



Institute for Empirical Research in Economics
University of Zurich

Working Paper Series
ISSN 1424-0459

Working Paper No. 515

**Nonlinear Shrinkage Estimation of
Large-Dimensional Covariance Matrices**

Olivier Ledoit and Michael Wolf

October 2010

Nonlinear Shrinkage Estimation of Large-Dimensional Covariance Matrices

Olivier Ledoit

Inst. for Emp. Research in Economics
University of Zurich
CH-8006 Zurich, Switzerland
oledoit@iew.uzh.ch

Michael Wolf*

Inst. for Emp. Research in Economics
University of Zurich
CH-8006 Zurich, Switzerland
mwolf@iew.uzh.ch

October 2010

Abstract

Many applied problems require an estimate of a covariance matrix and/or its inverse. When the matrix dimension is large compared to the sample size, which happens frequently, the sample covariance matrix is known to perform poorly and may suffer from ill-conditioning. There already exists an extensive literature concerning improved estimators in such situations. In the absence of further knowledge about the structure of the true covariance matrix, the most successful approach so far, arguably, has been shrinkage estimation. Shrinking the sample covariance matrix to a multiple of the identity, by taking a weighted average of the two, turns out to be equivalent to linearly shrinking the sample eigenvalues to their grand mean, while retaining the sample eigenvectors. Our paper extends this approach by considering nonlinear transformations of the sample eigenvalues. We show how to construct an estimator which is asymptotically equivalent to an oracle estimator suggested in previous work. As demonstrated in extensive Monte Carlo simulations, the resulting *bona fide* estimator can result in sizeable improvements over the sample covariance matrix and also over linear shrinkage.

KEY WORDS: Large-dimensional asymptotics, nonlinear shrinkage, rotation equivariance.

JEL CLASSIFICATION NOS: C13.

*Research has been supported by the NCCR Finrisk project “New Methods in Theoretical and Empirical Asset Pricing”.

1 Introduction

Many statistical applications require an estimate of a covariance matrix and/or of its inverse when the matrix dimension, p , is large compared to the sample size, n . It is well-known that in such situations, the usual estimator — the sample covariance matrix — performs poorly. It tends to be far from the population covariance matrix and ill-conditioned. The goal then becomes to find estimators which outperform the sample covariance matrix, both in finite samples and asymptotically. For the purposes of asymptotic analyses, to reflect the fact that p is large compared to n , one has to employ large-dimensional asymptotics where p is allowed to go to infinity together with n . In contrast, standard asymptotics would assume that p remains fixed while n tends to infinity.

One way to come up with improved estimators is to incorporate additional knowledge in the estimation process, such as sparseness, a graph model, or a factor model; for example, see Bickel and Levina (2008), Rajaratnam et al. (2008), Khare and Rajaratnam (2010), Fan et al. (2008), and the references therein.

However, not always is such additional knowledge available or trustworthy. In this general case, it is reasonable to require that covariance matrix estimators be rotation-equivariant. This means that rotating the data by some orthogonal matrix rotates the estimator in exactly the same way. In terms of the well-known decomposition of a matrix into eigenvectors and eigenvalues, an estimator is rotation-equivariant if and only if it has the same eigenvectors as the sample covariance matrix. Therefore, it can only differentiate itself by its eigenvalues.

Ledoit and Wolf (2004) demonstrate that the largest sample eigenvalues are systematically biased upwards, and the smallest ones downwards. It is advantageous to correct this bias by pulling down the largest eigenvalues and pushing up the smallest ones, towards the grand mean of all sample eigenvalues. This is an application of the general shrinkage principle, going back to Stein (1955). Working under large-dimensional asymptotics, Ledoit and Wolf (2004) derive the optimal *linear* shrinkage formula (when the loss is defined as the Frobenius norm of the difference between the estimator and the true covariance matrix). The same shrinkage intensity is applied to all sample eigenvalues, regardless of their positions. For example, if the linear shrinkage intensity is 0.5, then every sample eigenvalue is moved half-way towards the grand mean of all sample eigenvalues. Ledoit and Wolf (2004) both derive asymptotic optimality properties of the resulting estimator of the covariance matrix and demonstrate that it has desirable finite-sample properties via simulation studies.

A cursory glance at the Marčenko and Pastur (1967) equation, which governs the relationship between sample and population eigenvalues under large-dimensional asymptotics, shows that linear shrinkage is the first-order approximation to a fundamentally nonlinear problem. How good is this approximation? Ledoit and Wolf (2004) are very clear about this. Depending on the situation at hand, the improvement over the sample covariance matrix can either be

to use nonlinear shrinkage.

Many applied problems require an estimate of the precision matrix, which is the inverse of the covariance matrix, instead of (or in addition to) an estimate of the covariance matrix itself. Of course, one possibility is to simply take the inverse of the nonlinear shrinkage estimate of the covariance matrix itself. However, this would be *ad hoc*. The superior approach is to estimate the inverse covariance matrix directly by nonlinearly shrinking the inverses of the sample eigenvalues. This gives quite different and — according to both theory and Monte-Carlo simulations — better results. We provide a detailed, in-depth resolution of this important question as well.

The remainder of the paper is organized as follows. Section 2 defines our framework for large-dimensional asymptotics and reviews some fundamental results from the corresponding literature. Section 3 presents the oracle shrinkage estimator which motivates our *bona fide* nonlinear shrinkage estimator. Sections 4 and 5 show that the *bona fide* estimator is consistent for the oracle estimator. Section 6 examines finite-sample behavior. Finally, Section 7 concludes. All mathematical proofs are collected in the appendix.

2 Large-Dimensional Asymptotics

2.1 Basic Framework

Let n denote the sample size and $p \equiv p(n)$ the number of variables, with $p/n \rightarrow c \in (0, 1)$ as $n \rightarrow \infty$. This framework is known as large-dimensional asymptotics. The restriction to the case $c < 1$ which we make here somewhat simplifies certain mathematical results as well as the implementation of our routines in software. The case $c > 1$, where the sample covariance matrix is singular, could be handled by similar methods, but is left to future research.

The following set of assumptions will be maintained throughout the paper.

- (A1) The population covariance matrix Σ_n is a nonrandom p -dimensional positive definite matrix.
- (A2) Let X_n be an $n \times p$ matrix of real independent and identically distributed (i.i.d.) random variables with zero mean and unit variance. One only observes $Y_n \equiv X_n \Sigma_n^{1/2}$, so neither X_n nor Σ_n are observed on their own.
- (A3) Let $((\tau_{n,1}, \dots, \tau_{n,p}); (v_{n,1}, \dots, v_{n,p}))$ denote a system of eigenvalues and eigenvectors of Σ_n . The empirical distribution function (e.d.f.) of the population eigenvalues is defined as: $\forall t \in \mathbb{R}, H_n(t) \equiv p^{-1} \sum_{i=1}^p \mathbf{1}_{[\tau_{n,i}, +\infty)}(t)$, where $\mathbf{1}$ denotes the indicator function of a set. We assume $H_n(t)$ converges to some limit $H(t)$ at all points of continuity of H .
- (A4) $\text{Supp}(H)$, the support of H , is the union of a finite number of closed intervals, bounded away from zero and infinity. Furthermore, there exists a compact interval in $(0, +\infty)$

that contains $\text{Supp}(H_n)$ for all n large enough.

Let $((\lambda_{n,1}, \dots, \lambda_{n,p}); (u_{n,1}, \dots, u_{n,p}))$ denote a system of eigenvalues and eigenvectors of the sample covariance matrix $S_n \equiv n^{-1}Y_n'Y_n = n^{-1}\Sigma_n^{1/2}X_n'X_n\Sigma_n^{1/2}$. We can assume that eigenvalues are sorted in increasing order without loss of generality (w.l.o.g.). The first subscript, n , will be omitted when no confusion is possible. The e.d.f. of the sample eigenvalues is defined as: $\forall \lambda \in \mathbb{R}, F_n(\lambda) \equiv p^{-1} \sum_{i=1}^p \mathbf{1}_{[\lambda_i, +\infty)}(\lambda)$.

In the remainder of the paper, we shall use the notations $\text{Re}(z)$ and $\text{Im}(z)$ for the real and imaginary parts, respectively, of a complex number z , so that:

$$\forall z \in \mathbb{C} \quad z = \text{Re}(z) + i \cdot \text{Im}(z) .$$

The Stieltjes transform of a nondecreasing function G is defined by

$$\forall z \in \mathbb{C}^+ \quad m_G(z) \equiv \int_{-\infty}^{+\infty} \frac{1}{\lambda - z} dG(\lambda) ,$$

where \mathbb{C}^+ is the half-plane of complex numbers with strictly positive imaginary part. The Stieltjes transform has a well-known inversion formula:

$$G(b) - G(a) = \lim_{\eta \rightarrow 0^+} \frac{1}{\pi} \int_a^b \text{Im}[m_G(\xi + i\eta)] d\xi ,$$

which holds if G is continuous at a and b . Thus, the Stieltjes transform of the e.d.f. of sample eigenvalues is:

$$\forall z \in \mathbb{C}^+ \quad m_{F_n}(z) = \frac{1}{p} \sum_{i=1}^p \frac{1}{\lambda_i - z} = \frac{1}{p} \text{Tr} [(S_n - zI)^{-1}] ,$$

where I denotes a conformable identity matrix.

2.2 Marčenko-Pastur Equation and Reformulations

Marčenko and Pastur (1967) and others have proven that $F_n(\lambda)$ converges almost surely (a.s.) to some nonrandom limit $F(\lambda)$ at all points of continuity of F under certain sets of assumptions. Furthermore, Marčenko and Pastur discovered the equation that relates m_F to H . The most convenient expression of the Marčenko-Pastur equation is the one found in Silverstein (1995, Equation (1.4)):

$$\forall z \in \mathbb{C}^+ \quad m_F(z) = \int_{-\infty}^{+\infty} \frac{1}{\tau [1 - c - cz m_F(z)] - z} dH(\tau) . \quad (2.1)$$

This version of the Marčenko-Pastur equation is the one that we start out with. In addition, Silverstein and Choi (1995) showed that: $\forall \lambda \in \mathbb{R} - \{0\}$, $\lim_{z \in \mathbb{C}^+ \rightarrow \lambda} m_F(z) \equiv \check{m}_F(\lambda)$ exists, and that F has a continuous derivative $F' = \pi^{-1} \text{Im}[\check{m}_F]$ on all of \mathbb{R} with $F' \equiv 0$ on $(-\infty, 0]$.

For purposes that will become apparent later, it is useful to reformulate the Marčenko-Pastur equation.

The limiting e.d.f. of the eigenvalues of $n^{-1}Y_n'Y_n = n^{-1}\Sigma_n^{1/2}X_n'X_n\Sigma_n^{1/2}$ was defined as F . In addition, define the limiting e.d.f. of the eigenvalues of $n^{-1}Y_nY_n' = n^{-1}X_n\Sigma_nX_n'$ as \underline{F} . It then holds:

$$\begin{aligned}\forall x \in \mathbb{R} \quad \underline{F}(x) &= (1-c) \mathbf{1}_{[0,+\infty)}(x) + cF(x) \\ \forall x \in \mathbb{R} \quad F(x) &= \frac{c-1}{c} \mathbf{1}_{[0,+\infty)}(x) + \frac{1}{c} \underline{F}(x) \\ \forall z \in \mathbb{C}^+ \quad m_{\underline{F}}(z) &= \frac{c-1}{z} + c m_F(z) \\ \forall z \in \mathbb{C}^+ \quad m_F(z) &= \frac{1-c}{cz} + \frac{1}{c} m_{\underline{F}}(z).\end{aligned}$$

With this notation, Equation (1.3) of Silverstein and Choi (1995) rewrites the Marčenko-Pastur equation as: for each $z \in \mathbb{C}^+$, $m_{\underline{F}}(z)$ is the unique solution in \mathbb{C}^+ to the equation:

$$m_{\underline{F}}(z) = - \left[z - c \int_{-\infty}^{+\infty} \frac{\tau}{1 + \tau m_{\underline{F}}(z)} dH(\tau) \right]^{-1}. \quad (2.2)$$

Now introduce $u_{\underline{F}}(z) \equiv -1/m_{\underline{F}}(z)$. Notice that: $u_{\underline{F}}(z) \in \mathbb{C}^+ \iff m_{\underline{F}}(z) \in \mathbb{C}^+$. The mapping from $u_{\underline{F}}(z)$ to $m_{\underline{F}}(z)$ is one-to-one on \mathbb{C}^+ .

With this change of variable, Equation (2.2) is equivalent to saying that: for each $z \in \mathbb{C}^+$, $u_{\underline{F}}(z)$ is the unique solution in \mathbb{C}^+ to the equation:

$$u_{\underline{F}}(z) = z + c u_{\underline{F}}(z) \int_{-\infty}^{+\infty} \frac{\tau}{\tau - u_{\underline{F}}(z)} dH(\tau). \quad (2.3)$$

Let the linear operator L transform any c.d.f. G into:

$$LG(x) \equiv \int_{-\infty}^x \tau dG(\tau).$$

Combining L with the Stieltjes transform, we get the following:

$$m_{LG}(z) = \int_{-\infty}^{+\infty} \frac{\tau}{\tau - z} dG(\tau) = 1 + z m_G(z).$$

Thus, we can rewrite Equation (2.3) more concisely as:

$$u_{\underline{F}}(z) = z + c u_{\underline{F}}(z) m_{LH}(u_{\underline{F}}(z)). \quad (2.4)$$

As Silverstein and Choi (1995, Equation (1.4)) explain, the function defined in Equation (2.2) is invertible. Thus, we can define the inverse function:

$$z_{\underline{F}}(m) \equiv -\frac{1}{m} + c \int_{-\infty}^{+\infty} \frac{\tau}{1 + \tau m} dH(\tau). \quad (2.5)$$

We can do the same thing for Equation (2.4) and define the inverse function:

$$\tilde{z}_F(u) \equiv u - c u m_{LH}(u). \quad (2.6)$$

Equations (2.1), (2.2), (2.4), (2.5), and (2.6) are all completely equivalent to one another: solving any one of them means having solved them all. They are all just reformulations of the Marčenko-Pastur equation.

As will be detailed in Section 3, the oracle nonlinear shrinkage estimator of Σ_n involves the quantity $\check{m}_F(\lambda)$, for various inputs λ . The following section describes how this quantity can be found in the hypothetical case that F and H are actually known. This will then allow us later to discuss consistent estimation of $\check{m}_F(\lambda)$ in the realistic case when F and H are unknown.

2.3 Solving the Marčenko-Pastur Equation

Silverstein and Choi (1995) explain how the support of F , denoted by $\text{Supp}(F)$ is determined. Let $B \equiv \{u \in \mathbb{R} : u \neq 0, u \in \text{Supp}^b(H)\}$. Then plot the function $\tilde{z}_F(u)$ of (2.6) on the set B . Find the extreme values on each interval. Delete these points and everything in between on the real line. Do this for all increasing intervals. What is left is just $\text{Supp}(F)$; see Figure 1 of Bai and Silverstein (1998) for an illustration.

To simplify, we will assume from here on that $\text{Supp}(F)$ is a single compact interval, bounded away from zero, with $F' > 0$ in the interior of this interval. But if $\text{Supp}(F)$ is the union of a finite number of such intervals, the arguments presented in this section as well as in the remainder of the paper apply separately to each interval. In particular, our consistency results presented in subsequent sections can be easily extended to this more general case. On the other hand, the even more general case of $\text{Supp}(F)$ being the union of an infinite number of such intervals or being a non-compact interval is ruled out by Assumption (A4). By our assumption then, $\text{Supp}(F)$ is given by the compact interval $[\tilde{z}_F(u_1), \tilde{z}_F(u_2)]$ for some $u_1 < u_2$. To keep the notation shorter in what follows, let $\tilde{z}_1 \equiv \tilde{z}_F(u_1)$ and $\tilde{z}_2 \equiv \tilde{z}_F(u_2)$.

We know that for every λ in the interior of $\text{Supp}(F)$, there exists a unique $v \in \mathbb{C}^+$, denoted by v_λ , such that:

$$v_\lambda - c v_\lambda m_{LH}(v_\lambda) = \lambda. \quad (2.7)$$

We further know that:

$$F'(\lambda) = \frac{1}{c} \underline{F}'(\lambda) = \frac{1}{c\pi} \text{Im}[\check{m}_F(\lambda)] = \frac{1}{c\pi} \text{Im} \left[-\frac{1}{v_\lambda} \right].$$

The converse is also true. Since $\text{Supp}(F) = [\tilde{z}_F(u_1), \tilde{z}_F(u_2)]$, for every $x \in (u_1, u_2)$, there exists a unique $y > 0$, denoted by y_x , such that:

$$(x + iy_x) - c(x + iy_x) m_{LH}(x + iy_x) \in \mathbb{R}.$$

In other words, y_x is the unique value of $y > 0$ for which $\text{Im}[(x+iy) - c(x+iy)m_{LH}(x+iy)] = 0$. Also, if λ_x denotes the value of λ for which we have $(x+iy_x) - c(x+iy_x)m_{LH}(x+iy_x) = \lambda$, then, by definition, $z_{\lambda_x} = x + iy_x$.

Once we find a way to consistently estimate y_x for any $x \in [u_1, u_2]$, then we have an estimate of the (asymptotic) solution to the Marčenko-Pastur equation. For example: $\text{Im}[-1/(x+iy_x)]/(c\pi)$ is the value of the density F' evaluated at $\text{Re}[(x+iy_x) - c(x+iy_x)m_{LH}(x+iy_x)] = (x+iy_x) - c(x+iy_x)m_{LH}(x+iy_x)$.

