

A More Powerful Two-Sample Test in High Dimensions using Random Projection

*Miles Lopes
Laurent Jacob
Martin Wainwright*



Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2014-28

<http://www.eecs.berkeley.edu/Pubs/TechRpts/2014/EECS-2014-28.html>

April 9, 2014

Copyright © 2014, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Acknowledgement

MEL gratefully acknowledges the support of the DOE CSGF Fellowship, under grant number DE-FG02-97ER25308.

A More Powerful Two-Sample Test in High Dimensions using Random Projection

Miles E. Lopes¹ Laurent J. Jacob¹ Martin J. Wainwright^{1,2}
mlopes@stat.berkeley.edu laurent@stat.berkeley.edu wainwrig@stat.berkeley.edu

Departments of Statistics¹ and EECS²
University of California, Berkeley

Abstract

We consider the hypothesis testing problem of detecting a shift between the means of two multivariate normal distributions in the high-dimensional setting, allowing for the data dimension p to exceed the sample size n . Specifically, we propose a new test statistic for the two-sample test of means that integrates a random projection with the classical Hotelling T^2 statistic. Working under a high-dimensional framework with $(p, n) \rightarrow \infty$, we first derive an asymptotic power function for our test, and then provide sufficient conditions for it to achieve greater power than other state-of-the-art tests. Using ROC curves generated from synthetic data, we demonstrate superior performance against competing tests in the parameter regimes anticipated by our theoretical results.

1 Introduction

Application domains such as molecular biology and fMRI [e.g., 1, 2, 3, 4] have stimulated considerable interest in two-sample hypothesis testing problems in the high-dimensional setting, where two samples of data $\{X_1, \dots, X_{n_1}\}$ and $\{Y_1, \dots, Y_{n_2}\}$ are subsets of \mathbb{R}^p , and $n_1, n_2 \ll p$. The problem of discriminating between two data-generating distributions becomes difficult in this context as the cumulative effect of variance in many variables can “explain away” the correct hypothesis. In transcriptomics, for instance, p gene expression measures on the order of hundreds or thousands may be used to investigate differences between two biological conditions, and it is often difficult to obtain sample sizes n_1 and n_2 larger than several dozen in each condition. For problems such as these, classical methods may be ineffective, or not applicable at all. Likewise, there has been growing interest in developing testing procedures that are better suited to deal with the effects of dimension [e.g., 5, 6, 7, 8, 9].

A fundamental instance of the general two-sample problem is the two-sample test of means with Gaussian data. In this case, two independent sets of samples $\{X_1, \dots, X_{n_1}\}$ and $\{Y_1, \dots, Y_{n_2}\} \subset \mathbb{R}^p$ are generated in an i.i.d. manner from p -dimensional multivariate normal distributions $N(\mu_1, \Sigma)$ and $N(\mu_2, \Sigma)$ respectively, where the mean vectors μ_1 and μ_2 , and positive-definite covariance matrix $\Sigma \succ 0$, are all fixed and unknown. The hypothesis testing problem of interest is

$$\mathbf{H}_0 : \mu_1 = \mu_2 \quad \text{versus} \quad \mathbf{H}_1 : \mu_1 \neq \mu_2. \tag{1}$$

The most well-known test statistic for this problem is the Hotelling T^2 statistic, defined by

$$T^2 := \frac{n_1 n_2}{n_1 + n_2} (\bar{X} - \bar{Y})^\top \hat{\Sigma}^{-1} (\bar{X} - \bar{Y}), \tag{2}$$

where $\bar{X} := \frac{1}{n_1} \sum_{j=1}^{n_1} X_j$ and $\bar{Y} := \frac{1}{n_2} \sum_{j=1}^{n_2} Y_j$ are the sample means, and $\widehat{\Sigma}$ is the pooled sample covariance matrix, given by $\widehat{\Sigma} := \frac{1}{n} \sum_{j=1}^{n_1} (X_j - \bar{X})(X_j - \bar{X})^\top + \frac{1}{n} \sum_{j=1}^{n_2} (Y_j - \bar{Y})(Y_j - \bar{Y})^\top$, where we define $n := n_1 + n_2 - 2$ for convenience.

When $p > n$, the matrix $\widehat{\Sigma}$ is singular, and the Hotelling test is not well-defined. Even when $p \leq n$, the Hotelling test is known to perform poorly if p is nearly as large as n . This was shown in an important paper of Bai and Saranadasa (abbreviated BS) [5], who studied the performance of the Hotelling test under $(p, n) \rightarrow \infty$ with $p/n \rightarrow 1 - \epsilon$, and showed that the asymptotic power of the test suffers for small values of $\epsilon > 0$. Consequently, several improvements on the Hotelling test have been proposed in the high-dimensional setting in past years [e.g., 5, 6, 7, 8].

Due to the well-known degradation of $\widehat{\Sigma}$ as an estimate of Σ in high dimensions, the line of research on extensions of Hotelling test for problem (1) has focused on replacing $\widehat{\Sigma}$ in the definition of T^2 with other estimators of Σ . In the paper [5], BS proposed a test statistic based on the quantity $(\bar{X} - \bar{Y})^\top (\bar{X} - \bar{Y})$, which can be viewed as replacing $\widehat{\Sigma}$ with $I_{p \times p}$. It was shown by BS that this statistic achieves non-trivial asymptotic power whenever the ratio p/n converges to a constant $c \in (0, \infty)$. This statistic was later refined by Chen and Qin [8] (CQ for short) who showed that the same asymptotic power can be achieved without imposing any explicit restriction on the limit of p/n . Another direction was considered by Srivastava and Du [6, 7] (SD for short), who proposed a test statistic based on $(\bar{X} - \bar{Y})^\top \widehat{D}^{-1} (\bar{X} - \bar{Y})$, where \widehat{D} is the diagonal matrix associated with $\widehat{\Sigma}$, i.e. $\widehat{D}_{ii} = \widehat{\Sigma}_{ii}$. This choice ensures that \widehat{D} is invertible for all dimensions p with probability 1. Srivastava and Du demonstrated that their test has superior asymptotic power to the tests of BS and CQ under a certain parameter setting and local alternative when $n = \mathcal{O}(p)$. To the best of our knowledge, the procedures of CQ and SD represent the state-of-the-art among tests for problem (1)¹ with a known asymptotic power function under the scaling $(p, n) \rightarrow \infty$.

In this paper, we propose a new testing procedure for problem (1) in the high-dimensional setting, which involves randomly projecting the p -dimensional samples into a space of lower dimension $k \leq \min\{n, p\}$, and then working with the Hotelling test in \mathbb{R}^k . Allowing $(p, n) \rightarrow \infty$, we derive an asymptotic power function for our test and show that it outperforms the tests of BS, CQ, and SD in terms of asymptotic relative efficiency under certain conditions. Our comparison results are valid with p/n tending to a constant or infinity. Furthermore, whereas the mentioned testing procedures can only offer approximate level- α critical values, our procedure specifies *exact* level- α critical values for general multivariate normal data. Our test is also very easy to implement, and has a computational cost of order $\mathcal{O}(n^2 p)$ operations when k scales linearly with n , which is modest in the high-dimensional setting.

From a conceptual point of view, the procedure studied here is most distinct from past approaches in the way that covariance structure is incorporated into the test statistic. As stated above, the test statistics of BS, CQ, and SD are essentially based on versions of the Hotelling T^2 with diagonal estimators of Σ . Our analysis and simulations show that this limited estimation of Σ sacrifices power when the data variables are correlated, or when most of the variance can be captured in a small number of variables. In this regard, our procedure is motivated by the idea that covariance structure may be used more effectively by testing with projected samples in a space of lower dimension. The use of projection-based test statistics has also been considered previously in Jacob et al. [10] and Cl  m  n  on et al. [9].

¹The tests of BS, CQ, and SD actually extend somewhat beyond the problem (1) in that their asymptotic power functions have been obtained under data-generating distributions more general than Gaussian, e.g. satisfying simple moment conditions.

The remainder of this paper is organized as follows. In Section 2, we discuss the intuition for our testing procedure, and then formally define the test statistic. Section 3 is devoted to a number of theoretical results about the performance of the test. Theorem 1 in Section 3.1 provides an asymptotic power function, and Theorems 2 and 3 in Sections 3.4 and 3.5 give sufficient conditions for achieving greater power than the tests of CQ and SD in the sense of asymptotic relative efficiency. In Sections 4.1 and 4.2, we use synthetic data to make performance comparisons with ROC and calibration curves against the mentioned tests, as well as some recent non-parametric procedures such as maximum mean discrepancy (MMD) [11], kernel Fisher discriminant analysis (KFDA) [12], and a test based on area-under-curve maximization, denoted TreeRank [9]. These simulations show that our test outperforms competing tests in the parameter regimes anticipated by our theoretical results.

Notation. We use $\delta := \mu_1 - \mu_2$ to denote the *shift vector* between the distributions $N(\mu_1, \Sigma)$ and $N(\mu_2, \Sigma)$. For a positive-definite covariance matrix Σ , let D_σ be the diagonal matrix obtained by setting the off-diagonal entries of Σ to 0, and also define the associated correlation matrix $R := D_\sigma^{-1/2} \Sigma D_\sigma^{-1/2}$. Let $z_{1-\alpha}$ denote the $1 - \alpha$ quantile of the standard normal distribution, and let Φ be its cumulative distribution function. If A is a matrix in $\mathbb{R}^{p \times p}$, let $\|A\|_2$ denote its spectral norm (maximum singular value), and define the Frobenius norm $\|A\|_F := \sqrt{\sum_{i,j} A_{ij}^2}$. When all the eigenvalues of A are real, we denote them by $\lambda_{\min}(A) = \lambda_p(A) \leq \dots \leq \lambda_1(A) = \lambda_{\max}(A)$. If A is positive-definite, we write $A \succ 0$, and $A \succeq 0$ if A is positive semidefinite. We use the notation $f(n) \lesssim g(n)$ if there is some absolute constant $c \in (0, \infty)$ such that the inequality $f(n) \leq cn$ holds for all large n . If both $f(n) \lesssim g(n)$ and $g(n) \lesssim f(n)$ hold, then we write $f(n) \asymp g(n)$. The notation $f(n) = o(g(n))$ means $f(n)/g(n) \rightarrow 0$ as $n \rightarrow \infty$. For two random variables X and Y , equality in distribution is written as $X \stackrel{d}{=} Y$.

2 Random projection method

For the remainder of the paper, we retain the setup for the two-sample test of means (1) with Gaussian data given in Section 1. In particular, our procedure can be implemented with $p > n$ or $p \leq n$, as long as k is chosen such that $k \leq \min\{n, p\}$. In Section 3.3, we demonstrate an optimality property of the choice $k = \lfloor n/2 \rfloor$, which is valid in moderate or high-dimensions, i.e., $p \geq \lfloor n/2 \rfloor$, and we restrict our attention to this case in Theorems 2 and 3.

2.1 Intuition for random projection method

At a high level, our method can be viewed as a two-step procedure. First, a single random projection $P_k^\top \in \mathbb{R}^{k \times p}$ is drawn, and is used to map the samples from the high-dimensional space \mathbb{R}^p to a low-dimensional space \mathbb{R}^k . Second, the Hotelling T^2 test is applied to a new hypothesis testing problem, denoted $\mathbf{H}_{0,\text{proj}}$ versus $\mathbf{H}_{1,\text{proj}}$, in the projected space. A decision is then pulled back to the high-dimensional problem (1) by simply rejecting the original null hypothesis \mathbf{H}_0 whenever the Hotelling test rejects $\mathbf{H}_{0,\text{proj}}$ in the projected space.

To provide some intuition for our method, it is possible to consider the problem (1) in terms of a competition between the dimension p , and the “statistical distance” separating \mathbf{H}_0 and \mathbf{H}_1 . On one hand, the accumulation of variance from a large number of variables makes it difficult to discriminate between the hypotheses, and thus, it is desirable to reduce the dimension of the data. On the other

hand, methods for reducing dimension also tend to bring \mathbf{H}_0 and \mathbf{H}_1 “closer together,” making them harder to distinguish. Mindful of the fact that the Hotelling T^2 measures the separation of \mathbf{H}_0 and \mathbf{H}_1 in terms of the Kullback-Leibler divergence $D_{\text{KL}}(N(\mu_1, \Sigma) \| N(\mu_2, \Sigma)) = \frac{1}{2} \delta^\top \Sigma^{-1} \delta$, with $\delta = \mu_1 - \mu_2$,² we see that the relevant statistical distance is driven by the length of δ . Consequently, we seek to transform the data in a way that reduces dimension and preserves most of the length of δ upon passing to the transformed distributions. From this geometric point of view, it is natural to exploit the fact that random projections can simultaneously reduce dimension and approximately preserve length with high probability [14].

