \max_{ij} p_{ij}$, and let $\widetilde{\mathbf{M}}$ be the matrix obtained from $\mathbf{M}$ by zeroing out all the rows and columns having more than $40np'$ positive entries. Then there exists some constant $C > 0$ such that with probability at least $1 - \frac{1}{n^2}$,

$$\|\widetilde{\mathbf{M}}\|_{\text{op}} \le C\sqrt{np'}$$

Proof: The claim follows from [3, Lemma 9] with σ^2 therein set to p'. ■

Let $\mathbf{W}^{\text{up}}$ be obtained from $\mathbf{W}$ by zeroing out its lower triangular entries. To bound S_2, we observe that

$$\begin{aligned} S_2 &= 2\left\langle \mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}}), \mathbf{W}^{\text{up}}\right\rangle \\ &= 2\left\langle \mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}}), \widetilde{\mathbf{W}^{\text{up}}}\right\rangle + 2\left\langle \mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}}), \mathbf{W}^{\text{up}} - \widetilde{\mathbf{W}^{\text{up}}}\right\rangle \\ &\overset{(a)}{\le} 2\,\mathrm{Tr}\left[\mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}})\right]\cdot\|\widetilde{\mathbf{W}^{\text{up}}}\|_{\text{op}} \\ &\quad + 2\|\mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}})\|_\infty\|\mathbf{W}^{\text{up}} - \widetilde{\mathbf{W}^{\text{up}}}\|_1 \\ &\overset{(b)}{\le} 2\frac{\gamma}{n}\|\widetilde{\mathbf{W}^{\text{up}}}\|_{\text{op}} + 8\|\mathbf{W}^{\text{up}} - \widetilde{\mathbf{W}^{\text{up}}}\|_1, \end{aligned}$$

where step (a) follows from noting that $\mathcal{P}_{T^\perp}(\widehat{\mathbf{Y}}) \succeq 0$ (by Fact 18(a)) and applying inequality (24), and step (b) holds by Facts 18(a) and 18(b). We then apply Lemma 32 to $\widetilde{\mathbf{W}^{\text{up}}}$ to bound $\|\widetilde{\mathbf{W}^{\text{up}}}\|_{\text{op}}$, and apply Lemma 30 to $\mathbf{W}^{\text{up}}$ and $(\mathbf{W}^{\text{up}})^\top$ to bound $\|\mathbf{W}^{\text{up}} - \widetilde{\mathbf{W}^{\text{up}}}\|_1$.[16] Note that the assumption $p' \ge \frac{C}{n}$ of Lemma 30 is satisfied by the assumption of this proposition that $nI^* \ge C_{I^*}$ for some large enough $C_{I^*} > 0$ (since Fact 11(b) implies $I^* \lesssim \frac{(p-q)^2}{p} \le p$). We conclude that with probability at least $1 - \frac{1}{n^2} - \frac{2}{\sqrt{n}}$,

$$S_2 \le C_0\frac{\gamma}{n}\sqrt{np} + C_1 n^2 p e^{-5np} =: C_0 Q_1 + C_1 Q_2$$

for some constants $C_0, C_1 > 0$. It remains to control Q_1 and Q_2 above.

For Q_1, note that Fact 11(b) implies $\sqrt{I^*} \le C'\frac{p-q}{\sqrt{p}}$ for some constant $C' > 0$ and therefore

$$Q_1 = \gamma\frac{p-q}{p-q}\sqrt{\frac{p}{n}} \le \frac{1}{C'}\gamma(p-q)\sqrt{\frac{1}{nI^*}}.$$

Bounding Q_2 involves some elementary manipulation.

b) Controlling Q_2: We record an elementary inequality.

Lemma 33: There exists a constant $C_{I^*} \ge 1$ such that if $nI^* \ge C_{I^*}$, then $pe^{-5pn/2} \le (p-q)e^{-5nI^*/4}$.

Proof: Note that $pn \ge (p-q)n \ge \frac{(p-q)^2}{p}n \ge \frac{1}{C'}nI^* \ge \frac{1}{C'}C_{I^*}$ for some constant $C' > 0$ by Fact 11(b). As long as C_{I^*} is sufficiently large, we have $\frac{pn}{2} \le e^{5pn/4}$. These inequalities imply that

$$\frac{p}{p-q} \le \frac{pn}{2} \le e^{5pn/4} \le e^{5(2p-I^*)n/4}.$$

Multiplying both sides by $(p-q)e^{-5pn/2}$ yields the claimed inequality. ■

Equipped with the above bound, we are ready to bound Q_2. Let $\kappa := e^{-(1-\xi)nI^*/2}$ where $\xi := C_e\sqrt{\frac{1}{nI^*}}$ for some constant $C_e > 0$ such that $\lfloor n\kappa \rfloor > 0$. If $\frac{\gamma}{n} = 0$ or $\frac{\gamma}{n} \le \lfloor n\kappa \rfloor$, then the first inequality in Proposition 20 holds and we are done. It remains to consider the case $\frac{\gamma}{n} > \lfloor n\kappa \rfloor > 0$. We have that $\lfloor n\kappa \rfloor$ is a positive integer and $\gamma > n\lfloor n\kappa \rfloor \ge \frac{1}{2}n^2\kappa$.

[16] Here, we assume that $\mathbf{A}$ has zero diagonal and therefore $\mathbf{W}$ also has zero diagonal. This assumption is inconsequential to our proof as mentioned in Section III-B.

We therefore have

$$\begin{aligned} Q_2 &\le pn^2 e^{-5pn/2} \\ &\overset{(a)}{\le} (p-q)n^2 e^{-5nI^*/4} \\ &\le (p-q)e^{-nI^*/2} \cdot n^2 e^{-(1-\xi)nI^*/2} \\ &\le 2(p-q)e^{-nI^*/2} \cdot \gamma \\ &\overset{(b)}{\le} 2\gamma(p-q)\sqrt{\frac{1}{nI^*}}, \end{aligned}$$

where step (a) is due to Lemma 33, and step (b) holds because $nI^* \ge C_{I^*}$ for C_{I^*} sufficiently large.

c) Putting together: Combining the above bounds for Q_1 and Q_2, we obtain that

$$S_2 \le C_2\gamma(p-q)\sqrt{\frac{1}{nI^*}}$$

for some constant $C_2 > 0$. Under the assumption $0 < c_0 p \le q < p \le 1 - c_1$, we have $p - q \le \frac{C' I^*}{t^*}$ for a constant $C' > 0$ by Facts 11(c) and 11(b). This completes the proof of Proposition 20 for SBM.

E. Proof of Lemma 21

In this section, we establish the pilot bound in Lemma 21 under each of the three models.

