

High-dimensional general linear hypothesis tests via non-linear spectral shrinkage

HAORAN LI¹, ALEXANDER AUE^{2,*} and DEBASHIS PAUL^{2,†}

¹*Department of Statistics, Columbia University, 1255 Amsterdam Avenue, New York, NY 10027, USA.
E-mail: hl3369@columbia.edu*

²*Department of Statistics, University of California at Davis, One Shields Avenue, Davis, CA 95616, USA.
E-mail: *aaue@ucdavis.edu; †debpaull@ucdavis.edu*

We are interested in testing general linear hypotheses in a high-dimensional multivariate linear regression model. The framework includes many well-studied problems such as two-sample tests for equality of population means, MANOVA and others as special cases. A family of rotation-invariant tests is proposed that involves a flexible spectral shrinkage scheme applied to the sample error covariance matrix. The asymptotic normality of the test statistic under the null hypothesis is derived in the setting where dimensionality is comparable to sample sizes, assuming the existence of certain moments for the observations. The asymptotic power of the proposed test is studied under various local alternatives. The power characteristics are then utilized to propose a data-driven selection of the spectral shrinkage function. As an illustration of the general theory, we construct a family of tests involving ridge-type regularization and suggest possible extensions to more complex regularizers. A simulation study is carried out to examine the numerical performance of the proposed tests.

Keywords: general linear hypothesis; local alternatives; random matrix theory; ridge shrinkage; spectral shrinkage

1. Introduction

In multivariate analysis, one of the fundamental inferential problems is to test a hypothesis involving a linear transformation of regression coefficients under a linear model. Suppose \mathbf{Y} is a $p \times N$ matrix of observations modeled as

$$\mathbf{Y} = \mathbf{B}\mathbf{X} + \Sigma_p^{1/2}\mathbf{Z}, \quad (1.1)$$

where (i) \mathbf{B} is a $p \times k$ matrix of regression coefficients; (ii) \mathbf{X} is a $k \times N$ design matrix of rank k ; (iii) \mathbf{Z} is a $p \times N$ matrix with i.i.d. entries having zero mean and unit variance; and (iv) Σ_p , a $p \times p$ nonnegative definite matrix, is the population covariance matrix of the errors, with $\Sigma_p^{1/2}$ a “square-root” of Σ_p so that $\Sigma_p = \Sigma_p^{1/2}(\Sigma_p^{1/2})^T$. General linear hypotheses involving the linear model (1.1) are of the form

$$H_0 : \mathbf{B}\mathbf{C} = \mathbf{0} \quad \text{vs.} \quad H_a : \mathbf{B}\mathbf{C} \neq \mathbf{0} \quad (1.2)$$

for an arbitrary $k \times q$ “constraints matrix” \mathbf{C} , subject to the requirement that $\mathbf{B}\mathbf{C}$ is estimable. Without loss of generality, \mathbf{C} is taken to be of rank q . Throughout, we assume that q and k

are fixed, even as observation dimension p and sample size N increase to infinity. Henceforth, $n = N - k$ is used to denote the effective sample size, which is also the degree of freedom associated with the sample error covariance matrix.

With various choices of X and C , the testing formulation incorporates many hypotheses of interest. For example, multivariate analysis of variance (MANOVA) is a special case. When the sample size N is substantially larger than the dimension p of the observations, this problem is well-studied. Anderson [1] and Muirhead [25] are among standard references. Various classical inferential procedures involve the matrices

$$\widehat{\Sigma}_p = \frac{1}{n} \mathbf{Y} (I - X^T (X X^T)^{-1} X) \mathbf{Y}^T, \quad (1.3)$$

$$\widehat{\mathbf{H}}_p = \frac{1}{n} \mathbf{Y} X^T (X X^T)^{-1} C [C^T (X X^T)^{-1} C]^{-1} C^T (X X^T)^{-1} X \mathbf{Y}^T, \quad (1.4)$$

so that $\widehat{\Sigma}_p$ is the residual covariance of the full model, an estimator of Σ_p , while $\widehat{\mathbf{H}}_p$ is the hypothesis sums of squares and cross products matrix, scaled by n^{-1} . In a one-way MANOVA set-up, $\widehat{\Sigma}_p$ and $\widehat{\mathbf{H}}_p$ are, respectively, the within-group and between-group sums of squares and products matrices, scaled by n^{-1} . In the rest of the paper, we shall refer to $\widehat{\Sigma}_p$ as the sample covariance matrix.

The testing problem (1.2) is well-studied in the classical multivariate analysis literature. Three standard test procedures are the likelihood ratio test (LR), Lawley–Hotelling trace test (LH) and Bartlett–Nanda–Pillai trace (BNP) test. They are called *invariant tests*, since under Gaussianity the null distributions of the test statistics are invariant with respect to Σ_p . One common feature is that all test statistics are linear functionals of the spectrum of $\widehat{\mathbf{H}}_p \widehat{\Sigma}_p^{-1}$. Since this matrix is asymmetric, for convenience, a standard transformation is applied, giving the expressions of the invariant tests as follows. Define

$$\begin{aligned} Q_n &= X^T (X X^T)^{-1} C [C^T (X X^T)^{-1} C]^{-1/2}, \\ \mathbf{M}_0 &= \frac{1}{n} Q_n^T \mathbf{Y}^T \widehat{\Sigma}_p^{-1} \mathbf{Y} Q_n. \end{aligned} \quad (1.5)$$

The matrix $Q_n Q_n^T$ is the “hat matrix” of the reduced model under the null hypothesis. Note that the non-zero eigenvalues of $\widehat{\mathbf{H}}_p \widehat{\Sigma}_p^{-1} = n^{-1} \mathbf{Y} Q_n Q_n^T \mathbf{Y}^T \widehat{\Sigma}_p^{-1}$ are the same as those of \mathbf{M}_0 . The test statistics for the LR, LH and BNP tests can be expressed as

$$\begin{aligned} T_0^{\text{LR}} &= \sum_{i=1}^q \log \{1 + \lambda_i(\mathbf{M}_0)\}, \\ T_0^{\text{LH}} &= \sum_{i=1}^q \lambda_i(\mathbf{M}_0), \\ T_0^{\text{BNP}} &= \sum_{i=1}^q \lambda_i(\mathbf{M}_0) / \{1 + \lambda_i(\mathbf{M}_0)\}. \end{aligned}$$

The symbol $\lambda_i(\cdot)$ denotes the i -th largest eigenvalue of a symmetric matrix, further using the convention that $\lambda_{\max}(\cdot)$ and $\lambda_{\min}(\cdot)$ indicate the largest and smallest eigenvalue.

In contemporary statistical research and applications, high-dimensional data whose dimension is at least comparable to the sample size is ubiquitous. In this paper, focus is on the interesting boundary case when dimension and sample sizes are comparable. Primarily due to inconsistency of conventional estimators of model parameters – such as $\widehat{\Sigma}_p$ –, classical test procedures for the hypothesis (1.2) – such as the LR, LH and BNP tests – perform poorly in such settings. When the dimension p is larger than the degree of freedom n , the invariant tests are not even well-defined because $\widehat{\Sigma}_p$ is singular. Even when p is strictly less than n , but the ratio $\gamma_n = p/n$ is close to 1, these tests are known to have poor power behavior. Asymptotic results when $\gamma_n \rightarrow \gamma \in (0, 1)$ were obtained in [16] under Gaussianity of the populations, and more recently in [3] under more general settings that only require the existence of certain moments. LR, LH and BNP tests can be generalized as linear spectral statistics of large-dimensional F -matrices, whose CLT is studied in [9,33,34].

Pioneering work on modifying the classical solutions in high dimension is in [4], who corrected the scaling of the LR statistic when $n \geq p$ but p , k and q are proportional to n . The corrected LR statistic was shown to have significantly more power than its classical counterpart. In contrast, in this paper, we focus on the setting where k and q are fixed even as n , $p \rightarrow \infty$ so that $\gamma_n = p/n \rightarrow \gamma \in (0, \infty)$. In the multivariate regression problem, this corresponds to a situation where the response is high-dimensional, while the predictor is finite-dimensional. In the MANOVA problem, this framework corresponds to high-dimensional observations belonging to one of a finite number of populations.

To the best of our knowledge, when $n < p$, the linear hypothesis testing problem has been studied in depth only for specific submodels of (1.1), primarily for the important case of two-sample tests for equality of population means. For the latter tests, a widely used idea is to construct modified statistics based on replacing $\widehat{\Sigma}_p^{-1}$ with an appropriate substitute. This approach was pioneered in [5] and further developed in [15]. Various extensions to one-way MANOVA [18, 28,31] and a general multi-sample Behrens–Fisher problem under heteroscedasticity [35] exist. Other notable works for the two-sample problem include [8,12,14,17,24,29,30]. A second approach aims to regularize $\widehat{\Sigma}_p$ to address the issue of its near-singularity in high dimensions; see [13] and [22] for ridge-type penalties in two-sample settings. Finally, another alternative line of attack consists of exploiting sparsity; see [10,11]. Other related works include [36].

In this paper, we seek to regularize the spectrum of $\widehat{\Sigma}_p$ by flexible shrinkage functions. For a symmetric $p \times p$ matrix A and a function $g(\cdot)$ on \mathbb{R} , define

$$g(A) = R_A \text{diag}(g(\lambda_1(A)), \dots, g(\lambda_p(A))) R_A^T,$$

where R_A is the matrix of eigenvectors associated with the ordered eigenvalues of A . Now, consider any real-valued function $f(\cdot)$ on \mathbb{R} that is analytic over a specific domain associated with the limiting behavior of the eigenvalues of $\widehat{\Sigma}_p$, as elaborated in Section 2. The proposed statistics are functionals of eigenvalues of the regularized quadratic forms

$$\mathbf{M}(f) = \frac{1}{n} Q_n^T \mathbf{Y}^T f(\widehat{\Sigma}_p) \mathbf{Y} Q_n.$$

Specifically, we propose regularized versions of LR, LH and BNP test criteria, namely

$$\begin{aligned} T^{\text{LR}}(f) &= \sum_{i=1}^q \log\{1 + \lambda_i(\mathbf{M}(f))\}, \\ T^{\text{LH}}(f) &= \sum_{i=1}^q \lambda_i(\mathbf{M}(f)), \\ T^{\text{BNP}}(f) &= \sum_{i=1}^q \lambda_i(\mathbf{M}(f)) / \{1 + \lambda_i(\mathbf{M}(f))\}. \end{aligned}$$

These test statistics are designed to capture possible departures from the null hypothesis, when $\widehat{\Sigma}_p$ is replaced by $f(\widehat{\Sigma}_p)$, while suitable choices of the regularizer f allow for getting around the problem of singularity or near-singularity when p is comparable to n .

Notice that $\mathbf{M}(f)$ has the same non-zero eigenvalues as $f(\widehat{\Sigma}_p)\widehat{\mathbf{H}}_p$. Thus, the proposed test family is a generalization of the classical statistics based on $\widehat{\Sigma}_p^{-1}\widehat{\mathbf{H}}_p$. Importantly, $\mathbf{M}(f)$ – and consequently the proposed statistics – is *rotation-invariant*, which means if a linear transformation is applied to the observations with an arbitrary orthogonal matrix, the statistic remains unchanged. It is a desirable property when not much additional knowledge about Σ_p and BC is available. It should be noted that the two-sample mean tests by Bai and Saranadasa [5] and Li et al. [22], together with their generalization to MANOVA, are special cases of the proposed family with $f(x) = 1$ and $f(x) = 1/(x + \lambda)$, $\lambda > 0$, respectively.

The present work builds on the work by Li et al. [22]. The theoretical analysis also involves an extension of the analytical framework adopted by Pan and Zhou [26] in their study of the asymptotic behavior of Hotelling's T^2 statistic for non-Gaussian observations. However, the current work goes well beyond the existing literature in several aspects. We highlight these as the key contributions of this manuscript: (a) We propose new families of rotation-invariant tests for general linear hypotheses for multivariate regression problems involving high-dimensional response and fixed-dimensional predictor variables that incorporate a flexible regularization scheme to account for the dimensionality of the observations growing proportional to the sample size. (b) Unlike Li et al. [22], who assumed sub-Gaussianity, here only the existence of finite fourth moments of the observations is required. (c) Unlike Pan and Zhou [26], who assumed $\Sigma_p = I_p$, Σ_p is allowed to be fairly arbitrary and subjected only to some standard conditions on the limiting behavior of its spectrum. (d) We carry out a detailed analysis of the power characteristics of the proposed tests. The proposal of a class of local alternatives enables a clear interpretation of the contributions of different parameters in the performance of the test. (e) We develop a data-driven test procedure based on the principle of maximizing asymptotic power under appropriate local alternatives. This principle leads to the definition of a composite test that combines the optimal tests associated with a set of different kinds of local alternatives. The latter formulation is an extension of the data-adaptive test procedure designed by Li et al. [22] for the two-sample testing problem.

The rest of the paper is organized as follows. Section 2 introduces the asymptotics of the proposed test family both under the null hypothesis and under a class of local alternatives. Using

these local alternatives, in Section 3 a data-driven shrinkage selection methodology based on maximizing asymptotic power is developed. In Section 4, an application of the asymptotic theory and the shrinkage selection method is given for the ridge-regularization family. An extension of ridge-regularization to higher orders is also discussed. The results of a simulation study are reported in Section 5. In the Appendix, a proof outline of the main theorem is presented, while technical details and proofs of other theorems are collected in the Supplementary Material [21].

2. Asymptotic theory

After giving necessary preliminaries on *Random Matrix Theory* (RMT), the asymptotic theory of the proposed tests under the null hypothesis and under various local alternative models is presented in this section. For any $p \times p$ symmetric matrix A , define the *Empirical Spectral Distribution* (ESD) F^A of A by

$$F^A(\tau) = p^{-1} \sum_{i=1}^p \mathbb{1}_{\{\lambda_i(A) \leq \tau\}}.$$

In the following, $\|\cdot\|_{\max}$ stands for the maximum absolute value of the entries of a matrix. The following assumptions are employed.

C1 (Moment conditions) \mathbf{Z} has i.i.d. entries z_{ij} such that $\mathbb{E}z_{ij} = 0$, $\mathbb{E}z_{ij}^2 = 1$, $\mathbb{E}z_{ij}^4 < \infty$.

C2 (High-dimensional setting) k and q are fixed, while $p, n \rightarrow \infty$ such that $\gamma_n = p/n \rightarrow \gamma \in (0, \infty)$ and $\sqrt{n}|\gamma_n - \gamma| \rightarrow 0$.

C3 (Boundedness of spectral norm) Σ_p is non-negative definite; $\limsup_p \lambda_{\max}(\Sigma_p) < \infty$.

