Published in final edited form as:

J Econom. 2019 January ; 208(1): 5–22. doi:10.1016/j.jeconom.2018.09.003.

Robust Covariance Estimation for Approximate Factor Models

Jianqing Fan^{*,†}, Weichen Wang^{*}, and Yiqiao Zhong^{*}

^{*}Department of Operations Research and Financial Engineering, Sherrerd Hall, Princeton University, Princeton, NJ 08544, USA

[†]School of Data Science, Fudan University

Abstract

In this paper, we study robust covariance estimation under the approximate factor model with observed factors. We propose a novel framework to first estimate the initial joint covariance matrix of the observed data and the factors, and then use it to recover the covariance matrix of the observed data. We prove that once the initial matrix estimator is good enough to maintain the element-wise optimal rate, the whole procedure will generate an estimated covariance with desired properties. For data with only bounded fourth moment, we propose to use adaptive Huber loss minimization to give the initial joint covariance estimation. This approach is applicable to a much wider class of distributions, beyond sub-Gaussian and elliptical distributions. We also present an asymptotic result for adaptive Huber's M-estimator with a diverging parameter. The conclusions are demonstrated by extensive simulations and real data analysis.

Keywords

Robust covariance matrix; Approximate factor model; M-estimator

1 Introduction

The problem of estimating a covariance matrix and its inverse has been fundamental in many areas of statistics and econometrics, including principal component analysis (PCA) and undirected graphical models for instance. The intense research in high dimensional statistics has contributed a stream of papers related to covariance matrix estimation, including sparse principal component analysis (Johnstone and Lu, 2009; Amini and Wainwright, 2008; Vu and Lei, 2012; Birnbaum et al., 2013; Berthet and Rigollet, 2013; Ma, 2013; Cai et al., 2013), sparse covariance estimation (Bickel and Levina, 2008; Cai and Liu, 2011; Cai et al., 2010; Lam and Fan, 2009; Ravikumar et al., 2011) and factor model analysis (Stock and Watson, 2002; Bai, 2003; Fan et al., 2008, 2013, 2016; Onatski, 2012). A strong interest in precision matrix estimation (undirected graphical model) has also emerged in the statistics community following the pioneering works in Meinshausen and Bühlmann (2006) and Friedman et al. (2008). In the application aspect, many areas such as portfolio allocation (Fan et al., 2008), have benefited from this continuing research.

In the high dimensional setting, the number of variables p is comparable or greater than the sample size n . This dimensionality poses a challenge to the estimation of covariance matrices. It has been shown in Johnstone and Lu (2009) that the empirical covariance matrix

behaves poorly, and sparsity of leading eigenvectors circumvents this issue. Following this work, a flourishing literature on sparse PCA has developed in-depth analysis and refined algorithms; see Vu and Lei (2012); Berthet and Rigollet (2013); Ma (2013). Taking a different route, Bickel and Levina (2008) advocated thresholding as a regularization approach to estimate a sparse matrix, in the sense that most entries of the matrix are close to zero and this approach was used independently in (Fan et al., 2008) for estimating covariance matrix with factor structure.

Another challenge in high-dimensional statistics is that measurements may not have light tails. For example, large scale datasets are often obtained by using bio-imaging technology (e.g., fMRI and microarrays) that often leads to heavy-tailed measurement errors (Dinov et al., 2005). Moreover, it is well known that financial returns exhibit heavy tails. These invalidate the fundamental assumptions in high-dimensional statistics that data have sub-Gaussian or sub-exponential tails, popularly imposed in most of the aforementioned papers. Significant relaxation of the assumption requires some new ideas and forms the subject of this paper.

Recently, motivated by Fama-French model (Fama and French, 1993) from financial econometrics, Fan et al. (2008) and Fan et al. (2013) considered the covariance structure of the *static approximate factor model*, which models the covariance matrix by a low-rank signal matrix and a sparse noise matrix. The same model will also be the focus of the current paper. The model assumes existence of several low-dimensional factors that drives a large panel data $\{y_{it}\}_{i=1, \dots, p, t=1, \dots, n}$ that is

$$y_{it} = b_i^T f_t + u_{it}, \quad i \leq p, t \leq n, \quad (1.1)$$

where f_t 's are the common factors, which are observed; and b_i 's are their corresponding factor loadings, which are considered unknown but fixed parameters in this work. The noises u_{it} 's, known as the idiosyncratic component, are uncorrelated with the factors $f_t \in \mathbf{R}^r$. Here r is relatively small compared with p and n . We will treat r as fixed independent of p and n throughout this paper. When the factors are known, this model subsumes the well-known CAPM model (Sharpe, 1964; Lintner, 1965) and Fama-French model (Fama and French, 1993). When f_t is unobserved, the model tries to recover the underlying factors for the movements of the whole panel data. Here the “approximate” factor model indicates that the covariance Σ_u of $u_t = (u_{1t}, \dots, u_{pt})'$ is sparse, including the strict factor model in which Σ_u is diagonal as a special case. In addition, “static” is a specific case of the dynamic model which takes into account the time lag and allows more general infinite dimensional representations (Forni et al., 2000; Forni and Lippi, 2001).

The covariance matrix of the outcome $y_t = (y_{1t}, \dots, y_{pt})'$ from model (1.1) can be written as

$$\Sigma = B \Sigma_f B^T + \Sigma_u, \quad (1.2)$$

where $B_{p \times r}$ consisting of b'_i in each row is the loading matrix, Σ_f is the covariance of f_t and Σ_u is the sparse covariance matrix for u_t . Here we assume the process of (f_t, u_t) is stationary so that Σ_f, Σ_u do not change over time. When factors are unknown, Fan et al. (2013) proposed applying PCA to obtain an estimate of the low rank part and sparse part Σ_u . The crucial assumption is that the factors are *pervasive*, meaning that the factors have non-negligible effects on a large amount of dimensions of the outcomes. Wang and Fan (2017) gave more explanation from random matrix theories and aims to relax the pervasiveness assumption in applications such as risk management and estimation of the false discovery proportion. See Onatski (2012) for more discussions on strong and weak factors.

In this paper, we consider estimating Σ with known factors. Unknown factors pose more difficulties for robust estimation, which will be explored in future works. The main focus of the paper is on robustness instead of factor recovery. Under exponential tails of the factors and noises, Fan et al. (2011) proposed the idea of performing thresholding on the estimate of Σ_u , obtained from the sample covariance of the residuals of multiple regression (1.1). The legitimacy of this approach hinges on the assumption that the tails of the factor and error distributions decay exponentially, which is likely to be violated in practice, especially in the financial applications. Thus, the need to extend the applicability of this approach beyond well-behaved noise has driven further research such as Fan et al. (2015b), in which they assume that y_t has an elliptical distribution (Fang et al., 1990).

This paper studies model (1.1) under a much more relaxed condition: the random variables f_t and u_{it} only have finite fourth moment. The main observation that motivates our method is that, the joint covariance matrix of $(y_t^T, f_t^T)^T$ supplies sufficient information to estimate $B\Sigma_f B^T$ and Σ_u . To estimate the joint covariance matrix in a robust way, the classical idea that dates back to Huber (1964) proves to be vital and effective. The novelty here is that we let the parameter diverge in order to control the bias in high-dimensional applications. The Huber loss function with a diverging parameter, together with other similar functions, has been shown to produce concentration bounds for M-estimators, when the random variables have fat tails; see for example Catoni (2012) and Fan et al. (2017). This point will be clarified in Sections 2 and 3. The M-estimators considered here have additional merits in asymptotic analysis, which is studied in Section 3.3.

This paper can be placed in the broader context of low rank plus sparse representation. In the past few years, robust principal component analysis has received much attention among statisticians, applied mathematicians and computer scientists. Their focus is on identifying the low rank component and sparse component from a corrupted matrix (Chandrasekaran et al., 2011; Candès et al., 2011; Xu et al., 2010). However, the matrices considered therein do not come from random samples, and as a result, neither estimation nor inference are involved. Agarwal et al. (2012) does consider the noisy decomposition, but still focuses more on identifying and separating the low rank part and sparse part. In spite of connections with the robust PCA literature, such as the incoherence condition (see Section 2), this paper and its predecessors are more engaged in disentangling “true signal” from noise, in order to improve estimation of covariance matrices. In this respect, they bear more similarity with the literature of covariance matrix estimation.

We make a few notational definitions before presenting the main results. For a general matrix M , the max-norm of M , or the entry-wise maximum, is denoted as $\|M\|_\infty = \max_{i,j} |M_{ij}|$. The operator norm of M is $\|M\| = \lambda_{\max}^{1/2}(M^T M)$ whereas the Frobenius norm is $\|M\|_F = \sqrt{\sum_{i,j} M_{ij}^2}$. If M is furthermore symmetric, we denote $\lambda_j(M)$ as the j^{th} largest eigenvalue, $\lambda_{\max}(M)$ as the largest one, and $\lambda_{\min}(M)$ as the smallest one. In the paper, C is a generic constant that may differ from line to line in the assumptions and also derivation of our theories.

The paper is organized as follows. In Section 2, we present the procedure for robust covariance estimation when only finite fourth moment is assumed for both factors and noises without specific distributional assumption. The theoretical justification will be provided in Section 3. Simulations will be carried out in Section 4 to demonstrate the effectiveness of the proposed procedure. We also conduct real data analysis on portfolio risk of S&P stocks via Fama-French model in Section 5. Technical proofs will be delayed to the appendix.

2 Robust covariance estimation

Consider the factor model (1.1) again with observed factors. It can be written in the vector form as

$$y_t = B f_t + u_t, \quad (2.1)$$

where $y_t = (y_{1t}, \dots, y_{pt})^T$, $f_t \in \mathbf{R}^r$ are the factors for $t = 1, \dots, T$, $B = (b_1, \dots, b_p)^T$ is the fixed unknown loading matrix and $u_t = (u_{1t}, \dots, u_{pt})^T$ is uncorrelated with the factors. We assume that (u_t^T, f_t^T) have zero mean and independent for $t = 1, \dots, T$. A motivating example from economic and financial studies is the classical Fama-French model, where y_{it} 's represent excess returns of stocks in the market and f_t 's are interpreted as common factors driving the market. It is more natural to allow for weak temporal dependence such as α -mixing as in the work of Fan et al. (2016). Though possible, we assume independence in this paper for the sake of simplicity of analysis.