In addition to reducing dimension in a way that tends to preserve statistical distance between \mathbf{H}_0 and \mathbf{H}_1 , random projections have two other interesting properties with regard to the design of test statistics. Note that when the Hotelling test statistic is constructed from the projected samples in a space of dimension $k \leq \min\{n, p\}$, it is proportional to $[P_k^\top (\bar{X} - \bar{Y})]^\top (P_k^\top \widehat{\Sigma} P_k)^{-1} [P_k^\top (\bar{X} - \bar{Y})]$.³ Thus, whereas the tests of BS, CQ, and SD replace $\widehat{\Sigma}$ in the definition of T^2 with a diagonal estimator, our procedure uses $P_k^\top \widehat{\Sigma} P_k$ as a $k \times k$ surrogate for $\widehat{\Sigma}$. The key advantage is that $P_k^\top \widehat{\Sigma} P_k$ retains some information about the off-diagonal entries of Σ . Another benefit offered by random projection concerns the robustness of critical values. In the classical setting where $p \leq n$, the critical values of the Hotelling test are exact in the presence of Gaussian data. It is also well-known from the projection pursuit literature that the empirical distribution of randomly projected data tends to be approximately Gaussian [15]. Our procedure leverages these two facts by first “inducing Gaussianity” and then applying a test that has exact critical values for Gaussian data. Consequently, we expect that the critical values of our procedure may be accurate even when the p -dimensional data are not Gaussian, and this idea is illustrated by a simulation with data generated from a mixture model, as well as an example with real transcriptomic data in Section ??.

2.2 Formal testing procedure

For an integer $k \in \{1, \dots, \min\{n, p\}\}$, let $P_k^\top \in \mathbb{R}^{k \times p}$ denote a random matrix with i.i.d. $N(0, 1)$ entries,⁴ drawn independently of the data. Conditioning on a given draw of P_k^\top , the projected samples $\{P_k^\top X_1, \dots, P_k^\top X_{n_1}\}$ and $\{P_k^\top Y_1, \dots, P_k^\top Y_{n_2}\}$ are distributed i.i.d. according to $N(P_k^\top \mu_i, P_k^\top \Sigma P_k)$ respectively, with $i = 1, 2$. Since the projected data are Gaussian and lie in a space of dimension no larger than n , it is natural to consider applying the Hotelling test to the following two-sample problem in the projected space \mathbb{R}^k :

$$\mathbf{H}_{0,\text{proj}} : P_k^\top \mu_1 = P_k^\top \mu_2 \quad \text{versus} \quad \mathbf{H}_{1,\text{proj}} : P_k^\top \mu_1 \neq P_k^\top \mu_2. \quad (3)$$

For this projected problem, the Hotelling test statistic takes the form

$$T_k^2 := \frac{n_1 n_2}{n_1 + n_2} [P_k^\top (\bar{X} - \bar{Y})]^\top (P_k^\top \widehat{\Sigma} P_k)^{-1} [P_k^\top (\bar{X} - \bar{Y})],$$

²When $p \leq n$, the distribution of the Hotelling T^2 under both \mathbf{H}_0 and \mathbf{H}_1 is given by a scaled noncentral F distribution $\frac{pn}{n-p+1} F_{p, n-p-1}(\eta)$, with noncentrality parameter $\eta := \frac{n_1 n_2}{n_1 + n_2} \delta^\top \Sigma^{-1} \delta$. The expected value of T^2 grows linearly with η , e.g., see Muirhead [13, p. 216, p. 25].

³For the choice of P_k^\top given in Section 2.2, the matrix $P_k^\top \widehat{\Sigma} P_k$ is invertible with probability 1.

⁴We refer to P_k^\top as a projection, even though it is not a projection in the strict sense of being idempotent. Also, we do not normalize P_k^\top by $1/\sqrt{k}$ (which is commonly used for Gaussian matrices [14]) because our statistic T_k^2 is invariant with respect to this scaling.

where \bar{X} , \bar{Y} , and $\widehat{\Sigma}$ are as stated in the introduction. Note that $P_k^\top \widehat{\Sigma} P_k$ is invertible with probability 1 when P_k^\top has i.i.d. $N(0, 1)$ entries, which ensures that T_k^2 is well-defined, even when $p > n$.

When conditioned on a draw of P_k^\top , the T_k^2 statistic has an $\frac{kn}{n-k+1} F_{k, n-k+1}$ distribution under $\mathbf{H}_{0, \text{proj}}$, since it is an instance of the Hotelling test statistic [13, p. 216]. Inspection of the formula for T_k^2 also shows that its distribution is the same under both \mathbf{H}_0 and $\mathbf{H}_{0, \text{proj}}$. Consequently, if we let $t_\alpha := \frac{kn}{n-k+1} F_{k, n-k+1}^{1-\alpha}$, where $F_{k, n-k+1}^{1-\alpha}$ is the $1-\alpha$ quantile of the $F_{k, n-k+1}$ distribution, then the condition $T_k^2 \geq t_\alpha$ is a level- α decision rule for rejecting the null hypothesis in both the projected problem (3) and the original problem (1). Accordingly, we *define* this as the condition for rejecting \mathbf{H}_0 at level α in our procedure for (1). We summarize the implementation of our procedure below.

Implementation of random projection-based test at level α for problem (1).

1. Generate a single random matrix $P_k^\top \in \mathbb{R}^{k \times p}$ with i.i.d. $N(0, 1)$ entries.
 2. Compute T_k^2 , using P_k^\top and the two sets of samples.
 3. If $T_k^2 \geq t_\alpha$, reject \mathbf{H}_0 ; otherwise accept \mathbf{H}_0 .
- (★)

3 Main results and their consequences

This section is devoted to the statement and discussion of our main theoretical results, including an asymptotic power function for our test (Theorem 1), and comparisons of asymptotic relative efficiency with state-of-the-art tests proposed in past work (Theorems 2 and 3).

3.1 Asymptotic power function

Our first main result characterizes the asymptotic power of the T_k^2 test statistic in the high-dimensional setting. As is standard in high-dimensional asymptotics, we consider a sequence of hypothesis testing problems indexed by n , allowing the dimension p , sample sizes n_1 and n_2 , mean vectors μ_1 and μ_2 and covariance matrix Σ to implicitly vary as functions of n , with n tending to infinity. We also make another type of asymptotic assumption, known as a *local alternative*, which is commonplace in hypothesis testing (e.g., see van der Vaart [16§14.1]). The idea lying behind a local alternative assumption is that if the difficulty of discriminating between \mathbf{H}_0 and \mathbf{H}_1 is “held fixed” with respect to n , then it is often the case that most testing procedures have power tending to 1 under \mathbf{H}_1 as $n \rightarrow \infty$. In such a situation, it is not possible to tell if one test has greater asymptotic power than another. Consequently, it is standard to derive asymptotic power results under the extra condition that \mathbf{H}_0 and \mathbf{H}_1 become harder to distinguish as n grows. This theoretical device aids in identifying the conditions under which one test is more powerful than another. The following local alternative (**A0**), and balancing assumption (**A1**), are the same as those used by Bai and Saranadasa [5] to study the asymptotic power of the classical Hotelling test under $(n, p) \rightarrow \infty$. In particular, the local alternative (**A0**) means that the Kullback-Leibler divergence between the p -dimensional sampling distributions, $D_{\text{KL}}(N(\mu_1, \Sigma) \| N(\mu_2, \Sigma)) = \frac{1}{2} \delta^\top \Sigma^{-1} \delta$, tends to 0 as $n \rightarrow \infty$.

(A0) (Local alternative.) The shift vector and covariance matrix satisfy $\delta^\top \Sigma^{-1} \delta = o(1)$.

(A1) There is a constant $b \in (0, 1)$ such that $n_1/n \rightarrow b$.

(A2) There is a constant $y \in (0, 1)$ such that $k/n \rightarrow y$.

To set some notation for our asymptotic power result in Theorem 1, let $\theta := (\delta, \Sigma)$ be an ordered pair containing the relevant parameters for problem (1), and define Δ_k^2 as twice the Kullback-Leibler divergence between the projected sampling distributions,

$$\Delta_k^2 := 2 D_{\text{KL}} \left(N(P_k^\top \mu_1, P_k^\top \Sigma P_k) \parallel N(P_k^\top \mu_2, P_k^\top \Sigma P_k) \right) = \delta^\top P_k (P_k^\top \Sigma P_k)^{-1} P_k^\top \delta. \quad (4)$$

When interpreting the statement of Theorem 1 below, it is important to notice that each time the procedure (\star) is implemented, a draw of P_k^\top induces a new test statistic T_k^2 . Making this dependence on P_k^\top explicit, let $\beta(\theta; P_k^\top)$ denote the exact (non-asymptotic) power function of the T_k^2 statistic at level α for problem (1), conditioned on a given draw of P_k^\top , as in procedure (\star) .

Theorem 1. *Assume conditions **(A0)**, **(A1)**, and **(A2)**. Then, for almost all sequences of projections P_k^\top , the power function $\beta(\theta; P_k^\top)$ satisfies*

$$\beta(\theta; P_k^\top) - \Phi \left(-z_{1-\alpha} + b(1-b) \sqrt{\frac{1-y}{2y}} \Delta_k^2 \sqrt{n} \right) \rightarrow 0 \quad \text{as } n \rightarrow \infty. \quad (5)$$

Remarks. Notice that if $\Delta_k^2 = 0$ (e.g. under \mathbf{H}_0), then $\Phi(-z_{1-\alpha} + 0) = \alpha$, which corresponds to blind guessing at level α . Consequently, the second term $b(1-b) \sqrt{\frac{1-y}{2y}} \Delta_k^2 \sqrt{n}$ determines the advantage of our procedure over blind guessing. Since Δ_k^2 is twice the KL-divergence between the projected sampling distributions, these observations conform to the intuition from Section 2 that the KL-divergence measures the discrepancy between \mathbf{H}_0 and \mathbf{H}_1 .

Proof of Theorem 1. Let $\beta_{\text{H}}(\theta; P_k^\top)$ denote the exact power of the Hotelling test for the projected problem (3) at level α . As a preliminary step, we verify that

$$\beta(\theta; P_k^\top) = \beta_{\text{H}}(\theta; P_k^\top), \quad (6)$$

for almost all P_k^\top . To see this, first recall from Section 2 that the condition $T_k^2 \geq t_\alpha$ is a level- α rejection criterion in both the procedure (\star) for the original problem (1), and the Hotelling test for the projected problem (3). Next, note that if $\mathbf{H}_1 : \delta \neq 0$ holds, then $\mathbf{H}_{1,\text{proj}} : P_k^\top \delta \neq 0$ holds with probability 1, since $P_k^\top \delta$ is distributed as $N(0, \|\delta\|_2^2 I_{k \times k})$. Consequently, for almost all P_k^\top , the level- α decision rule $T_k^2 \geq t_\alpha$ has the same power against the alternative in both the original and the projected problems, which verifies (6). This establishes a technical link that allows results on the power of the classical Hotelling test to be transferred to the high-dimensional problem (1).

In order to complete the proof, we use a result of Bai and Saranadasa [5, Theorem 2.1]⁵, which asserts that if $\Delta_k^2 = o(1)$ holds for a fixed sequence of projections P_k^\top , and assumptions **(A1)** and **(A2)** hold, then $\beta_{\text{H}}(\theta; P_k^\top)$ satisfies

$$\beta_{\text{H}}(\theta; P_k^\top) - \Phi \left(-z_{1-\alpha} + b(1-b) \sqrt{\frac{1-y}{2y}} \Delta_k^2 \sqrt{n} \right) \rightarrow 0 \quad \text{as } n \rightarrow \infty. \quad (7)$$

To ensure $\Delta_k^2 = o(1)$, we appeal to a deterministic matrix inequality that follows from the proof of Lemma 3 in Jacob et al. [10]. Namely, for any full rank matrix $M^\top \in \mathbb{R}^{k \times p}$, and any $\delta \in \mathbb{R}^p$,

$$\delta^\top M (M^\top \Sigma M)^{-1} M^\top \delta \leq \delta^\top \Sigma^{-1} \delta.$$

⁵To prevent confusion, note that the notation in BS [5] for δ differs from ours.

Since P_k^\top is full rank with probability 1, we see that $\Delta_k^2 \leq \delta^\top \Sigma^{-1} \delta \rightarrow 0$ for almost all sequences of P_k^\top under the local alternative **(A0)**, as needed. Thus, the proof of Theorem 1 is completed by combining equation (6) with the limit (7). \square

3.2 Asymptotic relative efficiency (ARE)

Having derived an asymptotic power function in Theorem 1, we are now in position to provide a detailed comparison with the tests of CQ [8] and SD [6, 7]. We denote the asymptotic power function of our level- α random projection-based test (RP) by

$$\beta_{\text{RP}}(\theta; P_k^\top) := \Phi \left(-z_{1-\alpha} + b(1-b) \sqrt{\frac{1-y}{2y}} \Delta_k^2 \sqrt{n} \right), \quad (8)$$

where we recall $\theta := (\delta, \Sigma)$. The asymptotic power functions for the level- α testing procedures of CQ [8] and SD [6, 7] are given by

$$\beta_{\text{CQ}}(\theta) := \Phi \left(-z_{1-\alpha} + \frac{b(1-b)}{\sqrt{2}} \frac{\|\delta\|_2^2 n}{\|\Sigma\|_F} \right), \quad \text{and} \quad (9a)$$

$$\beta_{\text{SD}}(\theta) := \Phi \left(-z_{1-\alpha} + \frac{b(1-b)}{\sqrt{2}} \frac{\delta^\top D_\sigma^{-1} \delta n}{\|R\|_F} \right), \quad (9b)$$

where D_σ denotes the matrix formed by setting the off-diagonal entries of Σ to 0, and R denotes the correlation matrix associated to Σ . The functions β_{CQ} and β_{SD} are derived under local alternatives and asymptotic assumptions that are similar to the ones used here to obtain β_{RP} . In particular, all three functions can be obtained allowing p/n to tend to an arbitrary positive constant, or to infinity.

A standard method of comparing asymptotic power functions is through the concept of *asymptotic relative efficiency*, or ARE for short (e.g., see van der Vaart [16, ch. 14-15]). Since the term added to $-z_{1-\alpha}$ inside the Φ function is what controls power, the relative efficiency of tests is defined by the ratio of such terms. More explicitly, we define

$$\text{ARE}(\beta_{\text{CQ}}; \beta_{\text{RP}}) := \left(\frac{\|\delta\|_2^2 n}{\|\Sigma\|_F} / \sqrt{\frac{1-y}{y}} \Delta_k^2 \sqrt{n} \right)^2, \quad \text{and} \quad (10a)$$

$$\text{ARE}(\beta_{\text{SD}}; \beta_{\text{RP}}) := \left(\frac{\delta^\top D_\sigma^{-1} \delta n}{\|R\|_F} / \sqrt{\frac{1-y}{y}} \Delta_k^2 \sqrt{n} \right)^2. \quad (10b)$$

Whenever the ARE is less than 1, our procedure is considered to have greater asymptotic power than the competing test—with our advantage being greater for smaller values of the ARE. Consequently, we seek sufficient conditions in Theorems 2 and 3 below for ensuring that the ARE is small.

In classical analyses of asymptotic relative efficiency, the ARE is usually a deterministic quantity that does not depend on n . However, in the current context, our use of high-dimensional asymptotics, as well as a randomly constructed test statistic, lead to an ARE that varies with n and is random. (In other words, the ARE specifies a sequence of random variables indexed by n .) Moreover, the dependence of the ARE on Δ_k^2 implies that the ARE is affected by the orientation of the shift vector δ .⁶ To consider an average-case scenario, where no single orientation of δ is of

⁶In fact, $\text{ARE}(\beta_{\text{CQ}}; \beta_{\text{RP}})$ and $\text{ARE}(\beta_{\text{SD}}; \beta_{\text{RP}})$ are invariant with respect to scaling of δ , and so the orientation $\delta/\|\delta\|_2$ is the only part of the shift vector that is relevant for comparing power.

particular importance, we place a prior on δ , and assume that it follows a spherical distribution⁷ with $\mathbb{P}(\delta = 0) = 0$. This implies that the orientation $\delta/\|\delta\|_2$ of the shift follows the uniform (Haar) distribution on the unit sphere. We emphasize that our procedure (\star) does not rely on this choice of prior, and that it is only a device for making an average-case comparison against CQ and SD in Theorems 2 and 3. Lastly, we point out that a similar assumption was considered by Srivastava and Du [6], who let δ be a deterministic vector with all coordinates equal to the same value, in order to compare with the results of BS [5].

To be clear about the meaning of Proposition 1 and Theorems 2 and 3 below, we henceforth regard the ARE as a function of two random objects, P_k^\top and δ , and our probability statements are made with this understanding. We complete the preparation for our comparison theorems by stating Proposition 1 and several limiting assumptions with $n \rightarrow \infty$.

(A3) The shift δ has a spherical distribution with $\mathbb{P}(\delta = 0) = 0$, and is independent of P_k^\top .

(A4) There is a constant $a \in [0, 1)$ such that $k/p \rightarrow a$.

(A5) Assume $\frac{1}{\sqrt{k}} \frac{\text{tr}(\Sigma)}{p \lambda_{\min}(\Sigma)} = o(1)$.

(A6) Assume $\frac{\|D_\sigma^{-1}\|_2}{\text{tr}(D_\sigma^{-1})} = o(1)$.

As can be seen from the formulas for β_{RP} and the ARE, the performance of the T_k^2 statistic is determined by the random quantity Δ_k^2 . The following proposition provides interpretable upper and lower bounds on Δ_k^2 that hold with high-probability. This proposition is the main technical tool needed for our comparison results in Theorems 2 and 3. A proof is given in Appendix B.

Proposition 1. *Under conditions (A3), (A4), and (A5), let c be any positive constant strictly less than $(1 - \sqrt{a})^2$, and let C be any constant strictly greater than $\frac{(1 + \sqrt{a})^2}{(1 - \sqrt{a})^2}$. Then, as $n \rightarrow \infty$, we have*

$$\mathbb{P}\left(\frac{\Delta_k^2}{\|\delta\|_2^2} \geq \frac{ck}{\text{tr}(\Sigma)}\right) \rightarrow 1, \quad \text{and} \quad (11a)$$

$$\mathbb{P}\left(\frac{\Delta_k^2}{\|\delta\|_2^2} \leq \frac{Ck}{p \lambda_{\min}(\Sigma)}\right) \rightarrow 1. \quad (11b)$$

Remarks. Although we have presented upper and lower bounds in an asymptotic manner, our proof specifies non-asymptotic bounds on $\Delta_k^2/\|\delta\|_2^2$. Due to the fact that Proposition 1 is a tool for making asymptotic comparisons of power in Theorems 2 and 3, it is sufficient and simpler to state the bounds in this asymptotic form. Note that if the condition $\text{tr}(\Sigma) \asymp p \lambda_{\min}(\Sigma)$ holds, then Proposition 1 is sharp in the sense the upper and lower bounds (11a) and (11b) match up to constants.

3.3 Choice of projection dimension $k = \lfloor n/2 \rfloor$

We now demonstrate an optimality property of the choice of projected dimension $k = \lfloor n/2 \rfloor$. Note that this choice implicitly assumes $p \geq \lfloor n/2 \rfloor$, but this does not affect the applicability of procedure

⁷i.e. $\delta \stackrel{d}{=} U\delta$ for any orthogonal matrix U .

(\star) in moderate or high-dimensions. Letting $k/n \rightarrow y \in (0, 1)$ as in assumption **(A2)**, recall that the asymptotic power function from Theorem 1 is

$$\Phi\left(-z_{1-\alpha} + b(1-b)\sqrt{\frac{1-y}{2y}}\Delta_k^2\sqrt{n}\right).$$

Since Proposition 1 indicates that Δ_k^2 scales linearly in k up to random fluctuations, we see that formally replacing k with yn leads to maximizing the function $f(y) := \sqrt{\frac{1-y}{2y}}y$. The fact that f is maximized at $y = 1/2$ suggests that in certain cases, $k = \lfloor n/2 \rfloor$ may be asymptotically optimal in a suitable sense. Considering a simple case where $\Sigma = \sigma^2 I_{p \times p}$ for some absolute constant $\sigma^2 > 0$, it can be shown⁸ that under assumptions **(A2)**, **(A3)**, and integrability of $\|\delta\|_2^2$,

$$\frac{p}{n} \frac{\mathbb{E}(\Delta_k^2)}{\mathbb{E}(\|\delta\|_2^2)} \rightarrow y/\sigma^2, \quad (12)$$

for all $y \in (0, 1)$, as $n \rightarrow \infty$. The following proposition is an immediate extension of this observation, and shows that $k = \lfloor n/2 \rfloor$ is optimal in a precise sense for parameter settings that include $\Sigma = \sigma^2 I_{p \times p}$ as a special case. Namely, as $n \rightarrow \infty$, the quantity $\sqrt{\frac{1-y}{2y}}\Delta_k^2$ is largest on average for $k = \lfloor n/2 \rfloor$ among all choices of k , under the conditions stated below.

Proposition 2. *In addition to assumptions **(A2)** and **(A3)**, suppose that $\|\delta\|_2^2$ is integrable. Also assume that for some absolute constant $\sigma^2 > 0$, the limit (12) holds for any $y \in (0, 1)$. Let $y^* = 1/2$, and $k^* = \lfloor n/2 \rfloor$. Then, for any $y \in (0, 1)$,*

$$\lim_{n \rightarrow \infty} \frac{\sqrt{\frac{1-y^*}{2y^*}} \mathbb{E}(\Delta_{k^*}^2)}{\sqrt{\frac{1-y}{2y}} \mathbb{E}(\Delta_k^2)} = \frac{1}{2\sqrt{y(1-y)}} \geq 1. \quad (13)$$

Remarks. The ROC curves in Figure 1 illustrate several choices of projection dimension, with $k = \lfloor yn \rfloor$ and $y = 0.1, 0.3, 0.5, 0.7, 0.9$, under two different parameter settings. In Setting (1), $\Sigma = \sigma^2 I_{p \times p}$ with $\sigma^2 = 50$, and in Setting (2), the matrix Σ was constructed with a rapidly decaying spectrum, and a matrix of eigenvectors drawn from the uniform (Haar) distribution on the orthogonal group, as in panel (d) of Figure 3 (see Section 4.1 for additional details). The curves in both settings were generated by sampling $n_1 = n_2 = 50$ data at points from each of the distributions $N(\mu_1, \Sigma)$ and $N(\mu_2, \Sigma)$ in $p = 200$ dimensions, and repeating the process 2000 times under both \mathbf{H}_0 and \mathbf{H}_1 . For the experiments under \mathbf{H}_1 , the shift δ was drawn uniformly from a sphere of radius 3 for Setting (1), and radius 1 for Setting (2)—in accordance with assumption **(A3)** in Proposition 2. Note that $k = \lfloor n/2 \rfloor$ gives the best ROC curve for Setting (1) in Figure 1, which agrees with the fact that $\Sigma = \sigma^2 I_{p \times p}$ satisfies the conditions of Proposition 2. In Setting (2), we see that the choice $k = \lfloor n/2 \rfloor$ is not far from optimal, even when Σ is very different from $\sigma^2 I_{p \times p}$.

⁸Note that $\|\delta\|_2$ and $\delta/\|\delta\|_2$ are independent, and $\mathbb{E}\left[\frac{\delta^\top}{\|\delta\|_2} A \frac{\delta}{\|\delta\|_2}\right] = \text{tr}(A)/p$ for any $A \in \mathbb{R}^{p \times p}$, under **(A3)**; see [13, p. 38].

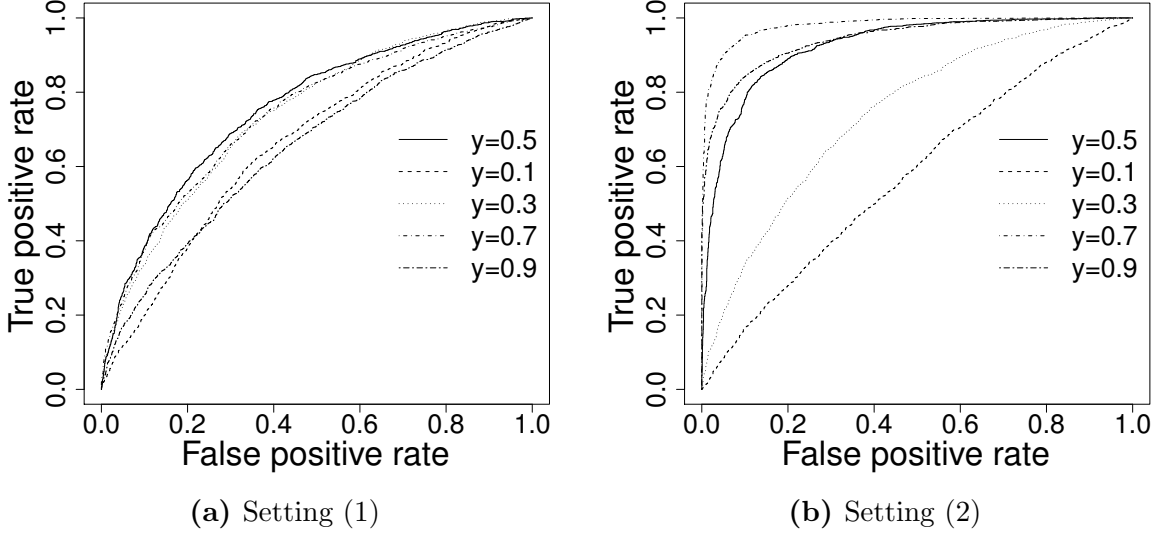


Figure 1. Setting (1) corresponds to $\Sigma = \sigma^2 I_{p \times p}$ with $\sigma^2 = 50$, and Setting (2) involves a covariance matrix Σ with randomly selected eigenvectors and a rapidly decaying spectrum. The ROC curves indicate that $k = \lfloor n/2 \rfloor$ is optimal, or nearly optimal, among the five choices of y in the two settings.

3.4 Power comparison with CQ

The next result provides a sufficient condition for the T_k^2 statistic to be asymptotically more powerful than the test of CQ. A proof is given at the end of this section (3.4).

Theorem 2. *Under the conditions of Proposition 1, suppose that we use a projection dimension $k = \lfloor n/2 \rfloor$, where we assume $p \geq \lfloor n/2 \rfloor$. Fix a number $\epsilon_1 > 0$, and let $c_1(\epsilon_1)$ be any constant strictly greater than $\frac{4}{\epsilon_1(1-\sqrt{a})^4}$. If the condition*

$$n \geq c_1(\epsilon_1) \frac{\text{tr}(\Sigma)^2}{\|\Sigma\|_F^2}, \quad (14)$$

holds for all large n , then $\mathbb{P}[\text{ARE}(\beta_{\text{CQ}}; \beta_{\text{RP}}) \leq \epsilon_1] \rightarrow 1$ as $n \rightarrow \infty$.

