



Published in final edited form as:

J Stat Plan Inference. 2021 July ; 213: 16–32. doi:10.1016/j.jspi.2020.10.006.

An Orthogonally Equivariant Estimator of the Covariance Matrix in High Dimensions and for Small Sample Sizes

Samprit Banerjee,

Division of Biostatistics, Weill Medical College of Cornell University

Stefano Monni

Department of Mathematics, American University of Beirut

Abstract

We introduce an estimation method of covariance matrices in a high-dimensional setting, i.e., when the dimension of the matrix, p , is larger than the sample size n . Specifically, we propose an orthogonally equivariant estimator. The eigenvectors of such estimator are the same as those of the sample covariance matrix. The eigenvalue estimates are obtained from an adjusted profile likelihood function derived by approximating the integral of the density function of the sample covariance matrix over its eigenvectors, which is a challenging problem in its own right. Exact solutions to the approximate likelihood equations are obtained and employed to construct estimates that involve a tuning parameter. Bootstrap and cross-validation based algorithms are proposed to choose this tuning parameter under various loss functions. Finally, comparisons with two well-known orthogonally equivariant estimators are given, which are based on Monte-Carlo risk estimates for simulated data and misclassification errors in real data analyses.

Keywords

adjusted profile likelihood; high-dimensional inference; covariance matrix estimation; singular Wishart distribution

1. Introduction

Many multivariate methods require an estimate of the covariance matrix. In this paper, we are interested in the problem of estimating the covariance matrix of a multivariate normal distribution, $\mathcal{N}(0, \Sigma)$, using a sample of mutually independent draws X_1, \dots, X_m from it, when n is less than the dimension p of Σ . This problem has received much attention in the recent past because of an increasing number of applications where measurements are collected on a large number of variables, often greater than the available experimental units. The sample covariance matrix is not a good estimator in this case. In the general framework where both p and n go to infinity in such a way that their ratio p/n converges to a positive

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

finite constant (often referred to as the large-dimensional asymptotic regime), the sample covariance matrix, its eigenvalues and its eigenvectors cease to be consistent. Some alternative estimators have thus been proposed in the literature. Ledoit and Wolf (2015) propose estimators of the eigenvalues of the covariance matrix in the large-dimensional framework that are consistent, in the sense that the mean squared deviation between the estimated eigenvalues and the population eigenvalues converges to zero almost surely. Their method is based on a particular discretization of a version of the Marcenko-Pastur equation that links the limiting spectral distribution of the sample eigenvalues and that of the population eigenvalues (Ledoit and Wolf, 2017). This method is then used to derive estimators of the covariance matrix itself that are asymptotically optimal with respect to a given loss function in the space of orthogonally equivariant estimators (Ledoit and Wolf, 2018). Additional results by these same authors have very recently appeared when our paper was under review (Ledoit and Wolf, 2020). Estimators that are derived in the large-dimensional asymptotic regime are proposed also by El Karoui (2008), Mestre (2008), Yao et al. (2012), among others. Estimators that deal with the case $p > n$ and are derived in a decision-theoretic framework are those of Konno (2009), and, more recently, Tsukuma (2016). There is a vast literature on estimation of Σ where structural assumptions on Σ are made such as ordering or sparsity, for example Bickel and Levina (2008); Bien and Tibshirani (2011); Naul and Taylor (2017); Won et al. (2013).

In this paper, we propose an estimator for the covariance matrix that is equivariant under orthogonal transformations. In particular, these transformations include rotations of the variable axes. Equivariant estimation of the covariance matrix under the orthogonal group has been studied extensively (e.g., Dey and Srinivasan (1985); Ledoit and Wolf (2015); Takemura (1984)) since the pioneering work of Stein (1975, 1986). In this study, we follow our previous work Banerjee et al. (2016), where we describe estimators that are valid when $n > p$, and extend it to the case when $p > n$. Because of the property of equivariance, the eigenvectors of the covariance matrix are estimated by those of $S = \sum_{i=1}^n X_i X_i^T$, to which we refer as the sample covariance matrix in this paper. Thus, the estimation problem is completed by providing estimates of the eigenvalues. These estimates are obtained from an adjusted profile likelihood function of the population eigenvalues, which is derived by approximating the integral of the density function of S over its eigenvectors (corresponding to the non-zero eigenvalues). This approximation is however not the large- n (Laplace) approximation of such integral, which results in the modified profile likelihood of Barndorff-Nielsen (1983), but it is an approximation suggested in Hikami and Brézin (2006) useful for large p . Our estimates are a mixture $\hat{\lambda}^\kappa = (1 - \kappa)\hat{\lambda}^0 + \kappa\hat{\lambda}^1$ of an exact critical point $\hat{\lambda}^0$ of such a likelihood function, which is in fact a maximum when some conditions are satisfied, and an approximate critical point $\hat{\lambda}^1$ whose components are a modification of the non-zero sample eigenvalues by terms of order $1/p$. The tuning parameter κ is determined from the data and controls the shrinkage of the eigenvalues. We will describe two algorithms to determine κ : one based on bootstrap and one based on cross-validation. High-dimensional estimators are generally derived under an asymptotic regime in which both n and p increase in such a way that their ratio tends to a constant. In our case, n is kept fixed, and the high-

dimensionality of the estimator comes into play because we consider a large- p approximation of the marginal density of the eigenvalues of S .

In a variety of finite-sample simulation scenarios, we compare our estimator to two Ledoit-Wolf estimators, which are also orthogonally equivariant and have previously been shown to better many other estimators under some loss functions. Figures 3 and 4 summarize the results of our comparison, in terms of risk evaluated with respect to nine loss functions. To assess how our covariance matrix estimator estimates the eigenvalues of the population covariance matrix, we also compare the eigenvalues of our estimator with Ledoit-Wolf consistent estimator of the population eigenvalues themselves, under loss functions that depend only on the eigenvalues. Figure 5 displays such a comparison. Finally, comparison of covariance matrix estimators is also carried out using two real data applications of breast cancer data and leukemia data: in a linear discriminant analysis of these data, we use plugged-in estimates of the covariance matrix in the classifier and demonstrate that our estimator leads to lower misclassification errors in the breast cancer data and similar misclassification errors in the leukemia data. The two Ledoit-Wolf covariance matrix estimators are optimal asymptotically under two loss functions, but we show that finite sample improvements and improvements under other loss functions are indeed possible. Since the tuning parameter κ of our estimator can be chosen by minimizing risk estimates with respect to any loss function, our estimator can be used with any loss function appropriate to a statistical application.

This paper is organized as follows. In Section 2, we introduce the adjusted profile likelihood that is used to obtain our estimator, which is introduced and discussed in Section 3. Section 4 and Section 5 present some numeric assessment of the performance of our estimator in simulated and real data respectively.

2. Marginal Density and Likelihood Function

In this section, we introduce some notations, review the singular Wishart distribution, derive an approximation to the marginal density of the sample eigenvalues and then obtain an adjusted profile likelihood for the eigenvalues of the population covariance matrix. Consider the case of mutually independent draws X_1, \dots, X_n from a multivariate p -dimensional normal distribution $\mathcal{N}(0, \Sigma)$, with Σ a $p \times p$ maximum-rank positive definite matrix. Assume $p > n$, and let S be the $p \times p$ sample covariance matrix $X^\top X$, with X the matrix whose rows are the vectors X_i^\top . S is positive semi-definite and of maximum rank, with distinct positive eigenvalues: $\ell_1 > \ell_2 > \dots > \ell_n > 0$. Geometrically, S is an interior point of an $(n+1)n/2$ -dimensional face of the closed convex cone of semi-positive definite $p \times p$ symmetric matrices (Barvinok, 2002). Uhlig (1994) showed that S has a distribution specified by the density

$$p(S)(dS) = K(\det \Sigma)^{-n/2} \text{etr}(-\Sigma^{-1}S/2)(\det L)^{(n-p-1)/2}(dS), \quad (1)$$

where $K = \pi^{(-pn+n^2)/2} 2^{-pn/2} / \Gamma_n(n/2)$, $L = \text{diag}(\ell_1, \dots, \ell_n)$ is the diagonal matrix of the non-zero eigenvalues of S , $\text{etr}(\cdot) = \exp(\text{tr}(\cdot))$, and (dS) is the volume element on the space of

positive semi-definite $p \times p$ symmetric matrices of rank n , with n distinct positive eigenvalues. This distribution, which extends the usual ($n > p$) Wishart distribution, is often called (non-central) singular Wishart distribution, but some authors (Srivastava and Khatri, 1979) prefer the name non-singular pseudo-Wishart distribution. It corresponds to the case 7 of the classification scheme of Díaz-García et al. (1997) (described in Table 1 therein), where generalizations are considered that include the cases when Σ and S have non-maximum rank and when the samples are not independent. The method we will present can be also extended to the case when S does not have maximum rank. For example, in some applications, one may wish to center the observations to have mean zero. The resulting matrix S constructed from the centered data will have rank less than n .

Consider now the singular value decomposition of $X = UL^{1/2}H_1^\top$ with $U \in O(n)$ an orthogonal $n \times n$ matrix, $L = \text{diag}(\ell_1, \dots, \ell_n)$ as defined above and H_1 the matrix whose n columns are the corresponding n eigenvectors of S . These n eigenvectors are uniquely determined up to column multiplication by ± 1 . The formulae below assume that one of these 2^n choices has been made. H_1 is a point in the Stiefel manifold, $V_n(\mathbb{R}^p)$, of all orthonormal n -frames in \mathbb{R}^p . The joint density of H_1 and L is

$$p(H_1, L) \left(H_1^\top dH_1 \right) \wedge (dL)$$

where

$$p(H_1, L) = 2^{-n} K(\det \Sigma)^{-n/2} \text{etr} \left(-\Sigma^{-1} H_1 L H_1^\top / 2 \right) (\det L)^{(p-n-1)/2} \prod_{i < j}^n (\ell_i - \ell_j),$$

\wedge is the exterior product, $(dL) = d\ell_1 \wedge \dots \wedge d\ell_p$, and $(H_1^\top dH_1)$ is the Haar measure of $V_n(\mathbb{R}^p)$ normalized as follows

$$\int_{V_n(\mathbb{R}^p)} (H_1^\top dH_1) = \text{Vol}(V_n(\mathbb{R}^p)) = \frac{2^n \pi^{pn/2}}{\Gamma_n(1/2^p)}.$$

Even if we have written the densities of S and of H_1, L as differential forms, we need not keep track of the sign of the form, if we define the integrals to be positive. In the following, when the measure is clear, we will revert to using the term density for the scalar part.