1) Proof of Lemma 21 for Model 1 (Z2): Since $\widehat{\mathbf{Y}} \in \mathcal{Y}(\mathbf{A})$, we have

$$\begin{aligned} 0 &\le \left\langle \widehat{\mathbf{Y}} - \mathbf{Y}^*, \mathbf{A} \right\rangle \\ &= \left\langle \widehat{\mathbf{Y}} - \mathbf{Y}^*, \mathbb{E}\mathbf{A} \right\rangle + \left\langle \widehat{\mathbf{Y}} - \mathbf{Y}^*, \mathbf{A} - \mathbb{E}\mathbf{A} \right\rangle. \end{aligned}$$

By Fact 17 with $\mathbf{Y} = \widehat{\mathbf{Y}}$ and the fact that $\mathbb{E}\mathbf{A} = \mathbf{Y}^*$, we have $\left\langle \widehat{\mathbf{Y}} - \mathbf{Y}^*, \mathbb{E}\mathbf{A} \right\rangle = -\gamma$. Combining, we have the bound $\gamma \le \left\langle \widehat{\mathbf{Y}} - \mathbf{Y}^*, \mathbf{A} - \mathbb{E}\mathbf{A} \right\rangle$. We proceed by controlling the RHS as

$$\gamma \le \mathrm{Tr}(\widehat{\mathbf{Y}}) \cdot \|\mathbf{W}\|_{\mathrm{op}} + \mathrm{Tr}(\mathbf{Y}^*) \cdot \|\mathbf{W}\|_{\mathrm{op}} = 2n\|\mathbf{W}\|_{\mathrm{op}},$$

which follows inequality (24) applied to positive semidefinite matrices $\widehat{\mathbf{Y}}$ and $\mathbf{Y}^*$ satisfying $\mathrm{Tr}(\widehat{\mathbf{Y}}) = \mathrm{Tr}(\mathbf{Y}^*) = n$. Applying the spectral norm bound in Lemma 25, we obtain that with probability at least $1 - e^{-n/2}$,

$$\gamma \le 8n\sqrt{\tau^2 n} = 4\sqrt{2}\sqrt{\frac{n^3}{I^*}},$$

where the last step follows from the definition of I^* in Equation (5). The proof is completed.

2) Proof of Lemma 21 for Model 2 (CBM): Recall that we have introduced the shorthands $p := \alpha(1-\epsilon)$ and $q := \alpha\epsilon$. Since $\widehat{\mathbf{Y}} \in \mathcal{Y}(\mathbf{A})$, we have

$$\begin{aligned} 0 &\le \left\langle \widehat{\mathbf{Y}} - \mathbf{Y}^*, \mathbf{A} \right\rangle \\ &= \left\langle \widehat{\mathbf{Y}} - \mathbf{Y}^*, \mathbb{E}\mathbf{A} \right\rangle + \left\langle \widehat{\mathbf{Y}} - \mathbf{Y}^*, \mathbf{A} - \mathbb{E}\mathbf{A} \right\rangle. \end{aligned}$$

By Fact 17 and the fact that $\mathbb{E}\mathbf{A} = (p-q)\mathbf{Y}^*$, we have $\left\langle \widehat{\mathbf{Y}} - \mathbf{Y}^*, \mathbb{E}\mathbf{A} \right\rangle = -(p-q)\gamma$. Combining, we have the bound $\gamma \le \frac{1}{p-q}\left\langle \widehat{\mathbf{Y}} - \mathbf{Y}^*, \mathbf{A} - \mathbb{E}\mathbf{A} \right\rangle$. To control the RHS, we compute

$$\left\langle \widehat{\mathbf{Y}} - \mathbf{Y}^*, \mathbf{A} - \mathbb{E}\mathbf{A} \right\rangle \le 2 \sup_{\mathbf{Y}\succeq 0,\ \mathrm{diag}(\mathbf{Y}) \le \mathbf{1}} \left|\langle \mathbf{Y}, \mathbf{A} - \mathbb{E}\mathbf{A} \rangle\right|.$$

Grothendieck's inequality [63], [64] guarantees that

$$\sup_{\mathbf{Y}\succeq 0,\ \mathrm{diag}(\mathbf{Y}) \le \mathbf{1}} \left|\langle \mathbf{Y}, \mathbf{A} - \mathbb{E}\mathbf{A} \rangle\right| \le K_G \|\mathbf{A} - \mathbb{E}\mathbf{A}\|_{\infty\to 1}$$

where K_G denotes Grothendieck's constant ($0 < K_G \le 1.783$) and

$$\begin{aligned} \|\mathbf{A} - \mathbb{E}\mathbf{A}\|_{\infty\to 1} &:= \sup_{\mathbf{x}:\|\mathbf{x}\|_\infty \le 1} \|(\mathbf{A} - \mathbb{E}\mathbf{A})\mathbf{x}\|_1 \\ &= \sup_{\mathbf{y},\mathbf{z}\in\{\pm 1\}^n} \left|\mathbf{y}^\top(\mathbf{A} - \mathbb{E}\mathbf{A})\mathbf{z}\right|. \end{aligned}$$

Set $v^2 := \sum_{1\le i\le j\le n} \mathrm{Var}(A_{ij})$ and note that $|A_{ij} - \mathbb{E}A_{ij}| \le 2$ for $i, j \in [n]$. For each pair of fixed vectors $\mathbf{y}, \mathbf{z} \in \{\pm 1\}^n$, the Bernstein inequality ensures that for each number $t \ge 0$,

$$\mathbb{P}\left\{\left|\mathbf{y}^\top(\mathbf{A} - \mathbb{E}\mathbf{A})\mathbf{z}\right| > t\right\} \le 2\exp\left\{-\frac{t^2}{2v^2 + 4t/3}\right\}.$$

Setting $t = \sqrt{16nv^2} + \frac{8}{3}n$ gives

$$\mathbb{P}\left\{\left|\mathbf{y}^\top(\mathbf{A} - \mathbb{E}\mathbf{A})\mathbf{z}\right| > \sqrt{16nv^2} + \frac{8}{3}n\right\} \le 2e^{-2n}.$$

Applying the union bound and using the fact that $v^2 \le \alpha\frac{n^2+n}{2} = \frac{p}{1-\epsilon}\cdot\frac{n^2+n}{2}$, we obtain that with probability at least $1 - 2^{2n}\cdot 2e^{-2n} = 1 - 2(e/2)^{-2n}$,

$$\|\mathbf{A} - \mathbb{E}\mathbf{A}\|_{\infty\to 1} \le 2\sqrt{2\frac{p}{1-\epsilon}(n^3+n^2)} + \frac{8}{3}n.$$

Combining pieces, we conclude that with probability at least $1 - 2(e/2)^{-2n}$,

$$\left\langle \widehat{\mathbf{Y}} - \mathbf{Y}^*, \mathbf{A} - \mathbb{E}\mathbf{A} \right\rangle \le 8\sqrt{2\frac{p}{1-\epsilon}(n^3+n^2)} + \frac{32}{3}n;$$

whence

$$\begin{aligned} \gamma &\le \frac{1}{p-q}\left(8\sqrt{2\frac{p}{1-\epsilon}(n^3+n^2)} + \frac{32}{3}n\right) \\ &\overset{(a)}{\le} \frac{45}{p-q}\sqrt{\frac{p}{1-\epsilon}n^3} \overset{(b)}{\le} \frac{45}{C'}\sqrt{\frac{n^3}{I^*(1-\epsilon)}}, \end{aligned}$$

for some constant $C' > 0$, where step (a) holds by our assumption $p := \alpha(1-\epsilon) \ge \frac{1-\epsilon}{n}$, and step (b) follows from Fact 10(b). The proof is completed in view of the assumption of Model 2 that ϵ is a constant.

