

# Learning a Factor Model via Regularized PCA

Yi-Hao Kao · Benjamin Van Roy

Received: 15 July 2012 / Accepted: 30 March 2013 / Published online: 20 April 2013  
©The Authors

**Abstract** We consider the problem of learning a linear factor model. We propose a regularized form of principal component analysis (PCA) and demonstrate through experiments with synthetic and real data the superiority of resulting estimates to those produced by pre-existing factor analysis approaches. We also establish theoretical results that explain how our algorithm corrects the biases induced by conventional approaches. An important feature of our algorithm is that its computational requirements are similar to those of PCA, which enjoys wide use in large part due to its efficiency.

**Keywords** Principal component analysis · Factor model · High-dimensional data · Covariance matrix estimation · Regularization

## 1 Introduction

Linear factor models have been widely used for a long time and with notable success in economics, finance, medicine, psychology, and various other natural and social sciences (Harman, 1976). In such a model, each observed variable is a linear combination of unobserved common factors plus idiosyncratic noise, and the collection of random variables is jointly Gaussian. We consider in this paper the problem of learning a factor model from a training set of vector observations. In particular, our learning problem entails simultaneously estimating the loadings of each factor and the residual variance of each variable. We seek an estimate of these parameters that best explains out-of-sample data. For this purpose, we consider the likelihood of test data that is independent of the training data. As such, our goal is to design a learning algorithm that maximizes the likelihood of a test set that is not used in the learning process.

---

Editor: Csaba Szepesvari

---

Y.-H. Kao, Stanford University, CA, USA  
e-mail: yhkao@alumni.stanford.edu  
B. Van Roy, Stanford University, CA, USA  
e-mail: bvr@stanford.edu

A common approach to factor model learning involves application of principal component analysis (PCA). If the number of factors is known and residual variances are assumed to be uniform, PCA can be applied to efficiently compute model parameters that maximize likelihood of the training data (Tipping and Bishop, 1999). In order to simplify analysis, we begin our study with a context for which PCA is ideally suited. In particular, before treating more general models, we will restrict attention to models in which residual variances are uniform. As a baseline among learning algorithms, we consider applying PCA together with cross-validation, computing likelihood-maximizing parameters for different numbers of factors and selecting the number of factors that maximizes likelihood of a portion of the training data that is reserved for validation. We will refer to this baseline as *uniform-residual rank-constrained maximum-likelihood* (URM) estimation.

To improve on URM, we propose *uniform-residual trace-penalized maximum-likelihood* (UTM) estimation. Rather than estimating parameters of a model with a fixed number of factors and iterating over the number of factors, this approach maximizes likelihood across models without restricting the number of factors but instead penalizing the trace of a matrix derived from the model’s covariance matrix. This trace penalty serves to regularize the model and naturally selects a parsimonious set of factors. The coefficient of this penalty is chosen via cross-validation, similarly with the way in which the number of factors is selected by URM. Through a computational study using synthetic data, we demonstrate that UTM results in better estimates than URM. In particular, we find that UTM requires as little as two-thirds of the quantity of data used by URM to match its performance. Further, leveraging recent work on random matrix theory, we establish theoretical results that explain how UTM corrects the biases induced by URM.

We then extend UTM to address the more general and practically relevant learning problem in which residual variances are not assumed to be uniform. To evaluate the resulting algorithm, which we refer to as *scaled trace-penalized maximum-likelihood* (STM) estimation, we carry out experiments using both synthetic data and real stock price data. The computational results demonstrate that STM leads to more accurate estimates than alternatives available from prior art. We also provide an analysis to illustrate how these alternatives can suffer from biases in this nonuniform residual variance setting.

Aside from the aforementioned empirical and theoretical analyses, an important contribution of this paper is in the design of algorithms that make UTM and STM efficient. The UTM approach is formulated as a convex semidefinite program (SDP), which can be solved by existing algorithms such as interior-point methods or alternating direction method of multipliers (see, e.g., Boyd et al. (2011)). However, when the data dimension is large, as is the case in many relevant contexts, such algorithms can take too long to be practically useful. This exemplifies a recurring obstacle that arises in the use of SDP formulations to study large data sets. We propose an algorithm based on PCA that solves the UTM formulation efficiently. In particular, we found this method to typically require three orders of magnitude less compute time than the alternating direction method of multipliers. Variations of PCA such as URM have enjoyed wide use to a large extent because of their efficiency, and the computation time required for UTM is essentially the same as that of URM. STM requires additional computation but remains in reach for problems where the computational costs of URM are acceptable.

Our formulation is related to that of Chandrasekaran et al. (2012), which estimates a factor model using a similar trace penalty. There are some important differences, however, that distinguish our work. First, the analysis of Chandrasekaran et al. (2012) focuses on establishing perfect recovery of structure in an asymptotic regime, whereas our work makes the point that this trace penalty reduces nonasymptotic bias. Second, our approach to dealing with nonuniform residual variances is distinctive and we demonstrate through computational and theoretical analysis that this difference reduces bias. Third, Chandrasekaran et al. (2012) treats the problem as a semidefinite program, whose solution is often computationally demanding when data dimension is large. We provide an algorithm based on PCA that efficiently solves our problem. The algorithm can also be adapted to solve the formulation of Chandrasekaran et al. (2012), though that is not the aim of our paper.

In addition, there is another thread of research on regularized maximum-likelihood estimation for covariance matrices that relates loosely to this paper. Along this line, Banerjee et al. (2008) regularizes

maximum-likelihood estimation by the  $\ell_1$  norm of the inverse covariance matrix in order to recover a sparse graphical model. An efficient algorithm called graphical Lasso was then proposed by Friedman et al. (2008) for solving this formulation. Similar formulations can also be found in Yuan and Lin (2007) and Ravikumar et al. (2011), who instead penalize the  $\ell_1$  norm of off-diagonal elements of the inverse covariance matrix when computing maximum-likelihood estimates. For a detailed survey, see Pourahmadi (2011). Although our approach shares some of the spirit represented by this line of research in that we also regularize maximum-likelihood estimation by an  $\ell_1$ -like penalty, the settings are fundamentally different: while ours focuses on a factor model, theirs are based on sparse graphical models. We propose an approach that corrects the bias induced by conventional factor analysis, whereas their results are mainly concerned with accurate recovery of the topology of an underlying graph. As such, their work does not address biases in covariance estimates. On the algorithmic front, we develop a simple and efficient solution method that builds on PCA. On the contrary, their algorithms are more complicated and computationally demanding.<sup>1</sup>

## 2 Problem Formulation

We consider the problem of learning a factor model without knowledge of the number of factors. Specifically, we want to estimate a  $M \times M$  covariance matrix  $\Sigma_*$  from samples  $\mathbf{x}_{(1)}, \dots, \mathbf{x}_{(N)} \sim \mathcal{N}(0, \Sigma_*)$ , where  $\Sigma_*$  is the sum of a symmetric matrix  $\mathbf{F}_* \succeq 0$  and a diagonal matrix  $\mathbf{R}_* \succeq 0$ . These samples can be thought of as generated by a factor model of the form  $\mathbf{x}_{(n)} = \mathbf{F}_*^{\frac{1}{2}} \mathbf{z}_{(n)} + \mathbf{w}_{(n)}$ , where  $\mathbf{z}_{(n)} \sim \mathcal{N}(0, \mathbf{I})$  represents a set of common factors and  $\mathbf{w}_{(n)} \sim \mathcal{N}(0, \mathbf{R}_*)$  represents residual noise. The number of factors is represented by  $\text{rank}(\mathbf{F}_*)$ , and it is usually assumed to be much smaller than the dimension  $M$ .

Our goal is to produce based on the observed samples a factor loadings matrix  $\mathbf{F} \succeq 0$  and a residual variance matrix  $\mathbf{R} \succeq 0$  such that the resulting factor model best explains out-of-sample data. In particular, we seek a pair of  $(\mathbf{F}, \mathbf{R})$  such that the covariance matrix  $\Sigma = \mathbf{F} + \mathbf{R}$  maximizes the average log-likelihood of out-of-sample data:

$$L(\Sigma, \Sigma_*) \triangleq \mathbb{E}_{\mathbf{x} \sim \mathcal{N}(0, \Sigma_*)} [\log p(\mathbf{x} | \Sigma)] = -\frac{1}{2} (M \log(2\pi) + \log \det(\Sigma) + \text{tr}(\Sigma^{-1} \Sigma_*)).$$

This is also equivalent to minimizing the Kullback-Leibler divergence between  $\mathcal{N}(0, \Sigma_*)$  and  $\mathcal{N}(0, \Sigma)$ .

## 3 Learning Algorithms

Given our objective, one simple approach is to choose an estimate  $\Sigma$  that maximizes in-sample log-likelihood:

$$\log p(\mathcal{X} | \Sigma) = -\frac{N}{2} (M \log(2\pi) + \log \det(\Sigma) + \text{tr}(\Sigma^{-1} \Sigma_{\text{SAM}})), \quad (1)$$

where  $\mathcal{X} = \{\mathbf{x}_{(1)}, \dots, \mathbf{x}_{(N)}\}$ , and we use  $\Sigma_{\text{SAM}} = \sum_{n=1}^N \mathbf{x}_{(n)} \mathbf{x}_{(n)}^T / N$  to denote the sample covariance matrix. Here, the maximum likelihood estimate is simply given by  $\Sigma = \Sigma_{\text{SAM}}$ .

The problem with maximum likelihood estimation in this context is that in-sample log-likelihood does not accurately predict out-of-sample log-likelihood unless the number of samples  $N$  far exceeds the dimension  $M$ . In fact, when the number of samples  $N$  is smaller than the dimension  $M$ ,  $\Sigma_{\text{SAM}}$  is ill-conditioned and the out-of-sample log-likelihood is negative infinity. One remedy to such poor generalization involves exploiting factor structure, as we discuss in this section.

<sup>1</sup>The code of our algorithms can be downloaded at: <http://www.yhkao.com/RPCA-code.zip>.

### 3.1 Uniform Residual Variances

We begin with a simplified scenario in which the residual variances are assumed to be identical. As we will later see, such simplification facilitates theoretical analysis. This assumption will be relaxed in the next subsection.

#### 3.1.1 Constraining the Number of Factors

Given a belief that the data is generated by a factor model with few factors, one natural approach is to employ maximum likelihood estimation with a constraint on the number of factors. Now suppose the residual variances in the generative model are identical, and as a result we impose an additional assumption that  $\mathbf{R}$  is a multiple  $\sigma^2\mathbf{I}$  of the identity matrix. This leads to an optimization problem

$$\begin{aligned} \max_{\mathbf{F} \in \mathbb{S}_+^M, \sigma^2 \in \mathbb{R}_+} \quad & \log p(\mathcal{X}|\boldsymbol{\Sigma}) \\ \text{s.t.} \quad & \boldsymbol{\Sigma} = \mathbf{F} + \sigma^2\mathbf{I} \\ & \text{rank}(\mathbf{F}) \leq K \end{aligned} \tag{2}$$

where  $\mathbb{S}_+^M$  denote the set of all  $M \times M$  positive semidefinite symmetric matrices, and  $K$  is the exogenously specified number of factors. In this case, we can efficiently compute an analytical solution via principal component analysis (PCA), as established in Tipping and Bishop (1999). This involves first computing an eigendecomposition of the sample covariance matrix  $\boldsymbol{\Sigma}_{\text{SAM}} = \mathbf{B}\mathbf{S}\mathbf{B}^T$ , where  $\mathbf{B} = [\mathbf{b}_1 \ \dots \ \mathbf{b}_M]$  is orthonormal and  $\mathbf{S} = \text{diag}(s_1, \dots, s_M)$  with  $s_1 \geq \dots \geq s_M$ . The solution to (2) is then given by

$$\begin{aligned} \hat{\sigma}^2 &= \frac{1}{M-K} \sum_{i=K+1}^M s_i \\ \hat{\mathbf{F}} &= \sum_{k=1}^K (s_k - \hat{\sigma}^2) \mathbf{b}_k \mathbf{b}_k^T. \end{aligned} \tag{3}$$

In other words, the estimate for residual variance equals the average of the last  $M - K$  sample eigenvalues, whereas the estimate for factor loading matrix is spanned by the top  $K$  sample eigenvectors with coefficients  $s_k - \hat{\sigma}^2$ . We will refer to this method as *uniform-residual rank-constrained maximum-likelihood* estimation, and use  $\boldsymbol{\Sigma}_{\text{URM}}^K = \hat{\mathbf{F}} + \hat{\sigma}^2\mathbf{I}$  to denote the covariance matrix resulting from this procedure. It is easy to see that the eigenvalues of  $\boldsymbol{\Sigma}_{\text{URM}}^K$  are  $s_1, \dots, s_K, \hat{\sigma}^2, \dots, \hat{\sigma}^2$ , as illustrated in Figure 1(a).

A number of methods have been proposed for estimating the number of factors  $K$  (Akaike, 1987; Bishop, 1998; Minka, 2000; Hirose et al., 2011). Cross-validation provides a conceptually simple approach that in practice works at least about as well as any other. To obtain best performance from such a procedure, one would make use of so-called  $n$ -fold cross-validation. To keep things simple in our study and comparison of estimation methods, for all methods we will consider, we employ a version of cross-validation that reserves a single subset of data for validation and selection of  $K$ . Details of the procedure we used can be found in the appendix. Through selection of  $K$ , this procedure arrives at a covariance matrix which we will denote by  $\boldsymbol{\Sigma}_{\text{URM}}$ .