2.1 Assumptions

We now state the main assumptions of the model. Let Σ_f be the covariance of f_b and Σ_u the covariance of u_t . A covariance decomposition shows that Σ , the covariance of y_b comprises two parts,

$$\Sigma = B \Sigma_f B^T + \Sigma_u. \quad (2.2)$$

We assume that Σ_u is sparse and the sparsity level is measured through

$$m_q = \max_{i \leq p} \sum_{j \leq p} |(\sum_u)_{ij}|^q, \text{ for some } q \in [0, 1]. \quad (2.3)$$

If $q = 0$, m_q is defined to be $\max_{i \leq p} \sum_{j \leq p} \mathbb{1}((\sum_u)_{ij} \neq 0)$, i.e. the exact sparsity. An intuitive justification of the sparsity measurement stems from modeling of the covariance structure: after taking out the common factors, the rest only has weak cross-sectional dependence. In addition, we assume that $\|\sum_u\|$, as well as $\|\sum_f\|$, is bounded away from 0 and ∞ . In the case of degenerate $\sum_f$ we can always consider rescaling the factors and reduce the number of observed factors to meet the requirement of non-vanishing minimum eigenvalue of $\sum_f$. This leads to our first assumption.

Assumption 2.1—There exists a constant $C > 0$ such that $C^{-1} \leq \|\sum_u\| \leq C$ and $C^{-1} \leq \|\sum_f\| \leq C$, where $\sum_f$ is a $r \times r$ matrix with r being a fixed number.

Here assuming a fixed r is just for simplicity of presentation. It can be allowed to grow with n and p . Then we would need to keep track of r in the theoretical analysis and impose certain growth condition on r .

Another important feature of the factor model, observed by Stock and Watson (2002), is that the factors are *pervasive* in the sense that the low rank part of (2.2) is the dominant component of Σ ; more specifically, the top r eigenvalues grow linearly as p . This motivates the following assumption.

Assumption 2.2

- i. There exists a constant $c > 0$ such that $\lambda_r(\Sigma) > cp$.
- ii. The elements of B are uniformly bounded by a constant C .

First note assumption (ii) implies that

$\lambda_1(\Sigma) \leq \lambda_1(B \sum_f B^T) + \|\sum_u\| \leq \lambda_1(\sum_f) \lambda_1(B^T B) + \|\sum_u\| = O(p)$. So together with (i), the above assumption requires leading eigenvalues to grow with an order of p . This assumption is satisfied by the approximate factor model, since by Weyl's inequality,

$\lambda_i(\Sigma)/p = \lambda_i(B \sum_f B^T)/p + o(1)$ if the main term is bounded from below. Furthermore, for illustrative purposes, if we additionally assume (though not needed in this paper) that each entry of B is iid with a finite second moment, it is not hard to see

$\lambda_i(B \sum_f B^T)/p = \lambda_i(\sum_f (B^T B/p))$ satisfies such a condition with probability tending to one.

Consequently, it is natural to assume $\lambda_i(\Sigma)/p$ is lower bounded for $i \leq r$. Note that B is considered to be deterministic throughout the paper.

Assumption (ii) is related to the matrix incoherence condition. In fact, when $\lambda_{\max}(\Sigma)$ grows linearly with p , the condition of bounded $\|B\|_\infty$ is equivalent to the incoherence of

eigenvectors of Σ being bounded, which is standard in the matrix completion literature (Candès and Recht, 2009) and the robust PCA literature (Chandrasekaran et al., 2011).

We now consider the moment assumption of random variables in model (1.1).

Assumption 2.3— (f_t, u_t) is iid with mean zero and bounded fourth moment. That is, there exists a constant $C > 0$ such that $\max_k E f_{kt}^4 < C$ and $\max_i E u_{it}^4 < C$.

The independence assumption can be relaxed to mixing conditions, but we do not pursue this direction in the current paper. Note that our main Theorem 3.1 is essentially deterministic. So under certain mixing condition such as that used by Fan et al. (2011), as long as we achieve a max-norm error bound (3.2) in Corollary 3.1, all conclusions in Theorem 3.1 follow immediately. Please find more details in Section 3.

We are going to establish our results based on the above general distribution family with only bounded fourth moments.

2.2 Robust estimation procedure

The basic idea we propose is to estimate the covariance matrix of the joint vector (y_t, f_t) instead of just that of y_t , although it is our target. The covariance of the concatenated $p + r$ dimensional vector $z_t^T = (y_t^T, f_t^T)$ contains all the information we need to recover the low-ranks and sparse structure. Observe that the covariance matrix $\Sigma_z := \text{Cov}(z_t)$ can be expressed as

$$\Sigma_z = \begin{pmatrix} B \Sigma_f B^T + \Sigma_u & B \Sigma_f \\ \Sigma_f B^T & \Sigma_f \end{pmatrix} = \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix}.$$

Any method which yields an estimate of Σ_z as an initial estimator or estimates of $\hat{\Sigma}_{11}$, $\hat{\Sigma}_{12}$, $\hat{\Sigma}_{21}$, $\hat{\Sigma}_{22}$ could be used to infer the unknown B , Σ_f and Σ_u . Specifically, using the estimator $\hat{\Sigma}_z$, we can readily obtain an estimator of $B \Sigma_f B^T$ through the identity

$$B \Sigma_f B^T = \Sigma_{12} \Sigma_{22}^{-1} \Sigma_{21}.$$

Subsequently, we can subtract the estimator of $B \Sigma_f B^T$ from $\hat{\Sigma}_{11}$ to obtain $\hat{\Sigma}_u$. With the sparsity structure of Σ_u assumed in Section 2.1, the well-studied thresholding (Bickel and Levina, 2008; Rothman et al., 2009; Cai and Liu, 2011) can be employed. Applying thresholding to $\hat{\Sigma}_u$, we obtain a thresholded matrix $\hat{\Sigma}_u^{\mathcal{T}}$ with guaranteed error in terms of max-norm and operator norm. The final step is to add up $\hat{\Sigma}_u^{\mathcal{T}}$ with the estimator of $B \Sigma_f B^T$ from $\hat{\Sigma}_z$ to produce the final $\hat{\Sigma}^{\mathcal{T}}$ for Σ .

Due to the fact that we only have bounded fourth moments for factors and errors, we estimate the covariance matrix Σ_z through robust methodology. For the sake of simplicity, we assume z_t has zero mean, so the covariance matrix of z_t takes the form $Ez_t z_t^T$. We shall use the M-estimator proposed in Catoni (2012) and Fan et al. (2017), where the authors proved the concentration property in the estimation of population mean of a random variable with only finite second moment. Here the variable of interest is $z_t z_t^T$, and naturally we will need bounded fourth moment of z_t .

In essence, minimizing a suitable loss function, say Huber loss, yields an estimator of the population mean with deviation of order $n^{-1/2}$. The Huber loss reads

$$l_\alpha(x) = \begin{cases} 2\alpha|x| - \alpha^2, & |x| > \alpha, \\ x^2, & |x| \leq \alpha. \end{cases} \quad (2.4)$$

Choosing $\alpha = \sqrt{(nv^2)/\log(\varepsilon^{-1})}$, $\varepsilon \in (0, 1)$ where v is an upper bound of the standard deviation of the iid random variables X_j of interest, Fan et al. (2017) showed that the minimizer $\hat{\mu} = \operatorname{argmin}_\mu \sum_{i=1}^n l_\alpha(X_i - \mu)$ satisfies

$$P\left(|\hat{\mu} - \mu| \leq 4u\sqrt{\frac{\log(\varepsilon^{-1})}{n}}\right) \geq 1 - 2\varepsilon, \quad (2.5)$$

when $n \geq 8 \log(\varepsilon^{-1})$ where $\mu = EX_j$. This finite sample result holds for any distributions with bounded second moments, including asymmetric distributions generated by Z^2 . This assumption of bounded second moments for mean estimation translates into a fourth moments assumption for our covariance estimation, because covariances are products of two random variables. When applying (2.5), we will take X_j to be the square of a random variable or products of two random variables. The diverging parameter α is chosen to reduce the biases of the M-estimator for asymmetric distributions. When applying this method to estimate Σ_z element-wisely, we expect $\hat{\Sigma}_{11}^*$, $\hat{\Sigma}_{12}^*$, $\hat{\Sigma}_{21}^*$, $\hat{\Sigma}_{22}^*$ to achieve element-wise maximal error of $Op(\sqrt{\log p/n})$, where the logarithmic term is incurred when we bound the errors uniformly. The formal result will be given in Section 3.

In an earlier work, Catoni (2012) proposed solving the equation $\sum_{i=1}^n h[\alpha^{-1}(\mu - \hat{\mu})] = 0$, where the strictly increasing $h(x)$ satisfies $-\log(1 - x + x^2/2) \leq h(x) \leq \log(1 + x + x^2/2)$. For $\varepsilon \in (0, 1)$ and $n > 2 \log(\varepsilon^{-1})$, Catoni (2012) proved that

$$P\left(|\hat{\mu} - \mu| \leq u\sqrt{\frac{2 \log(\varepsilon^{-1})}{n - 2 \log(\varepsilon^{-1})}}\right) \geq 1 - 2\varepsilon,$$

when $n \geq 4 \log(\epsilon^{-1})$ and $\alpha = \sqrt{nv^2(1 + \frac{2 \log(\epsilon^{-1})}{n - 2 \log(\epsilon^{-1})}) / \{2 \log(\epsilon^{-1})\}}$, where v is an upper bound of the standard deviation. This M -estimator can also be used for covariance estimation, though it usually has a larger bias as shown in Fan et al. (2017).

The whole procedure can be presented in the following steps:

Step 1 For each entry of the covariance matrix Σ_z , obtain a robust estimator by solving a convex minimization problem (through, for example, Newton-Rapson method):

$$(\widehat{\Sigma}_z^R)_{ij} = \operatorname{argmin}_x \sum_{t=1}^n l_\alpha(z_{it}z_{jt} - x), \quad (2.6)$$

where α is chosen as discussed above and $\widehat{\Sigma}_z = \widehat{\Sigma}_z^R = \begin{pmatrix} \widehat{\Sigma}_{11}^R & \widehat{\Sigma}_{12}^R \\ \widehat{\Sigma}_{21}^R & \widehat{\Sigma}_{22}^R \end{pmatrix}$.

