OPTIMAL RATES OF CONVERGENCE FOR COVARIANCE MATRIX ESTIMATION

BY T. TONY CAI[1], CUN-HUI ZHANG[2] AND HARRISON H. ZHOU[3]

University of Pennsylvania, Rutgers University and Yale University

Covariance matrix plays a central role in multivariate statistical analysis. Significant advances have been made recently on developing both theory and methodology for estimating large covariance matrices. However, a minimax theory has yet been developed. In this paper we establish the optimal rates of convergence for estimating the covariance matrix under both the operator norm and Frobenius norm. It is shown that optimal procedures under the two norms are different and consequently matrix estimation under the operator norm is fundamentally different from vector estimation. The minimax upper bound is obtained by constructing a special class of tapering estimators and by studying their risk properties. A key step in obtaining the optimal rate of convergence is the derivation of the minimax lower bound. The technical analysis requires new ideas that are quite different from those used in the more conventional function/sequence estimation problems.

1. Introduction. Suppose we observe independent and identically distributed p-variate random variables $\mathbf{X}_1, \ldots, \mathbf{X}_n$ with covariance matrix $\Sigma_{p \times p}$ and the goal is to estimate the unknown matrix $\Sigma_{p \times p}$ based on the sample $\{\mathbf{X}_i : i = 1, \ldots, n\}$. This covariance matrix estimation problem is of fundamental importance in multivariate analysis. A wide range of statistical methodologies, including clustering analysis, principal component analysis, linear and quadratic discriminant analysis, regression analysis, require the estimation of the covariance matrices. With dramatic advances in technology, large high-dimensional data are now routinely collected in scientific investigations. Examples include climate studies, gene expression arrays, functional magnetic resonance imaging, risk management and portfolio allocation and web search problems. In such settings, the standard and most natural estimator, the sample covariance matrix, often performs poorly. See, for example, Muirhead (1987), Johnstone (2001), Bickel and Levina (2008a, 2008b) and Fan, Fan and Lv (2008).

Regularization methods, originally developed in nonparametric function estimation, have recently been applied to estimate large covariance matrices. These

include banding method in Wu and Pourahmadi (2009) and Bickel and Levina (2008a), tapering in Furrer and Bengtsson (2007), thresholding in Bickel and Levina (2008b) and El Karoui (2008), penalized estimation in Huang et al. (2006), Lam and Fan (2007) and Rothman et al. (2008), regularizing principal components in Johnstone and Lu (2009) and Zou, Hastie and Tibshirani (2006). Asymptotic properties and convergence results have been given in several papers. In particular, Bickel and Levina (2008a, 2008b), El Karoui (2008) and Lam and Fan (2007) showed consistency of their estimators in operator norm and even obtained explicit rates of convergence. However, it is not clear whether any of these rates of convergence are optimal.

Despite recent progress on covariance matrix estimation there has been remarkably little fundamental theoretical study on optimal estimation. In this paper, we establish the optimal rate of convergence for estimating the covariance matrix as well as its inverse over a wide range of classes of covariance matrices. Both the operator norm and Frobenius norm are considered. It is shown that optimal procedures for these two norms are different and consequently matrix estimation under the operator norm is fundamentally different from vector estimation. In addition, the results also imply that the banding estimator given in Bickel and Levina (2008a) is sub-optimal under the operator norm and the performance can be significantly improved.

We begin by considering optimal estimation of the covariance matrix Σ over a class of matrices that has been considered in Bickel and Levina (2008a). Both minimax lower and upper bounds are derived. We write $a_n \asymp b_n$ if there are positive constants c and C independent of n such that $c \le a_n/b_n \le C$. For a matrix A its operator norm is defined as $\|A\| = \sup_{\|x\|_2=1} \|Ax\|_2$. We assume that $p \le \exp(\gamma n)$ for some constant $\gamma > 0$. Combining the results given in Section 3, we have the following optimal rate of convergence for estimating the covariance matrix under the operator norm.

THEOREM 1. *The minimax risk of estimating the covariance matrix Σ over the class $\mathcal{P}_\alpha$ given in* (3) *satisfies*

$$(1) \qquad \inf_{\hat{\Sigma}} \sup_{\mathcal{P}_\alpha} \mathbb{E}\|\hat{\Sigma} - \Sigma\|^2 \asymp \min\left\{ n^{-2\alpha/(2\alpha+1)} + \frac{\log p}{n}, \frac{p}{n} \right\}.$$

The minimax upper bound is obtained by constructing a class of tapering estimators and by studying their risk properties. It is shown that the estimator with the optimal choice of the tapering parameter attains the optimal rate of convergence. In comparison to some existing methods in the literature, the proposed procedure does not attempt to estimate each row/column optimally as a vector. In fact, our procedure does not optimally trade bias and variance for each row/column. As a vector estimator, it has larger variance than squared bias for each row/column. In other words, it is undersmoothed as a vector.

A key step in obtaining the optimal rate of convergence is the derivation of the minimax lower bound. The lower bound is established by using a testing argument, where at the core is a novel construction of a collection of least favorable multivariate normal distributions and the application of Assouad's lemma and Le Cam's method. The technical analysis requires ideas that are quite different from those used in the more conventional function/sequence estimation problems.

In addition to the asymptotic analysis, we also carry out a small simulation study to investigate the finite sample performance of the proposed estimator. The tapering estimator is easy to implement. The numerical performance of the estimator is compared with that of the banding estimator introduced in Bickel and Levina (2008a). The simulation study shows that the proposed estimator has good numerical performance; it nearly uniformly outperforms the banding estimator.

The paper is organized as follows. In Section 2, after basic notation and definitions are introduced, we propose a tapering procedure for the covariance matrix estimation. Section 3 derives the optimal rate of convergence for estimation under the operator norm. The upper bound is obtained by studying the properties of the tapering estimators and the minimax lower bound is obtained by a testing argument. Section 4 considers optimal estimation under the Frobenius norm. The problem of estimating the inverse of a covariance matrix is treated in Section 5. Section 6 investigates the numerical performance of our procedure by a simulation study. The technical proofs of auxiliary lemmas are given in Section 7.

2. Methodology. In this section we will introduce a tapering procedure for estimating the covariance matrix $\Sigma_{p \times p}$ based on a random sample of p-variate observations $\mathbf{X}_1, \ldots, \mathbf{X}_n$. The properties of the tapering estimators under the operator norm and Frobenius norm are then studied and used to establish the minimax upper bounds in Sections 3 and 4.

Given a random sample $\{\mathbf{X}_1, \ldots, \mathbf{X}_n\}$ from a population with covariance matrix $\Sigma = \Sigma_{p \times p}$, the sample covariance matrix is

$$\frac{1}{n-1} \sum_{l=1}^{n} (\mathbf{X}_l - \bar{\mathbf{X}})(\mathbf{X}_l - \bar{\mathbf{X}})^T,$$

which is an unbiased estimate of Σ, and the maximum likelihood estimator of Σ is

$$(2) \qquad \Sigma^* = (\sigma_{ij}^*)_{1 \leq i, j \leq p} = \frac{1}{n} \sum_{l=1}^{n} (\mathbf{X}_l - \bar{\mathbf{X}})(\mathbf{X}_l - \bar{\mathbf{X}})^T,$$

when $\mathbf{X}_l$'s are normally distributed. These two estimators are close to each other for large n. We shall construct estimators of the covariance matrix Σ by tapering the maximum likelihood estimator Σ^*.

Following Bickel and Levina (2008a) we consider estimating the covariance matrix $\Sigma_{p\times p} = (\sigma_{ij})_{1\leq i,j\leq p}$ over the following parameter space:

(3)
$$\mathcal{F}_\alpha = \mathcal{F}_\alpha(M_0, M) = \left\{\Sigma : \max_j \sum_i \{|\sigma_{ij}| : |i - j| > k\} \leq Mk^{-\alpha}\right.$$

$$\left. \text{for all } k, \text{ and } \lambda_{\max}(\Sigma) \leq M_0\right\},$$

where $\lambda_{\max}(\Sigma)$ is the maximum eigenvalue of the matrix Σ, and $\alpha > 0$, $M > 0$ and $M_0 > 0$. Note that the smallest eigenvalue of any covariance matrix in the parameter space F_α is allowed to be 0 which is more general than the assumption in (5) of Bickel and Levina (2008a). The parameter α in (3), which essentially specifies the rate of decay for the covariances σ_{ij} as they move away from the diagonal, can be viewed as an analog of the smoothness parameter in nonparametric function estimation problems. The optimal rate of convergence for estimating Σ over the parameter space $\mathcal{F}_\alpha(M_0, M)$ critically depends on the value of α. Our estimators of the covariance matrix Σ are constructed by tapering the maximum likelihood estimator (2) as follows.

Estimation procedure. For a given even integer k with $1 \leq k \leq p$, we define a tapering estimator as

(4)
$$\hat{\Sigma} = \hat{\Sigma}_k = (w_{ij}\sigma_{ij}^*)_{p\times p},$$

where σ_{ij}^* are the entries in the maximum likelihood estimator Σ^* and the weights

(5)
$$w_{ij} = k_h^{-1}\{(k - |i - j|)_+ - (k_h - |i - j|)_+\},$$

where $k_h = k/2$. Without loss of generality we assume that k is even. Note that the weights w_{ij} can be rewritten as

$$w_{ij} = \begin{cases} 1, & \text{when } |i - j| \leq k_h, \\ 2 - \dfrac{|i - j|}{k_h}, & \text{when } k_h < |i - j| < k, \\ 0, & \text{otherwise.} \end{cases}$$

See Figure 1 for a plot of the weights w_{ij} as a function of $|i - j|$.

The tapering estimators are different from the banding estimators used in Bickel and Levina (2008a). It is important to note that the tapering estimator given in (4) can be rewritten as a sum of many small block matrices along the diagonal. This simple but important observation is very useful for our technical arguments. Define the block matrices

$$M_l^{*(m)} = (\sigma_{ij}^* I\{l \leq i < l + m, l \leq j < l + m\})_{p\times p}$$

and set

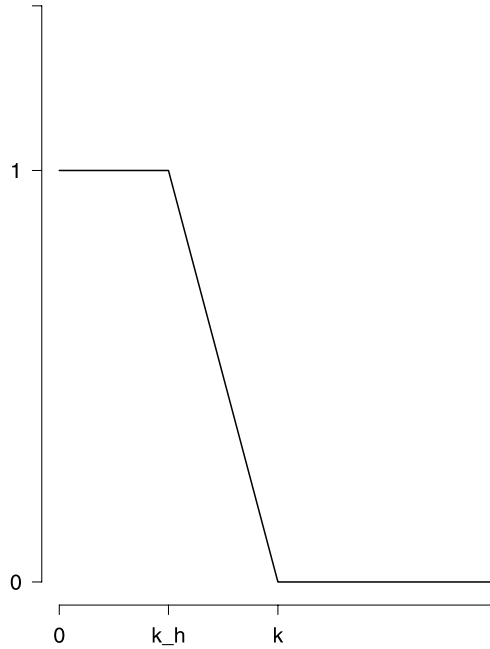$$S^{*(m)} = \sum_{l=1-m}^{p} M_l^{*(m)}$$

FIG. 1.   *The weights as a function of* $|i - j|$.

for all integers $1 - m \le l \le p$ and $m \ge 1$.