We are interested in obtaining an estimator for Σ that belongs to the class of orthogonally equivariant estimators. This equivariance is intended in the usual meaning. Namely, consider the action of the orthogonal group $O(p)$ on the sample space that is defined by $X \mapsto XG^\top$, or equivalently, by $S \mapsto GSG^\top$. We require $\hat{\Sigma}(S) \mapsto G\hat{\Sigma}(S)G^\top$, which is the same way as Σ transforms $\Sigma \mapsto G\Sigma G^\top$. Under such action, $p(S)$ is invariant and so is the measure (dS) . This equivariance implies that the eigenvectors of $\hat{\Sigma}$ are the same as those of S . This is to say that the estimators are of the form $\hat{\Sigma} = H\hat{\Lambda}H^\top$, where the elements of the diagonal matrix $\hat{\Lambda}$ are

functions of the non-zero eigenvalues of S and the orthogonal matrix $H = [H_1 : H_2]$ is that of the eigenvectors of S (Stein, 1986). In the case in which $p > n$, the zero eigenvalue has multiplicity $(p - n)$, so that the corresponding eigenvectors of S given by the $(p - n)$ columns $H_2 = (h_{n+1}, \dots, h_p)$, are unique only up to an orthogonal transformation of the last $(p - n)$ coordinate axes. Namely, we can consider H_2 or $H_2 \cdot P$ with any orthogonal matrix $P \in O(p - n)$. In general each different choice (of P) will lead to a different estimator of Σ . However, if the estimates of the smallest $(p - n)$ eigenvalues of Σ are identical, all such choices will be immaterial, in the sense that they will lead to the same $\hat{\Sigma}$. The estimates of the last $(p - n)$ eigenvalues that we propose are indeed identical, so that one can use as estimates of all the eigenvectors of Σ whatever representative of the class of the eigenvectors of S a numerical routine outputs.

To find the eigenvalue estimates, we follow our previous paper Banerjee et al. (2016), and consider the marginal density of the sample roots of $S = X^T X = H_1 L H_1^T$:

$$p(\ell_1, \dots, \ell_n) = \frac{2^{-np/2} \pi^{(-pn + n^2)/2}}{2^n \Gamma_n(n/2)} (\det \Sigma)^{-n/2} \prod_{i=1}^n \ell_i^{(p-n-1)/2} \prod_{i < j}^n (\ell_i - \ell_j) \cdot J_n$$

where

$$J_n = \int_{V_n(\mathbb{R}^p)} \text{etr} \left(-\frac{1}{2} \Sigma^{-1} H_1 L H_1^T \right) (H_1^T dH_1).$$

The integral J_n cannot be computed in closed form. However, in Appendix A, we derive a useful approximation. Namely,

Proposition 1.

For large p , the integral J_n is approximated by the following expression

$$J_n \approx \exp \left(-\frac{1}{2} \sum_{i=1}^n \frac{\hat{\ell}_i}{\lambda_i} \right) \cdot \prod_{i < j}^p \left(1 + \frac{1}{p} \left(\frac{\lambda_i - \lambda_j}{\lambda_i \lambda_j} \right) (\hat{\ell}_i - \hat{\ell}_j) \right)^{-\frac{1}{2}}$$

with $\hat{\ell}_i = \ell_i$ for $1 \leq i \leq n$ and $\hat{\ell}_i = 0$ for $i > n$.

Remark.

The following condition $0 < (I_j - I_i)(\lambda_i - \lambda_j)/\lambda_i \lambda_j < p$ is employed in the proof of the proposition.

Remark.

This approximate formula is valid even when some of the population eigenvalues are equal. It is also interesting to notice that when $\Sigma = \mu I_p$ is proportional to the identity matrix, the value of J_n obtained using such formula is equal to the exact value of the integral, which can be computed to be $\exp(-\sum_{i=1}^n \ell_i/2\mu)$ times the volume of the Stiefel manifold, which is

omitted in the formula above (see Appendix A). In the following, however, we assume the population eigenvalues to be distinct except when stated otherwise.

Employing such an approximation, we then obtain an approximate log-likelihood function for the true eigenvalues λ

$$\begin{aligned} \mathcal{L}(\lambda) = & -\frac{n}{2} \sum_{i=1}^p \ln \lambda_i - \frac{1}{2} \sum_{a=1}^n \frac{\ell_a}{\lambda_a} \\ & - \frac{1}{2} \sum_{1 \leq a < b \leq n} \ln \left(1 + \frac{1}{p} \left(\frac{\lambda_a - \lambda_b}{\lambda_a \lambda_b} \right) (\ell_a - \ell_b) \right) \\ & - \frac{1}{2} \sum_{1 \leq a \leq n < r \leq p} \ln \left(1 + \frac{1}{p} \left(\frac{\lambda_a - \lambda_r}{\lambda_a \lambda_r} \right) \ell_a \right). \end{aligned} \tag{2}$$

The first two terms of this function are the profile log-likelihood function for the parameters λ , which is the partially maximized log-likelihood function of (λ, V) , where V is replaced by the maximum likelihood estimator \hat{V}_λ for fixed λ . We show in Appendix C that \hat{V}_λ is a solution to the equation $V^T H_1 = M$, where M is a $p \times n$ matrix such that $M_{ij} = \pm \delta_{ij}$, with δ_{ij} the Kronecker delta. The other terms in \mathcal{L} can thus be interpreted as an adjustment to the profile log-likelihood.

3. The Proposed Estimator

In this section, we derive an estimator for the eigenvalues of the population covariance matrix and discuss its properties.

Our starting point is $\mathcal{L}(\lambda)$, the function given in (2), which can be considered as a pseudo-(log)likelihood, of which our goal is to find the maximum points. We note that this function of $\lambda \in \mathbb{R}_+^p$ is not concave on the whole domain \mathbb{R}_+^p , for all given values of (ℓ_1, \dots, ℓ_n) .

The critical points are the solutions to the following equations

$$n\lambda_i = \ell_i - \frac{1}{p} \sum_{b=1}^n \frac{\ell_i - \ell_b}{1 + \frac{1}{p} \left(\frac{\lambda_i - \lambda_b}{\lambda_i \lambda_b} \right) (\ell_i - \ell_b)} - \frac{\ell_i}{p} \sum_{r=n+1}^p \frac{1}{1 + \frac{1}{p} \left(\frac{\lambda_i - \lambda_r}{\lambda_i \lambda_r} \right) \ell_i} \tag{3}$$

where $\ell_i = 0$ when $i = n + 1, \dots, p$. Exact solutions to (3) satisfy what we can call a trace condition: $n \sum_{i=1}^p \lambda_i = \sum_{a=1}^n \ell_a$, which is desirable since $E(S) = n\Sigma$.

A solution to (3) is seen to be given by $\hat{\lambda}_i^0 = \sum_{a=1}^n \ell_a / np$, for $i = 1, \dots, p$. This results in a diagonal estimator for the covariance matrix $\hat{\Sigma}^0 = \text{tr}(S) \cdot I_p / np$. However there is no guarantee that such a solution may be a maximum of \mathcal{L} for all given values of (ℓ_1, \dots, ℓ_n) . Indeed, for some data (sample eigenvalues) the Hessian matrix evaluated at this solution can have positive eigenvalues. However, one can notice that $\hat{\lambda}^0$ is a genuine maximum of the likelihood function at order 1 in $1/p$. Namely, it is a maximum of

$$\mathcal{L}' = -n \sum_{i=1}^p \ln \lambda_i - \sum_{a=1}^n \frac{\ell_a}{\lambda_a} - \frac{1}{p} \sum_{1 \leq a < b \leq n} \frac{\lambda_a - \lambda_b}{\lambda_a \lambda_b} (\ell_a - \ell_b) - \frac{1}{p} \sum_{1 \leq a \leq n < r \leq p} \frac{\lambda_a - \lambda_r}{\lambda_a \lambda_r} \ell_a,$$

which is obtained by expanding to first order the logarithm in (2). Furthermore, $\hat{\lambda}^0$ is also the maximum of the exact likelihood function when the true covariance matrix is proportional to the identity matrix, which can be computed exactly, without the need of any approximation, as remarked earlier. The advantage of the solution $\hat{\lambda}^0$ is that it shrinks the highest and pushes up the lowest eigenvalue. Indeed, the general theorem of van der Vaart (1961) tells us that the highest eigenvalue $\hat{\ell}_n$ is upward biased and (obviously in this $p > n$ case) the lowest eigenvalue downward biased. The disadvantage is that the shrinkage may be too extreme. It is perhaps not surprising that the eigenvalue estimates are degenerate. In our derivation, the sample size n is fixed and p is large, thus there may not be enough information in the sample to obtain a different estimate for each eigenvalue. To deal with the degeneracy of the estimates, we construct approximate solutions to the equations (3). We look for $\hat{\lambda}_i$ with an expansion of the form $a_i + f_i/p + O(p^{-2})$, for $i = 1, \dots, n$. There is no reason a priori why solutions should take this form. However, the solution $\hat{\lambda}_i^0$ is of this form, with $a_i = 0$ and f_i the same for all $i = 1, \dots, n$. Our goal is to perturb the exact solution $\hat{\lambda}^0$ away from having all components equal, keeping the resulting estimates ordered and satisfying the trace condition. We find