3) Proof of Lemma 21 for Model 3 (SBM): The proof follows similar arguments as those in Appendix D-E.2. Since $\widehat{\mathbf{Y}} \in \mathcal{Y}(\mathbf{A})$, we have

$$\begin{aligned} 0 &\le \left\langle \widehat{\mathbf{Y}} - \mathbf{Y}^*, \mathbf{A} \right\rangle \\ &\overset{(a)}{=} \left\langle \widehat{\mathbf{Y}} - \mathbf{Y}^*, \mathbf{A} - \frac{p+q}{2}\mathbf{J} \right\rangle \\ &= \left\langle \widehat{\mathbf{Y}} - \mathbf{Y}^*, \mathbb{E}\mathbf{A} - \frac{p+q}{2}\mathbf{J} \right\rangle + \left\langle \widehat{\mathbf{Y}} - \mathbf{Y}^*, \mathbf{A} - \mathbb{E}\mathbf{A} \right\rangle \end{aligned}$$

where step (a) holds since $\left\langle \mathbf{J}, \widehat{\mathbf{Y}} \right\rangle = \langle \mathbf{J}, \mathbf{Y}^* \rangle$. By Fact 17 and the fact that $\mathbb{E}\mathbf{A} - \frac{p+q}{2}\mathbf{J} = \frac{p-q}{2}\mathbf{Y}^*$, we have

$$\left\langle \widehat{\mathbf{Y}} - \mathbf{Y}^*, \mathbb{E}\mathbf{A} - \frac{p+q}{2}\mathbf{J} \right\rangle = -\frac{p-q}{2}\gamma.$$

Therefore, we have the bound $\gamma \le \frac{2}{p-q}\left\langle \widehat{\mathbf{Y}} - \mathbf{Y}^*, \mathbf{A} - \mathbb{E}\mathbf{A} \right\rangle$. To control the RHS, we apply Grothendieck's inequality [63], [64] to obtain

$$\begin{aligned}\left\langle \widehat{\mathbf{Y}} - \mathbf{Y}^*, \mathbf{A} - \mathbb{E}\mathbf{A} \right\rangle &\le 2 \sup_{\mathbf{Y}\succeq 0,\ \mathrm{diag}(\mathbf{Y})\le \mathbf{1}} |\langle \mathbf{Y}, \mathbf{A} - \mathbb{E}\mathbf{A}\rangle| \\ &\le 2K_G \|\mathbf{A} - \mathbb{E}\mathbf{A}\|_{\infty\to 1},\end{aligned}$$

where K_G is Grothendieck's constant $(0 < K_G \le 1.783)$ and

$$\begin{aligned}\|\mathbf{A} - \mathbb{E}\mathbf{A}\|_{\infty\to 1} &:= \sup_{\mathbf{x}:\|\mathbf{x}\|_\infty \le 1} \|(\mathbf{A} - \mathbb{E}\mathbf{A})\mathbf{x}\|_1 \\ &= \sup_{\mathbf{y},\mathbf{z}\in\{\pm 1\}^n} \left|\mathbf{y}^\top(\mathbf{A} - \mathbb{E}\mathbf{A})\mathbf{z}\right|.\end{aligned}$$

Set $v^2 := \sum_{1\le i<j\le n} \mathrm{Var}(A_{ij})$. For each pair of fixed vectors $\mathbf{y}, \mathbf{z} \in \{\pm 1\}^n$, the Bernstein inequality ensures that for each number $t \ge 0$,

$$\mathbb{P}\left\{\left|\mathbf{y}^\top(\mathbf{A} - \mathbb{E}\mathbf{A})\mathbf{z}\right| > t\right\} \le 2\exp\left\{-\frac{t^2}{2v^2 + 4t/3}\right\}.$$

Setting $t = \sqrt{16nv^2} + \frac{8}{3}n$ gives

$$\mathbb{P}\left\{\left|\mathbf{y}^\top(\mathbf{A} - \mathbb{E}\mathbf{A})\mathbf{z}\right| > \sqrt{16nv^2} + \frac{8}{3}n\right\} \le 2e^{-2n}.$$

Applying the union bound and using the fact that $v^2 \le p(n^2 + n)/2$, we obtain that with probability at least $1 - 2^{2n}\cdot 2e^{-2n} = 1 - 2(e/2)^{-2n}$,

$$\|\mathbf{A} - \mathbb{E}\mathbf{A}\|_{\infty\to 1} \le 2\sqrt{2p(n^3 + n^2)} + \frac{8}{3}n.$$

Combining pieces, we conclude that with probability at least $1 - 2(e/2)^{-2n}$,

$$\begin{aligned}\gamma &\le \frac{2}{p-q}\cdot 2K_G \cdot \left(2\sqrt{2p(n^3+n^2)} + \frac{8}{3}n\right) \\ &\overset{(a)}{\le} \frac{45\sqrt{pn^3}}{p-q} \le \frac{45}{C'}\sqrt{\frac{n^3}{I^*}},\end{aligned}$$

for some constant $C' > 0$, where step (a) holds by our assumption $p \ge \frac{1}{n}$ and the last step follows from Fact 11(b). The proof is completed.

F. Proof of Lemma 22

If $\phi_i \ge 0$ for all $i \in [\lceil\beta\rceil]$, then the result follows immediately. Now we assume that at least one of $\{\phi_i\}$ is negative. Define $w := \arg\min\{i \in [\lceil\beta\rceil] : \phi_i < 0\}$ to be the smallest index of negative ϕ_i. If $\beta_0 \in [1, w-1]$, we have $V(\beta_0) \ge 0$ since $\phi_i \ge 0$, $\forall i \in [1, w-1]$. If $\beta_0 \in [w-1, \beta]$, we note that V is decreasing on $[w-1, \beta]$ since $\phi_i < 0$, $\forall i \in [w, \lceil\beta\rceil]$, hence $V(\beta_0) \ge V(\beta) \ge 0$. The proof is completed.

G. Proof of Lemma 23

Recall that $X_i := -D_{ii} = -\sigma_i^* \sum_{j\in[n]} A_{ij}^0 \sigma_j^*$. For clarity of exposition, we define the shorthands

$$\begin{aligned}L_m &:= \max_{\mathcal{M}\subset[n],\ |\mathcal{M}|=m} \left[\sum_{i\in\mathcal{M}} X_i\right], \quad \text{for } m \in [\lfloor M\rfloor], \\ L_{m,\mathcal{M}} &:= \sum_{i\in\mathcal{M}} X_i, \quad \text{for } \mathcal{M}\subset[n] \text{ with } |\mathcal{M}| = m, \\ R_m &:= \frac{1}{t^*}\left((1+\eta)m\log\left(\frac{ne}{m}\right) - (1-2\eta)m\bar{n}I^*\right), \\ &\quad\ \text{for } m \in [\lfloor M\rfloor], \\ P_{m,\mathcal{M}} &:= \mathbb{P}\left(L_{m,\mathcal{M}} \ge R_m\right), \\ &\quad\ \text{for } \mathcal{M}\subset[n] \text{ with } |\mathcal{M}| = m, \\ P_m &:= \mathbb{P}\left(\exists\mathcal{M}\subset[n], |\mathcal{M}| = m : L_{m,\mathcal{M}} \ge R_m\right), \\ P &:= \mathbb{P}\left(\exists m \in [\lfloor M\rfloor] : L_m \ge R_m\right).\end{aligned}$$

Our goal is to show that $P \le 3\exp\left(-\sqrt{\log n}\right)$. We start the proof by controlling $P_{m,\mathcal{M}}$ for a fixed $\mathcal{M}\subset[n]$ with $|\mathcal{M}| = m$.