C4 (Asymptotic stability of ESD) There exists a distribution L^Σ with compact support in $[0, \infty)$, non-degenerate at zero, such that $\sqrt{n}\mathcal{D}_W(F^{\Sigma_p}, L^\Sigma) \rightarrow 0$, as $n, p \rightarrow \infty$, where $\mathcal{D}_W(F_1, F_2)$ denotes the *Wasserstein distance* between distributions F_1 and F_2 , defined as

$$\mathcal{D}_W(F_1, F_2) = \sup_f \left\{ \left| \int f dF_1 - \int f dF_2 \right| : f \text{ is 1-Lipschitz} \right\}.$$

C5 (Asymptotically full rank) X is of full rank and $n^{-1}XX^T$ converges to a positive definite $k \times k$ matrix. Moreover, $\limsup_{n \rightarrow \infty} \|X\|_{\max} < \infty$.

C6 (Asymptotically estimable) $\liminf_{n \rightarrow \infty} \lambda_{\min}(C^T(n^{-1}XX^T)^{-1}C) > 0$.

Remark 2.1. The conditions are mild: **C5** and **C6** are commonly made in multivariate analysis for asymptotic results of regression models under non-Gaussianity. The $o(n^{-1/2})$ convergence rate in **C2** and **C4** is unnecessary for proving the asymptotic normality of the proposed tests introduced in Section 2.2 and 2.3. The assumption can be dropped if $m(\mathbb{Z})$ introduced in (2.1) is replaced by the solution to the equation $m_p^0(\mathbb{Z}) = \int \{\tau(1 - \gamma_n - \gamma_n \mathbb{Z} m_p^0(\mathbb{Z}) - \mathbb{Z})^{-1} dF^{\Sigma_p}(\tau)\}$. However, such a modification will significantly complicate mathematical expressions. In order to emphasize readability and succinctness, the $o(n^{-1/2})$ convergence rate is adopted. Moreover, for the purpose of deriving validity of the data-driven selection of shrinkage functions introduced in

Section 3, the $o(n^{-1/2})$ convergence rate is necessary. Notably, the convergence rate assumption is practically not overly restrictive. First, it imposes little constraint on the observations. Secondly, we always use γ_n and F^{Σ_p} to estimate γ and L^Σ in the proposed inferential procedure.

2.1. Preliminaries on random matrix theory

Recall that the Stieltjes transform $m_G(\cdot)$ of any function G of bounded variation on \mathbb{R} is defined by

$$m_G(z) = \int_{-\infty}^{\infty} \frac{dG(x)}{x - z}, \quad z \in \mathbb{C}^+ := \{u + iv : v > 0\}.$$

Minor modifications of a standard RMT result imply that, under Conditions **C1**–**C6**, the ESD $F^{\widehat{\Sigma}_p}$ converges almost surely to a nonrandom distribution F^∞ at all points of continuity of F^∞ . This limit is determined in such a way that for any $z \in \mathbb{C}^+$, the Stieltjes transform $m(\cdot) = m_{F^\infty}(\cdot)$ of F^∞ is the unique solution in \mathbb{C}^+ of the equation

$$m(z) = \int \frac{dL^\Sigma(\tau)}{\tau(1 - \gamma - \gamma z m(z)) - z}. \quad (2.1)$$

Equation (2.1) is often referred to as the Marčenko–Pastur equation. Moreover, pointwise almost surely for $z \in \mathbb{C}^+$, $m_{F^{\widehat{\Sigma}_p}}(z)$ converges to $m_{F^\infty}(z)$. The convergence holds even when $z \in \mathbb{R}_-$ (negative reals) with a smooth extension of m_{F^∞} to \mathbb{R}_- . Readers may refer to [6] and [27] for more details. From now on, for notational simplicity, we shall write $m_{F^\infty}(z)$ as $m(z)$ and write $m_{F^{\widehat{\Sigma}_p}}(z)$ as $m_{n,p}(z)$. Note that

$$m_{n,p}(z) = p^{-1} \operatorname{tr}(\widehat{\Sigma}_p - zI_p)^{-1}$$

and define

$$\Theta(z, \gamma) = \{1 - \gamma - \gamma z m(z)\}^{-1}. \quad (2.2)$$

It is known that $(\widehat{\Sigma}_p - zI_p)^{-1}$, for any fixed $z \in \mathbb{C}^+$, has a *deterministic equivalent* [6,22,23], given by

$$\{\Theta^{-1}(z, \gamma)\Sigma_p - zI\}^{-1},$$

in the sense that for symmetric matrices A bounded in operator norm, as $n \rightarrow \infty$,

$$p^{-1} \operatorname{tr}[(\widehat{\Sigma}_p - zI_p)^{-1}A] - p^{-1} \operatorname{tr}[\{\Theta^{-1}(z, \gamma)\Sigma_p - zI\}^{-1}A] \rightarrow 0, \quad \text{with probability 1.}$$

Resolvent and deterministic equivalent will be used frequently in this paper. They will appear for example as Cauchy kernels in contour integrals in various places.

2.2. Asymptotics under the null hypothesis

To begin with, for $q \geq 1$, denote by $\mathbf{W} = [w_{ij}]_{i,j=1}^q$ the *Gaussian Orthogonal Ensemble (GOE)* defined by (1) $w_{ij} = w_{ji}$; (2) $w_{ii} \sim \mathcal{N}(0, 1)$, $w_{ij} \sim \mathcal{N}(0, 1/2)$, $i \neq j$; (3) w_{ij} 's are jointly independent for $1 \leq i \leq j \leq q$. Throughout this paper, $f(\cdot)$ is assumed to be analytic in an open interval containing

$$\mathcal{X} := \left[0, \limsup_{p \rightarrow \infty} \lambda_{\max}(\Sigma_p)(1 + \sqrt{\gamma})^2\right].$$

Let \mathcal{C} to be a closed contour enclosing \mathcal{X} such that $f(\cdot)$ has a complex extension to the interior of \mathcal{C} . Further use \mathcal{C}^2 to denote $\mathcal{C} \otimes \mathcal{C} = \{\mathbb{z}_1, \mathbb{z}_2 : \mathbb{z}_1, \mathbb{z}_2 \in \mathcal{C}\}$.

Theorem 2.1. Suppose **C1–C6** hold. Under the null hypothesis $H_0 : BC = 0$,

$$\sqrt{n}\{\mathbf{M}(f) - \Omega(f, \gamma)I_q\} \Longrightarrow \Delta^{1/2}(f, \gamma)\mathbf{W},$$

where \Longrightarrow denotes weak convergence and $\Omega(f, \gamma)$ and $\Delta(f, \gamma)$ are as follows.

$$\Omega(f, \gamma) = \frac{-1}{2\pi i} \oint_{\mathcal{C}} f(\mathbb{z})(\Theta(\mathbb{z}, \gamma) - 1) d\mathbb{z}.$$

See (2.2) for the definition of $\Theta(\mathbb{z}, \gamma)$. For any two analytic functions f_1 and f_2 ,

$$\Delta(f_1, f_2, \gamma) = \frac{2}{(2\pi i)^2} \iint_{\mathcal{C}^2} f_1(\mathbb{z}_1) f_2(\mathbb{z}_2) \delta(\mathbb{z}_1, \mathbb{z}_2, \gamma) d\mathbb{z}_1 d\mathbb{z}_2,$$

and $\Delta(f, f, \gamma)$ is written as $\Delta(f, \gamma)$ for simplicity. The kernel $\delta(\mathbb{z}_1, \mathbb{z}_2, \gamma)$ is such that

$$\begin{aligned} \delta(\mathbb{z}_1, \mathbb{z}_2, \gamma) &= \Theta(\mathbb{z}_1, \gamma)\Theta(\mathbb{z}_2, \gamma) \left[\frac{\mathbb{z}_1\Theta(\mathbb{z}_1, \gamma) - \mathbb{z}_2\Theta(\mathbb{z}_2, \gamma)}{\mathbb{z}_1 - \mathbb{z}_2} - 1 \right], \\ \delta(\mathbb{z}, \mathbb{z}, \gamma) &= \lim_{\mathbb{z}_2 \rightarrow \mathbb{z}} \delta(\mathbb{z}, \mathbb{z}_2, \gamma) = \Theta^2(\mathbb{z}, \gamma) \left[\frac{\partial \mathbb{z} \Theta(\mathbb{z}, \gamma)}{\partial \mathbb{z}} - 1 \right] \\ &= \gamma \{1 + \mathbb{z}m(\mathbb{z})\} \Theta^3(\mathbb{z}, \gamma) + \gamma \mathbb{z} \{m(\mathbb{z}) + \mathbb{z}m'(\mathbb{z})\} \Theta^4(\mathbb{z}, \gamma). \end{aligned}$$

The contour integral is taken counter-clockwise.

Using knowledge of the eigenvalues of the **GOE** leads to the following statement.

Corollary 2.1. Under the conditions of Theorem 2.1, assume further that $\Delta(f, \gamma) > 0$. Let

$$\tilde{\lambda}_i = \frac{\sqrt{n}}{\Delta^{1/2}(f, \gamma)} \{\lambda_i(\mathbf{M}(f)) - \Omega(f, \gamma)\}, \quad i = 1, \dots, q.$$

Then, the limiting joint density function of $(\tilde{\lambda}_1, \dots, \tilde{\lambda}_q)$ at $y_1 \geq y_2 \geq \dots \geq y_q$ is given by

$$\left(2^{q/2} \prod_{i=1}^q \Gamma(i/2)\right)^{-1} \prod_{i < j} (y_i - y_j) \exp\left(-\frac{1}{2} \sum_{i=1}^q y_i^2\right).$$

Although without closed forms, $\Omega(f, \gamma)$ and $\Delta(f, \gamma)$ do not depend on the choice of \mathcal{C} used to compute the contour integral. With the resolvent as kernel $\mathbf{M}(f)$ can be expressed as the integral of $f(\mathbb{z})n^{-1}Q_n^T \mathbf{Y}^T (\widehat{\Sigma}_p - \mathbb{z}I_p)^{-1} \mathbf{Y} Q_n$ on any contour \mathcal{C} , up to a scaling factor. The quadratic form $n^{-1}Q_n^T \mathbf{Y}^T (\widehat{\Sigma}_p - \mathbb{z}I_p)^{-1} \mathbf{Y} Q_n$ is then shown to concentrate around $[\Theta(\mathbb{z}, \gamma) - 1]I_q$, which consequently serves as the integral kernel in $\Omega(f, \gamma)$. The kernel $\delta(\mathbb{z}_1, \mathbb{z}_2, \gamma)$ of $\Delta(f, \gamma)$ is the limit of $\mathbb{E}[n^{-1} \text{tr}\{(\widehat{\Sigma}_p - \mathbb{z}_1 I_p)^{-1} \Sigma_p (\widehat{\Sigma}_p - \mathbb{z}_2 I_p)^{-1} \Sigma_p\}]$.

Remark 2.2. Two sufficient conditions for $\Delta(f, \gamma) > 0$ are

- (1) $f(x) > 0$ for $x \in \mathcal{X}$;
- (2) $f(x) \geq 0$ for $x \in \mathcal{X}$, with $f(x) \neq 0$ for some $x \in \mathcal{X}$, and $\liminf \lambda_{\min}(\Sigma_p) > 0$.

It would be convenient if $\Omega(f, \gamma)$ and $\Delta(f, \gamma)$ had closed forms in order to avoid computational inefficiencies. Closed forms are available for special cases as shown in the following lemma.

Lemma 2.1. When $f(x, \ell) = (x - \ell)^{-1}$ with $\ell \in \mathbb{R}^-$, the contour integrals in Theorem 2.1 have closed forms, namely, for $j, j_1, j_2 = 0, 1, 2, \dots$,

$$\begin{aligned} \frac{-1}{2\pi i} \oint_{\mathcal{C}} \frac{\partial^j f(\mathbb{z}, \ell)}{\partial \ell^j} (\Theta(\mathbb{z}, \gamma) - 1) d\mathbb{z} &= \frac{\partial^j (\Theta(\ell, \gamma) - 1)}{\partial \ell^j}, \\ \frac{1}{(2\pi i)^2} \oint\!\!\!\oint_{\mathcal{C}^2} \frac{\partial^{j_1} f(\mathbb{z}_1, \ell_1)}{\partial \ell_1^{j_1}} \frac{\partial^{j_2} f(\mathbb{z}_2, \ell_2)}{\partial \ell_2^{j_2}} \delta(\mathbb{z}_1, \mathbb{z}_2, \gamma) d\mathbb{z}_1 d\mathbb{z}_2 &= \frac{\partial^{j_1+j_2} \delta(\ell_1, \ell_2, \gamma)}{\partial \ell_1^{j_1} \partial \ell_2^{j_2}}. \end{aligned}$$

The results continue to hold when $\ell \in \mathbb{C} \setminus \mathcal{X}$.

Lemma 2.1 indicates that it is possible to have convenient and accurate estimators of the asymptotic mean and variance of $\mathbf{M}(f)$ under ridge-regularization. The result easily generalizes to the setting when $f(x)$ is a linear combination of functions of the form $(x - \ell_j)^{-1}$, for any finite collection of ℓ_j 's. We elaborate on this in Section 4.

To conduct the tests, consistent estimators of $\Omega(f, \gamma)$ and $\Delta(f, \gamma)$ are needed.

Lemma 2.2. Let $\widehat{\Theta}(\mathbb{z}, \gamma_n)$ and $\widehat{\delta}(\mathbb{z}_1, \mathbb{z}_2, \gamma_n)$ be the plug-in estimators of $\Theta(\mathbb{z}, \gamma)$ and $\delta(\mathbb{z}_1, \mathbb{z}_2, \gamma)$, with $(m(\mathbb{z}), \gamma)$ estimated by $(m_{n,p}(\mathbb{z}), \gamma_n)$. For general f, f_1, f_2 , we can estimate $\Omega(f, \gamma)$ and $\Delta(f_1, f_2, \gamma)$ by replacing $\Theta(\mathbb{z}, \gamma)$ and $\delta(\mathbb{z}_1, \mathbb{z}_2, \gamma)$ with $\widehat{\Theta}(\mathbb{z}, \gamma_n)$ and

$\widehat{\delta}(\mathbb{Z}_1, \mathbb{Z}_2, \gamma_n)$. Denote the resulting estimators by $\widehat{\Omega}(f, \gamma_n)$ and $\widehat{\Delta}(f_1, f_2, \gamma_n)$. Then,

$$\sqrt{n}|\widehat{\Omega}(f, \gamma_n) - \Omega(f, \gamma)| \xrightarrow{P} 0,$$

$$\sqrt{n}|\widehat{\Delta}(f_1, f_2, \gamma_n) - \Delta(f_1, f_2, \gamma)| \xrightarrow{P} 0,$$

where \xrightarrow{P} indicates convergence in probability. Again, we write $\widehat{\Delta}(f, f, \gamma_n)$ as $\widehat{\Delta}(f, \gamma_n)$.