LEMMA 1.   *The tapering estimator* $\hat{\Sigma}_k$ *given in* (4) *can be written as*

$$(6) \qquad \hat{\Sigma}_k = k_h^{-1}\big(S^{*(k)} - S^{*(k_h)}\big).$$

It is clear that the performance of the estimator $\hat{\Sigma}_k$ depends on the choice of the tapering parameter k. The optimal choice of k critically depends on the norm under which the estimation error is measured. We will study in the next two sections the rate of convergence of the tapering estimator under both the operator norm and Frobenius norm. Together with the minimax lower bounds derived in Sections 3 and 4, the results show that a tapering estimator with the optimal choice of k attains the optimal rate of convergence under these two norms.

3. Rate optimality under the operator norm.   In this section we will establish the optimal rate of convergence under the operator norm. For $1 \le q \le \infty$, the matrix ℓ_q-norm of a matrix A is defined by $\|A\|_q = \max_{\|x\|_q=1}\|Ax\|_q$. The commonly used operator norm $\|\cdot\|$ coincides with the matrix ℓ_2-norm $\|\cdot\|_2$. For a symmetric matrix A, it is known that the operator norm $\|A\|$ is equal to the largest magnitude of eigenvalues of A. Hence it is also called the spectral norm. We will

establish Theorem 1 by deriving a minimax upper bound using the tapering estimator and a matching minimax lower bound by a careful construction of a collection of multivariate normal distributions and the application of Assouad's lemma and Le Cam's method. We shall focus on the case $p \geq n^{1/(2\alpha+1)}$ in Sections 3.1 and 3.2. The case of $p < n^{1/(2\alpha+1)}$, which will be discussed in Section 3.3, is similar and slightly easier.

3.1. *Minimax upper bound under the operator norm.* We derive in this section the risk upper bound for the tapering estimators defined in (6) under the operator norm. Throughout the paper we denote by C a generic positive constant which may vary from place to place but always depends only on indices α, M_0 and M of the matrix family. We shall assume that the distribution of the X_i's is sub-Gaussian in the sense that there is $\rho > 0$ such that

$$(7) \qquad \mathbb{P}\{|\mathbf{v}^T(\mathbf{X}_1 - \mathbb{E}\mathbf{X}_1)| > t\} \leq e^{-t^2\rho/2} \qquad \text{for all } t > 0 \text{ and } \|\mathbf{v}\|_2 = 1.$$

Let $\mathcal{P}_\alpha = \mathcal{P}_\alpha(M_0, M, \rho)$ denote the set of distributions of $\mathbf{X}_1$ that satisfy (3) and (7).

THEOREM 2. *The tapering estimator $\hat{\Sigma}_k$, defined in* (6), *of the covariance matrix $\Sigma_{p \times p}$ with $p \geq n^{1/(2\alpha+1)}$ satisfies*

$$(8) \qquad \sup_{\mathcal{P}_\alpha} \mathbb{E}\|\hat{\Sigma}_k - \Sigma\|^2 \leq C\frac{k + \log p}{n} + Ck^{-2\alpha}$$

for $k = o(n)$, $\log p = o(n)$ and some constant $C > 0$. In particular, the estimator $\hat{\Sigma} = \hat{\Sigma}_k$ with $k = n^{1/(2\alpha+1)}$ satisfies

$$(9) \qquad \sup_{\mathcal{P}_\alpha} \mathbb{E}\|\hat{\Sigma} - \Sigma\|^2 \leq Cn^{-2\alpha/(2\alpha+1)} + C\frac{\log p}{n}.$$

From (8) it is clear that the optimal choice of k is of order $n^{1/(2\alpha+1)}$. The upper bound given in (9) is thus rate optimal among the class of the tapering estimators defined in (6). The minimax lower bound derived in Section 3.2 shows that the estimator $\hat{\Sigma}_k$ with $k = n^{1/(2\alpha+1)}$ is in fact rate optimal among all estimators.

PROOF OF THEOREM 2. Note that Σ^* is translation invariant and so is $\hat{\Sigma}$. We shall thus assume $\mu = 0$ for the rest of the paper. Write

$$\Sigma^* = \frac{1}{n}\sum_{l=1}^{n}(\mathbf{X}_l - \bar{\mathbf{X}})(\mathbf{X}_l - \bar{\mathbf{X}})^T = \frac{1}{n}\sum_{l=1}^{n}\mathbf{X}_l\mathbf{X}_l^T - \bar{\mathbf{X}}\bar{\mathbf{X}}^T,$$

where $\bar{\mathbf{X}}\bar{\mathbf{X}}^T$ is a higher order term (see Remark 1 at the end of this section). In what follows we shall ignore this negligible term and focus on the dominating term $\frac{1}{n}\sum_{l=1}^{n}\mathbf{X}_l\mathbf{X}_l^T$.

Set $\tilde{\Sigma} = \frac{1}{n} \sum_{l=1}^{n} \mathbf{X}_l \mathbf{X}_l^T$ and write $\tilde{\Sigma} = (\tilde{\sigma}_{ij})_{1 \le i, j \le p}$. Let

$$(10) \qquad \check{\Sigma} = (\check{\sigma}_{ij})_{1 \le i, j \le p} = (w_{ij} \tilde{\sigma}_{ij})_{1 \le i, j \le p}$$

with w_{ij} given in (5). Let $\mathbf{X}_l = (X_1^l, X_2^l, \ldots, X_p^l)^T$. We then write $\tilde{\sigma}_{ij} = \frac{1}{n} \sum_{l=1}^{n} X_i^l X_j^l$. It is easy to see

$$(11) \qquad \mathbb{E}\tilde{\sigma}_{ij} = \sigma_{ij},$$

$$(12) \quad \operatorname{Var}(\tilde{\sigma}_{ij}) = \frac{1}{n} \operatorname{Var}(X_i^l X_j^l) \le \frac{1}{n} \mathbb{E}(X_i^l X_j^l)^2 \le \frac{1}{n} [\mathbb{E}(X_i^l)^4]^{1/2} [\mathbb{E}(X_j^l)^4]^{1/2} \le \frac{C}{n},$$

that is, $\tilde{\sigma}_{ij}$ is an unbiased estimator of σ_{ij} with a variance $O(1/n)$.

We will first show that the variance part satisfies

$$(13) \qquad \mathbb{E}\|\check{\Sigma} - \mathbb{E}\check{\Sigma}\|^2 \le C \frac{k + \log p}{n}$$

and the bias part satisfies

$$(14) \qquad \|\mathbb{E}\check{\Sigma} - \Sigma\|^2 \le Ck^{-2\alpha}.$$

It then follows immediately that

$$\mathbb{E}\|\check{\Sigma} - \Sigma\|^2 \le 2\mathbb{E}\|\check{\Sigma} - \mathbb{E}\check{\Sigma}\|^2 + 2\|\mathbb{E}\check{\Sigma} - \Sigma\|^2 \le 2C\left(\frac{k + \log p}{n} + k^{-2\alpha}\right).$$

This proves (8) and equation (9) then follows. Since $p \ge n^{1/(2\alpha+1)}$, we may choose

$$(15) \qquad k = n^{1/(2\alpha+1)}$$

and the estimator $\hat{\Sigma}$ with k given in (15) satisfies

$$\mathbb{E}\|\hat{\Sigma} - \Sigma\|^2 \le 2C\left(n^{-2\alpha/(2\alpha+1)} + \frac{\log p}{n}\right).$$

Theorem 2 is then proved. $\square$

We first prove the risk upper bound (14) for the bias part. It is well known that the operator norm of a symmetric matrix $A = (a_{ij})_{p \times p}$ is bounded by its ℓ_1 norm, that is,

$$\|A\| \le \|A\|_1 = \max_{i=1,\ldots,p} \sum_{j=1}^{p} |a_{ij}|$$

[see, e.g., page 15 in Golub and Van Loan (1983)]. This result was used in Bickel and Levina (2008a, 2008b) to obtain rates of convergence for their proposed procedures under the operator norm (see discussions in Section 3.3). We bound the operator norm of the bias part $\mathbb{E}\check{\Sigma} - \Sigma$ by its ℓ_1 norm. Since $\mathbb{E}\tilde{\sigma}_{ij} = \sigma_{ij}$, we have

$$\mathbb{E}\check{\Sigma} - \Sigma = ((w_{ij} - 1)\sigma_{ij})_{p \times p},$$

where $w_{ij} \in [0, 1]$ and is exactly 1 when $|i - j| \le k$, then

$$\|\mathbb{E}\check{\Sigma} - \Sigma\|^2 \le \left[\max_{i=1,\ldots,p} \sum_{j\,:\,|i-j|>k} |\sigma_{ij}|\right]^2 \le M^2 k^{-2\alpha}.$$

Now we establish (13) which is relatively complicated. The key idea in the proof is to write the whole matrix as an average of matrices which are sum of a large number of small disjoint block matrices, and for each small block matrix the classical random matrix theory can be applied. The following lemma shows that the operator norm of the random matrix $\check{\Sigma} - \mathbb{E}\check{\Sigma}$ is controlled by the maximum of operator norms of p number of $k \times k$ random matrices. Let $M_l^{(m)} = (\tilde{\sigma}_{ij} I\{l \le i < l + m, l \le j < l + m\})_{p \times p}$. Define

$$N_l^{(m)} = \max_{1 \le l \le p-m+1} \|M_l^{(m)} - \mathbb{E}M_l^{(m)}\|.$$

LEMMA 2. *Let $\check{\Sigma}$ be defined as in* (6). *Then*

$$\|\check{\Sigma} - \mathbb{E}\check{\Sigma}\| \le 3N_l^{(m)}.$$

For each small $m \times m$ random matrix with $m = k$, we control its operator norm as follows.

LEMMA 3. *There is a constant $\rho_1 > 0$ such that*

(16) $$\mathbb{P}\{N_l^{(m)} > x\} \le 2p5^m \exp(-nx^2\rho_1)$$

for all $0 < x < \rho_1$ and $1 - m \le l \le p$.