$$\begin{aligned} \hat{\lambda}_a^1 &= \frac{\ell_a}{n} - \frac{1}{n} \sum_{b=1}^n \frac{\ell_a - \ell_b}{p + n \left(\frac{1}{\ell_b} - \frac{1}{\ell_a} \right) (\ell_a - \ell_b)}, \quad a = 1, \dots, n \\ \hat{\lambda}_r^1 &= 0, \quad r = n + 1, \dots, p. \end{aligned}$$

which are a modification of the eigenvalues $\hat{\ell}_n$ of the (usual) sample covariance matrix S/n (in our conventions $\hat{\ell}$ are the eigenvalues of $S = X^T X$) with a correction term of order $1/p$. We do not use such an approximate solution as the estimate of our eigenvalues (it would lead to a non-invertible estimator of the covariance matrix, for one thing). We employ $\hat{\lambda}^1$ as a perturbation of the true solution $\hat{\lambda}^0$. In fact, what we propose as an estimator is a linear combination of $\hat{\lambda}^0$ and $\hat{\lambda}^1$, controlled by a tuning parameter κ . Namely,

$$\begin{aligned} \hat{\lambda}_a^\kappa &= \frac{\kappa}{n} \left(\ell_a - \sum_{b=1}^n \frac{\ell_a - \ell_b}{p + n \left(\frac{1}{\ell_b} - \frac{1}{\ell_a} \right) (\ell_a - \ell_b)} \right) + (1 - \kappa) \frac{\sum_{a=1}^n \ell_a}{np}, \quad a = 1, \dots, n \\ \hat{\lambda}_r^\kappa &= (1 - \kappa) \frac{\sum_{a=1}^n \ell_a}{np}, \quad r = n + 1, \dots, p \end{aligned} \tag{4}$$

where $0 < \kappa < 1$. The parameter κ controls the shrinkage of the eigenvalue estimate and is to be determined from the data. When κ is zero, the shrinkage is highest, and we recover the solution to the ML equations, with all the eigenvalues being equal and, accordingly, $\hat{\Sigma} = \hat{\Sigma}^0$

is proportional to the identity I_p . When κ tends to one we get $\hat{\lambda}^1$ with distinct estimates of the first n eigenvalues. There is no guarantee that $\hat{\lambda}^\kappa$ is a maximizer for (2). For these reason, one could try to maximize (2) numerically. We did consider numerical solutions to (3) using Newton's method and a constraint of positivity on the solutions. The resulting roots were always found to be close to $\hat{\lambda}^1$ with the last $(p - n)$ values negligible. Furthermore, when used in place of $\hat{\lambda}^1$, in the estimator (4), these numerical components had similar estimates of risk compared to our estimator (4) in the simulation study conducted in Section 4.2 (results not shown), but added an un-necessary computational step.

The last $(p - n)$ estimates of the eigenspectrum (4) are all equal. As observed in Section 2, this property guarantees that any chosen basis for the eigenspace corresponding to the zero eigenvalues of S will give rise to the same estimator $\hat{\Sigma}^\kappa$. In fact, this property should be required of any orthogonally equivariant estimators of Σ in the $p > n$ setting, although it has not been explicitly mentioned before, and it also holds for the non-linear estimators of Ledoit and Wolf (2015). Our proposed estimators for the true eigenvalues have the following additional properties proven in Appendix B.

Proposition 2 (The properties of the eigenvalue estimates).

For κ , such that $0 < \kappa < 1$, the estimates $\hat{\lambda}^\kappa$ given in (4) have the following properties

- a. $\hat{\lambda}_1^\kappa > \hat{\lambda}_2^\kappa > \dots > \hat{\lambda}_n^\kappa > \hat{\lambda}_{n+1}^\kappa = \dots = \hat{\lambda}_p^\kappa > 0$
- b. $n \sum_{i=1}^p \hat{\lambda}_i^\kappa = \sum_{a=1}^n \ell_a$

Thus the corresponding estimator $\hat{\Sigma}_\kappa$ of Σ will be positive definite. In addition, because of the ordering of the estimates, there is no additional step, such as re-ordering or isotonization, that often is necessary. The computational burden of obtaining the proposed estimates only stems from finding the singular value decomposition of the data matrix X or the eigenspectrum of S , and by the evaluation of the parameter κ , which we discuss in Section 3.1.

The formulae presented so far have been obtained under the assumption that the data matrix X or $S = X^T X$ were of maximum rank n . In some applications one may wish to center the data, $Y_i = X_i - \bar{X}$, and consider the matrix $\sum_i Y_i Y_i^T$ in place of S . All formulae can be applied to these situations, if we replace n with the rank q of the rescaled matrix in the corresponding maximum rank equations. A sketch of their derivation is given in Appendix D.

3.1. Selecting the Tuning Parameter κ

The tuning parameter κ , $0 < \kappa < 1$, of $\hat{\lambda}^\kappa$ needs to be determined from data. Selection of tuning parameters in an unsupervised setting is a difficult problem, and there is no method which is always satisfactory. In the context of covariance estimation, tuning parameters are often determined by minimizing estimates of risks (see for example Bickel and Levina (2008) for a cross-validation approach and Yi and Zou (2013) for an approach using Stein's

unbiased estimate of risk). This is also the approach we follow to choose κ , although our estimates of risk differ. Namely, we consider some loss function $L(\hat{\Sigma}, \Sigma)$ and compute the corresponding risk as follows:

$$R(\hat{\Sigma}, \Sigma) = E(L(\hat{\Sigma}, \Sigma)), \quad (5)$$

where the expectation is over the data distribution. When $\hat{\Sigma} = \hat{\Sigma}^\kappa$, the risk can be seen as a function of κ . The ‘‘oracle’’ κ' is then $\kappa' = \underset{\kappa}{\operatorname{argmin}} R(\kappa)$.

It is noteworthy that estimating the risk and estimating the tuning parameter that minimizes that risk, are not necessarily the same problem. The estimation of the risk is complicated by the fact that the true population matrix Σ is unknown in practice. We propose two methods to estimate the risk of the estimator $\hat{\Sigma}^\kappa$ under a loss L : one method relies on a bootstrap re-sampling scheme and the other on cross-validation.

κ -selection via bootstrap—To estimate the risk via bootstrap, we randomly choose n rows with replacement from our data matrix X . Let X_b be such a sample and $S_b = X_b^\top X_b$ be the corresponding sample covariance matrix. We then compute the reduced rank estimator, $\hat{\Sigma}_b^\kappa$, as described in Appendix D and evaluate the loss $L(\hat{\Sigma}_b^\kappa, \bar{\Sigma})$ with respect to a reference estimator $\bar{\Sigma}$ for a grid of values of $\kappa \in [0, 1)$. This procedure is iterated B times. The risk estimate is taken to be

$$\hat{R}(\kappa) = \frac{1}{B} \sum_{b=1}^B L(\hat{\Sigma}_b^\kappa, \bar{\Sigma})$$

and the optimal κ determined as

$$\hat{\kappa} = \underset{\kappa \in [0, 1)}{\operatorname{argmin}} \hat{R}(\kappa).$$

The choice of the reference estimator $\bar{\Sigma}$ in our proposed κ -selection procedure requires discussion. We have considered using S and $\hat{\Sigma}^1$ (matrix estimator corresponding to $\hat{\lambda}^1$) as reference estimators. However, since these estimators are singular, they may not be used when computing loss functions that require their inversions (such as Stein’s loss function or the quadratic loss function, see Section 4). In these cases, we have flipped the role of the reference estimator and the estimator at hand when computing the loss functions. As an alternative approach, we have used as reference estimator a non-singular extension of $\hat{\Sigma}^1$, which we call $\hat{\Sigma}_{NS}^1$, where the zero eigenvalues are replaced with the smallest non-zero eigenvalue estimate. Our simulation studies (not shown here) indicate that the second strategy of using a non-singular estimator performs better than the first approach in selecting κ . We should emphasize that the choice of κ closest to the κ' depends largely on the reference estimator and a better reference estimator will lead to a much improved estimator.

κ -selection via cross-validation—A second method estimates the risk using cross-validation (CV). Different losses lead to different estimates, and we have implemented this method for Stein's, quadratic and Frobenius loss functions (see Section 4.1 for definition of these loss functions). First, consider the Frobenius loss, or actually its square, for reasons that we will be readily apparent:

$$frob(\hat{\Sigma}^{\kappa}, \Sigma)^2 = \text{tr}\left((\hat{\Sigma}^{\kappa} - \Sigma)(\hat{\Sigma}^{\kappa} - \Sigma)^{\top}\right) = \text{tr}(\hat{\Sigma}^{\kappa} \hat{\Sigma}^{\kappa}) - 2\text{tr}(\Sigma \hat{\Sigma}^{\kappa}) + \text{tr}(\Sigma^2).$$