For the special case of $f^{(j)}(x, \ell) = \partial^j(x - \ell)^{-1}/\partial \ell^j$, $j = 0, 1, 2, \dots$ and $\ell \in \mathbb{C} \setminus \mathcal{X}$, using Lemma 2.1, natural estimators in closed forms are

$$\widehat{\Omega}(f^{(j)}(x, \ell), \gamma_n) = \frac{\partial^j(\widehat{\Theta}(\ell, \gamma_n) - 1)}{\partial \ell^j},$$

$$\widehat{\Delta}(f^{(j_1)}(x, \ell_1), f^{(j_2)}(x, \ell_2), \gamma_n) = \frac{\partial^{j_1+j_2} 2\widehat{\delta}(\ell_1, \ell_2, \gamma_n)}{\partial \ell_1^{j_1} \partial \ell_2^{j_2}}.$$

In particular, for $j, j_1, j_2 = 0$,

$$\widehat{\Omega}(f(x, \ell), \gamma_n) = \widehat{\Theta}(\ell, \gamma_n) - 1,$$

$$\widehat{\Delta}(f(x, \ell_1), f(x, \ell_2), \gamma_n) = 2\widehat{\delta}(\ell_1, \ell_2, \gamma_n).$$

The estimators are consistent, for any fixed j and ℓ . Given the eigenvalues of $\widehat{\Sigma}_p$, the computational complexity of calculating the above estimators is $O(p)$.

Recall the definitions of $T^{\text{LR}}(f)$, $T^{\text{LH}}(f)$ and $T^{\text{BNP}}(f)$ from Section 1.

Theorem 2.2. Suppose **C1–C6** hold and $\Delta(f, \gamma) > 0$. Under $H_0 : BC = 0$,

$$\widehat{T}^{\text{LR}}(f) := \frac{\sqrt{n}\{1 + \widehat{\Omega}(f, \gamma_n)\}}{q^{1/2}\widehat{\Delta}^{1/2}(f, \gamma_n)} [T^{\text{LR}}(f) - q \log\{1 + \widehat{\Omega}(f, \gamma_n)\}] \implies \mathcal{N}(0, 1),$$

$$\widehat{T}^{\text{LH}}(f) := \frac{\sqrt{n}}{q^{1/2}\widehat{\Delta}^{1/2}(f, \gamma_n)} \{T^{\text{LH}}(f) - q\widehat{\Omega}(f, \gamma_n)\} \implies \mathcal{N}(0, 1),$$

$$\widehat{T}^{\text{BNP}}(f) := \frac{\sqrt{n}\{1 + \widehat{\Omega}(f, \gamma_n)\}^2}{q^{1/2}\widehat{\Delta}^{1/2}(f, \gamma_n)} \left\{ T^{\text{BNP}}(f) - q \frac{\widehat{\Omega}(f, \gamma_n)}{1 + \widehat{\Omega}(f, \gamma_n)} \right\} \implies \mathcal{N}(0, 1).$$

For any of the three tests, the null hypothesis is rejected at asymptotic level α , if $\widehat{T}(f) > \xi_\alpha$, where ξ_α is the $1 - \alpha$ quantile of the standard normal distribution.

2.3. Asymptotic power under local alternatives

This subsection deals with the behavior of the proposed family of tests under a host of local alternatives. We start with deterministic alternatives, a framework commonly used in the literature to

study the asymptotic power of inferential procedures. Next, we consider a Bayesian framework, using a class of priors that characterize the structure of the alternatives. Because the results to follow simultaneously hold for $\hat{T}^{\text{LR}}(f)$, $\hat{T}^{\text{LH}}(f)$ and $\hat{T}^{\text{BNP}}(f)$, the unifying notation $\hat{T}(f)$ will be used to refer to each of the test statistics.

2.3.1. Deterministic local alternatives

Consider a sequence of BC such that, on an open subset of \mathbb{C} containing \mathcal{X} ,

$$\sqrt{n}C^T B^T \{\Theta^{-1}(\mathbb{z}, \gamma)\Sigma_p - \mathbb{z}I\}^{-1} BC \longrightarrow D(\mathbb{z}, \gamma) \quad \text{pointwise, as } n, p \rightarrow \infty. \quad (2.3)$$

Observe that $\mathbf{Y}Q_n = \sqrt{n}BC[C^T(n^{-1}XX^T)^{-1}C]^{-1/2} + \Sigma_p^{1/2}\mathbf{Z}Q_n$ and define

$$\mathcal{H}(D, f) = T^{-1/2} \left[\frac{-1}{2\pi i} \oint_{\mathcal{C}} f(\mathbb{z}) D(\mathbb{z}, \gamma) d\mathbb{z} \right] T^{-1/2} \quad \text{where} \quad (2.4)$$

$$T = \lim_{n \rightarrow \infty} C^T (n^{-1}XX^T)^{-1} C. \quad (2.5)$$

Note that T exists and is non-singular under **C5** and **C6**. If further $f(x) \geq 0$ for any $x \in \mathcal{X}$, $\mathcal{H}(D, f)$ is non-negative definite.

Theorem 2.3. Suppose **C1–C6** and (2.3) hold, and $\Delta(f, \gamma) > 0$. Then, as $n \rightarrow \infty$,

$$\frac{\sqrt{n}}{\Delta^{1/2}(f, \gamma)} \{\mathbf{M}(f) - \Omega(f, \gamma)I_q\} \Longrightarrow \mathbf{W} + \frac{\mathcal{H}(D, f)}{\Delta^{1/2}(f, \gamma)}.$$

Denote the power functions of $\hat{T}(f)$ at asymptotic level α , conditional on BC , by

$$\Upsilon(BC, f) = \mathbb{P}(\hat{T}(f) > \xi_\alpha \mid BC).$$

The asymptotic behavior of the power functions is described in the following corollary.

Corollary 2.2. Under the assumptions of Theorem 2.3, as $n \rightarrow \infty$,

$$\Upsilon(BC, f) \longrightarrow \Phi\left(-\xi_\alpha + \frac{\text{tr}(\mathcal{H}(D, f))}{q^{1/2}\Delta^{1/2}(f, \gamma)}\right),$$

where Φ is the standard normal CDF.

Remark 2.3. Corollary 2.2 indicates the three proposed statistics have identical asymptotic powers under the assumed local alternatives. This is because the first-order Taylor expansions of x , $\log(1+x)$ and $x/(1+x)$ coincide at 0. However, the respective empirical powers may differ considerably for moderate sample sizes.

The following remark provides a sufficient condition under which (2.3) is satisfied. Denoting the columns of BC by $[\mu_1, \dots, \mu_q]$, it follows that

$$\sqrt{n}C^T B^T \{\Theta^{-1}(\mathbb{Z}, \gamma)\Sigma_p - \mathbb{Z}I\}^{-1} BC = \sqrt{n}[\mu_i^T \{\Theta^{-1}(\mathbb{Z}, \gamma)\Sigma_p - \mathbb{Z}I\}^{-1} \mu_j]_{i,j=1}^q.$$

Remark 2.4. (a) Let $\mathbf{E}_{m,p}$ denote the eigen-projection associated with $\lambda_{m,p} = \lambda_m(\Sigma_p)$. Suppose that there exists a sequence (in p) of mappings $[\mathfrak{B}_{ij;p}]_{i,j=1}^q$ from $[0, \infty)^{q^2}$ to $[0, \infty)^{q^2}$, satisfying $\mathfrak{B}_{ij;p}(\lambda_{m,p}) = \sqrt{n}p\mu_i^T \mathbf{E}_{m,p} \mu_j$, $m = 1, \dots, p$, and a mapping $[\mathfrak{B}_{ij;\infty}]_{i,j=1}^q$ continuous on $[0, \infty)^{q^2}$ such that, as $p \rightarrow \infty$ and for $1 \leq i, j \leq q$,

$$\int |\mathfrak{B}_{ij;p}(x) - \mathfrak{B}_{ij;\infty}(x)| dF^{\Sigma_p}(x) \rightarrow 0.$$

Then, under **C4**, it follows that (2.3) holds with $D(\mathbb{Z}, \gamma) = [d_{ij}(\mathbb{Z}, \gamma)]_{i,j=1}^q$ and

$$d_{ij}(\mathbb{Z}, \gamma) = \int \frac{\mathfrak{B}_{ij;\infty}(x) dL^{\Sigma}(x)}{x\Theta^{-1}(\mathbb{Z}, \gamma) - \mathbb{Z}} = \int \frac{\mathfrak{B}_{ij;\infty}(x) dL^{\Sigma}(x)}{x\{1 - \gamma - \gamma\mathbb{Z}m(\mathbb{Z})\} - \mathbb{Z}}.$$

(b) If $\Sigma_p = I_p$, then (2.3) is satisfied if $\sqrt{n}\mu_i^T \mu_j \rightarrow \mathcal{K}_{ij}$, for some constants \mathcal{K}_{ij} , $1 \leq i, j \leq q$. In this case, $D(\mathbb{Z}, \gamma) = (\Theta^{-1}(\mathbb{Z}, \gamma) - \mathbb{Z})^{-1}[\mathcal{K}_{ij}]_{i,j=1}^q$.

2.3.2. Probabilistic local alternatives

While deterministic local alternatives provide useful information, they are somewhat restrictive for the purpose of a systematic investigation of the power characteristics. Therefore, probabilistic alternatives are considered in the form of a sequence of prior distributions for BC . This has the added advantage of providing flexibility for incorporating structural information about the regression parameters and the constraints matrices. The proposed formulation of probabilistic alternatives can be seen as an extension of the proposal adopted by Li et al. [22] in the context of two-sample tests for equality of means. One challenge associated with formulating meaningful alternatives to the hypothesis (1.2), when compared to the two-sample testing problem, is that there are many more plausible ways in which the null hypothesis can be violated. Considering this, we propose a class of alternatives, that on one hand can incorporate a multitude of structures of the parameter BC , while on the other hand retains analytical tractability in terms of providing interpretable expressions for the local asymptotic power.

Assume the following prior model of BC with separable covariance

$$BC = n^{-1/4} p^{-1/2} \mathcal{R} \mathcal{V} \mathcal{S}^T, \quad (2.6)$$

where \mathcal{V} is a $p \times m$ stochastic matrix ($m \geq 1$ fixed) with independent elements v_{ij} such that $\mathbb{E}[v_{ij}] = 0$, $\mathbb{E}[|v_{ij}|^2] = 1$ and $\max_{ij} \mathbb{E}[|v_{ij}|^4] \leq p^{c_v}$ for some $c_v \in (0, 1)$; \mathcal{R} is a $p \times p$ deterministic matrix and \mathcal{S} is a fixed $q \times m$ matrix. Moreover, let $\|\mathcal{R}\|_2 \leq \mathcal{K}_1 < \infty$ and suppose there is a nonrandom function $h(\mathbb{Z}, \gamma)$ such that, as $p \rightarrow \infty$, on an open subset of \mathbb{C} containing \mathcal{X} ,

$$p^{-1} \text{tr}\{(\Theta^{-1}(\mathbb{Z}, \gamma)\Sigma_p - \mathbb{Z}I)^{-1} \mathcal{R} \mathcal{R}^T\} \rightarrow h(\mathbb{Z}, \gamma) \quad \text{pointwise.} \quad (2.7)$$

Recalling that $(\Theta^{-1}(\mathbb{Z}, \gamma)\Sigma_p - \mathbb{Z}I)^{-1}$ is the deterministic equivalent of the resolvent $(\widehat{\Sigma}_p - \mathbb{Z}I)^{-1}$, existence of the limit (2.7) also implies that $p^{-1} \text{tr}\{(\widehat{\Sigma}_p - \mathbb{Z}I)^{-1}\mathcal{R}\mathcal{R}^T\}$ converges pointwise in probability to $h(\mathbb{Z}, \gamma)$. Notice also that $p^{-1} \text{tr}\{(\widehat{\Sigma}_p - \mathbb{Z}I)^{-1}\mathcal{R}\mathcal{R}^T\}$ is the Stieltjes transform of a measure supported on the eigenvalues of $\widehat{\Sigma}_p$.

Model (2.6) leads to a fairly broad covariance design for multi-dimensional random elements, encompassing structures commonly encountered in many application domains, especially in spatio-temporal statistics. We give some representative examples by considering various functional forms of the matrix \mathcal{S} . Denote by μ_j the columns of BC and by V_j the columns of \mathcal{V} .

Example 2.1. In all that follows j takes values in $1, \dots, q$.

- (a) *Independent:* $\mu_j = n^{-1/4} p^{-1/2} \mathcal{R} V_j$.
- (b) *Longitudinal:* $\mu_j = n^{-1/4} p^{-1/2} \mathcal{R} (V_1 + V_2 j + \dots + V_m j^{m-1})$.
- (c) *Moving average:* $\mu_j = n^{-1/4} p^{-1/2} \mathcal{R} [V_{j+t} + \theta_1 V_{j+t-1} + \dots + \theta_t V_j]$ for constants $\theta_1, \dots, \theta_t$.

Taking the MANOVA problem to illustrate, suppose that the columns of B represent group mean vectors, and suppose C is the matrix that determines successive contrasts among them. Then, μ_j is the difference between the means of group j and group $j + 1$. Parts (a)–(c) of Example 2.1 correspond then to μ_1, \dots, μ_q respectively following an independent, a longitudinal and a moving average process. The row-wise covariance structure is assumed to be such that each μ_j has a covariance matrix proportional to $n^{-1/2} p^{-1} \mathcal{R} \mathcal{R}^T$. The factor $n^{-1/2} p^{-1}$ provides the scaling for the tests to have non-trivial local power.

A sufficient condition that leads to (2.7), similar to Remark 2.4, is to postulate the existence of functions \mathfrak{B}_p satisfying $\mathfrak{B}_p(\lambda_{j,p}) = \text{tr}\{\mathbf{E}_{j,p} \mathcal{R} \mathcal{R}^T\}$, $j = 1, \dots, p$, and

$$\int |\tilde{\mathfrak{B}}_p(x) - \tilde{\mathfrak{B}}_\infty(x)| dF^{\Sigma_p}(x) \rightarrow 0$$

for some function $\tilde{\mathfrak{B}}_\infty$ continuous on $[0, \infty)$, where $\lambda_{j,p}$ is the j th eigenvalue of Σ_p and $\mathbf{E}_{j,p}$ is the eigen-projection associated with $\lambda_{j,p}$. Then

$$h(\mathbb{Z}, \gamma) = \int \frac{\tilde{\mathfrak{B}}_\infty(x) dL^\Sigma(x)}{x\{1 - \gamma - \gamma \mathbb{Z} m(\mathbb{Z})\} - \mathbb{Z}}. \quad (2.8)$$

Equations (2.7) and (2.8) indicate that $h(\mathbb{Z}, \gamma)$ effectively captures the distribution of the total spectral mass of $\mathcal{R} \mathcal{R}^T$ across the spectral coordinates of Σ_p , also taking into account the dimensionality effect through the aspect ratio γ . Later, we shall discuss specific classes of the matrices \mathcal{R} that lead to analytically tractable expressions for $h(\mathbb{Z}, \gamma)$, with the structure of \mathcal{R} linking the parameter BC under the alternative through (2.6) to the structure of Σ_p .

Another important feature of the probabilistic model is that it incorporates both dense and sparse alternatives through different specifications of the innovation variables v_{ij} . We consider two special cases.

1. *Dense alternative:* $v_{ij} \sim \mathcal{N}(0, 1)$.
2. *Sparse alternative:* $v_{ij} \sim G_\eta$, for some $\eta \in (0, 1)$, where G_η is the discrete probability distribution assigning mass $1 - p^{-\eta}$ to 0 and mass $(1/2)p^{-\eta}$ to the points $\pm p^{\eta/2}$.