With Lemmas 2 and 3 we are now ready to show the variance bound (13). By Lemma 2 we have

$$\mathbb{E}\|\check{\Sigma} - \mathbb{E}\check{\Sigma}\|^2 \le 9\mathbb{E}(N_l^{(m)})^2 = 9\mathbb{E}(N_l^{(m)})^2[I(N_l^{(m)} \le x) + I(N_l^{(m)} > x)]$$
$$\le 9[x^2 + \mathbb{E}(N_l^{(m)})^2 I(N_l^{(m)} > x)].$$

Note that $\|\mathbb{E}\check{\Sigma}\| \le \|\Sigma\|$, which is bounded by a constant, and $\|\check{\Sigma}\| \le \|\check{\Sigma}\|_F$. The Cauchy–Schwarz inequality then implies

$$\mathbb{E}\|\check{\Sigma} - \mathbb{E}\check{\Sigma}\|^2 \le C_1[x^2 + \mathbb{E}(\|\check{\Sigma}\|_F^2 + C)I(N_l^{(m)} > x)]$$
$$\le C_1\left[x^2 + \sqrt{\mathbb{E}(\|\check{\Sigma}\|_F + C)^4}\sqrt{\mathbb{P}(N_l^{(m)} > x)}\right].$$

Set $x = 4\sqrt{\frac{\log p + m}{n\rho_1}}$. Then x is bounded by ρ_1 as $n \to \infty$. From Lemma 3 we obtain

(17)
$$\mathbb{E}\|\check{\Sigma} - \mathbb{E}\check{\Sigma}\|^2 \le C\left[\frac{\log p + m}{n} + p^2 \cdot (p5^m \cdot p^{-8}e^{-8m})^{1/2}\right]$$
$$\le C_1\left(\frac{\log p + m}{n}\right).$$

REMARK 1.   In the proof of Theorem 2, the term $\bar{\mathbf{X}}\bar{\mathbf{X}}^T$ was ignored. It is not difficult to see that this term has negligible contribution after tapering. Let $H = \bar{\mathbf{X}}\bar{\mathbf{X}}^T$ and $H = (h_{ij})_{p \times p}$. Define

$$H_l^{(m)} = (h_{ij} I\{l \leq i < l + m, l \leq j < l + m\})_{p \times p}.$$

Similarly to Lemma 3, it can be shown that

$$(18) \qquad \mathbb{P}\Big\{ \max_{1 \leq l \leq p-m+1} \| H_l^{(m)} - \mathbb{E}H_l^{(m)} \| > t \Big\} \leq 2p5^m \exp(-nt\rho_2)$$

for all $0 < t < \rho_2$ and $1 - m \leq l \leq p$. Note that $\mathbb{E}H = \frac{1}{n}\Sigma$, then

$$\mathbb{E}\|H\|^2 \leq 2\mathbb{E}\|H - \mathbb{E}H\|^2 + 2\|\mathbb{E}H\|^2 \leq 2\mathbb{E}\|H - \mathbb{E}H\|^2 + 2M_0^2/n^2.$$

Let $t = 16\frac{\log p + m}{n\rho_2}$. From (18) we have

$$\mathbb{E}\|H - \mathbb{E}H\|^2 \leq t^2 + \mathbb{E}\|H - \mathbb{E}H\|^2 I\Big(\max_{1 \leq l \leq p-m+1} \| H_l^{(m)} - \mathbb{E}H_l^{(m)} \| > t \Big)$$

$$= t^2 + o(t^2) \leq C\Big(\frac{\log p + m}{n} \Big)^2$$

by similar arguments as for (17). Therefore H has a negligible contribution to the risk.

3.2. *Lower bound under the operator norm.*   Theorem 2 in Section 3.1 shows that the optimal tapering estimator attains the rate of convergence $n^{-2\alpha/(2\alpha+1)} + \frac{\log p}{n}$. In this section we shall show that this rate of convergence is indeed optimal among all estimators by showing that the upper bound in equation (9) cannot be improved. More specifically we shall show that the following minimax lower bound holds.

THEOREM 3.   *Suppose $p \leq \exp(\gamma n)$ for some constant $\gamma > 0$. The minimax risk for estimating the covariance matrix Σ over $\mathcal{P}_\alpha$ under the operator norm satisfies*

$$\inf_{\hat{\Sigma}} \sup_{\mathcal{P}_\alpha} \mathbb{E}\|\hat{\Sigma} - \Sigma\|^2 \geq cn^{-2\alpha/(2\alpha+1)} + c\frac{\log p}{n}.$$

The basic strategy underlying the proof of Theorem 3 is to carefully construct a finite collection of multivariate normal distributions and calculate the total variation affinity between pairs of probability measures in the collection.

We shall now define a parameter space that is appropriate for the minimax lower bound argument. For given positive integers k and m with $2k \leq p$ and $1 \leq m \leq k$, define the $p \times p$ matrix $B(m, k) = (b_{ij})_{p \times p}$ with

$$b_{ij} = I\{i = m \text{ and } m + 1 \leq j \leq 2k, \text{ or } j = m \text{ and } m + 1 \leq i \leq 2k\}.$$

Set $k = n^{1/(2\alpha+1)}$ and $a = k^{-(\alpha+1)}$. We then define the collection of 2^k covariance matrices as

$$(19) \qquad \mathcal{F}_{11} = \left\{ \Sigma(\theta) : \Sigma(\theta) = I_p + \tau a \sum_{m=1}^{k} \theta_m B(m, k), \theta = (\theta_m) \in \{0, 1\}^k \right\},$$

where I_p is the $p \times p$ identity matrix and $0 < \tau < 2^{-\alpha-1} M$. Without loss of generality we assume that $M_0 > 1$ and $\rho > 1$. Otherwise we replace I_p in (19) by εI_p for $0 < \varepsilon < \min\{M_0, \rho\}$. For $0 < \tau < 2^{-\alpha-1} M$ it is easy to check that $\mathcal{F}_{11} \subset \mathcal{F}_\alpha(M_0, M)$ as $n \to \infty$. In addition to $\mathcal{F}_{11}$ we also define a collection of diagonal matrices

$$(20) \qquad \mathcal{F}_{12} = \left\{ \Sigma_m : \Sigma_m = I_p + \left(\sqrt{\frac{\tau}{n} \log p_1} I\{i = j = m\} \right)_{p \times p}, 0 \le m \le p_1 \right\},$$

where $p_1 = \min\{p, e^{n/2}\}$ and $0 < \tau < \min\{(M_0 - 1)^2, (\rho - 1)^2, 1\}$. Let $\mathcal{F}_1 = \mathcal{F}_{11} \cup \mathcal{F}_{12}$. It is clear that $\mathcal{F}_1 \subset \mathcal{F}_\alpha(M_0, M)$.

We shall show below separately that the minimax risks over multivariate normal distributions with covariance matrix in (19) and (20) satisfy

$$(21) \qquad \inf_{\hat{\Sigma}} \sup_{\mathcal{F}_{11}} \mathbb{E} \| \hat{\Sigma} - \Sigma \|^2 \ge c n^{-2\alpha/(2\alpha+1)}$$

and

$$(22) \qquad \inf_{\hat{\Sigma}} \sup_{\mathcal{F}_{12}} \mathbb{E} \| \hat{\Sigma} - \Sigma \|^2 \ge c \frac{\log p}{n}$$

for some constant $c > 0$. Equations (21) and (22) together imply

$$(23) \qquad \inf_{\hat{\Sigma}} \sup_{\mathcal{F}_1} \mathbb{E} \| \hat{\Sigma} - \Sigma \|^2 \ge \frac{c}{2} \left(n^{-2\alpha/(2\alpha+1)} + \frac{\log p}{n} \right)$$

for multivariate normal distributions and this proves Theorem 3. We shall establish the lower bound (21) by using Assouad's lemma in Section 3.2.1 and the lower bound (22) by using Le Cam's method and a two-point argument in Section 3.2.2.

3.2.1. *A lower bound by Assouad's lemma.* The key technical tool to establish equation (21) is Assouad's lemma in Assouad (1983). It gives a lower bound for the maximum risk over the parameter set $\Theta = \{0, 1\}^k$ to the problem of estimating an arbitrary quantity $\psi(\theta)$, belonging to a metric space with metric d. Let $H(\theta, \theta') = \sum_{i=1}^k |\theta_i - \theta_i'|$ be the Hamming distance on $\{0, 1\}^k$, which counts the number of positions at which θ and θ' differ. For two probability measures P and Q with density p and q with respect to any common dominating measure μ, write the total variation affinity $\|P \wedge Q\| = \int p \wedge q \, d\mu$. Assouad's lemma provides a minimax lower bound for estimating $\psi(\theta)$.

LEMMA 4 (Assouad). *Let $\Theta = \{0, 1\}^k$ and let T be an estimator based on an observation from a distribution in the collection $\{P_\theta, \theta \in \Theta\}$. Then for all $s > 0$*

$$\max_{\theta \in \Theta} 2^s \mathbb{E}_\theta d^s(T, \psi(\theta)) \geq \min_{H(\theta,\theta') \geq 1} \frac{d^s(\psi(\theta), \psi(\theta'))}{H(\theta, \theta')} \cdot \frac{k}{2} \cdot \min_{H(\theta,\theta')=1} \|\mathbb{P}_\theta \wedge \mathbb{P}_{\theta'}\|.$$

Assouad's lemma is connected to multiple comparisons. In total there are k comparisons. The lower bound has three factors. The first factor is basically the minimum cost of making a mistake per comparison, and the last factor is the lower bound for the total probability of making type I and type II errors for each comparison, and $k/2$ is the expected number of mistakes one makes when $\mathbb{P}_\theta$ and $\mathbb{P}_{\theta'}$ are not distinguishable from each other when $H(\theta, \theta') = 1$.

We now prove the lower bound (21). Let $\mathbf{X}_1, \ldots, \mathbf{X}_n \overset{\text{i.i.d.}}{\sim} N(0, \Sigma(\theta))$ with $\Sigma(\theta) \in \mathcal{F}_{11}$. Denote the joint distribution by P_θ. Applying Assouad's lemma to the parameter space $\mathcal{F}_{11}$, we have

(24)
$$\inf_{\hat{\Sigma}} \max_{\theta \in \{0,1\}^k} 2^2 E_\theta \|\hat{\Sigma} - \Sigma(\theta)\|^2$$

$$\geq \min_{H(\theta,\theta') \geq 1} \frac{\|\Sigma(\theta) - \Sigma(\theta')\|^2}{H(\theta, \theta')} \frac{k}{2} \min_{H(\theta,\theta')=1} \|P_\theta \wedge P_{\theta'}\|.$$

We shall state the bounds for the the first and third factors on the right-hand side of (24) in two lemmas. The proofs of these lemmas are given in Section 7.

LEMMA 5. *Let $\Sigma(\theta)$ be defined as in (19). Then for some constant $c > 0$*

$$\min_{H(\theta,\theta') \geq 1} \frac{\|\Sigma(\theta) - \Sigma(\theta')\|^2}{H(\theta, \theta')} \geq cka^2.$$

LEMMA 6. *Let $\mathbf{X}_1, \ldots, \mathbf{X}_n \overset{\text{i.i.d.}}{\sim} N(0, \Sigma(\theta))$ with $\Sigma(\theta) \in \mathcal{F}_{11}$. Denote the joint distribution by P_θ. Then for some constant $c > 0$*

$$\min_{H(\theta,\theta')=1} \|P_\theta \wedge P_{\theta'}\| \geq c.$$

It then follows from Lemmas 5 and 6 together, with the fact $k = n^{1/(2\alpha+1)}$,

$$\max_{\Sigma(\theta) \in \mathcal{F}_{11}} 2^2 E_\theta \|\hat{\Sigma} - \Sigma(\theta)\|^2 \geq \frac{c^2}{2} k^2 a^2 \geq c_1 n^{-2\alpha/(2\alpha+1)}$$

for some $c_1 > 0$.

3.2.2. *A lower bound using Le Cam's method.* We now apply Le Cam's method to derive the lower bound (22) for the minimax risk. Let X be an observation from a distribution in the collection $\{P_\theta, \theta \in \Theta\}$ where $\Theta = \{\theta_0, \theta_1, \ldots, \theta_{p_1}\}$. Le Cam's method, which is based on a two-point testing argument, gives a lower bound for the maximum estimation risk over the parameter set Θ. More specifically, let L be the loss function. Define $r(\theta_0, \theta_m) = \inf_t [L(t, \theta_0) + L(t, \theta_m)]$ and $r_{\min} = \inf_{1 \le m \le p_1} r(\theta_0, \theta_m)$, and denote $\bar{\mathbb{P}} = \frac{1}{p_1} \sum_{m=1}^{p_1} \mathbb{P}_{\theta_m}$.