1) A Closer Look at $L_{m,\mathcal{M}}$: For fixed m and $\mathcal{M}$, the quantity $L_{m,\mathcal{M}}$ is the sum of mn random variables: $L_{m,\mathcal{M}} = \sum_{j\in[mn]} V_j$. A technicality is that due to the symmetry of the matrix $\mathbf{A}$, there may exist some $j \ne j' \in [mn]$ such that V_j and $V_{j'}$ identify the same random variable. Let us define a set to group together all such random variables. We set

$$\mathcal{J} := \{j \in [mn] : \exists j' \in [mn]\backslash\{j\} \text{ s.t. } V_j = V_{j'}\}.$$

We also define the complement of $\mathcal{J}$ as $\mathcal{J}^\complement := [mn]\backslash\mathcal{J}$. Note that

$$\begin{aligned}m_1 &:= \left|\mathcal{J}^\complement\right| = mn - m^2 + m, \quad \text{and} \\ m_2 &:= \frac{1}{2}|\mathcal{J}| = \frac{m(m-1)}{2}.\end{aligned}$$

It is not hard to see that $\{V_j : j \in \mathcal{J}^\complement\}$ and half of $\{V_j : j \in \mathcal{J}\}$ are independent. Now we can write

$$L_{m,\mathcal{M}} = \sum_{j\in\mathcal{J}} V_j + \sum_{j\in\mathcal{J}^\complement} V_j.$$

2) Controlling $P_{m,\mathcal{M}}$: Recall that t^* defined in Equation (15) satisfies $t^* > 0$. Using the Chernoff bound, we have

$$\begin{aligned}P_{m,\mathcal{M}} &= \mathbb{P}\left(L_{m,\mathcal{M}} \ge R_m\right) \\ &= \mathbb{P}\left(\exp\left(t^* L_{m,\mathcal{M}}\right) \ge \exp\left(t^* R_m\right)\right) \\ &\le \exp\left(-t^* R_m\right)\cdot\left[\mathbb{E}\exp\left(t^*\sum_{j\in\mathcal{J}^\complement} V_j\right)\right] \\ &\quad\cdot\left[\mathbb{E}\exp\left(t^*\sum_{j\in\mathcal{J}} V_j\right)\right] \\ &=: Q_1 Q_2 Q_3.\end{aligned}$$

It suffices to control Q_1, Q_2 and Q_3. By definition of R_m, we have

$$Q_1 = \exp\left[-(1+\eta)m\log\left(\frac{ne}{m}\right) + (1-2\eta)m\bar{n}I^*\right].$$

As our main step, we show that the following bounds for Q_2 and Q_3 hold for all three models.

Lemma 34: Under the assumption in Lemma 23, we have

$$Q_2 \le \exp\left[-(1-\eta)m\bar{n}I^*\right] \qquad \text{and} \qquad Q_3 = 1.$$

The proof of the lemma is model-dependent, and is given in Appendices D-G.4, D-G.5 and D-G.6 for Models 1, 2 and 3, respectively. Combining the above bounds for Q_1, Q_2 and Q_3, we obtain

$$\begin{aligned} P_{m,\mathcal{M}} &\le \exp\left[-(1+\eta)m\log\left(\frac{ne}{m}\right) + (1-2\eta)m\bar{n}I^*\right] \\ &\quad\cdot \exp\left[-(1-\eta)m\bar{n}I^*\right]\cdot 1 \\ &= \exp\left[-(1+\eta)m\log\left(\frac{ne}{m}\right) - \eta m\bar{n}I^*\right]. \end{aligned}$$

3) Controlling P_m and P: Using the above bound on $P_{m,\mathcal{M}}$ and applying the union bound, we have

$$\begin{aligned} P_m &\le \sum_{\mathcal{M}\subset[n]:|\mathcal{M}|=m} P_{m,\mathcal{M}} \\ &\le \binom{n}{m}\exp\left[-(1+\eta)m\log\left(\frac{ne}{m}\right) - \eta m\bar{n}I^*\right] \\ &\overset{(a)}{\le} \exp\left[-\eta m\log\left(\frac{ne}{m}\right) - \eta m\bar{n}I^*\right] \\ &\overset{(b)}{\le} \left\{\exp\left[-C\sqrt{\frac{1}{nI^*}}\log\left(\frac{ne}{m}\right) - \frac{C}{2}\sqrt{nI^*}\right]\right\}^m, \end{aligned}$$

where step (a) holds due to $\binom{n}{m} \le \left(\frac{en}{m}\right)^m$, and step (b) holds due to the definition that $\eta := C\sqrt{\frac{1}{nI^*}}$ and the fact that $\bar{n} \ge \frac{n}{2}$. We proceed by considering two cases: (i) If $m \le \sqrt{n}$, then

$$\begin{aligned} &C\sqrt{\frac{1}{nI^*}}\log\left(\frac{ne}{m}\right) + \frac{C}{2}\sqrt{nI^*} \\ &\quad\ge \frac{C}{2}\sqrt{2\log\left(\frac{ne}{m}\right)} \ge \sqrt{\log n}, \end{aligned}$$

where the last step holds since $C \ge 2$. We hence have $P_m \le \left[\exp\left(-\sqrt{\log n}\right)\right]^m < \frac{1}{2}$. (ii) If $m > \sqrt{n}$, then

$$C\sqrt{\frac{1}{nI^*}}\log\left(\frac{ne}{m}\right) + \frac{C}{2}\sqrt{nI^*} \ge \frac{C}{2}\sqrt{nI^*} \ge \log 10$$

under the assumption in Lemma 23 that $C \ge 2\sqrt{2}$ and $nI^* \ge C_{I^*} \ge 4$. We hence have $P_m \le \left[\exp\left(-\log 10\right)\right]^m = \frac{1}{10^m}$. Combining the two cases and applying union bound, we conclude that

$$\begin{aligned} P &\le \sum_{1\le m\le\sqrt{n}} P_m + \sum_{\sqrt{n}<m\le n} P_m \\ &\le \sum_{1\le m<\infty}\left[\exp\left(-\sqrt{\log n}\right)\right]^m + n\cdot\frac{1}{10^{\sqrt{n}}} \\ &\le \frac{\exp\left(-\sqrt{\log n}\right)}{1-\exp\left(-\sqrt{\log n}\right)} + \frac{1}{n} \\ &\le 2\exp\left(-\sqrt{\log n}\right) + \exp\left(-\log n\right) \\ &\le 3\exp\left(-\sqrt{\log n}\right) \end{aligned}$$

as desired. This completes the proof of Lemma 23.

4) Proof of Lemma 34 for Model 1 (Z2): Recall the random variable $H \sim N(1,\tau^2)$ defined in Appendix A. We need the following fact.

Fact 35: Under Model 1, we have the following identities

$$\mathbb{E}e^{t^*(-H)} = e^{-I^*} \qquad \text{and} \qquad \mathbb{E}e^{2t^*(-H)} = 1.$$

Proof: Recall the definitions $I^* := (2\tau^2)^{-1}$ and $t^* := \frac{1}{\tau^2}$ in Equations (5) and (15). The results follow from direct calculation:

$$\begin{aligned} \mathbb{E}e^{t^*(-H)} &= \exp\left[-t^* + \frac{1}{2}\tau^2\left(t^*\right)^2\right] = e^{-I^*}, \qquad \text{and} \\ \mathbb{E}e^{2t^*(-H)} &= \exp\left[-2t^* + 2\tau^2\left(t^*\right)^2\right] = 1. \end{aligned}$$

■

To proceed, note that each of $\{V_j : j \in [mn]\}$ is distributed as $-H$.

Controlling Q_2: We have

$$\begin{aligned} Q_2 &:= \mathbb{E}\exp\left[t^*\sum_{j\in\mathcal{J}^\complement} V_j\right] \\ &\overset{(a)}{=} \exp\left[-m_1 I^*\right] \\ &= \exp\left[-(mn-m^2+m)I^*\right] \end{aligned}$$

where step (a) follows from Fact 35. If $1 \ge C\sqrt{\frac{n}{I^*}}$, then we must have $m = \lfloor M\rfloor = 1$ and

$$Q_2 = \exp\left[-mnI^*\right] \le \exp\left[-(1-\eta)mnI^*\right].$$

If $1 \le C\sqrt{\frac{n}{I^*}}$, then we have $1 \le M \le C\sqrt{\frac{n}{I^*}}$ and

$$Q_2 \le \exp\left[-(mn-m^2)I^*\right] \le \exp\left[-(1-\eta)mnI^*\right]$$

where the last step holds since $m \le M \le C\sqrt{\frac{n}{I^*}} = n\eta$. Either way, we have the desired inequality.