#### 3.1.2 Penalizing the Trace

Although (2) can be elegantly solved via PCA, it is unclear that imposing a hard constraint on the number of factors will lead to an optimal estimate. In particular, one might suspect a “softer” regularization could improve estimation accuracy. Motivated by this idea, we propose penalizing the trace instead of constraining the rank of the factor loading matrix. As we shall see in the experiment results and theoretical analysis, such an approach indeed improves estimation accuracy significantly.

Nevertheless, naively replacing the rank constraint of (2) by a trace constraint  $\text{tr}(\mathbf{F}) \leq t$  will result in a non-convex optimization problem, and it is not clear to us whether it can be solved efficiently. Let us explore a related alternative. Some straightforward matrix algebra shows that if  $\mathbf{\Sigma} = \mathbf{F} + \sigma^2 \mathbf{I}$  with  $\mathbf{F} \in \mathbb{S}_+^M$  and  $\sigma^2 > 0$ , then the matrix defined by  $\mathbf{G} = \sigma^{-2} \mathbf{I} - \mathbf{\Sigma}^{-1}$  is in  $\mathbb{S}_+^M$ , with  $\text{rank}(\mathbf{G}) = \text{rank}(\mathbf{F})$ . This observation, together with the well-known fact that the log-likelihood of  $\mathcal{X}$  is concave in the *inverse* covariance matrix  $\mathbf{\Sigma}^{-1}$ , motivates the following convex program:

$$\begin{aligned} \max_{\mathbf{G} \in \mathbb{S}_+^M, v \in \mathbb{R}_+} \quad & \log p(\mathcal{X} | \mathbf{\Sigma}) \\ \text{s.t.} \quad & \mathbf{\Sigma}^{-1} = v \mathbf{I} - \mathbf{G} \\ & \text{tr}(\mathbf{G}) \leq t. \end{aligned}$$

Here, the variable  $v$  represents the reciprocal of residual variance. Pricing out the trace constraint leads to a closely related problem in which the trace is penalized rather than constrained:

$$\begin{aligned} \max_{\mathbf{G} \in \mathbb{S}_+^M, v \in \mathbb{R}_+} \quad & \log p(\mathcal{X} | \mathbf{\Sigma}) - \lambda \text{tr}(\mathbf{G}) \\ \text{s.t.} \quad & \mathbf{\Sigma}^{-1} = v \mathbf{I} - \mathbf{G}. \end{aligned} \quad (4)$$

We will consider the trace penalized problem instead of the trace constrained problem because it is more convenient to design algorithms that address the penalty rather than the constraint. Let  $(\hat{\mathbf{G}}, \hat{v})$  be an optimal solution to (4), and let  $\mathbf{\Sigma}_{\text{UTM}}^\lambda = (\hat{v} \mathbf{I} - \hat{\mathbf{G}})^{-1}$  denote the covariance matrix estimate derived from it. Here, the ‘‘U’’ indicates that residual variances are assumed to be *uniform* across variables and ‘‘T’’ stands for *trace-penalized*.

It is easy to see that (4) is a semidefinite program. As such, the problem can be solved in polynomial time by existing algorithms such as interior-point methods or alternating direction method of multipliers (ADMM). However, when the number of variables  $M$  is large, as is the case in many contexts of practical import, such algorithms can take too long to be practically useful. One contribution of this paper is an efficient method for solving (4), which we now describe. The following result motivates the algorithm we will propose for computing  $\mathbf{\Sigma}_{\text{UTM}}^\lambda$ :

**Theorem 1**  $\mathbf{\Sigma}_{\text{SAM}}$  and  $\mathbf{\Sigma}_{\text{UTM}}^\lambda$  share the same trace and eigenvectors, and letting the eigenvalues of the two matrices, sorted in decreasing order, be denoted by  $s_1, \dots, s_M$  and  $h_1, \dots, h_M$ , respectively, we have

$$h_m = \max \left\{ s_m - \frac{2\lambda}{N}, \frac{1}{\hat{v}} \right\}, \text{ for } m = 1, \dots, M. \quad (5)$$

This theorem suggests an algorithm for computing  $\mathbf{\Sigma}_{\text{UTM}}^\lambda$ . First, we compute the eigendecomposition of  $\mathbf{\Sigma}_{\text{SAM}} = \mathbf{B} \mathbf{S} \mathbf{B}^T$ , where  $\mathbf{B}$  and  $\mathbf{S}$  are as defined in Section 3.1.1. This provides the eigenvectors and trace of  $\mathbf{\Sigma}_{\text{UTM}}^\lambda$ . To obtain its eigenvalues, we only need to determine the value of  $\hat{v}$  such that the eigenvalues given by (5) sum to the desired trace. This is equivalent to determining the largest integer  $K$  such that

$$s_K - \frac{2\lambda}{N} > \frac{1}{M-K} \left( K \cdot \frac{2\lambda}{N} + \sum_{m=K+1}^M s_m \right).$$

To see this, note that setting

$$\begin{aligned} \hat{v}^{-1} &= \frac{1}{M-K} \left( K \cdot \frac{2\lambda}{N} + \sum_{m=K+1}^M s_m \right) \\ h_m &= \begin{cases} s_m - \frac{2\lambda}{N}, & m = 1, \dots, K \\ \hat{v}^{-1}, & m = K+1, \dots, M \end{cases} \end{aligned}$$

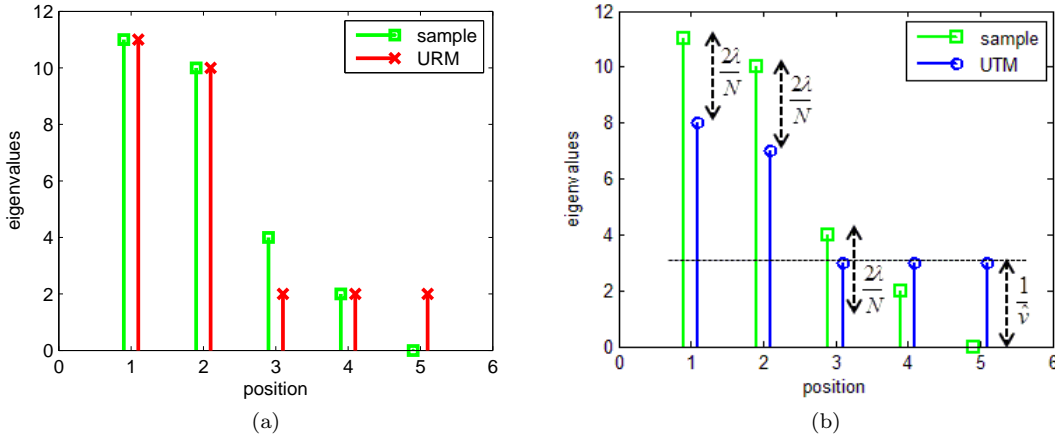


Figure 1: (a) An example of sample eigenvalues and the corresponding eigenvalues of URM estimate, with  $M = 5$  and  $K = 2$ . URM essentially preserves the top eigenvalues and averages the remaining ones as residual variance. (b) An example of sample eigenvalues and the corresponding eigenvalues of UTM estimate. With a particular choice of  $\lambda$ , this estimate has two outstanding eigenvalues, but their magnitudes are  $2\lambda/N$  below the sample ones. Its residual variance  $1/\hat{v}$  is determined in a way that ensures the summation of UTM eigenvalues equal to the sample one.

uniquely satisfies (5) and ensures  $\sum_{m=1}^M h_m = \sum_{m=1}^M s_m$ . Algorithm 1 presents this method in greater detail. In our experiments, we found this method to typically require three orders of magnitude less compute time than ADMM. For example, it can solve a problem of dimension  $M = 1000$  within seconds on a workstation, whereas ADMM requires hours to attain the same level of accuracy.

Also note that for reasonably large  $\lambda$ , this algorithm will flatten most sample eigenvalues and allow only the largest eigenvalues to remain outstanding, effectively producing a factor model estimate. Figure 1(b) illustrates this effect. Comparing Figure 1(a) and Figure 1(b), it is easy to see that URM and UTM primarily differ in the largest eigenvalues they produce: while URM simply retains the largest sample eigenvalues, UTM subtracts a constant  $2\lambda/N$  from them. As we shall see in the theoretical analysis, this subtraction indeed corrects the bias incurred in sample eigenvalues.

Like URM, the most computationally expensive step of UTM lies in the eigendecomposition.<sup>2</sup> Beyond that, the evaluation of UTM eigenvalues for any given  $\lambda$  takes  $O(M)$ , and is generally negligible. In our implementation, this regularization parameter  $\lambda$  is chosen by cross-validation from a range around  $M\hat{\sigma}^2$ , whose reasons will become apparent in Section 5.1. We denote by  $\Sigma_{\text{UTM}}$  the covariance matrix resulting from this cross-validation procedure.

### 3.2 Nonuniform Residual Variances

We now relax the assumption of uniform residual variances and discuss several methods for the general case. As in the previous subsection, the hyper-parameters of these methods will be selected by cross-validation.

<sup>2</sup>In fact, a full eigendecomposition is not required, as we only need the top  $K$  eigenvectors to compute the estimate.

---

**Algorithm 1** Procedure for computing  $\Sigma_{\text{UTM}}^\lambda$

**Input:**  $\mathcal{X}, \lambda$

**Output:**  $\Sigma_{\text{UTM}}^\lambda$

---

Compute eigendecomposition  $\Sigma_{\text{SAM}} = \mathbf{B}\mathbf{S}\mathbf{B}^\top$   
 $v_k^{-1} \leftarrow \frac{1}{M-k} \left( k \cdot \frac{2\lambda}{N} + \sum_{m=k+1}^M s_m \right), \quad \forall k = 0, 1, \dots, M-1$   
 $K \leftarrow \max \{k : s_k - \frac{2\lambda}{N} > v_k^{-1}\}$  // define  $s_0 = \infty$   
 $\hat{v} \leftarrow v_K$   
 $h_m \leftarrow \begin{cases} s_m - \frac{2\lambda}{N} & \text{if } m \leq K \\ \hat{v}^{-1} & \text{otherwise} \end{cases}, \quad \forall m = 1, \dots, M$   
 $\Sigma_{\text{UTM}}^\lambda \leftarrow \sum_{k=1}^K (h_k - \hat{v}^{-1}) \mathbf{b}_k \mathbf{b}_k^\top + \hat{v}^{-1} \mathbf{I}$

---

### 3.2.1 Constraining the Number of Factors

Without the assumption of uniform residual variances, the ranked-constrained maximum-likelihood formulation can be written as

$$\begin{aligned} \max_{\mathbf{F} \in \mathbb{S}_+^M, \mathbf{R} \in \mathbb{D}_+^M} \quad & \log p(\mathcal{X}|\Sigma) \\ \text{s.t.} \quad & \Sigma = \mathbf{F} + \mathbf{R} \\ & \text{rank}(\mathbf{F}) \leq K \end{aligned} \quad (6)$$

where  $\mathbb{D}_+^M$  denote the set of all  $M \times M$  positive semidefinite diagonal matrices. Unlike (2), this formulation is generally hard to solve, and therefore we consider two widely-used approximate solutions.

The expectation-maximization (EM) algorithm (Rubin and Thayer, 1982) is arguably the most conventional approach to solving (6), though there is no guarantee that this will result in a global optimal solution. The algorithm generates a sequence of iterates  $\mathbf{F}^{\frac{1}{2}} \in \mathbb{R}^{M \times K}$  and  $\mathbf{R} \in \mathbb{D}_+^M$ , such that the covariance matrix  $\Sigma = \mathbf{F}^{\frac{1}{2}} \mathbf{F}^{\frac{1}{2}\top} + \mathbf{R}$  increases the log-likelihood of  $\mathcal{X}$  with each iteration. Each iteration involves an estimation step in which we assume the data are generated according to the covariance matrix  $\Sigma = \mathbf{F}^{\frac{1}{2}} \mathbf{F}^{\frac{1}{2}\top} + \mathbf{R}$ , and compute expectations  $\mathbb{E}[\mathbf{z}_{(n)}|\mathbf{x}_{(n)}]$  and  $\mathbb{E}[\mathbf{z}_{(n)}\mathbf{z}_{(n)}^\top|\mathbf{x}_{(n)}]$  for  $n = 1, \dots, N$ . A maximization step then updates  $\mathbf{F}$  and  $\mathbf{R}$  based on these expectations. In our implementation, the initial  $\mathbf{F}$  and  $\mathbf{R}$  are selected by the MRH algorithm described in the next paragraph. We will denote the estimate produced by the EM algorithm by  $\Sigma_{\text{EM}}^K$  and that resulting from further selection of  $K$  through cross-validation by  $\Sigma_{\text{EM}}$ .

A common heuristic for approximately solving (6) without entailing iterative computation is to first compute  $\Sigma_{\text{URM}}^K$  by PCA and then take the factor matrix estimate to be the  $\hat{\mathbf{F}}$  defined in (3) and the residual variances to be  $\hat{\mathbf{R}}_{m,m} = \left( \Sigma_{\text{SAM}} - \hat{\mathbf{F}} \right)_{m,m}$ , for  $m = 1, \dots, M$ . In other words,  $\hat{\mathbf{R}}_{m,m}$  is selected so that the diagonal elements of the estimated covariance matrix  $\hat{\Sigma} = \hat{\mathbf{F}} + \hat{\mathbf{R}}$  are equal to those of the sample covariance matrix. We will refer to this method as *marginal-variance-preserving rank-constrained heuristic* and denote the resulting estimates by  $\Sigma_{\text{MRH}}^K$  and  $\Sigma_{\text{MRH}}$ .