LEMMA 7. *Let T be an estimator of θ based on an observation from a distribution in the collection $\{P_\theta, \theta \in \Theta = \{\theta_0, \theta_1, \ldots, \theta_{p_1}\}\}$, then*

$$\sup_\theta \mathbb{E}L(T, \theta) \ge \frac{1}{2} r_{\min} \|\mathbb{P}_{\theta_0} \wedge \bar{\mathbb{P}}\|.$$

We refer to Yu (1997) for more detailed discussions on Le Cam's method.

To apply Le Cam's method, we need to first construct a parameter set. For $1 \le m \le p_1$, let Σ_m be a diagonal covariance matrix with $\sigma_{mm} = 1 + \sqrt{\tau \frac{\log p_1}{n}}$, $\sigma_{ii} = 1$ for $i \ne m$, and let Σ_0 be the identity matrix. Let $\mathbf{X}_l = (X_1^l, X_2^l, \ldots, X_p^l)^T \sim N(0, \Sigma_m)$, and denote the joint density of $\mathbf{X}_1, \ldots, \mathbf{X}_n$ by f_m, $0 \le m \le p_1$ with $p_1 = \max\{p, e^{n/2}\}$, which can be written as follows:

$$f_m = \prod_{1 \le i \le n, 1 \le j \le p, j \ne m} \phi_1(x_j^i) \cdot \prod_{1 \le i \le n} \phi_{\sigma_{mm}}(x_m^i),$$

where ϕ_σ, $\sigma = 1$ or σ_{mm}, is the density of $N(0, \sigma^2)$. Denote by f_0 the joint density of $\mathbf{X}_1, \ldots, \mathbf{X}_n$ when $\mathbf{X}_l \sim N(0, \Sigma_0)$.

Let $\theta_m = \Sigma_m$ for $0 \le m \le p_1$ and the loss function L be the squared operator norm. It is easy to see $r(\theta_0, \theta_m) = \frac{1}{2} \tau \frac{\log p_1}{n}$ for all $1 \le m \le p_1$. Then the lower bound (22) follows immediately from Lemma 7 if there is a constant $c > 0$ such that

$$(25) \qquad \qquad \|\mathbb{P}_{\theta_0} \wedge \bar{\mathbb{P}}\| \ge c.$$

Note that for any two densities q_0 and q_1, $\int q_0 \wedge q_1 \, d\mu = 1 - \frac{1}{2} \int |q_0 - q_1| \, d\mu$, and Jensen's inequality implies

$$\left[\int |q_0 - q_1| \, d\mu \right]^2 = \left(\int \left| \frac{q_0 - q_1}{q_1} \right| q_1 \, d\mu \right)^2 \le \int \frac{(q_0 - q_1)^2}{q_1} \, d\mu = \int \frac{q_0^2}{q_1} \, d\mu - 1.$$

Hence $\int q_0 \wedge q_1 \, d\mu \ge 1 - \frac{1}{2} (\int \frac{q_0^2}{q_1} \, d\mu - 1)^{1/2}$. To establish equation (25), it thus suffices to show that $\int (\frac{1}{p_1} \sum_{m=1}^{p_1} f_m)^2 / f_0 \, d\mu - 1 \to 0$, that is,

$$(26) \qquad \frac{1}{p_1^2} \sum_{m=1}^{p_1} \int \frac{f_m^2}{f_0} \, d\mu + \frac{1}{p_1^2} \sum_{m \ne j} \int \frac{f_m f_j}{f_0} \, d\mu - 1 \to 0.$$

We now calculate $\int \frac{f_m f_j}{f_0} d\mu$. For $m \neq j$ it is easy to see

$$\int \frac{f_m f_j}{f_0} d\mu - 1 = 0.$$

When $m = j$, we have

$$\int \frac{f_m^2}{f_0} d\mu = \frac{(\sqrt{2\pi}\sigma_{mm})^{-2n}}{(\sqrt{2\pi})^{-n}} \prod_{1 \leq i \leq n} \int \exp\left[(x_m^i)^2\left(-\frac{1}{\sigma_{mm}} + \frac{1}{2}\right)\right] dx_m^i$$

$$= [1 - (1 - \sigma_{mm})^2]^{-n/2} = \left(1 - \tau \frac{\log p_1}{n}\right)^{-n/2}.$$

Thus

$$\int \left(\frac{1}{p_1}\sum_{m=1}^{p_1} f_m\right)^2 / f_0 \, d\mu - 1$$

$$= \frac{1}{p_1^2}\sum_{m=1}^{p_1}\left(\int \frac{f_m^2}{f_0} d\mu - 1\right)$$

(27)

$$\leq \frac{1}{p_1}\left(1 - \tau\frac{\log p_1}{n}\right)^{-n/2} - \frac{1}{p_1}$$

$$= \exp\left[-\log p_1 - \frac{n}{2}\log\left(1 - \tau\frac{\log p_1}{n}\right)\right] - \frac{1}{p_1} \to 0$$

for $0 < \tau < 1$, where the last step follows from the inequality $\log(1 - x) \geq -2x$ for $0 < x < 1/2$. Equation (27), together with Lemma 7, now immediately implies the lower bound given in (22).

REMARK 2. In covariance matrix estimation literature, it is commonly assumed that $\frac{\log p}{n} \to 0$. See, for example, Bickel and Levina (2008a). The lower bound given in this section implies that this assumption is necessary for estimating the covariance matrix consistently under the operator norm.

3.3. *Discussion.* Theorems 2 and 3 together show that the minimax risk for estimating the covariance matrices over the distribution space $\mathcal{P}_\alpha$ satisfies, for $p \geq n^{1/(2\alpha+1)}$,

(28)
$$\inf_{\hat{\Sigma}} \sup_{\mathcal{P}_\alpha} \mathbb{E}\|\hat{\Sigma} - \Sigma\|^2 \asymp n^{-2\alpha/(2\alpha+1)} + \frac{\log p}{n}.$$

The results also show that the tapering estimator $\hat{\Sigma}_k$ with tapering parameter $k = n^{1/(2\alpha+1)}$ attains the optimal rate of convergence $n^{-2\alpha/(2\alpha+1)} + \frac{\log p}{n}$.

A few interesting points can be made on the optimal rate of convergence $n^{-2\alpha/(2\alpha+1)} + \frac{\log p}{n}$. When the dimension p is relatively small, that is, $\log p =$

$o(n^{1/(2\alpha+1)})$, p has no effect on the convergence rate and the rate is purely driven by the "smoothness" parameter α. However, when p is large, that is, $\log p \gg n^{1/(2\alpha+1)}$, p plays a significant role in determining the minimax rate.

We should emphasize that the optimal choice of the tapering parameter $k \asymp n^{1/(2\alpha+1)}$ is different from the optimal choice for estimating the rows/columns as vectors under mean squared error loss. Straightforward calculation shows that in the latter case the best cutoff is $k \asymp n^{1/(2(\alpha+1))}$ so that the tradeoff between the squared bias and the variance is optimal. With $k \asymp n^{1/(2\alpha+1)}$, the tapering estimator has smaller squared bias than the variance as a vector estimator of each row/column.

It is also interesting to compare our results with those given in Bickel and Levina (2008a). A banding estimator with bandwidth $k = (\frac{\log p}{n})^{1/(2(\alpha+1))}$ was proposed and the rate of convergence $(\frac{\log p}{n})^{\alpha/(\alpha+1)}$ was proved. It is easy to see that the banding estimator given in Bickel and Levina (2008a) is not rate optimal. Take, for example, $\alpha = 1/2$ and $p = e^{\sqrt{n}}$. Their rate is $n^{-1/6}$, while the optimal rate in Theorem 1 is $n^{-1/2}$.

It is instructive to take a closer look at the motivation behind the construction of the banding estimator in Bickel and Levina (2008a). Let the banding estimator be

$$(29) \qquad \hat{\Sigma}_B = (\sigma_{ij}^* I\{|i-j| \le k\})$$

and denote $\hat{\Sigma}_B - \mathbb{E}\hat{\Sigma}_B$ by V, and let $V = (v_{ij})$. An important step in the proof of Theorem 1 in Bickel and Levina (2008a) is to control the operator norm by the ℓ_1 norm as follows:

$$\mathbb{E}\|\hat{\Sigma}_B - \mathbb{E}\hat{\Sigma}_B\|^2 \le \mathbb{E}\|\hat{\Sigma}_B - \mathbb{E}\hat{\Sigma}_B\|_1^2 = \mathbb{E}\left(\max_{j=1,\dots,p}\sum_i |v_{ij}|\right)^2$$

$$\le C\left(\frac{k}{\sqrt{n}}\sqrt{\log p}\right)^2 = C\frac{k^2 \log p}{n}.$$

Note that $\mathbb{E}[|v_{ij}|I\{|i-j| \le k\}] \asymp 1/\sqrt{n}$, then $\mathbb{E}\sum_i |v_{ij}| \asymp k/\sqrt{n}$. It is then expected that $\mathbb{E}(\max_{j=1,\dots,p}\sum_i |v_{ij}|)^2 \le C(\frac{k}{\sqrt{n}}\sqrt{\log p})^2$ [see Bickel and Levina (2008a) for details] and so

$$\mathbb{E}\|\breve{\Sigma} - \Sigma\|_1^2 \le C\frac{k^2 \log p}{n} + Ck^{-2\alpha}.$$

An optimal tradeoff of k is then $(\frac{\log p}{n})^{1/(2(\alpha+1))}$ which implies a rate of $(\frac{\log p}{n})^{-\alpha/(\alpha+1)}$ in Theorem 1 in Bickel and Levina (2008a). This rate is slower than the optimal rate $n^{-2\alpha/(2\alpha+1)} + \frac{\log p}{n}$ in Theorem 1.

We have considered the parameter space $\mathcal{F}_\alpha$ defined in (3). Other similar parameter spaces can also be considered. For example, in time series analysis it is

often assumed the covariance $|\sigma_{ij}|$ decays at the rate $|i - j|^{-(\alpha+1)}$ for some $\alpha > 0$. Consider the collection of positive-definite symmetric matrices satisfying the following conditions:

(30)
$$\begin{aligned}\mathcal{G}_\alpha &= \mathcal{G}_\alpha(M_0, M_1) \\ &= \{\Sigma : |\sigma_{ij}| \leq M_1 |i - j|^{-(\alpha+1)} \text{ for } i \neq j \text{ and } \lambda_{\max}(\Sigma) \leq M_0\},\end{aligned}$$

where $\lambda_{\max}(\Sigma)$ is the maximum eigenvalues of the matrix Σ. Note that $\mathcal{G}_\alpha(M_0, M_1)$ is a subset of $\mathcal{F}_\alpha(M_0, M)$ as long as $M_1 \leq \alpha M$. Using virtually identical arguments one can show that

$$\inf_{\hat{\Sigma}} \sup_{\mathcal{P}_\alpha'} \mathbb{E}\|\hat{\Sigma} - \Sigma\|^2 \asymp n^{-2\alpha/(2\alpha+1)} + \frac{\log p}{n}.$$

Let $\mathcal{P}_\alpha' = \mathcal{P}_\alpha'(M_0, M, \rho)$ denote the set of distributions of $\mathbf{X}_1$ that satisfies (7) and (30).

