


## Cleaning large-dimensional covariance matrices for correlated samples

Zdzislaw Burda <sup>\*</sup>

AGH University of Science and Technology, Faculty of Physics and Applied Computer Science,  
al. Mickiewicza 30, 30-059 Kraków, Poland

Andrzej Jarosz <sup>†</sup>

CRISIL Irevna Poland Sp z o.o., ul. Świętego Mikołaja 7, 50-529 Wrocław, Poland



(Received 8 July 2021; revised 1 February 2022; accepted 24 February 2022; published 24 March 2022)

We elucidate the problem of estimating large-dimensional covariance matrices in the presence of correlations between samples. To this end, we generalize the Marčenko-Pastur equation and the Ledoit-Péché shrinkage estimator using methods of random matrix theory and free probability. We develop an efficient algorithm that implements the corresponding analytic formulas based on the Ledoit-Wolf kernel estimation technique. We also provide an associated open-source Python library, called SHRINKAGE, with a user-friendly API to assist in practical tasks of estimation of large covariance matrices. We present an example of its usage for synthetic data generated according to exponentially decaying autocorrelations.

DOI: [10.1103/PhysRevE.105.034136](https://doi.org/10.1103/PhysRevE.105.034136)

### I. INTRODUCTION

In experimental research, when working with large data sets, one often faces a generic problem of determining correlations between multiple entities of interest, on the basis of observed data. This problem, in its minimalistic form, can be formulated as follows: One performs  $T$  measurements of a statistical system with  $N$  degrees of freedom,  $Y_i$ ,  $i = 1, \dots, N$ , and collects the observations in a matrix  $Y = [y_{it}]$  of dimensions  $N \times T$ , where  $y_{it}$  is the  $t$ th measured value of the  $i$ th entity. The challenge lies in estimating from this data set the underlying two-point covariances,

$$C_{ij} = \langle Y_i Y_j \rangle, \quad (1)$$

collected in the so-called *population* (a.k.a. *true* or *signal*) *covariance matrix*,  $C = [C_{ij}]$ . (We assume throughout the paper that the data has already been centered,  $\langle Y_i \rangle = 0$ .)

This setting is very general, commonly encountered in physics, finance, genomics and bioinformatics, signal processing and acoustics, image recognition, speech recognition, cancer research, climatology, neuroscience, and many other areas. Below we sketch several such occurrences, also in each case providing relevant orders of magnitude (marked with the symbol  $\sim$ ) for  $N$  and  $T$ ; we will see that, crucially, both these numbers are typically quite large and of comparable value.

For instance, in lattice QCD, one may consider [1] neutron-antineutron two- and three-point correlation functions, described using  $N$  ( $\sim$  a few dozen) parameters related to the number of dynamics time steps. One simulates field configurations on the lattice, thereby obtaining  $T$  samples (in one version of the procedure, one sample per gauge field

configuration, thus  $T \sim$  a few dozen) for these correlation functions. These samples are used to estimate the statistical correlation matrix between the  $N$  degrees of freedom, necessary for a least-squares fit to the correlation functions.

In cosmology [2], when investigating cosmic shear, i.e., the distortion of images of distant galaxies by weak gravitational lensing caused by the large-scale structure, the basic observable is the shear two-point correlation function, computed from products of ellipticities of galaxy pairs located within given angular bins,  $N \sim$  a hundred to thousands. This correlation function can be estimated from mock simulations, only  $T \sim$  a few thousands, as they are computationally expensive.

In finance, one needs to know correlations between  $N$  investment assets, e.g., S&P 500 stocks ( $N \sim$  a few hundred), to construct optimal, well-diversified portfolios, in the spirit of Markowitz theory [3–5]. We estimate these correlations from  $T$  daily stock returns ( $T \sim$  hundreds to thousands, due to stationarity requirements).

In genomics [6], one analyzes expression profiles of  $N$  genes ( $N \sim$  thousands, from which we may select a subset of differentially expressed genes, say  $N \sim$  a hundred) to describe similarities and functional groupings of them. Each gene is sampled  $T$  times in a microarray experiment ( $T \sim$  a few observation times), and the resulting data matrix  $Y$ , having chosen some suitable gene expression measure, is called a gene expression matrix.

In signal processing and acoustics [7], one considers an array of  $N$  ( $\sim$  a dozen) directional sensors, and performs so-called beamforming, combining elements of an antenna array in such a way as to have constructive or destructive interference along given directions. But (frequency-domain) beamforming algorithms, such as the Capon beamformer, require knowledge of the correlation matrix between the array outputs, which in turn is estimated from  $T$  ( $\sim$  a few dozen) samples.

\*zdzislaw.burda@agh.edu.pl

†Corresponding author: jedrekjarosz@gmail.com

In anomaly detection in hyperspectral images [8], one has a sensor image comprised of  $N$  spectral bands ( $\sim$  a hundred), i.e., collected narrow ranges of the electromagnetic spectrum. When the goal is detecting anomalous pixels in such an image, one needs the Mahalanobis distance between each ( $N$ -dimensional) pixel and their background mean that in turn requires the correlation matrix between the bands, estimated from a local sliding window of pixels, of length  $T$  which is not too large to cover a homogeneous area of space ( $\sim$  hundreds).

**II. COVARIANCE MATRIX ESTIMATORS**

**A. Classical regime**

This problem of estimating the covariance matrix is thus of paramount importance in very diverse fields of science, and has accordingly inspired a large body of research. Actually, a situation that is even more typical than what the above examples show is that of  $N \ll T$ , meaning that one is able to collect many more samples than there are correlated entities. This is the textbook regime of *classical multivariate statistics*:

$$N \text{ finite and } T \rightarrow \infty. \tag{2}$$

In this setting, the usual *sample covariance estimator*,

$$E = \frac{1}{T} Y Y^\top, \tag{3}$$

converges almost surely to  $C$ , making it a strongly consistent estimator of the population covariance matrix. In particular, the expected value of the Frobenius norm of the difference ( $E - C$ ), i.e., the mean squared error (MSE),

$$L(E) = \langle \|E - C\|^2 \rangle = \langle \text{Tr}(E - C)^2 \rangle, \tag{4}$$

tends to zero in this limit (2). From the point of view of estimation theory,  $L$  is a *loss function* and measures the quality of an estimator, here  $E$ . In the classical statistical regime, the loss is zero for the sample estimator  $E$  (3).

**B. Big data regime and noise cleaning**

Yet, quite often, like in the examples outlined above, one rather faces a different limit, that of both  $N$  and  $T$  large and of comparable magnitude,

$$N, T \rightarrow \infty \text{ and } \frac{N}{T} \rightarrow q, \tag{5}$$

with  $q$  some positive constant. In this *big data regime*, the simultaneous estimation of  $\sim N^2$  elements of  $C$  from  $T \sim N$  samples leads to substantial noise (variance) in the sample estimator  $E$ , rendering the whole process close to meaningless. To put it in a practical context of one of the examples, if a financial manager wanted to allocate assets according to a Markowitz-type procedure, and used  $E$  in the relevant formulas (which seems a very natural thing to do!), they would be optimizing their portfolio for noise! Clearly, more powerful statistical methods are required [3].

In mathematical terms, the MSE loss (4) is no longer zero in the big data regime. The challenge is then to construct other estimators  $\Xi$  of  $C$  which have a lower value of the loss as compared to the sample estimator  $E$ . Such novel estimators would have some of the statistical noise *cleaned*, rendering them

more useful in practice. (For instance, our portfolio manager could more safely use such  $\Xi$  in the Markowitz formulas.) Due to the prevalence of big data scenarios in scientific and industrial applications and the generic nature of the estimation problem, this challenge has been met in recent years with considerable effort. We will now outline some key results, as well as add a new brick to the construction.

**C. Bayesian approach and linear shrinkage**

One way to search for estimators that are optimal in the regime (5), i.e., at finite  $q$ , is through an approach generally known as *shrinkage*. The idea is to construct estimators with higher bias, yet much lower variance, so the MSE (4) (which can be decomposed as  $\text{bias}^2 + \text{variance}$ ) is in effect reduced (shrunk).

This concept of shrinkage has been first introduced in the problem of estimating the mean  $\bar{\mu}$  of an  $N$ -dimensional normally-distributed variable  $\bar{Y}$  with given covariance matrix  $\sigma^2 \mathbb{I}$ , based on a single observation  $\bar{y}$ . The standard least-squares (or maximum-likelihood) estimator is simply

$$\bar{\mu}_{\text{LS}} = \bar{y}. \tag{6}$$

However, James-Stein (JS) [9] demonstrated that a biased estimator of the mean,

$$\bar{\mu}_{\text{JS}} = \left( 1 - \frac{(N-2)\sigma^2}{\|\bar{y}\|^2} \right) \bar{y}, \tag{7}$$

in dimension  $N \geq 3$ , has in fact a lower MSE than (6), precisely due to reduced variance, which more than compensates for increased bias. The existence of the JS estimator is often expressed as *Stein's paradox*, that when estimating three or more parameters simultaneously, there exists a combined estimator with better MSE than any method handling them separately. The paradox's resolution lies in noticing that we are reducing the *total* MSE, not that of individual components, and so the JS estimator should really be used only when this total error is of interest rather than the individual errors.

The JS estimator can be naturally understood within the *empirical Bayes* approach (cf. [4] for an extensive discussion):  $\bar{\mu}$  itself is considered a random variable, with a prior distribution assumed Gaussian of zero mean and some given variance. Minimizing the MSE (4), where now averaging is with respect to the joint distribution of  $\bar{Y}$  and  $\bar{\mu}$ , leads to expression (7). The choice of the prior is dictated by tractability: it is a *conjugate prior* with respect to the Gaussian likelihood, which means that the posterior distribution is analytically calculable; thus the minimization problem of the loss  $L$  can be carried out explicitly.

