



HAL
open science

High-dimensional analysis of ridge regression for non-identically distributed data with a variance profile

Jérémie Bigot, Issa-Mbenard Dabo, Camille Male

► **To cite this version:**

Jérémie Bigot, Issa-Mbenard Dabo, Camille Male. High-dimensional analysis of ridge regression for non-identically distributed data with a variance profile. 2024. hal-04559313v2

HAL Id: hal-04559313

<https://hal.science/hal-04559313v2>

Preprint submitted on 13 May 2024

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

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



Distributed under a Creative Commons Attribution 4.0 International License

High-dimensional analysis of ridge regression for non-identically distributed data with a variance profile

Jérémie Bigot & Issa-Mbenard Dabo & Camille Male
Institut de mathématiques de Bordeaux & CNRS (UMR 5251)
Université de Bordeaux, France

May 13, 2024

Abstract

High-dimensional linear regression has been thoroughly studied in the context of independent and identically distributed data. We propose to investigate high-dimensional regression models for independent but non-identically distributed data. To this end, we suppose that the set of observed predictors (or features) is a random matrix with a variance profile and with dimensions growing at a proportional rate. Assuming a random effect model, we study the predictive risk of the ridge estimator for linear regression with such a variance profile. In this setting, we provide deterministic equivalents of this risk and of the degree of freedom of the ridge estimator. For certain class of variance profile, our work highlights the emergence of the well-known double descent phenomenon in high-dimensional regression for the minimum norm least-squares estimator when the ridge regularization parameter goes to zero. We also exhibit variance profiles for which the shape of this predictive risk differs from double descent. The proofs of our results are based on tools from random matrix theory in the presence of a variance profile that have not been considered so far to study regression models. Numerical experiments are provided to show the accuracy of the aforementioned deterministic equivalents on the computation of the predictive risk of ridge regression. We also investigate the similarities and differences that exist with the standard setting of independent and identically distributed data.

Keywords: High-dimensional linear ridge regression; Non-identically distributed data; Degrees of freedom; Double descent; Variance profile; Heteroscedasticity; Random Matrices; Deterministic equivalents.

1 Introduction

High-dimensionality is a subject of interest in the field of statistics, especially in regression problems, driven by the advent of massive data sets. This context gives rise to unexpected phenomena and contradictions with established statistical heuristics when the dimension p of the predictors is fixed and the number n of observations tends to infinity. These phenomena particularly appear in the context of linear regression. Indeed, as the sample size and dimension of acquired data increase, the study of this model is different from the classical framework. In the asymptotic regime where $\min(n, p) \rightarrow +\infty$ and $\frac{p}{n} \rightarrow c > 0$, one can notably mention the occurrence of the double descent phenomenon corresponding to estimators that both interpolates the data and show good generalization performances [BHMM19]. In this asymptotic setting, using tools from random matrix theory (RMT), many authors have therefore focused on the consequences of high-dimensionality on linear regression, see e.g. [DW18, Bac24, HMRT22, LC18]

and references therein. In this paper, we focus on the linear regression model

$$Y_n = X_n \theta_* + \varepsilon_n, \quad (1.1)$$

where X_n is $n \times p$ matrix of random predictors, $\varepsilon_n \in \mathbb{R}^n$ is a noise vector independent of X_n with $\mathbb{E}[\varepsilon_n] = 0$ and $\mathbb{E}[\varepsilon_n \varepsilon_n^t] = \sigma^2 I_n$, $\theta_* \in \mathbb{R}^p$ is a vector of unknown parameters, and $Y_n \in \mathbb{R}^n$ is the vector of observed responses.

Classically, the predictors are assumed to be independent and identically distributed (iid) data, meaning that the rows of the matrix X_n are independent vectors sampled from the same probability distribution. In this paper, we propose to depart from this assumption by considering the setting where the rows of X_n are independent but non-identically distributed. To this end, we suppose that X_n is expressed in the following form

$$X_n = \Upsilon_n \circ Z_n,$$

where \circ denotes the Hadamard product between two matrices, $Z_n = (Z_{ij})$ has iid centered entries with variance one, and $\Upsilon_n = (\gamma_{ij}^{(n)})$ is a deterministic matrix. To simplify the notation, we shall sometimes write $\gamma_{ij}^{(n)} = \gamma_{ij}$ and thus drop the (possibly) dependence of γ_{ij} on n . The matrix

$$\Gamma_n = (\gamma_{ij}^2) \in \mathbb{R}^{n \times p}$$

governs the variance of the entries of X_n , and it is called a *variance profile*.

In the RMT literature, there exist various works on the analysis of the spectrum of large random matrices with a variance profile [Shl96, HLN07, ACD⁺21, BM20, EYY12, AEK17a]. In particular, we rely on results from [HLN07] to obtain a deterministic equivalent of the spectral distribution of a data matrix $X_n = \Upsilon_n \circ Z_n$ with a variance profile matrix Γ_n . The motivation for studying linear regression using such a variance profile is to consider the setting where one has n independent pairs of observations $(Y_i, X_i)_{1 \leq i \leq n}$ (with $X_i = (X_{ij})_{1 \leq j \leq p}$) that are not necessarily identically distributed. Note that in the standard setting of iid data, one that

$$\gamma_{ij} = \gamma_j \quad \text{for all } 1 \leq i \leq n, \text{ and } 1 \leq j \leq p.$$

The main goal of this paper is then to understand how assuming such a variance profile for X_n influences the statistical properties of ridge regression in the linear model (1.1) when compared to the standard assumption of iid observations. In this setting, our approach also allows to analyze the performances of the minimum norm least-squares estimator when the ridge regularization parameter goes to zero.

We consider the high-dimensional context (with p growing to infinity at a rate proportional to n) for which the least squares estimator is possibly not uniquely defined. Thus, we focus our analysis on the ridge regression estimator that is the minimizer of the following loss function

$$\hat{\theta}_\lambda = \arg \min_{\theta \in \mathbb{R}^p} \frac{1}{n} \|Y_n - X_n \theta\|^2 + \lambda \|\theta\|^2,$$

for some regularization parameter $\lambda > 0$. Regardless of the ratio between n and p , this estimator has the following explicit expression

$$\hat{\theta}_\lambda = (X_n^t X_n + n\lambda I_p)^{-1} X_n^t Y_n = X_n^t (X_n X_n^t + n\lambda I_n)^{-1} Y_n. \quad (1.2)$$

Our analysis also includes the study of the minimum least-square estimator defined as

$$\hat{\theta} = \arg \min_{\theta \in \mathbb{R}^p} \left\{ \|\theta\| : \theta \text{ minimizes } \frac{1}{n} \|Y_n - X_n \theta\|^2 \right\},$$

to which the ridge regression estimator converges when λ tends to zero. The estimator $\hat{\theta}$ is also known to be the solution found by gradient descent when initialized to zero, see e.g. [HMRT22][Proposition 1].

To study the statistical performances of the ridge regression estimator, we analyse its predictive risk defined as

$$\hat{r}_\lambda(X_n) = \mathbb{E}[(\tilde{Y} - \tilde{X}^t \hat{\theta}_\lambda)^2 | X_n], \quad (1.3)$$

where $(\tilde{Y}, \tilde{X}) \in \mathbb{R} \times \mathbb{R}^p$ is independent from (Y_n, X_n) and satisfies

$$\tilde{Y} = \tilde{X}^t \theta_* + \tilde{\varepsilon}, \text{ with } \mathbb{E}[\tilde{\varepsilon}] = 0, \mathbb{E}[\tilde{\varepsilon}^2] = \sigma^2 \text{ and } \tilde{X}, \tilde{\varepsilon} \text{ independent.}$$

In the above formula, $\tilde{X} = \tilde{S}_p^{1/2} \tilde{Z}$ with $\tilde{Z} \in \mathbb{R}^p$ a random vector with iid centered entries and variance one, and

$$\tilde{S}_p = \mathbb{E}[\tilde{X} \tilde{X}^t] = \text{diag}(\tilde{\gamma}_1^2, \dots, \tilde{\gamma}_p^2)$$

denotes the variance profile of \tilde{X} . Note that the risk $\hat{r}_\lambda(X_n)$ is conditioned on the predictors X_n , and it is thus a random variable.

Following [DW18], we focus on a random-effect hypothesis by assuming that the components of the vector θ_* are drawn independently at random. As argued in [DW18], this assumption corresponds to an average case analysis over a set of dense regression coefficients as opposed to the ‘‘sparsity hypothesis’’ [HTW15] or the ‘‘manifold hypothesis’’ [LHT23] that are other popular assumptions in high-dimensional linear regression. More precisely, the following random coefficients assumption is made throughout the paper.

Assumption 1.1. *The vector θ_* of regression coefficients is random, independent from $X_n, \tilde{X}, \varepsilon_n$ and $\tilde{\varepsilon}$, with $\mathbb{E}[\theta_*] = 0$ and*

$$\mathbb{E}[\theta_* \theta_*^t] = \frac{\alpha^2}{p} I_p.$$

The above coefficient $\alpha > 0$ represents the average amount of signal strength in model (1.1). Under Assumption 1.1, the expectation in (1.3) used to define the predictive risk is thus taken with respect to both the randomness of the vector of coefficients θ_* , the vector \tilde{X} and the additive noise ε_n and $\tilde{\varepsilon}$.

1.1 Main contributions

Recall that the estimation of Y_n by ridge regression is

$$\hat{Y}_\lambda = X_n \hat{\theta}_\lambda = A_\lambda Y_n, \quad \text{where } A_\lambda = X_n (X_n^t X_n + n\lambda I_p)^{-1} X_n^t.$$

Then, the so-called degrees of freedom (DOF) of the estimator $\hat{\theta}_\lambda$ that is defined as

$$\hat{df}_1(\lambda) = \text{Tr}[A_\lambda] = \text{Tr}[\hat{\Sigma}_n (\hat{\Sigma}_n + \lambda I_p)^{-1}], \text{ where } \hat{\Sigma}_n = \frac{1}{n} X_n^t X_n,$$

represents the so-called effective dimension of the linear estimator \hat{Y}_λ . The DOF is widely used in statistics to define various criteria for model selection among a collection of estimators, see e.g. [Efr04].

Inspired by recent results from [Bac24] in the setting of iid data, a first contribution of this work is to prove the following deterministic equivalence of the DOF

$$\hat{df}_1(\lambda) \sim df_1(\lambda), \quad \text{where } df_1(\lambda) = \text{Tr}[\Sigma_n (\Sigma_n + \kappa(\lambda))^{-1}], \quad (1.4)$$

where

$$\Sigma_n = \mathbb{E}[\widehat{\Sigma}_n] = \frac{1}{n} \text{diag} \left(\sum_{i=1}^n \gamma_{i1}^2, \dots, \sum_{i=1}^n \gamma_{ip}^2 \right),$$

and $\kappa(\lambda)$ is diagonal matrix that depends upon the regularization parameter λ and the variance profile matrix Γ_n .

Throughout the paper, the meaning of the equivalence notation $A_n \sim B_n$ between two random variables is

$$\lim_{n \rightarrow \infty, p/n \rightarrow c} |A_n - B_n| = 0, \text{ almost surely.}$$

Hence, the equivalence relation (1.4) indicates that the DOF of the ridge regression estimator for the empirical covariance matrix $\widehat{\Sigma}_n$ corresponds to the DOF computed with its expected version Σ_n (the usual population covariance matrix for iid data), and another additive regularization structure than λI_p that is given by the diagonal matrix $\kappa(\lambda)$ whose explicit expression is given in Section 3.

Then, the second and main contribution of the paper is to derive a deterministic equivalent of the predictive risk $\hat{r}_\lambda(X_n)$ in the case where the number of samples n and the dimension p tend to infinity at a proportional rate, that is $\lim_{n \rightarrow +\infty} \frac{p}{n} = c > 0$. This deterministic equivalent allows us to understand the influence of the ratio c on the predictive risk and to also analyze the effect of the signal strength α . We also study the convergence of the predictive risk as λ tends to 0 to analyze the statistical properties the minimum norm least square estimator. In this setting, it appears a phenomenon arising from the curse of dimensionality that is commonly known as double descent for iid data. This phenomenon contradicts the consensus heuristic that, when a model becomes over-parameterized, then the predictive risk increases due to overfitting of the training data and the model is no longer capable of generalizing. This double descent has been thoroughly studied in the case of high-dimensional linear regression using tool from RMT, see e.g. [HMRT22, Bac24, BHX20] and references therein. In this paper, we show that it also occurs for non iid data with a variance profile. Our deterministic equivalent of the predictive risk also allows to derive the asymptotic behavior of an optimal choice for the regularization parameter and to compare it to the one obtained for iid data in [DW18].

As a third contribution, using synthetic data and various illustrative examples of variance profile, we conduct numerical experiments to verify the accuracy of our deterministic equivalent of the predictive risk using finite samples. We also investigate the similarities and differences that exist between the standard setting of iid data and the one of non-identically distributed data with a variance profile.

For example, if the variance profile is assumed to be quasi doubly stochastic in the sense that

$$\frac{1}{n} \sum_{j=1}^p \gamma_{ij}^2 = \frac{p}{n}, \text{ for all } 1 \leq i \leq n, \quad \text{and} \quad \frac{1}{n} \sum_{i=1}^n \gamma_{ij}^2 = 1, \text{ for all } 1 \leq j \leq p,$$

then, we prove in this paper that

$$\lim_{n \rightarrow \infty, p/n \rightarrow c} \hat{r}_0(X_n) = \begin{cases} \sigma^2 \frac{c}{1-c} + \sigma^2 & \text{if } c < 1 \\ \alpha^2 \left(1 - \frac{1}{c}\right) + \sigma^2 \frac{1}{c-1} & \text{if } c > 1 \end{cases}, \text{ almost surely.} \quad (1.5)$$

The above result corresponds to the known asymptotic limit of the predictive risk of the minimum norm least squares estimator for iid data when the entries of X_n are independent centered random variables with variance 1, see e.g. [HMRT22], which is referred to as a constant variance profile in this paper (that is $\gamma_{ij} = 1$). In this setting, the predictive risk of $\hat{\theta}$ is thus increasing with p up to $p/n < 1$ and then decreasing for $p > n$ (that is beyond the interpolation threshold $p = n$) which is classically referred to as the double descent phenomenon as illustrated in Figure 1(a).

