INNOVATED SCALABLE EFFICIENT ESTIMATION IN ULTRA-LARGE GAUSSIAN GRAPHICAL MODELS[1]

BY YINGYING FAN AND JINCHI LV

University of Southern California

Large-scale precision matrix estimation is of fundamental importance yet challenging in many contemporary applications for recovering Gaussian graphical models. In this paper, we suggest a new approach of innovated scalable efficient estimation (ISEE) for estimating large precision matrix. Motivated by the innovated transformation, we convert the original problem into that of large covariance matrix estimation. The suggested method combines the strengths of recent advances in high-dimensional sparse modeling and large covariance matrix estimation. Compared to existing approaches, our method is scalable and can deal with much larger precision matrices with simple tuning. Under mild regularity conditions, we establish that this procedure can recover the underlying graphical structure with significant probability and provide efficient estimation of link strengths. Both computational and theoretical advantages of the procedure are evidenced through simulation and real data examples.

1. Introduction. The surge of big data in an unprecedented scale has brought us an enormous amount of information about individuals in a spectrum of contemporary applications including social networks, online marketing and modern healthcare. It is often of practical interest to uncover the underlying network formed by a large number of individuals that are sparsely related. Graphical models provide a flexible way to specify the conditional independence structure among a set of nodes. See, for example, [29, 42] for detailed accounts and applications of such models. In Gaussian graphical models, the conditional independence structure is fully characterized by the zero entries in the precision (inverse covariance) matrix. For instance, the nonzero entries of a precision matrix estimated from genomic data detect interactions among genes or proteins of potential interest. The

precision matrix also appears in many other applications such as classification and portfolio management.

The problem of identifying zeros in the precision matrix was termed as covariance selection in [9], which serves as a parsimonious way to simplify the model on the covariance structure. A stepwise estimation procedure was proposed therein based on the rule that the covariance matrix estimator is positive definite and matches the sample one on a set of entries, while its inverse has zeros in the remaining entries. In the Gaussian setting, it was shown that such a covariance model attains maximum entropy (simplicity) and the proposed covariance matrix estimator has the appealing property of being the restricted maximum likelihood estimate. Such a procedure works for the case when the number of variables p is low but becomes computationally expensive as p increases.

Large precision matrix estimation has attracted much recent attention of many researchers. Broadly speaking, existing methods can be classified into two classes: the penalized likelihood or empirical risk methods, and the penalized regression or Dantzig selector type optimization methods. The former class includes, for example, [13, 22, 37, 45, 47]. These methods share a common feature that the precision matrix is estimated by maximizing the penalized Gaussian likelihood or minimizing the penalized empirical risk. The latter class includes, for instance, [5, 6, 35, 36, 38, 44]. Such methods convert the problem of precision matrix estimation into a nodewise or pairwise regression, or optimization problem and then apply the technique of high-dimensional regularization using the Lasso or Dantzig selector type methods. In particular, optimal rates of convergence for estimating sparse precision matrix have been established in [6]. The aforementioned methods are efficient in estimating precision matrix in moderate dimensions, but may become computationally inefficient when dealing with a huge number of nodes.

To address the important issue of scalability that is crucial to uncovering ultra-large Gaussian graphical models, in this paper we suggest a new method, called the innovated scalable efficient estimation (ISEE), for large precision matrix estimation. Our approach is motivated by the idea of the innovated transformation, which is a linear transformation of the p-variate random vector for the p nodes using the precision matrix; see (3) for formal definition. A simple observation is that the covariance matrix of the transformed p-variate random vector is exactly the precision matrix of the original p-variate random vector. Aided by such a transformation, we convert the original problem of large precision matrix estimation into that of large covariance matrix estimation. To estimate the innovated data matrix, the so-called oracle empirical matrix [see (8) for formal definition], which is unavailable to practitioners, we exploit the scaled Lasso regression in [40] applied p times based on a partition of all the p nodes. After obtaining such a data matrix, we treat it as a "sample" from the innovated p-variate random vector, and apply the approach of thresholding in [3] to construct a sparse precision matrix estimator.

The innovated transformation is related to the term "innovation" used in the time series literature [23] and has been utilized by other researchers in various

contexts. For example, it was proposed and exploited in [23] for detecting sparse signals when the noises are correlated. It was used in [18] for high-dimensional optimal classification with correlated features. See also [27] for a discussion of the innovated transformation in the multiple testing setting.

The suggested ISEE method combines the strengths of recent advances in both fields of high-dimensional sparse modeling and large covariance matrix estimation. The scaled Lasso is a convex regularization method that is tuning free and admits efficient implementation, while the thresholding method for large covariance matrix estimation is easy to implement and powered by appealing theoretical properties. As a consequence, there is only one tuning parameter for ISEE which is the threshold. To select such a threshold, we adapt the method of the cross-validation in [2, 3] for large covariance matrix estimation. Since we apply the cross-validation to the estimated oracle empirical matrix, not the original data matrix, there is no need to repeat the sparse regression step, and thus the ISEE enjoys computational efficiency. As such, ISEE is scalable and can deal with much larger precision matrices with simple tuning, compared to existing approaches. In addition to the computational advantage, we have also shown that the suggested procedure can recover the underlying graphical structure with significant probability and provide efficient estimation of link strengths under mild regularity conditions.

The rest of the paper is organized as follows. Section 2 introduces the suggested approach of ISEE for large Gaussian graphical models, and discusses its computation in large or ultra-large scale. We present the asymptotic efficiency of the new method in Section 3. Section 4 details some examples of applications for our method. We provide several numerical examples in Section 5. Section 6 discusses some extensions of the suggested method to a few settings. The proofs of some main results are relegated to the Appendix. Additional proofs of main results and technical details are provided in the Supplementary Material [21].

2. Innovated scalable efficient estimation in ultra-large Gaussian graphical models.

2.1. *Model setting.* Consider the Gaussian graphical model $G = (V, E)$ for a p-variate random vector

$$\mathbf{x} = (X_1, \ldots, X_p)^T \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma}), \tag{1}$$

where $\boldsymbol{\mu}$ is a p-dimensional mean vector, $\boldsymbol{\Sigma} = (\sigma_{jk})$ is a $p \times p$ covariance matrix, and G is an undirected graph associated with $\mathbf{x}$ with $V = \{X_1, \ldots, X_p\}$ the set of vertices (or nodes) and $E = \{(j, k)\}$ the set of edges (or links) between the vertices. In this model, the lack of an edge (j, k) between a pair of vertices X_j and X_k is characterized by the probabilistic property that these two components are independent conditional on the remaining $p - 2$ vertices. In other words, the existence of an edge amounts to conditional dependence between the two vertices given all other ones. Denote by $\boldsymbol{\Omega} = (\omega_{jk})$ the precision matrix, that is, the inverse

$\mathbf{\Sigma}^{-1}$ of the covariance matrix $\mathbf{\Sigma}$. It is well known in the Gaussian graphical model theory that there is an edge (j, k) between a pair of vertices X_j and X_k if and only if the corresponding entry ω_{jk} of the precision matrix $\mathbf{\Omega}$ is nonzero. See, for example, [29] for a detailed account of graphical models. Such a characterization of the edge set shows that the problem of recovering the Gaussian graph G is equivalent to recovering the support

(2) $\text{supp}(\mathbf{\Omega}) = \{(j, k) : \omega_{jk} \neq 0\}$ modulo symmetry,

meaning the equivalence of links between nodes j and k in undirected graphs, of the precision matrix $\mathbf{\Omega}$. In particular, the strength of each link (j, k) is characterized by the magnitude of the corresponding entry ω_{jk}.

Suppose $(\mathbf{x}_i)_{i=1}^n$ is an independent and identically distributed (i.i.d.) sample from the Gaussian graphical model (1). Without loss of generality, assume that the mean vector $\boldsymbol{\mu} = \mathbf{0}$ throughout the paper. One natural and important question is how to efficiently recover the graphical structure and infer about the link strengths in large scale, that is, when the number of nodes p is large compared to the sample size n. We will address this problem in the remaining part of the paper.

2.2. *Innovated scalable efficient estimation.* Estimating the precision matrix $\mathbf{\Omega}$ associated with the Gaussian graph G is challenging even in moderate dimensionality p. Directly inverting the sample covariance matrix is infeasible since it is singular when $p > n$. To overcome this difficulty, various methods have been proposed. As discussed in the Introduction, a common limitation of these methods is that they are computationally intensive which can restrain their applications when estimating very large graphs.

To address these challenges, we propose a new procedure called the innovated scalable efficient estimation (ISEE) for effective and efficient large precision matrix estimation. The main idea of our approach is to convert the problem of estimating large precision matrix $\mathbf{\Omega}$ to that of estimating large covariance matrix. Our method is motivated by the following linear transformation:

(3) $\widetilde{\mathbf{x}} = \mathbf{\Omega}\mathbf{x}.$

Observe that the p-variate transformed random vector $\widetilde{\mathbf{x}}$ in (3) still has a Gaussian distribution with mean $\mathbf{0}$ and covariance matrix

(4) $\text{cov}(\widetilde{\mathbf{x}}) = \mathbf{\Omega}\,\text{cov}(\mathbf{x})\mathbf{\Omega} = \mathbf{\Omega}\mathbf{\Sigma}\mathbf{\Omega} = \mathbf{\Omega}.$

Thus, if the transformed vector $\widetilde{\mathbf{x}}$ were observable, then estimating the precision matrix $\mathbf{\Omega}$ could be achieved by estimating the covariance matrix of the p-variate Gaussian random vector $\widetilde{\mathbf{x}}$. Our new view of this problem naturally provides flexible alternative ways of Gaussian graph estimation powered by recent developments in large covariance matrix estimation. See, for example, [2–4, 7, 12, 28, 39], among others.

The transformation (3) with the precision matrix $\boldsymbol{\Omega}$ is termed as innovation in the time series literature. We thus refer to (3) as the innovated transformation and incorporate the word "innovated" in the name of ISEE. As mentioned in the Introduction, such a transformation has also been used in other settings. The innovated transformation (3) is, however, not directly applicable for large precision matrix estimation because the transformed vector $\widetilde{\mathbf{x}}$ is unobservable. Estimating $\widetilde{\mathbf{x}}$ by the two parts according to (3) is infeasible since it depends on the unknown precision matrix $\boldsymbol{\Omega}$ which is our estimation target. We overcome this difficulty by breaking the long vector $\widetilde{\mathbf{x}}$ into small subvectors and then estimating each one as a whole with the representation (3), which we describe in detail as follows.

We start with introducing some notation that will be used repeatedly in our presentation. For any subsets $A, B \subset \{1, \ldots, p\}$, denote by $\mathbf{x}_A$ a subvector of $\mathbf{x}$ formed by its components with indices in A, and $\boldsymbol{\Omega}_{A,B} = (\omega_{jk})_{j \in A, k \in B}$ a submatrix of $\boldsymbol{\Omega}$ with rows in A and columns in B. We also use the shorthand notation $\boldsymbol{\Omega}_A$ for $\boldsymbol{\Omega}_{A,A}$ for convenience. Note that by the definition of $\widetilde{\mathbf{x}}$, we can write the subvector $\widetilde{\mathbf{x}}_A$ in the following form:

$$\widetilde{\mathbf{x}}_A = \boldsymbol{\Omega}_{A,A} \boldsymbol{\eta}_A, \tag{5}$$

where $\boldsymbol{\eta}_A = \mathbf{x}_A + \boldsymbol{\Omega}_{A,A}^{-1} \boldsymbol{\Omega}_{A,A^c} \mathbf{x}_{A^c}$ with A^c the complement of set A.

The estimation of the two terms on the right-hand side of (5) is interrelated and can be achieved simultaneously and effectively through linear regression techniques. The essence of our proposal comes from a simple yet useful fact in Gaussian graphical model theory. Recall that in the Gaussian graphical model (1), it holds for any subset $A \subset \{1, \ldots, p\}$ that

$$\mathbf{x}_A | \mathbf{x}_{A^c} \sim N\left(-\boldsymbol{\Omega}_{A,A}^{-1} \boldsymbol{\Omega}_{A,A^c} \mathbf{x}_{A^c}, \boldsymbol{\Omega}_{A,A}^{-1}\right). \tag{6}$$

The conditional distribution (6) suggests a multivariate linear regression model

$$\mathbf{x}_A = \mathbf{C}_A^T \mathbf{x}_{A^c} + \boldsymbol{\eta}_A, \tag{7}$$

where $\mathbf{C}_A = -\boldsymbol{\Omega}_{A^c,A} \boldsymbol{\Omega}_{A,A}^{-1}$ is a matrix of regression coefficients, and $\boldsymbol{\eta}_A$ is the vector of model errors which takes the form introduced in (5) and has a multivariate Gaussian distribution $N(\mathbf{0}, \boldsymbol{\Omega}_{A,A}^{-1})$.

The representation of the subvector $\mathbf{x}_A$ in (7) suggests that regression techniques can be exploited to estimate the unknown subvector $\widetilde{\mathbf{x}}_A$. To see this, let $\widehat{\boldsymbol{\eta}}_A$ be the residual vector obtained by using some regression technique to fit model (7). Then the unknown matrix $\boldsymbol{\Omega}_{A,A}$ can be estimated as the inverse of the sample covariance matrix of the model residual vector $\widehat{\boldsymbol{\eta}}_A$. Denote by $\widehat{\boldsymbol{\Omega}}_A$ the resulting estimator. Then we can estimate the subvector $\widetilde{\mathbf{x}}_A$ in (5) as $\widehat{\mathbf{x}}_A = \widehat{\boldsymbol{\Omega}}_A \widehat{\boldsymbol{\eta}}_A$.

