



HAL
open science

Shrinking the eigenvalues of M-estimators of covariance matrix

Esa Ollila, Daniel Palomar, Frédéric Pascal

► **To cite this version:**

Esa Ollila, Daniel Palomar, Frédéric Pascal. Shrinking the eigenvalues of M-estimators of covariance matrix. IEEE Transactions on Signal Processing, 2020, pp.256 - 269. <10.1109/TSP.2020.3043952>. <hal-03105531>

HAL Id: hal-03105531

<https://centralesupelec.hal.science/hal-03105531v1>

Submitted on 11 Jan 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Shrinking the eigenvalues of M-estimators of covariance matrix

Esa Ollila, *Member, IEEE*, Daniel P. Palomar, *Fellow, IEEE* and Frédéric Pascal, *Senior Member, IEEE*

Abstract—A highly popular regularized (shrinkage) covariance matrix estimator is the shrinkage sample covariance matrix (SCM) which shares the same set of eigenvectors as the SCM but shrinks its eigenvalues toward the grand mean of the eigenvalues of the SCM. In this paper, a more general approach is considered in which the SCM is replaced by an M-estimator of scatter matrix and a fully automatic data adaptive method to compute the optimal shrinkage parameter with minimum mean squared error is proposed. Our approach permits the use of any weight function such as Gaussian, Huber’s, Tyler’s, or t weight functions, all of which are commonly used in M-estimation framework. Our simulation examples illustrate that shrinkage M-estimators based on the proposed optimal tuning combined with robust weight function do not loose in performance to shrinkage SCM estimator when the data is Gaussian, but provide significantly improved performance when the data is sampled from an unspecified heavy-tailed elliptically symmetric distribution. Also, real-world and synthetic stock market data validate the performance of the proposed method in practical applications.

Index Terms—M-estimators, sample covariance matrix, shrinkage, regularization, elliptically symmetric distributions

I. INTRODUCTION

Consider a sample of p -dimensional vectors $\{\mathbf{x}_i\}_{i=1}^n$ sampled from a distribution of a random vector \mathbf{x} with mean vector equal to zero (i.e., $\mathbb{E}[\mathbf{x}] = \mathbf{0}$). One of the first task in the analysis of high-dimensional data is to estimate the covariance matrix. The most commonly used estimator is the sample covariance matrix (SCM), $\mathbf{S} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i^\top$, but its main drawbacks are its loss of efficiency when sampling from distributions which have heavier tails than the multivariate normal (MVN) distribution and its sensitivity to outliers. Although being unbiased estimator of the covariance matrix $\text{cov}(\mathbf{x}) = \mathbb{E}[\mathbf{x}\mathbf{x}^\top]$ for any sample length $n \geq 1$, it is well-known that the eigenvalues are poorly estimated when n is not orders of magnitude larger than p . In such cases, one commonly uses a *regularized SCM (RSCM)* with a linear shrinkage towards a scaled identity matrix, defined as

$$\mathbf{S}_\beta = \beta \mathbf{S} + (1 - \beta) \frac{\text{tr}(\mathbf{S})}{p} \mathbf{I}, \quad (1)$$

where $\beta \in [0, 1]$ is the regularization parameter. The RSCM \mathbf{S}_β shares the same set of eigenvectors as the SCM \mathbf{S} , but its eigenvalues are shrunked towards the grand mean of the eigenvalues of the SCM \mathbf{S} . That is, if d_1, \dots, d_p denote the eigenvalues of \mathbf{S} , then $\beta d_j + (1 - \beta)\bar{d}$ are the eigenvalues of

\mathbf{S}_β , where $\bar{d} = p^{-1} \sum_{j=1}^p d_j$. Optimal computation of β such that \mathbf{S}_β has minimum mean squared error (MMSE) has been developed for example in [1], [2] or in [3] for certain structured target matrices. A Bayesian approach has been considered in [4].

However, the estimator in (1) remains sensitive to outliers and non-Gaussianity. M-estimators of scatter [5] are popular robust alternatives to SCM. We consider the situation where $n > p$ and hence a conventional M-estimator of scatter $\hat{\Sigma}$ exists under mild conditions on the data (see [6]) and can thus be used in place of the SCM \mathbf{S} in (1). We then propose a fully automatic data adaptive method to compute the optimal shrinkage parameter β . First, we derive an approximation for the optimal parameter β_o attaining the minimum MSE and then propose a data adaptive method for its computation. The benefit of our approach is that it can be easily applied to any M-estimator using any weight function $u(t)$. Our simulation examples illustrate that a shrinkage M-estimator using the proposed tuning and a robust weight function does not loose in performance to optimal shrinkage SCM estimator when the data is Gaussian, but is able to provide significantly improved performance in the case of heavy-tailed data and in the presence of outliers.

Earlier works, e.g., in [7], [8], [9], [10], [11], [12], [13], [14], [15], proposed regularized M-estimators of scatter matrix either by adding a penalty function to M-estimation objective function or a diagonal loading term to the respective first-order solution (M-estimating equation). We consider a simpler approach that uses a conventional M-estimator and shrinks its eigenvalues to the grand mean of the eigenvalues. Our approach permits computation of the optimal shrinkage parameter for any M-estimation weight function. Preliminary study of the proposed estimators has appeared in a conference proceeding [16].

Finally, we note some related but different approaches to what is pursued in this paper. For example, covariance matrix estimation in the low sample size large dimensionality setting commonly arises in radar signal processing, where often some constrained, mismatched or structured estimation framework of the covariance matrix is exploited. See e.g., [17], [18], [19], [20], [21], [22]. On the other hand, there are also other approaches for parameter tuning of regularized covariance matrix estimators such as cross-validation or expected likelihood approach [23], [24], [25].

The paper is structured as follows. Section II introduces the proposed shrinkage M-estimator framework while Section III discusses automatic computation of the optimal shrinkage parameter under the assumption of sampling from unspecified elliptical distribution. Extension to complex case is discussed

E. Ollila is with the Department of Signal Processing and Acoustics, Aalto University, P.O. Box 15400, FI-00076 Aalto, Finland. D. Palomar is with the Hong Kong University of Science and Technology, Hong Kong. F. Pascal is with Université Paris-Saclay, CNRS, CentraleSupélec, Laboratoire des signaux et systèmes, 91190, Gif-sur-Yvette, France.

Manuscript received November xx, 202X; revised xx, 202X.

in subsection III-A. Section IV addresses the most commonly used M-estimators, namely, the Gaussian weight function, Huber's weight function, Tyler's weight and t -distribution weight function. We provide simulation studies in Section V and experimental results in Section VI. In Section VI we validate the promising performance of the proposed approach both in the case of real-world and synthetic stock market data. Finally, Section VII concludes, while proofs of theorems and lemmas are kept in the Appendix.

II. SHRINKAGE M-ESTIMATORS OF SCATTER

In this paper, we assume that $n > p$ (except for Gaussian loss) and consider an M-estimator of scatter matrix [5] that solves an estimating equation

$$\hat{\Sigma} = \frac{1}{n} \sum_{i=1}^n u(\mathbf{x}_i^\top \hat{\Sigma}^{-1} \mathbf{x}_i) \mathbf{x}_i \mathbf{x}_i^\top, \quad (2)$$

where $u : [0, \infty) \rightarrow [0, \infty)$ is a non-increasing *weight function*. An M-estimator is a sort of adaptively weighted SCM with weights determined by function $u(\cdot)$. To guarantee existence of the solution, it is required that the data verifies the condition stated in [6]. An M-estimator of scatter which shrinks the eigenvalues towards the grand mean of the eigenvalues is then defined as:

$$\hat{\Sigma}_\beta = \beta \hat{\Sigma} + (1 - \beta) \frac{\text{tr}(\hat{\Sigma})}{p} \mathbf{I}. \quad (3)$$

Thus $\hat{\Sigma}_\beta$ is indexed by the shrinkage parameter $\beta \in [0, 1]$. If $\beta = 1$, then $\hat{\Sigma}_\beta$ coincides with the conventional M-estimator in (2) while if $\beta = 0$, then $\hat{\Sigma}_\beta$ equals an identity matrix scaled by mean of the eigenvalues of $\hat{\Sigma}$. Next we discuss the commonly used weight functions u .

The RSCM \mathbf{S}_β in (1) is obtained when one uses the *Gaussian weight function* $u_G(t) = 1 \forall t$. Terminology ‘Gaussian’ stems from the fact that $\hat{\Sigma} = \mathbf{S}$ is the maximum likelihood estimate (MLE) of the covariance matrix of MVN distribution. We return to this in Section III. *Huber's weight function* is defined as

$$u_h(t; c) = \begin{cases} 1/b, & \text{for } t \leq c^2 \\ c^2/(tb), & \text{for } t > c^2 \end{cases} \quad (4)$$

where $c > 0$ is a user defined tuning constant that determines the robustness and efficiency of the estimator and b is a scaling factor; see subsection IV-B for more details. Another popular choice is *MVT-weight function* [6]:

$$u_\tau(t; \nu) = \frac{p + \nu}{\nu + t} \quad (5)$$

in which case the corresponding M-estimator $\hat{\Sigma}$ is also the MLE of the scatter matrix parameter of the *multivariate t (MVT) distribution* with $\nu > 0$ degrees of freedom (d.o.f.). We return to this estimator in subsection IV-D. Finally, another classic choice, with nice robustness properties, is Tyler's [26] M-estimator, in which case the weight function is

$$u_{\text{Ty}}(t) = \frac{p}{t}. \quad (6)$$

Both Huber's and MVT-weight function yield Tyler's weight function as special cases; namely, for $\nu = 0$, one notices that $u_\tau(t; \nu = 0) = u_{\text{Ty}}(t)$ and in the limit case as $c \rightarrow 0$ Huber's weight function tends to Tyler's weight function.

We would like to stress that $n > p$ is a necessary assumption for all but Gaussian weight functions for a solution $\hat{\Sigma}$ to (2) to exist. We do not include the limit case $n = p$ since any affine equivariant M-estimator $\hat{\Sigma}$ when $n = p$ and data is in general position is just a scalar multiple of the SCM \mathbf{S} , that is, $\hat{\Sigma} = \gamma \mathbf{S}$ for some $\gamma > 0$ [27]. For example, M-estimator based on Huber's or t -weights are affine equivariant. Moreover, note that Theorem 2 in [28] ensures similar results in the large sample regime. Namely, the following convergence is proved

$$\|\hat{\Sigma} - \hat{\mathbf{S}}_{np}\| \xrightarrow[n, p \rightarrow \infty]{a.s.} 0$$

with $n/p \rightarrow c \in (0, 1)$ and $\hat{\mathbf{S}}_{np}$ is an appropriate weighted SCM, and the norm denotes the spectral norm.

An M-estimator is a consistent estimator of the M-functional of scatter matrix, defined as

$$\Sigma_0 = \mathbb{E}[u(\mathbf{x}^\top \Sigma_0^{-1} \mathbf{x}) \mathbf{x} \mathbf{x}^\top]. \quad (7)$$

If the population M-functional Σ_0 is known, then by defining a *1-step estimator*

$$\mathbf{C} = \frac{1}{n} \sum_{i=1}^n u(\mathbf{x}_i^\top \Sigma_0^{-1} \mathbf{x}_i) \mathbf{x}_i \mathbf{x}_i^\top, \quad (8)$$

we can compute

$$\mathbf{C}_\beta = \beta \mathbf{C} + (1 - \beta) [\text{tr}(\mathbf{C})/p] \mathbf{I}, \quad (9)$$

which serves as a proxy for $\hat{\Sigma}_\beta$. Naturally, such an estimator is fictional, since the initial value Σ_0 is unknown. The 1-step estimator \mathbf{C} is, by its construction, an unbiased estimator of Σ_0 , i.e., $\mathbb{E}[\mathbf{C}] = \Sigma_0$.