From the above arguments, it follows that:

$$\forall \lambda \in (\tilde{z}_1, \tilde{z}_2) \quad \check{m}_{\underline{F}}(\lambda) = -\frac{1}{v_\lambda} \quad \text{and so} \quad \check{m}_F(\lambda) = \frac{1-c}{c\lambda} - \frac{1}{c} \frac{1}{v_\lambda} . \quad (2.8)$$

3 Oracle Estimator

3.1 Covariance Matrix

In the absence of specific information about the true covariance matrix Σ_n , it appears reasonable to restrict attention to the class of estimators which are equivariant with respect to rotations of the observed data. To be more specific, let W be an arbitrary p -dimensional orthogonal matrix. Let $\hat{\Sigma}_n \equiv \hat{\Sigma}_n(Y_n)$ be an estimator of Σ_n . Then the estimator is said to be *rotation-equivariant* if it satisfies $\hat{\Sigma}_n(Y_n W) = \hat{\Sigma}_n(Y_n)W$. In other words, the estimate based on the rotated data equals the rotation of the estimate based on the original data. The class of rotation-equivariant estimators of the covariance matrix is constituted of all the estimators that have the same eigenvectors as the sample covariance matrix; see Perlman (2007, Section 5.4). Every rotation-equivariant estimator is thus of the form:

$$U_n D_n U_n' \quad \text{where} \quad D_n \equiv \text{Diag}(d_1, \dots, d_p) \text{ is diagonal,}$$

and where U_n is the matrix whose i^{th} column is the sample eigenvector $u_i \equiv u_{n,i}$. This is the class we consider.

The starting objective is to find the matrix in this class that is closest to Σ_n . In order to measure distance, we choose the Frobenius norm defined as:

$$\|A\| \equiv \sqrt{\text{Tr}(AA')/r} \quad \text{for any matrix } A \text{ of dimension } r \times m .$$

(Dividing by the dimension of the square matrix AA' inside the root is not standard, but we do this for asymptotic purposes so that the Frobenius norm remains constant equal to one for the identity matrix regardless of the dimension; see Ledoit and Wolf (2004).) As a result, we end up with the following minimization problem:

$$\min_{D_n} \|U_n D_n U_n' - \Sigma_n\| .$$

Elementary matrix algebra shows that its solution is:

$$D_n^* \equiv \text{Diag}(d_1^*, \dots, d_p^*) \quad \text{where} \quad d_i^* \equiv u_i' \Sigma_n u_i \quad \text{for } i = 1, \dots, p. \quad (3.1)$$

The interpretation of d_i^* is that it captures how the i^{th} sample eigenvector u_i relates to the population covariance matrix Σ_n as a whole. As a result, the finite-sample optimal estimator is given by:

$$S_n^* \equiv U_n D_n^* U_n' \quad \text{where} \quad D_n^* \text{ is defined as in (3.1) .} \quad (3.2)$$

By generalizing the Marčenko-Pastur equation (2.1), Ledoit and Pécché (2011) show that d_i^* can be approximated by the quantity:

$$d_i^{or} \equiv \frac{\lambda_i}{|1 - c - c \lambda_i \check{m}_F(\lambda_i)|^2} \quad \text{for } i = 1, \dots, p, \quad (3.3)$$

from which they deduce their oracle estimator:

$$S_n^{or} \equiv U_n D_n^{or} U_n' \quad \text{where} \quad D_n^{or} \equiv \text{Diag}(d_1^{or}, \dots, d_n^{or}). \quad (3.4)$$

The key difference between D_n^* and D_n^{or} is that the former depends on the unobservable population covariance matrix, whereas the latter depends on the limiting distribution of sample eigenvalues, which makes it amenable to estimation, as explained below.

Note that S_n^{or} constitutes a nonlinear shrinkage estimator: since the value of the denominator of d_i^{or} varies with λ_i , the shrunk eigenvalues d_i^{or} are obtained by applying a nonlinear transformation to the sample eigenvalues λ_i ; see Figure 2 for an illustration. Ledoit and P  ch   (2011) also illustrate in some (limited) simulations that this oracle estimator can provide a magnitude of improvement over the linear shrinkage estimator of Ledoit and Wolf (2004).

Of course, the estimator (3.4) is not available in practice, since $\check{m}_F(\lambda)$ is an unknown population quantity. However, being able to consistently estimate $\check{m}_F(\lambda)$, uniformly in λ , will allow us to construct a *bona fide* estimator which converges to the oracle estimator almost surely (in the sense that the Frobenius norm of the difference between the two estimators converges to zero almost surely). This is the topic of Sections 4 and 5. Needless to say, the constant c is also not known in practice, but a trivial consistent estimator is given by $\hat{c}_n \equiv p/n$.

3.2 Precision Matrix

Often times an estimator of the inverse of the covariance matrix, or the precision matrix, Σ_n^{-1} is required. A reasonable strategy would be to first estimate Σ_n and to then simply take the inverse of the resulting estimate. However, such a strategy will generally not be optimal.

By arguments analogous to those leading up to (3.2), among the class of rotation-equivariant estimators, the finite-sample optimal estimator of Σ_n^{-1} with respect to the Frobenius norm is given by:

$$P_n^* \equiv U_n A_n^* U_n' \quad \text{where} \quad a_i^* \equiv u_i' \Sigma_n^{-1} u_i \quad \text{for } i = 1, \dots, p. \quad (3.5)$$

In particular, note that $P_n^* \neq (S_n^*)^{-1}$ in general.

Studying the asymptotic behavior of the diagonal matrix A_n^* led Ledoit and P  ch   (2011) to the following oracle estimator:

$$P_n^{or} \equiv U_n A_n^{or} U_n' \quad \text{where} \quad a_i^{or} \equiv \lambda_i^{-1} (1 - c - 2c\lambda_i \operatorname{Re}[\check{m}_F(\lambda_i)]) \quad \text{for } i = 1, \dots, p. \quad (3.6)$$

In particular, note that $P_n^{or} \neq (S_n^{or})^{-1}$ in general.

It is seen that both oracle estimators S_n^{or} and P_n^{or} involve the unknown quantities $\check{m}_F(\lambda_i)$, for $i = 1, \dots, p$. The next section explains how to construct a uniformly consistent estimator of $\check{m}_F(\lambda)$, for arbitrary λ , based on a consistent estimator of H , the limiting spectral distribution of the population eigenvalues. In the section thereafter, we will discuss how to construct a consistent estimator of H from the observed data.

4 Estimation of $\check{m}_F(\lambda)$

Fix $x \in [u_1 + \eta, u_2 - \eta]$, where $\eta > 0$ is some small number. From the previous discussion in Section 2, it follows that the equation

$$\text{Im}[x + iy - c(x + iy) m_{LH}(x + iy)] = 0$$

has a unique solution $y \in (0, +\infty)$, called y_x . Since $u_1 < x < u_2$, it follows that $y_x > 0$; for $x = u_1$ or $x = u_2$, we would have $y_x = 0$ instead. The goal is to consistently estimate y_x , uniformly in $x \in [u_1 + \eta, u_2 - \eta]$.

Define for any c.d.f. G and for any $d > 0$, the real function

$$g_{G,d}(y,x) \equiv \left| \text{Im} [x + iy - d(x + iy) m_{LG}(x + iy)] \right|.$$

With this notation, y_x is the unique minimizer in $(0, +\infty)$ of $g_{H,c}(y, x)$ then. In particular, $g_{H,c}(y_x, x) = 0$.

In the remainder of the paper, the symbol \Rightarrow denotes weak convergence (or convergence in distribution).

Proposition 4.1.

- (i) Let $\{\widehat{H}_n\}$ be a sequence of probability measures with $\widehat{H}_n \Rightarrow H$. Let $\{\widehat{c}_n\}$ be a sequence of positive real numbers with $\widehat{c}_n \rightarrow c$. Let $K \subseteq (0, \infty)$ be a compact interval satisfying $\{y_x : x \in [u_1 + \eta, u_2 - \eta]\} \subseteq K$. For a given $x \in [u_1 + \eta, u_2 - \eta]$, let $\widehat{y}_{n,x} \equiv \min_{y \in K} g_{\widehat{H}_n, \widehat{c}_n}(y, x)$. It then holds that $\widehat{y}_{n,x} \rightarrow y_x$ uniformly in $x \in [u_1 + \eta, u_2 - \eta]$.
- (ii) In case of $\widehat{H}_n \Rightarrow H$ a.s., it holds that $\widehat{y}_{n,x} \rightarrow y_x$ a.s. uniformly in $x \in [u_1 + \eta, u_2 - \eta]$.

It should be pointed out that the assumption $\{y_x : x \in [u_1 + \eta, u_2 - \eta]\} \subseteq K$ is not really restrictive, since one can choose $K \equiv [\varepsilon, 1/\varepsilon]$ for ε arbitrarily small.

We also need to solve the ‘inverse’ estimation problem, namely starting with λ and recovering the corresponding v_λ . Fix $\lambda \in [\tilde{z}_1 + \tilde{\delta}, \tilde{z}_2 - \tilde{\delta}]$, where $\tilde{\delta} > 0$ is some small number. From the previous discussion, it follows that the equation

$$v - c v m_{LH}(v) = \lambda$$

has a unique solution $v \in \mathbb{C}^+$, called v_λ . The goal is to consistently estimate v_λ , uniformly in $\lambda \in [\tilde{z}_1 + \tilde{\delta}, \tilde{z}_2 - \tilde{\delta}]$.

Define for any c.d.f. G and for any $d > 0$, the real function

$$h_{G,d}(v, \lambda) \equiv |v - d v m_{LG}(v) - \lambda|.$$

With this notation, v_λ is the unique minimizer in \mathbb{C}^+ of $h_{H,c}(v, \lambda)$ then. In particular, $h_{H,c}(v_\lambda, \lambda) = 0$.

Proposition 4.2.

(i) Let $\{\hat{H}_n\}$ be a sequence of probability measures with $\hat{H}_n \Rightarrow H$. Let $\{\hat{c}_n\}$ be a sequence of positive real numbers with $\hat{c}_n \rightarrow c$. Let $K \subseteq \mathbb{C}^+$ be a compact set satisfying $\{v_\lambda : \lambda \in [\tilde{z}_1 + \tilde{\delta}, \tilde{z}_2 - \tilde{\delta}]\} \subseteq K$. For a given $\lambda \in [\tilde{z}_1 + \tilde{\delta}, \tilde{z}_2 - \tilde{\delta}]$, let $\hat{v}_{n,\lambda} \equiv \min_{v \in K} h_{\hat{H}_n, \hat{c}_n}(v, \lambda)$. It then holds that $\hat{v}_{n,\lambda} \rightarrow v_\lambda$ uniformly in $\lambda \in [\tilde{z}_1 + \tilde{\delta}, \tilde{z}_2 - \tilde{\delta}]$.

(ii) In case of $\hat{H}_n \Rightarrow H$ a.s., it holds that $\hat{v}_{n,\lambda} \rightarrow v_\lambda$ a.s. uniformly in $\lambda \in [\tilde{z}_1 + \tilde{\delta}, \tilde{z}_2 - \tilde{\delta}]$.

Being able to find consistent estimators of v_λ , uniformly in λ , now allows us to find consistent estimators of $\check{m}_F(\lambda)$, uniformly in λ , based on (2.8). Our estimator of $\check{m}_F(\lambda)$ is given by:

$$\check{m}_{F_{\hat{H}_n, \hat{c}_n}}(\lambda) \equiv \frac{1 - \hat{c}_n}{\hat{c}_n \lambda} - \frac{1}{\hat{c}_n \hat{v}_{n,\lambda}}. \quad (4.1)$$

This, in return, provides us with a consistent estimator of S_n^{or} , the oracle nonlinear shrinkage estimator of Σ_n . Define:

$$\hat{S}_n \equiv U_n \hat{D}_n U_n' \quad \text{where} \quad \hat{d}_i \equiv \frac{\lambda_i}{|1 - \hat{c}_n - \hat{c}_n \lambda_i \check{m}_{F_{\hat{H}_n, \hat{c}_n}}(\lambda_i)|^2} \quad \text{for } i = 1, \dots, p. \quad (4.2)$$

It also provides us with a consistent estimator of P_n^{or} , the oracle nonlinear shrinkage estimator of Σ_n^{-1} . Define:

$$\hat{P}_n \equiv U_n \hat{A}_n U_n' \quad \text{where} \quad \hat{a}_i \equiv \lambda_i^{-1} (1 - \hat{c}_n - 2 \hat{c}_n \lambda_i \operatorname{Re}[\check{m}_{F_{\hat{H}_n, \hat{c}_n}}(\lambda_i)]) \quad \text{for } i = 1, \dots, p. \quad (4.3)$$

In particular, note that $\hat{P}_n \neq \hat{S}_n^{-1}$ in general.

Proposition 4.3.

(i) Let $\{\hat{H}_n\}$ be a sequence of probability measures with $\hat{H}_n \Rightarrow H$. Let $\{\hat{c}_n\}$ be a sequence of positive real numbers with $\hat{c}_n \rightarrow c$. It then holds that:

- (a) $\check{m}_{F_{\hat{H}_n, \hat{c}_n}}(\lambda) \rightarrow \check{m}_F(\lambda)$ uniformly in $\lambda \in [\tilde{z}_1 + \tilde{\delta}, \tilde{z}_2 - \tilde{\delta}]$
- (b) $\|\hat{S}_n - S_n^{or}\| \rightarrow 0$
- (c) $\|\hat{P}_n - P_n^{or}\| \rightarrow 0$

(ii) In case of $\hat{H}_n \Rightarrow H$ a.s., it holds that:

- (a) $\check{m}_{F_{\hat{H}_n, \hat{c}_n}}(\lambda) \rightarrow \check{m}_F(\lambda)$ uniformly in $\lambda \in [\tilde{z}_1 + \tilde{\delta}, \tilde{z}_2 - \tilde{\delta}]$ a.s.
- (b) $\|\hat{S}_n - S_n^{or}\| \rightarrow 0$ a.s.
- (c) $\|\hat{P}_n - P_n^{or}\| \rightarrow 0$ a.s.

5 Estimation of H

As described before, consistent estimation of the oracle estimators of Ledoit and P  ch   (2011) requires (uniformly) consistent estimation of $\check{m}_F(\lambda)$. Since $\text{Im}[\check{m}_F(\lambda)] = \pi F'(\lambda)$, one possible approach could be to take an off-the-shelf density estimator for F' , based on the observed sample eigenvalues λ_i . There exists a large literature on density estimation; for example, see Silverman (1986). The real part of $\check{m}_F(\lambda_i)$ could be estimated in a similar manner.

However, the sample eigenvalues do not satisfy any of the regularity conditions usually invoked for the underlying data. It really is not clear at all whether an off-the-shelf density estimator applied to the sample eigenvalues would result in consistent estimation of F' .

Even if this issue was somehow resolved, using such a generic procedure would not exploit the specific features of the problem. Namely: F is not just any distribution, it is a distribution of sample eigenvalues. It is the solution to the Mar  chenko-Pastur equation for some H . This is valuable information that narrows down considerably the set of possible distributions F . Therefore an estimation procedure specifically designed to incorporate this *a priori* knowledge would be better suited to the problem at hand. This is the approach we select.

In a nutshell: our estimator of F is the c.d.f. that is closest to F_n among the c.d.f.s that are a solution to the Mar  chenko-Pastur equation for some \tilde{H} and for $\tilde{c} \equiv p/n$. The ‘underlying’ distribution \tilde{H} which produces the thus obtained estimator of F is, in return, our estimator of H . If we can show that this estimator of H is consistent, then the results of the previous section demonstrate that the implied estimator of $\check{m}_F(\lambda)$ is uniformly consistent.

Section 5.1 derives theoretical properties of this approach, while Section 5.2 discusses various issues concerning the practical implementation.

5.1 Consistency Results

For a grid of real numbers $Q \equiv \{\dots, t_{-1}, t_0, t_1, \dots\} \subseteq \mathbb{R}$, with $t_{k-1} < t_k$, define the corresponding grid size γ as

$$\gamma \equiv \sup_k (t_k - t_{k-1}) .$$

A grid Q is said to cover a compact interval $[a, b] \subseteq \mathbb{R}$ if there exists at least one $t_k \in Q$ with $t_k \leq a$ and at least another $t_{k'} \in Q$ with $b \leq t_{k'}$. A sequence of grids $\{Q_n\}$ is said to eventually cover a compact interval $[a, b]$ if for every $\phi > 0$ there exist $N \equiv N(\phi)$ such that Q_n covers the compact interval $[a + \phi, b - \phi]$ for all $n \geq N$.

For any probability measure \tilde{H} on the real line and for any $\tilde{c} > 0$, let $F_{\tilde{H}, \tilde{c}}$ denote the c.d.f. on the real line induced by the corresponding solution of the Marčenko-Pastur equation. More specifically, for each $z \in \mathbb{C}^+$, $m_{F_{\tilde{H}, \tilde{c}}}(z)$ is the unique solution for $m \in \mathbb{C}^+$ to the equation:

$$m = \int_{-\infty}^{+\infty} \frac{1}{\tau [1 - \tilde{c} - \tilde{c} z m] - z} d\tilde{H}(\tau) .$$

In this notation, we then have $F = F_{H, c}$.