We can now ignore terms that do not depend on κ since we wish to minimize with respect to κ . Observing further that

$$\text{tr}(\Sigma \hat{\Sigma}^{\kappa}) = E^*(\mathbf{x}_*^{\top} \hat{\Sigma}^{\kappa} \mathbf{x}_*)$$

where the expectation E^* is taken with respect to the distribution of \mathbf{x}_* , an independent sample from $\mathcal{N}(0, \Sigma)$, we obtain the following leave-one-out cross-validation estimate of the risk

$$f\hat{r}_{ob\kappa}^2 = \text{tr}(\hat{\Sigma}^{\kappa} \hat{\Sigma}^{\kappa}) - \frac{2}{n} \sum_{i=1}^n X_i^{\top}(\hat{\Sigma}_{\setminus i}^{\kappa})X_i$$

where $\hat{\Sigma}_{\setminus i}^{\kappa}$ is the estimator obtained removing the i -th row X_i^{\top} from X . If n is not very small, one can consider K -fold CV instead of leave-one-out CV, to ease the computational burden. When the loss function is quadratic or Stein's, we reverse the role of Σ and $\hat{\Sigma}$ and consider $L(\Sigma, \hat{\Sigma}^{\kappa})$ rather than $L(\hat{\Sigma}^{\kappa}, \Sigma)$. We can express the traces involving Σ and the inverse of $\hat{\Sigma}^{\kappa}$ as expectation with respect to samples of $\mathcal{N}(0, \Sigma)$, and we are able to obtain the estimates of risks, or, more accurately, of functions that have the same minimum as the risks, since terms that do not depend on κ can be ignored. More precisely, the quadratic loss with the reversed role of Σ and $\hat{\Sigma}$ is

$$\begin{aligned} q(\Sigma, \hat{\Sigma}) &= \text{tr}\left((\Sigma \hat{\Sigma}^{-1} - I)^2\right) = p - 2\text{tr}(\Sigma \hat{\Sigma}^{-1}) + \text{tr}(\Sigma \hat{\Sigma}^{-1} \Sigma \hat{\Sigma}^{-1}) \\ &= p - 2E^*(\mathbf{x}_*^{\top} \hat{\Sigma}^{-1} \mathbf{x}_*) + \frac{1}{2}E^*(\mathbf{x}_*^{\top} \hat{\Sigma}^{-1} \mathbf{x}_* \cdot \mathbf{x}_*^{\top} \hat{\Sigma}^{-1} \mathbf{x}_*) \\ &\quad - \frac{1}{2}\left(E^*(\mathbf{x}_*^{\top} \hat{\Sigma}^{-1} \mathbf{x}_*)\right)^2. \end{aligned}$$

Ignoring terms that do not depend on κ , the leave-one-out cross-validation estimate of the risk under the quadratic loss is

$$\hat{q}_{\kappa} = -\frac{2}{n} \sum_{i=1}^n X_i^{\top}(\hat{\Sigma}_{\setminus i}^{\kappa})^{-1} X_i + \frac{1}{2n} \sum_{i=1}^n \left(X_i^{\top}(\hat{\Sigma}_{\setminus i}^{\kappa})^{-1} X_i\right)^2 - \frac{1}{2}\left(\frac{1}{n} \sum_{i=1}^n X_i^{\top}(\hat{\Sigma}_{\setminus i}^{\kappa})^{-1} X_i\right)^2.$$

The Stein's loss $st(\Sigma, \hat{\Sigma})$, which is now twice the KL-divergence of normal densities with covariance matrix Σ and $\hat{\Sigma}$, is equal, up to terms that do not depend on κ , to

$$\mathcal{L}_\kappa = \frac{1}{2} E_* \left(\mathbf{x}_*^\top \hat{\Sigma}^{-1} \mathbf{x}_* \right) + \frac{1}{2} \ln \det \hat{\Sigma}^\kappa,$$

from which one obtains the leave-one-out cross-validation risk estimate

$$\hat{st}_\kappa = \frac{1}{2n} \sum_{i=1}^n X_i^\top (\hat{\Sigma}_i^\kappa)^{-1} X_i + \frac{1}{2} \ln \det \hat{\Sigma}^\kappa.$$

Numerical comparison of κ -selecting methods—We conducted a simulation study (see Section 4.2 for details) to evaluate these two strategies and compare the corresponding values of κ with the “oracle” κ' . We considered nine different loss functions (see Section 4.1), eight different covariance structures and $p = 50, 100, 500, 1000$ with $\gamma = \frac{p}{n} = 1.25, 2, 5$.

The results of this study are shown in Figures 1 and 2 where each panel corresponds to a loss function, the horizontal axis of each panel represents various combinations of the true covariance structure Σ , γ and p , and the vertical axis of each panel represents the chosen κ for the two proposed methods (CV in dashed line and bootstrap in dotted line) and the “oracle” κ' (in solid line) which is determined using the true matrix Σ . Figure 1 has true covariance structures Σ_1 to Σ_4 and Figure 2 has true covariance structures Σ_5 to Σ_8 . Note that the leave-one-out CV is only available for three out of the nine loss functions (i.e., Frobenius, Stein's and quadratic). The leave-one-out CV estimates the oracle κ' almost perfectly for the Frobenius loss and also works well for Stein's and quadratic loss functions. The bootstrap method on the other hand works reasonably well in choosing κ for most loss functions except the ones that depend on the smallest eigenvalues entirely. We have observed in our simulations that when the bootstrap estimate $\hat{\kappa}$ is quite different from the oracle κ' , the risk curves are quite flat, i.e., $\hat{R}(\hat{\kappa})$ and $\hat{R}(\kappa')$ are quite similar and there is little risk improvement with different choices of κ .

As noted earlier, the choice of κ depends largely on the reference estimator. When analyzing real data, we recommend the analyst to choose a loss function that is appropriate for the applied problem. If any of Stein's, quadratic or Frobenius (i.e., *st*, *q* or *frob*) loss functions is an appropriate choice, we recommend using the CV-method to select κ as it does not depend on a reference estimator. For any other loss function, we would recommend using other estimators as reference (e.g., the estimators of Ledoit and Wolf) in addition to $\hat{\Sigma}_{NS}^1$.

4. Numerical Risk Comparisons with Other Estimators

In this section we perform Monte Carlo simulations to evaluate our proposed estimator with respect to various loss functions. Specifically, we compare it with the two non-linear shrinkage estimators of Ledoit and Wolf, one which is asymptotically optimal under Frobenius loss (LW1) and one which is asymptotically optimal under Stein's loss (LW2). Ledoit and Wolf's nonlinear shrinkage estimators are of the form $H\hat{D}H^\top$, and thus

orthogonally equivariant, with $\widehat{D} = \text{diag}(\widehat{\varphi}^*(\ell_1), \dots, \widehat{\varphi}^*(\ell_p))$. The function $\widehat{\varphi}^*$ is the nonlinear function responsible for shrinking the sample eigenvalues. Its form depends on which loss function is asymptotically minimized (Ledoit and Wolf, 2018). We refer the reader to Section 3.2 of Ledoit and Wolf (2015) for the specific form of $\widehat{\varphi}^*$ when the loss is Frobenius and to Sections 5 and 6 of Ledoit and Wolf (2018) for its explicit form when the loss is Stein's. In addition, we compare our eigenvalue estimates with those of the Ledoit-Wolf consistent estimator of the population eigenvalues (Ledoit and Wolf, 2015), which we call LW3. The comparison is carried out under the loss functions that only depend on eigenvalues: loss functions 3, 4, 7, 8, and 9 in the Section 4.1 below.

4.1. Loss Functions

The comparison of the estimators is carried out using the Monte Carlo estimates of risk as defined in eq. (5) with respect to the nine loss functions defined below. We choose a variety of loss functions, some of which depend on the complete covariance matrix and its estimate, and others depend only on the eigenvalues and their estimates. To an applied statistician, the type of loss function is determined by the context of the application which could involve the estimation of eigenvalues, an important characteristic of the covariance matrix. Thus, the reader will get a sense of overall performance of an estimator under various loss functions. We also refer the reader to our previous paper Banerjee et al. (2016) for a detailed description of most of these loss functions.

1. Stein's (entropy) loss $st(\widehat{\Sigma}, \Sigma) = \text{tr}(\widehat{\Sigma}\Sigma^{-1} - I) - \ln \det(\widehat{\Sigma}\Sigma^{-1})$;
2. the quadratic loss $q(\widehat{\Sigma}, \Sigma) = \text{tr}(\widehat{\Sigma}\Sigma^{-1} - I)^2$;
3. L_1 eigenvalue loss $evl1(\widehat{\Sigma}, \Sigma) = \sum_{i=1}^p |\widehat{\lambda}_i - \lambda_i|/p$;
4. L_2 eigenvalue loss $evl2(\widehat{\Sigma}, \Sigma) = \sum_{i=1}^p (\widehat{\lambda}_i - \lambda_i)^2/p$;
5. Frobenius loss $frob(\widehat{\Sigma}, \Sigma) = \|\widehat{\Sigma} - \Sigma\|_F$, with $\|A\|_F^2 = \text{tr}(AA^T)$;
6. Matrix L_1 -norm, the max of the L_1 norm of the columns of $|\widehat{\Sigma} - \Sigma|$ or $\|\widehat{\Sigma} - \Sigma\|_{1,1}$,
 $Onenorm(\widehat{\Sigma}, \Sigma) = \max_{1 \leq j \leq p} \sum_{i=1}^p |\widehat{\sigma}_{ij} - \sigma_{ij}|$;
7. L_1 loss on the largest eigenvalue $TopEV(\widehat{\Sigma}, \Sigma) = |\widehat{\lambda}_1 - \lambda_1|$;
8. L_1 loss on the smallest eigenvalue $LastEV(\widehat{\Sigma}, \Sigma) = |\widehat{\lambda}_p - \lambda_p|$;
9. L_1 loss on the smallest quartile of the eigenvalues
 $EVS(\widehat{\Sigma}, \Sigma) = \sum_{i=\lceil 3p/4 \rceil}^p |\widehat{\lambda}_i - \lambda_i|$.

We notice that, for our estimator $\widehat{\Sigma}^{\kappa}$, the oracle κ' is chosen as described in Section 3.1.