Ideally we would like to find the value of $\beta \in [0, 1]$ for which the corresponding estimator $\hat{\Sigma}_\beta$ attains the minimum MSE, that is,

$$\beta_o = \arg \min_{\beta} \left\{ \text{MSE}(\hat{\Sigma}_\beta) = \mathbb{E} \left[\|\hat{\Sigma}_\beta - \Sigma_0\|_F^2 \right] \right\}, \quad (10)$$

where $\|\cdot\|_F$ denotes the Frobenius matrix norm (i.e., $\|\mathbf{A}\|_F^2 = \text{tr}(\mathbf{A}^\top \mathbf{A})$ and $\|\mathbf{A}\|_F^2 = \text{tr}(\mathbf{A}^H \mathbf{A})$ for real-valued and complex-valued matrices, respectively, where $(\cdot)^H$ denotes the Hermitian transpose). Since solving (10) is not feasible due to the implicit form of M-estimators, we instead solve the following much simpler problem:

$$\beta_o^{\text{app}} = \arg \min_{\beta} \left\{ \text{MSE}(\mathbf{C}_\beta) = \mathbb{E} \left[\|\mathbf{C}_\beta - \Sigma_0\|_F^2 \right] \right\}. \quad (11)$$

Such approach was also used in [29] to derive an optimal parameter for the shrinkage Tyler's M-estimator of scatter proposed by the authors,

Before stating the expression for β_o^{app} we introduce a *sphericity* measure of scatter, defined as

$$\gamma = \frac{p \text{tr}(\Sigma_0^2)}{\text{tr}(\Sigma_0)^2}. \quad (12)$$

Sphericity γ [30], [31] measures how close Σ_0 is to a scaled identity matrix: $\gamma \in [1, p]$ where $\gamma = 1$ if and only if $\Sigma_0 \propto \mathbf{I}$ and $\gamma = p$ if Σ_0 has rank equal to 1.

Theorem 1. *Suppose $\mathbf{x}_1, \dots, \mathbf{x}_n$ is an i.i.d. random sample from any p -variate distribution, and u is a weight function for which the expectation $\mathbb{E}[\text{tr}(\mathbf{C}^2)]$ exists. The oracle parameter β_o^{app} in (11) is*

$$\beta_o^{\text{app}} = \frac{\|\Sigma_0 - \eta_o \mathbf{I}\|_{\text{F}}^2}{\mathbb{E}\left[\|\mathbf{C} - (\text{tr}(\mathbf{C})/p)\mathbf{I}\|_{\text{F}}^2\right]} \quad (13)$$

$$= \frac{p(\gamma - 1)\eta_o^2}{\mathbb{E}[\text{tr}(\mathbf{C}^2)] - p^{-1}\mathbb{E}[\text{tr}(\mathbf{C})]^2} \quad (14)$$

where $\eta_o = \text{tr}(\Sigma_0)/p$ and γ is defined in (12). Note that $\beta_o^{\text{app}} \in [0, 1)$ and the value of the MSE at the optimum is

$$\text{MSE}(\mathbf{C}_{\beta_o^{\text{app}}}) = \frac{\mathbb{E}[\text{tr}(\mathbf{C}^2)] - \text{tr}(\Sigma_0)^2}{p} + (1 - \beta_o^{\text{app}})\|\Sigma_0 - \eta_o \mathbf{I}\|_{\text{F}}^2. \quad (15)$$

Proof. The proof is postponed to Appendix A. \square

In the next section, we derive a more explicit form of β_o^{app} by assuming that the data is generated from an unspecified elliptically symmetric distribution.

III. SHRINKAGE PARAMETER FOR ELLIPTICAL SAMPLES

Maronna [5] developed M-estimators of scatter matrices originally within the framework of elliptically symmetric distributions [32], [33]. The probability density function (p.d.f.) of centered (zero mean) elliptically distributed random vector, denoted by $\mathbf{x} \sim \mathcal{E}_p(\mathbf{0}, \Sigma, g)$, is

$$f(\mathbf{x}) = C_{p,g} |\Sigma|^{-1/2} g(\mathbf{x}^\top \Sigma^{-1} \mathbf{x}), \quad (16)$$

where Σ is a positive definite symmetric matrix parameter, called the scatter matrix, $g : [0, \infty) \rightarrow [0, \infty)$ is the *density generator*, which is a fixed function that is independent of \mathbf{x} and Σ , and $C_{p,g}$ is a normalizing constant ensuring that $f(\mathbf{x})$ integrates to 1. The density generator g determines the elliptical distribution. For example, the MVN distribution $\mathcal{N}_p(\mathbf{0}, \Sigma)$ is obtained when $g(t) = \exp(-t/2)$ and the t -distribution with ν d.o.f., denoted $\mathbf{x} \sim t_\nu(\mathbf{0}, \Sigma)$, is obtained when $g(t) = (1 + t/\nu)^{-(p+\nu)/2}$. Then the weight function for the MLE of scatter matrix corresponds to the case that the weight function is of the form $u(t) \propto -g'(t)/g(t)$. This yields (5) as the weight function for the MLE of scatter for t -distribution. If the second moments of \mathbf{x} exists, then g can be defined so that Σ represents the covariance matrix of \mathbf{x} , i.e., $\Sigma = \text{cov}(\mathbf{x})$; see [33] for details.

When $\mathbf{x} \sim \mathcal{E}_p(\mathbf{0}, \Sigma, g)$, then the M-functional Σ_0 in (7) is related to underlying scatter matrix parameter Σ via the relationship

$$\Sigma_0 = \sigma \Sigma, \quad (17)$$

where $\sigma > 0$ is a solution to an equation

$$\mathbb{E}\left[\psi\left(\frac{r^2}{\sigma}\right)\right] = p, \quad (18)$$

where $\psi(t) = u(t)t$ and $r = \|\Sigma^{-1/2} \mathbf{x}\|$. Often σ needs to be solved numerically from (18) but in some cases an analytic

expression can be derived. If $\mathbf{x} \sim \mathcal{E}_p(\mathbf{0}, \Sigma, g)$ and the used weight function matches with the data distribution, so $u(t) \propto -g'(t)/g(t)$, then $\sigma = 1$.

Next we derive expressions for $\mathbb{E}[\text{tr}(\mathbf{C}^2)]$ and $\mathbb{E}[\text{tr}(\mathbf{C}^2)]$ appearing in the denominator of β_o^{app} in (14). These depend on a constant ψ_1 (which depends on weight function u via $\psi(t) = u(t)t$) as follows:

$$\psi_1 = \frac{1}{p(p+2)} \mathbb{E}\left[\psi\left(\frac{r^2}{\sigma}\right)^2\right], \quad (19)$$

where the statistical expectation is again computed w.r.t. distribution of the positive random variable $r = \|\Sigma^{-1/2} \mathbf{x}\|$.

Lemma 1. *Suppose $\mathbf{x}_1, \dots, \mathbf{x}_n$ is an i.i.d. random sample from $\mathcal{E}_p(\mathbf{0}, \Sigma, g)$. Then*

$$\mathbb{E}[\text{tr}(\mathbf{C}^2)] = \left(1 + \frac{2\psi_1 - 1}{n}\right) \text{tr}(\Sigma_0^2) + \frac{\psi_1}{n} \text{tr}(\Sigma_0)^2$$

and

$$\mathbb{E}[\text{tr}(\mathbf{C})^2] = \frac{2\psi_1}{n} \text{tr}(\Sigma_0^2) + \left(1 + \frac{\psi_1 - 1}{n}\right) \text{tr}(\Sigma_0)^2$$

given that expectation (19) exists.

Proof. The proof is given in Appendix B. \square

This then yields the following main result.

Theorem 2. *Let $\mathbf{x}_1, \dots, \mathbf{x}_n$ denote an i.i.d. random sample from an elliptical distribution $\mathcal{E}_p(\mathbf{0}, \Sigma, g)$. Then the oracle parameter β_o^{app} that minimizes the MSE in Theorem 1 is*

$$\beta_o^{\text{app}} = \frac{n(\gamma - 1)}{n(\gamma - 1)(1 - 1/n) + \psi_1(1 - 1/p)(2\gamma + p)} \quad (20)$$

where γ is defined in (12) and ψ_1 in (19).

Proof. Follows from Theorem 1 after substituting the values for $\mathbb{E}[\text{tr}(\mathbf{C}^2)]$ and $\mathbb{E}[\text{tr}(\mathbf{C})^2]$ derived in Lemma 1 into the denominator of β_o^{app} in (14). \square

A closely related result is derived in [34, Theorem 1]. Namely, [34] considers oracle estimator as in (9) but using a shrinkage target equal to identity matrix \mathbf{I} instead of $[\text{tr}(\mathbf{C})/p]\mathbf{I}$ as in this paper. This is due to the fact that [34] assumes that $\text{tr}(\Sigma) = p$. Another difference is that [34] assumes that $\Sigma_0 = \Sigma$ (so $\sigma = 1$), i.e., that the used M-estimator is consistent to the scatter matrix of the underlying elliptical population. This assumption implies knowledge of the underlying elliptical distribution in which case it is natural to use the ML-weight $u_{\text{ML}}(t) = -2g(t)/g'(t)$ as was done in [34]. Thus Theorem 2 compared to [34, Theorem 1] holds in the more general case when the scale $\text{tr}(\Sigma)$ is not known *a priori* and no assumption on the knowledge of the elliptical distribution is imposed. Furthermore, in the next subsection, we extend the result to the complex-valued case which was not considered in [34].

Lemma 1 also allows to construct an unbiased estimate of $\vartheta = \text{tr}(\Sigma_0^2)/p$ as is shown next.

Theorem 3. Suppose $\mathbf{x}_1, \dots, \mathbf{x}_n$ is an i.i.d. random sample from $\mathcal{E}_p(\mathbf{0}, \Sigma, g)$. Then an unbiased estimate of $\vartheta = \text{tr}(\Sigma_0^2)/p$ for any finite n and any p is

$$\hat{\vartheta} = b_n \left(\frac{\text{tr}(\mathbf{C}^2)}{p} - \psi_1 a_n \frac{p}{n} \left[\frac{\text{tr}(\mathbf{C})}{p} \right]^2 \right), \quad (21)$$

where

$$a_n = \frac{n}{n + \psi_1 - 1} \quad \text{and} \quad b_n = \frac{n}{n - 1} \left(\frac{n - 1 + \psi_1}{n - 1 + 3\psi_1} \right) \quad (22)$$

given that expectation (19), defining ψ_1 , exists.

Proof. First note that

$$\mathbb{E}[\hat{\vartheta}] = b_n (\mathbb{E}[\text{tr}(\mathbf{C}^2)]/p - a_n (p/n) \mathbb{E}[\text{tr}(\mathbf{C})^2]/p^2). \quad (23)$$

Then substituting the values of $\mathbb{E}[\text{tr}(\mathbf{C}^2)]$ and $\mathbb{E}[\text{tr}(\mathbf{C})^2]$ from Lemma 1 into (23) yields $\mathbb{E}[\hat{\vartheta}] = \text{tr}(\Sigma_0^2)/p$. \square

It is instructive to consider in more detail the case of Gaussian loss. In this case, \mathbf{C} equals the SCM, $\mathbf{C} = \mathbf{S}$, and the statistic $\hat{\vartheta}$ no longer depends on the unknown Σ_0 . Furthermore, if data is generated from a Gaussian distribution $\mathcal{N}_p(\mathbf{0}, \Sigma)$, then $\psi_1 = 1$, $\Sigma_0 = \Sigma$ and $\hat{\vartheta}$ in (21) reduces to the estimator that is identical to one proposed by [31, Lemma 2.1] in the case that location is known ($\mu = \mathbf{0}$); see also [3, Theorem 2]. Moreover, [2, Theorem 4] is obtained in the general elliptical case, again assuming known location parameter ($\mu = \mathbf{0}$).

