Title	Locally Most Powerful Naive Test for the Identity of the Covariance Matrix
Manuscript ID	SS-2026-0001
URL	http://www.stat.sinica.edu.tw/statistica/
DOI	10.5705/ss.202026.0001
Complete List of Authors	Xiaona Lu, Jiang Hu, Zhidong Bai and Haiyan Song
Corresponding Authors	Jiang Hu
E-mails	huj156@nenu.edu.cn
Notice: Accepted author version.	

Statistica Sinica

Locally most powerful naive test for the identity of the covariance matrix

Xiaona Lu, Jiang Hu*, Zhidong Bai# and Haiyan Song

Northeast Normal University

Abstract: We develop a distribution-free, nonparametric family of optional naive tests for high-dimensional covariance identity testing. Within the naive test framework of Hu and Bai (2016), we consider an operator-indexed family of optional naive tests defined by a positive-definite operator $\mathbf{W}$ acting on the covariance structure. Classical U-statistic-based identity tests, such as that of Chen et al. (2010), arise as isotropic choices of $\mathbf{W}$, thereby embedding standard Frobenius-norm-type procedures as special cases within a single operator-based framework. We establish high-dimensional limiting distributions under the null and alternatives and, to our knowledge, provide the first rigorous formalization of a locally most powerful naive test (LMPNT) for high-dimensional covariance identity testing. We propose a spectral optimization criterion for selecting $\mathbf{W}$ by aligning its eigenvectors and eigenvalues with the departure spectrum; at the oracle level, it yields an explicit operator $\mathbf{W}_{\mathbf{\Sigma}_{1p},\mathrm{opt}}$ that is LMPNT over a structured class of alternatives. Extensive Monte Carlo experiments under Gaussian

* Corresponding author. https://orcid.org/0000-0001-5603-2798

https://orcid.org/0000-0002-5300-5513

and non-Gaussian designs show accurate size control and substantial power gains over existing methods.

Key words and phrases: High-dimensional covariance identity testing, Locally most powerful naive test, Optional naive test, Spectral optimization, U-statistic.

1. Introduction

Testing hypotheses about covariance matrices is a classical and well-established topic in multivariate analysis. Traditional multivariate methods are typically developed under the asymptotic regime where the sample size n tends to infinity while the dimension p remains fixed; see, for example, Anderson (2003). However, modern applications in genomics, signal processing and finance routinely involve data sets whose dimensionality is comparable to or even much larger than the sample size. In such high-dimensional settings, many classical procedures either break down theoretically or suffer from severe loss of power, and substantial efforts have been devoted to constructing new tests together with an appropriate asymptotic theory.

In this paper, we consider the high-dimensional covariance identity testing problem. Let $\mathbf{x}_1, \ldots, \mathbf{x}_n$ be independent and identically distributed p-dimensional observations with mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. We are interested in

$$H_0 : \mathbf{\Sigma} = \mathbf{I}_p \quad \text{vs.} \quad H_1 : \mathbf{\Sigma} = \mathbf{\Sigma}_{1p} \neq \mathbf{I}_p, \tag{1.1}$$

where $\mathbf{I}_p$ denotes the $p \times p$ identity matrix and $\mathbf{\Sigma}_{1p}$ is a positive-definite covariance matrix under the alternative. More generally, if the null hypothesis is $H_0 : \mathbf{\Sigma} = \mathbf{\Sigma}_0$ for a fixed positive-definite matrix $\mathbf{\Sigma}_0$, then applying the linear transformation $\mathbf{\Sigma}_0^{-1/2}\mathbf{x}_i$ reduces the problem to (1.1). We therefore focus on the identity case without loss of generality.

In particular, hypotheses of the form (1.1) arise naturally in the analysis of gene-expression microarray data and the associated large-scale multiple testing. A common working simplification in such settings is that gene-expression measurements are approximately independent across genes, or at least exhibit a relatively simple covariance structure. It is therefore statistically prudent to examine whether the covariance structure of the data is compatible with the identity model, for example by testing $H_0 : \mathbf{\Sigma} = \mathbf{I}_p$, before applying high-dimensional inferential procedures that rely on such assumptions.

In the classical low-dimensional Gaussian setting, with $\mathbf{S}_n$ denoting the unbiased centered sample covariance matrix, the likelihood ratio statistic

$$L_n = \mathrm{tr}(\mathbf{S}_n) - \log|\mathbf{S}_n| - p$$

is a natural candidate for testing covariance identity, and its scaled version is asymptotically chi-square; see Anderson (2003), Nagao (1973) and Wilks (1932). However, when p grows with n, the chi-square calibration breaks down and the uncorrected likelihood ratio test may suffer from severe size distortion and power loss; see Bai et al. (2009).

To address this breakdown of classical procedures, a substantial literature has developed high-dimensional identity tests, which can be broadly grouped into three lines. The first line consists of spectral procedures calibrated by random matrix theory. Early developments along this direction include corrected likelihood ratio tests (CLRT) and their variants; see Bai et al. (2009), Jiang and Yang (2013), Jiang and Qi (2015) and Zheng et al. (2015), among others. These works establish asymptotic normality of suitably centered and scaled LRT-type statistics when $p/n \to c \in (0, 1)$, but the resulting procedures are not directly applicable in the $p > n$ regime. More recently, Ding and Wang (2025) established global and local central limit theorems for linear spectral statistics (LSS) of general sample covariance matrices in ultra-high-dimensional regimes where the dimension may be much larger than the sample size, and used these results to develop tests based on linear spectral statistics. Relatedly, Ding et al. (2025) studied two-sample covariance testing in ultra-high dimensions. These recent developments further highlight the importance of random matrix theory in high-dimensional and ultra-high-dimensional covariance testing. A second line of research develops tests based on quadratic functionals of $\mathbf{\Sigma}-\mathbf{I}_p$, typically measured by the squared Frobenius norm. Early contributions include Nagao's trace-type statistic (Nagao, 1973) and its Gaussian refinement by Ledoit and Wolf (2002). A prominent milestone is the unbiased U-statistic of Chen et al. (2010), which estimates $\|\mathbf{\Sigma}-\mathbf{I}_p\|_F^2$ and yields a simple, powerful and nonparametric procedure that remains valid in high-dimensional settings where both the dimension p and the sample size n may be large, without imposing explicit constraints on their relative growth. Further extensions and related developments can be found in Fisher (2012), Srivastava et al. (2014) and the optimality-oriented work of Cai and Ma (2013), among others. A third line focuses on maximum-type tests, which are particularly effective against sparse alternatives. Xiao and Wu (2013) proposed a test based on the maximum standardized entrywise deviation of the sample covariance matrix from the identity and established its limiting extreme-value distribution. Such procedures complement Frobenius-type tests, but may be less sensitive to dense or moderately dense alternatives. Further related developments include Wu and Li (2020) and Chen et al. (2023).

These developments have substantially advanced our understanding of high-dimensional covariance testing in terms of size control, detection boundaries and minimax optimality. Nevertheless, there remains a noticeable gap between the high-dimensional literature and the classical Neyman–Pearson theory. In particular, Frobenius-norm-based tests, for example those proposed by Ledoit and Wolf (2002), Chen et al. (2010) and Srivastava et al. (2014), among many others, are typically evaluated in terms of detection rates and asymptotic power under specific classes of alternatives, and have not, in general, been developed within an explicitly defined class of competing tests or equipped with a systematic theory of most powerful tests analogous to that in low-dimensional statistics. To the best of our knowledge, a rigorous theory characterizing naive tests that are most powerful within structured families in high-dimensional covariance identity testing has so far received little attention in the statistics literature.

The present paper aims to bridge this gap by viewing high-dimensional covariance identity testing through the lens of the nonparametric naive test framework of Hu and Bai (2016) and by developing an operator-indexed family of naive tests together with an accompanying local optimality theory. The main contributions of this paper are threefold. First, working within this framework, we introduce a positive-definite operator $\mathbf{W}$ acting on the covariance structure and use it to index a distribution-free optional naive test family for testing covariance identity. For this operator-indexed family we establish high-dimensional central limit theorems for the associated test statistics under both the null and alternative hypotheses, valid for arbitrary combinations of dimension p and sample size n. In this formulation, classical U-statistic-based identity tests, such as that of Chen et al. (2010), arise as isotropic choices of $\mathbf{W}$, so that standard Frobenius-norm-type procedures are embedded as special cases within a single unified framework. Second, within this operator-indexed family we propose a spectral optimization principle for selecting the indexing operator $\mathbf{W}$ by maximizing the local asymptotic power at a fixed alternative covariance matrix $\boldsymbol{\Sigma}_{1p}$. At the oracle level, this optimization yields an explicit operator $\mathbf{W}_{\boldsymbol{\Sigma}_{1p},\text{opt}}$ whose spectral structure can be expressed in closed form in terms of the eigen-decomposition of $(\boldsymbol{\Sigma}_{1p} - \mathbf{I}_p)^2$, thereby aligning both the eigenvectors and eigenvalues of $\mathbf{W}$ with the principal directions of the covariance departure. Third, building on these distributional and optimization results, we formally introduce the notion of a locally most powerful naive test (LMPNT) for high-dimensional covariance identity testing and show that the test based on $\mathbf{W}_{\boldsymbol{\Sigma}_{1p},\text{opt}}$ is LMPNT within the optional naive test family under a suitably structured class of local alternatives. Extensive Monte Carlo experiments under Gaussian and non-Gaussian designs show that suitable choices of $\mathbf{W}$ maintain accurate size control and provide substantial power gains over existing competing procedures. Related evidence from high-dimensional inference also suggests that explicitly modelling covariance or precision structure can improve test sensitivity; see, for example, Yang et al. (2024) in two-sample mean testing via structured precision estimation.