Controlling Q_3: Fact 35 directly implies the desired equality:

$$Q_3 := \mathbb{E}\exp\left[t^*\sum_{j\in\mathcal{J}} V_j\right] = \left(\mathbb{E}e^{2t^*V_1}\right)^{m_2} = 1.$$

5) Proof of Lemma 34 for Model 2 (CBM): Recall the definition of the random variable H in Appendix A. We need the following fact, whose proof is deferred to the end of this section.

Fact 36: Under Model 2, we have the following identities

$$\mathbb{E}e^{t^*(-H)} = 1 - I^* \qquad \text{and} \qquad \mathbb{E}e^{2t^*(-H)} = 1.$$

To proceed, note that each of $\{V_j : j \in [mn]\}$ is distributed as $-H$.

Controlling Q_2: We have

$$\begin{aligned} Q_2 &:= \mathbb{E}\exp\left[t^*\sum_{j\in\mathcal{J}^\complement} V_j\right] \\ &\overset{(a)}{=} (1-I^*)^{m_1} \\ &\overset{(b)}{=} \exp\left[(mn-m^2+m)\log(1-I^*)\right] \\ &\overset{(c)}{\le} \exp\left[-(mn-m^2+m)I^*\right] \\ &\overset{(d)}{\le} \exp\left[-(1-\eta)mnI^*\right], \end{aligned}$$

where step (a) follows from Fact 36, step (b) follows from the fact that $m_1 = mn - m^2 + m$, step (c) holds since $\log(1-x) \le -x, \forall x < 1$, and step (d) holds since $m \le M \le C\sqrt{\frac{n}{I^*}} = n\eta$ when $1 \le C\sqrt{\frac{n}{I^*}}$, or $m = \lfloor M \rfloor = 1$ when $1 \ge C\sqrt{\frac{n}{I^*}}$. We thus obtain the desired bound on Q_2.

Controlling Q_3: Fact 36 directly implies the desired equality:

$$Q_3 := \mathbb{E}\exp\left[t^* \sum_{j\in\mathcal{J}} V_j\right] = \left(\mathbb{E}e^{2t^*(-H)}\right)^{m_2} = 1.$$

Proof of Fact 36: Recall the shorthands $p := \alpha(1-\epsilon)$ and $q := \alpha\epsilon$ introduced for Model 2; note that $\alpha = p + q$. Also recall the definitions $I^* := (\sqrt{\alpha(1-\epsilon)} - \sqrt{\alpha\epsilon})^2 = (\sqrt{p} - \sqrt{q})^2$ and $t^* := \frac{1}{2}\log\frac{1-\epsilon}{\epsilon}$ in Equations (5) and (15). The results follow from direct calculation:

$$\begin{aligned}\mathbb{E}e^{t^*(-H)} &= (1-\alpha) + pe^{-t^*} + qe^{t^*} \\ &= (1-\alpha) + 2\sqrt{pq} = 1 - I^*,\end{aligned}$$

and

$$\begin{aligned}\mathbb{E}e^{2t^*(-H)} &= (1-\alpha) + pe^{-2t^*} + qe^{2t^*} \\ &= (1-\alpha) + q + p = 1.\end{aligned}$$

■

6) Proof of Lemma 34 for Model 3 (SBM): We record the following fact, whose proof is deferred to the end of this section.

Fact 37: Let $Z \sim \mathrm{Ber}(q)$ and $Y \sim \mathrm{Ber}(p)$. We have the following identities

$$\begin{aligned}\mathbb{E}e^{t^*Z}\mathbb{E}e^{-t^*Y} &= e^{-I^*}, \\ \left(\mathbb{E}e^{t^*Z}\right)^{\frac{1}{2}}\left(\mathbb{E}e^{-t^*Y}\right)^{-\frac{1}{2}} e^{-t^*\lambda^*} &= 1, \\ \mathbb{E}e^{2t^*Z}\mathbb{E}e^{-2t^*Y} &= 1, \\ \left(\mathbb{E}e^{2t^*Z}\right)^{\frac{1}{2}}\left(\mathbb{E}e^{-2t^*Y}\right)^{-\frac{1}{2}} e^{-2t^*\lambda^*} &= 1.\end{aligned}$$

Let $\{Z_j\}, \{Z'_j\} \overset{\text{i.i.d.}}{\sim} \mathrm{Ber}(q)$ and independently $\{Y_j\}, \{Y'_j\} \overset{\text{i.i.d.}}{\sim} \mathrm{Ber}(p)$. Note that each of $\{V_j : j \in [mn]\}$ is distributed as either $Z_1 - \lambda^*$ or $-Y_1 + \lambda^*$. We define the quantities

$$\begin{aligned}m_p &:= \left|\left\{j \in \mathcal{J}^{\complement} : V_j \sim -Y_1 + \lambda^*\right\}\right|, \\ m_q &:= \left|\left\{j \in \mathcal{J}^{\complement} : V_j \sim Z_1 - \lambda^*\right\}\right|, \\ m'_p &:= \frac{1}{2}\left|\{j \in \mathcal{J} : V_j \sim -Y_1 + \lambda^*\}\right|, \\ m'_q &:= \frac{1}{2}\left|\{j \in \mathcal{J} : V_j \sim Z_1 - \lambda^*\}\right|.\end{aligned}$$

Note that $m_p + m_q = \left|\mathcal{J}^{\complement}\right| = m_1 := mn - m^2 + m$ and $m'_p + m'_q = \frac{1}{2}\left|\mathcal{J}\right|$.

Controlling Q_2: Expanding the definition of Q_2, we have

$$\begin{aligned}&Q_2 \\ &= \mathbb{E}\exp\left(t^* \sum_{j\in[m_q]} (Z_j - \lambda^*) - t^* \sum_{j\in[m_p]} (Y_j - \lambda^*)\right) \\ &= e^{-t^*\lambda^*(m_q - m_p)}\left(\mathbb{E}e^{t^*Z_1}\right)^{m_q}\left(\mathbb{E}e^{-t^*Y_1}\right)^{m_p} \\ &= \left(\mathbb{E}e^{t^*Z_1}\mathbb{E}e^{-t^*Y_1}\right)^{\frac{1}{2}m_p + \frac{1}{2}m_q} \\ &\quad \times \left(\left(\frac{\mathbb{E}e^{t^*Z_1}}{\mathbb{E}e^{-t^*Y_1}}\right)^{\frac{1}{2}} e^{-t^*\lambda^*}\right)^{m_q - m_p}.\end{aligned}$$

By Fact 37, we can continue to write

$$\begin{aligned}Q_2 &\le \exp\left(-\left(\frac{1}{2}m_p + \frac{1}{2}m_q\right)I^*\right) \\ &\le \exp\left(-\frac{1}{2}(mn - m^2 + m)I^*\right) \\ &\le \exp\left(-(1-\eta)\frac{mn}{2}I^*\right),\end{aligned}$$

where the last step holds since $m \le M \le C\sqrt{\frac{n}{I^*}} = n\eta$ when $1 \le C\sqrt{\frac{n}{I^*}}$, or $m = \lfloor M \rfloor = 1$ when $1 \ge C\sqrt{\frac{n}{I^*}}$. We thus obtain the desired bound on Q_2.