Step 2 Derive an estimator of Σ_u through the algebraic manipulation

$$\widehat{\Sigma}_u = \widehat{\Sigma}_{11} - \widehat{\Sigma}_{12} \widehat{\Sigma}_{22}^{-1} \widehat{\Sigma}_{21},$$

and then apply adaptive thresholding of Cai and Liu (2011). That is,

$$(\widehat{\Sigma}_u^{\mathcal{T}})_{ij} = \begin{cases} (\widehat{\Sigma}_u)_{ij}, & i = j \\ s_{ij}((\widehat{\Sigma}_u)_{ij}) \mathbb{1}(|(\widehat{\Sigma}_u)_{ij}| \geq \mathcal{T}_{ij}), & i \neq j \end{cases}$$

where $s_{ij}(\cdot)$ is the generalized shrinkage function (Antoniadis and Fan, 2001; Rothman et al., 2009) and $\mathcal{T}_{ij} = \mathcal{T}((\widehat{\Sigma}_u)_{ii}(\widehat{\Sigma}_u)_{jj})^{1/2}$ is an entry-dependent threshold.

Step 3 Produce the final estimator for Σ :

$$\widehat{\Sigma}^{\mathcal{T}} = \widehat{\Sigma}_{12} \widehat{\Sigma}_{22}^{-1} \widehat{\Sigma}_{21} + \widehat{\Sigma}_u^{\mathcal{T}}.$$

Note in the above steps, the choice of the parameters v (in the definition of α) and τ_{ij} are not yet specified and will be discussed in Section 3.

There are $p(p+1)/2$ adaptive Huber estimators (2.6) that we need to compute in Step 1. This can be very computationally intensive when p is large, even though only univariate convex optimization problems are solved. The fact that we solve so many similar optimization problems (2.6) gives us some advantages to implement more efficiently. The key is to pick a

good initial value. A simple method is to use the sample covariance $s_{ij} = n^{-1} \sum_{t=1}^n z_{it} z_{jt}$ as the initial value for (2.6). A better alternative is to use the interpolation of computed $\{(s_{ij}, (\Sigma_z^R)_{ij})\}$ as the initial value. More precisely, to get an initial value of the robust covariance of the pair (i_0, j_0) , we find the interval $[s_{i_1, j_1}, s_{i_2, j_2}]$ to be the smallest interval that contains s_{i_0, j_0} , and then take the linear interpolation of $(\Sigma_z^R)_{i_1, j_1}$ and $(\Sigma_z^R)_{i_2, j_2}$ as the initial value for solving the adaptive Huber problem for the pair (i_0, j_0) .

Before delving into the analysis of the procedure, we first deviate to look at a technical issue. Recall that $\hat{\Sigma}_{22}$ is an estimator of Σ_f by Weyl's inequality,

$$|\lambda_i(\hat{\Sigma}_{22}) - \lambda_i(\Sigma_f)| \leq \|\hat{\Sigma}_{22} - \Sigma_f\|.$$

Since both matrices are of low dimensionality, as long as we are able to estimate every entry of Σ_f accurate enough (see Lemma 3.1 below), $\|\hat{\Sigma}_{22} - \Sigma_f\|$ vanishes with high probability as n diverges. Therefore $\hat{\Sigma}_{22}$ is invertible with high probability, and there is no major issue implementing the procedure. In cases where positive semidefinite matrix is expected, we replace the matrix by its nearest positive semidefinite version. We can do this projection for either $\hat{\Sigma}_u$ or $\hat{\Sigma}_z$. For example, for $\hat{\Sigma}_u$, we solve the following optimization problem:

$$\tilde{\Sigma}_u = \operatorname{argmin}_{\Sigma_u \geq 0} \|\hat{\Sigma}_u - \Sigma_u\|_{\infty}, \quad (2.7)$$

and simply employ $\tilde{\Sigma}_u$ as a surrogate of $\hat{\Sigma}_u$. Observe that

$$\|\tilde{\Sigma}_u - \Sigma_u\|_{\infty} \leq \|\tilde{\Sigma}_u - \hat{\Sigma}_u\|_{\infty} + \|\hat{\Sigma}_u - \Sigma_u\|_{\infty} \leq 2\|\hat{\Sigma}_u - \Sigma_u\|_{\infty}.$$

Thus, apart from a different constant, $\tilde{\Sigma}_u$ inherits all the desired properties of $\hat{\Sigma}_u$, as we will see in section 3 that those properties only rely on the maximal error bound. Hence we are able to safely replace $\hat{\Sigma}_u$ with $\tilde{\Sigma}_u$ without modifying our estimation procedure.

Moreover, (2.7) can be cast into the semidefinite programming problem below,

$$\min_{t, \Sigma_u \geq 0} \text{ s.t. } \|\hat{\Sigma}_u - \Sigma_u\|_{ij} \leq t, \quad (2.8)$$

which can be easily solved by a semidefinite programming solver, e.g. Grant et al. (2008).

3 Theoretical analysis

In this section, we will show the theoretical properties of our robust estimator under bounded fourth moment. We will also show that when the data are known to be generated from more restricted families (e.g. sub-Gaussian), commonly used estimators such as sample covariance estimator suffices as the initial estimator in Step 1.

3.1 General theoretical properties

From the above discussion on M-estimators and their concentration results, it is immediate to have the following lemma.

Lemma 3.1—Suppose that a d -dimensional random vector X is centered and has finite fourth moment, i.e. $EX = 0$, $\max_i EX_i^4 < +\infty$ for $i = 1, 2, \dots, p$. Let $\sigma_{ij} = E(X_i X_j)$ and $\hat{\sigma}_{ij}$ be Huber's estimator with parameter $\alpha = \sqrt{nv^2 / \log(p^2/\delta)}$, then there exists a universal constant C such that for any $\delta \in (0, 1)$ and $n \geq C \log(p/\delta)$, with probability $1 - \delta$,

$$\max_{ij} |\hat{\sigma}_{ij} - \sigma_{ij}| \leq C v \sqrt{\frac{\log p + \log(1/\delta)}{n}}, \quad (3.1)$$

where v is a pre-determined parameter satisfying $v^2 \geq \max_{i,j} \text{Var}(X_i X_j)$.

In practice, we do not know any of the fourth moments in advance. To pick up a good v , one possibility is Lepski's adaptation method (Lepskii, 1992) where a sequence of geometrically increasing v is tried and the estimated v is picked up as the middle of the smallest confidence interval intersecting all the larger ones. Please refer the details to Catoni (2012). Alternatively, we may simply use the empirical variance to give a rough bound of v similar to Fan et al. (2015b).

Recall that z_t is a $p + r$ dimensional vector concatenating y_t and f_t . From Assumption 2.3, there is a constant C_0 as a uniform bound for Ez_{it}^4 . This leads to the following result.

Corollary 3.1—Suppose that $\hat{\Sigma}_z$ is an estimator of covariance matrix Σ_z , whose entries are Huber's estimators with parameter $\alpha = \sqrt{nv^2 / \log((p+r)^2/\delta)}$. Then there exists a universal constant C such that for any $\delta \in (0, 1)$ and $n \geq C \log(p/\delta)$, with probability $1 - \delta$,

$$\|\hat{\Sigma}_z - \Sigma_z\|_{\infty} \leq C v \sqrt{\frac{\log p + \log(1/\delta)}{n}}, \quad (3.2)$$

where v is a pre-determined parameter satisfying $v^2 \geq C_0$.

After Step 1 of the proposed procedure, we obtain an estimator $\hat{\Sigma}_z$ that achieves optimal rate of element-wise convergence. With $\hat{\Sigma}_z$, we proceed to establish convergence rates for

both $\hat{\Sigma}_u^{\mathcal{T}}$ and $\hat{\Sigma}^{\mathcal{T}}$. The key theorem that links the estimation error under element-wise max-norm with other metrics is stated below.

Theorem 3.1—Under Assumptions 2.1–2.3, if we have estimator $\hat{\Sigma}_z^{\mathcal{T}}$ satisfying

$$\|\hat{\Sigma}_z^{\mathcal{T}} - \Sigma_z\|_{\infty} = Op(\sqrt{\log p/n}), \quad (3.3)$$

then the three-step procedure in Section 2.2 with $\mathcal{T} \asymp \sqrt{\log p/n}$ generates $\hat{\Sigma}_u^{\mathcal{T}}$ and $\hat{\Sigma}^{\mathcal{T}}$ satisfying

$$\|\hat{\Sigma}_u^{\mathcal{T}} - \Sigma_u\|_2 = \|(\hat{\Sigma}_u^{\mathcal{T}})^{-1} - \Sigma_u^{-1}\|_2 = Op\left(m_p\left(\frac{\log p}{n}\right)^{(1-q)/2}\right), \quad (3.4)$$

and furthermore

$$\|\hat{\Sigma}^{\mathcal{T}} - \Sigma\|_{\infty} = Op\left(\sqrt{\frac{\log p}{n}}\right), \quad (3.5)$$

$$\|\hat{\Sigma}^{\mathcal{T}} - \Sigma\|_{\Sigma} = Op\left(\frac{\sqrt{p \log p}}{n} + m_p\left(\frac{\log p}{n}\right)^{(1-q)/2}\right), \quad (3.6)$$

$$\|(\hat{\Sigma}^{\mathcal{T}})^{-1} - \Sigma^{-1}\| = Op\left(m_p\left(\frac{\log p}{n}\right)^{(1-q)/2}\right), \quad (3.7)$$

where $\|A\|_{\Sigma} = p^{-1/2} \|\Sigma^{-1/2} A \Sigma^{-1/2}\|_F$ is the relative Frobenius norm defined in Fan et al. (2008), if n is large enough so that $m_p(\log p/n)^{(1-q)/2}$ is bounded.

Theorem 3.1 provides a nice interface connecting max-norm guarantee with the desired convergence rates. Therefore, any robust method that attains the element-wise optimal rate as in Corollary 3.1 can be used in Step 1 instead of the current M-estimator approach.

3.2 Estimators under more restricted distributional assumptions

We analyzed theoretical properties of the robust procedure in the previous subsection under the assumption of bounded fourth moment. Theorem 3.1 shows that any estimator that achieves the optimal max-norm convergence rate could serve as an initial pilot estimator for Σ_z to be used in Step 2 and Step 3 of our procedure. Thus the procedure depends on the distributional assumption 2.3 only through Step 1 where a proper estimator $\hat{\Sigma}_z^{\mathcal{T}}$ is proposed.

Sometimes, we do have more information on the shapes of the distributions of factors and noises. For example, if the distribution of $z_t = (f_t^T, u_t^T)^T$ has a sub-Gaussian tail, the sample covariance matrix $\hat{\Sigma}_z^S = n^{-1} \sum_{t=1}^n z_t z_t^T$ attains the optimal element-wise maximal rate for estimating Σ_z .