A. An extension to the complex-valued case

Consider the case that \mathbf{x}_i , $i = 1, \dots, n$ are complex-valued (i.e., $\mathbf{x}_i \in \mathbb{C}^p$) and represent a random sample from a (circular) complex elliptically symmetric (CES) distribution [33]. The p.d.f. of a random vector $\mathbf{x} \in \mathbb{C}^p$ with centered (zero mean) CES distribution, denoted using same notation $\mathbf{x} \sim \mathcal{E}_p(\mathbf{0}, \Sigma, g)$, is

$$f(\mathbf{x}) = C_{p,g} |\Sigma|^{-1} g(\mathbf{x}^H \Sigma^{-1} \mathbf{x}), \quad \mathbf{x} \in \mathbb{C}^p,$$

where Σ is the positive definite hermitian (PDH) matrix parameter, called the scatter matrix, $g: [0, \infty) \rightarrow [0, \infty)$ is the density generator, which is a fixed function that is independent of \mathbf{x} and Σ , and $C_{p,g}$ is a normalizing constant ensuring that $f(\mathbf{x})$ integrates to 1.

An M-estimator of scatter matrix $\hat{\Sigma}$ is a PDH matrix that solves an estimating equation

$$\hat{\Sigma} = \frac{1}{n} \sum_{i=1}^n u(\mathbf{x}_i^H \hat{\Sigma}^{-1} \mathbf{x}_i) \mathbf{x}_i \mathbf{x}_i^H, \quad (24)$$

where $u: [0, \infty) \rightarrow [0, \infty)$ is a non-increasing weight function. Again, $u(t) = 1$ gives the SCM $\mathbf{S} = \frac{1}{n} \sum_i \mathbf{x}_i \mathbf{x}_i^H$, Huber's and Tyler's weight functions are as earlier, stated in (4) and (6), respectively, whereas $u_t(t; \nu) = \frac{2p+\nu}{\nu+2t}$ corresponds to MLE of the scatter matrix parameter when sampling from a p -variate complex t -distribution with ν d.o.f. We refer to [33], [35] for more details. We may now define the shrinkage M-estimator as in (3). Definitions (7)-(9) hold also in the complex-valued case with obvious modifications (replacing transpose by the Hermitian transpose).

Theorem 1 did not require an assumption that random vectors are real-valued, i.e., it holds also when $\mathbf{x}_1, \dots, \mathbf{x}_n$ are i.i.d. complex-valued random vectors. This means that we only need to derive expectations in Lemma 1 in the case of random sampling from a CES distribution. First, we define the parameter ψ_1 in the complex-valued case as

$$\psi_1 = \frac{1}{p(p+1)} \mathbb{E} \left[\psi \left(\frac{r^2}{\sigma} \right)^2 \right], \quad (25)$$

where the expectation is w.r.t. $r = \|\Sigma^{-1/2} \mathbf{x}\|$, where $\mathbf{x} \sim \mathcal{E}_p(\mathbf{0}, \Sigma, g)$. The analog of Lemma 1 to complex case is given next.

Lemma 2. Suppose $\mathbf{x}_1, \dots, \mathbf{x}_n$ is an i.i.d. random sample from a (circular) complex elliptically symmetric distribution $\mathcal{E}_p(\mathbf{0}, \Sigma, g)$. Then

$$\mathbb{E}[\text{tr}(\mathbf{C}^2)] = \left(1 + \frac{\psi_1 - 1}{n} \right) \text{tr}(\Sigma_0^2) + \frac{\psi_1}{n} \text{tr}(\Sigma_0)^2$$

and

$$\mathbb{E}[\text{tr}(\mathbf{C})^2] = \frac{\psi_1}{n} \text{tr}(\Sigma_0^2) + \left(1 + \frac{\psi_1 - 1}{n} \right) \text{tr}(\Sigma_0)^2$$

given that expectation (25) exists.

Proof. The proof is given in Appendix C. \square

Plugging in the expectations above into β_o^{app} derived in Theorem 1 yields the following result.

Theorem 4. Let $\mathbf{x}_1, \dots, \mathbf{x}_n$ denote an i.i.d. random sample from a (circular) p -variate complex elliptically symmetric distribution $\mathcal{E}_p(\mathbf{0}, \Sigma, g)$ and assume that expectation (25) exists. Then the oracle parameter β_o^{app} that minimizes the MSE in Theorem 1 is

$$\beta_o^{\text{app}} = \frac{n(\gamma - 1)}{n(\gamma - 1)(1 - 1/n) + \psi_1(1 - 1/p)(\gamma + p)}, \quad (26)$$

where γ is defined in (12) and ψ_1 in (25).

Furthermore, $\hat{\vartheta}$ defined as in (21) with

$$a_n = \frac{n}{n + \psi_1 - 1} \quad \text{and} \quad b_n = \frac{n}{n - 1} \left(\frac{n - 1 + \psi_1}{n - 1 + 2\psi_1} \right)$$

is an unbiased estimate of $\vartheta = \text{tr}(\Sigma_0^2)/p$ for any finite n and any p .

B. Computing the shrinkage parameter

In order to construct a data-adaptive estimate of β_o^{app} (either in real- or complex-valued cases), all we need to estimate is the sphericity γ and the constant ψ_1 . An estimate $\hat{\psi}_1$ of ψ_1 is constructed separately for each weight function (Gaussian, Huber's, MVT- and Tyler's weight function) in Section IV. However, it is also possible to use an empirical (sample mean) version of (19). Next we discuss computation of the sphericity estimator $\hat{\gamma}$.

As an estimator $\hat{\gamma}$ we use the estimate derived in [14], which was named as Ell1-estimator in [2], and defined as

$$\hat{\gamma}^{\text{Ell1*}} = \frac{n}{n-1} \left(p \text{tr} \left(\frac{1}{n} \sum_{i=1}^n \frac{\mathbf{x}_i \mathbf{x}_i^T}{\|\mathbf{x}_i\|^2} \right) - \frac{p}{n} \right) \quad (27)$$

which for complex-valued case is defined analogously (transpose replaced with the Hermitian transpose).

Next recall that $\hat{\nu}$ defined in (21) is an unbiased estimator of $\text{tr}(\mathbf{\Sigma}_0^2)/p$. This statistics depends on \mathbf{C} and ψ_1 which are unknown, but a plug-in estimate of $\hat{\nu}$ can be constructed by replacing \mathbf{C} and ψ_1 with $\hat{\mathbf{\Sigma}}$ and $\hat{\psi}_1$, respectively. Dividing this plug-in statistic further by $(\text{tr}(\hat{\mathbf{\Sigma}})/p)^2$, leads to another estimator of sphericity, named as Ell2-estimator, and defined as

$$\hat{\gamma}^{\text{Ell2}^*} = \hat{b}_n \left(\frac{p \text{tr}(\hat{\mathbf{\Sigma}}^2)}{\text{tr}(\hat{\mathbf{\Sigma}})^2} - \hat{\psi}_1 \hat{a}_n \frac{p}{n} \right), \quad (28)$$

where the constants \hat{a}_n and \hat{b}_n are as defined in Theorem 3 (and Theorem 4 in complex case) but with ψ_1 replaced by its estimate $\hat{\psi}_1$. When one uses Gaussian weight function, then $\hat{\mathbf{\Sigma}} = \mathbf{S}$ and $\hat{\psi}_1 = 1 + \hat{\kappa}$, where $\hat{\kappa}$ is an estimate of elliptical kurtosis (*cf.* subsection IV-A). In this case, $\hat{\gamma}^{\text{Ell2}^*}$ corresponds to Ell2-estimator of sphericity proposed in [2, Sect. IV.B].

In order to guarantee that the estimators remain in the valid interval, $1 \leq \gamma \leq p$, we use

$$\hat{\gamma}^{\text{Ell1}} = \min(p, \max(1, \hat{\gamma}^{\text{Ell1}^*})) \quad (29)$$

as our final estimator (and similarly for Ell2-estimator). The related shrinkage parameter can then computed as

$$\beta = \beta_o^{\text{app}}(\hat{\gamma}^{\text{Ell1}}, \hat{\psi}_1),$$

and again similarly for Ell2-estimator.

IV. IMPORTANT SPECIAL CASES

A. Regularized SCM (RSCM) estimator

If one uses Gaussian weight function $u_G(t) \equiv 1$, then a necessary assumption is that the underlying elliptical distribution possesses finite 4th-order moments. As discussed earlier, one may then assume w.l.o.g. that the scatter matrix parameter equals the covariance matrix, i.e., $\mathbf{\Sigma} = \text{cov}(\mathbf{x})$. When $u_G(t) \equiv 1$, one has that $\hat{\mathbf{\Sigma}} = \mathbf{S}$ and $\mathbf{C}_\beta = \mathbf{S}_\beta$ and hence $\beta_o = \beta_o^{\text{app}}$, meaning that the approximate MMSE solution is exact in this case. Finally, we may relate ψ_1 with an elliptical kurtosis [36] parameter κ :

$$\psi_1 = 1 + \kappa = \begin{cases} \frac{\mathbb{E}[r^4]}{p(p+2)}, & \text{real case} \\ \frac{\mathbb{E}[r^4]}{p(p+1)}, & \text{complex case} \end{cases} \quad (30)$$

where the expectation is over the distribution of the random variable $r = \|\mathbf{\Sigma}^{-1/2}\mathbf{x}\|$. The elliptical kurtosis parameter is defined as a generalization of the kurtosis parameter to the vector case, and as such it vanishes (so $\kappa = 0$) when \mathbf{x} has MVN distribution (denoted $\mathbf{x} \sim \mathcal{N}_p(\mathbf{0}, \mathbf{\Sigma})$). Since \mathbf{S} exists for any $n \geq 1$, we can drop the assumption that $n > p$ in this case.

Corollary 1. *Let $\mathbf{x}_1, \dots, \mathbf{x}_n$ denote an i.i.d. random sample from an (real or complex) elliptical distribution $\mathcal{E}_p(\mathbf{0}, \mathbf{\Sigma}, g)$ with finite 4th order moments and covariance matrix $\mathbf{\Sigma} =$*

cov(\mathbf{x}). Then for the shrinkage SCM estimator \mathbf{S}_β in (1) one has that

$$\beta_o = \arg \min_{\beta} \mathbb{E}[\|\mathbf{S}_\beta - \mathbf{\Sigma}\|_{\text{F}}^2] = \frac{n(\gamma - 1)}{n(\gamma - 1) + p + a}, \quad (31)$$

where

$$a = \begin{cases} \kappa(2\gamma(1 - 1/p) + p - 1) + \gamma(1 - 2/p), & \text{real case} \\ \kappa(\gamma(1 - 1/p) + p - 1) - \gamma/p, & \text{complex case} \end{cases}$$

Proof. The result follows from Theorem 2 and Theorem 4 since $\mathbf{C}_\beta = \mathbf{S}_\beta$ and the M-functional for Gaussian loss is $\mathbf{\Sigma}_0 = \text{cov}(\mathbf{x}) = \mathbf{\Sigma}$ and $\sigma = 1$. Since for Gaussian loss, $\psi(t) = t$, we notice from (19) and (25) hat

$$\psi_1 = 1 + \kappa. \quad (32)$$

Plugging $\psi_1 = 1 + \kappa$ into β_o^{app} in Theorem 2 and Theorem 4 yields the stated expressions, respectively. \square

The elliptical kurtosis parameter κ can be easily estimated using the following relationship to kurtosis even in the cases when $p > n$. First, recall that kurtosis of a random variable x in the real and complex case is defined as

$$\text{kurt}(x) = \frac{\mathbb{E}[x^4]}{(\mathbb{E}[x^2])^2} - 3 \quad \text{and} \quad \text{kurt}(x) = \frac{\mathbb{E}[|x|^4]}{(\mathbb{E}[|x|^2])^2} - 2, \quad (33)$$

respectively. Kurtosis vanishes when the random variable has real or complex Gaussian distribution with variance $\mathbb{E}[|x|^2]$. The following result establishes the relationship of elliptical kurtosis parameter with marginal kurtosis.

Lemma 3. *Assume that \mathbf{x} is a random vector from real or complex elliptically symmetric distribution with covariance matrix $\mathbf{\Sigma} = \text{cov}(\mathbf{x})$ possessing finite 4th order moments. Then*

$$\kappa = \begin{cases} \frac{1}{3} \text{kurt}(x_j), & \text{real case} \\ \frac{1}{2} \text{kurt}(x_j), & \text{complex case} \end{cases} \quad (34)$$

where x_j is any j th component of \mathbf{x} ($j \in \{1, \dots, p\}$).