### 3.2.2 Penalizing the Trace

We now develop an extension of Algorithm 1 that applies when residual variances are nonuniform. One formulation that may seem natural involves replacing  $v\mathbf{I}$  in (4) with a diagonal matrix  $\mathbf{V} \in \mathbb{D}_+^M$ . That is,

$$\begin{aligned} \max_{\mathbf{G} \in \mathbb{S}_+^M, \mathbf{V} \in \mathbb{D}_+^M} \quad & \log p(\mathcal{X}|\Sigma) - \lambda \text{tr}(\mathbf{G}) \\ \text{s.t.} \quad & \Sigma^{-1} = \mathbf{V} - \mathbf{G}. \end{aligned} \quad (7)$$

Indeed, a closely related formulation is proposed in Chandrasekaran et al. (2012). However, as we will see in Sections 4 and 5, solutions to this formulation suffer from bias and do not compete well against the method we will propose next. That said, let us denote the estimates resulting from solving this formulation by  $\hat{\Sigma}_{\text{TM}}^\lambda$  and  $\hat{\Sigma}_{\text{TM}}$ , where “T” stands for *trace-penalized*. Also note that this formulation can be efficiently solved by a straightforward generalization of Theorem 1, though we will not elaborate on this.

Our approach involves componentwise scaling of the data. Consider an estimate  $\hat{\Sigma}$  of  $\Sigma_*$ . Recall that we evaluate the quality of the estimate using the expected log-likelihood  $L(\hat{\Sigma}, \Sigma_*)$  of out-of-sample data. If we multiply each data sample by a matrix  $\mathbf{T} \in \mathbb{R}^{M \times M}$ , the data set becomes  $\mathbf{T}\mathcal{X} \triangleq \{\mathbf{T}\mathbf{x}_{(1)}, \dots, \mathbf{T}\mathbf{x}_{(N)}\}$ , where  $\mathbf{T}\mathbf{x}_{(n)} \sim \mathcal{N}(0, \mathbf{T}\Sigma_*\mathbf{T}^\text{T})$ . If we also change our estimate accordingly to  $\mathbf{T}\hat{\Sigma}\mathbf{T}^\text{T}$  then the new expected log-likelihood becomes

$$L(\mathbf{T}\hat{\Sigma}\mathbf{T}^\text{T}, \mathbf{T}\Sigma_*\mathbf{T}^\text{T}) = L(\hat{\Sigma}, \Sigma_*) - \log \det \mathbf{T}.$$

Therefore, as long as we constrain  $\mathbf{T}$  to have unit determinant,  $L(\mathbf{T}\hat{\Sigma}\mathbf{T}^\text{T}, \mathbf{T}\Sigma_*\mathbf{T}^\text{T})$  will be equal to  $L(\hat{\Sigma}, \Sigma_*)$ , suggesting that if  $\hat{\Sigma}$  is a good estimate of  $\Sigma_*$  then  $\mathbf{T}\hat{\Sigma}\mathbf{T}^\text{T}$  is a good estimate of  $\mathbf{T}\Sigma_*\mathbf{T}^\text{T}$ . This motivates the following optimization problem:

$$\begin{aligned} \max_{\mathbf{G} \in \mathbb{S}_+^M, v \in \mathbb{R}_+, \mathbf{T} \in \mathbb{D}_+^M} \quad & \log p(\mathbf{T}\mathcal{X}|\Sigma) - \lambda \text{tr}(\mathbf{G}) \\ \text{s.t.} \quad & \Sigma^{-1} = v\mathbf{I} - \mathbf{G}. \\ & \log \det \mathbf{T} \geq 0. \end{aligned} \tag{8}$$

The solution to this problem identifies a componentwise-scaling matrix  $\mathbf{T} \in \mathbb{D}_+^M$  that allows the data to be best-explained by a factor model with uniform residual variances. Given an optimal solution,  $1/\mathbf{T}_{i,i}^2$  should be approximately proportional to the residual variance of the  $i$ th variable, so that scaling by  $\mathbf{T}_{i,i}$  makes residual variances uniform. Note that the optimization problem constrains  $\log \det \mathbf{T}$  to be nonnegative rather than zero. This makes the feasible region convex, and this constraint is binding at the optimal solution. Denote the optimal solution to (8) by  $(\hat{\mathbf{G}}, \hat{v}, \hat{\mathbf{T}})$ . Our estimate is thus given by  $\hat{\mathbf{T}}^{-1}(\hat{v}\mathbf{I} - \hat{\mathbf{G}})^{-1}\hat{\mathbf{T}}^{-\text{T}}$ .

The objective function of (8) is not concave in  $(\mathbf{G}, v, \mathbf{T})$ , but is biconcave in  $(\mathbf{G}, v)$  and  $\mathbf{T}$ . We solve it by coordinate ascent, alternating between optimizing  $(\mathbf{G}, v)$  and  $\mathbf{T}$ . This procedure is guaranteed convergence. In our implementation, we initialize  $\mathbf{T}$  by  $\mathbf{I}$ . We will denote the resulting estimates by  $\Sigma_{\text{STM}}$ , where “ST” stands for *scaled and trace-penalized*.

## 4 Experiments

We carried out two sets of experiments to compare the performance of aforementioned algorithms. The first is based on synthetic data, whereas the second uses historical prices of stocks that make up the S&P 500 index.

### 4.1 Synthetic Data

We generated two kinds of synthetic data. The first was generated by a model in which each residual has unit variance. This data was sampled according to the following procedure, which takes as input the number of factors  $K_*$ , the dimension  $M$ , the factor variances  $\sigma_f^2$ , and the number of samples  $N$ :

1. Sample  $K_*$  orthonormal vectors  $\phi_1, \phi_2, \dots, \phi_{K_*} \in \mathbb{R}^M$  isotropically.
2. Sample  $f_1, f_2, \dots, f_{K_*} \sim \mathcal{N}(0, \sigma_f^2)$ .
3. Let  $\mathbf{F}_*^{\frac{1}{2}} = [f_1\phi_1 \quad f_2\phi_2 \quad \dots \quad f_{K_*}\phi_{K_*}]$ .
4. Let  $\Sigma_* = \mathbf{F}_*^{\frac{1}{2}}\mathbf{F}_*^{\frac{\text{T}}{2}} + \mathbf{I}$ .






Figure 2: (a) The average out-of-sample log-likelihood delivered by URM and UTM, when residual variances are identical. (b) The average portion of data required by UTM to match the performance of URM. The error bars denote 95% confidence interval.

5. Sample  $\mathbf{x}_{(1)}, \dots, \mathbf{x}_{(N)}$  iid from  $\mathcal{N}(0, \Sigma_*)$ .

We repeated this procedure one hundred times for each  $N \in \{50, 100, 200, 400\}$ , with  $M = 200$ ,  $K_* = 10$ , and  $\sigma_f = 5$ . We applied to this data URM and UTM, since they are methods designed to treat such a scenario with uniform residual variances. Regularization parameters  $K$  and  $\lambda$  were selected via cross-validation, where about 70% of each data set was used for training and 30% for validation. Figure 2(a) plots out-of-sample log-likelihood delivered by the two algorithms. Performance is plotted as a function of the log-ratio of the number of samples to the number of variables, which represents the availability of data relative to the number of variables. We expect this measure to drive performance differences. UTM outperforms URM in all scenarios. The difference is largest when data is scarce. When data is abundant, both methods work about as well. This should be expected since both estimation methods are consistent.

To interpret this result in a more tangible way, we also plot the *equivalent data requirement* of UTM in Figure 2(b). This metric is defined as the portion of training data required by UTM to match the performance of URM. As we can see, UTM needs as little as 67% of the data used by the URM to reach the same estimation accuracy.

Our second type of synthetic data was generated using an entirely similar procedure except step 4 was replaced by

$$\Sigma_* = \mathbf{F}_*^{\frac{1}{2}} \mathbf{F}_*^{\frac{T}{2}} + \text{diag}(e^{r_1}, e^{r_2}, \dots, e^{r_M}),$$

where  $r_1, \dots, r_M$  were sampled iid from  $\mathcal{N}(0, \sigma_r^2)$ . Note that  $\sigma_r$  effectively controls the variation among residual variances. Since these residual variances are nonuniform, EM, MRH, TM, and STM were applied. Figure 3 plots the results for the cases  $\sigma_r = 0.5$  and  $\sigma_r = 0.8$ , corresponding to moderate and large variation among residual variances, respectively. In either case, STM outperforms the alternatives. Figure 4 further gives the equivalent data requirement of STM with respect to each alternative. It is worth pointing out that the performance of MRH and TM degrades significantly as the variation among residual variances grows, while EM is less susceptible to such change. We will elaborate on this phenomenon in Section 5.2.




Figure 3: The average out-of-sample log-likelihood delivered by EM, MRH, TM, and STM, when residuals have independent random variances with (a)  $\sigma_r = 0.5$  and (b)  $\sigma_r = 0.8$ .

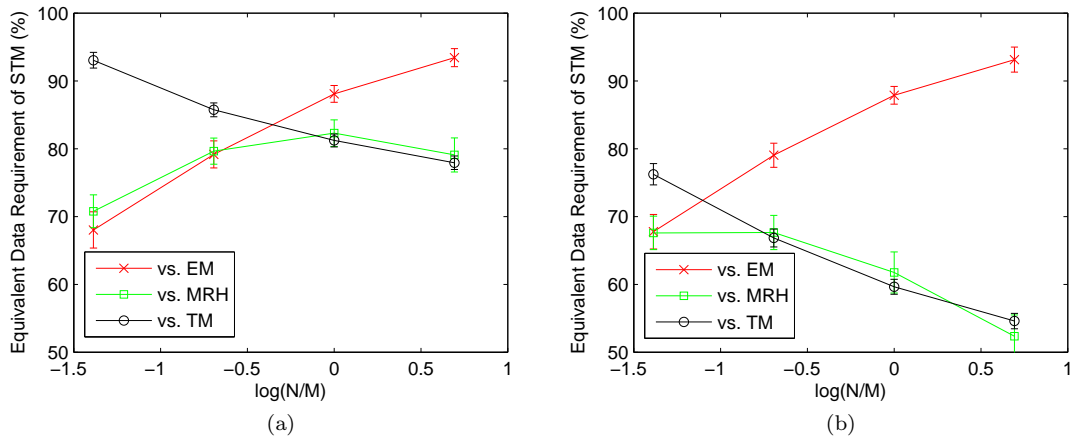


Figure 4: The equivalent data requirement of STM with respect to EM, MRH, and TM when (a)  $\sigma_r = 0.5$  and (b)  $\sigma_r = 0.8$ . The error bars denote 95% confidence interval.

## 4.2 S&P 500 Data

An important application area of factor analysis is finance, where return covariances are used to assess risks and guide diversification (Markowitz, 1952). The experiments we will now describe involve estimation of such covariances from historical daily returns of stocks represented in the S&P 500 index as of March, 2011. We use price data collected from the period starting November 2, 2001, and ending August 9, 2007. This period was chosen to avoid the erratic market behavior observed during the bursting of the dot-com bubble in 2000 and the financial crisis that began in 2008. Normalized daily log-returns were computed from closing prices through a process described in detail in the appendix. Over this duration, there were 1400 trading days and 453 of the stocks under consideration were active. This produced a data set  $\mathcal{Y} = \{\mathbf{y}_{(1)}, \dots, \mathbf{y}_{(1400)}\}$ , in which the  $i$ th component of  $\mathbf{y}_{(t)} \in \mathbb{R}^{453}$  represents the normalized log-daily-return of stock  $i$  on day  $t$ .

We generated estimates corresponding to each among a subset of the 1400 days. As would be done in real-time application, for each such day  $t$  we used  $N$  data points  $\{\mathbf{y}_{(t-N+1)}, \dots, \mathbf{y}_{(t)}\}$  that would have been available on that day to compute the estimate and subsequent data to assess performance. In particular, we generated estimates every ten days beginning on day 1200 and ending on day 1290. For each of these days, we evaluated average log-likelihood of log-daily-returns over the next ten days. Algorithm 2 formalizes this procedure.

---

### Algorithm 2 Testing Procedure $\mathcal{T}$

**Input:** learning algorithm  $\mathcal{U}$ , regularization parameter  $\theta$ , window size  $N$ , time point  $t$

**Output:** test-set log-likelihood

---

$\mathcal{X} \leftarrow \{\mathbf{y}_{(t-N+1)}, \dots, \mathbf{y}_{(t)}\}$  (training set)

$\mathcal{X}' \leftarrow \{\mathbf{y}_{(t+1)}, \dots, \mathbf{y}_{(t+10)}\}$  (test set)

$\hat{\Sigma} \leftarrow \mathcal{U}(\mathcal{X}, \theta)$

**return**  $\log p(\mathcal{X}' | \hat{\Sigma})$

---

These tests served the purpose of sliding-window cross-validation, as we tried a range of regularization parameters over this time period and used test results to select a regularization parameter for each algorithm. More specifically, for each algorithm  $\mathcal{U}$ , its regularization parameter was selected by

$$\hat{\theta} = \operatorname{argmax}_{\theta} \sum_{j=0}^9 \mathcal{T}(\mathcal{U}, \theta, N, 1200 + 10j).$$

On days 1300, 1310,  $\dots$ , 1390, we generated one estimate per day per algorithm, in each case using the regularization parameter selected earlier and evaluating average log-likelihood over the next ten days. For each algorithm  $\mathcal{U}$ , we took the average of these ten ten-day averages to be its out-of-sample performance, defined as

$$\frac{1}{100} \sum_{j=0}^9 \mathcal{T}(\mathcal{U}, \hat{\theta}, N, 1300 + 10j).$$

Figure 5 plots the performance delivered EM, MRH, TM, and STM with  $N \in \{200, 300, \dots, 1200\}$ . STM is the dominant solution. It is natural to ask why the performance of each algorithm improves then degrades as  $N$  grows. If the time series were stationary, one would expect performance to monotonically improve with  $N$ . However, this is a real time series and is not necessarily stationary. We believe that the distribution changes enough over about a thousand trading days so that using historical data collected further back worsens estimates. This observation points out that in real applications STM is likely to generate superior results even when all aforementioned algorithms are allowed to use all available data. This is in contrast with the experiments of Section 4.1 involving synthetic data, which may have led to an impression that the performance difference could be made small by using more data.