In an earlier work, Fan et al. (2011) proposed to simply regress observations y_t on f_t in order to obtain

$$\hat{B} = Y^T F (F^T F)^{-1}, \quad (3.8)$$

where $Y = (y_1, \dots, y_n)^T$ and $F = (f_1, \dots, f_n)^T$. Then they threshold the matrix $\hat{\Sigma}_u = \hat{\Sigma} - \hat{B} \hat{\Sigma}_f \hat{B}^T$ where $\hat{\Sigma} = n^{-1} Y Y^T$ and $\hat{\Sigma}_f = n^{-1} F^T F$. This regression-based method is equivalent to applying $\hat{\Sigma}_z^S$ directly in Step 1 and also equivalent to solving a least-square minimization problem, and thus suffers from robustness issue when the data come from heavy-tailed distributions. All the convergence rates achieved in Theorem 3.1 are identical with Fan et al. (2011) where exponentially decayed tails are assumed.

As we explained, if z_t is sub-Gaussian distributed, $\hat{\Sigma}_z^S$ instead of $\hat{\Sigma}_z^R$ can be used. If f_t and u_t exhibit heavy tails, another widely used assumption is multivariate t-distribution, which is included in the elliptical distribution family. The elliptical distribution is defined as follows. Let $\mu \in \mathbb{R}^p$ and $\Sigma \in \mathbb{R}^{p \times p}$ with $\text{rank}(\Sigma) = q < p$. A p -dimensional random vector y has an elliptical distribution, denoted by $y \sim ED_p(\mu, \Sigma, \zeta)$, if it has a stochastic representation (Fang et al., 1990)

$$y \stackrel{d}{=} \mu + \zeta A U, \quad (3.9)$$

where U is a uniform random vector on the unit sphere in $\mathbb{R}^q$, $\zeta \geq 0$ is a scalar random variable independent of U , $A \in \mathbb{R}^{p \times q}$ is a deterministic matrix satisfying $A A^T = \Sigma$. To make the representation (3.9) identifiable, we require $\mathbb{E} \zeta^2 = q$ so that $\text{Cov}(y) = \Sigma$. Here we also assume continuous elliptical distributions with $\mathbb{P}(\zeta = 0) = 0$.

If f_t and u_t are uncorrelated and jointly elliptical, i.e., $z_t = (f_t^T, u_t^T)^T \sim ED_p(0, \text{diag}(\Sigma_f, \Sigma_u), \zeta)$, then a well-known good estimator for the correlation matrix R of z_t is the marginal Kendall's tau. Kendall's tau correlation coefficient is defined as

$$\hat{\mathcal{T}}_{jk} = \frac{2}{n(n-1)} \sum_{i < i'} \text{sgn}((z_{ij} - z_{i'j})(z_{ik} - z_{i'k})), \quad (3.10)$$

whose population counterpart is

$$\mathcal{T}_{jk} := \mathbb{P}((z_{1j} - z_{2j})(z_{1k} - z_{2k}) > 0) - \mathbb{P}((z_{1j} - z_{2j})(z_{1k} - z_{2k}) < 0). \quad (3.11)$$

For the elliptical family, the key identity $r_{jk} = \sin(\pi \tau_{jk}/2)$ relates Pearson correlation with Kendall's correlation (Fang et al., 1990). Using $\hat{r}_{jk} = \sin(\pi \hat{\mathcal{T}}_{jk}/2)$, Han and Liu (2014) showed that $\hat{R}$ is an accurate estimate of R , achieving $\|\hat{R} - R\|_{\infty} = Op(\sqrt{\log p/n})$. Let $\Sigma_Z = DRD$ where R is the correlation matrix and $D = \text{diag}(\sigma_1, \dots, \sigma_p)$ is a diagonal matrix consisting of standard deviations for each dimension. We construct Σ_Z^K by separately estimating D and R . As before, if the fourth moment exists, we estimate D by only considering $i = j$ in Step 1, namely by using the adaptive Huber method.

Therefore, if Z_i is elliptically distributed, Σ_Z^K can be used as the initial pilot estimator for Σ_Z in Step 1. Note that Σ_Z^K is more computationally efficient than Σ_Z^R . This difference can be reduced by using However, for general heavy-tailed distributions, there is no simple way to connect the Pearson correlation with Kendall's correlation. Thus we should favor Σ_Z^R instead. We will compare the three estimators Σ_Z^S , Σ_Z^K and Σ_Z^R thoroughly through simulations in Section 4.

3.3 Asymptotics of robust mean estimators

In this section we look further into robust mean estimators. Though the result we shall present is asymptotic and not essential for our main Theorem 3.1, it is interesting in its own right and deserves some treatment.

Perhaps the best known result of Huber's mean estimator is the asymptotic minimax theory. Huber (1964) considered the so-called ε -contamination model:

$$\mathcal{P}_{\varepsilon} = \{F \mid F(x) = (1 - \varepsilon)G(x - \theta) + \varepsilon H(x), H \in \mathcal{F}, \theta \in \mathbf{R}\},$$

where G is a known distribution, ε is fixed and $\mathcal{F}$ is the family of symmetric distributions. Let T_n be the minimizer of $\sum_{i=1}^n \rho_H(x_i - \mu)$, where $\rho_H(x) = x^2/2$ for $|x| < a$, and $\rho_H(x) = a|x| - a^2/2$ for $|x| \geq a$, where a is fixed. In the special case where G is Gaussian, Huber's result showed that with appropriate choice of a , Huber's estimator minimizes the maximal asymptotic variance among all translation invariant estimators, the maximum being taken over $\mathcal{P}_{\varepsilon}$.

One problem with ε -contamination model is that it makes sense only when we assume symmetry of H , if θ is the quantity we are interested in. In contrast, Catoni (2012) and Fan et al. (2017) studied a different family, in which distributions have finite second moment. Bickel (1976) called them 'local' and 'global' models respectively, and offered a detailed discussion.

This paper, along with the preceding two papers (Catoni, 2012; Fan et al., 2017), studies robustness in the sense of the second model. The technical novelty primarily lies in the nice concentration property, which is fundamental to high dimensional statistics. This requires the parameter α of ρ_H to grow with n , versus being kept fixed, such that the condition in Corollary 3.1 is satisfied. It turns out that, in addition to the concentration property, we can establish results regarding its asymptotic behaviors in an exact manner.

Let $\rho_n(x) = x^2/2$ for $|x| < \alpha_n$ and $\rho_n(x) = \alpha_n|x| - \alpha_n^2/2$ for $|x| \geq \alpha_n$; its derivative $\psi_n = \rho'_n$. Let us write $\lambda_n(t) = E\psi_n(X - t)$. Denote t_n as a solution of $\lambda_n(t) = 0$, which is unique when n is sufficiently large, and T_n a solution of $\sum_{i=1}^n \psi_n(x_i - t) = 0$. We have the following theorem.

Theorem 3.2—Suppose that $x_1, \dots, x_n$ is drawn from some distribution F with mean μ and finite variance σ^2 . Suppose $\{\alpha_n\}$ is any sequence with $\lim_{n \rightarrow \infty} \alpha_n = \infty$. Then, as $n \rightarrow \infty$,

$$\sqrt{n}(T_n - t_n) \xrightarrow{d} N(0, \sigma^2),$$

and moreover

$$\frac{t_n - \mu}{E\psi_n(X - \mu)} \rightarrow 1.$$

Theorem 3.2 gives a decomposition of error $T_n - \mu$ into two components: variance and bias. The rate of bias $E\psi_n(X - \mu)$ depends on the distribution F and $\{\alpha_n\}$. When the distribution is either symmetric or $\liminf_n \alpha_n/\sqrt{n} > 0$, the second component $t_n - \mu$ is $o(1/\sqrt{n})$, a negligible quantity compared with the asymptotic variance. While Huber's approach needs the symmetric restriction, there is no need for our estimator. This theorem also lends credibility to the bias-variance tradeoff we observed in the simulation (see Section 4.1).

It is worth comparing the above Huber loss minimization with another candidate for robust mean estimation called “median-of-means” given by Hsu and Sabato (2014). The method, as its name suggests, first divides samples into k subgroups and calculates means for each subgroup, then take the median of those means as the final estimator. The first step basically symmetrizes the distribution by the central limit theorem and the second step is to robustify the procedure. According to Hsu and Sabato (2014), if we choose $k = 4.5 \log(p/\delta)$ and element-wisely estimate Σ_z similar to (2.5), with probability $1 - \delta$, we have

$$\|\widehat{\Sigma}_z - \Sigma_z\|_\infty \leq 3\sqrt{3}v\sqrt{\frac{\log p + \log(1/\delta)}{n}}.$$

Although “median-of-means” has the desired concentration property, unlike our estimator here, its asymptotic behavior differs from the empirical mean estimator, and as a consequence, it is not asymptotically efficient when the distribution F is Gaussian. Therefore, regarding efficiency, we prefer our proposed procedure in Section 2.2.

4 Simulations

We now present simulation results to demonstrate improvement of the proposed robust method over the least-square based method (Fan et al., 2008, 2011) and Kendall's tau based method (Han and Liu, 2014; Fan et al., 2015b) when factors and errors are elliptically distributed and generally heavy-tailed.

However, one must be cautious of the choice of the tuning parameter α , since it plays an important role in the quality of the robust estimates. Out of this concern, we shall discuss the intricacy of choosing parameter α before presenting the performance of robust estimates of covariance matrices.

4.1 Robust estimates of variances and covariances

For random variables $X_1, \dots, X_p$ with zero mean that may potentially exhibit heavy-tailed behavior, the sample mean of $v_{ij} = E(X_i X_j)$ is not good enough for our estimation purpose. Though being unbiased, in the high dimensional setting, there is no guarantee that multiple sample means stay close to the true values simultaneously.

As shown in theoretical analysis, this problem is alleviated for robust estimators constructed through M-estimators, whose influence functions grow slowly at extreme values. The desired concentration property in (3.2) depends on the choice of parameter α , which decides the range outside which large values cease to become more influential. However, in practice, we have to make a good guess of $\text{Var}(X_i X_j)$ as the theory suggests; even so, we may be too conservative in the choice of α .

To show this, we plot in Figure 1 the histograms of our estimates of $v = \text{Var}(X_i)$ in 1000 runs, where X_i is generated from a t-distribution with degree of freedom $\nu = 4.2$. The first three histograms show the estimates constructed from Huber's M-estimator, with parameter

$$\alpha = \beta \sqrt{\frac{n \text{Var}(X_i^2)}{2}}, \quad (4.1)$$

where β is 0.2, 1, 5 respectively, and the last histogram is the usual sample estimate (or $\beta = \infty$). The quality of estimates ranges from large biases to large variances. We also plot in Figure 2 the histograms of estimates of $v = \text{Cov}(X_i, X_j)$, where (X_i, X_j) , $i \neq j$ is generated from a multivariate t-distribution with $\nu = 4.2$ and an identity scale matrix. The only difference is that in (4.1), the variance of X_i^2 is replaced by the variance of $X_i X_j$.

From Figure 1, we observe a bias-variance tradeoff phenomenon as α varies. This is also consistent with the theory in Section 3.3. When α is small, the robust method underestimate the variance, yielding a large bias due to the asymmetry of the distribution of X_i^2 . As α increases, a larger variance is traded for a smaller bias, until $\alpha = \infty$, in which case the robust estimator simply becomes the sample mean.

For the covariance estimation, Figure 2 exhibits a different phenomenon. Since the distribution of $X_i X_j$ is symmetric for $i = j$, there is no bias incurred when α is small. Since the variance is smaller when α is smaller, we have a net gain in terms of the quality of estimates. In the extreme case where α is zero, we are actually estimating the median. Fortunately, under distributional symmetry, the mean and the median are the same.

The simple simulations help us to understand how to choose α in practice: if the distribution is close to a symmetric one, one can choose α aggressively, i.e. making α smaller; otherwise, a conservative α is preferred.

4.2 Covariance matrix estimation

We implemented the robust estimation procedure with three initial pilot estimators $\hat{\Sigma}_z^S$, $\hat{\Sigma}_z^K$ and $\hat{\Sigma}_z^R$. We simulated n samples of $z_t = (f_t^T, u_t^T)^T$ from a multivariate t-distribution with covariance matrix $\text{diag}\{I_p, 5I_p\}$ and various degrees of freedom. Each row of B is independently sampled from a standard normal distribution. The population covariance matrix of $y_t = Bf_t + u_t$ is $\Sigma = BB^T + 5I_p$. For p running from 200 to 900 and $n = p/2$, we calculated errors of the robust procedure in different norms. As suggested by the experiments in the previous section, we chose a larger parameter α to estimate the diagonal elements of Σ_z and a smaller one to estimate its off-diagonal elements. We used the thresholding parameter $\mathcal{T} = 2\sqrt{\log p/n}$.

The estimation errors are gauged in the following norms: $\|\hat{\Sigma}_u^{\mathcal{T}} - \Sigma_u\|$, $\|(\hat{\Sigma}^{\mathcal{T}})^{-1} - \Sigma^{-1}\|$ and $\|\hat{\Sigma}^{\mathcal{T}} - \Sigma\|_{\Sigma}$ as shown in Theorem 3.1. We considered two different settings: (1) z_t is generated from multivariate t distribution with very heavy ($\nu = 3$), medium heavy ($\nu = 5$), and light ($\nu = \infty$ or Gaussian) tail; (2) z_t is element-wise iid one-dimensional t distribution with degree of freedom $\nu = 3, 5$ and ∞ . They are separately plotted in Figures 3 and 4. The estimation errors of applying sample covariance matrix $\hat{\Sigma}_z^S$ are used as the baseline for comparison. For example, if $\|\hat{\Sigma}^{\mathcal{T}} - \Sigma\|_{\Sigma}$ is used to measure performance, the blue curve represents ratio $\|(\hat{\Sigma}^{\mathcal{T}})^R - \Sigma\|_{\Sigma} / \|(\hat{\Sigma}^{\mathcal{T}})^S - \Sigma\|_{\Sigma}$ while the black curve represents ratio $\|(\hat{\Sigma}^{\mathcal{T}})^K - \Sigma\|_{\Sigma} / \|(\hat{\Sigma}^{\mathcal{T}})^S - \Sigma\|_{\Sigma}$ where $(\hat{\Sigma}^{\mathcal{T}})^R$, $(\hat{\Sigma}^{\mathcal{T}})^K$, $(\hat{\Sigma}^{\mathcal{T}})^S$ are respectively estimators given by the robust procedure with initial pilot estimators $\hat{\Sigma}_z^R$, $\hat{\Sigma}_z^K$, $\hat{\Sigma}_z^S$ for Σ_z . Therefore if the ratio curve moves below 1, the method is better than naive sample estimator given in Fan et al. (2011) and vice versa. The more it gets below 1, the more robust the procedure is against heavy-tailed randomness.

The first setting (Figure 3) represents a heavy-tailed elliptical distribution, where we expect the two robust methods work better than the sample covariance based method, especially in the case of extremely heavy tails (solid lines for $\nu = 3$). As expected, both black curves and blue curves under the three measures behave visibly better (smaller than 1). On the other hand, if data are indeed Gaussian (dotted line for $\nu = \infty$), the method with sample

covariance performs better under most measures (greater than 1). Nevertheless, our robust method still performs comparably with the sample covariance method, as the median error ratio stays around 1 whereas Kendall's tau method can be much worse than the sample covariance method. A plausible explanation is that the variance reduced compensates for the bias incurred in our procedure. In addition, the IQR plots tell us the proposed robust method is indeed more stable than Kendall's tau.

The second setting (Figure 4) provides an example of non-elliptical distributed heavy-tailed data. We can see that the performance of the robust method dominates the other two methods, which verifies the approach in this paper especially when data comes from a general heavy-tailed distribution. While our method is able to deal with more general distributions, Kendall's tau method does not apply to distributions outside the elliptical family, which excludes the element-wise iid t distribution in this setting. This explains why under various measures, our robust method is better than Kendall's tau method by a clear margin. Note that even in the first setting where the data are indeed elliptical, with proper tuning, the proposed robust methods can still outperform Kendall's tau.

5 Real data analysis

In this section, we look into financial historical data during 2005–2013, and assess to what extent our factor model characterizes the data.

The dataset we used in our analysis consists of daily returns of 393 stocks, all of which are large market capitalization constituents of S&P 500 index, collected without missing values from 2005 to 2013. This dataset has also been used in Fan et al. (2016), where they investigated how covariates (e.g. size, volume) could be utilized to help estimate factors and factor loadings, whereas the focus of the current paper is to develop robust methods in the presence of heavy tailed data.

In addition, we collected factors data for the same period, where the factors are calculated according to Fama-French three-factor model (Fama and French, 1993). After centering, the panel matrix we will use for analysis, is a 393 by 2265 matrix Y , in addition to a factor matrix F of size 2265 by 3. Here 2265 is the number of daily returns and 393 is the number of stocks.

5.1 Tail-heaviness

First, we look at how the daily returns are distributed. Especially, we are interested in the behaviors of their tails. In Figure 5, we made Q-Q plots that compare the distribution of all y_{it} with either Gaussian distribution or t -distributions with varying degree of freedom, ranging from $df = 2$ to $df = 6$. We also fit a line in each plot, showing how much the return data deviate from the base distribution. It is clear that the data has a tail heavier than that of a Gaussian distribution, and that t -distribution with $df = 4$ is almost in alignment with the return data. Similarly, we made the Q-Q plots for the factors in Figure 6. The plots also show that t -distribution is better in terms of fitting the data; however, the tails are even heavier, and t -distribution t_2 seems to best fit the data.

5.2 Spiked covariance structure

We now consider how the covariance matrix of returns looks like, since a spiked covariance structure would justify the pervasiveness assumption. To find the spectral structure, we calculated eigenvalues of the sample covariance matrix YY^T/n , and made a histogram based on logarithmic scale (see the left panel in Figure 7). In the histogram, the counts in the rightmost four bins are 5, 1, 0 and 1, representing only a few large eigenvalues, which is a strong signal of a spiked structure. We also plotted the proportion of residue eigenvalues $\sum_{i=K+1}^P \lambda_i / \sum_{i=1}^P \lambda_i$, against K in the right panel of Figure 7. The top 3 eigenvalues account for a major part of the variances, which lends weight to the pervasive assumption.

The spiked covariance structure has been studied in Paul (2007), Johnstone and Lu (2009) and many other papers, but under their regime, the top eigenvalues or “spiked” eigenvalues do not grow with the dimension. In this paper, the spiked eigenvalues have stronger signals, and thus are easier to be separated from the rest of eigenvalues. In this respect, the connotation of “spiked covariance structure” is closer to that in Wang and Fan (2017). As empirical evidence, this phenomenon also buttresses the motivation of study in Wang and Fan (2017).

5.3 Portfolio risk estimation

We consider portfolio risk estimation. To be specific, for a portfolio with weight vector $w \in \mathbf{R}^P$ on all the market assets, its risk is measured by quantity $w^T \Sigma w$ where Σ is the true covariance of excess returns of all the assets. Note that Σ is time varying. Here we consider a class of weights with gross exposure $c = 1$, that is $\sum_i w_i = 1$ and $\sum_j |w_j| = c$. We consider four scenarios $c = 1, 1.4, 1.8, 2.2$. Note that $c = 1$ represents the case of no short selling and $(c - 1)/2$ represents the level of exposure to short selling.

To assess how well our robust estimator performs compared with sample covariance, we calculated the covariance estimators $\widehat{\Sigma}_t^R$ and $\widehat{\Sigma}_t^S$, using the daily data of preceding 12 months, where $\widehat{\Sigma}_t^R$ is our robust covariance estimator and $\widehat{\Sigma}_t^S$ is the sample covariance, for every trading day from 2006 to 2013. We indexed those dates by t where t runs from 1 to 2013 (from 2006–01–01 to 2013–12–31, it happens to contain 2013 trading days, so here 2013 is the total number of trading days instead of a year indicator). Let γ_{t+1} be the excess return of the following trading day after t . For a weight vector w , the error we used to gauge the two approaches is

$$R^R(w) = \frac{1}{2013} \sum_{t=1}^{2013} \|w^T \widehat{\Sigma}_t^R w - (w^T \gamma_{t+1})^2\|, \quad R^S(w) = \frac{1}{2013} \sum_{t=1}^{2013} \|w^T \widehat{\Sigma}_t^S w - (w^T \gamma_{t+1})^2\|.$$

Note the bias-variance decomposition

$$E \|w^T \widehat{\Sigma}_t w - (w^T \gamma_{t+1})^2\|^2 = E \|(w^T \gamma_{t+1})^2 - w^T \Sigma_t w\|^2 + E \|w^T \widehat{\Sigma}_t w - w^T \Sigma_t w\|^2,$$

where $\sum_t = E\gamma_{t+1}\gamma_{t+1}^T$. The first term measures the size of the stochastic error that cannot be reduced while the second term is the estimation error for the risk of portfolio w .

To generate multiple random weights w with gross exposure c , we adopted the strategy used in Fan et al. (2015a), which aims to generate a uniform distribution on the simplex $\{w: \sum_i w_i = 1, \sum_i |w_i| = c\}$: (1) for each index $i = 1, \dots, p$ let $\eta_i = 1$ (long) with probability $(c+1)/2c$ and $\eta_i = -1$ (short) with probability $(c-1)/2c$; (2) generate iid ξ_i by exponential distribution; (3) for $\eta_i = 1$, let $w_i = \frac{c+1}{2} \cdot \xi_i / \sum_{\eta_i=1} \xi_i$ and for $\eta_i = -1$, let $w_i = -\frac{c-1}{2} \cdot \xi_i / \sum_{\eta_i=-1} \xi_i$. We made a set of scatter plots in Figure 8, in which the x-axis represents $R^R(w)$ and the y-axis $R^S(w)$. In addition, we highlighted in the first plot the point with uniform weights (i.e. $w_i = 1/p$), which serves as a benchmark for comparison. The dashed line shows where the two approaches have the same performance. Clearly, for all w the robust approach has smaller risk errors, and therefore has better empirical performance in estimating portfolio risks.

Acknowledgments

The research was partially supported by NSF grants DMS-1206464 and DMS-1406266 and NIH grants R01-GM072611-11 and NIH R01GM100474-04.

A Appendix

Proof of Theorem 3.1

Since we have robust estimator $\widehat{\Sigma}_z$ such that $\|\widehat{\Sigma}_z - \Sigma_z\|_\infty = O_p(\sqrt{\log p/n})$, we clearly know

$\widehat{\Sigma}_{11}, \Sigma_{12}, \widehat{\Sigma}_{21}, \widehat{\Sigma}_{22}$, achieve the same rate. Using this, let us first prove

$\|\widehat{\Sigma}_u - \Sigma_u\|_\infty = O_p(\sqrt{\log p/n})$. Obviously,

$$\|\widehat{\Sigma}_{12}\widehat{\Sigma}_{22}^{-1}\widehat{\Sigma}_{21}^T - B\Sigma_f B^T\|_\infty = \|\widehat{\Sigma}_{12}\widehat{\Sigma}_{22}^{-1}\widehat{\Sigma}_{21}^T - \Sigma_{12}\Sigma_{22}^{-1}\Sigma_{21}^T\|_\infty = O_p(\sqrt{\log p/n}), \quad (\text{A.1})$$

because the multiplication is along the fixed dimension r and each element is estimated with the rate of convergence $O_p(\sqrt{\log p/n})$. Also $\|\widehat{\Sigma}_{11} - \Sigma\|_\infty = O_p(\sqrt{\log p/n})$, therefore

$\widehat{\Sigma}_u = \widehat{\Sigma}_{11} - \widehat{\Sigma}_{12}\widehat{\Sigma}_{22}^{-1}\widehat{\Sigma}_{21}^T$ is good enough to estimate $\Sigma_u = \Sigma - B\Sigma_f B^T$ with error $O_p(\sqrt{\log p/n})$ in max-norm.

Once the max error of sparse matrix Σ_u is controlled, it is not hard to show the adaptive procedure in Step 2 gives $\widehat{\Sigma}_u^{\mathcal{T}}$ such that the spectral error $\|\widehat{\Sigma}_u^{\mathcal{T}} - \Sigma_u\| = O_p(m_p w_n^{1-q})$ (Fan et al., 2011; Cai and Liu, 2011; Rothman et al., 2009) where we define $w_n = \sqrt{\log p/n}$.

Furthermore, $\|(\widehat{\Sigma}_u^{\mathcal{T}})^{-1} - \Sigma_u^{-1}\| \leq \|(\widehat{\Sigma}_u^{\mathcal{T}})^{-1}\| \|\widehat{\Sigma}_u^{\mathcal{T}} - \Sigma_u\| \|\Sigma_u^{-1}\|$. So $\|(\widehat{\Sigma}_u^{\mathcal{T}})^{-1} - \Sigma_u^{-1}\|$ is also $O_p(m_p w_n^{1-q})$ due to the lower boundedness of $\|\Sigma_u\|$. So (3.4) is valid.

Proving (3.5) is trivial. $\|\widehat{\Sigma}_u^{\mathcal{T}} - \Sigma_u\|_{\infty} \leq \|\widehat{\Sigma}_u^{\mathcal{T}} - \widehat{\Sigma}_u\|_{\infty} + \|\widehat{\Sigma}_u - \Sigma_u\|_{\infty} = O_P(\mathcal{T} + w_n) = O_P(w_n)$ when τ is chosen as the same order as w_n and thus

$$\|\widehat{\Sigma}^{\mathcal{T}} - \Sigma\|_{\infty} \leq \|\widehat{\Sigma}_{12}\widehat{\Sigma}_{22}^{-1}\widehat{\Sigma}_{21}^T - B\Sigma_f B^T\|_{\infty} + \|\widehat{\Sigma}_u^{\mathcal{T}} - \Sigma_u\|_{\infty} = O_P(w_n).$$

Next let us take a look at the relative Frobenius convergence (3.6) for $\|\widehat{\Sigma}^{\mathcal{T}} - \Sigma\|_{\Sigma}$.

$$\begin{aligned} \|\widehat{\Sigma}^{\mathcal{T}} - \Sigma\|_{\Sigma} &\leq \|\widehat{\Sigma}_{12}\widehat{\Sigma}_{22}^{-1}\widehat{\Sigma}_{21}^T - \Sigma_{12}\Sigma_{22}^{-1}\Sigma_{21}^T\|_{\Sigma} + \|\widehat{\Sigma}_u^{\mathcal{T}} - \Sigma_u\|_{\Sigma} \\ &\leq \|(\widehat{\Sigma}_{12} - \Sigma_{12})\widehat{\Sigma}_{22}^{-1}(\widehat{\Sigma}_{21} - \Sigma_{21})^T\|_{\Sigma} + 2\|(\widehat{\Sigma}_{12} - \Sigma_{12})\widehat{\Sigma}_{22}^{-1}\Sigma_{21}^T\|_{\Sigma} \\ &\quad + \|\Sigma_{12}(\widehat{\Sigma}_{22}^{-1} - \Sigma_{22}^{-1})\Sigma_{21}^T\|_{\Sigma} + \|\widehat{\Sigma}_u^{\mathcal{T}} - \Sigma_u\|_{\Sigma} =: \Delta_1 + 2\Delta_2 + \Delta_3 + \Delta_4. \end{aligned} \quad (\text{A.2})$$

We bound the four terms one by one. The last term is the most straightforward,

$$\Delta_4 \leq p^{-1/2}\|\Sigma_u^{\mathcal{T}} - \Sigma_u\|_F \|\Sigma^{-1}\| = O_P(\|\Sigma_u^{\mathcal{T}} - \Sigma_u\|) = O_P(m_p w_n^{1-q}).$$

Bound for Δ_1 uses the fact that $\|\widehat{\Sigma}_{22}^{-1}\|$ and $\|\Sigma^{-1}\|$ are $O_P(1)$ and

$\|\widehat{\Sigma}_{12} - \Sigma_{12}\|_F = O_P(\sqrt{p \log p/n})$. So

$$\Delta_1 \leq p^{-1/2}\|\widehat{\Sigma}_{12} - \Sigma_{12}\|_F^2 \|\widehat{\Sigma}_{22}^{-1}\| \|\Sigma^{-1}\| = O_P(\frac{\sqrt{p \log p}}{n});$$

Bound for Δ_3 needs additional conclusion that

$\|\Sigma_{21}^T \Sigma^{-1} \Sigma_{12}\| \leq \|B^T \Sigma^{-1} B\| \|\Sigma_{22}\|^2 \leq 2\|\Sigma_{22}\| = O(1)$, where $B = \Sigma_{12}\Sigma_{22}^{-1}$ and the last inequality is shown in Fan et al. (2008). So

$$\begin{aligned} \Delta_3 &= p^{-1/2} \text{tr}^{1/2} \left((\widehat{\Sigma}_{22}^{-1} - \Sigma_{22}^{-1}) \Sigma_{21}^T \Sigma^{-1} \Sigma_{12} (\widehat{\Sigma}_{22}^{-1} - \Sigma_{22}^{-1}) \Sigma_{21}^T \Sigma^{-1} \Sigma_{12} \right) \\ &\leq p^{-1/2} \|(\widehat{\Sigma}_{22}^{-1} - \Sigma_{22}^{-1}) \Sigma_{21}^T \Sigma^{-1} \Sigma_{12}\|_F \leq p^{-1/2} \|\widehat{\Sigma}_{22}^{-1} - \Sigma_{22}^{-1}\|_F \|\Sigma_{21}^T \Sigma^{-1} \Sigma_{12}\| = O_P(\sqrt{\log p/(np)}). \end{aligned}$$

Lastly, by similar trick, we have

$$\begin{aligned} \Delta_2 &= p^{-1/2} \text{tr}^{1/2} \left((\widehat{\Sigma}_{12} - \Sigma_{12}) \widehat{\Sigma}_{22}^{-1} \Sigma_{21}^T \Sigma^{-1} \Sigma_{21} \widehat{\Sigma}_{22}^{-1} (\widehat{\Sigma}_{12} - \Sigma_{12}) \Sigma^{-1} \right) \\ &\leq p^{-1/2} \|\widehat{\Sigma}_{12} - \Sigma_{12}\|_F \|\widehat{\Sigma}_{22}^{-1}\| \|\Sigma^{-1}\|^{1/2} \|\Sigma_{21}^T \Sigma^{-1} \Sigma_{12}\|^{1/2} = O_P(\sqrt{\log p/n}). \end{aligned}$$

Combining results above, by (A.2), we conclude that

$$\|\widehat{\Sigma}^{\mathcal{T}} - \Sigma\|_{\Sigma} = O_P(\sqrt{p \log p/n} + m_p (\log p/n)^{(1-q)/2}).$$

Finally we show the rate of convergence for $\|(\widehat{\Sigma}^{\mathcal{T}})^{-1} - \Sigma^{-1}\|$. By Woodbury formula,

$$\Sigma^{-1} = \Sigma_u^{-1} - \Sigma_u^{-1} \Sigma_{12} [\Sigma_{22} + \Sigma_{12}^T \Sigma_u^{-1} \Sigma_{21}]^{-1} \Sigma_{21}^T \Sigma_u^{-1}.$$

Thus, let $A = \Sigma_{22} + \Sigma_{12}^T \Sigma_u^{-1} \Sigma_{21}$, $\hat{A} = \widehat{\Sigma}_{22} + \widehat{\Sigma}_{12}^T (\widehat{\Sigma}_u^{\mathcal{T}})^{-1} \widehat{\Sigma}_{21}$ and $D = \Sigma_u^{-1} \Sigma_{12}$,

$\hat{D} = (\widehat{\Sigma}_u^{\mathcal{T}})^{-1} \widehat{\Sigma}_{12}$, we have the following bound similar to (A.2):

$$\begin{aligned} \|(\widehat{\Sigma}^{\mathcal{T}})^{-1} - \Sigma^{-1}\| &\leq \|\widehat{D}\hat{A}^{-1}\widehat{D}^T - DA^{-1}D^T\| + \|(\widehat{\Sigma}_u^{\mathcal{T}})^{-1} - \Sigma_u^{-1}\| \\ &\leq \|(\widehat{D} - D)\hat{A}^{-1}(\widehat{D} - D)^T\| + 2\|(\widehat{D} - D)\hat{A}^{-1}D^T\| + \|D(\hat{A}^{-1} - A^{-1})D^T\| \\ &\quad + \|(\widehat{\Sigma}_u^{\mathcal{T}})^{-1} - \Sigma_u^{-1}\| =: \widetilde{\Delta}_1 + 2\widetilde{\Delta}_2 + \widetilde{\Delta}_3 + \widetilde{\Delta}_4. \end{aligned} \quad (\text{A.3})$$

From (3.4), $\widetilde{\Delta}_4 = O_p(m_p \omega_n^{1-q})$. For the remaining terms, we need to find the rates for

$\|\widehat{D} - D\|$, $\|\hat{A}^{-1}\|$, $\|D\|$ and $\|\hat{A}^{-1} - A^{-1}\|$ separately. Note that

$\|\Sigma_{12}\| = \|B\Sigma_{22}\| \leq \|B\|\|\Sigma_{22}\| = O_p(\sqrt{p})$ by Assumption 2.2 (ii). So $\|D\| = O_p(\sqrt{p})$ and

$$\|\widehat{D} - D\| \leq \|(\widehat{\Sigma}_u^{\mathcal{T}})^{-1}\| \|\widehat{\Sigma}_{12} - \Sigma_{12}\| + \|\Sigma_{12}\| \|(\widehat{\Sigma}_u^{\mathcal{T}})^{-1} - \Sigma_u^{-1}\| = O_p(\sqrt{pm_p} \omega_n^{1-q}).$$

In addition, it is not hard to show $\|\hat{A} - A\| = O_p(pm_p \omega_n^{1-q})$. In addition, we claim $\|A^{-1}\| = O_p(p^{-1})$ since $\lambda_{\min}(A) \geq \lambda_{\min}(\Sigma_{12}^T \Sigma_u^{-1} \Sigma_{21}) \geq \lambda_{\min}(\Sigma_u^{-1}) \lambda_{\min}(\Sigma_f) \lambda_r(B \Sigma_f B^T)$ and by Weyl's inequality, $\lambda_r(B \Sigma_f B^T) \geq \lambda_r(\Sigma) - \|\Sigma\|$ c_p by Assumption 2.2 (i). Therefore, $\|\hat{A}^{-1} - A^{-1}\| \leq \|A^{-1}\| \|\hat{A} - A\|$ implies $\|\hat{A}^{-1} - A^{-1}\| = O_p(p^{-1} m_p \omega_n^{1-q})$, and furthermore $\|\hat{A}^{-1}\| = O_p(p^{-1})$. Finally we incorporate the above rates together and conclude

$$\widetilde{\Delta}_1 = O_p(p^{-1} \|\widehat{D} - D\|^2) = O_p(m_p^2 \omega_n^{2(1-q)}),$$

$$\widetilde{\Delta}_2 = O_p(p^{-1/2} \|\widehat{D} - D\|) = O_p(m_p \omega_n^{1-q}),$$

$$\widetilde{\Delta}_3 = O_p(p \|\hat{A}^{-1} - A^{-1}\|) = O_p(m_p \omega_n^{1-q}).$$

So combining rates for $\widetilde{\Delta}_i$, $i = 1, 2, 3, 4$, we show (3.7) is true. The proof is now complete. $\square$

Proof of Theorem 3.2

Without loss of generality we can assume $\mu = 0$. By dominated converge theorem we know that for all t , $\lim_n \lambda_n(t) = -t$, that $\lambda_n(t)$ is differentiable, that $\lambda'_n(t) = -E\psi'_n(X - t)$, and that $\lim_n \lambda'_n(t) = -1$. With Taylor's expansion, we have

$$\lambda_n(t) = \lambda_n(0) + \lambda'_n(0)t + \Delta_n(t), \quad (\text{A.4})$$

where $|\Delta_n(t)| \leq |t| \sup\{|\lambda'_n(s) - \lambda'_n(0)| : 0 \leq s \leq t\}$. Observe that

$$|\lambda'_n(s) - \lambda'_n(0)| = |P(|X - s| \leq \alpha_n) - P(|X| \leq \alpha_n)| \leq P(|X - s| > \alpha_n) + P(|X| > \alpha_n).$$

By Markov's inequality,

$$\sup\{|\lambda'_n(s) - \lambda'_n(0)| : 0 \leq s \leq t\} \leq \frac{1}{\alpha_n}(2E|X| + |t|).$$

For any $\varepsilon \in (0, 1)$, there exists $N > 0$, such that for all $n > N$,

$$|\lambda_n(0)| \leq 2, \quad \frac{1 + \varepsilon/2}{1 + \varepsilon} \leq -\lambda'_n(0) \leq \frac{1 - \varepsilon/2}{1 - \varepsilon}, \quad \frac{1}{\alpha_n}(2E|X| + 4) \leq \frac{\varepsilon}{4(1 + \varepsilon)}.$$

Plugging $t = (1 + \varepsilon)\lambda_n(0)$ into (A.4),

$$\lambda_n((1 + \varepsilon)\lambda_n(0)) = \lambda_n(0) + \lambda'_n(0)(1 + \varepsilon)\lambda_n(0) + \Delta_n(t),$$

where $|\Delta_n(t)| \leq (1 + \varepsilon)|\lambda_n(0)| \frac{\varepsilon}{4(1 + \varepsilon)} = \varepsilon|\lambda_n(0)|/4$. Equivalently,

$$\lambda_n((1 + \varepsilon)\lambda_n(0)) = \lambda_n(0)(1 + \lambda'_n(0)(1 + \varepsilon) + \beta'_n),$$

where $|\beta'_n| \leq \varepsilon/4$. Similarly,

$$\lambda_n((1 - \varepsilon)\lambda_n(0)) = \lambda_n(0)(1 + \lambda'_n(0)(1 - \varepsilon) + \beta'_n),$$

where $|\beta'_n| \leq \varepsilon/4$. Also we have $1 + \lambda'_n(0)(1 + \varepsilon) + \beta'_n < 0$ and $1 + \lambda'_n(0)(1 - \varepsilon) + \beta'_n > 0$.

Multiplying both sides of the equations, we deduce that

$$\lambda_n((1 + \varepsilon)\lambda_n(0)) \cdot \lambda_n((1 - \varepsilon)\lambda_n(0)) \leq 0.$$

If $\lambda_n(0) = 0$, equation $\lambda_n(t) = 0$ has one zero $t = 0$; and in fact it is the unique one for sufficiently large n , since $\lambda_n(t)$ is nonincreasing and $\lambda'_n(0) \neq 0$ for n large enough. If $\lambda_n(0) \neq 0$,

0, at least one zero lies in the interval with endpoints $(1 + \varepsilon)\lambda_n(0)$ and $(1 - \varepsilon)\lambda_n(0)$. Since $\lambda_n(0) \rightarrow 0$, for any zero t'_n in this interval we have $t'_n \rightarrow 0$, which implies $\lambda'_n(t'_n) \rightarrow -1$. It follows that such zero is unique for sufficiently large n . This leads to $t_n/\lambda_n(0) \rightarrow 1$, thus proving the second claim in the theorem.

The proof of the first claim is similar in spirit to that of Huber (1964). Let us denote

$$T_n^- = \sup\{t: \sum_{i=1}^n \psi_n(x_i - t) > 0\},$$

$$T_n^+ = \inf\{t: \sum_{i=1}^n \psi_n(x_i - t) < 0\}.$$

By monotonicity, $T_n \in [T_n^-, T_n^+]$. Since

$$P(T_n^- < t) = P(\sum_{i=1}^n \psi_n(x_i - t) \leq 0),$$

it follows that for any fixed $z \in \mathbf{R}$,

$$\begin{aligned} P(\sqrt{n}(T_n^- - t_n) < z) &= P(T_n^- < t_n + z/\sqrt{n}) = P(\sum_{i=1}^n \psi_n(x_i - u_n) \leq 0) \\ &= P\left(\frac{1}{\sqrt{n}} \sum_{i=1}^n \frac{\psi_n(x_i - u_n) - \lambda_n(u_n)}{\sigma_n(u_n)} \leq -\frac{\sqrt{n}\lambda_n(u_n)}{\sigma_n(u_n)}\right), \end{aligned}$$

where we denote $u_n = t_n + z/\sqrt{n}$ and $\sigma_n(u) = E\psi_n(X - u)^2 - \lambda_n(u)^2$.

By dominate convergence theorem, $\lambda'_n(t_n) \rightarrow -1$ and $\sigma_n(u_n)^2 \rightarrow \sigma^2$. By Taylor expansion of $\lambda_n(u_n)$ at t_n ,

$$\lambda_n(u_n) = \lambda_n(t_n) + z/\sqrt{n}\lambda'_n(t_n) + \Delta_z^n,$$

where $|\Delta_z^n| \leq n^{-1/2}|z|\sup\{|\lambda'_n(t_n + s) - \lambda'_n(t_n)| : 0 \leq s \leq z/\sqrt{n}\}$. A similar argument shows that

$$\sup\{|\lambda'_n(t_n + s) - \lambda'_n(t_n)| : 0 \leq s \leq z/\sqrt{n}\} \leq \frac{1}{\alpha_n}(2E(X) + 2|t_n| + |z|/\sqrt{n}) = o(1).$$

This leads to $\lambda_n(u_n) = z/\sqrt{n}(\lambda'_n(t_n) + o(1)) = z/\sqrt{n}(-1 + o(1))$, and thus $\sqrt{n}\lambda_n(u_n) \rightarrow -z$.

Let us write

$$\xi_i = \frac{\psi_n(x_i - u_n) - \lambda_n(u_n)}{\sigma_n(u_n)}$$

for the centered variance ξ_i with unit variance. If we can show

$$\frac{1}{\sqrt{n}} \sum_{i=1}^n \xi_i \xrightarrow{d} N(0, 1), \quad (\text{A.5})$$

then by continuity of Φ , standard normal distribution function, we have

$$P\left(\frac{1}{\sqrt{n}} \sum_{i=1}^n \xi_i \leq -\frac{\sqrt{n}\lambda_n(u_n)}{\sigma_n(u_n)}\right) \rightarrow \Phi\left(\frac{z}{\sigma}\right),$$

which gives $P(\sqrt{n}(T_n^- - t_n) < z) \rightarrow \Phi(z/\sigma)$. It is similar to show that

$P(\sqrt{n}(T_n^+ - t_n) < z) \rightarrow \Phi(z/\sigma)$. At this point, we are able to conclude that the first claim in the theorem holds, i.e. $\sqrt{n}(T_n - t_n) \xrightarrow{d} N(0, \sigma^2)$.

To prove (A.5), it suffices to check Lindeberg's condition:

$$E(\xi_i^2 \mathbf{1}\{|\xi_i| > \sqrt{n}\varepsilon\}) \rightarrow 0$$

for any $\varepsilon > 0$. Notice that $\lambda_n(u_n) \rightarrow 0$ and $\sigma_n(u_n) \rightarrow \sigma$, we only need to show

$$E(\psi_n^2(X - u_n) \mathbf{1}\{|\psi_n(X - u_n)| > \sqrt{n}\varepsilon\}) \rightarrow 0.$$

This is true due to

$$\psi_n^2(X - u_n) \leq |X - u_n|^2 \leq 2|X|^2 + 2u_n^2$$

and dominated convergence theorem. $\square$

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. High-dimensional analysis of semidefinite relaxations for sparse principal components. *Information Theory, 2008 ISIT 2008 IEEE International Symposium on IEEE*. 2008
- Antoniadis A, Fan J. Regularization of wavelet approximations. *Journal of the American Statistical Association*. 2001; 96
- Bai J. Inferential theory for factor models of large dimensions. *Econometrica*. 2003; 71:135–171.

- Berthet Q, Rigollet P. Optimal detection of sparse principal components in high dimension. *The Annals of Statistics*. 2013; 41:1780–1815.
- Bickel PJ. Another look at robustness: a review of reviews and some new developments. *Scandinavian Journal of Statistics*. 1976:145–168.
- Bickel PJ, Levina E. Covariance regularization by thresholding. *The Annals of Statistics*. 2008:2577–2604.
- Birnbaum A, Johnstone IM, Nadler B, Paul D. Minimax bounds for sparse pca with noisy high-dimensional data. *Annals of statistics*. 2013; 41:1055. [PubMed: 25324581]
- Cai T, Liu W. Adaptive thresholding for sparse covariance matrix estimation. *Journal of the American Statistical Association*. 2011; 106:672–684.
- Cai T, Ma Z, Wu Y. Optimal estimation and rank detection for sparse spiked covariance matrices. *Probability Theory and Related Fields*. 2013; 161:781–815.
- Cai TT, Zhang CH, Zhou HH. Optimal rates of convergence for covariance matrix estimation. *The Annals of Statistics*. 2010; 38:2118–2144.
- Candès EJ, Li X, Ma Y, Wright J. Robust principal component analysis? *Journal of the ACM (JACM)*. 2011; 58:11.
- Candès EJ, Recht B. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*. 2009; 9:717–772.
- Catoni O. *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*. Vol. 48. Institut Henri Poincaré; 2012. Challenging the empirical mean and empirical variance: a deviation study.
- Chandrasekaran V, Sanghavi S, Parrilo PA, Willsky AS. Rank-sparsity incoherence for matrix decomposition. *SIAM Journal on Optimization*. 2011; 21:572–596.
- Dinov ID, Boscardin JW, Mega MS, Sowell EL, Toga AW. A wavelet-based statistical analysis of fmri data. *Neuroinformatics*. 2005; 3:319–342. [PubMed: 16284415]
- Fama EF, French KR. Common risk factors in the returns on stocks and bonds. *Journal of financial economics*. 1993; 33:3–56.
- Fan J, Fan Y, Lv J. High dimensional covariance matrix estimation using a factor model. *Journal of Econometrics*. 2008; 147:186–197.
- Fan J, Li Q, Wang Y. Estimation of high dimensional mean regression in the absence of symmetry and light tail assumptions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2017; 79:247–265. [PubMed: 28479862]
- Fan J, Liao Y, Mincheva M. High dimensional covariance matrix estimation in approximate factor models. *Annals of statistics*. 2011; 39:3320. [PubMed: 22661790]
- 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*. 2015a; 186:367–387. [PubMed: 26195851]
- Fan J, Liao Y, Wang W. Projected principal component analysis in factor models. *Annals of statistics*. 2016; 44:219. [PubMed: 26783374]
- Fan J, Liu H, Wang W. Large covariance estimation through elliptical factor models. *arXiv preprint arXiv:1507.08377*. 2015b
- Fang KT, Kotz S, Ng KW. *Symmetric multivariate and related distributions*. Chapman and Hall; 1990.
- Forni M, Hallin M, Lippi M, Reichlin L. The generalized dynamic-factor model: Identification and estimation. *Review of Economics and statistics*. 2000; 82:540–554.
- Forni M, Lippi M. The generalized dynamic factor model: representation theory. *Econometric theory*. 2001; 17:1113–1141.
- Friedman J, Hastie T, Tibshirani R. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*. 2008; 9:432–441. [PubMed: 18079126]
- Grant M, Boyd S, Ye Y. *Cvx: Matlab software for disciplined convex programming*. 2008
- Han F, Liu H. Scale-invariant sparse PCA on high-dimensional meta-elliptical data. *Journal of the American Statistical Association*. 2014; 109:275–287. [PubMed: 24932056]
- Hsu D, Sabato S. Heavy-tailed regression with a generalized median-of-means. *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*. 2014

- Huber PJ. Robust estimation of a location parameter. *The Annals of Mathematical Statistics*. 1964; 35:73–101.
- 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. URL <http://amstat.tandfonline.com/doi/abs/10.1198/jasa.2009.0121>. [PubMed: 20617121]
- Lam C, Fan J. Sparsistency and rates of convergence in large covariance matrix estimation. *Annals of statistics*. 2009; 37:4254. [PubMed: 21132082]
- Lepskii O. Asymptotically minimax adaptive estimation. i: Upper bounds. optimally adaptive estimates. *Theory of Probability & Its Applications*. 1992; 36:682–697.
- Lintner J. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *The review of economics and statistics*. 1965:13–37.
- Ma Z. Sparse principal component analysis and iterative thresholding. *The Annals of Statistics*. 2013; 41:772–801.
- Meinshausen N, Bühlmann P. High-dimensional graphs and variable selection with the lasso. *The Annals of Statistics*. 2006:1436–1462.
- 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.
- Ravikumar P, Wainwright MJ, Raskutti G, Yu B, et al. High-dimensional covariance estimation by minimizing ℓ_1 -penalized log-determinant divergence. *Electronic Journal of Statistics*. 2011; 5:935–980.
- Rothman AJ, Levina E, Zhu J. Generalized thresholding of large covariance matrices. *Journal of the American Statistical Association*. 2009; 104:177–186.
- Sharpe WF. Capital asset prices: A theory of market equilibrium under conditions of risk*. *The journal of finance*. 1964; 19:425–442.
- Stock J, Watson M. Forecasting using principal components from a large number of predictors. 2002; 97:1167–1179.
- Vu VQ, Lei J. Minimax rates of estimation for sparse pca in high dimensions. arXiv preprint arXiv: 1202.0786. 2012
- Wang W, Fan J. Asymptotics of empirical eigen-structure for high dimensional spiked covariance. *The Annals of statistics* (to appear). 2017
- Xu H, Caramanis C, Sanghavi S. Robust pca via outlier pursuit. *Advances in Neural Information Processing Systems*. 2010

Figure 1.

The histograms show the estimates of $\text{Var}(X_j)$ with different parameters α , parametrized by β via (4.1), in 1000 runs. $X_j \sim t_{4,2}$ so that the true variance $\text{Var}(X_j) = 1.909$. The sample size $n = 100$.

Figure 2.

The histograms show the estimates of $Cov(X_i, X_j)$ with different parameters α in 1000 runs. The true covariance $Cov(X_i, X_j) = 0$. $n = 100$ and the degree of freedom is 4.2.

Figure 3.

Errors of robust estimates against varying p . Blue line represents ratio of errors with $\sum_z^R$ over errors with $\sum_z^S$, while black line represents ratio of errors with $\sum_z^K$ over errors with $\sum_z^S$. z_t is generated by multivariate t -distribution with $df=3$ (solid), 5 (dashed) and ∞ (dotted). The median errors and their IQR over 100 simulations are reported.

Figure 4.

Errors of robust estimates against varying p . Blue line represents ratio of errors with $\hat{\Sigma}_z^R$ over errors with $\hat{\Sigma}_z^S$, while black line represents ratio of errors with $\hat{\Sigma}_z^K$ over errors with $\hat{\Sigma}_z^S$. z_t is generated by element-wise *iid* t -distribution with $df=3$ (solid), 5 (dashed) and ∞ (dotted). The median errors and their IQR over 100 simulations are reported.

Figure 5.

Q-Q plot of excess returns y_{it} for all i and t against Gaussian distribution and t -distribution with degree of freedom 2, 4 and 6. For each plot, a line is fitted by connecting points at first and third quartile.

Figure 6.

Q-Q plot of factor f_{it} against Gaussian distribution and t -distribution with degree of freedom 2, 4 and 6. For each plot, a line is fitted by connecting points at first and third quartile.

Figure 7.

Left panel: Histogram of eigenvalues of sample covariance matrix YY^T/n . The histogram is plotted on the logarithmic scale, i.e. each bin counts the number of $\log \lambda_i$ in a given range.

Right panel: Proportion of residue eigenvalues $\sum_{i=K+1}^p \lambda_i / \sum_{i=1}^p \lambda_i$ against varying K , where λ_i is the i^{th} largest eigenvalue of sample covariance matrix YY^T/n .

Figure 8. ($R^R(w)$, $R^S(w)$) for multiple randomly generated w . The four plots compare the errors of the two methods under different settings (upper left: no short selling; upper right: exposure $c = 1.4$; lower left: exposure $c = 1.8$; lower right: exposure $c = 2.2$). The red diamond in the first plot corresponds to uniform weights. The dashed line is the 45 degree line representing equal performance. Our robust method gives smaller errors.