



Published in final edited form as:

Ann Stat. 2017 June ; 45(3): 1342–1374. doi:10.1214/16-AOS1487.

Asymptotics of empirical eigenstructure for high dimensional spiked covariance

Weichen Wang[†] and Jianqing Fan^{†,‡}

[†]Princeton University

[‡]Fudan University

Abstract

We derive the asymptotic distributions of the spiked eigenvalues and eigenvectors under a generalized and unified asymptotic regime, which takes into account the magnitude of spiked eigenvalues, sample size, and dimensionality. This regime allows high dimensionality and diverging eigenvalues and provides new insights into the roles that the leading eigenvalues, sample size, and dimensionality play in principal component analysis. Our results are a natural extension of those in Paul (2007) to a more general setting and solve the rates of convergence problems in Shen et al. (2013). They also reveal the biases of estimating leading eigenvalues and eigenvectors by using principal component analysis, and lead to a new covariance estimator for the approximate factor model, called shrinkage principal orthogonal complement thresholding (S-POET), that corrects the biases. Our results are successfully applied to outstanding problems in estimation of risks of large portfolios and false discovery proportions for dependent test statistics and are illustrated by simulation studies.

Keywords and phrases

Asymptotic distributions; Principal component analysis; Spiked covariance model; Diverging eigenvalues; Approximate factor model; Relative risk management; False discovery proportion

1. Introduction

Principal Component Analysis (PCA) is a powerful tool for dimensionality reduction and data visualization. Its theoretical properties such as consistency and asymptotic distributions of the empirical eigenvalues and eigenvectors are challenging especially in high dimensional regime. For the past half century, substantial amount of efforts have been devoted to understanding empirical eigen-structure. An early effort came from Anderson (1963) who established the asymptotic normality of sample eigenvalues and eigenvectors under the

Correspondence to: Jianqing Fan.

Dept of Operations Research & Financial Engineering, Sherrerd Hall, Princeton University, Princeton, NJ 08544, USA, weichenw@princeton.edu

School of Data Science, Fudan University, Shanghai, China, jqfan@princeton.edu

SUPPLEMENTARY MATERIAL

Supplement: Technical proofs Wang and Fan (2015) (). This document contains technical lemmas for Section 3 and the comparison of assumptions and theoretical proofs for Section 4.

classical regime with large sample size n and fixed dimension p . However, when dimensionality diverges at the same rate as the sample size, sample covariance matrix is a notoriously bad estimator with dramatically different eigen-structure from the population one. A lot of recent literature makes the endeavor to understand the behavior of the empirical eigenvalues and eigenvectors under high dimensional regime where both n and p go to infinity. See for example Baik, Ben Arous and P  ch   (2005); Bai (1999); Paul (2007); Johnstone and Lu (2009); Onatski (2012); Shen et al. (2013) and many related papers. For additional developments and references, see Bai and Silverstein (2009).

Among different structures of population covariance, the spiked covariance model is of great interest. It typically assumes several eigenvalues larger than the remaining, and focuses on recovering only these leading eigenvalues and their associated eigenvectors. The spiked part is of importance, as we are usually interested in the directions that explain the most variations of the data. In this paper, we consider a high dimensional spiked covariance model with the leading eigenvalues larger than the rest. We provide new understanding on how the spiked empirical eigenvalues and eigenvectors fluctuate around their theoretical counterparts and what their asymptotic biases are. Three quantities play an essential role in determining the asymptotic behavior of empirical eigen-structure: the sample size n , the dimension p , and the magnitude of leading eigenvalues $\{\lambda_j\}_{j=1}^m$. The natural question to ask is how the asymptotics of empirical eigen-structure depends on the interplay of those quantities. We will give a unified answer to this important question in the principal component analysis. Theoretical properties of PCA have been investigated from three different perspectives: (i) random matrix theories, (ii) sparse PCA and (iii) approximate factor model.

The first angle to analyze PCA is through random matrix theories, where it is typically assumed $p/n \rightarrow \gamma \in (0, \infty)$ with bounded spike sizes. It is well known that if the true covariance matrix is identity, the empirical spectral distribution converges almost surely to the Marcenko-Pastur distribution (Bai, 1999) and when $\gamma < 1$ the largest and smallest eigenvalues converge almost surely to $(1 + \sqrt{\gamma})^2$ and $(1 - \sqrt{\gamma})^2$ respectively (Bai and Yin, 1993; Johnstone, 2001). If the true covariance structure takes the form of a spiked matrix, Baik, Ben Arous and P  ch   (2005) showed that the asymptotic distribution of the top empirical eigenvalue exhibits an $n^{2/3}$ scaling when the eigenvalue lies below a threshold $1 + \sqrt{\gamma}$, and an $n^{1/2}$ scaling when it is above the threshold (named BBP phase transition after the authors). The phase transition is further studied by Benaych-Georges and Nadakuditi (2011) and Bai and Yao (2012) under more general assumptions. For the case where we have the regular scaling, Paul (2007) investigated the asymptotic behavior of the corresponding empirical eigenvectors and showed that the major part of an eigenvector is normally distributed with a regular scaling $n^{1/2}$. The convergence of principal component scores under this regime was considered by Lee, Zou and Wright (2010). The same random matrix regime has also been considered by Onatski (2012) in studying the principal component estimator for high-dimensional factor models. More recently, Koltchinskii and Lounici (2014a,b) revealed a profound link of concentration bounds of empirical eigen-structure with the effective rank defined as $\bar{r} = \text{tr}(\mathbf{\Sigma})/\lambda_1$ (Vershynin, 2010). Their results extend the regime of bounded eigenvalues to a more general setting, although the asymptotic results in most cases still rely on the assumption $\bar{r} = o(n)$, which essentially requires a low dimensionality, i.e. p/n

$\rightarrow 0$, if λ_1 is bounded. In this paper, we consider the general regime of bounded $p/(n\lambda_1)$, which implies $\bar{r} = O(n)$ and allows diverging λ_1 . More discussions will be given in Section 3.

A second line of efforts is through sparse PCA. According to Johnstone and Lu (2009), PCA does not generate consistent estimators for leading eigenvectors if $p/n \rightarrow \gamma \in (0, 1)$ with bounded eigenvalues. This motivates the development of sparse PCA, which leverages the extra assumption on the sparsity of eigenvectors. A large amount of literature has contributed to the topic of sparse PCA, for example Amini and Wainwright (2008); Vu and Lei (2012); Birnbaum et al. (2013); Berthet and Rigollet (2013); Ma (2013). Specifically, Vu and Lei (2012) derived optimal bound for the minimax estimation error of the first sparse leading eigenvector, while Cai, Ma and Wu (2015) conducted a more thorough study on the minimax optimal rates for estimating top eigenvalues and eigenvectors of spiked covariance matrices with jointly k -sparse eigenvectors. This type of work assumes bounded eigenvalues, which ignore the contributions of the strong signals from the data in many real applications. To make the problem solvable, sparsity assumptions on the eigenvectors are imposed. In contrast, driven by applications such as genomics, economics, and finance, this paper studies the contributions of the diverging eigenvalues (signals) to the estimation of their associated eigenvectors, without relying on sparsity assumptions on the eigenvectors.

In order to illustrate the third perspective, let us briefly review the approximate factor model (Bai, 2003; Fan, Liao and Mincheva, 2013) and see how the spiked eigenvalues arise naturally from the model. Consider the following data generating model:

$$\mathbf{y}_t = \mathbf{B}\mathbf{f}_t + \boldsymbol{\varepsilon}_t, \quad \text{for } t=1, \dots, T,$$

where \mathbf{y}_t is a p -dimensional vector observed at time t , $\mathbf{f}_t \in \mathbb{R}^m$ is the vector of latent factors that drive the cross-sectional dependence at time t , \mathbf{B} is the matrix of the corresponding loading coefficients, and $\boldsymbol{\varepsilon}_t$ is the idiosyncratic part that can not be explained by the factors. Assume without loss of generality that $\text{var}(\mathbf{f}_t) = \mathbf{I}_m$, the $m \times m$ identity matrix. Then, the model implies $\boldsymbol{\Sigma} = \text{var}(\mathbf{y}_t) = \mathbf{B}\mathbf{B}' + \boldsymbol{\Sigma}_e$, where $\boldsymbol{\Sigma}_e = \text{var}(\boldsymbol{\varepsilon})$. It admits a low-rank plus sparse structure, when $\boldsymbol{\Sigma}_e$ is assumed to be sparse (Fan, Fan and Lv, 2008; Fan, Liao and Mincheva, 2013). The recovery of the low-rank and sparse matrices was considered thoroughly by Candès et al. (2011) and Chandrasekaran et al. (2011) under the incoherence condition in the noise-less setting and by Agarwal, Negahban and Wainwright (2012) in the noisy case. If the factor loadings $\{\mathbf{b}_j\}_j$ (the transpose of rows of \mathbf{B}) are i.i.d. samples from a population with mean zero and covariance $\boldsymbol{\Sigma}_b$ [this is a pervasive assumption commonly used in the factor models (Fan, Liao and Mincheva, 2013)], then by the law of large numbers,

$p^{-1}\mathbf{B}'\mathbf{B} = p^{-1}\sum_{i=1}^p \mathbf{b}_i\mathbf{b}_i' \rightarrow \boldsymbol{\Sigma}_b$, as $p \rightarrow \infty$. In other words, the eigenvalues of $\mathbf{B}\mathbf{B}'$ are approximately

$$p\lambda_1(\boldsymbol{\Sigma}_b)(1+o(1)), \dots, p\lambda_m(\boldsymbol{\Sigma}_b)(1+o(1)), 0, \dots, 0,$$

where $\lambda_j(\Sigma_b)$ is the j^{th} eigenvalue of Σ_b . Then, by Weyl's theorem, we conclude that the eigenvalues of Σ

$$\lambda_j = p\lambda_j(\Sigma_b)(1+o(1)), \quad \text{for } j=1, \dots, m, \quad (1.1)$$

and the remaining are bounded, if $\|\Sigma_b\|$ is bounded. Therefore, the factor model implies a spiked covariance with diverging leading eigenvalues. Fan, Liao and Mincheva (2013) showed that if the leading eigenvalues grow linearly with the dimension, then the corresponding eigenvectors can be consistently estimated as long as sample size goes to infinity. See Section 4 for more details.

Deviating from the classical random matrix and sparse PCA literature, we consider the high dimensional regime, allowing $p/n \rightarrow \infty$. To take into account the contributions of the signals for PCA, we allow $\lambda_j \rightarrow \infty$ for the first m leading eigenvalues. This leads to the third perspective for understanding PCA from this high dimensional setting. Shen et al. (2013) adopted this point of view and considered the regime of $p/(n\lambda_j) \rightarrow \gamma_j$ where $0 < \gamma_j < \infty$ for leading eigenvalues. This is more general than the bounded eigenvalue condition. Specifically if eigenvalues are bounded, we require the ratio p/n converges to a bounded constant as in the random matrix regime. On the other hand, if the dimension is much larger than the sample size, we offset the dimensionality by assuming increased signals or sample size, without additional sparse eigenvector assumption as in sparse PCA regime. In particular, as shown in (1.1), the strong (or pervasive) factors considered in financial applications corresponds to $\gamma_j = 0$ with the leading eigenvalues $\lambda_j \asymp p$; see for example Stock and Watson (2002); Bai (2003); Bai and Ng (2002); Fan, Liao and Mincheva (2013); Fan, Liao and Wang (2016). The weak or semi-strong factors considered by De Mol, Giannone and Reichlin (2008) and Onatski (2012) also imply bounded $p/(n\lambda_1)$, with p/n bounded and $\lambda_j \asymp p^\theta$ for some $\theta \in [0, 1)$.

Hall, Marron and Neeman (2005); Jung and Marron (2009) started the research of high dimension low sample size (HDLSS) regime. With n fixed, Jung and Marron (2009) concluded that consistency of leading eigenvalues and eigenvectors is granted if $\lambda_j \asymp p^\theta$ for $\theta > 1$, which also corresponds to $\gamma_j = 0$. Shen et al. (2013) revealed an interesting fact that when $\gamma_j = 0$, spiked sample eigenvalues almost surely converge to a biased quantity of the true eigenvalues; furthermore the corresponding sample eigenvectors show an asymptotic conical structure. However, their work focuses only on the consistency problem. In this study, we will consider the same regime as theirs, but focus more on the rates of convergence and the asymptotic distributions of the empirical eigen-structure, and under more relaxed conditions. Our results can be viewed as a natural extension of Paul (2007) to the high dimensional setting.

We would like to emphasize more on the scope and importance of our contributions here. Firstly, the regime we consider in this paper is $p/(n\lambda_j) \rightarrow \gamma_j \in [0, \infty)$ for $j = 1, \dots, m$, which permits high dimensionality $p/n \rightarrow \infty$ and diverging eigenvalues without specifying their divergence rates. As we have argued, this encompasses many situations considered in the existing literature. It puts into the same framework of the typical random matrix regime with

bounded eigenvalues and HDLSS analysis with fixed sample size. Secondly, the contributions of diverging eigenvalues are indeed recognized and accounted for in our theoretical developments. This avoids the restrictive assumptions on sparse eigenvectors. PCA without sparsity assumptions has been widely used in the diverging fields such as population association study (Yamaguchi-Kabata et al., 2008), genome-wide association study (Ringnér, 2008), microarray data (Landgrebe, Wurst and Welzl, 2002; Price et al., 2006), fMRI data (Thomas, Harshman and Menon, 2002), and financial returns (Chen and Shimerda, 1981; Chamberlain and Rothschild, 1983). Our efforts contribute to theoretical understanding why such a plain PCA works in these diverse fields. Finally, by allowing certain generality, we gain theoretical insights into how n , p and signal strength λ_j interplay.

The results are useful in two ways. On the one hand, they help quantify the biases of empirical eigen-structure and explain where they come from. Specifically, in Theorem 3.1, the bias of the j^{th} sample eigenvalue ($j = m$) is quantified by $p/(n\lambda_j)$, which is also showed by Yata and Aoshima (2012, 2013) under different assumptions of the spiked covariance model. Our novel contribution lies in Theorem 3.2, revealing the bias of the j^{th} sample eigenvector ($j = m$). In (3.7), we provide a decomposition of each empirical eigenvector into a spiked part, which converges to the true eigenvector with a deflation factor also quantified by $p/(n\lambda_j)$, and a non-spiked part, which creates a random bias distributed uniformly on an ellipse. More details will be presented in Section 3. On the other hand, the theoretical results provide new technical tools for analyzing factor models, which motivate the study. As we have seen, although it is natural to assume eigenvalues grow linearly with dimension, the assumption imposes a strong signal. Note that when $p/(n\lambda_j) \rightarrow 0$, no biases will occur. So in Section 4, we consider to relax the order of spikes to slightly faster than \sqrt{p} . By correcting the biases, we propose a new method called Shrinkage Principal Orthogonal complement Thresholding (S-POET) and employ it to two applications: risk assessment of large portfolios (Pesaran and Zaffaroni, 2008; Fan, Liao and Shi, 2015) and false discovery proportion estimation for dependent test statistics (Leek and Storey, 2008; Fan, Han and Gu, 2012). Existing methodologies for those two problems reply on rather strong signal level, but we are able to relax it with the help of S-POET.

The paper is organized as follows. Section 2 introduces the notations, assumptions, and an important fact which serves as basis of our proofs. Sections 3.1 and 3.2 devote to the theoretical results for the sample eigenvalues and eigenvectors of the spiked covariance matrix. In Section 4, we discuss several applications of the theories in the previous section. Simulations are conducted in Section 5 to demonstrate the theoretical results at the finite sample and the performance of S-POET. Section 6 provides concluding remarks. The proofs for Section 3 are provided in the appendix and those for Section 4 are relegated to the supplementary material.

2. Assumptions and a simple fact

Assume that $\{Y_i\}_{i=1}^n$ is a sequence of i.i.d. random variables with zero mean and covariance matrix $\Sigma_{p \times p}$. Let $\lambda_1, \dots, \lambda_p$ be the eigenvalues of Σ in descending order. We consider the spiked covariance model as follows.

Assumption 2.1

$\lambda_1 > \lambda_2 > \dots > \lambda_m > \lambda_{m+1} \dots \lambda_p > 0$, where the non-spiked eigenvalues are bounded, i.e. $c_0 \leq \lambda_j \leq C_0, j > m$ for constants $c_0, C_0 > 0$ and the spiked eigenvalues are well separated, i.e. $\exists \delta_0 > 0$ such that $\min_{j \leq m} (\lambda_j - \lambda_{j+1}) / \lambda_j \geq \delta_0$.

The eigenvalues are divided into the spiked ones and bounded non-spiked ones. We do not need to specify the order of the leading eigenvalues nor require them to diverge. Thus, our results in Section 3 are applicable to both bounded and diverging leading eigenvalues. For simplicity, we only consider distinguishable eigenvalues (multiplicity 1) for the largest m eigenvalues and a fixed number m , independent of n and p .

The factor model $\mathbf{y} = \mathbf{B}\mathbf{f} + \mathbf{e}$ with pervasive factors considered in Fan, Liao and Mincheva (2013) implies a spiked covariance model with $\lambda_j \asymp p$ in (1.1) and satisfies the above assumption. For the interplay of the sample size n , dimension p and the spikes λ_j 's, the following relationship is assumed as in Shen et al. (2013).