Note that the usual notion of sparsity corresponds to the setting where in addition, $\mathcal{R} = I_p$. More generally, the second specification above formulates a prior model for BC that is sparse in the coordinate system determined by \mathcal{R} . In particular, if $\mathcal{R}\mathcal{R}^T$ is a polynomial in Σ_p (see Section 3.2 for a discussion), BC can be seen as sparse in the spectral coordinates of Σ_p .

Theorem 2.4. *Suppose that C1–C6 hold and $\Delta(f, \gamma) > 0$. Also suppose that, under H_a , BC has a prior distribution given by (2.6). Then, the power function of each of the three test statistics satisfies*

$$\Upsilon(BC, f) \xrightarrow{L_1} \Phi\left(-\xi_\alpha + \frac{\text{tr}(\mathcal{S}\mathcal{S}^T T^{-1})}{q^{1/2} \Delta^{1/2}(f, \gamma)} \oint_{\mathcal{C}} \frac{-1}{2\pi i} f(\mathbb{z}) h(\mathbb{z}, \gamma) d\mathbb{z}\right), \quad (2.9)$$

as $n \rightarrow \infty$, where T is as in (2.5) and $\xrightarrow{L_1}$ indicates L_1 -convergence (with respect to the prior measure of BC).

Remark 2.5. Even if the quantity $h_p(\mathbb{z}, \gamma) = p^{-1} \text{tr}\{(\Theta^{-1}(\mathbb{z}, \gamma) \Sigma_p - \mathbb{z}I)^{-1} \mathcal{R}\mathcal{R}^T\}$ does not converge, it can be verified that the difference between the left- and right-hand sides of (2.9) still converges to zero in L_1 if $h(\mathbb{z}, \gamma)$ is replaced by $h_p(\mathbb{z}, \gamma)$.

Observe that the matrices \mathcal{R} and \mathcal{S} decouple in the expression (2.9) for the asymptotic power. Dependence on the unknown error covariance matrix Σ_p , besides $\Delta^{1/2}(f, \gamma)$, is only through the function $h(\mathbb{z}, \gamma)$, which incorporates the structure of the matrix $\mathcal{R}\mathcal{R}^T$. It is also noticeable that distributional characteristics of the variables v_{ij} do not affect the asymptotic power. Indeed, the proposed tests have the same local asymptotic power under both sparse and dense alternatives.

3. Data-driven selection of shrinkage

In this section, we introduce a data-driven procedure to select the “optimal” f from a parametric family \mathfrak{F} of shrinkage functions. The strategy is to maximize the local power function $\Upsilon(BC, f)$ over f , given a class of probabilistic local alternatives as in (2.6). In designing the classes of alternatives, we focus our attention only on the specification of \mathcal{R} . This is because, as the expression (2.9) shows, the dependence on the matrix \mathcal{S} is only through a multiplier involving a “known” matrix T , while the effect of the unknown covariance Σ_p (and its interaction with \mathcal{R}) manifests itself through the function $h(\mathbb{z}, \gamma)$. Another reason for focusing on \mathcal{R} is that the choice of \mathcal{S} is closely related to the specific type of linear model being considered, while the choice of \mathcal{R} is associated with the structure of the error distribution.

We present some settings of BC for which $h(\mathbb{z}, \gamma)$ can be computed explicitly. We also verify that the standardized test statistic with the data-driven selection of f is still asymptotically standard normal under suitable conditions. Hence, the Type 1 error rate of the tests is asymptotically not inflated, although the same data is used for both shrinkage selection and testing. Lastly, we present a composite test procedure that combines the optimal tests corresponding to different prior models of BC and thereby improves adaptivity to various kinds of alternatives.

3.1. Shrinkage family

Suppose the family of shrinkage functions is such that

$$\mathfrak{F} = \{f_\ell : \ell \in \mathcal{L}\},$$

- (i) \mathcal{L} is a compact subset of \mathbb{R}^r , $r \in \mathbb{N}^+$.
- (ii) There is a closed, connected subset \mathcal{Z} of \mathbb{C} such that $\mathcal{X} = [0, \limsup_p \lambda_{\max}(\Sigma_p) \times (1 + \sqrt{\gamma})^2] \subset \mathcal{Z}$, and the third-order partial derivatives of f_ℓ with respect to ℓ are continuous on $\mathcal{L} \otimes \mathcal{Z}$.
- (iii) The gradient $\nabla_\ell f_\ell$ and the Hessian $\nabla_\ell^2 f_\ell$ of f_ℓ with respect to ℓ have analytic extensions to \mathcal{Z} for all $\ell \in \mathcal{L}$.
- (iv) $\inf_{\ell \in \mathcal{L}} \Delta(f_\ell, \gamma) > 0$.

Under the probabilistic prior model (2.6) with $h(\mathbb{z}, \gamma)$ in (2.7) given, define

$$\Xi(\ell, h, \gamma) = \frac{-1}{2\pi i \Delta^{1/2}(f_\ell, \gamma)} \oint_{\mathcal{C}} f_\ell(\mathbb{z}) h(\mathbb{z}, \gamma) d\mathbb{z}.$$

Theorem 2.4 suggests that ℓ should be chosen such that $\Xi(\ell, h, \gamma)$ is maximized, that is,

$$\ell_{\text{opt}} = \arg \max_{\ell \in \mathcal{L}} \Xi(\ell, h, \gamma).$$

The test with the selected shrinkage will then be the locally most powerful test under the alternatives specified by (2.6) and (2.7) for any given choice of \mathcal{S} . Since $\Xi(\ell, h, \gamma)$ is continuous with respect to ℓ under condition (i)–(iv), ℓ_{opt} exists. Importantly, $\Xi(\ell, h, \gamma)$ does not rely on \mathcal{S} . In other words, different column-wise covariance structures of BC are uniform in terms of selecting the optimal shrinkage. This significantly simplifies the selection procedure.

Recall that $h(\mathbb{z}, \ell)$ is the limit of $p^{-1} \text{tr}\{(\Theta^{-1}(\mathbb{z}, \gamma) \Sigma_p - \mathbb{z}I)^{-1} \mathcal{R} \mathcal{R}^T\}$. We next present two possible settings of $\mathcal{R} \mathcal{R}^T$ under which $h(\mathbb{z}, \gamma)$ and consequently $\Xi(\ell, h, \gamma)$ can be accurately estimated:

(1) Suppose $\mathcal{R} \mathcal{R}^T$ is specified. Then, $h(\mathbb{z}, \gamma)$ is estimated by $\widehat{h}(\mathbb{z}, \gamma_n) = p^{-1} \text{tr}\{(\widehat{\Sigma}_p - \mathbb{z}I)^{-1} \mathcal{R} \mathcal{R}^T\}$ and

$$\widehat{\Xi}(\ell, \widehat{h}, \gamma_n) := \frac{-1}{2\pi i \widehat{\Delta}^{1/2}(f_\ell, \gamma_n)} \oint_{\mathcal{C}} f_\ell(\mathbb{z}) \widehat{h}(\mathbb{z}, \gamma_n) d\mathbb{z}$$

is a consistent estimator of $\Xi(f, h, \gamma)$. As an example of this scenario, assume that the p components of μ_j admit a natural ordering such that the dependence between their coordinates is a function of the difference between their indexes. Then we may set $\mathcal{R} \mathcal{R}^T$ to be a Toeplitz matrix (stationary auto-covariance structure).

(2) Only the spectral mass distribution of $\mathcal{R} \mathcal{R}^T$ in the form of $\tilde{\mathfrak{B}}_\infty$ described in (2.8) is specified.

The remainder of this section is devoted to dealing with the second scenario.

3.2. Polynomial alternatives

Even if $\tilde{\mathfrak{B}}_\infty$ is given, the estimation of $h(\mathbb{Z}, \gamma)$ is still challenging since it involves the unknown limiting spectral distribution L^Σ . In order to estimate $h(\mathbb{Z}, \gamma)$, it is convenient to have it in a closed form. It is feasible if $\tilde{\mathfrak{B}}_\infty$ is a polynomial, which is true if $\mathcal{R}\mathcal{R}^T$ is a matrix polynomial in Σ_p . Since any smooth function can be approximated by polynomials, this formulation is quite flexible and practically beneficial. Assume therefore that

$$\mathcal{R}\mathcal{R}^T = \sum_{j=0}^s t_j \Sigma_p^j, \quad (3.1)$$

where t_0, \dots, t_s are pre-specified weights such that $\sum_{j=0}^s t_j \Sigma_p^j$ is nonnegative definite. Under the model,

$$h(\mathbb{Z}, \gamma) = \lim_{p \rightarrow \infty} p^{-1} \operatorname{tr} \left[\left(\Theta^{-1}(\mathbb{Z}, \gamma) \Sigma_p - \mathbb{Z} I \right)^{-1} \sum_{j=0}^s t_j \Sigma_p^j \right] = \sum_{j=0}^s t_j \rho_j(\mathbb{Z}, \gamma),$$

where the functions $\rho_j(\mathbb{Z}, \gamma)$ satisfy the recursive formula (see [19])

$$\rho_0(\mathbb{Z}, \gamma) = m(\mathbb{Z}), \quad \rho_{j+1}(\mathbb{Z}, \gamma) = \Theta(\mathbb{Z}, \gamma) \left[\int x^j dL^\Sigma(x) + \mathbb{Z} \rho_j(\mathbb{Z}, \gamma) \right].$$

For any $j \in \mathbb{N}$, $\int x^j dL^\Sigma(x)$, and consequently $\rho_j(\mathbb{Z}, \gamma)$, can be estimated consistently [2], Lemma 1. Specifically, $p^{-1} \operatorname{tr}(\widehat{\Sigma}_p)$ is a consistent estimator of $\int x dL^\Sigma(x)$.

In practice, we restrict to the case $s = 2$. There are several considerations that guided this choice of s as stated in [22]. First, for $s = 2$, all quantities involved can be computed explicitly without requiring knowledge of higher-order moments of the observations. Also, the corresponding estimating equations for $h(\mathbb{Z}, \gamma)$ are more stable as they do not involve higher-order spectral moments. Second, the choice of $s = 2$ yields a significant, yet nontrivial, concentration of the prior covariance of μ_j , $j = 1, \dots, q$, (that is $\mathcal{R}\mathcal{R}^T$ up to a scaling factor) in the directions of the leading eigenvectors of Σ_p . Finally, the choice $s = 2$ allows for both convex and concave shapes of the spectral mass distribution $\tilde{\mathfrak{B}}_\infty$ since the latter becomes a quadratic function.

With $s = 2$, we estimate $\rho_0(\mathbb{Z}, \gamma)$, $\rho_1(\mathbb{Z}, \gamma)$, $\rho_2(\mathbb{Z}, \gamma)$, and $h(\mathbb{Z}, \gamma)$ by

$$\begin{aligned} \widehat{\rho}_0(\mathbb{Z}, \gamma_n) &= m_{n,p}(\mathbb{Z}), \\ \widehat{\rho}_1(\mathbb{Z}, \gamma_n) &= \widehat{\Theta}(\mathbb{Z}, \gamma_n) [1 + \mathbb{Z} m_{n,p}(\mathbb{Z})], \\ \widehat{\rho}_2(\mathbb{Z}, \gamma_n) &= \widehat{\Theta}(\mathbb{Z}, \gamma_n) [p^{-1} \operatorname{tr}(\widehat{\Sigma}_p) + \mathbb{Z} \widehat{\rho}_1(\mathbb{Z}, \gamma_n)], \\ \widehat{h}(\mathbb{Z}, \gamma_n) &= \sum_{j=0}^2 t_j \widehat{\rho}_j(\mathbb{Z}, \gamma_n). \end{aligned} \quad (3.2)$$

The algorithm for the data-driven shrinkage selection is stated next.

Algorithm 3.1 (Data-driven shrinkage selection).

1. Specify prior weights $\tilde{t} = (t_0, t_1, t_2)$. The canonical choices are $(1, 0, 0)$, $(0, 1, 0)$, $(0, 0, 1)$.
2. Compute $\hat{h}(\mathbb{Z}, \gamma_n) = \sum_{j=0}^2 t_j \hat{\rho}_j(\mathbb{Z}, \gamma_n)$.
3. For any $\ell \in \mathcal{L}$, numerically compute the integral

$$\widehat{\Xi}(\ell, \hat{h}, \gamma_n) = \frac{-1}{2\pi i \widehat{\Delta}^{1/2}(f_\ell, \gamma_n)} \oint_{\mathcal{C}} f_\ell(\mathbb{Z}) \hat{h}(\mathbb{Z}, \gamma_n) d\mathbb{Z}.$$

4. Select $\ell_{\text{opt}}(\tilde{t}) = \arg \max_{\ell \in \mathcal{L}} \widehat{\Xi}(\ell, \hat{h}, \gamma_n)$.

The behavior of the tests applied with the data-driven shrinkage selection is described in the following theorem.

Theorem 3.1. *Suppose C1–C6 hold and \mathfrak{F} satisfies conditions (i)–(iv). Then,*

- (1) $\sup_{\ell \in \mathcal{L}} \sqrt{n} |\widehat{\Xi}(\ell, \hat{h}, \gamma_n) - \Xi(\ell, h, \gamma)| \xrightarrow{P} 0$ as $n \rightarrow \infty$.
- (2) Let ℓ^* be any local maximizer of $\Xi(\ell, h, \gamma)$ in the interior of \mathcal{L} . Assume there exists a neighborhood of ℓ^* such that for all feasible points $\ell \in \mathcal{L}$ within the neighborhood, there exists a constant $\mathcal{K} > 0$ such that

$$\Xi(\ell, h, \gamma) - \Xi(\ell^*, h, \gamma) \leq -\mathcal{K} \|\ell - \ell^*\|_2^2. \quad (3.3)$$

Then, there exists a sequence $(\ell_n^* : n \in \mathbb{N})$ of local maximizers of $(\widehat{\Xi}(\ell, \hat{h}, \gamma_n) : n \in \mathbb{N})$ satisfying

$$n^{1/4} \|\ell_n^* - \ell^*\|_2 = O_p(1) \quad (n \rightarrow \infty). \quad (3.4)$$

Further, recalling notation in Section 2, under the null hypothesis,

$$\frac{\sqrt{n}}{\widehat{\Delta}^{1/2}(f_{\ell_n^*}, \gamma_n)} \{ \mathbf{M}(f_{\ell_n^*}) - \widehat{\Omega}(f_{\ell_n^*}, \gamma_n) I_q \} \Longrightarrow \mathbf{W}. \quad (3.5)$$

- (3) Let ℓ^* be any local maximizer of $\Xi(\ell, h, \gamma)$ on the boundary of \mathcal{L} . Assume there exists a neighborhood of ℓ^* such that for all feasible points $\ell \in \mathcal{L}$ within the neighborhood, there is a constant $\mathcal{K}' > 0$ satisfying

$$\Xi(\ell, h, \gamma) - \Xi(\ell^*, h, \gamma) \leq -\mathcal{K}' \|\ell - \ell^*\|_2. \quad (3.6)$$

Then, (3.4) and (3.5) still hold.