REMARK 3.   Both the tapering estimator proposed in this paper and banding estimator given in Bickel and Levina (2008a) are not necessarily positive-semidefinite. A practical proposal to avoid this would be to project the estimator $\hat{\Sigma}$ to the space of positive-semidefinite matrices under the operator norm. More specifically, one may first diagonalize $\hat{\Sigma}$ and then replace negative eigenvalues by 0. The resulting estimator is then positive-semidefinite.

3.3.1. *The case of $p < n^{1/(2\alpha+1)}$.*   We have focused on the case $p \geq n^{1/(2\alpha+1)}$ in Sections 3.1 and 3.2. The case of $p < n^{1/(2\alpha+1)}$ can be handled in a similar way. The main difference is that in this case we no longer have a tapering estimator $\hat{\Sigma}_k$ with $k = n^{1/(2\alpha+1)}$ because $k > p$. Instead the maximum likelihood estimator Σ^* can be used directly. It is easy to show in this case

(31)
$$\sup_{\mathcal{P}_\alpha} \mathbb{E}\|\Sigma^* - \Sigma\|^2 \leq C\frac{p}{n}.$$

The lower bound can also be obtained by the application of Assouad's lemma and by using a parameter space that is similar to $\mathcal{F}_{11}$. To be more specific, for an integer $1 \leq m \leq p/2$, define the $p \times p$ matrix $B_m = (b_{ij})_{p \times p}$ with

$$b_{ij} = I\{i = m \text{ and } m + 1 \leq j \leq p, \text{ or } j = m \text{ and } m + 1 \leq i \leq p\}.$$

Define the collection of $2^{p/2}$ covariance matrices as

(32) $$\mathcal{F}^* = \left\{\Sigma(\theta) : \Sigma(\theta) = I_p + \tau \frac{1}{\sqrt{np}} \sum_{m=1}^{p/2} \theta_m B(m, k), \theta = (\theta_m) \in \{0, 1\}^{p/2}\right\}.$$

Since $p < n^{1/(2\alpha+1)}$, then $\frac{1}{\sqrt{np}} < 2^{\alpha+1/2} p^{-(\alpha+1)}$. Again it is easy to check $\mathcal{F}^* \subset \mathcal{F}_\alpha(M_0, M)$ when $0 < \tau < 2^{-\alpha-1}M$. The following lower bound then follows from the same argument as in Section 3.2.1:

$$(33) \qquad \inf_{\hat{\Sigma}} \sup_{\mathcal{F}^*} \mathbb{E}\|\hat{\Sigma} - \Sigma\|^2 \geq cp\left(\frac{1}{\sqrt{np}}\right)^2 \cdot \frac{p}{2} \cdot c_1 \geq c_2 \frac{p}{n}.$$

Equations (31) and (33) together yield the minimax rate of convergence for the case $p \leq n^{1/(2\alpha+1)}$,

$$(34) \qquad \inf_{\hat{\Sigma}} \sup_{\mathcal{P}_\alpha} \mathbb{E}\|\hat{\Sigma} - \Sigma\|^2 \asymp \frac{p}{n}.$$

This, together with equation (28), gives the optimal rate of convergence:

$$(35) \qquad \inf_{\hat{\Sigma}} \sup_{\mathcal{P}_\alpha} \mathbb{E}\|\hat{\Sigma} - \Sigma\|^2 \asymp \min\left\{ n^{-2\alpha/(2\alpha+1)} + \frac{\log p}{n}, \frac{p}{n} \right\}.$$

4. Rate optimality under the Frobenius norm.

In addition to the operator norm, the Frobenius norm is another commonly used matrix norm. The Frobenius norm is used in defining the numerical rank of a matrix which is useful in many applications, such as the principle component analysis. See, for example, Rudelson and Vershynin (2007). The Frobenius norm has also been used in the literature for measuring the accuracy of a covariance matrix estimator. See, for example, Lam and Fan (2007) and Ravikumar et al. (2008). In this section we consider the optimal rate of convergence for covariance matrix estimation under the Frobenius norm. The Frobenius norm of a matrix $A = (a_{ij})$ is defined as the ℓ_2 vector norm of all entries in the matrix

$$\|A\|_F = \sqrt{\sum_{i,j} a_{ij}^2}.$$

This is equivalent to treating the matrix A as a vector of length p^2. It is easy to see that the operator norm is bounded by the Frobenius norm, that is, $\|A\| \leq \|A\|_F$.

The following theorem gives the minimax rate of convergence for estimating the covariance matrix Σ under the Frobenius norm based on the sample $\{\mathbf{X}_1, \ldots, \mathbf{X}_n\}$.

THEOREM 4. *The minimax risk under the Frobenius norm satisfies*

$$(36) \qquad \begin{aligned} \inf_{\hat{\Sigma}} \sup_{\mathcal{P}_\alpha} \mathbb{E}\frac{1}{p}\|\hat{\Sigma} - \Sigma\|_F^2 &\asymp \inf_{\hat{\Sigma}} \sup_{\mathcal{P}'_\alpha} \mathbb{E}\frac{1}{p}\|\hat{\Sigma} - \Sigma\|_F^2 \\ &\asymp \min\left\{ n^{-(2\alpha+1)/(2(\alpha+1))}, \frac{p}{n} \right\}. \end{aligned}$$

We shall establish below separately the minimax upper bound and minimax lower bound.

4.1. *Upper bound under the Frobenius norm.* We will only prove the upper bound for the distribution set $\mathcal{P}'_\alpha$ given in (30). The proof for the parameter space $\mathcal{P}_\alpha$ is slightly more involved by thresholding procedures as in Wavelet estimation. The minimax upper bound is derived by again considering the tapering estimator (4). Under the Frobenius norm the risk function is separable. The risk of the tapering estimator can be bounded separately under the squared ℓ_2 loss for each row/column. This method has been commonly used in nonparametric function estimation using orthogonal basis expansions. Since

$$\mathbb{E}\tilde{\sigma}_{ij} = \sigma_{ij} \quad \text{and} \quad \text{Var}(\tilde{\sigma}_{ij}) \leq \frac{C}{n}$$

for the tapering estimator (4), we have

$$\mathbb{E}(w_{ij}\tilde{\sigma}_{ij} - \sigma_{ij})^2 \leq (1 - w_{ij})^2\sigma_{ij}^2 + w_{ij}^2\frac{C}{n}.$$

It can be seen easily that

$$\frac{1}{p}\mathbb{E}\|\check{\Sigma} - \Sigma\|_F^2 \leq \frac{1}{p}\sum_{\{(i,j):k_h<|i-j|\}}\sigma_{ij}^2 + \frac{1}{p}\sum_{\{(i,j):|i-j|\leq k\}}\left[(1 - w_{ij})^2\sigma_{ij}^2 + w_{ij}^2\frac{C}{n}\right]$$

$$\equiv R_1 + R_2.$$

The assumption $\lambda_{\max}(\Sigma) \leq M_0$ implies that $\sigma_{ii} \leq M_0$ for all i. Since $|\sigma_{ij}|$ is also uniformly bounded for all $i \neq j$ from assumption (30), we immediately have $R_2 \leq C\frac{k}{n}$.

It is easy to show that

$$(37) \qquad \frac{1}{p}\sum_{\{(i,j):k<|i-j|\}}\sigma_{ij}^2 \leq Ck^{-2\alpha-1},$$

where $|\sigma_{ij}| \leq C_1|i-j|^{-(\alpha+1)}$ for all $i \neq j$. Thus

$$(38) \qquad \mathbb{E}\frac{1}{p}\|\check{\Sigma} - \Sigma\|_F^2 \leq Ck^{-2\alpha-1} + C\frac{k}{n} \leq C_2 n^{-(2\alpha+1)/(2(\alpha+1))}$$

by choosing

$$(39) \qquad k = n^{1/(2(\alpha+1))}$$

if $n^{1/(2(\alpha+1))} \leq p$, which is different from the choice of k for the operator norm in (15). If $n^{1/(2(\alpha+1))} > p$, we will choose $k = p$, then the bias part is 0 and consequently

$$\mathbb{E}\frac{1}{p}\|\check{\Sigma} - \Sigma\|_F^2 \leq C\frac{p}{n}.$$

REMARK 4. For the parameter space $\mathcal{P}'_\alpha$, under the Frobenius norm the optimal tapering parameter k is of the order $n^{1/(2(\alpha+1))}$. The rate of convergence of the

tapering estimator with $k \asymp n^{1/(2(\alpha+1))}$ under the operator norm is

$$\frac{\log p}{n} + n^{-\alpha/(\alpha+1)},$$

which is slower than $n^{-2\alpha/(2\alpha+1)} + \frac{\log p}{n}$ in (1). Similarly, the optimal procedure under the operator norm is not rate optimal under the Frobenius norm. Therefore, the optimal choice of the tapering parameter k critically depends on the norm under which the estimation accuracy is measured.

REMARK 5. Similarly for $\mathcal{P}'_\alpha$, it can be shown that under the Frobenius norm the banding estimator with $k \asymp n^{1/(2(\alpha+1))}$ is rate optimal. Under the operator norm, Bickel and Levina (2008a) chose $k \asymp (\frac{\log p}{n})^{1/(2(\alpha+1))}$ for the banding estimator which is close to $n^{1/(2(\alpha+1))}$ up to a logarithmic factor of p. On the other hand, it can be shown that for the parameter space $\mathcal{P}_\alpha$ no linear estimator can achieve the optimal convergence rate under the Frobenius norm.

4.2. *Lower bound under the Frobenius norm.* It is sufficient to establish the lower bound for the parameter space $\mathcal{P}'_\alpha$ given in (30). Again the argument for $\mathcal{P}_\alpha$ is similar. As in the case of estimation under the operator norm, we need to construct a finite collection of multivariate normal distributions with a parameter space $\mathcal{G}_2 \subset \mathcal{G}_\alpha$ such that

$$\inf_{\hat{\Sigma}} \sup_{\mathcal{G}_2} \mathbb{E} \frac{1}{p} \|\hat{\Sigma} - \Sigma\|_F^2 \geq c\frac{k}{n}$$

for some $c > 0$ when $k = \min\{n^{1/(2(\alpha+1))}, p/2\}$.

We construct $\mathcal{G}_2$ as follows. Let $0 < \tau < M$ be a constant. Define

$$\mathcal{G}_2 = \{\Sigma(\theta) : \Sigma(\theta) = I + (\theta_{ij}\tau n^{-1/2}I\{1 \leq |i-j| \leq k\})_{p \times p},$$

$$\text{for } \theta_{ij} = \theta_{ji} = 0 \text{ or } 1\}.$$

It is easy to verify that $\mathcal{G}_2 \subset \mathcal{G}_\alpha$ as $n \to \infty$. Note that $\theta \in \Theta = \{0, 1\}^{kp-k(k+1)/2}$.