Assumption 2.2

Assume $p > n$. For the spiked part $1 \leq j \leq m, c_j = p/(n\lambda_j)$ is bounded, and for the non-spiked part, $(p - m)^{-1} \sum_{j=m+1}^p \lambda_j = \bar{c} + o(n^{-1/2})$.

We allow $p/n \rightarrow \infty$ in any manner, though λ_j also needs to grow fast enough to ensure bounded c_j . In particular, $c_j = \alpha(1)$ is allowed as in the factor model. We do not assume the non-spiked eigenvalues are identical, as in most spiked covariance model literature (e.g. Paul (2007); Johnstone and Lu (2009)).

By spectral decomposition, $\Sigma = \Gamma\Lambda\Gamma'$, where the orthonormal matrix Γ is constructed by the eigenvectors of Σ and $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_p)$. Let $\mathbf{X}_i = \Gamma' \mathbf{Y}_i$. Since the empirical eigenvalues are invariant and the empirical eigenvectors are equivariant under an orthonormal transformation, we focus the analysis on the transformed domain of \mathbf{X}_i and then translate the results into those of the original data. Note that $\text{var}(\mathbf{X}_i) = \Lambda$. Let $\mathbf{Z}_i = \Lambda^{-1/2} \mathbf{X}_i$ be the elementwise standardized random vector.

Assumption 2.3

$\{\mathbf{Z}_i\}_{i=1}^n$ are i.i.d copies of \mathbf{Z} . The standardized random vector $\mathbf{Z} = (Z_1, \dots, Z_p)$ is sub-Gaussian with independent entries of mean zero and variance one. The sub-Gaussian norms of all components are uniformly bounded: $\max_j \|Z_j\|_{\Psi_2} \leq C_0$, where $\|Z_j\|_{\Psi_2} = \sup_{q \geq 1} q^{-1/2} (E|Z_j|^q)^{1/q}$.

Since $\text{Var}(\mathbf{X}_i) = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_p)$, the first m population eigenvectors are simply unit vectors $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_m$ with only one nonvanishing element. Denote the n by p transformed data matrix by $\mathbf{X} = (\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_n)'$. Then the sample covariance matrix is

$$\hat{\Sigma}_{p \times p} = \frac{1}{n} \mathbf{X}' \mathbf{X} = \frac{1}{n} \sum_{i=1}^n \mathbf{X}_i \mathbf{X}_i'$$

whose eigenvalues are denoted as $\hat{\lambda}_1, \hat{\lambda}_2, \dots, \hat{\lambda}_p$ ($\hat{\lambda}_j = 0$ for $j > n$) with corresponding eigenvectors $\hat{\xi}_1, \hat{\xi}_2, \dots, \hat{\xi}_p$. Note that the empirical eigenvectors of data \mathbf{Y}_i 's are $\hat{\xi}_j^{(Y)} = \Gamma \hat{\xi}_j$.

Let \mathbf{Z}_j be the j^{th} column of the standardized \mathbf{X} . Then each \mathbf{Z}_j has i.i.d sub-Gaussian entries with zero mean and unit variance. Exchanging the roles of rows and columns, we get the n by n Gram matrix

$$\tilde{\Sigma}_{n \times n} = \frac{1}{n} \mathbf{X} \mathbf{X}' = \frac{1}{n} \sum_{j=1}^p \lambda_j \mathbf{Z}_j \mathbf{Z}_j',$$

with the same nonzero eigenvalues $\hat{\lambda}_1, \hat{\lambda}_2, \dots, \hat{\lambda}_n$ as $\hat{\Sigma}$ and the corresponding eigenvectors $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n$. It is well known that for $i = 1, 2, \dots, n$

$$\hat{\xi}_i = (n \hat{\lambda}_i)^{-1/2} \mathbf{X}' \mathbf{u}_i \quad \text{and} \quad \mathbf{u}_i = (n \hat{\lambda}_i)^{-1/2} \mathbf{X} \hat{\xi}_i, \quad (2.1)$$

while the other eigenvectors of $\hat{\Sigma}$ constitute a $(p-n)$ -dimensional orthogonal complement of $\hat{\xi}_1, \dots, \hat{\xi}_n$.

By using this simple fact, for the specific case with $c_0 = C_0 = 1$ in Assumption 2.1, $\lambda_j = 1$ for $j > m$ in Assumption 2.2, and Gaussian data in Assumption 2.3, Shen et al. (2013) showed that

$$\frac{\hat{\lambda}_j}{\lambda_j} \xrightarrow{\text{a.s.}} 1 + c_j, \quad 1 \leq j \leq m;$$

and

$$\left| \langle \hat{\xi}_j, \mathbf{e}_j \rangle \right| \xrightarrow{\text{a.s.}} (1 + c_j)^{-\frac{1}{2}},$$

where $\langle \mathbf{a}, \mathbf{b} \rangle$ denotes the inner product of two vectors. However, they fail to establish any results on convergence rates or asymptotic distributions of the empirical eigen-structure. This motivates the current paper.

The aim of this paper is to establish the asymptotic normality of the empirical eigenvalues and eigenvectors under more relaxed conditions. Our results are a natural extension of Paul (2007) to a more general setting with new insights, where the asymptotic normality of sample eigenvectors is derived using complicated random matrix techniques for Gaussian data under the regime of $p/n \rightarrow \gamma \in [0, 1)$. In comparison, our proof, based on the relationship (2.1), is much simpler and insightful for understanding the behavior of high dimensional PCA.

Here are some notations that we will use in the paper. For a general matrix \mathbf{M} , we denote its matrix entry-wise max norm as $\|\mathbf{M}\|_{\max} = \max_{i,j} \{|M_{i,j}|\}$ and define the quantities

$\|\mathbf{M}\| = \lambda_{\max}^{1/2}(\mathbf{M}'\mathbf{M})$, $\|\mathbf{M}\|_F = (\sum_{i,j} M_{i,j}^2)^{1/2}$, $\|\mathbf{M}\|_{\infty} = \max_i \sum_j |M_{i,j}|$ to be its spectral, Frobenius and induced ℓ_{∞} norms. If \mathbf{M} is symmetric, we define $\lambda_j(\mathbf{M})$ to be the j^{th} largest eigenvalue of \mathbf{M} and $\lambda_{\max}(\mathbf{M})$, $\lambda_{\min}(\mathbf{M})$ to be the maximal and minimal eigenvalues respectively. We denote $\text{tr}(\mathbf{M})$ as the trace of \mathbf{M} . For any vector \mathbf{v} , its ℓ_2 norm is represented by $\|\mathbf{v}\|$ while ℓ_1 norm is written as $\|\mathbf{v}\|_1$. We use $\text{diag}(\mathbf{v})$ to denote the diagonal matrix with the same diagonal entries as \mathbf{v} . For two random vectors \mathbf{a} , \mathbf{b} of the same length, we say $\mathbf{a} = \mathbf{b} + O_P(\delta)$ if $\|\mathbf{a} - \mathbf{b}\| = O_P(\delta)$ and $\mathbf{a} = \mathbf{b} + o_P(\delta)$ if $\|\mathbf{a} - \mathbf{b}\| = o_P(\delta)$. We denote $\mathbf{a} \stackrel{d}{\rightarrow} \mathcal{L}$ for some distribution \mathcal{L} if there exists $\mathbf{b} \sim \mathcal{L}$ such that $\mathbf{a} = \mathbf{b} + o_P(1)$. Throughout the paper, C is a generic constant that may differ from line to line.

3. Asymptotic behavior of empirical eigen-structure

3.1. Asymptotic normality of empirical eigenvalues

Let us first study the behavior of the leading m empirical eigenvalues of $\hat{\Sigma}$. Denote by $\lambda_j(\mathbf{A})$ the j^{th} largest eigenvalue of matrix \mathbf{A} and recall that $\hat{\lambda}_j = \lambda_j(\hat{\Sigma})$. We have the following asymptotic normality of $\hat{\lambda}_j$.

Theorem 3.1—*Under Assumptions 2.1 – 2.3, $\{\hat{\lambda}_j\}_{j=1}^m$'s have independent limiting distributions. In addition,*

$$\sqrt{n} \left\{ \frac{\hat{\lambda}_j}{\lambda_j} - \left(1 + \bar{c}c_j + O_P(\lambda_j^{-1} \sqrt{p/n}) \right) \right\} \stackrel{d}{\rightarrow} N(0, \kappa_j - 1), \tag{3.1}$$

where κ_j is the kurtosis of X_j .

The theorem shows that the bias of $\hat{\lambda}_j/\lambda_j$ is $\bar{c}c_j + O_P(\lambda_j^{-1} \sqrt{p/n})$. The second term is dominated by the first term since $p > n$ and it is of order $O_P(n^{-1/2})$ if $\sqrt{p} = o(\lambda_j)$. The latter assumption is satisfied by the strong factor model in Fan, Liao and Mincheva (2013) and a part of weak or semi-strong factor model in Onatski (2012). The theorem reveals the bias is controlled by a term of rate $p/(n\lambda_j)$. To get the asymptotically unbiased estimate, it requires $c_j = p/(n\lambda_j) \rightarrow 0$ for $j = 1, \dots, m$. This result is more general than that of Shen et al. (2013) and sheds a similar light to that of Koltchinskii and Lounici (2014a,b) i.e. $\|\hat{\Sigma} - \Sigma\|/\|\Sigma\| \rightarrow 0$ almost surely if and only if the effective rank $\bar{r} = \text{tr}(\Sigma)/\lambda_1$ is of order $\alpha(n)$, which is true when $c_1 = \alpha(1)$. Our result here holds for each individual spike. Yata and Aoshima (2012, 2013) employed a similar technical trick and gave a comprehensive study on the asymptotic consistency and distributions of the eigenvalues. They got similar results under different conditions from ours. Our framework is more general here. If $c_j \rightarrow 0$, bias reduction can also be made; see Section 4.2, where an estimator for \bar{c} is proposed. Under the bounded spiked covariance model considered in Baik, Ben Arous and P  ch   (2005), Johnstone and Lu (2009) and Paul (2007), it is assumed $\lambda_j = c_0 = C_0$, $j > m$ so that $\bar{c} = c_0$, the minimum eigenvalue of the population covariance matrix. Our result is also consistent with Anderson (1963)'s result that

$$\sqrt{n}(\hat{\lambda}_j - \lambda_j) \xrightarrow{d} N(0, 2\lambda_j^2),$$

for Gaussian data and fixed p and λ_j 's, where the non-spiked part does not exist and thus the bias $O_p(\lambda_j^{-1} \sqrt{p/n})$ disappears. The proof is relegated to the appendix.

3.2. Behavior of empirical eigenvectors

Let us consider the asymptotic distribution of the empirical eigenvectors $\hat{\boldsymbol{\xi}}_j$'s corresponding to $\hat{\lambda}_j, j = 1, 2, \dots, m$. As in Paul (2007), each $\hat{\boldsymbol{\xi}}_j$ is divided into two parts corresponding to the spiked and non-spiked components, i.e. $\hat{\boldsymbol{\xi}}_j = (\hat{\boldsymbol{\xi}}_{jA}', \hat{\boldsymbol{\xi}}_{jB}')'$ where $\hat{\boldsymbol{\xi}}_{jA}$ is of length m .

Theorem 3.2—*Under Assumptions 2.1 – 2.3, we have*

- i. *For the spiked part, if $m = 1$,*

$$\frac{2(1+\bar{c}c_1)}{\bar{c}c_1} \sqrt{n} \left(\sqrt{1+\bar{c}c_1} \hat{\xi}_{1A} - 1 + O_p \left(\sqrt{\frac{p}{n\lambda_1^2}} \right) \right) \xrightarrow{d} N(0, \kappa_1 - 1), \tag{3.2}$$

while if $m > 1$,

$$\sqrt{n} \left(\frac{\hat{\boldsymbol{\xi}}_{jA}}{\|\hat{\boldsymbol{\xi}}_{jA}\|} - \mathbf{e}_{jA} + O_p \left(\sqrt{\frac{p}{n\lambda_j^2}} \right) \right) \xrightarrow{d} N_m(\mathbf{0}, \boldsymbol{\Sigma}_j), \tag{3.3}$$

for $j = 1, 2, \dots, m$, with

$$\boldsymbol{\Sigma}_j = \sum_{k \in [m] \setminus j} a_{jk}^2 \mathbf{e}_{kA} \mathbf{e}_{kA}',$$

where $[m] = \{1, \dots, m\}$, \mathbf{e}_{kA} is the first m elements of the unit vector \mathbf{e}_k , and

$a_{jk} = \lim_{\lambda_j, \lambda_k \rightarrow \infty} \sqrt{\lambda_j \lambda_k} / (\lambda_j - \lambda_k)$, which is assumed to exist.

- ii. *For the non-spiked part, if we further assume the data is Gaussian, there exists $p - m$ dimensional vector $\mathbf{h}_0 \sim \text{Unif}(B_{p-m}(1))$ such that*

$$\left\| \mathbf{D}_0 \frac{\hat{\boldsymbol{\xi}}_{jB}}{\|\hat{\boldsymbol{\xi}}_{jB}\|} - \mathbf{h}_0 \right\| = O_p \left(\sqrt{\frac{n}{p}} \right) + o_p \left(\frac{1}{\sqrt{n}} \right), \tag{3.4}$$

where $\mathbf{D}_0 = \text{diag}(\sqrt{\bar{c}/\lambda_{m+1}}, \dots, \sqrt{\bar{c}/\lambda_p})$ is a diagonal matrix and $\text{Unif}(B_k(r))$ denotes the uniform distribution over the centered sphere of radius r . In addition, the max norm of $\hat{\boldsymbol{\xi}}_{jB}$ satisfies

$$\|\hat{\xi}_{jB}\|_{\max} = O_P \left(p/(n\lambda_j^{3/2}) + \sqrt{\log p/(n\lambda_j)} \right). \quad (3.5)$$

iii. *Furthermore, $\|\hat{\xi}_{jA}\| = (1 + \bar{c}c_j)^{-1/2} + O_P(\lambda_j^{-1} \sqrt{p/n + p/(n^{3/2}\lambda_j)})$ and $\|\hat{\xi}_{jB}\| = (\frac{\bar{c}c_j}{1 + \bar{c}c_j})^{1/2} + O_P(\sqrt{1/\lambda_j} + \sqrt{p/(n^2\lambda_j)})$. Together with (i), this implies the inner product between empirical eigenvector and the population one converges to $(1 + \bar{c}c_j)^{-1/2}$ in probability and*

$$\langle \hat{\xi}_j, \mathbf{e}_j \rangle - \frac{1}{\sqrt{1 + \bar{c}c_j}} = O_P \left(\lambda_j^{-1} \sqrt{p/n + p/(n^{3/2}\lambda_j)} \right) + O_P(n^{-1})I_{\{m > 1\}}. \quad (3.6)$$

In the above theory, we assume $a_{jk} = \lim_{\lambda_j, \lambda_k \rightarrow \infty} \frac{\sqrt{\lambda_j \lambda_k}}{\lambda_j - \lambda_k}$ exists. This is not restrictive if eigenvalues are well separated i.e. $\min_{j \neq k} |\lambda_j - \lambda_k|/\lambda_j \geq \delta_0$ from assumption 2.1. The assumption obviously holds for the pervasive factor model, in which

$$a_{jk} = \sqrt{\lambda_j(\Sigma_b)\lambda_k(\Sigma_b)/(\lambda_k(\Sigma_b) - \lambda_j(\Sigma_b))}.$$

Theorem 3.2 is an extension of random matrix results into high dimensional regime. Its proof sheds light on how to use the smaller $n \times n$ matrix $\tilde{\Sigma}$ as a tool to understand the behavior of the larger $p \times p$ covariance matrix $\hat{\Sigma}$. Specifically, we start from $\tilde{\Sigma}\mathbf{u}_j = \hat{\lambda}_j\mathbf{u}_j$ or identity (A.3) and then use the simple fact (2.1) to get a relationship (A.4) of eigenvector $\hat{\xi}_j$. Then (A.4) is rearranged as (A.6) which gives a clear separation of the dominating term, that is asymptotically normal, and the error term. This makes the whole proof much simpler in comparison with Paul (2007) who showed a similar type of result through a complicated representation of $\hat{\xi}_j$ and $\hat{\lambda}_j$ under more restricted assumptions. From this simple trick, we can understand deeply how some important high and low dimensional quantities link together and differ from each other.

