

TESTING INDEPENDENCE WITH HIGH-DIMENSIONAL CORRELATED SAMPLES

BY XI CHEN [†], WEIDONG LIU ^{*,‡}

New York University[†] and Shanghai Jiao Tong University[‡]

Testing independence among a number of (ultra) high-dimensional random samples is a fundamental and challenging problem. By arranging n identically distributed p -dimensional random vectors into a $p \times n$ data matrix, we investigate the problem of testing independence among columns under the matrix-variate normal modeling of data. We propose a computationally simple and tuning-free test statistic, characterize its limiting null distribution, analyze the statistical power and prove its minimax optimality. As an important by-product of the test statistic, a ratio-consistent estimator for the quadratic functional of a covariance matrix from correlated samples is developed. We further study the effect of correlation among samples to an important high-dimensional inference problem — large-scale multiple testing of Pearson’s correlation coefficients. Indeed, blindly using classical inference results based on the assumed independence of samples will lead to many false discoveries, which suggests the need for conducting independence testing before applying existing methods. To address the challenge arising from correlation among samples, we propose a “sandwich estimator” of Pearson’s correlation coefficient by de-correlating the samples. Based on this approach, the resulting multiple testing procedure asymptotically controls the overall false discovery rate at the nominal level while maintaining good statistical power. Both simulated and real data experiments are carried out to demonstrate the advantages of the proposed methods.

1. Introduction. The independence among samples is a fundamental assumption in most statistical modeling upon which numerous estimation and inference methods and theories have been developed. Indeed, from classical statistical inference (e.g., student’s t -test) to popular topics in modern statistics (e.g., high-dimensional problems, such as regression, matrix estimation and inference), this assumption of independence occurs widely. Consider n samples $\mathbf{X}_1, \dots, \mathbf{X}_n$, where each sample is a p -dimensional vector

*Research supported by NSFC, Grants No.11322107 and No.11431006, the Program for Professor of Special Appointment (Eastern Scholar) at Shanghai Institutions of Higher Learning, Shanghai Shuguang Program and 973 Program (2015CB856004).

MSC 2010 subject classifications: Primary 62F05; secondary 62H10

Keywords and phrases: Independence test, multiple testing of correlations, false discovery rate, matrix-variate normal, quadratic functional estimation, high-dimensional sample correlation matrix

from the same population distribution with mean $\boldsymbol{\mu} \in \mathbb{R}^p$ and covariance $\boldsymbol{\Sigma} = (\sigma_{ij})_{p \times p}$. It is often convenient to pool n samples together to form a $p \times n$ data matrix $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_n)$. More specifically, for example in microarray data, \mathbf{X} is an expression level matrix for p genes measured on n subjects. Such data are usually high-dimensional; thus, we mainly consider the setting where p is much larger than n . Most existing works in high-dimensional literature make the independence assumption among columns of \mathbf{X} , serving as the starting point of methodology development and technical analysis. However, recent studies have shown that there are correlation structures among subjects in various microarray datasets (see, e.g., [Teng and Huang \(2009\)](#); [Efron \(2009\)](#); [Allen and Tibshirani \(2012\)](#); [Kim et al. \(2012\)](#)), demonstrating the potential risk of making the seemingly natural assumption of independence. Therefore, given a data matrix \mathbf{X} , it is important to first test whether the samples are indeed independent before applying any method that assumes independence.

A data matrix $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_n)$ is known as *transposable* data when both rows and columns are potentially correlated ([Lazzeroni and Owen, 2002](#); [Allen and Tibshirani, 2012](#)). For a transposable data matrix \mathbf{X} , it is commonly assumed that \mathbf{X} follows a matrix-variate normal distribution, which has been widely applied to model microarray data (see, e.g., [Teng and Huang \(2009\)](#); [Efron \(2009\)](#); [Muralidharan \(2010\)](#); [Allen and Tibshirani \(2012\)](#); [Yin and Li \(2012\)](#); [Kim et al. \(2012\)](#); [Zhou \(2014\)](#)). The matrix-variate normal distribution is a natural generalization of familiar vector-variate normal distribution ([Dawid, 1981](#)). In particular, let $\text{vec}(\mathbf{X}) \in \mathbb{R}^{np \times 1}$ be the vectorization of matrix \mathbf{X} obtained by stacking the columns of \mathbf{X} on top of each other. We say $\mathbf{X} \in \mathbb{R}^{p \times n}$ follows a matrix-variate normal distribution with the mean matrix $\mathbf{M} \in \mathbb{R}^{p \times n}$ and covariance matrix $\boldsymbol{\Sigma} \otimes \boldsymbol{\Psi} \in \mathbb{R}^{np \times np}$ (denoted by $\mathbf{X} \sim N(\mathbf{M}, \boldsymbol{\Sigma} \otimes \boldsymbol{\Psi})$) if and only if $\text{vec}(\mathbf{X}') \sim N(\text{vec}(\mathbf{M}'), \boldsymbol{\Sigma} \otimes \boldsymbol{\Psi})$. Here, \mathbf{X}' denotes the transpose of \mathbf{X} , \otimes is the Kronecker product, and $\boldsymbol{\Psi} = (\psi_{ij})_{n \times n} \in \mathbb{R}^{n \times n}$ is the covariance matrix of *row* vectors of \mathbf{X} . Given a matrix-variate normal $\mathbf{X} \sim N(\mathbf{M}, \boldsymbol{\Sigma} \otimes \boldsymbol{\Psi})$, each column $\mathbf{X}_i \sim N(\mathbf{M}_i, \psi_{ii}\boldsymbol{\Sigma})$ for $1 \leq i \leq n$, where \mathbf{M}_i is the i -th column of the mean matrix \mathbf{M} . Recall our problem setup: each \mathbf{X}_i follows the same population distribution with mean vector $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$. Thus, we have $\mathbf{M} = \boldsymbol{\mu}\mathbf{1}'$ where $\mathbf{1}$ is the n -dimensional all one column vector and $\psi_{ii} = 1$ for $1 \leq i \leq n$. Under the matrix-variate normal modeling of the data, the independence testing problem is equivalent to the global test of whether $\boldsymbol{\Psi}$ is a diagonal matrix, i.e.,

$$(1) \quad H_0 : \quad \psi_{ij} = 0 \quad \text{for all } 1 \leq i < j \leq n.$$

The testing problem in (1) is closely related to the following correlation test problem

$$(2) \quad H_0 : \quad \rho_{ij} = 0 \quad \text{for all } 1 \leq i < j \leq p,$$

where $\rho_{ij} = \sigma_{ij} / \sqrt{\sigma_{ii}\sigma_{jj}}$ is the Pearson's correlation coefficient. The testing problem in (2) is a classical problem in multivariate analysis (Nagao, 1973; Anderson, 2003). It has also been extensively studied in the past decade under the high-dimensional setting (e.g., Johnstone (2001), Ledoit and Wolf (2002), Jiang (2004), Schott (2005), Liu, Lin and Shao (2008), Bai et al. (2009), Cai and Jiang (2011), Jiang and Yang (2013), Han and Liu (2014)). However, the reported results are based on the assumption that samples are independent. In fact, our problem in (1) is equivalent to the testing problem (2) with correlated samples. To see this, note that when treating each row of \mathbf{X} as an individual sample, the role of Ψ and Σ interchanges since $\mathbf{X}' \sim N(\mathbf{1}\boldsymbol{\mu}', \Psi \otimes \Sigma)$, i.e., the matrix Σ models the correlations among row samples while Ψ becomes the population covariance matrix. For many types of data (e.g., genetic data, financial data), there exists a complicated correlation structure among p variables. Thus, Σ will not be a diagonal matrix and row vectors are not independent. Our problem in (1) essentially tests the correlation among row vectors when samples are correlated. The correlation among samples makes our problem more challenging; and the aforementioned methods for testing (2), which are based on the assumption of sample independence, cannot be applied to our problem.

The classical methods for testing independence among samples commonly assume p is fixed and are usually designed only for time series data. It is also known as serial independence test, see Hong (1998) and the references therein. In such a framework, the methods require that the samples under alternatives come from some time series. These samples satisfy an ordering structure such that the dependence between two samples decays as the distance of their indices increases. In our setting, there is no structural assumption among samples. Without any structural assumption, we will show in Theorem 2.7 that any test will not have the power tending to 1 uniformly over a large class of alternatives when the dimension p is small (e.g., fixed constant or $p = o(\log n)$). On the other hand, for $p \geq c \log n$ but is small compared to n , the independence test is relatively easy. In fact, if Σ is known, the data matrix can be transformed as $\Sigma^{-1/2}\mathbf{X} \sim N(\Sigma^{-1/2}\boldsymbol{\mu}\mathbf{1}', \mathbf{I}_{p \times p} \otimes \Psi)$; and thus the independence test can be directly carried out using existing approaches (e.g., Jiang (2004); Liu, Lin and Shao (2008)). One can apply such an approach with a plug-in estimator $\widehat{\Sigma}^{-1}$. However, as we will explain later in Section 5, when $p \geq cn$, even the optimal convergence rate of the

estimator $\widehat{\Sigma}^{-1}$ is not fast enough to solve this problem. In fact, although we have more information (i.e., row samples) as p becomes larger, the number of unknown parameters in Σ increases accordingly, which makes the problem challenging. Therefore, the high-dimensional setting is the most interesting case, and will be the main focus of the paper.

Although the testing of independence among high-dimensional samples is an important and fundamental problem, few existing works have done so. Based on matrix-variate normal modeling of the data, some inference approaches were proposed by Efron (2009) and Muralidharan (2010). However, these works do not explore the limiting null distributions as well as the validity and power of the test. Pan, Gao and Yang (2014) proposed a statistic for this problem based on random matrix theory. However, it requires the condition that p is proportionally as large as n (i.e., $0 < \lim_{n \rightarrow \infty} \frac{p}{n} < \infty$), and thus cannot be applied to cases where $p = n^r$ with $r > 1$ or, as in the *ultra high-dimensional* setting, where $p = \exp(n^\gamma)$ for some $0 < \gamma < 1$; both scenarios are common in genetic applications. Further, the method in Pan, Gao and Yang (2014) requires splitting n samples into two parts and differences in splitting could lead to different test results.

In this paper, we consider the (ultra) high-dimensional setup and propose a minimax optimal test procedure in terms of the statistical power for the testing problem in (1). We show that the distribution of the proposed max-type test statistic converges to a type I extreme value distribution under the null (Theorem 2.4). Therefore, the proposed test has the pre-specified significance level asymptotically. We also investigate the statistical power. Roughly speaking, we show that under some very mild conditions on off-diagonal elements of Ψ , the power will converge to 1. Further, we prove that the proposed test is minimax rate-optimal over a large class of Ψ (Theorems 2.5 and 2.6).

Our construction of the test statistic combines a bias correction and a variance correlation based on the sample covariance matrix $(\hat{\psi}_{ij})_{n \times n}$, where we treat each row of \mathbf{X} as a sample. The bias correction technique allows us to handle the ultra high-dimensional case. Moreover, the variance correlation technique deals with the correlation structure among “row samples” of \mathbf{X} , which is specified by Σ . To characterize the strength of correlation among row samples, we identify a key quantity $A_p = \frac{p \|\Sigma\|_F^2}{(\text{tr}(\Sigma))^2}$, which comes from the asymptotic variance of our bias-corrected statistic. Here, $\|\cdot\|_F$ denotes the Frobenius norm and $\text{tr}(\Sigma)$ is the trace of Σ . To simultaneously control the type I error under null and maintain the minimax rate-optimal statistical power, we need a ratio consistent estimator of A_p regardless of the correlation among samples. Therefore, the remaining task essentially reduces to the

problem of estimating $\|\Sigma\|_{\text{F}}^2$ from correlated samples.

It is noteworthy that estimating $\|\Sigma\|_{\text{F}}^2$ itself is an important problem, which is known as quadratic functional estimation of Σ (see, e.g., [Bai and Saranadasa \(1996\)](#); [Chen and Qin \(2010\)](#); [Fan, Rigollet and Wang \(2015\)](#)). Most existing works are based on the assumption that samples are independent and identically distributed (*i.i.d.*) and, thus, cannot be directly applied to our problem. Motivated by the thresholding estimator in [Fan, Rigollet and Wang \(2015\)](#), we propose a plugin estimator for $\|\Sigma\|_{\text{F}}^2$ based on a thresholded sample covariance matrix but we relax the independence assumption among samples. Further, we propose a *definite* threshold level, which is adaptive to the amount of correlations among samples and guarantees the consistency of the resulting estimator. Our simulation results demonstrate the superior performance of the proposed estimator of $\|\Sigma\|_{\text{F}}^2$ over the existing approaches, which leads to a significant improvement in statistical power.