It follows from Silverstein and Choi (1995) again that: $\forall \lambda \in \mathbb{R} - \{0\}$, $\lim_{z \in \mathbb{C}^+ \rightarrow \lambda} m_{F_{\tilde{H}, \tilde{c}}}(z) \equiv \check{m}_{F_{\tilde{H}, \tilde{c}}}(\lambda)$ exists, and that $F_{\tilde{H}, \tilde{c}}$ has a continuous derivative $F'_{\tilde{H}, \tilde{c}} = \pi^{-1} \text{Im} [\check{m}_{F_{\tilde{H}, \tilde{c}}}]$ on $(0, +\infty)$. In the case $\tilde{c} < 1$, $F_{\tilde{H}, \tilde{c}}$ has a continuous derivative on all of \mathbb{R} with $F'_{\tilde{H}, \tilde{c}} \equiv 0$ on $(-\infty, 0]$.

For a grid Q on the real line and for two c.d.f.s G_1 and G_2 , define:

$$\|G_1 - G_2\|_Q \equiv \sup_{t \in Q} |G_1(t) - G_2(t)|$$

The following theorem shows that both F and H can be estimated consistently via an idealized algorithm.

Theorem 5.1. *Let $\{Q_n\}$ be a sequence of grids on the real line eventually covering the support of F with corresponding grid sizes $\{\gamma_n\}$ satisfying $\gamma_n \rightarrow 0$. Let $\{\hat{c}_n\}$ be a sequence of positive real numbers with $\hat{c}_n \rightarrow c$. Let \hat{H}_n be defined as:*

$$\hat{H}_n \equiv \underset{\tilde{H}}{\text{argmin}} \|F_{\tilde{H}, \hat{c}_n} - F_n\|_{Q_n} , \quad (5.1)$$

where \tilde{H} is a probability measure.

Then we have: (i) $F_{\hat{H}_n, \hat{c}_n} \Rightarrow F$ a.s.; and (ii) $\hat{H}_n \Rightarrow H$ a.s.

The algorithm used in the theorem is not practical for two reasons. First, it is not possible to optimize over all probability measures \tilde{H} . But similar to El Karoui (2008), we can show that it is sufficient to optimize over all probability measures which are sums of atoms, the location of which is restricted to a fixed-size grid, with the grid size vanishing asymptotically.

Corollary 5.1. *Let $\{Q_n\}$ be a sequence of grids on the real line eventually covering the support of F with corresponding grid sizes $\{\gamma_n\}$ satisfying $\gamma_n \rightarrow 0$. Let $\{\hat{c}_n\}$ be a sequence of positive real numbers with $\hat{c}_n \rightarrow c$. Let \mathcal{P}_n denote the set of all probability measures which are sums of atoms belonging to the grid $\{J_n/T_n, (J_n+1)/T_n, \dots, K_n/T_n\}$ with $T_n \rightarrow \infty$, J_n being the largest integer satisfying $J_n/T_n \leq \lambda_1$, and K_n being the smallest integer satisfying $K_n/T_n \geq \lambda_p$. Let \hat{H}_n be defined as:*

$$\hat{H}_n \equiv \underset{\tilde{H} \in \mathcal{P}_n}{\text{argmin}} \|F_{\tilde{H}, \hat{c}_n} - F_n\|_{Q_n} , \quad (5.2)$$

Then we have: (i) $F_{\hat{H}_n, \hat{c}_n} \Rightarrow F$ a.s.; and (ii) $\hat{H}_n \Rightarrow H$ a.s.

But even restricting the optimization over a manageable set of probability measures is not quite practical yet for a second reason. Namely, to compute $F_{\tilde{H}, \tilde{c}_n}$ exactly for a given \tilde{H} , one would have to (numerically) solve the Marčenko-Pastur equation for an infinite number of points. In practice, we can only afford to solve the equation for a finite number of points and then approximate $F_{\tilde{H}, \tilde{c}_n}$ by trapezoidal integration. Fortunately, this approximation does not negatively affect the consistency of our estimators.

Let G be a c.d.f. with continuous density g and compact support $[a, b]$. For a grid $Q \equiv \{\dots, t_{-1}, t_0, t_1, \dots\}$ covering the support of G , the approximation to G via trapezoidal integration over the grid Q , denoted by \hat{G}_Q , is obtained as follows. For $t \in [a, b]$, let $J_{lo} \equiv \max\{k : t_k \leq a\}$ and $J_{hi} \equiv \min\{k : t < t_k\}$. Then:

$$\hat{G}_Q(t) \equiv \sum_{k=J_{lo}}^{J_{hi}-1} \frac{(t_{k+1} - t_k)[g(t_k) + g(t_{k+1})]}{2}. \quad (5.3)$$

Now turn to the special case $G \equiv F_{\tilde{H}, \tilde{c}}$ and $Q \equiv Q_n$. In this case, we denote the approximation to $F_{\tilde{H}, \tilde{c}}$ via trapezoidal integration over the grid Q_n by $\hat{F}_{\tilde{H}, \tilde{c}; Q_n}$.

Corollary 5.2. *Assume the same assumptions as in Corollary 5.1. Let \hat{H}_n be defined as:*

$$\hat{H}_n \equiv \operatorname{argmin}_{\tilde{H} \in \mathcal{P}_n} \|\hat{F}_{\tilde{H}, \tilde{c}; Q_n} - F_n\|_{Q_n}, \quad (5.4)$$

Let $\check{m}_{F_{\hat{H}_n, \hat{c}_n}}(\lambda)$, \hat{S}_n , and \hat{P}_n be defined as in (4.1), (4.2), and (4.3), respectively. Then:

$$(i) \quad F_{\hat{H}_n, \hat{c}_n} \Rightarrow F \text{ a.s.}$$

$$(ii) \quad \hat{H}_n \Rightarrow H \text{ a.s.}$$

$$(iii) \quad \text{For any } \tilde{\delta} > 0, \check{m}_{F_{\hat{H}_n, \hat{c}_n}}(\lambda) \rightarrow \check{m}_F(\lambda) \text{ a.s. uniformly in } \lambda \in [\tilde{z}_1 - \tilde{\delta}, \tilde{z}_2 + \tilde{\delta}].$$

$$(iv) \quad \|\hat{S}_n - S_n^{or}\| \rightarrow 0 \text{ a.s.}$$

$$(v) \quad \|\hat{P}_n - P_n^{or}\| \rightarrow 0 \text{ a.s.}$$

5.2 Implementation Details

Decomposition of the c.d.f. of Population Eigenvalues As discussed before, it is not practical to search over the set of all possible c.d.f.s \tilde{H} . Following El Karoui (2008), we project H onto a certain basis of c.d.f.s $(M_k)_{k=1, \dots, K}$, where K goes to infinity along with n and p . The projection of H onto this basis is given by the nonnegative weights w_1, \dots, w_K , where:

$$\forall t \in \mathbb{R} \quad H(t) \approx \tilde{H}(t) \equiv \sum_{k=1}^K w_k M_k(t) \quad \text{and} \quad \sum_{k=1}^K w_k = 1. \quad (5.5)$$

Thus, our estimator for F will be a solution to the Marčenko-Pastur equation for \tilde{H} given by Equation (5.5) for some $(w_k)_{k=1,\dots,K}$, and for $\tilde{c} \equiv p/n$. It is just a matter of searching over all sets of nonnegative weights summing up to one.

Choice of Basis We base the c.d.f.s $(M_k)_{k=1,\dots,K}$ on a grid of p equally spaced points on the interval $[\lambda_1, \lambda_p]$:

$$x_i \equiv \lambda_1 + \frac{i-1}{p}(\lambda_p - \lambda_1) \quad \text{for } i = 1, \dots, p. \quad (5.6)$$

Thus, $x_1 = \lambda_1$ and $x_p = \lambda_p$. We then form the basis $\{M_1, \dots, M_K\}$ as the union of three families of c.d.f.s:

1. the indicator functions $\mathbf{1}_{[x_i, +\infty)}$ ($i = 1, \dots, p$);
2. the c.d.f.s whose derivatives are linearly increasing on the interval $[x_{i-1}, x_i]$ and zero everywhere else ($i = 2, \dots, p$);
3. the c.d.f.s whose derivatives are linearly decreasing on the interval $[x_{i-1}, x_i]$ and zero everywhere else ($i = 2, \dots, p$).

This list yields a basis $(M_k)_{k=1,\dots,K}$ of dimension $K = 3p - 2$. Notice that by the theoretical results of Section 5.1, it would be sufficient to use the first family only. Including the second and third families in addition cannot make the approximation to H any worse.

Trapezoidal Integration For a given $\tilde{H} \equiv \sum_{k=1}^K w_k M_k$, it is computationally too expensive (in the context of an optimization procedure) to solve the Marčenko-Pastur equation for $m_F(z)$ over all $z \in \mathbb{C}^+$. It is more efficient to solve the Marčenko-Pastur equation only for $\check{m}_F(x_i)$ ($i = 1, \dots, p$), and to use the trapezoidal approximation formula to deduce from it $F(x_i)$ ($i = 1, \dots, p$). The trapezoidal rule gives:

$$\begin{aligned} \forall i = 1, \dots, p \quad F(x_i) &= \sum_{j=1}^{i-1} \frac{x_{j+1} - x_{j-1}}{2} F'(x_j) + \frac{x_i - x_{i-1}}{2} F'(x_i) \\ &= \sum_{j=1}^{i-1} \frac{(x_{j+1} - x_{j-1}) \operatorname{Im} [\check{m}_F(x_j)]}{2\pi} + \frac{(x_i - x_{i-1}) \operatorname{Im} [\check{m}_F(x_i)]}{2\pi}, \end{aligned} \quad (5.7)$$

with the convention $x_0 \equiv 0$.

Objective Function The objective function measures the distance between F_n and the F that solves the Marčenko-Pastur equation for $\tilde{H} \equiv \sum_{k=1}^K w_k M_k$ and for $\tilde{c} \equiv p/n$. Traditionally, F_n is defined as càdlàg, that is : $F_n(\lambda_1) = 1/p$ and $F_n(\lambda_p) = 1$. However, there is a certain degree of arbitrariness in this convention: why is $F_n(\lambda_p)$ equal to one but $F_n(\lambda_1)$ not equal to zero? By symmetry, there is no *a priori* justification for specifying that the largest eigenvalue

is closer to the supremum of the support of F than the smallest to its infimum. Therefore, a different convention might be more appropriate in this case, which leads us to the following definition:

$$\forall i = 1, \dots, p \quad \widehat{F}_n(\lambda_i) \equiv \frac{i}{p} - \frac{1}{2p}. \quad (5.8)$$

This choice restores a certain element of symmetry to the treatment of the smallest vs. the largest eigenvalue. From Equation (5.8), we deduce $\widehat{F}_n(x_i)$, for $i = 2, \dots, p-1$, by linear interpolation. With a sup-norm error penalty, this leads to the following objective function:

$$\max_{i=1, \dots, p} \left| F(x_i) - \widehat{F}_n(x_i) \right|, \quad (5.9)$$

where $F(x_i)$ is given by Equation (5.7) for $i = 1, \dots, p$. Using Equation (5.7), we can rewrite this objective function as:

$$\max_{i=1, \dots, p} \left| \sum_{j=1}^{i-1} \frac{(x_{j+1} - x_{j-1}) \operatorname{Im} [\check{m}_F(x_j)]}{2\pi} + \frac{(x_i - x_{i-1}) \operatorname{Im} [\check{m}_F(x_i)]}{2\pi} - \widehat{F}_n(x_i) \right|.$$

Optimization Program We now have all the ingredients needed to state the optimization program that will extract the estimator of $\check{m}_F(x_1), \dots, \check{m}_F(x_p)$ from the observations $\lambda_1, \dots, \lambda_p$. It is the following:

$$\min_{\substack{m_1, \dots, m_p \\ w_1, \dots, w_K}} \max_{i=1, \dots, p} \left| \sum_{j=1}^{i-1} \frac{(x_{j+1} - x_{j-1}) \operatorname{Im} [m_j]}{2\pi} + \frac{(x_i - x_{i-1}) \operatorname{Im} [m_i]}{2\pi} - \widehat{F}_n(x_i) \right|$$

subject to:

$$\forall j = 1, \dots, p \quad m_j = \sum_{k=1}^K \int_{-\infty}^{+\infty} \frac{w_k}{t [1 - (p/n) - (p/n) x_j m_j] - x_j} dM_k(t) \quad (5.10)$$

$$\sum_{k=1}^K w_k = 1$$

$$\forall j = 1, \dots, p \quad m_j \in \mathbb{C}^+$$

$$\forall k = 1, \dots, K \quad w_k \geq 0.$$

The key is to introduce the variables $m_j \equiv \check{m}_F(x_j)$, for $j = 1, \dots, p$. The constraint in Equation (5.10) imposes that m_j is the solution to the Marčenko-Pastur equation evaluated as $z \in \mathbb{C}^+ \rightarrow x_j$ when $\widetilde{H} = \sum_{k=1}^K w_k M_k$.

Real Optimization Program In practice, most optimizers only accept real variables, therefore it is necessary to decompose m_j into its real and imaginary parts: $a_j \equiv \operatorname{Re}[m_j]$ and $b_j \equiv \operatorname{Im}[m_j]$. Then we can optimize separately over the two sets of real variables a_j and b_j for $j = 1, \dots, p$. The Marčenko-Pastur constraint in Equation (5.10) splits into two constraints:

one for the real part and the other for the imaginary part. The reformulated optimization program is:

$$\min_{\substack{a_1, \dots, a_p \\ b_1, \dots, b_p \\ w_1, \dots, w_K}} \max_{i=1, \dots, p} \left| \sum_{j=1}^{i-1} \frac{(x_{j+1} - x_{j-1}) b_j}{2\pi} + \frac{(x_i - x_{i-1}) b_i}{2\pi} - \widehat{F}_n(x_i) \right| \quad (5.11)$$

subject to:

$$\forall j = 1, \dots, p \quad a_j = \sum_{k=1}^K \int_{-\infty}^{+\infty} \operatorname{Re} \left\{ \frac{w_k}{t [1 - (p/n) - (p/n) x_j (a_j + ib_j)] - x_j} \right\} dM_k(t) \quad (5.12)$$

$$\forall j = 1, \dots, p \quad b_j = \sum_{k=1}^K \int_{-\infty}^{+\infty} \operatorname{Im} \left\{ \frac{w_k}{t [1 - (p/n) - (p/n) x_j (a_j + ib_j)] - x_j} \right\} dM_k(t) \quad (5.13)$$

$$\sum_{k=1}^K w_k = 1 \quad (5.14)$$

$$\forall j = 1, \dots, p \quad b_j \geq 0 \quad (5.15)$$

$$\forall k = 1, \dots, K \quad w_k \geq 0. \quad (5.16)$$

Remark 5.1. Since the theory of Sections 4 and 5.1 partly assumes that m_j belongs to a compact set in \mathbb{C}^+ bounded away from the real line, we might want to add to the real optimization program the constraints that $-1/\varepsilon \leq a_j \leq 1/\varepsilon$ and that $\varepsilon \leq b_j \leq 1/\varepsilon$, for some small $\varepsilon > 0$. Our simulations indicate that for a small value of ε such as $\varepsilon = 10^{-6}$, this makes no difference in practice.

Sequential Linear Programming While the optimization program defined in Equations (5.11)–(5.16) may appear daunting at first sight, it is in fact solved quickly and efficiently by off-the-shelf optimization software implementing Sequential Linear Programming (SLP). The key is to linearize Equations (5.12)–(5.13), the two constraints that embody the Marčenko-Pastur equation, around an approximate solution point. Once they are linearized, the optimization program (5.11)–(5.16) becomes a standard Linear Programming (LP) problem, which can be solved very quickly. Then we linearize again Equations (5.12)–(5.13) around the new point, and this generates a new LP problem; hence the name: *Sequential* Linear Programming. The software iterates until a satisfactory degree of convergence is achieved. All of this is handled automatically by the SLP optimizer. The user only needs to specify the problem (5.11)–(5.16), as well as an adequate starting point, and then launch the SLP optimizer. For our SLP optimizer, we selected a standard off-the-shelf commercial software: SNOPTTM Version 7.2-5; see Gill et al. (2002). While SNOPTTM was originally designed for Sequential Quadratic Programming, it also handles SLP, since Linear Programming can be viewed as a particular case of Quadratic Programming with no quadratic term. On average, it converges in only about five seconds on a desktop Mac for a problem with $p = 100$ variables.

Starting Point A neutral way to choose the starting point is to place equal weights on all the c.d.f.s in our basis: $w_k = 1/K$ ($k = 1, \dots, K$). Then it is necessary to solve the Marčenko-Pastur equation numerically once *before* launching the SLP optimizer, in order to compute the values of $\check{m}_F(x_j)$ ($j = 1, \dots, p$) that correspond to this initial choice of $\tilde{H} = \sum_{k=1}^K M_k/K$. The initial values for a_j are taken to be $\text{Re}[\check{m}_F(x_j)]$, and $\text{Im}[\check{m}_F(x_j)]$ for b_j ($j = 1, \dots, p$). Any good starting point should satisfy all the constraints, in particular those corresponding to the Marčenko-Pastur equation.

Estimating the Covariance Matrix Once the SLP optimizer has converged, it generates optimal values (a_1^*, \dots, a_p^*) , (b_1^*, \dots, b_p^*) , and (w_1^*, \dots, w_K^*) . The first two sets of variables at the optimum are used to estimate the oracle shrinkage factors. From the reconstructed $\check{m}_F^*(x_j) = a_j^* + ib_j^*$, we deduce by linear interpolation $\check{m}_F^*(\lambda_j)$ for $j = 1, \dots, p$. Our estimator of the covariance matrix \hat{S}_n is built by keeping the same eigenvectors as the sample covariance matrix, and dividing each sample eigenvalue λ_j by the following correction factor:

$$\left| 1 - \frac{p}{n} - \frac{p}{n} \lambda_j \check{m}_F^*(\lambda_j) \right|^2.$$

Corollary 5.2 assures us that the resulting *bona fide* nonlinear shrinkage estimator is asymptotically equivalent to the oracle estimator S_n^{or} . Also, we can see that, as the concentration $\hat{c}_n = p/n$ gets closer to zero, that is, as we get closer to fixed-dimension asymptotics, the magnitude of the correction becomes smaller. This makes sense because under fixed-dimension asymptotics the sample covariance matrix is a consistent estimator of the population covariance matrix.