Several remarks are in order. Firstly, since $\hat{\xi}_j^{(Y)} = \Gamma\hat{\xi}_j$ is the j^{th} empirical eigenvector based on observed data \mathbf{Y} , we have decomposition

$$\hat{\xi}_j^{(Y)} = \Gamma_A \hat{\xi}_{jA} + \Gamma_B \hat{\xi}_{jB}, \quad (3.7)$$

where $\Gamma = (\Gamma_A, \Gamma_B)$. Note that $\Gamma_A \hat{\xi}_{jA}$ converges to the true eigenvector deflated by a factor of $\sqrt{1 + \bar{c}c_j}$ with the convergence rate $O_P(\sqrt{p/(n\lambda_j^2)} + p/(n^{3/2}\lambda_j) + n^{-1/2})$ while $\Gamma_B \hat{\xi}_{jB}$ creates a random bias, which is distributed uniformly on an ellipse of $(p - m)$ dimension and projected into the p dimensional space spanned by Γ_B . The two parts intertwine in such a way that correction for the biases of estimating eigenvectors is almost impossible. More details are discussed in Section 4 for the factor models. Secondly, it is clearly as in the eigenvalue

case, the bias term $\lambda_j^{-1} \sqrt{p/n}$ Theorem 3.2 (i) disappears when $\sqrt{p} = o(\lambda_j)$. In particular, for the stronger factor given by (1.1), $\hat{\xi}_j^{(Y)}$ is a consistent estimator. Thirdly, the situations $m = 1$ and $m > 1$ have slight difference in that multiple spikes could interact with each other. Especially this is reflected in the convergence of the angle between the empirical eigenvector and its population counterpart: the angle converges to $(1 + \bar{c}_j)^{-1/2}$ with an extra rate $O_p(1/n)$ which stems from estimating $\hat{\xi}_{jk}$ for $j \neq k \leq m$ (see proof of Theorem 3.2 (iii)). The difference will only be seen when the spike magnitude is higher than the order $\sqrt{pn} \vee pn^{-1/2}$. We will verify this by a simple simulation in Section 5. Finally, it is the first time that the max norm bound of the non-spiked part is derived. This bound will be useful for analyzing factor models in Section 4.

Theorem 3.2 again implies the results of Shen et al. (2013). It also generalizes the asymptotic distribution of non-spiked part from pure orthogonal invariant case of Paul (2007) to a more general setting. In particular, when $p/n \rightarrow \infty$, the asymptotic distribution of the normalized non-spiked component is not uniform over a sphere any more, but over an ellipse. In addition, our result can be compared with the low dimensional case, where Anderson (1963) showed that

$$\sqrt{n} (\hat{\xi}_j - \mathbf{e}_j) \xrightarrow{d} N_p \left(\mathbf{0}, \sum_{k \in [m] \setminus j} \frac{\lambda_j \lambda_k}{(\lambda_j - \lambda_k)^2} \mathbf{e}_k \mathbf{e}_k' \right), \quad (3.8)$$

for fixed p and λ_j 's. Under our assumptions, since the spiked eigenvalues may go to infinity, the constants in the asymptotic covariance matrix are replaced by the limits a_{jk} 's. Similar to the behavior of eigenvalues, the spiked part $\hat{\xi}_{jA}$ preserves the normality property except for a bias factor $1/(1 + \bar{c}_j)$ caused by the high dimensionality. Also, recent work of Koltchinskii and Lounici (2014b) provides general asymptotic results for the empirical eigenvectors from a spectral projector point of view, but they mainly focus on the regime of $p/n\lambda_j \rightarrow 0$ or $\bar{r} = \alpha(n)$. Last but not least, it has been shown by Johnstone and Lu (2009) that PCA generates consistent eigenvector estimation if and only if $p/n \rightarrow 0$ when the spike sizes are fixed. This motivates the study of sparse PCA. We take the spike magnitude into account and provide additional insights by showing that PCA consistently estimate eigenvalues and eigenvectors if and only if $p/(n\lambda_j) \rightarrow 0$. This explains why Fan, Liao and Mincheva (2013) can consistently estimate the eigenvalues and eigenvectors while Johnstone and Lu (2009) can not.

4. Applications to factor models

In this section, we propose a method named Shrinkage Principal Orthogonal complement Thresholding (S-POET) for estimating large covariance matrices induced by the approximate factor models. The estimator is based on correction of the bias of the empirical eigenvalues as specified in (3.1). We derive for the first time the bound for the relative estimation errors of covariance matrices under the spectral norm. The results are then

applied to assessing large portfolio risks and estimating false discovery proportions, where the conditions in existing literature are significantly relaxed.

4.1. Approximate factor models

Factor models have been widely used in various disciplines. For example, it is used to extract information from financial market for sufficient forecasting of other time series (Stock and Watson, 2002; Fan, Xue and Yao, 2015) and to adjust heterogeneity for biological data aggregation of multiple sources (Leek et al., 2010; Fan et al., 2016). Consider the approximate factor model

$$y_{it} = \mathbf{b}_i' \mathbf{f}_t + u_{it}, \quad (4.1)$$

where y_{it} is the observed data for the i^{th} ($i = 1, \dots, p$) individual (e.g. returns of stocks) or component (e.g. expressions of genes) at time $t = 1, \dots, T$; \mathbf{f}_t is an $m \times 1$ vector of latent common factors and \mathbf{b}_i is the factor loadings for the i^{th} individual or component; u_{it} is the idiosyncratic error, uncorrelated with the common factors. In genomics application, t can also index repeated experiments. For simplicity we assume there is no time dependency.

The factor model can be written into a matrix form as follows:

$$\mathbf{Y} = \mathbf{B}\mathbf{F}' + \mathbf{U}, \quad (4.2)$$

where $\mathbf{Y}_{p \times T}$, $\mathbf{B}_{p \times m}$, $\mathbf{F}_{T \times m}$, $\mathbf{U}_{p \times T}$ are respectively the matrix form of the observed data, the factor loading matrix, the factor matrix, and the error matrix. For identifiability, we impose the condition that $\text{cov}(\mathbf{f}_t) = \mathbf{I}$. Thus, the covariance matrix is given by

$$\Sigma = \mathbf{B}\mathbf{B}' + \Sigma_u, \quad (4.3)$$

where Σ_u is the covariance matrix of the idiosyncratic error at any time t .

Under the assumption that $\Sigma_u = (\sigma_{u,ij})_{i,j \leq p}$ is sparse with its eigenvalues bounded away from zero and infinity, the population covariance exhibits a “low-rank plus sparse” structure. The sparsity is measured by the following quantity

$$m_p = \max_{i \leq p} \sum_{j \leq p} |\sigma_{u,ij}|^q,$$

for some $q \in [0, 1]$ (Bickel and Levina, 2008). In particular, with $q = 0$, m_p equals the maximum number of nonzero elements in each row of Σ_u .

In order to estimate the true covariance matrix with the above factor structure, Fan, Liao and Mincheva (2013) proposed a method called “POET” to recover the unknown factor matrix as well as the factor loadings. The idea is simply to first decompose the sample covariance

matrix into the spiked and non-spiked part and estimate them separately. Specifically, recall $\hat{\Sigma} = T^{-1}\mathbf{Y}\mathbf{Y}'$ and let $\{\hat{\lambda}_j\}$ and $\{\hat{\xi}_j\}$ be its corresponding eigenvalues and eigenvectors. They define

$$\hat{\Sigma}^\top = \sum_{j=1}^m \hat{\lambda}_j \hat{\xi}_j \hat{\xi}_j' + \hat{\Sigma}_u^\top, \tag{4.4}$$

where $\hat{\Sigma}_u^\top$ is the matrix after applying thresholding method (Bickel and Levina, 2008) to

$$\hat{\Sigma}_u = \hat{\Sigma} - \sum_{j=1}^m \hat{\lambda}_j \hat{\xi}_j \hat{\xi}_j'$$

They showed that the above estimation procedure is equivalent to the least square approach that minimizes

$$(\hat{\mathbf{B}}, \hat{\mathbf{F}}) = \arg \min_{\mathbf{B}, \mathbf{F}} \|\mathbf{Y} - \mathbf{B}\mathbf{F}'\|_F^2 \text{ s.t. } \frac{1}{T}\mathbf{F}'\mathbf{F} = \mathbf{I}_m, \mathbf{B}'\mathbf{B} \text{ is diagonal.} \tag{4.5}$$

The columns of $\hat{\mathbf{F}}/\sqrt{T}$ are the eigenvectors corresponding to the m largest eigenvalues of the $T \times T$ matrix $T^{-1}\mathbf{Y}'\mathbf{Y}$ and $\hat{\mathbf{B}} = T^{-1}\mathbf{Y}\hat{\mathbf{F}}$. After \mathbf{B} and \mathbf{F} are estimated, the sample covariance of $\hat{\mathbf{U}} = \mathbf{Y} - \hat{\mathbf{B}}\hat{\mathbf{F}}'$ can be formed: $\hat{\Sigma}_u = T^{-1}\hat{\mathbf{U}}\hat{\mathbf{U}}'$. Finally thresholding is applied to

$\hat{\Sigma}_u$ to generate $\hat{\Sigma}_u^\top = (\hat{\sigma}_{u,ij}^\top)_{p \times p}$, where

$$\hat{\sigma}_{u,ij}^\top \begin{cases} \hat{\sigma}_{u,ij}, & i=j; \\ s_{ij}(\hat{\sigma}_{u,ij})I(|\hat{\sigma}_{u,ij}| \geq \tau_{ij}), & i \neq j. \end{cases} \tag{4.6}$$

Here $s_j(\cdot)$ is the generalized shrinkage function (Antoniadis and Fan, 2001; Rothman, Levina and Zhu, 2009) and $\tau_{ij} = \tau(\hat{\sigma}_{u,i}\hat{\sigma}_{u,ij})^{1/2}$ is the entry-dependent threshold. The above adaptive threshold corresponds to applying thresholding with parameter τ to the correlation matrix of $\hat{\Sigma}_u$. The positive parameter τ will be determined later.

Fan, Liao and Mincheva (2013) showed that under Assumptions ?? - ?? listed in Appendix ?? in the supplementary material (Wang and Fan, 2015),

$$\|\hat{\Sigma}^\top - \Sigma\|_{\Sigma, F} = O_p \left(\frac{\sqrt{p} \log p}{T} + m_p \left(\frac{\log p}{T} + \frac{1}{p} \right)^{(1-q)/2} \right), \tag{4.7}$$

where $\|\mathbf{A}\|_{\Sigma, F} = p^{-1/2}\|\Sigma^{-1/2}\mathbf{A}\Sigma^{-1/2}\|_F$ and $\|\cdot\|_F$ is the Frobenius norm. Note that

$$\|\hat{\Sigma}^\top - \Sigma\|_{\Sigma, F} = p^{-1/2}\|\Sigma^{-1/2}\hat{\Sigma}^\top \Sigma^{-1/2} - \mathbf{I}_p\|_F,$$

which measures the relative error in Frobenius norm. A more natural metric is relative error under the operator norm $\|\mathbf{A}\|_{\Sigma} = \|\Sigma^{-1/2}\mathbf{A}\Sigma^{-1/2}\|$, which can not be obtained by using the technical device of Fan, Liao and Mincheva (2013). Note $\|\mathbf{A}\|_{\Sigma, F} \leq \|\mathbf{A}\|_{\Sigma}$. Via our new results in the last section, we will establish a result under those two relative norms, under weaker conditions than their pervasiveness assumption. Note that the relative error convergence is particularly meaningful for spiked covariance matrix, as eigenvalues are in different scales.

4.2. Shrinkage POET under relative spectral norm

The discussion above reveals several drawbacks of POET. First, the spike size has to be of order p which rules out relatively weak factors. Second, it is well known that the empirical eigenvalues are inconsistent if the spiked eigenvalues do not significantly dominate the non-spiked part. Therefore, a proper correction or shrinkage is needed. See a recent paper by Donoho, Gavish and Johnstone (2014) for optimal shrinkage of eigenvalues.

Regarding to the first drawback, we relax the assumption $\|p^{-1}\mathbf{B}'\mathbf{B} - \mathbf{\Omega}_0\| = o(1)$ in Assumption ?? to the following weaker assumption.

Assumption 4.1— $\|\Lambda_A^{-1/2}\mathbf{B}'\mathbf{B}\Lambda_A^{-1/2} - \mathbf{\Omega}_0\| = o(1)$ for some $\mathbf{\Omega}_0$ with eigenvalues bounded from above and below, where $\Lambda_A = \text{diag}(\lambda_1, \dots, \lambda_m)$. In addition, we assume $\lambda_m \rightarrow \infty$, λ_1/λ_m is bounded from above and below.

This assumption does not require the first m eigenvalues of Σ to take on any specific rate. They can still be much smaller than p , although for simplicity we require them to diverge and share the same diverging rate. Since $\|\Sigma_p\|$ is assumed to be bounded, the assumption $\lambda_m \rightarrow \infty$ is also imposed to avoid the issue of identifiability. When λ_m does not diverge, more sophisticated condition is needed for identifiability (Chandrasekaran et al., 2011).

In order to handle the second drawback, we propose the Shrinkage POET (S-POET) method. Inspired by (3.1), the shrinkage POET modifies the first part in POET estimator (4.4) as follows:

$$\hat{\Sigma}^S = \sum_{j=1}^m \hat{\lambda}_j^S \hat{\xi}_j \hat{\xi}_j' + \hat{\Sigma}_u^{\top}, \tag{4.8}$$

where $\hat{\lambda}_j^S = \max\{\hat{\lambda}_j - \bar{c}p/n, 0\}$, a simple soft thresholding correction. Obviously if $\hat{\lambda}_j$ is sufficiently large, $\hat{\lambda}_j^S/\lambda_j = \hat{\lambda}_j/\lambda_j - \bar{c}c_j = 1 + o_p(1)$. Since \bar{c} is unknown, a natural estimator \hat{c} is such that the total of the eigenvalues remains unchanged:

$$\text{tr}(\hat{\Sigma}) = \sum_{j=1}^m (\hat{\lambda}_j - \hat{c}p/n) + (p - m)\hat{c}$$

or $\hat{c} = (\text{tr}(\hat{\Sigma}) - \sum_{j=1}^m \hat{\lambda}_j) / (p - m - pm/n)$. It has been shown by Lemma 7 of Yata and Aoshima (2012) that

$$(\hat{c} - \bar{c}) \frac{p}{n\lambda_j} = O_P \left(\frac{\text{tr}(\hat{\Sigma}) - \sum_{j=1}^m \hat{\lambda}_j}{(n - m)\lambda_m} - \frac{\bar{c}p}{n\lambda_m} \right) = O_P(n^{-1}).$$

Thus, replacing \bar{c} by \hat{c} , we have $\hat{\lambda}_j^S / \lambda_j - 1 = O_P(\lambda_j^{-1} \sqrt{p/n + n^{-1/2}})$, i.e. the estimation error in \hat{c} is negligible. From Lemma 3.1, we can easily obtain the asymptotic normality,

$$\sqrt{n}(\hat{\lambda}_j^S / \lambda_j - 1) \xrightarrow{d} N(0, \kappa_j - 1) \text{ if } \sqrt{p} = o(\lambda_j).$$

To get the convergence of relative errors under the operator norm, we also need the following additional assumptions:

Assumption 4.2

- i. $\{\mathbf{u}_i, \mathbf{f}_t\}_{t=1}^T$ are independently and identically distributed with $\mathbb{E}[u_{it}] = \mathbb{E}[u_{it}f_{jt}] = 0$ for all $i = 1, \dots, p, j = 1, \dots, m$ and $t = 1, \dots, T$.
- ii. There exist positive constants c_1 and c_2 such that $\lambda_{\min}(\Sigma_u) > c_1$, $\|\Sigma_u\|_{\infty} < c_2$, and $\min_{i,j} \text{Var}(u_{it}u_{jt}) > c_1$.
- iii. There exist positive constants r_1, r_2, b_1 and b_2 such that for $s > 0, i = 1, \dots, p, j = 1, \dots, m$,

$$\mathbb{P}(|u_{it}| > s) \leq \exp(-(s/b_1)^{r_1}) \text{ and } \mathbb{P}(|f_{jt}| > s) \leq \exp(-(s/b_2)^{r_2}).$$

- iv. There exists $M > 0$ such that for all $i = 1, \dots, p, j = 1, \dots, m, |b_{ij}| \leq M \sqrt{\lambda_j/p}$.
- v. $\sqrt{p}(\log T)^{1/r_2} = o(\lambda_m)$.

The first three conditions are common in factor model literature. If we write $\mathbf{B} = (\tilde{\mathbf{b}}_1, \dots, \tilde{\mathbf{b}}_m)$, by Weyl's inequality we have $\max_{1 \leq j \leq m} \|\tilde{\mathbf{b}}_j\|^2 / \lambda_j = 1 + \|\Sigma_u\| / \lambda_j = 1 + o(1)$. Thus it is

reasonable to assume the magnitude $|b_{ij}|$ of factor loadings is of order $\sqrt{\lambda_j/p}$ in the fourth condition. The last condition is imposed to ease technical presentation.