Remarks. The case of $\epsilon_1 = 1$ serves as the reference for equal asymptotic performance, with values $\epsilon_1 < 1$ corresponding to the T_k^2 statistic being asymptotically more powerful than the test of CQ. To interpret the result, note that Jensen's inequality implies that the ratio $\text{tr}(\Sigma)^2 / \|\Sigma\|_F^2$ lies between 1 and p , for any choice of Σ . As such, it is reasonable to interpret this ratio as a measure of the *effective dimension* of the covariance structure.⁹ The message of Theorem 2 is that as long as the sample size n grows faster than the effective dimension, then our projection-based test is asymptotically superior to the test of CQ.

The ratio $\text{tr}(\Sigma)^2 / \|\Sigma\|_F^2$ can also be viewed as measuring the *decay rate* of the spectrum of Σ , with the condition $\text{tr}(\Sigma)^2 / \|\Sigma\|_F^2 \ll p$ indicating rapid decay. This condition means that the data

⁹This ratio has also been studied as an effective measure of matrix rank in the context of low-rank matrix reconstruction [17].

has low variance in “most” directions in \mathbb{R}^p , and so projecting onto a random set of k directions will likely map the data into a low-variance subspace in which it is harder for chance variation to explain away the correct hypothesis, thereby resulting in greater power.

Example 1. One instance of spectrum decay occurs when the top s eigenvalues of Σ contain most of the mass in the spectrum. When Σ is diagonal, this has the interpretation that s variables capture most of the total variance in the data. For simplicity, assume $\lambda_1 = \dots = \lambda_s > 1$ and $\lambda_{s+1} = \dots = \lambda_p = 1$, which is similar to the *spiked covariance model* introduced by Johnstone [18]. If the top s eigenvalues contain half of the total mass of the spectrum, then $s \lambda_1 = (p - s)$, and a simple calculation shows that

$$\frac{\text{tr}(\Sigma)^2}{\|\Sigma\|_F^2} = \frac{4 \lambda_1^2}{\lambda_1^2 + \lambda_1} s \leq 4s. \quad (15)$$

This again illustrates the idea that condition (14) is satisfied as long as n grows at a faster rate than the effective number of variables s . It is straightforward to check that this example satisfies assumption **(A5)** of Theorem 2 when, for instance, $\lambda_1 = o(\sqrt{k})$.

Example 2. Another example of spectrum decay can be specified by $\lambda_i(\Sigma) \propto i^{-\nu}$, for some absolute proportionality constant, a rate parameter $\nu \in (0, \infty)$, and $i = 1, \dots, p$. This type of decay arises in connection with the Fourier coefficients of functions in Sobolev ellipsoids [19§7.2]. Noting that $\text{tr}(\Sigma) \asymp \int_1^p x^{-\nu} dx$ and $\|\Sigma\|_F^2 \asymp \int_1^p x^{-2\nu} dx$, direct computation of the integrals shows that

$$\frac{\text{tr}(\Sigma)^2}{\|\Sigma\|_F^2} \asymp \begin{cases} 1 & \text{if } \nu > 1 \\ \log^2 p & \text{if } \nu = 1 \\ p^{2(1-\nu)} & \text{if } \nu \in (\frac{1}{2}, 1) \\ p / \log p & \text{if } \nu = \frac{1}{2} \\ p & \text{if } \nu \in (0, \frac{1}{2}) \end{cases}.$$

Thus, a decay rate given by $\nu \geq 1$ is easily sufficient for condition (14) to hold unless the dimension grows exponentially with n . On the other hand, decay rates associated to $\nu \leq 1/2$ are too slow for condition (14) to hold when $n \ll p$, and rates corresponding to $\nu \in (\frac{1}{2}, 1)$ lead to a more nuanced competition between p and n . Assumption **(A5)** of Theorem 2 holds for all $\nu \in (0, 1)$, but when $\nu = 1$ or $\nu > 1$, the dimension p must satisfy the extra conditions $\log p = o(\sqrt{k})$ or $p^{\nu-1} = o(\sqrt{k})$ respectively.¹⁰

The proof of Theorem 2 is a direct application of Proposition 1.

Proof of Theorem 2. Recalling $\text{ARE}(\beta_{\text{CQ}}; \beta_{\text{RP}}) = \left(\frac{n \|\delta\|_2^2}{\|\Sigma\|_F^2} / \sqrt{n} \Delta_k^2 \right)^2$, with $k = \lfloor n/2 \rfloor$ and $y = 1/2$, the event of interest,

$$\text{ARE}(\beta_{\text{CQ}}; \beta_{\text{RP}}) \leq \epsilon_1, \quad (16)$$

is the same as

$$\frac{n}{\|\Sigma\|_F^2} \frac{1}{\epsilon_1} \leq \left(\frac{\Delta_k^2}{\|\delta\|_2^2} \right)^2.$$

¹⁰It may be possible to relax **(A5)** with a more refined analysis of the proof of Proposition 1.

By Proposition 1, we know that for any positive constant c strictly less than $(1 - \sqrt{a})^2$, the probability of the event

$$\frac{ck}{\text{tr}(\Sigma)} \leq \frac{\Delta_k^2}{\|\delta\|_2^2}. \quad (17)$$

tends to 1 as $n \rightarrow \infty$. Consequently, as long as the inequality

$$\frac{n}{\|\Sigma\|_F^2} \frac{1}{\epsilon_1} \leq \left(\frac{ck}{\text{tr}(\Sigma)} \right)^2, \quad (18)$$

holds for all large n , then the event (16) of interest will also have probability tending to 1. Replacing k with $\frac{n}{2} \cdot [1 - o(1)]$, the last condition is the same as

$$n \geq \frac{\text{tr}(\Sigma)^2}{\|\Sigma\|_F^2} \cdot \frac{4}{\epsilon_1 c^2 [1 - o(1)]^2}. \quad (19)$$

Thus, for a given choice of $c_1(\epsilon_1)$ in the statement of the theorem, it is possible to choose a positive $c < (1 - \sqrt{a})^2$ so that inequality (18) is implied by the claimed sufficient condition (14) for all large n . \square

3.5 Power comparison with SD

We now give a sufficient condition for our procedure to be asymptotically more powerful than SD.

Theorem 3. *In addition to the conditions of Theorem 2, assume that (A6) holds. Fix a number $\epsilon_1 > 0$, and let $c_1(\epsilon_1)$ be any constant strictly greater than $\frac{4}{\epsilon_1(1-\sqrt{a})^4}$. If the condition*

$$n \geq c_1(\epsilon_1) \left(\frac{\text{tr}(\Sigma)}{p} \right)^2 \left(\frac{\text{tr}(D_\sigma^{-1})}{\|R\|_F} \right)^2 \quad (20)$$

holds for all large n , then $\mathbb{P}[\text{ARE}(\beta_{\text{SD}}; \beta_{\text{RP}}) \leq \epsilon_1] \rightarrow 1$ as $n \rightarrow \infty$.

Remarks. Unlike the comparison against CQ, the correlation matrix R plays a large role in determining relative performance of our test against SD. Correlation enters in two different ways. First, the Frobenius norm $\|R\|_F$ is larger when the data variables are more correlated. Second, if Σ has a large number of small eigenvalues, then $\text{tr}(D_\sigma^{-1})$ is very large when the variables are uncorrelated, i.e. when Σ is diagonal. Letting $U\Lambda U^\top$ be a spectral decomposition of Σ , with u_i being the i th column of U^\top , note that $(D_\sigma)_{ii} = u_i^\top \Lambda u_i$. When the data variables are correlated, the vector u_i will have many nonzero components, which will give $(D_\sigma)_{ii}$ a contribution from some of the larger eigenvalues of Σ , and prevent $(D_\sigma)_{ii}$ from being too small. For example, if u_i is uniformly distributed on the unit sphere, as in Example 4 below, then on average $\mathbb{E}[(D_\sigma)_{ii}] = \text{tr}(\Sigma)/p$. Therefore, correlation has the effect of mitigating the growth of $\text{tr}(D_\sigma^{-1})$. Since the SD test statistic [6] can be thought of as a version of the Hotelling T^2 with a diagonal estimator of Σ , the SD test statistic makes no essential use of correlation structure. By contrast, our T_k^2 statistic *does* take correlation into account, and so it is understandable that correlated data enhance the performance of our test relative to SD.

Example 3. Suppose the correlation matrix $R \in \mathbb{R}^{p \times p}$ has a block-diagonal structure, with m identical blocks $B \in \mathbb{R}^{d \times d}$ along the diagonal:

$$R = \begin{pmatrix} B & & \\ & \ddots & \\ & & B \end{pmatrix}. \quad (21)$$

Note that $p = m \cdot d$. Fix a number $\rho \in (0, 1)$, and let B have diagonal entries equal to 1, and off-diagonal entries equal to ρ , i.e. $B = (1 - \rho)I_{d \times d} + \rho \mathbf{1}\mathbf{1}^\top$, where $\mathbf{1} \in \mathbb{R}^d$ is the all-ones vector. Consequently, R is positive-definite, and we may consider $\Sigma = R$ for simplicity. Since $\|B\|_F^2 = d + 2\rho^2 \binom{d}{2}$, and $\|R\|_F^2 = m \|B\|_F^2$, it follows that

$$\|R\|_F^2 = [1 + \rho^2(d - 1)]p.$$

Also, in this example we have $\text{tr}(\Sigma) = \text{tr}(D_\sigma^{-1}) = p$ and $p/d = m$, which implies

$$\left(\frac{\text{tr}(\Sigma)}{p}\right)^2 \left(\frac{\text{tr}(D_\sigma^{-1})}{\|R\|_F}\right)^2 = \frac{p}{1 + \rho^2(d - 1)} \leq \frac{m}{\rho^2}. \quad (22)$$

Under these conditions, we conclude that the sufficient condition (20) in Theorem 3 is satisfied when n grows at a faster rate than the number of blocks m . Note too that the spectrum of Σ consists of m copies of $\lambda_{\max}(\Sigma) = (1 - \rho) + \rho d$ and $(p - m)$ copies of $\lambda_{\min}(\Sigma) = 1 - \rho$, which means when ρ is not too small, the number of blocks is the same as the number of dominant eigenvalues—revealing a parallel with Example 1. From these observations, it is straightforward to check that this example satisfies assumptions **(A5)** and **(A6)** of Theorem 3. The simulations in Section 4.1 give an example where R has the form in line (21) and the variables corresponding to each block are highly correlated.

Example 4. To consider the performance of our test in a case where Σ is not constructed deterministically, Section 4.1 illustrates simulations involving *randomly selected* matrices Σ for which T_k^2 is more powerful than the tests of BS, CQ, and SD. Random correlation structure can be generated by sampling the matrix of eigenvectors of Σ from the uniform (Haar) distribution on the orthogonal group, and then imposing various decay constraints on the eigenvalues of Σ . Additional details are provided in Section 4.1.

Example 5. It is possible to show that the sufficient condition (20) *requires* non-trivial correlation in the high-dimensional setting. To see this, consider an example where the data are completely free of correlation, i.e., where $R = I_{p \times p}$. Then, $\|R\|_F = \sqrt{p}$, and Jensen's inequality implies that $\text{tr}(D_\sigma^{-1}) \geq p^2 / \text{tr}(D_\sigma) = p^2 / \text{tr}(\Sigma)$, giving $\left(\frac{\text{tr}(\Sigma)}{p}\right)^2 \left(\frac{\text{tr}(D_\sigma^{-1})}{\|R\|_F}\right)^2 \geq p$. Altogether, this shows if the data exhibits very low correlation, then (20) cannot hold when p grows faster than n (in the presence of a uniformly oriented shift δ). This is confirmed by the simulations of Section 4.1. Similarly, it is shown in the paper [6] that the SD test statistic is asymptotically superior to the CQ test statistic¹¹ when Σ is diagonal and δ is a deterministic vector with all coordinates equal to the same value.

¹¹Although the work in SD (2008) [6] was published prior to that of CQ (2010) [8], the asymptotic power function of CQ for problem (1) is the same as that of the method proposed in BS (1996) [5], and SD offer a comparison against the method of BS.

The proof of Theorem 3 makes use of concentration bounds for Gaussian quadratic forms, which are stated below in Lemma 1 (see Appendix A for proof). These bounds are similar to results in the papers of Bechar, and Laurent and Massart [20, 21] (c.f. Lemma 3 in Appendix A), but have error terms involving the spectral norm as opposed to the Frobenius norm, and hence Lemma 1 may be of independent interest.

Lemma 1. *Let $A \in \mathbb{R}^{p \times p}$ be a positive semidefinite matrix with $\|A\|_2 > 0$, and let $Z \sim N(0, I_{p \times p})$. Then, for any $t > 0$,*

$$\mathbb{P} \left[Z^\top A Z \geq \text{tr}(A) \left(1 + t \sqrt{\frac{\|A\|_2}{\text{tr}(A)}} \right)^2 \right] \leq \exp(-t^2/2), \quad (23)$$

and for any $t \in (0, \sqrt{\frac{\text{tr}(A)}{\|A\|_2}} - 1)$, we have

$$\mathbb{P} \left[Z^\top A Z \leq \text{tr}(A) \left(\sqrt{1 - \frac{\|A\|_2}{\text{tr}(A)}} - t \sqrt{\frac{\|A\|_2}{\text{tr}(A)}} \right)^2 \right] \leq \exp(-t^2/2). \quad (24)$$

Equipped with this lemma, we can now prove Theorem 3.

Proof of Theorem 3. We proceed along the lines of the proof of Theorem 2. Let us define the event of interest, $\mathcal{E}_n := \{\text{ARE}(\beta_{\text{SD}}; \beta_{\text{RP}}) \leq \epsilon_1\}$, where we recall $\text{ARE}(\beta_{\text{SD}}; \beta_{\text{RP}}) = \left(\frac{n \delta^\top D_\sigma^{-1} \delta}{\|R\|_F^2} / \sqrt{n} \Delta_k^2 \right)^2$ with $k = \lfloor n/2 \rfloor$ and $y = 1/2$. The event \mathcal{E}_n holds if and only if

$$\frac{n}{\|R\|_F^2} \frac{1}{\epsilon_1} \leq \left(\frac{\Delta_k^2}{\|\delta\|_2^2} \right)^2 \left(\frac{\|\delta\|_2^2}{\delta^\top D_\sigma^{-1} \delta} \right)^2. \quad (25)$$

We consider the two factors on the right hand side of (25) separately. By Proposition 1, for any constant $c \in (0, (1 - \sqrt{a})^2)$, the first factor $\frac{\Delta_k^2}{\|\delta\|_2^2}$ satisfies

$$\mathbb{P} \left(\frac{ck}{\text{tr}(\Sigma)} \leq \frac{\Delta_k^2}{\|\delta\|_2^2} \right) \rightarrow 1 \quad \text{as } n \rightarrow \infty. \quad (26)$$

Turning to the second factor $\frac{\|\delta\|_2^2}{\delta^\top D_\sigma^{-1} \delta}$ in line (25), we note that $\delta/\|\delta\|_2$ is uniformly distributed on the unit sphere of \mathbb{R}^p , and so $\delta/\|\delta\|_2 \stackrel{d}{=} Z/\|Z\|_2$, where $Z \sim N(0, I_{p \times p})$. Next, using Lemma 1, we see that assumption **(A6)** implies

$$\frac{Z^\top D_\sigma^{-1} Z}{\text{tr}(D_\sigma^{-1})} \rightarrow 1 \quad \text{in probability.}$$

Since $\|Z\|_2^2/p \rightarrow 1$ almost surely, we obtain the limit

$$\frac{\delta^\top D_\sigma^{-1} \delta}{\|\delta\|_2^2} \frac{p}{\text{tr}(D_\sigma^{-1})} \stackrel{d}{=} \frac{Z^\top D_\sigma^{-1} Z}{\text{tr}(D_\sigma^{-1})} \frac{p}{\|Z\|_2^2} \rightarrow 1 \quad \text{in probability.} \quad (27)$$

Consequently, for any $\tilde{c} \in (0, 1)$, the random variable $\frac{\|\delta\|_2^2}{\delta^\top D_\sigma^{-1} \delta}$ is greater than $\frac{\tilde{c}p}{\text{tr}(D_\sigma^{-1})}$ with probability tending to 1 as $n \rightarrow \infty$. Applying this observation to line (25), and using the limit (26), we conclude that $\mathbb{P}(\mathcal{E}_n) \rightarrow 1$ as long as the inequality

$$\frac{n}{\|R\|_F^2} \frac{1}{\epsilon_1} \leq \left(\frac{ck}{\text{tr}(\Sigma)} \right)^2 \left(\frac{\tilde{c}p}{\text{tr}(D_\sigma^{-1})} \right)^2 \quad (28)$$

holds for all large n . Replacing k with $\frac{n}{2} \cdot [1 - o(1)]$, the last condition is equivalent to

$$n \geq \left(\frac{\text{tr}(\Sigma)}{p} \right)^2 \left(\frac{\text{tr}(D_\sigma^{-1})}{\|R\|_F} \right)^2 \cdot \frac{4}{\epsilon_1 c^2 \tilde{c}^2 [1 - o(1)]^2}. \quad (29)$$

Thus, for a given choice of $c_1(\epsilon_1)$ in the statement of the theorem, it is possible to choose $c < (1 - \sqrt{a})^2$ and $\tilde{c} < 1$ so that the claimed sufficient condition (20) implies the inequality (28) for all large n , which completes the proof. \square

4 Performance comparisons on real and synthetic data

In this section, we compare our procedure to a broad collection of competing methods on synthetic data, illustrating the effects of the different factors involved in Theorems 2 and 3. Sections 4.1 and 4.2 consider ROC curves and calibration curves respectively.

4.1 ROC curves on synthetic data

Using multivariate normal data, we generated ROC curves (see Figure 3) in five distinct parameter settings. For each ROC curve, we sampled $n_1 = n_2 = 50$ data points from each of the distributions $N(\mu_1, \Sigma)$ and $N(\mu_2, \Sigma)$ in $p = 200$ dimensions, and repeated the process 500 times with $\delta = \mu_1 - \mu_2 = 0$ under \mathbf{H}_0 , and 500 times with $\|\delta\|_2 = 1$ under \mathbf{H}_1 . For each simulation under \mathbf{H}_1 , the shift δ was sampled as $Z/\|Z\|_2$ for $Z \sim N(0, I_{p \times p})$, so as to be drawn uniformly from the unit sphere, and satisfy assumption **(A3)** in Theorems 2 and 3. Letting $U\Lambda U^\top$ denote a spectral decomposition of Σ , we specified the first four parameter settings by choosing Λ to have a spectrum with slow or fast decay, and choosing U to be $I_{p \times p}$ or a randomly drawn matrix from the uniform (Haar) distribution on the orthogonal group [22]. Note that $U = I_{p \times p}$ gives a diagonal covariance matrix Σ , whereas a randomly chosen U induces correlation among the variables. To consider two rates of spectral decay, we selected p equally spaced eigenvalues $\lambda_1, \dots, \lambda_p$ between 10^{-2} and 1, and raised them to the power 15 for fast decay, and the power 6 for slow decay. We then added 10^{-3} to each eigenvalue to control the condition number of Σ , and rescaled them so that $\|\Sigma\|_F = \sqrt{\lambda_1^2 + \dots + \lambda_p^2} = 50$ in each of the first four settings (fixing a common amount of variance). Plots of the resulting spectra are shown in Figure 2. The fifth setting was specified by choosing the correlation matrix R to have a block-diagonal structure, corresponding to 40 groups of highly correlated variables. Specifically, the matrix R was constructed to have 40 identical blocks $B \in \mathbb{R}^{5 \times 5}$ along its diagonal, with the diagonal entries of B equal to 1, and the off-diagonal entries of B equal to $\rho := \frac{1}{1.01}$ (c.f. Example 3). The matrix Σ was then formed by setting $D_\sigma = \frac{1}{\rho} I_{p \times p}$, and $\Sigma = D_\sigma^{1/2} R D_\sigma^{1/2}$.

In addition to our random projection (RP)-based test, we implemented the methods of BS [5], SD [6], and CQ [8], which are all designed specifically for problem (1) in the high-dimensional

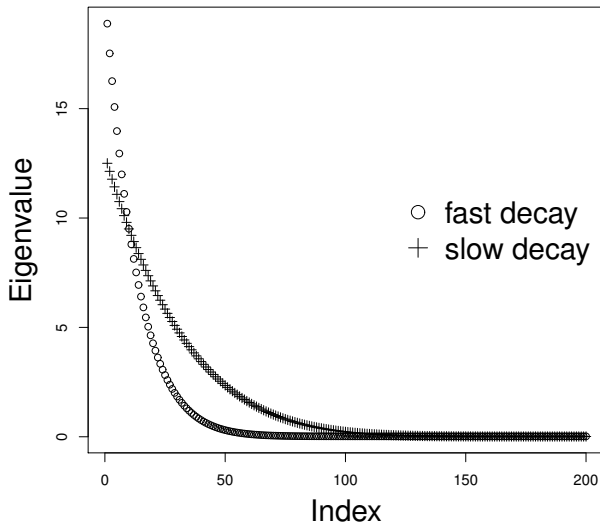


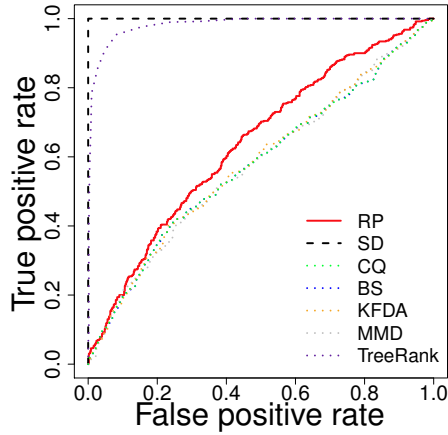
Figure 2. Plots of two sets of eigenvalues $\lambda_1, \dots, \lambda_p$, with slow and fast decay, both satisfying $\|\Sigma\|_F = \sqrt{\lambda_1^2 + \dots + \lambda_p^2} = 50$. To interpret the number of non-negligible eigenvalues, there are 29 eigenvalues greater than $\frac{1}{10}\lambda_{\max}(\Sigma)$ in the case of fast decay, and there are 65 eigenvalues greater than $\frac{1}{10}\lambda_{\max}(\Sigma)$ in the case of slow decay.

setting. For the sake of completeness, we also show comparisons against two recent non-parametric procedures that are based on kernel methods: maximum mean discrepancy (MMD) [11], and kernel Fisher discriminant analysis (KFDA) [12], as well as a test based on area-under-curve maximization, denoted TreeRank [9]. Overall, the ROC curves in Figure 3 show that in each of the five settings, either our test, or the test of SD, perform the best within this collection of procedures.

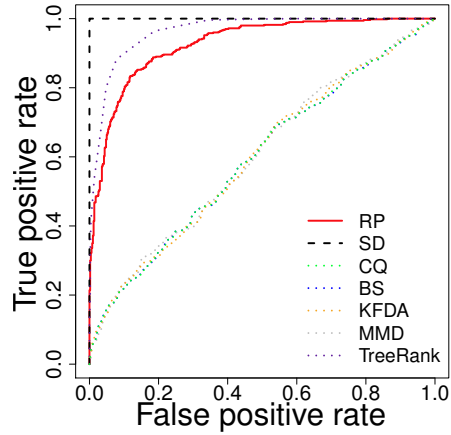
On a qualitative level, Figure 3 reveals some striking differences between our procedure and the competing tests. Comparing independent variables versus correlated variables, i.e. panels (a) and (b), with panels (c) and (d), we see that the tests of SD and TreeRank lose power in the presence of correlated data. Meanwhile, the ROC curve of our test is essentially unchanged when passing from independent variables to correlated variables. Similarly, our test also exhibits a large advantage when the correlation structure is prescribed in a block-diagonal manner in panel (e). The agreement of this effect with Theorem 3 is explained in the remarks and examples after that theorem. Comparing slow spectral decay versus fast spectral decay, i.e. panels (a) and (c), with panels (b) and (d), we see that the competing tests are essentially insensitive to the change in spectrum, whereas our test is able to take advantage of low-dimensional covariance structure. The remarks and examples of Theorem 2 give a theoretical justification for this observation.

It is also possible to offer a more quantitative assessment of the ROC curves in light of Theorems 2 and 3. Table 1 summarizes approximate values of $\text{tr}(\Sigma)^2 / \|\Sigma\|_F^2$ and $\left(\frac{\text{tr}(\Sigma)}{p}\right)^2 \left(\frac{\text{tr}(D_\Sigma^{-1})}{\|R\|_F}\right)^2$ from Theorems 2 and 3 in the five settings described above.¹² The table shows that our theory is consistent with Figure 3 in the sense that the only settings for which our test yields an inferior ROC

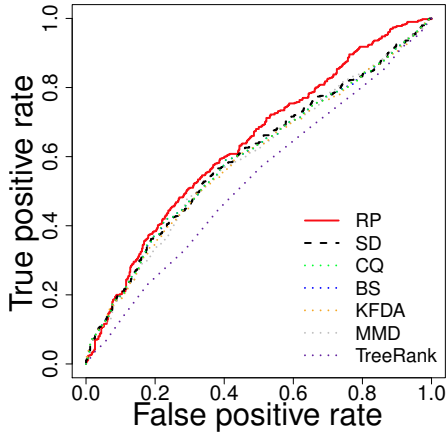
¹²For the case of randomly selected Σ , the quantities are obtained as the average from 500 draws.



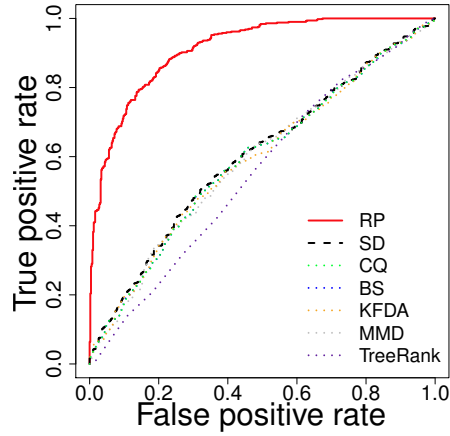
(a) diagonal Σ , slow decay



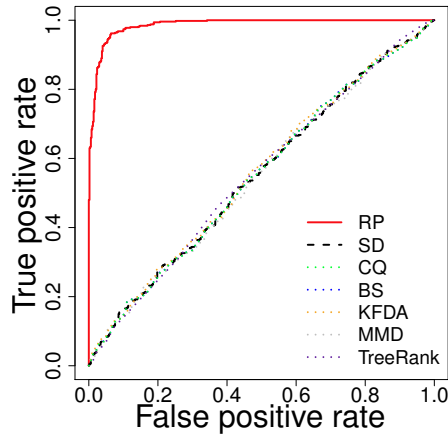
(b) diagonal Σ , fast decay



(c) random Σ , slow decay



(d) random Σ , fast decay



(e) block-diagonal correlation

Figure 3. ROC curves of several test statistics for five different settings of correlation structure and spectral decay of Σ : (a) Diagonal covariance / slow decay, (b) Diagonal covariance / fast decay, (c) Random covariance / slow decay, and (d) Random covariance / fast decay. (e) Block-diagonal correlation.

curve are those for which the quantity $\left(\frac{\text{tr}(\Sigma)}{p}\right)^2 \left(\frac{\text{tr}(D_\sigma^{-1})}{\|R\|_F}\right)^2$ is drastically larger than $n = 50 + 50 - 2 = 98$. (In all of the settings where $\text{tr}(\Sigma)^2 / \|\Sigma\|_F^2$ and $\left(\frac{\text{tr}(\Sigma)}{p}\right)^2 \left(\frac{\text{tr}(D_\sigma^{-1})}{\|R\|_F}\right)^2$ are less than n , our test yields the best ROC curve against the competitors.) However, if the entries in the table are multiplied by a choice of the constant $c_1(\epsilon_1) > \frac{4}{\epsilon_1(1-\sqrt{a})^4}$ from Theorems 2 and 3, we see that our asymptotic conditions (14) and (20) are somewhat conservative at the finite sample level. Considering $n = 98$, the table shows that $c_1(\epsilon_1)$ would need to be roughly equal to 1.5 so that the inequalities (14) and (20) hold in all the settings for which our method has a better ROC curve than the relevant competitor. In the basic case that $\epsilon_1 = 1$ and $a = 0$, we have $\frac{4}{\epsilon_1(1-\sqrt{a})^4} = 4$, which means that the constant $\frac{4}{\epsilon_1(1-\sqrt{a})^4}$ needs to be improved by roughly a factor of $4/1.5 \simeq 3$ or better. We expect that such improvement is possible with a more refined analysis of the proof of Proposition 1.

Table 1. Approximate values of the quantities $\text{tr}(\Sigma)^2 / \|\Sigma\|_F^2$ and $\left(\frac{\text{tr}(\Sigma)}{p}\right)^2 \left(\frac{\text{tr}(D_\sigma^{-1})}{\|R\|_F}\right)^2$ in the five parameter settings of the synthetic data experiments. Theorems 2 and 3 assert that these quantities determine the relative performance of our test against CQ and SD respectively.

	diagonal Σ , slow decay	diagonal Σ , fast decay	random Σ , slow decay	random Σ , fast decay	block-diagonal correlation
(Thm. 2 vs. CQ) $\text{tr}(\Sigma)^2 / \ \Sigma\ _F^2$	54	25	54	25	41
(Thm. 3 vs. SD) $\left(\frac{\text{tr}(\Sigma)}{p}\right)^2 \left(\frac{\text{tr}(D_\sigma^{-1})}{\ R\ _F}\right)^2$	4.6×10^5	3.5×10^5	58	30	41

4.2 Calibration curves on synthetic data

Figure 4 contains calibration plots resulting from the simulations described in Section 4.1—showing how well the observed false positive rates (FPR) of the various tests compare against the nominal level α . (Note that these plots only reflect simulations under \mathbf{H}_0 .) Ideally, when testing at level α , the observed FPR should be as close to α as possible, and a thin diagonal grey line is used here as a reference for perfect calibration. Figures 4 (a) and (b) correspond respectively to the settings from Section 4.1 where Σ is diagonal, with a slowly decaying spectrum, and where Σ has random eigenvectors and a rapidly decaying spectrum. In these cases the tests of BS, CQ, and SD are reasonably well-calibrated, and our test is nearly on top of the optimal diagonal line. To consider robustness of calibration, we repeated the simulation from panel (a), but replaced the sampling distributions $N(\mu_i, \Sigma)$, $i = 1, 2$, with the mixtures $0.2 N(\mu_i + d_{1,i}, \Sigma) + 0.3 N(\mu_i + d_{2,i}, \Sigma) + 0.5 N(\mu_i + d_{3,i}, \Sigma)$, where $0.2 d_{1,i} + 0.3 d_{2,i} + 0.5 d_{3,i} = 0$, and $d_{1,1}, d_{2,1}, d_{1,2}, d_{2,2}$ were drawn independently and uniformly from a sphere of radius $\|\Sigma\|_F^2$. The resulting calibration plot in Figure 4 (c) shows that our test deviates slightly from the diagonal in this case, but the calibration of the other three tests degrades to a much more noticeable extent. Experiments on other non-Gaussian distributions (e.g. with heavy tails) gave similar results, suggesting that the critical values of our procedure may be generally more robust (see also the discussions of robustness in Section 2.1).

5 Conclusion

We have proposed a novel testing procedure for the two-sample test of means in high dimensions. This procedure can be implemented in a simple manner by first projecting a dataset with a single

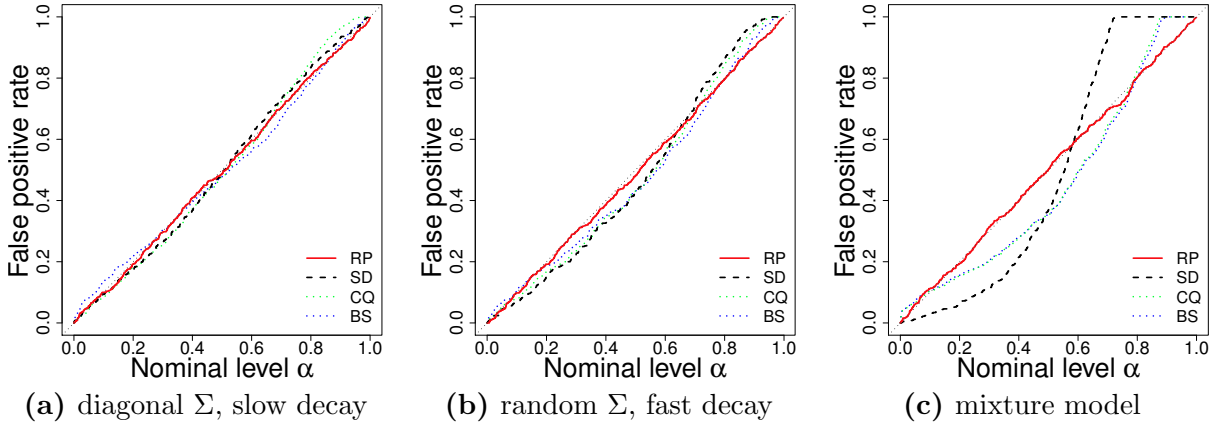


Figure 4. Calibration plots on synthetic data under three different data-generating distributions. The grey line is a reference for optimal calibration.

randomly drawn matrix, and then applying the standard Hotelling T^2 test in the projected space. In addition to deriving an asymptotic power function for this test, we have provided interpretable conditions on the covariance and correlation matrices for achieving greater power than competing tests in the sense of asymptotic relative efficiency. Specifically, our theoretical comparisons show that our test is well-suited to interesting regimes where the data variables are correlated, or where most of the variance can be captured in a small number of variables. Furthermore, in the realistic case of $(n, p) = (100, 200)$, these types of conditions were shown to correspond to favorable performance of our test against several competitors in ROC curve comparisons on synthetic data. Extensions of this work may include more refined applications of random projection to other high-dimensional testing problems.

Acknowledgements. The authors thank Sandrine Dudoit, Anne Biton, and Peter Bickel for helpful discussions. MEL gratefully acknowledges the support of the DOE CSGF Fellowship, under grant number DE-FG02-97ER25308, and LJJ the support of Stand Up to Cancer. MJW was partially supported by NSF grant DMS-0907632.

A Matrix and Concentration Inequalities

This appendix is devoted to a number of matrix and concentration inequalities used at various points in our analysis. We also prove Lemma 1, which is stated in the main text in Section 3.5.

Lemma 2. *If A and B are square real matrices of the same size with $A \succeq 0$ and $B = B^\top$, then*

$$\lambda_{\min}(B) \operatorname{tr}(A) \leq \operatorname{tr}(AB) \leq \lambda_{\max}(B) \operatorname{tr}(A). \quad (30)$$

Proof. The upper bound is an immediate consequence of Fan's inequality [28, p.10], which states that any two symmetric matrices $A, B \in \mathbb{R}^{p \times p}$ satisfy $\operatorname{tr}(AB) \leq \sum_{i=1}^p \lambda_i(A) \lambda_i(B)$, where $\lambda_{i+1}(\cdot) \leq \lambda_i(\cdot)$. Replacing A with $-A$ yields the lower bound. \square

See the papers of Bechar [20], or Laurent and Massart [21] for proofs of the following concentration bounds for Gaussian quadratic forms.

Lemma 3. *Let $A \in \mathbb{R}^{p \times p}$ with $A \succeq 0$, and $Z \sim N(0, I_{p \times p})$. Then for any $t > 0$, we have*

$$\mathbb{P} \left[Z^\top A Z \geq \operatorname{tr}(A) + 2 \|A\|_F \sqrt{t} + 2 \|A\|_2 t \right] \leq \exp(-t), \quad \text{and} \quad (31a)$$

$$\mathbb{P} \left[Z^\top A Z \leq \operatorname{tr}(A) - 2 \|A\|_F \sqrt{t} \right] \leq \exp(-t). \quad (31b)$$

The following result on the extreme eigenvalues of Wishart matrices is given in Davidson and Szarek [29, Theorem II.13].

Lemma 4. *For $k \leq p$, let $P_k^\top \in \mathbb{R}^{k \times p}$ be a random matrix with i.i.d. $N(0, 1)$ entries. Then, for all $t > 0$, we have*

$$\mathbb{P} \left[\lambda_{\max} \left(\frac{1}{p} P_k^\top P_k \right) \geq \left(1 + \sqrt{k/p} + t \right)^2 \right] \leq \exp(-pt^2/2), \quad \text{and} \quad (32a)$$

$$\mathbb{P} \left[\lambda_{\min} \left(\frac{1}{p} P_k^\top P_k \right) \leq \left(1 - \sqrt{k/p} - t \right)^2 \right] \leq \exp(-pt^2/2). \quad (32b)$$

Proof of Lemma 1. Note that the function $f(Z) := \sqrt{Z^\top A Z} = \|A^{1/2} Z\|_2$ has Lipschitz constant $\|A^{1/2}\|_2 = \sqrt{\|A\|_2}$ with respect to the Euclidean norm on \mathbb{R}^p . By the Gaussian isoperimetric inequality [30], we have for any $s > 0$,

$$\mathbb{P} [f(Z) \leq \mathbb{E}[f(Z)] - s] \leq \exp \left(\frac{-s^2}{2\|A\|_2} \right). \quad (33)$$

From the Poincaré inequality for Gaussian measures [31], the variance of $f(Z)$ is bounded above as $\operatorname{Var}[f(Z)] \leq \|A\|_2$. Noting that $\mathbb{E}[f(Z)^2] = \operatorname{tr}(A)$, we see that the expectation of $f(Z)$ is lower bounded as

$$\mathbb{E}[f(Z)] \geq \sqrt{\operatorname{tr}(A) - \|A\|_2}.$$

Substituting this lower bound into the concentration inequality (33) yields

$$\mathbb{P} \left[f(Z) \leq \sqrt{\operatorname{tr}(A) - \|A\|_2} - s \right] \leq \exp \left(\frac{-s^2}{2\|A\|_2} \right).$$

Finally, letting $t \in \left(0, \sqrt{\frac{\text{tr}(A)}{\|A\|_2} - 1}\right)$, and choosing $s^2 = t^2 \|A\|_2$ yields the claim (24).

The Gaussian isoperimetric inequality also implies $\mathbb{P}[f(Z) \geq \mathbb{E}f(Z) + s] \leq \exp\left(\frac{-s^2}{2\|A\|_2}\right)$. By Jensen's inequality, we have

$$\mathbb{E}[f(Z)] = \mathbb{E}\sqrt{Z^\top AZ} \leq \sqrt{\mathbb{E}[Z^\top AZ]} = \sqrt{\text{tr}(A)},$$

from which we obtain $\mathbb{P}\left[f(Z) \geq \sqrt{\text{tr}(A)} + s\right] \leq \exp\left(\frac{-s^2}{2\|A\|_2}\right)$, and setting $s^2 = t^2 \|A\|_2$ for $t > 0$ yields the claim (23).

B Proof of Proposition 1

The proof of Proposition 1 is based on Lemmas 5 and 6, which we state and prove below in Section B.1. We then prove Proposition 1 in two parts, by first proving the lower bound (11a), and then the upper bound (11b) in sections B.2 and B.3 respectively.

B.1 Two auxiliary lemmas

Note that the following two lemmas only deal with the randomness in the $k \times p$ matrix P_k^\top , and they can be stated without reference to the sample size n .

Lemma 5. *Let $P_k^\top \in \mathbb{R}^{k \times p}$ have i.i.d. $N(0, 1)$ entries, where $k \leq p$. Assume there is a constant $a \in [0, 1)$ such that $k/p \rightarrow a$ as $(k, p) \rightarrow \infty$. Then, there is a sequence of numbers $c_k \rightarrow (1 - \sqrt{a})^2$ such that*

$$\mathbb{P}\left[\frac{1}{p} \text{tr}(P_k(P_k^\top \Sigma P_k)^{-1} P_k^\top) \geq \frac{k}{\text{tr}(\Sigma)} c_k\right] \rightarrow 1 \text{ as } (k, p) \rightarrow \infty.$$

Proof. By the cyclic property of trace and Lemma 2, we have

$$\frac{1}{p} \text{tr}\left(P_k(P_k^\top \Sigma P_k)^{-1} P_k^\top\right) = \frac{1}{p} \text{tr}\left((P_k^\top \Sigma P_k)^{-1} P_k^\top P_k\right) \quad (34)$$

$$\geq \frac{1}{p} \text{tr}\left((P_k^\top \Sigma P_k)^{-1}\right) \lambda_{\min}\left(P_k^\top P_k\right). \quad (35)$$

For a general positive-definite matrix $A \in \mathbb{R}^{k \times k}$, Jensen's inequality implies $\text{tr}(A^{-1}) \geq k^2/\text{tr}(A)$. Combining this with the lower-bound on $\lambda_{\min}(P_k^\top P_k)$ from Lemma (4) leads to

$$\frac{1}{p} \text{tr}\left(P_k(P_k^\top \Sigma P_k)^{-1} P_k^\top\right) \geq \frac{k^2}{\text{tr}(P_k^\top \Sigma P_k)} \cdot \underbrace{[1 - \sqrt{k/p} - t_2]^2}_{=: 1 - r_2(t_2)}, \quad (36)$$

with probability at least $1 - \exp(-pt_2^2/2)$.

We now obtain a high-probability upper bound on $\text{tr}(P_k^\top \Sigma P_k)$ that is of order $k \text{tr}(\Sigma)$. First let $\Sigma = U\Lambda U^\top$ be a spectral decomposition of Σ . Writing $P_k^\top \Sigma P_k$ as $(P_k^\top U)\Lambda(U^\top P_k)$, and recalling that the columns of P_k are distributed as $N(0, I_{p \times p})$, we see that $P_k^\top \Sigma P_k$ is distributed as $P_k^\top \Lambda P_k$. Hence, we may work under the assumption that Σ and Λ are interchangeable. Let $0 < \lambda_1 \leq \dots \leq \lambda_p$

be the eigenvalues of Σ , with $\lambda_i = \Lambda_{ii}$, and let $\mathbf{Z} \in \mathbb{R}^{(pk) \times 1}$ be a concatenated column vector of k independent and identically distributed $N(0, I_{p \times p})$ vectors. Likewise, let $\mathbf{\Lambda} \in \mathbb{R}^{pk \times pk}$ be a diagonal matrix obtained by arranging k copies of Λ along the diagonal, i.e.

$$\mathbf{\Lambda} := \begin{pmatrix} \Lambda & & \\ & \ddots & \\ & & \Lambda \end{pmatrix}. \quad (37)$$

By considering the diagonal entries of $P_k^\top \mathbf{\Lambda} P_k$, it is straightforward to verify that $\text{tr}(P_k^\top \mathbf{\Lambda} P_k) \stackrel{d}{=} \mathbf{Z}^\top \mathbf{\Lambda} \mathbf{Z}$. Applying Lemma 3 to the quadratic form $\mathbf{Z}^\top \mathbf{\Lambda} \mathbf{Z}$, and noting that $\|\mathbf{\Lambda}\|_F / \text{tr}(\mathbf{\Lambda})$ and $\|\mathbf{\Lambda}\|_2 / \text{tr}(\mathbf{\Lambda})$ are at most 1, we have

$$\begin{aligned} \text{tr}(P_k^\top \mathbf{\Lambda} P_k) &\leq \text{tr}(\mathbf{\Lambda}) + 2\sqrt{t_3} \|\mathbf{\Lambda}\|_F + 2t_3 \|\mathbf{\Lambda}\|_2 \\ &= k \text{tr}(\Lambda) + 2\sqrt{t_3} \sqrt{k} \|\Lambda\|_F + 2t_3 \|\Lambda\|_2 \\ &\leq k \text{tr}(\Lambda) \underbrace{\left(1 + \frac{2\sqrt{t_3}}{\sqrt{k}} + \frac{2t_3}{k}\right)}_{=: 1+r_3(t_3)}, \end{aligned} \quad (38)$$

with probability at least $1 - \exp(-t_3)$, giving the desired upper bound on $\text{tr}(P_k^\top \mathbf{\Lambda} P_k)$. In order to combine the last bound with (36), define the event

$$E_k := \left\{ \frac{1}{p} \text{tr} \left(P_k (P_k^\top \Sigma P_k)^{-1} P_k^\top \right) \geq \frac{k}{\text{tr}(\Sigma)} \frac{1 - r_2(t_2)}{1 + r_3(t_3)} \right\},$$

and then observe that $\mathbb{P}(E_k) \geq 1 - \exp(-pt_2^2/2) - \exp(-t_3)$ by the union bound. Choosing $t_2 = 1/p^{1/4}$ and $t_3 = \sqrt{k}$, we ensure that $\mathbb{P}(E_k) \rightarrow 1$ as $(k, p) \rightarrow \infty$, and moreover, that

$$\frac{1 - r_2(t_2)}{1 + r_3(t_3)} \rightarrow (1 - \sqrt{a})^2,$$

which completes the proof. \square

Lemma 6. *Assume the conditions of Lemma 5. Then for any $C > \frac{(1+\sqrt{a})^2}{(1-\sqrt{a})^2}$, we have*

$$\mathbb{P} \left[\left\| P_k (P_k^\top \Sigma P_k)^{-1} P_k^\top \right\|_F \leq \frac{C\sqrt{k}}{\lambda_{\min}(\Sigma)} \right] \rightarrow 1 \quad \text{as } (k, p) \rightarrow \infty. \quad (39)$$

Proof. By the relation $\|A\|_F^2 = \text{tr}(A^2)$ for symmetric matrices A , and the cyclic property of trace,

$$\left\| P_k (P_k^\top \Sigma P_k)^{-1} P_k^\top \right\|_F^2 = \text{tr} \left(\left(P_k (P_k^\top \Sigma P_k)^{-1} P_k^\top \right)^2 \right) = \text{tr} \left(\left((P_k^\top \Sigma P_k)^{-1} P_k^\top P_k \right)^2 \right).$$

Letting $\rho(\cdot)$ denote the spectral radius of a matrix, we use the fact that $|\text{tr}(A)| \leq k\rho(A) \leq k\|A\|_2$ for all real $k \times k$ matrices A (see [32, p. 297]) to obtain

$$\left\| P_k (P_k^\top \Sigma P_k)^{-1} P_k^\top \right\|_F^2 \leq k \left\| \left((P_k^\top \Sigma P_k)^{-1} P_k^\top P_k \right)^2 \right\|_2.$$

Using the submultiplicative property of $\|\cdot\|_2$ twice in succession,

$$\begin{aligned}
\left\| P_k (P_k^\top \Sigma P_k)^{-1} P_k^\top \right\|_F^2 &\leq k \left\| (P_k^\top \Sigma P_k)^{-1} P_k^\top P_k \right\|_2^2 \\
&\leq k \left\| (P_k^\top \Sigma P_k)^{-1} \right\|_2^2 \cdot \left\| P_k^\top P_k \right\|_2^2 \\
&= k \frac{1}{\lambda_{\min}^2(P_k^\top \Sigma P_k)} \cdot \lambda_{\max}^2(P_k^\top P_k).
\end{aligned} \tag{40}$$

Next, by Lemma 4, we have the bound

$$\lambda_{\max}(P_k^\top P_k) \leq p \cdot \underbrace{[1 + \sqrt{k/p} + t_4]^2}_{=: 1+r_4(t_4)}, \tag{41}$$

with probability at least $1 - \exp(-pt_4^2/2)$.

By the variational characterization of eigenvalues, followed by Lemma 4, we have

$$\begin{aligned}
\lambda_{\min}(P_k^\top \Sigma P_k) &= \inf_{\|x\|_2=1} \left(x^\top P_k^\top \Sigma P_k x \right) \\
&\geq \inf_{\|y\|_2=1} \left(y^\top \Sigma y \right) \inf_{\|x\|_2=1} \|P_k x\|_2^2 \\
&= \lambda_{\min}(\Sigma) \cdot \lambda_{\min}(P_k^\top P_k) \\
&\geq \lambda_{\min}(\Sigma) \cdot p \cdot (1 - r_2(t_2)),
\end{aligned} \tag{42}$$

with probability at least $1 - \exp(-pt_2^2/2)$, and $r_2(t_2)$ defined as in line (36).

Substituting the bounds (41) and (42) into line (40), we obtain

$$\left\| P_k (P_k^\top \Sigma P_k)^{-1} P_k^\top \right\|_F^2 \leq \frac{k}{\lambda_{\min}^2(\Sigma)} \frac{(1 + r_4(t_4))^2}{(1 - r_2(t_2))^2}. \tag{43}$$

with probability at least $1 - \exp(-pt_2^2/2) - \exp(-pt_4^2/2)$, where we have used the union bound.

Setting $t_2 = t_4 = 1/p^{1/4}$, the probability of the event (43) tends to 1 as $(k, p) \rightarrow \infty$. Furthermore,

$$\frac{(1 + r_4(t_4))^2}{(1 - r_2(t_2))^2} \rightarrow \frac{(1 + \sqrt{a})^4}{(1 - \sqrt{a})^4},$$

and so we may take C in the statement of the lemma to be any constant strictly greater than $\frac{(1+\sqrt{a})^2}{(1-\sqrt{a})^2}$. \square

B.2 Proof of lower bound (11a) in Proposition 1

By the assumption on the distribution of δ , we may write $\delta/\|\delta\|_2$ as $Z/\|Z\|_2$ where $Z \sim N(0, I_{p \times p})$. Furthermore, because $\|Z\|_2/\sqrt{p} \rightarrow 1$ almost surely as $n \rightarrow \infty$, it is possible to replace $\delta/\|\delta\|_2$ with Z/\sqrt{p} , and work under the assumption that $\frac{\Delta_k^2}{\|\delta\|_2^2} = \frac{1}{p} Z^\top P_k (P_k^\top \Sigma P_k)^{-1} P_k^\top Z$. Noting that we

may take Z to be independent of P_k^\top , the concentration inequality for Gaussian quadratic forms in Lemma 2 gives a lower bound on the conditional probability

$$\mathbb{P} \left[\frac{\Delta_k^2}{\|\delta\|_2^2} \geq \frac{1}{p} \operatorname{tr} (P_k(P_k^\top \Sigma P_k)^{-1} P_k^\top) - \psi(t_1) \mid P_k^\top \right] \geq 1 - \exp(-t_1), \quad (44)$$

where $\psi(t_1) := \frac{2\sqrt{t_1}}{p} \left\| \left\| P_k(P_k^\top \Sigma P_k)^{-1} P_k^\top \right\| \right\|_F$ is a random error term, and t_1 is a positive real number that may vary with n . Now that the randomness from δ has been separated out in (44), we study the randomness from P_k^\top by defining the event

$$\mathcal{E}_n := \left\{ \frac{1}{p} \operatorname{tr} (P_k(P_k^\top \Sigma P_k)^{-1} P_k^\top) - \psi(t_1) \geq L_n \right\}, \quad (45)$$

where L_n is a real number whose dependence on n will be specified below. To see the main line of argument toward the statement of the proposition, we integrate the conditional probability in line (44) with respect to P_k^\top , and obtain

$$\mathbb{P} \left(\frac{\Delta_k^2}{\|\delta\|_2^2} \geq L_n \right) \geq [1 - \exp(-t_1)] \mathbb{P}(\mathcal{E}_n). \quad (46)$$

The rest of the proof proceeds in two parts. First, we lower-bound $\operatorname{tr} (P_k(P_k^\top \Sigma P_k)^{-1} P_k^\top) / p$ on an event \mathcal{E}'_n with $\mathbb{P}(\mathcal{E}'_n) \rightarrow 1$ as $n \rightarrow \infty$. Second, we upper-bound $\psi(t_1)$ on an event \mathcal{E}''_n with $\mathbb{P}(\mathcal{E}''_n) \rightarrow 1$. Then we choose L_n so that $\mathcal{E}_n \supset \mathcal{E}'_n \cap \mathcal{E}''_n$, and take $t_1 \rightarrow \infty$ so that (46) implies $\mathbb{P}(\Delta_k^2 / \|\delta\|_2^2 \geq L_n) \rightarrow 1$ as $n \rightarrow \infty$.

For the first step of lower-bounding $\operatorname{tr} (P_k(P_k^\top \Sigma P_k)^{-1} P_k^\top) / p$, Lemma 5 asserts that there is a sequence of numbers $c_n \rightarrow (1 - \sqrt{a})^2$ such that the event

$$\mathcal{E}'_n := \left\{ \frac{1}{p} \operatorname{tr} (P_k(P_k^\top \Sigma P_k)^{-1} P_k^\top) \geq \frac{k}{\operatorname{tr}(\Sigma)} c_n \right\} \quad (47)$$

satisfies $\mathbb{P}(\mathcal{E}'_n) \rightarrow 1$ as $n \rightarrow \infty$.

Next, for the second step of upper-bounding the error $\psi(t_1)$, Lemma 6 guarantees that for any constant C strictly greater than $\frac{(1+\sqrt{a})^2}{(1-\sqrt{a})^2}$, the event

$$\mathcal{E}''_n := \left\{ \frac{2}{p} \left\| \left\| P_k(P_k^\top \Sigma P_k)^{-1} P_k^\top \right\| \right\|_F \leq \frac{2C\sqrt{k}}{p\lambda_{\min}(\Sigma)} \right\} \quad (48)$$

satisfies $\mathbb{P}(\mathcal{E}''_n) \rightarrow 1$ as $n \rightarrow \infty$.

Now, with consideration to \mathcal{E}'_n and \mathcal{E}''_n , define the deterministic quantity

$$L_n := \frac{k}{\operatorname{tr}(\Sigma)} \left[c_n - \sqrt{t_1} \frac{2C}{\sqrt{k}} \frac{\operatorname{tr}(\Sigma)}{p\lambda_{\min}(\Sigma)} \right], \quad (49)$$

which ensures $\mathcal{E}_n \supset \mathcal{E}'_n \cap \mathcal{E}''_n$ for all choices of t_1 . Consequently, $\mathbb{P}(\mathcal{E}_n) \rightarrow 1$, and it remains to choose t_1 appropriately so that the probability in line (46) tends to 1. If we let $t_1 = \sqrt{k} \frac{p\lambda_{\min}(\Sigma)}{\operatorname{tr}(\Sigma)}$, then $t_1 \rightarrow \infty$ by assumption **(A5)**, and the second term inside the brackets in line (49) vanishes as $n \rightarrow \infty$. Altogether, we have shown that

$$L_n \frac{\operatorname{tr}(\Sigma)}{k} \rightarrow (1 - \sqrt{a})^2, \quad \text{and} \quad \mathbb{P} \left(\frac{\Delta_k^2}{\|\delta\|_2^2} \geq L_n \right) \rightarrow 1.$$

It follows that $\mathbb{P} \left(\frac{\Delta_k^2}{\|\delta\|_2^2} \geq \frac{ck}{\operatorname{tr}(\Sigma)} \right) \rightarrow 1$ for any positive constant $c < (1 - \sqrt{a})^2$, which completes the proof of the lower bound (11a). \square

B.3 Proof of upper bound (11b) in Proposition 1

As in the proof of the lower bound 11a in Appendix B.2, we may reduce to the case that $\frac{\Delta_k^2}{\|\delta\|_2^2} = \frac{1}{p} Z^\top P_k (P_k^\top \Sigma P_k)^{-1} P_k^\top Z$. Conditioning on P_k^\top , Lemma 3 gives a lower bound on the conditional probability

$$\mathbb{P} \left[\frac{\Delta_k^2}{\|\delta\|_2^2} \leq \frac{1}{p} \text{tr} (P_k (P_k^\top \Sigma P_k)^{-1} P_k^\top) + \psi(s_1) + \phi(s_1) \mid P_k^\top \right] \geq 1 - \exp(-s_1), \quad (50)$$

where s_1 is a positive real number that may vary with n , and we define

$$\psi(s_1) := \frac{2\sqrt{s_1}}{p} \left\| P_k (P_k^\top \Sigma P_k)^{-1} P_k^\top \right\|_F, \quad \phi(s_1) := \frac{2s_1}{p} \left\| P_k (P_k^\top \Sigma P_k)^{-1} P_k^\top \right\|_2. \quad (51)$$

Clearly, $\phi(s_1) \leq \sqrt{s_1} \psi(s_1)$. Again, as in the proof of the lower bound (11a), we let U_n denote an upper bound whose dependence on n will be specified below, and we define an event

$$\mathcal{D}_n := \left\{ \frac{1}{p} \text{tr} (P_k (P_k^\top \Sigma P_k)^{-1} P_k^\top) + (1 + \sqrt{s_1}) \psi(s_1) \leq U_n \right\}, \quad (52)$$

and integrate with respect to P_k^\top to obtain

$$\mathbb{P} \left(\frac{\Delta_k^2}{\|\delta\|_2^2} \leq U_n \right) \geq [1 - \exp(-s_1)] \mathbb{P}(\mathcal{D}_n). \quad (53)$$

Continuing along the parallel line of reasoning, we upper-bound $\frac{1}{p} \text{tr} (P_k (P_k^\top \Sigma P_k)^{-1} P_k^\top)$ on an event \mathcal{D}'_n (defined below) with $\mathbb{P}(\mathcal{D}'_n) \rightarrow 1$, and re-use the upper bound of $\psi(s_1)$ on the event \mathcal{E}''_n (see line (48)), which was shown to satisfy $\mathbb{P}(\mathcal{E}''_n) \rightarrow 1$. Then, we choose U_n so that $\mathcal{D}_n \supset \mathcal{D}'_n \cap \mathcal{E}''_n$, yielding $\mathbb{P}(\mathcal{D}_n) \rightarrow 1$. Lastly, we take $s_1 \rightarrow \infty$ at an appropriate rate so that the probability in line (53) tends to 1.

To define the event \mathcal{D}'_n for upper-bounding $\frac{1}{p} \text{tr} (P_k (P_k^\top \Sigma P_k)^{-1} P_k^\top)$, note that for a symmetric matrix A with rank k , Jensen's inequality implies $\text{tr}(A) \leq \sqrt{k \text{tr}(A^2)}$, regardless of the size of A . Considering $A = P_k (P_k^\top \Sigma P_k)^{-1} P_k^\top$, and $\sqrt{\text{tr}(A^2)} = \|A\|_F$, we see that we may choose $\mathcal{D}'_n = \mathcal{E}''_n$ from line (48), and on this set we have the inequality,

$$\frac{1}{p} \text{tr} (P_k (P_k^\top \Sigma P_k)^{-1} P_k^\top) \leq \sqrt{k} \frac{C\sqrt{k}}{p \lambda_{\min}(\Sigma)}, \quad (54)$$

with probability tending to 1 as $n \rightarrow \infty$, as long as C is strictly greater than $\frac{(1+\sqrt{a})^2}{(1-\sqrt{a})^2}$. In order to guarantee the inclusion $\mathcal{D}_n \supset \mathcal{D}'_n \cap \mathcal{E}''_n$, we define

$$U_n := \frac{Ck}{p \lambda_{\min}(\Sigma)} \left[1 + (s_1 + \sqrt{s_1}) \frac{2}{\sqrt{k}} \right]. \quad (55)$$

Note that $k = \lfloor n/2 \rfloor$ implies $k \rightarrow \infty$ as $n \rightarrow \infty$, so choosing $s_1 = k^{1/4}$ ensures that $s_1 \rightarrow \infty$ and the second term inside the brackets in line (55) vanishes. Combining lines (53) and (55), we have

$$\mathbb{P} \left(\frac{\Delta_k^2}{\|\delta\|_2^2} \leq U_n \right) \rightarrow 1, \quad \text{and} \quad U_n \frac{p \lambda_{\min}(\Sigma)}{Ck} \rightarrow 1.$$

It follows that $\mathbb{P} \left(\frac{\Delta_k^2}{\|\delta\|_2^2} \leq \frac{Ck}{p \lambda_{\min}(\Sigma)} \right) \rightarrow 1$ for any constant C strictly greater than $\frac{(1+\sqrt{a})^2}{(1-\sqrt{a})^2}$, which completes the proof of the upper bound (11b). \square

References

- [1] Y. Lu, P. Liu, P. Xiao, and H. Deng. Hotelling’s T2 multivariate profiling for detecting differential expression in microarrays. *Bioinformatics*, 21(14):3105–3113, Jul 2005.
- [2] J. J. Goeman and P. Bühlmann. Analyzing gene expression data in terms of gene sets: methodological issues. *Bioinformatics*, 23(8):980–987, Apr 2007.
- [3] D. V. D. Ville, T. Blue, and M. Unser. Integrated wavelet processing and spatial statistical testing of fmri data. *Neuroimage*, 23(4):1472–1485, 2004.
- [4] U. Ruttimann et al. Statistical analysis of functional mri data in the wavelet domain. *IEEE Transactions on Medical Imaging*, 17(2):142–154, 1998.
- [5] Z. Bai and H. Saranadasa. Effect of high dimension: by an example of a two sample problem. *Statistica Sinica*, 6:311,329, 1996.
- [6] M. S. Srivastava and M. Du. A test for the mean vector with fewer observations than the dimension. *Journal of Multivariate Analysis*, 99:386–402, 2008.
- [7] M. S. Srivastava. A test for the mean with fewer observations than the dimension under non-normality. *Journal of Multivariate Analysis*, 100:518–532, 2009.
- [8] S. X. Chen and Y. L. Qin. A two-sample test for high-dimensional data with applications to gene-set testing. *Annals of Statistics*, 38(2):808–835, Feb 2010.
- [9] S. Cléménçon, M. Depecker, and Vayatis N. AUC optimization and the two-sample problem. In *Advances in Neural Information Processing Systems (NIPS 2009)*, 2009.
- [10] L. Jacob, P. Neuvial, and S. Dudoit. Gains in power from structured two-sample tests of means on graphs. Technical Report arXiv:q-bio/1009.5173v1, arXiv, 2010.
- [11] A. Gretton, K. M. Borgwardt, M. Rasch, B. Schölkopf, and A.J. Smola. A kernel method for the two-sample-problem. In B. Schölkopf, J. Platt, and T. Hoffman, editors, *Advances in Neural Information Processing Systems 19*, pages 513–520. MIT Press, Cambridge, MA, 2007.
- [12] Z. Harchaoui, F. Bach, and E. Moulines. Testing for homogeneity with kernel Fisher discriminant analysis. In John C. Platt, Daphne Koller, Yoram Singer, and Sam T. Roweis, editors, *NIPS*. MIT Press, 2007.
- [13] R. J. Muirhead. *Aspects of Multivariate Statistical Theory*. John Wiley & Sons, inc., 1982.
- [14] S. S. Vempala. *The Random Projection Method*. DIMACS Series in Discrete Mathematics and Theoretical Computer Science. American Mathematical Society, 2004.
- [15] P. Diaconis and D. Freedman. Asymptotics of graphical projection pursuit. *Annals of Statistics*, 12(3):793–815, 1984.
- [16] A. W. van der Vaart. *Asymptotic Statistics*. Cambridge, 2007.
- [17] G. Tang and A. Nehorai. The stability of low-rank matrix reconstruction: a constrained singular value view. *arXiv:1006.4088, submitted to IEEE Trans. Information Theory*, 2010.

- [18] I. Johnstone. On the distribution of the largest eigenvalue in principal components analysis. *Annals of Statistics*, 29(2):295–327, 2001.
- [19] L. Wasserman. *All of Non-Parametric Statistics*. Springer Series in Statistics. Springer-Verlag, New York, NY, 2006.
- [20] I. Bechar. A Bernstein-type inequality for stochastic processes of quadratic forms of Gaussian variables. Technical Report arXiv:0909.3595v1, arXiv, 2009.
- [21] B. Laurent and P. Massart. Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics*, 28(5):1302–1338, 2000.
- [22] G. W. Stewart. The Efficient Generation of Random Orthogonal Matrices with an Application to Condition Estimators. *SIAM Journal on Numerical Analysis*, 17(3):403–409, 1980.
- [23] R. W. Tothill et al. Novel molecular subtypes of serous and endometrioid ovarian cancer linked to clinical outcome. *Clin Cancer Res*, 14(16):5198–5208, Aug 2008.
- [24] J. Moreaux et al. A high-risk signature for patients with multiple myeloma established from the molecular classification of human myeloma cell lines. *Haematologica*, 96(4):574–582, Apr 2011.
- [25] R. N. Jorissen et al. Metastasis-associated gene expression changes predict poor outcomes in patients with dukes stage b and c colorectal cancer. *Clin Cancer Res*, 15(24):7642–7651, Dec 2009.
- [26] A. Subramanian et al. Gene set enrichment analysis: a knowledge-based approach for interpreting genome-wide expression profiles. *Proc. Natl. Acad. Sci. USA*, 102(43):15545–15550, Oct 2005.
- [27] T. Beissbarth and T. P. Speed. Gostat: find statistically overrepresented gene ontologies within a group of genes. *Bioinformatics*, 20(9):1464–1465, Jun 2004.
- [28] A. S. Lewis J. M. Borwein. *Convex Analysis and Nonlinear Optimization Theory and Examples*. CMS Bookks in Mathematics. Canadian Mathematical Society, 2000.
- [29] K. R. Davidson and S. J. Szarek. Local operator theory, random matrices, and Banach spaces. in *Handbook of Banach Spaces*, 1, 2001.
- [30] P. Massart. *Concentration Inequalities and Model Selection*. Lecture Notes in Mathematics: Ecole d’Eté de Probabilités de Saint-Flour XXXIII-2003. Springer, Berlin, Heidelberg, 2007.
- [31] W. Beckner. A generalized Poincaré inequality for Gaussian measures. *Proceedings of the American Mathematical Society*, 105(2):397–400, 1989.
- [32] R. Horn and C. Johnson. *Matrix Analysis*. Cambridge University Press, 22nd printing edition, 2009.