Applying Assouad's lemma with d the Frobenius norm and $s = 2$ to the parameter space $\mathcal{G}_2$, we have

$$\max_{\theta \in \mathcal{G}_2} 2^2 E_\theta \frac{1}{p} \|\hat{\Sigma} - \Sigma(\theta)\|_F^2$$

$$\geq \min_{H(\theta,\theta')\geq 1} \frac{1/p\|\Sigma(\theta) - \Sigma(\theta')\|_F^2}{H(\theta,\theta')} \frac{kp - k(k+1)/2}{2} \min_{H(\theta,\theta')=1} \|P_\theta \wedge P_{\theta'}\|.$$

Note that

$$\min_{H(\theta,\theta')\geq 1} \frac{1}{p} \frac{\|\Sigma(\theta) - \Sigma(\theta')\|_F^2}{H(\theta,\theta')} = \min_{H(\theta,\theta')\geq 1} \frac{1}{p} \frac{[\tau n^{-1/2}]^2 \sum |\theta_{ij} - \theta'_{ij}|^2}{H(\theta,\theta')}$$

$$= \frac{\tau^2}{p}n^{-1}.$$

It is easy to see that

$$\frac{kp - k(k+1)/2}{2} \asymp kp.$$

LEMMA 8. *Let P_θ be the joint distribution of $\mathbf{X}_1, \ldots, \mathbf{X}_n \overset{i.i.d.}{\sim} N(0, \Sigma(\theta))$ with $\Sigma(\theta) \in \mathcal{G}_2$. Then for some constant $c_1 > 0$ we have*

$$\min_{H(\theta, \theta')=1} \| P_\theta \wedge P_{\theta'} \| \geq c_1.$$

We omit the proof of this lemma. It is very similar to and simpler than the proof of Lemma 6.

From Lemma 8 we have for some $c > 0$

$$(40) \qquad \min_{H(\theta, \theta')=1} \| P_\theta \wedge P_{\theta'} \| \geq c$$

thus

$$\max_{\theta \in \mathcal{G}_2} 2^2 E_\theta \frac{1}{p} \| \hat{\Sigma} - \Sigma(\theta) \|_F^2 \geq c \min \left\{ n^{-(2\alpha+1)/(2(\alpha+1))}, \frac{p}{n} \right\},$$

which implies that the rate obtained in (38) is optimal.

5. Estimation of the inverse covariance matrix. The inverse of the covariance matrix Σ^{-1} is of significant interest in many statistical applications. The results and analysis given in Section 3 can be used to derive the optimal rate of convergence for estimating Σ^{-1} under the operator norm.

For estimating the inverse covariance matrix Σ^{-1} we require the minimum eigenvalue of Σ to be bounded away from zero. For $\delta > 0$, we define

$$(41) \qquad L_\delta = \{ \Sigma : \lambda_{\min}(\Sigma) \geq \delta \}.$$

Let $\tilde{\mathcal{P}}_\alpha = \tilde{\mathcal{P}}_\alpha(M_0, M, \rho, \delta)$ denote the set of distributions of $\mathbf{X}_1$ that satisfy (3), (7) and (41), and similarly, distributions in $\tilde{\mathcal{P}}_\alpha' = \tilde{\mathcal{P}}_\alpha'(M_0, M, \rho, \delta)$ satisfy (7), (30) and (41).

The following theorem gives the minimax rate of convergence for estimating Σ^{-1}.

THEOREM 5. *The minimax risk of estimating the inverse covariance matrix Σ^{-1} satisfies*

$$(42) \qquad \inf_{\hat{\Sigma}} \sup_{\tilde{\mathcal{P}}} \mathbb{E} \| \hat{\Sigma}^{-1} - \Sigma^{-1} \|^2 \asymp \min \left\{ n^{-2\alpha/(2\alpha+1)} + \frac{\log p}{n}, \frac{p}{n} \right\},$$

where $\tilde{\mathcal{P}}$ denotes either $\tilde{\mathcal{P}}_\alpha$ or $\tilde{\mathcal{P}}_\alpha'$.

PROOF. We shall focus on the case $p \geq n^{1/(2\alpha+1)}$. The proof for the case of $p < n^{1/(2\alpha+1)}$ is similar. To establish the upper bound, note that

$$\hat{\Sigma}^{-1} - \Sigma^{-1} = \hat{\Sigma}^{-1}(\Sigma - \hat{\Sigma})\Sigma^{-1},$$

then

$$\|\hat{\Sigma}^{-1} - \Sigma^{-1}\|^2 = \|\hat{\Sigma}^{-1}(\Sigma - \hat{\Sigma})\Sigma^{-1}\|^2 \leq \|\hat{\Sigma}^{-1}\|^2 \|\Sigma - \hat{\Sigma}\|^2 \|\Sigma^{-1}\|^2.$$

It follows from assumption (3) that $\|\Sigma^{-1}\|^2 \leq \delta^{-2}$. Note that $\mathbb{P}\{\|\check{\Sigma} - \mathbb{E}\check{\Sigma}\|^2 > \epsilon\} \leq 4p5^m \exp(-n\epsilon^2 \rho_1)$ for any $\epsilon > 0$ which decays faster than any polynomial of n as shown in the proof of Lemmas 2 and 3. Let $\lambda_{\min}(\check{\Sigma})$ and $\lambda_{\min}(\mathbb{E}\check{\Sigma})$ be the smallest eigenvalues of $\check{\Sigma}$ and $\mathbb{E}\check{\Sigma}$, respectively. Then $\mathbb{P}(\lambda_{\min}(\check{\Sigma}) \leq \lambda_{\min}(\mathbb{E}\check{\Sigma}) - \epsilon^{1/2}) \geq \mathbb{P}(|\lambda_{\min}(\check{\Sigma}) - \lambda_{\min}(\mathbb{E}\check{\Sigma})| \geq \epsilon^{1/2})$ decays faster than any polynomial of n. Let $0 < \epsilon < [\lambda_{\min}(\mathbb{E}\check{\Sigma})/2]^2$ and $c = 1/[\lambda_{\min}(\mathbb{E}\check{\Sigma}) - \epsilon^{1/2}]$, then $\mathbb{P}(\|\hat{\Sigma}^{-1}\| \geq c)$ decays faster than any polynomial of n. Therefore,

$$\mathbb{E}\|\hat{\Sigma}^{-1} - \Sigma^{-1}\|^2 \leq \left(\frac{c}{\delta}\right)^2 \mathbb{E}\|\Sigma - \hat{\Sigma}\|^2$$

$$+ \mathbb{E}[\|\hat{\Sigma}^{-1}\|^2 \|\Sigma - \hat{\Sigma}\|^2 \|\Sigma^{-1}\|^2 I(\|\hat{\Sigma}^{-1}\| \geq c)]$$

$$\leq C \min\left\{ n^{-2\alpha/(2\alpha+1)} + \frac{\log p}{n}, \frac{p}{n} \right\}.$$

The proof of the lower bound is almost identical to that of Theorem 1 except that here we need to show

$$\min_{H(\theta, \theta') \geq 1} \frac{\|\Sigma^{-1}(\theta) - \Sigma^{-1}(\theta')\|^2}{H(\theta, \theta')} \geq cka^2$$

instead of Lemma 5. For a positive definite matrix A, let $\lambda_{\min}(A)$ denote the minimum eigenvalue of A. Since

$$\Sigma^{-1}(\theta) - \Sigma^{-1}(\theta') = \Sigma^{-1}(\theta')(\Sigma(\theta) - \Sigma(\theta'))\Sigma^{-1}(\theta),$$

we have

$$\|\Sigma^{-1}(\theta) - \Sigma^{-1}(\theta')\| \geq \lambda_{\min}(\Sigma^{-1}(\theta))\lambda_{\min}(\Sigma^{-1}(\theta'))\|\Sigma(\theta) - \Sigma(\theta')\|.$$

Note that

$$\lambda_{\min}(\Sigma^{-1}(\theta)) > 1/M_0, \qquad \lambda_{\min}(\Sigma^{-1}(\theta')) > 1/M_0,$$

then Lemma 5 implies

$$\min_{H(\theta, \theta') \geq 1} \frac{\|\Sigma^{-1}(\theta) - \Sigma^{-1}(\theta')\|^2}{H(\theta, \theta')} \geq M_0^{-4} \min_{H(\theta, \theta') \geq 1} \frac{\|\Sigma(\theta) - \Sigma(\theta')\|^2}{H(\theta, \theta')} \geq cka^2$$

for some constant $c > 0$. □

6. Simulation study. We now turn to the numerical performance of the proposed tapering estimator and compare it with that of the banding estimator of Bickel and Levina (2008a). In the numerical study, we shall consider estimating a covariance matrix in the parameter space $\mathcal{F}_\alpha$ defined in (3). Specifically, we consider the covariance matrix $\Sigma = (\sigma_{ij})_{1 \le i, j \le p}$ of the form

$$(43) \qquad \sigma_{ij} = \begin{cases} 1, & 1 \le i = j \le p, \\ \rho |i - j|^{-(\alpha+1)}, & 1 \le i \neq j \le p. \end{cases}$$

Note that this is a Toeplitz matrix. But we do not assume that the structure is known and do not use the information in any estimation procedure.

The banding estimator in (29) depends on the choice of k. An optimal tradeoff of k is $k \asymp (n/\log p)^{1/(2\alpha+2)}$ as discussed in Section 3.3. See Bickel and Levina (2008a). The tapering estimator (6) also depends on k for which the optimal tradeoff is $k \asymp n^{1/(2\alpha+1)}$. In our simulation study, we choose $k = \lfloor (n/\log p)^{1/(2\alpha+2)} \rfloor$ for the banding estimator and $k = \lfloor n^{1/(2\alpha+1)} \rfloor$ for the tapering estimator.

A range of parameter values for α, n and p are considered. Specifically, α ranges from 0.1 to 0.5, the sample size n ranges from 250 to 3000 and the dimension p goes from 250 to 3000. We choose the value of ρ to be $\rho = 0.6$ so that all matrices are nonnegative definite and their smallest eigenvalues are close to 0. Table 1 reports the average errors under the spectral norm over 100 replications for the two procedures. The cases where the tapering estimator underperforms the banding estimator are highlighted in boldface. Figure 2 plots the ratios of the average errors of the banding estimator to the corresponding average errors of the tapering estimator for $\alpha = 0.1, 0.2, 0.3$ and 0.5. The case of $\alpha = 0.4$ is similar to the case of $\alpha = 0.3$.

It can be seen from Table 1 and Figure 2 that the tapering estimator outperforms the banding estimator in 121 out of 125 cases. For the given dimension p, the ratio of the average error of the banding estimator to the corresponding average error of the tapering estimator tends to increase as the sample size n increases. The tapering estimator fails to outperform the banding estimator only when $\alpha = 0.5$ and $n = 250$ in which case the values of k are small for both estimators.

REMARK 6. We have also carried out additional simulations for larger values of α with the same sample sizes and dimensions. The performance of the tapering and abnding estimators are similar. This is mainly dur to the fact that the values of k for both estimators are very small for large α when n and p are only moderately large.