We focus on local alternatives, as violations of the local condition typically lead to covariance matrices $\mathbf{\Sigma}$ that deviate substantially from the identity matrix $\mathbf{I}_p$, rendering the two hypotheses easily distinguishable. In such asymptotic settings, the power of any reasonable test approaches one as the sample size n increases, making it difficult to meaningfully compare the relative performance of different procedures. From a practical standpoint, it is more natural to consider scenarios in which the power converges to a finite value β with $\alpha < \beta < 1$. In this context, the framework developed in this paper is expected to play a significant role.

To make the following discussion precise, we detail some notation. The notation $\|\cdot\|$ denotes the spectral norm of a matrix or the Euclidean norm of a vector. The symbol $\xrightarrow{\text{a.s.}}$ denotes almost sure convergence. Hereafter, we use $\xrightarrow{p}$ and $\xrightarrow{d}$ for convergence in probability and in distribution, respectively. Given any matrix $\mathbf{A} = (a_{ij})_{m\times n}$, we denote its trace by $\mathrm{tr}(\mathbf{A})$ when $\mathbf{A}$ is square, its transpose by $\mathbf{A}'$, its Frobenius norm by $\|\mathbf{A}\|_F = (\sum_{i=1}^m \sum_{j=1}^n |a_{ij}|^2)^{1/2} = \sqrt{\mathrm{tr}\,(\mathbf{A}'\mathbf{A})}$, and its largest eigenvalue by $\lambda_{\max}(\mathbf{A})$ when $\mathbf{A}$ is symmetric. Lastly, we write $A_{n,r} = \frac{n!}{(n-r)!}$ and use $\sum^*$ to denote summation over mutually distinct indices. For example, $\sum^*_{i,j,k}$ means summation over $\{(i,j,k) : i \neq j,\ j \neq k,\ k \neq i\}$.

The remainder of the paper is organized as follows. Section 2 introduces the optional naive test family indexed by a positive-definite operator $\mathbf{W}$ and derives the limiting distributions of the corresponding test statistics under both the null and alternative hypotheses. Section 3 develops the spectral optimization criterion and identifies the oracle operator $\mathbf{W}_{\mathbf{\Sigma}_{1p},\mathrm{opt}}$. Section 4 establishes the existence of a locally most powerful naive test under suitable structural constraints. Section 5 presents simulation studies that support the optimization theory in Sections 3 and 4, and Section 6 reports a real-data analysis comparing the proposed test with several existing methods. Section 7 provides a discussion. Detailed proofs, additional simulation results, and further details for the real-data analysis are provided in the Supplementary Material.

2. The optional naive test family and its asymptotic distribution

For the hypothesis testing problem described in (1.1), we construct a family of test statistics indexed by a $p \times p$ positive-definite operator $\mathbf{W}$. We consider a class of naive tests discussed in Hu and Bai (2016) that utilizes the squared Frobenius norm. The test is established based on the central limit theorem (CLT), where $\alpha_n h(\hat{\theta}_n) \to N(0, \sigma^2)$, and $h(\hat{\theta}_n)$ is a consistent estimator of $h(\theta_n)$. Here, $h(\theta_n)$ satisfies $h(\theta_n) = 0$ under the null hypothesis and $h(\theta_n) > 0$ under the alternative hypothesis. Note that

$$\frac{1}{p}\mathrm{tr}[(\mathbf{\Sigma} - \mathbf{I}_p)\mathbf{W}(\mathbf{\Sigma} - \mathbf{I}_p)] = \frac{1}{p}\mathrm{tr}(\mathbf{W}\mathbf{\Sigma}^2) - \frac{2}{p}\mathrm{tr}(\mathbf{W}\mathbf{\Sigma}) + \frac{\mathrm{tr}(\mathbf{W})}{p}, \tag{2.2}$$

When $\mathbf{W} = \mathbf{I}_p$, the quantity in (2.2) reduces to the squared Frobenius norm $p^{-1}\|\mathbf{\Sigma} - \mathbf{I}_p\|_F^2$, and the resulting naive test coincides with the U-statistic-based procedure of Chen et al. (2010). Allowing for general positive-definite $\mathbf{W}$ therefore embeds the classical Frobenius-type identity tests into a broader operator-indexed family, within which we will later optimize both the eigenvectors and the eigenvalues of $\mathbf{W}$.

Following Chen et al. (2010), we construct unbiased estimators of the terms in (2.2) via U-statistics. Define

$$U_n^1 = \frac{1}{A_{n,1}} \sum_{i=1}^{n} \mathbf{x}_i' \mathbf{W} \mathbf{x}_i, \quad U_n^2 = \frac{1}{A_{n,2}} \sum_{i \neq j} \mathbf{x}_i' \mathbf{x}_j \, \mathbf{x}_i' \mathbf{W} \mathbf{x}_j, \quad U_n^3 = \frac{1}{A_{n,2}} \sum_{i \neq j} \mathbf{x}_i' \mathbf{W} \mathbf{x}_j,$$

$$U_n^4 = \frac{1}{A_{n,3}} \sum_{i,j,k}^{*} \mathbf{x}_i' \mathbf{x}_j \, \mathbf{x}_j' \mathbf{W} \mathbf{x}_k, \quad U_n^5 = \frac{1}{A_{n,4}} \sum_{i,j,k,l}^{*} \mathbf{x}_i' \mathbf{x}_j \, \mathbf{x}_k' \mathbf{W} \mathbf{x}_l,$$

Remark 1. For efficient numerical implementation, U_n^4 and U_n^5 can be computed without directly enumerating all triples and quadruples. Since $T_n(\mathbf{W})$ is translation invariant, we work with the centered data matrix $\mathbf{X}_c = (\mathbf{x}_1 - \bar{\mathbf{x}}, \ldots, \mathbf{x}_n - \bar{\mathbf{x}})$, where $\bar{\mathbf{x}} = n^{-1} \sum_{i=1}^{n} \mathbf{x}_i$. Let $\mathbf{G} = \mathbf{X}_c^\top \mathbf{X}_c$ and $\mathbf{H} = \mathbf{X}_c^\top \mathbf{W} \mathbf{X}_c$. A straightforward calculation gives $A_{n,3} U_n^4 = -\mathrm{tr}(\mathbf{HG}) + 2\mathrm{tr}(\mathbf{H} \circ \mathbf{G})$ and $A_{n,4} U_n^5 = \mathrm{tr}(\mathbf{H})\mathrm{tr}(\mathbf{G}) + 2\mathrm{tr}(\mathbf{HG}) - 6\mathrm{tr}(\mathbf{H} \circ \mathbf{G})$. Hence, after forming $\mathbf{G}$ and $\mathbf{H}$, the additional cost of evaluating U_n^4 and U_n^5 is only $O(n^2)$. For a general $\mathbf{W}$, the overall computational cost is $O(np^2 + n^2 p)$, which reduces to $O(n^2 p)$ when $\mathbf{W} = \mathbf{I}_p$.

Then unbiased estimators of $\mathrm{tr}(\mathbf{W\Sigma})$ and $\mathrm{tr}(\mathbf{W\Sigma}^2)$ are given respectively by

$$U_n(1) = U_n^1 - U_n^3 \quad \text{and} \quad U_n(2) = U_n^2 - 2U_n^4 + U_n^5.$$

Consequently, we define the family of test statistics

$$T_n(\mathbf{W}) = \frac{1}{p} U_n(2) - \frac{2}{p} U_n(1) + \frac{\mathrm{tr}(\mathbf{W})}{p}, \tag{2.3}$$

and refer to the collection $\{T_n(\mathbf{W}) : \mathbf{W} \succ 0\}$ as the operator-indexed family of test statistics, where $\mathbf{W}$ ranges over positive-definite operators. The corresponding optional naive test family is generated by these statistics.

Below, we list the main assumptions:

Assumption 1. The observations $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n$ admit the following model:

$$\mathbf{x}_i = \boldsymbol{\mu} + \boldsymbol{\Gamma}\mathbf{z}_i, \quad i = 1, 2, \ldots, n,$$

where $\boldsymbol{\mu} \in \mathbb{R}^p$ denotes the mean vector, $\boldsymbol{\Gamma} \in \mathbb{R}^{p\times m}$ is a matrix with $m \geq p$ such that $\boldsymbol{\Gamma}\boldsymbol{\Gamma}' = \boldsymbol{\Sigma}$, and $\{\mathbf{z}_i = (z_{i1}, \ldots, z_{im})'\}_{i=1}^n$ are i.i.d. m-dimensional random vectors with $\mathrm{E}(\mathbf{z}_1) = \mathbf{0}$ and $\mathrm{Var}(\mathbf{z}_1) = \mathbf{I}_m$. Moreover, we assume that $\sup_\ell \mathrm{E}(z_{1\ell}^8) < \infty$, $\mathrm{E}(z_{1\ell}^4) = 3 + \Delta$ for some constant Δ, and for any integers $\ell_v \geq 0$ with $\sum_{v=1}^q \ell_v = 8$,

$$\mathrm{E}\left(\mathbf{z}_{1i_1}^{\ell_1}\mathbf{z}_{1i_2}^{\ell_2}\cdots\mathbf{z}_{1i_q}^{\ell_q}\right) = \mathrm{E}\left(\mathbf{z}_{1i_1}^{\ell_1}\right)\mathrm{E}\left(\mathbf{z}_{1i_2}^{\ell_2}\right)\cdots\mathrm{E}\left(\mathbf{z}_{1i_q}^{\ell_q}\right),$$

whenever $i_1, i_2, \ldots, i_q$ are distinct indices.