In summary, we propose a simple max-type test statistic to conduct the global test of independence among high-dimensional random samples in (1). Our approach has the following advantages:

1. Our construction is direct and computationally attractive, which only requires the row sample covariance matrix $(\hat{\psi}_{ij})_{n \times n}$ and a threshold estimator of $\|\Sigma\|_{\text{F}}$. Further, our test statistic is completely tuning free.
2. The limiting null distribution is characterized and thus the type I error is controlled asymptotically. Further, our test procedure is minimax rate-optimal over a sufficiently large class of Ψ , which is enough for most practical purposes.
3. As an important by-product, we provide a ratio-consistent estimator for estimating quadratic functional of covariance matrix from correlated samples.

We would like to note that we only focus on the matrix-variate normal distribution, which is a common assumption for studying a transposable data matrix and widely used for modeling correlated microarray data. It is of interest to investigate the independence test for more general distributions, e.g., a matrix elliptical distribution ([Dawid, 1977](#); [Fang and Zhang, 1990](#)) or $\mathbf{X} = \Sigma^{1/2} \mathbf{Z} \Psi^{1/2}$, where entries of $\mathbf{Z} = (Z_{ij})_{p \times n}$ are *i.i.d.* random variables with unit variance. We leave the extension to such distributions of data matrices for future work.

After we conduct the independence test, if the samples are indeed correlated, many classical inference approaches cannot be directly applied. We use the multiple testing problem of Pearson's correlation coefficients to illustrate the effect of the correlation among samples, demonstrate the reason

why the classical approach will fail when samples are correlated and further develop a new method to de-correlate the samples. In particular, we consider the following large-scale multiple testing problem, for $1 \leq i < j \leq p$,

$$(3) \quad H_{0ij} : \rho_{ij} = 0 \quad \text{versus} \quad H_{1ij} : \rho_{ij} \neq 0.$$

Problem (3) is a natural extension of the global test of independence in (2). In fact, the hypothesis that Σ is a diagonal matrix is a strong null hypothesis, which will be rejected in most real data applications (e.g., microarray data, stock data). In contrast, the goal of the multiple testing problem (3) is to identify the pairs of correlated variables and thus find many applications in real data analysis, e.g., gene coexpression network analysis (Lee, Hsu and Sajdak, 2004; Carter et al., 2004; Zhu et al., 2005; Hirai, 2007), and brain connectivity analysis (Shaw, 2006). The goal of the testing problem in (3) is consistent with the goal of *support recovery* of a sparse Σ . The latter problem has been extensively studied in recent years (e.g., see Rothman, Levina and Zhu (2009), Lam and Fan (2009), Cai and Liu (2011), Bien and Tibshirani (2011)). These works establish consistency results of support recovery from independent samples under certain conditions, for example, all the absolute values of nonzero ρ_{ij} are lower bounded by $C\sqrt{\frac{\log p}{n}}$, which might be hard to hold in practice. Instead of trying to achieve the perfect support recovery, the multiple testing problem (3) has a more refined control of the type I error rate in support recovery under weaker assumptions. In particular, it usually aims to control the false discovery rate (FDR), which is a useful measure for evaluating the performance of support recovery. We also note that Cai and Liu (2015) recently studied problem (3) in a high-dimensional setting but it still requires the independence assumption.

For correlated samples from a matrix-variate normal distribution, we first establish the following result on the limiting distribution of the sample correlation coefficient $\hat{\rho}_{ij}$ (see Proposition 3.1):

$$(4) \quad \frac{\sqrt{n}(\hat{\rho}_{ij} - \rho_{ij})}{\sqrt{B_n(1 - \rho_{ij}^2)}} \Rightarrow N(0, 1),$$

where $B_n = \frac{\|\Psi\|_{\mathbb{F}}^2}{n}$, which quantifies the strength of the correlation among samples. Eq. (4) subsumes as a special case the classical results on the limiting distribution of $\hat{\rho}_{ij}$ when samples are *i.i.d.* ($B_n = 1$ in (4)) (see Theorem 4.2.4 in Anderson (2003)). When the correlation is strong to a certain extent such that $B_n > 1 + c$ for some constant $c > 0$, directly using sample correlation coefficient $\sqrt{n}\hat{\rho}_{ij}$ or Fisher's z statistic will lead to many false positives; this is verified by our simulations in Section 4.2. In fact, even if

B_n is known and one uses the correct limiting null distribution $N(0, \frac{B_n}{n})$ of $\hat{\rho}_{ij}$, the variance of $\hat{\rho}_{ij} - \rho_{ij}$ becomes larger as B_n increases, which leads to a lower power of the test.

To overcome the side effect of correlation among samples, we propose a “sandwich estimator” of ρ_{ij} by de-correlating the samples, which has the limiting distribution $N(\rho_{ij}, \frac{1}{n}(1 - \rho_{ij}^2)^2)$. The corresponding asymptotical variance does not depend on B_n and is smaller than that of the naïve estimator $\hat{\rho}_{ij}$. Therefore the proposed “sandwich estimator” has an improved statistical power especially when the correlation among samples is strong. Based on the proposed “sandwich estimator” of ρ_{ij} , the standard multiple testing procedure (Benjamini and Hochberg, 1995) is proven to asymptotically control the FDR at the nominal level (see Theorem 3.2).

Finally, we introduce some necessary notations. For a positive integer p , $[p] := \{1, \dots, p\}$. For a square matrix \mathbf{A} , let $tr(\mathbf{A})$ denote the trace of \mathbf{A} , $\lambda_{\max}(\mathbf{A})$ the maximum eigenvalue of \mathbf{A} , and $\lambda_{\min}(\mathbf{A})$ the minimum eigenvalue of \mathbf{A} . Let $I\{B\}$ be the indicator function that takes value one when the event B is true and zero otherwise. For a given set \mathcal{H} , let $Card(\mathcal{H})$ be the cardinality of \mathcal{H} . For any two real numbers a and b , let $a \vee b = \max(a, b)$ and $a \wedge b = \min(a, b)$. We use $\overline{\lim}$ and $\underline{\lim}$ to denote limit superior and limit inferior, respectively. Throughout the paper, we use $\mathbf{I}_{p \times p}$ to denote the $p \times p$ identity matrix, and use C, c, c_1 , etc. to denote constants for which values might change from place to place and do not depend on n and p .

The rest of the paper is organized as follows. In Section 2, we study the global test in (1). The test statistic is proposed in Section 2.1. In Section 2.2, we provide the ratio consistent estimator of A_p and $\|\Sigma\|_{\text{F}}$ from correlated samples. The estimation error is characterized in Theorem 2.1. We further provide the limiting null distribution of the test statistic and the power analysis (Theorems 2.4–2.7). Section 3 studies the multiple testing of correlations in (3) from correlated samples. Experimental results are given in Section 4 followed by discussion in Section 5. Due to space limitations, the proofs of our results as well as some additional experimental results are provided in the supplementary material (Chen and Liu, 2015).

2. Sample independence test. We study the global testing problem of sample independence in (1) given the $p \times n$ data matrix $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_n) \sim N(\boldsymbol{\mu}\mathbf{1}', \Sigma \otimes \Psi)$.

2.1. *Construction of the test statistic.* Recall that $\mathbf{X}_i = (X_{i1}, \dots, X_{ip})'$ denotes the i -th sample for $1 \leq i \leq n$ and let $\bar{\mathbf{X}} = \frac{1}{n} \sum_{i=1}^n \mathbf{X}_i =: (\bar{X}_1, \dots, \bar{X}_p)'$.

Define

$$(5) \quad \hat{\psi}_{ij} = \frac{1}{p} \sum_{k=1}^p (X_{ik} - \bar{X}_k)(X_{jk} - \bar{X}_k), \quad 1 \leq i, j \leq n.$$

In fact, from the proof, the statistic $(\hat{\psi}_{ij})_{n \times n}$ is the sample covariance coefficient corresponding to $\frac{\text{tr}(\mathbf{\Sigma})}{p} \psi_{ij}$. Further, under the null H_0 , we can show that

$$(6) \quad \hat{\psi}_{ij} = \frac{1}{p} \sum_{k=1}^p (X_{ik} - \mu_k)(X_{jk} - \mu_k) - \frac{1}{np} \sum_{k=1}^p \sigma_{kk} + O_{\mathbb{P}}\left(\frac{1}{\sqrt{np}}\right).$$

The first term $\frac{1}{p} \sum_{k=1}^p (X_{ik} - \mu_k)(X_{jk} - \mu_k)$ has mean $\frac{\text{tr}(\mathbf{\Sigma})}{p} \psi_{ij}$ and variance $\frac{\|\mathbf{\Sigma}\|_{\mathbb{F}}^2}{p^2} (\psi_{ii}\psi_{jj} + \psi_{ij}^2)$. The bias term $\frac{1}{np} \sum_{k=1}^p \sigma_{kk}$ comes from the centralization statistics $\{\bar{X}_k\}_{k=1}^p$ in (5). When $p = o(n^2)$, we have $\frac{1}{np} \sum_{k=1}^p \sigma_{kk} = o(1/\sqrt{p})$ and $\sqrt{p}\hat{\psi}_{ij}$ can be shown to converge to a normal distribution. However, as we are interested in the ultra high-dimensional case where p can be as large as $\exp(o(n^\gamma))$ for some $0 < \gamma < 1$, when p becomes larger such that $n^2 = o(p)$, $\sqrt{p}\hat{\psi}_{ij} \rightarrow -\infty$ in probability under the null. To enable the applicability of our test statistic in the ultra high-dimensional setting, we first propose the following bias corrected quantity:

$$(7) \quad T_{ij} := \hat{\psi}_{ij} + \frac{1}{np} \sum_{k=1}^p \hat{\sigma}_{kk},$$

where $\hat{\sigma}_{kk} = \frac{1}{n-1} \sum_{j=1}^n (X_{jk} - \bar{X}_k)^2$ is the sample variance corresponding to σ_{kk} . Since the first term in (6) has variance $\frac{\|\mathbf{\Sigma}\|_{\mathbb{F}}^2}{p^2} (\psi_{ii}\psi_{jj} + \psi_{ij}^2)$, the asymptotic variance of T_{ij} is $\left(\frac{\text{tr}(\mathbf{\Sigma})}{p}\right)^2 \frac{A_p}{p}$, where

$$(8) \quad A_p = \frac{p\|\mathbf{\Sigma}\|_{\mathbb{F}}^2}{(\text{tr}(\mathbf{\Sigma}))^2}$$

quantifies the strength of correlations among row vectors of \mathbf{X} .

Given A_p in (8), we will show that under the null as $(n, p) \rightarrow \infty$,

$$(9) \quad \mathbb{P}\left(\frac{p}{A_p} \max_{1 \leq i < j \leq n} \frac{T_{ij}^2}{\hat{\psi}_{ii}\hat{\psi}_{jj}} - 4 \log n + \log \log n \leq t\right) \rightarrow \exp\left(-\frac{1}{\sqrt{8\pi}} \exp\left(-\frac{t}{2}\right)\right)$$

for $t \in \mathbb{R}$, where the term A_p plays the role of variance correction for T_{ij} . The remaining task is to develop a ratio consistent estimator \hat{A}_p for A_p .

In addition, to maintain the statistical power, the estimator \hat{A}_p should also be consistent for correlated samples. In Section 2.2, we will develop such an estimator for A_p . Given the estimator \hat{A}_p (see (14)), we propose the following test statistic for the independence test in (1),

$$(10) \quad \hat{T}_{n,p} = \frac{p}{\hat{A}_p} \max_{1 \leq i < j \leq n} \frac{T_{ij}^2}{\hat{\psi}_{ii}\hat{\psi}_{jj}}.$$

2.2. *Estimation of A_p and $\|\Sigma\|_F^2$ from correlated samples.* The estimation of $\|\Sigma\|_F^2$ finds many applications and has been studied in several works (Bai and Saranadasa, 1996; Chen and Qin, 2010; Fan, Rigollet and Wang, 2015). However, all these works rely on the sample independence assumption. In particular, Fan, Rigollet and Wang (2015) proved that the simple plug-in procedure based on threshold estimators are minimax optimal over a large class of covariance matrices. Moreover, the threshold level in Fan, Rigollet and Wang (2015) takes the form of $C\sqrt{\frac{\log p}{n}}$, where the constant C needs to be carefully tuned to achieve good performance in practice. A cross-validation (CV) procedure was suggested; however, there is no theoretical justification for such a CV procedure. In this section, we introduce a threshold estimator for $\|\Sigma\|_F^2$ with an explicit threshold level, which is completely data-driven without any tuning and automatically adaptive to the correlation among samples. We will show in Theorem 2.1 that the obtained estimator is ratio-consistent for correlated samples.