Estimating the Precision Matrix The output of the same optimization process can also be used to estimate the oracle shrinkage factors for the precision matrix. Our estimator of the precision matrix Σ_n^{-1} is built by keeping the same eigenvectors as the sample covariance matrix, and multiplying the inverse λ_j^{-1} of each sample eigenvalue by the following correction factor:

$$1 - \frac{p}{n} - 2 \frac{p}{n} \lambda_j \text{Re}[\check{m}_F^*(\lambda_j)] .$$

Corollary 5.2 assures us that the resulting *bona fide* nonlinear shrinkage estimator is asymptotically equivalent to the oracle estimator P_n^{or} .

Estimating H We point out that the optimal values (w_1^*, \dots, w_K^*) generated from the SLP optimizer yield a consistent estimate of H in the following fashion:

$$H^* \equiv \sum_{k=1}^K w_k^* M_k .$$

This estimator could be considered an alternative to the estimator introduced by El Karoui (2008). The most salient difference between the two optimization algorithms is that our objective function tries to match F_n on \mathbb{R} , whereas his objective function tries to match (a function of) m_{F_n} on \mathbb{C}^+ . However, since the estimation of H is not of prime interest to us, we will not pursue this topic further or investigate potential (dis)advantages of these two approaches.

6 Monte Carlo Simulations

In this section, we present the results of various sets of Monte Carlo simulations designed to illustrate the finite-sample properties of the nonlinear shrinkage estimator of the covariance matrix. As detailed in Section 3, the finite-sample optimal estimator in the class of rotation-equivariant estimators is given by S_n^* as defined in (3.2). Thus, the improvement of the shrinkage estimator \widehat{S}_n over the sample covariance matrix will be measured by how closely this estimator approximates S_n^* relative to the sample covariance matrix. More specifically, we report the Percentage Relative Improvement in Average Loss (PRIAL), which is defined as:

$$\text{PRIAL} \equiv \text{PRIAL}(\widehat{\Sigma}_n) \equiv 100 \times \left\{ 1 - \frac{\mathbb{E}[\|\widehat{\Sigma}_n - S_n^*\|^2]}{\mathbb{E}[\|S_n - S_n^*\|^2]} \right\} \% , \quad (6.1)$$

where $\widehat{\Sigma}_n$ is an arbitrary estimator of Σ_n . By definition, the PRIAL of S_n is 0% while the PRIAL of S_n^* is 100%.

Most of the simulations will be designed around a population covariance matrix Σ_n that has 20% of its eigenvalues equal to 1, 40% equal to 3, and 40% equal to 10. This is a particularly interesting and difficult example introduced and analyzed in great detail by Bai and Silverstein (1998). For concentration values such as $c = 1/3$ and below, it displays ‘spectral separation’, that is, the support of the distribution of sample eigenvalues is the union of three disjoint intervals, each one corresponding to a Dirac of population eigenvalues. Detecting this pattern and handling it correctly is a real challenge for any covariance matrix estimation method.

6.1 Convergence

The first set of Monte Carlo simulations shows how our nonlinear shrinkage estimator \widehat{S}_n behaves as the matrix dimension p and the sample size n go to infinity together. We assume that the concentration ratio $\widehat{c}_n = p/n$ remains constant and equal to $1/3$. For every value of p (and hence n), we run 1,000 simulations with normally distributed variables. The PRIAL is plotted in Figure 1. For the sake of comparison, we also report the PRIALs of the oracle S_n^{or} , and of the optimal linear shrinkage estimator \overline{S}_n developed by Ledoit and Wolf (2004).

We can see that the performance of the nonlinear shrinkage estimator \widehat{S}_n converges quickly towards that of the oracle and of S_n^* . Even for relatively small matrices of dimension $p = 30$, it realizes 88% of the possible gains over the sample covariance matrix. The optimal linear

shrinkage estimator \overline{S}_n performs also well relative to the sample covariance matrix, but the improvement is limited: in general, it does not converge to 100% under large-dimensional asymptotics. This is because there are strong nonlinear effects in the optimal shrinkage of sample eigenvalues. These effects are clearly visible in Figure 2, which plots a typical simulation result for $p = 100$.

We can see that the nonlinear shrinkage estimator \widehat{S}_n shrinks the eigenvalues of the sample covariance matrix almost as if it ‘knew’ the correct shape of the distribution of population eigenvalues. In particular, the various curves and gaps of the oracle nonlinear shrinkage formula are well picked up and followed by this estimator. By contrast, the linear shrinkage estimator can only use the best linear approximation to this highly nonlinear transformation. We also plot the 45-degrees line as a visual reference to show what would happen if no shrinkage was applied to the sample eigenvalues, that is, if we simply used S_n .

6.2 Concentration

The next set of Monte Carlo simulations shows how the PRIAL of the shrinkage estimators varies as a function of the concentration ratio $\widehat{c}_n = p/n$ if we keep the product $p \times n$ constant and equal to 9,000. We keep the same population covariance matrix Σ_n as in Section 6.1. For every value of p/n , we run 1,000 simulations with normally distributed variables. The respective PRIALs of S_n^{or} , \widehat{S}_n and \overline{S}_n are plotted in Figure 3.

We can see that the nonlinear shrinkage estimator performs well across the board, closely in line with the oracle, and always achieves at least 90% of the possible improvement over the sample covariance matrix. By contrast, the linear shrinkage estimator achieves relatively little improvement over the sample covariance matrix when the concentration is low. This is because, when the sample size is large relative to the matrix dimension, there is a lot of precise information about the optimal nonlinear way to shrink the sample eigenvalues that is waiting to be extracted by a suitable nonlinear procedure. By contrast, when the sample size is not so large, information about the population covariance matrix is relatively fuzzy, therefore a simple linear approximation can achieve up to 93% of the potential gains.

6.3 Dispersion

The third set of Monte Carlo simulations shows how the PRIAL of the shrinkage estimators varies as a function of the dispersion of population eigenvalues. We take a population covariance matrix Σ_n with 20% of its eigenvalues equal to 1, 40% equal to $1 + 2d/9$, and 40% equal to $1 + d$, where the dispersion parameter d varies from 0 to 20. Thus, for $d = 0$, Σ_n is the identity matrix and, for $d = 9$, Σ_n is the same matrix as in Section 6.1. The sample size is $n = 300$ and the matrix dimension is $p = 100$. For every value of d , we run 1,000 simulations with normally distributed variables. The respective PRIALs of S_n^{or} , \widehat{S}_n and \overline{S}_n are plotted in Figure 4.

We can see that the linear shrinkage estimator \overline{S}_n beats the nonlinear shrinkage estimator \widehat{S}_n for very low dispersion levels. For example, when $d = 0$, that is, when the population covariance matrix is equal to the identity matrix, \overline{S}_n realizes 99.9% of the possible improvement over the sample covariance matrix, while \widehat{S}_n realizes ‘only’ 99.4% of the possible improvement. This is because, in this case, linear shrinkage is optimal or (when d is strictly positive but still small) nearly optimal, hence there is nothing to little to be gained by resorting to a nonlinear shrinkage method. However, as dispersion increases, linear shrinkage delivers less and less improvement over the sample covariance matrix, while nonlinear shrinkage retains a PRIAL above 96% and close to that of the oracle.

6.4 Fat Tails

We also have some results on the effect of non-normality on the performance of the shrinkage estimators. We take the same population covariance matrix as in Section 6.1, that is, Σ_n has 20% of its eigenvalues equal to 1, 40% equal to 3, and 40% equal to 10. The sample size is $n = 300$, and the matrix dimension is $p = 100$. We compare two types of random variates: a Student t distribution with $\text{df} = 3$ degrees of freedom, and a Student t distribution with $\text{df} = \infty$ degrees of freedom (which is the Gaussian distribution). For each number of degrees of freedom df , we run 1,000 simulations. The respective PRIALs of S_n^{or} , \widehat{S}_n and \overline{S}_n are summarized in Table 1.

	Average Squared Frobenius Loss		PRIAL	
	df = 3	df = ∞	df = 3	df = ∞
Sample Covariance Matrix	5.856	5.837	0%	0%
Linear Shrinkage Estimator	1.883	1.883	67.84%	67.74%
Nonlinear Shrinkage Estimator	0.128	0.133	97.81%	97.71%
Oracle	0.043	0.041	99.27%	99.30%

Table 1: Effect of Non-normality. 20% of population eigenvalues are equal to 1, 40% are equal to 3, and 40% are equal to 10. 1,000 Monte Carlo simulations with $p = 100$ and $n = 300$.

We can see from this table that departure from normality does not seem to have any noticeable effect on performance.

6.5 Precision Matrix

The next set of Monte Carlo simulations focuses on estimating the precision matrix Σ_n^{-1} . We take the same population eigenvalues as in Section 6.1. The concentration ratio $\widehat{c}_n = p/n$ is

set to the value $1/3$. For various values of p between 30 and 200, we run 1,000 simulations with normally distributed variables. The respective PRIALs of P_n^{or} , \hat{P}_n , \hat{S}_n^{-1} and \bar{S}_n^{-1} (with respect to S_n^{-1} and P_n^* now instead of S_n and S_n^*) are plotted in Figure 5.

We observe that the nonlinear shrinkage method seems to be just as effective for the purpose of estimating the precision matrix as it is for the purpose of estimating the covariance matrix itself. Moreover, there is a clear benefit in directly estimating the precision matrix by means of \hat{P}_n as opposed to the indirect estimation by means of \hat{S}_n^{-1} (which on its own significantly outperforms \bar{S}_n^{-1}).

6.6 Shape

Next, we study how the nonlinear shrinkage estimator performs for a wide variety of shapes of population spectral densities. This requires using a family of distributions with bounded support and which, for various parameter values, can take on different shapes. The best-suited family for this purpose is the beta distribution. The c.d.f. of the beta distribution with parameters (α, β) is:

$$\forall x \in [0, 1] \quad F_{(\alpha, \beta)}(x) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} \int_0^x t^{\alpha-1}(1-t)^{\beta-1} dt.$$

While the support of the beta distribution is $[0, 1]$, we shift it to the interval $[1, 10]$ by applying a linear transformation. Thanks to the flexibility of the beta family of densities, selecting different parameters (α, β) enables us to generate eight different shapes for the population spectral density: rectangular $(1, 1)$, linearly decreasing triangle $(1, 2)$, linearly increasing triangle $(2, 1)$, circular $(1.5, 1.5)$, U-shaped $(0.5, 0.5)$, bell-shaped $(5, 5)$, left-skewed $(5, 2)$ and right-skewed $(2, 5)$; see Figure 6 for a graphical illustration.

For every one of these eight beta densities, we take the population eigenvalues to be equal to:

$$1 + 9 F_{(\alpha, \beta)}^{-1} \left(\frac{i}{p} - \frac{1}{2p} \right), \quad i = 1, \dots, p.$$

The concentration ratio $\hat{c}_n = p/n$ is equal to $1/3$. For various values of p between 30 and 200, we run 1,000 simulations with normally distributed variables. The PRIAL of the nonlinear shrinkage estimator \hat{S}_n is plotted in Figure 7.

As in all the other simulations presented above, the PRIAL of the nonlinear shrinkage estimator always exceeds 88%, and more often than not exceeds 95%. In order to preserve the clarity of the picture, we do not report the PRIALs of the oracle and of the linear shrinkage estimator, but as usual the nonlinear shrinkage estimator ranked between them.

6.7 Fixed-Dimension Asymptotics

Finally, we report a set of Monte-Carlo simulations that departs from the large-dimensional asymptotics assumption under which the nonlinear shrinkage estimator \hat{S}_n was derived. The

goal is to compare it against the sample covariance matrix S_n in the setting where S_n is known to have certain optimality properties (at least in the normal case): traditional asymptotics, that is, when the number of variables p remains fixed while the sample size n goes to infinity. This gives as much advantage to the sample covariance matrix as it can possibly have. We fix the dimension $p = 100$ and let the sample size n vary from $n = 125$ to $n = 10,000$. In practice, very few applied researchers are fortunate enough to have as many as $n = 10,000$ i.i.d. observations, or a concentration ratio $c = p/n$ as low as 0.01. The respective PRIALs of S_n^{or} , \hat{S}_n and \bar{S}_n are plotted in Figure 8.

One crucial difference with all the previous simulations is that the target for the PRIAL is no longer S_n^* , but instead the population covariance matrix Σ itself, because now Σ can be consistently estimated. Note that, since the matrix dimension is fixed, Σ_n does not change with n ; therefore, we can drop the subscript n . Thus, in this subsection only:

$$\text{PRIAL} \equiv \text{PRIAL}(\hat{\Sigma}_n) \equiv 100 \times \left\{ 1 - \frac{\mathbb{E}[\|\hat{\Sigma}_n - \Sigma\|^2]}{\mathbb{E}[\|S_n - \Sigma\|^2]} \right\} \% ,$$

where $\hat{\Sigma}_n$ is an arbitrary estimator of the population covariance matrix Σ . By definition, the PRIAL of S_n is 0% while the PRIAL of Σ is 100%.

In this setting, Ledoit and Wolf (2004) acknowledge that the improvement of the linear shrinkage estimator over the sample covariance matrix vanishes asymptotically, because the optimal linear shrinkage intensity vanishes. Therefore it should be no surprise that the PRIAL of \bar{S}_n appears to go to zero in Figure 8. Perhaps more surprising is the continued ability of the oracle and the nonlinear shrinkage estimator to improve by approximately 60% over the sample covariance matrix, even for a sample size as large as $n = 10,000$, and with no sign of abating as n goes to infinity. This is an encouraging result, as our simulation gave every possible advantage to the sample covariance matrix by placing it in the asymptotic conditions where it possesses well-known optimality properties, and where the earlier linear shrinkage estimator of Ledoit and Wolf (2004) is most disadvantaged.

Intuitively, this is because the oracle shrinkage formula becomes more and more nonlinear as n goes to infinity for fixed p . Bai and Silverstein (1998) show that the sample covariance matrix exhibits ‘spectral separation’ when the concentration ratio p/n is sufficiently small. It means that the sample eigenvalues coalesce into clusters, each cluster corresponding to a Dirac of population eigenvalues. Within a given cluster, the smallest sample eigenvalues need to be nudged upwards, and the largest ones downwards, to the average of the cluster. In other words: full shrinkage within clusters, and no shrinkage between clusters. This is illustrated in Figure 9, which plots a typical simulation result for $n = 10,000$.¹

¹For enhanced ability to distinguish linear shrinkage from the sample covariance matrix, we plot the two uninterrupted lines, even though the sample eigenvalues lie in three disjoint intervals (as can be seen from nonlinear shrinkage).

By detecting this intricate pattern automatically, that is, by discovering where to shrink and where not to shrink, the nonlinear shrinkage estimator \widehat{S}_n showcases its ability to generate substantial improvements over the sample covariance matrix even for very low concentration ratios.

7 Conclusion

Estimating a large-dimensional covariance matrix is a very important and challenging problem. In the absence of additional information concerning the structure of the true covariance matrix, a successful approach consists of appropriately shrinking the sample eigenvalues, while retaining the sample eigenvectors. In particular, such shrinkage estimators enjoy the desirable property of being rotation equivariant.

In this paper, we have extended the linear approach of Ledoit and Wolf (2004) by applying a nonlinear transformation to the sample eigenvalues. The specific transformation suggested is motivated by the oracle estimator of Ledoit and P  ch   (2011) which, in turn, was derived by studying the asymptotic behavior of the finite-sample optimal rotation equivariant estimator (that is, the estimator with the rotation equivariant property which is closest to the true covariance matrix when distance is measured by the Frobenius norm).

The oracle estimator involves the Stieltjes transform of the limiting spectral distribution of the sample eigenvalues, evaluated at various points on the real line. By finding a way to consistently estimate these quantities, in a uniform sense, we have been able to construct a *bona fide* nonlinear shrinkage estimator which is asymptotically equivalent to the oracle.

Extensive Monte Carlo studies have demonstrated the improved finite-sample properties of our nonlinear shrinkage estimator compared to the sample covariance matrix and the linear shrinkage estimator of Ledoit and Wolf (2004), as well as its fast convergence to the performance of the oracle. In particular, when the sample size is very large compared to the dimension, or the population eigenvalues are very dispersed, the nonlinear shrinkage estimator still yields a significant improvement over the sample covariance matrix, while the linear shrinkage estimator no longer does.

Many applied problems require an estimator of the inverse of the covariance matrix, which is called the precision matrix. We have modified our nonlinear shrinkage approach to this alternative problem, thereby constructing a direct estimator of the precision matrix. Monte Carlo studies have confirmed that this estimator yields a sizeable improvement over the indirect method of simply inverting the nonlinear shrinkage estimator of the covariance matrix itself.

References

- Bai, Z. and Silverstein, J. W. (1998). No eigenvalues outside the support of the limiting spectral distribution of large-dimensional random matrices. *Annals of Probability*, 26(1):316–345.
- Bai, Z. D. and Silverstein, J. W. (1999). Exact separation of eigenvalues of large-dimensional sample covariance matrices. *Annals of Probability*, 27(3):1536–1555.
- Bickel, P. J. and Levina, E. (2008). Regularized estimation of large covariance matrices. *Annals of Statistics*, 36(1):199–227.
- El Karoui, N. (2008). Spectrum estimation for large dimensional covariance matrices using random matrix theory. *Annals of Statistics*, 36(6):2757–2790.
- Fan, J., Fan, Y., and Lv, J. (2008). High dimensional covariance matrix estimation using a factor model. *Journal of Econometrics*, 147(1):186–197.
- Geronimo, J. S. and Hill, T. P. (2003). Necessary and sufficient condition that the limit of Stieltjes transforms is a Stieltjes transform. *Journal of Approximation Theory*, 121:54–60.
- Gill, P. E., Murray, W., and Saunders, M. A. (2002). SNOPT: An SQP algorithm for large-scale constrained optimization. *SIAM Journal on Optimization*, 12(4):979–1006.
- Khare, K. and Rajaratnam, B. (2010). Wishart distributions for decomposable covariance graph models. *Annals of Statistics*. Forthcoming.
- Ledoit, O. and Péché, S. (2011). Eigenvectors of some large sample covariance matrix ensembles. *Probability Theory and Related Fields*, 150(1–2). Forthcoming.
- Ledoit, O. and Wolf, M. (2004). A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis*, 88(2):365–411.
- Marčenko, V. A. and Pastur, L. A. (1967). Distribution of eigenvalues for some sets of random matrices. *Sbornik: Mathematics*, 1(4):457–483.
- Mestre, X. (2008). On the asymptotic behavior of the sample estimates of eigenvalues and eigenvectors of covariance matrices. *IEEE Transactions on Signal Processing*, 56(11):5353–5368.
- Perlman, M. D. (2007). *STAT 542: Multivariate Statistical Analysis*. University of Washington (On-Line Class Notes), Seattle, Washington.
- Pollard, D. (1984). *Convergence of Stochastic Processes*. Springer-Verlag, New York.

- Rajaratnam, B., Massam, H., and Carvalho, C. M. (2008). Flexible covariance estimation in graphical Gaussian models. *Annals of Statistics*, 36(6):2818–2849.
- Silverman, B. W. (1986). *Density Estimation for Statistics and Data Analysis*. Chapman & Hall/CRC, Boca Raton.
- Silverstein, J. W. (1995). Strong convergence of the empirical distribution of eigenvalues of large-dimensional random matrices. *Journal of Multivariate Analysis*, 55:331–339.
- Silverstein, J. W. and Choi, S. I. (1995). Analysis of the limiting spectral distribution of large-dimensional random matrices. *Journal of Multivariate Analysis*, 54:295–309.
- Stein, C. (1955). Inadmissibility of the usual estimator for the mean of a multivariate normal distribution. In *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability*, pages 197–206. University of California Press.

A Proofs

Before proving Proposition 4.1, it is instructive to first state and prove a simpler result only claiming pointwise convergence of the estimated solutions. We will then see that this simpler proof can be extended relatively easily to also cover the more general claim of uniform convergence.

Proposition A.1. *Let $\{\hat{H}_n\}$ be a sequence of probability measures with $\hat{H}_n \Rightarrow H$. Let $\{\hat{c}_n\}$ be a sequence of positive real numbers with $\hat{c}_n \rightarrow c$. Let $K \subseteq (0, \infty)$ be a compact interval satisfying $y_x \in K$. Let $\hat{y}_{n,x} \equiv \min_{y \in K} g_{\hat{H}_n, \hat{c}_n}(y, x)$. It then holds that $\hat{y}_{n,x} \rightarrow y_x$.*

PROOF. Assume $K = [k_1, k_2]$. Define $B \equiv \{x + iy : x \in [u_1, u_2], y \in K\}$, which implies $B \subseteq \mathbb{C}^+$.

We first claim that:

$$m_{L\hat{H}_n}(z) \rightarrow m_{LH}(z) \quad \text{uniformly in } z \in B. \quad (\text{A.1})$$

Recalling that for any c.d.f. G , we have $m_{LG}(z) = 1 + z m_G(z)$ and by the compactness of the set B , this results will follow from:

$$m_{\hat{H}_n}(z) \rightarrow m_H(z) \quad \text{uniformly in } z \in B, \quad (\text{A.2})$$

which we establish now.

For fixed $z \in B$, consider the function:

$$h_z(\tau) \equiv \frac{\tau}{\tau - z}.$$

Then it is easy to see that there exist two finite constants d_1, d_2 , depending only on $k_1 > 0$ but not on z , such that:

$$|h_z(\tau_1) - h_z(\tau_2)| \leq d_1 |\tau_1 - \tau_2| \quad \text{and} \quad \sup_{\tau} |h_z(\tau)| \leq d_2. \quad (\text{A.3})$$

The fact that convergence in distribution of \hat{H}_n to H is equivalent to convergence to zero of the bounded-Lipschitz metric between \hat{H}_n and H then implies (A.2); for example, see Example 22 of Pollard (1984). In turn, we have thus established (A.1) as well. But (A.1) immediately implies:

$$g_{\hat{H}_n, \hat{c}_n}(y, x) \rightarrow g_{H, c}(y, x) \quad \text{uniformly in } y \in K. \quad (\text{A.4})$$

We note the following two facts:

$$\forall \varepsilon > 0 \quad \exists \delta > 0 \quad \text{such that} \quad \inf_{y \in K, |y - y_x| \geq \varepsilon} g_{H, c}(y, x) \geq \delta \quad (\text{A.5})$$

and

$$g_{\hat{H}_n, \hat{c}_n}(\hat{y}_{n,x}, x) = o(1), \quad (\text{A.6})$$

where (A.6) follows from $g_{\widehat{H}_n, \widehat{c}_n}(\widehat{y}_{n,x}, x) \leq g_{\widehat{H}_n, \widehat{c}_n}(y_x, x)$, (A.4), and $g_{H,c}(y_x, x) = 0$.

By the triangular inequality:

$$\begin{aligned} g_{H,c}(\widehat{y}_{n,x}, x) &\leq |g_{H,c}(\widehat{y}_{n,x}) - g_{\widehat{H}_n, \widehat{c}_n}(\widehat{y}_{n,x})| + |g_{\widehat{H}_n, \widehat{c}_n}(\widehat{y}_{n,x})| \\ &= |g_{H,c}(\widehat{y}_{n,x}) - g_{\widehat{H}_n, \widehat{c}_n}(\widehat{y}_{n,x})| + o(1) \quad \text{by (A.6)} \\ &= o(1) + o(1) \quad \text{by (A.4)} \\ &= o(1) . \end{aligned}$$

This last result together with (A.5) now imply $\widehat{y}_{n,x} \rightarrow y_x$. ■

PROOF OF PROPOSITION 4.1. We start with part (i). Assume $K = [k_1, k_2]$. Define $B \equiv \{x + iy : x \in [u_1, u_2], y \in K\}$, which implies $B \subseteq \mathbb{C}^+$.

By the same arguments leading up to (A.4) we can more generally establish that:

$$g_{\widehat{H}_n, \widehat{c}_n}(z) \rightarrow g_{H,c}(z) \quad \text{uniformly in } z \in B . \quad (\text{A.7})$$

We note the following two facts:

$$\forall \varepsilon > 0 \quad \exists \delta > 0 \text{ such that } \inf_{x \in [u_1 + \eta, u_2 - \eta]} \left\{ \inf_{y \in K, |y - y_x| \geq \varepsilon} g_{H,c}(y, x) \right\} \geq \delta \quad (\text{A.8})$$

and

$$\sup_{x \in [u_1 + \eta, u_2 - \eta]} g_{\widehat{H}_n, \widehat{c}_n}(\widehat{y}_{n,x}, x) = o(1) , \quad (\text{A.9})$$

where (A.9) follows from $g_{\widehat{H}_n, \widehat{c}_n}(\widehat{y}_{n,x}, x) \leq g_{\widehat{H}_n, \widehat{c}_n}(y_x, x)$, (A.7), and $g_{H,c}(y_x, x) = 0$.

To simplify the notation, let $I \equiv [u_1 + \eta, u_2 - \eta]$. By the triangular inequality:

$$\begin{aligned} \sup_{x \in I} g_{H,c}(\widehat{y}_{n,x}, x) &\leq \sup_{x \in I} |g_{H,c}(\widehat{y}_{n,x}) - g_{\widehat{H}_n, \widehat{c}_n}(\widehat{y}_{n,x})| + \sup_{x \in I} |g_{\widehat{H}_n, \widehat{c}_n}(\widehat{y}_{n,x})| \\ &= \sup_{x \in I} |g_{H,c}(\widehat{y}_{n,x}) - g_{\widehat{H}_n, \widehat{c}_n}(\widehat{y}_{n,x})| + o(1) \quad \text{by (A.9)} \\ &= o(1) + o(1) \quad \text{by (A.7)} \\ &= o(1) . \end{aligned}$$

This last result together with (A.9) now imply $\widehat{y}_{n,x} \rightarrow y_x$ uniformly in $x \in I = [u_1 + \eta, u_2 - \eta]$.

Part (ii) is proven analogously to part (i) by restricting attention to the set of probability one on which $\widehat{H}_n \Rightarrow H$ happens. ■

PROOF OF PROPOSITION 4.2. The proof is similar to the proof of Proposition 4.1. The details are left to the reader. ■

PROOF OF PROPOSITION 4.3. We start with part (i)(a). Fix $\lambda \in [\widetilde{z}_1 + \widetilde{\delta}, \widetilde{z}_2 - \widetilde{\delta}]$. Consider:

$$\left| \check{m}_{F_{\widehat{H}_n, \widehat{c}_n}}(\lambda) - \check{m}_F(\lambda) \right| = \left| \frac{1 - \widehat{c}_n}{\widehat{c}_n \lambda} - \frac{1}{\widehat{c}_n \widehat{v}_{n,\lambda}} - \left(\frac{1 - c}{c \lambda_x} - \frac{1}{c v_\lambda} \right) \right| .$$

The function mapping λ onto v_λ is continuous, and therefore uniformly continuous, in $\lambda \in [\tilde{z}_1, \tilde{z}_2]$. As λ varies in $[\tilde{z}_1 + \tilde{\delta}, \tilde{z}_2 - \tilde{\delta}]$, the resulting v_λ varies in a compact region in \mathbb{C}^+ . Therefore, for any $\xi > 0$, there exists $\kappa > 0$ such that:

$$|\check{m}_{F_{\hat{H}_n, \hat{c}_n}}(\lambda) - \check{m}_F(\lambda)| < \xi \quad \text{as long as} \quad \max\{|\hat{c}_n - c|, |\hat{v}_{n, \lambda} - v_\lambda|\} < \kappa.$$

First, we can find N_1 such that $|\hat{c}_n - c| < \kappa$ for all $n \geq N_1$. Second, by part (i) of Proposition 4.2, we can find N_2 such that $|\hat{v}_{n, \lambda} - v_\lambda| < \kappa$ for all $n \geq N_2$, uniformly in $\lambda \in [\tilde{z}_1 + \tilde{\delta}, \tilde{z}_2 - \tilde{\delta}]$. Define $N \equiv \max\{N_1, N_2\}$. Then for all $n \geq N$, it holds that:

$$|\check{m}_{F_{\hat{H}_n, \hat{c}_n}}(\lambda) - \check{m}_F(\lambda)| < \xi, \quad \text{uniformly in } \lambda \in [\tilde{z}_1 + \tilde{\delta}, \tilde{z}_2 - \tilde{\delta}].$$

Since ξ can be chosen arbitrarily small, part (i)(a) obtains.

We now turn to part (i)(b). For any $\tilde{\delta} > 0$, it holds:

$$\begin{aligned} \|\hat{S}_n - S_n^{or}\|^2 &= \frac{1}{p} \sum_{i=1}^p \left(\frac{\lambda_i}{|1 - \hat{c}_n - \hat{c}_n \lambda_i \check{m}_{F_{\hat{H}_n, \hat{c}_n}}(\lambda_i)|^2} - \frac{\lambda_i}{|1 - c - c \lambda_i \check{m}_F(\lambda_i)|^2} \right)^2 \\ &= \frac{1}{p} \sum_{\lambda_i \in [\tilde{z}_1 + \tilde{\delta}, \tilde{z}_2 - \tilde{\delta}]} \left(\frac{\lambda_i}{|1 - \hat{c}_n - \hat{c}_n \lambda_i \check{m}_{F_{\hat{H}_n, \hat{c}_n}}(\lambda_i)|^2} - \frac{\lambda_i}{|1 - c - c \lambda_i \check{m}_F(\lambda_i)|^2} \right)^2 \\ &\quad + \frac{1}{p} \sum_{\lambda_i \notin [\tilde{z}_1 + \tilde{\delta}, \tilde{z}_2 - \tilde{\delta}]} \left(\frac{\lambda_i}{|1 - \hat{c}_n - \hat{c}_n \lambda_i \check{m}_{F_{\hat{H}_n, \hat{c}_n}}(\lambda_i)|^2} - \frac{\lambda_i}{|1 - c - c \lambda_i \check{m}_F(\lambda_i)|^2} \right)^2 \\ &\equiv A + B. \end{aligned}$$

By our general set of assumptions, in particular Assumption (A4), combined with the results of Bai and Silverstein (1998) and Mestre (2008, Section II), there exist two finite, non-zero constants $\kappa_1 < \kappa_2$ such that $\kappa_1 \leq \lambda_i \leq \kappa_2$ for all $i = 1, \dots, p$ and for all n large enough.

Fix $\varepsilon > 0$. First, we can pick $\tilde{\delta}$ small enough to achieve $B \leq \varepsilon/2$ eventually. To appreciate why, denote by $\mu(\tilde{\delta})$ the mass which F assigns to the set $[\tilde{z}_1, \tilde{z}_1 + \tilde{\delta}] \cup [\tilde{z}_2 - \tilde{\delta}, \tilde{z}_2]$, satisfying $\mu(\tilde{\delta}) \rightarrow 0$ as $\tilde{\delta} \rightarrow 0$. Then it is not too difficult to see that there exists a finite constant Δ , possibly depending on H and c , such that $B \leq \Delta \mu(\tilde{\delta})$, for n sufficiently large. The reason, in addition to $\kappa_1 \leq \lambda_i \leq \kappa_2$, is that also the correction factors $|1 - \hat{c}_n - \hat{c}_n \lambda_i \check{m}_{F_{\hat{H}_n, \hat{c}_n}}(\lambda_i)|^2$ and $1/|1 - c - c \lambda_i \check{m}_F(\lambda_i)|^2$ are bounded away from infinity. Then, choose $\tilde{\delta}$ small enough so that $\mu(\tilde{\delta}) \leq (2/\varepsilon)\Delta$.

Having chosen and fixed $\tilde{\delta}$, the first half of the assertion ensures that $A \leq \varepsilon/2$ eventually. Again, we use here that $\kappa_1 \leq \lambda_i \leq \kappa_2$ and that also the correction factors $1/|1 - c - c \lambda_i \check{m}_F(\lambda_i)|^2$ are bounded away from infinity. This demonstrates part (i)(b).

Part(i)(c) can be handled in a very similar fashion.

Part (ii) is proven analogously to part (i) by focusing on the set of probability one on which $\hat{H}_n \Rightarrow H$ happens. ■

Before proving Theorem 5.1, we need to establish some auxiliary results.

Recall the following notation. For a grid Q on the real line and for two c.d.f.s G_1 and G_2 , define:

$$\|G_1 - G_2\|_Q \equiv \sup_{t \in Q} |G_1(t) - G_2(t)|$$

Lemma A.1. *Let $\{G_n\}$ and G be c.d.f.s on the real line, with the support of G being compact. Let $\{Q_n\}$ be a sequence of grids on the real line, asymptotically covering the support of G , with grid sizes $\{\gamma_n\}$ satisfying $\gamma_n \rightarrow 0$.*

If G is continuous, then $G_n \Rightarrow G$. In particular, $\sup_t |G_n(t) - G(t)| \rightarrow 0$.