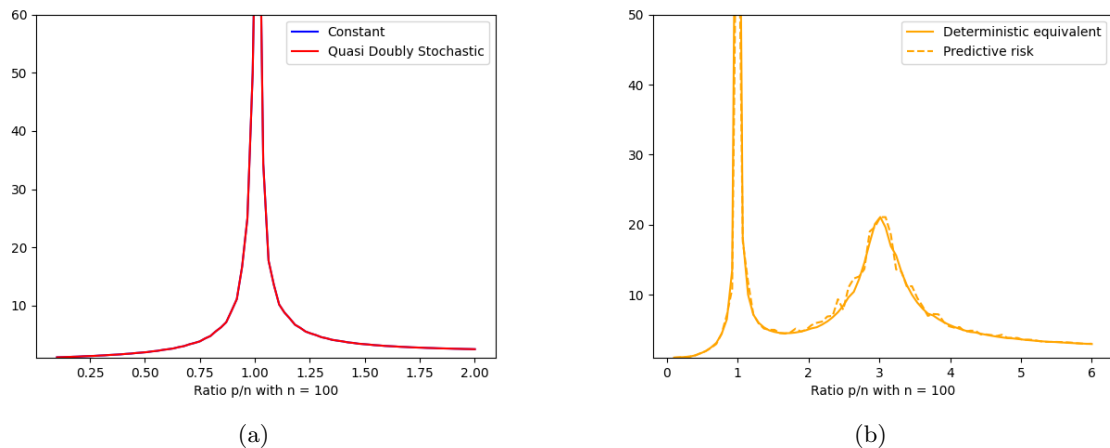


Figure 1: Predictive risk for several variance profiles with $\lambda = 0$. (a) Comparison of constant and quasi doubly stochastic variance profiles with $\alpha = \sigma = 1$, $n = 100$ and p varying from 10 to 200. (b) Piecewise constant variance profile with $\gamma_1^2 = 0.0005$, $\gamma_2^2 = 1$, $\alpha = \sigma = 1$, $n = 100$ and p varying from 10 to 600.

By going beyond the quasi doubly stochastic assumption (1.5), the results of this paper, on the asymptotic limit of the predictive risk, also allow to exhibit variance profiles for which the predictive risk has a shape that differs from double descent. This is illustrated in Figure 1(b), where the deterministic equivalent of the predictive risk has a triple descent behavior (as a function of p/n) for the following piecewise constant variance profile

$$\Gamma_n = \begin{bmatrix} \gamma_1^2 \mathbb{1}_{n/4} \mathbb{1}_{p/4}^t & \gamma_2^2 \mathbb{1}_{n/4} \mathbb{1}_{3p/4}^t \\ \gamma_2^2 \mathbb{1}_{3n/4} \mathbb{1}_{p/4}^t & \gamma_1^2 \mathbb{1}_{3n/4} \mathbb{1}_{3p/4}^t \end{bmatrix},$$

where $\mathbb{1}_q$ denotes the vector of length q with all entries equal to one, and γ_1, γ_2 are positive constant such that $\gamma_2 \gg \gamma_1$. Note that Figure 1(b) also illustrates the accuracy of the deterministic equivalent of the predictive risk that is proposed in this paper.

1.2 Organisation of the paper

In Section 2, we review various works related to the analysis of high-dimensional linear regression and the study of random matrices with a variance profile. The main results are presented and discussed in Section 3. Numerical experiments are then reported in Section 4. A conclusion and some perspectives are proposed in Section 5. All proofs using tools from the theory of random matrices and operator-valued Stieltjes transforms are deferred to an Appendix where we also discuss the use of random matrices with a variance profile in free probability.

1.3 Publicly available source code

For the sake of reproducible research, a Python code is available at the following address:

<https://github.com/Issoudab/RidgeRegressionVarianceProfile>

to implement the experiments carried out in this paper.

2 Related works

In this section, we review the literature on the analysis of high-dimensional regression using tools from RMT. We also discuss existing works on linear regression for non-iid data, and the use of random matrices with a variance profile in statistics and RMT.

2.1 High-dimensional linear regression from the random matrix perspective

When the sample size is comparable to the dimensionality of the observations, recent advances in RMT have been successfully applied to various inference problems in high-dimensional multivariate statistics, see e.g. [NPW21] for a recent overview. Many works have considered the high-dimensional analysis of the linear model using tools from RMT for iid data with a general covariance structure $\Sigma \in \mathbb{R}^{p \times p}$ (assumed to be a positive semi-definite matrix) that is for

$$X_n = Z_n \Sigma^{1/2}, \text{ for an } n \times p \text{ matrix } Z_n \text{ with iid centered entries having variance one.}$$

In particular, for such data, the study of the minimum norm least-square estimator and the double descent behavior of the predictive risk has been considered in [HMRT22, Bac24, BHX20, RMR21]. The analysis of the predictive risk of ridge regression using iid data with a general covariance structure has been studied in [DW18, Bac24], while previous works on the statistical analysis of ridge regression from the RMT perspective include [EK18, Dic16] and [CD11, TV04] for applications in wireless communication.

2.2 Linear regression for independent but non-identically distributed data

While the statistical analysis of linear regression for iid data with a general covariance structure is very well understood, the literature on the study of the linear model for non-identically distributed predictors appears to be scarcer. A first analysis of maximum likelihood estimation in standard models (including linear regression) for independent but non-identically distributed data dates back to [Ber82]. More recent works [BBK⁺19, KBB⁺20], on statistical inference in linear regression in the so-called model-free framework, allow to consider the setting on non-identically distributed predictors. However, to the best of our knowledge, the high-dimensional analysis of the linear model using non-identically distributed data has not considered so far.

2.3 The use of variance profile in RMT

RMT allows to describe the asymptotic distribution of the eigenvalues of large matrices with random entries, see e.g. [BS10]. In particular, the well-known Marchenko-Pastur theorem characterizes the limiting spectral distribution of the covariance matrix $\widehat{\Sigma}_n = \frac{1}{n} X_n^t X_n$ for a data matrix $X_n = Z_n \Sigma^{1/2}$ with iid rows in the asymptotic setting $\lim_{n \rightarrow +\infty} \frac{p}{n} = c > 0$.

However, in many applications, such as photon imaging [SHDW14], network traffic analysis [BMG13], ecology [AGB⁺15, AT15], neurosciences [ASS15, ARS15] or genomics for microbiome studies [CZL19], the amount of variance in the observed data matrix may significantly change from one sample to the other, that is between the rows of X_n . The literature on statistical inference from high-dimension matrices with heteroscedasticity has thus recently been growing [BM20, BDF17, LDS18, UHZB16, ZCW22, JFL22]. Modeling data as a random matrix $X_n = \Upsilon_n \circ Z_n$ with a variance profile to handle the setting of non-iid data has also found applications in the analysis of the performances of wireless digital communication channels [HLN07]. In the RMT literature, Hermitian random matrices with centered entries but non-equal distribution are

referred to as generalized Wigner matrices for which many asymptotic properties are now well understood. For example, for Hermitian random matrices with a variance profile that is doubly stochastic (namely its rows and columns elements sum up to one), bulk universality at optimal spectral resolution for local spectral statistics have been established in [EYY12] and they are shown to converge to those of a standard Wigner matrix (that is with iid sub-diagonal entries). The case of a generalized Wigner matrix with a variance profile that is not necessarily doubly stochastic has been studied in [AEK17a], and non-hermitian random matrices with a variance profile have been considered in [CHNR18, HLN07, HLN06] using the notion of deterministic equivalent that consists in approximating the spectral distribution of a random matrix by a deterministic function. A recent work [BHX23] considers the analysis of approximate message passing for statistical estimation problems involving a data matrix with a variance profile, and an application to ridge regression to characterize the non-asymptotic distribution of the ridge estimator is proposed.

Let us now recall the key notion of Stieltjes transform.

Definition 2.1. *Let μ be a probability measure supported on \mathbb{R} . Then, its Stieltjes transform is defined as*

$$g_\mu(z) = \int_{\mathbb{R}} \frac{1}{t-z} d\mu(t), \quad \text{for } t \in \mathbb{C}^+,$$

where $\mathbb{C}^+ = \{z \in \mathbb{C}, \Im(z) > 0\}$ and $\Im(\cdot)$ denotes the imaginary part of a complex number.

Then, we build upon results from [HLN07] to construct deterministic equivalents (when $\lim_{n \rightarrow +\infty} p/n \rightarrow c > 0$) of the Stieltjes transforms

$$g_{\hat{\mu}_n}(z) = \frac{1}{p} \text{Tr}[(\hat{\Sigma}_n - zI_p)^{-1}], \quad \text{for } z \in \mathbb{C} \setminus \mathbb{R}^+,$$

and

$$g_{\tilde{\mu}_n}(z) = \frac{1}{n} \text{Tr}[(\tilde{\Sigma}_n - zI_n)^{-1}], \quad \text{for } z \in \mathbb{C} \setminus \mathbb{R}^+,$$

of the empirical eigenvalue distribution $\hat{\mu}_n$ of $\hat{\Sigma}_n$, and the empirical eigenvalue distribution $\tilde{\mu}_n$ of $\tilde{\Sigma}_n = \frac{1}{n} X_n X_n^t$ respectively.

To obtain these deterministic equivalents, the following assumptions will be needed, on mild moment conditions on the Z_{ij} 's and boundedness hypotheses on the entries of the variance profile matrix.

Assumption 2.1. *There exists $\delta > 0$ such that, $\mathbb{E}[|Z_{ij}|^{4+\delta}] < +\infty$, for all $1 \leq i \leq n$ and $1 \leq j \leq p$.*

Assumption 2.2. *There exists $\gamma_{\max} > 0$ such that, $\sup_{n \geq 1} \max_{\substack{1 \leq i \leq n \\ 1 \leq j \leq p}} |\gamma_{ij}^{(n)}| < \gamma_{\max}$.*

Assumption 2.3. *There exists $\gamma_{\min} > 0$ such that, $\forall n \geq 1, \min_{\substack{1 \leq i \leq n \\ 1 \leq j \leq n}} |\gamma_{ij}^{(n)}| \geq \gamma_{\min}$.*

Assumption 2.4. *There exists $d_* > 0$ such that, $|\frac{p}{n} - 1| \geq d_*$ for all values of n and p .*

Throughout the paper, we also assume that $\tilde{X} = \tilde{S}_p^{1/2} \tilde{Z}$ with

$$\mathbb{E}[|\tilde{Z}_j|^{4+\delta}] < +\infty, \quad \max_{1 \leq j \leq p} |\tilde{\gamma}_j^{(n)}| \leq \gamma_{\max} \text{ and } \min_{1 \leq j \leq n} |\tilde{\gamma}_j^{(n)}| \geq \gamma_{\min}$$

and for all $n \geq 1$. Then, using [HLN07][Theorem 2.5], the following result holds.

Proposition 2.1. *Under Assumptions 2.1-2.4, the following limit holds true almost surely*

$$\lim_{n \rightarrow \infty, p/n \rightarrow c} \left(g_{\hat{\mu}_n}(z) - \frac{1}{p} \text{Tr}[T_p(z)] \right) = 0, \quad \text{for all } z \in \mathbb{C} \setminus \mathbb{R}^+,$$

$$\lim_{n \rightarrow \infty, p/n \rightarrow c} \left(g_{\tilde{\mu}_n}(z) - \frac{1}{n} \text{Tr}[\tilde{T}_n(z)] \right) = 0, \quad \text{for all } z \in \mathbb{C} \setminus \mathbb{R}^+,$$

where

$$T_p(z) = \text{diag}(T_p^{(1)}(z), \dots, T_p^{(p)}(z)) \quad \text{and} \quad \tilde{T}_n(z) = \text{diag}(\tilde{T}_n^{(1)}(z), \dots, \tilde{T}_n^{(n)}(z)),$$

are diagonal matrices of size $p \times p$ and $n \times n$ respectively, whose diagonal elements are the unique solutions of the deterministic system of $p + n$ equations

$$T_p^{(j)}(z) = \frac{-1}{z(1 + (1/n) \text{Tr}[\tilde{D}_j \tilde{T}_n(z)])} \quad \text{for } 1 \leq j \leq p, \quad (2.1)$$

$$\tilde{T}_n^{(i)}(z) = \frac{-1}{z(1 + (1/n) \text{Tr}[D_i T_p(z)])} \quad \text{for } 1 \leq i \leq n, \quad (2.2)$$

where

$$\tilde{D}_j = \text{diag}(\gamma_{1j}^2, \dots, \gamma_{nj}^2) \quad \text{and} \quad D_i = \text{diag}(\gamma_{i1}^2, \dots, \gamma_{ip}^2).$$

Moreover, $\frac{1}{p} \text{Tr}[T_p(z)]$ and $\frac{1}{n} \text{Tr}[\tilde{T}_n(z)]$ are the Stieltjes transforms of probability measures denoted as ν_n and $\tilde{\nu}_n$ respectively.