Now we are ready to investigate $\|\hat{\Sigma}^S - \Sigma\|_{\Sigma}$. Suppose the SVD decomposition of Σ ,

$$\Sigma = \left(\Gamma_{p \times m} \Omega_{p \times (p-m)} \right) \begin{pmatrix} \Lambda_{m \times m} & \\ & \Theta_{(p-m) \times (p-m)} \end{pmatrix} \begin{pmatrix} \Gamma' \\ \Omega' \end{pmatrix}.$$

Then obviously

$$\|\hat{\Sigma}^S - \Sigma\|_{\Sigma} \leq \left\| \Sigma^{-\frac{1}{2}} (\hat{\Gamma} \hat{\Lambda}^S \hat{\Gamma}' - \mathbf{B} \mathbf{B}') \Sigma^{-\frac{1}{2}} \right\| + \left\| \Sigma^{-\frac{1}{2}} (\hat{\Sigma}_u^{\top} - \Sigma_u) \Sigma^{-\frac{1}{2}} \right\| =: \Delta_L + \Delta_S, \tag{4.9}$$

and

$$\Delta_S \leq \|\Sigma^{-1}\| \|\hat{\Sigma}_u^\top - \Sigma_u\| \leq C \|\hat{\Sigma}_u^\top - \Sigma_u\|. \quad (4.10)$$

It can be shown

$$\Delta_L = \left\| \begin{pmatrix} \Lambda^{-\frac{1}{2}} \Gamma' \\ \Theta^{-\frac{1}{2}} \Omega' \end{pmatrix} (\hat{\Gamma} \hat{\Lambda}^S \hat{\Gamma}' - \mathbf{B} \mathbf{B}') (\Gamma \Lambda^{-\frac{1}{2}} \Omega \Theta^{-\frac{1}{2}}) \right\| \leq \Delta_{L1} + \Delta_{L2}, \quad (4.11)$$

where $\Delta_{L1} = \|\Lambda^{-\frac{1}{2}} \Gamma' (\hat{\Gamma} \hat{\Lambda}^S \hat{\Gamma}' - \mathbf{B} \mathbf{B}') \Gamma \Lambda^{-\frac{1}{2}}\|$ and $\Delta_{L2} = \|\Theta^{-\frac{1}{2}} \Omega' (\hat{\Gamma} \hat{\Lambda}^S \hat{\Gamma}' - \mathbf{B} \mathbf{B}') \Omega \Theta^{-\frac{1}{2}}\|$. Thus in order to find the convergence rate of relative spectral norm, we need to consider the terms $L1$, $L2$ and S separately. Notice that $L1$ measures the relative error of the estimated spiked eigenvalues, $L2$ reflects the goodness of the estimated eigenvectors, and S controls the error of estimating the sparse idiosyncratic covariance matrix. To bound the relative Frobenius norm $\|\hat{\Sigma}_S - \Sigma\|_{\Sigma, F}$, we define similar quantities $\tilde{L1}$, $\tilde{L2}$, \tilde{S} which replace the spectral norm by Frobenius norm multiplied by $p^{-1/2}$. Note that (4.9) – (4.11) also hold for relative Frobenius norm with $\tilde{L1}$, $\tilde{L2}$, \tilde{S} . The following theorem reveals the rate of each term. Its proof will be provided in Appendix ?? of the supplementary material (Wang and Fan, 2015).

Theorem 4.1—*Under Assumptions 2.1, 2.2, 2.3, 4.1 and 4.2, if $p \log p > \max\{T(\log T)^{4/2}, T(\log(pT))^{2/\tau_1}\}$, we have*

$$\tilde{\Delta}_{L1} \leq \Delta_{L1} = O_p\left(T^{-1/2}\right), \Delta_{L2} = O_p\left(\frac{p}{T}\right), \tilde{\Delta}_{L2} = O_p\left(\frac{\sqrt{p}}{T}\right),$$

and by the applying adaptive thresholding estimator (4.6) with

$$\tau_{ij} = C \omega_T (\hat{\sigma}_{u,ii} \hat{\sigma}_{u,jj})^{1/2}, \quad \text{and } \omega_T = \sqrt{\log p/T} + \sqrt{1/p},$$

we have

$$\tilde{\Delta}_S \leq \Delta_S = O_p\left(m_p \omega_T^{1-q}\right).$$

Combining the three terms, $\|\hat{\Sigma}^S - \Sigma\|_{\Sigma} = O_p(p/T + m_p \omega_T^{1-q})$ and

$$\|\hat{\Sigma}^S - \Sigma\|_{\Sigma, F} = O_p(\sqrt{p}/T + m_p \omega_T^{1-q}).$$

The relative error convergence characterizes the accuracy of estimation for spiked covariance matrix. In contrast with the result on relative Frobenius norm, this is the first time that the relative rate under spectral norm is derived. As long as λ_m is slightly above \sqrt{p} , we

reach the same rates of convergence. Therefore, we can conclude S-POET is effective even under a much weaker signal level. Comparing the rate with (4.7), under relative Frobenius norm, we achieve a better rate without the artificial log term, thanks to the new asymptotic results.

4.3. Portfolio risk management

The risk of a given portfolio with allocation weight \mathbf{w} is conventionally measured by its variance $\mathbf{w}'\Sigma\mathbf{w}$, where Σ is the volatility (covariance) matrix of the returns of underlying assets. To estimate large portfolio's risks, it needs to estimate a large covariance matrix Σ and factor models are frequently used to reduce the dimensionality. This was the idea of Fan, Liao and Shi (2015) in which they used POET estimator to estimate Σ . However, the basic method for bounding the risk error $|\mathbf{w}'\hat{\Sigma}\mathbf{w} - \mathbf{w}'\Sigma\mathbf{w}|$ in their paper is

$$|\mathbf{w}'\hat{\Sigma}\mathbf{w} - \mathbf{w}'\Sigma\mathbf{w}| \leq \|\mathbf{w}\|_1^2 \|\hat{\Sigma} - \Sigma\|_{\max}.$$

They assumed that the gross exposure of the portfolio is bounded, i.e. $\|\mathbf{w}\|_1 = O(1)$. Technically, when p is large, $\mathbf{w}'\Sigma\mathbf{w}$ can be small. What an investor cares mostly is the relative risk error $\text{RE}(\mathbf{w}) = |\mathbf{w}'\hat{\Sigma}\mathbf{w}/\mathbf{w}'\Sigma\mathbf{w} - 1|$. Often \mathbf{w} is a data-driven investment strategy, which depends on the past data. Regardless of what \mathbf{w} is,

$$\max_{\mathbf{w}} \text{RE}(\mathbf{w}) = \|\hat{\Sigma} - \Sigma\|_{\Sigma},$$

which does not converge by Theorem 4.1 for $p > T$. Thus the question of interest is what kind of portfolio \mathbf{w} will make the relative error converge. Decompose \mathbf{w} as a linear combination of the eigenvectors of Σ , namely $\mathbf{w} = (\Gamma, \Omega)\boldsymbol{\eta}$ and $\boldsymbol{\eta} = (\boldsymbol{\eta}'_A, \boldsymbol{\eta}'_B)'$. We have the following useful result for risk management.

Theorem 4.2—Under Assumptions 2.1, 2.2, 4.1, 4.2 and the factor model (4.1) with Gaussian noises and factors, if there exists $C_1 > 0$ such that $\|\boldsymbol{\eta}_B\|_1 \leq C_1$, and assume $\lambda_j \propto p^\alpha$ for $j = 1, \dots, m$ and $T \leq Cp^\beta$ for $\alpha > 1/2$, $0 < \beta < 1$, $\alpha + \beta > 1$, then the relative risk error is of order

$$\text{RE}(\mathbf{w}) = \left| \frac{\mathbf{w}'\hat{\Sigma}\mathbf{w}}{\mathbf{w}'\Sigma\mathbf{w}} - 1 \right| = O_p \left(T^{-\min\{\frac{2(\alpha+\beta-1)}{\beta}, \frac{1}{2}\}} + m_p w_T^{1-q} \right),$$

for $\alpha < 1$. If $\alpha = 1$ or $\|\boldsymbol{\eta}_A\|_1 \leq C_2$, $\text{RE}(\mathbf{w}) = O_p(m_p w_T^{1-q})$.

The condition $\|\boldsymbol{\eta}_B\|_1 \leq C_1$ is generally weaker than $\|\mathbf{w}\|_1 = O(1)$. It does not limit the total exposure of investor's position, but only put constraint on investment of the non-spiked section. Note that under the conditions of Theorem 4.2, $p/(T\lambda_j) \rightarrow 0$, and S-POET and POET are approximately the same. So the stated result is valid for POET too.

4.4. Estimation of false discovery proportion

Another important application of the factor model is the estimation of false discovery proportion. For simplicity, we assume Gaussian data $\mathbf{X}_j \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ with an unknown correlation matrix $\boldsymbol{\Sigma}$ and wish to test which coordinates of $\boldsymbol{\mu}$ are nonvanishing. Consider the test statistic $\mathbf{Z} = \sqrt{n}\bar{\mathbf{X}}$ where $\bar{\mathbf{X}}$ is the sample mean of all data. Then $\mathbf{Z} \sim \mathcal{N}(\boldsymbol{\mu}^*, \boldsymbol{\Sigma})$ with $\boldsymbol{\mu}^* = \sqrt{n}\boldsymbol{\mu}$. The problem is to test

$$H_{0j}:\mu_j^*=0 \quad \text{v.s.} \quad H_{1j}:\mu_j^* \neq 0.$$

Define the number of discoveries $R(t) = \#\{j : P_j \leq t\}$ and the number of false discoveries $V(t) = \#\{\text{true null} : P_j \leq t\}$, where P_j is the p-value associated with the j^{th} test. Note that $R(t)$ is observable while $V(t)$ needs to be estimated. The false discovery proportion (FDP) is defined as $\text{FDP}(t) = V(t)/R(t)$.

Fan and Han (2013) proposed to employ the factor structure

$$\boldsymbol{\Sigma} = \mathbf{B}\mathbf{B}' + \mathbf{A}, \quad (4.12)$$

where $\mathbf{B} = (\sqrt{\lambda_1}\boldsymbol{\xi}_1, \dots, \sqrt{\lambda_m}\boldsymbol{\xi}_m)$, λ_j and $\boldsymbol{\xi}_j$ are respectively the j^{th} eigenvalue and eigenvector of $\boldsymbol{\Sigma}$ as before. Then \mathbf{Z} can be stochastically decomposed as

$$\mathbf{Z} = \boldsymbol{\mu}^* + \mathbf{B}\mathbf{W} + \mathbf{K},$$

where $\mathbf{W} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_m)$ are m common factors and $\mathbf{K} \sim \mathcal{N}(\mathbf{0}, \mathbf{A})$, independent of \mathbf{W} , are the idiosyncratic errors. For simplicity, assume the maximal number of nonzero elements of each row of \mathbf{A} is bounded. In Fan and Han (2013), they argued that the asymptotic upper bound

$$\text{FDP}_A(t) = \sum_{i=1}^p [\Phi(a_i(z_{t/2} + \eta_i)) + \Phi(a_i(z_{t/2} - \eta_i))] / R(t) \quad (4.13)$$

of $\text{FDP}(t)$ should be a realistic target to estimate for dependence tests, where $z_{t/2}$ is the $t/2$ -quantile of the standard normal distribution $a_j = (1 - \|\mathbf{b}_j\|^2)^{-1/2}$, $\eta_i = \mathbf{b}'_i \mathbf{W}$ and \mathbf{b}'_i is the i^{th} row of \mathbf{B} .

Realized factors \mathbf{W} and the loading matrix \mathbf{B} are typically unknown. If a generic estimator $\hat{\boldsymbol{\Sigma}}$ is provided, then we are able to estimate \mathbf{B} and thus \mathbf{b}_j from its empirical eigenvalues and eigenvectors $\hat{\lambda}_j$'s and $\hat{\boldsymbol{\xi}}_j$'s. \mathbf{W} can be estimated by the least-squares estimate $\hat{\mathbf{W}} = (\hat{\mathbf{B}}'\hat{\mathbf{B}})^{-1}\hat{\mathbf{B}}'\mathbf{Z}$. Fan and Han (2013) proposed the following estimator for $\text{FDP}_A(t)$:

$$\widehat{\text{FDP}}_{U'}(t) = \sum_{i=1}^p [\Phi(\hat{a}_i(z_{t/2} + \hat{\eta}_i)) + \Phi(\hat{a}_i(z_{t/2} - \hat{\eta}_i))] / R(t), \tag{4.14}$$

where $\hat{a}_i = (1 - \|\hat{\mathbf{b}}_i\|^2)^{-1/2}$ and $\hat{\eta}_i = \hat{\mathbf{b}}_i' \hat{\mathbf{W}}$. The following assumptions are in their paper.

Assumption 4.3—*There exists a constant $h > 0$ such that (i) $R(t)/p > hp^{-\theta}$ for $h > 0$ and $\theta > 0$ as $p \rightarrow \infty$ and (ii) $\hat{a}_i \geq h$, $\hat{\eta}_i \leq h$ for all $i = 1, \dots, p$.*

They showed that if $\hat{\Sigma}$ is based on the POET estimator with a spike size $\lambda_m \asymp p$, under Assumptions ?? - ??, on the event that Assumption 4.3 holds,

$$|\widehat{\text{FDP}}_{U, \text{POET}}(t) - \text{FDP}_A(t)| = O_P \left(p^\theta \left(\sqrt{\frac{\log p}{T}} + \frac{\|\boldsymbol{\mu}^*\|}{\sqrt{p}} \right) \right). \tag{4.15}$$

Again we can relax the assumption on the spike magnitude from order p to much weaker Assumption 4.1. Since Σ is a correlation matrix, $\lambda_1 = \text{tr}(\Sigma) = p$. This, together with Assumption 4.1, leads us to consider leading eigenvalues of order p^α for $1/2 < \alpha < 1$.

Now we apply the proposed S-POET method to obtain $\hat{\Sigma}^S$ and use it for FDP estimation. The following theorem shows the estimation error.

Theorem 4.3—*If Assumptions 2.1, 2.2, 4.1, and 4.2 are applied to Gaussian independent data $\mathbf{X}_j \sim N(\boldsymbol{\mu}, \Sigma)$, and $\lambda_j \propto p^\alpha$ for $j = 1, \dots, m$, $T \asymp Cp^\beta$ for $1/2 < \alpha < 1$, $0 < \beta < 1$, $\alpha + \beta > 1$, on the event that Assumption 4.3 holds, we have*

$$|\widehat{\text{FDP}}_{U, \text{SPOET}}(t) - \text{FDP}_A(t)| = O_P \left(p^\theta \left(\|\boldsymbol{\mu}^*\| p^{-\frac{1}{2}} + T^{-\min\{\frac{\alpha+\beta-1}{\beta}, \frac{1}{2}\}} \right) \right).$$

Comparing the result with (4.15), this convergence rate attained by S-POET is more general than the rate achieved before. The only difference is the second term, which is $O(T^{-1/2})$ if

$\alpha + \frac{1}{2}\beta \geq 1$ and $T^{-(\alpha + \beta - 1)/\beta}$ otherwise. So we relax the condition from $\alpha = 1$ in Fan and Han (2013) to $\alpha \in (1/2, 1]$. This means a weaker signal than order p is actually allowed to obtain a consistent estimate of false discovery proportion.

5. Simulations

We conducted some simulations to demonstrate the finite sample behavior of empirical eigen-structure, the performance of S-POET, and validity of applying it to estimate false discovery proportion.

5.1. Eigen-structure

In this simulation, we set $n = 50$, $p = 500$ and $\Sigma = \text{diag}(50, 20, 10, 1, \dots, 1)$, which has three spike eigenvalues ($m = 3$) $\lambda_1 = 50$, $\lambda_2 = 20$, $\lambda_3 = 10$ and correspondingly $c_1 = 0.2$, $c_2 = 0.5$, $c_3 = 1$. Data was generated from multivariate Gaussian. The number of simulations is 1000.

The histograms of the standardized empirical eigenvalues $\sqrt{n/2}(\hat{\lambda}_j/\lambda_j - 1 - c_j)$, and their associated asymptotic distributions (standard normal) are plotted in Figure 1. The approximations are very good even for this low sample size $n = 50$.

Figure 2 shows the histograms of $\sqrt{n}(\hat{\xi}_{jA}/\|\hat{\xi}_{jA}\| - e_{jA})$ for the first three elements (the spiked part) of the first three eigenvectors. On the one hand, according to the asymptotic results, the values in the diagonal position should stochastically converge to 0 as observed. On the other hand, plots in the off-diagonal positions should converge in distribution to $N(0, 1)$ for $k \neq j$ after standardization, which is indeed the case. We also report the correlations between the first three elements for the three eigenvectors based on those 1000 repetitions in Table 1. The correlations are all quite close to 0, which is consistent with the theory.

For the normalized non-spiked part $\hat{\xi}_{jB}/\|\hat{\xi}_{jB}\|$, it should be distributed uniformly over the unit sphere. This can be tested by the results of Cai, Fan and Jiang (2013). For any n data points $\mathbf{X}_1, \dots, \mathbf{X}_n$ on a p -dimensional sphere, define the normalized empirical distribution of angles of each pair of vectors as

$$\mu_{n,p} = \frac{1}{\binom{n}{2}} \sum_{1 \leq i < j \leq n} \delta_{\sqrt{p-2}(\pi/2 - \Theta_{ij})},$$

where $\Theta_{ij} \in [0, \pi]$ is the angle between vectors \mathbf{X}_i and \mathbf{X}_j . When the data are generated uniformly from a sphere, $\mu_{n,p}$ converges to the standard normal distribution with probability 1.