Proof. The proof is given in Appendix D. \square

Since all marginal variables possess the same kurtosis, an estimate $\hat{\kappa}$ can be formed simply as the mean of marginal sample kurtosis statistics. This is the same estimate of the elliptical kurtosis proposed in [2]. Note that [2] only considered the real-valued case, and thus Corollary 1 allows us to extend the RSCM estimator in [2] to complex-valued case.

In the sequel, we use acronym **RSCM-Ell1** to refer to estimator \mathbf{S}_β with β computed as $\beta = \beta_o(\hat{\kappa}, \hat{\gamma}^{\text{Ell1}})$ with β_o given by (31) and $\hat{\gamma}^{\text{Ell1}}$ being the estimate of sphericity defined in (29) and $\hat{\kappa}$ an estimate of elliptical kurtosis described above. An **RSCM-Ell2** estimator is defined similarly but now using Ell2-estimator of sphericity.

A natural competitor for RSCM-Ell1 or RSCM-Ell2 estimators (at least in the real-valued case) is the estimator proposed by Ledoit and Wolf [1], referred to as **LWE**. We note that LWE also uses RSCM \mathbf{S}_β , but the parameter β is computed in a different manner. An extra benefit of our approach is that an estimator of the optimal shrinkage parameter can be computed for real- or complex-valued observations while LWE assumes real-valued observations.

B. Regularized Huber's M-estimator (RHub)

Next consider the Huber's weight function $u_H(t; c)$ in (4). Note that $b > 0$ is a scaling constant; if $\hat{\Sigma}$ is Huber M-estimator of scatter when $b = 1$, then the Huber M-estimator of scatter when $b = b_o$ is simply $b_o \hat{\Sigma}$. The scaling constant b is usually chosen so that the resulting scatter estimator is Fisher consistent for the covariance matrix at MVN distribution, i.e., $\sigma = 1$ when $\mathbf{x} \sim \mathcal{N}_p(\mathbf{0}, \Sigma)$. In the real case, this holds when

$$b = F_{\chi_{p+2}^2}(c^2) + c^2(1 - F_{\chi_p^2}(c^2))/p,$$

where $F_{\chi_p^2}(\cdot)$ denotes the cumulative distribution function (c.d.f.) of chi-squared distribution with p d.o.f. Since $r^2 = \|\Sigma^{-1/2}\mathbf{x}\|^2$ has a χ_p^2 -distribution when $\mathbf{x} \sim \mathcal{N}_p(\mathbf{0}, \Sigma)$, the tuning constant c^2 is chosen as q th upper quantile of χ_p^2 -distribution:

$$q = \Pr(r^2 \leq c^2) \Leftrightarrow F_{\chi_p^2}^{-1}(q) = c^2 \quad (35)$$

for some $q \in (0, 1]$. Tuning constant c and scaling factor b can be determined similarly in the complex-valued case; see [33], [35] for details.

Let us define a winsorized observation \mathbf{w} based on $\mathbf{x} \sim \mathcal{E}_p(\mathbf{0}, \Sigma, g)$ as

$$\mathbf{w} = \text{wins}(\mathbf{x}; \Sigma, c) = \frac{1}{\sqrt{b}} \begin{cases} \mathbf{x}, & \|\Sigma^{-1/2}\mathbf{x}\|^2 \leq c^2 \\ c \frac{\mathbf{x}}{\|\Sigma^{-1/2}\mathbf{x}\|}, & \|\Sigma^{-1/2}\mathbf{x}\|^2 > c^2 \end{cases}$$

where c is the threshold c of Huber's weight function and b is the respective scaling factor. The winsorized r.v. \mathbf{w} also has an elliptically symmetric distribution since the contours remain elliptical in shape (so the p.d.f. is still defined by (16) but for a truncated density generator g) and thus it shares the properties of elliptical random vectors.

If we take $\sigma = 1$ (which holds at least when \mathbf{x} has MVN distribution), then the constant ψ_1 can be written as

$$\begin{aligned} \psi_1 &= \frac{\mathbb{E}[\psi_H^2(\|\Sigma^{-1/2}\mathbf{x}\|^2; c)]}{p(p+2)} = \frac{\mathbb{E}[\|\Sigma^{-1/2}\mathbf{w}\|^4]}{p(p+2)} \\ &= 1 + \kappa_{\mathbf{w}} \end{aligned}$$

where $\kappa_{\mathbf{w}}$ is the elliptical kurtosis parameter (cf. Lemma 3) of an elliptical random vector \mathbf{w} . An estimate $\hat{\psi}_1$ of ψ_1 can be then calculated similarly as ψ_1 for RSCM-Ell1 or RSCM-Ell2 estimators defined earlier (recall relation (32)). The only difference is that κ is now computed for winsorized data $\{\mathbf{w}_i\}_{i=1}^n$, where $\mathbf{w}_i = \text{wins}(\mathbf{x}_i; \hat{\Sigma}, c)$ and $\hat{\Sigma}$ denotes the Huber's M-estimator.

In the sequel, we use acronym **RHub-Ell1** or **RHub-Ell2** to refer to shrinkage M-estimator $\hat{\Sigma}_\beta$ that uses Huber's weight $u(\cdot) = u_H(\cdot; c)$ with threshold c^2 determined from (35) for user specified q and shrinkage parameter $\beta = \beta_o^{\text{app}}(\hat{\gamma}^{\text{Ell1}}, \hat{\psi}_1)$. or $\beta = \beta_o^{\text{app}}(\hat{\gamma}^{\text{Ell2}}, \hat{\psi}_1)$, respectively.

C. Regularized Tyler's M-estimator (RTyl)

Let \mathbf{V} denote a *shape matrix* (normalized scatter matrix), defined as

$$\mathbf{V} = p\Sigma / \text{tr}(\Sigma),$$

where Σ denotes the scatter matrix parameter of the ES distribution. Note that $\text{tr}(\mathbf{V}) = p$. If one uses Tyler's weight function in (6), then (17) holds with $\sigma = p/\text{tr}(\Sigma)$, i.e., $\Sigma_0 = \mathbf{V}$, that is, Tyler's M-estimator is an estimate of the shape matrix. The following result hence follows at once from Theorem 2 and Theorem 4.

Corollary 2. *Let $\mathbf{x}_1, \dots, \mathbf{x}_n$ denote an i.i.d. random sample from a (real or complex) elliptical distribution $\mathcal{E}_p(\mathbf{0}, \Sigma, g)$. When using Tyler's weight (6), it holds that*

$$\begin{aligned} \beta_o^{\text{app}} &= \arg \min_{\beta} \mathbb{E}[\|\mathbf{C}_\beta - \mathbf{V}\|_{\text{F}}^2] \\ &= \begin{cases} \frac{n(\gamma-1)}{n(\gamma-1)(1-\frac{1}{n}) + \frac{p-1}{p+2}(2\gamma+p)}, & \text{real case} \\ \frac{n(\gamma-1)}{n(\gamma-1)(1-\frac{1}{n}) + \frac{p-1}{p+1}(\gamma+p)}, & \text{complex case} \end{cases} \end{aligned}$$

Proof. The real-valued case follows by noting that ψ_1 in (19) is equal to $\psi_1 = p/(p+2)$ for all random vectors $\mathbf{x} \sim \mathcal{E}_p(\mathbf{0}, \Sigma, g)$ regardless of g (i.e., the functional form of the density generator) and that $\Sigma_0 = \mathbf{V}$. Plugging $\psi_1 = p/(p+2)$ into (20) yields the stated expression. The complex-valued case follows similarly. \square

Since Tyler's M-estimator verifies $\text{tr}(\hat{\Sigma}) = p$, the shrinkage estimator in (3) simplifies to

$$\hat{\Sigma}_\beta = \beta \hat{\Sigma} + (1 - \beta)\mathbf{I}, \quad (36)$$

where $\hat{\Sigma}$ is Tyler's M-estimator, i.e., an M-estimator based on weight (6). By **RTyl-Ell1** we refer to (36), where the shrinkage parameter is computed as $\beta = \beta_o^{\text{app}}(\hat{\gamma}^{\text{Ell1}})$ with β_o^{app} given by Corollary 2 and $\hat{\gamma}^{\text{Ell1}}$ by (29).

A related regularized Tyler's estimator was proposed by [7] as the limit of the algorithm

$$\begin{aligned} \Sigma_{k+1} &\leftarrow \beta \frac{p}{n} \sum_{i=1}^n \frac{\mathbf{x}_i \mathbf{x}_i^{\text{T}}}{\mathbf{x}_i^{\text{T}} \mathbf{V}_k^{-1} \mathbf{x}_i} + (1 - \beta)\mathbf{I} \\ \mathbf{V}_{k+1} &\leftarrow p \Sigma_{k+1} / \text{tr}(\Sigma_{k+1}), \end{aligned} \quad (37)$$

where $\beta \in (0, 1)$ is a fixed shrinkage parameter. This algorithm represents a diagonally loaded version of the fixed-point algorithm given for Tyler's M-estimator. Uniqueness and convergence of the recursive algorithm has been later derived in [29], [12]. By **CWH** estimator we now refer to estimator obtained by iterating (37) using same value $\beta = \beta_o^{\text{app}}(\hat{\gamma}^{\text{Ell1}})$ as for RTyl-Ell1. An interesting question then is how different is RTyl-Ell1 in its performance from CWH. We explore this by simulation studies later. This is interesting as the former is simply shrinking the eigenvalues of Tyler's M-estimator towards its grand mean where as the latter does not have an explicit connection to Tyler's M-estimator $\hat{\Sigma}$ for any $\beta \in (0, 1)$.

D. Regularized M-estimator for MVT distribution (RMVT)

We assume that the data is arising from a MVT distribution $t_\nu(\mathbf{0}, \Sigma)$ but the d.o.f. parameter ν is unknown and is adaptively estimated from the data using Algorithm 1 explained below. Once $\hat{\nu}$ is found, we use function $u(\cdot) = u_\tau(\cdot; \hat{\nu})$

to compute the underlying M-estimator $\hat{\Sigma}$ for the postulated MVT distribution.

Yet, we need to address the question of how the constant ψ_1 is computed. Due to data adaptive estimation of ν , we can assume that $\sigma \approx 1$ since the scaling factor σ equals unity for an MLE of the scatter matrix parameter. We use the fact that for MVT distribution (i.e., when $\mathbf{x} \sim t_\nu(\mathbf{0}, \Sigma)$), the ψ_1 parameter is¹

$$\psi_1 = \begin{cases} \frac{p + \nu}{2 + p + \nu}, & \text{real case} \\ \frac{2p + \nu}{2 + 2p + \nu}, & \text{complex case} \end{cases}.$$

Hence, a natural estimate is $\hat{\psi}_1 = (p + \hat{\nu}) / (2 + p + \hat{\nu})$ in the real case. An estimate of $\hat{\psi}_1$ is constructed similarly in the complex case. We use acronym **RMVT-Ell1** to refer to shrinkage M-estimator $\hat{\Sigma}_\beta$ that uses $u(\cdot; \hat{\nu})$ with shrinkage parameter calculated by $\beta = \beta_{\text{opt}}^{\text{app}}(\hat{\gamma}^{\text{Ell1}}, \hat{\psi}_1)$. **RMVT-Ell2** is constructed similarly, but now Ell2-estimator of sphericity $\hat{\gamma}^{\text{Ell2}}$ is used.