4.2. Simulation Study

We construct eight covariance structures to represent typical applications. The matrix Σ_1 has widely spaced eigenvalues, Σ_2 has one large eigenvalue and mimics a typical principal

components analysis covariance structure, Σ_3 is a time series example, Σ_4 is a spiked covariance structure, Σ_5 is the identity matrix, Σ_6 has eigenvalues drawn from a U-shaped beta distribution, Σ_7 has eigenvalues drawn from a linearly decreasing beta density and Σ_8 has eigenvalues of 1, 3 and 10 distributed with a frequency of 20%, 40% and 40% respectively. Namely,

1. $\Sigma_1 = \text{diag}(p^2, \dots, 2^2, 1^2)$;
2. $\Sigma_2 = \text{diag}(\lambda_{(1)}^*, \lambda_{(2)}, \dots, \lambda_{(p)})$, where $\lambda_i \sim U(1, p/2)$, $\lambda_{(j)}$ are the ordered λ_i 's, $\lambda_{(1)}^* = \lambda_{(1)}^2$, with $U(a, b)$ being the uniform distribution over the $[a, b]$ interval;
3. $\Sigma_3 = AR(1)$, the first-order autoregressive covariance matrix, where $\sigma_{ij} = 4 \times 0.7^{|i-j|}$ for $i = j$ and $\sigma_{ij} = 4^2$ for $i = 1, \dots, p$;
4. $\Sigma_4 = \text{diag}(2p, p, 1, \dots, 1)$;
5. $\Sigma_5 = I_p$ where I_p is the p -dimensional identity matrix;
6. $\Sigma_6 = \text{diag}(\lambda_i)$ where $\lambda_i = 1 + 9F_{(0.5, 0.5)}^{-1}(\frac{i}{p} - \frac{1}{2p})$ $i = 1, \dots, p$ and $F(\alpha, \beta)$ is the cumulative distribution function (c.d.f) of a beta distribution with parameters (α, β) ; the choice of $\alpha = 0.5, \beta = 0.5$ draws eigenvalues from a U-shaped density with more mass on large (around 10) and small (around 1) eigenvalues as the support of the beta density has been shifted to $[1, 10]$;
7. Σ_7 is similar to Σ_6 with $(\alpha = 1, \beta = 2)$ for the shape parameters of the beta c.d.f reflecting a linearly decreasing triangle with the highest density at 1 and lowest density at 10;
8. Σ_8 has 20% of its eigenvalues equal to 1, 40% equal to 3 and the remaining 40% equal to 10.

For each covariance structure Σ_j ($j = 1, 2, 3, 4, 5, 6, 7, 8$) four values of p are considered $p = 50, 100, 500, 1000$. For all eight covariance structures and their various dimensions, three values of n are chosen corresponding to $p/n = \gamma = 1.25, 2, 5$. We generate n vectors $X_i \sim \mathcal{N}(0, \Sigma_j)$, $i = 1, \dots, n$, for the first seven covariance structures. For the eighth case, we generate n vectors from a p -variate Student t distribution with four degrees of freedom so that Σ_8 is used to test robustness to deviations from normality. We evaluate all nine loss functions, denoted henceforth by $st, q, ev1, ev2, frob, Onenorm, topev, lastev$ and evs , on our estimator and compare the latter with LW1 (nonlinear shrinkage estimator that is optimal under Frobenius loss) and LW2 (nonlinear shrinkage estimator optimal under Stein's loss) (see Section 4.1). Additionally, we evaluate the five loss functions ($ev1, ev2, topev, lastev, evs$) that depend only on eigenvalues to compare our estimator with LW3 (the Ledoit-Wolf consistent estimator of population eigenvalues). Risk estimates are based on 1000 repetitions for each simulation scenario.

As a measure of comparison, we use the Proportion Reduction in Integrated Average Loss (PRIAL) of our estimator $\hat{\Sigma}$ over $\hat{\Sigma}^{LW_k}$, which is LW1 ($k = 1$) or LW2 ($k = 2$). Namely, for a loss function $L(\cdot)$, PRIAL is defined similarly to that in Lin and Perlman (1985):

$$\left(\sum_i L(\hat{\Sigma}_i^{LW_k}, \Sigma) - \sum_i L(\hat{\Sigma}_i, \Sigma) \right) / \sum_i L(\hat{\Sigma}_i^{LW_k}, \Sigma) \quad (6)$$

where the sum (over i) is over all datasets. For LW3, we replace $\hat{\Sigma}^{LW3}$ with $\hat{\lambda}^{LW3}$, $\hat{\Sigma}$ with $\hat{\lambda}$, and Σ with λ in (6) to compute the PRIAL. Figure 3, 4 and 5 are heatmaps of PRIAL with respect to LW1, LW2 and LW3 respectively. The rows correspond to various simulation scenarios (covariance structure, p and γ) just described and the columns correspond to various loss functions, described in Section 4.1. The red (blue) shades in the heatmap indicate positive (negative) PRIAL and the white represents zero PRIAL. The more intense the hue is, the larger the absolute value of the PRIAL. A positive PRIAL means our estimator compares favorably with LW1/LW2/LW3, negative PRIAL means the opposite and zero PRIAL means the estimators are comparable. Our estimator has positive PRIAL in 40.4%, 45.7% and 36.9% of all simulation scenarios when compared to LW1, LW2 and LW3 respectively.

It seems difficult to arrive at a general conclusion on which estimator (LW1, LW2, LW3 or ours) is preferable for which covariance structure or which loss function although they perform similarly for the majority of scenarios. Since LW1 is asymptotically optimal under the Frobenius loss function, one should expect LW1 to be better than our estimator with the Frobenius loss. Figure 3 confirms this, except in the case in which the true covariance matrix is the identity (Σ_5) and in the case of t -distributed data with covariance matrix Σ_8 when our estimator has an advantage. The exception with the identity matrix is also expected because our approximate solution with $\kappa = 0$ is a true maximum of the marginal likelihood function. For $p = 500$, LW1 outperforms our estimator when the true structure is spiked (Σ_4) under all loss functions and when the true structure represents a PCA-like situation (Σ_2) under *frob*, *Onenorm* and *topev* loss functions. It is also expected that LW2, which is optimal under Stein's loss function, would outperform our estimator using Stein's loss. However, we observe in Figure 4 that, under the Stein's loss function, LW2 is better than our estimator only for the spiked covariance structure (Σ_4). On the other hand, our estimator tends to outperform LW2 with Σ_8 (robust case) and Σ_5 (identity) under multiple loss functions. If we compare estimators with respect to the true covariance matrix over all loss functions, LW1 and LW2 are better than our estimator for the spiked covariance structure (Σ_4), while our estimator is better than LW1 and LW2 when data are generated from a non-normal distribution with fat-tails (Σ_8) and it seems preferable to LW1 and LW2, when the true covariance matrix is the identity (Σ_5). When we compare our eigenvalue estimators with LW3 (Figure 5), we see a similar pattern: the estimators perform comparably for the majority of scenarios, with LW3 outperforming our eigenvalue estimates for the spiked covariance structure (Σ_4) when $p = 500$ under all eigenvalue loss functions and for Σ_2 when $p = 500$ under the *evl2* and *topev*.

5. Linear Discriminant Analysis (LDA) on Breast Cancer and Leukemia Data

In this section we apply our estimator to two-class classification problems using LDA in breast cancer and leukemia data. Specifically, we plug-in our estimator (and also LW1 and LW2) for the common covariance matrix of both classes in the discriminant function of LDA. In the first application we consider a breast cancer dataset. Hess et al. (2006) proposed a 31-probeset multi-gene predictor of pathologic complete response (pCR) to chemotherapy in an experiment with 133 patients with stage I-III breast cancer. Following Hess et al. (2006), we split the samples into a training set of size 82 and a test set of size 51. We develop our classifier on a subset of the training data so that $n < p = 31$ by randomly selecting 20 patients and preserving the ratio of the two classes. We then evaluate and compare discrimination metrics, such as misclassification rate (MCR), sensitivity (Sens) and specificity (Spec), of the classifier that uses our estimator with those of the classifiers that use LW1 and LW2 as the plug-in estimators for the common-covariance matrix. The comparisons are presented in Table 1. To choose κ for our estimator $\hat{\Sigma}^\kappa$, we follow the procedures described in Section 3.1 with nine loss functions and $\hat{\Sigma}_{N,S}^1$ as the reference estimator for the bootstrap-based method and the three loss functions for the CV-based method. The combination of the κ -selection method and the corresponding loss function that was chosen is described in Table 1. In addition, we also use LW1 and LW2 as the reference estimators in determining κ for the bootstrap-based method. Our estimator has comparable but higher MCR and lower sensitivity and specificity compared to LW1 and LW2 when we use $\hat{\Sigma}_{N,S}^1$ as the reference estimator to choose κ . However, when we use LW1 as the reference estimator our estimator improves MCR by 4%, sensitivity by 7.7 % and specificity by 2.6% compared with LW1. Similarly, when using LW2 as the reference estimator in choosing κ , our estimator improves MCR by 15.7% and specificity by 21%. The estimator identified by the κ that minimizes the criterion \hat{q}_κ via cross-validation as described in Section 3.1 also performs well.

In the second application we analyze the leukemia data set of Golub et al. (1999), consisting of gene expression measurements on 72 leukemia patients, 47 ALL and 25 AML. We retain $p = 3571$ of the 7148 gene expression levels for the analysis. This smaller data set can be downloaded from T. Hastie's website https://web.stanford.edu/~hastie/CASI_files/DATA/leukemia.html. A training set and a validation set are then obtained by randomly selecting samples from the two classes but preserving the proportion of the classes as in the original data, with the training set comprising 31 ALL and 17 AML patients. We consider all LDA-classifiers as described in the breast cancer example, which employ our estimator, LW1, and LW2. All these estimators achieve perfect classification of the samples in the validation set.

6. Summary and Conclusion

Estimation of the covariance matrix is encountered in many statistical problems and has received much attention recently. When p is comparable to n or even greater, the sample covariance matrix is a poor and ill-conditioned estimator primarily due to an overspread

eigenspectrum. Several alternative estimators have been considered in the literature for such scenarios, some of which are asymptotically optimal with respect to certain loss functions and others are derived under strong structural assumptions on the covariance matrix (e.g., sparsity). Often, estimators are valid in the regime in which both n and p go to infinity in such a way that their ratio is finite.