Figure 3 shows the empirical distributions of all pairwise angles of the realized $\hat{\xi}_{jB}/\|\hat{\xi}_{jB}\|$ ($j = 1, 2, 3$) in 1000 simulations. Since number of such pairwise angles is $\binom{1000}{2}$, the empirical distributions and the asymptotic distributions $N(0, 1)$ are almost identical. The normality holds even for a small subset of the angles.

Lastly, we did simulations to verify the rate difference of $\langle \hat{\xi}_j, e_j \rangle$ for $m = 1$ and $m > 1$, revealed in Theorem 3.2 (iii). We choose $n = \lceil 10 \times 1.2^l \rceil$ for $l = 0, \dots, 9$, $p = \lceil n^3/100 \rceil$, where $\lceil \cdot \rceil$ represents rounding. We set $\lambda_j = 1$ for $j \geq 3$ and consider two situations: (1) $\lambda_1 = p$, $\lambda_2 = 1$, (2) $\lambda_1 = 2\lambda_2 = p$. Under both cases, simulations were carried out 500 times and the corresponding angle of the empirical eigenvector and its truth was calculated for each simulation. The logarithm of the median absolute error of $\langle \hat{\xi}_1, e_1 \rangle - 1/\sqrt{1+c_1}$ was plotted against $\log(n)$. Under the two situations, the rates of convergence are $O_P(n^{-3/2})$ and $O_P(n^{-1})$ respectively. Thus the slope of the curves should be $-3/2$ for a single spike and -1 for two spikes, which is indeed the case as shown in Figure 4.

In short, all the simulation results match well with the theoretical results for the high dimensional regime.

5.2. Performance of S-POET

We demonstrate the effectiveness of S-POET in comparison with POET. A similar setting to the last section is used, i.e. $m = 3$ and $c_1 = 0.2$, $c_2 = 0.5$, $c_3 = 1$. The sample size T ranges from 50 to 150 and $p = \lceil T^{3/2} \rceil$. Note that when $T = 150$, $p \approx 1800$. The spiked eigenvalues are determined from $p(\mathcal{T}\lambda_j) = c_j$ so that λ_j is of order \sqrt{T} , which is much smaller than p . For each pair of T and p , the following steps are used to generate observed data from the factor model for 200 times.

1. Each row of \mathbf{B} is simulated from the standard multivariate normal distribution and the j^{th} column is normalized to have norm λ_j for $j = 1, 2, 3$.
2. Each row of \mathbf{F} is simulated from standard multivariate normal distribution.
3. Set $\Sigma_u = \text{diag}(\sigma_1^2, \dots, \sigma_p^2)$ where σ_i 's are generated from Gamma(α , β) with $\alpha = \beta = 100$ (mean 1, standard deviation 0.1). The idiosyncratic error \mathbf{U} is simulated from $\mathcal{N}(\mathbf{0}, \Sigma_u)$.
4. Compute the observed data $\mathbf{Y} = \mathbf{B}\mathbf{F}' + \mathbf{U}$.

Both S-POET and POET are applied to estimate the covariance matrix $\Sigma = \mathbf{B}\mathbf{B}' + \Sigma_u$. Their mean estimation errors over 200 simulations, measured in relative spectral norm $\|\hat{\Sigma} - \Sigma\|_{\Sigma}$, relative Frobenius norm $\|\hat{\Sigma} - \Sigma\|_{\Sigma, F}$, spectral norm $\|\hat{\Sigma} - \Sigma\|$ and max norm $\|\hat{\Sigma} - \Sigma\|_{\max}$, are reported in Figure 5. The errors for sample covariance matrix are also depicted for comparison. First notice that no matter in what norm, S-POET uniformly outperforms POET and the sample covariance. It affirms the claim that shrinkage of spiked eigenvalues is necessary to maintain good performance when the spikes are not sufficiently large. Since the low rank part is not shrunk for POET, its error under the spectral norm is comparable and even slightly larger than that of the sample covariance matrix. The errors under max norm and relative Frobenius norm as expected decrease as T and p increase. However the error under the relative spectral norm does not converge: our theory shows it should increase in the order $p/T = \sqrt{T}$.

5.3. FDP estimation

In this section, we report simulation results on FDP estimation by using both POET and S-POET. The data are simulated in a similar way as in Section 5.2 with $p = 1000$ and $n = 100$. The first $m = 3$ eigenvalues have spike size proportional to p/\sqrt{n} which corresponds to $\alpha = \beta = 2/3$ in Theorem 4.3. The true FDP is calculated by using $\text{FDP}(t) = V(t)/R(t)$ with $t = 0.01$. The approximate FDP, $\text{FDP}_A(t)$, is calculated as in (4.13) with known \mathbf{B} but estimated \mathbf{W} given by $\hat{\mathbf{W}} = (\mathbf{B}\mathbf{B}')^{-1}\mathbf{B}'\mathbf{Z}$. This $\text{FDP}_A(t)$ based on a known covariance matrix serves as a benchmark for our estimated covariance matrix to compare with. We employed POET and S-POET to get $\widehat{\text{FDP}}_{U, \text{POET}}(t)$ and $\widehat{\text{FDP}}_{U, \text{SPOET}}(t)$.

In Figure 6, three scatter plots are drawn to compare $FDP_A(t)$, $\widehat{FDP}_{U,POET}(t)$ and $\widehat{FDP}_{U,SPOET}(t)$ with the true $FDP(t)$. The points are basically aligned along the 45 degree line, meaning that all of them are quite close to the true FDP. With the semi-strong signal $\lambda \propto p/\sqrt{n}$, although much weaker than order p , POET accomplishes the task as well as S-POET. Both estimators perform as well as if we know the covariance matrix Σ , the benchmark.

6. Conclusions

In this paper, we studied two closely related problems: the asymptotic behavior of empirical eigenvalues and eigenvectors under a general regime of bounded $p/(n\lambda_j)$ and the large covariance estimation for factor models with relaxed signal level of $\sqrt{p}=o(\lambda_j)$.

The first study provides new technical tools for the derivation of error bounds for large covariance estimation under relative Frobenius norm (with better rate) and relative spectral norm (for the first time). The results motivate the new proposed covariance estimator S-POET for the second problem by correcting biases of the estimated leading eigenvalues. S-POET is demonstrated to have better sampling properties than POET, and this is convincingly verified in the simulation study. In addition, we are able to apply S-POET to two important applications, risk management and false discovery control, and relax the required signal to \sqrt{p} . Those conclusions shed new lights for applications of factor models.

On the other hand, the second problem is a key motivation for us to study the empirical engen-structure in a more general high dimensional regime. We aim to understand why PCA works for pervasive factor models but fails classical random matrix problems, without sparsity assumptions. What are the fundamental limit for the PCA in high dimension? We clearly showed that for both empirical eigenvalues and vectors, consistency is granted once $p/(n\lambda_j) \rightarrow 0$. Further, our theories give a fine-grained characterization of the asymptotic behavior under the generalized and unified regime, which includes the situation of bounded eigenvalues, HDLSS, and pervasive factor models, especially for empirical eigenvectors. The asymptotic rate of convergence is obtained as long as $p/(n\lambda_j)$ is bounded, while the asymptotic distribution is fully described when $\sqrt{p}=o(\lambda_j)$. Some interesting phenomena, such as interaction between multiple spikes, are also revealed in our results. Our proofs are novel in that we clearly identify terms that keep the low-dimensional asymptotic normality and terms that generate the random biases. In sum, our results serve as a necessary complement of the random matrix literature when the signal diverges with dimensionality.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The research was partially supported by NSF grants DMS-1206464 and DMS-1406266 and NIH grants R01-GM072611-11 and NIH R01GM100474-04. We would like to thank the Editor, associate editor, and anonymous referees for constructive comments that lead to substantial improvement of the presentation and results of the paper.

APPENDIX A

PROOFS FOR SECTION 3

A.1. Proof of Theorem 3.1

We first provide three useful lemmas for the proof. Lemma A.1 provides non-asymptotic upper and lower bound for the eigenvalues of weighted Wishart matrix for sub-Gaussian distributions.

Lemma A.1—Let $\mathbf{A}_1, \dots, \mathbf{A}_n$'s be n independent p dimensional sub-Gaussian random vectors with zero mean and identity variance, and the sub-Gaussian norms bounded by a constant C_0 . Then for every $t > 0$, with probability at least $1 - 2 \exp(-ct^2)$, one has

$$\bar{w} - \max\{\delta, \delta^2\} \leq \lambda_p \left(\frac{1}{n} \sum_{i=1}^n w_i \mathbf{A}_i \mathbf{A}_i' \right) \leq \lambda_1 \left(\frac{1}{n} \sum_{i=1}^n w_i \mathbf{A}_i \mathbf{A}_i' \right) \leq \bar{w} + \max\{\delta, \delta^2\}.$$

where $\delta = C \sqrt{p/n + t} / \sqrt{n}$ for constants $C, c > 0$, depending on C_0 . Here $|w_j|$'s is bounded for all i and $\bar{w} = n^{-1} \sum_{i=1}^n w_i$.

The above lemma is the extension of the classical Davidson-Szarek bound [Theorem II.7 of Davidson and Szarek (2001)] to the weighted sample co-variance with sub-Gaussian distribution. It was shown by Vershynin (2010) that the conclusion holds with $w_j = 1$ for all i . With similar techniques to those developed in Vershynin (2010), we can obtain the above lemma for general bounded weights. The details are omitted.

Now in order to prove the theorem, let us define two quantities and treat them separately in the following two lemmas. Let

$$\mathbf{A} = n^{-1} \sum_{j=1}^m \lambda_j \mathbf{Z}_j \mathbf{Z}_j', \quad \text{and} \quad \mathbf{B} = n^{-1} \sum_{j=m+1}^p \lambda_j \mathbf{Z}_j \mathbf{Z}_j',$$

where \mathbf{Z}_j is columns of $\mathbf{X} \boldsymbol{\Lambda}^{-\frac{1}{2}}$. Then,

$$\tilde{\boldsymbol{\Sigma}} = \frac{1}{n} \sum_{j=1}^p \lambda_j \mathbf{Z}_j \mathbf{Z}_j' = \mathbf{A} + \mathbf{B}. \tag{A.1}$$

Lemma A.2—Under Assumptions 2.1 – 2.3, as $n \rightarrow \infty$,

$$\sqrt{n} (\lambda_j(\mathbf{A}) / \lambda_j - 1) \xrightarrow{d} N(0, \kappa_j - 1), \quad \text{for } j=1, \dots, m.$$

In addition, they are asymptotically independent.

Lemma A.3—Under Assumptions 2.1 – 2.3, for $j = 1, \dots, m$, we have

$$\lambda_k(\mathbf{B})/\lambda_j = \bar{c}c_j + O_p\left(\lambda_j^{-1}\sqrt{p/n}\right) + o_p(n^{-\frac{1}{2}}), \text{ for } k=1, 2, \dots, n.$$

The proofs of the above two lemmas will be given in Appendix?? in the supplementary material (Wang and Fan, 2015).

Proof of Theorem 3.1—By Wely’s Theorem, $\lambda_j(\mathbf{A}) + \lambda_n(\mathbf{B}) \hat{\lambda}_j = \lambda_j(\mathbf{A}) + \lambda_1(\mathbf{B})$: Therefore from Lemma A.3,

$$\frac{\hat{\lambda}_j}{\lambda_j} = \frac{\lambda_j(\mathbf{A})}{\lambda_j} + \bar{c}c_j + O_p\left(\lambda_j^{-1}\sqrt{\frac{p}{n}}\right) + o_p(c_j n^{-1/2}),$$

By Lemma A.2 and Slutsky’s theorem, we conclude that

$\sqrt{n}\left(\hat{\lambda}_j/\lambda_j - \left(1 + \bar{c}c_j + O_p(\lambda_j^{-1}\sqrt{p/n})\right)\right)$ converges in distribution to $N(0, \kappa_j - 1)$ and the limiting distributions of the first m eigenvalues are independent.

A.2. Proofs of Theorem 3.2

The proof of Theorem 3.2 is mathematically involved. The basic idea for proving part (i) is outlined in Section 2. We relegate less important technical lemmas A.4 – A.6 to Appendix ?? in the supplementary material (Wang and Fan, 2015) in order not to distract the readers. The proof of part (ii) utilizes the invariance of standard Gaussian distribution under orthogonal transformations.

Proof of Theorem 3.2—(i) Let us start by proving the asymptotic normality of $\hat{\xi}_{jA}$ for the case $m > 1$. Write

$$\mathbf{X} = (\mathbf{Z}_A \Lambda_A^{\frac{1}{2}}, \mathbf{Z}_B \Lambda_B^{\frac{1}{2}}) = (\sqrt{\lambda_1} \mathbf{Z}_1, \dots, \sqrt{\lambda_m} \mathbf{Z}_m, \sqrt{\lambda_{m+1}} \mathbf{Z}_{m+1}, \dots, \sqrt{\lambda_p} \mathbf{Z}_p),$$

where each \mathbf{Z}_j follows a sub-Gaussian distribution with mean $\mathbf{0}$ and identity variance \mathbf{I}_D . Then by the eigenvalue relationship of equation (2.1), we have

$$\hat{\xi}_{jA} = \frac{\Lambda_A^{\frac{1}{2}} \mathbf{Z}'_A \mathbf{u}_j}{\sqrt{n \hat{\lambda}_j}} \text{ and } \mathbf{u}_j = \frac{\mathbf{X} \hat{\xi}_j}{\sqrt{n \hat{\lambda}_j}} = \frac{\mathbf{Z}_A \Lambda_A^{\frac{1}{2}} \hat{\xi}_{jA}}{\sqrt{n \hat{\lambda}_j}} + \frac{\mathbf{Z}_B \Lambda_B^{\frac{1}{2}} \hat{\xi}_{jB}}{\sqrt{n \hat{\lambda}_j}}. \tag{A.2}$$

Recall \mathbf{u}_j is the eigenvector of the matrix $\tilde{\Sigma}$, that is, $\frac{1}{n} \mathbf{X} \mathbf{X}' \mathbf{u}_j = \hat{\lambda}_j \mathbf{u}_j$. Using

$\mathbf{X} = (\mathbf{Z}_A \Lambda_A^{\frac{1}{2}}, \mathbf{Z}_B \Lambda_B^{\frac{1}{2}})$, we obtain

$$\left(\mathbf{I}_n - \frac{1}{n} \mathbf{Z}'_A \frac{\Lambda_A}{\lambda_j} \mathbf{Z}'_A \right) \mathbf{u}_j = \mathbf{D} \mathbf{u}_j - \Delta \mathbf{u}_j, \tag{A.3}$$

where we denote $\mathbf{D} = (n\lambda_j)^{-1} \mathbf{Z}'_B \Lambda_B \mathbf{Z}'_B - \bar{c}c_j \mathbf{I}_n$, $= \hat{\lambda}_j / \lambda_j - (1 + \bar{c}c_j)$. We then left-multiply equation (A.3) by $\Lambda_A^{\frac{1}{2}} \mathbf{Z}'_A / \sqrt{n\hat{\lambda}_j}$ and employ relationship (A.2) to replace \mathbf{u}_j by $\hat{\xi}_{jA}$ and $\hat{\xi}_{jB}$ as follows:

$$\left(\mathbf{I}_m - \frac{\Lambda_A}{\lambda_j} \right) \hat{\xi}_{jA} = \frac{\Lambda_A^{\frac{1}{2}} (\frac{1}{n} \mathbf{Z}'_A \mathbf{Z}_A - \mathbf{I}_m) \Lambda_A^{\frac{1}{2}}}{\lambda_j} \hat{\xi}_{jA} + \frac{\Lambda_A^{\frac{1}{2}} \mathbf{Z}'_A \mathbf{D} \mathbf{Z}_A \Lambda_A^{\frac{1}{2}}}{n\hat{\lambda}_j} \hat{\xi}_{jA} + \frac{\Lambda_A^{\frac{1}{2}} \mathbf{Z}'_A \mathbf{D} \mathbf{Z}_B \Lambda_B^{\frac{1}{2}}}{n\hat{\lambda}_j} \hat{\xi}_{jB} - \Delta \hat{\xi}_{jA}. \tag{A.4}$$

Further define

$$\mathbf{R} = \sum_{k \in [m] \setminus j} \frac{\lambda_j}{\lambda_j - \lambda_k} \mathbf{e}_{kA} \mathbf{e}'_{kA}.$$

Then we have $\mathbf{R}(\mathbf{I} - \Lambda_A / \lambda_j) = \mathbf{I}_m - \mathbf{e}_{jA} \mathbf{e}'_{jA}$. Note that \mathbf{R} is only well defined if $m > 1$. Therefore, by left multiplying \mathbf{R} to equation (A.4),

$$\hat{\xi}_{jA} - \langle \hat{\xi}_{jA}, \mathbf{e}_{jA} \rangle \mathbf{e}_{jA} = \mathbf{R} \left(\frac{\Lambda_A}{\lambda_j} \right)^{\frac{1}{2}} \mathbf{K} \left(\frac{\Lambda_A}{\lambda_j} \right)^{\frac{1}{2}} \hat{\xi}_{jA} + \mathbf{R} \frac{\Lambda_A^{\frac{1}{2}} \mathbf{Z}'_A \mathbf{D} \mathbf{Z}_B \Lambda_B^{\frac{1}{2}}}{n\hat{\lambda}_j} \hat{\xi}_{jB} - \Delta \mathbf{R} \hat{\xi}_{jA}, \tag{A.5}$$

where $\mathbf{K} = n^{-1} \mathbf{Z}'_A \mathbf{Z}_A - \mathbf{I}_n + \lambda_j (n\hat{\lambda}_j)^{-1} \mathbf{Z}'_A \mathbf{D} \mathbf{Z}_A$. Dividing both side by $\|\hat{\xi}_{jA}\|$, we are able to write

$$\frac{\hat{\xi}_{jA}}{\|\hat{\xi}_{jA}\|} - \mathbf{e}_{jA} = \mathbf{R} \left(\frac{\Lambda_A}{\lambda_j} \right)^{\frac{1}{2}} \mathbf{K} \left(\frac{\Lambda_A}{\lambda_j} \right)^{\frac{1}{2}} \mathbf{e}_{jA} + \mathbf{r}_n, \tag{A.6}$$

where

$$\begin{aligned} \mathbf{r}_n = & \left(\frac{\langle \hat{\boldsymbol{\xi}}_{jA}, \mathbf{e}_{jA} \rangle}{\|\hat{\boldsymbol{\xi}}_{jA}\|} - 1 \right) \mathbf{e}_{jA} + \mathbf{R} \left(\frac{\Lambda_A}{\lambda_j} \right)^{\frac{1}{2}} \mathbf{K} \left(\frac{\Lambda_A}{\lambda_j} \right)^{\frac{1}{2}} \left(\frac{\hat{\boldsymbol{\xi}}_{jA}}{\|\hat{\boldsymbol{\xi}}_{jA}\|} - \mathbf{e}_{jA} \right) \\ & + \mathbf{R} \frac{\Lambda_A^{\frac{1}{2}} \mathbf{Z}'_A \mathbf{D} \mathbf{Z}_B \Lambda_B^{\frac{1}{2}} \hat{\boldsymbol{\xi}}_{jB}}{n \hat{\lambda}_j \|\hat{\boldsymbol{\xi}}_{jA}\|} - \Delta \mathbf{R} \left(\frac{\hat{\boldsymbol{\xi}}_{jA}}{\|\hat{\boldsymbol{\xi}}_{jA}\|} - \mathbf{e}_{jA} \right). \end{aligned} \tag{A.7}$$

Lemma A.4—As $n \rightarrow \infty$, $\|\mathbf{r}_n\| = O_P(\lambda_j^{-1} \sqrt{p/n} + 1/n)$.