The two conditions (3.3) and (3.6) ensure that the parameter ℓ^* is locally identifiable in a neighborhood of ℓ^* . In general, the two conditions depend on the structure of L^Σ .

3.3. Combination of prior models

An extensive simulation analysis revealed that there is considerable variation in the shape of the power functions and the values of $\tilde{t} = (t_0, t_1, t_2)$, especially when the condition number of Σ_p is relatively large. In this subsection, we consider a convenient collection of priors that are representative of certain structural scenarios. A composite test, called \hat{T}_{\max} , is defined as the maximum of the standardized statistics $\hat{T}(f_{\ell_i^*})$ where ℓ_i^* is obtained from Algorithm 3.1 under prior \tilde{t}_i , $i = 1, \dots, m$. The following strategy is applicable to LR, LH and BNP. We therefore continue to use $\hat{T}(f)$ to denote the general test statistic. In summary, we propose to test the hypothesis by rejecting for large values of the statistic

$$\hat{T}_{\max} = \max_{\tilde{t} \in \tilde{\Pi}} \hat{T}(f_{\ell_i^*}),$$

where $\tilde{\Pi} = \{\tilde{t}_1, \dots, \tilde{t}_m\}$, $m \geq 1$, is a pre-specified finite class of weights. A simple but effective choice of $\tilde{\Pi}$ consists of the three canonical weights $\tilde{t}_1 = (1, 0, 0)$, $\tilde{t}_2 = (0, 1, 0)$, $\tilde{t}_3 = (0, 0, 1)$.

Theorem 3.2. Suppose C1–C6 hold and \mathfrak{F} satisfies condition (i)–(iv). For each $i = 1, \dots, m$, assume that ℓ_{in}^* is a sequence of local maximizers of the empirical power function $\hat{\Xi}(\ell, \hat{h}, \gamma_n)$ under prior model with weight \tilde{t}_i such that

$$n^{1/4} \|\ell_{in}^* - \ell_i^*\|_2 = O_p(1).$$

(See (3.4).) Then, under the null hypothesis $H_0 : BC = 0$,

$$(\hat{T}(f_{\ell_{1n}^*}), \dots, \hat{T}(f_{\ell_{mn}^*})) \implies \mathcal{N}(0, \Delta_*),$$

where Δ_* is an $m \times m$ matrix with diagonal entries 1 and (i, j) -th off-diagonal entry

$$\Delta^{-1/2}(f_{\ell_i^*}, \gamma) \Delta(f_{\ell_i^*}, f_{\ell_j^*}, \gamma) \Delta^{-1/2}(f_{\ell_j^*}, \gamma).$$

Theorem 3.2 shows that \hat{T}_{\max} has a non-degenerate limiting distribution under H_0 . It is worth mentioning that LR, LH and BNP share the covariance matrix Δ_* . Theorem 3.2 can be used to determine the cut-off values of the test by deriving analytical formulas for the quantiles of the limiting distribution. Aiming to avoid complex calculations, a parametric bootstrap procedure is applied to approximate the cut-off values. Specifically, Δ_* is first estimated by $\hat{\Delta}_*$, and then bootstrap replicates are generated by simulating from $\mathcal{N}(0, \hat{\Delta}_*)$, thereby providing an approximation of the null distribution of \hat{T}_{\max} . Replacing $\Delta(f_{\ell_i^*}, f_{\ell_j^*}, \gamma)$ with $\hat{\Delta}(f_{\ell_i^*}, f_{\ell_j^*}, \gamma_n)$ yields the natural estimator.

Remark 3.1. Observe that $\hat{\Delta}_*$ defined above may not be nonnegative definite even though it is symmetric. If such a case occurs, the resulting estimator can be projected onto its closest non-negative definite matrix simply by setting the negative eigenvalues to zero. This covariance matrix estimator is denoted by $\hat{\Delta}_*^+$ and it is used for generating the bootstraps samples.

4. Ridge and higher-order regularizers

4.1. Ridge regularization

One of the most commonly used shrinkage procedures in statistics is ridge regularization, corresponding to choosing $f_\ell(x) = 1/(x - \ell)$, $\ell < 0$, so that $f_\ell(\widehat{\Sigma}_p) = (\widehat{\Sigma}_p - \ell I_p)^{-1}$. It is an effective way to shift $\widehat{\Sigma}_p$ away from singularity by adding a ridge term $-\ell I_p$. In this subsection, we apply the results of Sections 2 and 3 using the ridge-shrinkage family

$$\mathfrak{F}_{\text{ridge}} := \{f_\ell(x) = (x - \ell)^{-1}, \ell \in [\underline{\ell}, \bar{\ell}]\}, \quad -\infty < \underline{\ell} < \bar{\ell} < 0.$$

In the literature, ridge-regularization was applied to high-dimensional one- and two-sample mean tests in [13] and [22]. Hence, this subsection is a generalization of their methods to general linear hypotheses.

From the aspect of population covariance estimation, ridge-regularization can be viewed as an order-one estimation where Σ_p is estimated by a weighted average of $\widehat{\Sigma}_p$ and I_p , namely $\alpha_0 I_p + \alpha_1 \widehat{\Sigma}_p$. The estimator is equivalent to ridge-regularization with $\ell = -\alpha_0/\alpha_1$ for testing purposes. Within a restricted region of (α_1, α_2) , the large eigenvalues of $\widehat{\Sigma}_p$ are shrunk down and the small ones are lifted upward. It is a desired property since in high-dimensional settings, large sample eigenvalues are systematically biased upward and small sample eigenvalues downwards.

An important advantage of ridge regularization is that the test procedure is computationally efficient due to the fact that $\Omega(f_\ell, \gamma)$ and $\Delta(f_\ell, \gamma)$ admit closed forms as shown in Lemma 2.1. These quantities can be estimated by $\widehat{\Omega}_\ell(\gamma_n) = \widehat{\Theta}(\ell, \gamma_n) - 1$ and $\widehat{\Delta}_\ell(\gamma_n) = 2\widehat{\delta}(\ell, \ell, \gamma_n)$, respectively. A closed-form estimator $\widehat{\Xi}_\ell(\widehat{h}, \gamma_n)$ is then also available for $\Xi(\ell, h, \gamma)$. This leads to the following algorithm.

Algorithm 4.1 (Ridge-regularized test procedure).

1. Specify prior weights $\tilde{t} = (t_0, t_1, t_2)$.
2. With $m_{n,p}(\ell) = p^{-1} \text{tr}(\widehat{\Sigma}_p - \ell I_p)^{-1}$, compute, for any $\ell \in [\underline{\ell}, \bar{\ell}]$,

$$\begin{aligned} \widehat{\Theta}(\ell, \gamma_n) &= \{1 - \gamma_n - \gamma_n \ell m_{n,p}(\ell)\}^{-1}, \\ \widehat{\Omega}_\ell(\gamma_n) &= \widehat{\Theta}(\ell, \gamma_n) - 1, \\ \widehat{\Delta}_\ell(\gamma_n) &= 2\gamma_n \{1 + \ell m_{n,p}(\ell)\} \widehat{\Theta}^3(\ell, \gamma_n) + 2\gamma_n \ell \{m_{n,p}(\ell) + \ell m'_{n,p}(\ell)\} \widehat{\Theta}^4(\ell, \gamma_n). \end{aligned}$$

3. For any $\ell \in [\underline{\ell}, \bar{\ell}]$, compute $\widehat{h}(\ell, \gamma_n) = \sum_{j=0}^2 t_j \widehat{\rho}_j(\ell, \gamma_n)$ as defined in (3.2) and

$$\widehat{\Xi}_\ell(\widehat{h}, \gamma_n) = \frac{\widehat{h}(\ell, \gamma_n)}{\widehat{\Delta}_\ell^{1/2}(\gamma_n)}.$$

4. Select $\ell^* = \arg \max_{\ell \in [\underline{\ell}, \bar{\ell}]} \widehat{\Xi}_\ell(\widehat{h}, \gamma_n)$.

5. Use one of the standardized statistics

$$\begin{aligned}\widehat{T}^{\text{LR}}(\ell^*) &:= \frac{\sqrt{n}\{1 + \widehat{\Omega}_{\ell^*}(\gamma_n)\}}{q^{1/2}\widehat{\Delta}_{\ell^*}^{1/2}(\gamma_n)} [T^{\text{LR}}(\ell^*) - q \log\{1 + \widehat{\Omega}_{\ell^*}(\gamma_n)\}], \\ \widehat{T}^{\text{LH}}(\ell^*) &:= \frac{\sqrt{n}}{q^{1/2}\widehat{\Delta}_{\ell^*}^{1/2}(\gamma_n)} [T^{\text{LH}}(\ell^*) - q\widehat{\Omega}_{\ell^*}(\gamma_n)], \\ \widehat{T}^{\text{BNP}}(\ell^*) &:= \frac{\sqrt{n}\{1 + \widehat{\Omega}_{\ell^*}(\gamma_n)\}^2}{q^{1/2}\widehat{\Delta}_{\ell^*}^{1/2}(\gamma_n)} \left[T^{\text{BNP}}(\ell^*) - \frac{q\widehat{\Omega}_{\ell^*}(\gamma_n)}{1 + \widehat{\Omega}_{\ell^*}(\gamma_n)} \right],\end{aligned}$$

where

$$T^{\text{LR}}(\ell^*) = \sum_{i=1}^q \log(1 + \lambda_i), \quad T^{\text{LH}}(\ell^*) = \sum_{i=1}^q \lambda_i, \quad T^{\text{BNP}}(\ell^*) = \sum_{i=1}^q \frac{\lambda_i}{1 + \lambda_i},$$

and $\lambda_1, \dots, \lambda_q$ are the eigenvalues of $n^{-1}Q_n^T \mathbf{Y}^T (\widehat{\Sigma}_p - \ell^* I_p)^{-1} \mathbf{Y} Q_n$. Reject the null at asymptotic level α if the test statistic value exceeds ξ_α .

Although in theory any negative ℓ^* is allowed in the test procedure, in practice, meaningful lower and upper bounds $\underline{\ell}$ and $\bar{\ell}$ are needed to ensure stability of the test statistics when $p \approx n$ or $p > n$ and also to carry out the search for optimal ℓ at a low computational cost. In our simulation settings, we use $\bar{\ell} = -p^{-1} \text{tr}(\widehat{\Sigma}_p)/100$ and $\underline{\ell} = -20\lambda_{\max}(\widehat{\Sigma}_p)$, which generally lead to quite robust performance.

The composite test procedure with ridge-regularization is summarized below.

Algorithm 4.2 (Composite ridge-regularized test procedure).

1. Select prior weights $\tilde{\Pi} = (\tilde{t}_1, \dots, \tilde{t}_m)$. The canonical choice is $((1, 0, 0), (0, 1, 0), (0, 0, 1))$.
2. For each \tilde{t}_j in $\tilde{\Pi}$, run Algorithm 4.1, get the standardized test statistic $\widehat{T}(\ell_j^*)$ and compute $\widehat{T}_{\max} = \max_{1 \leq j \leq m} \widehat{T}(\ell_j^*)$.
3. With the selected tuning parameters $(\ell_1^*, \ell_2^*, \ell_3^*)$ compute the matrix $\widehat{\Delta}_*$ whose diagonal elements are equal to one and whose (i, j) -th entry for $i \neq j$ is

$$\widehat{\Delta}_{\ell_i^*}^{-1/2}(\gamma_n) \widehat{\Delta}_{\ell_i^*, \ell_j^*}(\gamma_n) \widehat{\Delta}_{\ell_j^*}^{-1/2}(\gamma_n),$$

where $\widehat{\Delta}_{\ell_i^*}(\gamma_n)$ is defined in Step 2 of Algorithm 4.1 and

$$\widehat{\Delta}_{\ell_i^*, \ell_j^*}(\gamma_n) = 2\widehat{\Theta}(\ell_i^*, \gamma_n)\widehat{\Theta}(\ell_j^*, \gamma_n) \left[\frac{\ell_i^* \widehat{\Theta}(\ell_i^*, \gamma_n) - \ell_j^* \widehat{\Theta}(\ell_j^*, \gamma_n)}{\ell_i^* - \ell_j^*} - 1 \right].$$

4. Project $\widehat{\Delta}_*$ to its closest non-negative definite matrix $\widehat{\Delta}_*^+$ by setting the negative eigenvalues to zero. Generate $\varepsilon_1, \dots, \varepsilon_G$ with $\varepsilon_b = \max_{1 \leq i \leq m} Z_i^{(b)}$ with $Z^{(b)} = [Z_i^{(b)}]_{i=1}^m \sim \mathcal{N}(0, \widehat{\Delta}_*^+)$.
5. Compute the p -value as $G^{-1} \sum_{b=1}^G \mathbb{1}\{\varepsilon_b > \widehat{T}_{\max}\}$.

4.2. Extension to higher-order regularizers

Through an extensive simulation study in a MANOVA setting, it is shown in Section 5 that the ridge-regularized tests compare favorably against a host of existing test procedures. This is consistent with the findings in [22] in the two-sample mean test framework. Ridge-shrinkage rescales $\widehat{\mathbf{H}}_p$ by $(\widehat{\Sigma}_p - \ell I_p)^{-1}$ instead of $\widehat{\Sigma}_p^{-1}$. Broader classes of scaling matrices have been studied extensively (see [20] for an overview). They can be set up in the form $f(\widehat{\Sigma}_p)$. When $f(\cdot)$ is analytic, such scaling falls within the class of the proposed tests.

The flexibility provided by a larger class of scaling matrices can be useful to design test procedures for detecting a specific kind of alternative. The choice of the test procedure may for example be guided by questions such as *Which f leads to the best asymptotic power under a specific sequence of local alternatives, if H_0 is rejected based on large eigenvalues of $\mathbf{M}(f)$?* While a full characterization of this question is beyond the scope of this paper, a partial answer may be provided by restricting to functions f in the higher-order class

$$\mathfrak{F}_{\text{high}} = \left\{ f_\ell(x) = \left[\sum_{j=0}^{\kappa} l_j x^j \right]^{-1} : \ell = (l_0, \dots, l_\kappa)^T \in \mathcal{G} \right\},$$

where \mathcal{G} is such that f_ℓ is uniformly bounded and monotonically decreasing on \mathcal{X} , for any $\ell \in \mathcal{G}$. These higher-order shrinkage functions are weighted averages of ridge-type shrinkage functions. To see this, suppose the polynomial $\sum_{j=0}^{\kappa} l_j x^j$ has roots $r_1, \dots, r_{\kappa_0} \in \mathbb{C} \setminus \mathcal{X}$ with multiplicity $s_1, \dots, s_{\kappa_0} \in \mathbb{N}^+$. Via basic algebra, f_ℓ can be expressed as

$$f_\ell(x) = \left[\sum_{j=0}^{\kappa} l_j x^j \right]^{-1} = \sum_{j=1}^{\kappa_0} \sum_{i=1}^{s_{\kappa_0}} w_{ji} (x - r_j)^{-i}, \quad (4.1)$$

with some weights $w_{ji} \in \mathbb{C}$. If all roots are simple, f_ℓ is a weighted average of ridge-regularization with κ different parameters. Heuristically, it is expected that a higher order f_ℓ yields tests more robust against unfavorable selection of ridge shrinkage parameter.

The design of \mathcal{G} is not easy when κ is large. Here, we select $\kappa = 3$, which is the minimum degree that allows f_ℓ^{-1} to be both locally convex and concave. In this case, the complexity of selecting the optimal regularizer is significantly higher than for ridge-regularization. Due to space limitations, we move the design of \mathcal{G} and the test procedure when $\kappa = 3$ to Section S.1 of the Supplementary Material [21].

5. Simulations

In this section, the proposed tests are compared by means of a simulation study to two representative existing methods in the literature, [35] (ZGZ) and [11] (CX). We focus on one-way MANOVA, a set-up for which both competing methods are applicable. It is worth mentioning that CX requires a good estimator of the precision matrix Σ_p^{-1} , that is typically unavailable

when both Σ_p and Σ_p^{-1} are dense. In the simulations, the true Σ_p^{-1} is utilized for CX, thus making it an oracle procedure. In the following, LR_{ridge} , LH_{ridge} , and $\text{BNP}_{\text{ridge}}$ denote the ridge-regularized tests presented in Algorithm 4.1. LR_{high} , LH_{high} , and BNP_{high} denote the tests with higher-order shrinkage introduced in Section 4.2 with $\kappa = 3$. LR_{comp} , LH_{comp} and BNP_{comp} denote the composite ridge-regularized tests of Algorithm 4.2 with the canonical choice of $\tilde{\Pi} = ((1, 0, 0), (0, 1, 0), (0, 0, 1))$.

5.1. Settings

The observation matrix \mathbf{Y} was generated as in (1.1) with normally distributed \mathbf{Z} . Specifically, we selected $k = 3$ or 5 , and $N = 300$. For $k = 3$, the three groups had 75, 90 and 135 observations, respectively. For $k = 5$, the design was balanced with each group containing 60 observations. The dimension p was 150, 600, 3000, so that $\gamma_n = p/n \approx 0.5, 2$ and 10 . The columns of B were the k group mean vectors. Accordingly, the columns of X were the group index indicators of observation subjects. We selected C to be the successive contrast matrix of order $q = k - 1$. This is a standard one-way MANOVA setting.

Under the null, B is the zero matrix. Under the alternative, for each setting of the parameters and each replicate, B is generated using one of the following models.

(i) *Dense alternative*: The entries of B are i.i.d. $\mathcal{N}(0, c^2)$ with $c = O(n^{-1/4} p^{-1/2})$ used to tune signal strength to a non-trivial level.

(ii) *Sparse alternative*: $B = c\mathcal{R}\mathbf{V}$ with $c = O(n^{-1/4} p^{-1/2})$, where \mathcal{R} is a diagonal $p \times p$ matrix with 10% randomly and uniformly selected diagonal entries being $\sqrt{10}$ and the remaining 90% being equal to 0, and \mathbf{V} is a $p \times p$ matrix with i.i.d. standard normal entries.

The following four models for the covariance matrix $\Sigma = \Sigma_p$ were considered. All models were further scaled so that $\text{tr}(\Sigma_p) = p$.

(i) *Identity matrix (ID)*: $\Sigma = I_p$.

(ii) *Dense case* Σ_{den} : Here $\Sigma = P\Sigma_{(1)}P^T$ with a unitary matrix P randomly generated from the Haar measure and resampled for each different setting, and a diagonal matrix $\Sigma_{(1)}$ whose eigenvalues are given by $\lambda_j = (0.1 + j)^6 + 0.05p^6$, $j = 1, \dots, p$. The eigenvalues of Σ decay slowly, so that no dominating leading eigenvalue exists.

(iii) *Toeplitz case* Σ_{toep} : Here Σ is a Toeplitz matrix with the (i, j) -th element equal to $0.5^{|i-j|}$. It is a setting where Σ^{-1} is sparse but Σ is dense.

(iv) *Discrete case* Σ_{dis} : Here $\Sigma = P\Sigma_{(2)}P^T$ with P generated in the same way as in (ii), and $\Sigma_{(2)}$ is a diagonal matrix with 40% eigenvalues 1, 40% eigenvalues 3 and 20% eigenvalues 10.

All tests were conducted at significance level $\alpha = 0.05$. Empirical sizes for the various tests are shown in Tables 1 and 2. Empirical power curves versus expected signal strength $n^{1/4} p^{1/2} c$ are reported in Figures 1–3. To better compare the power of each test, curves are displayed after size adjustment where the tests utilize the size-adjusted cut-off values based on the actual null distribution computed by simulations. Counterparts of Figures 1–3 that utilize asymptotic (approximate) cut-off values are reported in Section S.12 of the Supplementary Material [21]. The difference between the two types is limited. LR, LH and BNP criteria behave similarly across

Table 1. Empirical sizes at level 5%. $\Sigma = \text{ID}$ and Σ_{den} ; $\tilde{t}_1 = (1, 0, 0)$, $\tilde{t}_2 = (0, 1, 0)$, $\tilde{t}_3 = (0, 0, 1)$

$n = 300, p =$		$\Sigma = I_p$						$\Sigma = \Sigma_{\text{den}}$					
		$k = 3$			$k = 5$			$k = 3$			$k = 5$		
		150	600	3000	150	600	3000	150	600	3000	150	600	3000
LR _{ridge}	\tilde{t}_1	5.4	5.2	5.1	5.2	5.1	5.1	4.9	4.4	4.7	4.4	3.3	4.2
	\tilde{t}_2	5.4	5.2	5.1	5.2	5.1	5.1	4.9	5.2	4.9	4.4	4.9	4.7
	\tilde{t}_3	5.3	5.2	5.1	5.2	5.1	5.1	5.8	5.9	5.1	5.3	5.2	4.9
LH _{ridge}	\tilde{t}_1	5.4	5.2	5.1	5.3	5.1	5.2	6.2	7.2	5.7	6.2	7.7	6.0
	\tilde{t}_2	5.4	5.2	5.1	5.3	5.1	5.2	6.2	5.9	5.2	6.2	5.9	5.1
	\tilde{t}_3	5.3	5.2	5.1	5.3	5.1	5.2	5.8	5.9	5.2	5.4	5.2	5.0
BNP _{ridge}	\tilde{t}_1	5.3	5.2	5.0	5.2	5.0	5.0	4.0	2.5	3.7	2.9	1.3	3.1
	\tilde{t}_2	5.4	5.2	5.0	5.2	5.0	5.0	4.0	4.7	4.6	2.9	3.9	4.4
	\tilde{t}_3	5.3	5.2	5.0	5.2	5.0	5.0	5.8	5.8	5.0	5.3	5.1	4.7
LR _{high}	\tilde{t}_1	6.5	6.3	5.3	6.5	5.3	5.5	6.0	5.8	5.1	6.5	5.9	4.5
	\tilde{t}_2	6.5	6.3	5.3	6.5	5.3	5.5	8.3	6.8	5.5	8.4	7.2	5.2
	\tilde{t}_3	6.6	6.3	5.3	6.6	5.3	5.5	6.7	6.7	5.5	6.4	7.1	5.2
LH _{high}	\tilde{t}_1	6.7	6.4	5.4	6.8	5.5	5.7	6.1	5.9	5.7	6.7	6.2	5.5
	\tilde{t}_2	6.7	6.4	5.4	6.8	5.4	5.7	8.3	6.8	5.6	8.5	7.3	5.5
	\tilde{t}_3	6.7	6.4	5.4	6.8	5.4	5.7	6.7	6.7	5.6	6.5	7.2	5.5
BNP _{high}	\tilde{t}_1	6.2	6.3	5.2	6.1	5.3	5.2	5.9	5.7	4.6	6.4	5.5	3.7
	\tilde{t}_2	6.3	6.3	5.2	6.1	5.2	5.2	8.3	6.7	5.3	8.3	7.0	4.9
	\tilde{t}_3	6.3	6.3	5.1	6.1	5.2	5.2	6.6	6.6	5.3	6.4	6.9	4.9
LR _{comp}		5.1	5.1	5.0	5.4	5.3	5.0	6.0	5.1	5.5	5.6	5.0	5.1
LH _{comp}		5.1	5.1	5.1	5.5	5.3	5.1	6.7	5.8	5.9	6.9	6.2	5.7
BNP _{comp}		5.1	5.0	5.0	5.4	5.2	5.0	5.4	4.5	5.1	4.7	4.4	4.6
ZGZ		5.6	5.7	5.2	5.6	4.8	5.2	5.9	5.5	5.4	5.4	5.4	5.3
CX (Oracle)		5.6	6.3	7.0	7.3	6.9	8.6	5.8	5.9	6.8	6.0	7.2	9.0

simulation settings, as indicated by Theorem 2.4. Therefore, only one of them is displayed in each figure for ease of visualization. More figures can be found in Section S.11 of the Supplementary Material [21]. Note that, in some of the settings, several of the power curves nearly overlap, creating an occlusion effect. Then, power curves corresponding to the composite tests are plotted as the top layer.

5.2. Summary of simulation results

Tables 1 and 2 show the empirical sizes of the proposed tests are mostly controlled under 7.5%. The slight oversize is caused by the fact that $\mathbf{M}(f)$ behaves like a quadratic form, therefore the finite sample distribution is skewed. LR and BNP tests are more conservative than LH tests

Table 2. Empirical sizes at level 5%. $\Sigma = \Sigma_{\text{dis}}$ and Σ_{toep} ; $\tilde{t}_1 = (1, 0, 0)$, $\tilde{t}_2 = (0, 1, 0)$, $\tilde{t}_3 = (0, 0, 1)$

$n = 300, p =$		$\Sigma = \Sigma_{\text{dis}}$						$\Sigma = \Sigma_{\text{toep}}$					
		$k = 3$			$k = 5$			$k = 3$			$k = 5$		
		150	600	3000	150	600	3000	150	600	3000	150	600	3000
LR _{ridge}	\tilde{t}_1	4.8	5.0	4.6	4.7	4.5	5.0	5.4	4.4	4.8	4.5	4.6	4.6
	\tilde{t}_2	5.1	5.2	4.9	5.2	4.6	5.1	5.4	4.9	4.9	4.9	4.8	5.0
	\tilde{t}_3	5.6	5.5	5.1	5.7	5.3	5.3	5.8	5.2	5.0	5.7	5.4	5.1
LH _{ridge}	\tilde{t}_1	5.8	6.0	5.2	6.6	6.3	5.6	6.4	5.3	5.2	6.2	6.3	5.3
	\tilde{t}_2	5.7	5.7	5.1	6.3	5.6	5.5	5.9	5.3	5.0	5.8	5.6	5.3
	\tilde{t}_3	5.6	5.5	5.2	5.8	5.3	5.4	5.8	5.3	5.1	5.7	5.4	5.2
BNP _{ridge}	\tilde{t}_1	3.9	4.1	4.3	3.1	3.1	4.1	4.4	3.7	4.4	3.2	3.4	3.9
	\tilde{t}_2	4.6	4.8	4.8	4.1	4.0	4.9	4.9	4.4	4.8	4.1	4.3	4.7
	\tilde{t}_3	5.5	5.5	5.0	5.7	5.2	5.1	5.8	5.2	5.0	5.6	5.4	5.1
LR _{high}	\tilde{t}_1	6.3	6.4	4.8	5.9	7.0	5.5	7.1	7.0	5.3	7.5	6.9	5.2
	\tilde{t}_2	7.9	6.5	4.8	8.3	7.1	5.5	7.6	7.2	5.3	7.8	7.0	5.2
	\tilde{t}_3	6.1	5.6	4.8	6.4	6.1	5.5	6.7	6.5	5.3	6.6	6.4	5.2
LH _{high}	\tilde{t}_1	6.6	6.5	5.0	6.2	7.2	5.7	7.2	7.2	5.5	7.7	7.0	5.5
	\tilde{t}_2	8.0	6.6	5.0	8.5	7.2	5.7	7.8	7.2	5.5	8.0	7.1	5.5
	\tilde{t}_3	6.2	5.6	5.0	6.5	6.2	5.7	6.7	6.5	5.5	6.7	6.5	5.5
BNP _{high}	\tilde{t}_1	6.1	6.3	4.7	5.6	6.8	5.3	7.1	7.0	5.2	7.2	6.8	5.1
	\tilde{t}_2	7.9	6.4	4.7	8.2	7.0	5.3	7.5	7.1	5.2	7.7	7.0	5.1
	\tilde{t}_3	6.1	5.5	4.7	6.4	6.0	5.3	6.6	6.4	5.2	6.5	6.3	5.1
LR _{comp}		6.2	5.2	5.0	5.2	5.3	5.5	5.9	5.0	5.1	5.5	4.9	4.9
LH _{comp}		7.0	5.9	5.3	6.5	6.4	6.0	6.6	5.6	5.3	6.6	5.7	5.3
BNP _{comp}		5.5	4.6	4.8	4.4	4.6	5.0	5.4	4.6	4.9	4.8	4.4	4.6
ZGZ		5.5	4.7	4.6	5.7	5.1	5.3	6.0	5.5	5.0	5.9	5.6	5.0
CX (Oracle)		5.3	5.9	6.6	6.8	7.2	8.6	5.3	6.2	6.8	6.8	7.2	8.4

because the former two calibrate the statistics by transforming eigenvalues of $\mathbf{M}(f)$. Ridge-regularized tests are slightly more conservative under higher-order shrinkage.

Note that in both simulation settings, B consists of independent entries. Therefore, $\tilde{t}_1 = (1, 0, 0)$ is considered as a correctly specified prior, while $\tilde{t}_2 = (0, 1, 0)$ and $\tilde{t}_3 = (0, 0, 1)$ are considered as moderately and severely misspecified, respectively. The composite tests combine \tilde{t}_1 , \tilde{t}_2 and \tilde{t}_3 , and are therefore considered as consistently capturing the correct prior. We shall treat the composite tests as a baseline to study the effect of prior misspecification, by comparing them to tests using a single \tilde{t} .

For each simulation configuration considered in this study, the proposed procedures are as powerful as the procedure with the best performance, except for the cases when B is sparse, p is small, and priors are severely misspecified in the proposed tests; see Figure S.11.6 in the




Figure 1. Size-adjusted power with $\Sigma = \Sigma_{\text{den}}$, $k = 5$. Rows (top to bottom): $B = \text{Dense}$ and Sparse ; Columns (left to right): $p = 150, 600, 3000$. BNP_{comp} (red, solid); ZGZ (green, solid); oracle CX (purple, solid); $\text{BNP}_{\text{ridge}}$ (black, dashed) and BNP_{high} (blue, dotted-dashed) with $\tilde{\tau} = (1, 0, 0)$.

Supplementary Material [21]. We highlight the following observations based on the simulation results.

(1) The composite tests are slightly less efficient than $\text{BNP}_{\text{ridge}}$ and BNP_{high} when the correct prior $\tilde{\tau}_1$ is used, as in Figure 1. However, as in Figure 2, when the prior is severely misspecified, the composite test is significantly more powerful. It suggests that the composite tests are robust against prior misspecification, although losing some efficiency against tests with correctly specified priors.