Let $(A_l)_{l=1}^L$ be a partition of the index set $\{1, \ldots, p\}$, that is, $\bigcup_{l=1}^L A_l = \{1, \ldots, p\}$ and $A_l \cap A_m = \varnothing$ for any $1 \le l \ne m \le L$. Although the ideas of our approach are applicable to the case of general $|A_l|$, to simplify the presentation we focus our attention on the case of $|A_l| = 2$ when the number of nodes p is

even, and the case of $|A_l| = 2$ or 3 when p is an odd number. Without loss of generality, throughout the paper we consider the specific partition $A_l = \{2l - 1, 2l\}$ for $1 \leq l \leq L - 1$ and $A_L = \{2L - 1, \ldots, p\}$ with $L = \lfloor p/2 \rfloor$ the integer part of $p/2$. The ISEE repeats the above procedure for each A_l with $1 \leq l \leq L$ to obtain estimated subvectors $\widehat{\mathbf{x}}_{A_l}$'s, and then stacks all these subvectors together to form an estimate $\widehat{\mathbf{x}}$ of the oracle innovated vector $\widetilde{\mathbf{x}} = \mathbf{\Omega}\mathbf{x}$. By doing so, the problem of estimating the precision matrix based on the original vector $\mathbf{x}$ reduces to that of estimating the covariance matrix based on the estimated transformed vector $\widehat{\mathbf{x}}$.

By its nature, the ISEE breaks large-scale precision matrix estimation into smaller-scale linear regression problems, each of which can be solved effectively and efficiently. Thanks to the scalability of ISEE, it has advantages over existing methods in estimating very large precision matrices. Detailed comparisons of ISEE with existing methods are given in Section 2.4.

2.3. *Estimation procedure by ISEE.* We now discuss in detail the implementation of the ISEE procedure. To ease the presentation, we introduce some matrix notation. Denote by $\mathbf{X} = (\mathbf{x}_1, \ldots, \mathbf{x}_n)^T$ the $n \times p$ data matrix. We refer to the innovated data matrix

$$(8) \qquad \widetilde{\mathbf{X}} = (\widetilde{\mathbf{x}}_1, \ldots, \widetilde{\mathbf{x}}_n)^T = \mathbf{X}\mathbf{\Omega}$$

as the oracle empirical matrix, which is unavailable to practitioners. Using matrix notation, the multivariate linear regression model (7) can be written as

$$(9) \qquad \mathbf{X}_A = \mathbf{X}_{A^c}\mathbf{C}_A + \mathbf{E}_A,$$

where $\mathbf{X}_A$ and $\mathbf{X}_{A^c}$ are the submatrices of $\mathbf{X}$ with columns in A and its complement A^c, respectively, and $\mathbf{E}_A$ is an $n \times |A|$ model error matrix with rows as i.i.d. copies of $\boldsymbol{\eta}_A^T$. Then the corresponding submatrix $\widetilde{\mathbf{X}}_A$ can be written as

$$(10) \qquad \begin{aligned} \widetilde{\mathbf{X}}_A = (\mathbf{X}\mathbf{\Omega})_A &= \mathbf{X}_A\mathbf{\Omega}_{A,A} + \mathbf{X}_{A^c}\mathbf{\Omega}_{A^c,A} \\ &= (\mathbf{X}_A + \mathbf{X}_{A^c}\mathbf{\Omega}_{A^c,A}\mathbf{\Omega}_{A,A}^{-1})\mathbf{\Omega}_{A,A} = \mathbf{E}_A\mathbf{\Omega}_{A,A}. \end{aligned}$$

The representation in (10) provides the foundation for the estimation of the oracle empirical matrix $\widetilde{\mathbf{X}}$.

Many existing methods can be used to fit the Gaussian linear regression model (7) and obtain the estimates for $\mathbf{\Omega}_{A,A}$ and $\boldsymbol{\eta}_A$. To avoid the issue of overfitting caused by high dimensionality, some kind of regularization, however, needs to be applied to control model complexity. There is a large body of literature on regularization methods; see, for example, [8, 15, 17, 33, 41, 46, 48], among many others. See also [20] for the connections and differences for a wide class of regularization methods in high dimensions, and [32] for characterizations of the impacts of high dimensionality in finite samples. For our implementation, we suggest to use the scaled Lasso method proposed in [40]. We opt to work with this method for two main reasons. First, scaled Lasso is a natural likelihood-based extension of

the Lasso [41] that is tuning free and admits efficient implementation; see (12) for details about its tuning-free feature. The efficient implementation of scaled Lasso greatly reduces the computational cost of ISEE. Second, as seen from (5), we are interested in the prediction property (i.e., the estimation of $\boldsymbol{\eta}_A$) instead of the variable selection property (i.e., the estimation of $\mathbf{C}_A$) when fitting (7). The sampling properties of scaled Lasso as revealed in [40] guarantee the accuracy in estimating $\boldsymbol{\Omega}_{A,A}$ and $\boldsymbol{\eta}_A$, and thus the scaled Lasso is sufficient for our purpose. We also remark that alternatively one can also exploit regularization methods for multivariate linear regression models instead of fitting one response at a time as in the scaled Lasso.

For each node j in the index set A, let us consider the univariate linear regression model for response $\mathbf{X}_j$, which is the jth column of the data matrix $\mathbf{X}$, given by the multivariate linear regression model (9)

$$(11) \qquad \mathbf{X}_j = \mathbf{X}_{A^c}\boldsymbol{\beta}_j + \mathbf{E}_j,$$

where the $(p - |A|)$-dimensional vector $\boldsymbol{\beta}_j$ is the column of the regression coefficient matrix $\mathbf{C}_A$ corresponding to node j and the n-dimensional error vector $\mathbf{E}_j$ is the corresponding column of the error matrix $\mathbf{E}_A$. In model (11), node j is regressed on all nodes in the complement set A^c. As mentioned before, in contrast to the conventional setting, our object of interest now is on the error vector $\mathbf{E}_j$, instead of directly on the regression coefficient vector $\boldsymbol{\beta}_j$. Thus we treat the regression coefficient vector $\boldsymbol{\beta}_j$ as a nuisance parameter, and estimate it along with the error standard deviation using the penalized least squares with the scaled Lasso

$$(12) \qquad (\widehat{\boldsymbol{\beta}}_j, \widehat{\theta}_j^{1/2}) = \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^{p-|A|}, \sigma \geq 0} \left\{ \frac{\|\mathbf{X}_j - \mathbf{X}_{A^c}\boldsymbol{\beta}\|_2^2}{2n\sigma} + \frac{\sigma}{2} + \lambda \|\boldsymbol{\beta}_*\|_1 \right\},$$

where $\boldsymbol{\beta}_*$ is the Hadamard (componentwise) product of two $(p - |A|)$-dimensional vectors $\boldsymbol{\beta}$ and $(n^{-1/2}\|\mathbf{X}_k\|_2)_{k \in A^c}$ with $\mathbf{X}_k$ the kth column of $\mathbf{X}$, $\lambda \geq 0$ is a regularization parameter associated with the weighted L_1-penalty, and $\|\mathbf{v}\|_q$ denotes the L_q-norm of a given vector $\mathbf{v}$ for $q \geq 1$. Here, the minimizer $\widehat{\theta}_j^{1/2}$, which is over σ, provides an estimator of the error standard deviation $\theta_j^{1/2} = \mathrm{var}^{1/2}(\eta_j)$, where η_j is a component of $\boldsymbol{\eta}_A$ corresponding to node j. The tuning-free feature of the scaled Lasso is entailed by the fact that the theoretical choice of the regularization parameter $\lambda = C\{(2 \log p)/n\}^{1/2}$ with $C > 1$ some constant which can be made free of the noise level in the linear regression model; see [40] for more details. Hereafter, we fix such a universal choice of λ for scaled Lasso in (12), and discuss an automatic empirical choice for λ, which is indeed tuning free, in Section 5.1. The use of the scale vector $(n^{-1/2}\|\mathbf{X}_k\|_2)_{k \in A^c}$ amounts to rescaling each column of the design matrix $\mathbf{X}_{A^c}$ to have L_2-norm $n^{1/2}$, matching that of the constant covariate $\mathbf{1}$ for the intercept, which is standard in the studies for regularization methods.

Based on the regression step, for each node j in the index set A we define

$$(13) \qquad \widehat{\mathbf{E}}_j = \mathbf{X}_j - \mathbf{X}_{A^c}\widehat{\boldsymbol{\beta}}_j \quad \text{and} \quad \widehat{\mathbf{E}}_A = (\widehat{\mathbf{E}}_j)_{j \in A},$$

where $\widehat{\boldsymbol{\beta}}_j$ is defined in (12) and $\widehat{\mathbf{E}}_A$ is an $n \times |A|$ matrix consisting of columns $\widehat{\mathbf{E}}_j$ with nodes j in the index set A. Clearly, the residual vector $\widehat{\mathbf{E}}_j$ is a natural estimate of the error vector $\mathbf{E}_j$, and thus $\widehat{\mathbf{E}}_A$ is a natural estimate of the error matrix $\mathbf{E}_A$. In view of (9) and (7), the $|A| \times |A|$ matrix $n^{-1}\widehat{\mathbf{E}}_A^T\widehat{\mathbf{E}}_A$ is a natural estimator of the error covariance matrix $\boldsymbol{\Omega}_{A,A}^{-1}$. This observation motivates us to construct a natural estimator

$$(14) \qquad \widehat{\boldsymbol{\Omega}}_A = \left(n^{-1}\widehat{\mathbf{E}}_A^T\widehat{\mathbf{E}}_A\right)^{-1}$$

for the principal submatrix $\boldsymbol{\Omega}_{A,A}$ of the precision matrix $\boldsymbol{\Omega}$ given by the index set A. These observations suggest a simple plug-in estimator $\widehat{\mathbf{E}}_A\widehat{\boldsymbol{\Omega}}_A$ for the unobservable submatrix $\widetilde{\mathbf{X}}_A$ in (10).

When A ranges over a partition $(A_l)_{l=1}^L$ of the index set $\{1, \ldots, p\}$, the ISEE estimates the oracle empirical matrix $\widetilde{\mathbf{X}}$ as the $n \times p$ matrix

$$(15) \qquad \widehat{\mathbf{X}} = (\widehat{\mathbf{X}}_{A_l})_{1 \le l \le L},$$

where the submatrix of $\widehat{\mathbf{X}}$ with columns in the index set A_l is given by $\widehat{\mathbf{X}}_{A_l} = \widehat{\mathbf{E}}_{A_l}\widehat{\boldsymbol{\Omega}}_{A_l}$ as constructed before. Then the ISEE proceeds as follows:

(a) (*Recovery of graph*). First, calculate the initial ISEE estimator as the sample covariance matrix

$$(16) \qquad \widehat{\boldsymbol{\Omega}}_{\text{ISEE,ini}} = n^{-1}\widehat{\mathbf{X}}^T\widehat{\mathbf{X}}.$$

Then for a given threshold $\tau \ge 0$, define

$$(17) \qquad \widehat{\boldsymbol{\Omega}}_{\text{ISEE},g} = T_\tau(\widehat{\boldsymbol{\Omega}}_{\text{ISEE,ini}}),$$

where $T_\tau(\mathbf{B}) = (b_{jk}1_{\{|b_{jk}|\ge\tau\}})$ denotes the matrix $\mathbf{B} = (b_{jk})$ thresholded at τ. Estimate the graphical structure E, the set of links, as $\widehat{E}_{\text{ISEE}} = \text{supp}(\widehat{\boldsymbol{\Omega}}_{\text{ISEE},g})$.

(b) (*Estimation of link strength*). For each link (j, k) in the recovered graph $\widehat{E}_{\text{ISEE}}$ with nodes j and k from different index sets A_l's, update the corresponding entry of $\widehat{\boldsymbol{\Omega}}_{\text{ISEE},g}$ as the off-diagonal entry of the 2×2 matrix $\widehat{\boldsymbol{\Omega}}_{A_l}$ given in (14) with A_l replaced by $\{j, k\}$. This yields a refined sparse precision matrix estimator $\widehat{\boldsymbol{\Omega}}_{\text{ISEE}}$ for the link strength.

We refer to the former $\widehat{\boldsymbol{\Omega}}_{\text{ISEE},g}$ as the ISEE estimator for the graph, and the latter $\widehat{\boldsymbol{\Omega}}_{\text{ISEE}}$ as the ISEE estimator with refinement throughout the paper. In particular, it is easy to see that the principal submatrix of the initial ISEE estimator $\widehat{\boldsymbol{\Omega}}_{\text{ISEE,ini}}$ given by each index set A_l is simply the matrix $\widehat{\boldsymbol{\Omega}}_{A_l}$ given in (14).

The choice of the threshold τ in (17) is important for practical implementation. We adapt the cross-validation method proposed in [2, 3] for large covariance matrix estimation. Specifically, we randomly split the sample of n rows of the estimated oracle empirical matrix $\widehat{\mathbf{X}}$ into two subsamples of sizes n_1 and n_2, and repeat this N_1 times. Denote by $\widehat{\boldsymbol{\Omega}}_{\text{ISEE,ini}}^{1,v}$ and $\widehat{\boldsymbol{\Omega}}_{\text{ISEE,ini}}^{2,v}$ the corresponding sample

covariance matrices as defined in (16) based on these two subsamples, respectively, for the νth split. The threshold τ can be chosen to minimize

$$R(\tau) = N_1^{-1} \sum_{\nu=1}^{N_1} \| T_\tau(\widehat{\boldsymbol{\Omega}}_{\text{ISEE,ini}}^{1,\nu}) - \widehat{\boldsymbol{\Omega}}_{\text{ISEE,ini}}^{2,\nu} \|^2, \tag{18}$$

where $\| \cdot \|$ denotes the Frobenius norm of a given matrix.