Assumption 2. As $n \to \infty$, $p = p(n) \to \infty$, $\mathrm{tr}(\boldsymbol{\Sigma}^2) \to \infty$ and $\lambda_{\max}(\boldsymbol{\Sigma}^2) = o\{\mathrm{tr}(\boldsymbol{\Sigma}^2)\}$.

Remark 2. Assumption 1 incorporates a linear representation of $\mathbf{x}_i$ along with a pseudo-independence condition, which includes the multivariate Gaussian distribution among others, as discussed in Bai and Saranadasa (1996) and Chen et al. (2010). Assumption 2 is necessary for establishing the asymptotic normality of the test statistic. This condition can be viewed as a completely equivalent reformulation of the condition proposed by Chen et al. (2010), yet it offers a more intuitive expression. A similar assumption was also utilized in Wang et al. (2015). It is important to note that both conditions—Assumption 2 and the condition in Assumption 2 of Chen et al. (2010)—are relatively weak. They hold trivially when all the eigenvalues of $\boldsymbol{\Sigma}$ are bounded away from zero and infinity.

Theorem 1. *Under Assumptions 1–2, as $n, p \to \infty$,*

$$\sigma_{1,n}^{-1}\Big\{p\,T_n(\mathbf{W}) - \operatorname{tr}\big[\mathbf{W}(\boldsymbol{\Sigma}-\mathbf{I}_p)^2\big]\Big\} \xrightarrow{d} N(0,1),$$

where

$$\begin{aligned}\sigma_{1,n}^2 =& \frac{2}{n^2}\operatorname{tr}\left((\mathbf{W}\boldsymbol{\Sigma})^2\right)\operatorname{tr}(\boldsymbol{\Sigma}^2) + \frac{2}{n^2}\left\{\operatorname{tr}(\mathbf{W}\boldsymbol{\Sigma}^2)\right\}^2 \\ &+ \frac{4\Delta}{n}\operatorname{tr}\big[(\mathbf{A}\boldsymbol{\Gamma}'\mathbf{W}\boldsymbol{\Gamma} - \boldsymbol{\Gamma}'\mathbf{W}\boldsymbol{\Gamma})\circ(\mathbf{A}\boldsymbol{\Gamma}'\mathbf{W}\boldsymbol{\Gamma} - \boldsymbol{\Gamma}'\mathbf{W}\boldsymbol{\Gamma})\big] \\ &+ \frac{4}{n}\operatorname{tr}(\mathbf{W}\boldsymbol{\Sigma}\mathbf{W}\boldsymbol{\Sigma}^3 + \mathbf{W}\boldsymbol{\Sigma}^2\mathbf{W}\boldsymbol{\Sigma}^2 - 4\mathbf{W}\boldsymbol{\Sigma}\mathbf{W}\boldsymbol{\Sigma}^2 + 2\mathbf{W}\boldsymbol{\Sigma}\mathbf{W}\boldsymbol{\Sigma}),\end{aligned}$$

$\mathbf{A} = \boldsymbol{\Gamma}'\boldsymbol{\Gamma}$*, and $\circ$ denotes the Hadamard product. In particular, under the null hypothesis $\boldsymbol{\Sigma} = \mathbf{I}_p$,*

$$\sigma_{0,n}^{-1}\,p\,T_n(\mathbf{W}) \xrightarrow{d} N(0,1),$$

where $\sigma_{0,n}^2 = \frac{2p}{n^2}\operatorname{tr}(\mathbf{W}^2) + \frac{2}{n^2}\left\{\operatorname{tr}(\mathbf{W})\right\}^2$.

Theorem 1 thus provides a high-dimensional central limit theorem for the operator-indexed statistic $T_n(\mathbf{W})$. In the special case $\mathbf{W} = \mathbf{I}_p$, the asymptotic variances simplify to

$$\sigma_{1,n}^2 = \frac{4}{n^2}\left\{\mathrm{tr}(\mathbf{\Sigma}^2)\right\}^2 + \frac{8}{n}\mathrm{tr}\left((\mathbf{\Sigma}^2 - \mathbf{\Sigma})^2\right) + \frac{4\Delta}{n}\mathrm{tr}\big[(\mathbf{A}^2 - \mathbf{A}) \circ (\mathbf{A}^2 - \mathbf{A})\big],$$
$$\sigma_{0,n}^2 = \frac{4p^2}{n^2}.$$

These expressions coincide with those derived in Chen et al. (2010).

Remark 3 (Proof strategy for Theorem 1). We briefly outline the proof strategy for Theorem 1 here. First, after decomposing $pT_n(\mathbf{W})$ into its leading stochastic component and remainder terms arising from the U-statistic construction, we identify the leading variance term and show that the remainders are negligible after normalization. This step uses moment bounds for quadratic forms and trace inequalities under Assumptions 1–2. Second, the leading component is written as a sum of martingale differences with respect to the observations. After verifying the convergence of the conditional variance and the Lyapunov condition, the martingale central limit theorem yields the asymptotic normality of the normalized statistic. The detailed proof is given in the Supplementary Material.

The next step is to determine how the operator $\mathbf{W}$ should be chosen so that the test based on $T_n(\mathbf{W})$ can achieve improved local power under high-dimensional alternatives.

3. Design and optimization of test operators

3.1 Oracle construction of the test operator

This section addresses the central design problem of our framework: within the optional naive test family introduced in Section 2, how should a positive-definite operator $\mathbf{W}$ be chosen so that the resulting statistic $T_n(\mathbf{W})$ achieves high local power in the high-dimensional regime? The question is tackled by analyzing the asymptotic power function under local alternatives and, for each fixed alternative covariance matrix $\mathbf{\Sigma}_{1p}$, deriving an oracle operator $\mathbf{W}_{\mathbf{\Sigma}_{1p},\mathrm{opt}}$ that maximizes this power within the optional naive test family.

To formalize the local regime in which different procedures are compared, the following notion of local alternatives is introduced.

Definition 1 (Local alternatives)**.** Under the asymptotic regime $n, p \to \infty$, an alternative hypothesis H_1 is termed a *local alternative* if it satisfies

$$\sigma_{1,n}^2 = \sigma_{0,n}^2\big[1 + o(1)\big] \qquad \text{as } n, p \to \infty. \tag{3.4}$$

Remark 4. Definition 1 captures the situation where H_1 is asymptotically close to H_0. The analysis in this paper focuses on such local regimes. Without such restrictions, the covariance matrix $\mathbf{\Sigma}$ typically departs substantially from $\mathbf{I}_p$ as $n, p \to \infty$, and the powers of all reasonable tests converge to one, making meaningful comparison impossible. From a practical perspective, it is more informative to study scenarios in which the power converges to a non-degenerate limit β with $\alpha < \beta < 1$.

We now derive the oracle choice of $\mathbf{W}$ by maximizing the asymptotic power under a fixed alternative covariance matrix $\mathbf{\Sigma}_{1p}$. The construction consists of three steps: first, formulating the optimization objective via the power function; second, aligning the eigenvectors of $\mathbf{W}$ with those of $(\mathbf{\Sigma}_{1p} - \mathbf{I}_p)^2$; and third, determining the optimal eigenvalues of $\mathbf{W}$.

Step 1: Constructing the objective via the power function.

Let z_α denote the upper α-quantile of the standard normal distribution. According to Theorem 1 and Definition 1, the asymptotic power function of the test with significance level $\alpha \in (0, 1)$ is

$$\begin{aligned}
\beta(\mathbf{\Sigma}_{1p}) &= \Pr\left(\sigma_{0,n}^{-1} p T_n(\mathbf{W}) \geq z_\alpha \,\middle|\, \mathbf{\Sigma} = \mathbf{\Sigma}_{1p} \neq \mathbf{I}_p\right) \\
&= \Phi\left(-\frac{z_\alpha \sigma_{0,n}}{\sigma_{1,n}} + \frac{\operatorname{tr}\{\mathbf{W}(\mathbf{\Sigma}_{1p} - \mathbf{I}_p)^2\}}{\sigma_{1,n}}\right) + o(1) \\
&= \Phi\left(-\frac{z_\alpha}{1 + o(1)} + \frac{\operatorname{tr}\{\mathbf{W}(\mathbf{\Sigma}_{1p} - \mathbf{I}_p)^2\}}{\sigma_{0,n}[1 + o(1)]}\right) + o(1),
\end{aligned}$$

where Φ is the distribution function of the standard normal distribution $N(0, 1)$.

Motivated by this expression, the oracle optimal operator is defined as

$$\mathbf{W}_{\boldsymbol{\Sigma}_{1p},\mathrm{opt}} = \arg\max_{\mathbf{W}} \frac{\left\{\mathrm{tr}\big(\mathbf{W}(\boldsymbol{\Sigma}_{1p} - \mathbf{I}_p)^2\big)\right\}^2}{\sigma_{0,n}^2}. \tag{3.5}$$

which maximizes the asymptotic power under the given alternative $\boldsymbol{\Sigma}_{1p}$ over the optional naive test family indexed by $\mathbf{W}$. Since $\boldsymbol{\Sigma}_{1p}$ is unknown in practice, $\mathbf{W}_{\boldsymbol{\Sigma}_{1p},\mathrm{opt}}$ should be interpreted as a theoretical benchmark rather than a directly available operator. In applications, scientifically meaningful structural information about the alternative, together with data-based estimates of the unknown covariance quantities, can guide the construction of a feasible operator toward this oracle target.

Step 2: Optimizing the eigenvectors of W.

Let $\lambda_1 \geq \cdots \geq \lambda_p > 0$ and $\mathbf{u}_1, \ldots, \mathbf{u}_p$ denote the eigenvalues and associated orthonormal eigenvectors of $\mathbf{W}_{\boldsymbol{\Sigma}_{1p},\mathrm{opt}}$, respectively. Similarly, denote the eigenvalues of $(\boldsymbol{\Sigma}_{1p} - \mathbf{I}_p)^2$ by $q_1 \geq \cdots \geq q_p \geq 0$ and the corresponding eigenvectors by $\mathbf{v}_1, \ldots, \mathbf{v}_p$. Using the spectral decompositions of $\mathbf{W}$ and $(\boldsymbol{\Sigma}_{1p} - \mathbf{I}_p)^2$, the maximization problem in (3.5) can be rewritten as

$$\max_{\mathbf{W}} \frac{\left\{\mathrm{tr}\big(\mathbf{W}(\boldsymbol{\Sigma}_{1p} - \mathbf{I}_p)^2\big)\right\}^2}{\sigma_{0,n}^2} = \max_{\lambda_1 \geq \cdots \geq \lambda_p > 0,\, \mathbf{u}_1,\ldots,\mathbf{u}_p} \frac{\left[\sum_{i,j} \lambda_i q_j (\mathbf{u}_i' \mathbf{v}_j)^2\right]^2}{\frac{2p}{n^2}\sum_{i=1}^p \lambda_i^2 + \frac{2}{n^2}\left(\sum_{i=1}^p \lambda_i\right)^2}. \tag{3.6}$$

To optimize the eigenvectors $\mathbf{u}_1, \ldots, \mathbf{u}_p$, define

$$\boldsymbol{\Lambda}_\lambda = \mathrm{diag}(\lambda_1, \ldots, \lambda_p), \quad \boldsymbol{\Lambda}_q = \mathrm{diag}(q_1, \ldots, q_p), \quad \mathbf{D} = (d_{ij}) = (\mathbf{u}_i' \mathbf{v}_j).$$

By von Neumann's trace inequality and the orthogonality of $\mathbf{D}$,

$$\sum_{i,j} \lambda_i q_j (\mathbf{u}_i' \mathbf{v}_j)^2 = \mathrm{tr}(\boldsymbol{\Lambda}_\lambda \mathbf{D} \boldsymbol{\Lambda}_q \mathbf{D}') \leq \sum_{i=1}^{p} s_i(\boldsymbol{\Lambda}_\lambda \mathbf{D})\, s_i(\boldsymbol{\Lambda}_q \mathbf{D}') = \sum_{i=1}^{p} \lambda_i q_i,$$

with equality if and only if $\mathbf{u}_i = \mathbf{v}_i$ for $i = 1, \ldots, p$, where $s_i(\cdot)$ denotes the i-th singular value. Hence, (3.6) reduces to

$$\max_{\substack{\lambda_1 \geq \cdots \geq \lambda_p > 0, \\ \mathbf{u}_1, \ldots, \mathbf{u}_p}} \frac{\left[\sum_{i,j} \lambda_i q_j (\mathbf{u}_i' \mathbf{v}_j)^2\right]^2}{\frac{2p}{n^2} \sum_{i=1}^{p} \lambda_i^2 + \frac{2}{n^2} \left(\sum_{i=1}^{p} \lambda_i\right)^2} = \max_{\lambda_1 \geq \cdots \geq \lambda_p > 0} \frac{\left(\sum_{i=1}^{p} \lambda_i q_i\right)^2}{\frac{2p}{n^2} \sum_{i=1}^{p} \lambda_i^2 + \frac{2}{n^2} \left(\sum_{i=1}^{p} \lambda_i\right)^2}. \tag{3.7}$$

Thus the oracle operator $\mathbf{W}_{\boldsymbol{\Sigma}_{1p},\mathrm{opt}}$ shares its eigenvectors with $(\boldsymbol{\Sigma}_{1p} - \mathbf{I}_p)^2$. The optimization over directions is completely characterized by the alignment $\mathbf{u}_i = \mathbf{v}_i$, and the remaining task is to determine the optimal eigenvalues $\lambda_1, \ldots, \lambda_p$.

Step 3: Optimizing the eigenvalues of W.

Note that if $(\lambda_1, \ldots, \lambda_p)$ is replaced by $c(\lambda_1, \ldots, \lambda_p)$ for any positive constant c, the value of the objective function in (3.7) remains unchanged. The overall scale is therefore fixed by imposing the normalization $\sigma_{0,n}^2 = \frac{p^2}{n^2}$. Under this constraint, define

$$A = \sum_{i=1}^{p} \lambda_i^2, \quad B = \sum_{i=1}^{p} \lambda_i, \quad C = \sum_{i=1}^{p} \lambda_i q_i, \quad Q_1 = \sum_{i=1}^{p} q_i, \quad Q_2 = \sum_{i=1}^{p} q_i^2. \tag{3.8}$$

The optimization of (3.7) is then considered in two cases.

Case 1: $pQ_2 - Q_1^2 = 0$.

This condition implies that $q_1 = \cdots = q_p$. By Jensen's inequality,

$$\frac{\left(\sum_{i=1}^p \lambda_i q_i\right)^2}{\frac{2p}{n^2}\sum_{i=1}^p \lambda_i^2 + \frac{2}{n^2}\left(\sum_{i=1}^p \lambda_i\right)^2} \le \frac{q_1^2\left(\sum_{i=1}^p \lambda_i\right)^2}{\frac{4}{n^2}\left(\sum_{i=1}^p \lambda_i\right)^2} = \frac{1}{4}n^2 q_1^2,$$

and equality holds if and only if $\lambda_1 = \cdots = \lambda_p = \frac{1}{2}$.

Thus, when $(\mathbf{\Sigma}_{1p} - \mathbf{I}_p)^2$ is spherical, the oracle operator is a scalar multiple of the identity, that is, $\mathbf{W}_{\mathbf{\Sigma}_{1p},\mathrm{opt}} \propto \mathbf{I}_p$. Under the normalization $\sigma_{0,n}^2 = p^2/n^2$, this corresponds to $\lambda_1 = \cdots = \lambda_p = \frac{1}{2}$. If one instead chooses the scale so that $\mathbf{W}_{\mathbf{\Sigma}_{1p},\mathrm{opt}} = \mathbf{I}_p$, then $\sigma_{0,n}^2 = 4p^2/n^2$, which coincides with the null variance derived for the U-statistic in Chen et al. (2010). In this sense, the statistic of Chen et al. (2010) arises as the isotropic special case $\mathbf{W} = \mathbf{I}_p$ within the optional naive test family indexed by the positive-definite operator $\mathbf{W}$.

Case 2: $pQ_2 - Q_1^2 \neq 0$.

We now use the method of Lagrange multipliers to determine the optimal spectrum. Let

$$f(\lambda_1, \ldots, \lambda_p, \mu) = \left(\sum_{i=1}^p \lambda_i q_i\right)^2 + \mu\left[\frac{2}{p}\sum_{i=1}^p \lambda_i^2 + \frac{2}{p^2}\left(\sum_{i=1}^p \lambda_i\right)^2 - 1\right],$$

where μ denotes the Lagrange multiplier. To find the critical points, we take the partial derivatives of f with respect to λ_j and μ and set them to zero:

$$\frac{\partial f}{\partial \lambda_j} = 2\Big(\sum_{i=1}^{p} \lambda_i q_i\Big) q_j + \mu \left[\frac{4}{p}\lambda_j + \frac{4}{p^2}\sum_{i=1}^{p}\lambda_i\right] = 0, \quad j = 1, \ldots, p,$$
$$\frac{\partial f}{\partial \mu} = \frac{2}{p}\sum_{i=1}^{p}\lambda_i^2 + \frac{2}{p^2}\Big(\sum_{i=1}^{p}\lambda_i\Big)^2 - 1 = 0.$$

Using the definitions of A, B, and C in (3.8), this system of equations yields

$$\lambda_j = \frac{-2\mu B - p^2 C q_j}{2p\mu}, \quad j = 1, \ldots, p. \tag{3.9}$$

Combining this with (3.8) gives

$$A = \frac{-2\mu B^2 - p^2 C^2}{2p\mu}, \quad \mu = -\frac{pCQ_1}{4B}, \quad C = \frac{B(2pQ_2 - Q_1^2)}{pQ_1}, \tag{3.10}$$

together with the normalization $\sigma_{0,n}^2 = \frac{2p}{n^2}A + \frac{2}{n^2}B^2 = \frac{p^2}{n^2}$. It then follows that

$$8pQ_2B^2 - 4Q_1^2B^2 - p^2Q_1^2 = 0.$$

Since $B > 0$, we obtain

$$B = \frac{pQ_1}{2\sqrt{2pQ_2 - Q_1^2}}.$$

Substituting this result into (3.9) and (3.10), we obtain the eigenvalues that maximize (3.7), denoted by $\lambda_1^* \geq \cdots \geq \lambda_p^*$:

$$\lambda_j^* = \frac{B}{p}\left(\frac{2p}{Q_1}q_j - 1\right), \quad j = 1, \ldots, p. \tag{3.11}$$

By the Cauchy–Schwarz inequality, one verifies that this extremum is indeed the global maximum of (3.7).

To ensure that each eigenvalue λ_j^* in (3.11) is strictly positive, we introduce the following assumption.

Assumption 3. The minimum eigenvalue q_p of $(\mathbf{\Sigma}_{1p} - \mathbf{I}_p)^2$ satisfies $q_p > \frac{Q_1}{2p}$, where $Q_1 = \sum_{i=1}^p q_i$.

Remark 5. Assumption 3 prevents the smallest eigenvalue of $(\mathbf{\Sigma}_{1p} - \mathbf{I}_p)^2$ from being too small relative to the empirical average Q_1/p. In particular, Assumption 3 guarantees that the explicit solution in (3.11) satisfies $\lambda_j^* > 0$ for all $j = 1, \ldots, p$, so that the resulting operator $\mathbf{W}_{\mathbf{\Sigma}_{1p},\mathrm{opt}}$ is positive-definite. This condition is naturally suited to relatively dense alternatives and provides a convenient setting for establishing the optimality properties of the proposed optional naive test based on $\mathbf{W}_{\mathbf{\Sigma}_{1p},\mathrm{opt}}$.

Finally, expanding and simplifying (3.9) yields the more transparent representation

$$\lambda_j^* = \frac{p}{\sqrt{2pQ_2 - Q_1^2}} q_j - \frac{Q_1}{2\sqrt{2pQ_2 - Q_1^2}}, \quad j = 1, \ldots, p. \tag{3.12}$$

Equation (3.12) reveals an explicit linear relationship between the eigenvalues λ_j^* of $\mathbf{W}_{\mathbf{\Sigma}_{1p},\mathrm{opt}}$ and the eigenvalues q_j of $(\mathbf{\Sigma}_{1p} - \mathbf{I}_p)^2$. The linear coefficient

$$\frac{p}{\sqrt{2pQ_2 - Q_1^2}}$$

is strictly positive (since $pQ_2 - Q_1^2 > 0$ in Case 2) and depends only on the departure spectrum $\{q_j\}_{j=1}^p$, reflecting a clear spectral correspondence between the test operator and the covariance structure. Consequently, $\mathbf{W}_{\mathbf{\Sigma}_{1p},\mathrm{opt}}$ adapts to the eigen-structure of $(\mathbf{\Sigma}_{1p} - \mathbf{I}_p)^2$: the eigenvalue λ_j^* increases linearly with the departure level q_j. In particular, the mapping $q_j \mapsto \lambda_j^*$ is strictly increasing, so the eigenvalue ordering of $(\mathbf{\Sigma}_{1p} - \mathbf{I}_p)^2$ is preserved by $\mathbf{W}_{\mathbf{\Sigma}_{1p},\mathrm{opt}}$. This construction jointly optimizes the eigenvectors and eigenvalues of $\mathbf{W}$, accentuating directions along which $\mathbf{\Sigma}_{1p}$ deviates more strongly from $\mathbf{I}_p$ and thereby enhancing the discriminatory power of the resulting test.

3.2 Summary of the oracle operator construction

Based on the above analysis, we summarize the construction of the oracle test operator $\mathbf{W}_{\mathbf{\Sigma}_{1p},\mathrm{opt}}$ in Algorithm 1.

Theorem 2. *Under Assumptions 1, 2 and 3, consider the testing problem* (1.1) *and a fixed alternative covariance matrix $\mathbf{\Sigma}_{1p} \in \mathbb{R}^{p\times p}$. Then the operator $\mathbf{W}_{\mathbf{\Sigma}_{1p},\mathrm{opt}}$ constructed in Algorithm 1 maximizes the asymptotic power, as characterized in Theorem 1, over the optional naive test family indexed*

Algorithm 1: Oracle construction of the test operator $\mathbf{W}_{\boldsymbol{\Sigma}_{1p},\mathrm{opt}}$

Input: Alternative covariance matrix $\boldsymbol{\Sigma}_{1p} \in \mathbb{R}^{p\times p}$

Output: Oracle test operator $\mathbf{W}_{\boldsymbol{\Sigma}_{1p},\mathrm{opt}} \in \mathbb{R}^{p\times p}$

Perform spectral decomposition $(\boldsymbol{\Sigma}_{1p} - \mathbf{I}_p)^2 = \sum_{i=1}^{p} q_i \mathbf{v}_i \mathbf{v}_i'$;

Compute $Q_1 = \sum_{j=1}^{p} q_j$ and $Q_2 = \sum_{j=1}^{p} q_j^2$;

if $q_1 = q_p$ *(spherical case)* **then**

  Set $\lambda_1^* = \cdots = \lambda_p^* = \frac{1}{2}$ under the normalization $\sigma_{0,n}^2 = p^2/n^2$;

else

  Set $\lambda_j^* = \dfrac{p}{\sqrt{2pQ_2 - Q_1^2}} q_j - \dfrac{Q_1}{2\sqrt{2pQ_2 - Q_1^2}}, \quad j = 1, \ldots, p;$

Construct the oracle operator $\mathbf{W}_{\boldsymbol{\Sigma}_{1p},\mathrm{opt}} = \sum_{i=1}^{p} \lambda_i^* \mathbf{v}_i \mathbf{v}_i'$;

by $\mathbf{W}$.

Remark 6. When Assumption 3 is not satisfied, the explicit solution in (3.12) may yield some nonpositive eigenvalues λ_j^*, typically corresponding to directions with extremely weak signals. In practice, we retain the positive eigenvalues and replace only the nonpositive ones by a small positive constant ε, that is, we set $\tilde{\lambda}_j = \lambda_j^*$ if $\lambda_j^* > 0$, and $\tilde{\lambda}_j = \varepsilon$ otherwise, for $j = 1, \ldots, p$. The resulting regularized operator is defined as

$$\mathbf{W}_{\boldsymbol{\Sigma}_{1p},\mathrm{trunc}} = \sum_{j=1}^{p} \tilde{\lambda}_j \mathbf{v}_j \mathbf{v}_j'.$$

In our implementation, we take $\varepsilon = 10^{-10}$. We also checked $\varepsilon = 10^{-5}$, which gave essentially unchanged empirical powers. This regularization preserves the oracle eigendirections and the ordering induced by the departure spectrum $\{q_j\}$, while ensuring that the implemented operator remains positive-definite. Since Theorem 1 is stated for a general positive-definite operator $\mathbf{W}$, the corresponding asymptotic null distribution and the associated local power expression continue to apply to the statistic constructed from $\mathbf{W}_{\mathbf{\Sigma}_{1p},\mathrm{trunc}}$. By contrast, the oracle optimality result in Theorem 2 and the LMPNT conclusion in Section 4 are stated for $\mathbf{W}_{\mathbf{\Sigma}_{1p},\mathrm{opt}}$ under Assumption 3, and are not claimed for the regularized operator when Assumption 3 fails. Nevertheless, $\mathbf{W}_{\mathbf{\Sigma}_{1p},\mathrm{trunc}}$ remains a practically useful choice. The additional simulation results in the Supplementary Material indicate that this regularized operator remains effective and still provides power gains even when Assumption 3 fails.

4. The locally most powerful naive test

In this section, we show how the pointwise oracle optimality result obtained in Section 3 can be extended to a classwise local optimality statement. The key step is to identify a structured alternative space on which the oracle operator remains invariant.

Let $\phi_{\text{opt},\boldsymbol{\Sigma}_{1p}}(x)$ denote the test function corresponding to the optional naive test statistic constructed from the oracle operator $\mathbf{W}_{\boldsymbol{\Sigma}_{1p},\text{opt}}$ in Section 3. We begin by specifying the structured class of alternatives under which the local optimality result will be established.

Definition 2 (Structured alternative space Θ_1)**.** Let $\boldsymbol{\Sigma}_{1p}$ be a fixed alternative covariance matrix for the testing problem (1.1), and write $(\boldsymbol{\Sigma}_{1p} - \mathbf{I}_p)^2$ for its squared departure matrix. The structured alternative space Θ_1 is defined as

$$\Theta_1 = \left\{ \boldsymbol{\Sigma} \;\middle|\; (\boldsymbol{\Sigma} - \mathbf{I}_p)^2 = k\,(\boldsymbol{\Sigma}_{1p} - \mathbf{I}_p)^2,\ k > 0 \right\}, \tag{4.13}$$

where k is a positive scalar that controls the overall magnitude of the departure from the identity.

The class Θ_1 thus forms a *local ray* of alternatives: for $\boldsymbol{\Sigma} \in \Theta_1$, the quadratic departure $(\boldsymbol{\Sigma} - \mathbf{I}_p)^2$ has the same eigen-directions as $(\boldsymbol{\Sigma}_{1p} - \mathbf{I}_p)^2$, and its eigenvalues are uniformly scaled by k.

As shown in Section 3, the oracle operator $\mathbf{W}_{\boldsymbol{\Sigma}_{1p},\text{opt}}$ is spectrally aligned with $(\boldsymbol{\Sigma}_{1p} - \mathbf{I}_p)^2$, and its eigenvalues $\{\lambda_j^*\}_{j=1}^p$ depend linearly on the departure eigenvalues $\{q_j\}_{j=1}^p$ via (3.12). Substituting the relation $(\boldsymbol{\Sigma} - \mathbf{I}_p)^2 = k(\boldsymbol{\Sigma}_{1p} - \mathbf{I}_p)^2$ for $\boldsymbol{\Sigma} \in \Theta_1$ into (3.12) therefore shows that the optimizer, and hence the oracle operator, remains unchanged along Θ_1:

$$\mathbf{W}_{\boldsymbol{\Sigma},\mathrm{opt}} = \mathbf{W}_{\boldsymbol{\Sigma}_{1p},\mathrm{opt}} \qquad \forall\, \boldsymbol{\Sigma} \in \Theta_1.$$

Thus the spectral structure of the oracle operator is invariant on Θ_1, while the scalar k indexes the strength of the departure from H_0. Within this class, $\mathrm{tr}\{\mathbf{W}_{\boldsymbol{\Sigma}_{1p},\mathrm{opt}}(\boldsymbol{\Sigma}-\mathbf{I}_p)^2\}$ is proportional to k, so the asymptotic power is an increasing function of k. In particular, $k = 0$ corresponds to $(\boldsymbol{\Sigma}-\mathbf{I}_p)^2 = \mathbf{0}$, that is $H_0 : \boldsymbol{\Sigma} = \mathbf{I}_p$.