Figure 5: The average log-likelihood of test set, delivered by EM, MRH, TM, and STM, over different training-window sizes  $N$ .

## 5 Analysis

In this section, we explain why UTM and STM are expected to outperform alternatives as we have seen in our experimental results.

### 5.1 Uniform Residual Variances

Let us start with the simpler context in which residual variances are identical. In other words, let  $\Sigma_* = \mathbf{F}_* + \sigma^2 \mathbf{I}$  for a low rank matrix  $\mathbf{F}_*$  and uniform residual variance  $\sigma^2$ . We will begin our analysis with two desirable properties of UTM, and then move on to the comparison between UTM and URM.

It is easy to see that  $\Sigma_{\text{UTM}}^\lambda$  is a consistent estimator of  $\Sigma_*$  for any  $\lambda > 0$ , since  $\lim_{N \rightarrow \infty} \frac{2\lambda}{N} = 0$  and by Theorem 1 we have

$$\lim_{N \rightarrow \infty} \Sigma_{\text{UTM}}^\lambda = \lim_{N \rightarrow \infty} \Sigma_{\text{SAM}} \xrightarrow{a.s.} \Sigma_*.$$

Another important property of UTM is the fact that the trace of UTM estimate is the same as that of sample covariance matrix. This preservation is desirable as suggested by the following result.

**Proposition 1** For any fixed  $N, K$ , scalars  $\ell_1 \geq \ell_2 \geq \dots \geq \ell_K \geq \sigma^2 > 0$ , and any sequence of covariance matrices  $\Sigma_*^{(M)} \in \mathbb{S}_+^M$  with eigenvalues  $\ell_1, \dots, \ell_K, \sigma^2, \dots, \sigma^2$ , we have

$$\frac{\text{tr} \Sigma_{\text{SAM}}}{\text{tr} \Sigma_*} \xrightarrow{a.s.} 1,$$

as  $M \rightarrow \infty$  and

$$\Pr \left( \left| \frac{\text{tr} \Sigma_{\text{SAM}}}{\text{tr} \Sigma_*} - 1 \right| \geq \epsilon \right) \leq 2 \exp(-N \epsilon^2 \Omega(M)). \quad (9)$$

Note that, for any fixed number of samples  $N$ , the right-hand-side of (9) diminishes towards 0 as data dimension  $M$  grows. In other words, as long as the data dimension is large compared to the number of factors  $K$ , the sample trace is usually a good estimate of the true one, even when we have very limited data samples.

Now we would like to understand why UTM outperforms URM. Recall that, given an eigendecomposition  $\Sigma_{\text{SAM}} = \mathbf{B}\mathbf{S}\mathbf{B}^T$  of the sample covariance matrix, estimates generated by URM and UTM admit eigendecompositions  $\Sigma_{\text{URM}} = \mathbf{B}\mathbf{H}_{\text{URM}}\mathbf{B}^T$  and  $\Sigma_{\text{UTM}} = \mathbf{B}\mathbf{H}_{\text{UTM}}\mathbf{B}^T$ , deviating from the sample eigenvalues  $\mathbf{S}$  but not the eigenvectors  $\mathbf{B}$ . Hence, URM and UTM differ only in the way they select eigenvalues: URM takes each eigenvalue to be either a constant or the corresponding sample eigenvalue, while UTM takes each eigenvalue to be either a constant or the corresponding sample eigenvalue less another constant. Thus, large eigenvalues produced by UTM are a constant offset less than those produced by URM, as illustrated in Figure 1. We now explain why such subtraction lends UTM an advantage over URM in high-dimensional cases.

Given the eigenvectors  $\mathbf{B} = [\mathbf{b}_1 \cdots \mathbf{b}_M]$ , let us consider the optimal eigenvalues that maximize out-of-sample log-likelihood of the estimate. Specifically, let us define

$$\mathbf{H}^* \triangleq \operatorname{argmax}_{\mathbf{H} \in \mathbf{D}_+^M} L(\mathbf{B}\mathbf{H}\mathbf{B}^T, \Sigma_*).$$

With some straightforward algebra, we can show that  $\mathbf{H}^* = \operatorname{diag}(h_1^*, \dots, h_M^*)$ , where  $h_i^* = \mathbf{b}_i^T \Sigma_* \mathbf{b}_i$ , for  $i = 1, \dots, M$ . Let each  $i$ th sample eigenvalue be denoted by  $s_i = \mathbf{S}_{i,i}$ , and let the  $i$ th largest eigenvalue of  $\Sigma_*$  be denoted by  $\ell_i$ . The following theorem, whose proof relies on two results from random matrix theory found in Baik and Silverstein (2006) and Paul (2007), relates sample eigenvalues  $s_i$  to optimal eigenvalues  $h_i^*$ .

**Theorem 2** For all  $K$ , scalars  $\ell_1 > \ell_2 > \dots > \ell_K > \sigma^2 > 0$ ,  $\rho \in (0, 1)$ , sequences  $N_{(M)}$  such that  $|M/N_{(M)} - \rho| = o(1/\sqrt{N_{(M)}})$ , covariance matrices  $\Sigma_*^{(M)} \in \mathbb{S}_+^M$  with eigenvalues  $\ell_1, \dots, \ell_K, \sigma^2, \dots, \sigma^2$ , and  $i$  such that  $\ell_i > (1 + \sqrt{\rho})\sigma^2$ , there exists  $\epsilon_i \in (0, 2\sigma^2/(\ell_i - \sigma^2))$  such that  $h_i^* \xrightarrow{p} s_i - (2 + \epsilon_i)\rho\sigma^2$  as  $M \rightarrow \infty$ .

Consider eigenvalues  $\ell_i$  that are large relative to  $\sigma^2$  so that  $\epsilon_i$  is negligible. In such cases, when in the asymptotic regime identified by Theorem 2, we have  $h_i^* \approx s_i - 2\rho\sigma^2$ . This observation suggests that, when the number of factors  $K$  is relatively small compared to data dimension  $M$ , and when  $M$  and number of samples  $N$  scale proportionally to large numbers, the way in which UTM subtracts a constant from large sample eigenvalues should improve performance relative to URM, which does not modify large sample eigenvalues. Furthermore, comparing Theorem 1 and 2, we can see that the correction term should satisfy  $\frac{2\lambda}{N} \simeq 2\rho\sigma^2$ , or equivalently  $\lambda \simeq M\sigma^2$ . This relation can help us narrow the search range of  $\lambda$  in cross-validation.

It is worth pointing out that the over-shooting effect of sample eigenvalues is well known in statistics literature (see, e.g. Johnstone (2001)). Our contribution, however, is to quantify this effect for factor models, and show that the large eigenvalues are not only biased high, but biased high by the same amount.

## 5.2 Nonuniform Residual Variances

Comparing Figure 2, 3, and 4, we can see that the relation between URM and UTM is analogous to that between EM and STM. Specifically, the equivalent data requirement of UTM versus URM behaves very similarly as that of STM versus EM. This should not be surprising, as we now explain.

To develop an intuitive understanding of this phenomenon, let us consider an idealized, analytically tractable context in which both EM and STM successfully estimate the relative magnitudes of residual variances. In particular, suppose we impose an additional constraint  $\mathbf{R} \propto \mathbf{R}_*$  into (6) and an additional constraint  $\mathbf{T} \propto \mathbf{R}_*^{-\frac{1}{2}}$  into (8).<sup>3</sup> In this case, it is straightforward to show that EM is equivalent to URM with data scaled by  $\mathbf{R}_*^{-\frac{1}{2}}$ , and STM is equivalent to UTM with the same scaled data. Therefore, by the argument given in Section 5.1, it is natural to expect STM outperforms EM.

A question that remains, however, is why MRH and TM are not as effective as STM. We believe the reason to be that they tend to select factor loadings that assign larger values than appropriate to variables with large

<sup>3</sup>Here we use the notation  $\mathbf{A} \propto \mathbf{B}$  to mean that there exists  $\gamma \geq 0$  such that  $\mathbf{A} = \gamma\mathbf{B}$ .

residual variances. Indeed, such disadvantage has been observed in our synthetic data experiment: when the variation among residual variances increases, the performances of MRH and TM degrade significantly, as shown in Figure 3.

Again, let us illustrate this phenomenon through an idealized context. Specifically, consider a case in which the sample covariance matrix  $\Sigma_{\text{SAM}}$  turns out to be identical to  $\Sigma_* = \mathbf{F}_* + \mathbf{R}_*$ , with  $\mathbf{R}_* = \text{diag}(r, 1, 1, \dots, 1)$  and  $\mathbf{F}_* = \mathbf{1}\mathbf{1}^\top$ , where  $\mathbf{1}$  is a vector with every component equal to 1. Recall that MRH uses the eigenvectors of  $\Sigma_{\text{SAM}}$  corresponding to the largest eigenvalues as factor loading vectors. One would hope that factor loading estimates are insensitive to underlying residual variances. However, the following proposition suggests that, as  $r$  grows, the first component of the first eigenvector of  $\Sigma_{\text{SAM}}$  dominates other components by an unbounded ratio.

**Proposition 2** *Suppose  $\mathbf{R}_* = \text{diag}(r, 1, 1, \dots, 1)$  and  $\mathbf{F}_* = \mathbf{1}\mathbf{1}^\top$ ,  $r > 1$ . Let  $\mathbf{f} = [f_1 \ \dots \ f_M]^\top$  be the top eigenvector of  $\Sigma_*$ . Then we have  $f_1/f_i = \Omega(r), \forall i > 1$ .*

As such, the factor estimated by this top eigenvector can be grossly misrepresented, implying MRH is not preferable when residual variances differ significantly from each other.

TM suffers from a similar problem, though possibly to a lesser degree. The matrix  $\mathbf{V}$  in the TM formulation (7) represents an estimate of  $\mathbf{R}_*^{-1}$ . For simplicity, let us consider an idealized TM formulation which further incorporates a constraint  $\mathbf{V} = \mathbf{R}_*^{-1}$ . That is,

$$\begin{aligned} \max_{\mathbf{V} \in \mathbb{D}_+^M, \mathbf{G} \in \mathbb{S}_+^M} \quad & \log p(\mathcal{X}|\Sigma) - \lambda \text{tr}(\mathbf{G}) \\ \text{s.t.} \quad & \Sigma^{-1} = \mathbf{V} - \mathbf{G} \\ & \Sigma_{\text{SAM}} = \Sigma_* \\ & \mathbf{V} = \mathbf{R}_*^{-1}. \end{aligned} \tag{10}$$

Using the same setting as in Proposition 2, we can show that when this idealized TM algorithm produces an estimate of exactly one factor as desired, the first component of the estimated factor loading vector is strictly larger than the other components, as formally described in the following proposition.

**Proposition 3** *Suppose  $\mathbf{R}_*$  and  $\mathbf{F}_*$  are given as in Proposition 2, and let the estimate resulting from (10) be  $\hat{\Sigma} = \mathbf{R}_* + \hat{\mathbf{F}}$ . Then for all  $\lambda > 0$  we have*

1.  $\text{rank}(\hat{\mathbf{F}}) = 1$  if and only if  $\lambda < MN/2$ .
2. In that case, if we rewrite  $\hat{\mathbf{F}}$  as  $\mathbf{f}\mathbf{f}^\top$ , where  $\mathbf{f} = [f_1 \ \dots \ f_M]^\top$ , then  $\forall i > 1$ ,  $f_1/f_i$  is greater than 1 and monotonically increasing with  $r$ . Furthermore, if  $\lambda > (M-1)N/2$ , then  $f_1/f_i = \Omega(r)$ .

Again, this represents a bias that overemphasizes the variable with large residual variance, even when we incorporate additional information into the formulation. On the contrary, it is easy to see that STM can accurately recover all major factors if similar favorable constraints are incorporated into its formulation (ie., if we set  $\Sigma_{\text{SAM}} = \Sigma_*$  and  $\mathbf{T} \propto \mathbf{R}_*^{-\frac{1}{2}}$  in (8)).

## 6 Conclusion

We proposed factor model estimates UTM and STM, both of which are regularized versions of those that would be produced via PCA. UTM deals with contexts where residual variances are assumed to be uniform, whereas STM handles variation among residual variances. Our algorithm for computing the UTM estimate is as efficient as conventional PCA. For STM, we provide an iterative algorithm with guaranteed convergence. Computational experiments involving both synthetic and real data demonstrate that the estimates produced

by our approach are significantly more accurate than those produced by pre-existing methods. Further, we provide a theoretical analysis that elucidates the way in which UTM and STM corrects biases induced by alternative approaches.

Let us close by mentioning a few possible directions for further research. Our analysis has relied on data being generated by a Gaussian distribution. It would be useful to understand how things change if this assumption is relaxed. Further, in practice estimates are often used to guide subsequent decisions. It would be interesting to study the impact of STM on decision quality and whether there are other approaches that fare better in this dimension. Our recent paper on directed principle component analysis (Kao and Van Roy, 2012) relates to this. In some applications, PCA is used to identify a subspace for dimension reduction. It would be interesting to understand if and when the subspace identified by STM is more suitable. Finally, there is a growing body of research on robust variations of factor analysis and PCA. These include the pursuit of sparse factor loadings (Jolliffe et al., 2003; Zou et al., 2006; D’Aspremont et al., 2004; Johnstone and Lu, 2009; Amini and Wainwright, 2009), and the methods that are resistant to corrupted data (Pison et al., 2003; Candès et al., 2009; Xu et al., 2010). It would be interesting to explore connections to this body of work.

**Acknowledgements** This work was supported in part by Award CMMI-0968707 from the National Science Foundation.

## A Proofs

We first prove a main lemma that will be used in the proof of Theorem 1 and Proposition 3.

**Lemma 1** Fixing  $\mathbf{V} \in \mathbb{D}_{++}^M$ , consider the optimization problem

$$\begin{aligned} \max_{\mathbf{G} \in \mathbb{S}_+^M} \quad & \log p(\mathcal{X}|\boldsymbol{\Sigma}) - \lambda \text{tr}(\mathbf{G}) \\ \text{s.t.} \quad & \boldsymbol{\Sigma}^{-1} = \mathbf{V} - \mathbf{G}. \end{aligned} \quad (11)$$

Let  $\mathbf{G}_{\mathbf{V}}$  be the solution to (11),  $\lambda' = 2\lambda/N$ , and  $\mathbf{U}\mathbf{D}\mathbf{U}^T$  be an eigendecomposition of matrix  $\mathbf{V}^{\frac{1}{2}}(\boldsymbol{\Sigma}_{\text{SAM}} - \lambda'\mathbf{I})\mathbf{V}^{\frac{1}{2}}$  with  $\mathbf{U}$  orthonormal. Then we have  $(\mathbf{V} - \mathbf{G}_{\mathbf{V}})^{-1} = \mathbf{V}^{-\frac{1}{2}}\mathbf{U}\mathbf{L}\mathbf{U}^T\mathbf{V}^{-\frac{1}{2}}$ , where  $\mathbf{L}$  is a diagonal matrix with entries  $\mathbf{L}_{i,i} = \max\{\mathbf{D}_{i,i}, 1\}$ ,  $\forall i = 1, \dots, M$ .

**Proof** We can rewrite (11) as

$$\begin{aligned} \min_{\mathbf{G}} \quad & -\log \det(\mathbf{V} - \mathbf{G}) + \text{tr}((\mathbf{V} - \mathbf{G})\boldsymbol{\Sigma}_{\text{SAM}}) + \lambda' \text{tr}(\mathbf{G}) \\ \text{s.t.} \quad & \mathbf{G} \in \mathbb{S}_+^M \end{aligned}$$

Now associate a Lagrange multiplier  $\boldsymbol{\Omega} \in \mathbb{S}_+^M$  with the  $\mathbf{G} \in \mathbb{S}_+^M$  constraint and write down the Lagrangian as

$$\mathcal{L}(\mathbf{G}, \boldsymbol{\Omega}) = -\log \det(\mathbf{V} - \mathbf{G}) + \text{tr}((\mathbf{V} - \mathbf{G})\boldsymbol{\Sigma}_{\text{SAM}}) + \lambda' \text{tr}(\mathbf{G}) - \text{tr}(\boldsymbol{\Omega}\mathbf{G}).$$

Let  $\boldsymbol{\Omega}^*$  denote the dual solution. By KKT conditions we have:

$$\nabla_{\mathbf{G}} \mathcal{L} \Big|_{\mathbf{G}_{\mathbf{V}}, \boldsymbol{\Omega}^*} = (\mathbf{V} - \mathbf{G}_{\mathbf{V}})^{-1} - \boldsymbol{\Sigma}_{\text{SAM}} + \lambda'\mathbf{I} - \boldsymbol{\Omega}^* = 0 \quad (12)$$

$$\boldsymbol{\Omega}^*, \mathbf{G}_{\mathbf{V}} \in \mathbb{S}_+^M \quad (13)$$

$$\text{tr}(\boldsymbol{\Omega}^* \mathbf{G}_{\mathbf{V}}) = 0. \quad (14)$$

Recall that

$$(\mathbf{V} - \mathbf{G}_{\mathbf{V}})^{-1} = (\mathbf{V}^{\frac{1}{2}}\mathbf{V}^{\frac{1}{2}} - \mathbf{G}_{\mathbf{V}})^{-1} = \mathbf{V}^{-\frac{1}{2}}(\mathbf{I} - \mathbf{V}^{-\frac{1}{2}}\mathbf{G}_{\mathbf{V}}\mathbf{V}^{-\frac{1}{2}})^{-1}\mathbf{V}^{-\frac{1}{2}}. \quad (15)$$

Plugging this into (12) we get

$$\mathbf{V}^{-\frac{1}{2}}(\mathbf{I} - \mathbf{V}^{-\frac{1}{2}}\mathbf{G}_V\mathbf{V}^{-\frac{1}{2}})^{-1}\mathbf{V}^{-\frac{1}{2}} = \boldsymbol{\Sigma}_{\text{SAM}} - \lambda'\mathbf{I} + \boldsymbol{\Omega}^*$$

and so

$$\begin{aligned} \left(\mathbf{I} - \mathbf{V}^{-\frac{1}{2}}\mathbf{G}_V\mathbf{V}^{-\frac{1}{2}}\right)^{-1} &= \mathbf{V}^{\frac{1}{2}}(\boldsymbol{\Sigma}_{\text{SAM}} - \lambda'\mathbf{I} + \boldsymbol{\Omega}^*)\mathbf{V}^{\frac{1}{2}} \\ &= \mathbf{V}^{\frac{1}{2}}(\boldsymbol{\Sigma}_{\text{SAM}} - \lambda'\mathbf{I})\mathbf{V}^{\frac{1}{2}} + \mathbf{V}^{\frac{1}{2}}\boldsymbol{\Omega}^*\mathbf{V}^{\frac{1}{2}}. \end{aligned} \quad (16)$$

By (13), both  $\mathbf{V}^{-\frac{1}{2}}\mathbf{G}_V\mathbf{V}^{-\frac{1}{2}}$  and  $\mathbf{V}^{\frac{1}{2}}\boldsymbol{\Omega}^*\mathbf{V}^{\frac{1}{2}}$  are in  $\mathbb{S}_+^M$ . Let an eigendecomposition of  $\mathbf{V}^{-\frac{1}{2}}\mathbf{G}_V\mathbf{V}^{-\frac{1}{2}}$  be  $\mathbf{A}\mathbf{Q}\mathbf{A}^T$  for which  $\mathbf{A} = [\mathbf{a}_1 \ \dots \ \mathbf{a}_M]$  is orthonormal and  $\mathbf{Q}_{i,i} \geq 0, i = 1, \dots, M$ . Using (14) and the fact that trace is invariant under cyclic permutations, we have

$$\begin{aligned} 0 &= \text{tr}(\boldsymbol{\Omega}^*\mathbf{G}_V) = \text{tr}(\boldsymbol{\Omega}^*\mathbf{V}^{\frac{1}{2}}\mathbf{V}^{-\frac{1}{2}}\mathbf{G}_V\mathbf{V}^{-\frac{1}{2}}\mathbf{V}^{\frac{1}{2}}) \\ &= \text{tr}\left(\left(\mathbf{V}^{\frac{1}{2}}\boldsymbol{\Omega}^*\mathbf{V}^{\frac{1}{2}}\right)\left(\mathbf{V}^{-\frac{1}{2}}\mathbf{G}_V\mathbf{V}^{-\frac{1}{2}}\right)\right) = \text{tr}\left(\left(\mathbf{V}^{\frac{1}{2}}\boldsymbol{\Omega}^*\mathbf{V}^{\frac{1}{2}}\right)\mathbf{A}\mathbf{Q}\mathbf{A}^T\right) \\ &= \text{tr}\left(\mathbf{A}^T\left(\mathbf{V}^{\frac{1}{2}}\boldsymbol{\Omega}^*\mathbf{V}^{\frac{1}{2}}\right)\mathbf{A}\mathbf{Q}\right) = \sum_{i=1}^M \mathbf{Q}_{i,i}\mathbf{a}_i^T\left(\mathbf{V}^{\frac{1}{2}}\boldsymbol{\Omega}^*\mathbf{V}^{\frac{1}{2}}\right)\mathbf{a}_i. \end{aligned}$$

Since  $\mathbf{Q}_{i,i} \geq 0$  and  $\mathbf{V}^{\frac{1}{2}}\boldsymbol{\Omega}^*\mathbf{V}^{\frac{1}{2}} \in \mathbb{S}_+^M$ , we can deduce

$$\mathbf{a}_i^T\left(\mathbf{V}^{\frac{1}{2}}\boldsymbol{\Omega}^*\mathbf{V}^{\frac{1}{2}}\right)\mathbf{a}_i = 0 \text{ if } \mathbf{Q}_{i,i} > 0, \quad \forall i = 1, \dots, M.$$

Let  $\mathcal{I}_+ = \{i : \mathbf{Q}_{i,i} > 0\}$ . Because  $\mathbf{V}^{\frac{1}{2}}\boldsymbol{\Omega}^*\mathbf{V}^{\frac{1}{2}}$  is positive semidefinite, for all  $i_0 \in \mathcal{I}_+$  we also have  $\mathbf{V}^{\frac{1}{2}}\boldsymbol{\Omega}^*\mathbf{V}^{\frac{1}{2}}\mathbf{a}_{i_0} = 0$ . Furthermore, since

$$\left(\mathbf{I} - \mathbf{V}^{-\frac{1}{2}}\mathbf{G}_V\mathbf{V}^{-\frac{1}{2}}\right)^{-1} = \mathbf{A}\text{diag}\left(\frac{1}{1 - \mathbf{Q}_{1,1}}, \dots, \frac{1}{1 - \mathbf{Q}_{M,M}}\right)\mathbf{A}^T$$

multiplying both sides of (16) by  $\mathbf{a}_{i_0}$  leads to

$$\frac{\mathbf{a}_{i_0}}{1 - \mathbf{Q}_{i_0,i_0}} = \mathbf{V}^{\frac{1}{2}}(\boldsymbol{\Sigma}_{\text{SAM}} - \lambda'\mathbf{I})\mathbf{V}^{\frac{1}{2}}\mathbf{a}_{i_0}$$

which shows  $\mathbf{a}_{i_0}$  is an eigenvector of  $\mathbf{V}^{\frac{1}{2}}(\boldsymbol{\Sigma}_{\text{SAM}} - \lambda'\mathbf{I})\mathbf{V}^{\frac{1}{2}}$ . Recall that  $\mathbf{V}^{\frac{1}{2}}(\boldsymbol{\Sigma}_{\text{SAM}} - \lambda'\mathbf{I})\mathbf{V}^{\frac{1}{2}} = \mathbf{U}\mathbf{D}\mathbf{U}^T$ . Without loss of generality, we can take  $\mathbf{a}_i = \mathbf{u}_i$  for all  $i \in \mathcal{I}_+$ . Indeed, since  $\mathcal{I}_+ = \{i : \mathbf{Q}_{i,i} > 0\}$ , we have

$$\mathbf{A}\mathbf{Q}\mathbf{A}^T = \sum_{i=1}^M \mathbf{Q}_{i,i}\mathbf{a}_i\mathbf{a}_i^T = \sum_{i \in \mathcal{I}_+} \mathbf{Q}_{i,i}\mathbf{a}_i\mathbf{a}_i^T + \sum_{j \notin \mathcal{I}_+} 0 \cdot \mathbf{a}_j\mathbf{a}_j^T = \sum_{i \in \mathcal{I}_+} \mathbf{Q}_{i,i}\mathbf{u}_i\mathbf{u}_i^T + \sum_{j \notin \mathcal{I}_+} 0 \cdot \mathbf{u}_j\mathbf{u}_j^T$$

which implies we can further take  $\mathbf{A} = \mathbf{U}$ . This gives

$$\left(\mathbf{I} - \mathbf{V}^{-\frac{1}{2}}\mathbf{G}_V\mathbf{V}^{-\frac{1}{2}}\right)^{-1} = \mathbf{U}\mathbf{L}\mathbf{U}^T, \quad (17)$$

where  $\mathbf{L}$  is a diagonal matrix with entries  $\mathbf{L}_{i,i} = \frac{1}{1 - \mathbf{Q}_{i,i}}$ , for  $i = 1, \dots, M$ . Plugging this into (16) we have

$$\mathbf{U}\mathbf{L}\mathbf{U}^T = \mathbf{U}\mathbf{D}\mathbf{U}^T + \mathbf{V}^{\frac{1}{2}}\boldsymbol{\Omega}^*\mathbf{V}^{\frac{1}{2}}. \quad (18)$$

For any  $i_0 \in \mathcal{I}_+$ , multiplying the both sides of the above equation by  $\mathbf{u}_{i_0}$  results in

$$\mathbf{L}_{i_0,i_0}\mathbf{u}_{i_0} = \mathbf{D}_{i_0,i_0}\mathbf{u}_{i_0} + 0$$



which implies  $\mathbf{L}_{i_0, i_0} = \mathbf{D}_{i_0, i_0}$ , or more generally

$$\mathbf{L}_{i,i} = \begin{cases} \mathbf{D}_{i,i} & \text{if } \mathbf{Q}_{i,i} > 0 \\ 1 & \text{otherwise} \end{cases}, \quad i = 1, \dots, M.$$

Since  $\mathbf{L}_{i,i} = \frac{1}{1-\mathbf{Q}_{i,i}} \geq 1$ , to see  $\mathbf{L}_{i,i} = \max\{\mathbf{D}_{i,i}, 1\}$ , it remains to show  $\mathbf{L}_{i,i} \geq \mathbf{D}_{i,i}$  for all  $i$ . This follows by rearranging (18)

$$\begin{aligned} \mathbf{ULU}^T - \mathbf{UDU}^T &= \mathbf{V}^{\frac{1}{2}} \mathbf{\Omega}^* \mathbf{V}^{\frac{1}{2}} \succeq 0 \\ \Rightarrow \mathbf{L} - \mathbf{D} &\succeq 0 \Rightarrow \mathbf{L}_{i,i} \geq \mathbf{D}_{i,i}, \quad \forall i = 1, \dots, M. \end{aligned}$$

Finally, plugging (17) into (15) completes the proof.  $\blacksquare$

**Theorem 1**  $\Sigma_{\text{SAM}}$  and  $\Sigma_{\text{UTM}}^\lambda$  share the same trace and eigenvectors, and letting the eigenvalues of the two matrices, sorted in decreasing order, be denoted by  $s_1, \dots, s_M$  and  $h_1, \dots, h_M$ , respectively, we have

$$h_m = \max \left\{ s_m - \frac{2\lambda}{N}, \frac{1}{\hat{v}} \right\}, \quad \text{for } m = 1, \dots, M.$$

**Proof** Let  $\mathbf{U} \text{diag}(s_1, \dots, s_M) \mathbf{U}^T$  be an eigendecomposition of  $\Sigma_{\text{SAM}}$  such that  $\mathbf{U}$  is orthonormal. Define  $\mathbf{V} = \hat{v} \mathbf{I}$ , and note that an eigendecomposition of matrix

$$\mathbf{V}^{\frac{1}{2}} \left( \Sigma_{\text{SAM}} - \frac{2\lambda}{N} \mathbf{I} \right) \mathbf{V}^{\frac{1}{2}} = \hat{v} \left( \Sigma_{\text{SAM}} - \frac{2\lambda}{N} \mathbf{I} \right)$$

can be written as  $\mathbf{UDU}^T$ , where  $\mathbf{D}_{i,i} = \hat{v}(s_i - 2\lambda/N)$ ,  $i = 1, \dots, M$ . By Lemma 1 we have

$$\Sigma_{\text{UTM}}^\lambda = (\hat{v} \mathbf{I} - \hat{\mathbf{G}})^{-1} = \mathbf{V}^{-\frac{1}{2}} \mathbf{ULU}^T \mathbf{V}^{-\frac{1}{2}} = \frac{1}{\hat{v}} \mathbf{ULU}^T$$

where  $\mathbf{L} \in \mathbb{D}_+^M$ , and  $\mathbf{L}_{i,i} = \max\{\mathbf{D}_{i,i}, 1\} = \max\{\hat{v}(s_i - 2\lambda/N), 1\} = \hat{v} \max\{s_i - 2\lambda/N, 1/\hat{v}\} = \hat{v} h_i$ ,  $\forall i = 1, \dots, M$ . Therefore,

$$\Sigma_{\text{UTM}}^\lambda = \frac{1}{\hat{v}} \mathbf{ULU}^T = \mathbf{UHU}^T.$$

where  $\mathbf{H} = \text{diag}(h_1, \dots, h_M)$ , as desired.

Furthermore, recall that we impose no constraint on  $v$  when solving UTM, and as a result the partial derivative of the objective function with respect to  $v$  should vanish at  $\hat{v}$ . That is,

$$\frac{\partial}{\partial v} (\log p(\mathcal{X}|\Sigma) - \lambda \text{tr}(\mathbf{G})) \Big|_{\hat{\mathbf{G}}, \hat{v}} = -\frac{N}{2} \left( \text{tr}(\Sigma_{\text{SAM}}) - \text{tr}((\hat{v} \mathbf{I} - \hat{\mathbf{G}})^{-1}) \right) = 0$$

which implies  $\text{tr}(\Sigma_{\text{SAM}}) = \text{tr}((\hat{v} \mathbf{I} - \hat{\mathbf{G}})^{-1}) = \text{tr}(\Sigma_{\text{UTM}}^\lambda)$ .  $\blacksquare$

**Theorem 2** For all  $K$ , scalars  $\ell_1 > \ell_2 > \dots > \ell_K > \sigma^2 > 0$ ,  $\rho \in (0, 1)$ , sequences  $N_{(M)}$  such that  $|M/N_{(M)} - \rho| = o(1/\sqrt{N_{(M)}})$ , covariance matrices  $\Sigma_*^{(M)} \in \mathbb{S}_+^M$  with eigenvalues  $\ell_1, \dots, \ell_K, \sigma^2, \dots, \sigma^2$ , and  $i$  such that  $\ell_i > (1 + \sqrt{\rho})\sigma^2$ , there exists  $\epsilon_i \in (0, 2\sigma^2/(\ell_i - \sigma^2))$  such that  $h_i^* \xrightarrow{P} s_i - (2 + \epsilon_i)\rho\sigma^2$ .

**Proof** Let  $\mathbf{A}\mathbf{L}\mathbf{A}^T$  be an eigendecomposition of  $\Sigma_*$ , where  $\mathbf{A} = [\mathbf{a}_1 \ \dots \ \mathbf{a}_M]$  is orthonormal and  $\mathbf{L} = \text{diag}(\ell_1, \dots, \ell_K, \sigma^2, \dots, \sigma^2)$ . Recall that

$$\begin{aligned} h_i^* &= \mathbf{b}_i^T \Sigma_* \mathbf{b}_i = \mathbf{b}_i^T \left( \sum_{j=1}^K \mathbf{a}_j \ell_j \mathbf{a}_j^T + \sum_{j=K+1}^M \mathbf{a}_j \sigma^2 \mathbf{a}_j^T \right) \mathbf{b}_i \\ &= \sum_{j=1}^K \ell_j (\mathbf{b}_i^T \mathbf{a}_j)^2 + \sigma^2 \sum_{j=K+1}^M (\mathbf{b}_i^T \mathbf{a}_j)^2. \end{aligned} \quad (19)$$

Using Theorem 4 in Paul (2007), we have

$$(\mathbf{b}_i^T \mathbf{a}_i)^2 \xrightarrow{\text{a.s.}} \left( 1 - \frac{\rho\sigma^4}{(\ell_i - \sigma^2)^2} \right) / \left( 1 + \frac{\rho\sigma^2}{\ell_i - \sigma^2} \right). \quad (20)$$

Furthermore, decomposing  $\mathbf{b}_i$  into  $\tilde{\mathbf{b}}_i + \tilde{\mathbf{b}}_i^\perp$ , where  $\tilde{\mathbf{b}}_i \in \text{span}(\mathbf{a}_1, \dots, \mathbf{a}_K)$  and  $\tilde{\mathbf{b}}_i^\perp \in \text{span}(\mathbf{a}_{K+1}, \dots, \mathbf{a}_M)$ , by Theorem 5 in Paul (2007) we have  $\frac{\tilde{\mathbf{b}}_i}{\|\tilde{\mathbf{b}}_i\|} \xrightarrow{\text{P}} \mathbf{a}_i$ . This implies if  $1 \leq j \leq K, j \neq i$ ,

$$\mathbf{b}_i^T \mathbf{a}_j = \tilde{\mathbf{b}}_i^T \mathbf{a}_j \xrightarrow{\text{P}} \|\tilde{\mathbf{b}}_i\| \mathbf{a}_i^T \mathbf{a}_j = 0 \quad (21)$$

and for  $K < j \leq M$ ,

$$\sum_{j=K+1}^M (\mathbf{b}_i^T \mathbf{a}_j)^2 = 1 - \left\| \sum_{j=1}^K \mathbf{b}_i^T \mathbf{a}_j \right\|^2 \xrightarrow{\text{P}} 1 - (\mathbf{b}_i^T \mathbf{a}_i)^2. \quad (22)$$

Plugging (20), (21), and (22) into (19) we arrive at

$$\begin{aligned} h_i^* &\xrightarrow{\text{P}} \ell_i \left( 1 - \frac{\rho\sigma^4}{(\ell_i - \sigma^2)^2} \right) / \left( 1 + \frac{\rho\sigma^2}{\ell_i - \sigma^2} \right) + \sigma^2 \left( 1 - \left( 1 - \frac{\rho\sigma^4}{(\ell_i - \sigma^2)^2} \right) / \left( 1 + \frac{\rho\sigma^2}{\ell_i - \sigma^2} \right) \right) \\ &= \frac{\ell_i}{1 + \rho\sigma^2/(\ell_i - \sigma^2)} \\ &= \ell_i - \left( 1 + \frac{(1-\rho)\sigma^2}{\ell_i - (1-\rho)\sigma^2} \right) \rho\sigma^2. \end{aligned} \quad (23)$$

By Theorem 1.1 in Baik and Silverstein (2006) we have

$$s_i \xrightarrow{\text{a.s.}} \ell_i + \frac{\rho\ell_i\sigma^2}{\ell_i - \sigma^2} = \ell_i + \left( 1 + \frac{\sigma^2}{\ell_i - \sigma^2} \right) \rho\sigma^2. \quad (24)$$

Finally, combining (23) and (24) yields

$$h_i^* \xrightarrow{\text{P}} s_i - (2 + \epsilon_i) \rho\sigma^2$$

where  $\epsilon_i = \frac{(1-\rho)\sigma^2}{\ell_i - (1-\rho)\sigma^2} + \frac{\sigma^2}{\ell_i - \sigma^2}$ . It is easy to see  $0 < \epsilon_i < \frac{\sigma^2}{\ell_i - \sigma^2} + \frac{\sigma^2}{\ell_i - \sigma^2}$ , as desired.  $\blacksquare$

**Proposition 1** For any fixed  $N, K$ , scalars  $\ell_1 \geq \ell_2 \geq \dots \geq \ell_K \geq \sigma^2 > 0$ , and any sequence of covariance matrices  $\Sigma_*^{(M)} \in \mathbb{S}_+^M$  with eigenvalues  $\ell_1, \dots, \ell_K, \sigma^2, \dots, \sigma^2$ , we have

$$\frac{\text{tr} \Sigma_{\text{SAM}}}{\text{tr} \Sigma_*} \xrightarrow{\text{a.s.}} 1$$

as  $M \rightarrow \infty$  and

$$\Pr \left( \left| \frac{\text{tr} \Sigma_{\text{SAM}}}{\text{tr} \Sigma_*} - 1 \right| \geq \epsilon \right) \leq 2 \exp(-N\epsilon^2\Omega(M)).$$

**Proof** Let  $\mathbf{A}\mathbf{L}\mathbf{A}^T$  be an eigendecomposition of  $\Sigma_*$ , where  $\mathbf{A} = [\mathbf{a}_1 \ \dots \ \mathbf{a}_M]$  is orthonormal and  $\mathbf{L} = \text{diag}(\ell_1, \dots, \ell_K, \sigma^2, \dots, \sigma^2)$ . Thus,  $\text{tr}\Sigma_* = \text{tr}\mathbf{L} = \sum_{k=1}^K \ell_k + (M - K)\sigma^2$ . Note that

$$\text{tr}\Sigma_{\text{SAM}} = \text{tr}(\mathbf{A}\mathbf{A}^T\Sigma_{\text{SAM}}) = \text{tr}(\mathbf{A}^T\Sigma_{\text{SAM}}\mathbf{A}) = \sum_{i=1}^M \mathbf{a}_i^T \Sigma_{\text{SAM}} \mathbf{a}_i = \sum_{i=1}^M \frac{1}{N} \sum_{n=1}^N (\mathbf{a}_i^T \mathbf{x}_{(n)})^2.$$

Since  $\mathbf{x}_{(n)} \sim \mathcal{N}(0, \Sigma_*)$ , we can think of each  $\mathbf{x}_{(n)}$  as generated by  $\mathbf{x}_{(n)} = \mathbf{A}\mathbf{z}_{(n)}$ , where  $\mathbf{z}_{(n)}$  is sampled iid from  $\mathcal{N}(0, \mathbf{L})$ . This leads to

$$\text{tr}\Sigma_{\text{SAM}} = \sum_{i=1}^M \frac{1}{N} \sum_{n=1}^N \mathbf{z}_{(n),i}^2 = \sum_{i=1}^M \frac{\mathbf{L}_{i,i}}{N} w_i^2,$$

where  $w_i^2$ 's are i.i.d. samples from  $\chi_N^2$ . Therefore,

$$\lim_{M \rightarrow \infty} \frac{\text{tr}\Sigma_{\text{SAM}}}{\text{tr}\Sigma_*} = \lim_{M \rightarrow \infty} \frac{\sum_{i=1}^M \frac{\mathbf{L}_{i,i}}{N} w_i^2}{\sum_{k=1}^K \ell_k + (M - K)\sigma^2} = \lim_{M \rightarrow \infty} \frac{\sum_{k=1}^K \ell_k \frac{w_i^2}{N} + \sum_{i=K+1}^M \sigma^2 \frac{w_i^2}{N}}{\sum_{k=1}^K \ell_k + (M - K)\sigma^2}.$$

Since the first terms in the denominator and the numerator are bounded and do not scale with  $K$ , we can drop them in the limit and rewrite

$$\lim_{M \rightarrow \infty} \frac{\text{tr}\Sigma_{\text{SAM}}}{\text{tr}\Sigma_*} = \lim_{M \rightarrow \infty} \frac{\sum_{i=K+1}^M \sigma^2 \frac{w_i^2}{N}}{(M - K)\sigma^2} = \lim_{M \rightarrow \infty} \frac{1}{M - K} \sum_{i=K+1}^M \frac{w_i^2}{N} = 1 \text{ (w. p. 1)}$$

due to the strong law of large numbers and the fact that  $\mathbb{E}[w_i^2] = N$ .

To prove the second part of the proposition, let us rewrite  $w_i^2 = \sum_{n=1}^N \tilde{w}_{i,n}^2$ , where  $\tilde{w}_{i,n}$  are i.i.d samples from  $\mathcal{N}(0, 1)$ . Therefore,

$$\text{tr}\Sigma_{\text{SAM}} - \text{tr}\Sigma_* = \sum_{i=1}^M \sum_{n=1}^N \frac{\mathbf{L}_{i,i}}{N} \tilde{w}_{i,n}^2 - \sum_{i=1}^M \mathbf{L}_{i,i} = \sum_{i=1}^M \sum_{n=1}^N \frac{\mathbf{L}_{i,i}}{N} (\tilde{w}_{i,n}^2 - 1).$$

By the exponential inequality for chi-square distributions (Laurent and Massart, 2000), we have

$$\Pr(|\text{tr}\Sigma_{\text{SAM}} - \text{tr}\Sigma_*| \geq 2\xi\sqrt{\tau}) \leq 2\exp(-\tau), \quad \forall \tau > 0,$$

where  $\xi = \sqrt{\sum_{i=1}^M \sum_{n=1}^N \left(\frac{\mathbf{L}_{i,i}}{N}\right)^2} = \sqrt{\frac{1}{N} \sum_{i=1}^M \mathbf{L}_{i,i}^2}$ . Taking  $\tau = \left(\frac{\epsilon \text{tr}\Sigma_*}{2\xi}\right)^2$ , we can rewrite the above inequality as

$$\Pr\left(\left|\frac{\text{tr}\Sigma_{\text{SAM}}}{\text{tr}\Sigma_*} - 1\right| \geq \epsilon\right) \leq 2\exp\left(-\left(\frac{\epsilon \text{tr}\Sigma_*}{2\xi}\right)^2\right) = 2\exp\left(-N\epsilon^2 \frac{\left(\sum_{i=1}^M \mathbf{L}_{i,i}\right)^2}{4 \sum_{i=1}^M \mathbf{L}_{i,i}^2}\right).$$

The desired result then follows straightforwardly from the fact that

$$\frac{\left(\sum_{i=1}^M \mathbf{L}_{i,i}\right)^2}{\sum_{i=1}^M \mathbf{L}_{i,i}^2} = \frac{(M\sigma^2 + \kappa_1)^2}{M\sigma^4 + \kappa_2} = \Omega(M),$$

since  $\kappa_1 = \sum_{i=1}^K (\ell_i - \sigma^2)$  and  $\kappa_2 = \sum_{i=1}^K (\ell_i^2 - \sigma^4)$  are both constants. ■

**Proposition 2** Suppose  $\mathbf{R}_* = \text{diag}(r, 1, 1, \dots, 1)$  and  $\mathbf{F}_* = \mathbf{1}\mathbf{1}^\top$ ,  $r > 1$ . Let  $\mathbf{f} = [f_1 \ \dots \ f_M]^\top$  be the top eigenvector of  $\Sigma_*$ . Then we have  $f_1/f_i = \Omega(r), \forall i > 1$ .

**Proof** Note that  $f$  is the solution to the following optimization problem

$$\begin{aligned} \max_{\mathbf{u} \in \mathbb{R}^M} \quad & \mathbf{u}^\top (\mathbf{R}_* + \mathbf{F}_*) \mathbf{u} \\ \text{s.t.} \quad & \|\mathbf{u}\|_2 = 1 \end{aligned}$$

and the objective function can be written as  $r f_1^2 + \sum_{i=2}^M f_i^2 + \left(\sum_{i=1}^M f_i\right)^2$ . By symmetry, we have  $f_2 = f_3 = \dots = f_M$ . To simplify notation, let us represent  $\mathbf{f}$  as  $[x \ y \ y \ \dots \ y]^\top$ . Suppose the largest eigenvalue is  $q$ . By definition we have  $(\mathbf{R}_* + \mathbf{F}_*)\mathbf{f} = q\mathbf{f}$ , or equivalently

$$\begin{aligned} (r+1)x + (M-1)y &= qx \\ x + My &= qy. \end{aligned}$$

Solving the above equations leads to

$$q = \frac{1}{2} \left( M+r+1 + \sqrt{(M+r+1)^2 - 4(Mr+1)} \right) = \Omega(r),$$

and plugging this back to the above equations yields

$$x/y = q - M = \Omega(r).$$

■

**Proposition 3** Suppose  $\mathbf{R}_*$  and  $\mathbf{F}_*$  are given as in Proposition 2, and let the estimate resulting from (10) be  $\hat{\Sigma} = \mathbf{R}_* + \hat{\mathbf{F}}$ . Then for all  $\lambda > 0$  we have

1.  $\text{rank}(\hat{\mathbf{F}}) = 1$  if and only if  $\lambda < MN/2$ .
2. In that case, if we rewrite  $\hat{\mathbf{F}}$  as  $\mathbf{f}\mathbf{f}^\top$ , where  $\mathbf{f} = [f_1 \ \dots \ f_M]^\top$ , then  $\forall i > 1$ ,  $f_1/f_i$  is greater than 1 and monotonically increasing with  $r$ . Furthermore, if  $\lambda > (M-1)N/2$ , then  $f_1/f_i = \Omega(r)$ .

**Proof** Let  $\lambda' = 2\lambda/N$ ,  $\mathbf{C} = \mathbf{R}_*^{-\frac{1}{2}} (\Sigma_* - \lambda'\mathbf{I}) \mathbf{R}_*^{-\frac{1}{2}}$ ,  $\mathbf{U}\mathbf{D}\mathbf{U}^\top$  be an eigendecomposition of  $\mathbf{C}$ , and  $\mathbf{L}$  be the diagonal matrix with entries  $\mathbf{L}_{i,i} = \max\{\mathbf{D}_{i,i}, 1\}, \forall i = 1, \dots, M$ . Applying Lemma 1 with  $\Sigma_{\text{SAM}} = \Sigma_*$  and  $\mathbf{V} = \mathbf{R}_*^{-1}$  we have

$$\hat{\Sigma} = \mathbf{R}_*^{\frac{1}{2}} \mathbf{U} \mathbf{L} \mathbf{U}^\top \mathbf{R}_*^{\frac{1}{2}} = \mathbf{R}_*^{\frac{1}{2}} \mathbf{U} \mathbf{I} \mathbf{U}^\top \mathbf{R}_*^{\frac{1}{2}} + \mathbf{R}_*^{\frac{1}{2}} \mathbf{U} (\mathbf{L} - \mathbf{I}) \mathbf{U}^\top \mathbf{R}_*^{\frac{1}{2}} = \mathbf{R}_* + \mathbf{R}_*^{\frac{1}{2}} \mathbf{U} (\mathbf{L} - \mathbf{I}) \mathbf{U}^\top \mathbf{R}_*^{\frac{1}{2}}$$

and therefore

$$\hat{\mathbf{F}} = \mathbf{R}_*^{\frac{1}{2}} \mathbf{U} (\mathbf{L} - \mathbf{I}) \mathbf{U}^\top \mathbf{R}_*^{\frac{1}{2}}.$$

Since  $\mathbf{L}_{i,i} = \max\{\mathbf{D}_{i,i}, 1\}$ , we further have  $\text{rank}(\hat{\mathbf{F}}) = \text{rank}(\mathbf{L} - \mathbf{I}) = |\{i : \mathbf{D}_{i,i} > 1\}|$ . Recall that  $\mathbf{D}$  denotes the eigenvalues of matrix  $\mathbf{C}$ , which can be written as

$$\begin{aligned} \mathbf{R}_*^{-\frac{1}{2}} (\Sigma_* - \lambda'\mathbf{I}) \mathbf{R}_*^{-\frac{1}{2}} &= \mathbf{R}_*^{-\frac{1}{2}} \Sigma_* \mathbf{R}_*^{-\frac{1}{2}} - \lambda' \mathbf{R}_*^{-1} \\ &= \mathbf{R}_*^{-\frac{1}{2}} (\mathbf{R}_* + \mathbf{F}_*) \mathbf{R}_*^{-\frac{1}{2}} - \lambda' \left( \mathbf{I} - \left(1 - \frac{1}{r}\right) \mathbf{e}_1 \mathbf{e}_1^\top \right) \\ &= \mathbf{I} + \mathbf{a} \mathbf{a}^\top - \lambda' \mathbf{I} + \lambda' \left(1 - \frac{1}{r}\right) \mathbf{e}_1 \mathbf{e}_1^\top \\ &= (1 - \lambda') \mathbf{I} + \mathbf{a} \mathbf{a}^\top + \lambda' \left(1 - \frac{1}{r}\right) \mathbf{e}_1 \mathbf{e}_1^\top \end{aligned}$$

where  $\mathbf{a} = \left[ \frac{1}{\sqrt{r}} \ 1 \ \dots \ 1 \right]^T$  and  $\mathbf{e}_1 = [1 \ 0 \ \dots \ 0]^T$ . Let  $\mathbf{A} = \mathbf{a}\mathbf{a}^T + \lambda' \left(1 - \frac{1}{r}\right) \mathbf{e}_1\mathbf{e}_1^T$ . Since  $\mathbf{C} = (1 - \lambda')\mathbf{I} + \mathbf{A}$ , we know  $\mathbf{C}$  and  $\mathbf{A}$  share the same eigenvectors, and the corresponding eigenvalues differ by  $(1 - \lambda')$ . Thus, the number of the eigenvalues of  $\mathbf{C}$  that are greater than 1 is equal to the number of the eigenvalues of  $\mathbf{A}$  that are greater than  $\lambda'$ . However,  $\text{rank}(\mathbf{A}) = \text{rank}(\mathbf{a}\mathbf{a}^T + \lambda' \left(1 - \frac{1}{r}\right) \mathbf{e}_1\mathbf{e}_1^T) = 2$ , which implies  $\mathbf{A}$  has only 2 non-zero eigenvalues. Let  $q$  be one of them. By symmetry, we can denote the corresponding eigenvector by  $\mathbf{u} = [x \ y \ \dots \ y]^T$ . Then we have  $\mathbf{A}\mathbf{u} = q\mathbf{u}$ , which leads to

$$\begin{aligned} \frac{x}{r} + \frac{(M-1)y}{\sqrt{r}} + \left(\lambda' - \frac{\lambda'}{r}\right)x &= qx \\ \frac{x}{\sqrt{r}} + (M-1)y &= qy. \end{aligned} \quad (25)$$

After eliminating  $x$  and  $y$  we arrive at the following equation

$$rq^2 - ((M-1)r + (r-1)\lambda' + 1)q + \lambda'(M-1)(r-1) = 0.$$

Let  $s(q)$  denote the left-hand-side of the above equation. It is easy to see that its discriminant  $\Delta > 0$ , and therefore the equation  $s(q) = 0$  has two distinct real roots, each corresponding to one of the non-zero eigenvalues of  $\mathbf{A}$ . Recall that  $\text{rank}(\hat{\mathbf{F}})$  equals the number of the eigenvalues of  $\mathbf{A}$  that are greater than  $\lambda'$ , which is equal to the number of the roots of  $s(q) = 0$  that are greater than  $\lambda'$ . Thus,  $\text{rank}(\hat{\mathbf{F}}) = 1$  if and only if one root of  $s(q) = 0$  is greater than  $\lambda'$  and the other is less than  $\lambda'$ . This is equivalent to

$$s(\lambda') < 0 \iff \lambda'(\lambda' - M) < 0 \iff \lambda' < M \iff \lambda < \frac{MN}{2}.$$

To prove the second part of this proposition, let  $q_+$  be the greatest eigenvalue of  $\mathbf{A}$ , or equivalently the greater root of  $s(q) = 0$ , and  $\mathbf{u}_+$  be the corresponding eigenvector. That is,

$$q_+ = \frac{1}{2r} \left( (M-1)r + (r-1)\lambda' + 1 + \sqrt{\Delta} \right)$$

and  $\mathbf{u}_+ = [x_+ \ y_+ \ \dots \ y_+]^T$ , where  $(x_+, y_+)$  is a solution of  $(x, y)$  in (25) given  $q = q_+$ . Recall that

$$\hat{\mathbf{F}} = \mathbf{R}_*^{\frac{1}{2}} \mathbf{U}(\mathbf{L} - \mathbf{I})\mathbf{U}^T \mathbf{R}_*^{\frac{1}{2}} = \mathbf{R}_*^{\frac{1}{2}} \left( (q_+ - \lambda')\mathbf{u}_+\mathbf{u}_+^T \right) \mathbf{R}_*^{\frac{1}{2}},$$

which leads to  $\mathbf{f} = (q_+ - \lambda')^{\frac{1}{2}} \mathbf{R}_*^{\frac{1}{2}} \mathbf{u}_+ = (q_+ - \lambda')^{\frac{1}{2}} [\sqrt{r}x_+ \ y_+ \ \dots \ y_+]$ , and

$$\frac{f_1}{f_i} = \frac{\sqrt{r}x_+}{y_+} = r(q_+ + 1 - M), \quad \forall i > 1.$$

It is easy to show  $\lim_{\lambda' \rightarrow 0^+} q_+ = \frac{1}{r} + M - 1$  and as a result  $\lim_{\lambda' \rightarrow 0^+} \frac{f_1}{f_i} = 1$ . Therefore, to show  $\frac{f_1}{f_i}$  is greater than 1 and monotonically increasing with  $r$ , it is sufficient to show that

$$\frac{d\frac{f_1}{f_i}}{dr} > 0, \quad \forall r > 1, \lambda' \in (0, M).$$

By straight forward algebra we have

$$\frac{d\frac{f_1}{f_i}}{dr} = (\lambda' - M + 1) + \frac{1}{\sqrt{\Delta}} \left( (\lambda'r - \lambda' - Mr + r)(\lambda' - M + 1) + (M - 1 + \lambda') \right).$$

Now consider three cases:

1.  $0 < \lambda < \frac{(M-1)N}{2}$  : In this case  $0 < \lambda' < M - 1$  and

$$\begin{aligned} \frac{d\frac{f_1}{f_i}}{dr} > 0 &\iff ((\lambda'r - \lambda' - Mr + r)(\lambda' - M + 1) + (M - 1 + \lambda')) > (M - 1 - \lambda')\sqrt{\Delta} \\ &\iff \left(-\lambda'r + \lambda' + Mr - r + \frac{M - 1 + \lambda'}{M - 1 - \lambda'}\right)^2 > \Delta. \end{aligned}$$

Expanding the both sides of the last inequality yields the desired result.