2.4. *Comparisons with existing methods.* The ISEE is closely related to the methods of precision matrix estimation proposed in [38] and [44] in that all three methods are rooted in the regression formulation (7). For each $1 \le j \ne k \le p$, the ANT method in [38] estimates the (j, k)-entry of $\boldsymbol{\Omega}$ using the off-diagonal entry of $\widehat{\boldsymbol{\Omega}}_A$ defined in (14) with $A = \{j, k\}$. Thus, ANT needs to conduct $O(p^2)$ scaled Lasso regressions and can become more computationally expensive for large p. Based on the observation that the jth column of $\boldsymbol{\Omega}$ can be written as $(\boldsymbol{\Omega}_{A,A}, -\boldsymbol{\Omega}_{A,A} \mathbf{C}_A^T)^T$ with $A = \{j\}$ and $\mathbf{C}_A$ defined in (7), [44] proposed to exploit the Danzig selector [8] to estimate $\mathbf{C}_A$ and used a similar method as in ISEE to estimate $\boldsymbol{\Omega}_{A,A}$. So it is seen that both ISEE and ANT rely on the residual vector in the regression model (7), while the method in [44] relies on both the residual vector and the regression coefficient vector $\mathbf{C}_A$ whose estimation can suffer from the bias issue related to the Dantzig selector. In addition, the method in [44] requires to select a tuning parameter for each node and is thus more demanding in tuning.

The ISEE is also related to the neighborhood selection method in [35] and joint estimation method in [36] in the sense that all methods estimate the graph via Lasso-type regressions. The main difference between the methods in [35] and [36] is that the former conducts p nodewise Lasso regressions for graph recovery and needs tuning parameter selection for each node, while the latter exploits a single joint Lasso regression for precision matrix estimation with only one tuning parameter. Both methods in [35] and [36] require the irrepresentable-type condition for consistent graph recovery which can become stringent in large precision matrix estimation. The Lasso regularization for precision matrix estimation has also been exploited in [47], who proposed a Lasso penalized D-trace procedure in which the Lasso penalty is applied to a new quadratic loss with a positive-definiteness constraint. As a result, the obtained estimator enjoys the nice property of positive definiteness. The sparse recovery property was also established under the irrepresentable-type condition.

The CLIME [5] is another popularly used method for precision matrix estimation. It estimates the graphical structure node by node using a Dantzig selector type procedure. For each node, a tuning parameter needs to be selected. As pointed out in [38], in order to ensure consistency in graph recovery CLIME needs an additional threshold that depends on the L_1-norm of the true precision matrix $\boldsymbol{\Omega}$, which is unknown and can be large.

As mentioned in the Introduction, the penalized likelihood (e.g., [13, 22, 37, 45]) is a group of widely used methods for precision matrix estimation. In general, these methods are not scalable due to the complexity of the likelihood function. Thus they can be computationally expensive when the scale of the problem becomes large.

In summary, compared to those existing methods, the ISEE enjoys easy tuning and is scalable. As shown later in Section 3, it also has nice asymptotic properties under mild regularity conditions. We will also provide numerical comparisons of ISEE with some popularly used methods in Section 5.

2.5. *Computation.* In the new era of big data, designing procedures with scalability is key to powering contemporary applications. The ISEE method is naturally scalable since the main computational cost comes from the construction of the estimate $\widehat{\mathbf{X}}$ for the oracle empirical matrix $\widetilde{\mathbf{X}}$. Such an $n \times p$ matrix is constructed by running p penalized linear regression fittings. These univariate response problems and the use of different permutations of the set of nodes $\{1, \ldots, p\}$, which can help boost the power of detecting important links, are perfect for parallel and distributed computing. The nodes j in the same index set A can be allocated to a common processor. These computational advantages of ISEE make it ideal for cloud computing which becomes more prevalent nowadays, and thus appealing for uncovering ultra-large sparse graphs with big data.

3. Asymptotic efficiency of innovated scalable efficient estimation.

3.1. *Technical conditions.* For the technical analysis, we focus on the class of K-sparse Gaussian graphs with spectrum constraint

(19)
$$\mathcal{G}(M, K) = \big\{ \mathbf{\Omega} : \text{each row has at most } K \text{ nonzero off-diagonal}$$
$$\text{entries and } M^{-1} \leq \lambda_{\min}(\mathbf{\Omega}) \leq \lambda_{\max}(\mathbf{\Omega}) \leq M \big\},$$

where K is some positive integer that can grow with dimensionality p, $M \geq 1$ is some constant, and $\lambda_{\min}(\cdot)$ and $\lambda_{\max}(\cdot)$ denote the smallest and largest eigenvalues of a given symmetric matrix, respectively. For each graph in class (19), the number of links for each node is bounded by K from above and the precision matrix $\mathbf{\Omega}$ has bounded spectrum. A generalized concept of sparsity is considered in [38] to allow for the case when a portion of the links can be weak, that is, close to zero but not exactly zero. To simplify the technical presentation, we content ourselves with the class of K-sparse Gaussian graphs. For notational simplicity, all rates of convergence involving $\log p$ and probability bounds involving p are understood implicitly with p regarded as $\max(p, n)$. For each index set $S \subset \{1, \ldots, p\}$, denote by $\mathbf{u}_S$ and $\mathbf{u}_{S^c}$ the subvectors of $\mathbf{u} \in \mathbb{R}^p$ with components in S and its complement S^c, respectively.

CONDITION 1. *The Gaussian graph* (1) *belongs to class* $\mathcal{G}(M, K)$ *with* $K \leq c_0 n / (\log p)$ *for some sufficiently small constant* $c_0 > 0$, *the partition* $(A_l)_{l=1}^L$ *satisfies* $1 \leq \min_l |A_l| \leq \max_l |A_l| = O(1)$, *and* $\lambda = (1 + \varepsilon)\{2\delta(\log p)/n\}^{1/2} = o(1)$ *for any constants* $\delta \geq 2$ *and* $\varepsilon > 0$.

CONDITION 2. *There exist some constants* $0 \leq \alpha \leq 1/2$ *and* $\xi > 1$ *such that the* L_∞-*norm cone invertibility factor*

$$F_\infty = \inf\left\{\frac{\|\mathbf{\Sigma u}\|_\infty}{\|\mathbf{u}\|_\infty} : \|\mathbf{u}_{S^c}\|_1 \leq \xi\|\mathbf{u}_S\|_1 \neq 0\right.$$

(20)

$$\left. \textit{for some } S \subset \{1, \ldots, p\} \textit{ with } |S| \leq O(K) \right\}$$

of the covariance matrix $\mathbf{\Sigma} = \mathbf{\Omega}^{-1}$ *satisfies* $F_\infty^{-1} = O(K^\alpha)$.

PROPOSITION 1. *For any* $\mathbf{\Omega} \in \mathcal{G}(M, K)$, *it holds that* $\inf\{\|\mathbf{\Sigma u}\|_\infty / \|\mathbf{u}\|_\infty : \mathbf{u} \neq \mathbf{0}\} \geq (K + 1)^{-1/2} M^{-1}$ *with* $\mathbf{\Sigma} = \mathbf{\Omega}^{-1}$.

Condition 1 assumes the sparsity of the precision matrix and imposes an upper bound on the sparsity level K. The assumption of $\max_l |A_l| = O(1)$ is made to simplify the technical presentation and can be relaxed. Condition 2 puts a constraint on the cone invertibility factor F_∞. Proposition 1 above shows that the constant α in Condition 2 is indeed bounded from above by $1/2$. See, for example, [43] and [40] for more discussions on the cone invertibility factors under various norms. We remark that only Conditions 1 and 2 are needed for the theoretical development of ISEE approach alone.

3.2. *Main results.* Our first theorem establishes the entrywise infinity norm estimation bound for the initial ISEE estimator.

THEOREM 1. *Assume that Conditions 1–2 hold and* $K^{1+\alpha}\lambda = o(1)$. *Then with probability* $1 - o\{p^{-(\delta-2)}\}$ *tending to one the initial ISEE estimator* $\widehat{\mathbf{\Omega}}_{\text{ISEE,ini}}$ *in* (16) *satisfies that*

(21)
$$\|\widehat{\mathbf{\Omega}}_{\text{ISEE,ini}} - \mathbf{\Omega}\|_\infty = O(K^\alpha \lambda),$$

where $\|\cdot\|_\infty$ *denotes the entrywise* L_∞-*norm of a given matrix.*

From the proof of Theorem 1, we see that the rate of convergence for the initial ISEE estimator is the maximum of two components $O\{\max(K\lambda^2, \lambda)\}$ and $O(K^\alpha \lambda)$, corresponding to the block-diagonal and off-block-diagonal entries of $\widehat{\mathbf{\Omega}}_{\text{ISEE,ini}}$, respectively. Note that the block-diagonal entries are estimated directly from (14), while most of the off-block-diagonal ones are estimated from the cross

product terms $n^{-1}\widehat{\mathbf{X}}_{A_k}^T \widehat{\mathbf{X}}_{A_l}$ with $k \neq l$. The difference in the two estimation procedures results in the difference in two rates of convergence. Since it is assumed in Theorem 1 that $K^{1+\alpha}\lambda = o(1)$ with $\alpha \geq 0$, the rate of convergence $O(K^\alpha \lambda)$ dominates that of $O\{\max(K\lambda^2, \lambda)\}$, meaning that the block-diagonal entries are generally estimated more accurately than the off-block-diagonal ones.

As introduced in Section 2.3, we apply thresholding to obtain the ISEE estimator for the graph $\widehat{\mathbf{\Omega}}_{\mathrm{ISEE},g}$ defined in (17). For each identified link (j, k) in the recovered graph $\widehat{E}_{\mathrm{ISEE}} = \mathrm{supp}(\widehat{\mathbf{\Omega}}_{\mathrm{ISEE},g})$, the ISEE estimator with refinement $\widehat{\mathbf{\Omega}}_{\mathrm{ISEE}}$ updates its corresponding entry as the off-diagonal entry of the 2×2 matrix $\widehat{\mathbf{\Omega}}_{A_l}$ given in (14) with $A_l = \{j, k\}$. The following theorem shows that both sparse precision matrix estimators $\widehat{\mathbf{\Omega}}_{\mathrm{ISEE},g}$ and $\widehat{\mathbf{\Omega}}_{\mathrm{ISEE}}$ enjoy nice asymptotic properties.

THEOREM 2. *Assume that the conditions of Theorem 1 hold and $\omega_0 = \min\{|\omega_{jk}| : (j, k) \in \mathrm{supp}(\mathbf{\Omega})\} \geq \omega_0^* = CK^\alpha \lambda$ with $C > 0$ some sufficiently large constant. Then with probability $1 - o\{p^{-(\delta-2)}\}$ tending to one, it holds simultaneously that*:

(a) (Graph recovery). $\mathrm{supp}(\widehat{\mathbf{\Omega}}_{\mathrm{ISEE},g}) = \mathrm{supp}(\mathbf{\Omega})$ *for any $\tau \in [c\omega_0^*, \omega_0 - c\omega_0^*]$ with $0 < c < 1/2$ some constant*;

(b) (Graph screening). $\mathrm{supp}(\mathbf{\Omega}) \subset \mathrm{supp}(\widehat{\mathbf{\Omega}}_{\mathrm{ISEE},g})$ *for threshold τ chosen by cross-validation* (18) *with n_1/n_2 bounded away from 0 and ∞*;

(c) (Efficient estimation)

$$\|\widehat{\mathbf{\Omega}}_{\mathrm{ISEE}} - \mathbf{\Omega}\|_\infty = O(\lambda). \tag{22}$$

The first part of results in Theorem 2 is more of theoretical interest since the quantities ω_0 and ω_0^* are generally unknown in practice. The second part provides a theoretical backup for a fast practical approach to choosing threshold in large graph screening. Comparing (22) to (21), it is seen that ISEE with refinement has an improved rate of convergence for precision matrix estimation when $\alpha > 0$. Such an improvement occurs because the off-block-diagonal entries are estimated more accurately in the refinement step. We remark that the bound in (22) is obtained as $O\{\max(K\lambda^2, \lambda)\}$ which becomes $O(\lambda)$ since $K^{1+\alpha}\lambda = o(1)$ and $\alpha \geq 0$.

3.3. *A bias corrected initial ISEE estimator.* A comparison of the rates of convergence in Theorems 1 and 2 shows that the initial ISEE estimator is generally biased when $\alpha > 0$. As mentioned in the discussion after Theorem 1, such a bias stems from the estimation of the off-block-diagonal entries of the precision matrix $\mathbf{\Omega}$ using the cross product terms $n^{-1}\widehat{\mathbf{X}}_{A_k}^T \widehat{\mathbf{X}}_{A_l}$ with $k \neq l$. Motivated by the technical analysis of the initial ISEE estimator $\widehat{\mathbf{\Omega}}_{\mathrm{ISEE},\mathrm{ini}}$, we now define a bias corrected initial ISEE estimator $\widehat{\mathbf{\Omega}}_{\mathrm{ISEE},\mathrm{cini}}$ as

$$(\widehat{\mathbf{\Omega}}_{\mathrm{ISEE},\mathrm{cini}})_{A_l,A_l} = (\widehat{\mathbf{\Omega}}_{\mathrm{ISEE},\mathrm{ini}})_{A_l,A_l} \tag{23}$$

and

$$(24) \qquad (\widehat{\boldsymbol{\Omega}}_{\text{ISEE,cini}})_{A_l, A_m} = -\big[(\widehat{\boldsymbol{\Omega}}_{\text{ISEE,ini}})_{A_l, A_m} + \widehat{\mathbf{C}}_{A_l}^{A_m} \widehat{\boldsymbol{\Omega}}_{A_l} + \widehat{\mathbf{C}}_{A_m}^{A_l} \widehat{\boldsymbol{\Omega}}_{A_m}\big]$$

for each $1 \le l \ne m \le L$, where $\widehat{\mathbf{C}}_{A_l} = (\widehat{\boldsymbol{\beta}}_{j,l})_{j \in A_l}$ represents a $(p - |A_l|) \times |A_l|$ matrix of estimated regression coefficients with $\widehat{\boldsymbol{\beta}}_{j,l}$ as defined in (12), $\widehat{\mathbf{C}}_{A_l}^{A_m}$ denotes a submatrix of $\widehat{\mathbf{C}}_{A_l}$ consisting of rows with indices in A_m, and $\widehat{\boldsymbol{\Omega}}_{A_l}$ is given in (14). The following theorem shows that such a bias corrected precision matrix estimator indeed admits improved rate of convergence.