The measures ν_n and $\tilde{\nu}_n$, defined in Proposition 2.1, are called the deterministic equivalents of the empirical eigenvalue distributions $\hat{\mu}_n$ and $\tilde{\mu}_n$ respectively. By a slight abuse of notation, we may also denote by $T_p(z)$ the vector of size p whose entries are the coefficients $T_p^{(j)}(z)$ for $1 \leq j \leq p$ solutions of the fixed point equations (2.1). Currently, a classical method to numerically approximate the value of $T_p(z)$ is to solve the nonlinear system of deterministic equations (2.1) written in a vector form that is referred to as the Dyson equation in [AEK17a, AEK19, AEK17b, AEK18]. Indeed, as stated in [AEK17b][Theorem 2.1], the vector $T_p(z)$ is known to be the unique solution of the Dyson equation

$$\frac{1}{T_p(z)} = -zI_p + \frac{1}{n} \Gamma_n^t \frac{1}{1 + \frac{1}{n} \Gamma_n T_p(z)}, \quad (2.3)$$

that corresponds to equations (2.1) and (2.2) written in a vector form, where $1/v$ has to be understood as taking the inverse of the elements of the vector v entrywise. Similarly, the vector $\tilde{T}_n(z) = (\tilde{T}_n^{(i)}(z))_{1 \leq i \leq n}$ satisfies the Dyson equation

$$\frac{1}{\tilde{T}_n(z)} = -zI_n + \frac{1}{n} \Gamma_n \frac{1}{1 + \frac{1}{n} \Gamma_n^t \tilde{T}_n(z)}.$$

In this paper, we sometimes consider, as an illustrative example, a specific class of variance profiles said to be *quasi doubly stochastic* in the following sense.

Definition 2.2. *The variance profile Γ_n is called quasi doubly stochastic when its rows and columns sum up to constant values (that are necessarily different when $n \neq p$), namely*

$$\frac{1}{n} \sum_{j=1}^p \gamma_{ij}^2 = \frac{p}{n}, \quad \text{for all } 1 \leq i \leq n, \quad \text{and} \quad \frac{1}{n} \sum_{i=1}^n \gamma_{ij}^2 = 1, \quad \text{for all } 1 \leq j \leq p.$$

The fixed point equation (2.3) typically does not have an explicit expression, and one has to rely on numerical methods to solve it as done in Section 4. Nevertheless, when the variance profile is quasi doubly stochastic, the solution of the Dyson equation is a vector $T_p(z)$ having constant entries equal to $m_p(z) \in \mathbb{C}$ (or equivalently $T_p(z) = m_p(z)I_p$ is a scalar matrix) that satisfies

$$\frac{1}{m_p(z)} = -z + \frac{1}{1 + \frac{p}{n}m_p(z)}, \quad \text{for all } z \in \mathbb{C} \setminus \mathbb{R}^+. \quad (2.4)$$

The above equality corresponds to the well-known fixed point equation satisfied by the Stieltjes transform $m_p(z)$ of the Marchenko-Pastur distribution [BS10]. One also has that $\tilde{T}_n(z) = \tilde{m}_n(z)I_n$ with $\tilde{m}_n(z)$ satisfying

$$\frac{1}{\tilde{m}_n(z)} = -z + \frac{p/n}{1 + \tilde{m}_n(z)}, \quad \text{for all } z \in \mathbb{C} \setminus \mathbb{R}^+.$$

3 Main results

In this section, we derive deterministic equivalents for the DOF and the predictive risk of ridge regression. We also obtain a deterministic equivalent of the predictive risk of minimum norm least square estimation when the ridge regularization parameter tends to zero. We compare these results to those that are already known in the standard setting of iid data, and we highlight the emergence of the double descent phenomenon for non-iid data. Throughout this section, it is supposed that Assumptions 2.1-2.4 hold true.

3.1 Degrees of freedom

In this section, we prove that the quantity $df_1(\lambda)$ introduced in Section 1.1 is a deterministic equivalent of the DOF $\widehat{df}_1(\lambda)$.

Proposition 3.1. *For any $\lambda > 0$, the following holds*

$$\lim_{n \rightarrow \infty, p/n \rightarrow c} |\widehat{df}_1(\lambda) - df_1(\lambda)| = 0, \quad \text{almost surely,} \quad (3.1)$$

with $df_1(\lambda) = \text{Tr}[\Sigma_n(\Sigma_n + \kappa(\lambda))^{-1}]$, and

$$\kappa(\lambda) = \text{diag}(\kappa_1(\lambda), \dots, \kappa_p(\lambda)), \quad \text{where } \kappa_j(\lambda) = \frac{\text{Tr}[\tilde{D}_j]}{\text{Tr}[\tilde{D}_j \tilde{T}_n(-\lambda)]}, \quad \text{for } 1 \leq j \leq p. \quad (3.2)$$

When the variance profile is quasi doubly stochastic, one has that $\text{Tr}[\tilde{D}_j] = n$. Moreover, as remarked in Section 2.3, the matrix $\tilde{T}_n(-\lambda) = \tilde{m}_n(-\lambda)I_n$ is scalar, implying that

$$\text{Tr}[\tilde{D}_j \tilde{T}_n(-\lambda)] = n\tilde{m}_n(-\lambda), \quad \text{for all } 1 \leq j \leq p.$$

Moreover, in this setting, $\Sigma_n = I_p$. Consequently, if Γ_n is quasi doubly stochastic, one has that

$$\kappa(\lambda) = \frac{1}{\tilde{m}_n(-\lambda)}I_p,$$

and the deterministic equivalent of the DOF is

$$df_1(\lambda) = \text{Tr} \left[\Sigma_n (\Sigma_n + \kappa(\lambda))^{-1} \right] = p \left(1 + \frac{1}{\tilde{m}_n(-\lambda)} \right)^{-1}.$$

The above equality corresponds to the formula of a deterministic equivalent of the DOF derived in [Bac24] for iid data when the entries of the features matrix X_n are made of iid centered random variables with variances equal to 1.

3.2 Predictive risk

We first express the predictive risk in a more convenient way. To this end, for any square matrix A , we denote by $\Delta[A]$ the diagonal matrix whose diagonal entries are those of A .

Lemma 3.1. *For any $\lambda > 0$, the predictive risk $\hat{r}_\lambda(X_n) = \mathbb{E}[(\tilde{Y} - \tilde{X}^t \hat{\theta}_\lambda)^2 | X_n]$ has the following expression*

$$\hat{r}_\lambda(X_n) = \sigma^2 + \frac{\sigma^2}{n} \text{Tr}[\tilde{S}_p \Delta[Q_p(-\lambda)]] + \lambda \left(\frac{\lambda \alpha^2}{p} - \frac{\sigma^2}{n} \right) \text{Tr}[\tilde{S}_p \Delta[Q'_p(-\lambda)]]. \quad (3.3)$$

Moreover, we can exhibit the following Bias-Variance decomposition for $\hat{r}_\lambda(X_n)$:

$$\hat{r}_\lambda(X_n) = \sigma^2 + \text{Bias}(\hat{\theta}_\lambda) + \text{Var}(\hat{\theta}_\lambda),$$

with

$$\begin{aligned} \text{Bias}(\hat{\theta}_\lambda) &= \mathbb{E}[(\mathbb{E}[\hat{\theta}_\lambda | X_n, \theta_*] - \theta_*)^t \tilde{S}_p (\mathbb{E}[\hat{\theta}_\lambda | X_n, \theta_*] - \theta_*) | X_n] \\ &= \frac{\lambda^2 \alpha^2}{p} \text{Tr}[\tilde{S}_p \Delta[Q'_p(-\lambda)]], \end{aligned}$$

and

$$\begin{aligned} \text{Var}(\hat{\theta}_\lambda) &= \mathbb{E}[(\hat{\theta}_\lambda - \mathbb{E}[\hat{\theta}_\lambda | X_n, \theta_*])^t \tilde{S}_p (\hat{\theta}_\lambda - \mathbb{E}[\hat{\theta}_\lambda | X_n, \theta_*]) | X_n] \\ &= \frac{\sigma^2}{n} \text{Tr}[\tilde{S}_p \Delta[Q_p(-\lambda)] - \lambda \tilde{S}_p \Delta[Q'_p(-\lambda)]], \end{aligned}$$

where

$$Q_p(z) = (\hat{\Sigma}_n - zI_p)^{-1} \quad \text{for } z \in \mathbb{C} \setminus \mathbb{R}^+,$$

is the resolvent of the matrix $\hat{\Sigma}_n$, and Q'_p denotes the derivative of Q_p with respect to z .

We can now give a deterministic equivalent of the predictive risk $\hat{r}_\lambda(X_n)$ that is obtained in a simple way by replacing the diagonal matrix $\Delta[Q_p(z)]$ with the deterministic matrix $T_p(z)$ in the expression (3.3) of the predictive risk. All proofs of the following results are given in Appendix A.3.

Theorem 3.1. *A deterministic equivalent of the predictive risk is*

$$r_\lambda(\Upsilon_n) = \sigma^2 + \frac{\sigma^2}{n} \text{Tr}[\tilde{S}_p T_p(-\lambda)] + \lambda \left(\frac{\lambda \alpha^2}{p} - \frac{\sigma^2}{n} \right) \text{Tr}[\tilde{S}_p T'_p(-\lambda)], \quad (3.4)$$

in the sense that it satisfies

$$\lim_{n \rightarrow \infty, p/n \rightarrow c} |\hat{r}_\lambda(X_n) - r_\lambda(\Upsilon_n)| = 0, \quad \text{almost surely,}$$

where $T'_p(z)$ denotes the derivative of $T_p(z)$ with respect to z .

A key argument to derive the computation of $r_\lambda(\Upsilon_n)$ is to prove that the matrix $T_p(z)$ is a relevant deterministic equivalent of the diagonal matrix $\Delta[Q_p(z)]$ for an appropriate notion of asymptotic equivalence between matrices of growing size. This is the purpose of Appendix A.2 where we derive a stronger convergence result than the one stated in Proposition 2.1 which only shows that $\frac{1}{p} \text{Tr}[T_p(z)]$ is a deterministic equivalent of $g_{\hat{\mu}_n}(z) = \frac{1}{p} \text{Tr}[Q_p(z)]$.

Lemma 3.2. *If $\frac{p}{n} < 1$ then $T_p(-\lambda)$ and its derivative, $T'_p(-\lambda)$ admit a limit when λ tends to 0. Indeed, there exists a constant $\tau > 0$ such that*

$$\lim_{\lambda \rightarrow 0, \lambda > 0} T_p(-\lambda) = \int_{\tau}^{+\infty} \frac{\mu(dw)}{w} \quad \text{and} \quad \lim_{\lambda \rightarrow 0, \lambda > 0} T'_p(-\lambda) = \int_{\tau}^{+\infty} \frac{\mu(dw)}{w^2},$$

where $\mu = (\mu_{ij})$ is a positive $p \times p$ matrix valued-measure such that μ_{ii} is a probability measure and μ_{ij} is a null measure if $i \neq j$. We denote these limits by $T_p(0^-) = \lim_{\lambda \rightarrow 0} T_p(-\lambda)$ and $T'_p(0^-) = \lim_{\lambda \rightarrow 0} T'_p(-\lambda)$.

On the other hand, if $\frac{p}{n} > 1$ then $\kappa(\lambda)$ and its derivative $\kappa'(\lambda)$ also admit a limit when λ tends to 0 and we denote these limits by

$$\kappa(0^+) = \lim_{\lambda \rightarrow 0, \lambda > 0} \kappa(\lambda) \quad \text{and} \quad \kappa'(0^+) = \lim_{\lambda \rightarrow 0, \lambda > 0} \kappa'(\lambda),$$

where $\kappa(\lambda)$ is the diagonal matrix defined in Equation (3.2).

Corollary 3.1. *Suppose that the variance profile Γ_n is quasi doubly stochastic. Then, if $\frac{p}{n} < 1$, one has that*

$$T_p(0^-) = m_p(0)I_p,$$

and, if $\frac{p}{n} > 1$,

$$\kappa(0^+) = \frac{1}{\tilde{m}_n(0)}I_p \quad \text{and} \quad \kappa'(0^+) = \frac{\tilde{m}'_n(0)}{\tilde{m}_n^2(0)}I_p,$$

where $m_p(\cdot)$, resp. $\tilde{m}_n(\cdot)$, is the Stieltjes transform of the Marchenko-Pastur distribution with parameter p/n , resp. n/p .

Then, Theorem 3.1 allows us to understand the behavior of $r_\lambda(\Upsilon_n)$ when $\lambda \rightarrow 0$ and $\lambda \rightarrow +\infty$ through the following corollary.

Corollary 3.2. *As $\lambda \rightarrow +\infty$, the limit of the deterministic equivalent $r_\lambda(\Upsilon_n)$ of the predictive risk is*

$$\lim_{\lambda \rightarrow +\infty} r_\lambda(\Upsilon_n) = \frac{\alpha^2}{p} \text{Tr}[\tilde{S}_p] + \sigma^2.$$

Moreover, as $\lambda \rightarrow 0$, the limit of the deterministic equivalent $r_\lambda(\Upsilon_n)$ is as follows

- if $\frac{p}{n} < 1$ then

$$\lim_{\lambda \rightarrow 0, \lambda > 0} r_\lambda(\Upsilon_n) = \frac{\sigma^2}{n} \text{Tr}[\tilde{S}_p T_p(0^-)] + \sigma^2,$$

- if $\frac{p}{n} > 1$ then

$$\lim_{\lambda \rightarrow 0, \lambda > 0} r_\lambda(\Upsilon_n) = \frac{\alpha^2}{p} \text{Tr}[\tilde{S}_p \kappa(0^+) (\Sigma_n + \kappa(0^+))^{-1}] + \frac{\sigma^2}{n} \text{Tr}[\kappa'(0^+) \tilde{S}_p \Sigma_n (\Sigma_n + \kappa(0^+))^{-2}] + \sigma^2,$$

where $\kappa'(\lambda)$ denotes the derivative with respect to λ of the diagonal matrix $\kappa(\lambda)$ defined in Equation (3.2).

In the case $\lambda \rightarrow 0$, Corollary 3.2 yields a deterministic equivalent of the predictive risk of the minimum norm least-square estimator $\hat{\theta}$. Then, it can be seen that the risk $r_0(\Upsilon_n)$ exhibit two different behaviors depending on the value of the ratio p/n with respect to one.

If $p/n < 1$, then $\hat{\theta}$ is equivalent to the ordinary least-square estimator $(X_n^t X_n)^{-1} X_n^t Y_n$ which is known to be unbiased, and the risk $r_0(\Upsilon_n)$ is thus only composed of a variance term. Moreover, if the variance profile is assumed to be quasi-bistochastic, then, given that

$$T_p(0^-) = m_p(0)I_p \quad \text{and} \quad m_p(0) = \frac{1}{1 - p/n},$$

it follows that

$$r_0(\Upsilon_n) = \frac{\sigma^2}{1 - p/n} \times \frac{1}{n} \text{Tr}[\tilde{S}_p] + \sigma^2 = \sigma^2 \frac{p/n}{1 - p/n} + \sigma^2 \rightarrow \sigma^2 \frac{c}{1 - c} + \sigma^2, \text{ as } n \rightarrow \infty, p/n \rightarrow c,$$

which corresponds, when $c < 1$, to the known asymptotic limit of the predictive risk of least squares estimation for iid data when the entries of the features matrix X_n are iid centered random variables with variances equal to 1, see e.g. [HMRT22][Proposition 2].

If $p/n > 1$, then the deterministic equivalent of the predictive risk is composed of a bias term and a variance term. If the variance profile is assumed to be quasi-bistochastic, the values of these two terms can be made more explicitly as follows. In this setting, given that $\Sigma_n = I_p$, $\tilde{m}_p(0) = \frac{1}{p/n-1}$ and $\tilde{m}'_p(0) = \frac{p/n}{(p/n-1)^3}$, one has that

$$\kappa(0^+) = (p/n - 1)I_p \quad \text{and} \quad \kappa'(0^+) = \frac{p/n}{p/n - 1}I_p.$$

Hence, using that $\frac{1}{n} \text{Tr}[\tilde{S}_p] = p/n$, we finally obtain that

$$r_0(\Upsilon_n) = \alpha^2 \left(1 - \frac{n}{p}\right) + \sigma^2 \frac{1}{p/n - 1} \rightarrow \alpha^2 \left(1 - \frac{1}{c}\right) + \sigma^2 \frac{1}{c - 1}, \text{ as } n \rightarrow \infty, p/n \rightarrow c,$$

which corresponds, when $c > 1$, to known results on the bias-variance decomposition of the asymptotic limit of the predictive risk of the minimum norm least squares estimator for iid data when the entries of X_n are iid centered random variables with variance 1, see e.g. [HMRT22][Theorem 1].

Beyond the assumption of a quasi-stochastic variance profile, it is difficult to analytically determine the shape of the predictive risk $r_0(\Upsilon_n)$ as a function of the ratio p/n . Indeed, this requires to at least know upper and lower bounds on the magnitude of the diagonal elements of $T_p(0^-)$ and $\kappa(0^+)$ (and its derivative). As shown by Lemma 3.2 when $p/n < 1$, this issue amounts to finding upper and lower bounds of the support of the matrix-valued measure μ satisfying

$$T_p(-0) = \int_{\tau}^{+\infty} \frac{\mu(dw)}{w}.$$

Hence, upper and lower bounding $T_p(-0)$ is related to understanding the value of the constant τ and the size of the support of the limiting spectral distribution of the covariance matrix Σ_n which remains (to the best of our knowledge) an open problem for random matrices with an arbitrary variance profile.

Nevertheless, in Section 4, we use Corollary 3.2 and computational methods to evaluate $T_p(0^-)$ and $\kappa(0^+)$ to report numerical experiments illustrating that the double descent phenomenon for the predictive risk of $\hat{\theta}$ also holds in the high-dimensional model (1.1) with more general variance profile than a quasi-bistochastic one. In Section 4, we also exhibit variance profiles for which the predictive risk has a shape that differs from double descent.

Corollary 3.3. *The function $\lambda \mapsto \hat{r}_\lambda(X_n)$ reaches its minimum at $\lambda_* = \frac{\sigma^2 p}{\alpha^2 n}$ meaning that*

$$\hat{r}_{\lambda_*}(X_n) \leq \hat{r}_\lambda(X_n), \tag{3.5}$$

for any $\lambda > 0$.

Note that although $\hat{r}_\lambda(X_n)$ depends on the variance profile, this is not the case of the optimal value of λ_* that minimizes the predictive risk as shown by Corollary 3.3. Moreover, the optimal value λ_* is also the one minimizing $\lambda \mapsto \hat{r}_\lambda(X_n)$ in the framework of [DW18] with a matrix X_n of features made of iid rows.

4 Numerical experiments

In this section, we illustrate the results of this paper with numerical experiments. Although we do not have an explicit formula for $T_p(z)$, this matrix-valued function satisfies the fixed-point equation (2.3). This allows us to approximate $T_p(z)$ numerically using a fixed-point algorithm. In this manner, we are thus able to obtain a numerical approximation of the deterministic equivalent $r_\lambda(\Upsilon_n)$ of the predictive risk. In this section, we consider several variance profiles that we normalize such that

$$\frac{1}{n} \sum_{i=1}^n \sum_{j=1}^p \frac{\gamma_{ij}^2}{p} = 1.$$

Apart from the constant, quasi doubly stochastic, and piecewise constant variance profiles mentioned earlier, we will use the following examples of variance profiles:

- The alternated columns variance profile satisfying $\gamma_{ij} = \gamma_1$ if j is even and $\gamma_{ij} = \gamma_2$ if j is odd.
- The polynomial variance profile satisfying $\gamma_{ij}^2 = \left| \frac{i-j}{\min(n,p)} \right|^6 + \tau$ for some $\tau > 0$.
- The block variance profile

$$\Gamma_n = \begin{bmatrix} \gamma_1^2 \mathbb{1}_{n/4} \mathbb{1}_{p/4}^t & \gamma_2^2 \mathbb{1}_{n/4} \mathbb{1}_{p/3}^t & \mathbb{1}_{n/4} \mathbb{1}_{5p/12}^t \\ \gamma_2^2 \mathbb{1}_{n/3} \mathbb{1}_{p/4}^t & \gamma_1^2 \mathbb{1}_{n/3} \mathbb{1}_{p/3}^t & \gamma_3^2 \mathbb{1}_{n/3} \mathbb{1}_{5p/12}^t \\ \mathbb{1}_{5n/12} \mathbb{1}_{p/4}^t & \gamma_3^2 \mathbb{1}_{5n/12} \mathbb{1}_{p/3}^t & \gamma_1^2 \mathbb{1}_{5n/12} \mathbb{1}_{5p/12}^t \end{bmatrix},$$

for some sufficiently different constants $\gamma_1, \gamma_2, \gamma_3$.

- A variance profile referred to as the Berlin Photo for which γ_{ij} is equal to the value of the coordinate (i, j) of the pixel of the image depicted in Figure 2. The original image being of size 1250×850 , it is rescaled according to the values of n and p used in the numerical experiments thanks to the function `resize` from the Python module `PIL.Image`.

The values of the predictive risk and its deterministic equivalent are compared for various variance profiles in Figure 3 with λ ranging from 0.1 to 10, $n = 400$ and $p = 600$. The curves displayed in Figure 3 confirm that $r_\lambda(\Upsilon_n)$ is a very accurate estimator of $\hat{r}_\lambda(X_n)$ in high-dimension since the dashed and solid lines coincide. These variance profiles also provide curves of predictive risks having similar shapes (up to a vertical translation). The curves representing the cases of constant and quasi doubly stochastic variance profile coincide, confirming the comment made in Section 1.1 on the doubly stochastic profile and its similarities with the setting of iid data. Moreover, the minimum of $r_\lambda(\Upsilon_n)$ is indeed reached at the optimal value λ_* for any variance profile as shown by Corollary 3.3.

We illustrate the appearance of a double or triple descent phenomenon for several variance profiles in Figure 4 (a), (c) and (e). These figures represent the predictive risk and its approximation by a deterministic equivalent for various values of the ratio p/n with $\lambda = 0$. The solid lines depicts $r_0(\Upsilon_n)$ whereas the dashed lines represent $\hat{r}_0(X_n)$. For every variance profile, the solid and dashed lines coincide which confirms that $r_0(\Upsilon_n)$ is a relevant deterministic equivalent of $\hat{r}_0(X_n)$. For the constant and quasi doubly stochastic variance profiles, we observe the well

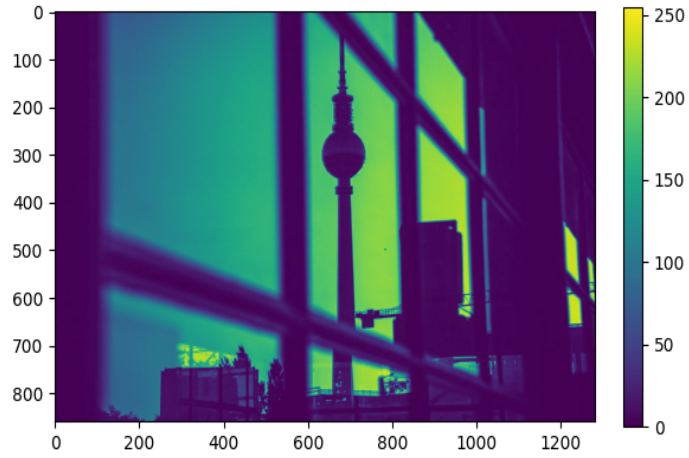


Figure 2: The Berlin Photo variance profile whose entries correspond to the the green channel of the pixels of a RGB picture taken in Berlin.

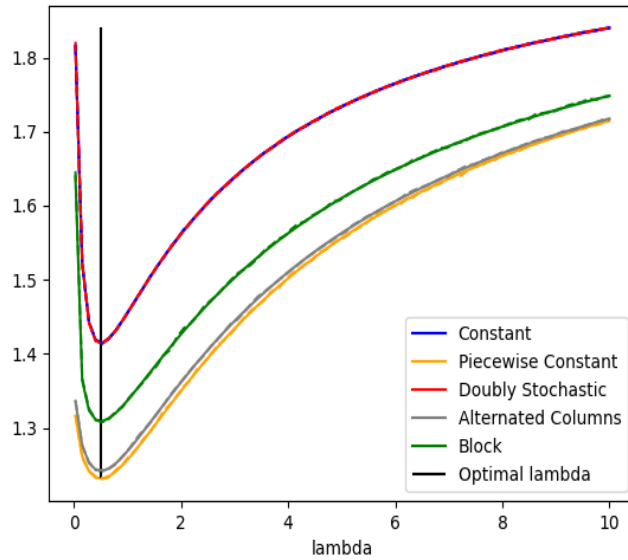


Figure 3: Comparison of $\hat{r}_\lambda(X_n)$ and $r_\lambda(\Upsilon_n)$ for several variance profiles , $\alpha = 1$, $\sigma = 1$, $n = 400$ and $p = 600$. The dashed lines correspond to the predictive risk while the solid lines correspond to the deterministic equivalent.

known double descent phenomenon as illustrated by Figure 4 (a) since the curves are increasing for $p/n < 1$ and decreasing for $p/n > 1$. This is related to Corollary 3.2 which states that the expression of $\lim_{\lambda \rightarrow 0} r_\lambda(\Upsilon_n)$ depends upon the value of the ratio p/n with respect to one. Nevertheless, for other variance profiles, the shape of the predictive risk can be very different from one variance profile to another, and it differs from the usual double descent. For some variance profiles a phenomenon of triple descent arises, notably for the piecewise constant and block cases, as shown in Figure 4 (c) and (e). We also remark the appearance of a quadruple descent in the case of the polynomial profile, as shown in Figure 4 (e).

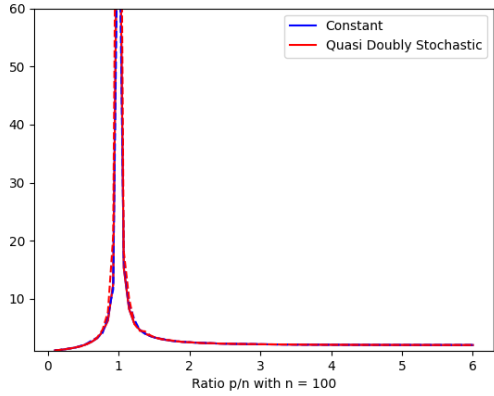
As already remarked at the end of Section 3, it is difficult to analytically explain these phenomenon for an arbitrary variance profile since we do not have an explicit formula for $T_p(z)$ in the general case. Nevertheless, as nicely discussed in [SKR⁺23], the double descent phenomenon is very much related to the distance from zero of the smallest non-negative eigenvalue of $\widehat{\Sigma}_n$. Having this eigenvalue close to zero causes double descent. In Figure 4 (b), (d) and (f), we thus represent the smallest non-zero eigenvalue $\tau_{\min}(\widehat{\Sigma}_n)$ of $\widehat{\Sigma}_n$ for many variance profiles as a function of p/n . It can be observed that all the curves $p/n \rightarrow \tau_{\min}(\widehat{\Sigma}_n)$ reach their minima at a value close to zero when $p = n$ where the double descent occurs, which is consistent with previous assertions in [SKR⁺23]. These curves in the constant and quasi doubly stochastic cases behave in a similar same way. However, one can remark from Figure 4 (f) that, for the piecewise constant and polynomial variance profiles, the value of $\tau_{\min}(\widehat{\Sigma}_n)$ stay much closer to 0 when p/n is between 0.25 and 3 than in the constant or quasi doubly stochastic cases. The fact that these curves have a plateau near zero for $p/n \in [0.25, 3]$ may be related to the appearance of the triple and quadruple descent phenomenon that is observed in Figure 4 (e).

5 Conclusion

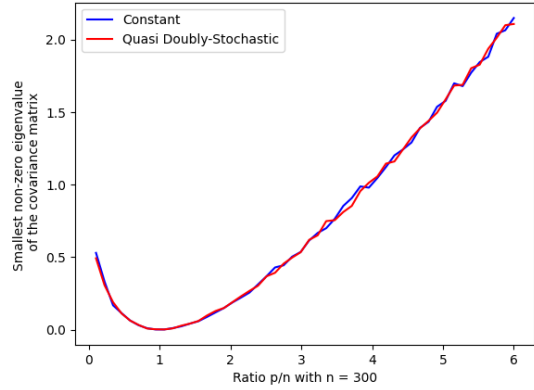
In this paper, we have derived a deterministic equivalent of the DOF and the predictive risk of ridge (less) regression in a high-dimensional framework with a variance profile to handle the setting of non-iid data. In order to achieve this result, we had to use ingredients from RMT to determine a deterministic equivalent of the diagonal of $Q_p(z)$. We have shown that results on the values of the DOF and the predictive that are already known for a constant variance profile, notably those of [DW18] and [Bac24], also hold in the case of a quasi doubly stochastic profile. However, there are variance profiles that cause a behavior of the predictive risk of the minimum norm least square estimator that differs from the double descent phenomenon as classically observed in the case of a constant variance profile. The numerical experiments that we have conducted confirm that our deterministic equivalent accurately estimates the predictive risk in high-dimension. Our results also allow to understand how assuming such a variance profile for the data influences the statistical properties of ridge regression when compared to the standard assumption of iid observations. We hope that our approach on the use of variance profiles may lead to further research works on the statistical analysis of other estimators than ridge regression in more complex models with non-iid data.

Appendix

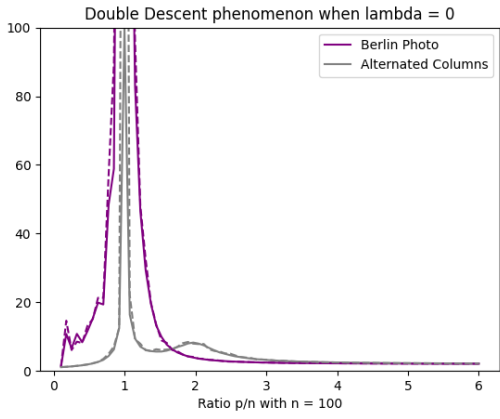
In this appendix, we give all the proofs of the results of the paper. We also discuss the link between random matrices with a variance profile and the notion of \mathcal{R} -transform in free probability.



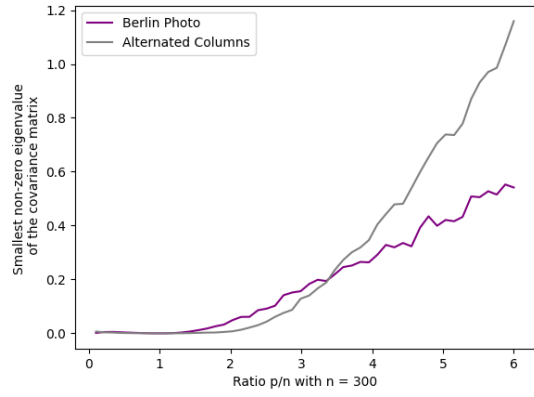
(a)



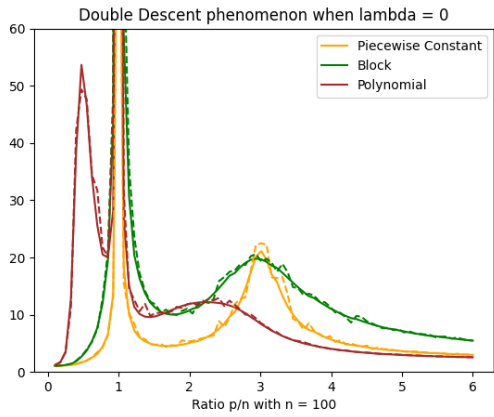
(b)



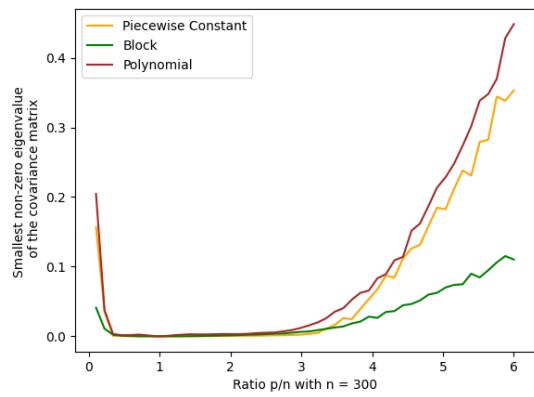
(c)



(d)



(e)



(f)

Figure 4: Left : Double descent phenomenon for several variance profiles with $\alpha = \sigma = 1$, $n = 100$ and p varying from 10 to 600. Right : Smallest non-zero eigenvalue of $\hat{\Sigma}_n$ for several variance profiles with $\alpha = \sigma = 1$, $n = 300$ and p varying from 30 to 1800.

A.1 Proofs of Proposition 3.1 and Lemma 3.1

Proof of Proposition 3.1. Recall that $\frac{1}{p}\widehat{df}_1(\lambda) = \frac{1}{p}\text{Tr}[\widehat{\Sigma}_n(\widehat{\Sigma}_n + \lambda I_p)^{-1}]$, for $\lambda > 0$. Thus by using $\widehat{\mu}_n$ the empirical eigenvalue distribution of $\widehat{\Sigma}_n$ we have this new expression of the degree of freedom

$$\frac{1}{p}\widehat{df}_1(\lambda) = \int_{\mathbb{R}} \frac{t}{t + \lambda} d\widehat{\mu}_n(t).$$

Moreover, $\widehat{\mu}_n$ being a probability measure one can notice from Definition 2.1 that

$$\frac{1}{p}\widehat{df}_1(\lambda) = \int_{\mathbb{R}} 1 - \frac{\lambda}{t + \lambda} d\widehat{\mu}_n(t) = 1 - \lambda g_{\widehat{\mu}_n}(-\lambda).$$

Thanks to Proposition 2.1, one has that

$$\frac{1}{p}\widehat{df}_1(\lambda) \sim 1 - \lambda g_{\nu_n}(-\lambda). \quad (\text{A.1})$$

Let us now remark that, by Proposition 2.1, one has that

$$g_{\nu_n}(z) = \frac{1}{p}\text{Tr}[T_p(z)] = -\frac{1}{p}\sum_{j=1}^p \frac{1}{z \left(1 + (1/n)\text{Tr}[\widetilde{D}_j \widetilde{T}_n(z)]\right)},$$

implying that

$$\begin{aligned} 1 - \lambda g_{\nu_n}(-\lambda) &= 1 - \frac{1}{p}\sum_{j=1}^p \frac{1}{\left(1 + (1/n)\text{Tr}[\widetilde{D}_j \widetilde{T}_n(-\lambda)]\right)} \\ &= \frac{1}{p}\sum_{j=1}^p \frac{(1/n)\text{Tr}[\widetilde{D}_j \widetilde{T}_n(-\lambda)]}{\left(1 + (1/n)\text{Tr}[\widetilde{D}_j \widetilde{T}_n(-\lambda)]\right)}. \end{aligned}$$

Then, recall that $\Sigma_n = \frac{1}{n}\text{diag}(\text{Tr}[\widetilde{D}_1], \dots, \text{Tr}[\widetilde{D}_p])$, which implies that

$$1 - \lambda g_{\nu_n}(-\lambda) = \frac{1}{p}\sum_{j=1}^p \frac{(1/n)\text{Tr}[\widetilde{D}_j]}{(1/n)\text{Tr}[\widetilde{D}_j] + \kappa_j(\lambda)} = \frac{1}{p}\text{Tr}[\Sigma_n(\Sigma_n + \kappa(\lambda))^{-1}], \quad (\text{A.2})$$

where

$$\kappa_j(\lambda) = \frac{\text{Tr}[\widetilde{D}_j]}{\text{Tr}[\widetilde{D}_j \widetilde{T}_n(-\lambda)]}, \quad \text{for } 1 \leq j \leq p,$$

and

$$\kappa(\lambda) = \text{diag}(\kappa_1(\lambda), \dots, \kappa_p(\lambda)).$$

Consequently, by combining Equations (A.1) and (A.2), we obtain the deterministic equivalent for the DOF stated in Equation (3.1), which completes the proof of Proposition 3.1. \square

Proof of Lemma 3.1. Let us consider the following decomposition of the predictive risk. Since \widetilde{X} and $\widetilde{\varepsilon}$ are independent from X_n , we can pursue the following computations to get a new expression for $\widehat{r}_\lambda(X_n)$:

$$\begin{aligned} \widehat{r}_\lambda(X_n) &= \mathbb{E}[(\widetilde{X}^t \widehat{\theta}_\lambda - \widetilde{X}^t \theta_* - \widetilde{\varepsilon})^2 | X_n] \\ &= \mathbb{E}[\text{Tr}[\widehat{\theta}_\lambda \widehat{\theta}_\lambda^t \widetilde{X} \widetilde{X}^t + \theta_* \theta_*^t \widetilde{X} \widetilde{X}^t + \widetilde{\varepsilon}^2 - 2\theta_* \widehat{\theta}_\lambda^t \widetilde{X} \widetilde{X}^t - 2\widetilde{\varepsilon} \widehat{\theta}_\lambda^t \widetilde{X} + 2\widetilde{\varepsilon} \theta_*^t \widetilde{X}] | X_n] \\ &= \sigma^2 + \text{Tr}[\mathbb{E}[(\widehat{\theta}_\lambda \widehat{\theta}_\lambda^t + \theta_* \theta_*^t - 2\theta_* \widehat{\theta}_\lambda^t) | X_n] \widetilde{S}_p] \\ &= \sigma^2 + \mathbb{E}[\text{Tr}[(\widehat{\theta}_\lambda - \theta_*)(\widehat{\theta}_\lambda - \theta_*)^t \widetilde{S}_p] | X_n] \\ &= \sigma^2 + \mathbb{E}[(\widehat{\theta}_\lambda - \theta_*)^t \widetilde{S}_p (\widehat{\theta}_\lambda - \theta_*) | X_n] = \sigma^2 + \widehat{R}(\widehat{\theta}_\lambda, \theta_*), \end{aligned}$$

with $\hat{R}(\hat{\theta}_\lambda, \theta_*) = \mathbb{E}[(\hat{\theta}_\lambda - \theta_*)^t \tilde{S}_p(\hat{\theta}_\lambda - \theta_*) | X_n]$. Then by using (1.1) and (1.2), we get

$$\hat{\theta}_\lambda - \theta_* = -\lambda(\hat{\Sigma}_n + \lambda I_p)^{-1} \theta_* + (\hat{\Sigma}_n + \lambda I_p)^{-1} \frac{X_n^t \varepsilon_n}{n}.$$

Therefore, we can rewrite

$$\begin{aligned} \hat{R}(\hat{\theta}_\lambda, \theta_*) &= \mathbb{E}[\lambda^2 \theta_*^t Q_p(-\lambda) \tilde{S}_p Q_p(-\lambda) \theta_* + \frac{\varepsilon_n^t X_n}{n} Q_p(-\lambda) \tilde{S}_p Q_p(-\lambda) \frac{X_n^t \varepsilon_n}{n} \\ &\quad - 2\lambda \theta_*^t Q_p(-\lambda) \tilde{S}_p Q_p(-\lambda) \frac{X_n^t \varepsilon_n}{n} | X_n] \\ &= \lambda^2 \text{Tr}[Q_p(-\lambda) \tilde{S}_p Q_p(-\lambda) \mathbb{E}[\theta_* \theta_*^t]] \\ &\quad + \frac{1}{n^2} \text{Tr}[\tilde{S}_p Q_p(-\lambda) X_n^t \mathbb{E}[\varepsilon_n \varepsilon_n^t] X_n Q_p(-\lambda)] \\ &= \frac{\lambda^2 \alpha^2}{p} \text{Tr}[Q_p(-\lambda) \tilde{S}_p Q_p(-\lambda)] + \frac{\sigma^2}{n} \text{Tr}[\tilde{S}_p Q_p(-\lambda) \hat{\Sigma}_n Q_p(-\lambda)] \\ &= \frac{\lambda^2 \alpha^2}{p} \text{Tr}[Q_p(-\lambda) \tilde{S}_p Q_p(-\lambda)] + \frac{\sigma^2}{n} \text{Tr}[\tilde{S}_p \hat{\Sigma}_n Q_p'(-\lambda)]. \end{aligned}$$

Note that $\hat{\Sigma}_n$ and $Q_p(-\lambda)$ are commuting, one has that $Q_p^2(-\lambda) = Q_p'(-\lambda)$, and moreover, $\hat{\Sigma}_n Q_p'(-\lambda) = Q_p(-\lambda) - \lambda Q_p'(\lambda)$. Then we get the following Bias-Variance decomposition of the predictive risk

$$\hat{R}(\hat{\theta}_\lambda, \theta_*) = \text{Bias}(\hat{\theta}_\lambda) + \text{Var}(\hat{\theta}_\lambda),$$

where

$$\text{Bias}(\hat{\theta}_\lambda) = \frac{\lambda^2 \alpha^2}{p} \text{Tr}[Q_p(-\lambda) \tilde{S}_p Q_p(-\lambda)] = \frac{\lambda^2 \alpha^2}{p} \text{Tr}[\tilde{S}_p Q_p'(-\lambda)]$$

and

$$\text{Var}(\hat{\theta}_\lambda) = \frac{\sigma^2}{n} \text{Tr}[\tilde{S}_p \hat{\Sigma}_n Q_p'(-\lambda)] = \frac{\sigma^2}{n} \text{Tr}[\tilde{S}_p (Q_p(-\lambda) - \lambda Q_p'(-\lambda))].$$

From these computations, we finally deduce the following formula of the predictive risk

$$\hat{r}_\lambda(X_n) = \sigma^2 + \frac{\sigma^2}{n} \text{Tr}[\tilde{S}_p Q_p(-\lambda)] + \lambda \left(\frac{\lambda \alpha^2}{p} - \frac{\sigma^2}{n} \right) \text{Tr}[\tilde{S}_p Q_p'(-\lambda)],$$

Then, Equation (3.3) directly follows from these computations since \tilde{S}_p is a diagonal matrix, which completes the proof of Lemma 3.1. \square

A.2 Deterministic equivalents of operator-valued Stieltjes transforms

One can remark that the deterministic equivalent given in Equation (3.4) is obtained by replacing $\Delta[Q_p(-\lambda)]$ with $T_p(-\lambda)$ in Equation (3.3). As argued in Section 3.2, even though Proposition 2.1 states that $\frac{1}{n} \text{Tr}[Q_p(-\lambda)] \sim \frac{1}{n} \text{Tr}[T_p(-\lambda)]$, this equivalence is not sufficient to prove Theorem 3.1. We need to find a deterministic matrix that is a good approximation of the diagonal of the resolvent $\Delta[Q_p(z)]$ in a high-dimensional sense. To this end, we first define the following equivalence relation in order to specify an appropriate notion of deterministic equivalent for the matrix $\Delta[Q_p(-\lambda)]$.

Definition A.1. Let $\mathbf{A} = (A_p)_{p \geq 1}$ and $\mathbf{B} = (B_p)_{p \geq 1}$ be a family of square complex random matrices such that for all $p \geq 1$, $A_p, B_p \in \mathbb{C}^{p \times p}$. Then, the two families of matrices \mathbf{A} and \mathbf{B} are said to be equivalent, denoted by $\mathbf{A} \sim \mathbf{B}$, if

$$\lim_{n \rightarrow \infty, p/n \rightarrow c} \left| \frac{1}{p} \text{Tr}[A_p U_p] - \frac{1}{p} \text{Tr}[B_p U_p] \right| = 0, \text{ almost surely,}$$

for all family of deterministic matrices $(U_p)_{p \geq 1}$ satisfying :

$$\forall p \geq 1, \quad U_p \in \mathbb{R}^{p \times p} \quad \text{and} \quad \exists K > 0, \quad \sup_{p \geq 1} \max_{i \in [p]} |U_p^{(i)}| \leq K, \quad (\text{A.3})$$

where $U_p^{(i)}$ denotes the i -th diagonal entry of U_p .

Then, using Definition A.1, we prove that $T_p(z)$ is a relevant deterministic equivalent of $Q_p(z)$ in the following sense.

Theorem A.1. *Let us define the following families of matrices $\Delta[\mathbf{Q}(z)] = (\Delta[Q_p(z)])_{p \geq 1}$, $\Delta[\mathbf{Q}'(z)] = (\Delta[Q'_p(z)])_{p \geq 1}$, $\mathbf{T}(z) = (T_p(z))_{p \geq 1}$ and $\mathbf{T}'(z) = (T'_p(z))_{p \geq 1}$. Then, the following equivalences hold:*

$$\Delta[\mathbf{Q}(z)] \sim \mathbf{T}(z) \quad \text{and} \quad \Delta[\mathbf{Q}'(z)] \sim \mathbf{T}'(z), \quad \text{for all } z \in \mathbb{C} \setminus \mathbb{R}^+.$$

Theorem A.1 is inspired by results from [HLN07], and it constitutes a key element in the proof of Theorem 3.1. In order to prove Theorem A.1, we introduce the diagonal matrices $R_p(z) = \text{diag}(R_p^{(j)}(z))_{1 \leq j \leq p}$ defined by

$$R_p^{(j)}(z) = \frac{-1}{z \left((1 + (1/n) \text{Tr}[\tilde{D}_j \tilde{Q}_n(z)]) \right)} \quad \text{for } z \in \mathbb{C} \setminus \mathbb{R}^+,$$

with $\tilde{Q}_n(z) = (\tilde{\Sigma}_n - zI_n)^{-1}$ is the resolvent matrix of $\tilde{\Sigma}_n = \frac{1}{n} X_n X_n^t$.

Lemma A.1. *Let us define $\Delta[\mathbf{Q}(z)]$ as in Theorem A.1 and denote by $\mathbf{R}(z)$ the family of matrices $(R_p(z))_{p \geq 1}$. Then, the following equivalences hold:*

$$\Delta[\mathbf{Q}(z)] \sim \mathbf{R}(z),$$

for all $z \in \mathbb{C}^+ = \{z \in \mathbb{C}, \Im(z) > 0\}$.

Proof of Lemma A.1. This proof is based on [HLN07][Lemma 6.1] which asserts that there exists a constant $K_1 > 0$ such that $\mathbb{E} \left[\left| \frac{1}{p} \text{Tr}[(Q_p(z) - R_p(z))U_p] \right|^{2+\delta/2} \right] \leq \frac{K_1}{p^{1+\delta/4}}$, where $(U_p)_{p \geq 1}$ denotes a family of deterministic matrices satisfying (A.3). We deduce from this inequality that

$$\mathbb{E} \left[\sum_{p \geq 1} \left| \frac{1}{p} \text{Tr}[(Q_p(z) - R_p(z))U_p] \right|^{2+\delta/2} \right] \leq \sum_{p \geq 1} \frac{K_1}{p^{1+\delta/4}} < +\infty.$$

The serie being finite, we obtain that $\mathbb{P} \left[\sum_{p \geq 1} \left| \frac{1}{p} \text{Tr}[(Q_p(z) - R_p(z))U_p] \right|^{2+\varepsilon/2} < +\infty \right] = 1$, which finally gives us :

$$\mathbb{P} \left[\frac{1}{p} \text{Tr}[(Q_p(z) - R_p(z))U_p] \xrightarrow{n \rightarrow +\infty} 0 \right] \geq \mathbb{P} \left[\sum_{p \geq 1} \left| \frac{1}{p} \text{Tr}[(Q_p(z) - R_p(z))U_p] \right|^{2+\varepsilon/2} < +\infty \right] = 1.$$

This concludes the proof of Lemma A.1. □

Then, Lemma A.1 needs to be combined with the following one in order to prove Theorem A.1.

Lemma A.2. *Let us define $\mathbf{T}(z)$ and $\mathbf{R}(z)$ as in Theorem A.1 and Lemma A.1. Then, the following equivalences hold:*

$$\mathbf{R}(z) \sim \mathbf{T}(z),$$

for all $z \in \mathcal{D} = \{z \in \mathbb{C}^+, \frac{|z|}{|\Im(z)|} < 2, |\Im(z)| > 4d\gamma_{max}^2\}$.

Proof of Lemma A.2. Since $T_p(z)$ and $R_p(z)$ are diagonal matrices, the following equations hold true

$$\mathrm{Tr}[(R_p(z) - T_p(z))U_p] = \sum_{i=1}^p (R_p^{(i)}(z) - T_p^{(i)}(z))U_p^{(i)},$$

where $(U_p)_{p \geq 1}$ denotes a family of deterministic matrices satisfying (A.3), and

$$(R_p^{(i)}(z) - T_p^{(i)}(z)) = R_p^{(i)}(z)T_p^{(i)}(z)(1/R_p^{(i)}(z) - 1/T_p^{(i)}(z)).$$

We know from [HLN07][Proposition 5.1], that $|R_p^{(i)}(z)T_p^{(i)}(z)| \leq \frac{1}{|\Im(z)|^2}$. Moreover [HLN07][Equation (6.15)] states that

$$\sup_{1 \leq i \leq p} \mathbb{E}[|(1/R_p^{(i)}(z) - 1/T_p^{(i)}(z))|^{2+\delta/2}] \leq \frac{K_2}{n^{1+\delta/4}}.$$

Therefore, we obtain from these inequalities and (A.3) that

$$\begin{aligned} \mathbb{E} \left[\sum_{p \geq 1} \left| \frac{1}{p} \mathrm{Tr}[(T_p(z) - R_p(z))U_p] \right|^{2+\delta/2} \right] &\leq \sum_{p \geq 1} \frac{1}{p} \sum_{i=1}^p |U_p^{(i)}|^{2+\delta/2} \mathbb{E} \left[|T_p^{(i)}(z) - R_p^{(i)}(z)|^{2+\delta/2} \right] \\ &\leq \sum_{p \geq 1} \frac{K^{2+\delta/2} K_2}{|\Im(z)|^2 p^{1+\delta/4}} \\ &\leq \sum_{p \geq 1} \frac{\tilde{K}}{p^{1+\delta/4}} \\ &< +\infty, \end{aligned}$$

with $\tilde{K} = \frac{K^{2+\delta/2} K_2}{16d^2 \gamma_{max}^4}$, since $z \in \mathcal{D}$. Therefore, arguing as in the proof of Lemma A.1, we obtain that

$$\mathbb{P} \left[\frac{1}{p} \mathrm{Tr}[(T_p(z) - R_p(z))U_p] \xrightarrow[n \rightarrow +\infty]{} 0 \right] = 1,$$

and this concludes the proof of Lemma A.2. \square

Proof of Theorem A.1. Since the equivalence relation introduced in Definition A.1 is transitive, we deduce from Lemma A.1 and Lemma A.2 that

$$\frac{1}{p} \mathrm{Tr}[Q_p(z)U_p] \sim \frac{1}{p} \mathrm{Tr}[T_p(z)U_p], \quad (\text{A.4})$$

for $z \in \mathcal{D}$. It remains to prove that this equivalence is true for $z \in \mathbb{C} \setminus \mathbb{R}^+$. To this end, let us define the function

$$f_p(z) = \frac{1}{p} \mathrm{Tr}[Q_p(z)U_p] - \frac{1}{p} \mathrm{Tr}[T_p(z)U_p]$$

for $z \in \mathbb{C} \setminus \mathbb{R}^+$. Proving that (A.4) holds for $z \in \mathbb{C} \setminus \mathbb{R}^+$ is equivalent to prove that $f_p(z)$ converges uniformly to zero. This function is analytic on $z \in \mathbb{C} \setminus \mathbb{R}^+$ since $Q_p(z)$ and $T_p(z)$ are analytic [HLN07][Proposition 5.1]. For $1 \leq i \leq p$, $T_p^{(i)}(z)$ is a Stieltjes transforms [HLN07][Theorem 2.4],

then we can state that $|T_p^{(i)}(z)| \leq \frac{1}{\text{dist}(z, \mathbb{R}^+)}$ with $z \in \mathbb{C} \setminus \mathbb{R}^+$, according to [HLN07][Proposition 2.2]. Moreover, we deduce from [HLN07][Proposition 2.3] and [HS12][Theorem 5.8] that $|Q_p^{(i)}(z)| \leq \|Q_p(z)\|_{sp} \leq \frac{1}{\text{dist}(z, \mathbb{R}^+)}$, where $\|\cdot\|_{sp}$ is the spectral norm, $Q_p^{(i)}(z)$ denotes the i -th diagonal entry of $Q_p(z)$ and $z \in \mathbb{C} \setminus \mathbb{R}^+$. Thus, using (A.3), we have proved that

$$\left| \frac{1}{p} \text{Tr}[Q_p(z)U_p] \right| \leq \frac{K}{\text{dist}(z, \mathbb{R}^+)} \quad \text{and} \quad \left| \frac{1}{p} \text{Tr}[T_p(z)U_p] \right| \leq \frac{K}{\text{dist}(z, \mathbb{R}^+)}. \quad (\text{A.5})$$

Hence, for each compact subset $C \subset \mathbb{C} \setminus \mathbb{R}^+$, f_p is uniformly bounded on C ,

$$|f_n(z)| \leq \frac{2K}{\delta_C},$$

where δ_C is the distance between C and \mathbb{R}^+ . Then, by the normal family Theorem [Rud87][Theorem 14.6] there exists a sub-sequence f_{p_k} which uniformly converges to f^* that is an analytical function on $\mathbb{C} \setminus \mathbb{R}^+$. Let $(z_k)_{k \in \mathbb{N}}$ be a sequence with an accumulation point in \mathcal{D} . Then, for each k , $f_p(z_k) \rightarrow 0$ with probability one as p tends to $+\infty$. This implies that $f^*(z_k) = 0$ for each k . We finally obtain that f^* is identically zero on $\mathbb{C} \setminus \mathbb{R}^+$. Therefore f_p converges uniformly to zero which proves that (A.4) holds for $z \in \mathbb{C} \setminus \mathbb{R}^+$.

Now, it remains to prove that

$$\frac{1}{p} \text{Tr}[Q'_p(z)U_p] \sim \frac{1}{p} \text{Tr}[T'(z)U_p]. \quad (\text{A.6})$$

The functions $Q_p(z)$ and $T_p(z)$ being analytic on $\mathbb{C} \setminus \mathbb{R}^+$ [HS12][Theorem 1.2], [HLN07][Proposition 5.1], the Cauchy integral formula yields that

$$T'_p(z) = \frac{1}{2\pi i} \int_{\rho} \frac{T_p(w)}{(w-z)^2} dw \quad \text{and} \quad Q'_p(z) = \frac{1}{2\pi i} \int_{\rho} \frac{Q_p(w)}{(w-z)^2} dw,$$

where ρ is a path around z in $\mathbb{C} \setminus \mathbb{R}^+$. Hence, we have that

$$\begin{aligned} \frac{1}{p} \text{Tr}[Q'_p(z)U_p] - \frac{1}{p} \text{Tr}[T'_p(z)U_p] &= \frac{1}{p} \text{Tr} \left[\int_{\rho} \frac{Q_p(w) - T_p(w)}{(w-z)^2} dw U_p \right] \\ &= \int_{\rho} \frac{1}{p} \frac{\text{Tr}[(Q_p(w) - T_p(w))U_p]}{(w-z)^2} dw. \end{aligned}$$

Since we have proved that (A.4) holds on $\mathbb{C} \setminus \mathbb{R}^+$, we obtain that for all $w \in \rho$

$$\lim_{n \rightarrow \infty, p/n \rightarrow c} \frac{1}{p} \text{Tr}[(Q_p(w) - T_p(w))U_p] = 0.$$

Moreover, for $w \in \mathbb{C} \setminus \mathbb{R}^+$, we have from (A.5) that

$$\frac{1}{p} \frac{|\text{Tr}[(Q_p(w) - T_p(w))U_p]|}{|w-z|^2} \leq \frac{2K}{\text{dist}(z, \mathbb{R}^+)|w-z|^2} \leq \frac{2K}{\text{dist}(\rho, \mathbb{R}^+)|w-z|^2}.$$

As $\frac{2K}{\text{dist}(\rho, \mathbb{R}^+)|w-z|^2}$ is integrable on ρ , using the dominated convergence theorem, we obtain that (A.6) holds for $z \in \mathbb{C} \setminus \mathbb{R}^+$ which completes the proof of Theorem A.1. \square

A.3 Proof of the main results

We have now all the ingredients needed to prove Theorem 3.1.

Proof of Theorem 3.1. Since Assumption 2.2 hold true for the variance profile \tilde{S}_p , it follows that $(\tilde{S}_p)_{p \geq 1}$ satisfy (A.3). Thus, we deduce from Theorem A.1 that

$$\frac{1}{p} \operatorname{Tr}[Q_p(z)\tilde{S}_p] \sim \frac{1}{p} \operatorname{Tr}[T_p(z)\tilde{S}_p] \quad \text{and} \quad \frac{1}{p} \operatorname{Tr}[Q'_p(z)\tilde{S}_p] \sim \frac{1}{p} \operatorname{Tr}[T'_p(z)\tilde{S}_p]. \quad (\text{A.7})$$

Hence Equation (3.4) directly follows from (A.7) and (3.3) which concludes the proof of Theorem 3.1. \square

Proof of Lemma 3.2. We assume that Assumptions 2.1 and 2.2 hold true in all this proof. According to [HLN07][Theorem 2.4], $T_p(-\lambda)$ can be expressed as follows for $\lambda > 0$

$$T_p(-\lambda) = \int_{\mathbb{R}^+} \frac{\mu(dw)}{w + \lambda}, \quad (\text{A.8})$$

where $\mu = (\mu_{ij})$ is a positive $p \times p$ matrix valued measure such that $\mu_{ij}(\mathbb{R}^+) = \delta_{ij}$. Moreover, $\lambda \mapsto \frac{1}{w/\lambda + 1}$ is differentiable for $w \in \mathbb{R}^+$, its derivative $w \mapsto \frac{w}{(w+\lambda)^2}$ is measurable for $\lambda > 0$ and $|\frac{w}{(w+\lambda)^2}| \leq 1$, thus $T_p(-\lambda)$ is differentiable and

$$T'_p(-\lambda) = \int_{\mathbb{R}^+} \frac{\mu(dw)}{(w + \lambda)^2}.$$

As stated in Proposition 2.1 $\frac{1}{p} \operatorname{Tr}[T_p(z)]$ is a Stieltjes transform

$$\frac{1}{p} \operatorname{Tr}[T_p(z)] = \int_{\mathbb{R}^+} \frac{\nu_p(dw)}{w + \lambda}, \quad (\text{A.9})$$

where there exists $\pi_* \in [0, 1]$ and $\pi : (0, +\infty) \rightarrow [0, +\infty)$ a locally Hölder-continuous function such that $\nu_p(dw) = \pi_* \delta_0(dw) + \pi(dw) \mathbf{1}_{w>0} dw$ [AEK17b][Theorem 2.1].

Under the additional Assumptions 2.3 and 2.4, [AEK17b][Theorem 2.9] asserts that there exists $\tau > 0$ such that $\pi((0, \tau]) = 0$ and if $p < n$ then $\pi_* = 0$. Hence we get from Equation (A.8) that

$$\int_0^\tau \frac{\frac{1}{p} \operatorname{Tr}(\mu(dw))}{w + \lambda} = \int_0^\tau \frac{\nu_p(dw)}{w + \lambda} = 0. \quad (\text{A.10})$$

Since $\frac{1}{w+\lambda} > 0$ for $\lambda > 0$ and $w \geq 0$, we deduce from Equation (A.10) that $\frac{1}{p} \sum_{i=1}^p \mu_{ii}((0, \tau]) = 0$. Which gives us $\mu_{ii}((0, \tau]) = 0$ for $1 \leq i \leq p$ because $(\mu_{ii})_{1 \leq i \leq p}$ are positive measures. Hence

$$T_p(-\lambda) = \frac{\pi_*}{\lambda} + \int_\tau^{+\infty} \frac{\mu(dw)}{w + \lambda}.$$

Let's focus on the case $p < n$, we then have

$$T_p(-\lambda) = \int_\tau^{+\infty} \frac{\mu(dw)}{w + \lambda}.$$

Moreover, $|\frac{1}{w+\lambda}| \leq \frac{1}{\tau}$ for $\lambda > 0$ and $w \geq \tau$, thus by the dominated convergence theorem $T_p(-\lambda)$ admits a limit when λ tends to 0, we denote $T_p(0^-) = \int_\tau^{+\infty} \frac{\mu(dw)}{w}$ this limit. A similar proof allows us to state that $T'_p(-\lambda)$ admits a limits when λ tends to 0, we denote $T'(0^-) = \int_\tau^{+\infty} \frac{\mu(dw)}{w^2}$ this limit. \square

Proof of Corollary 3.1. As stated in Section 2.3, in the case of a quasi doubly stochastic variance profile, $T_p(z)$ is a scalar matrix. Especially, $T_p(z) = m_p(z)I_p$ where $m_p(z)$ is the Stieltjes transform of the Marchenko-Pastur distribution. Moreover, following Equation (A.9) we get

$$m_p(z) = \int_{\mathbb{R}^+} \frac{\nu_p(dw)}{w + \lambda},$$

with $\nu_p(dw) = \pi_*\delta_0(dw) + \pi(dw)\mathbf{1}_{w>0}dw$. We have seen in the previous proof that there exists $\tau > 0$ such that $\nu_p((0, \tau]) = 0$ and if $p < n$ then $\pi_* = 0$. This implies that 0 belongs to the domain of definition of $m_p(\cdot)$. Hence $\lim_{\lambda \rightarrow 0} T_p(-\lambda) = \lim_{\lambda \rightarrow 0} m_p(-\lambda) = m_p(0)$ and consequently $T_p(0^-) = m_p(0)$.

We prove symmetrically that if $p > n$ then 0 belongs to the domain of definition of $\tilde{m}_n(\cdot)$. Moreover, we have seen in Section 3.1 that in the case of a quasi doubly stochastic variance profile $\kappa(z) = \frac{1}{\tilde{m}_n(z)}I_n$. Note that $\tilde{m}_n(0) = \int_{\tilde{\tau}}^{+\infty} \frac{\tilde{\nu}_n(dw)}{w}$ is positive. Hence we conclude that $\kappa(0^+) = \frac{1}{\tilde{m}_n(0)}I_n$. We prove $\kappa'(0^+) = \frac{\tilde{m}'_n(0)}{\tilde{m}_n^2(0)}I_n$ the same way. \square

Proof of Corollary 3.2. We assume that Assumptions 2.1 and 2.2 hold true in all this proof. We first focus on the limit with $\lambda \rightarrow +\infty$. According to [HLN07][Theorem 2.4], $T_p(-\lambda)$ can be expressed as follows for $\lambda > 0$

$$T_p(-\lambda) = \int_{\mathbb{R}^+} \frac{\mu(dw)}{w + \lambda} = \frac{1}{\lambda} \int_{\mathbb{R}^+} \frac{\mu(dw)}{w/\lambda + 1},$$

where $\mu = (\mu_{ij})$ is a $p \times p$ matrix valued measure such that $\mu(\mathbb{R}^+) = I_p$. Since $w \mapsto \frac{1}{w/\lambda + 1}$ is a measurable function for $\lambda > 0$, $\lim_{\lambda \rightarrow +\infty} \frac{1}{w/\lambda + 1} = 1$ for $w \in \mathbb{R}^+$ and $|\frac{1}{w/\lambda + 1}| \leq 1$, we deduce from the dominated convergence theorem that

$$\int_{\mathbb{R}^+} \frac{\mu(dw)}{w/\lambda + 1} \xrightarrow{\lambda \rightarrow +\infty} I_p.$$

Which finally gives us

$$T_p(-\lambda) \xrightarrow{\lambda \rightarrow +\infty} 0_p. \quad (\text{A.11})$$

Moreover, $\lambda \mapsto \frac{1}{w/\lambda + 1}$ is differentiable for $w \in \mathbb{R}^+$, its derivative $w \mapsto \frac{w}{(w + \lambda)^2}$ is measurable for $\lambda > 0$ and $|\frac{w}{(w + \lambda)^2}| \leq 1$, thus $T_p(-\lambda)$ is differentiable and

$$T'_p(-\lambda) = \int_{\mathbb{R}^+} \frac{\mu(dw)}{(w + \lambda)^2}.$$

We can prove the following limits the same way we proved (A.11)

$$\lambda T'_p(-\lambda) \xrightarrow{\lambda \rightarrow +\infty} 0_p \quad \text{and} \quad \lambda^2 T'_p(-\lambda) \xrightarrow{\lambda \rightarrow +\infty} I_p. \quad (\text{A.12})$$

Hence by combining (A.11) and (A.12), we get the limit of the predictive risk for large λ :

$$\lim_{\lambda \rightarrow +\infty} r_\lambda(\Upsilon_n) = \frac{\alpha^2}{p} \text{Tr}[\tilde{\mathcal{S}}_p] + \sigma^2.$$

Let's now compute the limit of the predictive risk for small λ . Considering the definition of $T_p(0^-)$ and $T'_p(0^-)$ given in Lemma 3.2, we directly get the limit of $r_\lambda(\Upsilon_n)$ when $p < n$

$$\lim_{\lambda \rightarrow 0, \lambda > 0} r_\lambda(\Upsilon_n) = \sigma^2 + \frac{\sigma^2}{n} \text{Tr}[\tilde{\mathcal{S}}_p T_p(0^-)].$$

Let's now suppose that $p > n$. We have from (2.1) that

$$T_p^{(j)}(-\lambda) = \frac{\kappa_j(\lambda)}{\lambda(\frac{1}{n}\text{Tr}[\tilde{D}_j] + \kappa_j(\lambda))}, \quad \text{for } 1 \leq j \leq p,$$

with $\kappa_j(\lambda) = \frac{\text{Tr}[\tilde{D}_j]}{\text{Tr}[\tilde{D}_j \tilde{T}_n(-\lambda)]}$. Let's denote $\kappa(\lambda) = \text{diag}(\kappa_j(\lambda))$, we then have a new expression of $T_p(-\lambda)$

$$T_p(-\lambda) = \frac{1}{\lambda} \kappa(\lambda) (\Sigma_n + \kappa(\lambda))^{-1}. \quad (\text{A.13})$$

This equation allows us to derive a new formula for the predictive risk

$$r_\lambda(\Upsilon_n) = \sigma^2 + \frac{\alpha^2}{p} \text{Tr}[\tilde{S}_p \kappa(\lambda) (\Sigma_n + \kappa(\lambda))^{-1}] + \left(\frac{\sigma^2}{n} - \frac{\lambda \alpha^2}{p} \right) \text{Tr}[\kappa'(\lambda) \tilde{S}_p \Sigma_n (\Sigma_n + \kappa(\lambda))^{-2}].$$

Moreover, according to Lemma 3.2 $\kappa(\lambda)$ and $\kappa'(\lambda)$ admit limits when λ tends to 0. Thus we finish this proof by getting considering the definitions of $\kappa(0^+)$ and $\kappa'(0^+)$ from Lemma 3.2.

$$\lim_{\lambda \rightarrow 0, \lambda > 0} r_\lambda(\Upsilon_n) = \frac{\alpha^2}{p} \text{Tr}[\tilde{S}_p \kappa(0^+) (\Sigma_n + \kappa(0^+))^{-1}] + \frac{\sigma^2}{n} \text{Tr}[\kappa'(0^+) \tilde{S}_p \Sigma_n (\Sigma_n + \kappa(0^+))^{-2}] + \sigma^2.$$

□

Proof of Corollary 3.3. Let us prove Equation (3.5). The function $g : \lambda \mapsto \hat{r}_\lambda(X_n)$ is differentiable for $\lambda > 0$, with

$$g'(\lambda) = 2 \left(\frac{\lambda \alpha^2}{p} - \frac{\sigma^2}{n} \right) \text{Tr}[\Sigma_n \hat{\Sigma}_n Q_p(\lambda)^3]$$

Since the symmetric and positive semi-definite matrices $\hat{\Sigma}_n$ and $Q_p(\lambda)^3$ are commuting, then $\Sigma_n \hat{\Sigma}_n Q_p(\lambda)^3$ is also symmetric positive semi-definite.

Thus $\text{Tr}[\Sigma_n \hat{\Sigma}_n Q_p(\lambda)^3] \geq 0$ which means that the sign of $g'(\lambda)$ only depends on $\left(\frac{\lambda \alpha^2}{p} - \frac{\sigma^2}{n} \right)$. Indeed, if $\lambda \leq \lambda_*$ then $g'(\lambda) \leq 0$, and if $\lambda \geq \lambda_*$ then $g'(\lambda) \geq 0$. This proves that Equation (3.5) holds true, and this concludes the proof of Corollary 3.3. □

A.4 Random matrices with a variance profile in free probability

We conclude this section by relating the computation of $T_p(z)$ to the notion of \mathcal{R} -transform in operator-valued free probability. The study of the spectral distribution of a generalized Wigner matrix (that is with a variance profile) dates back to [Shl96], where, using Voiculescu's notion of asymptotic freeness [VDN92], Shlyakhtenko proved that independent generalized Wigner matrices are asymptotically free with amalgamation over the diagonal. This property has then further been studied in [ACD⁺21, Mal20] for independent permutation invariant matrices with variance profiles.

More formally, let $\mathbb{D}_p(\mathbb{C})^+$ (resp. $\mathbb{D}_p(\mathbb{C})^-$) denotes the set of diagonal matrices \mathcal{Z} of size p with diagonal complex entries having positive (resp. negative) imaginary parts. Recall that, for any square matrix A , we denote by $\Delta[A]$ the diagonal matrix whose diagonal entries are those of A . The operator-valued Stieltjes transform \mathcal{G}_A of a Hermitian matrix A is then defined as the map

$$\mathcal{G}_A : \begin{array}{l} \mathbb{D}_p(\mathbb{C})^- \rightarrow \mathbb{D}_p(\mathbb{C})^+ \\ \mathcal{Z} \mapsto \Delta[(A - \mathcal{Z})^{-1}], \end{array} \quad (\text{A.14})$$

and it is sometimes defined as $-\mathcal{G}_A$. The operator-valued \mathcal{R} -transform \mathcal{R}_A of a Hermitian matrix A is the unique analytic map satisfying,

$$\mathcal{G}_A(\mathcal{Z}) = \left(-\mathcal{Z} + \mathcal{R}_A(-\mathcal{G}_A(\mathcal{Z})) \right)^{-1}, \quad (\text{A.15})$$

for all \mathcal{Z} in $\mathbb{D}_p(\mathbb{C})^-$ whose diagonal entries have imaginary parts large enough in absolute value. An efficient way to produce a deterministic equivalent for a random matrix A is then to find a simple approximation \mathcal{R}_A^\square of \mathcal{R}_A . In this manner, one can then approximate the operator-valued Stieltjes transform of A by the solution of the fixed point equation (A.15) where \mathcal{R}_A is replaced by \mathcal{R}_A^\square . This method also allows us to compute the spectrum of perturbations of A by independent matrices, see for instance [BM20].

For a generalized Wigner matrix W_p of size $p \times p$ with a symmetric variance profile Γ_p , Shlyakhtenko proves [Shl96] that a good approximation of \mathcal{R}_{W_p} is the deterministic linear map

$$\mathcal{R}_{W_p}^\square(\mathcal{Z}) = \text{deg} \left(\frac{1}{p} \Gamma_p \mathcal{Z} \right), \quad \text{for all } \mathcal{Z} \in \mathbb{D}_p(\mathbb{C})^-,$$

where for a matrix A , we denote by $\text{deg}(A)$ the diagonal matrix, whose k -diagonal element is the sum of the entries of the k -row of A .

To the best of our knowledge, apart from generalized Wigner matrices, there does not exist any other class of random matrices for which a simple approximation of the diagonal-valued \mathcal{R} -transform is known yet. Nevertheless, we can now remark that, for $\mathcal{Z} = zI_p$ with $z \in \mathbb{C} \setminus \mathbb{R}^+$, the diagonal matrix $T_p(z)$ solution of the Dyson equation (2.3) can be written as satisfying the fixed-point equation

$$T_p(z) = (-zI_p + \mathcal{R}_{\Sigma_n}^\square(-T_p(z)))^{-1},$$

where $\mathcal{R}_{\Sigma_n}^\square$ is the non-linear map

$$\mathcal{R}_{\Sigma_n}^\square(\mathcal{Z}) = \text{deg} \left(\frac{1}{n} \Gamma_n^t \left[I_n - \text{deg} \left(\frac{1}{n} \Gamma_n \mathcal{Z} \right) \right]^{-1} \right), \quad \text{for all } \mathcal{Z} \in \mathbb{D}_p(\mathbb{C})^-.$$

This suggests that $\mathcal{R}_{\Sigma_n}^\square$ is indeed a simple approximation of the operator-valued \mathcal{R} -transform of Σ_n .

References

- [ACD⁺21] Benson Au, Guillaume Cébron, Antoine Dahlqvist, Franck Gabriel, and Camille Male. Freeness over the diagonal for large random matrices. *The Annals of Probability*, 49(1):157 – 179, 2021.
- [AEK17a] Oskari H. Ajanki, László Erdős, and Torben Krüger. Universality for general Wigner-type matrices. *Probability Theory and Related Fields*, 169(3):667–727, 2017.
- [AEK17b] Johannes Alt, László Erdős, and Torben Krüger. Local law for random Gram matrices. *Electron. J. Probab.*, 22:41 pp., 2017.
- [AEK18] Johannes Alt, László Erdős, and Torben Krüger. Local inhomogeneous circular law. *Ann. Appl. Probab.*, 28(1):148–203, 02 2018.
- [AEK19] Oskari H. Ajanki, László Erdős, and Torben Krüger. Stability of the matrix Dyson equation and random matrices with correlations. *Probability Theory and Related Fields*, 173(1):293–373, 2019.

- [AGB⁺15] Stefano Allesina, Jacopo Grilli, György Barabás, Si Tang, Johnatan Aljadeff, and Amos Maritan. Predicting the stability of large structured food webs. *Nature communications*, 6(1):7842, 2015.
- [ARS15] Johnatan Aljadeff, David Renfrew, and Merav Stern. Eigenvalues of block structured asymmetric random matrices. *Journal of Mathematical Physics*, 56(10), 2015.
- [ASS15] Johnatan Aljadeff, Merav Stern, and Tatyana Sharpee. Transition to chaos in random networks with cell-type-specific connectivity. *Physical review letters*, 114(8):088101, 2015.
- [AT15] Stefano Allesina and Si Tang. The stability–complexity relationship at age 40: a random matrix perspective. *Population Ecology*, 57(1):63–75, 2015.
- [Bac24] Francis Bach. High-dimensional analysis of double descent for linear regression with random projections. *SIAM Journal on Mathematics of Data Science*, 6(1):26–50, 2024.
- [BBK⁺19] Andreas Buja, Lawrence Brown, Arun Kumar Kuchibhotla, Richard Berk, Edward George, and Linda Zhao. Models as Approximations II: A Model-Free Theory of Parametric Regression. *Statistical Science*, 34(4):545 – 565, 2019.
- [BDF17] J. Bigot, C. Deledalle, and D. Féral. Generalized sure for optimal shrinkage of singular values in low-rank matrix denoising. *Journal of Machine Learning Research*, 18(137):1–50, 2017.
- [Ber82] Rudolf Beran. Robust Estimation in Models for Independent Non-Identically Distributed Data. *The Annals of Statistics*, 10(2):415 – 428, 1982.
- [BHMM19] Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, 2019.
- [BHX20] Mikhail Belkin, Daniel Hsu, and Ji Xu. Two models of double descent for weak features. *SIAM Journal on Mathematics of Data Science*, 2(4):1167–1180, 2020.
- [BHX23] Zhigang Bao, Qiyang Han, and Xiacong Xu. A leave-one-out approach to approximate message passing, 2023.
- [BM20] Jérémie Bigot and Camille Male. Freeness over the diagonal and outliers detection in deformed random matrices with a variance profile. *Information and Inference: A Journal of the IMA*, 10(3):863–919, 07 2020.
- [BMG13] Juan Andres Bazerque, Gonzalo Mateos, and Georgios B. Giannakis. *Inference of Poisson count processes using low-rank tensor data*, pages 5989–5993. ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings, 10 2013.
- [BS10] Zhidong Bai and Jack W. Silverstein. *Spectral analysis of large dimensional random matrices*. Springer Series in Statistics. Springer, New York, second edition, 2010.
- [CD11] Romain Couillet and Mérouane Debbah. *Random Matrix Methods for Wireless Communications*. Cambridge University Press, 2011.
- [CHNR18] Nicholas Cook, Walid Hachem, Jamal Najim, and David Renfrew. Non-hermitian random matrices with a variance profile (i): deterministic equivalents and limiting esds. *Electron. J. Probab.*, 23:61 pp., 2018.

- [CZL19] Yuanpei Cao, Anru Zhang, and Hongzhe Li. Multisample estimation of bacterial composition matrices in metagenomics data. *Biometrika*, 107(1):75–92, 12 2019.
- [Dic16] Lee H. Dicker. Ridge regression and asymptotic minimax estimation over spheres of growing dimension. *Bernoulli*, 22(1):1 – 37, 2016.
- [DW18] Edgar Dobriban and Stefan Wager. High-dimensional asymptotics of prediction: Ridge regression and classification. *The Annals of Statistics*, 46(1):247 – 279, 2018.
- [Efr04] Bradley Efron. The estimation of prediction error. *Journal of the American Statistical Association*, 99(467):619–632, 2004.
- [EK18] Nouredine El Karoui. On the impact of predictor geometry on the performance on high-dimensional ridge-regularized generalized robust regression estimators. *Probability Theory and Related Fields*, 170:95–175, 02 2018.
- [EYY12] László Erdős, Horng-Tzer Yau, and Jun Yin. Bulk universality for generalized Wigner matrices. *Probability Theory and Related Fields*, 154(1-2):341–407, 2012.
- [HLN06] W. Hachem, P. Loubaton, and J. Najim. The empirical distribution of the eigenvalues of a gram matrix with a given variance profile. *Annales de l’Institut Henri Poincaré (B) Probability and Statistics*, 42(6):649 – 670, 2006.
- [HLN07] Walid Hachem, Philippe Loubaton, and Jamal Najim. Deterministic equivalents for certain functionals of large random matrices. *The Annals of Applied Probability*, 17(3):875 – 930, 2007.
- [HMRT22] Trevor Hastie, Andrea Montanari, Saharon Rosset, and Ryan J. Tibshirani. Surprises in high-dimensional ridgeless least squares interpolation. *The Annals of Statistics*, 50(2):949 – 986, 2022.
- [HS12] Peter D. Hislop and Israel Michael Sigal. *Introduction to spectral theory: With applications to Schrödinger operators*, volume 113. Springer Science & Business Media, 2012.
- [HTW15] Trevor Hastie, Robert Tibshirani, and Martin Wainwright. *Statistical Learning with Sparsity: The Lasso and Generalizations*. Chapman & Hall/CRC, 2015.
- [JFL22] Xiao Han Jianqing Fan, Yingying Fan and Jinchi Lv. Asymptotic theory of eigenvectors for random matrices with diverging spikes. *Journal of the American Statistical Association*, 117(538):996–1009, 2022.
- [KBB⁺20] Arun K. Kuchibhotla, Lawrence D. Brown, Andreas Buja, Junhui Cai, Edward I. George, and Linda H. Zhao. Valid post-selection inference in model-free linear regression. *The Annals of Statistics*, 48(5):2953 – 2981, 2020.
- [LC18] Zhenyu Liao and Romain Couillet. The dynamics of learning: A random matrix approach. In *International Conference on Machine Learning*, pages 3072–3081. PMLR, 2018.
- [LDS18] Lydia T. Liu, Edgar Dobriban, and Amit Singer. epca: High dimensional exponential family pca. *Ann. Appl. Stat.*, 12(4):2121–2150, 12 2018.
- [LHT23] Liangchen Liu, Juncai He, and Richard Tsai. Linear regression on manifold structured data: the impact of extrinsic geometry on solutions, 2023.

- [Mal20] Camille Male. Traffic distributions and independence: permutation invariant random matrices and the three notions of independence. *Mem. Amer. Math. Soc.*, 267(1300):v+88, 2020.
- [NPW21] Jamshid Namdari, Debashis Paul, and Lili Wang. High-dimensional linear models: A random matrix perspective. *Sankhya A: The Indian Journal of Statistics*, 83(2):645–695, 2021.
- [RMR21] Dominic Richards, Jaouad Mourtada, and Lorenzo Rosasco. Asymptotics of ridge(less) regression under general source condition. In Arindam Banerjee and Kenji Fukumizu, editors, *The 24th International Conference on Artificial Intelligence and Statistics, AISTATS 2021, April 13-15, 2021, Virtual Event*, volume 130 of *Proceedings of Machine Learning Research*, pages 3889–3897. PMLR, 2021.
- [Rud87] Walter Rudin. *Real and complex analysis*. McGraw-Hill Book Co., New York, third edition, 1987.
- [SHDW14] Joseph Salmon, Zachary T. Harmany, Charles-Alban Deledalle, and Rebecca Willett. Poisson noise reduction with non-local PCA. *Journal of Mathematical Imaging and Vision*, 48(2):279–294, 2014.
- [Shl96] Dimitri Shlyakhtenko. Random gaussian band matrices and freeness with amalgamation. *International Mathematics Research Notices*, 1996(20):1013–1025, 1996.
- [SKR⁺23] Rylan Schaeffer, Mikail Khona, Zachary Robertson, Akhilan Boopathy, Kateryna Pistunova, Jason W Rocks, Ila Rani Fiete, and Oluwasanmi Koyejo. Double descent demystified: Identifying, interpreting & ablating the sources of a deep learning puzzle. *arXiv preprint arXiv:2303.14151*, 2023.
- [TV04] Antonia Maria Tulino and Sergio Verdú. Random matrix theory and wireless communications. *Found. Trends Commun. Inf. Theory*, 1(1), 2004.
- [UHZB16] Madeleine Udell, Corinne Horn, Reza Zadeh, and Stephen Boyd. Generalized low rank models. *Foundations and Trends in Machine Learning*, 9(1):1–118, 2016.
- [VDN92] D. V. Voiculescu, K. J. Dykema, and A. Nica. *Free random variables*, volume 1 of *CRM Monograph Series*. American Mathematical Society, Providence, RI, 1992. A noncommutative probability approach to free products with applications to random matrices, operator algebras and harmonic analysis on free groups.
- [ZCW22] Anru R. Zhang, T. Tony Cai, and Yihong Wu. Heteroskedastic PCA: Algorithm, optimality, and applications. *The Annals of Statistics*, 50(1):53 – 80, 2022.