By Lemma A.4, \mathbf{r}_n is a smaller order term. Since $(\Lambda_A/\lambda_j)^{\frac{1}{2}} \mathbf{e}_{jA} = \mathbf{e}_{jA}$,

$$\sqrt{n} \left(\frac{\hat{\boldsymbol{\xi}}_{jA}}{\|\hat{\boldsymbol{\xi}}_{jA}\|} - \mathbf{e}_{jA} + O_P \left(\sqrt{\frac{p}{n\lambda_j^2}} \right) \right) = \sqrt{n} \mathbf{R} \left(\frac{\Lambda_A}{\lambda_j} \right)^{\frac{1}{2}} \mathbf{K} \mathbf{e}_{jA} + o_P(1). \tag{A.8}$$

Now let us derive normality of the right hand side of (A.8). According to the definition of \mathbf{R} ,

$$\mathbf{R} \left(\frac{\Lambda_A}{\lambda_j} \right)^{\frac{1}{2}} = \sum_{k \in [m] \setminus j} \frac{\sqrt{\lambda_j \lambda_k}}{\lambda_j - \lambda_k} \mathbf{e}_{kA} \mathbf{e}'_{kA} \rightarrow \sum_{k \in [m] \setminus j} a_{jk} \mathbf{e}_{kA} \mathbf{e}'_{kA}. \tag{A.9}$$

Let $\mathbf{W} = \sqrt{n} \mathbf{K} \mathbf{e}_{jA} = (W_1, \dots, W_m)$ and $\mathbf{W}^{(-j)}$ be the $(m-1)$ -dimensional vector without the j^{th} element in \mathbf{W} . Since the j^{th} diagonal element of \mathbf{R} is zero, $\mathbf{R}(\Lambda_A/\lambda_j)^{\frac{1}{2}} \mathbf{W}$ depends only on $\mathbf{W}^{(-j)}$.

Lemma A.5— $\mathbf{W}^{(-j)} + O_P(\lambda_j^{-1} \sqrt{p/n}) \xrightarrow{d} N(\mathbf{0}, \mathbf{I}_{m-1})$.

Therefore, by Lemma A.5 and Slutsky's theorem,

$$\sqrt{n} \mathbf{R} \left(\frac{\Lambda_A}{\lambda_j} \right)^{\frac{1}{2}} \mathbf{K} \mathbf{e}_{jA} + O_P \left(\sqrt{\frac{p}{n\lambda_j^2}} \right) \xrightarrow{d} N_m \left(\mathbf{0}, \sum_{k \in [m] \setminus j} a_{jk}^2 \mathbf{e}_{kA} \mathbf{e}'_{kA} \right).$$

Together with (A.8), we concludes (3.3) for the case $m > 1$.

Now let us turn to the case of $m = 1$. Since \mathbf{R} is not defined for $m = 1$, we need to find a different derivation. Equivalently, (A.3) can be written as

$$\frac{1}{n} \mathbf{Z}_1 \mathbf{Z}'_1 \mathbf{u}_1 + \frac{1}{n \lambda_1} \mathbf{Z}_B \Lambda_B \mathbf{Z}'_B \mathbf{u}_1 = \frac{\hat{\lambda}_1}{\lambda_1} \mathbf{u}_1.$$

Left-multiplying \mathbf{u}'_1 and using relationship (A.2), we obtain easily

$$\hat{\xi}_{1A}^2 = 1 - \frac{\bar{c}c_1}{\hat{\lambda}_1/\lambda_1} - \frac{\lambda_1}{\hat{\lambda}_1} \mathbf{u}'_1 \mathbf{D} \mathbf{u}_1 = 1 - \frac{\bar{c}c_1}{\hat{\lambda}_1/\lambda_1} + O_P(\lambda_1^{-1} \sqrt{p/n}),$$

where \mathbf{D} is defined as before and $\|\mathbf{D}\| = O_P(\lambda_1^{-1} \sqrt{p/n})$ according to the proof of Lemma A.

4. Expanding $\sqrt{1 - \bar{c}c_1/x}$ at the point of $(1 + \bar{c}c_1)$, we have

$$\hat{\xi}_{1A} = \frac{1}{\sqrt{1 + \bar{c}c_1}} + \frac{\bar{c}c_1}{2(1 + \bar{c}c_1)^{3/2}} (\hat{\lambda}_1/\lambda_1 - (1 + \bar{c}c_1)) + O_P\left(\sqrt{\frac{p}{n\lambda_1^2}} + c_1 n^{-1}\right).$$

Note that from Lemmas A.2 and A.3,

$\hat{\lambda}_1/\lambda_1 - (1 + \bar{c}c_1) = (\|\mathbf{Z}_1\|^2/n - 1) + O_P(\lambda_1^{-1}(p/n)^{1/2}) + O_P(c_j n^{-1/2})$. Therefore due to the fact $\sqrt{n}(\|\mathbf{Z}_1\|^2/n - 1)$ is asymptotically $N(0, \kappa_1 - 1)$, we conclude

$$\frac{2(1 + \bar{c}c_1)^{3/2}}{\bar{c}c_1} \sqrt{n} \left(\hat{\xi}_{1A} - \frac{1}{\sqrt{1 + \bar{c}c_1}} + O_P\left(\sqrt{\frac{p}{n\lambda_1^2}}\right) \right) \xrightarrow{d} N(0, \kappa_1 - 1).$$

This completes the first part of the proof.

(ii) We now prove the conclusion for the non-spiked part $\hat{\xi}_{jB}$. Recall that \mathbf{X}_j follows $N(\mathbf{0}, \Lambda)$.

Consider $\mathbf{X}_i^R = \text{diag}(\mathbf{I}_m, \mathbf{D}_0) \mathbf{X}_i$ where as defined in the theorem

$\mathbf{D}_0 = \text{diag}(\sqrt{\bar{c}/\lambda_{m+1}}, \dots, \sqrt{\bar{c}/\lambda_p})$. Here the superscript R indicates rescaled data by $\text{diag}(\mathbf{I}_m, \mathbf{D}_0)$. After rescaling, we have $\mathbf{X}_i^R \sim N(\mathbf{0}, \text{diag}(\Lambda_A, \bar{c}\mathbf{I}_{p-m}))$. Correspondingly, the $n \times p$ data matrix $\mathbf{X}^R = \mathbf{X} \text{diag}(\mathbf{I}_m, \mathbf{D}_0) = (\mathbf{X}_A, \mathbf{X}_B \mathbf{D}_0)$ where $\mathbf{X}_A = \mathbf{Z}_A \Lambda_A^{\frac{1}{2}}$ and $\mathbf{X}_B = \mathbf{Z}_B \Lambda_B^{\frac{1}{2}}$ as the notations before. Assume $\hat{\xi}_j^R$ and \mathbf{u}_j^R are eigenvectors given by $\hat{\Sigma}^R$ and $\tilde{\Sigma}^R$ of the rescaled data \mathbf{X}^R and $\hat{\xi}_j^R = (\hat{\xi}_{jA}^R, \hat{\xi}_{jB}^R)'$. It has been proved by Paul (2007) that $\mathbf{h}_0 := \hat{\xi}_{jB}^R / \|\hat{\xi}_{jB}^R\|$ is distributed uniformly over the unit sphere and is independent of $\|\hat{\xi}_{jB}^R\|$ due to the orthogonal invariance of the non-spiked part of $\hat{\xi}_{jB}^R$. Hence it only remains to link $\hat{\xi}_{jB}^R / \|\hat{\xi}_{jB}^R\|$ with \mathbf{h}_0 .

Note that $\tilde{\Sigma} = n^{-1} \mathbf{X} \mathbf{X}'$ and $\tilde{\Sigma}^R = n^{-1} \mathbf{X}^R \mathbf{X}^{R'}$, so

$$\|\tilde{\Sigma} - \tilde{\Sigma}^R\| = \left\| \frac{1}{n} \mathbf{X}_B (\mathbf{I} - \mathbf{D}_0^2) \mathbf{X}_B' \right\| = \left\| \frac{1}{n} \sum_{j=m+1}^p (\lambda_j - \bar{c}) \mathbf{Z}_j \mathbf{Z}_j' \right\|,$$

where the last term is of order $O_p(\sqrt{p/n})$ by Lemma A.1. Thus by the sin θ theorem of Davis and Kahan (1970), $\|\mathbf{u}_j - \mathbf{u}_j^R\| = O_p(\lambda_j^{-1} \sqrt{p/n})$. Next we convert from \mathbf{u}_j to $\hat{\boldsymbol{\xi}}_{jB}$ using the basic relationship (2.1). We have,

$$\begin{aligned} & \left\| \mathbf{D}_0 \frac{\hat{\boldsymbol{\xi}}_{jB}}{\|\hat{\boldsymbol{\xi}}_{jB}\|} - \frac{\hat{\boldsymbol{\xi}}_{jB}^R}{\|\hat{\boldsymbol{\xi}}_{jB}^R\|} \right\| = \left\| \frac{\mathbf{D}_0 \mathbf{X}'_B \mathbf{u}_j}{\sqrt{n \hat{\lambda}_j} \|\hat{\boldsymbol{\xi}}_{jB}\|} - \frac{\mathbf{D}_0 \mathbf{X}'_B \mathbf{u}_j^R}{\sqrt{n \hat{\lambda}_j^R} \|\hat{\boldsymbol{\xi}}_{jB}^R\|} \right\| \\ & \leq \left\| \frac{\mathbf{D}_0 \mathbf{X}'_B \mathbf{u}_j}{\sqrt{n \lambda_j}} \right\| \left| \sqrt{\frac{\lambda_j}{\hat{\lambda}_j \|\hat{\boldsymbol{\xi}}_{jB}\|^2}} - \sqrt{\frac{\lambda_j}{\hat{\lambda}_j^R \|\hat{\boldsymbol{\xi}}_{jB}^R\|^2}} \right| + \left\| \frac{\mathbf{D}_0 \mathbf{X}'_B}{\sqrt{n \lambda_j^R} \|\hat{\boldsymbol{\xi}}_{jB}^R\|} \right\| \|\mathbf{u}_j - \mathbf{u}_j^R\| \\ & =: I + II. \end{aligned}$$

First it is not hard to see $II = O_p(\lambda_j^{-1} \sqrt{p/n})$ since

$$\|\mathbf{u}_j - \mathbf{u}_j^R\| = O_p(\lambda_j^{-1} \sqrt{p/n}), \|\mathbf{X}'_B / \sqrt{n \lambda_j}\| = O_p(\sqrt{c_j}), \lambda_j / \hat{\lambda}_j^R = O_p(1) \text{ and } 1 / \|\hat{\boldsymbol{\xi}}_{jB}^R\| = O_p(1 / \sqrt{c_j}).$$

The last result is due to the following lemma.

Lemma A.6— $\|\hat{\boldsymbol{\xi}}_{jA}\| = (1 + \bar{c}c_j)^{-1/2} + O_p(\lambda_j^{-1} \sqrt{p/n} + c_j n^{-1/2})$ and

$$\|\hat{\boldsymbol{\xi}}_{jB}\| = \left(\frac{\bar{c}c_j}{1 + \bar{c}c_j}\right)^{1/2} + O_p(\sqrt{1/\lambda_j} + \sqrt{c_j} n^{-1/2}).$$

We claim $I = O_p(\sqrt{n/p}) + o_p(n^{-1/2})$. Actually from the proof of Lemma A.6, we have

$$\hat{\lambda}_j \|\hat{\boldsymbol{\xi}}_{jB}\|^2 / \lambda_j = \bar{c}c_j + O_p(\lambda_j^{-1} \sqrt{p/n}) + o_p(c_j n^{-1/2}).$$

Then some elementary calculation gives the rate of I . Therefore,

$$\|\mathbf{D}_0 \hat{\boldsymbol{\xi}}_{jB} / \|\hat{\boldsymbol{\xi}}_{jB}\| - \mathbf{h}_0\| = O_p(\sqrt{n/p}) + o_p(n^{-1/2}).$$

The conclusion (3.4) follows.

To prove the max norm bound (3.5) of $\|\hat{\boldsymbol{\xi}}_{jB}\|_{\max}$, we first show $\|\mathbf{h}_0\|_{\max} = O_p(\sqrt{\log p/p})$. Recall that \mathbf{h}_0 is uniformly distributed on the unit sphere of dimension $p - m$. This follows easily from its normal representation. Let \mathbf{G} to be $(p - m)$ -dimensional multivariate standard normal distributed, then $\mathbf{h}_0 \stackrel{d}{=} \mathbf{G} / \|\mathbf{G}\|$. It then follows

$$\|\mathbf{h}_0\|_{\max} = \max_{i \leq p-m} |G_i| / \|\mathbf{G}\| = O_p(\sqrt{\log p/p}).$$

From the derivation above,

$$\|\hat{\boldsymbol{\xi}}_{jB}\|_{\max} \leq \sqrt{\hat{\lambda}_j^R / \lambda_j} \|\mathbf{D}_0^{-1}\| \|\hat{\boldsymbol{\xi}}_{jB}^R\| (II + \|\mathbf{h}_0\|_{\max}),$$

which gives $O_P(\sqrt{c_j}(\sqrt{p/(n\lambda_j^2)} + \sqrt{\log p/p})) = O_P(p/n\lambda_j^{3/2} + \sqrt{\log p/(n\lambda_j)})$, given the fact that $\|\hat{\xi}_{jB}^R\| = O_P(\sqrt{c_j})$ by Lemma A.6. Thus we are done with the second part of the proof.

(iii) The proof for the convergence of $\|\hat{\xi}_{jA}\|$ and $\|\hat{\xi}_{jB}\|$ are given in Lemma A.6. If $m = 1$, the result for $\|\hat{\xi}_{jA}\|$ directly gives (3.6) with the same rate. For $m > 1$, from Lemma A.6 we have

$$\|\hat{\xi}_{jA}\|^2 = (1 + \bar{c}c_j)^{-1} + O_P\left(\sqrt{p/(n\lambda_j^2)} + c_j n^{-1/2}\right).$$

On the other hand, from Theorem 3.2 (i), $\hat{\xi}_{jk}^2 = O_P(p/(n\lambda_j^2) + 1/n)$ for $k \neq j$. So $\hat{\xi}_{j1}^2 = (1 + \bar{c}c_j)^{-1} + O_P(\sqrt{p/(n\lambda_j^2)} + c_j n^{-1/2} + 1/n)$, which implies (3.6).