In this paper, we consider the class of orthogonally equivariant estimators and propose an estimator that is valid when $p > n$. This work is an extension of our previous work on equivariant estimation when $p < n$. Equivariance under orthogonal transformations reduces the problem of estimating the covariance matrix to that of the estimation of its eigenvalues. To this end, we find a modification of the profile likelihood function of eigenvalues by integrating out the sample eigenvectors. The integration result is approximate and valid for large p . The critical point of this pseudo-likelihood function, a maximum under certain conditions, is in an estimator $\hat{\lambda}^0$ with all components equal, thereby resulting in extreme shrinkage. To get distinct eigenvalue estimates, we perturb $\hat{\lambda}^0$ by introducing an approximate solution $\hat{\lambda}^1$ to the likelihood equations along with a tuning parameter κ . The tuning parameter, $\kappa \in [0, 1)$, is selected by minimizing the risk, with respect to a loss function. We can find estimates of the risk using a bootstrap re-sampling scheme, which can be applied to any problem with any loss function. The κ selected with this method depends on the choice of a reference estimator, necessary to evaluate the loss function. Our estimator improves risk when a good estimator is employed as a reference estimator. Alternatively, a cross-validation estimate of the risk can be used, which was implemented for Frobenius, quadratic, and Stein's loss functions. We compare finite sample properties of our proposed estimator with two covariance matrix estimators of Ledoit and Wolf (Figure 3 and 4) using Monte Carlo estimates of risk with respect to nine different loss functions and eight different covariance structures. We also compare the estimates of the population eigenvalues obtained by our method with those of a consistent estimator for population eigenvalues proposed by Ledoit and Wolf (Figure 5). Furthermore, we demonstrate in a real breast cancer example that our estimator can substantially improve risk.

The encouraging finite sample properties of our estimator reported here suggest that our method of constructing an orthogonally equivariant estimator on the marginal distribution of the sample eigenvalues may provide improved estimation of the covariance matrix, which is needed in many statistical applications.

Acknowledgments

SB was supported by NIMH (National Institute of Mental Health) grant number P50 MH113838 for this work. Many thanks to Yiyuan Wu and Elizabeth Mauer for their assistance with the Figures.

Appendix A

In this appendix we prove Proposition 1. To do this, we first review a result that is fundamental in deriving the approximation. Consider the integral

$$I = \int_{O(p)} \text{etr}(H X H^{-1} Y) (dH)$$

where (dH) is the Haar measure of the group $O(p)$, X and Y are $p \times p$ symmetric matrices with eigenvalues $x = (x_1, \dots, x_p)$ and $y = (y_1, \dots, y_p)$ respectively. In Hikami and Brézin (2006), this integral and some of its generalizations (to the unitary and symplectic groups), known often as Harish-Chandra-Itzykson-Zuber integrals, are expressed in a form that involves the variables $\tau_{ij} = (x_i - x_j)(y_i - y_j)$, using the expansion about the saddle points of the integrand. Specifically, for the orthogonal group,

$$I = e^{\sum_{j=1}^p x_j y_j} \cdot f(\tau), \quad (7)$$

where the matrix τ is the symmetric matrix with entries τ_{ij} . The function f can be expanded as a power series of τ

$$f = 1 + f_1 + \dots + f_r + O(\tau^{r+1})$$

with the term f_s being of order s in τ . The solution (7) is obtained from a differential equation obeyed by the integral I . The power series in τ for the integral I is also obtained from the expression of I in terms of zonal polynomials. The term f_s of the expansion of f is a polynomial of degree s in the τ variables. Each monomial enjoys a graphical representation, as a graph that has p nodes and an edge that connects two nodes i, j if the variable τ_{ij} appears in the monomial. In particular, τ_{ij}^q is represented as an edge of multiplicity q , that is, as q lines between the nodes i, j . The total number of lines in such graph, and thus the sum of the multiplicities of all the edges, is s . The details can be found in Hikami and Brézin (2006), but a few example can make the discussion clearer. In the term of order 1,

$$f_1(\tau) = -\frac{1}{p} \sum_{i < j} \tau_{ij}.$$

each monomial τ_{ij} is represented by a graph in which two distinct nodes i, j are joined by a simple (i.e., with multiplicity 1) edge and the remaining nodes are singletons. Terms of order 2

$$f_2(\tau) = \frac{3}{2p(p+2)} \sum_{i < j} \tau_{ij}^2 + \frac{1}{p(2+p)} \sum_{i, j, k} \tau_{ij} \tau_{jk} + \frac{p+1}{(p-1)p(p+2)} \sum_{i, j, k, l} \tau_{ij} \tau_{lk},$$

are represented by three (types of) graphs with two edges: the first sum is associated with graphs that consist of $p-2$ singletons and two nodes joined by an edge of multiplicity two; the second sum with graphs with $p-3$ singletons and a connected component with 3 nodes and two simple edges; the third sum with graphs with $p-4$ singletons and two connected components, each having two nodes and one simple edge. The coefficients c_s of the

monomials of degree s in τ have different forms in general. However, for large p all coefficients are of the same form (Hikami and Brézin, 2006). Namely,

$$c_s = (-1)^s \frac{g}{\prod_{t=0}^{s-1} (p+2t)} \left(1 + O\left(\frac{2}{p}\right)\right),$$

where g is the degeneracy factor due to the multiplicities of the edges in the graph. For a graph with q edges with multiplicities (q_1, \dots, q_q) , the degeneracy is

$$g = \prod_{i=1}^q \frac{1}{q_i!} \prod_{m=1}^{q_i-1} (1+2m) = \prod_{i=1}^q \binom{2q_i}{q_i} \frac{1}{2^{q_i}}.$$

For example, the coefficients c_2 for the monomials of degree 2 in $f_2(\tau)$ above are in the large p limit $3/2p^2$ for the first and $1/p^2$ for the other two. Because of this formula, one can obtain the following large p expression for f

$$\prod_{i < j} \left(1 + \frac{2}{p} \tau_{ij}\right)^{-1/2}.$$

The details of its derivation are omitted in Hikami and Brézin (2006). We present them here. To make the notation more compact, let us use a bold symbol, \mathbf{a} , say, for an unordered pair of integers (i_a, j_a) . In other words, we are using a vector notation for the $P = p(p-1)/2$ distinct elements of τ , whose diagonal elements are zero. By applying the Taylor expansion of the inverse of the square root, we find

$$\begin{aligned} \prod_{i < j} \left(1 + \frac{2}{p} \tau_{ij}\right)^{-1/2} &= \prod_{\mathbf{a}=1}^P \left(1 + \frac{2}{p} \tau_{\mathbf{a}}\right)^{-1/2} = \prod_{\mathbf{a}=1}^P \sum_{k=0}^{\infty} (-1)^k \binom{2k}{k} \frac{1}{2^k} \left(\frac{\tau_{\mathbf{a}}}{p}\right)^k \\ &= \sum_{s=0}^{\infty} \frac{(-1)^s}{p^s} \frac{1}{2^s} \sum_{\mathbf{s}} \binom{2s_1}{s_1} \dots \binom{2s_P}{s_P} \tau_1^{s_1} \dots \tau_P^{s_P}, \end{aligned} \tag{8}$$

where the last sum is over the vectors $\mathbf{s} = (s_1, \dots, s_P)$ such that $0 \leq s_j \leq s$, with $\sum_{i=1}^P s_i = s$. Noticing that singletons in the graph just contribute 1, one now recognizes immediately that in (8) for each $s = 0, 1, \dots$, the sum is over all possible graphs with p nodes and with at most s edges whose multiplicities sum to s . These are exactly the graphs associated with the expansion of f . Since the Taylor series we have employed is convergent when $-1 < 2\tau_{ij}/p < 1$, the above result assumes $0 \leq 2\tau_{ij} < p$ (since τ_{ij} is non-negative). We can thus write (7) for large p as

$$I \approx e^{\sum_{i=1}^p x_i y_i} \cdot \prod_{i < j} \left(1 + \frac{2}{p} \tau_{ij}\right)^{-1/2}. \tag{9}$$

We use the approximation (9) in the following

Proof of Proposition 1.

We follow some steps as in Theorem 9.5.4 of Muirhead (2009). Consider first the integral

$$I(\ell) = \int_{O(p)} \text{etr}\left(-\frac{1}{2}\tilde{L}H^T\Sigma^{-1}H\right)(H^T dH)$$

where $H \in O(p)$, $(H^T dH)$ is the Haar measure in $O(p)$ (whose integral gives the volume of $O(p)$) and $\tilde{L} = \text{diag}(\ell_1, \dots, \ell_n, \ell, \dots, \ell)$ with ℓ repeated $(p-n)$ times. Partition $H = [H_1 : H_2]$, where $H_1 \in V_n(\mathbb{R}^p)$ and thus H_2 is a matrix of order $p \times (p-n)$ whose columns are orthonormal and orthogonal to those of H_1 . Since $H_2 H_2^T = I_p - H_1 H_1^T$, we have

$$\begin{aligned} \text{tr}(\tilde{L}H^T\Sigma^{-1}H) &= \text{tr}(LH_1^T\Sigma^{-1}H_1) + \ell \cdot \text{tr}(H_2^T\Sigma^{-1}H_2) \\ &= \text{tr}\left((L - \ell I_n)H_1^T\Sigma^{-1}H_1\right) + \ell \text{tr}\Sigma^{-1} \end{aligned}$$

where $L = \text{diag}(\ell_1, \dots, \ell_n)$. Using lemma 9.5.3 in Muirhead (2009), it follows

$$I(\ell) = \text{Vol}(O(p-n)) \text{etr}\left(-\frac{1}{2}\ell \cdot \Sigma^{-1}\right) \cdot J_n(\ell)$$

where $\text{Vol}(O(k)) = \frac{2^k \pi^{k^2/2}}{\Gamma_k(\frac{1}{2}k)}$ and

$$J_n(\ell) = \int_{V_n(\mathbb{R}^p)} \text{etr}\left(-\frac{1}{2}(L - \ell \cdot I_n)H_1^T\Sigma^{-1}H_1\right)(H_1^T dH_1).$$

Since the integrand function is bounded and $V_n(\mathbb{R}^p)$ is compact, in the limit $\ell \rightarrow 0$, we recover the integral of interest

$$J_n = \lim_{\ell \rightarrow 0} J_n(\ell) = \frac{1}{\text{Vol}(O(p-n))} \lim_{\ell \rightarrow 0} I(\ell)$$

with $\ell_1 > \ell_2 > \dots > \ell_n > 0$. Applying the large- p form (9) of the integral, we obtain

$$I(\ell) \approx \text{Vol}(O(p)) \cdot \exp\left(-\frac{1}{2} \sum_{i=1}^p \frac{\tilde{\ell}_i}{\lambda_i}\right) \cdot \prod_{i < j} \left(1 + \frac{1}{p} \left(\frac{\lambda_i - \lambda_j}{\lambda_i \lambda_j}\right) (\tilde{\ell}_i - \tilde{\ell}_j)\right)^{-1/2}$$

and hence the result, ignoring the ratio $\frac{\text{Vol}(O(p))}{\text{Vol}(O(p-n))} = \text{Vol}(V_n(\mathbb{R}^p))$. \square

Appendix B

Proof of Proposition 2 (The properties of the eigenvalue estimates).