Controlling Q_3: Similar to controlling Q_2, we compute

$$\begin{aligned}&Q_3 \\ &= \mathbb{E}\exp\left(2t^* \sum_{j\in[m'_q]} (Z'_j - \lambda^*) - 2t^* \sum_{j\in[m'_p]} (Y'_j - \lambda^*)\right) \\ &= e^{-2t^*\lambda^*(m'_q - m'_p)}\left(\mathbb{E}e^{2t^*Z'_1}\right)^{m'_q}\left(\mathbb{E}e^{-2t^*Y'_1}\right)^{m'_p} \\ &= \left(\mathbb{E}e^{2t^*Z'_1}\mathbb{E}e^{-2t^*Y'_1}\right)^{\frac{1}{2}m'_p + \frac{1}{2}m'_q} \\ &\quad \times \left(\left(\frac{\mathbb{E}e^{2t^*Z'_1}}{\mathbb{E}e^{-2t^*Y'_1}}\right)^{\frac{1}{2}} e^{-2t^*\lambda^*}\right)^{m'_q - m'_p} = 1,\end{aligned}$$

where the last step holds due to Fact 37.

Proof of Fact 37: Under Model 3, recall the definitions $I^* := -2\log\left[\sqrt{pq} + \sqrt{(1-p)(1-q)}\right]$, $t^* := \frac{1}{2}\log\frac{p(1-q)}{q(1-p)}$ and $\lambda^* := \frac{1}{2t^*}\log\frac{1-q}{1-p}$ in Equations (5), (15) and (16), respectively. For the first equation, we compute

$$\begin{aligned}&\mathbb{E}e^{t^*Z}\mathbb{E}e^{-t^*Y} \\ &= \left(qe^{t^*} + 1 - q\right)\left(pe^{-t^*} + 1 - p\right) \\ &= pq + (1-p)(1-q) + q(1-p)e^{t^*} + p(1-q)pe^{-t^*} \\ &= pq + (1-p)(1-q) + 2\sqrt{pq(1-p)(1-q)} \\ &= \left(\sqrt{pq} + \sqrt{(1-p)(1-q)}\right)^2 \\ &= e^{-I^*}.\end{aligned}$$

For the second equation, noting that $e^{2t^*\lambda^*} = \frac{1-q}{1-p}$, we compute

$$\frac{\mathbb{E}e^{t^*Z}}{\mathbb{E}e^{-t^*Y}} \cdot e^{-2t^*\lambda^*} = \frac{qe^{t^*} + 1 - q}{pe^{-t^*} + 1 - p} \cdot \frac{1-p}{1-q} = \frac{q\sqrt{\frac{p(1-q)}{q(1-p)}} + 1 - q}{p\sqrt{\frac{q(1-p)}{p(1-q)}} + 1 - p} \cdot \frac{1-p}{1-q} = 1$$

and then take the square root of both sides. Finally, the remaining two equations follow from $e^{2t^*\lambda^*} = \frac{1-q}{1-p}$ and the identities

$$\mathbb{E}e^{2t^*Z} = qe^{2t^*} + 1 - q = q\frac{p(1-q)}{q(1-p)} + 1 - q = \frac{1-q}{1-p},$$
$$\mathbb{E}e^{-2t^*Y} = pe^{-2t^*} + 1 - p = p\frac{q(1-p)}{p(1-q)} + 1 - p = \frac{1-p}{1-q}.$$

■

APPENDIX E
PROOF OF THE SECOND INEQUALITY IN THEOREM 4

Fix any $\widehat{\mathbf{Y}} \in \mathcal{Y}(\mathbf{A})$. Note that $\widehat{\mathbf{Y}}, \mathbf{Y}^* \in [-1,1]^{n\times n}$ by feasibility to the program (7) or (8). It follows that

$$\|\widehat{\mathbf{Y}} - \mathbf{Y}^*\|_F^2 \le \max_{i,j\in[n]} \left\{ \left|\widehat{Y}_{ij} - Y_{ij}^*\right| \right\} \cdot \sum_{i,j\in[n]} \left|\widehat{Y}_{ij} - Y_{ij}^*\right| = 2\|\widehat{\mathbf{Y}} - \mathbf{Y}^*\|_1.$$

Combining with the first inequality of Theorem 4, we obtain

$$\|\widehat{\mathbf{Y}} - \mathbf{Y}^*\|_F^2 \le n^2 \cdot 2\exp\left[-\left(1 - C_e\sqrt{\frac{1}{nI^*}} \right) \bar{n}I^* \right] =: n^2 \cdot \varepsilon.$$

Let $\hat{\mathbf{v}}$ be an eigenvector of $\widehat{\mathbf{Y}}$ corresponding to the largest eigenvalue with $\|\hat{\mathbf{v}}\|_2^2 = n$. It can be seen that the largest eigenvalue of $\mathbf{Y}^*$ is n with $\boldsymbol{\sigma}^*$ being the corresponding eigenvector, and that all the other eigenvalues are 0. Because $\widehat{\mathbf{Y}} = \mathbf{Y}^* + (\widehat{\mathbf{Y}} - \mathbf{Y}^*)$ and $\|\widehat{\mathbf{Y}} - \mathbf{Y}^*\|_F \le \sqrt{\varepsilon}n$, Davis-Kahan theorem (see, e.g., [65, Corollary 3]) implies that

$$\min_{g\in\{\pm1\}} \|g\hat{\mathbf{u}} - \mathbf{u}^*\|_2 = 2\left|\sin\left(\frac{\theta}{2}\right)\right| \le C\sqrt{\varepsilon}$$

for some absolute constant $C > 0$, where $\hat{\mathbf{u}}$ and $\mathbf{u}^*$ denote the unit-norm eigenvectors associated with the largest eigenvalues of $\widehat{\mathbf{Y}}$ and $\mathbf{Y}^*$, respectively, and $\theta \in [0, \frac{\pi}{2}]$ denotes the angle between these two vectors. By definition $\hat{\mathbf{v}} = \sqrt{n}\hat{\mathbf{u}}$ and $\boldsymbol{\sigma}^* = \sqrt{n}\mathbf{u}^*$, we obtain that

$$\min_{g\in\{\pm1\}} \|g\hat{\mathbf{v}} - \boldsymbol{\sigma}^*\|_2^2 \le C^2\varepsilon n.$$

We proceed by relating $\mathtt{err}(\widehat{\boldsymbol{\sigma}}^{\text{sdp}}, \boldsymbol{\sigma}^*)$ to $\min_{g\in\{\pm1\}} \|g\mathbf{v} - \boldsymbol{\sigma}^*\|_2^2$. Without loss of generality, assume that the minimum is attained by $g = 1$. Since $\widehat{\sigma}_i^{\text{sdp}} = \text{sign}(\hat{v}_i)$ by definition, we have the bound

$$C^2\varepsilon n \ge \|\hat{\mathbf{v}} - \boldsymbol{\sigma}^*\|_2^2 \ge \sum_{i\in[n]} (\hat{v}_i - \sigma_i^*)^2 \mathbb{I}\{\text{sign}(\hat{v}_i) \ne \sigma_i^*\} \ge \sum_{i\in[n]} \mathbb{I}\{\text{sign}(\hat{v}_i) \ne \sigma_i^*\} \ge n \cdot \mathtt{err}(\widehat{\boldsymbol{\sigma}}^{\text{sdp}}, \boldsymbol{\sigma}^*).$$

We divide both sides of the above equation by n, and note that the constant $2C^2$ can be absorbed into C_e under the assumption that $nI^* \ge C_{I^*}$ for C_{I^*} sufficiently large. The result follows.