PROOF. Denote the compact support of G by $[a, b]$. To prove the first part of the assertion, let $\varepsilon > 0$. Fix $\delta > 0$ such that for all $t < t'$ with $t' - t < \delta$, it holds $G(t') - G(t) < \varepsilon/4$. Also fix $\phi > 0$. First, there exists N_1 such that $\gamma_n < \delta$ for all $n \geq N_1$. Second, there exists N_2 such that $\sup_{t \in Q_n} |G_n(t) - G(t)| < \varepsilon/4$ for all $n \geq N_2$. Third, there exists N_3 such that Q_n covers $[a + \phi, b - \phi]$ for all $n \geq N_3$. Set $N \equiv \max\{N_1, N_2, N_3\}$. For an arbitrary $t \in [a + \phi, b - \phi]$ and for $n \geq N$, let $t_n \equiv \max\{\tilde{t} : \tilde{t} \in Q_n, \tilde{t} \leq t\}$ and $t'_n \equiv \min\{\tilde{t} : \tilde{t} \in Q_n, \tilde{t} \geq t\}$, which implies $t_n - t'_n < \delta$. Then for all $n \geq N$:

$$\begin{aligned} |G_n(t) - G(t)| &\leq |G_n(t_n) - G(t'_n)| + |G_n(t'_n) - G(t_n)| \\ &\leq |G_n(t_n) - G(t_n)| + |G_n(t'_n) - G(t'_n)| + \frac{\varepsilon}{4} + \frac{\varepsilon}{4} \\ &\leq \frac{\varepsilon}{4} + \frac{\varepsilon}{4} + \frac{\varepsilon}{4} + \frac{\varepsilon}{4} = \varepsilon. \end{aligned}$$

Therefore, $G_n(t)$ converges to $G(t)$ for all $t \in [a + \phi, b - \phi]$; and since ϕ can be chosen arbitrarily, $G_n(t)$ converges to $G(t)$ for all $t \in (a, b)$. By picking ϕ sufficiently small such that $|G(a + \phi)| \leq \varepsilon$ and $|G(b - \varepsilon)| \geq 1 - \varepsilon$, and by the monotonicity of c.d.f.s, it also follows that $|G_n(t)| \leq 2\varepsilon$ for all $t \leq a$ as well as $|G_n(t)| \geq 1 - 2\varepsilon$ for all $t \geq b$ as long as $n \geq N$ (where N of course is allowed to depend on ϕ .) Therefore, $G_n(t)$ converges to $G(t)$ for all t , which establishes $G_n \Rightarrow G$. The second part of the assertion follows immediately from the first part and Polya's Theorem. ■

Lemma A.2. *Let G be a probability measure with compact support contained in $(0, +\infty)$ and let $d > 0$. Let $\{\widehat{G}_n\}$ be a sequence of probability measures on the nonnegative real line with $\widehat{G}_n \Rightarrow G$ and let $\{\widehat{d}_n\}$ be a sequence of positive real numbers with $\widehat{d}_n \rightarrow d$. Also assume that there exists an interval $[a, b]$ contained in $(0, +\infty)$ such that $\text{Supp}(\widehat{G}_n) \subseteq [a, b]$ for all n large enough.*

Then $F_{\widehat{G}_n, \widehat{d}_n} \Rightarrow F_{G, d}$.

PROOF. Let $z_j \equiv i \cdot (1 + 1/j)$, for $j = 1, 2, \dots$. Then $\{z_j\}$ is an infinite sequence in \mathbb{C}^+ with limit point $z_0 \equiv i \in \mathbb{C}^+$. By Theorem 2 of Geronimo and Hill (2003), it is sufficient to show that for all z_j :

$$m_{F_{\widehat{G}_n, \widehat{d}_n}}(z_j) \rightarrow m_{F_{G, d}}(z_j). \quad (\text{A.10})$$

Recall the notation $m_{F_{\tilde{H}, \tilde{c}}}$ for the solution of the Marčenko-Pastur equation, for any probability measure \tilde{H} and for any $\tilde{c} > 0$. Namely, for each $z \in \mathbb{C}^+$, $m_{F_{\tilde{H}, \tilde{c}}}(z)$ is the unique solution for $m \in \mathbb{C}^+$ to the equation:

$$m = \int_{-\infty}^{+\infty} \frac{1}{\tau [1 - \tilde{c} - \tilde{c} z m] - z} d\tilde{H}(\tau) .$$

Also, define the function:

$$\forall m, z \in \mathbb{C} \quad f_{\tilde{H}, \tilde{c}}(m, z) \equiv \left| m - \int_{-\infty}^{+\infty} \frac{1}{\tau [1 - \tilde{c} - \tilde{c} z m] - z} d\tilde{H}(\tau) \right| .$$

In this notation, for a given $z \in \mathbb{C}^+$, $m_{F_{\tilde{H}, \tilde{c}}}(z)$ is the unique solution for $m \in \mathbb{C}^+$ to the equation $f_{\tilde{H}, \tilde{c}}(m, z) = 0$. Alternatively, $m_{F_{\tilde{H}, \tilde{c}}}(z)$ is the unique minimizer over $m \in \mathbb{C}^+$ of the function $f_{\tilde{H}, \tilde{c}}(\cdot, z)$. Note that the Stieltjes transform of any probability measure maps \mathbb{C}^+ onto \mathbb{C}^+ . So if $z \in \mathbb{C}^+$, then $m_{F_{\tilde{H}, \tilde{c}}}(z)$ is actually the unique minimizer over $m \in \mathbb{C}$ of the function $f_{\tilde{H}, \tilde{c}}(\cdot, z)$.

Fix z_j and use the following abbreviations: $\hat{m}_{n, z_j} \equiv m_{F_{\hat{G}_n, \hat{d}_n}}(z_j)$ and $m_{z_j} \equiv m_{F_{G, d}}(z_j)$. The goal then is to show that $\hat{m}_{n, z_j} \rightarrow m_{z_j}$.

We claim that there exists a compact set $S \subseteq \mathbb{C}$ such that $\hat{m}_{n, z_j} \in S$ for all n . The proof is by means of contradiction. Assume the claim does not hold. Then there exists a subsequence $\{n_k\}$ such that $|\hat{m}_{n_k, z_j}| \rightarrow \infty$. By the combined assumptions, we can then find $\Delta > 0$ such that for all n_k large enough and for all $\tau \in [a, b]$:

$$\frac{1}{|\tau [1 - \hat{d}_{n_k} - \hat{d}_{n_k} z_j \hat{m}_{n_k, z_j}] - z_j|} \leq \Delta$$

implying that for all n_k large enough:

$$\begin{aligned} |\hat{m}_{n_k, z_j}| &= \left| \int_{-\infty}^{+\infty} \frac{1}{\tau [1 - \hat{d}_{n_k} - \hat{d}_{n_k} z_j \hat{m}_{n_k, z_j}] - z_j} d\hat{G}_{n_k}(\tau) \right| \\ &= \left| \int_a^b \frac{1}{\tau [1 - \hat{d}_{n_k} - \hat{d}_{n_k} z_j \hat{m}_{n_k, z_j}] - z_j} d\hat{G}_{n_k}(\tau) \right| \\ &\leq \int_a^b \frac{1}{|\tau [1 - \hat{d}_{n_k} - \hat{d}_{n_k} z_j \hat{m}_{n_k, z_j}] - z_j|} d\hat{G}_{n_k}(\tau) \\ &\leq (b - a) \Delta . \end{aligned}$$

But this is in contradiction to $|\hat{m}_{n_k, z_j}| \rightarrow \infty$. We may assume w.l.o.g. that $m_{z_j} \in S$ as well; otherwise sufficiently enlarge S .

We may further assume that S is ‘doubly nonnegative’, that is, for all $m \in S$, it holds that $\operatorname{Re}(m) \geq 0$ as well as $\operatorname{Im}(m) \geq 0$. The reason is as follows. On the one hand, $\operatorname{Re}(\hat{m}_{n, z_j}) \geq 0$ for all n as well as $\operatorname{Re}(m_{z_j}) \geq 0$. For example, recalling that $\operatorname{Re}(z_j) = 0$:

$$\operatorname{Re}(m_{z_j}) = \operatorname{Re}(m_{F_{G, d}}(z_j)) = \int_{-\infty}^{\infty} \operatorname{Re} \left(\frac{1}{\lambda - z_j} \right) dF_{G, d}(\lambda) = \int_{-\infty}^{\infty} \frac{\lambda}{|\lambda - z_j|^2} dF_{G, d}(\lambda) ,$$

where $F_{G,d}$ places all its mass on $[0, +\infty)$. On the other hand, since $z_j \in \mathbb{C}^+$, also $\text{Im}(\hat{m}_{n,z_j}) > 0$ for all n as well as $\text{Im}(m_{z_j}) > 0$.

We next claim that:

$$f_{\hat{G}_{n,d}}(m, z_j) \rightarrow f_{G,d}(m, z_j) \quad \text{uniformly in } m \in S. \quad (\text{A.11})$$

To see why, for $m \in S$, consider the function:

$$h_{m,z_j}(\tau) \equiv \frac{1}{\tau [1 - d - d z_j m] - z_j}.$$

Since S is compact, $\min\{\text{Re}(m), \text{Im}(m)\} \geq 0$ for all $m \in S$, $\text{Re}(z_j) = 0$, and $\text{Im}(z_j) \geq 1$, there exist two finite constants d_1 and d_2 , allowed to depend on S , such that:

$$|h_{m,z_j}(\tau_1) - h_{m,z_j}(\tau_2)| \leq d_1 |\tau_1 - \tau_2| \quad \text{for } \tau_1, \tau_2 \in [0, +\infty) \quad (\text{A.12})$$

and

$$\sup_{\tau \in [0, +\infty)} |h_{m,z_j}(\tau)| \leq d_2. \quad (\text{A.13})$$

To see why, start with (A.13). It holds that:

$$\text{Im}(\tau [1 - d - d z_j m] - z_j) = -(\tau d [\text{Re}(z_j) \text{Im}(m) + \text{Im}(z_j) \text{Re}(m)] + \text{Im}(z_j)).$$

Under the stated conditions, $\text{Re}(z_j) \text{Im}(m) + \text{Im}(z_j) \text{Re}(m) \geq 0$ and $\text{Im}(z_j) \geq 1$. Therefore, as long as $\tau \geq 0$, it follows that:

$$|\tau [1 - d - d z_j m] - z_j| \geq |\text{Im}(\tau [1 - d - d z_j m] - z_j)| \geq 1,$$

implying that we may choose $d_2 \equiv 1$.

Moving on to (A.12), let $\Delta \equiv \max_{m \in S} |m|$ and note that $|z_j| \leq 2$. Therefore, for any $\tau_1, \tau_2 \in [0, +\infty)$:

$$\begin{aligned} |h_{m,z_j}(\tau_1) - h_{m,z_j}(\tau_2)| &= |\tau_1 - \tau_2| \left| \frac{1 - d - d z_j m}{(\tau_1 [1 - d - d z_j m] - z_j)(\tau_2 [1 - d - d z_j m] - z_j)} \right| \\ &= |\tau_1 - \tau_2| \frac{|1 - d - d z_j m|}{|\tau_1 [1 - d - d z_j m] - z_j| |\tau_2 [1 - d - d z_j m] - z_j|} \\ &= |\tau_1 - \tau_2| \frac{|1 - d - d z_j m|}{|\tau_1 [1 - d - d z_j m] - z_j| |\tau_2 [1 - d - d z_j m] - z_j|} \\ &\leq |\tau_1 - \tau_2| (1 + d + 2d\Delta), \end{aligned}$$

implying that we may choose $d_1 \equiv (1 + d + 2d\Delta)$.

Recall that convergence in distribution of \hat{G}_n to G is equivalent to convergence to zero of the bounded-Lipschitz metric between \hat{G}_n and G ; for example, see Pollard (1984, Example 22).

Furthermore, since \widehat{G}_n and G put all their mass on $[0, \infty)$, it is actually sufficient to start all integrals at $\tau = 0$ rather than at $\tau = -\infty$. Therefore:

$$\begin{aligned}
\int_{-\infty}^{+\infty} \frac{d\widehat{G}_n(\tau)}{\tau [1 - d - d z_j m] - z_j} &= \int_0^{+\infty} \frac{1}{\tau [1 - d - d z_j m] - z_j} d\widehat{G}_n(\tau) \\
&= \int_0^{\infty} h_{m, z_j}(\tau) d\widehat{G}_n(\tau) \\
&\rightarrow \int_0^{\infty} h_{m, z_j}(\tau) dG(\tau) \\
&= \int_0^{+\infty} \frac{1}{\tau [1 - d - d z_j m] - z_j} dG(\tau) \\
&= \int_{-\infty}^{+\infty} \frac{1}{\tau [1 - d - d z_j m] - z_j} dG(\tau) \quad \text{uniformly in } m \in S,
\end{aligned}$$

which establishes (A.11). But (A.11), combined with the compactness of S , further implies that also:

$$f_{\widehat{G}_n, \widehat{d}_n}(m, z_j) \rightarrow f_{G, d}(m, z_j) \quad \text{uniformly in } m \in S. \quad (\text{A.14})$$

Summing up, we have the following facts: First, there exists a compact set $S \subseteq C$ such that \widehat{m}_{n, z_j} is the unique minimizer of $f_{\widehat{G}_n, \widehat{d}_n}(\cdot, z_j)$ over $m \in S$ and m_{z_j} is the unique minimizer of $f_{G, d}(\cdot, z_j)$ over $m \in S$. Second, the function $f_{G, d}(\cdot, z_j)$ is continuous in m . Third, the uniform convergence (A.14).

With these facts, $\widehat{m}_{n, z_j} \rightarrow m_{z_j}$ follows from arguments very similar to those used in the proof of Proposition A.1. ■

PROOF OF THEOREM 5.1. We start with the proof of part (i). Since $c < 1$, it follows from Silverstein and Choi (1995) that F is continuously differentiable on all of \mathbb{R} . By Polya's Theorem it then follows that $\sup_t |F_n(t) - F(t)| \rightarrow 0$ a.s., implying that $\|F_n - F\|_{Q_n} \rightarrow 0$ a.s. Also, by construction, $\|F_{\widehat{H}_n, \widehat{c}_n} - F_n\|_{Q_n} \leq \|F_{H, \widehat{c}_n} - F_n\|_{Q_n}$. Therefore:

$$\begin{aligned}
\|F_{\widehat{H}_n, \widehat{c}_n} - F\|_{Q_n} &\leq \|F_{\widehat{H}_n, \widehat{c}_n} - F_n\|_{Q_n} + \|F_n - F\|_{Q_n} \\
&\leq \|F_{H, \widehat{c}_n} - F_n\|_{Q_n} + \|F_n - F\|_{Q_n} \\
&\leq \|F_{H, \widehat{c}_n} - F_{H, c}\|_{Q_n} + \|F_{H, c} - F_n\|_{Q_n} + \|F_n - F\|_{Q_n} \\
&= \|F_{H, \widehat{c}_n} - F\|_{Q_n} + 2\|F_n - F\|_{Q_n} \rightarrow 0 \quad \text{a.s.},
\end{aligned}$$

where Lemma A.2 in conjunction with Polya's Theorem is used to show that $\|F_{H, \widehat{c}_n} - F\|_{Q_n} \rightarrow 0$. The desired result now follows by Lemma A.1.

We now turn to proving part (ii). By Theorem 2 of Geronimo and Hill (2003), it is sufficient to show that there exists an infinite sequence $\{v_j\}$ in \mathbb{C}^+ with a limit point $v_0 \in \mathbb{C}^+$ such that:

$$m_{\widehat{H}_n}(v_j) \rightarrow m_H(v_j) \quad \text{a.s. } \forall j. \quad (\text{A.15})$$

Recall the notation $m_{F_{\tilde{H}, \tilde{c}}}$ for the solution of the Marčenko-Pastur equation, for any probability measure \tilde{H} and for any $\tilde{c} > 0$. Namely, for each $z \in \mathbb{C}^+$, $m_{F_{\tilde{H}, \tilde{c}}}(z)$ is the unique solution for $m \in \mathbb{C}^+$ to the equation:

$$m = \int_{-\infty}^{+\infty} \frac{1}{\tau [1 - \tilde{c} - \tilde{c} z m] - z} d\tilde{H}(\tau) .$$

Analogously, to Subsection 2.2, also let:

$$\forall x \in \mathbb{R} \quad \underline{F}_{\tilde{H}, \tilde{c}}(x) \equiv (1 - \tilde{c}) \mathbf{1}_{[0, +\infty)}(x) + \tilde{c} F_{\tilde{H}, \tilde{c}}(x)$$

and

$$\forall z \in \mathbb{C}^+ \quad m_{\underline{F}_{\tilde{H}, \tilde{c}}}(z) \equiv \frac{\tilde{c} - 1}{z} + \tilde{c} m_{F_{\tilde{H}, \tilde{c}}}(z) .$$

Hence, for each $z \in \mathbb{C}^+$, $m_{\underline{F}_{\tilde{H}, \tilde{c}}}(z)$ is the unique solution for $m \in \mathbb{C}^+$ to the equation:

$$m = - \left[z - \tilde{c} \int_{-\infty}^{+\infty} \frac{\tau}{1 + \tau m} d\tilde{H}(\tau) \right]^{-1} .$$

On \mathbb{C}^+ , $m_{\underline{F}_{\tilde{H}, \tilde{c}}}(z)$ has a unique inverse, given by:

$$\forall m \in m_{\underline{F}_{\tilde{H}, \tilde{c}}}(\mathbb{C}^+) \quad z_{\underline{F}_{\tilde{H}, \tilde{c}}}(m) \equiv -\frac{1}{m} + \tilde{c} \int_{-\infty}^{+\infty} \frac{\tau}{1 + \tau m} d\tilde{H}(\tau) .$$

Note that both $m_{\underline{F}_{\tilde{H}, \tilde{c}}}$ and $z_{\underline{F}_{\tilde{H}, \tilde{c}}}$ are continuous functions. Also in this notation, we have $\underline{F} = \underline{F}_{H, c}$, $m_{\underline{F}} = m_{\underline{F}_{H, c}}$, and $z_{\underline{F}} = z_{\underline{F}_{H, c}}$ then.

As Silverstein and Choi (1995) show:

$$\forall m \in m_{\underline{F}_{\tilde{H}, \tilde{c}}}(\mathbb{C}^+) \quad z_{\underline{F}_{\tilde{H}, \tilde{c}}}(m) = -\frac{1}{m} + \frac{\tilde{c}}{m} - \frac{\tilde{c}}{m^2} m_{\tilde{H}} \left(-\frac{1}{m} \right) ,$$

which, letting $v \equiv -1/m$, is equivalent to:

$$\forall v \in \mathbb{C}^+ \text{ such that } -\frac{1}{v} \in m_{\underline{F}_{\tilde{H}, \tilde{c}}}(\mathbb{C}^+) \quad m_{\tilde{H}}(v) = -\frac{1}{\tilde{c} v^2} \left[z_{\underline{F}_{\tilde{H}, \tilde{c}}} \left(-\frac{1}{v} \right) - v + \tilde{c} v \right] . \quad (\text{A.16})$$

For the special case of $\tilde{H} \equiv H$ and $\tilde{c} \equiv c$, this simplifies to

$$\forall v \in \mathbb{C}^+ \text{ such that } -\frac{1}{v} \in m_{\underline{F}}(\mathbb{C}^+) \quad m_H(v) = -\frac{1}{c v^2} \left[z_{\underline{F}} \left(-\frac{1}{v} \right) - v + c v \right] . \quad (\text{A.17})$$

Let $M \subseteq \mathbb{C}^+$ be a compact set contained in $m_{\underline{F}}(\mathbb{C}^+)$ and also contained in $m_{\underline{F}_{\hat{H}_n, \hat{c}_n}}(\mathbb{C}^+)$, at least for n large enough. Let $\{m_j\} \subseteq M$ be an infinite sequence with limit point $m_0 \in M$. Let $v_j \equiv -1/m_j$ and $v_0 \equiv -1/m_0$. Then $\{v_j\} \subseteq \mathbb{C}^+$ with limit point $v_0 \in \mathbb{C}^+$. Finally, let $z_j \equiv z_{\underline{F}}(m_j)$ and $z_0 \equiv z_{\underline{F}}(m_0)$.

Part (i) of the theorem implies that $\underline{F}_{\hat{H}_n, \hat{c}_n} \Rightarrow \underline{F}$ a.s. It then follows from Corollary 1 of Geronimo and Hill (2003) that:

$$m_{\underline{F}_{\hat{H}_n, \hat{c}_n}}(z_j) \rightarrow m_{\underline{F}}(z_j) \text{ a.s. } \forall j .$$

In particular, the proof of Corollary 1 of Geronimo and Hill (2003) uses that convergence in distribution of probability measures implies convergence of integrals of bounded and continuous functions. A completely analogous argument can therefore be invoked to show that also:

$$z_{\underline{F}_{\hat{H}_n, \hat{c}_n}}(m_j) \rightarrow z_{\underline{F}}(m_j) \quad \text{a.s. } \forall j$$

or, equivalently, that:

$$z_{\underline{F}_{\hat{H}_n, \hat{c}_n}}\left(-\frac{1}{v_j}\right) \rightarrow z_{\underline{F}}\left(-\frac{1}{v_j}\right) \quad \text{a.s. } \forall j .$$

Using relation (A.16), with $\tilde{H} \equiv \hat{H}_n$ and $\tilde{c} \equiv \hat{c}_n$, and relation (A.17), this implies that:

$$\begin{aligned} m_{\hat{H}_n}(v_j) &= -\frac{1}{\hat{c}_n v_j^2} \left[z_{\underline{F}_{\hat{H}_n, \hat{c}_n}}\left(-\frac{1}{v_j}\right) - v_j + \hat{c}_n v_j \right] \\ &\rightarrow -\frac{1}{c v_j^2} \left[z_{\underline{F}}\left(-\frac{1}{v_j}\right) - v_j + c v_j \right] = m_H(v_j) \quad \text{a.s. } \forall j , \end{aligned}$$

which completes the proof of part (ii) the theorem. ■

PROOF OF COROLLARY 5.1. We start with the proof of part (i). Following El Karoui (2008), we call H_{T_n} a discretization of H on the grid $\{J_n/T_n, (J_n + 1)/T_n, \dots, K_n/T_n\}$. For instance, we can choose H_{T_n} to be a step function with $H_{T_n}(x) \equiv H(x)$ if $x = l/T_n$, $l \in \mathbb{N}$, and H_{T_n} is constant on $[l/T_n, (l + 1)/T_n)$. If the support of H is given by $[h_1, h_2]$, say, then the support of H_{T_n} is contained in $[h_1 - 1/T_n, h_2 + 1/T_n]$. It is easy to see that for such a discretization H_{T_n} , it holds that $H_{T_n} \Rightarrow H$ as long as:

$$\exists b > 0 \text{ such that } \lambda_p \leq b \text{ for all } n \text{ sufficiently large} \quad \text{and} \quad (\text{A.18})$$

$$\exists \gamma > 0 \text{ such that } J_n/T_n \leq h_1 - \gamma \text{ and } K_n/T_n \geq h_2 + \gamma \text{ for all } n \text{ sufficiently large} . \quad (\text{A.19})$$

First, (A.18) holds a.s. as shown by Bai and Silverstein (1998) and Mestre (2008, Section II) given our set of assumptions, in particular Assumption (A4). Second, the support of F is denoted by $[\tilde{z}_1, \tilde{z}_2]$. On the one hand, it follows from Lemma 1.4 of Bai and Silverstein (1999) that $\tilde{z}_1 < h_1$ and $\tilde{z}_2 > h_2$. Therefore, it holds that $z_1 = h_1 - \delta_1$ and $z_2 = h_2 + \delta_2$ for some $\delta_1, \delta_2 > 0$. On the other hand, $F_n \Rightarrow F$ a.s., implying that $\lambda_1 \leq \tilde{z}_1 + \delta_1/2$ and $\lambda_p \geq \tilde{z}_2 - \delta_2/2$ for n sufficiently large a.s. So, letting $\gamma \equiv \min\{\delta_1/2, \delta_2/2\}$, condition (A.19) holds a.s. as well. Taken together, it follows that $H_{T_n} \Rightarrow H$ a.s.

By construction:

$$\|F_{\hat{H}_n, \hat{c}_n} - F_n\|_{Q_n} \leq \|F_{H_{T_n}, \hat{c}_n} - F_n\|_{Q_n} \leq \|F_{H_{T_n}, \hat{c}_n} - F\|_{Q_n} + \|F - F_n\|_{Q_n} .$$

We know that $\|F - F_n\|_{Q_n} \rightarrow 0$ a.s. So to establish part (i), it is sufficient to show that $\|F_{H_{T_n}, \hat{c}_n} - F\|_{Q_n} \rightarrow 0$ a.s. Since $H_{T_n} \Rightarrow H$ a.s. and $\hat{c}_n \rightarrow c$, it follows from Lemma A.2 and Polya's Theorem that $\sup_t |F_{H_{T_n}, \hat{c}_n}(t) - F(t)| \rightarrow 0$ a.s., implying that $\|F_{H_{T_n}, \hat{c}_n} - F\|_{Q_n} \rightarrow 0$ a.s.

But, having established part (i), part (ii) follows in exactly the same fashion as in the proof of Theorem 5.1. ■

PROOF OF COROLLARY 5.2. We start with some preliminary results, leading up to the proof of part (ii). Let G be a c.d.f. with continuous density g and compact support $[a, b]$. For a grid $Q \equiv \{\dots, t_{-1}, t_0, t_1, \dots\}$ covering the support of G , the approximation to G via trapezoidal integration is defined as in (5.3). Since g is Lipschitz-continuous on $[a, b]$, there exists a (smallest) finite $\varepsilon > 0$ such that $|g(t_1) - g(t_2)| \leq \varepsilon$ as long as $|t_1 - t_2| \leq \gamma$. Denote by \hat{g}_Q the density corresponding to \hat{G}_Q . By definition of the trapezoidal rule, \hat{g}_Q is piecewise linear and agrees with g at all points $t_k \in Q$. Since the grid size of Q is given by γ , we may infer that:

$$\sup_t |g(t) - \hat{g}_Q(t)| \leq 2\varepsilon \quad \text{and thus} \quad \sup_t |G(t) - \hat{G}_Q(t)| \leq 2\varepsilon(b - a + 2\gamma). \quad (\text{A.20})$$

We have assumed from the outset that $c < 1$. By construction:

$$\|\hat{F}_{\hat{H}_n, \hat{c}_n; Q_n} - F_n\|_{Q_n} \leq \|\hat{F}_{H_{T_n}, \hat{c}_n; Q_n} - F_n\|_{Q_n} \leq \|\hat{F}_{H_{T_n}, \hat{c}_n; Q_n} - F_{H_{T_n}, \hat{c}_n}\|_{Q_n} + \|F_{H_{T_n}, \hat{c}_n} - F_n\|_{Q_n}.$$

It follows from the proof of Corollary 5.1 that $\|F_{H_{T_n}, \hat{c}_n} - F_n\|_{Q_n} \rightarrow 0$ a.s. So if we can show that $\|\hat{F}_{H_{T_n}, \hat{c}_n; Q_n} - F_{H_{T_n}, \hat{c}_n}\|_{Q_n} \rightarrow 0$, it follows that $\|\hat{F}_{\hat{H}_n, \hat{c}_n; Q_n} - F_n\|_{Q_n} \rightarrow 0$ a.s.

For any probability measure \tilde{H} , any $\tilde{c} > 0$, and any $\lambda \in (0, +\infty)$, let:

$$\check{m}_{F_{\tilde{H}}, \tilde{c}}(\lambda) = \lim_{z \in \mathbb{C}^+ \rightarrow \lambda} m_{F_{\tilde{H}}, \tilde{c}}(z).$$

Also let $f_{\tilde{H}, \tilde{c}}(\lambda) \equiv \pi^{-1} \text{Im}[\check{m}_{F_{\tilde{H}}, \tilde{c}}(\lambda)]$ and define $f_{\tilde{H}, \tilde{c}}(0) \equiv 0$. Then:

$$\int_{-\infty}^t f_{\tilde{H}, \tilde{c}}(\lambda) d\lambda = \begin{cases} F_{\tilde{H}, \tilde{c}}(t) & \text{if } \tilde{c} < 1 \\ \tilde{c} \underline{F}_{\tilde{H}, \tilde{c}}(t) & \text{if } \tilde{c} > 1 \end{cases}.$$

We know that $f \equiv f_{H, c}$ is continuous, and therefore Lipschitz-continuous, on $[\tilde{z}_1, \tilde{z}_2]$ and constantly equal to zero outside $[\tilde{z}_1, \tilde{z}_2]$. Denote by f_{\max} the maximum value of f . Since $H_{T_n} \Rightarrow H$, it follows from part (i) of Proposition 4.2 that, for every $\tilde{\delta} > 0$:

$$f_{H_{T_n}, \hat{c}_n}(\lambda) \rightarrow f(\lambda) \quad \text{uniformly in } \lambda \in [\tilde{z}_1 + \tilde{\delta}, \tilde{z}_2 - \tilde{\delta}]. \quad (\text{A.21})$$

In particular, for every $\varepsilon > 0$, we can find N such that for all $n \geq N$:

$$|f_{H_{T_n}, \hat{c}_n}(\lambda) - f(\lambda)| < \varepsilon \quad \text{for all } \lambda \in [\tilde{z}_1 + \tilde{\delta}, \tilde{z}_2 - \tilde{\delta}].$$

For every n , the function $f_{H_{T_n}, \hat{c}_n}$ is monotonically increasing near the left boundary of its support and monotonically decreasing near the right boundary of its support; see Silverstein and Choi (1995, Section 5). The compact support of F is given by $[\tilde{z}_1, \tilde{z}_2]$. Lemma A.2 then implies that the support of $F_{H_{T_n}, \hat{c}_n}$ is contained in $[\tilde{z}_1 - \eta_n, \tilde{z}_2 + \eta_n]$ for some positive sequence $\eta_n \rightarrow 0$, so:

$$f_{H_{T_n}, \hat{c}_n}(\lambda) = 0 \quad \text{for } \lambda \notin [\tilde{z}_1 - \eta_n, \tilde{z}_2 + \eta_n]. \quad (\text{A.22})$$

And further, for η_n and $\tilde{\delta}$ sufficiently small and for n sufficiently large, we may assume that:

$$f_{H_{T_n}, \hat{c}_n}(\lambda) \leq 2f_{max} \quad \text{for all } \lambda \in [\tilde{z}_1 - \eta_n, \tilde{z}_1 + \tilde{\gamma}_n] \cup [\tilde{z}_2 - \tilde{\gamma}_n, \tilde{z}_2 + \eta_n] . \quad (\text{A.23})$$

Since f is Lipschitz-continuous on $[\tilde{z}_1, \tilde{z}_2]$, for $\varepsilon > 0$, there exists $\delta > 0$ such that $|f(\lambda_1) - f(\lambda_2)| \leq \varepsilon/2$ for all $\lambda_1, \lambda_2 \in [\tilde{z}_1, \tilde{z}_2]$ with $|\lambda_1 - \lambda_2| < \delta$. From (A.21) it then follows that for n large enough that:

$$|f_{H_{T_n}, \hat{c}_n}(\lambda_1) - f_{H_{T_n}, \hat{c}_n}(\lambda_2)| \leq \varepsilon \quad \text{for all } \lambda_1, \lambda_2 \in [\tilde{z}_1 + \tilde{\delta}, \tilde{z}_2 - \tilde{\delta}] \text{ with } |\lambda_1 - \lambda_2| \leq \delta .$$

Applying the previous discussion for a general c.d.f. G and a general grid Q leading to (A.20) to the special cases of $F_{H_{T_n}, \hat{c}_n}$ and Q_n , respectively, we thus obtain that for n large enough (in particular, satisfying $\gamma_n \leq \delta$):

$$\sup_{\lambda \in [\tilde{z}_1 + \tilde{\delta}, \tilde{z}_2 - \tilde{\delta}]} |f_{H_{T_n}, \hat{c}_n}(\lambda) - \hat{f}_{H_{T_n}, \hat{c}_n; Q_n}(\lambda)| \leq 2\varepsilon . \quad (\text{A.24})$$

Combining (A.22)–(A.24) yields, for ε and $\tilde{\delta}$ small enough and for n large enough:

$$\sup_{\lambda \in \mathbb{R}} |F_{H_{T_n}, \hat{c}_n}(\lambda) - \hat{F}_{H_{T_n}, \hat{c}_n; Q_n}(\lambda)| \leq 2\varepsilon (\tilde{z}_2 - \tilde{z}_1 + 2\tilde{\delta}) + 4f_{max}(\eta_n + \tilde{\delta}) . \quad (\text{A.25})$$

Since the right hand side of (A.25) can be made arbitrarily small, we have established that $\|\hat{F}_{H_{T_n}, \hat{c}_n; Q_n} - F_{H_{T_n}, \hat{c}_n}\|_{Q_n} \rightarrow 0$, which implies that $\|\hat{F}_{\hat{H}_n, \hat{c}_n; Q_n} - F_n\|_{Q_n} \rightarrow 0$ a.s., which, in return, implies that:

$$\|\hat{F}_{\hat{H}_n, \hat{c}_n; Q_n} - F\|_{Q_n} \leq \|\hat{F}_{\hat{H}_n, \hat{c}_n; Q_n} - F_n\|_{Q_n} + \|F_n - F\|_{Q_n} \rightarrow 0 \text{ a.s.} . \quad (\text{A.26})$$

Lemma A.1 then tells us that $\hat{F}_{\hat{H}_n, \hat{c}_n; Q_n} \Rightarrow F$ a.s.

We now show that this implies part (ii) of the corollary, namely that $\hat{H}_n \Rightarrow H$ a.s. by means of contradiction. To this end, assume that $\hat{H}_n \Rightarrow H$ a.s. is not the case. The sequence $\{\hat{H}_n\}$ is tight a.s. This is because the upper bound of the support of H_n is given by K_n/T_n which, by definition of K_n satisfies $K_n/T_n \leq \lambda_p + 1/T_n$; and we know from Bai and Silverstein (1998) that for any $\varepsilon > 0$, $\lambda_p \leq \tilde{z}_2 + \varepsilon$ for n large enough a.s. Similar for the lower bound, or simply use zero as very crude lower bound. Therefore, if $\hat{H}_n \Rightarrow H$ a.s. is not the case, there then exists a probability measure $H' \neq H$ and a subsequence $\{n_k\}$ such that on a set with positive probability, we have $\hat{H}_{n_k} \Rightarrow H'$.

Similarly to an argument used in the proof of part (i) of Corollary 5.1 — with \hat{H}_{n_k} and H' now playing the roles of H_{T_n} and H , respectively — it then follows that $\|F_{\hat{H}_{n_k}, \hat{c}_{n_k}} - F_{H', c}\|_{Q_{n_k}} \rightarrow 0$ on a set with positive probability. But it also holds that $\|\hat{F}_{\hat{H}_{n_k}, \hat{c}_{n_k}; Q_{n_k}} - F_{\hat{H}_{n_k}, \hat{c}_{n_k}}\|_{Q_{n_k}} \rightarrow 0$ similarly to an argument used above — with $F_{\hat{H}_{n_k}, \hat{c}_{n_k}}$ now playing the role of $F_{H_{T_n}, \hat{c}_n}$. Together, we obtain that $\|\hat{F}_{\hat{H}_{n_k}, \hat{c}_{n_k}; Q_{n_k}} - F_{H', c}\|_{Q_{n_k}} \rightarrow 0$ on a set with positive probability. Since we are working under the assumption that $c < 1$, both F_H and $F_{H'}$ are continuous. Lemma A.1 then

tells us that $\sup_t |\hat{F}_{\hat{H}_{n_k}, \hat{c}_{n_k}; Q_{n_k}}(t) - F_{H', c}(t)| \rightarrow 0$ on a set with positive probability. But this is in contradiction to $\sup_t |\hat{F}_{\hat{H}_n, \hat{c}_n; Q_n} - F(t)| \rightarrow 0$ a.s. So the proof of part (ii) is accomplished.