- a. Let $a = 1, \dots, n$, and define $\psi_{a,b} = p + n(\ell_a - \ell_b)^2 / \ell_a \ell_b$ and $\ell_a - \ell_{a+1} = d_a$. Clearly, $d_a > 0$ and $\psi_{ab} > p$. Then

$$\begin{aligned} n\hat{\lambda}_a^1 &= \ell_a - \sum_{b=1}^n \frac{\ell_a - \ell_b}{p + n\left(\frac{1}{\ell_b} - \frac{1}{\ell_a}\right)(\ell_a - \ell_b)} = \ell_{a+1} + d_a - \sum_{b=1}^n \frac{\ell_a - \ell_{a+1} + \ell_{a+1} - \ell_b}{\psi_{a,b}} \\ &> \ell_{a+1} + d_a \left(1 - \frac{n}{p}\right) - \sum_{b=1}^n \frac{\ell_{a+1} - \ell_b}{\psi_{a,b}} \\ &\geq \ell_{a+1} - \sum_{b=1}^a \frac{\ell_{a+1} - \ell_b}{\psi_{a,b}} - \sum_{b=a+2}^n \frac{\ell_{a+1} - \ell_b}{\psi_{a,b}} \\ &\geq \ell_{a+1} - \sum_{b=1}^a \frac{\ell_{a+1} - \ell_b}{\psi_{a+1,b}} - \sum_{b=a+2}^n \frac{\ell_{a+1} - \ell_b}{\psi_{a,b}} \\ &\geq \ell_{a+1} - \sum_{b=1}^n \frac{\ell_{a+1} - \ell_b}{\psi_{a+1,b}} + \sum_{b=a+2}^n \frac{\ell_{a+1} - \ell_b}{\psi_{a+1,b}} - \sum_{b=a+2}^n \frac{\ell_{a+1} - \ell_b}{\psi_{a,b}} \\ &= n\hat{\lambda}_{a+1}^1 + \sum_{b=a+2}^n (\ell_{a+1} - \ell_b) \left(\frac{1}{\psi_{a+1,b}} - \frac{1}{\psi_{a,b}} \right) \\ &\geq n\hat{\lambda}_{a+1}^1 > 0 \end{aligned}$$

where we have used the fact that $\psi_{a+1,b} > \psi_{a,b}$ for $b = a$ and $\psi_{a,b} > \psi_{a+1,b}$ for $b = a+1$.

Hence $\hat{\lambda}_a^\kappa > \hat{\lambda}_{a+1}^\kappa$ for all $\kappa \in [0, 1]$, and $\hat{\lambda}_n^\kappa > \hat{\lambda}_r^\kappa$, $r = n+1, \dots, p$.

- b. It follows immediately from the fact that $\sum_{b,a=1}^n \frac{\ell_a - \ell_b}{p + n\left(\frac{1}{\ell_b} - \frac{1}{\ell_a}\right)(\ell_a - \ell_b)} = 0$. \square

Appendix C

In this appendix, using standard arguments (Muirhead, 2009), we compute the profile log-likelihood for the eigenvalues λ of the covariance matrix Σ . Let $\Sigma = V\Lambda V^\top$ be the spectral decomposition of Σ . The log-likelihood function obtained from (1) is

$$\mathcal{L}(\Lambda, V) \equiv \mathcal{L}(\Lambda, V \mid L, H_1) = -\frac{n}{2} \sum_{i=1}^p \ln \lambda_i - \frac{1}{2} \text{tr} \left(\Lambda^{-1} V^\top H_1 L (V^\top H_1)^\top \right)$$

Since $\text{tr}(\Lambda^{-1} A \Lambda A^\top) \geq \sum_{i=1}^n \frac{\ell_i}{\lambda_i}$, when A satisfies the condition $A^\top A = I$, with equality when A is one of the 2^n matrices M of dimensions $p \times n$ with components $M_{ij} = \pm \delta_{ij}$, where δ_{ij} is the Kronecker delta, we obtain

$$\mathcal{L}(\Lambda, V) \leq -\frac{n}{2} \sum_{i=1}^p \ln \lambda_i - \frac{1}{2} \sum_{i=1}^n \frac{\ell_i}{\lambda_i}$$

and thus the expression on the right-hand side is the profile log-likelihood $\mathcal{L}_P(\Lambda) = \mathcal{L}(\Lambda, \hat{V}_\Lambda)$. Thus the maximizer \hat{V}_Λ of the log-likelihood for a fixed value of Λ is a solution to $V^\top H_1 = M$. Since V is orthogonal, then $\hat{V} = [H_1 M_n; H_2 P]$, where M_n is the $n \times n$ matrix of the first n rows of M and $P \in O(p-n)$ any orthogonal matrix.

Appendix D

In this appendix we extend our algorithm to the case in which $\text{rank}(S) = q < n$, with q distinct positive eigenvalues. In the rank- q case, the density of S and the measure are obtained from those in the maximum rank case by replacing n with q (Díaz-García et al., 1997). Namely,

$$p(S)(dS) = K_q (\det \Sigma)^{-q/2} \text{etr}(-\Sigma^{-1} S/2) (\det L)^{(q-p-1)/2} (dS),$$

where $S = H_1 L H_1^\top$, with $L = \text{diag}(\ell_1, \dots, \ell_q)$, and the volume form written in terms of the Haar measure $(H_1^\top dH_1)$ on $V_q(\mathbb{R}^p)$ is

$$(dS) = 2^{-q} \prod_{i=1}^q \ell_i^{p-q} \prod_{i < j}^q (\ell_i - \ell_j) (H_1^\top dH_1) \wedge d\ell_1 \wedge \dots \wedge d\ell_q.$$

Accordingly, all non-maximum rank formulae, from the marginal density of the eigenvalues to the ML equations, are obtained by replacing n with q in the corresponding maximum rank equations. Thus, an exact solution to the ML equations gives all estimates to be $\sum_{a=1}^q \ell_a / pq$, and our proposed estimates have the form

$$\hat{\lambda}_a^\kappa = \frac{\kappa}{q} \left(\ell_a - \sum_{b=1}^q \frac{\ell_a - \ell_b}{p+q \left(\frac{1}{\ell_b} - \frac{1}{\ell_a} \right) (\ell_a - \ell_b)} \right) + (1-\kappa) \frac{\sum_{b=1}^q \ell_b}{pq}, \quad a = 1, \dots, q$$

$$\hat{\lambda}_r^\kappa = (1-\kappa) \frac{\sum_{b=1}^q \ell_b}{pq}, \quad r = q+1, \dots, p$$

with $0 < \kappa < 1$. Such estimates can be shown to be positive and ordered by following the same steps as in the proof of Proposition 2.

References

- Banerjee S, Monni S, and Wells MT (2016). A regularized profile likelihood approach to covariance matrix estimation. *Journal of Statistical Planning and Inference*, 179:36–59.
- Barndorff-Nielsen OE (1983). On a formula for the conditional distribution of the maximum likelihood estimator. *Biometrika*, 70:343–365.