(2) Although ridge-shrinkage and higher-order shrinkage behave similarly under the correct prior, the latter outperforms the former when the prior is misspecified; see Figure 2. This pro-




Figure 2. Size-adjusted power with $\Sigma = \Sigma_{\text{den}}$, $k = 5$. Rows (top to bottom): $B = \text{Dense}$ and Sparse ; Columns (left to right): $p = 150, 600, 3000$. LH_{comp} (red, solid); ZGZ (green, solid); oracle CX (purple, solid); LH_{ridge} (black, dashed) and LH_{high} (blue, dotted-dashed) with $\tilde{\tau} = (0, 0, 1)$.




Figure 3. Size-adjusted power with $\Sigma = \Sigma_{\text{toep}}$, $k = 3$. Rows (top to bottom): $B = \text{Dense}$ and Sparse ; Columns (left to right): $p = 150, 600, 3000$. LR_{comp} (red, solid); ZGZ (green, solid); oracle CX (purple, solid); LR_{ridge} (black, dashed) and LR_{high} (blue, dotted-dashed) with $\tilde{t} = (0, 1, 0)$.

vides evidence for the robustness of high-order shrinkage against unfavorable ridge shrinkage parameter selection.

(3) ZGZ is a special case of the proposed test family with $f(x) = 1$ for all x , which amounts to replacing $\hat{\Sigma}_p$ with I_p . When $\Sigma_p = I_p$, ZGZ appears to be the reasonable option at least intuitively. Note, both $\mathfrak{F}_{\text{ridge}}$ and $\mathfrak{F}_{\text{high}}$ contain functions close to $f(x) = 1$. Figures for $\Sigma_p = I_p$ displayed in Section S.11 of the Supplementary Material [21] show that the proposed tests perform as well as ZGZ in that case. It may be viewed as evidence of the effectiveness of the data-driven shrinkage selection strategy detailed in Section 3.

(4) Comparing to ZGZ , when the eigenvalues of Σ_p are disperse, the proposed tests are significantly more powerful when $p = 150$ and 600 , but behave similarly as ZGZ when $p = 3000$. On the other hand, as in Figure 2, the ridge-regularized test with a severely misspecified prior \tilde{t}_3 , is close to ZGZ .

(5) CX is a test specifically designed for sparse alternatives. The procedure shows its advantage in favorable settings, especially when $p = 150$. Simulation results suggest that the proposed tests are still comparable to CX even under sparse BC and Σ_p^{-1} , as long as the prior in use is not severely misspecified. When p is large, the proposed tests are significantly better when $\Sigma_p = I_p$. Evidence may be found in Figures S.11.10, S.11.11 and S.11.12 of the Supplementary Material [21].

Appendix: Proof of Theorem 2.1

This appendix contains a proof outline of Theorem 2.1. Additional proofs of supporting lemmas and other theorems can be found in the Supplementary Material [21].

Recall $Q_n = X^T (XX^T)^{-1} C [C^T (XX^T)^{-1} C]^{-1/2}$. Introduce $Q_n = U_n V_n$ with

$$U_n = X^T (XX^T)^{-1/2} \quad \text{and} \quad V_n = (XX^T)^{-1/2} C [C^T (XX^T)^{-1} C]^{-1/2}.$$

This decomposition will aid the analysis of the correlation between $\mathbf{Y}Q_n$ and $\hat{\Sigma}_p$.

From now on, use $\Sigma_p^{T/2}$ to denote $(\Sigma_p^{1/2})^T$. Under the null hypothesis, the following representations hold:

$$\begin{aligned}\mathbf{M}(f) &= n^{-1} V_n^T U_n^T \mathbf{Z}^T \Sigma_p^{T/2} f(\widehat{\Sigma}_p) \Sigma_p^{1/2} \mathbf{Z} U_n V_n, \\ \widehat{\Sigma}_p &= n^{-1} \Sigma_p^{1/2} \mathbf{Z} (I - U_n U_n^T) \mathbf{Z}^T \Sigma_p^{T/2}.\end{aligned}$$

Observe that the joint asymptotic normality of entries in $\sqrt{n}\mathbf{M}(f)$ is equivalent to the asymptotic normality of

$$n^{-1/2} \alpha^T V_n^T U_n^T \mathbf{Z}^T \Sigma_p^{T/2} f(\widehat{\Sigma}_p) \Sigma_p^{1/2} \mathbf{Z} U_n V_n \eta$$

for arbitrary (but fixed) vectors α and $\eta \in \mathbb{R}^q$.

Recall that $\mathcal{X} = [0, \limsup_p \lambda_{\max}(\Sigma_p)(1 + \sqrt{\gamma})^2]$. Let \mathcal{C} be any contour enclosing \mathcal{X} such that $f(\cdot)$ is analytic on its interior. With slight modifications, all arguments in the following hold for arbitrary such \mathcal{C} . For convenience, select \mathcal{C} as rectangle with vertices $\underline{u} \pm i v_0$ and $\bar{u} \pm i v_0$, such that $v_0 > 0$; $\bar{u} > \limsup_p \lambda_{\max}(\Sigma_p)(1 + \sqrt{\gamma})^2$; $\underline{u} < 0$. Such a rectangle must exist.

By *Cauchy's integral formula*, if $\lambda_{\max}(\widehat{\Sigma}_p) < \bar{u}$,

$$\begin{aligned}n^{-1/2} \alpha^T V_n^T U_n^T \mathbf{Z}^T \Sigma_p^{T/2} f(\widehat{\Sigma}_p) \Sigma_p^{1/2} \mathbf{Z} U_n V_n \eta \\ = \frac{-1}{2\pi i} \oint_{\mathcal{C}} f(z) n^{-1/2} \alpha^T V_n^T U_n^T \mathbf{Z}^T \Sigma_p^{T/2} (\widehat{\Sigma}_p - zI)^{-1} \Sigma_p^{1/2} \mathbf{Z} U_n V_n \eta dz.\end{aligned}\quad (\text{A.1})$$

If $\lambda_{\max}(\widehat{\Sigma}_p) \geq \bar{u}$, the above equality may not hold. However, if we can show that $\mathbb{P}(\lambda_{\max}(\widehat{\Sigma}_p) \geq \bar{u})$ converges to 0, we can still acquire the weak limit of the left-hand side by deriving the weak limit of the right-hand side. [32], Theorem 3.1, implies that

$$\mathbb{P}(\lambda_{\max}(\widehat{\Sigma}_p) \geq \bar{u}) \rightarrow 0. \quad (\text{A.2})$$

Hence, it suffices to show the asymptotic normality of the process

$$\xi_n(z, \alpha, \eta) = n^{-1/2} \alpha^T V_n^T U_n^T \mathbf{Z}^T \Sigma_p^{T/2} (\widehat{\Sigma}_p - zI)^{-1} \Sigma_p^{1/2} \mathbf{Z} U_n V_n \eta, \quad z \in \mathcal{C}.$$

Clearly, $\xi(z, \alpha, \eta)$ is continuous with respect to z . All asymptotic results are derived in the space of continuous functions on \mathcal{C} with uniform topology. Results in Chapter 2 of [7] apply with Euclidean distance replaced by Frobenius norm of a matrix, that is $\|A\|_F = (\sum_{i=1}^m \sum_{j=1}^r |a_{ij}|^2)^{1/2}$, where $A = [a_{ij}]_{ij}$.

We may proceed to prove the asymptotic normality of $\xi_n(z, \alpha, \eta)$ on $z \in \mathcal{C}$ directly. However, several technical challenges need to be addressed. First, in view of the spectral norm of $(\widehat{\Sigma}_p - zI)^{-1}$ being unbounded when z is close to the real axis and extreme eigenvalues of $\widehat{\Sigma}_p$ exceed $\limsup_p \lambda_{\max}(\Sigma_p)(1 + \sqrt{\gamma})^2$, the tightness of the process $\xi_n(z, \alpha, \eta)$ is unclear. Secondly, $\widehat{\Sigma}_p$ is not a summation of independent terms, but contains $\mathbf{Z} U_n U_n^T \mathbf{Z}^T$, a component containing cross product terms between pairs of columns of \mathbf{Z} . These terms entangle the analysis of the correlation between $\widehat{\Sigma}_p$ and each single column of \mathbf{Z} . For these technical reasons, we avoid directly working on $\xi_n(z, \alpha, \eta)$ under **C1** on $z \in \mathcal{C}$, but start with

$n^{-1/2}U_n^T \mathbf{Z}^T \Sigma_p^{T/2}(\tilde{\Sigma}_p - \mathbb{Z}I)^{-1} \Sigma_p^{1/2} \mathbf{Z}U_n$, a component of $\xi_n(\mathbb{Z}, \alpha, \eta)$ with $\widehat{\Sigma}_p$ replaced by an uncentered counterpart

$$\tilde{\Sigma}_p = \frac{1}{n} \Sigma_p^{1/2} \mathbf{Z} \mathbf{Z}^T \Sigma_p^{T/2}. \quad (\text{A.3})$$

The relationship between $\tilde{\Sigma}_p$ and $\widehat{\Sigma}_p$ is given by $\widehat{\Sigma}_p = \tilde{\Sigma}_p - \frac{1}{n} \Sigma_p^{1/2} \mathbf{Z}U_n U_n^T \mathbf{Z}^T \Sigma_p^{T/2}$. Next, we modify the process and the distribution of \mathbf{Z} as follows.

Process smoothing. Select a sequence of $\rho_n > 0$ such that for some $\omega \in (1, 2)$

$$n\rho_n \downarrow 0, \quad \rho_n \geq n^{-\omega}.$$

Let $\mathcal{C}^+ = \mathcal{C} \cap \{u + iv : |v| \geq \rho_n\}$. Define

$$\begin{aligned} \tilde{Q}_n(\mathbb{Z}) &= n^{-1} U_n^T \mathbf{Z}^T \Sigma_p^{T/2}(\tilde{\Sigma}_p - \mathbb{Z}I)^{-1} \Sigma_p^{1/2} \mathbf{Z}U_n \quad \text{if } \mathbb{Z} \in \mathcal{C}^+, \\ \tilde{Q}_n(\mathbb{Z}) &= \frac{\rho_n - v}{2\rho_n} \tilde{Q}_n(u + i\rho_n) + \frac{v + \rho_n}{2\rho_n} \tilde{Q}_n(u - i\rho_n) \quad \text{if } \mathbb{Z} \in \mathcal{C} \setminus \mathcal{C}^+. \end{aligned}$$

To understand this definition better, note that if \mathbb{Z} is too close to the real axis, $\tilde{Q}_n(\mathbb{Z})$ is modified to be the linear interpolation of its values at $u + i\rho_n$ and $u - i\rho_n$. Observe that V_n appearing in $\xi_n(\mathbb{Z}, \alpha, \eta)$ was left out when defining $\tilde{Q}_n(\mathbb{Z})$. This trick helps transforming back to $\widehat{\Sigma}_p$ from $\tilde{\Sigma}_p$; see (A.5). Note also that V_n is a sequence of deterministic matrices of fixed dimensions, having a limit under **C5** and **C6**. The reason to smooth the process is to guarantee a bound of order $O(\rho_n^{-1})$ on the spectral norm of $(\tilde{\Sigma}_p - \mathbb{Z}I_p)^{-1}$. It is crucial in the proof of tightness.

Variable truncation. **C1** will be temporarily replaced by the following truncated variable condition. Select a positive sequence ε_n such that

$$\varepsilon_n \rightarrow 0 \quad \text{and} \quad \varepsilon_n^{-4} \mathbb{E}[z_{11}^4 \mathbb{1}(|z_{11}| \geq \varepsilon_n n^{1/2})] \rightarrow 0.$$

The existence of ε_n is shown in [32]. We then truncate z_{ij} to be $z_{ij} \mathbb{1}(|z_{ij}| \leq \varepsilon_n n^{1/2})$. The truncated variable is then standardized to maintain zero mean and unit variance. Since we will mostly work on the truncated variables in the following sections, for notational simplicity, we shall use z_{ij} to denote the truncated random variables and \check{z}_{ij} to denote the original random variable satisfying **C1**. That is,

$$z_{ij} = \frac{\check{z}_{ij} \mathbb{1}(|\check{z}_{ij}| \leq \varepsilon_n n^{1/2}) - \mathbb{E}[\check{z}_{ij} \mathbb{1}(|\check{z}_{ij}| \leq \varepsilon_n n^{1/2})]}{[\mathbb{E}[\check{z}_{ij} \mathbb{1}(|\check{z}_{ij}| \leq \varepsilon_n n^{1/2}) - \mathbb{E}[\check{z}_{ij} \mathbb{1}(|\check{z}_{ij}| \leq \varepsilon_n n^{1/2})]^2]^{1/2}}.$$

For some constant \mathcal{K} , when n is sufficiently large,

$$|z_{ij}| \leq \mathcal{K} \varepsilon_n n^{1/2}, \quad \mathbb{E}[z_{ij}] = 0, \quad \mathbb{E}[z_{ij}^2] = 1, \quad \mathbb{E}[z_{ij}^4] < \infty. \quad (\text{A.4})$$

The reason to truncate \check{z}_{ij} is to obtain a bound on the probability of extreme eigenvalues of $\widehat{\Sigma}_p$ exceeding $\limsup_p \lambda_{\max}(\Sigma_p)(1 + \sqrt{\gamma})^2$, which is crucial to be able to prove tightness of the smoothed random processes on \mathcal{C} . Under the original condition **C1**, although (A.2) holds, such a tail bound is not available. After the truncation, the following lemma shown in [6, 32] holds.

Lemma A.1. Suppose the entries of \mathbf{Z} satisfy (A.4). For any positive ℓ and any $\mathfrak{D} \in (\limsup_p \lambda_{\max}(\Sigma_p)(1 + \sqrt{\gamma})^2, \bar{u})$,

$$\mathbb{P}(\lambda_{\max}(\tilde{\Sigma}_p) \geq \mathfrak{D}) = o(n^{-\ell}).$$

It is argued later that the process smoothing and variable truncation steps do not change the weak limit of objects under consideration.

Theorem A.1. For arbitrary vectors a and $b \in \mathbb{R}^k$, define $G_n(\mathbb{z}, a, b) = a^T \tilde{Q}_n(\mathbb{z})b$. Suppose \mathbf{Z} satisfies (A.4) and suppose C2–C6 in Section 2 hold. Then,

$$n^{1/2}\{G_n(\mathbb{z}, a, b) - a^T b(1 - \Theta^{-1}(\mathbb{z}, \gamma))\} \xrightarrow{D} \Psi^{(1)}(\mathbb{z}), \quad \mathbb{z} \in \mathcal{C},$$

where \xrightarrow{D} denotes weak convergence in $C(\mathcal{C}, \mathbb{R}^2)$, and $\Psi^{(1)}(\mathbb{z})$ is a Gaussian process with zero mean and covariance function

$$\Gamma^{(1)}(\mathbb{z}_1, \mathbb{z}_2) = \delta(\mathbb{z}_1, \mathbb{z}_2, \gamma) \Theta^{-2}(\mathbb{z}_1, \gamma) \Theta^{-2}(\mathbb{z}_2, \gamma) [\|a\|^2 \|b\|^2 + (a^T b)^2].$$

See Section S.3 of the Supplementary Material [21] for proof of the theorem.