Let us define the (column) sample covariance matrix $\hat{\Sigma} = (\hat{\sigma}_{ij})_{1 \leq i, j \leq p}$ with $\hat{\sigma}_{ij} = \frac{1}{n-1} \sum_{k=1}^n (X_{ki} - \bar{X}_i)(X_{kj} - \bar{X}_j)$ and sample correlation coefficient $\hat{\rho}_{ij} = \hat{\sigma}_{ij} / \sqrt{\hat{\sigma}_{ii}\hat{\sigma}_{jj}}$ for $1 \leq i, j \leq p$. Further, define

$$(11) \quad B_n = \frac{\|\Psi\|_F^2}{n} = \frac{1}{n} \sum_{1 \leq i, j \leq n} \psi_{ij}^2,$$

which quantifies the average correlation among samples. It can be shown that $\frac{\hat{\rho}_{ij} - \rho_{ij}}{\sqrt{B_n(1 - \hat{\rho}_{ij}^2)}} \Rightarrow N(0, 1)$ (see Proposition 3.1 in Section 3 and note that $\hat{\rho}_{ij} \rightarrow \rho_{ij}$ in probability). We propose the following threshold estimator $\hat{\Sigma}_{thr} = (\hat{\sigma}_{ij,thr})_{1 \leq i, j \leq p}$, where

$$(12) \quad \hat{\sigma}_{ij,thr} = \hat{\sigma}_{ij} I \left\{ \frac{|\hat{\rho}_{ij}|}{1 - \hat{\rho}_{ij}^2} \geq \delta \sqrt{\frac{\hat{B}_n \log p}{n}} \right\} \quad \text{for } i \neq j, \quad \hat{\sigma}_{ii,thr} = \hat{\sigma}_{ii} \quad \text{for } 1 \leq i \leq p.$$

Here, \hat{B}_n is an estimator of B_n and δ can be any constant larger than $\sqrt{2}$. Let $\hat{\Psi} = (\frac{p}{\text{tr}(\hat{\Sigma})} \hat{\psi}_{ij})_{1 \leq i, j \leq n}$. Using the approach from Bai and Saranadasa

(1996), we construct

$$(13) \quad \hat{B}_n = \frac{1}{n} \left(\|\hat{\Psi}\|_{\mathbb{F}}^2 - \frac{1}{p} (\text{tr}(\hat{\Psi}))^2 \right).$$

Given the threshold estimator $\hat{\Sigma}_{thr} = (\hat{\sigma}_{ij,thr})_{1 \leq i, j \leq p}$ in (12), the $\|\Sigma\|_{\mathbb{F}}^2$ is estimated by $\|\hat{\Sigma}_{thr}\|_{\mathbb{F}}^2$ and A_p is estimated by

$$(14) \quad \hat{A}_p = \frac{p \|\hat{\Sigma}_{thr}\|_{\mathbb{F}}^2}{(\text{tr}(\hat{\Sigma}_{thr}))^2}.$$

Now we will show that $\|\hat{\Sigma}_{thr}\|_{\mathbb{F}}^2$ and \hat{A}_p are ratio-consistent estimators of $\|\Sigma\|_{\mathbb{F}}^2$ and A_p , respectively. We first make the following three assumptions throughout this section. Let $\lambda_{\min}(\Sigma) = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_p = \lambda_{\max}(\Sigma)$ be the eigenvalues of Σ and $\lambda_{\min}(\Psi) = \nu_1 \leq \nu_2 \leq \dots \leq \nu_n = \lambda_{\max}(\Psi)$ be eigenvalues of Ψ . We make the following standard assumption on eigenvalues:

(C1) We assume that $c^{-1} \leq \lambda_{\min}(\Sigma) \leq \lambda_{\max}(\Sigma) \leq c$ and $c^{-1} \leq \lambda_{\min}(\Psi) \leq \lambda_{\max}(\Psi) \leq c$ for some constant $c > 0$.

The condition **(C1)** is a typical eigenvalue assumption in high-dimensional covariance estimation literature (see the survey [Cai, Ren and Zhou \(2016\)](#) and references therein). This assumption is natural for many important classes of covariance matrices, e.g., bandable, Toeplitz, and sparse covariance matrices. There are cases that the assumption **(C1)** is violated, e.g., when the covariance matrix has equal correlation structure (i.e., $\Sigma = \rho \cdot \mathbf{1}\mathbf{1}' + (1 - \rho) \cdot \mathbf{I}_{p \times p}$ for some $\rho \in (0, 1)$). Our result will not hold for such a setting and please refer to [Figure 3](#) for the experimental illustrations.

We also note that this condition can be weakened by replacing the constant c by some $c_p \rightarrow \infty$ at a certain rate. However, for the sake of simplicity, we do not intend to seek the optimal rate of c_p . We only mention that this type of constraint on eigenvalues is needed in our problem. Without this type of constraints, T_{ij} in (7) will no longer be asymptotic normal because the Lindeberg's condition for the central limit theorem (CLT) of independent random variables (see the expression of $\hat{\psi}_{ij}$ in Eq. (67) in the supplementary material) is violated. Thus, our result on type I error rate control in [Proposition 2.3](#) will no longer hold.

The second condition is also a standard assumption on the norm of each row of Ψ and Σ .

(C2) For some $0 < \tau < 2$, assume that $\sum_{k=1}^n |\psi_{ik}|^{\tau} \leq C$ uniformly over each row $1 \leq i \leq n$ and $\sum_{k=1}^p |\sigma_{jk}|^{\tau} \leq C$ uniformly over each row $1 \leq j \leq p$.

Notably, the upper bounds on eigenvalues of Σ and Ψ in **(C1)** only imply the ℓ_2 -boundedness of each row of Ψ and Σ , i.e., $\sum_{k=1}^n |\psi_{ik}|^2 \leq c^2$ and $\sum_{k=1}^p |\sigma_{jk}|^2 \leq c^2$. The condition **(C2)** is stronger than this implication by noticing that $0 < \tau < 2$. Moreover, when $0 < \tau < 1$, this assumption becomes the typical *weak sparsity* assumption in high-dimensional covariance estimation.

The third assumption is on the relationship between n and p .

(C3) We assume that $p > cn$ for some universal constant $c > 0$ that does not depend on p and n . We further assume that $p = \exp(o(n^\gamma))$ with $\gamma = (1 - \epsilon) \wedge (\frac{2}{\tau} - 1)$ for some $\epsilon > 0$.

The first condition $p = p_n > cn$ is quite natural in a high-dimensional setting and the second condition $p = \exp(o(n^\gamma))$ allows us to deal with an ultra high-dimensional setting.

Under these three assumptions, we provide the following theorem, which establishes the ratio consistency of the estimators \hat{A}_p and $\|\hat{\Sigma}_{thr}\|_F^2$.

THEOREM 2.1. *Assume that (C1)-(C3) hold. For any $\delta > \sqrt{2}$, we have $\frac{\hat{A}_p}{A_p} = 1 + O_{\mathbb{P}}\left(\left(\sqrt{\frac{\log p}{n}}\right)^{\min(1, 2-\tau)}\right)$ and $\frac{\|\hat{\Sigma}_{thr}\|_F^2}{\|\Sigma\|_F^2} = 1 + O_{\mathbb{P}}\left(\left(\sqrt{\frac{\log p}{n}}\right)^{\min(1, 2-\tau)}\right)$.*

According to Theorem 2.1, we will simply set $\delta = 1.42$ in the estimator \hat{A}_p in our experiment. In fact, the experimental results are quite robust with respect to the choice of δ . As long as the δ is above $\sqrt{2}$ and does not take a too large value, the experimental results will not be affected.

Due to the term \hat{B}_n in the thresholding level, our estimator is adaptive to the correlations between the samples. We next show that, even when $\Sigma = \mathbf{I}_{p \times p}$, if we use the thresholding level designed for *i.i.d.* samples without \hat{B}_n as in [Fan, Rigollet and Wang \(2015\)](#), the resultant estimator \tilde{A}_p will over-estimate A_p and, hence, reduce the power. In particular, define the thresholding estimator

$$\hat{\Sigma}_1 = (\hat{\sigma}_{ij,1}), \quad \text{where} \quad \hat{\sigma}_{ij,1} = \hat{\sigma}_{ij} I\{|\hat{\sigma}_{ij}| \geq \lambda \sqrt{\frac{\log p}{n}}\}, \quad i \neq j,$$

and $\hat{\sigma}_{ii,1} = \hat{\sigma}_{ii}$. [Fan, Rigollet and Wang \(2015\)](#) showed that, under the *i.i.d.* assumption, for a large constant-valued λ (not depending on Ψ), $\sum_{i \neq j} \hat{\sigma}_{ij,1}^2$ attains the minimax-optimal rate for estimating $\sum_{i \neq j} \sigma_{ij}^2$. Let $\tilde{A}_p = \frac{p \|\hat{\Sigma}_1\|_F^2}{(\text{tr}(\hat{\Sigma}_1))^2}$. When the samples are correlated, $\|\hat{\Sigma}_1\|_F$ is no longer a ratio consistent estimator for $\|\Sigma\|_F$ and hence results in a poor estimator for A_p .

PROPOSITION 2.2. *Assume that $\Sigma = \mathbf{I}_{p \times p}$ and (C1)-(C3) hold. For any $\lambda > 0$ and $\nu > 0$, there is a class of covariance matrices Ψ with $B_n \geq 5\lambda^2/\nu$ such that $\mathbb{P}(\tilde{A}_p/A_p \geq 1 + cp^{1-\nu}/n) \rightarrow 1$ as $(n, p) \rightarrow \infty$.*

Proposition 2.2 shows that \tilde{A}_p will over-estimate A_p when $p \gg n$. If \tilde{A}_p is used to estimate A_p , then the resultant testing approach will be less powerful than the test with our estimator \hat{A}_p . We will further show the impact of \tilde{A}_p on the power in the simulation.

2.3. *Type I error rate control and optimality of statistical power.* The following proposition gives the limiting distribution of T_{ij} .

PROPOSITION 2.3. *Assume that $p \geq cn$ for some constant $c > 0$ (which does not depend on n and p) and (C1) holds. Under the null H_0 , for $t \in \mathbb{R}$, we have as $(n, p) \rightarrow \infty$*

$$\mathbb{P}\left(\frac{p}{A_p} \max_{1 \leq i < j \leq n} \frac{T_{ij}^2}{\hat{\psi}_{ii}\hat{\psi}_{jj}} - 4 \log n + \log \log n \leq t\right) \rightarrow \exp\left(-\frac{1}{\sqrt{8\pi}} \exp\left(-\frac{t}{2}\right)\right).$$

In Proposition 2.3, the test statistic $\frac{T_{ij}}{\sqrt{\hat{\psi}_{ii}\hat{\psi}_{jj}}}$ can be viewed as a sample correlation coefficient related with ψ_{ij} . We first note that Proposition 2.3 cannot be implied by Theorem 4 in Cai and Jiang (2011). Let us denote the sample correlation coefficient by

$$(15) \quad \hat{\rho}_{ij} = \frac{\sum_{k=1}^n (X_{ki} - \bar{X}_i)(X_{kj} - \bar{X}_j)}{\sqrt{\sum_{k=1}^n (X_{ki} - \bar{X}_i)^2 \sum_{k=1}^n (X_{kj} - \bar{X}_j)^2}}.$$

Cai and Jiang (2011) established the limiting distribution of $\max_{|i-j| \geq \tau} |\hat{\rho}_{ij}|$ for $\tau \geq 1$. Their result requires that n random vectors (X_{ki}, X_{kj}) for $1 \leq k \leq n$ in the sum $\sum_{k=1}^n (X_{ki} - \bar{X}_i)(X_{kj} - \bar{X}_j)$ in (15) are *i.i.d.* On the contrary, our statistic T_{ij} is based on $\sum_{k=1}^p X_{ik}X_{jk}$, which is a sum of p potentially correlated random variables, no matter under the null or alternatives.

In addition, it is worthwhile to note that Cai and Jiang (2012) revealed an interesting phase transition phenomenon in the limiting distribution of the largest off-diagonal entry of the sample correlation matrix. There are different regimes for large p , in which the limiting distributions are different. In contrast, in our problem, there is no such a phase transition phenomenon and the limiting distribution is unified in the high-dimensional setting when $p \geq cn$. To see this more clearly, let us assume that $\mathbf{X}^{(k)}$ for $k = 1, \dots, p$, are independent so that the results in Cai and Jiang (2012) are valid. Now, the

quantity p is the sample size and n is the dimension. According to Corollary 2.2 of Cai and Jiang (2012), there is a phase transition phenomenon for the distribution of the statistic in Proposition 2.3 (i.e., $\frac{p}{A_p} \max_{1 \leq i < j \leq n} \frac{T_{ij}^2}{\hat{\psi}_{ii}\hat{\psi}_{jj}} - 4 \log n + \log \log n$) between two regimes $\frac{1}{\sqrt{p}} \log n \rightarrow 0$ and $\frac{1}{\sqrt{p}} \log n \rightarrow \alpha \in (0, \infty)$. In our high-dimensional setting, we have $p \geq cn$, which belongs to the first regime $\frac{1}{\sqrt{p}} \log n \rightarrow 0$. Thus, there is no phase transition phenomenon in the high-dimensional setting.

Using Theorem 2.1, we provide the limiting null distribution of our test statistic $\hat{T}_{n,p}$ in the next theorem.

THEOREM 2.4. *Assume that (C1)-(C3) hold. Under the null H_0 , we have*

$$(16) \quad \mathbb{P}\left(\hat{T}_{n,p} - 4 \log n + \log \log n \leq t\right) \rightarrow \exp\left(-\frac{1}{\sqrt{8\pi}} \exp\left(-\frac{t}{2}\right)\right)$$

for $t \in \mathbb{R}$, as $(n, p) \rightarrow \infty$.

Remark: In Theorem 2.4, we need the additional assumption $p = \exp(o(n^\gamma))$ in (C3), which is used to obtain a ratio-consistent estimator of $\|\Sigma\|_{\mathbb{F}}^2$ and A_p . If we consider only the limiting distribution of the test statistic under the null (*i.i.d. samples*), one may use the method from Chen and Qin (2010) to estimate $\|\Sigma\|_{\mathbb{F}}$. The estimator from Chen and Qin (2010) does not require the condition $p = \exp(o(n^\gamma))$. However, in terms of statistical power, as we have shown in our simulations, their estimator will over-estimate $\|\Sigma\|_{\mathbb{F}}^2$ (see Figure C in the supplementary material) and reduce the power (see Figure 1) especially when the correlation among samples is strong. Our estimator is ratio-consistent for both null and alternative (see Theorem 2.1) under the extra condition $p = \exp(o(n^\gamma))$. For the thresholding estimator, such a condition on p is necessary. To see this, if $\log p$ is much larger than n , then the thresholding level in (12) is much larger than one. Thus, $\hat{\Sigma}_{thr}$ becomes $\text{diag}(\hat{\Sigma})$ and $\|\hat{\Sigma}_{thr}\|_{\mathbb{F}}^2$ will no longer be consistent. As a future direction, it would be interesting to construct a consistent estimator for $\|\Sigma\|_{\mathbb{F}}^2$ and A_p under the null and alternative simultaneously without the restriction on p .

According to Theorem 2.4, for a given significance level $0 < \alpha < 1$, we reject the null hypothesis whenever $\hat{T}_{n,p} \geq q_\alpha + 4 \log n - \log \log n$, where q_α is the $1 - \alpha$ quantile of the type I extreme value distribution with the cumulative distribution function (CDF) $\exp\left(-\frac{1}{\sqrt{8\pi}} \exp\left(-\frac{x}{2}\right)\right)$, i.e.,

$$(17) \quad q_\alpha = -\log(8\pi) - 2 \log \log(1 - \alpha)^{-1}.$$

Theorem 2.4 shows that the proposed test statistic controls the type I error rate at the nominal level asymptotically.

We now turn to the power analysis. For a given pair of $1 \leq i < j \leq n$, let us define

$$(18) \quad d_{ij, \Psi} = \psi_{ij} - \frac{\sum_{k=1, \neq i}^n \psi_{ik}}{n} - \frac{\sum_{k=1, \neq j}^n \psi_{jk}}{n} - \frac{\sum_{1 \leq i \neq j \leq n} \psi_{ij}}{n^2(n-1)}.$$

and

$$(19) \quad d_{n, \Psi} := \max_{1 \leq i < j \leq n} |d_{ij, \Psi}|.$$

The next theorem shows that for a large class of Ψ , the null hypothesis will be rejected by our test with probability tending to one.

THEOREM 2.5. *Assume that (C1)-(C3) hold and suppose that for some $\delta > 2$ and all large enough n and p ,*

$$(20) \quad d_{n, \Psi} \geq \delta \sqrt{\frac{A_p \log n}{p}}.$$

We have $\mathbb{P}(\hat{T}_{n,p} - 4 \log n + \log \log n \geq q_\alpha) \rightarrow 1$ as $(n, p) \rightarrow \infty$.

We next show that our test statistic is minimax rate optimal for statistical power even when μ and Σ are known. To this end, we introduce a class of covariance matrix for $\Psi \in \mathcal{F}(\delta)$ for some $\delta > 0$ as follows:

$$\mathcal{F}(\delta) = \left\{ \Psi \succ 0 : \psi_{ii} = 1, 1 \leq i \leq n \quad \text{and} \quad d_{n, \Psi} \geq \delta \sqrt{\frac{\log n}{p}} \right\}.$$

Let \mathcal{T}_α be the set of α -level tests with μ and Σ being known, i.e. $\mathcal{T}_\alpha = \{T_\alpha : \mathbb{P}(T_\alpha = 1 | H_0) \leq \alpha\}$. Here $T_\alpha = 1$ means the rejection of H_0 .

THEOREM 2.6. *Let $\alpha, \beta > 0$ and $\alpha + \beta < 1$. Assume that (C3) holds. For any $\delta < 2$, we have*

$$(21) \quad \overline{\lim}_{(n,p) \rightarrow \infty} \sup_{T_\alpha \in \mathcal{T}_\alpha} \inf_{\Psi \in \mathcal{F}(\delta)} \mathbb{P}(T_\alpha = 1) \leq 1 - \beta.$$

Theorem 2.6 shows that for any α -level test T_α and any $\delta < 2$, there must exist a covariance matrix $\Psi \in \mathcal{F}(\delta)$ such that the probability of rejecting the null is less than $\alpha + \varepsilon$ asymptotically for any $\varepsilon > 0$. Theorems 2.5 and 2.6 together show that the proposed test based on $\hat{T}_{n,p}$ is minimax rate optimal by noting that $1 \leq A_p \leq C$ for some constant $C > 0$ according to

the condition **(C1)**. In other words, the order of the lower bound $\sqrt{\frac{\log n}{p}}$ on $d_{n,\Psi}$ cannot be improved, which establishes the minimax-optimal rate for the test. Moreover, when $\Sigma = \mathbf{I}_{p \times p}$, we have $A_p = 1$ and, hence, our test statistic is also minimax constant optimal.

We further show that (20) is a rather wide class of Ψ in the sense that if (20) does not hold, it will be safe to assume the independence for some applications. In particular, assume that Ψ is a s_n sparse matrix, i.e., the number of nonzero elements in each row of Ψ is bounded from above by s_n . Then by (18),

$$d_{n,\Psi} \geq \left(1 - \frac{3s_n}{n}\right) \max_{1 \leq i < j \leq n} |\psi_{ij}|.$$

Thus, a sufficient condition for (20) to hold is $s_n = o(n)$ and

$$(22) \quad \max_{1 \leq i < j \leq n} |\psi_{ij}| \geq \delta \sqrt{\frac{A_p \log n}{p}} \quad \text{for some } \delta > 2.$$

Theorem 2.5 shows that under (22), the null hypothesis will be rejected with probability tending to 1. In fact, when (22) does not hold, the samples can be safely treated as independent for some applications. Let us take the multiple testing problem of correlations in (3) as an example. As we discussed in the introduction, the effect of the correlation among samples is quantified by $B_n = \frac{\|\Psi\|_F^2}{n}$ and when $B_n \rightarrow 1$, the limiting distribution of $\hat{\rho}_{ij}$ in (4) will be the same as the limiting distribution of $\hat{\rho}_{ij}$ estimated from independent samples. Indeed, when (22) does not hold and $p \geq cn^\gamma$ for some $\gamma > 1$, then we have $B_n \rightarrow 1$ (note that $\psi_{ii} = 1$ for $1 \leq i \leq n$). Thus, the correlation among samples is asymptotically negligible.

We next give a more general result on the relation between the lower bound of $\max_{1 \leq i < j \leq n} |\psi_{ij}|$, n and p . Here we only assume that $n \rightarrow \infty$ and p is a function of n (note that p can be a constant). Let

$$\mathcal{G}(a) = \left\{ \Psi \succ 0 : \psi_{ii} = 1, 1 \leq i \leq n \quad \text{and} \quad \max_{1 \leq i < j \leq n} |\psi_{ij}| \geq a \right\}.$$

THEOREM 2.7. *Let $\alpha, \beta > 0$ and $\alpha + \beta < 1$. For any a and p satisfying*

$$(23) \quad (1 - a^2)^{-p/2} = o(n^2) \quad \text{as } n \rightarrow \infty,$$

we have

$$\overline{\lim}_{n \rightarrow \infty} \sup_{T_\alpha \in \mathcal{T}_\alpha} \inf_{\Psi \in \mathcal{G}(a)} \mathbb{P}(T_\alpha = 1) \leq 1 - \beta.$$

Theorem 2.7 shows that when the dimension p is fixed, it is impossible to reject H_0 correctly for all $\Psi \in \mathcal{G}(a)$ with probability greater than $\alpha + \varepsilon$, even when the lower bound $\max_{1 \leq i < j \leq n} |\psi_{ij}|$ is close to one. It is easy to understand since the role of n and p is interchanged in our setting (we are testing an $n \times n$ covariance matrix Ψ with p row samples). It also indicates that the independence test problem (1) is essentially different from the serial independence test in time series analysis. When $a = c/\sqrt{n}$ for some constant $c > 0$, we must require $p \geq c_1 n \log n$ for some $c_1 > 0$ such that the independence testing problem (1) is solvable over $\mathcal{G}(a)$. Note that Pan, Gao and Yang (2014) requires $0 < \lim_{n \rightarrow \infty} \frac{p}{n} < \infty$, which means that their method fails to deal with the setting $a \leq c/\sqrt{n}$. On the other hand, by (22), such a setting of minimum signal $a \leq c/\sqrt{n}$ can be solved by the proposed test.

3. Multiple testing of correlations with correlated observations.

As we mentioned in the introduction, when the independence hypothesis in (1) is rejected, there is potential risk of using inference methods developed based on independence assumption. To illustrate the effect of sample correlations, we study an important high-dimensional problem —the large-scale multiple testing of correlations when the samples are correlated, i.e.,

$$(24) \quad H_{0ij} : \rho_{ij} = 0 \quad \text{versus} \quad H_{1ij} : \rho_{ij} \neq 0, \quad \text{for } 1 \leq i < j \leq p.$$

When the samples are *i.i.d.* and normally distributed, the following classical result from Anderson (2003) (Theorem 4.2.4) establishes the limiting distribution of the sample correlation coefficient $\hat{\rho}_{ij}$,

$$(25) \quad \frac{\sqrt{n}(\hat{\rho}_{ij} - \rho_{ij})}{1 - \rho_{ij}^2} \Rightarrow N(0, 1).$$

However, when the samples are correlated, the limiting distribution of $\hat{\rho}_{ij}$ in (25) does not hold. In fact, we can prove the following proposition.

PROPOSITION 3.1. *Assume that the condition (C1) holds. We have*

$$(26) \quad \frac{\sqrt{n}(\hat{\rho}_{ij} - \rho_{ij})}{\sqrt{B_n}(1 - \rho_{ij}^2)} \Rightarrow N(0, 1),$$

where $B_n = \frac{\|\Psi\|_F^2}{n}$.

The term B_n is same quantity as in (11), which represents the average correlation among n samples. When the sample correlation is strong enough to extent such that $B_n \geq 1 + c > 1$, the multiple testing procedure based on

(25) (e.g., Benjamini—Hochberg (BH) procedure, [Benjamini and Hochberg \(1995\)](#)) will lead to many false positives. In fact, even when the correct limiting distribution in (26) is used, the resulting test will lose statistical power. For simplicity, let us consider a single testing problem $H_{0ij} : \rho_{ij} = 0$. To control the type I error rate when the samples are correlated, we need a larger critical value for $\hat{\rho}_{ij}$, which is linear in $\sqrt{B_n}$. That is, the rejection region should be $\{\hat{\rho}_{ij} : \sqrt{n}|\hat{\rho}_{ij}| \geq \sqrt{B_n}\Phi^{-1}(1 - \alpha/2)\}$, where $\Phi(\cdot)$ is the standard normal CDF function. Plugging in a ratio-consistent estimator of B_n (e.g., using the method developed in Section 2.2 to estimate $\|\Psi\|_F^2$), we will obtain a test that controls the type I error rate asymptotically. However, such a test will lose statistical power since the length of the acceptance region grows with the strength of the correlation among samples.

In this section, we propose a multiple testing procedure that asymptotically controls the FDR at the nominal level while maintaining good statistical power. Our method is based on the construction of a “sandwich estimator” of ρ_{ij} by de-correlating the samples. In particular, first assume that μ and Ψ are known. We transform the data \mathbf{X} into $\mathbf{Y} = (\mathbf{Y}_1, \dots, \mathbf{Y}_n) := (\mathbf{X} - \mu\mathbf{1}')\Psi^{-1/2} \sim N(0, \Sigma \otimes \mathbf{I}_{n \times n})$ and columns $\mathbf{Y}_k \in \mathbb{R}^p$ for $1 \leq k \leq n$ are *i.i.d.* from $N(0, \Sigma)$. The corresponding “sample” covariance matrix of \mathbf{Y} is (“sample” is quoted here since μ and Ψ are unknown and thus $(\tilde{\sigma}_{ij,Y})_{p \times p}$ is not a real sample covariance matrix):

$$(27) \quad (\tilde{\sigma}_{ij,Y})_{p \times p} = \frac{1}{n} \sum_{k=1}^n \mathbf{Y}_k \mathbf{Y}_k' = \frac{1}{n} (\mathbf{X} - \mu\mathbf{1}')\Psi^{-1}(\mathbf{X} - \mu\mathbf{1})'$$

Let $\tilde{\rho}_{ij,Y} = \frac{\tilde{\sigma}_{ij,Y}}{\sqrt{\tilde{\sigma}_{ii,Y}\tilde{\sigma}_{jj,Y}}}$ be the “sample” correlation coefficient matrix. By (25), we have

$$(28) \quad \frac{\sqrt{n}(\tilde{\rho}_{ij,Y} - \rho_{ij})}{1 - \rho_{ij}^2} \Rightarrow N(0, 1),$$

which implies that the performance of the test statistic $\tilde{\rho}_{ij,Y}$ is the same as that of $\hat{\rho}_{ij}$ for independent samples. By comparing (28) and (26), the asymptotic variance of the sandwich estimator $\tilde{\rho}_{ij,Y}$ is always smaller than that of the sample correlation coefficient as $B_n \geq 1$. Therefore, even when B_n is bounded by a constant, the sandwich estimator is more powerful.

To obtain an estimate of $\tilde{\rho}_{ij,Y}$, we need to estimate μ and $\Psi^{-1} := (\gamma_{ij})_{n \times n}$. Let $\hat{\mu} = (\bar{X}_1, \dots, \bar{X}_p)'$ be the estimator of μ , where $\bar{X}_i = \frac{1}{n} \sum_{k=1}^n X_{ki}$ for $1 \leq i \leq p$. For estimating Ψ^{-1} , we adopt the CLIME estimator proposed in [Cai, Liu and Luo \(2011\)](#). In particular, following [Cai, Liu and Luo \(2011\)](#),

we assume that Ψ^{-1} is a weakly sparse matrix, which belongs to the class,

$$(29) \quad \mathcal{G} = \left\{ \Psi^{-1} : \|\Psi^{-1}\|_{l_1} \leq M_n, \max_{1 \leq i \leq n} \left| \sum_{j=1}^n \psi_{ij} \right| \leq N_n, \sum_{j=1}^n |\gamma_{ij}|^q \leq s_n \right\},$$

where $0 \leq q < 1/2$, $\|\Psi^{-1}\|_{l_1} = \max_{1 \leq j \leq n} \sum_{i=1}^n |\gamma_{ij}|$ and the relationship among M_n , N_n and s_n will be specified in the condition of Theorem 3.2. Let $\hat{\mathbb{R}}_{\Psi} = (\hat{\psi}_{ij})_{n \times n}$, where $\hat{\psi}_{ij}$ is defined in (5), and $\hat{\Gamma}^1 = (\hat{\gamma}_{ij}^1)_{n \times n}$ be any optimal solution of the following optimization problem,

$$(30) \quad \min_{\Gamma \in \mathbb{R}^{n \times n}} \|\Gamma\|_1 \quad \text{subject to} \quad \|\hat{\mathbb{R}}_{\Psi} \Gamma - \mathbf{I}_{n \times n}\|_{\infty} \leq \lambda_{n,p}.$$

Here, $\lambda_{n,p} = cM_n(\frac{N_n}{n} + \sqrt{\frac{\log n}{p}})$, c is a sufficiently large constant, $\|\Gamma\|_1 = \sum_{1 \leq i, j \leq n} |\gamma_{ij}|$ and $\|\mathbf{A}\|_{\infty} = \max_{1 \leq i, j \leq n} |a_{ij}|$ for matrix $\mathbf{A} = (a_{ij})_{n \times n}$. We note that in the estimation of Ψ^{-1} , each row of \mathbf{X} is treated as a sample, and thus the sample size is p and the dimensionality is n . The estimator of Ψ^{-1} , $\hat{\Gamma} = (\hat{\gamma}_{ij})_{n \times n}$, is obtained by a symmetrization of $\hat{\Gamma}^1$: $\hat{\gamma}_{ij} = \hat{\gamma}_{ij}^1 I\{|\hat{\gamma}_{ij}^1| \leq |\hat{\gamma}_{ji}^1|\} + \hat{\gamma}_{ji}^1 I\{|\hat{\gamma}_{ij}^1| > |\hat{\gamma}_{ji}^1|\}$. Based on the estimated $\hat{\boldsymbol{\mu}}$ and $\hat{\Gamma}$, we define the ‘‘sandwich estimator’’ of $(\hat{\sigma}_{ij,Y})_{p \times p}$, $(\hat{\sigma}_{ij,Y})_{p \times p} = \frac{1}{n}(\mathbf{X} - \hat{\boldsymbol{\mu}}\mathbf{1}')\hat{\Gamma}(\mathbf{X} - \hat{\boldsymbol{\mu}}\mathbf{1}')$ with each

$$\hat{\sigma}_{ij,Y} = \frac{1}{n}(\mathbf{X}_{\cdot,i} - \bar{X}_i\mathbf{1})'\hat{\Gamma}(\mathbf{X}_{\cdot,j} - \bar{X}_j\mathbf{1}),$$

where $\mathbf{X}_{\cdot,i} = (X_{1i}, \dots, X_{ni})'$ is the i -th column of \mathbf{X} . The corresponding correlation coefficient

$$(31) \quad \hat{\rho}_{ij,Y} = \frac{\hat{\sigma}_{ij,Y}}{\sqrt{\hat{\sigma}_{ii,Y}\hat{\sigma}_{jj,Y}}}.$$

We note that the ‘‘sandwich estimator’’ $\hat{\rho}_{ij,Y}$ is related to the Knorm correlation proposed by [Teng and Huang \(2009\)](#), which estimates Ψ^{-1} in $\tilde{\rho}_{ij,Y}$ by the inverse of maximum likelihood estimator (MLE) of Ψ . However, there is no closed-form solution for MLE of the matrix-variate normal distribution. So it is difficult to develop limiting distribution results for the Knorm correlation in high-dimensional settings.

In the proof of Theorem 3.2, we will show that $\sqrt{n} \max_{1 \leq i \leq j \leq p} |\hat{\rho}_{ij,Y} - \tilde{\rho}_{ij,Y}| = o_{\mathbb{P}}(1/\sqrt{\log p})$. Combining it with (28), we have $\frac{\sqrt{n}(\hat{\rho}_{ij,Y} - \rho_{ij})}{1 - \rho_{ij}^2} \Rightarrow N(0, 1)$. Therefore, for each single test problem $H_{0ij} : \rho_{ij} = 0$, we propose the test statistic,

$$(32) \quad \hat{T}_{ij} = \sqrt{n}\hat{\rho}_{ij,Y}.$$

and the null H_{0ij} is rejected when $|\hat{T}_{ij}| \geq t$ for some threshold level $t > 0$.

To implement the large-scale multiple testing of correlations, we adopt the popular BH method (Benjamini and Hochberg, 1995). In particular, we need to search for a threshold \hat{t} for $|\hat{T}_{ij}|$ that controls the false discovery proportion (FDP) and false discovery rate (FDR) defined as follows while rejecting as many hypotheses as possible,

$$\text{FDP} = \frac{\sum_{(i,j) \in \mathcal{H}_0} I\{|\hat{T}_{ij}| \geq \hat{t}\}}{\max\{\sum_{1 \leq i < j \leq p} I\{|\hat{T}_{ij}| \geq \hat{t}\}, 1\}} \quad \text{and} \quad \text{FDR} = \mathbb{E}(\text{FDP}),$$

where $\mathcal{H}_0 = \{(i, j) : \rho_{ij} = 0, 1 \leq i < j \leq p\}$ is the set of null. Therefore, an ideal choice of the threshold level for a pre-specified significance level $0 < \alpha < 1$ should be

$$(33) \quad \hat{t}_{orc} = \inf \left\{ t > 0 : \frac{\sum_{(i,j) \in \mathcal{H}_0} I\{|\hat{T}_{ij}| \geq t\}}{\max\{\sum_{1 \leq i < j \leq p} I\{|\hat{T}_{ij}| \geq t\}, 1\}} \leq \alpha \right\}.$$

The oracle threshold level \hat{t}_{orc} cannot be computed since \mathcal{H}_0 is unknown. Nevertheless, since $\hat{T}_{ij} \Rightarrow N(0, 1)$ under the null $\rho_{ij} = 0$, the numerator in (33), $\sum_{(i,j) \in \mathcal{H}_0} I\{|\hat{T}_{ij}| \geq t\}$, can be approximated by $2(1 - \Phi(t))\text{Card}(\mathcal{H}_0)$. The quantity $\text{Card}(\mathcal{H}_0)$ can be further bounded from above by $(p^2 - p)/2$ and such an upper bound is good when Σ is sparse, which is a common setup. Therefore, we propose the following threshold level \hat{t} and the corresponding multiple testing procedure.

<p>For a given $0 < \alpha < 1$, let</p> $(34) \quad \hat{t} = \inf \left\{ t \geq 0 : \frac{(1 - \Phi(t))(p^2 - p)}{\max\{\sum_{1 \leq i < j \leq p} I\{ \hat{T}_{ij} \geq t\}, 1\}} \leq \alpha \right\}.$ <p>For $1 \leq i < j \leq p$, we reject H_{0ij} if $\hat{T}_{ij} \geq \hat{t}$.</p>
--

The next theorem shows that the proposed procedure controls the FDP and FDR at level α asymptotically. Recall the definition of \mathcal{H}_0 . Let $h_0 = \text{Card}(\mathcal{H}_0)$, $\mathcal{H}_1 = \{(i, j) : \rho_{ij} \neq 0, 1 \leq i < j \leq p\}$, $h_1 = \text{Card}(\mathcal{H}_1)$ and $h = (p^2 - p)/2$. For a given $\gamma > 0$, we further define the following sets

$$(35) \quad \mathcal{A}_i(\gamma) = \{j : 1 \leq j \leq p, j \neq i, |\rho_{ij}| \geq (\log p)^{-2-\gamma}\}, \quad 1 \leq i \leq p.$$

THEOREM 3.2. *Assume that the condition (C1) holds, $p \leq n^r$ for some $r > 0$, and $\Psi^{-1} \in \mathcal{G}$ defined in (29) with*

$$(36) \quad \frac{p}{(ns_n^2(\log p)^3)^{1/(1-q)} M_n^4 \log n} \rightarrow \infty \quad \text{and} \quad s_n = o\left(\frac{n^{1/2-q}}{M_n^{2-2q} N_n^{1-q} (\log p)^{3/2}}\right).$$

Suppose that $h_1 \leq \kappa h$ for some $\kappa < 1$,

$$(37) \text{ Card}\{(i, j) : 1 \leq i < j \leq p, |\rho_{ij}| \geq 4\sqrt{\log p/n}\} \geq \sqrt{\log \log p},$$

and $\max_{1 \leq i \leq p} \text{Card}(\mathcal{A}_i(\gamma)) = O(p^\rho)$ for some $\rho < 1/2$ and $\gamma > 0$. We have

$$\lim_{(n,p) \rightarrow \infty} \frac{\text{FDR}}{\alpha h_0/h} = 1 \quad \text{and} \quad \frac{\text{FDP}}{\alpha h_0/h} \rightarrow 1 \quad \text{in probability as } (n, p) \rightarrow \infty.$$

We briefly comment on the condition in Theorem 3.2. We note that in the estimation of Ψ^{-1} , n plays the role of dimensionality and p plays the role of the sample size. The condition in (36) ensures that p is sufficiently large so that the estimation of Ψ^{-1} is accurate. On the other hand, the assumption that p is sufficiently large is also natural for high-dimensional applications (e.g., genetic studies). The assumption that $h_1 \leq \kappa h$ for some $\kappa < 1$ is necessary. Since if $h_0 = o(h)$, then almost all of ρ_{ij} are non-zeros and simply rejecting all the hypotheses will lead to $\text{FDR} \rightarrow 0$. The condition in (37), which is only slightly stronger than the condition that the number of true alternatives goes to infinity, is a nearly necessary condition. In fact, Proposition 2.1 in Liu and Shao (2014) shows that if the number of true alternatives is fixed, then it is impossible for the BH method to control the FDP with probability tending to one at any desired level. The condition on $\max_{1 \leq i \leq p} \text{Card}(\mathcal{A}_i(\gamma))$ is essentially a sparsity condition for Σ . In particular, when $p \geq n^{r_1}$ with $r_1 > 1$ and the number of nonzero entries in each row of Σ is on the order of \sqrt{n} (which is a common assumption for sparse Σ), then the condition on $\max_{1 \leq i \leq p} \text{Card}(\mathcal{A}_i(\gamma))$ automatically holds.

4. Numerical results. In this section, we provide numerical results to demonstrate the performance of the proposed test methods. Due to space constraints, some simulations and real experiments are provided in the supplementary material (Chen and Liu, 2015). Recall that the $p \times n$ data matrix \mathbf{X} follows a matrix-variate normal distribution $N(\boldsymbol{\mu}\mathbf{1}', \Sigma \otimes \Psi)$. The matrix Σ (and Ψ) is generated from one of the following classes of matrices:

1. Auto-correlation matrix where $\sigma_{ij} = \rho^{|i-j|}$ and ρ is set to 0.2, 0.5 or 0.8. The larger the parameter ρ is, the stronger the correlation.
2. Banded matrix (“band” for short) where $\sigma_{ii} = 1$, $\sigma_{i,i+1} = \sigma_{i+1,i} = 0.6$, $\sigma_{i,i+2} = \sigma_{i+2,i} = 0.3$, and $\sigma_{ij} = 0$ for $|i-j| \geq 3$.
3. Block diagonal matrix (“block” for short) where the main diagonal blocks are 10×10 square matrices and off-diagonal blocks are zeros matrices. A 10×10 main diagonal block $\mathbf{B} = (b_{ij})_{10 \times 10}$ has $b_{ii} = 1$ and $b_{ij} = 0.5$ when $i \neq j$.

In simulations, we fix $\boldsymbol{\mu} = 0$ and the level of significance $\alpha = 0.05$.

4.1. *Independence Test.* We consider the independence test problem in (1). All the reported empirical sizes and powers are averaged over 5,000 independent replications. In Table 1, we consider relatively large n and p and show the empirical type I error rate (a.k.a. the empirical size) of the proposed test statistics $\hat{T}_{n,p}$ in (10) under the null when $\Psi = \mathbf{I}_{n \times n}$. From Table 1, as the sample size n and dimension p increase, the empirical type I error rates get closer to the nominal level of 0.05, which verifies the validity of the proposed test statistics shown in Theorem 2.4.

Recalling in the construction of $\hat{T}_{n,p}$ (in particular, in the term \hat{A}_p), we threshold the sample covariance matrix $\hat{\Sigma}$ as in (12), where the threshold level involves the estimator \hat{B}_n of $B_n = \|\Psi\|_{\mathbb{F}}^2/n$. We compare the empirical sizes of the test statistics in the same form as $\hat{T}_{n,p}$ in (10) but using different estimators (listed as follows) of $\|\Sigma\|_{\mathbb{F}}^2$ in estimating $A_p = \frac{p\|\Sigma\|_{\mathbb{F}}^2}{(\text{tr}(\Sigma))^2}$:

1. CV (Fan, Rigollet and Wang, 2015): plugin estimator based on thresholded Σ with the threshold level tuned by cross validation (CV).
2. Bai: method proposed by Bai and Saranadasa (1996).
3. CQ: method proposed by Chen and Qin (2010).
4. \hat{B}_n : plugin estimator based on thresholded Σ as in (12) with the proposed estimator \hat{B}_n for setting the threshold level.

We would like to make it clear that for the ease of presentation, *the acronyms CV, Bai and CQ refer to the proposed test statistics in the form of $\hat{T}_{n,p}$ while using the corresponding method to construct the estimator of $\|\Sigma\|_{\mathbb{F}}^2$ in A_p .*

In Table 2, we show the comparison results when $n = 50$ or 100 and $p = 1000$. Of note, we only present smaller n and p cases since the computational cost of CV and CQ are expensive for large n and p and the case $n = 50/100$ and $p = 1000$ has been sufficient to demonstrate the points below. We show that the CV cannot control the type I error below the nominal level 0.05. The CQ leads the type I error rates that are closest to the nominal level. However, as we will show later, it has a lower statistical power. The empirical sizes of the proposed test statistics (with the thresholding level \hat{B}_n in estimating $\|\Psi\|_{\mathbb{F}}^2$) are below the nominal level, which shows that the proposed test statistic is conservative when n and p are *small*. This results from the slow rate of convergence in distribution for the max-type test statistics (Liu, Lin and Shao, 2008). For small n and p , one useful way to make the test less conservative is to adopt the critical value from a Monte-Carlo simulation instead of the one derived from the limiting distribution. In particular, we can generate M (e.g., $M = 10000$ in our simulation) replications of $p \times n$ data matrix, where each one is randomly drawn from $N(\mathbf{0}, \mathbf{I}_{p \times p} \otimes \mathbf{I}_{n \times n})$ under the null. We compute the corresponding test statistics $\hat{T}_{n,p}^{(i)}$, $1 \leq i \leq M$, for each

TABLE 1
Empirical type I error rates for testing independence based on 5,000 replications with $\alpha = 0.05$

n	Σ	Ψ	$p = 1,000$	$p = 2,000$	$p = 5,000$	$p = 10,000$
200	$0.2^{ i-j }$	$\mathbf{I}_{n \times n}$	0.046	0.046	0.042	0.043
	$0.5^{ i-j }$	$\mathbf{I}_{n \times n}$	0.040	0.049	0.049	0.050
	$0.8^{ i-j }$	$\mathbf{I}_{n \times n}$	0.045	0.048	0.055	0.058
	band	$\mathbf{I}_{n \times n}$	0.031	0.032	0.035	0.043
	block	$\mathbf{I}_{n \times n}$	0.014	0.025	0.030	0.035
500	$0.2^{ i-j }$	$\mathbf{I}_{n \times n}$	0.034	0.041	0.042	0.046
	$0.5^{ i-j }$	$\mathbf{I}_{n \times n}$	0.037	0.046	0.041	0.049
	$0.8^{ i-j }$	$\mathbf{I}_{n \times n}$	0.028	0.050	0.048	0.055
	band	$\mathbf{I}_{n \times n}$	0.032	0.035	0.038	0.040
	block	$\mathbf{I}_{n \times n}$	0.016	0.025	0.041	0.044
1000	$0.2^{ i-j }$	$\mathbf{I}_{n \times n}$	0.039	0.035	0.048	0.044
	$0.5^{ i-j }$	$\mathbf{I}_{n \times n}$	0.035	0.042	0.056	0.054
	$0.8^{ i-j }$	$\mathbf{I}_{n \times n}$	0.026	0.040	0.051	0.050
	band	$\mathbf{I}_{n \times n}$	0.029	0.037	0.040	0.045
	block	$\mathbf{I}_{n \times n}$	0.016	0.024	0.035	0.041

TABLE 2
Comparison of empirical type I error rates for testing independence when $p = 1000$ and $\alpha = 0.05$

n	Σ	Ψ	CV	Bai	CQ	\hat{B}_n
50	$0.2^{ i-j }$	$\mathbf{I}_{n \times n}$	0.029	0.001	0.038	0.002
	$0.5^{ i-j }$	$\mathbf{I}_{n \times n}$	0.129	0.004	0.039	0.008
	$0.8^{ i-j }$	$\mathbf{I}_{n \times n}$	0.157	0.010	0.022	0.037
	band	$\mathbf{I}_{n \times n}$	0.072	0.005	0.033	0.007
	block	$\mathbf{I}_{n \times n}$	0.249	0.011	0.028	0.017
100	$0.2^{ i-j }$	$\mathbf{I}_{n \times n}$	0.070	0.017	0.041	0.028
	$0.5^{ i-j }$	$\mathbf{I}_{n \times n}$	0.118	0.020	0.036	0.034
	$0.8^{ i-j }$	$\mathbf{I}_{n \times n}$	0.101	0.019	0.025	0.049
	band	$\mathbf{I}_{n \times n}$	0.066	0.017	0.034	0.025
	block	$\mathbf{I}_{n \times n}$	0.068	0.017	0.023	0.016

randomly generated data matrix and let c_α be the $(1 - \alpha)$ -quantile of the empirical distribution $\frac{1}{M} \sum_{i=1}^M I\{\hat{T}_{n,p}^{(i)} \leq t\}$. We reject the null whenever the our test statistic $\hat{T}_{n,p} \geq c_\alpha$ (note that the statistic is the same and only the critical value is changed). As shown in the additional experimental results in Section E.1 in the supplement, using a Monte-Carlo based critical value will push the empirical size closer to the nominal α when n and p are small.

Then, we compare statistical power of the proposed test procedure when using different estimators of $\|\Sigma\|_{\mathbb{F}}^2$ in our test statistic. In particular, we first consider $\Psi = (\rho^{|i-j|})_{n \times n}$ and vary the parameter ρ from 0.55 to 0.85. The larger the ρ is, the stronger the correlation among samples. For different

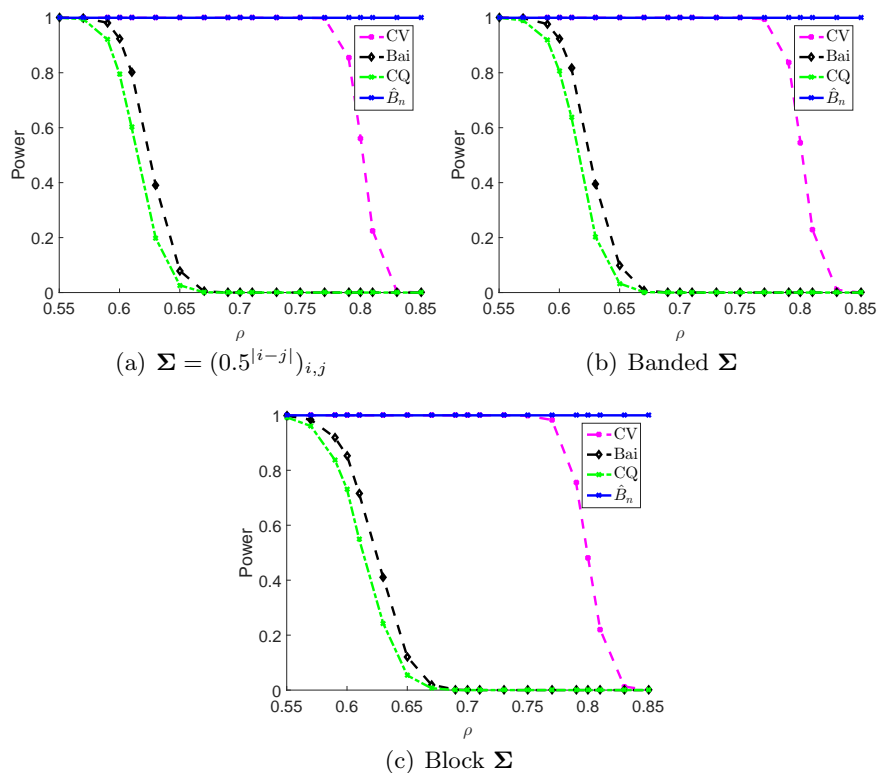


Fig 1: Comparison of empirical powers when using different estimators of $\|\Sigma\|_{\text{F}}^2$ in \hat{A}_p . The Ψ matrix is set to the auto-correlation matrix where $\psi_{ij} = \rho^{|i-j|}$ and we vary ρ (corresponding to x -axis in each figure) from 0.55 to 0.85. Here, $n = 50$, $p = 1000$ and $\alpha = 0.05$.

types of Σ , the empirical powers are all 100% for our method (Figure 1). The powers using Bai and CQ drop to zeros when ρ becomes larger than 0.7. Since both methods for estimating $\|\Sigma\|_{\text{F}}^2$ are developed under the *i.i.d.* assumption, when the sample correlation becomes stronger, the estimation of $\|\Sigma\|_{\text{F}}^2$ is inaccurate, which leads to inferior statistical powers. The CV-based thresholding method has maintained statistical power 100% for a wider range of ρ . However, we note that the CV fails to control the type I error rate as shown in Table 2. In Section E.2 in the supplementary material, we further demonstrate the superiority of using the proposed estimator for $\|\Sigma\|_{\text{F}}^2$ in terms of empirical powers when Ψ is a block diagonal matrix.

It is also of interest to investigate the performance of the proposed test statistics when n and p are comparable. We vary p from 50 to 2,000 and consider four settings for the sample size, $n = 0.5p$, $n = p$, $n = 2p$ and

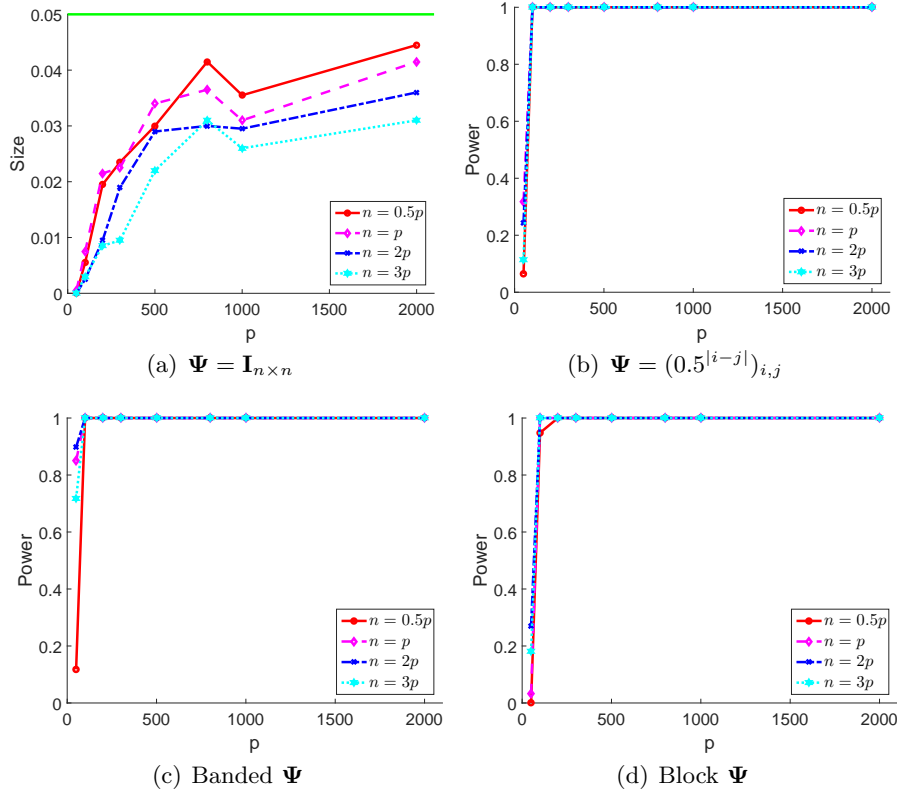


Fig 2: The empirical type I error rates and powers when $n = 0.5p$, $n = p$, $n = 2p$ and $n = 3p$ and p varies from 50 to 2,000. Panel 2(a) shows the empirical type I error rate and the green line indicates the nominal level $\alpha = 0.05$. Panels 2(b)-2(d) show the empirical powers for different Ψ .

$n = 3p$. We set $\Sigma = (0.5^{|i-j|})_{i,j}$ and show the empirical type I error rates and powers for different Ψ in Figure 2. As one can see from Figure 2(a), the empirical type I error rates are approaching the nominal level $\alpha = 0.05$ as p increases. Notably, when the ratio between n and p increases, the test statistic becomes more conservative. From Figure 2(b)-2(d), although the powers are low when p is very small (i.e., $p = 50$), they are 100% for moderate and large p . This simulation study suggests that the proposed independence test performs reasonably well when n and p are comparable.

Moreover, we consider the setting in which Σ does not satisfy the conditions (C1)-(C2). In particular, we choose an equicorrelation covariance matrix $\Sigma = 0.85 \cdot \mathbf{1}\mathbf{1}' + 0.15 \cdot \mathbf{I}_{p \times p}$, which enforces very strong correlation among every pairs of variables. It is easy to see that $\lambda_{\max}(\Sigma) = 0.85 \cdot p + 0.15$




Fig 3: The empirical type I error rates when $\Sigma = 0.85 \cdot \mathbf{1}\mathbf{1}' + 0.15 \cdot \mathbf{I}_{p \times p}$ for $n = 50, 100, 200$ and p varying from 1000 to 10,000. The green line indicates the nominal level $\alpha = 0.05$.

and $\sum_{k=1}^p |\sigma_{jk}|^\tau = 1 + 0.85^\tau \cdot (p - 1)$. Both quantities are linear in p and thus cannot be bounded by constants as p grows, which violates the assumptions **(C1)**-**(C2)**. Hence, it is expected that the results of type I error rate control in Proposition 2.3 and Theorem 2.4 will not hold for this model. This is verified by our simulation study. In particular, we vary n and p and show the type I error rates in Figure 3. When p grows and the conditions **(C1)**-**(C2)** no longer hold, the type I error rate exceeds the nominal level $\alpha = 0.05$ (represented by the green line).

Due to space limitations, we relegate the other settings of Ψ and some additional simulation studies for independence testing to Section E in the supplementary material, which includes:

1. We compare empirical powers when the Ψ is a block diagonal matrix and demonstrate the superiority of the proposed method.
2. To empirically verify the result in Theorem 2.5, we consider the case of extremely sparse Ψ where $\psi_{12} = \psi_{21} = \kappa \sqrt{\frac{\log n}{p}}$ and all the other off-diagonal elements are zeros. The experimental results show that for different types of Σ , the empirical powers all become 100% as κ increases, which demonstrates that our test statistic can successfully reject the null even when the Ψ is extremely sparse.
3. To provide more intuitive comparisons between different methods for estimating $\|\Sigma\|_{\text{F}}^2$, we directly show the relative estimation error under different settings. This experiment demonstrates that the proposed thresholding estimator greatly outperforms its competitors when the correlation among samples is large.

TABLE 3
Averaged FDP and power for testing correlations over 100 replications. Here, $\alpha = 0.05$
and $p = 1000$

n	Σ	Ψ	$\sqrt{n}\hat{\rho}_{ij,Y}$		$\sqrt{n}\hat{\rho}_{ij}$		$\frac{\sqrt{n}\hat{\rho}_{ij}}{\sqrt{B_n}}$		True Ψ^{-1}	
			FDP	Power	FDP	Power	FDP	Power	FDP	Power
50	band	$0.2^{ i-j }$	0.010	0.311	0.027	0.339	0.010	0.276	0.015	0.339
		$0.5^{ i-j }$	0.009	0.308	0.403	0.378	0.000	0.000	0.015	0.340
	block	$0.2^{ i-j }$	0.011	0.262	0.036	0.416	0.014	0.295	0.021	0.407
		$0.5^{ i-j }$	0.011	0.257	0.366	0.504	0.000	0.000	0.021	0.410
	block	$0.8^{ i-j }$	0.012	0.168	0.965	0.739	0.010	0.000	0.021	0.408
		$0.2^{ i-j }$	0.025	0.576	0.057	0.593	0.029	0.568	0.033	0.587
100	band	$0.5^{ i-j }$	0.025	0.575	0.581	0.650	0.011	0.448	0.032	0.587
		$0.8^{ i-j }$	0.025	0.556	0.986	0.795	0.000	0.000	0.032	0.587
	block	$0.2^{ i-j }$	0.028	0.942	0.061	0.961	0.035	0.945	0.038	0.966
		$0.5^{ i-j }$	0.028	0.935	0.454	0.942	0.017	0.646	0.038	0.966
	block	$0.8^{ i-j }$	0.030	0.820	0.963	0.924	0.000	0.000	0.038	0.966
		$0.2^{ i-j }$	0.036	0.839	0.072	0.852	0.039	0.820	0.041	0.854
200	band	$0.5^{ i-j }$	0.036	0.835	0.620	0.867	0.028	0.640	0.041	0.853
		$0.8^{ i-j }$	0.041	0.749	0.984	0.906	0.002	0.240	0.042	0.852
	block	$0.2^{ i-j }$	0.034	1.000	0.071	1.000	0.041	1.000	0.043	1.000
		$0.5^{ i-j }$	0.034	1.000	0.498	1.000	0.033	0.992	0.044	1.000
	block	$0.8^{ i-j }$	0.040	0.969	0.962	0.994	0.004	0.233	0.044	1.000

TABLE 4
Averaged FDP and power for testing correlations over 100 replications when samples are
i.i.d. Here, $\alpha = 0.05$ and $p = 1000$.

n	Σ	Ψ	$\sqrt{n}\hat{\rho}_{ij,Y}$		$\sqrt{n}\hat{\rho}_{ij}$	
			FDP	Power	FDP	Power
50	band	$\mathbf{I}_{n \times n}$	0.009	0.318	0.014	0.341
		block	$\mathbf{I}_{n \times n}$	0.012	0.293	0.021
100	band	$\mathbf{I}_{n \times n}$	0.025	0.578	0.032	0.587
		block	$\mathbf{I}_{n \times n}$	0.028	0.952	0.037
200	band	$\mathbf{I}_{n \times n}$	0.035	0.844	0.041	0.852
		block	$\mathbf{I}_{n \times n}$	0.035	1.000	0.043

4.2. *Large-scale Multiple Testing of Correlations.* In this section, we conduct both simulated and real data analysis to demonstrate performance of the proposed “sandwich” estimator in (32) for large-scale multiple testing of correlations in (24).

4.2.1. *FDP and Power of Simulated Results.* In simulated study, we compare the BH procedure based on four different estimators of ρ_{ij} :

1. The proposed sandwich estimator $\hat{T}_{ij}(\lambda_{n,p}) = \sqrt{n}\hat{\rho}_{ij,Y}$ in (32), where

Ψ^{-1} is estimated by CLIME (Cai, Liu and Luo, 2011). Further, we adopt the data-driven approach in Liu (2013) to tune the $\lambda_{n,p}$ in CLIME (see (30)). In particular, the parameter $\lambda_{n,p}$ is selected by,

$$\hat{\lambda}_{n,p} = \arg \min_{\lambda} \sum_{k=3}^9 \left(\frac{\sum_{1 \leq i \neq j \leq p} I \left\{ |\hat{T}_{i,j}(\lambda)| \geq \Phi^{-1}(1 - k/20) \right\}}{k(p^2 - p)/10} - 1 \right)^2.$$

2. The classical sample correlation estimator $\sqrt{n}\hat{\rho}_{ij}$ based on sample independence assumption.
3. The variance corrected sample correlation estimator $\frac{\sqrt{n}\hat{\rho}_{ij}}{\sqrt{B_n}}$, where true $B_n = \|\Psi\|_{\text{F}}^2/n$ is assumed to be known.
4. The proposed sandwich estimator in (32) with the true Ψ^{-1} , which serves as an oracle benchmark.

In Table 3, we report the averaged FDP and power over 100 replications. The matrix Σ is chosen to be either banded or block diagonal matrix, both of which are sparse. As we can see from Table 3, the FDPs of the BH procedure based on sandwich estimator are below $\alpha = 0.05$. The empirical powers get close to one as the sample size n increases and are only slightly worse than the powers of the oracle benchmark with true Ψ^{-1} . For the classical sample correlation estimator $\sqrt{n}\hat{\rho}_{ij}$, the FDP can be very large (e.g., around 50% when $\psi_{ij} = 0.5^{|i-j|}$ and more than 95% when $\psi_{ij} = 0.8^{|i-j|}$). This verifies our result showing that naïvely using the sample correlation estimator developed under the sample independence assumption will lead to many false positives. Using the variance corrected sample correlation estimator $\frac{\sqrt{n}\hat{\rho}_{ij}}{\sqrt{B_n}}$ will help reduce the number of false positives and control FDP as shown in Table 3, which is consistent with our result in Proposition 3.1. However, as we observe from Table 3, even when the true B_n is used, the powers of $\frac{\sqrt{n}\hat{\rho}_{ij}}{\sqrt{B_n}}$ are quite low, especially when the correlation among samples becomes stronger. The reason for this low power is explained in the paragraph below Proposition 3.1.

In Table 4, we consider the setting when the samples are *i.i.d.*, in which case the classical sample correlation estimator should be used as it is based on sample independence assumption. We also note that when samples are *i.i.d.*, both the variance corrected sample correlation estimator $\frac{\sqrt{n}\hat{\rho}_{ij}}{\sqrt{B_n}}$ ($B_n = 1$) and the sandwich estimator with true $\Psi^{-1} = \mathbf{I}_{n \times n}$ reduce to the classical sample correlation estimator. The power when using the sandwich estimator with the estimated Ψ^{-1} by CLIME is quite close to the power when using the benchmark sample correlation estimator (Table 4), which demonstrates the robustness of the proposed method.

We also conducted real experiments on correlation tests for yeast genomics data and stock data, which are detailed in Section E.5 in the supplementary material.

5. Discussion. This paper studies the sample/column independence test and multiple testing of Pearson’s correlation coefficients in a high-dimensional setting. The main difficulty in column independence test arises from the correlation among different variables, which is characterized by the covariance matrix Σ . If Σ is known, the data matrix can be transformed as $\Sigma^{-1/2}\mathbf{X} \sim N(\Sigma^{-1/2}\boldsymbol{\mu}\mathbf{1}', \mathbf{I}_{p \times p} \otimes \Psi)$, based on which the independence test can be directly carried out using existing approaches (e.g., Jiang (2004); Liu, Lin and Shao (2008)). However, the covariance matrix Σ is unknown. Although the problem of estimating Σ^{-1} has been well studied, the optimal convergence rate in matrix ℓ_1 -norm is known to be $O(s_p \|\Sigma^{-1}\|_{\ell_1} \sqrt{(\log p)/n})$, where s_p is the row sparsity level of Σ^{-1} (see, e.g., Cai, Liu and Zhou (2015)). However, such a rate is not fast enough for establishing a limiting null distribution of the test statistic based on the estimated Σ^{-1} . In particular, from the proof of Theorem 2.4, when using max-type test statistics, to eliminate the effect of the estimation error from Σ^{-1} and establish a limiting null distribution, the convergence rate needs to be $o_{\mathbb{P}}(1/\sqrt{p \log n})$. As p can be $\exp(o(n^\gamma))$ for some $\gamma > 0$ in an ultra high-dimensional setting, one cannot solve the independence test problem in (1) by simply plugging in the estimator of Σ^{-1} . On the other hand, when using the row sample correlation matrix $(\hat{\psi}_{ij})$ by treating each row of \mathbf{X} as a sample, we only need to estimate $\|\Sigma\|_{\mathbb{F}}^2$ instead of Σ^{-1} . The problem of estimating $\|\Sigma\|_{\mathbb{F}}^2$ from correlated samples has been successfully addressed in Section 2.2. We would also like to note that in the multiple testing problem of Pearson’s correlation coefficients, such a difficulty no longer exists. In fact, when estimating Ψ^{-1} from row samples of \mathbf{X} , the roles of n and p has interchanged (i.e., the sample size becomes p and the dimensionality becomes n) and, thus, the estimation problem is conducted in a relatively lower dimensional setting.

References.

- ALLEN, G. I. and TIBSHIRANI, R. (2012). Inference with transposable data: modelling the effects of row and column correlations. *J. R. Statist. Soc. B* **74** 721–734.
- ANDERSON, T. W. (2003). *An introduction to multivariate statistical analysis*, 3rd ed. Wiley.
- BAI, Z. and SARANADASA, H. (1996). Effect of high dimension: by an example of a two sample problem. *Statistica Sinica* **6** 311–329.
- BAI, Z., JIANG, D., YAO, J. and ZHENG, S. (2009). Corrections to LRT on large-dimensional covariance matrix by RMT. *Annals of Statistics* **37** 3822–3840.
- BENJAMINI, Y. and HOCHBERG, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. R. Statist. Soc. B* **57** 389–300.

- BIEN, J. and TIBSHIRANI, R. (2011). Sparse estimation of a covariance matrix. *Biometrika* **98** 807–820.
- CAI, T. and JIANG, T. (2011). Limiting laws of coherence of random matrices with applications to testing covariance structure and construction of compressed sensing matrices. *Annals of Statistics* **39** 1496–1525.
- CAI, T. and JIANG, T. (2012). Phase Transition in Limiting Distributions of Coherence of High-Dimensional Random Matrices. *Journal of Multivariate Analysis* **107** 24–39.
- CAI, T. and LIU, W. (2011). Adaptive thresholding for sparse covariance matrix estimation. *Journal of the American Statistical Association* **106** 672–684.
- CAI, T., LIU, W. and LUO, X. (2011). A constrained ℓ_1 minimization approach to sparse precision matrix estimation. *Journal of the American Statistical Association* **106** 594–607.
- CAI, T. and LIU, W. (2015). Large-scale multiple testing of correlations. *Journal of the American Statistical Association*.
- CAI, T., LIU, W. and ZHOU, H. (2015). Estimating structured high-dimensional covariance and precision matrices: optimal rates and adaptive estimation. *Annals of Statistics* (forthcoming).
- CAI, T., REN, Z. and ZHOU, H. (2016). Estimating Structured High-Dimensional Covariance and Precision Matrices: Optimal Rates and Adaptive Estimation. *Electronic Journal of Statistics* **10** 1–89.
- CARTER, S. L., BRECHBÜHLER, C. M., GRIFFIN, M. and BOND, A. T. (2004). Gene co-expression network topology provides a framework for molecular characterization of cellular state. *Bioinformatics* **20** 2242–2250.
- CHEN, X. and LIU, W. (2015). Supplement to “Testing independence with high-dimensional correlated samples”.
- CHEN, S. and QIN, Y. (2010). A two-sample test for high-dimensional data with applications to gene-set testing. *Annals of Statistics* **38** 808–835.
- DAWID, A. P. (1977). Spherical Matrix Distributions and a Multivariate Model. *Journal of the Royal Statistical Society, Series B (Methodological)* **39** 254–261.
- DAWID, A. P. (1981). Some matrix-variate distribution theory: notational considerations and a bayesian application. *Biometrika* **68** 265–274.
- EFRON, B. (2009). Are a set of microarrays independent of each other? *Annals of Applied Statistics* **3** 922–942.
- FAN, J., RIGOLLET, P. and WANG, W. (2015). Estimation of functionals Of sparse covariance matrices. *Annals of Statistics* **43** 2706–2737.
- FANG, K. T. and ZHANG, Y. T. (1990). *Generalized Multivariate Analysis*. Springer.
- HAN, F. and LIU, H. (2014). Distribution-free tests of independence with applications to testing more structures. arXiv:1410.4179v1.
- HIRAI, M. Y. E. A. (2007). Omics-based identification of Arabidopsis Myb transcription factors regulating aliphatic glucosinolate biosynthesis. *Proceedings of the National Academy of Sciences* **104** 6478–6483.
- HONG, Y. (1998). Testing for pairwise serial independence via the empirical distribution function. *Journal of the Royal Statistical Society, Series B* **60** 429–453.
- JIANG, T. (2004). The asymptotic distributions of the largest entries of sample correlation matrices. *The Annals of Applied Probability* **14** 865–880.
- JIANG, T. and YANG, F. (2013). Central limit theorems for classical likelihood ratio tests for high-dimensional normal distributions. *Ann. Statist.* **41** 2029–2074.
- JOHNSTONE, I. M. (2001). On the distribution of the largest eigenvalue in principal components analysis. *Annals of Statistics* **29** 295–327.
- KIM, K., JIANG, K., TENG, S. L., FELDMAN, L. J. and HUANG, H. (2012). Using biologi-

- cally interrelated experiments to identify pathway genes in Arabidopsis. *Bioinformatics* **28** 815–822.
- LAM, C. and FAN, J. (2009). Sparsistency and rates of convergence in large covariance matrix estimation. *Annals of Statistics* **37** 4254–4278.
- LAZZERONI, L. and OWEN, A. (2002). Plaid models for gene expression data. *Statistica Sinica* **12** 61–86.
- LEDOIT, O. and WOLF, M. (2002). Some hypothesis tests for the covariance matrix when the dimension is large compared to the sample size. *Annals of Statistics* **30** 1081–1102.
- LEE, H. K., HSU, A. K. and SAJDAK, J. (2004). Coexpression analysis of human genes across many microarray data sets. *Genome Research* **14** 1085–1094.
- LIU, W. (2013). Gaussian graphical model estimation with false discovery rate control. *Annals of Statistics* **41** 2948–2978.
- LIU, W. D., LIN, Z. Y. and SHAO, Q. M. (2008). The asymptotic distribution and Berry-Esseen bound of a new test for independence in high dimension with an application to stochastic optimization. *Annals of Applied Probability* **18** 2337–2366.
- LIU, W. and SHAO, Q. (2014). Phase transition and regularized bootstrap in large-scale t-tests with false discovery rate control. *Annals of Statistics* **42** 2003–2025.
- MURALIDHARAN, O. (2010). Detecting column dependence when rows are correlated and estimating the strength of the row correlation. *Electronic Journal of Statistics* **4** 1527–1546.
- NAGAO, H. (1973). On some test criteria for covariance matrix. *Annals of Statistics* **1** 700–709.
- PAN, G. M., GAO, J. and YANG, Y. R. (2014). Testing independence among a large number of high dimensional random vectors. *Journal of American Statistical Association* **109** 600–612.
- ROTHMAN, A. J., LEVINA, E. and ZHU, J. (2009). Generalized thresholding of large covariance matrices. *Journal of the American Statistical Association* **104** 177–186.
- SCHOTT, J. (2005). Testing for complete independence in high dimensions. *Biometrika* **92** 951–956.
- SHAW, E. A. P. (2006). Intellectual ability and cortical development in children and adolescents. *Nature* **440** 676–679.
- TENG, S. L. and HUANG, H. (2009). A statistical framework to infer functional gene associations from multiple biologically interrelated microarray experiments. *Journal of the American Statistical Association* **104** 465–473.
- YIN, J. and LI, H. (2012). Model selection and estimation in matrix normal graphical model. *Journal of Multivariate Analysis* **107** 119–140.
- ZHOU, S. (2014). Gemini: Graph estimation with matrix variate normal instances. *Annals of Statistics* **42** 532–562.
- ZHU, D., HERO, A. O., QIN, Z. S. and SWAROOP, A. (2005). High throughput screening of co-expressed gene pairs with controlled false discovery rate (FDR) and minimum acceptable strength (MAS). *Journal of Computational Biology* **12** 1029–1045.

STERN SCHOOL OF BUSINESS
 NEW YORK UNIVERSITY
 44 WEST 4TH STREET
 NEW YORK, NY, USA
 E-MAIL: xchen3@stern.nyu.edu

DEPARTMENT OF MATHEMATICS
 INSTITUTE OF NATURAL SCIENCES AND MOE-LSC
 SHANGHAI JIAO TONG UNIVERSITY
 SHANGHAI, CHINA
 E-MAIL: weidongli@sjtu.edu.cn