We now can establish that $\|\hat{F}_{\hat{H}_n, \hat{c}_n; Q_n} - F_{\hat{H}_n, \hat{c}_n}\|_{Q_n} \rightarrow 0$ a.s., knowing that $\hat{H}_n \Rightarrow H$ a.s., very much in the same way as we established before that $\|\hat{F}_{H_{T_n}, \hat{c}_n; Q_n} - F_{H_{T_n}, \hat{c}_n}\|_{Q_n}$, knowing that $H_{T_n} \Rightarrow H$. As a result, we obtain that $\|F_{\hat{H}_n, \hat{c}_n} - F\|_{Q_n} \rightarrow 0$ a.s. Invoking Lemma A.1 establishes part (i) then.

Parts (iii)–(iv) follow immediately from parts (i)–(ii) and Proposition 4.3, part (ii). ■

B Figures

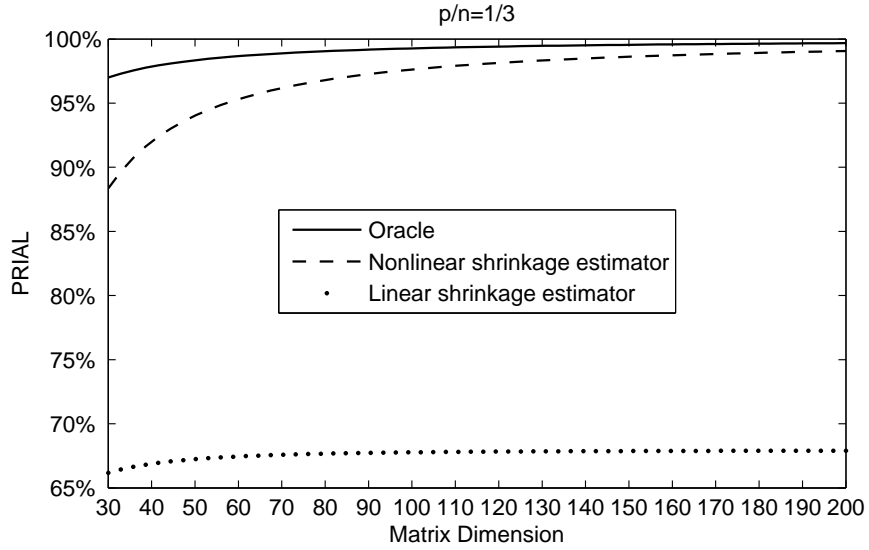


Figure 1: Comparison of the NonLinear vs. Linear Shrinkage Estimators. 20% of population eigenvalues are equal to 1, 40% are equal to 3, and 40% are equal to 10. Every point is the result of 1,000 Monte Carlo simulations.

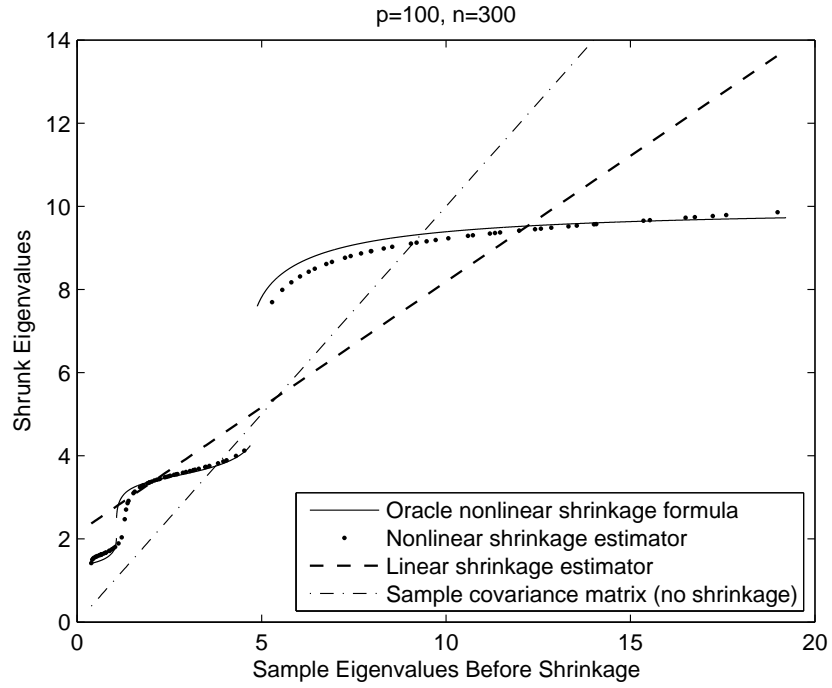


Figure 2: Nonlinearity of the Oracle Shrinkage Formula. 20% of population eigenvalues are equal to 1, 40% are equal to 3, and 40% are equal to 10. $p = 100$ and $n = 300$.

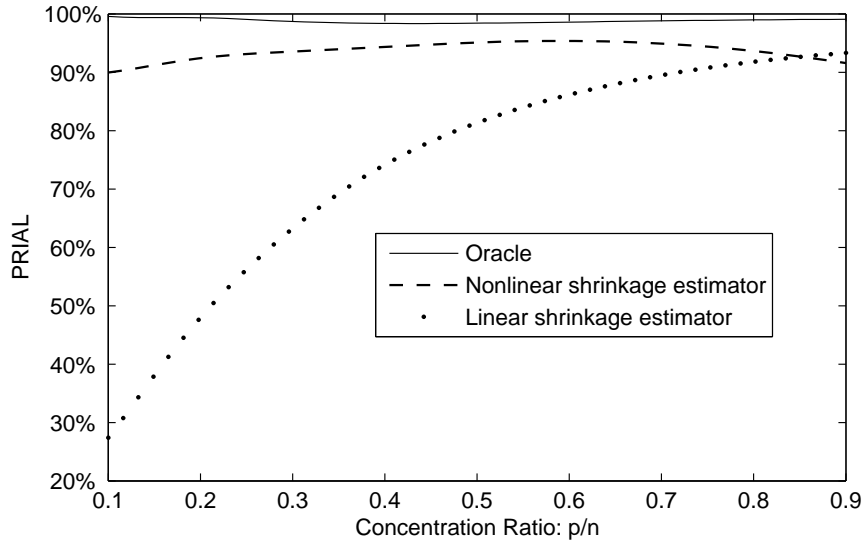


Figure 3: Effect of Varying the Concentration Ratio $\hat{c}_n = p/n$. 20% of population eigenvalues are equal to 1, 40% are equal to 3, and 40% are equal to 10. Every point is the result of 1,000 Monte Carlo simulations.

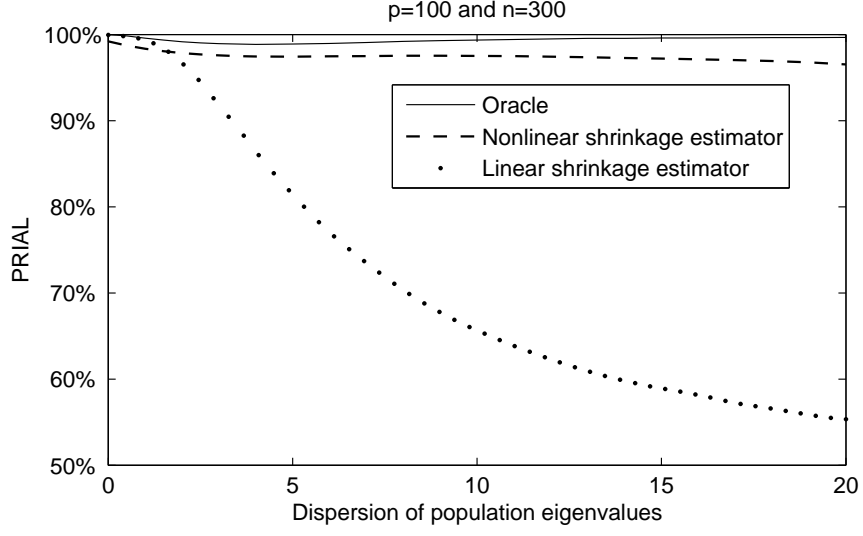


Figure 4: Effect of Varying the Dispersion of Population Eigenvalues. 20% of population eigenvalues are equal to 1, 40% equal to $1 + 2d/9$, and 40% equal to $1 + d$, where the dispersion parameter d varies from 0 to 20. $p = 100$ and $n = 300$. Every point is the result of 1,000 Monte Carlo simulations.

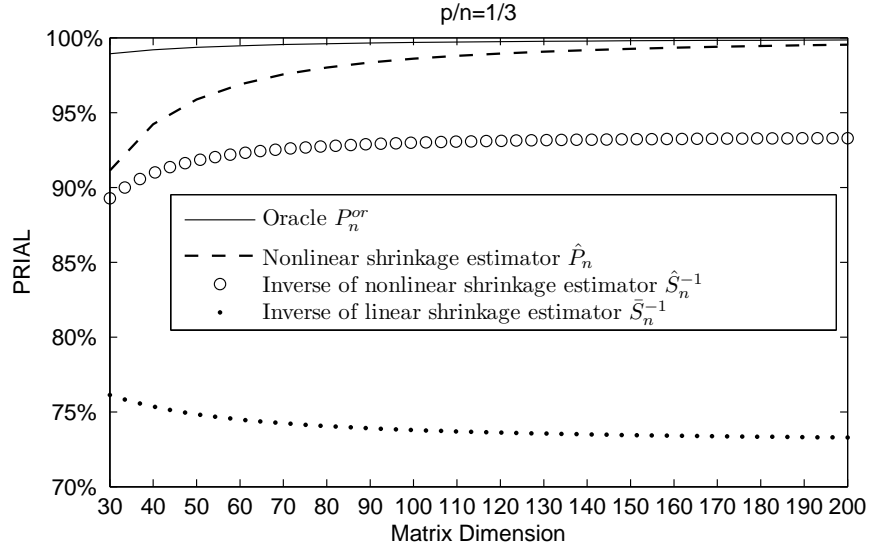


Figure 5: Estimating the Precision Matrix. 20% of population eigenvalues are equal to 1, 40% are equal to 3, and 40% are equal to 10. Every point is the result of 1,000 Monte Carlo simulations.

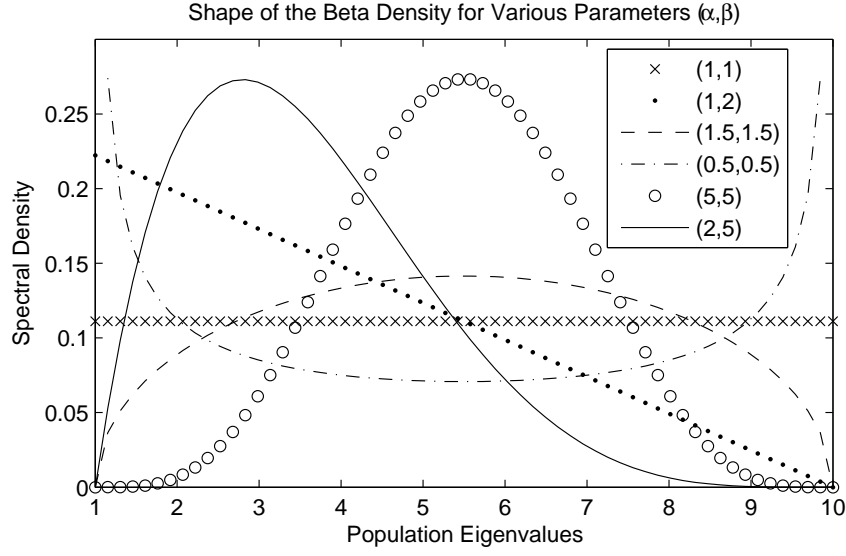


Figure 6: Shape of the Beta Density for Various Parameter Values. The support of the beta density has been shifted to the interval $[1, 10]$ by a linear transformation. In order to enhance clarity, the densities corresponding to the parameters $(2, 1)$ and $(5, 2)$ have been omitted, because they are symmetric to $(1, 2)$ and $(2, 5)$ respectively about the mid-point of the support.

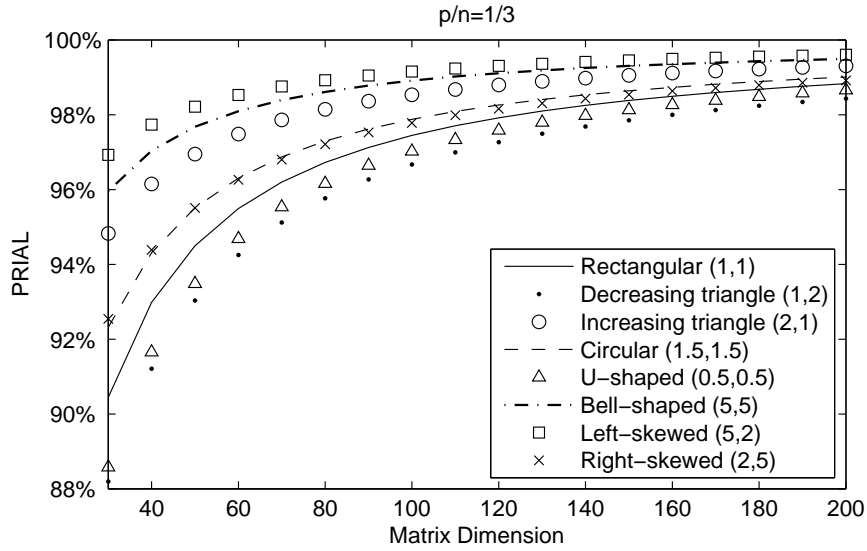


Figure 7: Performance of the Nonlinear Shrinkage with Beta Densities. The various curves correspond to different shapes of the population spectral density. The support of the population spectral density is $[1, 10]$.

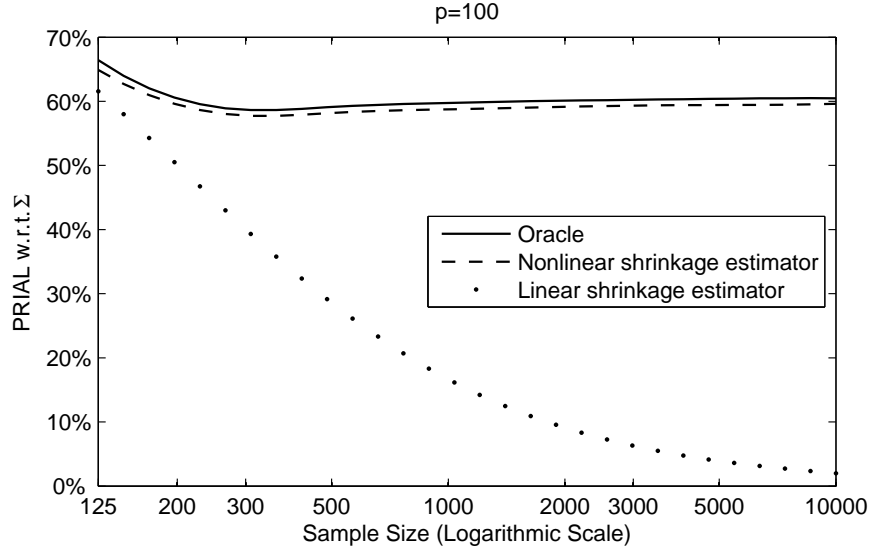


Figure 8: Fixed-Dimension Asymptotics. 20% of population eigenvalues are equal to 1, 40% are equal to 3, and 40% are equal to 10. Variables are normally distributed. Every point is the result of 1,000 Monte Carlo simulations.

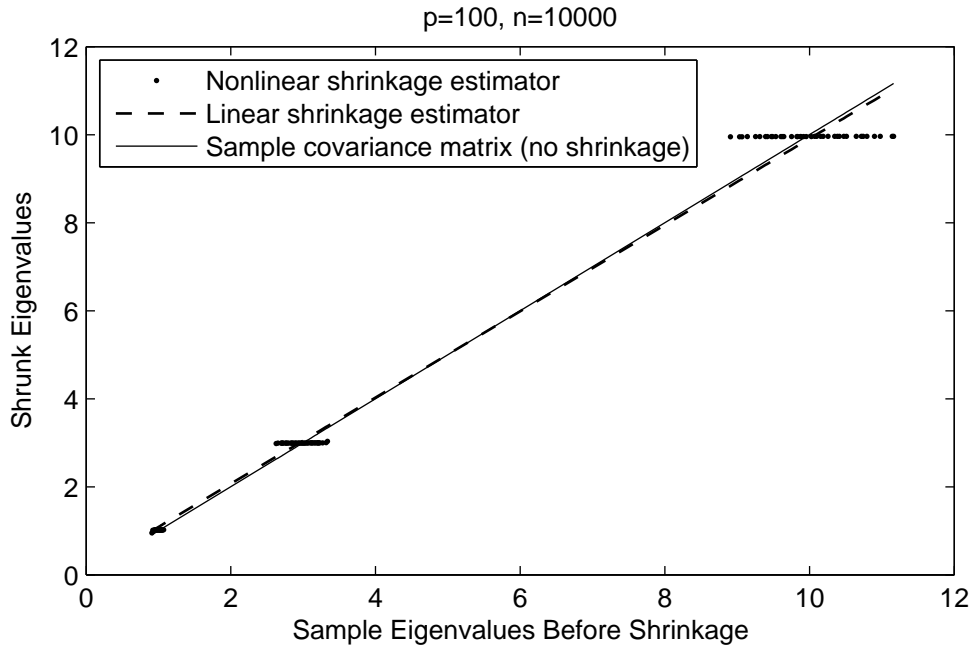


Figure 9: Nonlinear Shrinkage under Fixed-Dimension Asymptotics. 20% of population eigenvalues are equal to 1, 40% are equal to 3, and 40% are equal to 10. $p = 100$ and $n = 10,000$. The oracle is not shown because it is virtually identical to the nonlinear shrinkage estimator.