REFERENCES

[1] Y. Fei and Y. Chen, "Achieving the Bayes error rate in stochastic block model by SDP, robustly," in *Proc. Conf. Learn. Theory*, 2019, pp. 1235–1269.

[2] O. Guédon and R. Vershynin, "Community detection in sparse networks via Grothendieck's inequality," *Probab. Theory Rel. Fields*, vol. 165, nos. 3–4, pp. 1025–1049, Aug. 2016.

[3] Y. Fei and Y. Chen, "Exponential error rates of SDP for block models: Beyond Grothendieck's inequality," *IEEE Trans. Inf. Theory*, vol. 65, no. 1, pp. 551–571, Jan. 2019.

[4] A. Y. Zhang and H. H. Zhou, "Minimax rates of community detection in stochastic block models," *Ann. Statist.*, vol. 44, no. 5, pp. 2252–2280, Oct. 2016.

[5] U. Feige and J. Kilian, "Heuristics for semirandom graph problems," *J. Comput. Syst. Sci.*, vol. 63, no. 4, pp. 639–671, Dec. 2001.

[6] A. Moitra, W. Perry, and A. S. Wein, "How robust are reconstruction thresholds for community detection?" in *Proc. 48th Annu. ACM SIGACT Symp. Theory Comput. (STOC)*, 2016, pp. 828–841.

[7] C. Giraud and N. Verzelen, "Partial recovery bounds for clustering with the relaxed k means," 2018, *arXiv:1807.07547*. [Online]. Available: https://arxiv.org/abs/1807.07547

[8] B. Hajek, Y. Wu, and J. Xu, "Achieving exact cluster recovery threshold via semidefinite programming," *IEEE Trans. Inf. Theory*, vol. 62, no. 5, pp. 2788–2797, May 2016.

[9] A. S. Bandeira,"Random Laplacian matrices and convex relaxations," *Found. Comput. Math.*, vol. 18, no. 2, pp. 345–379, Apr. 2018.

[10] E. Abbe, "Community detection and stochastic block models: Recent developments," *J. Mach. Learn. Res.*, vol. 18, pp. 1–86, Jan. 2018.

[11] C. Moore, "The computer science and physics of community detection: Landscapes, phase transitions, and hardness," *Bull. Eur. Assoc. Theor. Comput. Sci.*, vol. 1, no. 121, Feb. 2017.

[12] X. Li, Y. Chen, and J. Xu, "Convex relaxation methods for community detection," 2018, *arXiv:1810.00315*. [Online]. Available: https://arxiv.org/abs/1810.00315

[13] A. S. Bandeira, N. Boumal, and A. Singer, "Tightness of the maximum likelihood semidefinite relaxation for angular synchronization," *Math. Program.*, vol. 163, nos. 1–2, pp. 145–167, May 2017.

[14] E. Abbe, J. Fan, K. Wang, and Y. Zhong, "Entrywise eigenvector analysis of random matrices with low expected rank," 2017, *arXiv:1709.09565*. [Online]. Available: https://arxiv.org/abs/1709.09565

[15] A. Perry, A. S. Wein, A. S. Bandeira, and A. Moitra, "Optimality and sub-optimality of pca for spiked random matrices and synchronization," 2016, *arXiv:1609.05573*. [Online]. Available: https://arxiv.org/abs/1609.05573

[16] M. Lelarge and L. Miolane, "Fundamental limits of symmetric low-rank matrix estimation," *Probab. Theory Rel. Fields*, vol. 173, nos. 3–4, pp. 859–929, Apr. 2019.

[17] E. Abbe, A. S. Bandeira, A. Bracher, and A. Singer, "Decoding binary node labels from censored edge measurements: Phase transition and efficient recovery," *IEEE Trans. Netw. Sci. Eng.*, vol. 1, no. 1, pp. 10–22, Jan. 2014.

[18] B. Hajek, Y. Wu, and J. Xu, "Exact recovery threshold in the binary censored block model," in *Proc. IEEE Inf. Theory Workshop-Fall (ITW)*, Oct. 2015, pp. 99–103.

[19] S. Heimlicher, M. Lelarge, and L. Massoulié, "Community detection in the labelled stochastic block model," 2012, *arXiv:1209.2910*. [Online]. Available: https://arxiv.org/abs/1209.2910

[20] M. Lelarge, L. Massoulie, and J. Xu, "Reconstruction in the labelled stochastic block model," *IEEE Trans. Netw. Sci. Eng.*, vol. 2, no. 4, pp. 152–163, Oct. 2015.

[21] P. Chin, A. Rao, and V. Vu, "Stochastic block model and community detection in sparse graphs: A spectral algorithm with optimal rate of recovery," in *Proc. 28th Conf. Learn. Theory (COLT)*, Paris, France, Jul. 2015, pp. 391–423. [Online]. Available: http://jmlr.org/proceedings/papers/v40/Chin15.html

[22] A. Saade, M. Lelarge, F. Krzakala, and L. Zdeborova, "Spectral detection in the censored block model," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jun. 2015, pp. 1184–1188.

[23] C. Gao, Z. Ma, A. Y. Zhang, and H. H. Zhou, "Achieving optimal misclassification proportion in stochastic block models," *J. Mach. Learn. Res.*, vol. 18, no. 60, pp. 1980–2024, Jan. 2017.

[24] C. Gao, Z. Ma, A. Y. Zhang, and H. H. Zhou, "Community detection in degree-corrected block models," *Ann. Statist.*, vol. 46, no. 5, pp. 2153–2185, Oct. 2018.

[25] M. Xu, V. Jog, and P.-L. Loh, "Optimal rates for community estimation in the weighted stochastic block model," 2017, *arXiv:1706.01175*. [Online]. Available: https://arxiv.org/abs/1706.01175

[26] S.-Y. Yun and A. Proutiere, "Accurate community detection in the stochastic block model via spectral algorithms," 2014, *arXiv:1412.7335*. [Online]. Available: https://arxiv.org/abs/1412.7335

[27] S.-Y. Yun and A. Proutiere, "Optimal cluster recovery in the labeled stochastic block model," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 965–973.

[28] A. Y. Zhang and H. H. Zhou, "Theoretical and computational guarantees of mean field variational inference for community detection," 2017, *arXiv:1710.11268*. [Online]. Available: https://arxiv.org/abs/1710.11268

[29] Z. Zhou and P. Li, "Non-asymptotic Chernoff lower bound and its application to community detection in stochastic block model," 2018, *arXiv:1812.11269*. [Online]. Available: https://arxiv.org/abs/1812.11269

[30] C. Gao and Z. Ma, "Minimax rates in network analysis: Graphon estimation, community detection and hypothesis testing," 2018, *arXiv:1811.06055*. [Online]. Available: https://arxiv.org/abs/1811.06055

[31] E. Abbe, A. S. Bandeira, and G. Hall, "Exact recovery in the stochastic block model," *IEEE Trans. Inf. Theory*, vol. 62, no. 1, pp. 471–487, Jan. 2016.

[32] E. Mossel, J. Neeman, and A. Sly, "Consistency thresholds for the planted bisection model," *Electron. J. Probab.*, vol. 21, no. 21, pp. 1–24, 2016, doi: 10.1214/16-EJP4185.

[33] E. Abbe and C. Sandon, "Community detection in general stochastic block models: Fundamental limits and efficient algorithms for recovery," in *Proc. IEEE 56th Annu. Symp. Found. Comput. Sci.*, Oct. 2015, pp. 670–688.

[34] E. Abbe and C. Sandon, "Recovering communities in the general stochastic block model without knowing the parameters," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 676–684.

[35] V. Jog and P.-L. Loh, "Information-theoretic bounds for exact recovery in weighted stochastic block models using the Renyi divergence," 2015, *arXiv:1509.06418*. [Online]. Available: https://arxiv.org/abs/1509.06418

[36] W. Perry and A. S. Wein, "A semidefinite program for unbalanced multisection in the stochastic block model," 2015, *arXiv:1507.05605*. [Online]. Available: https://arxiv.org/abs/1507.05605

[37] L. Massoulié, "Community detection thresholds and the weak Ramanujan property," in *Proc. 46th Annu. ACM Symp. Theory Comput. (STOC)*, 2014, pp. 694–703.

[38] E. Mossel, J. Neeman, and A. Sly, "Reconstruction and estimation in the planted partition model," *Probab. Theory Rel. Fields*, vol. 162, nos. 3–4, pp. 431–461, Aug. 2015.

[39] E. Abbe and C. Sandon, "Detection in the stochastic block model with multiple clusters: Proof of the achievability conjectures, acyclic BP, and the information-computation gap," 2015, *arXiv:1512.09080*. [Online]. Available: https://arxiv.org/abs/1512.09080

[40] E. Abbe, E. Boix, P. Ralli, and C. Sandon, "Graph powering and spectral robustness," 2018, *arXiv:1809.04818*. [Online]. Available: https://arxiv.org/abs/1809.04818

[41] D. Banerjee, "Contiguity and non-reconstruction results for planted partition models: The dense case," *Electron. J. Probab.*, vol. 23, 2018.

[42] C. Bordenave, M. Lelarge, and L. Massoulié, "Non-backtracking spectrum of random graphs: community detection and non-regular Ramanujan graphs," *Ann. Probab.*, vol. 46, no. 1, pp. 1–71, Jan. 2018.

[43] F. Caltagirone, M. Lelarge, and L. Miolane, "Recovering asymmetric communities in the stochastic block model," *IEEE Trans. Netw. Sci. Eng.*, vol. 5, no. 3, pp. 237–246, Jul. 2018.

[44] A. Coja-Oghlan, F. Krzakala, W. Perkins, and L. Zdeborová, "Information-theoretic thresholds from the cavity method," *Adv. Math.*, vol. 333, pp. 694–795, Jul. 2018.

[45] E. Mossel, J. Neeman, and A. Sly, "A proof of the block model threshold conjecture," *Combinatorica*, vol. 38, no. 3, pp. 665–708, Jun. 2018.

[46] L. Stephan and L. Massoulié, "Robustness of spectral methods for community detection," 2018, *arXiv:1811.05808*. [Online]. Available: https://arxiv.org/abs/1811.05808

[47] N. Agarwal, A. S. Bandeira, K. Koiliaris, and A. Kolla, "Multisection in the stochastic block model using semidefinite programming," in *Compressed Sensing and Its Applications*. Springer, 2017, pp. 125–162.

[48] A. Montanari and S. Sen, "Semidefinite programs on sparse random graphs and their application to community detection," in *Proc. 48th Annu. ACM SIGACT Symp. Theory Comput. (STOC)*, 2016, pp. 814–827.

[49] Y. Fei and Y. Chen, "Hidden integrality of SDP relaxation for sub-Gaussian mixture models," 2018, *arXiv:1803.06510*. [Online]. Available: https://arxiv.org/abs/1803.06510

[50] B. Hajek, Y. Wu, and J. Xu, "Achieving exact cluster recovery threshold via semidefinite programming: Extensions," *IEEE Trans. Inf. Theory*, vol. 62, no. 10, pp. 5918–5937, Oct. 2016.

[51] T. T. Cai and X. Li, "Robust and computationally feasible community detection in the presence of arbitrary outlier nodes," *Ann. Statist.*, vol. 43, no. 3, pp. 1027–1059, Jun. 2015. [Online]. Available: http://arxiv.org/abs/1404.6000

[52] K. Makarychev, Y. Makarychev, and A. Vijayaraghavan, "Learning communities in the presence of errors," in *Proc. 29th Annu. Conf. Learn. Theory*, 2016, pp. 1258–1291.

[53] P. W. Holland, K. B. Laskey, and S. Leinhardt, "Stochastic blockmodels: First steps," *Social Netw.*, vol. 5, no. 2, pp. 109–137, Jun. 1983.

[54] A. Blum and J. Spencer, "Coloring random and semi-random k-colorable graphs," *J. Algorithms*, vol. 19, no. 2, pp. 204–234, Sep. 1995.

[55] M. Gil, F. Alajaji, and T. Linder, "Rényi divergence measures for commonly used univariate continuous distributions," *Inf. Sci.*, vol. 249, pp. 124–131, Nov. 2013.

[56] Y. Chen, S. Sanghavi, and H. Xu, "Improved graph clustering," *IEEE Trans. Inf. Theory*, vol. 60, no. 10, pp. 6440–6455, Oct. 2014.

[57] E. Mossel, J. Neeman, and A. Sly, "A proof of the block model threshold conjecture," 2013, *arXiv:1311.4115*. [Online]. Available: https://arxiv.org/abs/1311.4115

[58] A. Dembo and O. Zeitouni, *Large Deviations Techniques and Applications* (Stochastic Modelling and Applied Probability). Berlin, Germany: Springer, 2010.

[59] K. R. Davidson and S. J. Szarek, "Local operator theory, random matrices and Banach spaces," in *Handbook of the Geometry of Banach Spaces*, vol. 1, W. Johnson and J. Lindenstrauss, Eds. Amsterdam, The Netherlands: Elsevier, 2001, pp. 317–366. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1874584901800103

[60] R. Vershynin, *High-Dimensional Probability: An Introduction With Applications in Data Science* (Cambridge Series in Statistical and Probabilistic Mathematics). Cambridge, U.K.: Cambridge Univ. Press, 2018.

[61] R. H. Keshavan, A. Montanari, and S. Oh, "Matrix completion from a few entries," *IEEE Trans. Inf. Theory*, vol. 56, no. 6, pp. 2980–2998, Jun. 2010.

[62] E. Rebrova and R. Vershynin, "Norms of random matrices: Local and global problems," *Adv. Math.*, vol. 324, pp. 40–83, Jan. 2018.

[63] A. Grothendieck, "Résumé de la théorie métrique des produits tensoriels topologiques," *Resenhas Inst. Mat. Estatistica Univ. São Paulo*, vol. 2, no. 4, pp. 401–481, 1953.

[64] J. Lindenstrauss and A. Pełczyński, "Absolutely summing operators in L_p-spaces and their applications," *Stud. Math.*, vol. 29, no. 3, pp. 275–326, 1968.

[65] V. Vu, "Singular vectors under random perturbation," *Random Struct. Algorithms*, vol. 39, no. 4, pp. 526–538, Dec. 2011.

Yingjie Fei received the B.A. degree in mathematics from New York University in 2015. He is currently pursuing the Ph.D. degree with the School of Operations Research and Information Engineering, Cornell University. His research interests include machine learning, high-dimensional and robust statistics, and optimization. His work has been awarded the Second Place in 2018 INFORMS George Nicholson Student Paper Competition.

Yudong Chen received the B.S. and M.S. degrees in control science and engineering from Tsinghua University and the Ph.D. degree in electrical and computer engineering from The University of Texas at Austin in 2013. He was a Post-Doctoral Researcher with the Electrical Engineering and Computer Sciences Department, University of California at Berkeley, Berkeley. He is currently an Assistant Professor with the School of Operations Research and Information Engineering, Cornell University. His research work lies in machine learning, high-dimensional statistics, and optimization, with applications in network scheduling, wireless communication, and financial systems. He has served on the senior program committees of AAAI and AISTATS.