2.  $\lambda = \frac{(M-1)N}{2}$  : In this case  $\lambda' = M - 1$  and  $\frac{d\frac{f_1}{f_i}}{dr} = \frac{2(M-1)}{\sqrt{\Delta}} > 0$ , as desired.
3.  $\frac{(M-1)N}{2} < \lambda < \frac{MN}{2}$  : In this case  $\lambda' - M + 1 \in (0, 1)$ , and we have

$$\begin{aligned} &(\lambda'r - \lambda' - Mr + r)(\lambda' - M + 1) + (M - 1 + \lambda') \\ &= r(\lambda' - M + 1)^2 - \lambda'(\lambda' - M + 1) + (\lambda' + M - 1) \\ &> -\lambda' + (\lambda' + M - 1) > 0 \end{aligned}$$

which implies  $\frac{d\frac{f_1}{f_i}}{dr} > (\lambda' - M + 1)$ . Since the derivative is bounded below by a positive constant, we conclude  $\frac{f_1}{f_i} = \Omega(r)$ . ■

## B Experiment Details

### B.1 Coordinate Ascent Algorithm for STM

Algorithm 3 describes our coordinate ascent method for solving STM.

---

**Algorithm 3** Procedure for solving STM

**Input:**  $\mathcal{X}, \lambda$

**Output:**  $\Sigma_{\text{STM}}^\lambda$

---

**T**  $\leftarrow$  **I**.

**repeat**

**$\Sigma$**   $\leftarrow$  UTM(**T** $\mathcal{X}$ ,  $\lambda$ )

**T**  $\leftarrow$   $\underset{\mathbf{T} \in \mathbb{D}_+^M}{\text{argmax}} \log p(\mathbf{T}\mathcal{X} | \Sigma)$ , s.t.  $\log \det \mathbf{T} \geq 0$

**until** converge

**$\Sigma_{\text{STM}}^\lambda$**   $\leftarrow$  **T**<sup>-1</sup> **$\Sigma$** **T**<sup>-T</sup>

---

### B.2 Cross Validation for Synthetic Data Experiment

For the synthetic data experiment, we select the regularization parameter via the following cross-validation procedure. Let  $\theta$  be the regularization parameter to be determined and  $\mathcal{U}(\mathcal{X}, \theta)$  be the learning algorithm that takes as input  $(\mathcal{X}, \theta)$  and returns a covariance matrix estimate. We randomly split  $\mathcal{X}$  into a partial training set  $\mathcal{X}_T$  and a validation set  $\mathcal{X}_V$ , whose sizes are roughly 70% and 30% of  $\mathcal{X}$ , respectively. For each candidate value of  $\theta$ ,  $\hat{\Sigma}_T^\theta = \mathcal{U}(\mathcal{X}_T, \theta)$  is computed and the likelihood  $p(\mathcal{X}_V | \hat{\Sigma}_T^\theta)$  of the validation set  $\mathcal{X}_V$

conditioned on the solution  $\hat{\Sigma}_T^\theta$  is evaluated. The value of  $\theta$  that maximizes this likelihood is then selected and fed into  $\mathcal{U}(\mathcal{X}, \theta)$  along with the full training set  $\mathcal{X}$ , resulting in our estimate  $\hat{\Sigma}^\theta$ .

In our synthetic data experiment, the  $K$  for URM/EM/MRH are selected from  $\{0, 1, \dots, 15\}$ , and the  $\lambda$  for UTM/TM/STM are selected from  $\{100, 120, \dots, 400\}$ . These ranges are chosen so that the selected values rarely fall on the extremes.

### B.3 Termination Criteria for Iterative Algorithms

In our implementation, the EM algorithm terminates when  $\max_i \left| \frac{\mathbf{R}_{i,i}^{\text{New}} - \mathbf{R}_{i,i}^{\text{Current}}}{\mathbf{R}_{i,i}^{\text{Current}}} \right| < 0.001$ . Similarly, the STM algorithm terminates when  $\max_i \left| \frac{\mathbf{T}_{i,i}^{\text{New}} - \mathbf{T}_{i,i}^{\text{Current}}}{\mathbf{T}_{i,i}^{\text{Current}}} \right| < 0.001$ .

### B.4 Equivalent Data Requirement

Consider two learning algorithms  $\mathcal{U}_1$  and  $\mathcal{U}_2$ , and  $N$  data samples  $\mathcal{X} = \{\mathbf{x}_{(1)}, \dots, \mathbf{x}_{(N)}\}$ . Denote the out-of-sample log-likelihood delivered by  $\mathcal{U}$  with  $\mathcal{X}$  by  $L(\mathcal{U}, \mathcal{X})$ . Suppose  $\mathcal{U}_2$  generally has better performance than  $\mathcal{U}_1$ . Algorithm 4 evaluates the equivalent data requirement of  $\mathcal{U}_2$  with respect to  $\mathcal{U}_1$ . In our implementation, we set step size  $\alpha = 2\%$  for uniform-residual experiment and  $\alpha = 10\%$  for nonuniform-residual ones.

---

**Algorithm 4** Procedure for evaluating equivalent data requirement

**Input:**  $\mathcal{X}, \mathcal{U}_1, \mathcal{U}_2$

**Output:**  $\gamma$

---

```

i ← 0
while 1 do
   $\gamma \leftarrow 1 - i\alpha$ 
   $\mathcal{X}_i \leftarrow \{\mathbf{x}_{(1)}, \dots, \mathbf{x}_{(\gamma N)}\}$ 
  if  $L(\mathcal{U}_2, \mathcal{X}_i) < L(\mathcal{U}_1, \mathcal{X})$  then
    if  $i > 0$  then
       $\gamma \leftarrow \gamma + \frac{L(\mathcal{U}_1, \mathcal{X}) - L(\mathcal{U}_2, \mathcal{X}_i)}{L(\mathcal{U}_2, \mathcal{X}_{i-1}) - L(\mathcal{U}_2, \mathcal{X}_i)} \alpha$  (interpolation)
    end if
    return  $\gamma$ 
  end if
   $i \leftarrow i + 1$ 
end while

```

---

### B.5 S&P500 Data Preprocessing

Define November 2, 2001 as trading day 1 and August 9, 2007 as trading day 1451. After deleting 47 constituent stocks that are not fully defined over this period, we compute for each stock the normalized log daily returns as follows:

1. Let  $y'_{i,j}$  be the adjusted close price of stock  $i$  on day  $j$ ,  $i = 1, \dots, 453$  and  $j = 1, \dots, 1451$ .
2. Compute the raw log-daily-return of stock  $i$  on day  $j$  by

$$y''_{i,j} = \log \frac{y'_{i,j+1}}{y'_{i,j}}, \quad i = 1, \dots, 453, \quad j = 1, \dots, 1450.$$

3. Let  $\bar{y}$  be the smallest number such that at least 99.5% of all  $y''_{i,j}$  are less than or equal to  $\bar{y}$ . Let  $\underline{y}$  be the largest number such that at least 99.5% of all  $y''_{i,j}$ 's are greater than or equal to  $\underline{y}$ . Clip all  $y''_{i,j}$  by the interval  $[\underline{y}, \bar{y}]$ .
4. Let the volatility of stock  $i$  on day  $j > 50$  be the 10-week rms  $\hat{\sigma}_{i,j} = \sqrt{\frac{1}{50} \sum_{t=1}^{50} y''_{i,j-t}{}^2}$ .
5. Set  $\mathbf{y}^{(n)} = \left[ \frac{y''_{1,n+50}}{\hat{\sigma}_{1,n+50}} \quad \dots \quad \frac{y''_{453,n+50}}{\hat{\sigma}_{453,n+50}} \right]^T$  for  $n = 1, \dots, 1400$ .

## B.6 Candidates for Regularization Parameter in Real Data Experiment

In our real data experiment, the  $K$  for EM/MRH are selected from  $\{0, 1, \dots, 40\}$ , and the  $\lambda$  for TM/STM are selected from  $\{200, 210, \dots, 600\}$ . These ranges are chosen so that the selected values never fall on the extremes.

## References

- Akaike, H. (1987). Factor analysis and AIC. *Psychometrika*, 52(3).
- Amini, A. A. and Wainwright, M. J. (2009). High-dimensional analysis of semidefinite relaxations for sparse principal components. *Ann. Statist.*, 37(5B):2877–2921.
- Baik, J. and Silverstein, J. W. (2006). Eigenvalues of large sample covariance matrices of spiked population models. *Journal of Multivariate Analysis*, 97:1382–1408.
- Banerjee, O., Ghaoui, L. E., and d’Aspremont, A. (2008). Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data. *Journal of Machine Learning Research*, 9:485–516.
- Bishop, C. M. (1998). Bayesian PCA. In Kearns, M. J., Solla, S. A., and Cohn, D. A., editors, *Advances in Neural Information Processing Systems 11*, pages 382–388. MIT Press.
- Boyd, S., Parikh, N., Chu, E., Peleato, B., and Eckstein, J. (2011). Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122.
- Candès, E. J., Li, X., Ma, Y., and Wright, J. (2009). Robust principal component analysis? *Journal of the ACM*, 58(1):1–37.
- Chandrasekaran, V., Parrilo, P. A., and Willsky, A. S. (2012). Latent variable graphical model selection via convex optimization. *Annals of Statistics*, 40(4):1935–2013.
- D’Aspremont, A., Ghaoui, L. E., Jordan, M. I., and Lanckriet, G. R. G. (2004). A direct formulation for sparse PCA using semidefinite programming. *SIAM Review*, 49(3):434.
- Friedman, J., Hastie, T., and Tibshirani, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441.
- Harman, H. H. (1976). *Modern Factor Analysis*. University of Chicago Press, third edition.
- Hirose, K., Kawano, S., Konishi, S., and Ichikawa, M. (2011). Bayesian information criterion and selection of the number of factors in factor analysis models. *Journal of Data Science*, 9(2).



- Johnstone, I. M. (2001). On the distribution of the largest eigenvalue in principal components analysis. *Annals of Statistics*, 29:295–327.
- Johnstone, I. M. and Lu, A. Y. (2009). Sparse principal components analysis. *Journal of the American Statistical Association*.
- Jolliffe, I. T., Trendafilov, N. T., and Uddin, M. (2003). A modified principal component technique based on the lasso. *Journal of Computational and Graphical Statistics*, 12:531–547.
- Kao, Y.-H. and Van Roy, B. (2012). Directed principle component analysis. Unpublished manuscript.
- Laurent, B. and Massart, P. (2000). Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics*, 28(5):1302–1338.
- Markowitz, H. M. (1952). Portfolio selection. *Journal of Finance*, 7:77–91.
- Minka, T. P. (2000). Automatic choice of dimensionality for PCA. In Leen, T. K., Dietterich, T. G., and Tresp, V., editors, *Advances in Neural Information Processing Systems 13*, pages 598–604. MIT Press.
- Paul, D. (2007). Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statistica Sinica*, 17(4):1617–1642.
- Pison, G., Rousseeuw, P. J., Filzmoser, P., and Croux, C. (2003). Robust factor analysis. *Journal of Multivariate Analysis*, 84(1):145 – 172.
- Pourahmadi, M. (2011). Covariance estimation: The GLM and regularization perspectives. *Statistical Science*, 26(3):369–387.
- Ravikumar, P., Raskutti, G., Wainwright, M. J., and Yu, B. (2011). High-dimensional covariance estimation by minimizing L1-penalized log-determinant. *Electronic Journal of Statistics*, 5:935–980.
- Rubin, D. B. and Thayer, D. T. (1982). EM algorithm for ML factor analysis. *Psychometrika*, 47(1):69–76.
- Tipping, M. E. and Bishop, C. M. (1999). Probabilistic principal component analysis. *Journal of the Royal Statistical Society, Series B*, 61:611–622.
- Xu, H., Caramanis, C., and Mannor, S. (2010). Principal component analysis with contaminated data: The high dimensional case. In *COLT*, pages 490–502.
- Yuan, M. and Lin, Y. (2007). Model selection and estimation in the gaussian graphical model. *Biometrika*, 94(1):19–35.
- Zou, H., Hastie, T., and Tibshirani, R. (2006). Sparse principal component analysis. *Journal of Computational and Graphical Statistics*, 15(2).