THEOREM 3. *Under the conditions of Theorem* 1, *the bias corrected initial ISEE estimator* $\widehat{\boldsymbol{\Omega}}_{\text{ISEE,cini}}$ *in* (23)–(24) *satisfies with probability* $1 - o\{p^{-(\delta-2)}\}$ *tending to one that*

$$(25) \qquad \|\widehat{\boldsymbol{\Omega}}_{\text{ISEE,cini}} - \boldsymbol{\Omega}\|_\infty = O(\lambda),$$

and graph recovery consistency in part (a) *of Theorem* 2, *with* $\widehat{\boldsymbol{\Omega}}_{\text{ISEE},g} = T_\tau(\widehat{\boldsymbol{\Omega}}_{\text{ISEE,cini}})$ *and* $\omega_0^* = C\lambda$ *for some sufficiently large constant* $C > 0$.

In light of Theorems 1–3, we see that both the ISEE estimator with refinement $\widehat{\boldsymbol{\Omega}}_{\text{ISEE}}$ and the bias corrected initial ISEE estimator $\widehat{\boldsymbol{\Omega}}_{\text{ISEE,cini}}$ enjoy the same rate of convergence which is generally faster than that for the initial ISEE estimator $\widehat{\boldsymbol{\Omega}}_{\text{ISEE,ini}}$ when $\alpha > 0$. We remark that our bias corrected initial ISEE estimator for the case of $|A_l| = 1$ for each $1 \le l \le L$ shares similar flavor to the bias corrected test statistics introduced in [31] for false discovery rate control in Gaussian graphical model estimation. In addition to the consistency result in Theorem 3, we can further show that the estimates for zero entries of the precision matrix $\boldsymbol{\Omega}$ can enjoy the asymptotic normality as in [31]. Due to space limitation, we do not pursue that direction in our current paper.

4. Applications of innovated scalable efficient estimation. As a byproduct, the ISEE procedure also provides a fast approach to estimating the innovated transformation (3), which is key to methods such as the ideas of multiple testing using the innovated higher criticism in [23], the optimal classification in sparse Gaussian graphic models in [18] and the interaction screening in high-dimensional quadratic discriminant analysis in [19]. With the aid of ISEE, these methods can be more effectively and efficiently applied for the analysis of big data. We next discuss some additional applications of ISEE.

4.1. *Dimension reduction.* Dimension reduction facilitates greatly large-scale data analysis by effectively reducing the intrinsic dimensions of the feature space. Among all dimension reduction approaches, the sliced inverse regression (SIR) [30] has been widely used. The SIR is based on the model

$$(26) \qquad Y = m(\boldsymbol{\beta}_1^T \mathbf{x}, \ldots, \boldsymbol{\beta}_{K_0}^T \mathbf{x}, \varepsilon),$$

where Y is the response variable, $\mathbf{x}$ is a p-dimensional covariate vector, $\boldsymbol{\beta}_1$, ..., $\boldsymbol{\beta}_{K_0}$ are unknown projection vectors with $1 \leq K_0 < p$ an unknown integer, $m : \mathbb{R}^{K_0+1} \to \mathbb{R}$ is an unknown function, and ε is the noise random variable with $E(\varepsilon|\mathbf{x}) = 0$. SIR aims at estimating the effective dimension reduction (EDR) space spanned by the EDR directions $\boldsymbol{\beta}_k$'s [24, 30]. The SIR algorithm begins with standardizing the covariate vectors by centering and rescaling using the square-root precision matrix $\boldsymbol{\Omega}^{1/2}$ of covariates, and produces an estimate of the EDR directions by multiplying the constructed eigenvectors by the same matrix.

The square-root precision matrix of covariates used in the SIR algorithm can be difficult and computationally expensive to estimate when p is much larger than n. This problem can be resolved using the innovated transformation $\widetilde{\mathbf{x}} = \boldsymbol{\Omega}\mathbf{x}$. Observe that by Theorem 3.1 in [30], the covariance matrix of $E(\widetilde{\mathbf{x}}|Y)$ is degenerate in any direction orthogonal to the linear subspace spanned by the K_0 vectors $\mathrm{cov}(\widetilde{\mathbf{x}})\boldsymbol{\Sigma}\boldsymbol{\beta}_k = \boldsymbol{\Omega}\boldsymbol{\Sigma}\boldsymbol{\beta}_k = \boldsymbol{\beta}_k$, by noting that $\boldsymbol{\Sigma}\boldsymbol{\beta}_k$ are the EDR directions for the transformed data $\widetilde{\mathbf{x}}_1, \ldots, \widetilde{\mathbf{x}}_n$. This suggests that the original EDR directions $\boldsymbol{\beta}_k$'s can be obtained by calculating the eigenvectors of $\mathrm{cov}\{E(\widetilde{\mathbf{x}}|Y)\}$. Since the oracle empirical matrix $\widetilde{\mathbf{X}} = (\widetilde{\mathbf{x}}_1, \ldots, \widetilde{\mathbf{x}}_n)^T$ can be estimated effectively and efficiently using the idea of ISEE, $\mathrm{cov}\{E(\widetilde{\mathbf{x}}|Y)\}$ can be estimated easily, and thus this alternative approach greatly reduces the computational cost of large-scale dimension reduction using SIR.

4.2. *Portfolio management.* The precision matrix also plays an pivotal role in optimal portfolio allocation. Let $\mathbf{y}_i = \boldsymbol{\mu} + \mathbf{x}_i$ be the return vector of p assets at time i, where $\boldsymbol{\mu} \in \mathbb{R}^p$ is the vector of mean returns of the p assets. Then $\boldsymbol{\Sigma} = \boldsymbol{\Omega}^{-1}$ is the covariance matrix of these p assets. Markowitz's mean-variance optimal portfolio [34] is defined as the solution to the following minimization problem:

$$(27) \qquad \min_{\boldsymbol{\xi} \in \mathbb{R}^p} \boldsymbol{\xi}^T \boldsymbol{\Omega}^{-1} \boldsymbol{\xi} \qquad \text{subject to } \boldsymbol{\xi}^T \mathbf{1} = 1 \quad \text{and} \quad \boldsymbol{\xi}^T \boldsymbol{\mu} = \gamma,$$

where $\mathbf{1}$ is a p-vector of ones and $\gamma > 0$ is the targeted return imposed on the portfolio $\boldsymbol{\xi}$. It is well known that Markowitz's optimal portfolio admits an explicit solution

$$\boldsymbol{\xi}_{\mathrm{opt}} = \frac{d_1 - \gamma d_2}{d_3 d_1 - d_2^2} \boldsymbol{\Omega}\mathbf{1} + \frac{\gamma d_3 - d_2}{d_3 d_1 - d_2^2} \boldsymbol{\Omega}\boldsymbol{\mu},$$

where $d_1 = \boldsymbol{\mu}^T \boldsymbol{\Omega}\boldsymbol{\mu}$, $d_2 = \mathbf{1}^T \boldsymbol{\Omega}\boldsymbol{\mu}$, and $d_3 = \mathbf{1}^T \boldsymbol{\Omega}\mathbf{1}$. With the ISEE estimate of the precision matrix $\boldsymbol{\Omega}$, the optimal portfolio from a large number of assets can be easily constructed.

4.3. *Multiple testing, feature screening and simultaneous confidence intervals.* Testing the significance of coefficients in a regression model is of particular importance in high dimensions, where feature selection is of interest in many applications. For simplicity, consider the linear regression model

$$(28) \qquad\qquad\qquad \mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon},$$

where $\mathbf{y}$ is an n-vector of response, $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_p)^T$ is a p-vector of regression coefficients and $\boldsymbol{\varepsilon}$ is an n-vector of i.i.d. random error with variance σ^2. There is a large literature on multiple testing with the false discovery rate (FDR) control [1]. It has been a convention to consider p marginal regression models and test each of the p marginal regression coefficients is equal to zero simultaneously. For example, [14] proposed the PFA method for high-dimensional multiple testing where the test statistics can have an arbitrary dependence structure. In contrast to testing the marginal effects of covariates, it is also interesting to test their joint effects

$$(29) \qquad H_{0j} : \beta_j = 0 \quad \text{versus} \quad H_{1j} : \beta_j \neq 0, \qquad j = 1, \ldots, p.$$

With the aid of the innovated transformation (3), the multiple testing problem (29) can be reduced to the scenario of marginal regression models linking the response and each of the p innovated covariates. To see this, note that

$$(30) \qquad n^{-1}\widetilde{\mathbf{X}}^T \mathbf{y} = n^{-1}\boldsymbol{\Omega}\mathbf{X}^T \mathbf{y} = \boldsymbol{\beta} + \widetilde{\boldsymbol{\varepsilon}},$$

where $\widetilde{\mathbf{X}} = \mathbf{X}\boldsymbol{\Omega}$ is the oracle empirical matrix and $\widetilde{\boldsymbol{\varepsilon}} = -(I_p - n^{-1}\boldsymbol{\Omega}\mathbf{X}^T\mathbf{X})\boldsymbol{\beta} + n^{-1}\boldsymbol{\Omega}\mathbf{X}^T\boldsymbol{\varepsilon}$. Observe that $n^{-1}\mathbf{X}^T\mathbf{X}$ is the sample estimate of the covariance matrix $\boldsymbol{\Sigma} = \boldsymbol{\Omega}^{-1}$. Thus, intuitively $I_p - n^{-1}\boldsymbol{\Omega}\mathbf{X}^T\mathbf{X}$ can be of a small order and thus the first term of $\widetilde{\boldsymbol{\varepsilon}}$ can also be of a small order. Similarly, the second term $n^{-1}\boldsymbol{\Omega}\mathbf{X}^T\boldsymbol{\varepsilon}$ has mean $\mathbf{0}$ and conditional covariance matrix $\text{cov}(n^{-1}\boldsymbol{\Omega}\mathbf{X}^T\boldsymbol{\varepsilon}|\mathbf{X}) = n^{-2}\sigma^2\boldsymbol{\Omega}\mathbf{X}\mathbf{X}^T\boldsymbol{\Omega}$ which can be in the order of $n^{-1}\sigma^2\boldsymbol{\Omega}$. Therefore, $\widetilde{\boldsymbol{\varepsilon}}$ can be treated similarly as a random error vector. The empirical version of (30) can be obtained by substituting $\widetilde{\mathbf{X}}$ with the estimate $\widehat{\mathbf{X}}$ in (15). The use of the innovated transformation has also been discussed in [27] to improve the performance of multiple testing using the correlation structure.

Feature screening with independence learning has been popularly used in both regression and classification problems. See, for example, [11, 16, 25], among many others. Intuitively, it is natural and appealing to exploit the joint information among the covariates. The innovated features given by the ISEE estimator pool such joint information and provide new features that can be used for ranking the importance of original features, as elucidated in (30). With the representation (30), one can also construct simultaneous confidence intervals for the p regression coefficients β_j's using asymptotic distributions or the bootstrap [10].

5. Numerical studies.

5.1. *Implementation of ISEE.* When implementing ISEE, we choose the regularization parameter λ in scaled Lasso following the suggestion of [38]; that is, we fix λ to be $B/(n - 1 + B^2)^{1/2}$, where $B = tq(1 - n^{1/2}/(2p \log p), n - 1)$ with $tq(\alpha, m)$ the αth quantile of a t-distribution with m degrees of freedom. The threshold τ is chosen adaptively using the random split method described in Section 2.3. In both our simulation study and real data analysis, we use 90% of the

sample to calculate $\widehat{\mathbf{\Omega}}_{\text{ISEE,ini}}^{1,\nu}$ and remaining 10% to calculate $\widehat{\mathbf{\Omega}}_{\text{ISEE,ini}}^{2,\nu}$, and then select τ from a grid of 20 values by minimizing the criterion (18) with the number of random splits set to $N_1 = 5$. Although the change in the computational cost of ISEE is negligible for a larger value of N_1 as discussed before, our choice of N_1 works well in empirical studies.

In our numerical studies, we observe that the ISEE estimator for the graph $\widehat{\mathbf{\Omega}}_{\text{ISEE},g}$ calculated using the above way of tuning tends to have the number of false positives very close to zero, while the number of false negatives, which can also be close to zero, tends to be slightly larger than the number of false positives. Since the first step of ISEE focuses on recovering the underlying graph, the sure screening property, that is, zero false negative and low false positives with significant probability, is desirable for this step. To reduce the number of false negatives, we borrow idea from the Bonferroni method. Specifically, we first randomly permute the columns of the $n \times p$ data matrix $\mathbf{X}$, then apply ISEE to the permuted data matrix to construct a sparse precision matrix estimator, and finally permute this sparse estimator back to obtain an estimate $\widehat{\mathbf{\Omega}}_{\text{ISEE},g}^{\pi}$ of the original precision matrix, where π denotes the corresponding permutation of $\{1, \ldots, p\}$ used in the estimate. We repeat this procedure N_2 times and construct the final estimate for the set of links of the graph as the union of supp($\widehat{\mathbf{\Omega}}_{\text{ISEE},g}^{\pi}$) over all N_2 permutations. For each identified link (j, k), we average all the nonzero estimates of ω_{jk} over the N_2 repetitions to construct its final estimate. Although this permutation method adds to the computational cost of ISEE, it reduces the number of false negatives in all our settings. Moreover, thanks to the efficiency of ISEE for each fixed permutation the computational cost of our procedure is still much lower than those of other comparison methods even after we include this additional step.