Let's return now to our current problem of covariance estimation. It has been precisely an analogous Bayesian reasoning that led to the first successful attempt at constructing a shrinkage estimator of  $C$ . Indeed, assuming the likelihood  $P(Y|C)$  to be multivariate Gaussian with covariance  $C$ , which is a common starting point in many applications, the so-called *inverse-Wishart* prior  $P(C)$  turns out to be conjugate to it. Hence, loss (4) minimization can again be carried out analytically, and the result is the *Ledoit-Wolf linear shrinkage estimator* [10]:

$$\Xi = \alpha_s E + (1 - \alpha_s) \mathbb{I}. \tag{8}$$

It interpolates (with some  $\alpha_s$  estimated from the data) between the sample estimator  $E$  and the null hypothesis, which here is the unit matrix; the null hypothesis can also be a more general prior matrix  $C_0$ , encoding a more specific prior belief, achieved by a generalized version of the inverse-Wishart prior. This simple result (8) has been the workhorse of large-dimensional estimation, and can be found in all the applications mentioned in the Introduction.

The choice of the inverse-Wishart conjugate pair is, however, dictated more by computational tractability than any insight gleaned from the data, and in fact the observed sample estimator  $E$  strongly constrains the prior distribution of  $C$ ; indeed, through the Marčenko-Pastur law (41), which we discuss later on. The recent strand of research [4,11–14], initiated by Ref. [15], thus approaches the problem of constructing shrinkage estimators in greater generality, and with marked attention to data. This we are now going to outline.

#### D. Rotationally invariant estimators and nonlinear shrinkage

Consider the eigendecomposition of the sample estimator:

$$E = \sum_{i=1}^N \lambda_i |\lambda_i\rangle \langle \lambda_i|. \quad (9)$$

As we have repeatedly stressed, in the big data regime (5), the *curse of dimensionality* means that both the eigenvalues and eigenvectors have high variance, thus concealing behind statistical noise the underlying nature of the true  $C$ . Suppose, in the first approximation, that we have no prior belief on the eigenvectors of  $C$ , so our goal is merely to clean the eigenvalues. In the absence of any preferred direction (bias) in the eigenvector space, the only available basis is that of  $E$ . In other words, we wish to seek a MSE-optimal estimator of  $C$  in the space of matrices with the eigenvectors  $|\lambda_i\rangle$ :

$$\Xi = \sum_{i=1}^N \xi_i |\lambda_i\rangle \langle \lambda_i|. \quad (10)$$

This form is called a *rotationally invariant estimator* (RIE) [4] because it can alternatively be obtained from a Bayesian argument with a prior on  $C$  invariant under orthogonal similarity transformations. Note that the linear shrinkage estimator (8) also has the same sample eigenvectors, so it belongs to the RIE class; its eigenvalues are simple linear functions of the sample eigenvalues,  $\xi_i = \alpha_s \lambda_i + 1 - \alpha_s$ . This Ledoit-Wolf estimator, despite being better in terms of the MSE (4) than the plain sample estimator, is nonetheless not optimal. Indeed, we will find out that the MSE-optimal *shrunk eigenvalues*  $\xi_i$  are nonlinear functions of  $\lambda_i$ , making (10) an example of *nonlinear shrinkage*.

#### E. Oracle RIE

Minimizing the MSE in this space (10) is a quadratic problem, and so its solution is straightforward,

$$\xi_i = \langle \lambda_i | C | \lambda_i \rangle = \frac{1}{N} \sum_{j=1}^N c_j O(\lambda_i, c_j), \quad (11)$$

where we introduce the eigendecomposition of the population covariance matrix,

$$C = \sum_{j=1}^N c_j |c_j\rangle \langle c_j|, \quad (12)$$

as well as the *overlaps* between the eigenvectors of  $E$  and  $C$ ,

$$O(\lambda_i, c_j) = N |\langle \lambda_i | c_j \rangle|^2. \quad (13)$$

Note that the scalar products  $\langle \lambda_i | c_j \rangle$  are of order  $1/\sqrt{N}$ , and so the factor of  $N$  makes the overlaps of order 1 in the large- $N$  limit. The solution (11) can be rewritten as an integral,

$$\xi_i = \int dc \rho_C(c) c O(\lambda_i, c), \quad (14)$$

making use of the *density of eigenvalues* of  $C$ ,

$$\rho_C(c) = \frac{1}{N} \sum_{j=1}^N \delta(c - c_j), \quad (15)$$

where  $\delta$  is the Dirac delta. This form is especially convenient in the big data limit (5).

The formal solution (11) is unfortunately not useful in practice as it depends on the knowledge of  $C$ , which is the very object we wish to estimate. After all, the purpose of this research is uncovering from data as much information about  $C$  as possible. Since it requires insight into the unknown, this solution is called the *oracle estimator*.

It is, however, a documented phenomenon that a MSE-optimal estimator may in fact happen to be independent of the parameter being estimated. A canonical example is estimating the population variance  $\sigma^2$  of a Gaussian random variable  $Y$ . If we restrict the search space to estimators of the form  $cV$ , where  $V = \frac{1}{T} \sum_{t=1}^T (Y_t - \bar{Y})^2$  is the sample variance estimator, and we are looking for a constant  $c$  which minimizes the MSE,  $\langle (cV - \sigma^2)^2 \rangle$ , then we find  $c = T/(T + 1)$ , independently of the true  $\sigma^2$ .

Indeed, a similar miracle occurs for our oracle estimator (11) in the big data limit (5). In this regime, it turns out that the right-hand side can be entirely expressed as a function of the sample estimator  $E$ , which can of course be easily computed based on the data at hand, without any recourse to  $C$ . This has been first demonstrated for uncorrelated samples by Ledoit and P  ch   (LP) [15], who derived an analytic expression (53) for the  $\xi_i$ 's, dependent only on  $E$  rather than  $C$ .

The purpose of the current paper is to extend this result to samples that are correlated. Our framework will be that of *random matrix theory* (RMT), as well as some ideas from the *free probability* calculus. In the next section, we will review the relevant notions of these theories.

### III. RANDOM MATRIX THEORY FOR COVARIANCE ESTIMATION

#### A. Random matrix theory transforms

Let's start by recalling a few definitions and relations which are useful in the derivation below. For a self-adjoint random matrix, an example of which is the sample covariance estimator  $E$ , the fundamental object is the *resolvent*, also

known as the *Green's function*. It is a matrix-valued function of a complex argument  $z$ :

$$G_E(z) = (z\mathbb{I} - E)^{-1}. \tag{16}$$

To fix attention, think of  $E$  as coming from the ensemble of  $N \times N$  matrices of the form (3), with  $Y$  a rectangular  $N \times T$  random matrix. We will also assume  $Y$  to have normally distributed entries. Our problem will then turn out to be analytically feasible, but at the same time it will be argued that the results are in fact representative of a wider class of matrices; hence, this assumption is less restrictive than it appears. Now, returning to the resolvent (16) in full generality, one may suppose that the randomness of  $E$  makes it a random function; however, the claim is that in the large- $N$  limit it is self-averaging, that is, independent of any realization of  $E$ , but converging to a deterministic matrix. With a slight abuse of notation, we will keep calling that deterministic function  $G_E(z)$ . In other words, the average of the resolvent over the ensemble of  $E$  can be replaced by the resolvent for a single realisation of  $E$  when  $N \rightarrow \infty$ .

Another important object is the normalized trace of the resolvent,

$$g_E(z) = \frac{1}{N} \text{Tr} G_E(z). \tag{17}$$

In the large- $N$  limit, it tends to the *Stieltjes transform* of the eigenvalue density  $\rho_E(\lambda)$ ,

$$g_E(z) = \int d\lambda \rho_E(\lambda) \frac{1}{z - \lambda}, \tag{18}$$

and thereby encodes the information about the spectrum of  $E$ . Indeed, evaluating  $g_E(z)$  near the real axis, that is, at  $z = \lambda - i\epsilon$ , with  $\epsilon \rightarrow 0^+$ , and applying the Sokhotski-Plemelj formula,

$$\frac{1}{x - i0^+} = \text{p.v.} \frac{1}{x} + i\pi \delta(x), \tag{19}$$

where p.v. stands for the Cauchy principal value, yields

$$g_E(\lambda - i0^+) = \pi h_E(\lambda) + i\pi \rho_E(\lambda), \tag{20}$$

where

$$h_E(\lambda) = \frac{1}{\pi} \text{p.v.} \int d\lambda' \rho_E(\lambda') \frac{1}{\lambda - \lambda'} \tag{21}$$

is the *Hilbert transform* of the eigenvalue density. The imaginary part of the above expression, on the other hand, is simply the eigenvalue density  $\rho_E(\lambda)$ .

A notion closely related to the resolvent is that of the *moment generating function* (MGF), also called the *M transform*. In its matrix form, analogous to (16), it is defined as

$$M_E(z) = zG_E(z) - \mathbb{I}. \tag{22}$$

Its normalized trace,

$$m_E(z) = \frac{1}{N} \text{Tr} M_E(z) = zg_E(z) - 1, \tag{23}$$

has a useful property of generating, in the  $1/z$  expansion, the *moments* of  $E$ ,

$$m_E(z) = \sum_{k=1}^{\infty} m_{E,k} z^{-k}, \tag{24}$$

where,

$$m_{E,k} = \left\langle \frac{1}{N} \text{Tr} E^k \right\rangle \xrightarrow{N \rightarrow \infty} \int d\lambda \rho_E(\lambda) \lambda^k. \tag{25}$$

In the free probability theory [16,17], one commonly uses instead of  $m_E(z)$  another generating function, the  $\psi$  transform,

$$\psi_E(z) = m_E(1/z), \tag{26}$$

since it produces the moments as coefficients of the  $z$  expansion rather than  $1/z$  as in (24).

Let's remark at this point that we have introduced our objects in pairs: a matrix version and its normalized trace version. In free probability, only the latter is relevant, and that is because there we only deal with distributions of *eigenvalues*. Here we are, however, also concerned about *eigenvectors*, and thus we need to encode the rotational information too, which we do in the more general matrix form. This is why we distinguish between *scalar* and *matrix* versions of the MGFs and related objects.

### B. Free multiplication and further transforms

The normalized-trace transforms of free probability, especially the  $\psi$ -transform (26), play a central role in *free multiplication* [16,17].