- Barvinok A (2002). A Course in Convexity, volume 54 of Graduate Studies in Mathematics. American Mathematical Society.
- Bickel PJ and Levina E (2008). Regularized estimation of large covariance matrices. *The Annals of Statistics*, 36:199–227.
- Bien J and Tibshirani R (2011). Sparse estimation of a covariance matrix. *Biometrika*, 98:807–820. [PubMed: 23049130]
- Dey DK and Srinivasan C (1985). Estimation of a covariance matrix under Stein’s loss. *The Annals of Statistics*, 13:1581–1591.
- Díaz-García JA, Jáimez RG, and Mardia K (1997). Wishart and pseudo-Wishart distributions and some applications to shape theory. *Journal of Multivariate Analysis*, 63:73–87.
- El Karoui N (2008). Spectrum estimation for large dimensional covariance matrices using random matrix theory. *The Annals of Statistics*, 36:2757–2790.
- Golub TR, Slonim DK, Tamayo P, Huard C, Gaasenbeek M, Mesirov JP, Coller H, Loh ML, Downing JR, Caligiuri MA, Bloomfield CD, and Lander ES (1999). Molecular classification of cancer: class discovery and class prediction by gene expression monitoring. *Science*, 286:531–537. [PubMed: 10521349]
- Hess KR, Anderson K, Symmans WF, Valero V, Ibrahim N, Mejia JA, Booser D, Theriault RL, Buzdar AU, Dempsey PJ, et al. (2006). Pharmacogenomic predictor of sensitivity to preoperative chemotherapy with paclitaxel and fluorouracil, doxorubicin, and cyclophosphamide in breast cancer. *Journal of Clinical Oncology*, 24(26):4236–4244. [PubMed: 16896004]
- Hikami S and Brézin E (2006). WKB-expansion of the HarishChandra-Itzykson-Zuber integral for arbitrary β . *Progress of Theoretical Physics*, 116:441–502.
- Konno Y (2009). Shrinkage estimators for large covariance matrices in multivariate real and complex normal distributions under an invariant quadratic loss. *Journal of Multivariate Analysis*, 100:2237–2253.
- Ledoit O and Wolf M (2015). Spectrum estimation: A unified framework for covariance matrix estimation and PCA in large dimensions. *Journal of Multivariate Analysis*, 139:360–384.
- Ledoit O and Wolf M (2017). Numerical implementation of the QuEST function. *Computational Statistics & Data Analysis*, 115:199–223.
- Ledoit O and Wolf M (2018). Optimal estimation of a large-dimensional covariance matrix under Stein’s loss. *Bernoulli*, 24(4B):3791–3832.
- Ledoit O and Wolf M (2020). Shrinkage estimation of large covariance matrices: Keep it simple, statistician? Technical Report ECON Working Paper 327, Department of Economics, University of Zurich Available at <https://www.econ.uzh.ch/en/research/workingpapers.html>.
- Lin SP and Perlman MD (1985). A Monte Carlo comparison of four estimators of a covariance matrix. *Multivariate Analysis VI*.
- Mestre X (2008). Improved estimation of eigenvalues and eigenvectors of covariance matrices using their sample estimates. *IEEE Transactions on Information Theory*, 54:5113–5129.
- Muirhead RJ (2009). *Aspects of Multivariate Statistical Theory* Wiley Series in Probability and Statistics. Wiley.
- Naul B and Taylor J (2017). Sparse Steinian covariance estimation. *Journal of Computational and Graphical Statistics*, 26:355–366.
- Srivastava MS and Khatri CG (1979). *An Introduction to Multivariate Statistics*. North-Holland, Amsterdam.
- Stein C (1975). Estimation of a covariance matrix Rietz lecture, 39th Annual Meeting, IMS.
- Stein C (1986). Lectures on the theory of estimation of many parameters. *Journal of Soviet Mathematics*, 34(1):1373–1403.
- Takemura A (1984). An orthogonally invariant minimax estimator of the covariance matrix of a multivariate normal population. *Tsukuba J. Math*, 8:367–376.
- Tsukuma H (2016). Estimation of a high-dimensional covariance matrix with the Stein loss. *Journal of Multivariate Analysis*, 148:1–17.
- Uhlig H (1994). On singular Wishart and singular multivariate beta distributions. *The Annals of Statistics*, 22:395–405.

- van der Vaart HR (1961). On certain characteristics of the distribution of the latent roots of a symmetric random matrix under general conditions. *The Annals of Mathematical Statistics*, 32:864–873.
- Won J-H, Lim J, Kim S-J, and Rajaratnam B (2013). Condition-number-regularized covariance estimation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 75(3):427–450. [PubMed: 23730197]
- Yao J, Kammoun A, and Najim J (2012). Eigenvalue estimation of parameterized covariance matrices of large dimensional data. *IEEE Transactions on Signal Processing*, 60:5893–5905.
- Yi F and Zou H (2013). SURE-tuned tapering estimation of large covariance matrices. *Computational Statistics & Data Analysis*, 58:339–351. The Third Special Issue on Statistical Signal Extraction and Filtering.

Loss functions panels for Σ, γ, p




Figure 1: κ -selection for Σ_1 to Σ_4 :

Chosen $\hat{\kappa}$ via bootstrap (dotted line) and CV (dashed line) are compared with oracle κ' (solid line) with nine loss functions (each panel) for various combinations of true Σ (Σ_1 to Σ_4), $\gamma = p/n$ and p (see Section 4).

Loss functions panels for Σ , γ , p

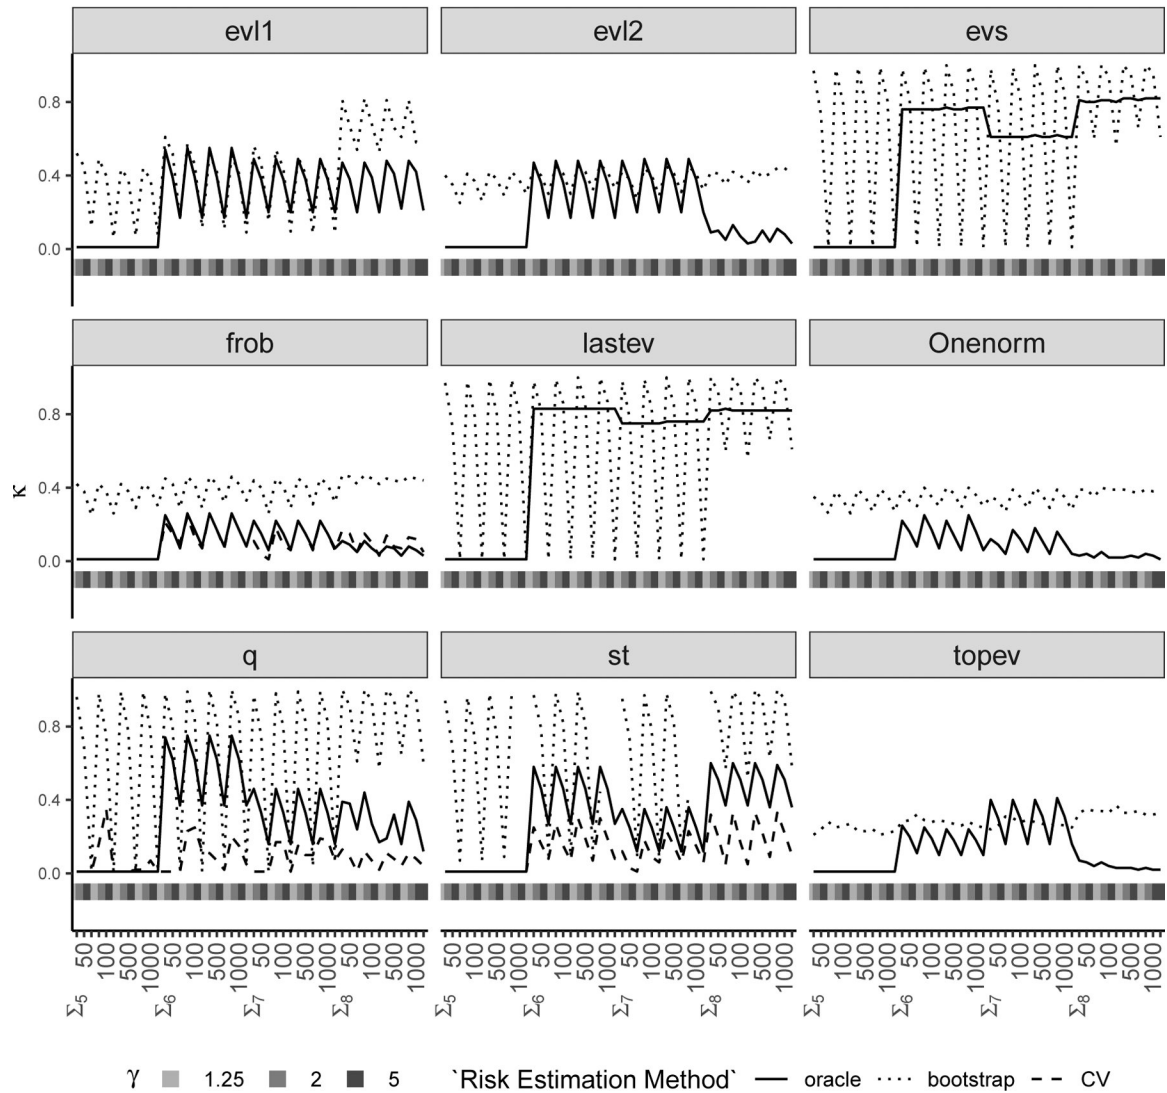


Figure 2: κ -selection for Σ_5 to Σ_8 :

Chosen $\hat{\kappa}$ via bootstrap (dotted line) and CV (dashed line) are compared with oracle κ' (solid line) with nine loss functions (each panel) for various combinations of true Σ (Σ_5 to Σ_8), $\gamma = p/n$ and p (see Section 4).




Figure 3: PRIAL comparison with LW1.

The heatmap shows values of PRIAL of $\hat{\Sigma}$ with respect to $\hat{\Sigma}^{LW1}$ for various simulation scenarios based on covariance structure, values of p and γ (rows) and various loss functions (columns). Top panel shows counts and histogram of various PRIAL values. PRIAL values are scaled column-wise for visual clarity. Red (blue) shades mean $PRIAL > 0$ ($PRIAL < 0$) indicating our estimator is better (worse) than LW1.




Figure 4: PRIAL comparison with LW2.

The heatmap shows values of PRIAL of $\hat{\Sigma}$ with respect to $\hat{\Sigma}^{LW_2}$ for various simulation scenarios based on covariance structure, values of p and γ (rows) and various loss functions (columns). Top panel shows counts and histogram of various PRIAL values. PRIAL values are scaled column-wise for visual clarity. Red (blue) shades mean $PRIAL > 0$ ($PRIAL < 0$) indicating our estimator is better (worse) than LW2.




Figure 5: PRIAL comparison with LW3.

The heatmap shows values of PRIAL of $\hat{\lambda}$ with respect to $\hat{\lambda}^{LW3}$ for various simulation scenarios based on covariance structure, values of p and γ (rows) and various loss functions (columns). Top panel shows counts and histogram of various PRIAL values. PRIAL values are scaled column-wise for visual clarity. Red (blue) shades mean $PRIAL > 0$ ($PRIAL < 0$) indicating our estimator is better (worse) than LW3.

Table 1:

Error rates for LDA analysis of breast cancer data

Estimator	Reference ($\bar{\Sigma}$)	MCR	Sens	Spec
LW1	-	0.275	0.385	0.842
LW2	-	0.333	0.538	0.711
$\tilde{\Sigma}$ (Boot-frob)	$\tilde{\lambda}^{NS}$	0.392	0.538	0.658
$\tilde{\Sigma}$ (Boot-st)	LW1	0.235	0.462	0.868
$\tilde{\Sigma}$ (Boot-OneNorm)	LW2	0.176	0.538	0.921
$\tilde{\Sigma}$ (CV-q)	-	0.216	0.538	0.868

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript