

# ScreeNOT: EXACT MSE-OPTIMAL SINGULAR VALUE THRESHOLDING IN CORRELATED NOISE

BY DAVID DONOHO<sup>1,a</sup>, MATAN GAVISH<sup>2,c</sup> AND ELAD ROMANOV<sup>1,b</sup>

<sup>1</sup>Department of Statistics, Stanford University, <sup>a</sup>[donoho@stanford.edu](mailto:donoho@stanford.edu), <sup>b</sup>[eromanov@stanford.edu](mailto:eromanov@stanford.edu)

<sup>2</sup>School of Computer Science and Engineering, Hebrew University of Jerusalem, <sup>c</sup>[gavish@cs.huji.ac.il](mailto:gavish@cs.huji.ac.il)

We derive a formula for optimal hard thresholding of the singular value decomposition in the presence of correlated additive noise; although it nominally involves unobservables, we show how to apply it even where the noise covariance structure is not a priori known or is not independently estimable. The proposed method, which we call *ScreeNOT*, is a mathematically solid alternative to Cattell’s ever-popular but vague scree plot heuristic from 1966. *ScreeNOT* has a surprising oracle property: it typically achieves *exactly*, in large finite samples, the lowest possible MSE for matrix recovery, on each given problem instance, that is, the specific threshold it selects gives exactly the smallest achievable MSE loss among all possible threshold choices for *that* noisy data set and *that* unknown underlying true low rank model. The method is computationally efficient and robust against perturbations of the underlying covariance structure. Our results depend on the assumption that the singular values of the noise have a limiting empirical distribution of compact support; this property, which is standard in random matrix theory, is satisfied by many models exhibiting either cross-row correlation structure or cross-column correlation structure, and also by many situations with more general, interelement correlation structure. Simulations demonstrate the effectiveness of the method even at moderate matrix sizes. The paper is supplemented by ready-to-use software packages implementing the proposed algorithm: package *ScreeNOT* in Python (via PyPI) and R (via CRAN).

**1. Introduction.** Across a wide variety of scientific and technical fields, practitioners have found many valuable applications of *singular value thresholding* (SVT). This procedure starts from the singular value decomposition (SVD), which represents the data matrix  $Y$  as

$$(1.1) \quad Y = \sum_{i=1}^{\min(n,p)} y_i \cdot \mathbf{u}_i \mathbf{v}_i^\top,$$

using the empirical singular values  $\{y_i\}_{i=1}^{\min(n,p)}$ , and the empirical left and right singular vectors of  $Y$ , denoted here  $\mathbf{u}_i$  and  $\mathbf{v}_i$ .

In such applications, it is generally claimed that the small singular values represent “noise” and the large singular values “signal”; practitioners attempt to separate signal from noise by setting a threshold  $\theta$  (say), and using, in place of  $Y$ , the partial reconstruction containing only would-be signal components:

$$(1.2) \quad \hat{X}_\theta = \sum_i y_i \mathbb{1}_{\{y_i > \theta\}} \cdot \mathbf{u}_i \mathbf{v}_i^\top.$$

How do practitioners determine the threshold  $\theta$ ? Often, by eye. They plot the ordered singular values and spot “elbows.” Sometimes, they give this a scholarly veneer by saying

---

Received May 2021; revised July 2022.

*MSC2020 subject classifications.* Primary 62C20, 62H25; secondary 90C25, 90C22.

*Key words and phrases.* Singular value thresholding, optimal threshold, scree plot, low-rank matrix denoising, high-dimensional asymptotics.

they are using the “scree plot method”; they might even formally cite the originator of this folk-tradition [12], which still gets more than 1000 citations yearly. According to the method prescribed in that paper, the practitioner plots the values  $\{y_i\}$  and uses her *eyes* to distinguish between “signal” and “noise” singular values of  $Y$ .

How *should* they determine the threshold? Relevant theory and methodology literature spans multiple disciplines over multiple decades; we mention only a few entry points, including: [1–3, 11, 13, 14, 17, 19, 20, 22, 24, 25, 28, 30, 31, 34]. Progress has been made in our understanding of the underlying problem, and many valuable quantitative approaches have been developed—to which we here add one more. Our contribution relies on recent advances in random matrix theory which point, we think convincingly, to the method introduced here. This method typically offers the exact optimal loss available on each specific, finite data set  $Y$ .

Our task formalization supposes that: (a) there is an underlying matrix  $X$  of fixed rank  $r$ —though  $X$  and even its rank  $r$  are unknown to us; (b) only a potentially loose upper bound on the signal rank  $r$  is known; (c) the data matrix  $Y$  has the signal + noise form  $Y = X + Z$ , where  $Z$  is a noise matrix with a general covariance structure—also unknown to us; (d) we use hard thresholding of singular values, exactly as in (1.2) above;<sup>1</sup> (e) we adopt squared error *loss*:<sup>2</sup>

$$(1.3) \quad \text{SE}[X|\theta] = \|\hat{X}_\theta - X\|_F^2.$$

As goal, we literally aim to choose a loss-minimizing value  $\theta_{\text{opt}} = \theta_{\text{opt}}(Y|X)$  solving

$$(1.4) \quad \text{SE}[X|\theta_{\text{opt}}] = \min_{\theta} \text{SE}[X|\theta].$$

Aiming for  $\theta_{\text{opt}}(Y|X)$  may seem overambitious, as we know only the data matrix  $Y$ , and not  $X$ ; wait and see.

Essentially this problem was studied previously by two of the authors in the special case of white noise. [19] supposed that the underlying noise  $Z$  matrix has i.i.d. Gaussian zero-mean entries and the problem is scaled so that the columns of  $Z$  have unit Euclidean squared norm in expectation, and considered a sequence of increasingly large problems. In the square case, when  $Y$  has as many rows as columns: the authors found results<sup>3</sup> which, in light of our results below, say that with eventually overwhelming probability, we have  $\theta_{\text{opt}} = 4/\sqrt{3}$ . Their analysis relied on then-recent advances in the “Johnstone spiked model” of random matrix theory [23]; they proposed a method for white noise with unknown variance, where the threshold formula became  $\theta_{\text{opt}} \approx 4/\sqrt{3} \cdot \frac{y_{\text{med}}}{\sqrt{n \cdot 0.6528}}$ , where  $y_{\text{med}}$  denotes the median empirical singular value of  $Y$ .<sup>4</sup>

Understanding the white noise case cannot be the end of the story. Practitioners ordinarily do not know that their noise is white, and in fact realistic noise models can include correlations between columns, rows or even general row-column combinations. Fortunately, a broad range of noise models can be studied using appropriate advances that have been made in random matrix theory. In this broader context, as we show a more general formula for the optimal threshold can be given, which of course reduces to  $4/\sqrt{3}$  in the above “square-matrix in white noise” case, but which is inevitably quite a bit more sophisticated in general.

<sup>1</sup>And not some variant, such as soft thresholding or a more general shrinkage.

<sup>2</sup> $\|X\|_F^2 = \sum_{i,j} X_{i,j}^2$  denotes the squared Frobenius norm.

<sup>3</sup>That is, the authors of [19] adopted a slightly different viewpoint involving asymptotic MSE, and showed that  $4/\sqrt{3}$  is optimal, whereas we consider here exact finite sample MSE loss, and show that with eventually overwhelming probability,  $4/\sqrt{3}$  is exactly optimal *on each typical realization*.

<sup>4</sup> $\sqrt{0.6528}$  is approximately the median of the standard quarter-circle law; see the original paper.

Section 2 below describes *ScreeNOT*, our proposed deployment of this formula on actual data. The acronym NOT stands for *Noise-adaptive Optimal Thresholding*; “adaptive” refers to the algorithm’s optimality across a wide range of unknown noise covariances. The prefix “Scree” reminds us that, still today, in many cases, the alternative would simply be “eyeballing” the scree plot [12]. Cattell and his many followers clearly believed that *something*, some visible feature, in the scree plot, namely, in the collection of data singular values  $\{y_i\}$ —could tell us where the noise stopped and the signal began. But what exactly? In a very concrete sense, the ScreeNOT algorithm shows that the information needed to separate signal from noise truly *is* there in the distribution of empirical singular values, where Cattell and his followers all hoped it would be. However, the ScreeNOT algorithm and the approach we develop here quantitatively identify this information as a specific *functional of the CDF of singular values*.

The method, once implemented, surprised us by the *finite-sample* optimality it exhibited; in simulations at reasonable problem sizes it typically achieves the exact minimal loss (1.4) for the given data set, even though the method is not entitled to know the underlying low-rank model  $X$  or specifics of the noise model on  $Z$ ; we initially expected a weaker and more “asymptotic” optimality property, perhaps similar to the one shown in [19]. Our analysis below proves typicality of such exact optimality in finite samples. This strong optimality is partly due to the penetrating nature of random matrix theory; but also to the very specific task: minimizing squared error loss (1.3) of singular value thresholding (1.2).

*Underlying analysis.* Hoping to make the paper helpful to prospective users of the proposed method, we have made the [Introduction](#) and also Section 2 mostly independent of the analysis to come; however, we now very briefly offer mathematically-oriented readers some insight about the approach being followed in later sections and the tools being developed there.

At heart, this paper concerns the asymptotic analysis of a sequence of matrix recovery problems where the problem sizes  $n$  and  $p$  grow to  $\infty$  in a proportional fashion. We assume that the matrix  $X$  has  $r$  nonzero singular values  $x_1, \dots, x_r$ , which are fixed independently of  $n$  and  $p$ . About the sequence of random noise matrices  $Z = Z_{n,p}$ , we assume that the sequence of empirical cumulative distribution functions (CDFs) of noise singular values converges to a compactly supported distribution  $F_Z$  with certain qualitative restrictions at the boundary of the support.

Using results of Benaych-Georges and Nadakuditi [10], we obtain an expression for an asymptotically optimal hard threshold, as a *functional*  $T(\cdot)$  of the limiting CDF of noise singular values  $F_Z$ . The functional is continuous and even differentiable in certain senses.

Admittedly, the limiting CDF of noise singular values  $F_Z$  is not observable to the statistician, as we only observe a sample of the signal + noise singular values mixed together. Performing a kind of amputation and prosthetic extension on the CDF  $F_Y$  of singular values of  $Y$ , which we do observe, we construct a modified empirical CDF  $\hat{F}_n$ , which consistently estimates the limiting CDF of noise-only singular values. Applying the hard threshold selection functional to this modified empirical CDF  $\hat{F}_n$  gives our proposed method, in the form  $\hat{\theta} = T(\hat{F}_n)$ . As we show in Section 2, there is a quite explicit and computationally tractable algorithm for computing  $T(\hat{F}_n)$ , which we label *ScreeNOT*.

Owing to the continuity of the hard threshold functional  $T(\cdot)$ , and the consistency of the constructed CDF, the resulting method is a consistent estimator of the underlying asymptotically optimal threshold  $T(F_Z)$ . We also prove a finite-sample optimality of the method. Specifically, the ScreeNOT algorithm is shown to be exactly optimal for squared error loss with high probability, in large enough finite samples, under very general model assumptions. For generic configurations of signal singular values  $(x_i)_{i=1}^r$ , there is, in large finite samples, an *optimal interval* of thresholds, all achieving the optimal MSE at that realization; the con-

sistency of the optimal threshold estimator implies that eventually for large enough  $n$ , with overwhelming probability, the proposed method achieves the exact optimal MSE loss.

*Contributions.* The approach we develop selects an optimal threshold for singular values, and thus selects the “signal” singular values, based on the principle of minimizing SE. Indeed, minimizing squared error loss is a ubiquitous goal in statistical theory, and we are not the first to consider it as a goal for threshold selection in the context of singular values. In addition to our own just-cited work [19], prior citable work on squared error loss includes Perry [30], Shabalin and Nobel [32] and Nadakuditi [27], although much of this concerns singular value shrinkage rather than thresholding.<sup>5</sup> While our approach is implemented for SE loss, we note that it could in principle be used to develop optimal thresholding rules for other loss criteria, such as the operator norm.

We especially point to [14]; in this work, Dobriban and Owen mainly study Parallel Analysis [18]—simulation-based significance testing for large singular values; they develop tools from random matrix theory to derandomize parallel analysis.<sup>6</sup> Beyond this, they also mention in a final section that their tools could be adapted to produce a threshold selector minimizing the asymptotic mean squared error of the resulting approximation; and their equation (6) provides a way to characterize such a functional.

In this paper, we make explicit (in equation (4.5) below) the functional  $T$  for threshold selection with minimal asymptotic SE, and show that it is well-defined; we offer (in Section 2.3) an explicit construction of a modified CDF  $\hat{F}_n$  to plug into  $T$ —the proposal involves singular value “ablation and prosthesis”; we develop a theoretical machinery, involving continuity properties of  $T$  and convergence properties of  $\hat{F}_n$ , and use the machinery to prove that  $T(\hat{F}_n)$  achieves not just asymptotic optimal loss (Theorem 1), but also (Theorem 2) that in large finite samples it achieves the exact minimal loss with overwhelming probability.

*Outline.* This paper is organized as follows. In Section 2, we offer a practical, succinct description of the ScreeNOT algorithm, for the convenience of prospective users. In Section 3, we introduce the signal + noise model used and survey relevant results from random matrix theory. In Section 4, we state our main results regarding the optimality and stability properties of ScreeNOT, both in finite matrix size and asymptotically as the matrix size grows to infinity. In Section 5, we demonstrate the mathematical results in various simulations and numerical examples; for space considerations, only a handful of figures are shown, with most simulation results deferred to the Supplementary Material and available in the code supplement [15]. The results are proved in Section 6, with some proofs referred to the Supplementary Material [16].

*Reproducibility advisory.* Implementation of the proposed algorithm, scripts generating all figures in this paper and many additional simulations have been permanently deposited and are available at the code supplement [15].

*Code packages.* Ready-to-use code packages, implementing the ScreeNOT procedure in various language are available. In Python: package ScreeNOT is available through PyPI; in R: package ScreeNOT is available through CRAN; and Matlab source code. For details, see the following GitHub repository: <https://github.com/eladromanov/ScreeNOT>. In addition, the source code has been permanently deposited and is available at the code supplement [15].

---

<sup>5</sup>Singular value shrinkage is considerably more involved as it changes the data singular values rather than selects them. The best-possible relative improvement of shrinkage over thresholding was studied in [19] in the white noise case.

<sup>6</sup>Potential users of ScreeNOT should note that in many scientific projects the goal is determining the number of statistically significant factors, without particular regard to the quality of squared-error approximation; it is possible that for such a goal parallel analysis or its derandomized version [14] is preferred.

**2. The ScreeNOT procedure: User-level description.** In this section, we give a brief self-contained description of our proposed procedure.

2.1. *Procedure API.* ScreeNOT selects a hard threshold for singular values, which can in finite samples give the optimal MSE approximation of a low-rank matrix from a noisy version; the noise may be correlated, and the threshold will adapt to that appropriately.

2.1.1. *Inputs.* The user provides these inputs to ScreeNOT:

$\mathbf{y}$ : the singular values  $y_1, \dots, y_{\min(n,p)}$  of the data matrix  $Y$ ;

$n, p$ : size parameters of the data matrix  $Y$ .

$k$ : upper bound on the rank  $r$  of the underlying unknown signal matrix  $X$ , which is to be recovered. This upper bound may be very loose.

2.1.2. *Outputs.* ScreeNOT returns  $\hat{\theta} = \hat{\theta}(Y)$ , the value to be used in singular value thresholding.

To use the threshold, the user should reconstruct an approximation to the underlying signal matrix  $X$  using the empirical singular values  $y_i$  and the empirical singular vectors  $\mathbf{u}_i$  and  $\mathbf{v}_i$  as follows:

$$\hat{X} = \sum_i y_i \mathbb{1}_{\{y_i > \hat{\theta}\}} \cdot \mathbf{u}_i \mathbf{v}_i^\top.$$

In this reconstruction, the singular values smaller than  $\hat{\theta}$  are judged to be noise and the corresponding singular decomposition components are ignored.

2.2. *Example in a stylized application.* We next construct a synthetic-data example, in which we know the ground truth for demonstration purposes. The synthetic data  $Y = X + Z$  and the invocation of ScreeNOT are based on these ingredients.

Signal  $X$ : The underlying signal matrix, unbeknownst to the hypothetical user, has rank 10, with singular values  $(x_{10}, \dots, x_1) = (1.0, 1.15, 1.3, \dots, 2.35)$ .

Noise  $Z$ : The underlying noise, unbeknownst to the hypothetical user, follows an AR(1) process in the row index, within each column. The AR(1) process has parameter  $\rho = 0.4$ , and additionally each entry is divided by  $\sqrt{n}$ , so to have variance<sup>7</sup>  $1/n$ .

Problem size:  $n = p = 1000$

Rank bound:  $k = 15$ . The user specifies a bound of  $k = 15$  on the possible rank of the signal.

Figure 1 shows a so-called *scree plot* [12] of the first 30 empirical singular values  $y_i$ . For this particular instance, it is verified (by exhaustive search) that the minimal loss is attained by retaining the first three principal components of  $Y$ ; in other words, thresholding at any point  $\theta \in (y_3, y_4)$  is optimal. The threshold  $\hat{\theta}$  returned by the ScreeNOT procedure is indicated by the green horizontal line and indeed it falls inside the optimal interval. The would-be “elbow” in the scree plot, determined subjectively by the authors, is indicated by the grey (lower) horizontal line; it corresponds to retaining the top 6 principal components of the data matrix. This rule attains strictly suboptimal SE: roughly 31.3, as opposed by 26.4 attained by ScreeNOT.

<sup>7</sup>That is, the columns of  $Z$  are independent and distributed as  $\mathbf{z}/\sqrt{n}$ , where the random vector  $\mathbf{z}$  has entries:  $z_1 = \varepsilon_1$  and  $z_i = \rho \cdot z_{i-1} + \sqrt{1 - \rho^2} \cdot \varepsilon_i$  for  $2 \leq i \leq p$ , where  $\varepsilon_1, \dots, \varepsilon_p \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$ .

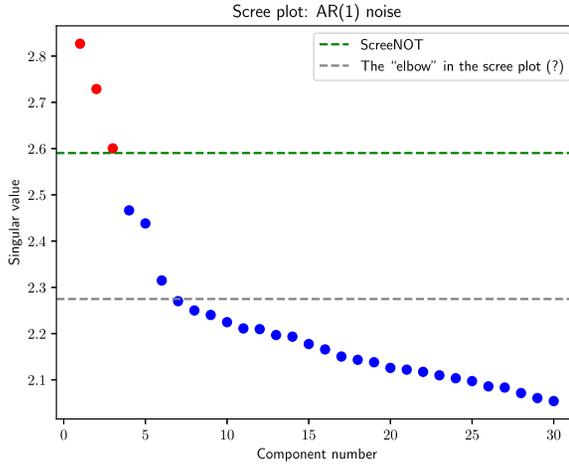


FIG. 1. Scree plot for the stylized example of Section 2.2. Horizontal axis: singular value index, where the singular values  $y_i$  of the data matrix  $Y$  are sorted in decreasing order. Vertical axis: singular values  $y_i$ . Dashed green (upper) line: the optimal threshold calculated by ScreeNOT (color online).

2.3. Internals of the procedure. We briefly describe the computational task performed by ScreeNOT.

Step 1. Sort the singular values in nonincreasing order:  $y_1 \geq \dots \geq y_p$ .

Step 2. Compute the “pseudo singular values”:

$$\tilde{y}_i = y_{k+1} + \frac{1 - \left(\frac{i-1}{k}\right)^{2/3}}{2^{2/3} - 1} (y_{k+1} - y_{2k+1}) \quad \text{for } i = 1, \dots, k,$$

and set  $\tilde{y}_i = y_i$  for  $i = k + 1, \dots, p$ .<sup>8</sup>

Step 3. Define the four scalar functions  $\varphi, \tilde{\varphi}, \varphi', \tilde{\varphi}'$  by

$$\varphi(y) = \frac{1}{p} \sum_{i=1}^n \frac{y}{y^2 - \tilde{y}_i^2}, \quad \varphi'(y) = -\frac{1}{p} \sum_{i=1}^n \frac{y^2 + \tilde{y}_i^2}{(y^2 - \tilde{y}_i^2)^2},$$

and

$$\tilde{\varphi}(y) = \gamma \varphi(y) + \frac{1 - \gamma}{y}, \quad \tilde{\varphi}'(y) = \gamma \varphi'(y) - \frac{1 - \gamma}{y^2}.$$

Now define

$$\Psi(y) = y \cdot \left( \frac{\varphi'(y)}{\varphi(y)} + \frac{\tilde{\varphi}'(y)}{\tilde{\varphi}(y)} \right).$$

Step 4. Assuming that  $\tilde{y}_1, \dots, \tilde{y}_p$  are not all zero, the function  $y \mapsto \Psi(y)$  can be shown to be continuous and strictly increasing for  $y > \tilde{y}_1$ . Moreover,  $\lim_{y \searrow \tilde{y}_1} \Psi = -\infty$  and  $\Psi(\infty) = -2$ . The computed hard threshold is the unique value  $\hat{\theta}$  satisfying

$$(2.1) \quad \Psi(\hat{\theta}) = -4.$$

This equation is then solved numerically, for example, by binary search.<sup>9</sup>

Step 5. The algorithm returns the value  $\hat{\theta}$ .

<sup>8</sup>We assume  $2k + 1 < p$ . Our proposed estimator is expected to perform poorly when  $k$  is large compared to  $p$ .

<sup>9</sup>We remark that the computational cost of this search is not considerable. For example, binary search requires a number of iterations, which is only logarithmic in the required precision.

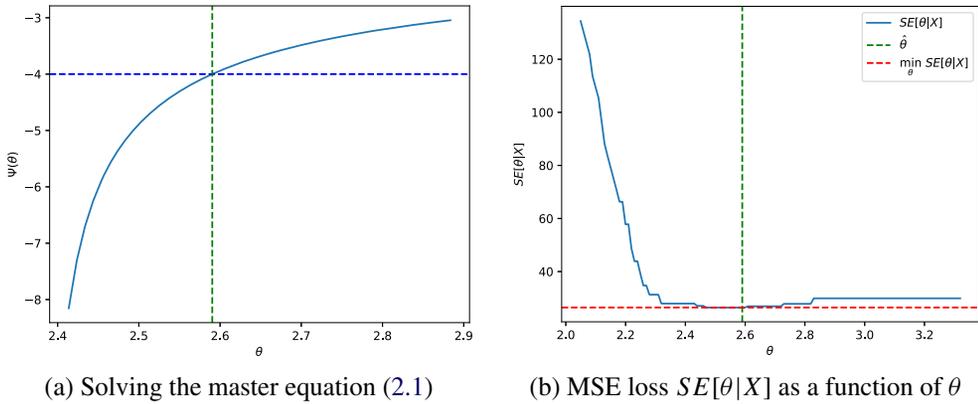


FIG. 2. Calculation of the optimal threshold  $\hat{\theta}$  on the stylized example of Section 2.2. (a) Left panel: Horizontal axis: candidate thresholds  $\theta$ . Vertical axis: The function  $\Psi(\theta)$ . The *ScreenNOT* algorithm solves the equation  $\Psi(\theta) = -4$  for  $\theta$ . Vertical green line shows the solution, denoted by  $\hat{\theta}$ . This is the value returned by *ScreenNOT*. (b) Right panel: The MSE loss function  $SE[\theta|X]$  for the stylized example of Section 2.2, plotted over candidate thresholds  $\theta$ . There is an interval of values  $\theta$  all achieving the lowest possible loss; the threshold  $\hat{\theta}$  returned by *ScreenNOT* (shown by the vertical green line) is located inside this optimal interval (color online).

Evidently, the procedure as stated costs  $O(n \log(n))$  flops; the dominant cost is sorting the singular values; ordinarily of course, sorting is performed anyway as part of a standard SVD. In that situation, the additional computational effort is  $O(n)$ , which is unimportant compared to the cost of the underlying SVD.

2.4. *How the procedure works on the stylized application.* Figure 2(a) shows a plot of  $\Psi(\theta)$  as a function of  $\theta$ . The horizontal blue line indicates the desired level  $-4$ . The vertical green line indicates the crossing point,  $\hat{\theta}$ , which is the value returned by *ScreenNOT*.

Figure 2(b) shows a plot of the loss  $SE[\theta|X]$  versus  $\theta$ . The red horizontal line shows the optimum achievable loss. The green vertical line shows the threshold selected by the procedure. It intersects the loss curve within the optimal level and the achieved loss is therefore optimal.

**3. Setup and background from random matrix theory.** The rest of this paper is dedicated to formal analysis of the *ScreenNOT* algorithm. To that end, we now define a precise signal + noise model and set up the necessary notation.

*An asymptotic model for low-rank matrices observed in additive noise.* To recap, let  $X_n$  be an unknown  $n$ -by- $p$  matrix, to be estimated. We observe a noisy measurement of  $X_n$ ,  $Y_n = X_n + Z_n$ , where  $Z_n$  is a noise matrix, which is statistically independent of  $X_n$ . Our analysis employs an asymptotic framework originating in random matrix theory, and considers a sequence of such problems  $n, p \rightarrow \infty$ , with the following generative assumptions.

1. *Limiting shape:* the dimensions  $n, p$  tend to infinity together at a fixed ratio  $p/n \rightarrow \gamma$ . More concretely, fix  $\gamma \in (0, 1]$  and set  $p = p_n = \lceil \gamma n \rceil$ . Denoting  $\gamma_n = p_n/n$ , of course,  $\gamma \leq \gamma_n < \gamma + \frac{1}{n}$  and  $\gamma_n \rightarrow \gamma$  as  $n \rightarrow \infty$ .

2. *Fixed signal rank and singular values:* The matrix  $X_n$  has fixed rank  $r = \text{rank}(X_n)$  and fixed singular values. Specifically, let  $r$  be constant, and fix  $r$  positive and distinct numbers  $x_1 > \dots > x_r > 0$ .  $X_n$  is the matrix

$$X_n = \sum_{i=1}^r x_i \mathbf{a}_{i,n} \mathbf{b}_{i,n}^\top,$$

where  $\mathbf{a}_{i,n} \in \mathbb{R}^n$  (resp.,  $\mathbf{b}_{i,n} \in \mathbb{R}^p$ ) for  $i = 1, \dots, r$  are sequences of left (resp., right) singular vectors of  $X_n$ , obeying a generative assumption as described next. We let  $\mathbf{x} = (x_1, \dots, x_r)$  denote the vector of singular values and we refer to either the matrix  $X$  or just  $\mathbf{x}$  as the *signal*.<sup>10</sup>

3. *Incoherent signal singular vectors.* The vectors  $\mathbf{a}_{1,n}, \dots, \mathbf{a}_{r,n}$  (resp.,  $\mathbf{b}_{i,n}$ ) constitute a random, uniformly distributed orthonormal  $r$ -frame in  $\mathbb{R}^n$  (resp., in  $\mathbb{R}^p$ ).<sup>11</sup>

4. *Compactly supported, limiting bulk distribution of noise singular values.* Each matrix  $Z_n$  is statistically independent of  $X_n$ . Let  $z_{1,n}, \dots, z_{p,n}$  denote its singular values, with empirical CDF  $F_{Z_n}$ ,  $F_{Z_n}(z) = p^{-1} \sum_{i=1}^p \mathbb{1}_{\{z_{i,n} \leq z\}}$ . There is a *limiting empirical CDF* (LECDF)  $F_Z$  such that  $F_{Z_n} \rightarrow F_Z$  a.s. at continuity points.

Moreover, we assume that  $F_Z$  is compactly supported<sup>12</sup> and denote the upper edge of the support (sometimes called the *noise bulk edge*) by

$$\mathcal{Z}_+ \equiv \mathcal{Z}_+(F_Z) = \sup\{z : F_Z(z) < 1\}.$$

We also assume that  $F_Z$  is nontrivial ( $dF_Z$  is not a single atom at  $z = 0$ ), in other words,  $\mathcal{Z}_+(F_Z) > 0$ . Note that neither the distribution  $F_Z$  nor its bulk edge  $\mathcal{Z}_+(F_Z)$  are assumed to be known to the statistician.

5. *No outliers straying from the bulk.* Asymptotically, no singular values of  $Z_n$  can be found above the bulk edge:

$$z_{1,n} = \|Z_n\| \xrightarrow{\text{a.s.}} \mathcal{Z}_+(F_Z).$$

6. *Thickness of the bulk edge.* The following condition holds:

$$\lim_{y \rightarrow \mathcal{Z}_+(F_Z)} \int (y - z)^{-2} dF_Z(z) = \infty,$$

where the limit is taken from the right. That is,  $F_Z$  puts “sufficient” mass near the upper edge of its support. Under this condition, when the signal singular values  $x_i$  are sufficiently small, the amount of “information” one can obtain about the corresponding singular vectors  $\mathbf{a}_{i,n} \mathbf{b}_{i,n}^\top$  from the leading singular vectors of  $Y_n$  also vanishes.

This assumption is by no means esoteric. For example, suppose that  $F_Z$  has a continuous density  $f_Z$  in a neighborhood of  $z_+ = \mathcal{Z}_+(F_Z)$ , where it behaves like  $f_Z(z) \sim C(z - z_+)^{\alpha}$  as  $z \rightarrow z_+$ ; here,  $\alpha > 0$  is some exponent. Then this condition holds whenever  $\alpha \leq 1$ . In Section 3.2, we mention a broad class of noise matrices  $Z_n$  for which this property holds with  $\alpha = 1/2$ .

*Class of estimators and performance measure.* Our goal is to estimate  $X_n$ . We consider the family of singular value hard-thresholding estimators:  $\hat{X}_\theta = \hat{X}_\theta(Y_n)$ , where

$$(3.1) \quad \hat{X}_\theta = \sum_{i=1}^p y_{i,n} \mathbb{1}_{\{y_{i,n} > \theta\}} \cdot \mathbf{u}_{i,n} \mathbf{v}_{i,n}^\top,$$

<sup>10</sup>The assumption that the  $x_i$ 's are all distinct is standard in the literature on singular value shrinkage in the spiked model. When there are multiplicities, the SVD of  $X_n$  is not uniquely defined and, consequently, some framing constructs we use in this paper become inapplicable. The reframings adapted to such degenerate situations are beyond our scope. At any rate, the distinctness condition is generic in the space of matrices.

<sup>11</sup>In other words,  $\mathbf{a}_{1,n}, \dots, \mathbf{a}_{r,n}$  are sampled from the  $O(n)$ -invariant distribution on the Stiefel manifold  $V_r(\mathbb{R}^n)$ . Equivalently, one can assume that  $\mathbf{a}_{i,n}$  and  $\mathbf{b}_{i,n}$  are any arbitrary sequences of orthonormal  $r$ -frames, and the distribution of  $Z_n$  is invariant to multiplication by  $O(n)$  to the left and by  $O(p)$  to the right.

<sup>12</sup>This assumption might seem unnatural to many statisticians when they first encounter it; but note that if  $Z_n$  is a standard Gaussian white noise, then even though the distribution of matrix entries is not compactly supported, the limiting bulk distribution of singular values is compactly supported in  $[(1 - \sqrt{\gamma}), (1 + \sqrt{\gamma})]$ .

where  $Y_n = \sum_{i=1}^p y_{i,n} \mathbf{u}_{i,n} \mathbf{v}_{i,n}^\top$  is a SVD. To ensure that the vectors  $\{\mathbf{u}_{i,n}, \mathbf{v}_{i,n}\}$  are well-defined, even when there are multiplicities in the spectrum, the right singular vectors  $\mathbf{v}_{i,n}$  corresponding to a degenerate singular value of  $Y_n$  are chosen to be a random (Haar distributed) orthonormal basis for the corresponding (right) singular subspace. (Accordingly, the corresponding  $\mathbf{u}_{i,n}$ 's constitute a random orthonormal basis for the corresponding left singular subspace.) We measure the error with respect to Frobenius norm (squared error),<sup>13</sup> where we denote

$$(3.2) \quad \text{SE}_n[\mathbf{x}|\theta] = \|X_n - \hat{X}_\theta(Y_n)\|_F^2.$$

Our task is to choose  $\theta$ , so as to make  $\text{SE}_n[\mathbf{x}|\theta]$  as small as possible, in an appropriate sense (note that  $\text{SE}_n[\mathbf{x}|\theta]$  is a random variable—we *do not* take the expectation of  $X_n$  and  $Y_n$ ). The best possible performance is given by the *oracle loss*,

$$(3.3) \quad \text{SE}_n^*[\mathbf{x}] = \min_{\theta \geq 0} \text{SE}_n[\mathbf{x}|\theta],$$

which is the best loss one can achieve over the family of singular value hard-threshold estimators, *even knowing the true signal*  $X_n$ . Our goal in this paper is to develop a threshold selector that “typically for large  $n$ ” attains the oracle loss  $\text{SE}_n^*[\mathbf{x}]$ . Note that the oracle loss  $\text{SE}_n^*[\mathbf{x}]$  is also a random variable, and it is not a priori clear how to estimate it. An important observation is that the (random) function  $\theta \mapsto \text{SE}_n[\mathbf{x}|\theta]$  is piecewise constant, with finitely many jumps (specifically, these are at the singular values of  $Y_n$ :  $y_{1,n}, \dots, y_{p,n}$ ). In particular, the minimum of  $\text{SE}_n[\mathbf{x}|\theta]$  is attained not strictly at a point, but on an *interval* (or a union of intervals).

### 3.1. Background from random matrix theory.

*The spiked model.* Our perspective on the matrix denoising problem extends the one proposed by Perry [30] and Shabalin and Nobel [32]. In the model they proposed, which was inspired by Johnstone’s spiked covariance model [23], one works under the same model  $Y_n = X_n + Z_n$  as described above, but specifically assumes that the noise matrix  $Z_n$  is column-normalized and white, namely, that its entries are properly scaled *i.i.d.* random variables. This model’s close sibling, the spiked model for high-dimensional covariance, has been extensively studied in the probability and statistics literature, to such an extent that we cannot point to all of the existing literature here. Seminal works such as [5, 9, 29] and others have shown that the randomness in the spiked model can be neatly described in terms of the so-called BBP phase transition, similar to the one discovered in [8] and of the displacement of the sample eigenvalues relative to the population eigenvalues, and of the rotation of the sample eigenvectors relative to the populations eigenvectors.

In the matrix denoising setup that we consider here, the model described by our assumptions above has been studied in [10], and the same three underlying phenomena were identified and quantified:

1. *BBP phase transition:* Let  $z_+ = \mathcal{Z}_+(F_Z)$  denote the noise bulk edge. There is a functional  $\mathcal{X}_+(F_Z, \gamma)$  that depends on the LECDF  $F_Z$  and the asymptotic shape  $\gamma$  that defines an important threshold phenomenon in the behavior of limiting empirical singular values. Setting  $x_+ = \mathcal{X}_+(F_Z, \gamma)$ , then for any  $i = 1, \dots, r$  where  $x_i \leq x_+$ ,

$$(3.4) \quad y_{i,n} \xrightarrow{\text{a.s.}} z_+, \quad n \rightarrow \infty.$$

In short, *sufficiently small signal singular values*  $x_i$  *do not produce outliers beyond the noise bulk edge*. As we are about to see, the situation for  $x_i > x_+$  is quite different. The split

<sup>13</sup>Recall that for a matrix  $A$ ,  $\|A\|_F^2 = \sum_{i,j} |A_{i,j}|^2$ .

between  $x_i \geq x_+$  is sometimes called the Baik–Ben Arous–Péché (BBP) phase transition, after the original example of this type [8].

2. *Limiting location of outlier singular values:* The limiting value of  $y_{i,n}$  is *not* its underlying population counterpart  $x_i$ . There is instead a functional  $\mathcal{Y}(x; F_Z, \gamma)$ , depending on  $F_Z$  and  $\gamma$ , describing this limiting behavior. The function of  $x$  obtained by fixing  $F_Z$  and  $\gamma$ — $\mathcal{Y}(x) \equiv \mathcal{Y}(x; F_Z, \gamma)$ —explains how the asymptotic limit varies with theoretical singular value  $x$ . For any  $i = 1, \dots, r$  where  $x_i \geq x_+ \equiv \mathcal{X}_+$ ,

$$(3.5) \quad y_{i,n} \xrightarrow{\text{a.s.}} y_{i,\infty} = \mathcal{Y}(x_i), \quad n \rightarrow \infty.$$

The function  $x \mapsto \mathcal{Y}(x)$  is strictly increasing and one-to-one between  $[x_+, \infty)$  and  $[z_+, \infty)$ .

3. *No limiting cross-correlation of noncorresponding principal subspaces:* For  $i \neq j$ , the empirical dyad  $\mathbf{u}_{n,i} \mathbf{v}_{n,i}^\top$  ultimately decorrelates from each of the noncorresponding population dyads  $\mathbf{a}_{n,j} \mathbf{b}_{n,j}^\top$ . For any  $i, j = 1, \dots, r$  such that  $i \neq j$ ,

$$(3.6) \quad \langle \mathbf{a}_{n,i}, \mathbf{u}_{n,j} \rangle \cdot \langle \mathbf{b}_{n,i}, \mathbf{v}_{n,j} \rangle \xrightarrow{\text{a.s.}} 0, \quad n \rightarrow \infty.$$

4. *Limiting cross-correlation of corresponding principal subspaces:* Suppose the signal singular values  $(x_i)_{i=1}^r$  are distinct. The empirical dyad  $\mathbf{u}_{n,i} \mathbf{v}_{n,i}^\top$  *does* correlate with its theoretical counterpart  $\mathbf{a}_{n,i} \mathbf{b}_{n,i}^\top$ , but not perfectly. The limit is described by a functional  $\mathcal{C}(x; F_Z, \gamma)$  depending on  $x$ ,  $F_Z$  and  $\gamma$ . Fixing once again  $F_Z$  and  $\gamma$ , we get a function of  $x$ ,  $\mathcal{C}(x) \equiv \mathcal{C}(x; F_Z, \gamma)$ , such that, with  $x_+ = \mathcal{X}_+(F_Z, \gamma)$ ,

$$(3.7) \quad \langle \mathbf{a}_{n,i}, \mathbf{u}_{n,i} \rangle \cdot \langle \mathbf{b}_{n,i}, \mathbf{v}_{n,i} \rangle \xrightarrow{\text{a.s.}} \begin{cases} \mathcal{C}(x_i) & x_i > x_+, \\ 0 & x_i \leq x_+. \end{cases}$$

We now give formulas for  $\mathcal{X}_+$  and the mappings  $\mathcal{Y}(\cdot)$  and  $\mathcal{C}(\cdot)$ , as computed in [10]. For a CDF  $H$ , let

$$(3.8) \quad \varphi(y; H) = \int \frac{y}{y^2 - z^2} dH(z),$$

which defines a smooth function on  $y > \mathcal{Z}_+(H)$ . Its derivative is

$$(3.9) \quad \varphi'(y; H) = - \int \frac{y^2 + z^2}{(y^2 - z^2)^2} dH(z).$$

Also define

$$(3.10) \quad \tilde{\varphi}_\gamma(y; H) = \gamma \varphi(y; H) + \frac{(1 - \gamma)}{y}, \quad \tilde{\varphi}'_\gamma(y; H) = \gamma \varphi'(y; H) - \frac{1 - \gamma}{y^2}.$$

Note that  $\tilde{\varphi}_\gamma(y; H)$  is simply  $\varphi(y; \tilde{H}_\gamma)$ , where  $\tilde{H}_\gamma(z) = \gamma H(z) + (1 - \gamma) \mathbb{1}_{\{z \geq 0\}}$ . This so-called *companion CDF*  $\tilde{H}_\gamma$  describes the same distribution of nonzero singular values as  $H$ , diluted by “zero padding” and has the following interpretation: if  $Z_n$  is a sequence of  $n$ -by- $p$  matrices with a limiting singular value distribution  $H$ , then  $Z_n^\top$  has a limiting singular value distribution  $\tilde{H}_\gamma$ .<sup>14</sup> Let

$$(3.11) \quad \begin{aligned} \mathcal{D}_\gamma(y; H) &\equiv \varphi(y; H) \cdot \tilde{\varphi}_\gamma(y; H), \\ \mathcal{D}'_\gamma(y; H) &\equiv \varphi'(y; H) \cdot \tilde{\varphi}_\gamma(y; H) + \varphi(y; H) \cdot \tilde{\varphi}'_\gamma(y; H). \end{aligned}$$

<sup>14</sup>Practitioners will recognize that computer software often offers two options for SVD outputs, a “fat” output with zero padding and a “thin” output with those superfluous zeros stripped away. If  $H$  denotes the LECDF of the “thin” output singular values, then  $\tilde{H}$  is the corresponding LECDF of the “fat” outputs.

To ease the notation in coming paragraphs, we put for short  $\mathcal{D}_\gamma(y) = \mathcal{D}_\gamma(y; F_Z, \gamma)$ , and similarly for  $\varphi(y)$ ,  $\tilde{\varphi}_\gamma(y)$ . Let  $z_+ = \mathcal{Z}_+(F_Z)$  denote the bulk edge. The BBP phase transition location  $x_+ = \mathcal{X}_+(F_Z, \gamma)$  is given by

$$(3.12) \quad x_+ = \lim_{y \rightarrow z_+} (\mathcal{D}_\gamma(y))^{-1/2},$$

equivalently,  $1/x_+^2 = \lim_{y \rightarrow z_+} \mathcal{D}_\gamma(y)$ . It is easy to verify that  $\varphi(y)$ ,  $\tilde{\varphi}_\gamma(y)$  and  $\mathcal{D}_\gamma(y)$  are nonnegative, strictly decreasing functions of  $y > z_+$ , each tending to 0 as  $y \rightarrow \infty$ . Thus,  $\mathcal{D}_\gamma(\cdot)$  maps the interval  $(z_+, \infty)$  bijectively into  $(x_+, 0)$ ; denote by  $\mathcal{D}_\gamma^{-1}(\cdot) \equiv \mathcal{D}_\gamma^{-1}(\cdot; F_Z)$  the inverse mapping.

We finally can give formulas for the fundamental phenomenological limits described earlier. The limiting empirical signal singular value  $y_{i,\infty} = \mathcal{Y}(x_i) \equiv \mathcal{Y}(x_i; F_Z, \gamma)$  obeys

$$(3.13) \quad \mathcal{Y}(x) = \mathcal{D}_\gamma^{-1}\left(\frac{1}{x^2}\right) \quad \text{for } x > x_+,$$

equivalently,  $\mathcal{D}_\gamma(\mathcal{Y}(x)) = 1/x^2$ . The asymptotic cosine  $\mathcal{C}(x) \equiv \mathcal{C}(x; F_Z, \gamma)$  is given by

$$(3.14) \quad \mathcal{C}(x) = -\frac{2}{x^3} \cdot \frac{1}{\mathcal{D}'_\gamma(\mathcal{Y}(x))} \quad \text{for } x > x_+.$$

One may readily verify that  $\mathcal{C}(x) \geq 0$  for all  $x > x_+$ .<sup>15</sup>

We sometimes adopt the implicit parameterization of  $\mathcal{C}(x)$  in terms of  $y = \mathcal{Y}(x)$ :

$$(3.15) \quad \mathcal{C}(x) = -2 \cdot \frac{(\mathcal{D}_\gamma(y))^{3/2}}{\mathcal{D}'_\gamma(y)} \quad \text{where } y = \mathcal{Y}(x) \text{ and } x > x_+.$$

*Existence of a BBP phase transition.* Recall that  $x_+ = \mathcal{X}_+(F_Z, \gamma)$  gives the threshold such that whenever  $x_i \leq x_+$ , one *does not* observe an outlier singular value away from the bulk of  $Y$ . Not all noise distributions display this phase transition phenomenon, that is, they may not exhibit  $x_+ > 0$ : indeed, by equation (3.12),  $\mathcal{X}_+ > 0$  if and only if  $\lim_{y \rightarrow \mathcal{Z}_+(F_Z)} \mathcal{D}_\gamma(y; F_Z) < \infty$ , equivalently,  $\lim_{y \rightarrow \mathcal{Z}_+(F_Z)} \int (y-z)^{-1} dF_Z(z) < \infty$ . This condition entails that near its own bulk edge,  $F_Z$  is not “thick.” For example, when  $F_Z$  has a density in a neighborhood of  $z_+ = \mathcal{Z}_+(F_Z)$  that behaves as  $f_Z(z) \sim C(z - z_+)^{\alpha}$ , this condition is satisfied whenever  $\alpha > 0$ . For example, the family of noise distributions described in Section 3.2 is of this type (with  $\alpha = 1/2$ ); they all display a BBP phase transition. Moreover, Assumption 6 gives  $\lim_{y \rightarrow z_+} \mathcal{D}'_\gamma(y; F_Z) = -\infty$ . From equation (3.15), this means that if  $x_+ \equiv \mathcal{X}_+(F_Z, \gamma) > 0$ , then  $\mathcal{C}(x) = 0$  as  $x \rightarrow x_+$  from the right. Curiously, when  $x_+ = 0$ , this does not have to be the case. For instance, when  $dF_Z = \delta_1$  and  $\gamma = 1$ , an easy computation shows  $x_+ = 0$  and  $\mathcal{C}(x) = \frac{y^3}{y(y^2+1)}$ , where  $y = \mathcal{Y}(x)$  and  $\mathcal{Z}_+(F_Z) = 1$ . We see that  $\lim_{x \rightarrow x_+} \mathcal{C}(x) = 1/2$ : this means that an *arbitrarily small* signal already creates a very strong bias in the direction of the principal singular vectors of  $Y_n$ .

<sup>15</sup>Note that in [10], equation (3.7) is only stated as  $|\langle \mathbf{a}_{n,i}, \mathbf{u}_{n,i} \rangle \cdot \langle \mathbf{b}_{n,i}, \mathbf{v}_{n,i} \rangle| \xrightarrow{\text{a.s.}} \mathcal{C}(x_i)$  (assuming  $x_i > x_+$ ), with the absolute value. One may readily verify that the limiting cross-correlation must, in fact, be nonnegative: Start with

$$y_{i,n} = \mathbf{u}_{i,n}^\top Y_n \mathbf{v}_{i,n} = \mathbf{u}_{i,n}^\top X_n \mathbf{v}_{i,n} + \mathbf{u}_{i,n}^\top Z_n \mathbf{v}_{i,n} \leq \mathbf{u}_{i,n}^\top X_n \mathbf{v}_{i,n} + \|Z_n\|.$$

By equation (3.6),  $\mathbf{u}_{i,n}^\top X_n \mathbf{v}_{i,n} \sim x_i \langle \mathbf{a}_{n,i}, \mathbf{u}_{n,i} \rangle \cdot \langle \mathbf{b}_{n,i}, \mathbf{v}_{n,i} \rangle$ , while  $\|Z_n\| \rightarrow z_+$ ,  $y_{i,n} \rightarrow y_{i,\infty} \geq z_+$ . The conclusion follows.

*Notation.* Throughout the paper, we use the notation

$$y_{i,\infty} = \begin{cases} \mathcal{Y}(x_i) & \text{when } x_i > \mathcal{X}_+, \\ \mathcal{Z}_+(F_Z) & \text{when } x_i \leq \mathcal{X}_+. \end{cases}$$

By the results of [10], the singular values of  $Y_n$ ,  $y_{1,n} \geq \dots \geq y_{p,n}$ , satisfy  $y_{i,n} \xrightarrow{\text{a.s.}} y_{i,\infty}$  for any fixed index  $i$  (for  $i > r$  this is an easy consequence of the interlacing inequality for singular values).

3.2. *Noise matrices with correlated columns.* We conclude this section by mentioning an important family of noise matrices satisfying our assumptions, namely, noise matrices with independent rows, having cross-column correlations. We consider noise matrices of the form  $Z_n = W_n S_n^{1/2}$ , where  $(W_n)$  and  $(S_n)$  are sequences of matrices obeying:

- $W_n$  is an  $n$ -by- $p$  matrix with i.i.d. elements. Specifically, let  $W$  denote a random variable with moments

$$\mathbb{E}(W) = 0, \quad \mathbb{E}(W^2) = 1, \quad \mathbb{E}(W^4) < \infty.$$

The entries of  $W_n$  are i.i.d., with law  $(W_n)_{ij} \stackrel{d}{=} n^{-1/2}W$ , that is, scaled to variance  $1/n$ . Finiteness of the fourth moment of  $W$  is essential; see [6].

- $(S_n)$  is a sequence of nonrandom  $p$ -by- $p$  matrices. Let  $\lambda_1(S_n) \geq \dots \geq \lambda_p(S_n)$  be the eigenvalues of  $S_n$ , and denote by  $F_{S_n}(\lambda) = p^{-1} \sum_{i=1}^p \mathbb{1}_{\{\lambda_i(S_n) \leq \lambda\}}$  the empirical CDF of its eigenvalues. We assume that the sequence  $(F_{S_n})$  converges to a compactly supported LECDF  $F_S$ . Moreover, denoting the upper and lower edges of the support by

$$\lambda_+(F_S) = \sup\{\lambda : F_S(\lambda) < 1\}, \quad \lambda_-(F_S) = \inf\{\lambda : F_S(\lambda) > 0\},$$

we assume that  $\lambda_1(S_n) \rightarrow \lambda_+(F_S)$  and  $\lambda_p(S_n) \rightarrow \lambda_-(F_S)$ .

We refer to a random matrix ensemble of the form above as a *noise matrix with correlated columns*. They appear, for example, in the following scenario: We observe  $n$  i.i.d.  $p$ -dimensional samples  $\mathbf{y}_i = \mathbf{x}_i + \mathbf{z}_i$ , where  $\mathbf{x}_i$  are instances of a signal vector, assumed to be supported in an  $r$ -dimensional subspace, and  $\mathbf{z}_i = S_n^{1/2} \mathbf{w}_i$  is a vector of correlated noise, with covariance  $\text{Cov}(\mathbf{w}_i) = S_n$ . Let  $Y_n$  be the  $n$ -by- $p$  matrix, whose rows are  $n^{-1/2} \mathbf{y}_i^\top$  (define  $X_n$ ,  $W_n$  and  $Z_n$  similarly). Then  $Y_n = X_n + Z_n = X_n + W_n S_n^{1/2}$ , where  $\text{rank}(X_n) \leq r$ , by assumption. For any estimator  $\hat{X} = \hat{X}(Y_n)$ , let  $\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_n$  be the rows of the matrix  $n \cdot \hat{X}$ . Then the Frobenius loss is just the average  $L^2$  loss in estimating the signal samples  $\mathbf{x}_i$  by the vectors  $\hat{\mathbf{x}}_i$ :  $\|X_n - \hat{X}(Y_n)\|_F^2 = n^{-1} \sum_{i=1}^n \|\mathbf{x}_i - \hat{\mathbf{x}}_i\|_2^2$ .

Much is known about the singular values of  $Z_n$ :

1. *Limiting singular value distribution:*  $F_{Z_n}$  converges weakly almost surely to a compactly supported law  $F_Z$ . This limiting law is defined in terms of its Stieltjes transform,<sup>16</sup>  $m(y) = \int (z^2 - y)^{-1} dF_Z(z)$ ;  $m(y)$  is the unique Stieltjes transform satisfying

$$m(y) = \int \frac{1}{t(1 - \gamma - \gamma y m(y)) - y} dF_S(t) \quad \text{for all } y \in \mathbb{C} \setminus \mathbb{R}.$$

<sup>16</sup> $m(y)$  is in fact the Stieltjes transform of the limiting eigenvalue distribution of  $Z^\top Z$ :

$$m(y) = \int (z - y)^{-1} dF_{Z^\top Z}(z) = \int (z^2 - y)^{-1} dF_Z(z).$$

2. *Extreme singular values:* The largest and smallest singular values of  $Z_n$  converge almost surely to the upper and lower edges of the support of the limiting law<sup>17</sup>  $F_Z$ :

$$\mathcal{Z}_+(F_{Z_n}) \xrightarrow{\text{a.s.}} \mathcal{Z}_+(F_Z), \quad \mathcal{Z}_-(F_{Z_n}) \xrightarrow{\text{a.s.}} \mathcal{Z}_-(F_Z).$$

3. *Behavior at the edge of the bulk:* On  $\mathbb{R} \setminus \{0\}$ , the limiting law  $F_Z$  is absolutely continuous with respect to Lebesgue measure. Denoting by  $f_Z$  the corresponding density, we have  $f_Z(z) \sim C \cdot \sqrt{|z - \mathcal{Z}_+(F_Z)|}$  as  $z \nearrow \mathcal{Z}_+(F_Z)$ . This is the same behavior as a Marčenko–Pastur law, corresponding to  $S_n = I$ . This edge behavior will motivate one of our strategies for estimating  $F_{Z_n}$  from the observed singular values  $F_{Y_n}$  (*imputation*, see Section 4.2); this is an important step in the *ScreenNOT* algorithm. Also, note that in particular, the limiting noise CDF  $F_Z$  satisfies Assumption 6.

4. *CLT for linear spectral statistics:* Denote

$$\underline{y} = (1 - \sqrt{\gamma})^2 \cdot \lambda_-(F_S), \quad \bar{y} = (1 + \sqrt{\gamma})^2 \cdot \lambda_+(F_S).$$

Note that  $\underline{y}^{1/2} \leq \mathcal{Z}_-(F_Z) \leq \mathcal{Z}_+(F_Z) \leq \bar{y}^{1/2}$ . Let  $g$  be analytic on an open domain in  $\mathbb{C}$  containing the closed interval  $[\underline{y}, \bar{y}]$ . Set

$$\Phi_n[g] = \int g(z^2)(dF_Z - dF_{Z_n})(z),$$

which is a random variable.<sup>18</sup> Then the sequence  $p \cdot \Phi_n[g]$  is tight. If, moreover,  $\mathbb{E}(W^4) = 3$ , then  $p \cdot \Phi_n[g]$  converges in law to a Gaussian random variable.

For properties (1) and (2), we refer to [6] and the references therein (see also the book [4]). Property (3) is proved in [33]. Property (4) is proved in [7]. In Figure 3 we plot the empirical singular value distribution of several noise matrices, sampled according to the model described above with different choices for the covariance matrix  $S_n$ .

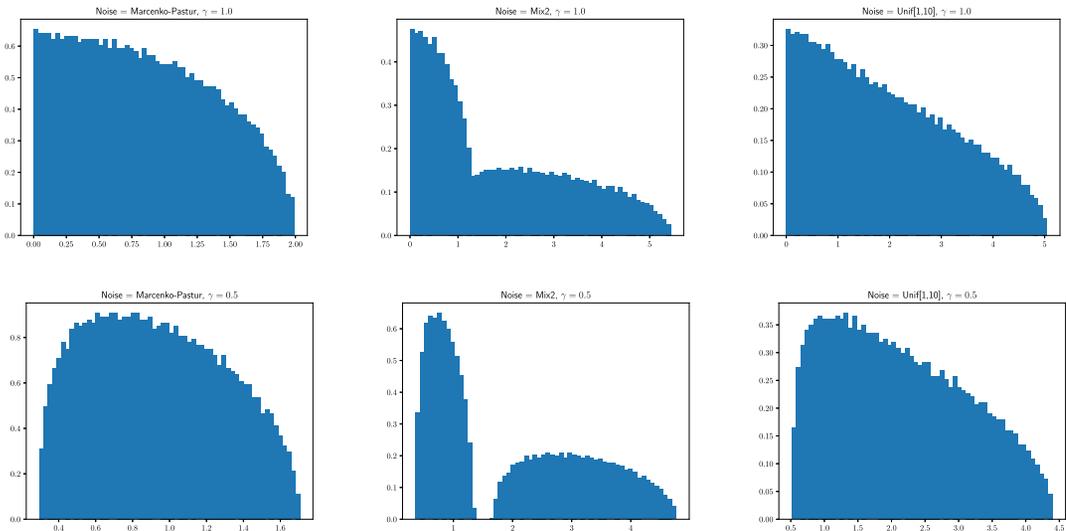


FIG. 3. Several empirical noise singular value distributions that come from the model in Section 3.2. Left to right: covariance eigenvalue distribution: (i)  $dF_S = \delta_1$  (giving a Marčenko–Pastur bulk); (ii) An equal mix of two atoms,  $dF_S = \frac{1}{2}\delta_1 + \frac{1}{2}\delta_{10}$ ; (iii)  $F_S$  uniform on  $[1, 10]$ . Top: shape  $\gamma = 1$ ; bottom:  $\gamma = 0.5$ . Each plot is the histogram of singular values from a single random  $n \times p$  matrix, with  $p = 3000$  and  $n = p/\gamma$ .

<sup>17</sup>Below,  $\mathcal{Z}_-(F_{Z_n})$  denotes the smallest singular value of  $Z_n$ . Recall that we assume  $p \leq n$ .

<sup>18</sup>A random variable of the form  $\int h(z)F_{Z_n}(z) = p^{-1} \sum_{i=1}^n h(z_{i,n})$  is called a linear spectral statistic.

As a final remark, we mention that when the covariance matrix  $S_n$  is invertible and known (or can be consistently estimated with respect to the operator norm), estimating  $X_n$  using the leading singular vectors of  $Y_n$  is suboptimal. Instead, it is better to first “whiten” the noise, that is, compute  $Y_n^w = Y_n S_n^{-1/2} = X_n S_n^{-1/2} + W_n$ . Letting  $\mathbf{u}_{i,n}^w$  and  $\mathbf{v}_{i,n}^w$  be respectively the left and right singular vectors of  $Y_n$ , we can “recolor” the right singular vectors,  $\mathbf{v}_{i,n}^c = S_n^{1/2} \mathbf{v}_{i,n}^w / \|S_n^{1/2} \mathbf{v}_{i,n}^w\|$ . Under a uniform prior on the signal singular vectors, as we assume in this paper, and when  $W_n$  is i.i.d. Gaussian, the correlations between the signal singular vectors and empirical singular vectors can be shown to be *stronger* in the whiten-then-recolor scheme (see [21, 26] for details),

$$\lim_{n \rightarrow \infty} \langle \mathbf{a}_{i,n}, \mathbf{u}_{i,n}^w \rangle \langle \mathbf{b}_{i,n}, \mathbf{v}_{i,n}^c \rangle \geq \lim_{n \rightarrow \infty} \langle \mathbf{a}_{i,n}, \mathbf{u}_{i,n} \rangle \langle \mathbf{b}_{i,n}, \mathbf{v}_{i,n} \rangle.$$

### 4. Results.

*Outline.* We start by developing a theory for optimal hard thresholding, under the assumption that the noise singular value distribution  $F_Z$  is known. We show that there is an asymptotically uniquely admissible hard threshold  $T_\gamma(F_Z)$ , which is given as a certain functional  $T_\gamma$  of the asymptotic aspect ratio  $\gamma$  and the limiting noise CDF  $F_Z$ . Relying on the fact that  $\text{SE}_n[\mathbf{x}|\theta]$  can only take a finite number of values as  $\theta$  varies, we show that thresholding at  $T_\gamma(F_Z)$  has rather strong optimality properties: it, in fact, attains oracle loss, at finite  $n$ , with probability increasing to 1 as  $n \rightarrow \infty$ . We then move onto the practical setting of interest, in which  $F_Z$  is unknown. We propose a method for consistently estimating  $T_\gamma(F_Z)$  from the observed data  $Y_n$ . We do this by applying the optimal threshold functional  $T_{p/n}(\cdot)$  on a judiciously transformed version of  $F_{Y_n}$ , the empirical singular value distribution of  $Y_n$ . The continuity of the functional with respect to the CDF and the shape parameter then implies that the resulting quantity is a consistent estimator for  $T_\gamma(F_Z)$ ; the optimality properties of the adaptive algorithm then follow from the previously developed theory. Unless otherwise stated, we always operate under assumptions (1)–(6) of Section 3.

4.1. *A theory for optimal singular value thresholding.* Consider the function  $\theta \mapsto \text{ASE}[\mathbf{x}|\theta]$  defined for  $\theta > 0$  by

$$(4.1) \quad \text{ASE}[\mathbf{x}|\theta] = \sum_{i=1}^r R(x_i|\theta) \quad \text{where } R(x_i|\theta) = \mathbb{1}_{\{y_{i,\infty} \leq \theta\}} \cdot R_0(x_i) + \mathbb{1}_{\{y_{i,\infty} > \theta\}} \cdot R_1(x_i),$$

and

$$(4.2) \quad R_0(x) = x^2, \quad R_1(x) = \begin{cases} x^2 + \mathcal{Y}(x)^2 - 2x\mathcal{Y}(x)\mathcal{C}(x) & \text{if } x > x_+, \\ x^2 + z_+^2 & \text{if } x \leq x_+. \end{cases}$$

Recall that  $y_{i,\infty} = z_+$  ( $\equiv \mathcal{Z}_+(F_Z)$ ) when  $x_i \leq x_+$  ( $\equiv \mathcal{X}_+(F_Z, \gamma)$ ) and  $y_{i,\infty} = \mathcal{Y}(x_i)$  when  $x > x_+$ . Define also

$$(4.3) \quad \text{ASE}^*[\mathbf{x}] = \sum_{i=1}^r R^*(x_i) \quad \text{where } R^*(x_i) = \min\{R_0(x_i), R_1(x_i)\}.$$

Clearly,  $R^*(x) \leq R(x|\theta)$  for any  $x$  and  $\theta$ , which means  $\text{ASE}^*[\mathbf{x}] \leq \text{ASE}[\mathbf{x}|\theta]$ .

It is easy to verify that for almost every  $\theta > \mathcal{Z}_+(F_Z)$ , the loss of thresholding at the fixed point  $\theta$  converges:  $\lim_{n \rightarrow \infty} \text{SE}_n[\mathbf{x}|\theta] = \text{ASE}[\mathbf{x}|\theta]$  almost surely; see Lemma 7 for a precise statement. We start by finding the threshold that attains minimum asymptotic loss.

DEFINITION 1 (Optimal threshold functional). For a compactly supported CDF  $H$  and  $\gamma \in (0, 1]$ , let

$$(4.4) \quad \Psi_\gamma(y; H) = y \cdot \frac{\mathcal{D}'_\gamma(y; H)}{\mathcal{D}_\gamma(y; H)} = y \cdot \left( \frac{\varphi'(y; H)}{\varphi(y; H)} + \frac{\tilde{\varphi}'_\gamma(y; H)}{\tilde{\varphi}_\gamma(y; H)} \right);$$

this is well-defined for  $y > \mathcal{Z}_+(H)$ . Define the functional of  $H$ ,

$$(4.5) \quad T_\gamma(H) = \inf\{y : y > \mathcal{Z}_+(H) \text{ and } \Psi_\gamma(y; H) \geq -4\}.$$

We call this the *optimal threshold functional*.

LEMMA 1. *The following holds:*

1. For any  $H$  and  $\gamma$ ,  $y \mapsto \Psi_\gamma(y; H)$  is negative and increasing, with  $\Psi_\gamma(\infty) = -2$ .
2. Assume that  $H$  is compactly supported and satisfies  $\lim_{y \rightarrow \mathcal{Z}_+(H)} \int (y - z)^{-2} dH(z) = \infty$  (note that, by assumption,  $H = F_Z$  satisfies this). Then  $T_\gamma(H)$  is the unique number  $> \mathcal{Z}_+(H)$  satisfying  $\Psi_\gamma(T_\gamma(H); H) = -4$ .
3. Thresholding at  $\theta^* = T_\gamma(F_Z)$  minimizes the asymptotic loss:

$$\text{ASE}[\mathbf{x}|\theta^*] = \min_{\theta} \text{ASE}[\mathbf{x}|\theta] = \text{ASE}^*[\mathbf{x}].$$

Moreover,  $\theta^*$  is the unique threshold for which the above holds universally, for all signals  $\mathbf{x}$ .

Lemma 1 is proved in Section 6.1.

Note that  $\theta \mapsto \text{ASE}[\mathbf{x}|\theta]$  is piecewise constant, with jumps at  $y_{1,\infty}, \dots, y_{r,\infty}$ . This means that its minimum is actually attained on an *interval*.

DEFINITION 2 (The asymptotic optimal interval). Let

$$(4.6) \quad \underline{\Theta}(\mathbf{x}) = \max\{y_{i,\infty} : y_{i,\infty} < T_\gamma(F_Z)\}, \quad \overline{\Theta}(\mathbf{x}) = \min\{y_{i,\infty} : y_{i,\infty} > T_\gamma(F_Z)\}.$$

Note that since  $T_\gamma(F_Z) > z_+ = \mathcal{Z}_+(F_Z, \gamma)$  and  $y_{r+1,\infty} = z_+$ , we always have, by definition,  $\underline{\Theta}(\mathbf{x}) \geq z_+$ . Moreover, if  $y_{1,\infty} \leq T_\gamma(F_Z)$  then we define  $\overline{\Theta}(\mathbf{x}) = \infty$ .

LEMMA 2.

1. Throughout the interval  $\theta \in (\underline{\Theta}(\mathbf{x}), \overline{\Theta}(\mathbf{x}))$ ,  $\text{ASE}[\mathbf{x}|\theta]$  is constant. Moreover, it attains its minimum there; if  $\theta_0 \in (\underline{\Theta}(\mathbf{x}), \overline{\Theta}(\mathbf{x}))$ , then

$$\text{ASE}[\mathbf{x}|\theta_0] = \min_{\theta \geq 0} \text{ASE}[\mathbf{x}|\theta] = \text{ASE}^*[\mathbf{x}].$$

2. Any  $\theta_1 > \mathcal{Z}_+(F_Z)$  outside  $[\underline{\Theta}(\mathbf{x}), \overline{\Theta}(\mathbf{x})]$  has

$$\text{ASE}[\mathbf{x}|\theta_1] > \text{ASE}^*[\mathbf{x}].$$

3. Unique asymptotic admissibility:  $T_\gamma(F_Z)$  is in the interior of the asymptotic optimal interval. In fact, it is the only threshold, which has optimal asymptotic loss simultaneously for all signals  $\mathbf{x}$ :

$$\bigcap_{\mathbf{x} \text{ signal}} (\underline{\Theta}(\mathbf{x}), \overline{\Theta}(\mathbf{x})) = \{T_\gamma(F_Z)\}.$$

Lemma 2 is proved in Section 6.1.

*The scree plot heuristic: A quantitatively interpretation.* Under our signal model, one could think of a “natural” quantification of Cattell’s scree plot heuristic. Roughly, it hopes to threshold the data singular values slightly above the pure-noise bulk edge. If this hope is fulfilled, the excess ASE incurred, compared to the minimal attainable ASE, is

$$\lim_{\delta \rightarrow 0} \text{ASE}[\mathbf{x}|z_+ + \delta] - \text{ASE}^*[\mathbf{x}] = \sum_{i: y_{i,\infty} \in (z_+, T_\gamma(F_Z))} (R_1(x_i) - R_0(x_i)).$$

The excess ASE is proportional to the number of barely/moderately emergent signal singular values, namely, such that  $y_{i,\infty} > z_+$  (so that they can be observed as outliers in the spectrum of  $Y_n$ ) but  $y_{i,\infty} < T_\gamma(F_Z)$  (meaning that the corresponding empirical singular vectors are too “noisy” so to be useful in estimating  $X_n$ ). Clearly, in the worst-case scenario, the signal  $\mathbf{x}$  consists entirely of barely emergent singular values, so that the excess ASE is proportional to  $r = \text{rank}(X_n)$ . For a concrete example, consider  $\mathbf{x}$  that consists of  $r$  distinct singular values, located *just slightly* above  $x_+$ ; in that case, the excess ASE is  $r \cdot (R_1(x_+) - R_0(x_+)) = r \cdot z_+^2$ .

It is clear at this point that thresholding at any point in the interior of the asymptotic optimal interval achieves the best asymptotic loss, among all other fixed hard thresholds. Our main result states that, remarkably, one *cannot* come up with a consistently better thresholding strategy, even if given access to the true unknown signal  $X_n$ .

#### THEOREM 1.

1. *Almost surely,*

$$\lim_{n \rightarrow \infty} \text{SE}_n^*[\mathbf{x}] = \text{ASE}^*[\mathbf{x}].$$

2. *Let  $\theta \in (\underline{\Theta}(\mathbf{x}), \overline{\Theta}(\mathbf{x}))$  be in the interior of the asymptotic optimal interval, and  $\theta_n$  be any sequence of thresholds (possibly depending on  $Y_n$ ) such that  $\theta_n \xrightarrow{\text{a.s.}} \theta$ . Then*

$$\text{SE}_n[\mathbf{x}|\theta_n] \xrightarrow{\text{a.s.}} \text{ASE}^*[\mathbf{x}].$$

Our next result states that thresholding inside the asymptotic optimal interval in fact achieves oracle risk with high probability, *for finite  $n$* .

THEOREM 2. *Suppose that  $T_\gamma(F_Z) \notin \{y_{1,\infty}, \dots, y_{r,\infty}\}$ .<sup>19</sup> Then:*

1. *Let  $\theta_0 \in (\underline{\Theta}(\mathbf{x}), \overline{\Theta}(\mathbf{x}))$  and  $\theta_n$  be a sequence with  $\theta_n \xrightarrow{\text{a.s.}} \theta_0$ . Then*

$$\mathbb{P}\{\exists N \text{ s.t. } \forall n \geq N : \text{SE}_n[\mathbf{x}|\theta_n] = \text{SE}_n^*[\mathbf{x}]\} = 1.$$

2. *Let  $\theta_1 \notin (\underline{\Theta}(\mathbf{x}), \overline{\Theta}(\mathbf{x}))$  and  $\theta_n \xrightarrow{\text{a.s.}} \theta_1$ . There exists  $\delta > 0$ ,  $\delta = \delta(\mathbf{x}; F_Z, \gamma)$  such that*

$$\mathbb{P}\{\exists N \text{ s.t. } \forall n \geq N : \text{SE}_n[\mathbf{x}|\theta_n] > \text{SE}_n^*[\mathbf{x}] + \delta\} = 1.$$

Theorems 1 and 2 are proved in Section 6.2.

<sup>19</sup>We need to exclude the case  $T_\gamma(F_Z) \in \{y_{1,\infty}, \dots, y_{r,\infty}\}$  for this reason: If  $y_{i,\infty} = T_\gamma(F_Z)$  for some  $i$ , then thresholding either slightly above or below  $y_{i,n}$  (but still inside the asymptotic optimal interval) will achieve the same (optimal) asymptotic risk. However, we cannot deduce that for finite  $n$ , one of those options is, necessarily, consistently better than the other, thereby achieving oracle risk exactly.

4.2. *The ScreeNOT algorithm.* In practice, the noise distribution  $F_Z$  is generally unknown to the statistician. Theorems 1 and 2, along with the unique admissibility property of Lemma 2, tell us that our goal should be to estimate the optimal threshold  $T_\gamma(F_Z)$ .

We start by showing that the functional  $(\gamma, H) \mapsto T_\gamma(H)$  is continuous with respect to weak convergence of CDFs, with the additional requirement that the edge of the support converges as well:

LEMMA 3 (Continuity of the optimal threshold functional). *Suppose that  $H$  is compactly supported and satisfies the condition  $\lim_{y \rightarrow \mathcal{Z}_+(H)} \int (y - z)^{-2} dH(z) = \infty$ . Let  $H_n$  be a sequence of CDFs such that:*

1.  $H_n$  converges weakly to  $H$ , denoted  $H_n \xrightarrow{d} H$ .
2.  $\mathcal{Z}_+(H_n) \rightarrow \mathcal{Z}_+(H)$ .

Then  $T_{p/n}(H_n) \rightarrow T_\gamma(H)$ .

The proof of Lemma 3 appears in the Supplementary Material, Section 4.

Recall that the empirical singular value distribution of the noise matrix,  $F_{Z_n}$ , converges, by assumption, weakly almost surely to  $F_Z$ , with  $\mathcal{Z}_+(F_{Z_n}) \xrightarrow{\text{a.s.}} \mathcal{Z}_+(F_Z)$ . The matrix noise  $Z_n$ , and consequently,  $F_{Z_n}$ , is of course unknown to the statistician. However, since  $Y_n$  is a rank- $r$  additive perturbation of  $Z_n$ , the interlacing inequalities for singular values imply, for example, the convergence of CDFs in Kolmogorov–Smirnov distance  $\|F_{Y_n} - F_{Z_n}\|_{\text{KS}} \rightarrow 0$  (see, e.g., the statement and proof of Lemma 4 below), and hence also in weak convergence. The obstacle preventing the would-be use of  $T_{p/n}(F_{Y_n})$  to estimate  $T_\gamma(F_Z)$  lies with the fact that  $T$  is not continuous in Kolmogorov–Smirnov metric convergence or other topologies involving CDF convergence such as weak convergence. More concretely,  $T_{p/n}(F_{Y_n})$  can be very different than  $T_\gamma(F_Z)$  because the top mass points of  $F_{Y_n}$  do not converge to the bulk edge  $F_Z$ .<sup>20</sup> Indeed, recall that  $\mathcal{Z}_+(F_{Y_n}) = y_{1,n} \xrightarrow{\text{a.s.}} y_{1,\infty}$ , which is  $> \mathcal{Z}_+(F_Z)$  when  $x_1 > \mathcal{X}_+$ .

To get a reasonable simulacrum of  $F_{Z_n}$  built from knowledge only of  $F_{Y_n}$ , we perform “surgery” on  $F_{Y_n}$ , “amputating” the top  $k$  mass points and fitting a “prosthesis” to replace them. Post-surgery, we get an estimate for the unknown empirical noise CDF  $F_{Z_n}$ .

As indicated in Section 2 above, the user of our proposed procedure supplies an upper bound (which can be potentially very loose)  $k \geq r$  on the rank of the unknown low-rank matrix.

We could, in principle, propose any one of the following “pseudo-noise” CDFs, derived from  $F_{Y_n}$ :

- *Transport to zero:* We construct a CDF,  $F_{n,k}^0$ , obtained by removing the  $k$  largest singular values of  $Y_n$ , and adding  $k$  additional zeros. That is,

$$F_{n,k}^0(y) = \frac{1}{p} \sum_{i=k+1}^p \mathbb{1}_{\{y_{i,n} \leq y\}} + \frac{k}{p} \mathbb{1}_{\{y \geq 0\}}.$$

- *“Winsorization” (clipping):* As in the previous construction, we remove the leading  $k$  singular values. Instead of adding  $k$  zeroes, we add  $k$  copies of  $y_{k+1,n}$ . Equivalently, we “clip” the large singular values of  $Y_n$  to be at most the size of  $y_{n,k+1}$ . That is,

$$F_{n,k}^w(y) = \frac{1}{p} \sum_{i=k+1}^p \mathbb{1}_{\{y_{i,n} \leq y\}} + \frac{k}{p} \mathbb{1}_{\{y_{k+1,n} \leq y\}}.$$

<sup>20</sup>Of course, this does not prevent convergence of ECDFs. Recall that  $F_{Y_n} \xrightarrow{d} F_Z$  means that for bounded and continuous functions  $f$ ,  $\int f(z) dF_{Y_n}(z) \rightarrow \int f(z) dF_Z(z)$ .

- “Imputation” (reconstruction of the missing upper tail): After removing the top  $k$  singular values of  $Y_n$ , we try to construct the noise tail in a principled way. Recall that when  $Z_n$  is a noise matrix with correlated columns, as described in Section 3.2,  $F_Z$  has a density near  $z_+ = \mathcal{Z}_+(F_Z)$  that behaves as  $f_Z(z) \sim C(z_+ - z)^{1/2}$  as  $z \rightarrow z_+$  ([33]). Using the heuristic,<sup>21</sup>

$$\frac{\ell - 1}{p} \approx \int_{z_{\ell,n}}^{\mathcal{Z}_+(F_Z)} f_Z(z) dz \approx \int_{z_{\ell,n}}^{z_+} C(z_+ - z)^{1/2} dz = C'(z_+ - z_{\ell,n})^{3/2},$$

we can estimate the distance between singular values in the upper tail as

$$z_{\ell,n} - z_{t,n} \approx C'' \left[ \left( \frac{t-1}{p} \right)^{2/3} - \left( \frac{\ell-1}{p} \right)^{2/3} \right].$$

Taking  $y_{\ell,n} \approx z_{\ell,n}$  for  $\ell \geq r + 1$ , we propose to estimate the unknown constant as

$$C'' = \frac{y_{2k+1,n} - y_{k+1}}{(2k/p)^{2/3} - (k/p)^{2/3}},$$

assuming  $2k + 1 < p$  (when  $k$  is not very small compared to  $p$ , there is no reason to believe this heuristic should give good results). We “reconstruct” the missing upper tail as

$$\tilde{y}_{i,n} = y_{k+1,n} + C'' \left[ \left( \frac{k}{p} \right)^{2/3} - \left( \frac{i-1}{p} \right)^{2/3} \right] = y_{k+1,n} + \frac{1 - \left( \frac{i-1}{k} \right)^{2/3}}{2^{2/3} - 1} (y_{2k+1,n} - y_{k+1,n}).$$

The CDF we use is then

$$F_{n,k}^i(y) = \frac{1}{p} \sum_{i=k+1}^p \mathbb{1}_{\{y_{i,n} \leq y\}} + \frac{1}{p} \sum_{i=1}^k \mathbb{1}_{\{\tilde{y}_{i,n} \leq y\}}.$$

The label  $i$  on  $F_{n,k}^i$  stands for “imputation,” a standard terminology in statistical practice for filling in utterly missing data with plausible pseudo-data. Numerical results in Section 5 suggest that in many cases, the “imputation” method gives significantly better results than truncation or Winsorization at finite  $n$ . It is also more psychologically “supportive,” which is why we recommended it to practitioners in Section 2 above. However, our formal results hold for all three methods. Importantly, the list of strategies above is by no means exhaustive. Indeed, any sequence of CDFs  $F_n^*$  that satisfies the conditions of Lemma 3 may be used instead, yielding an asymptotically consistent estimate of the optimal threshold. Furthermore, our “imputation” procedure is not claimed to be optimal; possibly, other strategies will outperform those we suggest in finite problem sizes.

LEMMA 4. Suppose that  $k = k_n$  satisfies  $k_n \geq r$  and  $k_n/p \rightarrow 0$  (in particular,  $k$  can be any constant  $\geq r$ ). Then for any choice  $\star \in \{0, w, i\}$ :

1. Almost surely,  $F_{n,k}^\star \xrightarrow{d} F_Z$ .
2.  $\mathcal{Z}_+(F_{n,k}^\star) \xrightarrow{\text{a.s.}} \mathcal{Z}_+(F_Z)$ .
3. We have the following bound on the Kolmogorov–Smirnov distance between  $F_{n,k}^\star$  and  $F_{Z_n}$ :

$$\|F_{n,k}^\star - F_Z\|_{\text{KS}} = \sup_z |F_{n,k}^\star(z) - F_{Z_n}(z)| \leq \frac{k}{p}.$$

<sup>21</sup>The exponent  $\alpha = 1/2$  was chosen as typical of bulk-edge distributions in random matrix theory. If there is reason to believe that the behavior at the bulk edge follows a different power law  $f_Z(z) \sim (z_+ - z)^\alpha$ , a correspondingly different exponent can be used instead.

The proof of Lemma 4 is deferred to the Supplementary Material, Section 4.

The following theorem states the optimality properties of the proposed ScreeNOT algorithm. It is an immediate corollary of Theorems 1, 2 and Lemma 4.

**THEOREM 3.** *Suppose that  $k = k_n$  satisfies  $k_n \geq r$  and  $k_n/p \rightarrow 0$ . For any  $\star \in \{0, w, i\}$ ,  $\hat{\theta}_n = T_{p/n}(F_{n,k}^\star)$  satisfies:*

1.  $\hat{\theta}_n \xrightarrow{\text{a.s.}} T_\gamma(F_Z)$ .
2.  $\text{SE}_n[\mathbf{x}|\hat{\theta}_n] \xrightarrow{\text{a.s.}} \text{ASE}^\star[\mathbf{x}]$ .
3. *Assume that  $T_\gamma(F_Z) \notin \{y_{1,\infty}, \dots, y_{r,\infty}\}$ . Then*

$$\mathbb{P}\{\exists N \text{ s.t. } \forall n \geq N : \text{SE}_n[\mathbf{x}|\theta_n] = \text{SE}_n^\star[\mathbf{x}]\} = 1.$$

Regarding the assumption in item 3 above, we note the following.

**LEMMA 5.** *The condition  $T_\gamma(F_Z) \notin \{y_{1,\infty}, \dots, y_{r,\infty}\}$  is generic, that is, in the space of possible singular value  $r$ -vectors  $\mathbf{x}$ , this condition holds on an open dense set.*

**PROOF.** Fix the noise bulk  $F_Z$ ; then  $\theta^\star = T_\gamma(F_Z)$  is a constant not varying as the underlying signal  $\mathbf{x}$  changes. Moreover, it always strictly exceeds the bulk edge  $\mathcal{Z}_+(F_Z)$ . So,  $x^\star = \mathcal{Y}^{-1}(\theta^\star; F_Z, \gamma)$  is a uniquely defined constant, which exceeds  $\mathcal{X}_+(F_Z, \gamma)$ . The set of vectors  $\mathbf{x}$  with all entries distinct from  $x^\star$  is open and dense.  $\square$

**4.3. Stability of ScreeNOT.** One wonders how fast  $T_{p/n}(F_{n,k}^\star)$  converges to the limit  $T_\gamma(F_Z)$ . We show that for noise matrices with correlated columns, the model described in Section 3.2, the typical deviations are of order  $\mathcal{O}(k/p)$ . We start with a “quantitative” version of Lemma 3.

**LEMMA 6.** *Adopt the setting of Lemma 3. Set*

$$\Delta_{1,n} = |\varphi(T_\gamma(H); H) - \varphi(T_\gamma(H); H_n)|, \quad \Delta_{2,n} = |\varphi'(T_\gamma(H); H) - \varphi'(T_\gamma(H); H_n)|,$$

where  $\varphi$  and  $\varphi'$  are given in equations (3.8) and (3.9). Then

$$|T_\gamma(H) - T_{p/n}(H_n)| = \mathcal{O}\left(\Delta_{1,n} + \Delta_{2,n} + \left|\frac{p}{n} - \gamma\right|\right).$$

Lemma 6, along with the Kolmogorov–Smirnov distance bound from Lemma 4 and the tightness result for linear spectral statistics from [7] (see Section 3.2), gives the following.

**PROPOSITION 1.** *Suppose that  $(Z_n)$  is a sequence of noise matrices with correlated columns, as described in Section 3.2, and let  $F_S$  denote the LECDF of eigenvalues of the cross-column covariances  $S_n$ . Assume, in addition, that  $T_\gamma(F_Z) > \bar{y}^{1/2} = (1 + \sqrt{\gamma}) \cdot \sqrt{\lambda_+(F_S)}$ .<sup>22</sup> Suppose that  $k \geq r$  with  $k/p \rightarrow 0$ . Then for any  $\star \in \{0, w, i\}$ ,*

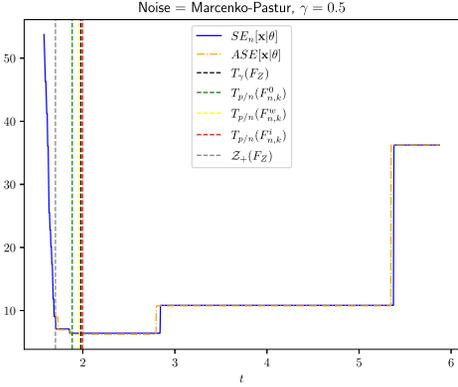
$$|T_\gamma(F_Z) - T_{p/n}(F_{n,k}^\star)| = \mathcal{O}_{\mathbb{P}}\left(\frac{k+1}{p}\right).$$

Lemma 6 and Proposition 1 are proved in the Supplementary Material, Section 4.

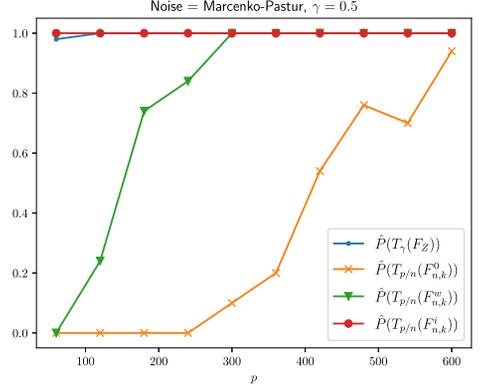
<sup>22</sup>This additional assumption is used due to a technical requirement in the results of [7]. We suspect that it can be removed.

**5. Numerical experiments.** The Supplementary Material contains comprehensive experiments conducted on a large variety of noise distributions. For space constraints, we include just a sample of these results—specifically, for white noise (Marčenko–Pastur LECDF) with  $\gamma = 0.5$ . Simulation results and code reproducing all figures here and in the Supplementary Material is available at [15]; see Section 5 of the Supplementary Material for full details on each experiment reported here.

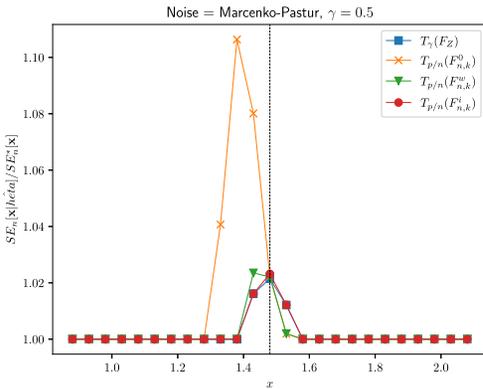
In Figure 4(a), we plot the function  $\theta \mapsto \text{SE}_n[\mathbf{x}|\theta]$  for a single fixed problem instance. The vertical lines correspond to thresholds  $\theta$ , taken to be either the true optimal threshold  $\theta = T_\gamma(F_Z)$ , its estimated versions  $\theta = T_\gamma(F_{n,k}^\star)$ ,  $\star \in \{0, w, i\}$  or the noise (asymptotic) bulk



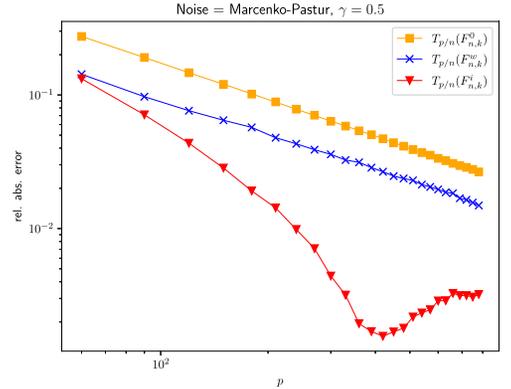
(a) A single problem instance, corresponding to the rank  $r = 5$  signal  $\mathbf{x} = (0.5, 1.0, 1.3, 2.5, 5.2)$ . Shown are the functions  $\text{SE}_n[\mathbf{x}|\theta]$  and  $\text{ASE}[\mathbf{x}|\theta]$  on top of each other. Here,  $\gamma = 0.5$ ,  $p = 500$  and  $n = p/\gamma$ . We indicate the locations of  $\mathcal{Z}_+$ ,  $T_\gamma(F_Z)$  and the estimates  $T_{p/n}(F_{n,k}^\star)$  for  $\star \in \{0, w, i\}$ , with  $k = 4r = 20$ .



(b) For various choices of  $p$  and 50 denoising experiments, shown are the fraction of experiments where each threshold  $\in \{T_\gamma(F_Z), T_{p/n}(F_{n,k}^0), T_{p/n}(F_{n,k}^w), T_{p/n}(F_{n,k}^i)\}$  attains oracle loss. Here,  $\mathbf{x} = (0.5, 1.0, 1.3, 2.5, 5.2)$  so that  $r = 5$  and  $k = 4r = 20$ .



(c) Oracle loss  $\text{SE}_n^*[\mathbf{x}]$  compared with  $\text{SE}_n[\mathbf{x}|\hat{\theta}]$  for the choices  $\hat{\theta} \in \{T_\gamma(F_Z), T_{p/n}(F_{n,k}^0), T_{p/n}(F_{n,k}^w), T_{p/n}(F_{n,k}^i)\}$ , for a single spike. We let the spike intensity  $x$  vary and plot  $\text{SE}_n[\mathbf{x}|\hat{\theta}]/\text{SE}_n^*[\mathbf{x}]$  for each choice of estimator. The ratios shown are averages across 20 experiments.



(d) Convergence rate of  $T_{p/n}(F_{n,k}^\star)$  toward  $T_\gamma(F_Z)$ . Here,  $r = 10$ ,  $x = (1, \dots, 10)$ ,  $k = 20$ . Shown is the relative absolute error  $|T_{p/n}(F_{n,k}^\star) - T_\gamma(F_Z)|/T_\gamma(F_Z)$ , plotted in logarithmic scale, as  $p$  increases and  $n = p/\gamma$ . Each point corresponds to the the average error across 50 experiments.

FIG. 4. Monte Carlo simulation results.

edge,  $\theta = \mathcal{Z}_+(F_Z)$ , which is the “natural” implementation of Cattell’s scree plot heuristic in the spiked model. The error landscape  $\text{SE}_n[\mathbf{x}|\theta]$  is seen to be a step function, and on this particular instance, all the proposed thresholding strategies fall inside the interval where it attains its global minimum; hence, they attain the oracle risk. In contrast, thresholding at the bulk edge results in a strictly suboptimal squared error.

Figures 4(b) and 4(c) compare the relative efficacy of the proposed threshold estimation strategies. In both experiments, thresholding at the exact optimal threshold  $\theta = T_\gamma(F_Z)$  (which is a priori unknown) yields the best results; among the proposed strategies, “imputation”  $\theta = T_\gamma(F_{n,k}^i)$  appears to give the best results for finite problem dimensions. Figure 4(d) demonstrates the superior finite- $n$  error of “imputation” in estimating the optimal threshold  $T_\gamma(F_Z)$ .

## 6. Proofs.

6.1. *The asymptotic loss at a fixed threshold.* To simplify notation, throughout this section  $F_Z$  and  $\gamma$  will be held fixed, and left implicit in the notation where possible. In particular, we set  $z_+ = \mathcal{Z}_+(F_Z)$ ,  $x_+ = \mathcal{X}_+(F_Z, \gamma)$  throughout, and we suppress mention of  $F_Z$  and  $\gamma$  in entities like  $\mathcal{C}$ ,  $\mathcal{Y}$ ,  $\mathcal{D}$ .

We start by investigating  $\lim_{n \rightarrow \infty} \text{SE}_n[\mathbf{x}|\theta]$  for fixed  $\theta$ . Note that when  $\theta < z_+$ , it is clear that  $\lim_{n \rightarrow \infty} \text{SE}_n[\mathbf{x}|\theta] = \infty$ ; the reason being that, for small enough  $\epsilon > 0$ , with probability 1,  $(1 - F_Z(\theta + \epsilon)) \cdot n = \Omega(n)$  empirical singular values  $y_{i,n}$  exceed the threshold  $\theta$ , so that  $\text{rank}(\hat{X}_\theta(Y_n))$  increases indefinitely. (This argument will be made more precise later.)

The following is an easy calculation.

LEMMA 7. *For any  $\theta > z_+$ , almost surely,*

1.

$$\liminf_{n \rightarrow \infty} \text{SE}_n[\mathbf{x}|\theta] \geq \text{ASE}^*[\mathbf{x}].$$

2. *If, in addition,  $\theta \notin \{y_{1,\infty}, \dots, y_{r,\infty}\}$ , then*

$$\lim_{n \rightarrow \infty} \text{SE}_n[\mathbf{x}|\theta] = \text{ASE}[\mathbf{x}|\theta].$$

*The quantities  $\text{ASE}[\mathbf{x}|\theta]$  and  $\text{ASE}^*[\mathbf{x}]$  appear in equations (4.1) and (4.3), respectively.*

PROOF. Since  $\theta > z_+$  and  $y_{r+1,n} \xrightarrow{\text{a.s.}} y_{r+1,\infty} = z_+$ , we see that with probability 1, for large enough  $n$ ,  $\hat{X}_\theta(Y_n) = \sum_{i=1}^r y_{i,n} \mathbb{1}_{\{y_{i,n} > \theta\}} \cdot \mathbf{u}_{i,n} \mathbf{v}_{i,n}^\top$ . Thus, for large enough  $n$ ,

$$\begin{aligned} \text{SE}_n[\mathbf{x}|\theta] &= \left\| \sum_{i=1}^r x_i \cdot \mathbf{a}_{i,n} \mathbf{b}_{i,n}^\top - \sum_{i=1}^r y_{i,n} \mathbb{1}_{\{y_{i,n} > \theta\}} \cdot \mathbf{u}_{i,n} \mathbf{v}_{i,n}^\top \right\|_F^2 \\ &= \sum_{i=1}^2 (x_i^2 + y_{i,n}^2 \mathbb{1}_{\{y_{i,n} > \theta\}}) + \sum_{i=1}^r \sum_{j=1}^r x_i y_{j,n} \mathbb{1}_{\{y_{i,n} > \theta\}} \cdot \langle \mathbf{a}_{i,n}, \mathbf{u}_{j,n} \rangle \langle \mathbf{b}_{i,n}, \mathbf{v}_{j,n} \rangle. \end{aligned}$$

The lemma follows by recalling (i) that  $y_{i,n} \xrightarrow{\text{a.s.}} y_{i,\infty}$  for all  $i = 1, \dots, r$ , where  $y_{i,\infty} = \mathcal{Y}(x_i)$  if  $x_i > x_+$  and  $y_{i,\infty} = z_+ < \theta$  whenever  $x_i \leq x_+$ , (ii) that

$$\langle \mathbf{a}_{i,n}, \mathbf{u}_{j,n} \rangle \langle \mathbf{b}_{i,n}, \mathbf{v}_{j,n} \rangle \rightarrow \begin{cases} \mathcal{C}(x_i) & \text{when } i = j \text{ and } x_i > x_+, \\ 0 & \text{otherwise} \end{cases}$$

and (iii) that  $\mathbb{1}_{\{y_{i,n} > \theta\}} \xrightarrow{\text{a.s.}} \mathbb{1}_{\{y_{i,\infty} > \theta\}}$  whenever  $\theta \neq y_{i,\infty}$ .  $\square$

Our goal for the moment is to characterize the minimum of  $\text{ASE}[\mathbf{x}|\theta]$  with respect to thresholds  $\theta$  strictly above the noise bulk edge,  $\theta > z_+$ . This will give us the optimal *fixed* threshold, in the sense of minimal asymptotic loss (though, at this point, we cannot exclude the possibility that thresholding precisely at  $\theta = z_+$  might achieve better asymptotic risk).

Recall, by equations (4.1) and (4.3), that the asymptotic loss decouples across the signal spikes as

$$\text{ASE}[\mathbf{x}|\theta] = \sum_{i=1}^r R(x_i|\theta), \quad \text{ASE}^*[\mathbf{x}] = \sum_{i=1}^r R^*(x_i).$$

Assuming that  $\theta > z_+$ , we have  $R(x|\theta) = R^*(x) = x^2$  when  $x \leq x_+$ , while for  $x > x_+$ ,

$$R(x|\theta) = \mathbb{1}_{\{\mathcal{Y}(x) \leq \theta\}} \cdot R_0(x) + \mathbb{1}_{\{\mathcal{Y}(x) > \theta\}} \cdot R_1(x), \quad R^*(x) = \min\{R_0(x), R_1(x)\},$$

with

$$R_0(x) = x^2, \quad R_1(x) = x^2 + \mathcal{Y}(x)^2 - 2x\mathcal{Y}(x)\mathcal{C}(x).$$

If we were able to find  $\theta > z_+$  such that  $R(x|\theta) = R^*(x)$  for all  $x > x_+$ , then clearly it achieves minimal asymptotic loss. To do that, it is convenient to introduce a reparameterization  $y = \mathcal{Y}(x)$ , where recall that  $\mathcal{Y}(\cdot)$  is an increasing bijection, mapping  $(x_+, \infty)$  to  $(z_+, \infty)$ . Using equations (3.13) and (3.14), assuming  $x > x_+$ , we get

$$x^2 = (\mathcal{D}(y))^{-1}, \quad \mathcal{C}(x) = -2 \frac{(\mathcal{D}(y))^{3/2}}{\mathcal{D}'(y)},$$

so that

$$R_1(x) - R_0(x) = y^2 - 2xy\mathcal{C}(x) = y^2 + 4y \cdot \frac{\mathcal{D}(y)}{\mathcal{D}'(y)} = y^2 \left( 1 + \frac{4}{\Psi_\gamma(y)} \right),$$

where

$$\Psi_\gamma(y) = y \cdot \frac{\mathcal{D}'(y)}{\mathcal{D}(y)}$$

is as defined in equation (4.4). Since  $\Psi_\gamma(\cdot)$  is negative ( $\mathcal{D}$  is positive and decreasing), we conclude that

$$(6.1) \quad R^*(x) = \mathbb{1}_{\{\Psi_\gamma(y) \leq -4\}} \cdot R_0(x) + \mathbb{1}_{\{\Psi_\gamma(y) > -4\}} \cdot R_1(x).$$

The next lemma establishes some essential properties of  $\Psi_\gamma(y)$ .

LEMMA 8. *Let  $H$  be a compactly supported CDF, with  $\mathcal{Z}_+(H) > 0$ . Let  $\gamma \in (0, 1]$ , and let  $\Psi_\gamma(y; H)$  be defined as in equation (4.4). Then:*

1. *The function  $y \mapsto \Psi_\gamma(y; H)$  is strictly increasing on  $y \in (\mathcal{Z}_+(H), \infty)$ , with  $\lim_{y \rightarrow \infty} \Psi_\gamma(y; H) = -2$ .*
2. *Assume that*

$$\lim_{y \rightarrow \mathcal{Z}_+(F_Z)} \int (y - z)^{-2} dH(z) = \infty.$$

*(This is Assumption 6 for  $H = F_Z$ .) Then  $\lim_{y \rightarrow \mathcal{Z}_+(H)} \Psi_\gamma(y; H) = -\infty$ , and there is a unique point  $y^* \in (\mathcal{Z}_+(F_Z), \infty)$  such that  $\Psi_\gamma(y; H) = -4$ .*

The proof of Lemma 8 appears in Section 1 of the Supplementary Material. An illustration of this lemma and its consequences appear in Figure 5 below.

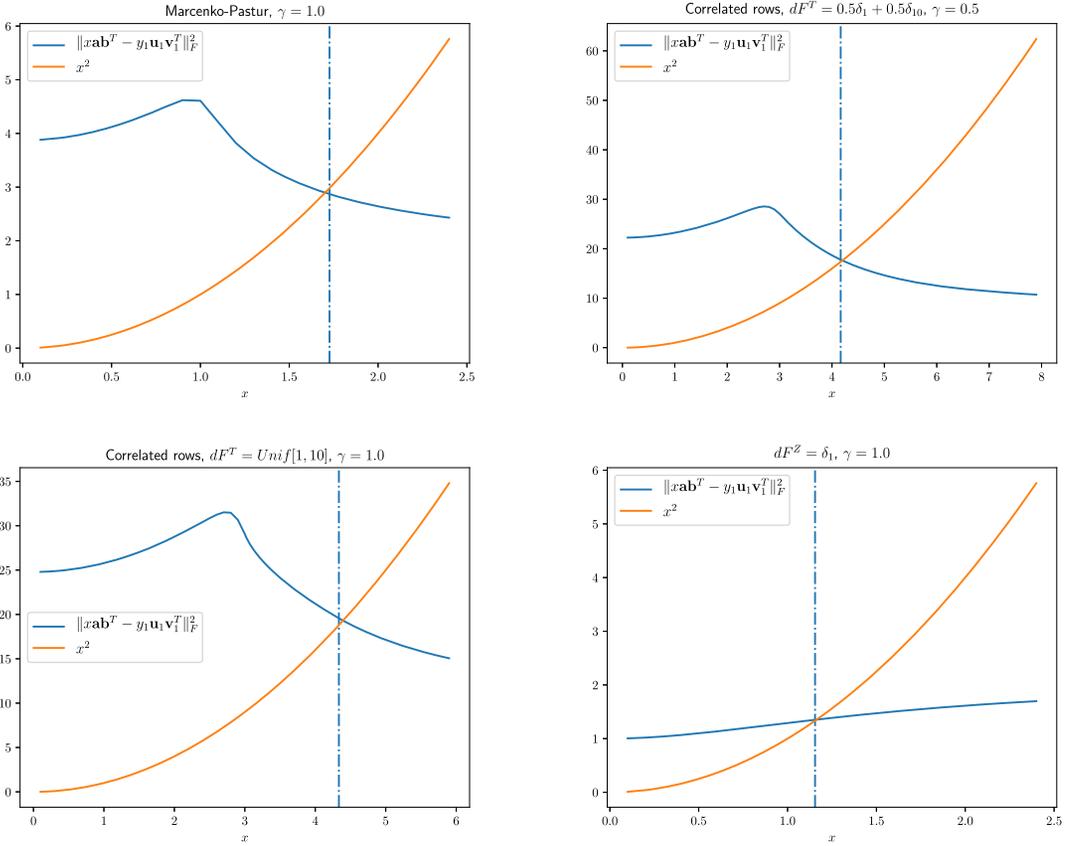


FIG. 5. A numerical illustration of Lemma 8 and its consequences. Assuming a rank-1 signal  $X = x \cdot \mathbf{a}_{1,n} \mathbf{b}_{1,n}^\top$ , set  $R_0(x) = \|X\|_F^2 = x^2$  and  $R_1(x) = \lim_{n \rightarrow \infty} \|X - y_{1,n} \mathbf{a}_{1,n} \mathbf{b}_{1,n}^\top\|_F^2$ . The point  $x^* = \mathcal{Y}^{-1}(T_\gamma(F_Z))$  is the unique crossing point  $R_0(x^*) = R_1(x^*)$ . When  $x < x^*$ , the principal components of  $Y$  are “too noisy,” so that estimating  $\hat{X} = 0$  gives better squared error; when  $x > x^*$ , the situation reverses. At each plot,  $R_0(x)$  and  $R_1(x)$  are plotted as  $x$  varies, for finite  $n$ , fixed signal components  $\mathbf{a}_{1,n} \mathbf{b}_{1,n}^\top$  and a single instance of  $Z_n$ . The dashed vertical line is an estimate of  $x^*$ , obtained by applying the functional  $T_{p/n}(\cdot)$  on  $F_{Z_n}$ , as well as computing the inverse map  $\mathcal{Y}^{-1}(T_{p/n}(F_{Z_n}))$  numerically from  $F_{Z_n}$ . In all cases,  $p = 500$  and  $n = p/\gamma$ . From left to right, top to bottom: (i) Marcenko–Pastur law with shape  $\gamma = 1$ ; (ii) noise matrix with correlated columns, with  $F_S = \frac{1}{2}\delta_1 + \frac{1}{2}\delta_{10}$  and  $\gamma = 0.5$ ; (iii) likewise, with  $F_S = \text{Unif}[1, 10]$ ; (iv)  $F_Z = \delta_1$  and  $\gamma = 1$  (specifically,  $Z_n = I$ ).

*Proof of Lemma 1.* The lemma follows as a straightforward corollary of Lemma 8. Lemma 8 implies that there is a unique number  $T_\gamma(F_Z) > \mathcal{Z}_+(F_Z)$  such that  $\Psi_\gamma(T_\gamma(F_Z); F_Z) = -4$ . Since  $y \mapsto \Psi_\gamma(y; F_Z)$  is increasing, plugging into equation (6.1),

$$(6.2) \quad R^*(x) = R(x|T_\gamma(F_Z)) = \mathbb{1}_{\{y \leq T_\gamma(F_Z)\}} \cdot R_0(x) + \mathbb{1}_{\{y > T_\gamma(F_Z)\}} \cdot R_1(x),$$

where  $x > \mathcal{X}_+$  and  $y = \mathcal{Y}(x)$ . We conclude that  $\text{ASE}^*[\mathbf{x}] = \text{ASE}[\mathbf{x}|T_\gamma(F_Z)]$ . This is the minimum of  $\text{ASE}[\mathbf{x}|\theta]$  over all  $\theta \geq 0$  since clearly  $\text{ASE}[\mathbf{x}|\theta] \geq \text{ASE}^*[\mathbf{x}]$  by definition. Moreover, for any  $\theta \neq T_\gamma(F_Z)$ , we can find some  $y > \mathcal{Z}_+(F_Z)$  such that either  $\theta < y < T_\gamma(F_Z)$  or  $T_\gamma(F_Z) < y < \theta$ . Taking  $x = \mathcal{Y}^{-1}(y)$ , we find that  $R(x|\theta) > R(x|T_\gamma(F_Z)) = R^*(x)$ , since there is a unique crossing point  $x > \mathcal{X}_+$  with  $R_0(x) = R_1(x)$  (because  $y \mapsto \Psi_\gamma(y; F_Z)$  is strictly increasing). Thus, we can construct a signal  $\mathbf{x}$  for which  $\text{ASE}[\mathbf{x}|\theta] > \text{ASE}^*[\mathbf{x}]$  and, therefore,  $T_\gamma(F_Z)$  is the unique threshold, which minimizes  $\text{ASE}[\mathbf{x}|\theta]$  universally for all  $\mathbf{x}$ .

*Proof of Lemma 2.* Part (1) of Lemma 2 follows from Lemma 1, along with the observation that if  $\mathcal{Y}(x) = T_\gamma(F_Z)$ , then  $R_0(x) = R_1(x) = R^*(x)$ ; this means that regardless of whether we threshold slightly above or below  $\mathcal{Y}(x)$ , we get the same asymptotic loss. Part (2) follows

by the same argument as in the proof of Lemma 1, in the paragraph above. Finally, part (3) follows right from the definition of  $\underline{\Theta}(\mathbf{x})$  and  $\overline{\Theta}(\mathbf{x})$ .

6.2. *Achieving oracle loss.* We move on to study the oracle loss  $\text{SE}_n^*[\mathbf{x}]$ . This random variable depends on both  $X_n$  and the noise  $Z_n$ . Denote

$$(6.3) \quad \hat{X}_{[k]} = \sum_{i=1}^k y_{n,i} \mathbf{u}_{i,n} \mathbf{v}_{i,n}^\top, \quad k = 0, \dots, p.$$

That is,  $\hat{X}_{[k]}$  is obtained from  $Y$  by keeping only the top  $k = 0, \dots, p$  singular values (equivalently, hard thresholding at  $t = y_{k+1,n}$ , in case one has  $y_{k+1,n} < y_{k,n}$ ). Any hard thresholding estimator  $\hat{X}_\theta$  obviously corresponds to some  $\hat{X}_{[k]}$  (however if there are multiplicities, possibly not every  $\hat{X}_{[k]}$  is representable by some threshold  $\theta$ ); thus,

$$\text{SE}_n^*[\mathbf{x}] \geq \min_{0 \leq k \leq p} \|X - \hat{X}_{[k]}\|_F^2.$$

We first show that keeping too many singular values is consistently suboptimal.

LEMMA 9. Set  $M = r + 1 + \lceil \frac{\text{ASE}_n^*[\mathbf{x}]}{z_+^2} \rceil$ . Then

$$\mathbb{P}\left\{\exists N \text{ s.t. } \forall n \geq N : \text{SE}_n^*[\mathbf{x}] < \min_{k \geq M} \|X - \hat{X}_{[k]}\|_F^2\right\} = 1.$$

The proof of Lemma 9 is deferred to the Supplementary Material, Section 2. Lemma 9 tells us that to study the oracle loss  $\text{SE}_n^*[\mathbf{x}]$  as  $n \rightarrow \infty$ , we only need, essentially, to study the risk of a *fixed* collection of estimators, the number of whom does not depend on  $n$ ; specifically,  $X_{[k]}$  for  $0 \leq k < M$  obtain formulas for  $\lim_{n \rightarrow \infty} \|X_n - \hat{X}_{[k]}\|_F^2$ . We need to compute the limiting correlations between the underlying signal dyads,  $\mathbf{a}_{1,n} \mathbf{b}_{1,n}^\top, \dots, \mathbf{a}_{r,n} \mathbf{b}_{r,n}^\top$  and the corresponding empirical dyads  $\mathbf{u}_{i,n} \mathbf{v}_{i,n}^\top$ , for all  $1 \leq i < M$ . For empirical spikes up to  $i = r$ , these limiting correlations are computed in [10] (recall equation (3.7)).

The next result shows that, as one would expect, the  $j$ th singular vectors of  $Y$ , for any bounded  $j \geq r + 1$ , are asymptotically uncorrelated with the signal singular vectors.

PROPOSITION 2. For any  $1 \leq i \leq r$  and fixed  $j \neq i$  (not necessarily  $j \leq r$ ), one has

$$\langle \mathbf{a}_{n,i}, \mathbf{u}_{n,j} \rangle \cdot \langle \mathbf{b}_{n,i}, \mathbf{v}_{n,j} \rangle \xrightarrow{\text{a.s.}} 0.$$

The proof of Proposition 2 is deferred to Section 3 of the Supplementary Material. Note that Proposition 2 implies that for any *fixed*  $M$  (meaning  $M$  cannot depend on  $n$ ), the event

$$\{\forall 1 \leq i \leq r, 1 \leq j \leq M, j \neq i : \langle \mathbf{a}_{n,i}, \mathbf{u}_{n,j} \rangle \cdot \langle \mathbf{b}_{n,i}, \mathbf{v}_{n,j} \rangle \longrightarrow 0\}$$

holds with probability 1. The following lemma is an immediate corollary.

LEMMA 10. For any  $k \geq r$ ,

$$\|X - \hat{X}_{[k]}\|_F^2 \xrightarrow{\text{a.s.}} \sum_{i=1}^r [x_i^2 + y_{i,\infty}^2 - 2x_i \cdot y_{i,\infty} \cdot \mathcal{C}(x_i)] + (k - r)z_+^2,$$

where, by way of notation, we use  $\mathcal{C}(x_i) = 0$  for  $x_i \leq x_+$ . In particular,

$$\mathbb{P}\left\{\exists N \text{ s.t. } \forall n \geq N : \text{SE}_n^*[\mathbf{x}] < \min_{k \geq r+1} \|X - \hat{X}_{[k]}\|_F^2\right\} = 1.$$

PROOF. The calculation is straightforward, as in the proof of Lemma 7. For the last part, simply recall that

$$\text{SE}_n[\mathbf{x}|T_\gamma(F_Z)] \xrightarrow{\text{a.s.}} \text{ASE}^*[\mathbf{x}] \leq \sum_{i=1}^r [x_i^2 + y_{i,\infty}^2 - 2x_i \cdot y_{i,\infty} \cdot \mathcal{C}(x_i)]. \quad \square$$

We are ready to prove Theorems 1 and 2.

*Proof of Theorem 1.* By Lemmas 9 and 10, almost surely, there exists  $N$  such that  $\forall n \geq N$ ,

$$\text{SE}_n^*[\mathbf{x}] \geq \min_{0 \leq k \leq r} \|X - \hat{X}_{[k]}\|_F^2.$$

Part (1) then follows from the observation that  $\min_{0 \leq k \leq r} \|X - \hat{X}_{[k]}\|_F^2 \xrightarrow{\text{a.s.}} \text{ASE}^*[\mathbf{x}]$ , as can be deduced from the calculations of Section 6.1, together with  $\text{SE}_n[\mathbf{x}|T_\gamma(F_Z)] \xrightarrow{\text{a.s.}} \text{ASE}^*[\mathbf{x}]$ . We now prove (2). Let us assume, for ease of notation, that  $T_\gamma(F_Z) \notin \{y_{1,\infty}, \dots, y_{r,\infty}\}$ . In that case, the asymptotic optimal interval is just the interval between two consecutive spikes, say,

$$\underline{\Theta}(\mathbf{x}) = y_{k^*+1,\infty}, \quad \overline{\Theta}(\mathbf{x}) = y_{k^*,\infty},$$

where  $y_{0,\infty} = \infty$ . Note that if  $T_\gamma(F_Z) = y_{k,\infty}$  for some  $k \geq 1$ , then  $\underline{\Theta}(\mathbf{x}) = y_{k+1,\infty}$ ,  $\overline{\Theta}(\mathbf{x}) = y_{k-1,\infty}$ . With probability one, for large enough  $n$ ,  $X_{\theta_n} = \hat{X}_{[k^*]}$ . Now, recall that  $\|X_n - \hat{X}_{[k^*]}\|_F^2 \xrightarrow{\text{a.s.}} \text{ASE}^*[\mathbf{x}]$ .

*Proof of Theorem 2.* Let  $k^*$  be as in the proof of Theorem 1. We know, from Lemmas 9, 10 and the definition of the asymptotic optimal interval, that  $\|X - \hat{X}_{[k^*]}\|_F^2 \xrightarrow{\text{a.s.}} \text{ASE}^*[\mathbf{x}]$  and that almost surely,  $\liminf_{n \rightarrow \infty} \min_{k \neq k^*} \|X - \hat{X}_{[k]}\|_F^2 > \text{ASE}^*[\mathbf{x}]$  (here is where we assume that there is no  $y_{i,\infty}$  that equals  $T_\gamma(F_Z)$ ). Thus,

$$\mathbb{P}\left\{\exists N \text{ s.t. } \forall n \geq N : \text{SE}_n^*[\mathbf{x}] = \|X - \hat{X}_{[k^*]}\|_F^2 < \min_{k \neq k^*} \|X - \hat{X}_{[k]}\|_F^2\right\} = 1.$$

The proof follows by noting that: (i) If  $\theta \in (\underline{\Theta}(\mathbf{x}), \overline{\Theta}(\mathbf{x}))$ , then with probability 1, for all large enough  $n$ ,  $\hat{X}_{\theta_n} = \hat{X}_{[k^*]}$ ; (ii) If  $\theta \notin [\underline{\Theta}(\mathbf{x}), \overline{\Theta}(\mathbf{x})]$ , then with probability 1, for large enough  $n$ ,  $\hat{X}_{\theta_n} \neq \hat{X}_{[k^*]}$ .

**Acknowledgments.** We are grateful to the anonymous reviewers for their thoughtful comments, which have helped improve this manuscript considerably.

**Funding.** DD was supported in part by NSF Grants DMS-1407813, 1418362 and 1811614. MG was supported in part by Israel Science Foundation grant 1523/16 and 871/22. This work was made possible by United States–Israel Binational Science Foundation (BSF) Grant 2016201 “Frontiers of Matrix Recovery.” ER was affiliated with the School of Computer Science and Engineering, the Hebrew University of Jerusalem, and supported in part by Israel Science Foundation grant 1523/16 and an Einstein–Kaye Fellowship from the Hebrew University of Jerusalem.

### SUPPLEMENTARY MATERIAL

**Proofs and additional empirical results** (DOI: [10.1214/22-AOS2232SUPP](https://doi.org/10.1214/22-AOS2232SUPP); .pdf). To keep within the space constraints, the proofs of several technical claims are deferred to the Supplementary Material. In addition, in the Supplementary Material we provide extensive numerical experiments validating different aspects of our results under various noise distributions.

## REFERENCES

- [1] ACHLIOPTAS, D. and MCSHERRY, F. (2001). Fast computation of low rank matrix approximations. In *Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing* 611–618. ACM, New York. MR2120364 <https://doi.org/10.1145/380752.380858>
- [2] ALTER, O., BROWN, P. O. and BOTSTEIN, D. (2000). Singular value decomposition for genome-wide expression data processing and modeling. *Proc. Natl. Acad. Sci. USA* **97** 10101–10106. <https://doi.org/10.1073/pnas.97.18.10101>
- [3] AZAR, Y., FIAT, A., KARLIN, A. R., MCSHERRY, F. and SAIA, J. (2001). Spectral analysis of data. In *Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing* 619–626.
- [4] BAI, Z. and SILVERSTEIN, J. W. (2010). *Spectral Analysis of Large Dimensional Random Matrices*, 2nd ed. *Springer Series in Statistics*. Springer, New York. MR2567175 <https://doi.org/10.1007/978-1-4419-0661-8>
- [5] BAI, Z. and YAO, J. (2008). Central limit theorems for eigenvalues in a spiked population model. *Ann. Inst. Henri Poincaré Probab. Stat.* **44** 447–474. MR2451053 <https://doi.org/10.1214/07-AIHP118>
- [6] BAI, Z. D. and SILVERSTEIN, J. W. (1998). No eigenvalues outside the support of the limiting spectral distribution of large-dimensional sample covariance matrices. *Ann. Probab.* **26** 316–345. MR1617051 <https://doi.org/10.1214/aop/1022855421>
- [7] BAI, Z. D. and SILVERSTEIN, J. W. (2004). CLT for linear spectral statistics of large-dimensional sample covariance matrices. *Ann. Probab.* **32** 553–605. MR2040792 <https://doi.org/10.1214/aop/1078415845>
- [8] BAIK, J., BEN AROUS, G. and PÉCHÉ, S. (2005). Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *Ann. Probab.* **33** 1643–1697. MR2165575 <https://doi.org/10.1214/009117905000000233>
- [9] BAIK, J. and SILVERSTEIN, J. W. (2006). Eigenvalues of large sample covariance matrices of spiked population models. *J. Multivariate Anal.* **97** 1382–1408. MR2279680 <https://doi.org/10.1016/j.jmva.2005.08.003>
- [10] BENAYCH-GEORGES, F. and NADAKUDITI, R. R. (2012). The singular values and vectors of low rank perturbations of large rectangular random matrices. *J. Multivariate Anal.* **111** 120–135. MR2944410 <https://doi.org/10.1016/j.jmva.2012.04.019>
- [11] BICKEL, P. J. and LEVINA, E. (2008). Covariance regularization by thresholding. *Ann. Statist.* **36** 2577–2604. MR2485008 <https://doi.org/10.1214/08-AOS600>
- [12] CATTELL, R. B. (1966). The Scree test for the number of factors. *Multivar. Behav. Res.* **1** 245–276. [https://doi.org/10.1207/s15327906mbr0102\\_10](https://doi.org/10.1207/s15327906mbr0102_10)
- [13] CHATTERJEE, S. (2015). Matrix estimation by universal singular value thresholding. *Ann. Statist.* **43** 177–214. MR3285604 <https://doi.org/10.1214/14-AOS1272>
- [14] DOBRIBAN, E. and OWEN, A. B. (2019). Deterministic parallel analysis: An improved method for selecting factors and principal components. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **81** 163–183. MR3904784 <https://doi.org/10.1111/rssb.12301>
- [15] DONOHO, D., GAVISH, M. and ROMANOV, E. (2020). Code supplement for “ScreeNOT: Exact MSE-optimal singular value thresholding in correlated noise”. Available at <https://purl.stanford.edu/py196rk3919>.
- [16] DONOHO, D., GAVISH, M. and ROMANOV, E. (2023). Supplement to “ScreeNOT: Exact MSE-optimal singular value thresholding in correlated noise.” <https://doi.org/10.1214/22-AOS2232SUPP>
- [17] EDFORS, O. and SANDELL, M. (1998). OFDM channel estimation by singular value decomposition. *IEEE Trans. Commun.* **46** 931–939.
- [18] FRANKLIN, S. B., GIBSON, D. J., ROBERTSON, P. A., POHLMANN, J. T. and FRALISH, J. S. (1995). Parallel analysis: A method for determining significant principal components. *J. Veg. Sci.* **6** 99–106.
- [19] GAVISH, M. and DONOHO, D. L. (2014). The optimal hard threshold for singular values is  $4/\sqrt{3}$ . *IEEE Trans. Inf. Theory* **60** 5040–5053. MR3245370 <https://doi.org/10.1109/TIT.2014.2323359>
- [20] HOFF, P. D. (2007). Model averaging and dimension selection for the singular value decomposition. *J. Amer. Statist. Assoc.* **102** 674–685. MR2325118 <https://doi.org/10.1198/016214506000001310>
- [21] HONG, D., BALZANO, L. and FESSLER, J. A. (2018). Asymptotic performance of PCA for high-dimensional heteroscedastic data. *J. Multivariate Anal.* **167** 435–452. MR3830656 <https://doi.org/10.1016/j.jmva.2018.06.002>
- [22] JACKSON, D. A. (1993). Stopping rules in principal components analysis: A comparison of heuristical and statistical approaches. *Ecology*.
- [23] JOHNSTONE, I. M. (2001). On the distribution of the largest eigenvalue in principal components analysis. *Ann. Statist.* **29** 295–327. MR1863961 <https://doi.org/10.1214/aos/1009210544>
- [24] JOLLIFFE, I. T. (2005). *Principal Component Analysis*. *Springer Series in Statistics*. Springer, New York. MR0841268 <https://doi.org/10.1007/978-1-4757-1904-8>

- [25] LAGERLUND, T. D., SHARBROUGH, F. W. and BUSACKER, N. E. (1997). Spatial filtering of multichannel electroencephalographic recordings through principal component analysis by singular value decomposition. *J. Clin. Neurophysiol.* **14** 73–82. <https://doi.org/10.1097/00004691-199701000-00007>
- [26] LEEB, W. and ROMANOV, E. (2021). Optimal spectral shrinkage and PCA with heteroscedastic noise. *IEEE Trans. Inf. Theory* **67** 3009–3037. MR4282401 <https://doi.org/10.1109/tit.2021.3055075>
- [27] NADAKUDITI, R. R. (2014). OptShrink: An algorithm for improved low-rank signal matrix denoising by optimal, data-driven singular value shrinkage. *IEEE Trans. Inf. Theory* **60** 3002–3018. MR3200641 <https://doi.org/10.1109/TIT.2014.2311661>
- [28] OWEN, A. B. and PERRY, P. O. (2009). Bi-cross-validation of the SVD and the nonnegative matrix factorization. *Ann. Appl. Stat.* **3** 564–594. MR2750673 <https://doi.org/10.1214/08-AOAS227>
- [29] PAUL, D. (2007). Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statist. Sinica* **17** 1617–1642. MR2399865
- [30] PERRY, P. O. (2009). Cross validation for unsupervised learning. PhD thesis, Stanford Univ.
- [31] PRICE, A. L., PATTERSON, N. J., PLENGE, R. M., WEINBLATT, M. E., SHADICK, N. A. and REICH, D. (2006). Principal components analysis corrects for stratification in genome-wide association studies. *Nat. Genet.* **38** 904–909. <https://doi.org/10.1038/ng1847>
- [32] SHABALIN, A. A. and NOBEL, A. B. (2013). Reconstruction of a low-rank matrix in the presence of Gaussian noise. *J. Multivariate Anal.* **118** 67–76. MR3054091 <https://doi.org/10.1016/j.jmva.2013.03.005>
- [33] SILVERSTEIN, J. W. and CHOI, S.-I. (1995). Analysis of the limiting spectral distribution of large-dimensional random matrices. *J. Multivariate Anal.* **54** 295–309. MR1345541 <https://doi.org/10.1006/jmva.1995.1058>
- [34] WOLD, S. (1978). Cross-validatory estimation of the number of components in factor and principal components models. *Technometrics* **20** 397–405.