7. Proofs of auxiliary lemmas. In this section we give proofs of auxiliary lemmas stated and used in Sections 3–5.

PROOF OF LEMMA 1. Without loss of generality we assume that $i \le j$. The set $\{i, j\}$ is contained in the set $\{l, \dots, l + k_h - 1\}$ if and only if $l \le i \le j \le l + k_h - 1$, that is, $j - k_h + 1 \le l \le i$. Note that Card$\{l : j - k_h + 1 \le l \le i\} =$

TABLE 1
The average errors under the spectral norm of the banding estimator (BL) and the tapering estimator (CZZ) over 100 replications. The cases where the tapering estimator underperforms the banding estimator are highlighted in italic

p	n	$\alpha = 0.1$		$\alpha = 0.2$		$\alpha = 0.3$		$\alpha = 0.4$		$\alpha = 0.5$	
		BL	CZZ	BL	CZZ	BL	CZZ	BL	CZZ	BL	CZZ
250	250	2.781	2.706	2.291	2.023	1.762	1.684	1.618	1.517	*1.325*	*1.507*
	500	2.409	2.302	1.898	1.575	1.562	1.204	1.361	1.185	1.080	0.822
	1000	2.029	1.685	1.631	1.361	1.289	1.018	1.056	0.795	0.911	0.859
	2000	1.706	1.153	1.369	1.122	1.106	0.908	0.878	0.655	0.715	0.542
	3000	1.522	0.926	1.242	0.896	0.983	0.798	0.810	0.658	0.645	0.482
500	250	3.277	2.914	2.609	2.097	1.961	1.788	1.745	1.610	*1.392*	*1.571*
	500	2.901	2.598	2.199	1.683	1.751	1.256	1.475	1.234	1.152	0.865
	1000	2.539	2.197	1.942	1.472	1.481	1.064	1.178	0.843	0.984	0.917
	2000	2.263	1.726	1.669	1.326	1.293	0.965	1.067	0.700	0.866	0.569
	3000	2.066	1.379	1.538	1.154	1.220	0.874	0.919	0.696	0.781	0.503
1000	250	3.747	3.086	2.873	2.223	2.385	1.842	1.833	1.694	*1.449*	*1.643*
	500	3.370	2.735	2.635	1.768	1.906	1.334	1.565	1.297	1.203	0.925
	1000	3.097	2.437	2.315	1.536	1.741	1.121	1.382	0.883	1.037	0.936
	2000	2.730	2.177	2.011	1.392	1.523	1.006	1.156	0.722	0.920	0.591
	3000	2.589	1.968	1.865	1.264	1.374	0.911	1.072	0.723	0.834	0.523
2000	250	4.438	3.177	3.107	2.300	2.511	1.956	1.903	1.744	*1.484*	*1.736*
	500	3.969	2.800	2.868	1.841	2.030	1.383	1.638	1.356	1.239	0.940
	1000	3.538	2.531	2.551	1.599	1.866	1.158	1.452	0.912	1.074	0.973
	2000	3.242	2.353	2.248	1.434	1.649	1.031	1.224	0.751	0.955	0.611
	3000	3.025	2.219	2.101	1.302	1.566	0.929	1.141	0.743	0.868	0.541
3000	250	4.679	3.219	3.230	2.358	2.576	1.995	1.931	1.797	*1.494*	*1.776*
	500	4.214	2.887	2.991	1.890	2.282	1.419	1.664	1.384	1.463	0.971
	1000	3.901	2.575	2.674	1.633	1.933	1.186	1.482	0.929	1.224	0.990
	2000	3.488	2.395	2.452	1.451	1.717	1.049	1.254	0.768	0.965	0.619
	3000	3.336	2.278	2.288	1.321	1.632	0.948	1.172	0.750	0.880	0.549

$(i - (j - k_h + 1) + 1)_+ = (k_h - |i - j|)_+$, then $\mathrm{Card}\{l : \{i, j\} \subset \{l, \ldots, l + k_h - 1\}\} = (k_h - |i - j|)_+$. Similarly, we have $\mathrm{Card}\{l : \{i, j\} \subset \{l, \ldots, l + k - 1\}\} = (k - |i - j|)_+$. Thus we have

$$kw_{ij} = (k - |i - j|)_+ - (k_h - |i - j|)_+$$
$$= \mathrm{Card}\{l : \{i, j\} \subset \{l, \ldots, l + k - 1\}\}$$
$$- \mathrm{Card}\{l : \{i, j\} \subset \{l, \ldots, l + k_h - 1\}\}. \qquad \square$$

PROOF OF LEMMA 2.   Without loss of generality we assume that p is divisible by m. Recall that $M_l^{(m)} = (\tilde{\sigma}_{ij} I\{l \leq i < l + m, l \leq j < l + m\})_{p \times p}$. Note that $M_l^{(m)}$
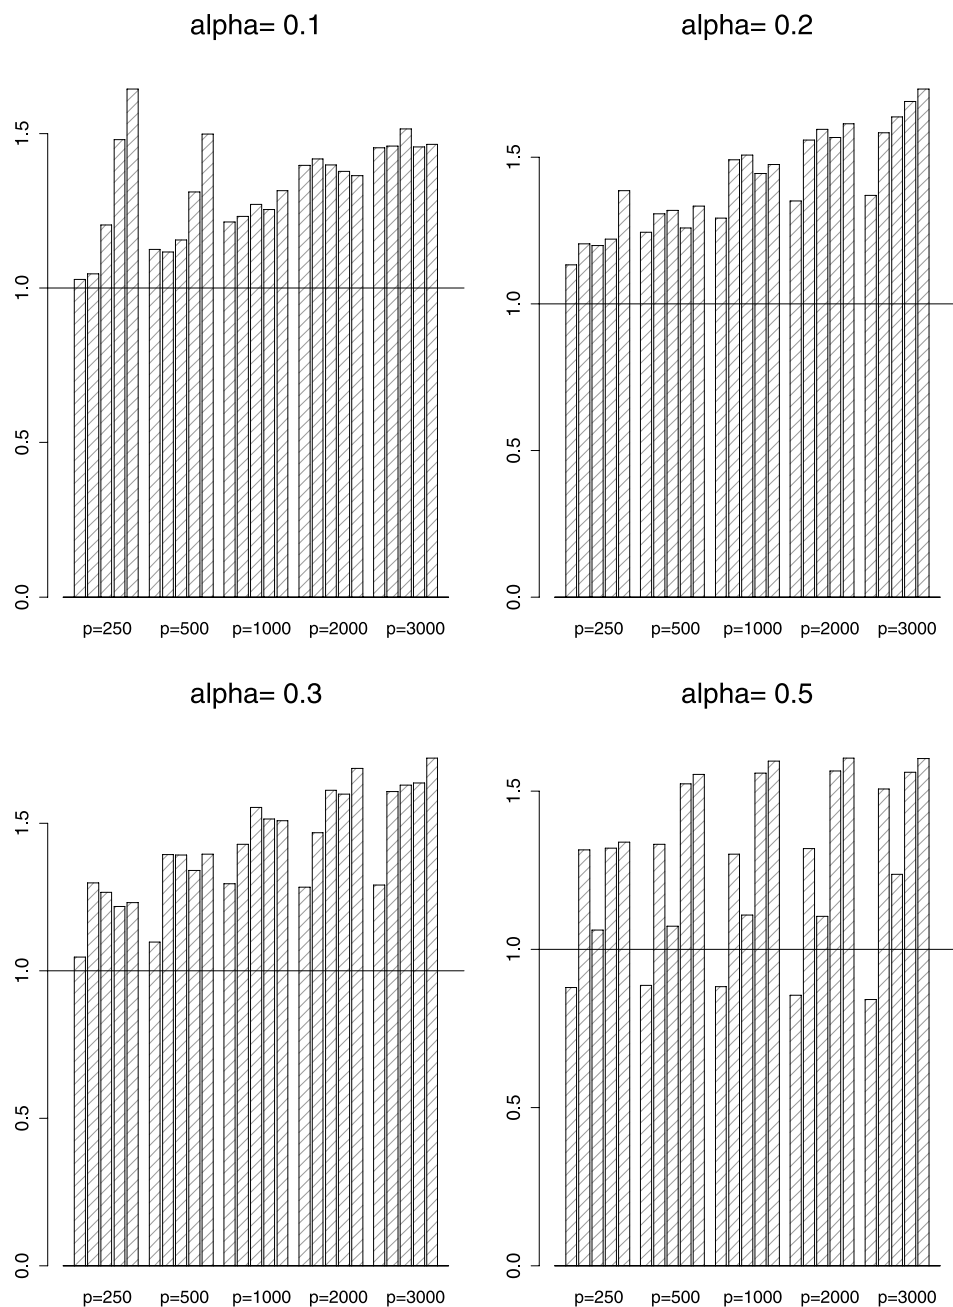
FIG. 2. *The vertical bars represent the ratios of the average error of the banding estimator to the corresponding average error of the tapering estimator. The higher the bar the better the relative performance of the tapering estimator. For each value of p the bars are ordered from left to right by the sample sizes (n = 250 to 3000).*

is empty when $l \leq 1 - m$, and has at least one nonzero entry when $l \geq 2 - m$. Set $\delta_l^{(m)} = M_l^{(m)} - \mathbb{E}M_l^{(m)}$ and $S^{(m)} = \sum_{l=2-m}^{p} M_l^{(m)}$. It follows from (6) that

$$
(44) \qquad \left\| S^{(m)} - \mathbb{E}S^{(m)} \right\| \leq \sum_{l=1}^{m} \left\| \sum_{-1 \leq j < p/m} \delta_{jm+l}^{(m)} \right\|.
$$

Since $\delta_{jm+l}^{(m)}$ are disjoint diagonal blocks over $-1 \leq j < p/m$, we have

$$
\left\| S^{(m)} - \mathbb{E}S^{(m)} \right\| \leq m \max_{1 \leq l \leq m} \left\| \sum_{-1 \leq j < p/m} \delta_{jm+l}^{(m)} \right\|
$$

$$
(45) \qquad\qquad\qquad\qquad \leq m \max_{1-m \leq l \leq p} \left\| \delta_l^{(m)} \right\|.
$$

Since $\delta_l^{(k_h)}$ and $\delta_l^{(k)}$ are all sub-blocks of certain matrix $\delta_l^{(k)}$ with $1 \leq l \leq p - k + 1$, Lemma 2 now follows immediately from equations (45) and (6). $\quad\square$

PROOF OF LEMMA 3.   For any $m \times m$ symmetric matrix A, we have

$$
|u^T A u| - |v^T A v| \leq |u^T A u - v^T A v| = |(u - v)^T A (u + v)|
$$

$$
\leq \|u - v\| \|A\| \|u + v\|.
$$

Let $S_{1/2}^{m-1}$ be a $1/2$ net of the unit sphere S^{m-1} in the Euclidean distance in R^m. We have

$$
\|A\| \leq \sup_{u \in S^{m-1}} |u^T A u| \leq \sup_{u \in S_{1/2}^{m-1}} |u^T A u| + \frac{1}{2} \|A\| \frac{3}{2}
$$

$$
= \sup_{u \in S_{1/2}^{m-1}} |u^T A u| + \frac{3}{4} \|A\|,
$$

which implies $\|A\| \leq 4 \sup_{u \in S_{1/2}^{m-1}} |u^T A u|$. Since we are allowed to pack $\mathrm{Card}(S_{1/2}^{m-1})$ balls of radius $1/4$ into a $1 + 1/4$ ball in R^m, volume comparison yields

$$
(1/4)^m \, \mathrm{Card}(S_{1/2}^{m-1}) \leq (5/4)^m,
$$

that is, $\mathrm{Card}(S_{1/2}^{m-1}) \leq 5^m$. Thus there exist $\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_{5^m} \in S^{m-1}$ such that

$$
\|A\| \leq 4 \sup_{j \leq 5^m} |v_j^T A v_j| \qquad \text{for all } m \times m \text{ symmetric } A.
$$

This one-step approximation argument is similar to the proof of Proposition 4.2(ii) in Zhang and Huang (2008).