Next we discuss our approach for estimating ν from the data. Assume $\mathbf{x} \sim t_\nu(\mathbf{0}, \Sigma, g)$ and denote $\eta = \text{tr}(\Sigma)/p$. Then,

$$\mathbf{R} = \text{cov}(\mathbf{x}) = (\nu/(\nu - 2))\Sigma$$

and hence $\text{tr}(\mathbf{R})/p = (\nu/(\nu - 2))\text{tr}(\Sigma)/p$. This means that

$$\frac{\nu}{\nu - 2} = \frac{\text{tr}(\mathbf{R})}{\text{tr}(\Sigma)} = \eta_{\text{ratio}}$$

from which we obtain the relation

$$\nu = \frac{2\eta_{\text{ratio}}}{\eta_{\text{ratio}} - 1}. \quad (38)$$

The above relation holds true in both real and complex cases. Then given an estimate $\hat{\Sigma}$ of Σ , we may compute an estimate $\hat{\eta}_{\text{ratio}} = \text{tr}(\mathbf{S})/\text{tr}(\hat{\Sigma})$ which in turn provides an estimate $\hat{\nu}$ via (38). This gives rise to an iterative algorithm to estimate ν detailed in Algorithm 1. In the simulations, the algorithm converged, but already 2 iterations are sufficient to yield accuracy to first decimal; see Figure 1 for an illustration. The initial estimate is $\nu_o = 2/(\max(0, \hat{\kappa}) + \delta) + 4$, where $\hat{\kappa}$ is an estimate of marginal kurtosis explained in subsection IV-A (see also [2] for more details) and $\delta > 0$ is a small number. The initial start ν_o is based on the following relationship with elliptical kurtosis parameter, $\kappa = 2/(\nu - 4)$, i.e., $\nu = 2/\kappa + 4$ which holds true both in real and complex cases. Again the estimate $\hat{\nu}$ in the complex-valued case is constructed similarly. Note that also other estimators of ν are proposed in the literature, for example in [34].

V. SIMULATION STUDIES

In the simulation study, we generate samples from real ES distributions with a scatter matrix Σ following an AR(1) structure, $(\Sigma)_{ij} = \tau \rho^{|i-j|}$, where $\rho \in (0, 1)$ and scale parameter $\tau = \text{tr}(\Sigma)/p = 10$. When $\rho \downarrow 0$, then Σ is close to an identity matrix scaled by τ , and when $\rho \uparrow 1$, Σ tends to a singular matrix of rank 1. The results are reported for the




Fig. 1. Average $\hat{\nu}$ by running the Algorithm 1 with different choices of T_{max} . Also shown is the initial estimate ν_0 . The samples are generated from a p -variate t -distribution with $\nu = 5$ d.o.f., where Σ follows the same AR(1) covariance matrix structure explained in the simulation set-up of Section V; $\rho = 0.6$ and $p = 40$. As can be noted, $\hat{\nu}$ converge to $\nu = 5$ as n increases, albeit the convergence is a bit slow.

Algorithm 1: Automatic data-adaptive computation of the d.o.f. parameter ν

Input : data matrix \mathbf{X} of size $n \times p$, maximum number of iterations T_{max}

Initialize: Compute $\nu_0 = 2/\max(0, \hat{\kappa}) + 4$, where $\hat{\kappa}$ is an estimate of κ explained in the text.

for $t = 0, 1, \dots, T_{max}$ **do**

Set $\hat{\Sigma}_t = \hat{\Sigma}$, where $\hat{\Sigma}$ denotes the t -MLE based on current estimate of d.o.f. parameter $\nu = \nu_t$, i.e., solving the M-estimating equation

$$\hat{\Sigma} = \frac{1}{n} \sum_{i=1}^n u(\mathbf{x}_i^\top \hat{\Sigma}^{-1} \mathbf{x}_i) \mathbf{x}_i \mathbf{x}_i^\top$$

with $u(\cdot) = u_t(\cdot; \nu_t)$ is the MVT-weight function.

Update the ratio $\hat{\eta}_t = \frac{\text{tr}(\mathbf{S})}{\text{tr}(\hat{\Sigma}_t)}$.

Update the d.o.f. parameter $\nu_{t+1} = \frac{2\hat{\eta}_t}{\hat{\eta}_t - 1}$.

if $|\nu_{t+1} - \nu_t|/\nu_t < 0.01$ **then**
| break

Output : $\hat{\nu} = \nu_{t+1}$

proposed shrinkage M-estimators using shrinkage parameter estimates based on Ell1-estimator of sphericity. However, for notational convenience, we drop the suffix -Ell1 from the proposed estimators. Thus the proposed estimators, described in Section IV are referred to as RSCM, RMVT, RHub, and RTyl. Furthermore, acronyms LW, CWH and RBLW are used to refer to estimators proposed in [1] (see also subsection IV-A), [29] (see also subsection IV-C and (37)) and [8], respectively. **RBLW** is the Rao-Blackwellized LW estimator, but unlike LW estimator, it assumes that the data distribution is Gaussian.

We also compare to RSCM estimator \mathbf{S}_β in (1), but now the shrinkage parameter β is chosen via k -fold cross-validation (CV), where as cross-validation fit we use $\|\mathbf{S}_{\beta, \text{tr}} - \mathbf{S}_{\text{val}}\|_F$, where \mathbf{S}_{val} is the SCM based on the validation set (data fold

¹Note that the t -weight in the complex case is [33]: $u_T(t; \nu) = \frac{2p+\nu}{\nu+2t}$.




Fig. 2. NMSE as a function of n when samples are drawn from a MVN distribution $\mathcal{N}_p(\mathbf{0}, \Sigma)$ with an AR(1) structure; $\rho = 0.6$ and $p = 40$.




Fig. 3. Shrinkage parameter β as a function of n when samples are drawn from a MVN distribution (left panel) and a t -distribution with $\nu = 5$ d.o.f. (right panel), where Σ has an AR(1) structure; $\rho = 0.6$ and $p = 40$.

that was left out) and $\mathbf{S}_{\beta, \text{tr}}$ is the RSCM computed on the training set (data based on remaining folds) for a given β . As a grid for β for the CV method we use a uniform grid in $[0, 1]$ with 0.05 increments and 5-fold cross-validation. We call this method as **RSCM-CV** or simply **CV**. All simulation results in this section are averaged over 2000 Monte-Carlo trials. Since $n > p$ is assumed for all estimators except for RSCM, we do not consider the case low sample regime, $n \leq p$, in our simulation studies. Furthermore, we adopt the MSE (squared Frobenius norm) as our performance metric as it is used in deriving the optimal shrinkage parameters in this paper. It is important to keep in mind, however, that in the low sample regime and for different performance metrics, the performance differences between estimators can often be noticeable and even quite different than in the $n > p$ regime that is considered here; See e.g., [8], [4], [2] for numerical illustrations and [37], [38] for different distances between covariances that could be used instead of the MSE metric.

A. Gaussian data

The data is generated from MVN distribution $\mathcal{N}_p(\mathbf{0}, \Sigma)$, where Σ has an AR(1) covariance matrix structure with $\rho = 0.6$. The dimension is $p = 40$ and n varies from 60 to 280. Value $q = 0.7$ determining the threshold c is used in (35) for Huber's weight. Since Huber's M-estimator is scaled to be

consistent to the covariance matrix for Gaussian samples, the underlying population parameter Σ_0 coincides with the covariance matrix Σ in this case. We also scaled the MVT-weight $u_T(t; \nu)$ so that it is consistent to Σ for Gaussian data. Figure 2 compares the normalized MSE (NMSE) $\|\hat{\Sigma}_\beta - \Sigma\|_F^2 / \|\Sigma\|_F^2$ of different estimators w.r.t. increasing sample length n . It can be noted that all estimators provide essentially equally good estimator of the covariance matrix Σ for Gaussian data; RSCM and RMVT are performing equally well, largely due to the effect of data-adaptive estimation of d.o.f. parameter ν . It should be noted that their performance difference to LW or RHub estimators are still marginal and differences can be spotted only by zooming in as in the sub-plot of Figure 2. As expected, RBLW estimator has a slight advantage over the other estimators in this case. The left panel of Figure 3 shows the (average) shrinkage parameter β as a function of n . As can be noted, the average shrinkage parameter of the proposed RSCM estimator can be seen to be roughly an average of CV and LW shrinkage parameters.

B. Heavy-tailed data

Next we computed the NMSE curves when the data is generated from a heavy-tailed t -distribution with $\nu = 5$ and $\nu = 3$ d.o.f. Note that NMSE of each estimator is now compared against the underlying population parameter Σ_0 of each M-estimator. Figure 4 displays the results. RBLW had a very poor performance which is due to its strict assumption of Gaussianity. It can be noted that CV method performs similarly, but slightly worse, than RSCM or LW. This can be partially attested to poor robustness properties of cross-validation. In the case of $\nu = 3$ d.o.f., also the non-robust RSCM and LW provided large NMSE and thus all non-robust estimators are not visible in the right panel of Figure 4. This was expected since t -distribution with $\nu = 3$ d.o.f. is very heavy-tailed with non-finite kurtosis. As can be noted, the proposed robust RHub and RMVT estimators provide significantly improved performance. We can also notice that RMVT estimator that adaptively estimates the d.o.f. ν from the data is able to outperform the regularized Huber's estimator (RHub).

The right panel of Figure 3 depicts the (average) shrinkage parameter β as a function of n in the case that samples are drawn from a t -distribution with $\nu = 5$ d.o.f. As can be noted the robust shrinkage estimators (RHub and RMVT) use larger shrinkage parameter value β than the non-robust RSCM and LW estimators. Compared to RSCM and LW, the RBLW (resp. CV) is seen to overestimate (resp. underestimate) the shrinkage parameter as it obtains much larger (resp. smaller) values.

Next we investigate how the estimators are able to estimate the shape matrix, i.e., the covariance matrix up to a scale. Figure 5 displays the NMSE, $\|\hat{\mathbf{V}} - \mathbf{V}\|_F^2 / \|\mathbf{V}\|_F^2$, of different shrinkage shape matrix estimators, defined as $\hat{\mathbf{V}} = p\hat{\Sigma}_\beta / \text{tr}(\hat{\Sigma}_\beta)$ when samples are generated from a t -distribution with $\nu = 5$ d.o.f. Note that such normalization is not necessary for CWH or RTy1 since they verify $\text{tr}(\hat{\Sigma}) = p$ in the first place. Figure 5 illustrate both the case when correlation parameter ρ of the AR(1) scatter matrix parameter




Fig. 4. NMSE as a function of n when samples are drawn from a p -variate t -distribution with $\nu = 5$ (left panel) and $\nu = 3$ (right panel) d.o.f. The scatter matrix follows an AR(1) structure; $\varrho = 0.6$ and $p = 40$.

Σ is fixed while n varies and the case that n is fixed while ϱ varies. As can be seen from the top panel of Figure 5, all robust shape estimators are performing well and very similarly. In fact, performance of RMVT and CWH is essentially the same. We can also observe that the two different approaches for shrinking Tyler's M-estimator, so CWH and the proposed RTyl are very similar. We can note from the bottom panels of Figure 5 that when $\varrho \approx 0$ (so Σ is close to a scaled identity matrix) all estimators perform similarly. This is because all estimators are shrunk heavily towards the scaled identity matrix (namely, $\beta \approx 0$ for all estimators). Similarly, when $\varrho \approx 1$ (so Σ is close to a singular matrix of rank 1), all estimators have a rather similar performance. This is because the true scatter matrix Σ is poorly conditioned ($\text{cond}(\Sigma) \approx 7000$) and all estimators share similar difficulties of capturing the subspace structure due to limited training data and no *a priori* information about such structure. Indeed biggest differences between estimators are observed when Σ has no particular structure, i.e., ϱ in the range $[0.4, 0.7]$.

C. Complex-valued data

Finally, we note that an important property of our shrinkage method is that it can be used for complex-valued data as well. Some other methods in the previous study, such as RBLW or LW assume real-valued data. In the supplementary material, we provide the results of a simulation study in the same set-up, but now the data being generated from circular complex Gaussian and heavy-tailed t -distribution, respectively. These are distributions in the class of CES distributions. In our study we also include empirical Bayes diagonal loading estimator (EBDL) [4] which was developed for complex circular Gaussian data. Results obtained for complex-valued data attest the validity of the findings in the real-valued case.

VI. APPLICATION TO FINANCIAL DATA AND PORTFOLIO DESIGN

A. Financial data

We use the S&P 500 stock market index (see Figure 6), which measures the stock performance of 500 large companies listed on stock exchanges in the United States, and its





Fig. 5. NMSE of different shrinkage estimators of shape matrix \mathbf{V} when samples are drawn a $p = 40$ variate t -distribution with $\nu = 5$ d.o.f. having an AR(1) structure: (a) $\varrho = 0.6$ and sample length n varies; (b) and (c) illustrate the case when ϱ varies while the sample length n is fixed.



Fig. 6. Log-prices of the S&P 500 index.

constituent stocks during the period 2016-01-01 to 2020-01-31.

We can easily observe from the returns shown in Figure 7 the effect of volatility clustering over time that is responsible for heavy tails. More concretely, if we assume that the data follows an MVT distribution, then we can compute the degrees of freedom ν on a rolling-window basis and verify that indeed the data has heavy tails with $\nu \approx 5$ (from mid-2017, ν varies between 4.5 and 6) as shown in Figure 8.

Factor model. Let \mathbf{x}_i denote the returns of the p stocks at time i . It turns out that the stock returns are largely driven by

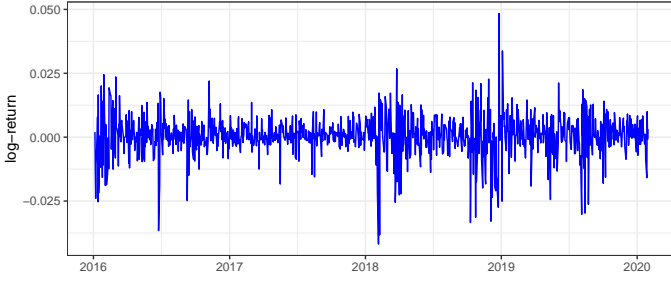



Fig. 7. Log-returns of the S&P 500 index.


 Fig. 8. Degrees of freedom ν of the S&P 500 index.

very few $k \ll p$ financial factors \mathbf{f}_i as

$$\mathbf{x}_i = \mathbf{B}\mathbf{f}_i + \boldsymbol{\epsilon}_i, \quad (39)$$

where $\mathbf{B} \in \mathbb{R}^{p \times k}$ is the factor loading matrix (which is very tall since $k \ll p$) and $\boldsymbol{\epsilon}_i$ is the residual idiosyncratic component. As a consequence, the covariance matrix of these data has the form (assuming normalized factors):

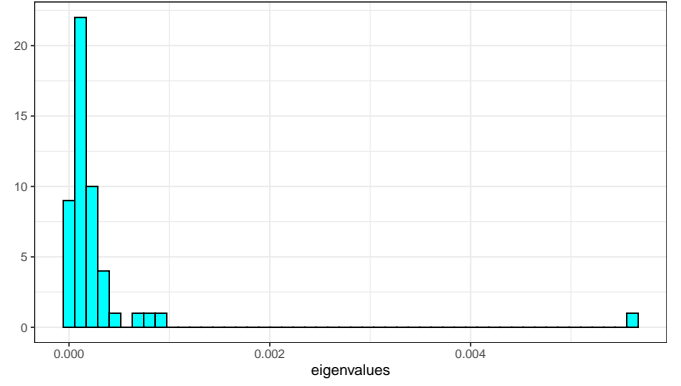

$$\boldsymbol{\Sigma} = \mathbf{B}\mathbf{B}^\top + \boldsymbol{\Psi} \quad (40)$$

where $\boldsymbol{\Psi}$ is the (diagonal) covariance matrix of the residuals. Typically, the term $\mathbf{B}\mathbf{B}^\top$ is much stronger than $\boldsymbol{\Psi}$ and this leads to the commonly used spike model in RMT (Random Matrix Theory) [39] which contains a few large eigenvalues and the rest small eigenvalues form the so-called bulk. Figure 9 shows the histogram of the empirical eigenvalues of the covariance matrix estimated from $p = 50$ stocks from the S&P 500 market data, where a very strong eigenvalue can be observed corresponding to the market factor.

B. Results in terms of MSE

Since our shrinkage estimators are derived to minimize the MSE in the estimation of the covariance matrix, we start by showing the obtained MSE in the context of financial data. We consider seven methods in our comparison:

- LWE is the Ledoit-Wolf estimator [1];
- RMVT-EI11 described in subsection IV-D;
- MVT: equals RMVT-EI11 with no shrinkage ($\beta = 1$);
- RSCM-EI11 estimator described in subsection IV-A;
- SCM: equals RSCM-EI11 with no shrinkage ($\beta = 1$);
- RHub-EI11 described in subsection IV-B; and
- RSCM-EI12, RMVT-EI12, and RHub-EI12 are as RSCM-EI11, RMVT-EI11, and RHub-EI11, respectively, but using $\hat{\gamma}^{\text{EI12}}$ estimator of sphericity γ (cf. subsection III-B).


 Fig. 9. Histogram of empirical eigenvalues obtained from the market data ($p = 50$) showing a strong market factor.

 Fig. 10. Normalized MSE of covariance matrix vs number of observations for Gaussian data (with $p = 50$).

To make sure that the robust estimators do not underperform the benchmarks when the data is not heavy-tailed, we start by generating synthetic MVN data following the empirical covariance matrix previously obtained from market data (see Figure 9). Figure 10 displays the normalized MSE of the covariance matrix $\mathbb{E}[\|\hat{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}\|_{\text{F}}^2]$ via 200 Monte-Carlo simulations. We do not observe any significant difference among the methods. Figure 11 shows the normalized MSE of the precision matrix $\mathbb{E}[\|\hat{\boldsymbol{\Sigma}}^{-1} - \boldsymbol{\Sigma}^{-1}\|_{\text{F}}^2]$ via 200 Monte-Carlo simulations. The main observation is that the two methods without shrinkage significantly underperform.

We now generate heavy-tailed synthetic data from a t -distribution following the empirical covariance matrix previously obtained from market data (see Figure 9) and with d.o.f. $\nu = 5$. Figure 12 shows the normalized MSE of the covariance matrix. We can clearly observe a striking difference showing the superiority of robust estimators (i.e., MVT, RMVT, RHub) over non-robust estimators. Among the robust estimators, the proposed shrinkage methods clearly outperform the non-shrinkage MVT in the low-sample regime, as expected. Figure 13 displays the normalized MSE of the precision matrix, with similar observations.




Fig. 11. Normalized MSE of precision matrix vs number of observations for Gaussian data (with $p = 50$).




Fig. 12. Normalized MSE of covariance matrix vs number of observations for t -distributed data (with $p = 50, \nu = 5$).




Fig. 13. Normalized MSE of precision matrix vs number of observations for t -distributed data (with $p = 50, \nu = 5$).

C. Results in terms of portfolio Sharpe ratio

After having shown the superiority of our proposed estimators, RMVT and RHub in terms of MSE in the estimation of the covariance matrix under heavy-tailed data, we now turn to assess the effects in terms of portfolio design. Note that an improvement of MSE in the covariance matrix may or may not translate into a significant improvement in terms of portfolio design; this depends on exactly what portfolio design is used and how it employs the estimated covariance matrix.




Fig. 14. Boxplot of Sharpe ratio obtained by the mean-variance portfolio according to different estimators for the covariance matrix.

For simplicity, we consider the most basic Markowitz portfolio design [40]:

$$\begin{aligned} & \underset{\mathbf{w}}{\text{minimize}} && \mathbf{w}^\top \Sigma \mathbf{w} \\ & \text{subject to} && \mathbf{w}^\top \boldsymbol{\mu} \geq \alpha \\ & && \mathbf{1}^\top \mathbf{w} = 1, \end{aligned}$$

where $\boldsymbol{\mu}$ is the expected return of the returns and α is the minimum return desired for the portfolio.

We perform our backtest during the market period 2016-01-01 to 2020-01-31 on a rolling-window basis with a window length of 378 days (1.5 years). To make sure that our results are realistic, rather than performing a single backtest, we use the R package portfolioBacktest [41] to randomly select a large number of 200 datasets from the market data in the following way: each dataset chooses randomly $p = 200$ stocks from the universe of 500, as well as a random period of 504 days (2 years) among the available period from 2016-01-01 to 2020-01-31.

Figure 14 shows a boxplot with the Sharpe ratio² obtained mean-variance portfolio according to different estimators for the covariance matrix (along with two benchmarks: the index and the $1/N$ or uniform portfolio). We can observe that the two methods without shrinkage underperform the shrinkage methods (in particular the SCM). Among the shrinkage methods, we can see that our robust estimators slightly outperform the others, although the improvement is not extremely significant.

D. Supplementary studies

Supplementary material to this paper also contain comparison of the proposed methods in the set-up described in [2], where the global mean variance portfolio (GMVP) is used as portfolio optimization strategy and the net returns correspond to $p = 50$ stocks that are currently included in the Hang Seng Index (HSI). Compared methods include GMVP weight vector based on LW estimator and an estimator proposed in [42] that uses robust regularized Tyler’s M-estimator with a

²The Sharpe ratio is defined as the expected return normalized with the volatility or standard deviation: $SR = \frac{\mathbf{w}^\top \boldsymbol{\mu} - r_f}{\sqrt{\mathbf{w}^\top \Sigma \mathbf{w}}}$, where r_f is the return of the risk-free asset.

tuning parameter selection optimized for GMVP strategy. We observed that GMVP based on the proposed RHub and RSCM are the best performing methods in terms of realized risk. We note that method of [42] was excluded in the study in subsection VI-C due to its high computational cost in the high-dimensional setting.

VII. CONCLUSIONS AND PERSPECTIVES

This work proposed an original and fully automatic approach to compute an optimal shrinkage parameter in the context of heavy-tailed distributions and/or in presence of outliers. It has been shown that the performance of the method is similar to optimal one when the data is Gaussian while it outperforms shrinkage Gaussian-based methods when the data distribution turns out to be non-Gaussian. One of the benefits of the proposed adaptive shrinkage parameter selection is that it permits using real-valued or complex-valued data. Furthermore, a MATLAB toolbox called ShrinkM is freely available at <http://users.spa.aalto.fi/esollila/shrinkM/> that includes functions to compute all of the proposed estimators (RHub, RTyl, RSCM, RMVT, and CV) as well as a script of one the simulation studies presented in the paper to reproduce the results. Furthermore, this paper opens several ways, notably considering the challenging cases where $p > n$ which is left for future work. Supplementary materials also provide additional examples illustrating the benefits of the proposed estimators.

APPENDIX

A. Proof of Theorem 1

Write $L(\beta) = \text{MSE}(\mathbf{C}_\beta) = \mathbb{E}[\|\mathbf{C}_\beta - \boldsymbol{\Sigma}_0\|_{\text{F}}^2]$. Then note that

$$\begin{aligned} L(\beta) &= \mathbb{E}[\|\beta\mathbf{C} + (1-\beta)p^{-1}\text{tr}(\mathbf{C})\mathbf{I} - \boldsymbol{\Sigma}_0\|_{\text{F}}^2] \\ &= \mathbb{E}[\|\beta(\mathbf{C} - \boldsymbol{\Sigma}_0) + (1-\beta)(p^{-1}\text{tr}(\mathbf{C})\mathbf{I} - \boldsymbol{\Sigma}_0)\|_{\text{F}}^2] \\ &= \beta^2 a_1 + (1-\beta)^2 a_2 + 2\beta(1-\beta)a_3 \end{aligned} \quad (41)$$

where $a_1 = \mathbb{E}[\|\mathbf{C} - \boldsymbol{\Sigma}_0\|_{\text{F}}^2] = \mathbb{E}[\text{tr}(\mathbf{C}^2)] - \text{tr}(\boldsymbol{\Sigma}_0^2)$, and

$$\begin{aligned} a_2 &= \mathbb{E}[\|p^{-1}\text{tr}(\mathbf{C})\mathbf{I} - \boldsymbol{\Sigma}_0\|_{\text{F}}^2] \\ &= a_3 + \text{tr}(\boldsymbol{\Sigma}_0^2) - p\eta_o^2 = a_3 + p(\gamma - 1)\eta_o^2 \\ a_3 &= p^{-1}\mathbb{E}[\text{tr}(\mathbf{C})\text{tr}(\mathbf{C} - \boldsymbol{\Sigma}_0)] = p^{-1}\mathbb{E}[\text{tr}(\mathbf{C}^2)] - \eta_o^2 p \end{aligned}$$

and $\eta_o = \text{tr}(\boldsymbol{\Sigma}_0)/p$. Note that $L(\beta)$ is a convex quadratic function in β with a unique minimum given by

$$\beta_o^{\text{app}} = \frac{a_2 - a_3}{(a_1 - a_3) + (a_2 - a_3)}. \quad (42)$$

Substituting the expressions for constants a_1 , a_2 and a_3 into β_o^{app} yields the stated result.