We are given two large ( $N \rightarrow \infty$ ) random matrices,  $A$  and  $B$ ; in particular, we know their respective eigenvalue densities,  $\rho_A(\lambda)$  and  $\rho_B(\lambda)$ . We assume they are *free* with respect to each other; freeness is a matrix analog of statistical independence of random variables. In such a case, it turns out it is possible to derive the eigenvalue density of their matrix product,  $\rho_{AB}(\lambda)$ , based solely on the individual densities. To be more precise,  $A$  and  $B$  should be self-adjoint, so their eigenvalues are real and the densities are real valued, but the product  $AB$  won't then typically be self-adjoint. If  $A$  is positive-definite, and  $B$  has trace different from zero, then the alternative product  $\sqrt{AB}\sqrt{A}$  has the same moments as  $AB$ , and is self-adjoint. This we will call the free product.

The prescription to do so makes use of a  $\chi$  transform, which is the functional inverse of the  $\psi$  transform (26):

$$\chi_E(\psi_E(z)) = \psi_E(\chi_E(z)) = z. \tag{27}$$

Furthermore, yet another useful object, the  $S$  transform, is related to the above as

$$S_E(z) = \frac{1+z}{z} \chi_E(z). \tag{28}$$

A justification for such a seemingly arbitrary construction is that in this language the law of free multiplication becomes very straightforward, namely, the  $S$  transforms are simply multiplicative under the free product of matrices:

$$S_{AB}(z) = S_A(z)S_B(z). \tag{29}$$

The logic here is that the known eigenvalue densities of the constituents  $A$  and  $B$  are used to calculate the respective Stieltjes transforms (18) of both matrices, then their  $\psi$  transforms through the simple formula (26), whereupon functional inversion gives the  $\chi$  transforms (27), and finally the  $S$  transforms (28). After multiplying the two according to (29), we follow the above path backward to eventually find the eigenvalue density of the free product  $\sqrt{AB}\sqrt{A}$ .



### C. The sandwich model

Remember that our present goal is to demonstrate that the oracle estimator (14), with the eigenvector overlaps (13), can be expressed entirely in terms of the observable  $E$ , rather than the true and unknown  $C$ . To this end, we need to somehow *model* the relationship between  $C$  and  $E$  or, equivalently, between  $C$  and the observed data  $Y$ . Such a model, simplistic as it may be, should be capable of capturing essential features of the data-generating process.

Recall that  $Y = [y_{it}]$  is an  $N \times T$  matrix of observations, all of which we now treat as random variables. The simplest possible model is that these entries are normally distributed, in such a way that the samples of a given entity are uncorrelated among themselves, while any two entities  $i$  and  $j$  are correlated according to the population matrix  $C_{ij}$ :

$$\langle y_{it}y_{js} \rangle = C_{ij}\delta_{ts}. \quad (30)$$

This structure has been extensively studied. The topic of the current publication is to take it one step further in a very natural direction [18–20]: by replacing the Kronecker delta in (30) with an arbitrary matrix  $A_{ts}$ , real symmetric, and positive semidefinite:

$$\langle y_{it}y_{js} \rangle = C_{ij}A_{ts}. \quad (31)$$

We also assume

$$\frac{1}{T}\text{Tr}A = 1, \quad (32)$$

so the decomposition on the right-hand side is unambiguous. You may think of the data set generated synthetically by a multivariate stochastic process (e.g., a multidimensional Ornstein-Uhlenbeck process or a VARMA process), alternatively the data set consisting of historic observations;  $A_{ts}$  describes *autocorrelations* between samples.

Note that, crucially, these autocorrelations are assumed identical for all the entities, i.e., decoupled from the cross correlations. It would be quite interesting to extend this work to a more general coupled structure,  $\langle y_{it}y_{js} \rangle = C_{ijts}$ .

For a *stationary* stochastic process, the autocorrelation matrix  $A$  is a Toeplitz matrix,

$$A_{ts} = a(|t - s|), \quad (33)$$

with  $a$  some function such that  $a(0) = 1$  (to ensure normalized trace). A simple and often physically justified case is that of exponentially decaying autocorrelations,

$$a(t) = e^{-t/\tau}, \quad (34)$$

with autocorrelation time  $\tau$ . This model is a good starting point for systems exhibiting *short-range* correlations between samples.

The model (31) is simple enough to allow for an analytical solution (below), yet, as in the central limit theorem, it is believed to describe in the large-dimensional limit a broader class of distributions (not necessarily Gaussian), only perhaps without fat tails. In fact, even fat tails may be accounted for: It is known that the maximum-likelihood estimator of correlations for fat-tailed random variables is the robust *Maronna estimator*; and it turns out [21] it can be expressed in the form (31) with a proper choice of  $A$ .

Yet another application [22] of the model (31) is that a proper choice of  $A$  allows us to describe the *exponentially weighted moving average* (EWMA) estimator of  $C$ . Essentially, instead of treating all the samples equally as in the sample estimator  $E$  (3), older samples are given lower weights than newer ones, thereby taking into account possible non-stationarity of the data. (We plan to address these cases in a separate publication.)

An attentive reader may wonder at this point how it is that while  $C$  is unknown (and we are trying to estimate it), we seem to presume the knowledge of  $A$ . Indeed, in the latter two applications, we mentioned (the Maronna and EWMA estimators),  $A$  plays more of a technical role: It turns out that the structures appearing in other contexts (covariance estimation under fat tails or samples getting obsolete due to nonstationarity, respectively), can be translated to the language of (31) with a certain  $A$ . In the first example, though, we should genuinely try to estimate  $A$  (e.g., via estimating  $\tau$ ) from the data. However, an alternative way of thinking is that  $\tau$  is an *effective parameter*, giving our model an additional degree of freedom; that added flexibility can then be used to fit the data better. A thorough investigation of this topic is left for another paper, but already at this point we announce the logic of the matter: the resulting RIE, toward which we are working, will obviously depend on  $\tau$ , and in a real-world scenario that would allow us to choose this effective autocorrelation time so the loss function (4) is lower than for  $A = \mathbb{I}$ . Interestingly, even though the MSE clearly depends on the unknown  $C$ , we will see a way of fitting  $\tau$  even without the knowledge of  $C$ .

To recapitulate, we assume a simple data-generating model, a Gaussian random matrix  $Y$  with the correlation structure (31). Since a linear combination of Gaussian variables remains Gaussian, it is easy to show that an equivalent formulation is that of a linear combination of standard normal variables,

$$Y = \sqrt{C}X\sqrt{A}, \quad (35)$$

with  $X = [x_{it}]$  an  $N \times T$  matrix consisting of independent identically distributed  $N(0, 1)$  random variables; in particular,  $\langle x_{it}x_{js} \rangle = \delta_{ij}\delta_{ts}$ . (The square roots are well-defined as  $C$  and  $A$  are positive semidefinite.) The modeled relation between the sample estimator  $E$  and the true  $C$  is therefore

$$E = \frac{1}{T}\sqrt{C}XAX^T\sqrt{C}. \quad (36)$$

Its structure is that of a *sandwich* consisting of several layers of matrices multiplied together.

### D. Marčenko-Pastur equation

Remember that we are on a quest to express the oracle estimator (11) through  $E$  rather than  $C$ . The first step has been to set up a model between the two matrices, tractable on one hand and realistic enough on the other. The sandwich model (36) has been our choice, allowing not only for correlated samples (a phenomenon observed in practical applications) but also technical extensions such as the Maronna or EWMA estimators.

The second step is to express this relationship mathematically through the random matrix theory transforms defined

above. We will see that this language will allow us to rewrite the eigenvector overlaps (13) entirely in terms of  $E$ .

There are two levels to this construction. First, we have already remarked that (36) is a product of matrices; in fact, we see that up to a certain reshuffling of factors it is a product of  $C$ ,  $A$ , and the standardized *Wishart random matrix*:

$$W = \frac{1}{T}XX^\top. \tag{37}$$

It turns out that from the point of view of *free probability*, this reshuffling does not matter, as we elucidate in Appendix A (it is essentially due to the cyclic property of trace), and so the free multiplication law (29) leads directly to [19]

$$S_E(z) = S_C(z)S_A(qz)\frac{1}{1+qz}. \tag{38}$$

The last term is the  $S$  transform of the Wishart matrix (37),  $S_W(z) = 1/(1+qz)$ . The additional factor  $q = N/T$  in the argument of  $S_A(qz)$  comes from the difference in dimensions between  $A$  and  $C$ .

Using elementary algebra and relations between the various transforms, (38) can be rewritten in a suggestive form:

$$m \equiv m_E(z) = m_C(Z), \quad \frac{Z(z)}{z} = \frac{S_A(qm)}{1+qm}. \tag{39}$$

In other words, the scalar (i.e., the normalized-trace version) MGF (23) of the sample estimator, evaluated at a complex argument  $z$ , which we denote for short  $m \equiv m_E(z)$ , is equal to the scalar MGF of the population covariance matrix but evaluated at a different complex argument  $Z$ ,  $m_C(Z)$ . This complex transformation  $z \rightarrow Z$  depends on  $m$ , and does so through the  $S$  transform of the autocorrelation matrix  $A$ . Note that we may equivalently express the right-hand side through the ransform (27),  $Z/z = \chi_A(qm)/qm$ , which will in fact be more convenient for us. We will refer to (39) as the (*scalar generalized Marčenko-Pastur equation*).

Let's also mention that (39) is sometimes spelled in the literature by explicitly writing the MGF of  $C$  as an integral (18):

$$\int d\lambda \rho_C(\lambda) \frac{\lambda}{Z-\lambda} = m. \tag{40}$$

In particular, this equation for  $A = \mathbb{I}$ , i.e., with  $Z = z/(1+qm)$ , is the classical Marčenko-Pastur law of 1967 [23]; that is why we call (39) generalized. We will, however, refrain from such a lengthy integral expression, preferring instead the brevity offered by the various transforms we have introduced.

One way of thinking about (39) is the following: Suppose one is given the true matrices  $C$  and  $A$ ; say, one generates the data from a stochastic process with these matrices as inputs. Consequently, the MGF of  $C$  and the  $S$  transform of  $A$  can be (at least in principle) calculated. Then (39) is a system of equations with unknown  $m$ . In some cases, this system is explicitly solvable; for instance, for  $C = \mathbb{I}_N$  and  $A = \mathbb{I}_T$ , we have  $m_C(u) = 1/(u-1)$  and  $S_A(u) = 1$ , directly from the definitions, and hence (39) becomes a quadratic equation for  $m$ ,  $qm^2 + m(1+q-z^2) + 1 = 0$ . Solving this equation, and using the Sokhotski-Plemelj formula (20), leads to the density of eigenvalues of  $E$ ,

$$\rho_E(\lambda) = \frac{1}{2\pi q\lambda} \sqrt{(\lambda_+ - \lambda)(\lambda - \lambda_-)}, \tag{41}$$

with  $\lambda_{\pm} \equiv (1 \pm \sqrt{q})^2$ . This is the famous *Marčenko-Pastur density* [23], and it beautifully demonstrates the very problem we are trying to solve in this paper: the true correlation eigenvalue is in this case  $c_i = 1$  (with multiplicity  $N$ ), but the observed sample eigenvalues  $\lambda_i$  are scattered around 1, smeared by statistical noise; the more so the greater  $q = N/T$  is, i.e., the more we are into the big data regime (and irrespective of how many samples  $T$  we have collected!); the underlying eigenvalue 1 is hidden inside this blob of sample eigenvalues,  $[\lambda_-, \lambda_+]$ .

Assuming the knowledge of  $C$  and  $A$ , and calculating  $m \equiv m_E(z)$  from them, is thus a useful perspective on the generalized Marčenko-Pastur equation (39), but this is *not* the problem we are trying to solve. Rather, we know  $m$ , which is straightforward to compute from the observed data through the sample estimator  $E$ . We also suppose we know  $A$ , which, as we have discussed above, is either a technical construct or can be treated as a set of effective parameters. The goal is to find  $C$ , which will turn the oracle estimator (11) into an observable quantity.

An attentive reader will now notice, though, that the oracle estimator depends on the overlaps (13) of the *eigenvectors* of  $C$  and  $E$ , which the *scalar* Marčenko-Pastur equation is not capable of capturing. Rather, we need an equation relating the *matrix* versions of the transforms in question. Interestingly enough, such an equation exists, and is in fact completely parallel to its scalar version (39),

$$M_E(z) = M_C(Z), \tag{42}$$

with complex numbers  $Z$  and  $z$  related in the exact same way as before. This is an  $N \times N$  matrix equation, expressed through the matrix MGFs (22). Of course, the previous scalar version follows by taking normalized trace of both sides. We will call (42) the *matrix generalized Marčenko-Pastur equation*. It has first been derived in Ref. [18] using diagrammatic methods; [4] presents another derivation via the replica trick and the low-rank orthogonal Harish-Chandra-Itzykson-Zuber integral.

#### IV. NONLINEAR SHRINKAGE FOR CORRELATED SAMPLES

##### A. Shrinkage estimator for the sandwich model

We are now ready to return to the problem at hand, which is to express the oracle estimator entirely through  $E$ . The key point is to rewrite the eigenvector overlaps (13) in the language of the above transforms and to apply the generalized Marčenko-Pastur equation (42) and (39) to translate all the references to transforms of  $C$  to those of  $E$  only.

To this end, we write the matrix resolvent (16) of  $E$  through the sample eigenvalues and eigenvectors (9) in an integral form weighted by the eigenvalue density,

$$G_E(z) = N \int d\lambda' \rho_E(\lambda') \frac{1}{z-\lambda'} |\lambda'\rangle\langle\lambda'|. \tag{43}$$

Evaluate this expression near the real axis, at  $z = \lambda - i0^+$ , and retain the imaginary part of the result; the Sokhotski-Plemelj

formula (19) implies

$$\frac{1}{N\pi\rho_E(\lambda)}\text{Im}G_E(\lambda-i0^+)=|\lambda\rangle\langle\lambda|. \quad (44)$$

Insert this equality inside a scalar product with an eigenvector of  $C$ :

$$\frac{1}{\pi\rho_E(\lambda)}\text{Im}\langle c|G_E(\lambda-i0^+)|c\rangle=O(\lambda,c). \quad (45)$$

We have thereby expressed the eigenvector overlap through the matrix resolvent of  $E$ ; still, however, there is an explicit reference to an eigenvector of  $C$ , which we will now remove in favor of some transform of  $C$ .

Indeed, the matrix resolvent of  $E$  evaluated at any complex  $z$  is equivalent, thanks to the matrix Marčenko-Pastur equation (42), to the matrix resolvent of  $C$  at the transformed  $Z=Z(z)$  (39), which has in turn a simple behavior when surrounding it with an eigenvector of  $C$ :

$$\langle c|G_E(z)|c\rangle=\frac{Z}{z}\langle c|G_C(Z)|c\rangle=\frac{Z}{z}\frac{1}{Z-c}. \quad (46)$$

Inserting this to (45),

$$\frac{1}{\pi\rho_E(\lambda)}\text{Im}\frac{Z(z)}{z}\frac{1}{Z(z)-c}\Big|_{z=\lambda-i0^+}=O(\lambda,c). \quad (47)$$

In the oracle estimator's generic eigenvalue  $\xi$  (14), the overlap appears multiplied by the corresponding eigenvalue  $c$ , and integrated over  $c$  with the proper eigenvalue density,  $\rho_C(c)$ . This operation, applied to the left-hand side of (47) produces, according to (23) and (18), the scalar MGF of  $C$  evaluated at the complex argument  $Z$ :

$$\frac{1}{\pi\rho_E(\lambda)}\text{Im}\frac{Z(z)}{z}m_C(Z(z))\Big|_{z=\lambda-i0^+}=\xi. \quad (48)$$

The final step is to replace in (48) the scalar MGF of  $C$  at argument  $Z$  by the scalar MGF of  $E$  at argument  $z$ , according to the scalar Marčenko-Pastur equation (39):

$$\frac{1}{\pi\rho_E(\lambda)}\text{Im}\frac{Z(z)}{z}m_E(z)\Big|_{z=\lambda-i0^+}=\xi. \quad (49)$$

And this is precisely what we set out to achieve! The oracle eigenvalue  $\xi$ , previously accessible only through the unknown  $C$ , is now expressed solely in terms of the observable  $E$ , through the scalar MGF,  $m=m_E(\lambda-i0^+)$ , present in (49) either explicitly, or inside  $Z/z=S_A(qm)/(1+qm)=\chi_A(qm)/qm$  (39).

Indeed, let's rewrite (49) even more appealingly. For any observed sample eigenvalue  $\lambda_i$ , for  $i=1,\dots,N$ , choose as a basic quantity,

$$u_i\equiv qm_E(\lambda_i-i0^+)=\alpha_i+i\beta_i, \quad (50)$$

where since we are close to the real axis, the Sokhotskiĭ-Plemelj formula implies (23), (20),

$$\alpha_i=q(\pi\lambda_i h_E(\lambda_i)-1), \quad \beta_i=q\pi\lambda_i\rho_E(\lambda_i), \quad (51)$$

through the directly observable Hilbert transform and density of eigenvalues of the sample estimator  $E$ .

In this notation, (49) becomes the main formula of this paper, the *nonlinear shrinkage for correlated samples*,

$$\frac{\xi_i}{\lambda_i}=\frac{\text{Im}\chi_A(u_i)}{\text{Im}u_i}, \quad (52)$$

where we recall the  $\chi$  transform (27) of  $A$ . We note also that since the  $S$  transform (28) is probably better known than the  $\chi$  transform, we may replace the latter in (52) by  $u_i S_A(u_i)/(1+u_i)$ .

We stress again the logic behind this formula:

(i) We perform  $T$  observations of some  $N$  given entities and collect these measurements in an  $N\times T$  data matrix  $Y$ .

(ii) From this data set, we calculate the standard sample estimator  $E$  (3), and diagonalize it, obtaining, in particular, a set of sample eigenvalues  $\lambda_i$ ,  $i=1,\dots,N$ . As we have extensively discussed, in the big data regime (5), these eigenvalues have a crucial component of statistical noise, making them unreliable for any estimation purposes.

(iii) From this set of sample eigenvalues, we first estimate the Hilbert transform  $h_E(\lambda)$  and density of eigenvalues  $\rho_E(\lambda)$ . We will discuss a relevant procedure below, based on the kernel method.

(iv) For any given sample eigenvalue  $\lambda_i$ , we thus readily calculate the corresponding  $\alpha_i$  and  $\beta_i$  (51), and so  $u_i=\alpha_i+i\beta_i$  (50).

(v) Having assumed some model of the autocorrelations  $A$  (see below), we know (at least in principle) the  $\chi$  (equivalently,  $S$ ) transform of  $A$ . Thus, (52) gives us explicitly the *cleaned* (a.k.a. *shrunk*) eigenvalue  $\xi_i$ . This is an eigenvalue of an optimal, with respect to the MSE loss function (4), RIE (10). Since the left-hand side of (52) is the ratio  $\xi_i/\lambda_i$ , the right-hand side may be called a *shrinkage factor*.

One sanity check is that for  $q\rightarrow 0$  we are moving from the big data regime (5) to the classical regime (2), in which case we should have  $\Xi=E$ , as the sample estimator is then optimal (and of course  $E$  belongs to the RIE class). Indeed, in this limit,  $u_i\rightarrow 0$ . An expansion of the  $S$  transform around zero is known [24], and when coupled with the assumed normalization condition of  $A$  (32), it implies  $S_A(u_i)\rightarrow 1$ , hence  $\chi_A(u_i)\rightarrow u_i$ , and so  $\xi_i\rightarrow \lambda_i$ , as expected.

## B. The Ledoit-Péché formula

Historically, the first nonlinear shrinkage formula, and a cornerstone of the theory we are elaborating here, has been the LP shrinkage [15] (cf. also Ref. [4]), derived for the sandwich model in the case of no correlations between samples (30).

To verify that our formula (52) reduces appropriately, set  $A=\mathbb{I}$ . Directly from the definitions (16), (22), (23), (26), the scalar  $\psi$  transform of  $A$  is  $\psi_A(z)=z/(1-z)$ , from which the  $\chi$  transform (27) follows by functional inversion,  $\chi_A(z)=z/(z+1)$ ; note that the  $S$  transform (28) is simply  $S_A(z)=1$ . The shrinkage factor (52) thus becomes  $\text{Im}(u_i/(1+u_i))/\text{Im}u_i$ , that is,

$$\frac{\xi_i}{\lambda_i}=\frac{1}{|1+u_i|^2}=\frac{1}{(\alpha_i+1)^2+\beta_i^2}, \quad (53)$$

which is indeed the celebrated LP formula.

**C. Exponentially decaying autocorrelations**

A nontrivial example of a model of correlations between the samples, often relevant for situations of short-ranged dependence between the steps of a sampling process, is that of an exponentially decaying function (34), with autocorrelation time  $\tau$ . It is a straightforward exercise [18,19] to derive the  $\chi$  transform of  $A$ ,

$$\chi_A(z) = \frac{1}{\gamma + \sqrt{\gamma^2 - 1 + 1/z^2}}, \quad (54)$$

with  $\gamma = \coth(1/\tau)$ . The shrinkage factor (52) for this case follows immediately.

There are other interesting models of autocorrelations, such as those generated by general VARMA( $r_1, r_2$ ) stochastic processes,

$$Y_{i,a} = \sum_{\beta=1}^{r_1} b_\beta Y_{i,a-\beta} + \sum_{\alpha=0}^{r_2} a_\alpha \epsilon_{i,a-\alpha}, \quad (55)$$

with  $\epsilon$  a matrix of i.i.d. standard Gaussian (or some other distribution, say Student-t) variables. Such a process has  $(r_2 + 1)$  moving average parameters  $a_\alpha$ , and  $r_1$  autoregressive parameters  $b_\beta$ . Note that the exponential decay is a case of VAR(1) with a particular set of parameters,  $b_1 = e^{-1/\tau}$  and  $a_0 = \sqrt{1 - b_1^2}$ .

Moreover, as we have alluded to above, specific models of  $A$  appear for the Maronna or EWMA estimators. These issues are left for another publication.

**D. Numerical considerations: Kernel density estimation**

As mentioned before, our shrinkage formula (52) for a general sandwich model (31) is complete except for one important component, that is, a numerical method of estimating from the observed data set the Hilbert transform and density of eigenvalues of  $E$ , which are inputs to the basic variable  $u_i$  (50) and (51).

The problem is that the eigenvalue density is approximated by a numerical histogram, and one cannot use it directly as an input for the Hilbert transform, which is an integral transform (21). Several ideas have been proposed to alleviate this difficulty.

A very simple procedure is to choose a small but finite  $\epsilon$  in (50), and approximate  $u_i \approx qm_E(\lambda_i - i\epsilon)$ . In particular, Ref. [4] uses this approach with  $\epsilon = N^{-1/2}$  and reports satisfactory behavior. In our numerical experiments, we have, however, found this algorithm very sensitive to the choice of  $\epsilon$ , and generally unstable.

For a time, a standard and comprehensive numerical solution (for uncorrelated samples) consisted of quite a complicated scheme termed *inverse QuEST* (quantized eigenvalues sampling transform) by Ledoit and Wolf [13]. In essence, it encompasses the following steps: First, suppose we know the true covariance eigenvalues  $c_1, \dots, c_N$ . The scalar Marčenko-Pastur relation (39), for  $z = \lambda - i0^+$ , can be manipulated to yield the sample eigenvalue density  $\rho_E(\lambda)$  in a parametric form, dependent on the  $c_i$ 's. [This is the same train of thought that led to the classical Marčenko-Pastur density formula (41).] Second, we integrate the density

numerically to find the cumulative distribution function,  $CDF(\lambda) = \int_0^\lambda d\lambda' \rho_E(\lambda')$ . Third, we invert it numerically to find the quantile function, i.e., from  $CDF(\lambda) = p$  to  $\lambda = Q(p)$ . Fourth, an assumption is made that the observed eigenvalues are distributed according to this quantile function, i.e., that we can approximate their positions by  $\tilde{\lambda}_i = Q(i/N)$ . The above algorithm (already quite involved) thus estimates the sample eigenvalues based on known population eigenvalues,  $\tilde{\lambda}_i(c_1, \dots, c_N)$ ; this is the QuEST function.

Now the inverse part of the prescription is the following: In reality, we do not know the  $c_i$ 's, rather we observe the  $\lambda_i$ 's. We estimate the former by minimizing the MSE between the QuEST estimates and the measured values:

$$(\tilde{c}_1, \dots, \tilde{c}_N) = \operatorname{argmin} \frac{1}{N} \sum_{i=1}^N (\tilde{\lambda}_i(c_1, \dots, c_N) - \lambda_i)^2. \quad (56)$$

Finally, we calculate  $\tilde{\lambda}_i(\tilde{c}_1, \dots, \tilde{c}_N)$ , from which  $u_i = qm_E(\tilde{\lambda}_i)$ , without a regularizer. As we see, the Ledoit-Wolf inverse QuEST is a complicated algorithm (and we have not mentioned additional numerical intricacies that need to be addressed) but provides a very solid and stable solution.

However, recently, Ledoit and Wolf [14] proposed a *kernel method* that is computationally straightforward, stable, and exceptionally easy to use in practice. Both the Hilbert transform and the eigenvalue density are approximated as a sum of  $N$  kernels set up around each  $\lambda_i$ . That is, for the density we write

$$\rho_E(\lambda) \approx \frac{1}{N} \sum_{i=1}^N \frac{1}{b_i \lambda_i} r\left(\frac{\lambda - \lambda_i}{b_i \lambda_i}\right). \quad (57)$$

Here  $r$  is a kernel function, assumed real, non-negative, normalized ( $\int dx r(x) = 1$ ), centered ( $\int dx x r(x) = 0$ ), and with unit width [ $\int dx x^2 r(x) = 1$ ]. The parameters  $b_i$  (*bandwidths*) are local scale factors, determining the width of the individual peaks. They should be chosen to make the neighboring peaks overlap. Ledoit and Wolf suggest selecting them identical for all eigenvalues (i.e., independent of  $i$ ), and equal to

$$b = T^{-1/3}. \quad (58)$$

It is a purely heuristic choice, but we too have found it stable and accurate.

The Hilbert transform (21) is a linear operation on the eigenvalue density, hence it transforms the sum (57) into another sum:

$$h_E(\lambda) \approx \frac{1}{N} \sum_{i=1}^N \frac{1}{b_i \lambda_i} h\left(\frac{\lambda - \lambda_i}{b_i \lambda_i}\right). \quad (59)$$

where  $h$  is the Hilbert transform of  $r$ .

The idea is to choose the latter such that its Hilbert transform is analytically tractable and easy to implement. A good candidate is the *Epanechnikov kernel*, essentially the non-negative part [denoted  $x^+ \equiv \max(x, 0)$ ] of a parabola,

$$r(x) = \frac{3}{4\sqrt{5}} \left(1 - \frac{x^2}{5}\right)^+. \quad (60)$$



Its Hilbert transform (21) reads

$$h(x) = \frac{3x}{10\pi} - \frac{3}{4\sqrt{5}\pi} \left(1 - \frac{x^2}{5}\right) \log \left| \frac{x - \sqrt{5}}{x + \sqrt{5}} \right|, \quad (61)$$

where the second term is understood to be zero for  $x = \pm\sqrt{5}$  and log is the natural logarithm.

To sum up, we have straightforward estimates of the sample eigenvalue density (57) and Hilbert transform (59), with the heuristic bandwidth (58), and the Epanechnikov kernel (60) and (61). They are inserted directly into the inputs (51) to the main shrinkage formula (52).

## V. ORACLE ESTIMATION BY CROSS-VALIDATION

### A. Moving-window cross validation

We have now completed the main derivation of the paper, with an essentially *analytical* formula for shrunk eigenvalues; it does require numerical components in the form of the Epanechnikov kernel estimation, as well as some further fitting procedures described below, needed in realistic situations, but the expression (52) itself is explicit. There exists, however, a purely *numerical* alternative construction that allows one to estimate the oracle eigenvalues directly from the data set  $Y$ , with no complicated random matrix theory in the process, and with excellent performance. We will now outline the method, and discuss its applicability. It has been introduced in Sec. 8.2. of Ref. [4], and is a certain modification of the method by Bartz [25].

Recall that we have started our journey from the oracle estimator (11), minimizing the Frobenius norm between the true  $C$  and the sought-for  $\Xi$  in the space of RIEs (10). The solution is simple (as the problem is quadratic),  $\xi_i = \langle \lambda_i | C | \lambda_i \rangle$ , but depends on  $C$ , which is in principle unknown. The machinery of random matrix theory in the big data limit (5) is then invoked to show that the actual dependence is solely on observable quantities, the sample eigenvalues  $\lambda_i$ , via one of the shrinkage formulas discussed in this paper.

This oracle formula can, however, be estimated by a numerical procedure of a *moving-window cross validation*. First, imagine we have collected some more samples, and now have  $T_{\text{total}} > T$  of them. Consider now a series of  $K$  pairs of consecutive moving windows. Each such pair consists of a training (in-sample) window of length  $T$ , and a testing (out-of-sample) window of some length  $T_{\text{out}}$ . This pair of windows of length  $(T + T_{\text{out}})$  we keep shifting  $K$  times through the whole data set; this, in particular, implies  $K = \lfloor (T_{\text{total}} - T) / T_{\text{out}} \rfloor$ . For the  $\mu$ th fold ( $\mu = 0, \dots, K - 1$ ), denote  $t_\mu = T + \mu T_{\text{out}} + 1$ ; following the convention used in the paper that the temporal index starts from 1, we see that the  $\mu$ th test fold has indices from  $t_\mu$  up to  $(t_\mu + T_{\text{out}} - 1)$  (both inclusive), while the  $\mu$ th train fold stretches between  $(t_\mu - T)$  and  $(t_\mu - 1)$ .

For each fold  $\mu$ , we denote by  $E^{\text{train},\mu}$  the sample estimator calculated on the current train fold and, in particular,  $|\lambda_i^{\text{train},\mu}\rangle$  its eigenvectors. On the other hand, by  $E^{\text{test},\mu}$  denote the sample estimator calculated on the current test fold.

The claim then is that the oracle eigenvalues can be estimated by the mean over all the folds of the average test sample

estimator in the state given by any train sample eigenvector:

$$\xi_i^{\text{mwc}} = \frac{1}{K} \sum_{\mu=0}^{K-1} \langle \lambda_i^{\text{train},\mu} | E^{\text{test},\mu} | \lambda_i^{\text{train},\mu} \rangle. \quad (62)$$

In other words, the numerical prescription consists of repeating  $K$  times a calculation of the (current) train sample estimator, its diagonalization, then calculation of the (current) test sample estimator, and finally taking the relevant scalar product.

An intuition behind this formula is that the unknown  $C$  in the oracle formula can be approximated by the *out of sample*, i.e., *unknown* from the point of view of an in-sample observer, estimator.

As a side note, let us make the following remark: The test sample estimator is derived from a data set of shape  $N \times T_{\text{out}}$ , i.e., with a different noise ratio than  $q = N/T$ . One might wonder if this mismatch leads then to correct estimation. An intuition why this should not matter is that the test sample estimator corresponds in the oracle formula to the true  $C$ , which, after all, knows nothing about  $q$ .

Let us also stress the importance of performing cross validation that preserves the time ordering, i.e., by moving windows. This is crucial, especially when we believe auto-correlations are present in the system. The method of Bartz, mentioned in the beginning, is exactly identical except that the train and test folds are chosen randomly from among the samples, thus leveling out any temporal dependence.

### B. Isotonic regression

The moving-window cross-validation-estimated oracle eigenvalues (62) turn out to have a significant variance; this will be visually clear in the figures below, where they form a broad cloud of scatter points.

The idea of Ref. [22] is extraordinarily simple and effective: fit *isotonic regression* to the estimated oracle eigenvalues, i.e., a monotonic function closest to the observations. We shall see that such a simple estimator is very hard to beat. Nonetheless, we shall also demonstrate that our VARMA shrinkage may be superior, once a proper search for its parameters is executed.

## VI. EXAMPLES

All the functionalities we have expounded on in this paper, including the Epanechnikov kernel estimation of  $h_E(\lambda)$  and  $\rho_E(\lambda)$ , and computing the shrunk or cleaned eigenvalues (52), under several sensible models of  $A$ , as well as the cross validation and isotonic regression procedures for the oracle eigenvalues, have been collected in an open-source Python module called SHRINKAGE. No URLs allowed in text. Move to references per journal requirements. It is versatile, extensible, and, we hope, easy to use [26].

Below we present a few examples produced by this library, with version 1.1.0. We work here with synthetic data, that is, we randomly generate the data matrix  $Y$  of dimensions  $N = 500$ ,  $T = 1000$  according to some given models of  $C$  and  $A$ .

Since the true  $C$  is known, we will be able to quantify the efficiency of various shrinkage methods. We will measure

TABLE I. Frobenius ratios for the three examples of synthetic data considered, and various shrinkage methods. (\*) means that while the fit of VARMA shrinkage parameters in examples 1 and 2 is that of an exp-decay, in example 3 we fit a three-parameter VARMA(1, 1) model.

	Example 1	Example 2	Example 3
Oracle mwcν	13%	23%	34%
Oracle mwcν isotonic	12%	22%	34%
LP	53%	33%	60%
Effective LP	33%	26%	37%
VARMA fit	12%	24%	27% (*)

it by the MSE (4); more precisely, for a given shrinkage estimator  $\Xi$  we will calculate its *Frobenius ratio*:

$$\mathcal{F} = \frac{\text{Tr}(\Xi - C)^2}{\text{Tr}(E - C)^2}. \tag{63}$$

The lower it is, the better the cleaning scheme; in other words,  $\Xi$  should be closer to  $C$  than  $E$  is. The Frobenius ratios found in the analysis to follow are summarized in Table I.

**A. Example 1: Gaussian + VAR(1)**

First, we assume a multivariate Gaussian distribution of  $Y$ , with the underlying true covariance matrix  $C$  having two distinct eigenvalues, 1 and 3, in 50% proportions, and with the autocorrelation matrix  $A$  given by the exponential decay model (34), with the autocorrelation time  $\tau = 3$ . This case encompasses the first group of Fig. 1.

Figures 1(a)–(c) show the same histogram of the sample eigenvalues  $\lambda_i$ 's (*peach*), as well as the Epanechnikov-kernel-approximated density of eigenvalues  $\rho_E(\lambda)$  (*orange*), enveloping the histogram. We see how broad this distribution is, completely obscuring behind statistical noise the true eigenvalues 1 and 3. We expect that shrinkage will narrow down this histogram, making it more localized about the true eigenvalues.

The first shrinkage approach we choose is the simplest but very powerful isotonic regression fitted to the cross-validation-estimated oracle eigenvalues (62) (with  $K = 10$  and  $T_{\text{out}} = 50$ ). Their histogram (*aqua*) and density (*teal*) in Fig. 1(a) reveal two clear peaks close to 1 and 3. The very low Frobenius ratio of 12% further confirms how effective this shrinkage prescription is.

In Fig. 1(b), we suppose for a moment that we are ignorant of the fact that there are autocorrelations in  $Y$ , and we choose to apply the LP shrinkage (53) [which recall is valid

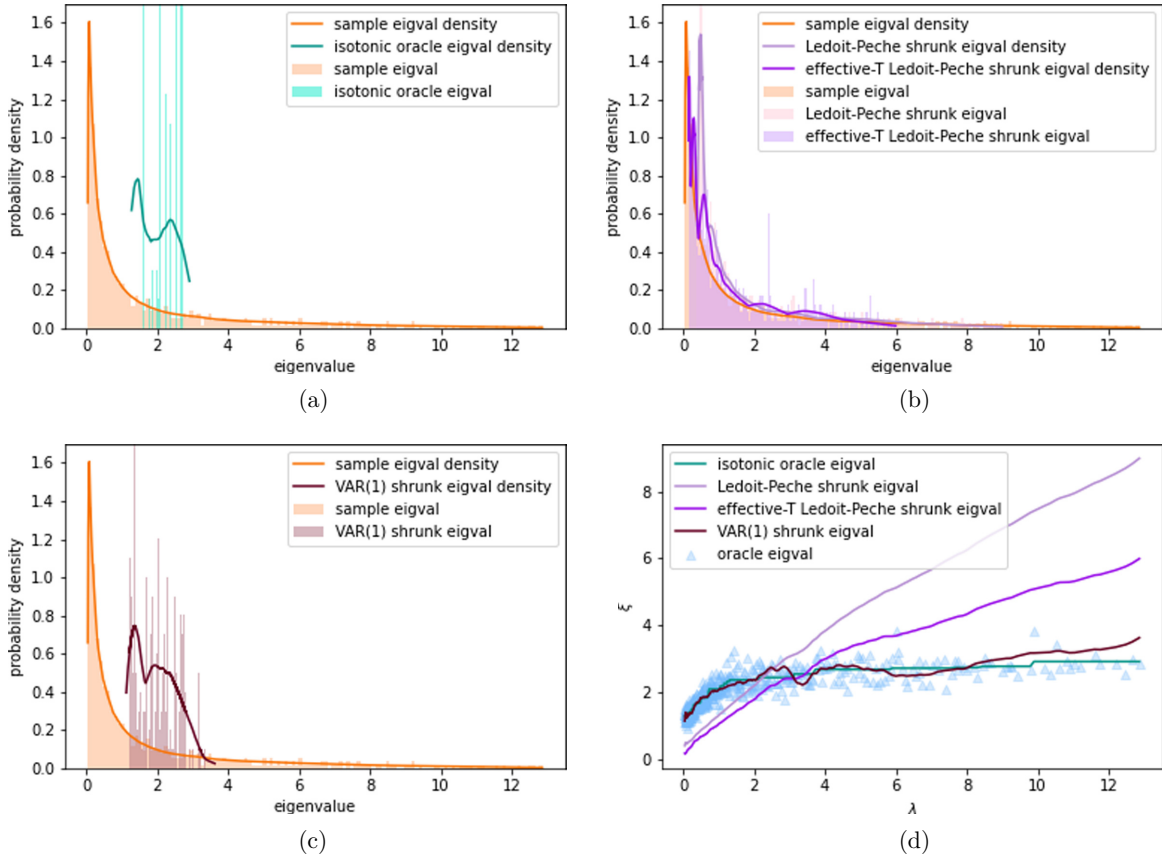


FIG. 1. Synthetic data model: Gaussian distribution,  $C$  with eigenvalues 1 and 3 in 50% proportions,  $A$  exp-decay with  $\tau = 3$ . Figures 1(a)–(c) are the empirical (bars) and kernel-density estimated (lines) distributions of the eigenvalues for the data and various estimators. Figure 1(d) contains plots of the shrinkage relationship  $\xi(\lambda)$  considered in the previous figures, as well as the cross-validation-estimated oracle eigenvalues (triangles).

for the sandwich model with no auto-correlations (30)]. The histogram (*pale pink*) and density (*pale purple*) of these  $\xi_i^{\text{LP}}$  is only slightly narrowed as compared to the sample spectrum. The Frobenius ratio of 53% is large.

One might argue at this point [28] that the reason why the LP formula does not shrink the eigenspectrum enough is that it underestimates the noise level  $q = N/T$ . Indeed, when autocorrelations with a characteristic time  $\tau_{\text{eff}}$  are present in the system, the effective number of samples is  $T_{\text{eff}} = T(1 - e^{-1/\tau_{\text{eff}}}) \sim T/\tau_{\text{eff}}$ , smaller than  $T$ . An interesting exercise would therefore be to find an effective number of samples (by varying  $\tau_{\text{eff}}$ ) such that the LP estimator performs better in terms of the Frobenius ratio. A caveat here is that it would be illegal to minimize the Frobenius ratio itself; it is, after all, an unknown quantity, dependent on  $C$ . One can, however, employ the following trick: fit rather to the scatter points of the cross-validation-estimated oracle eigenvalues (by minimizing the MSE between the two). Indeed, we thus get  $T_{\text{eff}} = 426$ , with the Frobenius ratio  $\mathcal{F} = 33\%$ . It is significantly lower than for the true- $T$  LP result, thus demonstrating the validity of this effective approach. Visually, the histogram (*lilac*) and density (*violet*) in Fig. 1(b) start slowly developing two peaks in the vicinity of 1 and 3.

Finally, apply our VARMA shrinkage formula (52). Since the data-generating process is known and, in particular,  $A$  is that of an exp-decay model with  $\tau = 3$ , one could be tempted to use the same VARMA parameters in the shrinkage formula. But one should immediately object that this information is not available in any real-world scenario. However, just like with the effective  $T_{\text{eff}}$  above, there is a way to *fit*  $\tau$  to the data even without knowing  $C$ ; indeed, to the cross-validation-estimated oracle eigenvalues. In this way,  $\tau$  becomes an effective parameter, giving our model greater flexibility than the LP case. Figure 1(c) shows the histogram (*mauve*) and density (*maroon*) of the shrunk eigenvalues  $\xi_i$ . The fit here is  $\tau = 2.9$  and we observe how the broad expansion of the sample eigenvalues is efficiently shrunk to two peaks close to 1 and 3, the true eigenvalues. Note, in passing, that the best fit is slightly different from the true  $\tau = 3$ , which may be an expression of a small bias present in the cross-validation-estimated oracle eigenvalues. The Frobenius ratio is much lower, too, at about 12%. This number is comparable to the (much simpler) isotonic regression estimator from Fig. 1(a).

### B. Example 2: Student-t + VMA(1)

To further simulate a real-world situation, now generate the data matrix  $Y$  from the Student-t distribution with  $\mu = 5$  degrees of freedom (i.e., with heavy tails), and with the auto-correlation matrix  $A$  of the VMA(1) (vector moving average) model with parameters  $a_0 = 0.8$ ,  $a_1 = 0.5$ , cf. (55). In other words, the distribution is quite far from the Gaussian, nor the correlations between samples are exponentially decaying. (The model for  $C$  stays as before, with two peaks at 1 and 3.)

Looking at Fig. 2, we see now that all four shrinkage estimators perform quite similarly, with the best isotonic regression at 22% Frobenius ratio, the LP at 33%, the effective LP (with fitted  $T_{\text{eff}} = 811$ ) at 26%, while our exp-decay shrinkage (with fitted  $\tau = 1.1$ ) has 24%. Two peaks are visible in all cases.

### C. Example 3: Student-t + VARMA(1, 1)

To make the data distribution even more complicated, now generate  $Y$  from the Student-t distribution with three degrees of freedom (so, very heavy tails),  $C$  with eigenvalues following the inverse-Wishart distribution with  $\kappa = 2$ , and  $A$  that of a VARMA(1, 1) model (55) with  $a_0 = 0.8$ ,  $a_1 = 0.5$ ,  $b_1 = 0.4$ . Recall [4] that for  $C$  to have random eigenvalues from the *inverse-Wishart* distribution with a parameter  $\kappa$ , we first calculate a noise ratio  $q_{\text{IW}} = 1/(1 + 2\kappa)$  and the corresponding number of samples,  $T_{\text{IW}} = \lfloor N/q_{\text{IW}} \rfloor$ . Generate an i.i.d. Gaussian random matrix  $R$  of shape  $N \times T_{\text{IW}}$  and form from it a Wishart random matrix,  $W = \frac{1}{T_{\text{IW}}} RR^T$ , i.e., its sample estimator. Finally,  $C = (1 - q_{\text{IW}})W^{-1}$ .

Figure 3 reveals the same set of graphs as in the previous two examples. There is one difference in the setup, though: Shown in Fig. 3(c) are the shrunk eigenvalues corresponding not to an exp-decay shrinkage with some optimal  $\tau$  fitted to the cross-validation-estimated oracle eigenvalues, as we did before, but a more general three-parameter VARMA(1, 1) shrinkage. We fit  $a_0$ ,  $a_1$ ,  $b_1$  by brute force, simply iterating over a certain grid of values (quite coarse, as this is simply a proof of concept) and selecting those for which the MSE distance to the estimated oracle eigenvalues is the smallest. The main finding from this analysis is that the Frobenius ratio of 27% of such an optimal fit outperforms the isotonic regression estimator, the latter with 34%.

In other words, despite severe model mismatch, the mere fact of incorporating autocorrelations into the shrinkage formula vastly improves the quality of eigenvalue cleaning, sometimes even better than the numerical isotonic regression scheme. This bodes well for empirical, real-world applications, where the underlying data-generating processes are genuinely unknown.

## VII. CONCLUSIONS

The main result of this paper is an exact formula (50)–(52) for the RIE [that is, sharing the eigenbasis with the standard sample estimator (3), but with different eigenvalues] (10) of the true covariance matrix (1), which is optimal with respect to the MSE (4). The result is crucial in the big-data regime (5), when the number of correlated entities  $N$  is large and of comparable magnitude with the number of collected samples  $T$ , since then the classical sample estimator contains mostly statistical noise instead of any information about actual correlations. The main novelty of our formula is that it allows for correlations between samples, thus extending the method of LP [15]. Such autocorrelations effectively decrease the number of samples  $T$ , thus making the estimation task even harder; a specifically tailored solution such as we present here is therefore of great importance.

We couple our analytical expression with a straightforward numerical scheme based on the kernel estimation method of Ledoit and Wolf. This means a numerically stable and effective pipeline which leads directly from the observed data set to the cleaned or shrunk eigenvalues. We implement this end-to-end in an open-source Python library, encouraging its use to practitioners of various fields of science and industry

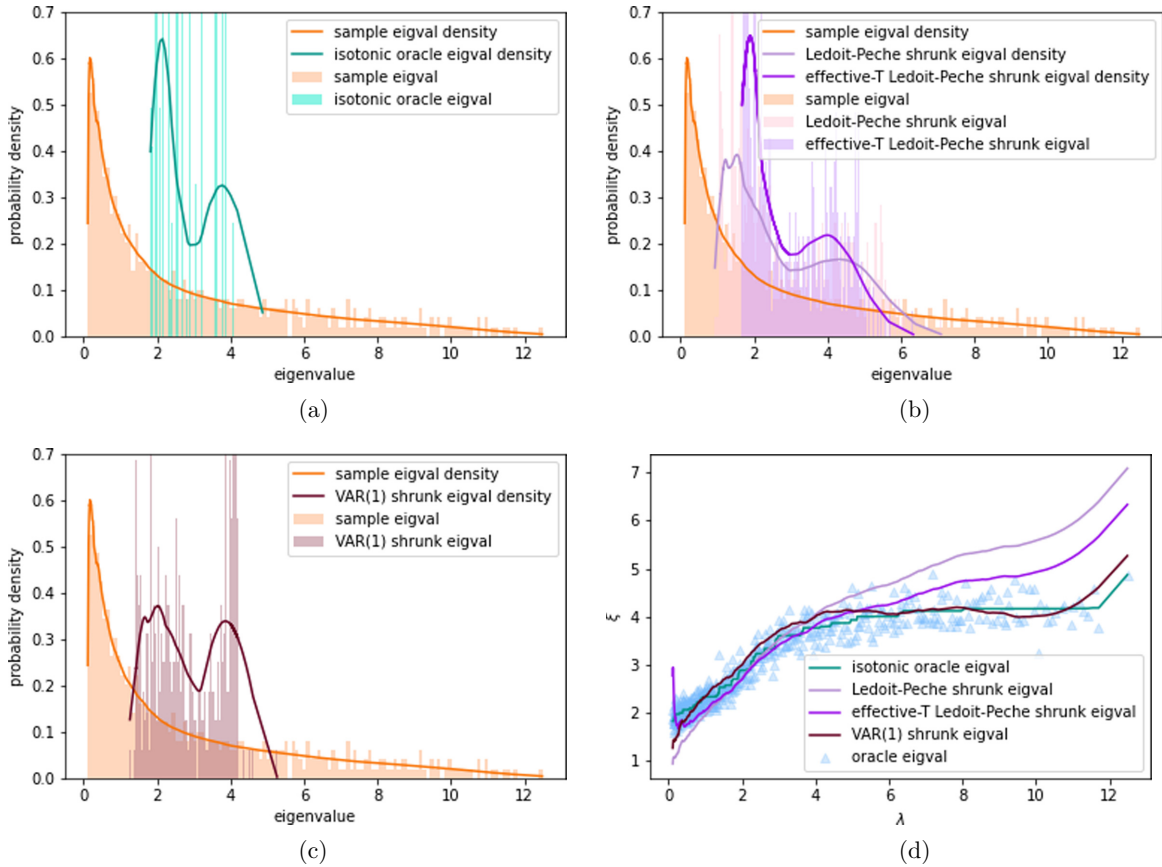


FIG. 2. Synthetic data model: Student-t distribution with five degrees of freedom,  $C$  with eigenvalues 1 and 3 in 50% proportions,  $A$  that of a VMA(1) model with  $a_0 = 0.8$ ,  $a_1 = 0.5$ . Figures 2(a)–(c) are the empirical (bars) and kernel-density estimated (lines) distributions of the eigenvalues for the data and various estimators. Figure 2(d) contains plots of the shrinkage relationship  $\xi(\lambda)$  considered in the previous figures, as well the cross-validation-estimated oracle eigenvalues (triangles).

where the problem of covariance estimation in the big data regime appears.

We leave for further work the following questions: One clear task is to apply our shrinkage formula in a realistic situation, for instance, a Markowitz-type portfolio optimization, on different markets, hopefully demonstrating that taking into account autocorrelations leads to better performance.

Another question is to consider more general, but still tractable, models of autocorrelations, by which we mean that their  $\chi$  transform is analytically accessible, akin to (54). We have already checked that some low-order VARMA models, in particular, VAR(1), VMA(1), VARMA(1, 1), VAR(2), VMA(2), belong to this class and, in fact, they are now available in the SHRINKAGE library. In fact, we have used here a VARMA(1, 1) model; it would be very beneficial to devise better methods of fitting parameters of such higher-order models (instead of a brute-force evaluation over a given grid), and interesting to see how they perform, especially compared to the isotonic regression estimator.

We have talked about how EWMA and Maronna estimators naturally fit within our framework of autocorrelations. One task would be to investigate the performance of an EWMA estimator in a situation of nonstationary data, such as on financial markets. On the other hand, the Maronna estimator has a much more involved form, and it would be interesting to see if it can be handled analytically.

This whole work falls within the class of RIEs, in which the key assumption is that we possess no belief about the eigenvectors of  $C$ . Once such a bias is available, though, for example, as a *factor model*, one could attempt to work out an appropriate optimal estimator.

Finally, and more speculatively, one could look at *non-symmetric* correlations, such as between different entities at *different* time moments, say with a time lag of one. Would it be possible to extend the matrix Marčenko-Pastur equation to this case, working out a proper generalization of the replica trick and the HCIZ integral?

ACKNOWLEDGMENT

We thank Artur Świąch, Vincent W.C. Tan, Stefan Zohren, and Christopher Wells for interesting discussions.

APPENDIX: SCALAR MARČENKO-PASTUR EQUATION AND FREE PRODUCTS OF RECTANGULAR MATRICES

In this Appendix, we sketch a derivation of the scalar Marčenko-Pastur equation (38). An important auxiliary result concerns multiplying *rectangular* random matrices. If  $W$  is  $N \times T$ , while  $V$  is  $T \times N$ , then  $WV$  is  $N \times N$ , and  $VW$  is  $T \times T$ . Since the moments (25) are normalized by the matrix dimension, we find, due to the cyclic property of trace,



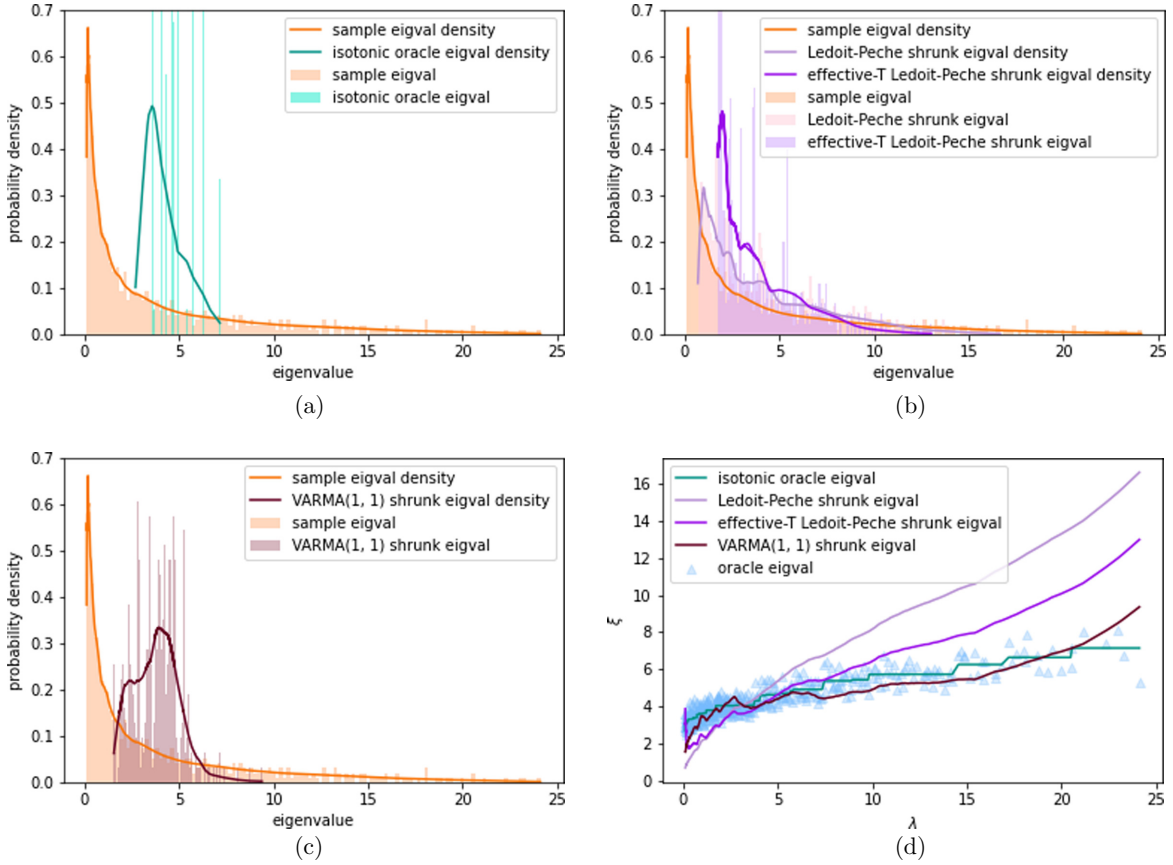


FIG. 3. Synthetic data model: Student-t distribution with three degrees of freedom,  $C$  with eigenvalues following the inverse-Wishart distribution with  $\kappa = 2$ ,  $A$  that of a VARMA(1, 1) model with  $a_0 = 0.8$ ,  $a_1 = 0.5$ ,  $b_1 = 0.4$ . Figures 3(a)–3(c) are the empirical (bars) and kernel-density estimated (lines) distributions of the eigenvalues for the data and various estimators. Figure 3(d) contains plots of the shrinkage relationship  $\xi(\lambda)$  considered in the previous figures, as well the cross-validation-estimated oracle eigenvalues (triangles).

$m_{VW,k} = \frac{1}{T} \text{Tr}(VW)^k = q \frac{1}{N} \text{Tr}(WV)^k = qm_{WV,k}$ . Thus, the scalar MGF (23) satisfies the same relation,  $m_{VW}(z) = qm_{WV}(z)$ . A short manipulation then leads to the corresponding relation between the  $S$  transforms (28) [19]:

$$S_{WV}(z) = \frac{q(1+z)}{1+qz} S_{VW}(qz). \quad (\text{A1})$$

For brevity, denote the prefactor here by  $f$ , as it won't be important in the following calculation; the important piece is the  $q$  in the argument on the right-hand side.

Consider now the  $S$  transform of the sample estimator in the sandwich model (36). Since the  $S$  transform depends on the matrix through its moments, we have as a consequence of the cyclic property of trace that

$$S_E(z) = S_{\frac{1}{T} X A X^T C}(z) = S_{\frac{1}{T} X A X^T}(z) S_C(z),$$

where the second equality comes from the free multiplication law (29). Now, in the first term, we move  $X^T$  to the front; this changes the dimension of the matrix, so we need to compensate according to (A1),

$$S_{\frac{1}{T} X A X^T}(z) = f S_{\frac{1}{T} X^T X A}(qz) = f S_{\frac{1}{T} X^T X}(qz) S_A(qz),$$

where again the second equality is the multiplication law. We now again move  $X$  to the front, in the first term, obtaining for the last expression the value of  $S_{\frac{1}{T} X X^T}(z) S_A(qz)$ , according to (A1). The first matrix here is the standard Wishart random matrix (37), and so we can write altogether

$$S_E(z) = S_C(z) S_A(qz) S_W(z),$$

with the  $q$  in the argument of  $S_A$  reminiscent of (A1). The  $S$  transform of the Wishart matrix is well-known [27],  $S_W(z) = 1/(1+qz)$ , which completes our proof of the scalar Marčenko-Pastur equation (38), based on the free multiplication law (29).

- [1] E. Rinaldi, S. Syritsyn, M.L. Wagman, M.I. Buchoff, C. Schroeder, and J. Wasem, *Phys. Rev. D* **99**, 074510 (2019).  
 [2] B. Joachimi, [arXiv:1612.00752](https://arxiv.org/abs/1612.00752).

- [3] R. Michaud, *Financ. Anal. J.* **45**, 31 (1989).  
 [4] J. Bun, J.-P. Bouchaud, and M. Potters, *Phys. Rep.* **666**, 1 (2017).

- [5] O. Ledoit and M. Wolf, *Rev. Financ. Stud.* **30**, 4349 (2017).
- [6] J. Schäfer and K. Strimmer, *Stat. Appl. Genet. Mol. Biol.* **4**, 32 (2005).
- [7] Y. Zhang, D. Sun, and D. Zhang, *Appl. Acoust.* **70**, 1029 (2009).
- [8] L. R. Bachege, J. Theiler, and C. A. Bouman, 2011 IEEE Applied Imagery Pattern Recognition Workshop (AIPR), 1 (2011).
- [9] W. James and C. Stein, Proc. Fourth Berkeley Symp. Math. Statist. Prob. **1**, 361 (1961).
- [10] O. Ledoit and M. Wolf, *J. Multivar. Anal.* **88**, 365 (2004).
- [11] O. Ledoit and M. Wolf, *Ann. Stat.* **40**, 1024 (2012).
- [12] O. Ledoit and M. Wolf, *J. Multivar. Anal.* **139**, 360 (2015).
- [13] O. Ledoit and M. Wolf, *Comput. Stat. Data Anal.* **115**, 199 (2017).
- [14] O. Ledoit and M. Wolf, *Ann. Stat.* **48**, 3043 (2020).
- [15] O. Ledoit and S. Pécché, *Probab. Theory Relat. Fields* **151**, 233 (2011).
- [16] D. V. Voiculescu, *J. Operator Theory* **18**, 223 (1987).
- [17] D. V. Voiculescu, K. J. Dykema, and A. Nica, *Free Random Variables*, CRM Monograph Series 1 (American Mathematical Society, Providence, RI, 1992).
- [18] Z. Burda, J. Jurkiewicz, and B. Waław, *Phys. Rev. E* **71**, 026111 (2005).
- [19] Z. Burda, A. Jarosz, J. Jurkiewicz, M. A. Nowak, G. Papp, and I. Zahed, *Quant. Finance* **11**, 1103 (2011).
- [20] Z. Burda, A. Jarosz, M. A. Nowak, and M. Snarska, *New J. Phys.* **12**, 075036 (2010).
- [21] R. Couillet, A. Kammoun, and F. Pascal, *J. Multivar. Anal.* **143**, 249 (2016).
- [22] V. W. C. Tan and S. Zohren, [arXiv:2012.05757](https://arxiv.org/abs/2012.05757).
- [23] V. A. Marčenko and L. A. Pastur, *Mat. Sb.* **72**, 507 (1967).
- [24] Z. Burda, *J. Phys.: Conf. Ser.* **473**, 012002 (2013).
- [25] D. Bartz, [arXiv:1611.00798](https://arxiv.org/abs/1611.00798).
- [26] A good entry point to the library is the *User's Guide* Jupyter notebook present in the repository, with detailed descriptions of its structure. For a quick look at the whole pipeline, the reader may consult its Sec. 4.3.6., in which an end-to-end problem is solved, from loading a real-world data set to producing plots such as shown in this paper.
- [27] R. A. Janik, M. A. Nowak, G. Papp, J. Wambach, and I. Zahed, *Phys. Rev. E* **55**, 4100 (1997).
- [28] We thank an anonymous referee for this suggestion.