To formulate the local optimality result, consider the composite testing problem

$$H_0 : \boldsymbol{\Sigma} = \mathbf{I}_p \quad \text{vs.} \quad H_1 : \boldsymbol{\Sigma} \in \Theta_1. \tag{4.14}$$

For a test function $\phi(\mathbf{x})$, write $\beta_\phi(\boldsymbol{\Sigma})$ for its power function. Throughout this section, the term "naive test" refers to a test based on $T_n(\mathbf{W})$ for some positive-definite operator $\mathbf{W}$, and hence belonging to the operator-indexed optional naive test family introduced in Section 2. The following definition formalizes the notion of a locally most powerful naive test.

Definition 3 (Locally most powerful naive test (LMPNT))**.** For the structured alternative space Θ_1 in Definition 2 and the testing problem (4.14), a naive test $\phi^*(x)$ is called a locally most powerful naive test (LMPNT) of asymptotic level α if:

---

1. its asymptotic level is α, that is,

$$\lim_{n,p\to\infty} \beta_{\phi^*}(\mathbf{I}_p) = \alpha;$$

2. for any other naive test ϕ with the same asymptotic level α,

$$\beta_{\phi^*}(\mathbf{\Sigma}) \geq \beta_{\phi}(\mathbf{\Sigma}), \quad \forall \mathbf{\Sigma} \in \Theta_1, \text{ for all sufficiently large } n, p.$$

By Theorem 1, the test $\phi_{\text{opt},\mathbf{\Sigma}_{1p}}(x)$ has asymptotic level α. Moreover, by Theorem 2, for each fixed $\mathbf{\Sigma} \neq \mathbf{I}_p$, the oracle operator $\mathbf{W}_{\mathbf{\Sigma},\text{opt}}$ maximizes the asymptotic power among all optional naive tests indexed by $\mathbf{W}$, so that

$$\beta_{\phi_{\text{opt},\mathbf{\Sigma}}}(\mathbf{\Sigma}) \ \geq \ \beta_{\phi}(\mathbf{\Sigma}), \quad \text{for all sufficiently large } n, p, \tag{4.15}$$

for any naive test ϕ with asymptotic level α.

For $\mathbf{\Sigma} \in \Theta_1$, the structural definition (4.13) and Algorithm 1 imply that the oracle operators coincide:

$$\mathbf{W}_{\mathbf{\Sigma},\text{opt}} = \mathbf{W}_{\mathbf{\Sigma}_{1p},\text{opt}}, \qquad \forall\, \mathbf{\Sigma} \in \Theta_1.$$

Consequently, the associated oracle tests also coincide, i.e. $\phi_{\text{opt},\mathbf{\Sigma}} \equiv \phi_{\text{opt},\mathbf{\Sigma}_{1p}}$ on Θ_1, and hence

$$\beta_{\phi_{\text{opt},\mathbf{\Sigma}_{1p}}}(\mathbf{\Sigma}) = \beta_{\phi_{\text{opt},\mathbf{\Sigma}}}(\mathbf{\Sigma}) \ \geq \ \beta_{\phi}(\mathbf{\Sigma}), \quad \forall \mathbf{\Sigma} \in \Theta_1, \tag{4.16}$$

for all sufficiently large n, p and any naive test ϕ with asymptotic level α.

5. Simulation studies

5.1 Oracle implementation

This subsection investigates the finite-sample performance of the proposed optional naive test indexed by the oracle operator $\mathbf{W}_{\mathbf{\Sigma}_{1p},\mathrm{opt}}$, referred to as T_{LMPNT}. The procedure is compared with several existing methods for testing the identity covariance matrix, including T_{CZZ} from Chen et al. (2010), T_{Fishera} and T_{Fisherb} from Fisher (2012), the corrected likelihood ratio test T_{CLRT} from Zheng et al. (2015), the maximum-type test T_{XW} from Xiao and Wu (2013), and the test T_{DW} based on the CLT in (IV.2) of Ding and Wang (2025), with the test function taken to be $h(x) = x^2$ in their LSS framework. Empirical sizes and powers are reported under a variety of high-dimensional regimes and structured departures from the identity. The aims are threefold: (i) to assess whether the proposed test attains the nominal level under the null; (ii) to illustrate its locally most powerful behavior under local alternatives of order $O(1/p)$; and (iii) to examine its robustness when Assumption 3 is mildly violated and when the signal strength is of fixed, nonlocal order.

The sample vectors $\mathbf{x}_1, \ldots, \mathbf{x}_n$ are generated from the independent-component model

$$\mathbf{x}_i = \mathbf{\Sigma}^{1/2}\mathbf{z}_i, \qquad i = 1, \ldots, n,$$

where $\mathbf{z}_i = (z_{i1}, \ldots, z_{ip})'$ and the latent coordinates $\{z_{ij}\}$ are i.i.d. with

$$z_{ij} \sim N(0,1) \quad \text{or} \quad z_{ij} \sim \Gamma(4,2) - 2.$$

For numerical implementation, we use centered observations throughout the simulation study. The alternatives are generated as eigen-structured perturbations of $\mathbf{I}_p$. For each scenario, let $\{(\widetilde{m}_i, \widetilde{q}_i)\}_{i=1}^p$ denote the initially generated pairs, where $\widetilde{m}_i = 1 \pm \sqrt{\widetilde{q}_i}$, with the choice of sign specified by the corresponding scenario. These pairs are rearranged jointly in decreasing order of $\widetilde{m}_i$, and the reordered pairs are denoted by $\{(m_i, q_i)\}_{i=1}^p$, so that $q_i = (m_i - 1)^2$. We then construct

$$\mathbf{\Sigma} = \sum_{i=1}^p m_i \mathbf{v}_i \mathbf{v}_i', \qquad (\mathbf{\Sigma} - \mathbf{I}_p)^2 = \sum_{i=1}^p q_i \mathbf{v}_i \mathbf{v}_i',$$

where $\mathbf{v}_1, \ldots, \mathbf{v}_p$ are the eigenvectors of the AR(1)-type correlation matrix $\mathbf{\Sigma}_{\mathrm{AR}} = (0.1^{|i-j|})_{p\times p}$, ordered according to the decreasing eigenvalues of $\mathbf{\Sigma}_{\mathrm{AR}}$. The following four scenarios specify the initially generated pairs $\{(\widetilde{m}_i, \widetilde{q}_i)\}_{i=1}^p$.

Scenario 1 (Balanced local perturbation). For $i = 1, \ldots, \lceil p/2 \rceil$, set $\tilde{q}_i = 1.6/p$ and $\tilde{m}_i = 1 - \sqrt{\tilde{q}_i}$; for $i = \lceil p/2 \rceil + 1, \ldots, p$, set $\tilde{q}_i = 4.4/p$ and $\tilde{m}_i = 1 + \sqrt{\tilde{q}_i}$.

Scenario 2 (Unbalanced local perturbation). For $i = 1, \ldots, \lceil p/2 \rceil$, set $\tilde{q}_i = 1/p$ and $\tilde{m}_i = 1 - \sqrt{\tilde{q}_i}$; for $i = \lceil p/2 \rceil + 1, \ldots, p$, set $\tilde{q}_i = 6/p$ and $\tilde{m}_i = 1 + \sqrt{\tilde{q}_i}$. This local scenario violates Assumption 3.

Scenario 3 (Multi-bulk local perturbation with multiple spikes). Let $b_p = 0.5 + 10p^{-1/2}$. The values of $\tilde{q}_i$ are

$$\frac{22}{p}, \quad \frac{14}{p}, \quad \frac{b_p + 0.20}{p}, \quad \frac{b_p + 0.10}{p}, \quad \frac{b_p}{p}, \quad \frac{8}{p},$$

with respective multiplicities

$$1, \quad 2, \quad p - \lfloor 0.7p \rfloor - 1, \quad \lfloor 0.7p \rfloor - \lfloor 0.3p \rfloor - 2, \quad \lfloor 0.3p \rfloor - 3, \quad 3.$$

For the first four groups, set $\tilde{m}_i = 1 + \sqrt{\tilde{q}_i}$, whereas for the last two groups, set $\tilde{m}_i = 1 - \sqrt{\tilde{q}_i}$. This scenario is motivated by the multi-bulk spectral structures discussed in Ding (2021).

Scenario 4 (Fixed-order perturbation). For $i = 1, 2$, set $\tilde{q}_i = 0.2$ and $\tilde{m}_i = 1 - \sqrt{\tilde{q}_i}$; for $i = 3, \ldots, p$, set $\tilde{q}_i = 0.015$ and $\tilde{m}_i = 1 + \sqrt{\tilde{q}_i}$.

Each scenario induces a population covariance matrix $\mathbf{\Sigma}$ and hence an oracle operator $\mathbf{W}_{\mathbf{\Sigma}_{1p},\mathrm{opt}}$. Scenarios 1–3 are local alternatives with $q_i = O(1/p)$. Scenarios 1 and 3 satisfy Assumption 3, whereas Scenario 2 provides a robustness check when this condition fails. Scenario 4 is a fixed-order, nonlocal alternative under which the powers of reasonable procedures are expected to approach one as n increases.

To explore a range of low- to high-dimensional regimes, we consider $p = 0.5n$, $p = n$, and $p = 1.5n$ with $n \in \{50, 100, 150, 200, 250, 300\}$, so that, under $O(1/p)$ local perturbations in Scenarios 1–3, the empirical powers do not collapse to either 0.05 or 1 over the grid of sample sizes considered.

The significance level is set at $\alpha = 0.05$, and all statistics are averaged over 5,000 independent replications. The empirical sizes under Scenarios 1–4 are reported in Tables S.1–S.8 in the Supplementary Material. These tables report the empirical sizes of seven different tests under four distinct covariance structures $\mathbf{\Sigma}$, corresponding to four choices of $\mathbf{W}_{\mathbf{\Sigma}_{1p},\mathrm{opt}}$. The results indicate that the empirical sizes of the proposed test T_{LMPNT} are well controlled around the nominal level $\alpha = 0.05$ across all scenarios, and that the other tests—T_{CZZ}, T_{CLRT}, T_{Fishera}, T_{Fisherb} and T_{DW} —exhibit broadly similar behavior. By contrast, T_{CLRT} is applicable only when $p < n$. Moreover, T_{Fisherb} performs well only under a Gaussian population, and its empirical sizes deviate markedly from 0.05 when z_{ij} follows a Gamma distribution. The maximum-type test T_{XW} also shows a notable departure from the nominal level 0.05, which can be attributed to the slow convergence rate of the Type I extreme value distribution; nevertheless, as $(p, n) \to \infty$, its empirical sizes tend to approach 0.05.

Figures 1–2 display the empirical powers of the seven tests under Scenarios 3–4. The corresponding power curves for Scenarios 1–2 are reported in Figures S.1–S.2 of the Supplementary Material.

Under local alternatives with $q_i = O(1/p)$, as in Scenario 3, the proposed T_{LMPNT} consistently achieves the highest power among all competitors across the considered (n, p) combinations. A clear power gain over T_{CZZ} is observed, with the improvement particularly pronounced in this heterogeneous multi-bulk setting, where multiple spikes in the covariance structure allow the procedure to concentrate more power on the most informative directions. This behavior is also in line with the LMPNT theory developed in Sections 3 and 4, which predicts that T_{LMPNT} is locally most powerful within the optional naive test family under suitably structured local alternatives.

In Scenario 4, where the perturbation parameters $\{q_i\}$ are of constant order, the alternative represents a stronger, nonlocal departure from the identity. As n increases, many procedures gradually approach power one, reflecting the stronger separation between the null and alternative. Nevertheless, T_{LMPNT} reaches high power at substantially smaller sample sizes than most competing methods, indicating that its advantage persists beyond the local asymptotic regime. We also observe that T_{DW} performs well in higher-dimensional regimes, especially when $p > n$.

Overall, the simulation results suggest that T_{LMPNT} attains accurate size and favorable local power under $O(1/p)$ alternatives, exhibits stable performance when Assumption 3 does not hold, and continues to perform well under fixed-order perturbations, often outperforming the classical test T_{CZZ} of Chen et al. (2010) in the settings considered.

5.2 Plug-in implementation with estimated $\mathbf{W}$

The simulations in Section 5.1 evaluate T_{LMPNT} using the oracle operator $\mathbf{W}_{\mathbf{\Sigma}_{1p},\mathrm{opt}}$, which depends on the population alternative covariance matrix and is therefore unavailable in applications. To examine practical implementability, we also consider a sample-splitting plug-in version, denoted by T_{SLMPNT}. In each Monte Carlo replication, the sample is randomly split into a training part and a testing part. A Ledoit–Wolf shrinkage covariance estimator of Ledoit and Wolf (2004) is computed from the training sample, Algorithm 1 is applied to construct a plug-in operator $\widehat{\mathbf{W}}$, and the testing sample is used only to evaluate $T_{n_{\mathrm{te}}}(\widehat{\mathbf{W}})$. Conditional on the training sample, $\widehat{\mathbf{W}}$ is fixed with respect to the testing sample. Hence, the fixed-$\mathbf{W}$ null calibration in Theorem 1 applies conditionally.

The detailed simulation setting and numerical results for this plug-in implementation are reported in Section S3.3 of the Supplementary Material. The corresponding results show that the empirical sizes of T_{SLMPNT} are

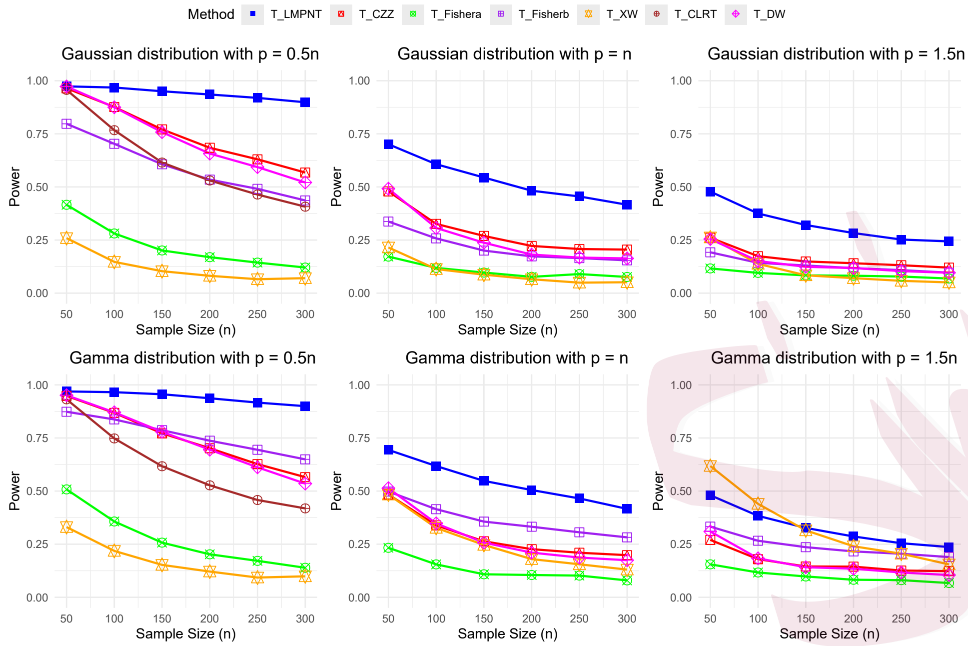

Figure 1: Empirical powers of seven tests at 5% significance in Scenario 3.

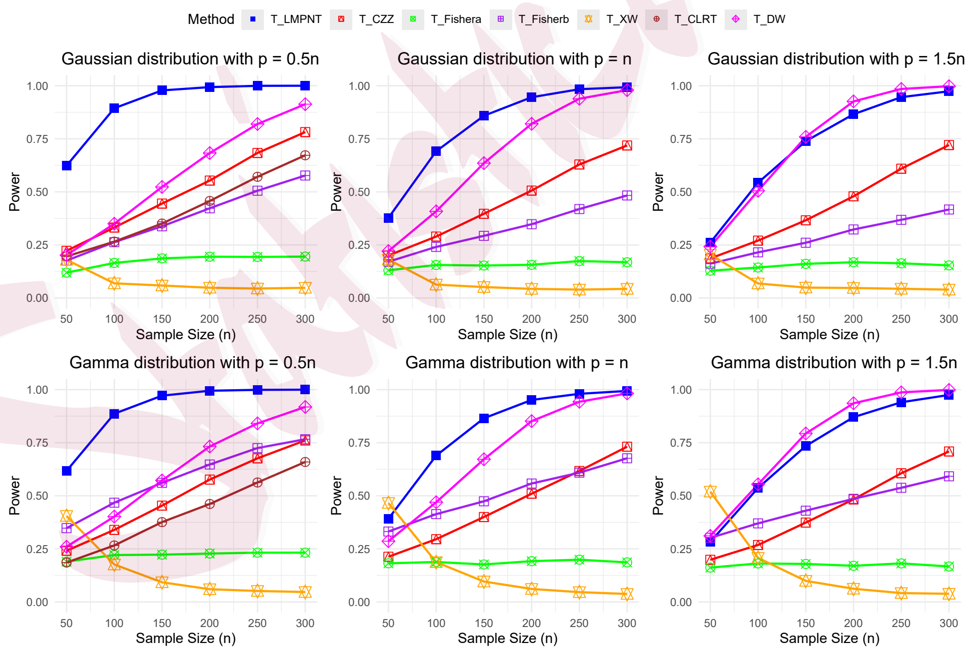

Figure 2: Empirical powers of seven tests at 5% significance in Scenario 4.

reasonably close to the nominal level and that its empirical power exceeds that of T_{CZZ} in several settings.

6. Real-data analysis

We illustrate the proposed procedure using two independent breast cancer gene-expression cohorts publicly available from the NCBI Gene Expression Omnibus and profiled on the same Affymetrix Human Genome U133A Array: the TRANSBIG cohort (GSE7390; Desmedt et al., 2007), containing $n = 198$ samples, and the NKI cohort (GSE2034; van de Vijver et al., 2002), containing $n = 286$ samples. Both cohorts contain the same 22,283 probe sets. We use the identity covariance model as a basic statistical benchmark and assess whether the covariance structure of the NKI gene-expression data is compatible with

$$H_0 : \mathbf{\Sigma} = \mathbf{I}_p \qquad \text{versus} \qquad H_1 : \mathbf{\Sigma} \neq \mathbf{I}_p.$$

The TRANSBIG cohort is used only for feature screening and operator construction, whereas the NKI cohort is reserved solely for the final hypothesis test. Consequently, with respect to the NKI testing sample, both the selected subset of probe sets and the learned operator $\mathbf{W}$ are determined externally and may be regarded as fixed.

After aligning the common probe sets across the two cohorts, we perform feature screening using only the original TRANSBIG expression matrix. Specifically, we retain the $p = 2000$ probe sets with the largest sample mean expression levels in the TRANSBIG cohort. Both cohorts are then restricted to the same set of selected probe sets and centered separately within each cohort. Based on the centered TRANSBIG data, we compute the Ledoit–Wolf shrinkage covariance estimator $\widehat{\boldsymbol{\Sigma}}_{\mathrm{LW}}$ (Ledoit and Wolf, 2004) and construct $\mathbf{W}$ according to Algorithm 1. The final covariance identity tests are applied only to the centered NKI data. Detailed information on data acquisition and preprocessing is provided in Section S4 of the Supplementary Material.

Table 1: Results of covariance identity testing for the NKI breast cancer dataset (GSE2034) with $p = 2000$.

	T_{LMPNT}	T_{CZZ}	T_{Fishera}	T_{Fisherb}	T_{XW}	T_{CLRT}	T_{DW}
Test statistic	2.91×10^{19}	7.45×10^{18}	4.91×10^{37}	2.90×10^{37}	1.85×10^{2}	NA	7.83×10^{18}
log(p-value)	-4.23×10^{38}	-2.78×10^{37}	-1.20×10^{75}	-4.19×10^{74}	-9.25×10^{1}	NA	-3.06×10^{37}

Table 1 reports the test statistics and the corresponding log-p-values for the competing methods. Since $p > n$, the CLRT method of Zheng et al. (2015) is not applicable in this setting. All remaining procedures yield extremely small p-values, indicating that $H_0 : \boldsymbol{\Sigma} = \mathbf{I}_p$ is rejected at any conventional significance level. For the selected $p = 2000$ probe sets, this suggests that the covariance structure of the centered NKI expression data deviates substantially from the identity matrix. This real-data example further illustrates the practical implementability of the proposed optional naive test in a high-dimensional setting.

7. Discussion

Beyond the setting considered in this paper, several directions merit further study. First, it would be natural to extend the operator-indexed naive framework and the associated local optimality perspective to other high-dimensional covariance testing problems, such as testing sphericity, bandedness, or the equality of two covariance matrices. Testing whether a correlation matrix is the identity is another important related problem. However, after standardizing the data using sample means and sample standard deviations, the resulting statistic involves additional randomness and dependence, so this problem is not directly covered by the covariance identity framework considered here and requires separate theoretical treatment.

Second, it would be worthwhile to extend the theoretical analysis of the proposed statistic to elliptical models. Recent results indicate that classical covariance U-tests may require additional corrections to retain valid asymptotic null distributions in elliptically distributed high-dimensional settings; see, for example, Ding and Xie (2025) and Xu et al. (2025). This suggests that such an extension is not merely a routine relaxation of the linear pseudo-independence condition in Assumption 1. Establishing the null distribution and local power theory of $T_n(\mathbf{W})$ under such models is therefore left for future work.

A remaining practical issue is the choice of $\mathbf{W}$ when the alternative covariance structure is unknown. Scientifically meaningful structural information, when available, can guide the construction of a feasible operator; otherwise, a general-purpose covariance estimator, such as a shrinkage estimator, can be used to estimate $\mathbf{\Sigma}_{1p}$ and construct a plug-in operator $\widehat{\mathbf{W}}$ through Algorithm 1. Since the null theory in Theorem 1 is established for an operator $\mathbf{W}$ that is fixed with respect to the final testing sample, any data-driven construction of $\widehat{\mathbf{W}}$ should rely on data independent of that sample. This motivates the sample-splitting procedure in Section 5.2 and the external-cohort analysis in Section 6. If this independence condition fails, the present fixed-$\mathbf{W}$ null theory is no longer directly applicable, and fully adaptive choices of $\mathbf{W}$ are left for future work.

Overall, this paper provides a first step toward formulating local optimality principles for high-dimensional covariance identity testing. We identify an operator-indexed optional naive test family and show that, within this family and under a structured class of local alternatives, the oracle operator yields a locally most powerful naive test. This perspective helps make precise the notion of "most powerful within a structured family" in high-dimensional covariance testing. We hope that this perspective will stimulate further work on optimality principles for high-dimensional inference.

Supplementary Material

The Supplementary Material contains detailed proofs, additional simulation studies, and further details for the real-data analysis.

Acknowledgements

Zhidong Bai was supported by the Jilin Province Science and Technology Development Plan Project (No.YDZJ202501ZYTS593).

References

Anderson, T. W. (2003). *An introduction to multivariate statistical analysis* (3rd ed.). New York: Wiley.

Bai, Z. D., D. D. Jiang, J. F. Yao, and S. R. Zheng (2009). Corrections to LRT on large-dimensional covariance matrix by RMT. *Annals of Statistics 37*(6B), 3822–3840.

Bai, Z. D. and H. Saranadasa (1996). Effect of high dimension: by an example of a two sample problem. *Statistica Sinica 6*, 311–329.

Cai, T. T. and Z. M. Ma (2013). Optimal hypothesis testing for high-dimensional covariance matrices. *Bernoulli 19*(5), 2359–2388.

Chen, S. X., Y. M. Qiu, and S. Y. Zhang (2023). Sharp optimality for high-dimensional covariance testing under sparse signals. *Annals of Statistics 51*(5), 1921–1945.

Chen, S. X., L. X. Zhang, and P. S. Zhong (2010). Tests for high-dimensional covariance matrices. *Journal of the American Statistical Association 105*(490), 810–819.

Desmedt, C., F. Piette, S. Loi, et al. (2007). Strong time dependence of the 76-gene prognostic signature for node-negative breast cancer patients in the transbig multicenter independent validation series. *Clinical Cancer Research 13*(11), 3207–3214.

Ding, X. C. (2021). Spiked sample covariance matrices with possibly multiple bulk components. *Random Matrices: Theory and Applications 10*(1), 2150014.

Ding, X. C., Y. C. Hu, and Z. G. Wang (2025). Two sample test for covariance matrices in ultra-high dimension. *Journal of the American Statistical Association 120*(552), 2210–2221.

Ding, X. C. and Z. G. Wang (2025). Global and local CLTs for linear spectral statistics of general sample covariance matrices when the dimension is much larger than the sample size with applications. *IEEE Transactions on Information Theory 71*(8), 6248–6296.

Ding, X. C. and J. H. Xie (2025). Tracy–Widom distribution for the edge eigenvalues of elliptical model. *Information and Inference: A Journal of the IMA 14*(2), iaaf004.

Fisher, T. J. (2012). On testing for an identity covariance matrix when the dimensionality equals or exceeds the sample size. *Journal of Statistical Planning and Inference 142*(1), 312–326.

Hu, J. and Z. D. Bai (2016). A review of 20 years of naive tests of significance for high-dimensional mean vectors and covariance matrices. *Science China Mathematics 59*, 2281–2300.

Jiang, T. F. and Y. C. Qi (2015). Likelihood ratio tests for high-dimensional normal distributions. *Scandinavian Journal of Statistics 42*(4), 988–1009.

Jiang, T. F. and F. Yang (2013). Central limit theorems for classical likelihood ratio tests for high-dimensional normal distributions. *Annals of Statistics 41*(4), 2029–2074.

Ledoit, O. and M. Wolf (2002). Some hypothesis tests for the covariance matrix when the dimension is large compared to the sample size. *Annals of Statistics 30*(4), 1081–1102.

Ledoit, O. and M. Wolf (2004). A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis 88*(2), 365–411.

Nagao, H. (1973). On some test criteria for covariance matrix. *Annals of Statistics 1*(4), 700–709.

Srivastava, M. S., H. Yanagihara, and T. Kubokawa (2014). Tests for covariance matrices in high dimension with less sample size. *Journal of Multivariate Analysis 130*, 289–309.

van de Vijver, M. J., Y. D. He, L. J. van 't Veer, et al. (2002). A gene-expression signature as a predictor of survival in breast cancer. *New England Journal of Medicine 347*(25), 1999–2009.

Wang, L., B. Peng, and R. Z. Li (2015). A high-dimensional nonparametric multivariate test for mean vector. *Journal of the American Statistical Association 110*(512), 1658–1669.

Wilks, S. S. (1932). Certain generalizations in the analysis of variance. *Biometrika 24*(3–4), 471–494.

Wu, T. L. and P. Li (2020). Projected tests for high-dimensional covariance matrices. *Journal of Statistical Planning and Inference 207*, 73–85.

Xiao, H. and W. B. Wu (2013). Asymptotic theory for maximum deviations of sample covariance matrix estimates. *Stochastic Processes and their Applications 123*(7), 2899–2920.

Xu, K., Y. Q. Zhou, and L. P. Zhu (2025). Adjusted location-invariant U-tests for the covariance matrix with elliptically high-dimensional data. *Scandinavian Journal of Statistics 52*(1), 249–269.

Yang, S. S., S. R. Zheng, and R. Z. Li (2024). A new test for high-dimensional two-sample mean problems with consideration of correlation structure. *Annals of Statistics 52*(5), 2217–2240.

Zheng, S. R., Z. D. Bai, and J. F. Yao (2015). Substitution principle for CLT of linear spectral statistics of high-dimensional sample covariance matrices with applications to hypothesis testing. *Annals of Statistics 43*(2), 546–591.

KLASMOE and School of Mathematics and Statistics, Northeast Normal University, Changchun 130022, China.

E-mail: luxn157@nenu.edu.cn

E-mail: huj156@nenu.edu.cn

E-mail: baizd@nenu.edu.cn

E-mail: songhy716@nenu.edu.cn