Next, transforming back to $\hat{\Sigma}_p$, define

$$\begin{aligned} \hat{Q}_n(\mathbb{z}) &= n^{-1} U_n^T \mathbf{Z}^T \Sigma_p^{T/2} (\hat{\Sigma}_p - \mathbb{z}I)^{-1} \Sigma_p^{1/2} \mathbf{Z} U_n, \quad \mathbb{z} \in \mathcal{C}^+, \\ \hat{Q}_n(\mathbb{z}) &= \frac{\rho_n - v}{2\rho_n} \hat{Q}_n(u + i\rho_n) + \frac{v + \rho_n}{2\rho_n} \hat{Q}_n(u - i\rho_n), \quad \mathbb{z} \in \mathcal{C} \setminus \mathcal{C}^+. \end{aligned}$$

Using the identity (A.3), and Lemma S.6 in the Supplementary Material [21], we get

$$\hat{Q}_n(\mathbb{z}) = \tilde{Q}_n(\mathbb{z}) [I_k - \tilde{Q}_n(\mathbb{z})]^{-1}. \quad (\text{A.5})$$

Notably, $(\Theta(\mathbb{z}, \gamma) - 1)/\Theta(\mathbb{z}, \gamma)$ is bounded away from 1 on \mathcal{C} . Since $\hat{Q}_n(\mathbb{z})$ is a smooth function of $\tilde{Q}_n(\mathbb{z})$, applying the delta-method, the following result is an immediate consequence of Theorem A.1.

Lemma A.2. Suppose \mathbf{Z} satisfies (A.4) and suppose C2–C6 in Section 2 hold. Then,

$$n^{1/2}\{\hat{Q}_n(\mathbb{z}) - \{\Theta(\mathbb{z}, \gamma) - 1\}I_k\} \xrightarrow{D} \Psi^{(2)}(\mathbb{z}), \quad \mathbb{z} \in \mathcal{C},$$

where \xrightarrow{D} denotes weak convergence in $C(\mathcal{C}, \mathbb{R}^{2k^2})$, and $\Psi^{(2)}(\mathbb{z}) = [\Psi^{(2)}(\mathbb{z})]_{ij}$ is a $k \times k$ symmetric Gaussian matrix process with zero mean and covariance

$$\begin{aligned} \mathbb{E}[\Psi^{(2)}(\mathbb{z}_1)]_{ii} [\Psi^{(2)}(\mathbb{z}_2)]_{ii} &= 2\delta(\mathbb{z}_1, \mathbb{z}_2, \gamma), \\ \mathbb{E}[\Psi^{(2)}(\mathbb{z}_1)]_{ij} [\Psi^{(2)}(\mathbb{z}_2)]_{ij} &= \delta(\mathbb{z}_1, \mathbb{z}_2, \gamma) \quad \text{if } i < j, \\ \mathbb{E}[\Psi^{(2)}(\mathbb{z}_1)]_{ij} [\Psi^{(2)}(\mathbb{z}_2)]_{i'j'} &= 0 \quad \text{if } i \neq i' \text{ or } j \neq j'. \end{aligned}$$

Define a smoothed version of $\xi_n(\mathbb{Z}, \alpha, \eta)$ as

$$\begin{aligned}\widehat{\xi}_n(\mathbb{Z}, \alpha, \eta) &= \xi_n(\mathbb{Z}, \alpha, \eta), \quad \mathbb{Z} \in \mathcal{C}^+, \\ \widehat{\xi}_n(\mathbb{Z}, \alpha, \eta) &= \frac{\rho_n - v}{2\rho_n} \xi_n(u + i\rho_n, \alpha, \eta) + \frac{v + \rho_n}{2\rho_n} \xi_n(u - i\rho_n, \alpha, \eta), \quad \mathbb{Z} \in \mathcal{C} \setminus \mathcal{C}^+.\end{aligned}$$

Lemma A.3. Suppose that **Z** satisfies (A.4) and **C2–C6** hold. Then,

$$\widehat{\xi}_n(\mathbb{Z}, \alpha, \eta) - n^{1/2}(\Theta(\mathbb{Z}, \gamma) - 1)\alpha^T \eta \xrightarrow{D} \Psi^{(3)}(\mathbb{Z}),$$

where \xrightarrow{D} denotes weak convergence in $C(\mathcal{C}, \mathbb{R}^2)$, and $\Psi^{(3)}(\mathbb{Z})$ is a Gaussian process with zero mean and covariance function $\Gamma^{(2)}(\mathbb{Z}_1, \mathbb{Z}_2) = \delta(\mathbb{Z}_1, \mathbb{Z}_2, \gamma)[\|\alpha\|^2\|\eta\|^2 + (\alpha^T \eta)^2]$.

The following result is an immediate consequence of the foregoing:

$$\oint_{\mathcal{C}} \frac{f(\mathbb{Z})\widehat{\xi}_n(\mathbb{Z}, \alpha, \eta)}{-2\pi i} d\mathbb{Z} - n^{1/2}\Omega(f, \gamma)\alpha^T \eta \implies \mathcal{N}(0, [\|\alpha\|^2\|\eta\|^2 + (\alpha^T \eta)^2]\Delta(f, \gamma)). \quad (\text{A.6})$$

To complete the proof of Theorem 2.1, we further need to show (A.6) still holds if (A.4) is extended to **C1** and $\widehat{\xi}_n(\mathbb{Z}, \alpha, \eta)$ is replaced by $\xi_n(\mathbb{Z}, \alpha, \eta)$. We present the extension of (A.6) in Section S.9 of the Supplementary Material [21].

Acknowledgements

H. Li's research was supported in part by NSF Grants DMS-1407530 and DMS-1811405.

D. Paul's research was supported in part by NSF Grants DMS-1407530, DMS-1713120, DMS-1811405 and DMS-1915894.

Supplementary Material

Supplementary material (DOI: [10.3150/19-BEJ1186SUPP](https://doi.org/10.3150/19-BEJ1186SUPP); .pdf). Supplementary Material includes additional simulation results and detailed proofs of the main theoretical results presented in this paper.

References

- [1] Anderson, T.W. (1958). *An Introduction to Multivariate Statistical Analysis*. Wiley Publications in Statistics. New York: Wiley; London: CRC Press. [MR0091588](#)
- [2] Bai, Z., Chen, J. and Yao, J. (2010). On estimation of the population spectral distribution from a high-dimensional sample covariance matrix. *Aust. N. Z. J. Stat.* **52** 423–437. [MR2791528](#) <https://doi.org/10.1111/j.1467-842X.2010.00590.x>

- [3] Bai, Z., Choi, K.P. and Fujikoshi, Y. (2018). Limiting behavior of eigenvalues in high-dimensional MANOVA via RMT. *Ann. Statist.* **46** 2985–3013. MR3851762 <https://doi.org/10.1214/17-AOS1646>
- [4] Bai, Z., Jiang, D., Yao, J. and Zheng, S. (2013). Testing linear hypotheses in high-dimensional regressions. *Statistics* **47** 1207–1223. MR3175743 <https://doi.org/10.1080/02331888.2012.708031>
- [5] Bai, Z. and Saranadasa, H. (1996). Effect of high dimension: By an example of a two sample problem. *Statist. Sinica* **6** 311–329. MR1399305
- [6] Bai, Z.D. and Silverstein, J.W. (2004). CLT for linear spectral statistics of large-dimensional sample covariance matrices. *Ann. Probab.* **32** 553–605. MR2040792 <https://doi.org/10.1214/aop/1078415845>
- [7] Billingsley, P. (1968). *Convergence of Probability Measures*. New York: Wiley. MR0233396
- [8] Biswas, M. and Ghosh, A.K. (2014). A nonparametric two-sample test applicable to high dimensional data. *J. Multivariate Anal.* **123** 160–171. MR3130427 <https://doi.org/10.1016/j.jmva.2013.09.004>
- [9] Bodnar, T., Dette, H. and Parolya, N. (2019). Testing for independence of large dimensional vectors. *Ann. Statist.* **47** 2977–3008. MR3988779 <https://doi.org/10.1214/18-AOS1771>
- [10] Cai, T.T., Liu, W. and Xia, Y. (2014). Two-sample test of high dimensional means under dependence. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **76** 349–372. MR3164870 <https://doi.org/10.1111/rssb.12034>
- [11] Cai, T.T. and Xia, Y. (2014). High-dimensional sparse MANOVA. *J. Multivariate Anal.* **131** 174–196. MR3252643 <https://doi.org/10.1016/j.jmva.2014.07.002>
- [12] Chang, J., Zheng, C., Zhou, W.-X. and Zhou, W. (2017). Simulation-based hypothesis testing of high dimensional means under covariance heterogeneity. *Biometrics* **73** 1300–1310. MR3744543 <https://doi.org/10.1111/biom.12695>
- [13] Chen, L.S., Paul, D., Prentice, R.L. and Wang, P. (2011). A regularized Hotelling’s T^2 test for pathway analysis in proteomic studies. *J. Amer. Statist. Assoc.* **106** 1345–1360. MR2896840 <https://doi.org/10.1198/jasa.2011.ap10599>
- [14] Chen, S.X., Li, J. and Zhong, P.-S. (2014). Two-sample tests for high dimensional means with thresholding and data transformation. Preprint. Available at [arXiv:1410.2848](https://arxiv.org/abs/1410.2848).
- [15] Chen, S.X. and Qin, Y.-L. (2010). A two-sample test for high-dimensional data with applications to gene-set testing. *Ann. Statist.* **38** 808–835. MR2604697 <https://doi.org/10.1214/09-AOS716>
- [16] Fujikoshi, Y., Himeno, T. and Wakaki, H. (2004). Asymptotic results of a high dimensional Manova test and power comparison when the dimension is large compared to the sample size. *J. Japan Statist. Soc.* **34** 19–26. MR2084057 <https://doi.org/10.14490/jjss.34.19>
- [17] Guo, B. and Chen, S.X. (2016). Tests for high dimensional generalized linear models. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **78** 1079–1102. MR3557190 <https://doi.org/10.1111/rssb.12152>
- [18] Hu, J., Bai, Z., Wang, C. and Wang, W. (2017). On testing the equality of high dimensional mean vectors with unequal covariance matrices. *Ann. Inst. Statist. Math.* **69** 365–387. MR3611524 <https://doi.org/10.1007/s10463-015-0543-8>
- [19] Ledoit, O. and P  ch  , S. (2011). Eigenvectors of some large sample covariance matrix ensembles. *Probab. Theory Related Fields* **151** 233–264. MR2834718 <https://doi.org/10.1007/s00440-010-0298-3>
- [20] Ledoit, O. and Wolf, M. (2012). Nonlinear shrinkage estimation of large-dimensional covariance matrices. *Ann. Statist.* **40** 1024–1060. MR2985942 <https://doi.org/10.1214/12-AOS989>
- [21] Li, H., Aue, A. and Paul, D. (2020). Supplement to “High-dimensional general linear hypothesis tests via non-linear spectral shrinkage.” <https://doi.org/10.3150/19-BEJ1186SUPP>
- [22] Li, H., Aue, A., Paul, D., Peng, J. and Wang, P. (2016). An adaptable generalization of Hotelling’s T^2 test in high dimension. Preprint. Available at [arXiv:1609.08725](https://arxiv.org/abs/1609.08725).
- [23] Liu, H., Aue, A. and Paul, D. (2015). On the Mar  enko–Pastur law for linear time series. *Ann. Statist.* **43** 675–712. MR3319140 <https://doi.org/10.1214/14-AOS1294>
- [24] Lopes, M., Jacob, L. and Wainwright, M.J. (2011). A more powerful two-sample test in high dimensions using random projection. In *Advances in Neural Information Processing Systems* 1206–1214.

- [25] Muirhead, R.J. (2009). *Aspects of Multivariate Statistical Theory*. Wiley Series in Probability and Mathematical Statistics New York: Wiley. MR0652932
- [26] Pan, G.M. and Zhou, W. (2011). Central limit theorem for Hotelling's T^2 statistic under large dimension. *Ann. Appl. Probab.* **21** 1860–1910. MR2884053 <https://doi.org/10.1214/10-AAP742>
- [27] Paul, D. and Aue, A. (2014). Random matrix theory in statistics: A review. *J. Statist. Plann. Inference* **150** 1–29. MR3206718 <https://doi.org/10.1016/j.jspi.2013.09.005>
- [28] Srivastava, M.S. and Fujikoshi, Y. (2006). Multivariate analysis of variance with fewer observations than the dimension. *J. Multivariate Anal.* **97** 1927–1940. MR2301265 <https://doi.org/10.1016/j.jmva.2005.08.010>
- [29] Srivastava, R., Li, P. and Ruppert, D. (2016). RAPTT: An exact two-sample test in high dimensions using random projections. *J. Comput. Graph. Statist.* **25** 954–970. MR3533647 <https://doi.org/10.1080/10618600.2015.1062771>
- [30] Wang, L., Peng, B. and Li, R. (2015). A high-dimensional nonparametric multivariate test for mean vector. *J. Amer. Statist. Assoc.* **110** 1658–1669. MR3449062 <https://doi.org/10.1080/01621459.2014.988215>
- [31] Yamada, T. and Himeno, T. (2015). Testing homogeneity of mean vectors under heteroscedasticity in high-dimension. *J. Multivariate Anal.* **139** 7–27. MR3349477 <https://doi.org/10.1016/j.jmva.2015.02.005>
- [32] Yin, Y.Q., Bai, Z.D. and Krishnaiah, P.R. (1988). On the limit of the largest eigenvalue of the large-dimensional sample covariance matrix. *Probab. Theory Related Fields* **78** 509–521. MR0950344 <https://doi.org/10.1007/BF00353874>
- [33] Zheng, S. (2012). Central limit theorems for linear spectral statistics of large dimensional F -matrices. *Ann. Inst. Henri Poincaré Probab. Stat.* **48** 444–476. MR2954263 <https://doi.org/10.1214/11-AIHP414>
- [34] Zheng, S., Bai, Z. and Yao, J. (2017). CLT for eigenvalue statistics of large-dimensional general Fisher matrices with applications. *Bernoulli* **23** 1130–1178. MR3606762 <https://doi.org/10.3150/15-BEJ772>
- [35] Zhou, B., Guo, J. and Zhang, J.-T. (2017). High-dimensional general linear hypothesis testing under heteroscedasticity. *J. Statist. Plann. Inference* **188** 36–54. MR3648316 <https://doi.org/10.1016/j.jspi.2017.03.005>
- [36] Zhu, Y. and Bradic, J. (2018). Linear hypothesis testing in dense high-dimensional linear models. *J. Amer. Statist. Assoc.* **113** 1583–1600. MR3902231 <https://doi.org/10.1080/01621459.2017.1356319>

Received September 2018 and revised October 2019