The expression for MSE of $\mathbf{C}_{\beta_o^{\text{app}}}$ then follows by substituting β_o^{app} into expression for $L(\beta)$ in (41) and using the relation, $(1 - \beta_o^{\text{app}})(a_2 - a_3) = \beta_o^{\text{app}}(a_1 - a_3)$, which follows from (42). This yields the MSE expression

$$L(\beta_o^{\text{app}}) = a_2 - \beta_o^{\text{app}}(a_2 - a_3) = a_3 + (1 - \beta_o^{\text{app}})(a_2 - a_3).$$

This gives the stated MSE expression after noting that $a_2 - a_3 = \|\boldsymbol{\Sigma}_0 - \eta_o\mathbf{I}\|_{\text{F}}^2$.

B. Proof of Lemma 1

First recall that $\boldsymbol{\Sigma}_0 = \sigma\boldsymbol{\Sigma}$, and hence

$$\begin{aligned} \mathbf{C} &= \frac{1}{n} \sum_{i=1}^n u(\mathbf{x}_i^\top \boldsymbol{\Sigma}_0^{-1} \mathbf{x}_i) \mathbf{x}_i \mathbf{x}_i^\top \\ &= \boldsymbol{\Sigma}_0^{1/2} \left\{ \frac{1}{n} \sum_{i=1}^n u\left(\frac{r_i^2}{\sigma}\right) \frac{r_i^2}{\sigma} \mathbf{v}_i \mathbf{v}_i^\top \right\} \boldsymbol{\Sigma}_0^{1/2}, \end{aligned} \quad (43)$$

where $\mathbf{v}_i = \boldsymbol{\Sigma}^{-1/2} \mathbf{x}_i / \|\boldsymbol{\Sigma}^{-1/2} \mathbf{x}_i\|$ and $r_i^2 = \|\boldsymbol{\Sigma}^{-1/2} \mathbf{x}_i\|^2$. Recall that stochastic representation theorem of elliptical random vectors states that r_i is independent of \mathbf{v}_i and \mathbf{v}_i -s are i.i.d. on a uniform distribution on the unit sphere $\mathcal{S}^{p-1} = \{\mathbf{v} : \mathbf{v}^\top \mathbf{v} = 1\}$. Then note that

$$\begin{aligned} &\mathbb{E}[\text{tr}(\mathbf{C}^2)] \\ &= \frac{1}{n^2} \mathbb{E} \left[\text{tr} \left\{ \boldsymbol{\Sigma}_0^{1/2} \sum_{i=1}^n u\left(\frac{r_i^2}{\sigma}\right) \frac{r_i^2}{\sigma} \mathbf{v}_i \mathbf{v}_i^\top \boldsymbol{\Sigma}_0^{1/2} \right. \right. \\ &\quad \left. \left. \cdot \boldsymbol{\Sigma}_0^{1/2} \sum_{j=1}^n u\left(\frac{r_j^2}{\sigma}\right) \frac{r_j^2}{\sigma} \mathbf{v}_j \mathbf{v}_j^\top \boldsymbol{\Sigma}_0^{1/2} \right\} \right] \\ &= \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \mathbb{E} \left[u\left(\frac{r_i^2}{\sigma}\right) \frac{r_i^2}{\sigma} u\left(\frac{r_j^2}{\sigma}\right) \frac{r_j^2}{\sigma} \right] \\ &\quad \times \mathbb{E} \left[\text{tr} \left\{ \underbrace{\mathbf{v}_i \mathbf{v}_i^\top \boldsymbol{\Sigma}_0 \mathbf{v}_j \mathbf{v}_j^\top \boldsymbol{\Sigma}_0}_{=(\mathbf{v}_i^\top \boldsymbol{\Sigma}_0 \mathbf{v}_j)^2} \right\} \right] \\ &= \frac{1}{n^2} \sum_{i=1}^n \mathbb{E} \left[u\left(\frac{r_i^2}{\sigma}\right)^2 \frac{r_i^4}{\sigma^2} \right] \mathbb{E}[(\mathbf{v}_i^\top \boldsymbol{\Sigma}_0 \mathbf{v}_i)^2] \\ &\quad + \frac{1}{n^2} \sum_{i \neq j} \mathbb{E} \left[u\left(\frac{r_i^2}{\sigma}\right) \frac{r_i^2}{\sigma} \right] \mathbb{E} \left[u\left(\frac{r_j^2}{\sigma}\right) \frac{r_j^2}{\sigma} \right] \mathbb{E}[(\mathbf{v}_i^\top \boldsymbol{\Sigma}_0 \mathbf{v}_j)^2] \\ &= \frac{1}{n} \mathbb{E} \left[u\left(\frac{r_1^2}{\sigma}\right)^2 \frac{r_1^4}{\sigma^2} \right] \mathbb{E}[(\mathbf{v}_1^\top \boldsymbol{\Sigma}_0 \mathbf{v}_1)^2] \\ &\quad + \left(1 - \frac{1}{n}\right) \left(\mathbb{E} \left[u\left(\frac{r_1^2}{\sigma}\right) \frac{r_1^2}{\sigma} \right] \right)^2 \mathbb{E}[(\mathbf{v}_1^\top \boldsymbol{\Sigma}_0 \mathbf{v}_2)^2]. \end{aligned} \quad (44)$$

In the second identity, we used the fact that r_i is independent of \mathbf{v}_i and $\text{tr}(\mathbf{A}\mathbf{B}) = \text{tr}(\mathbf{B}\mathbf{A})$. In the 3rd identity we used that r_i is independent of r_j and in the 4th identity, we used that \mathbf{v}_i -s and r_i -s are i.i.d.

Note that $\mathbb{E}[u(r_1^2/\sigma)(r_1^2/\sigma)] = p$ due to (18) and $\mathbb{E}[u(r_1^2/\sigma)^2(r_1^4/\sigma^2)] = \psi_1 p(p+2)$. Using these facts and the following results from [8]:

$$\mathbb{E}[(\mathbf{v}_1^\top \boldsymbol{\Sigma}_0 \mathbf{v}_1)^2] = \frac{2\text{tr}(\boldsymbol{\Sigma}_0^2) + \text{tr}(\boldsymbol{\Sigma}_0)^2}{p(p+2)} \quad (45)$$

$$\mathbb{E}[(\mathbf{v}_1^\top \boldsymbol{\Sigma}_0 \mathbf{v}_2)^2] = \frac{\text{tr}(\boldsymbol{\Sigma}_0^2)}{p^2}, \quad (46)$$

we get

$$\begin{aligned} &\mathbb{E}[\text{tr}(\mathbf{C}^2)] \\ &= \frac{\psi_1}{n} (2\text{tr}(\boldsymbol{\Sigma}_0^2) + \text{tr}(\boldsymbol{\Sigma}_0)^2) + \left(1 - \frac{1}{n}\right) \text{tr}(\boldsymbol{\Sigma}_0^2) \\ &= \left(1 + \frac{2\psi_1 - 1}{n}\right) \text{tr}(\boldsymbol{\Sigma}_0^2) + \frac{\psi_1}{n} \text{tr}(\boldsymbol{\Sigma}_0)^2. \end{aligned}$$

Next note that

$$\begin{aligned} \text{tr}(\mathbf{C})^2 &= \left\{ \frac{1}{n} \sum_{i=1}^n u\left(\frac{r_i^2}{\sigma}\right) \frac{r_i^2}{\sigma} \mathbf{v}_i^\top \boldsymbol{\Sigma}_0 \mathbf{v}_i \right\}^2 \\ &= \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n u\left(\frac{r_i^2}{\sigma}\right) \frac{r_i^2}{\sigma} u\left(\frac{r_j^2}{\sigma}\right) \frac{r_j^2}{\sigma} \mathbf{v}_i^\top \boldsymbol{\Sigma}_0 \mathbf{v}_j \mathbf{v}_j^\top \boldsymbol{\Sigma}_0 \mathbf{v}_i \end{aligned}$$

and hence

$$\begin{aligned} \mathbb{E}[\text{tr}(\mathbf{C})^2] &= \frac{1}{n^2} \sum_{i=1}^n \mathbb{E} \left[u\left(\frac{r_i^2}{\sigma}\right)^2 \frac{r_i^4}{\sigma^2} \right] \mathbb{E}[(\mathbf{v}_i^\top \boldsymbol{\Sigma}_0 \mathbf{v}_i)^2] \\ &\quad + \frac{1}{n^2} \sum_{i \neq j} \mathbb{E} \left[u\left(\frac{r_i^2}{\sigma}\right) \frac{r_i^2}{\sigma} \right] \mathbb{E} \left[u\left(\frac{r_j^2}{\sigma}\right) \frac{r_j^2}{\sigma} \right] \frac{\text{tr}(\boldsymbol{\Sigma}_0)^2}{p^2} \\ &= \frac{\psi_1 p(p+2)}{n} \mathbb{E}[(\mathbf{v}_i^\top \boldsymbol{\Sigma}_0 \mathbf{v}_i)^2] + \left(1 - \frac{1}{n}\right) \text{tr}(\boldsymbol{\Sigma}_0)^2. \end{aligned}$$

In the first identity we used that

$$\mathbb{E}[\mathbf{v}_i^\top \boldsymbol{\Sigma}_0 \mathbf{v}_i] = \text{tr}\{\mathbb{E}[\mathbf{v}_i \mathbf{v}_i^\top] \boldsymbol{\Sigma}_0\} = \text{tr}(\boldsymbol{\Sigma}_0)/p$$

as $\mathbb{E}[\mathbf{v}_i \mathbf{v}_i^\top] = (1/p)\mathbf{I}$ and the fact that r_i is independent of \mathbf{v}_i . In the second identity we used that samples are i.i.d. and $\mathbb{E}[u(r_i^2/\sigma)(r_i^2/\sigma)] = p$ due to (18) and $\mathbb{E}[u(r_i^2/\sigma)^2(r_i^4/\sigma^2)] = \psi_1 p(p+2)$. The result then follows by substituting (45) into the last equation.

C. Proof of Lemma 2

Using $\boldsymbol{\Sigma}_0 = \sigma \boldsymbol{\Sigma}$ and (43) one obtains as in Lemma 1 the following expression

$$\begin{aligned} \mathbb{E}[\text{tr}(\mathbf{C}^2)] &= \frac{1}{n} \mathbb{E} \left[u\left(\frac{r_1^2}{\sigma}\right)^2 \frac{r_1^4}{\sigma^2} \right] \mathbb{E}[(\mathbf{v}_1^\top \boldsymbol{\Sigma}_0 \mathbf{v}_1)^2] \\ &\quad + \left(1 - \frac{1}{n}\right) \left(\mathbb{E} \left[u\left(\frac{r_1^2}{\sigma}\right) \frac{r_1^2}{\sigma} \right] \right)^2 \mathbb{E}[|\mathbf{v}_1^\top \boldsymbol{\Sigma}_0 \mathbf{v}_2|^2]. \quad (47) \end{aligned}$$

Using the facts that $\mathbb{E}[u(r_1^2/\sigma)(r_1^2/\sigma)] = p$ due to (18) and $\mathbb{E}[u(r_1^2/\sigma)^2(r_1^4/\sigma^2)] = \psi_1 p(p+1)$ along with the facts that [29, cf. eq. (66), (67)]

$$\mathbb{E}[(\mathbf{v}^\top \boldsymbol{\Sigma} \mathbf{v})^2] = [p(p+1)]^{-1} (\text{tr}(\boldsymbol{\Sigma}^2) + \text{tr}(\boldsymbol{\Sigma})^2), \quad (48)$$

$$\mathbb{E}[(\mathbf{v}_1^\top \boldsymbol{\Sigma} \mathbf{v}_2)^2] = p^{-2} \text{tr}(\boldsymbol{\Sigma}^2), \quad (49)$$

we get

$$\begin{aligned} \mathbb{E}[\text{tr}(\mathbf{C}^2)] &= \frac{\psi_1}{n} (\text{tr}(\boldsymbol{\Sigma}_0^2) + \text{tr}(\boldsymbol{\Sigma}_0)^2) + \left(1 - \frac{1}{n}\right) \text{tr}(\boldsymbol{\Sigma}_0^2) \\ &= \left(1 + \frac{\psi_1 - 1}{n}\right) \text{tr}(\boldsymbol{\Sigma}_0^2) + \frac{\psi_1}{n} \text{tr}(\boldsymbol{\Sigma}_0)^2. \end{aligned}$$

Using similar proof as in proof of Lemma 1, we obtain

$$\mathbb{E}[\text{tr}(\mathbf{C})^2] = \frac{\psi_1 p(p+1)}{n} \mathbb{E}[(\mathbf{v}_i^\top \boldsymbol{\Sigma}_0 \mathbf{v}_i)^2] + \left(1 - \frac{1}{n}\right) \text{tr}(\boldsymbol{\Sigma}_0)^2. \quad (50)$$

The result then follows by substituting the expression from (49) into the last equation.

D. Proof of Lemma 3

We show the result in the real case only as the result follows similarly for complex-valued case. Write $\mathbf{z} = \boldsymbol{\Sigma}^{-1/2} \mathbf{x}$ and note that $\mathbf{z} \sim \mathcal{E}_p(\mathbf{0}, \mathbf{I}, g)$. The result

$$\kappa = \frac{\mathbb{E}[\|\mathbf{z}\|^4]}{p(p+2)} - 1 = \frac{1}{3} (\mathbb{E}[z_1^4] - 3) \quad (51)$$

follows by recalling the stochastic decomposition. Namely, $\mathbf{z} =_d r \mathbf{v}$, where $r =_d \|\mathbf{z}\|$ is independent of \mathbf{v} , and \mathbf{v} possesses a uniform distribution on the unit sphere S^{p-1} . Thus we have that

$$\mathbb{E}[z_i^4] = \mathbb{E}[\|\mathbf{z}\|^4] \mathbb{E}[v_i^4] = 3\psi_1$$

where we used that $\mathbb{E}[v_i^4] = 3(p(p+2))^{-1}$ (see e.g. [43, Lemma A.1.] and that $\psi_1 = \mathbb{E}[\|\mathbf{z}\|^4]/p(p+2)$. Furthermore, since $\mathbb{E}[z_i^2] = 1$, (51) states that $\kappa = (1/3)\text{kurt}(z_i)$. Then since $\text{kurt}(z_i) = \text{kurt}(x_i)$, we have the stated result that $\kappa = (1/3)\text{kurt}(x_i)$.

REFERENCES

- [1] O. Ledoit and M. Wolf, "A well-conditioned estimator for large-dimensional covariance matrices," *J. Mult. Anal.*, vol. 88, no. 2, pp. 365–411, 2004.
- [2] E. Ollila and E. Raninen, "Optimal shrinkage covariance matrix estimation under random sampling from elliptical distributions," *IEEE Trans. Signal Process.*, vol. 67, no. 10, pp. 2707–2719, May 2019.
- [3] J. Li, J. Zhou, and B. Zhang, "Estimation of large covariance matrices by shrinking to structured target in normal and non-normal distributions," *IEEE Access*, vol. 6, pp. 2158–2169, 2017.
- [4] A. Coluccia, "Regularized covariance matrix estimation via empirical bayes," *IEEE Signal Process. Lett.*, vol. 22, no. 11, pp. 2127–2131, 2015.
- [5] R. A. Maronna, "Robust M-estimators of multivariate location and scatter," *Ann. Stat.*, vol. 5, no. 1, pp. 51–67, 1976.
- [6] J. T. Kent and D. E. Tyler, "Re-descending M-estimates of multivariate location and scatter," *Ann. Stat.*, vol. 19, no. 4, pp. 2102–2119, 1991.
- [7] Y. I. Abramovich and N. K. Spencer, "Diagonally loaded normalised sample matrix inversion (LNSMI) for outlier-resistant adaptive filtering," in *Proc. IEEE International Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, Apr. 15–20 2007, pp. 1105–1108.
- [8] Y. Chen, A. Wiesel, Y. C. Eldar, and A. O. Hero, "Shrinkage algorithms for MMSE covariance estimation," *IEEE Trans. Signal Process.*, vol. 58, no. 10, pp. 5016–5029, 2010.
- [9] L. Du, J. Li, and P. Stoica, "Fully automatic computation of diagonal loading levels for robust adaptive beamforming," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 46, no. 1, pp. 449–458, 2010.
- [10] E. Ollila and D. E. Tyler, "Regularized M-estimators of scatter matrix," *IEEE Trans. Signal Process.*, vol. 62, no. 22, pp. 6059–6070, 2014.
- [11] F. Pascal, Y. Chitour, and Y. Quek, "Generalized robust shrinkage estimator and its application to STAP detection problem," *IEEE Trans. Signal Process.*, vol. 62, no. 21, pp. 5640–5651, 2014.
- [12] Y. Sun, P. Babu, and D. P. Palomar, "Regularized Tyler's scatter estimator: Existence, uniqueness, and algorithms," *IEEE Trans. Signal Process.*, vol. 62, no. 19, pp. 5143–5156, 2014.
- [13] R. Couillet and M. McKay, "Large dimensional analysis and optimization of robust shrinkage covariance matrix estimators," *J. Mult. Anal.*, vol. 131, pp. 99–120, 2014.
- [14] T. Zhang and A. Wiesel, "Automatic diagonal loading for Tyler's robust covariance estimator," in *IEEE Statistical Signal Processing Workshop (SSP'16)*, 2016, pp. 1–5.
- [15] M. Yi and D. E. Tyler, "Shrinking the covariance matrix using convex penalties on the matrix-log transformation," *Journal of Computational and Graphical Statistics*, no. just-accepted, pp. 1–10, 2020.
- [16] E. Ollila, D. P. Palomar, and F. Pascal, "M-estimators of scatter with eigenvalue shrinkage," in *IEEE International Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, May 4–8 2020, pp. 5305–5309.
- [17] A. Aubry, A. De Maio, L. Pallotta, and A. Farina, "Radar detection of distributed targets in homogeneous interference whose inverse covariance structure is defined via unitary invariant functions," *IEEE Trans. Signal Process.*, vol. 61, no. 20, pp. 4949–4961, 2013.

- [18] Y. Sun, A. Breloy, P. Babu, D. P. Palomar, F. Pascal, and G. Ginolhac, "Low-complexity algorithms for low rank clutter parameters estimation in radar systems," *IEEE Trans. Signal Process.*, vol. 64, no. 8, pp. 1986–1998, 2016.
- [19] A. Aubry, A. De Maio, and L. Pallotta, "A geometric approach to covariance matrix estimation and its applications to radar problems," *IEEE Trans. Signal Process.*, vol. 66, no. 4, pp. 907–922, 2018.
- [20] A. De Maio, L. Pallotta, J. Li, and P. Stoica, "Loading factor estimation under affine constraints on the covariance eigenvalues with application to radar target detection," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 55, no. 3, pp. 1269–1283, 2019.
- [21] M. Tang, Y. Rong, X. R. Li, and J. Zhou, "Invariance theory for adaptive detection in non-gaussian clutter," *IEEE Trans. Signal Process.*, vol. 68, pp. 2045–2060, 2020.
- [22] O. Besson, "Maximum likelihood covariance matrix estimation from two possibly mismatched data sets," *Signal Processing*, vol. 167, p. 107285, 2020.
- [23] Y. Abramovich and O. Besson, "Regularized covariance matrix estimation in complex elliptically symmetric distributions using the expected likelihood approach-part 1: The over-sampled case," *IEEE Trans. Signal Process.*, vol. 61, no. 23, pp. 5807–5818, 2013.
- [24] O. Besson and Y. Abramovich, "Regularized covariance matrix estimation in complex elliptically symmetric distributions using the expected likelihood approach-part 2: The under-sampled case," *IEEE Trans. Signal Process.*, vol. 61, no. 23, pp. 5819–5829, 2013.
- [25] Y. I. Abramovich and O. Besson, "On the expected likelihood approach for assessment of regularization covariance matrix," *IEEE Signal Process. Lett.*, vol. 22, no. 6, pp. 777–781, 2015.
- [26] D. E. Tyler, "A distribution-free M-estimator of multivariate scatter," *Ann. Stat.*, vol. 15, no. 1, pp. 234–251, 1987.
- [27] —, "A note on affine equivariant location and scatter statistics for sparse data," *Statist. and Prob. Letters*, vol. 80, pp. 1409–1413, 2010.
- [28] R. Couillet, F. Pascal, and J. W. Silverstein, "The Random Matrix Regime of Maronna's M -estimator with elliptically distributed samples," *J. Mult. Anal.*, vol. 139, pp. 56–78, July 2015. arXiv:1311.7034.
- [29] Y. Chen, A. Wiesel, and A. O. Hero, "Robust shrinkage estimation of high-dimensional covariance matrices," *IEEE Trans. Signal Process.*, vol. 59, no. 9, pp. 4097 – 4107, 2011.
- [30] O. Ledoit, M. Wolf *et al.*, "Some hypothesis tests for the covariance matrix when the dimension is large compared to the sample size," *Ann. Stat.*, vol. 30, no. 4, pp. 1081–1102, 2002.
- [31] M. S. Srivastava, "Some Tests Concerning the Covariance Matrix in High Dimensional Data," *Journal of the Japan Statistical Society*, vol. 35, no. 2, pp. 251–272, 2005.
- [32] K.-T. Fang, S. Kotz, and K.-W. Ng, *Symmetric Multivariate and Related Distributions*. London: Chapman and hall, 1990.
- [33] E. Ollila, D. E. Tyler, V. Koivunen, and H. V. Poor, "Complex elliptically symmetric distributions: survey, new results and applications," *IEEE Trans. Signal Process.*, vol. 60, no. 11, pp. 5597–5625, 2012.
- [34] K. Ashurbekova, A. Usseglio-Carleve, F. Forbes, and S. Achar, "Optimal shrinkage for robust covariance matrix estimators in a small sample size setting," *hal-02378034v2*, 2020.
- [35] A. M. Zoubir, V. Koivunen, E. Ollila, and M. Muma, *Robust Statistics for Signal Processing*. Cambridge, UK: Cambridge University Press, Nov. 2018.
- [36] R. J. Muirhead, *Aspects of Multivariate Statistical Theory*. New York: Wiley, 1982, 704 pages.
- [37] S. T. Smith, "Covariance, subspace, and intrinsic Cramér-Rao bounds," *IEEE Trans. Signal Process.*, vol. 53, no. 5, pp. 1610–1630, 2005.
- [38] A. Breloy, G. Ginolhac, A. Renaux, and F. Bouchard, "Intrinsic Cramér-Rao bounds for scatter and shape matrices estimation in CES distributions," *IEEE Signal Process. Lett.*, vol. 26, no. 2, pp. 262–266, 2019.
- [39] J. Bun, J.-P. Bouchaud, and M. Potters, "Cleaning correlation matrices," *Risk Management*, 2016.
- [40] H. Markowitz, "Portfolio selection," *J. Financ.*, vol. 7, no. 1, pp. 77–91, 1952.
- [41] D. P. Palomar and R. Zhou, *portfolioBacktest: Automated Backtesting of Portfolios over Multiple Datasets*, 2019, r package version 0.2.1. [Online]. Available: <https://CRAN.R-project.org/package=portfolioBacktest>
- [42] L. Yang, R. Couillet, and M. R. McKay, "A robust statistics approach to minimum variance portfolio optimization," *IEEE Trans. Signal Process.*, vol. 63, no. 24, pp. 6684–6697, 2015.
- [43] E. Ollila, H. Oja, and C. Croux, "The affine equivariant sign covariance matrix: asymptotic behavior and efficiencies," *J. Mult. Anal.*, vol. 87, no. 2, pp. 328–355, 2003.