REFERENCES

- Agarwal A, Negahban S, Wainwright MJ. Noisy matrix decomposition via convex relaxation: Optimal rates in high dimensions. *The Annals of Statistics*. 2012; 40:1171–1197.
- Amini, AA., Wainwright, MJ. Information Theory 2008. ISIT 2008. IEEE International Symposium on. IEEE; 2008. High-dimensional analysis of semidefinite relaxations for sparse principal components; p. 2454–2458.
- Anderson TW. Asymptotic theory for principal component analysis. *The Annals of Mathematical Statistics*. 1963; 34:122–148.
- Antoniadis A, Fan J. Regularization of wavelet approximations. *Journal of the American Statistical Association*. 2001; 96
- Bai Z. Methodologies in spectral analysis of large-dimensional random matrices, a review. *Statist. Sinica*. 1999; 9:611–677.
- Bai J. Inferential theory for factor models of large dimensions. *Econometrica*. 2003; 71:135–171.
- Bai J, Ng S. Determining the number of factors in approximate factor models. *Econometrica*. 2002; 70:191–221.
- Bai, Z., Silverstein, JW. Spectral analysis of large dimensional random matrices. 2nd. Springer; 2009.
- Bai Z, Yao J. on sample eigenvalues in a generalized spiked population model. *Journal of Multivariate Analysis*. 2012; 106:167–177.
- Bai Z, Yin Y. Limit of the smallest eigenvalue of a large dimensional sample covariance matrix. *The Annals of Probability*. 1993:1275–1294.
- Baik J, Ben Arous G, Péché S. Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *The Annals of Probability*. 2005:1643–1697.
- Benaych-Georges F, Nadakuditi RR. The eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices. *Advances in Mathematics*. 2011; 227:494–521.
- Berthet Q, Rigollet P. Optimal detection of sparse principal components in high dimension. *The Annals of Statistics*. 2013; 41:1780–1815.
- Bickel PJ, Levina E. Covariance regularization by thresholding. *The Annals of Statistics*. 2008:2577–2604.
- Birbaum A, Johnstone IM, Nadler B, Paul D. Minimax bounds for sparse PCA with noisy high-dimensional data. *The Annals of Statistics*. 2013; 41:1055. [PubMed: 25324581]
- Cai T, Fan J, Jiang T. Distributions of angles in random packing on spheres. *The Journal of Machine Learning Research*. 2013; 14:1837–1864. [PubMed: 25324693]
- Cai T, Ma Z, Wu Y. Optimal estimation and rank detection for sparse spiked covariance matrices. *Probability Theory and Related Fields*. 2015; 161:781–815. [PubMed: 26257453]

- Candès EJ, Li X, Ma Y, Wright J. Robust principal component analysis? *Journal of the ACM (JACM)*. 2011; 58:11.
- Chamberlain G, Rothschild M. Arbitrage, factor structure, and mean-variance analysis on large asset markets. *Econometrica*. 1983; 51:1305–1324.
- Chandrasekaran V, Sanghavi S, Parrilo PA, Willsky AS. Rank-sparsity incoherence for matrix decomposition. *SIAM Journal on Optimization*. 2011; 21:572–596.
- Chen KH, Shimerda TA. An empirical analysis of useful financial ratios. *Financial Management*. 1981:51–60.
- Davidson, KR., Szarek, SJ. Local operator theory, random matrices and Banach spaces. In: Johnson, WB., Lindenstrauss, J., editors. *Handbook of the Geometry of Banach Spaces*. Vol. 1. Elsevier Science BV; 2001. p. 317-366.
- Davis C, Kahan WM. The rotation of eigenvectors by a perturbation. III. *SIAM Journal on Numerical Analysis*. 1970; 7:1–46.
- De Mol C, Giannone D, Reichlin L. Forecasting using a large number of predictors: Is Bayesian shrinkage a valid alternative to principal components? *Journal of Econometrics*. 2008; 146:318–328.
- Donoho DL, Gavish M, Johnstone IM. Optimal shrinkage of eigenvalues in the Spiked Covariance Model. arXiv preprint arXiv:1311.0851. 2014
- Fan J, Fan Y, Lu J. High dimensional covariance matrix estimation using a factor model. *Journal of Econometrics*. 2008; 147:186–197.
- Fan J, Han X, Gu W. Estimating false discovery proportion under arbitrary covariance dependence. *Journal of the American Statistical Association*. 2012; 107:1019–1035. [PubMed: 24729644]
- Fan J, Han X. Estimation of false discovery proportion with unknown dependence. arXiv preprint arXiv:1305.7007. 2013
- Fan J, Liao Y, Mincheva M. Large covariance estimation by thresholding principal orthogonal complements. *Journal of the Royal Statistical Society: Series B*. 2013; 75:1–44.
- Fan J, Liao Y, Shi X. Risks of large portfolios. *Journal of Econometrics*. 2015; 186:367–387. [PubMed: 26195851]
- Fan J, Liao Y, Wang W. Projected Principal Component Analysis in Factor Models. *The Annals of Statistics*. 2016; 44:219–254. [PubMed: 26783374]
- Fan J, Xue L, Yao J. Sufficient Forecasting Using Factor Models. arXiv preprint arXiv:1505.07414. 2015
- Fan J, Liu H, Wang W, Zhu Z. Heterogeneity Adjustment with Applications to Graphical Model Inference. arXiv preprint arXiv:1602.05455. 2016
- Hall P, Marron J, Neeman A. Geometric representation of high dimension, low sample size data. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2005; 67:427–444.
- Johnstone IM. On the distribution of the largest eigenvalue in principal components analysis. *The Annals of Statistics*. 2001:295–327.
- Johnstone IM, Lu AY. On Consistency and Sparsity for Principal Components Analysis in High Dimensions. *Journal of the American Statistical Association*. 2009; 104:682–693. [PubMed: 20617121]
- Jung S, Marron J. PCA consistency in high dimension, low sample size context. *The Annals of Statistics*. 2009; 37:4104–4130.
- Koltchinskii V, Lounici K. Concentration Inequalities and Moment Bounds for Sample Covariance Operators. arXiv preprint arXiv:1405.2468. 2014a
- Koltchinskii V, Lounici K. Asymptotics and Concentration Bounds for Bilinear Forms of Spectral Projectors of Sample Covariance. arXiv preprint arXiv:1408.4643. 2014b
- Landgrebe J, Wurst W, Welzl G. Permutation-validated principal components analysis of microarray data. *Genome Biol*. 2002; 3:1–11.
- Lee S, Zou F, Wright FA. Convergence and prediction of principal component scores in high-dimensional settings. *The Annals of Statistics*. 2010; 38:3605. [PubMed: 21442047]
- Leek JT, Storey JD. A general framework for multiple testing dependence. *Proceedings of the National Academy of Sciences*. 2008; 105:18718–18723.

- Leek JT, Scharpf RB, Bravo HC, Simcha D, Langmead B, Johnson WE, Geman D, Baggerly K, Irizarry RA. Tackling the widespread and critical impact of batch effects in high-throughput data. *Nature Reviews Genetics*. 2010; 11:733–739.
- Ma Z. Sparse principal component analysis and iterative thresholding. *The Annals of Statistics*. 2013; 41:772–801.
- Onatski A. Asymptotics of the principal components estimator of large factor models with weakly influential factors. *Journal of Econometrics*. 2012; 168:244–258.
- Paul D. Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statistica Sinica*. 2007; 17:1617–1642.
- Pesaran MH, Zaffaroni P. Optimal asset allocation with factor models for large portfolios. 2008
- Price AL, Patterson NJ, Plenge RM, Weinblatt ME, Shadick NA, Reich D. Principal components analysis corrects for stratification in genome-wide association studies. *Nature Genetics*. 2006; 38:904–909. [PubMed: 16862161]
- Ringnér M. What is principal component analysis? *Nature Biotechnology*. 2008; 26:303–304.
- Rothman AJ, Levina E, Zhu J. Generalized thresholding of large covariance matrices. *Journal of the American Statistical Association*. 2009; 104:177–186.
- Shen D, Shen H, Zhu H, Marron J. Surprising asymptotic conical structure in critical sample eigen-directions. arXiv preprint arXiv:1303.6171. 2013
- Stock JH, Watson MW. Forecasting using principal components from a large number of predictors. *Journal of the American statistical association*. 2002; 97:1167–1179.
- Thomas CG, Harshman RA, Menon RS. Noise reduction in BOLD-based fMRI using component analysis. *Neuroimage*. 2002; 17:1521–1537. [PubMed: 12414291]
- Vershynin R. Introduction to the non-asymptotic analysis of random matrices. arXiv preprint arXiv:1011.3027. 2010
- Vu VQ, Lei J. Minimax rates of estimation for sparse PCA in high dimensions. arXiv preprint arXiv:1202.0786. 2012
- Wang W, Fan J. Supplementary appendix to the paper “Asymptotics of empirical eigen-structure for high dimensional spiked covariance”. 2015
- Yamaguchi-Kabata Y, Nakazono K, Takahashi A, Saito S, Hosono N, Kubo M, Nakamura Y, Kamatani N. Japanese population structure, based on SNP genotypes from 7003 individuals compared to other ethnic groups: effects on population-based association studies. *The American Journal of Human Genetics*. 2008; 83:445–456. [PubMed: 18817904]
- Yata K, Aoshima M. Effective PCA for high-dimension, low-sample-size data with noise reduction via geometric representations. *Journal of Multivariate Analysis*. 2012; 105:193–215.
- Yata K, Aoshima M. PCA consistency for the power spiked model in high-dimensional settings. *Journal of Multivariate Analysis*. 2013; 122:334–354.

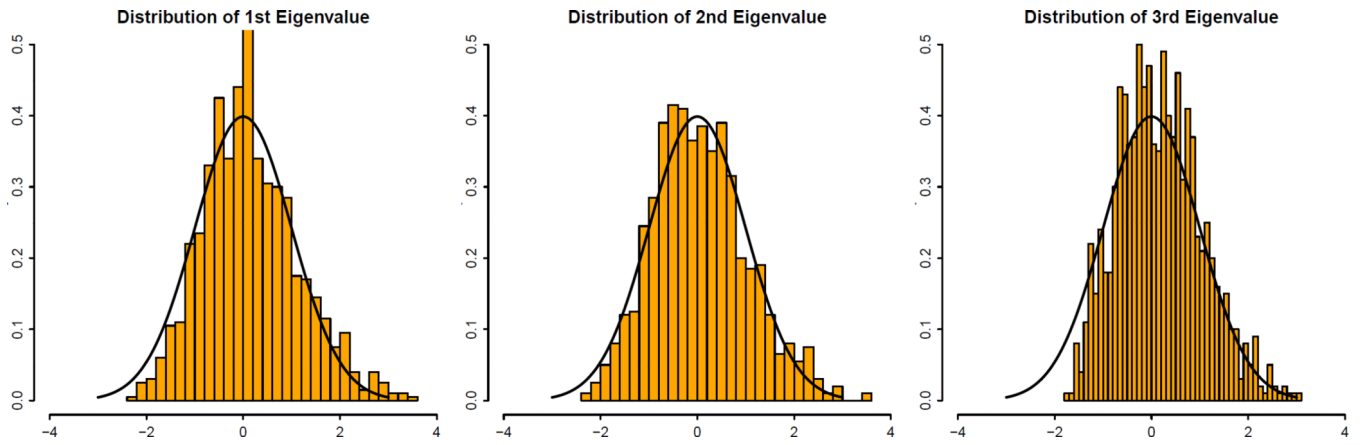


Fig 1.

Behavior of empirical eigenvalues. The empirical distributions of $\sqrt{n/2}(\hat{\lambda}_j/\lambda_j - 1 - c_j)$ for $j = 1, 2, 3$ are compared with their asymptotic distributions $\mathcal{N}(0, 1)$.

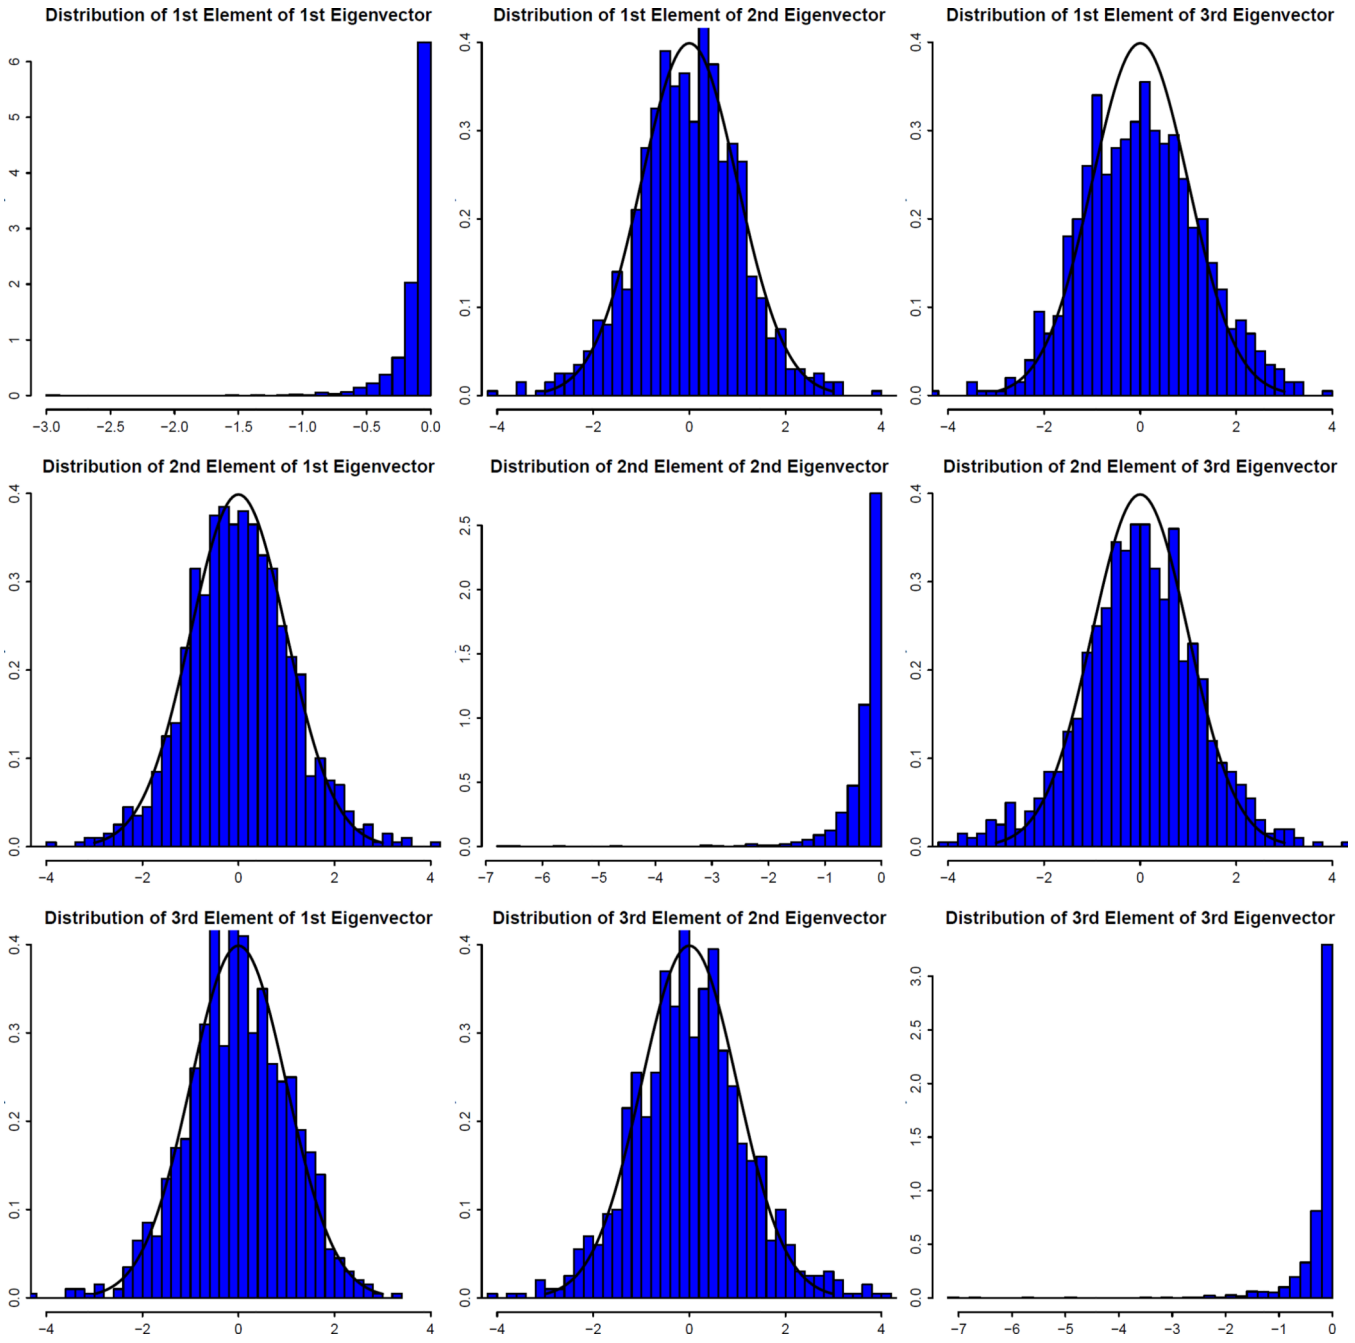


Fig 2. Behavior of empirical eigenvectors. The histogram of the k^{th} element of the j^{th} empirical vector is depicted in the location (k, j) for $k, j = 1, 2, 3$. Off-diagonal plots of values $\sqrt{n}\hat{\xi}_{jk}/\|\hat{\xi}_{jA}\|/\sqrt{\frac{c_j c_k}{(c_j - c_k)^2}}$ are compared to their asymptotic distributions $N(0, 1)$ for $k \neq j$ while diagonal plots of values $\sqrt{n}(\hat{\xi}_{jj}/\|\hat{\xi}_{jA}\| - 1)$ are compared to stochastically 0.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