We finally remark that in the simulation study, for a fair comparison with other methods ISEE is implemented without the refinement step. Thus, the ISEE estimator in our simulation examples refers to $\widehat{\mathbf{\Omega}}_{\text{ISEE},g}$ with the aforementioned way of tuning.

5.2. *Simulation examples.*

5.2.1. *Simulation example* 1. We start with a simulation example designed to compare the computational cost and accuracy of ISEE with some popularly used methods. We generate the precision matrix $\mathbf{\Omega}$ in two steps. First, we produce a band matrix $\mathbf{\Omega}_0$ with diagonal entries being one, $\mathbf{\Omega}_0(i, i + 1) = \mathbf{\Omega}_0(i + 1, i) = 0.5$ for $i = 1, \ldots, p - 1$, and all other entries being zero. Second, we randomly permute the rows and columns of $\mathbf{\Omega}_0$ to obtain the precision matrix $\mathbf{\Omega}$. Thus, in general, the final precision matrix $\mathbf{\Omega}$ no longer has the band structure. We then sample the rows of the $n \times p$ data matrix $\mathbf{X}$ as i.i.d. copies from the multivariate Gaussian distribution $N(\mathbf{0}, \mathbf{\Omega}^{-1})$. Throughout the simulation, we fix the sample size $n = 200$ and consider a range of dimensionality p.

The methods for comparison include the Glasso [22], CLIME [5] and ANT [38]. To implement Glasso and CLIME, we use the R packages `glasso` and `scio`, respectively. Both Glasso and CLIME have one tuning parameter, which is selected using fivefold cross-validation from a grid of 10 values. The ANT is implemented using the R package `ConditionalGGM` with the tuning parameters set to the default values, and is thus tuning free. Our ISEE approach is implemented in the way described in Section 5.1.

We generate 50 data sets. For each data set, we run the four comparison methods on a PC with 8 GB RAM and Intel(R) Core(TM) i5-2500 CPU (3.30 GHz). For each method in each repetition, the CPU time (in seconds) of obtaining a sparse precision matrix estimator is recorded. Three additional performance measures: the true positive rate (TPR), false positive rate (FPR) and estimation error under the Frobenius norm are also calculated. Here, the TPR and FPR are defined as

$$\text{TPR} = \frac{\text{\# of correctly identified edges}}{\text{\# of identified edges in total}},$$

$$\text{FPR} = \frac{\text{\# of falsely identified edges}}{\text{\# of identified nonedges in total}},$$

respectively.

The comparison results are summarized in Figure 1, with the x-axis indicating dimensionality p and y-axis showing the mean values of different performance measures over 50 repetitions. To make it easier to view, CPU time (top left) and estimation error under the Frobenius norm (bottom right) are both plotted under the common logarithmic scale. Due to high computational cost, the largest values of p in our simulation for Glasso, CLIME, and ANT are 250, 500 and 100, respectively. It is seen that ISEE is computationally much more efficient than all other methods. CLIME is the second best in terms of CPU time. When $p = 500$, ISEE is about 70 times faster than CLIME on average. Moreover, the accuracy of ISEE in support recovery and estimation is also among the best.

5.2.2. *Simulation example* 2. We now test the performance of ISEE in larger scales. We generate the precision matrix in two steps. First, we create a block-diagonal matrix $\mathbf{\Omega}_0$ whose diagonal blocks are matrices of size 20. The diagonal entries of $\mathbf{\Omega}_0$ are all equal to one. For each of the block matrix, the off-diagonal entries take value 0.5 with probability 0.3 and value 0 with probability 0.7. Since the matrix generated in this way may not be symmetric or positive definite, we first symmetrize it by forcing the lower triangular matrix to equal the upper triangular matrix, and then add a diagonal matrix $c I_{20}$ for some quantity c to make the smallest eigenvalue of each block matrix equal to 0.1, where I_{20} denotes an identity matrix of size 20. It is worth mentioning that in our example, the diagonal block matrices are generated independently of each other and are thus generally different. Second, we randomly permute the rows and columns of $\mathbf{\Omega}_0$ to construct
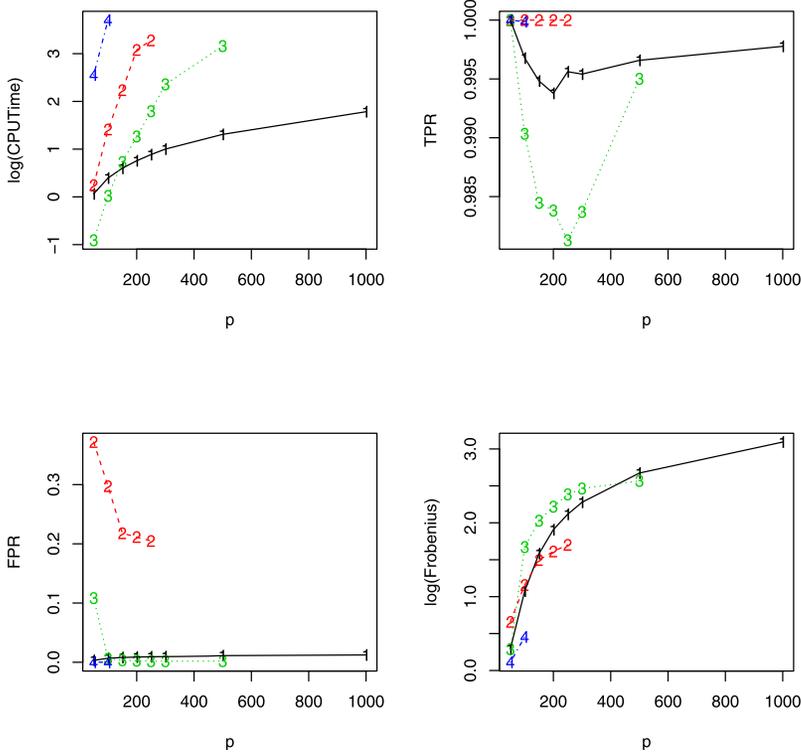
FIG. 1. *Comparisons of ISEE* (*marked as "1"*), *Glasso* (*marked as "2"*), *CLIME* (*marked as "3"*) *and ANT* (*marked as "4") in simulation example* 1. *Top left*: *the common logarithm of CPU time as a function of dimensionality p*; *Top right*: *TPR as a function of dimensionality p*; *Bottom left*: *FPR as a function of dimensionality p*; *Bottom right*: *the common logarithm of estimation error under the Frobenius norm as a function of dimensionality p*.

the final precision matrix $\boldsymbol{\Omega}$. Thus, our true precision matrix $\boldsymbol{\Omega}$ no longer has the block-diagonal structure. We consider two settings of dimensionality $p = 1000$ and 2000, with the same sample size n as in simulation example 1.

The same performance measures as in simulation example 1 are employed to evaluate the performance of ISEE. The means and standard errors over 100 repetitions are presented in Table 1, with Frob representing estimation error under the Frobenius norm. It is seen that even for these very large precision matrices, ISEE is still computationally efficient and performs well in graph recovery.

5.3. *Real data analysis*. We finally evaluate the performance of ISEE on a breast cancer data set analyzed in [26]. This data set consists of 22,283 gene expression levels of 130 breast cancer patients, among whom 33 patients had pathological complete response (pCR) and the remaining did not achieve pCR. Here,

p		Frob	TPR	FPR	CPU time
1000	Mean	3206.09	0.96799	0.05005	649.588
	SE	2.24128	0.00030	0.00006	0.70461
2000	Mean	7272.65	0.95867	0.03344	2287.34
	SE	3.42452	0.00023	0.00003	1.47541

pCR is defined as no evidence of viable, invasive tumor cells left in the surgical specimen, and thus has been regarded as a strong indicator of survival.

This breast cancer data set has been used in [5] and [13] to evaluate the accuracy of precision matrix estimation methods. We follow the steps therein to demonstrate the performance of ISEE. For completeness, we briefly list the data analysis procedure here. We first randomly split the data into training and test sets of sizes 109 and 21, respectively. Since the two classes have unbalanced sample size, a stratified sampling is used with 16 subjects randomly selected from pCR class and 5 subjects randomly selected from the other class to form the test set; the remaining subjects are used as the training set. Based on the training set, we conduct a two sample t-test and select the most significant $p = 400$ genes with the smallest p-values. We remark that both [5] and [13] kept only the most significant 110 genes. Thanks to the scalability of ISEE, we are able to deal with much larger precision matrix. We next conduct a gene-wise standardization by dividing the data matrix by the corresponding standard deviations. Then we estimate the $p \times p$ precision matrix $\boldsymbol{\Omega}$ using the ISEE approach based on the training set, and construct the linear discriminant analysis (LDA) rule on the test set. The LDA assumes that both classes have Gaussian distributions $N(\boldsymbol{\mu}_k, \boldsymbol{\Omega}^{-1})$ with different mean vectors $\boldsymbol{\mu}_1, \boldsymbol{\mu}_2$ and a common covariance matrix $\boldsymbol{\Sigma} = \boldsymbol{\Omega}^{-1}$. With the ISEE estimator $\widehat{\boldsymbol{\Omega}}_{\text{ISEE}}$ of the precision matrix, the discriminant function takes the form

$$(31) \qquad L(\mathbf{x}) = (\mathbf{x}^T - \bar{\boldsymbol{\mu}})\widehat{\boldsymbol{\Omega}}_{\text{ISEE}}\widehat{\boldsymbol{\mu}} + \log(n_1/n_2),$$

where $\bar{\boldsymbol{\mu}} = (\widehat{\boldsymbol{\mu}}_1 + \widehat{\boldsymbol{\mu}}_2)/2$ and $\widehat{\boldsymbol{\mu}} = \widehat{\boldsymbol{\mu}}_1 - \widehat{\boldsymbol{\mu}}_2$ with $\widehat{\boldsymbol{\mu}}_k$, $k = 1, 2$, the sample mean vectors, and n_1 and n_2 are the training sample sizes from classes 1 (pCR) and 2, respectively. For a new observation vector $\mathbf{x}$, LDA assigns it to class 1 if $L(\mathbf{x}) > 0$ and to class 2 otherwise. Such a procedure is repeated 100 times.

As pointed out and done in [5], an additional refit step may improve the accuracy of precision matrix estimation. We follow their suggestion and exploit a refitted ISEE estimator in calculating (31). There are different ways to refit the ISEE estimator. One option is the ISEE estimator with refinement described in Section 2.3. This approach of refitting can potentially suffer from growing computational cost for less sparse precision matrices. In our application, we adopt the

refitting procedure described in [18]. The main idea is to refit the ISEE estimator for the graph column by column after obtaining the support. Taking the first column as an example, ideally we would like to have

$$\widehat{\boldsymbol{\Sigma}} \widehat{\boldsymbol{\Omega}}(\cdot, 1) = \mathbf{e}_1, \tag{32}$$

where $\widehat{\boldsymbol{\Sigma}}$ is the sample covariance matrix, $\widehat{\boldsymbol{\Omega}}(\cdot, 1)$ denotes the first column of a precision matrix estimator $\widehat{\boldsymbol{\Omega}}$, and $\mathbf{e}_1$ is a p-vector with one in the first component and zero otherwise. Denote by $\mathcal{S} = \mathrm{supp}\{\widehat{\boldsymbol{\Omega}}(\cdot, 1)\}$ the recovered support of the first column. Then it follows from (32) that

$$\widehat{\boldsymbol{\Sigma}}_{\mathcal{S},\mathcal{S}} \widehat{\boldsymbol{\Omega}}_{\mathcal{S},\mathcal{S}}(\cdot, 1) = \mathbf{e}_{1,\mathcal{S}}. \tag{33}$$

Thus, we can refit on the support $\mathcal{S}$ by inverting the principal submatrix $\widehat{\boldsymbol{\Sigma}}_{\mathcal{S},\mathcal{S}}$ and taking out the first column, that is, $(\widehat{\boldsymbol{\Sigma}}_{\mathcal{S},\mathcal{S}})^{-1} \mathbf{e}_{1,\mathcal{S}}$. Recall that in this paper, we consider the class of sparse precision matrices $\mathcal{G}(M, K)$ with $K = O\{n/(\log p)\}$. As guaranteed by Theorem 2, ISEE enjoys nice graph recovery property, and thus the size of the support $\mathcal{S}$ can be much smaller than the sample size n with significant probability. So generally the inverse of the matrix $\widehat{\boldsymbol{\Sigma}}_{\mathcal{S},\mathcal{S}}$ can be obtained efficiently. Nevertheless, to enhance stability in real applications we suggest the use of the generalized inverse of matrix $\widehat{\boldsymbol{\Sigma}}_{\mathcal{S},\mathcal{S}}$ if $|\mathcal{S}|$ is close to or exceeds n.

To evaluate the performance of classification rule (31), we consider three measures: the specificity, sensitivity and Matthews correlation coefficient (MCC) which are defined as

$$\text{Specificity} = \frac{\text{TN}}{\text{TN} + \text{FP}}, \qquad \text{Sensitivity} = \frac{\text{TP}}{\text{TP} + \text{FN}},$$

$$\text{MCC} = \frac{\text{TP} \times \text{TN} - \text{FP} \times \text{FN}}{\sqrt{(\text{TP} + \text{FP})(\text{TP} + \text{FN})(\text{TN} + \text{FP})(\text{TN} + \text{FN})}} \tag{34}$$

with the TP, TN, FP and FN representing the true positives (pCR), true negatives, false positives and false negatives, respectively. For each of these three measures, the larger the value the better the classification performance.

As demonstrated in [11], feature selection can be crucial in high-dimensional classification since otherwise the noise accumulation caused by estimating a large number of parameters can dominate the signal, and thus deteriorate the classification power. The same phenomenon is observed in our study here. When estimating the two class mean difference vector $\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2$, we incorporate the feature selection component using the thresholded estimator $T_\tau(\widehat{\boldsymbol{\mu}})$ defined similarly as in (17). As the value of the threshold τ decreases, the number of nonzero components in $T_\tau(\widehat{\boldsymbol{\mu}})$ varies from 1 to p. Figure 2 reports the three measures defined in (34) as functions of threshold τ. To ease the presentation, we relabel the x-axis as the number of nonzero components in $T_\tau(\widehat{\boldsymbol{\mu}})$. The solid curves are the mean values across 100 repetitions and the dotted curves around them are one standard error away from the mean curves pointwise.
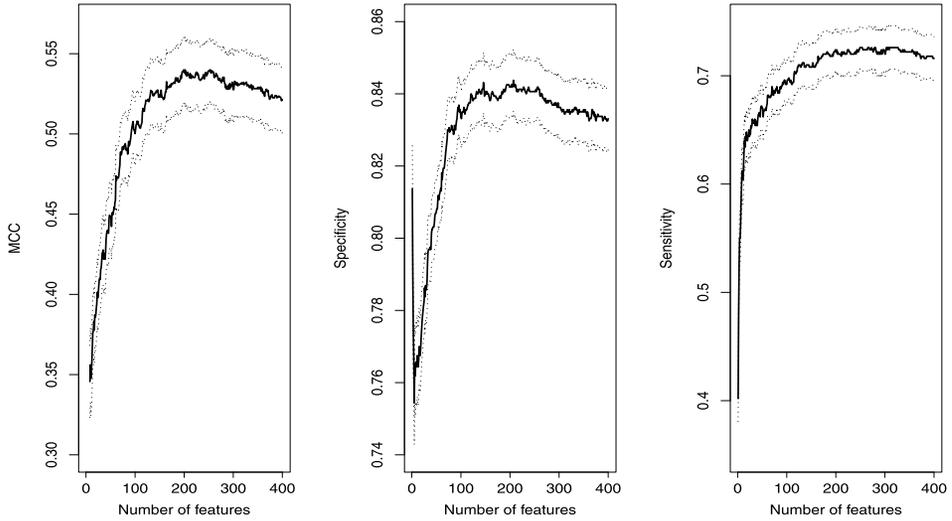
FIG. 2. *The classification results on breast cancer data with the x-axis representing the number of nonzero components in two-class mean difference estimate $T_\tau(\widehat{\boldsymbol{\mu}})$.*

The maximum value of MCC, which is 0.540, is achieved when 253 genes are used in LDA, and the corresponding standard error is 0.020. The values of MCC reported in [5] and [13] are 0.506 and 0.402, respectively, with standard error both equal to 0.020, when using only the 110 most significant genes. Our results show that using a larger number of genes and taking into account their correlation structure have potential to improve the classification results. The specificity and sensitivity reported in [13] are 0.794 (0.098) and 0.634 (0.220), respectively, with standard errors in parentheses, while the corresponding ones reported in [5] are 0.749 (0.005) and 0.806 (0.017), respectively. Comparing these results to Figure 2, it is seen that we have much improved specificity and comparable sensitivity over a large region of the threshold level τ.

6. Discussions. In this paper, we have introduced a new method ISEE for efficient estimation of ultra-large Gaussian graphs. Thanks to its scalability, ISEE provides an effective way of uncovering large sparse graphs with big data. The suggested method is ideal for parallel and distributed computing and cloud computing and has been shown to enjoy appealing theoretical properties. Both computational and theoretical advantages of ISEE have been demonstrated with empirical studies. The ISEE can further scale up along with the use of the SIS or ISIS in [16]; see Section B of the Supplementary Material [21] for detailed descriptions of such an extension as well as its theoretical properties.

It would be of interest to study several extensions of ISEE to different settings in future studies. For example, the idea of ISEE can be extended to the setting of large-scale multiple graphs comparison and estimation. Aided by ISEE, one can

estimate each graph individually and then conduct multiple testing to detect the difference and similarity of these graphs. Another possible extension of ISEE is the estimation of large latent variable Gaussian graphical models, where only a subset of the nodes are observable in practice. It is also interesting to extend ISEE to the estimation of large nonparanormal graphical models, where the original graph for the p-variate random vector $\mathbf{x}$ is nonnormal, but under some unknown nonlinear transformation $\mathbf{f}: \mathbb{R}^p \to \mathbb{R}^p$, $\mathbf{f}(\mathbf{x})$ becomes a normal random vector.

As discussed in Section 4, ISEE can be applied to such applications as dimension reduction, portfolio management and multiple testing, feature screening and simultaneous confidence intervals. It is interesting to investigate the performance of ISEE in these applications which demands future studies.

APPENDIX: PROOFS OF SOME MAIN RESULTS

We provide the proofs of Theorems 1–2 in this Appendix. The proofs of Theorem 3 and Proposition 1 and additional technical details are included in the Supplementary Material [21].

A.1. Proof of Theorem 1. Throughout the proof, we condition on the event $\mathcal{E}$ defined in (A.40) in the Supplementary Material with probability $1 - o\{p^{-(\delta-2)}\}$ tending to one, on which the bounds (A.23)–(A.25) in the Supplementary Material hold simultaneously and uniformly over all nodes j in the index sets A_l with $1 \leq l \leq L$, and the entrywise L_∞-norm bounds in (A.38) in the Supplementary Material hold uniformly as well. Observe that in view of (13)–(14) and (15)–(16), it is easy to see that the principal submatrix of the initial ISEE estimator $\widehat{\boldsymbol{\Omega}}_{\text{ISEE,ini}}$, which is the sample covariance matrix $n^{-1}\widehat{\mathbf{X}}^T\widehat{\mathbf{X}}$, given by each index set A_l is simply the matrix $\widehat{\boldsymbol{\Omega}}_{A_l}$ given in (14). Thus, the uniform entrywise L_∞-norm bound (A.29) in Lemma 2 in the Supplementary Material yields the bound

$$(35) \qquad \|n^{-1}\widehat{\mathbf{X}}_{A_l}^T\widehat{\mathbf{X}}_{A_l} - \boldsymbol{\Omega}_{A_l}\|_\infty \leq O\{\max(K\lambda^2, \lambda)\}$$

uniformly over the L blocks of principal submatrices of $\widehat{\boldsymbol{\Omega}}_{\text{ISEE,ini}}$ corresponding to the index sets A_l.

It remains to show that for each pair of index sets (A_l, A_m) with $l \neq m$, we have

$$(36) \qquad \|n^{-1}\widehat{\mathbf{X}}_{A_l}^T\widehat{\mathbf{X}}_{A_m}\|_\infty = O\{\max(K\lambda^2, K^\alpha\lambda)\}.$$

In light of (10)–(15) and (A.30) in the Supplementary Material, we have the following decomposition of the matrix:

$$
\begin{aligned}
(37) \qquad \widehat{\mathbf{X}}_{A_l} &= \mathbf{E}_{A_l}\boldsymbol{\Omega}_{A_l} + \mathbf{E}_{A_l}(\widehat{\boldsymbol{\Omega}}_{A_l} - \boldsymbol{\Omega}_{A_l}) - \mathbf{X}_{A_l^c}(\widehat{\mathbf{C}}_{A_l} - \mathbf{C}_{A_l})\widehat{\boldsymbol{\Omega}}_{A_l} \\
&= \widetilde{\mathbf{X}}_{A_l} + \mathbf{E}_{A_l}(\widehat{\boldsymbol{\Omega}}_{A_l} - \boldsymbol{\Omega}_{A_l}) - \mathbf{X}_{A_l^c}(\widehat{\mathbf{C}}_{A_l} - \mathbf{C}_{A_l})\widehat{\boldsymbol{\Omega}}_{A_l},
\end{aligned}
$$

where $\widehat{\mathbf{C}}_{A_l} = (\widehat{\boldsymbol{\beta}}_{j,l})_{j \in A_l}$ denotes a $(p - |A_l|) \times |A_l|$ matrix of estimated regression coefficients. It follows from (37) that

$$(38) \qquad n^{-1}\widehat{\mathbf{X}}_{A_l}^T\widehat{\mathbf{X}}_{A_m} = \boldsymbol{\eta}_1 + \boldsymbol{\eta}_2 + \boldsymbol{\eta}_3 + \boldsymbol{\eta}_4,$$

where the first term is $\boldsymbol{\eta}_1 = n^{-1}\widetilde{\mathbf{X}}_{A_l}^T \widetilde{\mathbf{X}}_{A_m}$, the second and third terms are $\boldsymbol{\eta}_2 = n^{-1}\widetilde{\mathbf{X}}_{A_l}^T[\mathbf{E}_{A_m}(\widehat{\boldsymbol{\Omega}}_{A_m} - \boldsymbol{\Omega}_{A_m}) - \mathbf{X}_{A_m^c}(\widehat{\mathbf{C}}_{A_m} - \mathbf{C}_{A_m})\widehat{\boldsymbol{\Omega}}_{A_m}]$ and $\boldsymbol{\eta}_3 = n^{-1}[\mathbf{E}_{A_l}(\widehat{\boldsymbol{\Omega}}_{A_l} - \boldsymbol{\Omega}_{A_l}) - \mathbf{X}_{A_l^c}(\widehat{\mathbf{C}}_{A_l} - \mathbf{C}_{A_l})\widehat{\boldsymbol{\Omega}}_{A_l}]^T \widetilde{\mathbf{X}}_{A_m}$, and the last term is $\boldsymbol{\eta}_4 = n^{-1}[\mathbf{E}_{A_l}(\widehat{\boldsymbol{\Omega}}_{A_l} - \boldsymbol{\Omega}_{A_l}) - \mathbf{X}_{A_l^c}(\widehat{\mathbf{C}}_{A_l} - \mathbf{C}_{A_l})\widehat{\boldsymbol{\Omega}}_{A_l}]^T[\mathbf{E}_{A_m}(\widehat{\boldsymbol{\Omega}}_{A_m} - \boldsymbol{\Omega}_{A_m}) - \mathbf{X}_{A_m^c}(\widehat{\mathbf{C}}_{A_m} - \mathbf{C}_{A_m})\widehat{\boldsymbol{\Omega}}_{A_m}]$. We will analyze these four terms separately.

Part 1. We start with the second and third terms $\boldsymbol{\eta}_2$ and $\boldsymbol{\eta}_3$. Since $\widetilde{\mathbf{X}}_{A_l} = \mathbf{E}_{A_l}\boldsymbol{\Omega}_{A_l}$, we can rewrite $\boldsymbol{\eta}_2$ as

$$(39) \qquad \begin{aligned} \boldsymbol{\eta}_2 &= n^{-1}\boldsymbol{\Omega}_{A_l}^T \mathbf{E}_{A_l}^T[\mathbf{E}_{A_m}(\widehat{\boldsymbol{\Omega}}_{A_m} - \boldsymbol{\Omega}_{A_m}) - \mathbf{X}_{A_m^c}(\widehat{\mathbf{C}}_{A_m} - \mathbf{C}_{A_m})\widehat{\boldsymbol{\Omega}}_{A_m}] \\ &= \mathbf{D}_1 - \mathbf{D}_2, \end{aligned}$$

where $\mathbf{D}_1 = \boldsymbol{\Omega}_{A_l}^T (n^{-1}\mathbf{E}_{A_l}^T \mathbf{E}_{A_m})(\widehat{\boldsymbol{\Omega}}_{A_m} - \boldsymbol{\Omega}_{A_m})$ and $\mathbf{D}_2 = \boldsymbol{\Omega}_{A_l}^T (n^{-1}\mathbf{X}_{A_m^c}^T \mathbf{E}_{A_l})^T \cdot (\widehat{\mathbf{C}}_{A_m} - \mathbf{C}_{A_m})\widehat{\boldsymbol{\Omega}}_{A_m}$. Note that the error matrices $\mathbf{E}_{A_l}$ and $\mathbf{E}_{A_m}$ with $l \neq m$ are independent of each other, and thus the mean of the random matrix $n^{-1}\mathbf{E}_{A_l}^T \mathbf{E}_{A_m}$ is $\mathbf{0}$. So the same concentration bound as in (A.34) in the Supplementary Material applies with $\boldsymbol{\xi}_1$ replaced by $n^{-1}\mathbf{E}_{A_l}^T \mathbf{E}_{A_m}$. Taking $t = [(\delta + 1)(\log p)/(cn)]^{1/2}$ in (A.34) and applying Bonferroni's inequality over $1 \leq l \neq m \leq L$ lead to

$$(40) \qquad P(\mathcal{E}_3) \geq 1 - p^2 \cdot O(e^{-cnt^2}) = 1 - O\{p^{-(\delta-1)}\} = 1 - o\{p^{-(\delta-2)}\},$$

where the event $\mathcal{E}_3$ is defined as

$$(41) \qquad \mathcal{E}_3 = \Big\{ \max_{1 \leq l \neq m \leq L} \|n^{-1}\mathbf{E}_{A_l}^T \mathbf{E}_{A_m}\|_\infty \leq t = O(\lambda) \Big\}.$$

From now on, we condition on the event $\mathcal{E} \cap \mathcal{E}_3$, which has the same asymptotic probability bound as $\mathcal{E}$ in view of (A.39) in the Supplementary Material and (40). As shown in the proof of Lemma 2 in the Supplementary Material, it holds that $\|\boldsymbol{\Omega}_{A_l}\|_\infty = O(1)$ and $\|\widehat{\boldsymbol{\Omega}}_{A_l}\|_\infty = O(1)$. Thus, in the event $\mathcal{E} \cap \mathcal{E}_3$, we have

$$(42) \qquad \|\mathbf{D}_1\|_\infty = O\{\lambda \max(K\lambda^2, \lambda)\}$$

in view of (41) and (A.29) in the Supplementary Material.

For the second term $\mathbf{D}_2$ in (39), consider the $|A_l| \times |A_m|$ matrix

$$(43) \qquad \mathbf{F} = (n^{-1}\mathbf{X}_{A_m^c}^T \mathbf{E}_{A_l})^T (\widehat{\mathbf{C}}_{A_m} - \mathbf{C}_{A_m}) = \mathbf{F}_1 + \mathbf{F}_2,$$

where $\mathbf{F}_1$ and $\mathbf{F}_2$ are defined through matrix multiplication by taking the rows of $n^{-1}\mathbf{X}_{A_m^c}^T \mathbf{E}_{A_l}$ and $\widehat{\mathbf{C}}_{A_m} - \mathbf{C}_{A_m}$ from nodes in index sets $A_m^c \cap A_l^c$ and A_l, respectively. In view of (A.25) and (A.23) in the Supplementary Material, we have

$$(44) \qquad \|\mathbf{F}_1\|_\infty \leq O(\lambda) \cdot O(K\lambda) = O(K\lambda^2).$$

Denote by $\mathbf{F}_3$ the $|A_l| \times |A_m|$ submatrix of $\widehat{\mathbf{C}}_{A_m} - \mathbf{C}_{A_m}$ given by rows corresponding to nodes in A_l. By Lemma 3 in the Supplementary Material, Theorem 3 of [43] applies to show that

$$(45) \qquad P(\mathcal{E}_4) = 1 - o\{p^{-(\delta-2)}\},$$

where $\mathcal{E}_4 = \{\max_{j \in A_l, 1 \le l \le L} \|\widehat{\boldsymbol{\beta}}_{j,l} - \boldsymbol{\beta}_{j,l}\|_\infty = O(K^\alpha \lambda)\}$. In view of (7)–(9), using similar arguments to those for proving (A.35) in the Supplementary Material with t chosen to be $[\delta(\log p)/(cn)]^{1/2}$ leads to

$$(46) \qquad P(\mathcal{E}_5) = 1 - o\{p^{-(\delta-2)}\},$$

where $\mathcal{E}_5 = \{\max_{1 \le l \le L} \|n^{-1}\mathbf{X}_{A_l}^T \mathbf{E}_{A_l} - \boldsymbol{\Omega}_{A_l}^{-1}\|_\infty \le O(\lambda)\}$.

Hereafter, we condition on the event $\mathcal{E} \cap (\bigcap_{3 \le i \le 5} \mathcal{E}_i)$, which has the same asymptotic probability bound as $\mathcal{E}$ in view of (A.39) in the Supplementary Material, (40), and (45)–(46). On this new event, it follows from (45)–(46) and the fact of $\|\boldsymbol{\Omega}_{A_l}^{-1}\|_\infty = O(1)$ that

$$(47) \qquad \|\mathbf{F}_2\|_\infty = \|(n^{-1}\mathbf{X}_{A_l}^T \mathbf{E}_{A_l})^T \mathbf{F}_3\|_\infty \le O(1) \cdot O(K^\alpha \lambda) = O(K^\alpha \lambda).$$

Combining (43)–(44) and (47) together with the facts of $\|\boldsymbol{\Omega}_{A_l}\|_\infty = O(1)$ and $\|\widehat{\boldsymbol{\Omega}}_{A_m}\|_\infty = O(1)$ gives

$$(48) \qquad \|\mathbf{D}_2\|_\infty = O\{\max(K\lambda^2, K^\alpha\lambda)\}.$$

Since $\boldsymbol{\eta}_3^T$ shares the same form as $\boldsymbol{\eta}_2$, putting (39), (42) and (48) together yields

$$(49) \qquad P\Big\{\max_{1 \le l \ne m \le L} \max(\|\boldsymbol{\eta}_2\|_\infty, \|\boldsymbol{\eta}_3\|_\infty) \le O\{\max(K\lambda^2, K^\alpha\lambda)\}\Big\}$$
$$= 1 - o\{p^{-(\delta-2)}\}.$$

Part 2. We next consider the fourth term $\boldsymbol{\eta}_4$. Let us decompose it into four terms as

$$(50) \qquad \boldsymbol{\eta}_4 = \mathbf{G}_1 - \mathbf{G}_2 - \mathbf{G}_3 + \mathbf{G}_4,$$

where the first term is $\mathbf{G}_1 = n^{-1}(\widehat{\boldsymbol{\Omega}}_{A_l} - \boldsymbol{\Omega}_{A_l})\mathbf{E}_{A_l}^T \mathbf{E}_{A_m}(\widehat{\boldsymbol{\Omega}}_{A_m} - \boldsymbol{\Omega}_{A_m})$, the second and third terms are $\mathbf{G}_2 = n^{-1}(\widehat{\boldsymbol{\Omega}}_{A_l} - \boldsymbol{\Omega}_{A_l})\mathbf{E}_{A_l}^T \mathbf{X}_{A_m^c}(\widehat{\mathbf{C}}_{A_m} - \mathbf{C}_{A_m})\widehat{\boldsymbol{\Omega}}_{A_m}$ and $\mathbf{G}_3 = n^{-1}\widehat{\boldsymbol{\Omega}}_{A_l}(\widehat{\mathbf{C}}_{A_l} - \mathbf{C}_{A_l})^T \mathbf{X}_{A_l^c}^T \mathbf{E}_{A_m}(\widehat{\boldsymbol{\Omega}}_{A_m} - \boldsymbol{\Omega}_{A_m})$, and the last term is $\mathbf{G}_4 = n^{-1}\widehat{\boldsymbol{\Omega}}_{A_l}(\widehat{\mathbf{C}}_{A_l} - \mathbf{C}_{A_l})^T \mathbf{X}_{A_l^c}^T \mathbf{X}_{A_m^c}(\widehat{\mathbf{C}}_{A_m} - \mathbf{C}_{A_m})\widehat{\boldsymbol{\Omega}}_{A_m}$. In view of (A.40) in the Supplementary Material and (41), we see that on the event $\mathcal{E} \cap \mathcal{E}_3$, it holds that

$$(51) \qquad \begin{aligned} \|\mathbf{G}_1\|_\infty &\le O\{\max(K\lambda^2, \lambda)\} \cdot O(\lambda) \cdot O\{\max(K\lambda^2, \lambda)\} \\ &= O\{\lambda[\max(K\lambda^2, \lambda)]^2\}. \end{aligned}$$

Note that $\mathbf{G}_2 = (\widehat{\boldsymbol{\Omega}}_{A_l} - \boldsymbol{\Omega}_{A_l})\mathbf{F}\widehat{\boldsymbol{\Omega}}_{A_m}$ in light of (43), and $\mathbf{G}_3^T$ shares the same form as $\mathbf{G}_2$. Thus on the event $\mathcal{E} \cap (\bigcap_{3 \le i \le 5} \mathcal{E}_i)$, combining (A.29) in the Supplementary Material, (43)–(44), and (47) along with the fact of $\|\widehat{\boldsymbol{\Omega}}_{A_m}\|_\infty = O(1)$ leads to

$$(52) \qquad \max(\|\mathbf{G}_2\|_\infty, \|\mathbf{G}_3\|_\infty) \le O\{\max(K\lambda^2, \lambda) \cdot \max(K\lambda^2, K^\alpha\lambda)\}.$$

For the last term $\mathbf{G}_4$, observe that an application of the Cauchy–Schwarz inequality together with bound (A.24) in the Supplementary Material entails

$$\|n^{-1}(\widehat{\mathbf{C}}_{A_l} - \mathbf{C}_{A_l})^T \mathbf{X}_{A_l^c}^T \mathbf{X}_{A_m^c}(\widehat{\mathbf{C}}_{A_m} - \mathbf{C}_{A_m})\|_\infty \leq O(K\lambda^2).$$

Since $\|\widehat{\mathbf{\Omega}}_{A_l}\|_\infty = O(1)$, it follows from the above bound that

$$(53) \qquad \|\mathbf{G}_4\|_\infty \leq O(K\lambda^2).$$

Thus, combining (50)–(53) results in

$$(54) \qquad P\Big\{\max_{1 \leq l \neq m \leq L} \|\boldsymbol{\eta}_4\|_\infty \leq O\{\max(K\lambda^2, \lambda)\}\Big\} = 1 - o\{p^{-(\delta-2)}\}.$$

Part 3. We finally consider the first term $\boldsymbol{\eta}_1$. In view of (4)–(8), $n^{-1}\widetilde{\mathbf{X}}^T\widetilde{\mathbf{X}}$ is the oracle sample covariance matrix estimator for the precision matrix $\mathbf{\Omega}$. Since $\widetilde{\mathbf{x}}$ defined in (3) is a p-variate Gaussian random vector, applying similar arguments to those for proving (A.35) in the Supplementary Material with t chosen to be $[(\delta + 1)(\log p)/(cn)]^{1/2}$ leads to

$$(55) \qquad P\{\|n^{-1}\widetilde{\mathbf{X}}^T\widetilde{\mathbf{X}} - \mathbf{\Omega}\|_\infty \leq O(\lambda)\} = 1 - o\{p^{-(\delta-2)}\},$$

which provides a uniform bound on $\boldsymbol{\eta}_1 = n^{-1}\widetilde{\mathbf{X}}_{A_l}^T\widetilde{\mathbf{X}}_{A_m}$.

Therefore, combining (38), (49) and (54)–(55) gives bound (36) uniformly over all pairs of index sets (A_l, A_m) with $l \neq m$. Observe that the order in (36) is in fact $O(K^\alpha\lambda)$ since the rate of convergence $O(K^\alpha\lambda)$ dominates that of $O(K\lambda^2)$ in light of the assumptions of $K^{1+\alpha}\lambda = o(1)$ and $\alpha \geq 0$. Then in view of (35), the proof of Theorem 1 concludes by noticing that all these uniform bounds hold simultaneously with significant probability $1 - o\{p^{-(\delta-2)}\}$.

A.2. Proof of Theorem 2. Since C in $\omega_0^* = CK^\alpha\lambda$ is some sufficiently large positive constant, Theorem 1 entails that with probability $1 - o\{p^{-(\delta-2)}\}$, it holds that

$$(56) \qquad \|\widehat{\mathbf{\Omega}}_{\text{ISEE,ini}} - \mathbf{\Omega}\|_\infty < c\omega_0^*,$$

where $c < 1/2$ is some positive constant. Thus in view of the assumption that $\omega_0 = \min\{|\omega_{jk}| : (j, k) \in \text{supp}(\mathbf{\Omega})\} \geq \omega_0^*$, by (56) we have $\text{supp}(\widehat{\mathbf{\Omega}}_{\text{ISEE},g}) \subset \text{supp}(\mathbf{\Omega})$ when $\tau \geq c\omega_0^*$, and $\text{supp}(\mathbf{\Omega}) \subset \text{supp}(\widehat{\mathbf{\Omega}}_{\text{ISEE},g})$ when $\tau \leq \omega_0 - c\omega_0^*$. This shows that $\text{supp}(\widehat{\mathbf{\Omega}}_{\text{ISEE},g}) = \text{supp}(\mathbf{\Omega})$ for any $\tau \in [c\omega_0^*, \omega_0 - c\omega_0^*]$, which proves part (a) of Theorem 2.

For part (b), we first make an important claim that the results of Theorem 1 hold for the initial ISEE estimator $\widehat{\mathbf{\Omega}}_{\text{ISEE,ini}}$ as defined in (16) but based on a subsample of n_0 rows of the estimated oracle empirical matrix $\widetilde{\mathbf{X}}$, where n_0/n is bounded away from 0. This claim follows from the same arguments as in the proof of Theorem 1, by noting that bounds (A.24)–(A.25) in the Supplementary Material hold when the subsample is used since n_0 is of the same order as n. Observe that both

n_1 and n_2 are of the same order as n by the assumption that n_1/n_2 is bounded away from 0 and ∞. Thus, with probability $1 - o\{p^{-(\delta-2)}\}$, the bound (56) also applies to both estimators $\widehat{\boldsymbol{\Omega}}_{\mathrm{ISEE,ini}}^{1,\nu}$ and $\widehat{\boldsymbol{\Omega}}_{\mathrm{ISEE,ini}}^{2,\nu}$, that is,

$$\text{(57)} \qquad \left\|\widehat{\boldsymbol{\Omega}}_{\mathrm{ISEE,ini}}^{i,\nu} - \boldsymbol{\Omega}\right\|_\infty < c\omega_0^*$$

for $i = 1, 2$.

As in [3], without loss of generality we work with the case of $N = 1$ in (18). In view of the proof for part (a), to prove the sure screening property $\mathrm{supp}(\boldsymbol{\Omega}) \subset \mathrm{supp}(\widehat{\boldsymbol{\Omega}}_{\mathrm{ISEE},g})$ it suffices to show that with probability $1 - o\{p^{-(\delta-2)}\}$, the threshold τ chosen by the cross-validation is bounded above by $\tau_0 = \omega_0 - c\omega_0^*$. Here, we use the convention that the smallest τ is preferred when the minimizer of $R(\tau)$ is not unique. To this end, we need only to show that $R(\tau) \geq R(\tau_0)$ whenever $\tau \geq \tau_0$.

Note that when $\tau = \tau_0$, we have $\mathrm{supp}\{T_\tau(\widehat{\boldsymbol{\Omega}}_{\mathrm{ISEE,ini}})\} = \mathrm{supp}(\boldsymbol{\Omega})$ by part (a), and $\mathrm{supp}\{T_\tau(\widehat{\boldsymbol{\Omega}}_{\mathrm{ISEE,ini}}^{i,\nu})\} = \mathrm{supp}(\boldsymbol{\Omega})$ for $i = 1, 2$ in light of (57). Thus, when τ increases from τ_0, the two matrices $T_\tau(\widehat{\boldsymbol{\Omega}}_{\mathrm{ISEE,ini}}^{1,\nu})$ and $T_{\tau_0}(\widehat{\boldsymbol{\Omega}}_{\mathrm{ISEE,ini}}^{1,\nu})$ can differ only over entries in $\mathrm{supp}(\boldsymbol{\Omega})$. Assume that M nonzero entries of $T_{\tau_0}(\widehat{\boldsymbol{\Omega}}_{\mathrm{ISEE,ini}}^{1,\nu})$ become zero in $T_\tau(\widehat{\boldsymbol{\Omega}}_{\mathrm{ISEE,ini}}^{1,\nu})$. Then by some simple algebra, it follows from (57) and the assumption of $\omega_0 = \min\{|\omega_{jk}| : (j,k) \in \mathrm{supp}(\boldsymbol{\Omega})\} \geq \omega_0^*$ that

$$\text{(58)} \qquad \begin{aligned} R(\tau) - R(\tau_0) &\geq M\big[(\omega_0 - c\omega_0^*)^2 - (2c\omega_0^*)^2\big] \\ &\geq M(1 - 3c)(1 + c)(\omega^*)_0^2 \geq 0, \end{aligned}$$

as long as we choose $c \leq 1/3$ in part (a). This completes the proof of part (b) of Theorem 2.

Finally for part (c), note that $\widehat{\boldsymbol{\Omega}}_{\mathrm{ISEE},g}$ from either of parts (a) and (b) satisfies that $\mathrm{supp}(\boldsymbol{\Omega}) \subset \mathrm{supp}(\widehat{\boldsymbol{\Omega}}_{\mathrm{ISEE},g})$. Using the same arguments as in the proof of Lemma 2 in the Supplementary Material with t chosen as $[(\delta + 1)(\log p)/(cn)]^{1/2}$ in (A.34), we can show that with probability $1 - o\{p^{-(\delta-2)}\}$, it holds uniformly over all pairs of nodes $A = \{j, k\}$ that $\|\widehat{\boldsymbol{\Omega}}_A - \boldsymbol{\Omega}_{A,A}\|_\infty = O\{\max(K\lambda^2, \lambda)\}$. This result along with $\mathrm{supp}(\boldsymbol{\Omega}) \subset \mathrm{supp}(\widehat{\boldsymbol{\Omega}}_{\mathrm{ISEE},g})$ yields the desired bound (22) in part (c), by noting that the order $O(\lambda)$ dominates $O(K\lambda^2)$ in view of the assumptions of $K^{1+\alpha}\lambda = o(1)$ and $\alpha \geq 0$. We conclude the proof of part (c) of Theorem 2 by showing the sure screening property which can be exploited to reduce the computational cost of the refinement step for estimating the link strength. When the ISEE estimator with refinement $\widehat{\boldsymbol{\Omega}}_{\mathrm{ISEE}}$ updates the (j, k)-entry of $\widehat{\boldsymbol{\Omega}}_{\mathrm{ISEE},g}$, two univariate linear regression models as defined in (11) with $A = \{j, k\}$ are considered for nodes j and k, respectively. In light of $\mathbf{C}_A = -\boldsymbol{\Omega}_{A^c,A}\boldsymbol{\Omega}_{A,A}^{-1}$ in model (9), it is easy to see that

$$\text{(59)} \qquad \mathrm{supp}(\boldsymbol{\beta}_j), \mathrm{supp}(\boldsymbol{\beta}_k) \subset \{m \in A^c : |\omega_{jm}| \text{ or } |\omega_{km}| \neq 0\}.$$

Denote by $\widehat{M}_{jk} = \{m \in A^c : |\widehat{\omega}_{jm}| \text{ or } |\widehat{\omega}_{km}| \neq 0\}$, where $\widehat{\boldsymbol{\Omega}}_{\text{ISEE},g} = (\widehat{\omega}_{jk})$. Thus, by (59) and $\text{supp}(\boldsymbol{\Omega}) \subset \text{supp}(\widehat{\boldsymbol{\Omega}}_{\text{ISEE},g})$, with probability $1 - o\{p^{-(\delta-2)}\}$ it holds uniformly over all pairs of nodes (j, k) that

$$(60) \qquad\qquad \text{supp}(\boldsymbol{\beta}_j), \text{supp}(\boldsymbol{\beta}_k) \subset \widehat{M}_{jk},$$

which gives the desired sure screening property for fitting model (11).

SUPPLEMENTARY MATERIAL

Supplement to "Innovated scalable efficient estimation in ultra-large Gaussian graphical models" (DOI: 10.1214/15-AOS1416SUPP; .pdf). Due to space constraints, the proofs of Theorem 3 and Proposition 1 and additional technical details are provided in the Supplementary Material [21].

REFERENCES

[1] BENJAMINI, Y. and HOCHBERG, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *J. Roy. Statist. Soc. Ser. B* **57** 289–300. MR1325392

[2] BICKEL, P. J. and LEVINA, E. (2008). Regularized estimation of large covariance matrices. *Ann. Statist.* **36** 199–227. MR2387969

[3] BICKEL, P. J. and LEVINA, E. (2008). Covariance regularization by thresholding. *Ann. Statist.* **36** 2577–2604. MR2485008

[4] CAI, T. and LIU, W. (2011). Adaptive thresholding for sparse covariance matrix estimation. *J. Amer. Statist. Assoc.* **106** 672–684. MR2847949

[5] CAI, T., LIU, W. and LUO, X. (2011). A constrained ℓ_1 minimization approach to sparse precision matrix estimation. *J. Amer. Statist. Assoc.* **106** 594–607. MR2847973

[6] CAI, T. T., LIU, W. and ZHOU, H. H. (2014). Estimating sparse precision matrix: Optimal rates of convergence and adaptive estimation. *Ann. Statist.* To appear.

[7] CAI, T. T. and YUAN, M. (2012). Adaptive covariance matrix estimation through block thresholding. *Ann. Statist.* **40** 2014–2042. MR3059075

[8] CANDES, E. and TAO, T. (2007). The Dantzig selector: Statistical estimation when p is much larger than n. *Ann. Statist.* **35** 2313–2351. MR2382644

[9] DEMPSTER, A. (1972). Covariance selection. *Biometrics* **28** 157–175.

[10] EFRON, B. (1979). Bootstrap methods: Another look at the jackknife. *Ann. Statist.* **7** 1–26. MR0515681

[11] FAN, J. and FAN, Y. (2008). High-dimensional classification using features annealed independence rules. *Ann. Statist.* **36** 2605–2637. MR2485009

[12] FAN, J., FAN, Y. and LV, J. (2008). High dimensional covariance matrix estimation using a factor model. *J. Econometrics* **147** 186–197. MR2472991

[13] FAN, J., FENG, Y. and WU, Y. (2009). Network exploration via the adaptive lasso and SCAD penalties. *Ann. Appl. Stat.* **3** 521–541. MR2750671

[14] FAN, J., HAN, X. and GU, W. (2012). Estimating false discovery proportion under arbitrary covariance dependence. *J. Amer. Statist. Assoc.* **107** 1019–1035. MR3010887

[15] FAN, J. and LI, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Amer. Statist. Assoc.* **96** 1348–1360. MR1946581

[16] FAN, J. and LV, J. (2008). Sure independence screening for ultrahigh dimensional feature space. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **70** 849–911. MR2530322

[17] FAN, J. and PENG, H. (2004). Nonconcave penalized likelihood with a diverging number of parameters. *Ann. Statist.* **32** 928–961. MR2065194

[18] FAN, Y., JIN, J. and YAO, Z. (2013). Optimal classification in sparse Gaussian graphic model. *Ann. Statist.* **41** 2537–2571. MR3161437

[19] FAN, Y., KONG, Y., LI, D. and ZHENG, Z. (2015). Innovated interaction screening for high-dimensional nonlinear classification. *Ann. Statist.* **43** 1243–1272. MR3346702

[20] FAN, Y. and LV, J. (2013). Asymptotic equivalence of regularization methods in thresholded parameter space. *J. Amer. Statist. Assoc.* **108** 1044–1061. MR3174683

[21] FAN, Y. and LV, J. (2016). Supplement to "Innovated scalable efficient estimation in ultra-large Gaussian graphical models." DOI:10.1214/15-AOS1416SUPP.

[22] FRIEDMAN, J., HASTIE, T. and TIBSHIRANI, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics* **9** 432–441.

[23] HALL, P. and JIN, J. (2010). Innovated higher criticism for detecting sparse signals in correlated noise. *Ann. Statist.* **38** 1686–1732. MR2662357

[24] HALL, P. and LI, K.-C. (1993). On almost linearity of low-dimensional projections from high-dimensional data. *Ann. Statist.* **21** 867–889. MR1232523

[25] HALL, P., TITTERINGTON, D. M. and XUE, J.-H. (2009). Tilting methods for assessing the influence of components in a classifier. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **71** 783–803. MR2750095

[26] HESS, K. R., ANDERSON, K., SYMMANS, W. F., VALERO, V., IBRAHIM, N., MEJIA, J. A., BOOSER, D., THERIAULT, R. L., BUZDAR, A. U., DEMPSEY, P. J., ROUZIER, R., SNEIGE, N., ROSS, J. S., VIDAURRE, T., GÓMEZ, H. L., HORTOBAGYI, G. N. and PUSZTAI, L. (2006). Pharmacogenomic predictor of sensitivity to preoperative chemotherapy with paclitaxel and fluorouracil, doxorubicin, and cyclophosphamide in breast cancer. *J. Clin. Oncol.* **24** 4236–4244.

[27] JIN, J. (2012). Comment: "Estimating false discovery proportion under arbitrary covariance dependence" [MR3010887]. *J. Amer. Statist. Assoc.* **107** 1042–1045. MR3010891

[28] LAM, C. and FAN, J. (2009). Sparsistency and rates of convergence in large covariance matrix estimation. *Ann. Statist.* **37** 4254–4278. MR2572459

[29] LAURITZEN, S. L. (1996). *Graphical Models.* Oxford Univ. Press, New York. MR1419991

[30] LI, K.-C. (1991). Sliced inverse regression for dimension reduction. *J. Amer. Statist. Assoc.* **86** 316–342. MR1137117

[31] LIU, W. (2013). Gaussian graphical model estimation with false discovery rate control. *Ann. Statist.* **41** 2948–2978. MR3161453

[32] LV, J. (2013). Impacts of high dimensionality in finite samples. *Ann. Statist.* **41** 2236–2262. MR3127865

[33] LV, J. and FAN, Y. (2009). A unified approach to model selection and sparse recovery using regularized least squares. *Ann. Statist.* **37** 3498–3528. MR2549567

[34] MARKOWITZ, H. M. (1952). Portfolio selection. *J. Finance* **7** 77–91.

[35] MEINSHAUSEN, N. and BÜHLMANN, P. (2006). High-dimensional graphs and variable selection with the lasso. *Ann. Statist.* **34** 1436–1462. MR2278363

[36] PENG, J., WANG, P., ZHOU, N. and ZHU, J. (2009). Partial correlation estimation by joint sparse regression models. *J. Amer. Statist. Assoc.* **104** 735–746. MR2541591

[37] RAVIKUMAR, P., WAINWRIGHT, M. J., RASKUTTI, G. and YU, B. (2011). High-dimensional covariance estimation by minimizing ℓ_1-penalized log-determinant divergence. *Electron. J. Stat.* **5** 935–980. MR2836766

[38] REN, Z., SUN, T., ZHANG, C.-H. and ZHOU, H. H. (2015). Asymptotic normality and optimalities in estimation of large Gaussian graphical models. *Ann. Statist.* **43** 991–1026. MR3346695

[39] ROTHMAN, A. J., BICKEL, P. J., LEVINA, E. and ZHU, J. (2008). Sparse permutation invariant covariance estimation. *Electron. J. Stat.* **2** 494–515. MR2417391

[40] SUN, T. and ZHANG, C.-H. (2012). Scaled sparse linear regression. *Biometrika* **99** 879–898. MR2999166

[41] TIBSHIRANI, R. (1996). Regression shrinkage and selection via the lasso. *J. Roy. Statist. Soc. Ser. B* **58** 267–288. MR1379242

[42] WAINWRIGHT, M. J. and JORDAN, M. I. (2008). Graphical models, exponential families, and variational inference. *Faund. Trends Mach. Learn.* **1** 1–305.

[43] YE, F. and ZHANG, C.-H. (2010). Rate minimaxity of the Lasso and Dantzig selector for the ℓ_q loss in ℓ_r balls. *J. Mach. Learn. Res.* **11** 3519–3540. MR2756192

[44] YUAN, M. (2010). High dimensional inverse covariance matrix estimation via linear programming. *J. Mach. Learn. Res.* **11** 2261–2286. MR2719856

[45] YUAN, M. and LIN, Y. (2007). Model selection and estimation in the Gaussian graphical model. *Biometrika* **94** 19–35. MR2367824

[46] ZHANG, C.-H. (2010). Nearly unbiased variable selection under minimax concave penalty. *Ann. Statist.* **38** 894–942. MR2604701

[47] ZHANG, T. and ZOU, H. (2014). Sparse precision matrix estimation via lasso penalized D-trace loss. *Biometrika* **101** 103–120. MR3180660

[48] ZOU, H. (2006). The adaptive lasso and its oracle properties. *J. Amer. Statist. Assoc.* **101** 1418–1429. MR2279469

DEPARTMENT OF DATA SCIENCES AND OPERATIONS
MARSHALL SCHOOL OF BUSINESS
UNIVERSITY OF SOUTHERN CALIFORNIA
LOS ANGELES, CALIFORNIA 90089
USA
E-MAIL: fanyingy@marshall.usc.edu
            jinchilv@marshall.usc.edu