Let $\mathbf{X}_1, \ldots, \mathbf{X}_n$ be i.i.d. p-vectors with $\mathbb{E}(\mathbf{X}_1 - \mu)(\mathbf{X}_1 - \mu)^T = \Sigma$. Under the sub-Gaussian assumption in (7) there exists $\rho > 0$ such that

$$\mathbb{P}\{\mathbf{v}^T (\mathbf{X}_i - \mathbb{E}\mathbf{X}_i)(\mathbf{X}_i - \mathbb{E}\mathbf{X}_i)^T \mathbf{v} > x\} \le e^{-x\rho/2} \qquad \text{for all } x > 0 \text{ and } \|\mathbf{v}\| = 1,$$

which implies $\mathbb{E}(t\mathbf{v}^T (\mathbf{X}_i - \mathbb{E}\mathbf{X}_i)(\mathbf{X}_i - \mathbb{E}\mathbf{X}_i)^T \mathbf{v}) < \infty$ for all $t < \rho/2$ and $\|\mathbf{v}\| = 1$, then there exists $\rho_1 > 0$ such that

$$\mathbb{P}\left\{\left|\frac{1}{n}\sum_{i=1}^n \mathbf{v}^T [(\mathbf{X}_i - \mathbb{E}\mathbf{X}_i)(\mathbf{X}_i - \mathbb{E}\mathbf{X}_i)^T - \Sigma]\mathbf{v}\right| > x\right\} \le e^{-nx^2\rho_1/2}$$

for all $0 < x < \rho_1$ and $\|\mathbf{v}\| = 1$. [See, e.g., Chapter 2 in Saulis and Statulevičius (1991).] Thus we have

$$\mathbb{P}\left\{\max_{1 \le l \le p-m+1} \|M_l^{(m)} - \mathbb{E}M_l^{(m)}\| > x\right\}$$

$$\le \sum_{1 \le l \le p-m+1} \mathbb{P}\{\|M_l^{(m)} - \mathbb{E}M_l^{(m)}\| > x\}$$

$$\le 2p5^m \sup_{\mathbf{v}_j, l} \mathbb{P}\{|\mathbf{v}_j^T (M_l^{(m)} - \mathbb{E}M_l^{(m)})\mathbf{v}_j| > x\}$$

$$\le 2p5^m \exp(-nx^2\rho_1/2). \qquad \square$$

PROOF OF LEMMA 5. Set $v = (1\{k_h \le i \le k\})$ and let

$$(w_i) = [\Sigma(\theta) - \Sigma(\theta')]v.$$

Note that there are exactly $H(\theta, \theta')$ number of w_i such that $|w_i| = \tau k_h a$, and $\|v\|_2^2 = k_h$. This implies

$$\|\Sigma(\theta) - \Sigma(\theta')\|^2 \ge \frac{\|[\Sigma(\theta) - \Sigma(\theta')]v\|_2^2}{\|v\|_2^2} \ge \frac{H(\theta, \theta') \cdot (\tau k a)^2}{k_h}$$

$$= H(\theta, \theta') \cdot \tau^2 k_h a^2. \qquad \square$$

PROOF OF LEMMA 6. When $H(\theta, \theta') = 1$, we will show

$$\|P_{\theta'} - P_\theta\|_1^2 \le 2K(P_{\theta'}|P_\theta)$$

$$= 2n\left[\frac{1}{2}\operatorname{tr}(\Sigma(\theta')\Sigma^{-1}(\theta)) - \frac{1}{2}\log\det(\Sigma(\theta')\Sigma^{-1}(\theta)) - \frac{p}{2}\right]$$

$$\le n \cdot cka^2$$

for some small $c > 0$, where $K(\cdot|\cdot)$ is the Kullback–Leibler divergence and the first inequality follows from the well-known Pinsker's inequality [see, e.g., Csiszár

(1967)]. This immediately implies the L_1 distance between two measures is bounded away from 1, and then the lemma follows. Write

$$\Sigma(\theta') = D_1 + \Sigma(\theta).$$

Then

$$\frac{1}{2}\operatorname{tr}(\Sigma(\theta')\Sigma^{-1}(\theta)) - \frac{p}{2} = \frac{1}{2}\operatorname{tr}(D_1\Sigma^{-1}(\theta)).$$

Let λ_i be the eigenvalues of $D_1\Sigma^{-1}(\theta)$. Since $D_1\Sigma^{-1}(\theta)$ is similar to the symmetric matrix $\Sigma^{-1/2}(\theta)D_1\Sigma^{-1/2}(\theta)$, and

$$\|\Sigma^{-1/2}(\theta)D_1\Sigma^{-1/2}(\theta)\| \leq \|\Sigma^{-1/2}(\theta)\|\|D_1\|\|\Sigma^{-1/2}(\theta)\|$$

$$\leq c_1\|D_1\| \leq c_1\|D_1\|_1 \leq c_2 ka,$$

then all eigenvalues λ_i's are real and in the interval $[-c_2 ka, c_2 ka]$, where $ka = k \cdot k^{-(\alpha+1)} = k^{-\alpha} \to 0$. Note that the Taylor expansion yields

$$\log\det(\Sigma(\theta')\Sigma^{-1}(\theta)) = \log\det(I + D_1\Sigma^{-1}(\theta)) = \operatorname{tr}(D_1\Sigma^{-1}(\theta)) - R_3,$$

where

$$R_3 \leq c_3\sum_{i=1}^{p}\lambda_i^2 \qquad \text{for some } c_3 > 0.$$

Write $\Sigma^{-1/2}(\theta) = UV^{1/2}U^T$, where $UU^T = I$ and V is a diagonal matrix. It follows from the fact that the Frobenius norm of a matrix remains the same after an orthogonal transformation that

$$\sum_{i=1}^{p}\lambda_i^2 = \|\Sigma^{-1/2}(\theta)D_1\Sigma^{-1/2}(\theta)\|_F^2 \leq \|V\|^2 \cdot \|U^T D_1 U\|_F^2$$

$$= \|\Sigma^{-1}(\theta)\|^2 \cdot \|D_1\|_F^2 \leq c_4 ka^2. \qquad \square$$

REFERENCES

ASSOUAD, P. (1983). Deux remarques sur l'estimation. *C. R. Acad. Sci. Paris Sér. I Math.* **296** 1021–1024. MR0777600

BICKEL, P. J. and LEVINA, E. (2008a). Regularized estimation of large covariance matrices. *Ann. Statist.* **36** 199–227. MR2387969

BICKEL, P. J. and LEVINA, E. (2008b). Covariance regularization by thresholding. *Ann. Statist.* **36** 2577–2604. MR2485008

CSISZÁR, I. (1967). Information-type measures of difference of probability distributions and indirect observation. *Studia Sci. Math. Hungar.* **2** 229–318. MR0219345

EL KAROUI, N. (2008). Operator norm consistent estimation of large dimensional sparse covariance matrices. *Ann. Statist.* **36** 2717–2756. MR2485011

GOLUB, G. H. and VAN LOAN, C. F. (1983). *Matrix Computations*. John Hopkins Univ. Press, Baltimore. MR0733103

FAN, J., FAN, Y. and LV, J. (2008). High dimensional covariance matrix estimation using a factor model. *J. Econometrics* **147** 186–197. MR2472991

FURRER, R. and BENGTSSON, T. (2007). Estimation of high-dimensional prior and posterior covariance matrices in Kalman filter variants. *J. Multivariate Anal.* **98** 227–255. MR2301751

HUANG, J., LIU, N., POURAHMADI, M. and LIU, L. (2006). Covariance matrix selection and estimation via penalised normal likelihood. *Biometrika* **93** 85–98. MR2277742

JOHNSTONE, I. M. (2001). On the distribution of the largest eigenvalue in principal components analysis. *Ann. Statist.* **29** 295–327. MR1863961

JOHNSTONE, I. M. and LU, A. Y. (2009). On consistency and sparsity for principal components analysis in high dimensions. *J. Amer. Statist. Assoc.* **104** 682–693.

LAM, C. and FAN, J. (2007). Sparsistency and rates of convergence in large covariance matrices estimation. Technical report, Princeton Univ.

MUIRHEAD, R. J. (1987). Developments in eigenvalue estimation. In *Advances in Multivariate Statistical Analysis* (A. K. Gupta, ed.) 277–288. Reidel, Dordrecht. MR0920436

RAVIKUMAR, P., WAINWRIGHT, M. J., RASKUTTI, G. and YU, B. (2008). High-dimensional covariance estimation by minimizing l_1-penalized log-determinant divergence. Technical report, Univ. California, Berkeley.

ROTHMAN, A. J., BICKEL, P. J., LEVINA, E. and ZHU, J. (2008). Sparse permutation invariant covariance estimation. *Electron. J. Stat.* **2** 494–515. MR2417391

RUDELSON, M. and VERSHYNIN, R. (2007). Sampling from large matrices: An approach through geometric functional analysis. *J. ACM* **54** Art. 21, 19 pp. (electronic). MR2351844

SAULIS, L. and STATULEVIČIUS, V. A. (1991). *Limit Theorems for Large Deviations.* Springer, Berlin.

WU, W. B. and POURAHMADI, M. (2009). Banding sample covariance matrices of stationary processes. *Statist. Sinica* **19** 1755–1768.

YU, B. (1997). Assouad, Fano and Le Cam. In *Festschrift for Lucien Le Cam* (D. Pollard, E. Torgersen and G. Yang, eds.) 423–435. Springer, Berlin. MR1462963

ZHANG, C.-H. and HUANG, J. (2008). The sparsity and bias of the Lasso selection in high-dimensional linear regression. *Ann. Statist.* **36** 1567–1594. MR2435448

ZOU, H., HASTIE, T. and TIBSHIRANI, R. (2006). Sparse principal components analysis. *J. Comput. Graph. Statist.* **15** 265–286. MR2252527

T. T. CAI
DEPARTMENT OF STATISTICS
THE WHARTON SCHOOL
UNIVERSITY OF PENNSYLVANIA
PHILADELPHIA, PENNSYLVANIA 19104-6302
USA
E-MAIL: tcai@wharton.upenn.edu

C.-H. ZHANG
DEPARTMENT OF STATISTICS
504 HILL CENTER
BUSCH CAMPUS
RUTGERS UNIVERSITY
PISCATAWAY, NEW JERSEY 08854-8019
USA
E-MAIL: cunhui@stat.rutgers.edu

H. H. ZHOU
DEPARTMENT OF STATISTICS
YALE UNIVERSITY
P.O. BOX 208290
NEW HAVEN, CONNECTICUT 06520-8290
USA
E-MAIL: huibin.zhou@yale.edu