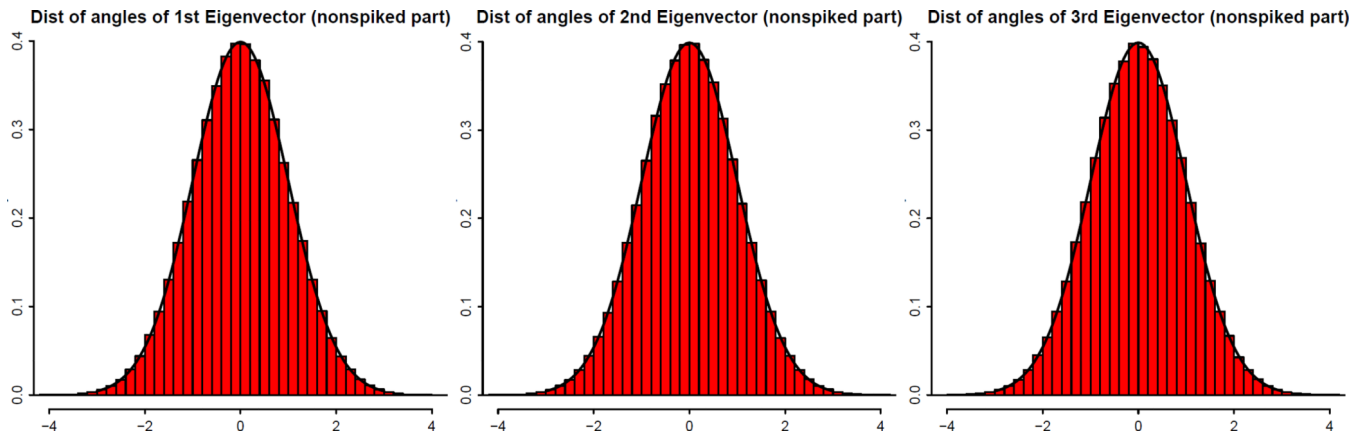


Fig 3. The empirical distributions of all pairwise angles of the 1000 realized $\hat{\xi}_{jB}/\|\hat{\xi}_{jB}\|$ ($j = 1, 2, 3$) compared with their asymptotic distributions $N(0, 1)$.

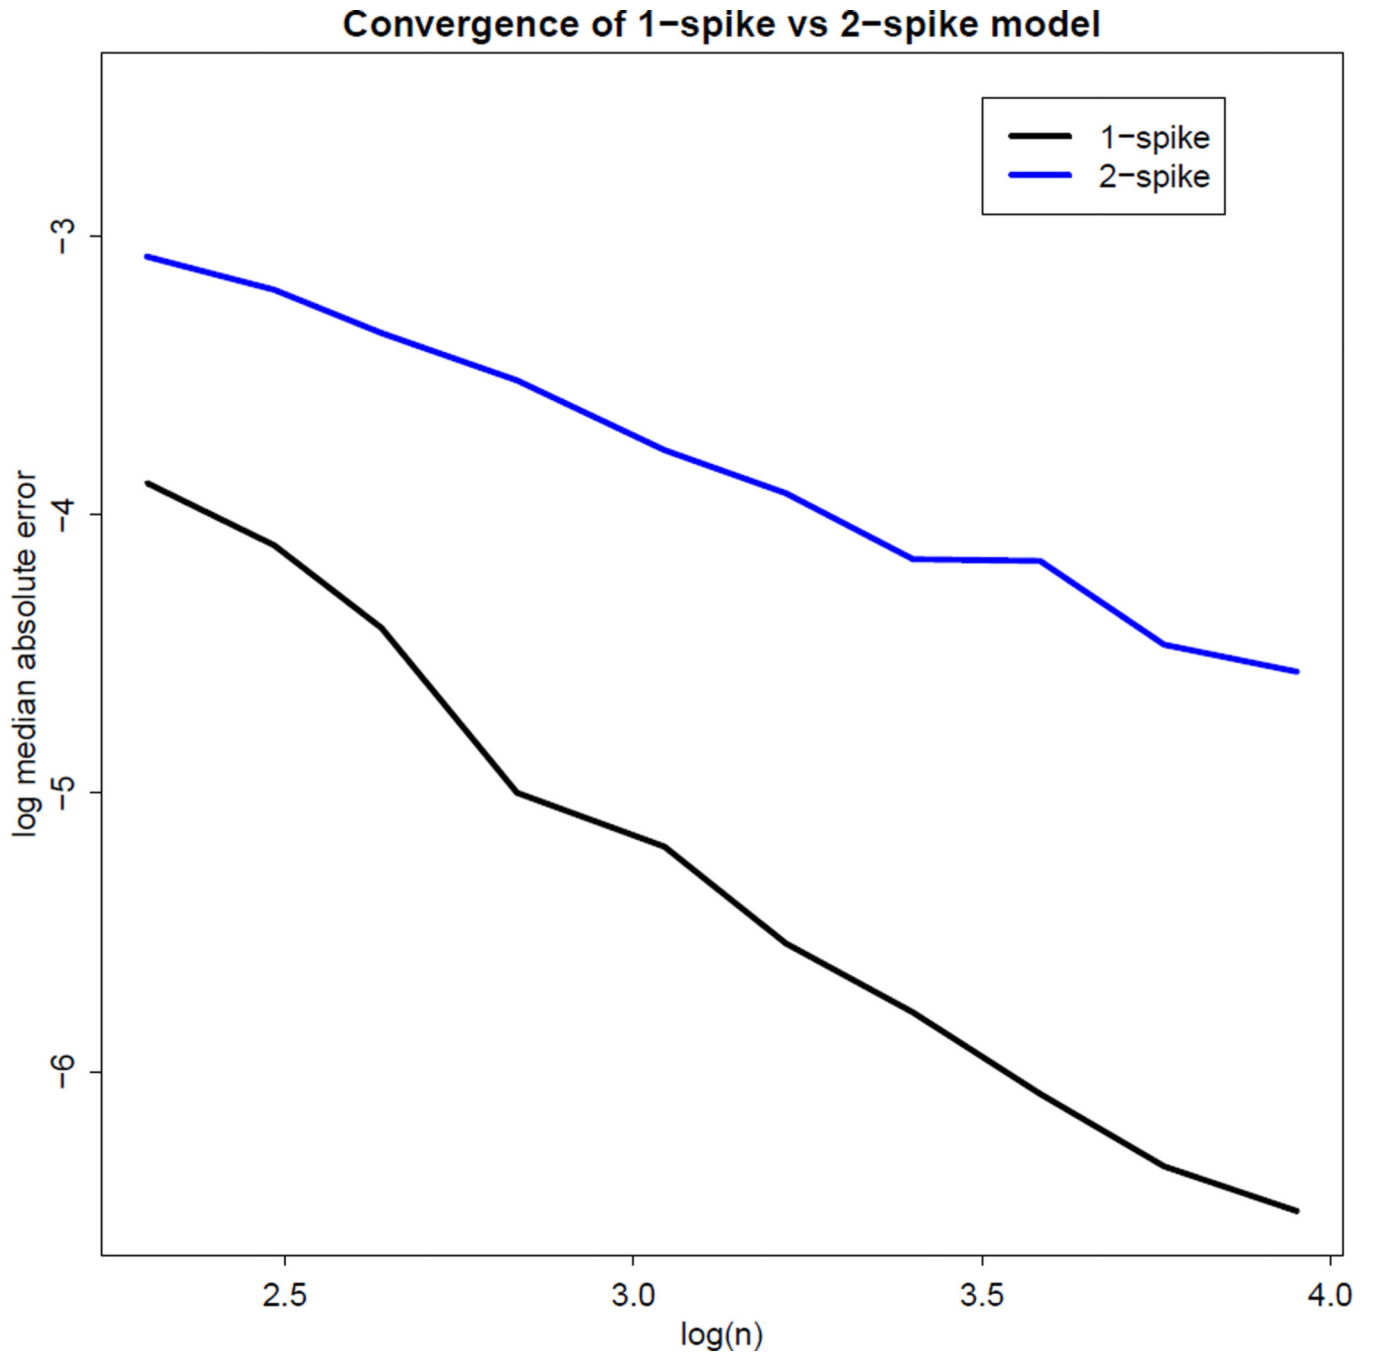


Fig 4.

Difference of convergence rate of $\langle \hat{\xi}_1, e_1 \rangle - 1/\sqrt{1+c_1}$ for models with a single spike and two spikes. The error should be expected to decrease at the rate of $O_P(n^{-3/2})$ and $O_P(n^{-1})$ respectively.

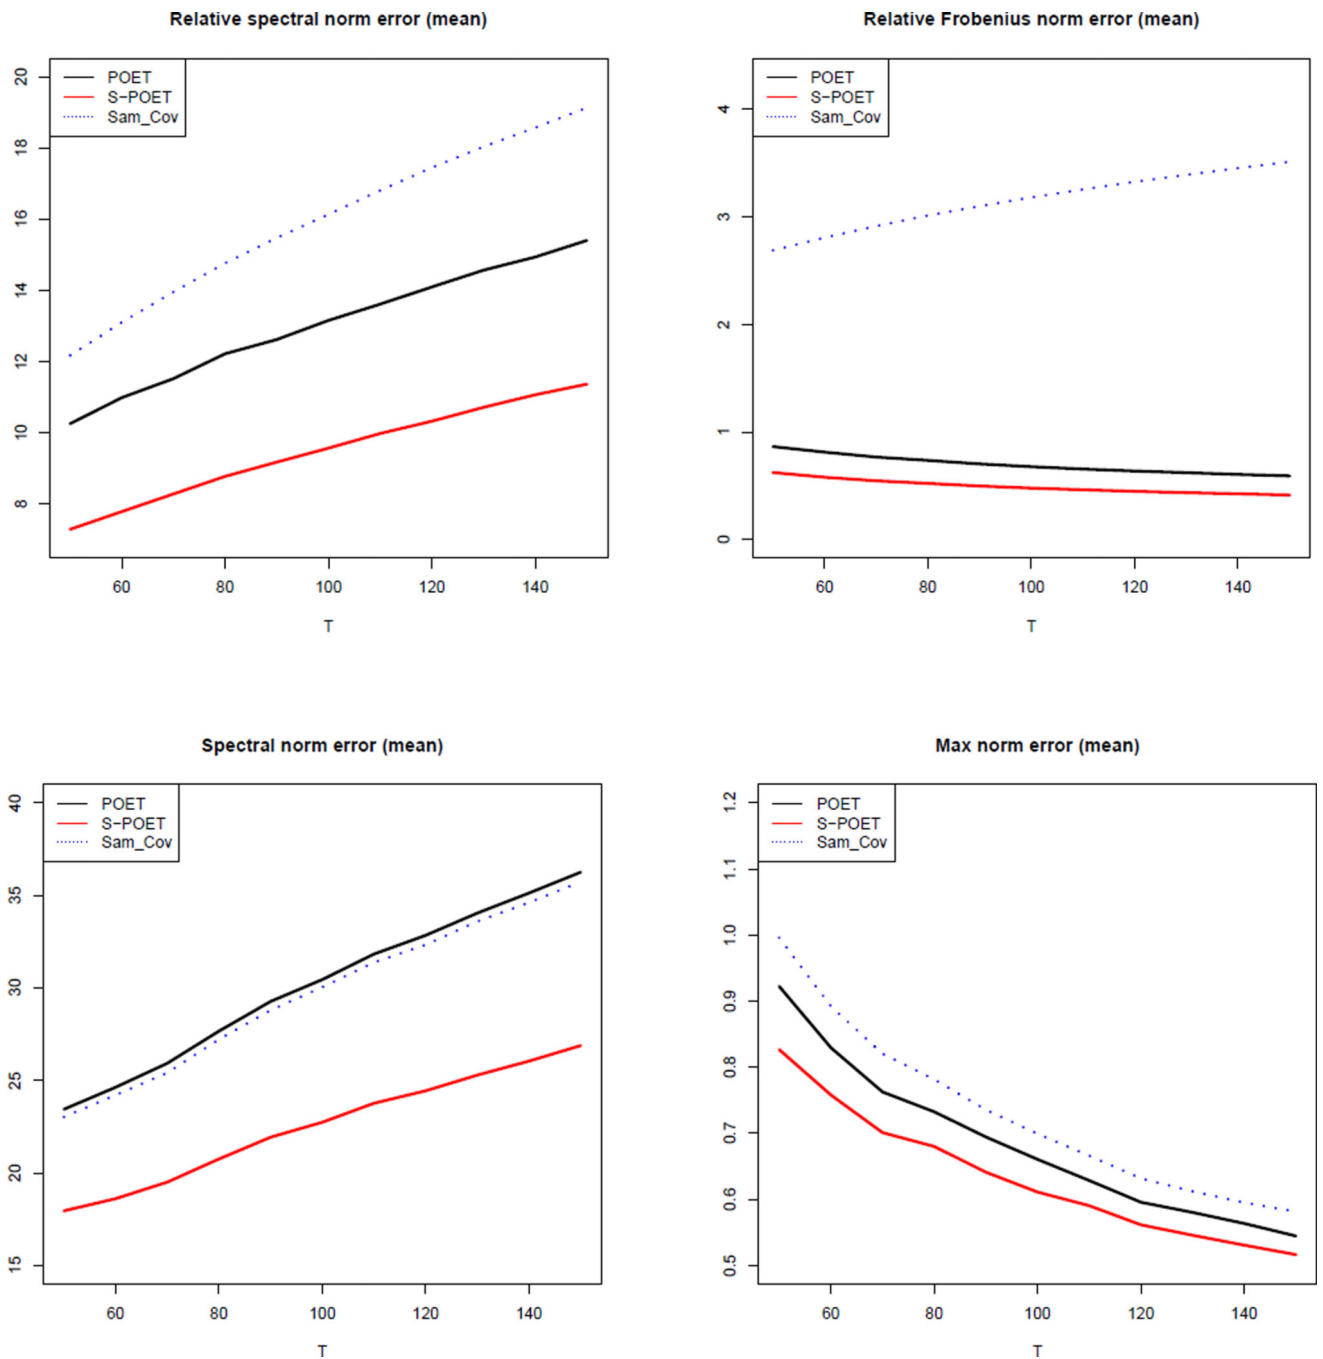


Fig 5. Estimation errors of covariance matrix under relative spectral, relative Frobenius, spectral and max norms using S-POET (red), POET (black) and sample covariance (blue).

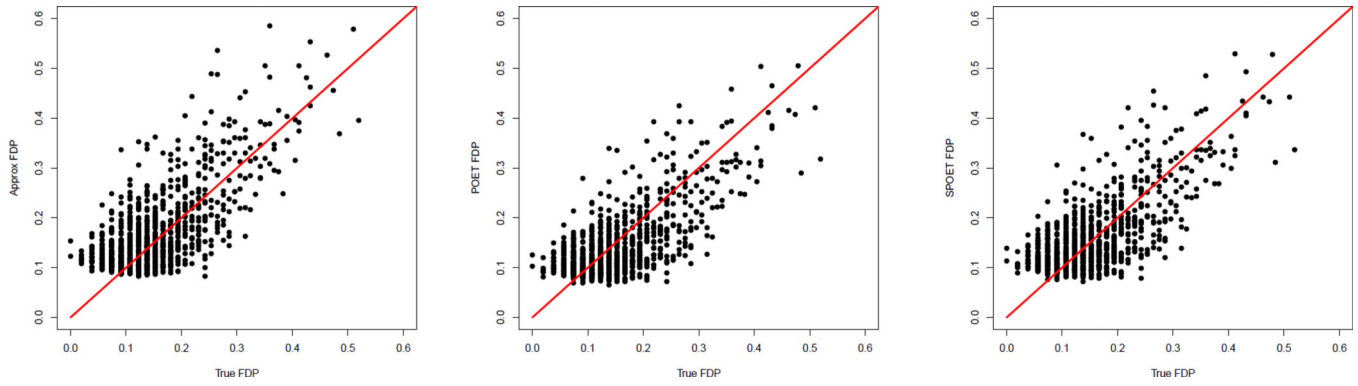


Fig 6. Comparison of estimated FDP's with true values. The left plot assumes knowledge of B , the middle and right ones are corresponding to POET and S-POET methods respectively. The results are aligned along the 45-degree line, indicating the accuracy of the estimated FDP.

Table 1

The correlations between the first three elements for each of the three empirical eigenvectors based on 1000 repetitions

	1st & 2nd elements	1st & 3rd elements	2nd & 3rd elements
1st Eigenvector	0.00156	-0.00192	-0.04112
2nd Eigenvector	-0.02318	-0.00403	0.01483
3rd Eigenvector	-0.02529	-0.04004	0.12524

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript