

Statistica Sinica Preprint No: SS-2025-0115

Title	Empirical Bayes Data Integration for Multi-Response Regression
Manuscript ID	SS-2025-0115
URL	http://www.stat.sinica.edu.tw/statistica/
DOI	10.5705/ss.202025.0115
Complete List of Authors	Antik Chakraborty and Fei Xue
Corresponding Authors	Fei Xue
E-mails	feixue@purdue.edu
Notice: Accepted author version.	

EMPIRICAL BAYES DATA INTEGRATION FOR MULTI-RESPONSE REGRESSION

Antik Chakraborty and Fei Xue

Purdue University

Abstract: Motivated by applications in tissue-wide association studies (TWAS), we develop a flexible and theoretically grounded empirical Bayes approach for integrating data obtained from different sources. We propose a linear shrinkage estimator that effectively shrinks singular values of a data matrix. This problem is closely connected to estimating covariance matrices under a specific loss, for which we develop asymptotically optimal estimators. The basic linear shrinkage estimator is then extended to a local linear shrinkage estimator, offering greater flexibility. Crucially, the proposed method works under sparse/dense or low-rank/non low-rank parameter settings unlike well-known sparse or reduced rank estimators in the literature. Furthermore, the empirical Bayes approach offers greater scalability in computation compared to intensive full Bayes procedures. The method is evaluated through an extensive set of numerical experiments, and applied to a real TWAS data obtained from the Genotype-Tissue Expression (GTEx) project.

Key words and phrases: Covariance matrix estimation, GTEx, reduced rank, shrinkage, TWAS.

1 Introduction

Genome-wide association studies (GWAS) aim to identify potential genotype markers associated with a particular phenotype, typically a disease. GWAS data usually involve a large number of genes, which caused the explosion of statistical methods that are able to handle many variables. More recently, genetic scientists are collecting gene expression data from

multiple tissues, e.g., brain tissues and heart tissues (Mai et al., 2023; Xue and Li, 2022). To leverage these multi-tissue gene expression data in identification of genotype markers, tissue-wide association studies (TWAS) prioritize genes that are functionally linked to the phenotype by associating genetically predicted gene expression with the phenotype (Wainberg et al., 2019), where we need to predict multi-tissue gene expression values based on genotype data.

Although individual analysis on the prediction of gene expression in each tissue is possible with existing methods, the resulting analysis does not integrate potential shared information across tissues. Indeed, when data from multiple sources with some commonality are available, a joint analysis across all the data sources allows for borrowing of strength. However, data obtained from TWAS studies might not adhere to sparsity or low-rank structure (Heap et al., 2009; Gresle et al., 2020), which is a crucial assumption for many of the available statistical methods (Velu and Reinsel, 2013). Such assumptions are also very hard to verify in practice. In this article, our aim is to develop methods that a) successfully integrate data across multiple sources (e.g., multiple tissues), b) are computationally scalable, and c) perform well regardless of specific structures within the parameter.

Specifically, we focus on the case where the ordinary least squares (OLS) estimate is available under a linear regression model with a response and a set of predictors in each data source (or tissue). Across different data sources, the predictors are the same but the response varies. Vector-valued outcomes from multiple sources are also encountered in other scientific disciplines including finance, bioinformatics, and growth curve models. For example, in finance stock prices of multiple companies are studied in relation to the same set of predictors.

Under the assumption of a linear relationship between the predictors and the response variables in all data sources, the interest centers on recovering the matrix of regression coefficients. Traditionally, this problem was studied through the lens of reduced-rank regression (Anderson, 1951; Izenman, 1975; Velu and Reinsel, 2013; Geweke, 1991). More recently, methods that are able to handle high-dimensional (many predictors) along with the reduced rank nature of the coefficient matrix have been developed (Yuan et al., 2007; Bunea et al., 2011, 2012; Chen and Huang, 2012) catering to modern applications, including denoising Gaussian matrices, a closely related problem. Within the frequentist framework, the reduced rank constraint and presence of effect of a subset of predictor variables are most naturally expressed in terms of penalized regression which can be interpreted as suitable priors over the parameter space leading to a maximum *a posteriori* (MAP) interpretation of the estimators. Full Bayes treatment of the problem has also been carried out (Bai and Ghosh, 2018; Chakraborty et al., 2020). The resulting procedures principally shrink the coefficient matrix towards low-rank structures.

A common thread between all penalty-based methods and Bayesian versions thereof is the assumption of an underlying structure (low-rank and/or row-sparse) in the coefficient matrix, which statistically is meaningful but is very hard to verify in practice. Here, we take a different view to the problem in line with the three objectives outlined earlier. Our solution is an empirical Bayes one, which is able to borrow information from multiple sources, avoids computationally intensive full Bayes procedures, and applies to situations where no specific structural information about the parameter is available.

In fact, the OLS estimates from each source can be treated as observed data under an additive model where Gaussian noise is added to the true matrix of coefficients. Efron and

Morris (1972) proposed an empirical Bayes estimator for this mean matrix estimation problem which induces linear shrinkage on the singular values of the observation matrix. More recently, Matsuda and Komaki (2015) developed superharmonic priors for singular value shrinkage for matrix-valued mean parameters. In their development, they closely follow Stein (1981) who showed optimality properties of Bayes estimates with a superharmonic prior distribution. These singular value superharmonic priors were then used in the context of matrix completion (Matsuda and Komaki, 2019), estimation under matrix quadratic loss Matsuda and Strawderman (2022). This class of priors also place increasing amount of mass near low-rank matrices, thus implicitly assuming such an underlying structure. Moreover, Wang and Zhao (2021) developed an empirical Bayes estimator for multivariate linear regression problems but they mainly focused on prediction.

Efron and Morris (1972, 1976) noted that the linear shrinkage estimator of the mean matrix can alternatively be interpreted as the posterior mean under Gaussian priors; the optimal decision in this context under the Frobenius loss. However, from a practical perspective, specification of Gaussian priors requires one to specify a prior covariance matrix which is not immediate, especially without certain structural assumption on the parameter. Efron and Morris (1976) subsequently show that the success of linear shrinkage estimators relies on the accurate estimation of the marginal covariance of the data under the relative savings loss for estimating covariance matrices. Efron and Morris (1976) considered rotation invariant estimators of the covariance matrix. Their initial suggestion was to use a linear shrinkage of eigenvalues for the covariance estimation problem. Linear shrinkage of eigenvalues was pioneered by Ledoit and Wolf (2004). In a series of papers the authors have developed general non-linear shrinkage estimators and studied their properties; see for ex-

ample Ledoit and Péché (2011); Ledoit and Wolf (2012, 2018, 2022). Although Ledoit and Wolf (2022) established an optimal (asymptotic) shrinkage rule under several loss functions within the class of rotation invariant estimators, they have not considered the relative savings loss which is important for the regression coefficient matrix estimation. Coincidentally, the need to accurately estimate the covariance matrix is also necessary for prediction problems. Banerjee et al. (2021) studies Bayes predictive estimators for multivariate Gaussian models. Here, the optimal Bayes rule involves a quadratic form in the unknown covariance. The authors assume a spiked covariance structure (Paul, 2007) for estimation of the unknown covariance, whereas, in this work, we take a loss-minimization-based approach.

In this paper, we consider the regression coefficient matrix estimation problem under the empirical Bayes framework, which mainly relies on the estimation of the covariance matrix of standardized OLS estimates. We first develop an asymptotically optimal shrinkage rule for estimating the covariance matrix under the relative savings loss, and then propose a linear shrinkage estimator of the regression coefficient matrix based on the estimated covariance matrix. In this way, our proposed coefficient estimator is optimal asymptotically within the class of linear shrinkage estimators of the regression coefficient matrix. The proposed estimators are derived under settings when the number of data sources available is larger than the number of variables, and when it is not.

The proposed shrinkage rule that minimizes the relative savings loss is defined in terms of a smoothing parameter. Our next contribution is to develop a data-dependent choice of this smoothing parameter, using the technique of unbiased risk estimation (SURE). Our numerical experiments reveal that such data-dependent tuning often results in improved risk results.

Finally, to gain further flexibility, we extend our estimator to a situation where the prior is a mixture of Gaussian, which results in an adaptively weighted local linear shrinkage rule for estimating the mean matrix. This is useful for capturing any complex prior structure in the parameter. However, a fundamental benefit of the proposed approach is that we can compute the proposed estimator without having to carefully devise an explicit prior that embeds this complex structure. Our numerical experiments reveal that the local linear shrinkage estimator has better or at par performance with estimators that are specifically designed for structural parameters, without resorting to such assumptions.

In Section 2 we introduce the problem and propose the estimator. Section 3 is devoted to the development of the covariance estimator and its data-dependent version. Section 4 describes a version of the proposed estimator under a mixture prior. In Sections 5 and 6 we evaluate the proposed estimator through numerical experiments, compare it with other approaches, and apply it to the Genotype-Tissue Expression (GTEx) data (Lonsdale et al., 2013).

2 Data integration by linear shrinkage

In this section, we introduce the problem setup and propose an empirical Bayesian framework for estimation of genotype effects on gene expression levels. For a certain gene, we let $y_i^{(t)}$ be its expression level for the i -th subject in the t -th tissue, and x_i be a fixed p -dimensional vector of single nucleotide polymorphisms (SNPs) for the i -th subject. Here $i = 1, \dots, N$ and $t = 1, \dots, n$. We assume the sample size $N > p$ but allow the tissue size n to be either larger or smaller than p . Consider a linear regression model for the t -th tissue

$$y_i^{(t)} = x_i^T \beta^{(t)} + \varepsilon_i^{(t)}, \quad (2.1)$$

where $\beta^{(t)}$ is a p -dimensional coefficient vector and $\varepsilon_i^{(t)} \sim N(0, \sigma^2)$ is the error term. We assume that $\varepsilon_1^{(t)}, \dots, \varepsilon_N^{(t)}$ are conditionally independent given the tissue-specific coefficient $\beta^{(t)}$.

The OLS estimator $\widehat{\beta}^{(t)} = (X^T X)^{-1} X^T y^{(t)}$ is a single-tissue estimator of $\beta^{(t)}$, where $X = (x_1, \dots, x_N)^T$ is the $N \times p$ fixed design matrix, and $y^{(t)}$ is a vector consisting of $y_i^{(t)}$ ($1 \leq i \leq N$). Since $\varepsilon_i^{(t)} \sim N(0, \sigma^2)$, we have $\widehat{\beta}^{(t)} \mid \beta^{(t)} \sim N(\beta^{(t)}, \sigma^2 (X^T X)^{-1})$. Let B be a $n \times p$ matrix with $\beta^{(t)}$ as the t -th row, and \widetilde{B} be a $n \times p$ matrix with $\widehat{\beta}^{(t)}$ as the t -th row.

By sufficiency of the OLS estimator, we can consider \widetilde{B} as our data matrix and aim to estimate B under the model

$$\widehat{\beta}^{(t)} = \beta^{(t)} + u^{(t)} \quad \text{with} \quad u^{(t)} \sim N(0, Q), \quad (2.2)$$

where $Q = \sigma^2 (X^T X)^{-1}$. To leverage shared information across multiple tissues, we propose to assign a common prior π to $\beta^{(t)}$ for $t = 1, \dots, n$. For an estimator \widetilde{B} of B , we consider the Frobenius loss

$$\mathbf{L}(B, \widetilde{B}) = \sum_{t=1}^n (\widetilde{\beta}^{(t)} - \beta^{(t)})^T (\widetilde{\beta}^{(t)} - \beta^{(t)}) = \text{tr}[(\widetilde{B} - B)(\widetilde{B} - B)^T] \quad (2.3)$$

where for any matrix A , we write $\text{tr}(A) = \sum_j A_{jj}$. The corresponding posterior expected loss is $\mathbb{E}_{B|\widetilde{B}}[\mathbf{L}(B, \widetilde{B})]$, and the minimizer of the $\mathbb{E}_{B|\widetilde{B}}[\mathbf{L}(B, \widetilde{B})]$ is the posterior expectation $\mathbb{E}(B \mid \widetilde{B})$. This is a vector-valued version of the canonical Normal means problem, which was considered by Efron and Morris (1972) as an extension of Stein's shrinkage idea to vector-valued observations. In the most general setup of the problem, σ^2 is unknown. However, it can be estimated from the tissue-specific regressions, and in this paper we use the average of those estimators as a fixed value of σ^2 .

In particular, a Gaussian prior distribution for $\beta^{(t)}$ yields a linear shrinkage decision rule.

Specifically, if π is $N(0, Q^{1/2}\Omega Q^{1/2})$, then $\mathbb{E}(\beta^{(t)} \mid \widehat{B}) = (\mathbf{I} - C)\widehat{\beta}^{(t)}$, where $C = Q^{1/2}(\mathbf{I} + \Omega)^{-1}Q^{-1/2}$. Here, we scale the prior with the observation noise Q which is quite common in the Bayesian literature (Park and Casella, 2008). In order to use this estimator, one needs to specify Ω , which is not straightforward without making structural assumptions on $\beta^{(t)}$. However, an empirical Bayes analysis could still be carried out without this specification by noting that $\Sigma = (\mathbf{I} + \Omega) = \text{Cov}(\widehat{\beta}_*^{(t)})$, where $\widehat{\beta}_*^{(t)} = Q^{-1/2}\widehat{\beta}^{(t)}$. That is, we can estimate Ω or Σ from the observed data.

Therefore, given an estimate $\widehat{\Sigma}^{-1}$ of Σ^{-1} , we propose to estimate the parameter $\beta^{(t)}$ using an estimated posterior mean $\mathbb{E}(\beta^{(t)} \mid \widehat{B}) = (\mathbf{I} - \widehat{C})\widehat{\beta}^{(t)}$, where $\widehat{C} = Q^{1/2}\widehat{\Sigma}^{-1}Q^{-1/2}$ is a plug-in estimate of C since Q is known. This implies that there is an intricate connection between estimating the parameter matrix B under the Frobenius loss and estimating the marginal covariance matrix Σ . In fact, by Efron and Morris (1976), estimating B under Frobenius loss, within the class of linear shrinkage estimators, i.e. $\widetilde{\beta}^{(t)} = (\mathbf{I} - \widehat{C})\widehat{\beta}^{(t)}$, is equivalent to the problem of estimating Σ^{-1} under a relative savings loss $\mathbf{L}(\Sigma^{-1}, \widehat{\Sigma}^{-1}) = \text{tr}((\Sigma^{-1} - \widehat{\Sigma}^{-1})^2 S)$, where $S = \sum_{t=1}^n \widehat{\beta}_*^{(t)} \widehat{\beta}_*^{(t)^\top}$. This equivalence is recorded in the following Proposition 1 for the sake of completeness. Its corresponding proof is provided in Section S.6.1 of the Supplement.

Proposition 1. Suppose $\beta^{(t)} \sim N(0, Q^{1/2}\Omega Q^{1/2})$ with known Q and $\widehat{\beta}^{(t)} \mid \beta^{(t)} \sim N(\beta^{(t)}, Q)$.

Consider the loss $\mathbf{L}(B, \widetilde{B}) = \text{tr}[(\widetilde{B} - B)(\widetilde{B} - B)^\top]$ and estimators $\widetilde{B} = \widehat{B}(\mathbf{I} - \widetilde{C})$, where \widetilde{C} is an estimator of $C = Q^{1/2}(\mathbf{I} + \Omega)^{-1}Q^{-1/2}$. Then

$$\mathbb{E}_{(B, \widetilde{B})}[\mathbf{L}(B, \widetilde{B})] = \mathbb{E}_{\widehat{B}}[\text{tr}\{(\widetilde{\Sigma}^{-1} - \Sigma^{-1})^2\} \widehat{B}^\top \widehat{B}] + \text{constant}.$$

Hence, the best linear shrinkage estimator of B should be based on the estimate $\widehat{\Sigma}^{-1}$ that is optimal in terms of the relative savings loss. We now turn our attention to the problem of estimating Σ^{-1} under the loss $\mathbf{L}(\Sigma^{-1}, \widehat{\Sigma}^{-1}) = \text{tr}((\Sigma^{-1} - \widehat{\Sigma}^{-1})^2 S)$. Efron and Morris

(1972, 1976) suggested two estimators for $\widehat{\Sigma}^{-1}$. The first is the natural unbiased estimator which is obtained by observing that $\widehat{\beta}_*^{(t)} \stackrel{iid}{\sim} N(0, \Sigma)$, and the standard multivariate Gaussian distribution theory yields $S^{-1} \sim \text{inv-Wishart}(\Sigma^{-1}, n)$. Thus, $\widehat{\Sigma}^{-1} = (n - p - 1)S^{-1}$ is an unbiased estimator of Σ^{-1} . The second involves a linear shrinkage estimator of Σ^{-1} . But linear shrinkage might not perform well under certain situations (Ledoit and Wolf, 2012).

3 Covariance shrinkage

In this section, we develop a rotation invariant estimator for Σ . Since this is an independent problem of interest, we consider a general setup: suppose that n p -dimensional independent and identically distributed observations are available with zero mean and covariance matrix $\Sigma_n = \Sigma$. Our following results include cases when $n > p$ and $n < p$. In this section, we use the subscript n to emphasize on the asymptotic framework we work in. The observed data is arranged in an $n \times p$ matrix Z_n . In the notation of the previous section, the rows of Z_n are given by $\widehat{\beta}_*^{(t)}$. Let $S_n = n^{-1}Z_n^T Z_n$ be the sample covariance matrix. Consider the spectral decomposition of $S_n = U_n \Lambda_n U_n^T = \sum_{i=1}^p \lambda_{n,i} u_{n,i} u_{n,i}^T$, where U_n is an orthogonal matrix with $u_{n,i}$ as the i -th column of U_n and Λ_n is a diagonal matrix with elements $\boldsymbol{\lambda}_n = (\lambda_{n,1}, \dots, \lambda_{n,p})^T$ as the corresponding eigenvalues arranged in a non-decreasing order.

We focus on the class of rotation invariant estimators $\widetilde{\Sigma}_n = U_n \widetilde{\Delta}_n U_n^T$, where $\widetilde{\Delta}_n = \text{diag}(\delta_n(\lambda_{n,1}), \dots, \delta_n(\lambda_{n,p}))$ and δ_n is a positive univariate function that may depend on S_n . Such estimators $\widetilde{\Sigma}_n$ are rotation invariant since multiplying the data Z_n by an orthogonal matrix with a determinant of one rotates the estimators accordingly. Modern high-dimensional methods often rely on low-dimensional structures of the data when considering the problem of estimating unknown covariance matrices, for instance, sparsity. While properties of these

estimators are well understood, relatively little is known about their performance when such assumptions do not hold. Instead, we pursue an estimator that principally shrinks eigenvalues without assuming any structure. As indicated earlier, we focus on the problem of estimation of Σ_n^{-1} under the loss $\mathbf{L}_n(\Sigma_n^{-1}, \tilde{\Sigma}_n^{-1}) = \text{tr}[(\Sigma_n^{-1} - \tilde{\Sigma}_n^{-1})^2 S_n]$.

In our study, we consider a general version of the relative savings loss. Specifically, we consider

$$\mathbf{L}_{m,n}(\Sigma_n^{-1}, \tilde{\Sigma}_n^{-1}) = \frac{1}{p} \text{tr}[(\Sigma_n^{-1} - \tilde{\Sigma}_n^{-1})^2 S_n^m]$$

for $m = 0, 1, 2, \dots$.

Remark 1. The case $m = 0$ corresponds to the inverse Frobenius loss; see also Ledoit and Wolf (2018); Haff (1979), whereas for $m = 1$, we recover the relative savings loss, the focus of our paper. Boukehil et al. (2021); Kubokawa and Srivastava (2008) studied the case $m = 2$. From a practical perspective, large values of m put increasingly larger weight on large sample eigenvalues in terms of their contribution to the loss. Another major motivation for studying this general class of loss functions is to draw a distinction between the cases $m = 0$ and $m \geq 1$. As it turns out, if $m = 0$, then the optimal shrinkage function depends on the population eigenvalues through the limiting population eigenvalue distribution H . This problem was addressed in Ledoit and Wolf (2018). Their solution was a numerical one, namely, the QuEST function. However, our results will show that for $m \geq 1$, the optimal solution does not involve H . In other words, an explicit solution is available.

Our goal here is to provide an optimal shrinkage rule under the loss function $\mathbf{L}_{m,n}$. We do so using three key steps - 1) we first find an almost sure non-random limit of $\mathbf{L}_{m,n}$, 2) then we find the shrinkage rule which minimizes this limit, and 3) find a consistent estimator of this optimum shrinkage rule.

To find the almost sure limit of $\mathbf{L}_{m,n}$, we introduce the following notations and rewrite the loss function. We let $F_n(x) = \frac{1}{p} \sum_{j=1}^p \mathbb{1}_{\{\lambda_{n,j} \leq x\}}$ and $F_n^*(x) = p^{-1} \sum_{j=1}^p \mathbb{1}_{\{\lambda_{n,j}^{-1} \leq x\}}$ be the empirical distribution functions of the sample eigenvalues and inverse eigenvalues, respectively. Clearly, for $x > 0$, $F_n^*(x) = 1 - F_n(1/x)$. Recall that $\tilde{\Sigma}_n = U_n \tilde{\Delta}_n U_n^T$, where $\tilde{\Delta}_n = \text{diag}(\delta_n(\lambda_{n,1}), \dots, \delta_n(\lambda_{n,p}))$. Let $\delta_{n,j} = \delta_n(\lambda_{n,j})$ for $j = 1, \dots, p$. We then have

$$\begin{aligned}\mathbf{L}_{m,n}(\Sigma_n^{-1}, \tilde{\Sigma}_n^{-1}) &= \frac{1}{p} \text{tr}[(\Sigma_n^{-1} - \tilde{\Sigma}_n^{-1})(\Sigma_n^{-1} - \tilde{\Sigma}_n^{-1}) S_n^m] \\ &= \frac{1}{p} \sum_{j=1}^p (u_{n,j}^T \Sigma_n^{-2} u_{n,j}) \lambda_{n,j}^m - \frac{2}{p} \sum_{j=1}^p (u_{n,j}^T \Sigma_n^{-1} u_{n,j}) \frac{\lambda_{n,j}^m}{\delta_{n,j}} + \frac{1}{p} \sum_{j=1}^p \frac{\lambda_{n,j}^m}{\delta_{n,j}^2} \\ &= \int_{-\infty}^{\infty} x^m d\Phi_n^{(-2)}(x) - 2 \int_{-\infty}^{\infty} \frac{x^m}{\delta_n(x)} d\Phi_n^{(-1)}(x) + \int_{-\infty}^{\infty} \frac{x^m}{\delta_n^2(x)} dF_n(x),\end{aligned}$$

where $\Phi_n^{(-l)}(x) = \frac{1}{p} \sum_{j=1}^p (u_{n,j}^T \Sigma_n^{-l} u_{n,j}) \mathbb{1}_{[\lambda_{n,j}, \infty)}(x)$ for $l = 1$ and 2 . Let $\Sigma_n = V_n \Gamma_n V_n^T$ be the spectral decomposition of the true covariance matrix Σ_n , $v_{n,k}$ be the k -th column of V_n , and $\gamma_{n,k}$ be the k -th diagonal element of Γ_n . Note that, for $l = 2$, we have $\Phi_n^{(-2)}(x) = \frac{1}{p} \sum_{i=1}^p \mathbb{1}_{[\lambda_{n,i}, \infty)}(x) \sum_{k=1}^p |u_{n,k}^T v_{n,k}|^2 \gamma_{n,k}^{-2}$. We also let $H_n(x) = p^{-1} \sum_{j=1}^p \mathbb{1}_{\{\gamma_{n,j} \leq x\}}$ be the empirical distribution of population eigenvalues.

We make the following assumptions on the data and population distribution.

Assumption 1. (Dimension) The concentration ratio $p/n \rightarrow c$ as $n \rightarrow \infty$, where $c > 0$.

We consider two scenarios:

- (a) The concentration ratio $c < 1$, and there is a compact interval in $(0, 1)$ that contains p/n for large n .
- (b) The concentration ratio $c > 1$, and there is a compact interval in $(1, \infty)$ that contains p/n for large n .

Assumption 2. (Population)

(a) The population covariance matrix Σ_n is a $p \times p$ nonrandom symmetric positive-definite matrix.

(b) We assume that H_n converges weakly to a limiting spectral distribution H , whose support, denoted by $\text{Supp}(H)$, is a finite union of closed intervals away from zero.

(c) There exists a closed interval in $(0, +\infty)$ that contains $\gamma_{n,1}, \dots, \gamma_{n,p}$ for all large n .

Assumption 3. (Data) The observed matrix $Z_n = W_n \Sigma_n^{1/2}$, where W_n is a $n \times p$ matrix of independent and identically distributed (i.i.d) random variables with mean zero, variance one, and a finite 12th moment.

Following Silverstein and Choi (1995), Silverstein and Bai (1995), and Silverstein (1995), under Assumptions 1(a), 2, 3, we have $F_n(x) \xrightarrow{a.s.} F(x)$ as $n \rightarrow \infty$ for any $x \in \mathbb{R}$, where $F(x)$ is a continuously differentiable limiting spectral distribution function. This in turn implies the strong convergence of $F_n^*(x)$ to $F^*(x)$ where $F^*(x) = 1 - F(1/x)$, $x > 0$. Moreover, by (Bai and Silverstein, 1998, Theorem 1.1), under the same assumptions, there exists a finite number $K \geq 1$ such that $\text{Supp}(F) = \bigcup_{k=1}^K [a_k, b_k]$, where $\text{Supp}(F)$ denotes the support set of F and $0 < a_1 < b_1 < a_2 < b_2 < \dots < a_K < b_K < \infty$. Since H_n also has a limit H , the limiting spectral distribution F_n of sample eigenvalues which is F , is uniquely determined by the concentration ratio and H .

Assumption 4 (Limiting shrinkage function). There exists a nonrandom continuously differentiable function δ defined on $\text{Supp}(F) = \bigcup_{k=1}^K [a_k, b_k]$ such that $\delta_n(x) \xrightarrow{a.s.} \delta(x)$ for $x \in \text{Supp}(F)$ as $n \rightarrow \infty$. In addition, the convergence is uniform for $x \in \bigcup_{k=1}^K [a_k + \eta, b_k - \eta]$ for any small $\eta > 0$. Furthermore, there exists a finite nonrandom constant M such that $|\delta_n(x)|$

is uniformly bounded away from zero by M almost surely for all $x \in \cup_{k=1}^K [a_k - \eta, b_k + \eta]$, large n , and small $\eta > 0$.

These assumptions are also adopted in Ledoit and Wolf (2022) and Ledoit and Wolf (2018). Assumption 4 requires that shrinkage functions δ_n to be well behaved asymptotically. For a bounded, non-decreasing function G , we define its Steiltjes transform by $m_G(z) = \int_{-\infty}^{\infty} (x - z)^{-1} dG(x)$ for $z \in \mathbb{C}^+$, i.e. $z = x + iy$ with some $x \in \mathbb{R}$ and $y \in \mathbb{R}^+$. For a complex valued function $g(z)$, let $\text{Re}[g(z)]$ and $\text{Im}[g(z)]$ denote the real and imaginary parts, respectively.

The following lemma is a key ingredient in finding the almost sure non-random limit of $\mathbf{L}_{m,n}$, where we obtain the limit of $\Phi_n^{(-l)}$ as $n \rightarrow \infty$.

Lemma 1. *Suppose Assumptions 1 (a) or (b), 2, 3 hold. For any integer l , we have that $\Phi_n^{(-l)}$ converges almost surely pointwisely to $\Phi^{(-l)}(x)$ as $n \rightarrow \infty$ for all x such that $\Phi^{(-l)}(x)$ is continuous, where*

$$\begin{aligned}\Phi^{(-l)}(x) &= \lim_{\eta \rightarrow 0^+} \frac{1}{\pi} \int_{-\infty}^x \text{Im}[\Theta^{(-l)}(\xi + i\eta)] d\xi, \text{ and} \\ \Theta^{(-l)}(z) &= \int_{-\infty}^{+\infty} \{\gamma[1 - c^{-1} - c^{-1}zm_F(z)] - z\}^{-1} \gamma^{-l} dH(\gamma).\end{aligned}$$

Proof. For $x \in \mathbb{R}$,

$$\Phi_n^{(-l)}(x) = \frac{1}{p} \sum_{i=1}^p \mathbb{1}_{[\lambda_{n,i}, \infty)}(x) \sum_{j=1}^p \frac{(u_{n,i}^T v_{n,j})^2}{\gamma_{n,j}^l}.$$

By (Ledoit and Péché, 2011, Lemma 6), we have $\Phi_n^{(-l)}(x) \xrightarrow{a.s.} \Phi^{(-l)}(x)$ as $n \rightarrow \infty$. □

Recall the Steiltjes transform $m_G(z)$ defined earlier. When G has derivative G' , $m_G(z)$ has an extension to the real line $\check{m}_G(x) = \lim_{z \in \mathbb{C}^+ \rightarrow x} m_G(z)$ for any $x \in \mathbb{R}$ even though

$\check{m}_G(x)$ could be complex valued. We are now ready to state the almost sure non-random limit of $\mathbf{L}_{m,n}$ and the corresponding minimizer.

Theorem 5. *Under Assumptions 1 (a), 2, and 3, the loss $\mathbf{L}_{m,n}(\Sigma_n^{-1}, \tilde{\Sigma}_n^{-1}) = \text{tr}[(\Sigma_n^{-1} - \tilde{\Sigma}_n^{-1})^2 S_n^m]/p$ has the following almost sure limit:*

$$\begin{aligned}\mathbf{L}_m &= \int x^m d\Phi^{(-2)}(x) - 2 \sum_{k=1}^K \int_{a_k}^{b_k} \frac{x^m}{\delta(x)} d\Phi^{(-1)}(x) + \sum_{k=1}^K \int_{a_k}^{b_k} \frac{x^m}{\delta^2(x)} dF(x) \\ &= \int x^m d\Phi^{(-2)}(x) - 2 \sum_{k=1}^K \int_{a_k}^{b_k} \frac{x^m}{\delta(x)} \phi^{(-1)}(x) dF(x) + \sum_{k=1}^K \int_{a_k}^{b_k} \frac{x^m}{\delta^2(x)} dF(x)\end{aligned}\quad (3.4)$$

as $n \rightarrow \infty$, where

$$\phi^{(-1)}(x) = \begin{cases} 0 & \text{if } x \leq 0 \\ \frac{1 - c - 2cx \text{Re}[\check{m}_F(x)]}{x} & \text{if } x > 0. \end{cases}$$

The proof is provided in the Appendix. We also provide simulations for the convergence of the loss $\mathbf{L}_{1,n}$ in Section S.1 of the supplement.

Since we focus on rotation invariant estimators $\tilde{\Sigma}_n = U_n \tilde{\Delta}_n U_n^T$ with $\tilde{\Delta}_n = \text{diag}(\delta_n(\lambda_{n,1}), \dots, \delta_n(\lambda_{n,p}))$, minimizing the loss $\mathbf{L}_{m,n}(\Sigma_n^{-1}, \tilde{\Sigma}_n^{-1})$ with respect to the estimator $\tilde{\Sigma}_n$ is equivalent to minimizing the loss with respect to the shrinkage function δ_n . Note that the limit of δ_n is δ by Assumption 4, and that the limit of the loss $\mathbf{L}_{m,n}$ is \mathbf{L}_m by Theorem 5. Thus, minimizing the limiting loss \mathbf{L}_m with respect to the shrinkage rule δ could lead to the best rotation invariant estimator for Σ in terms of the limiting loss, which is provided in the following corollary.

Corollary 1. *The limit \mathbf{L}_m in Theorem 5 is minimized at*

$$\delta^*(x) = \frac{1}{\phi^{(-1)}(x)} = \frac{x}{1 - c - 2cx \text{Re}[\check{m}_F(x)]}, \quad \forall x \in \text{Supp}(F).$$

Therefore, the oracle best rotation invariant estimator that is based on the true F and minimizes the limiting loss is $\tilde{\Sigma}_n^* = U_n \Delta^* U_n^T$, where Δ^* is a diagonal matrix with diagonal elements δ^* as defined above. The denominator in the definition of δ^* is bounded away from zero under the same set of assumptions by Ledoit and Wolf (2018, Proposition 3.2). Then it remains to construct a consistent estimator of δ^* .

Theorem 6. *Under the assumptions of Theorem 5, $\delta_n^*(x)$ defined as*

$$\begin{aligned}\delta_n^*(x) &= \left[\left(1 - \frac{p}{n}\right) x^{-1} + \frac{p}{n} x^{-1} 2g_n^*(x^{-1}) \right]^{-1} \text{ with} \\ g_n^*(x) &= \frac{1}{p} \sum_{k=1}^p \lambda_{n,k}^{-1} \frac{\lambda_{n,k}^{-1} - x}{(\lambda_{n,k}^{-1} - x)^2 + h_n^2 \lambda_{n,k}^{-2}},\end{aligned}\tag{3.5}$$

$h_n \sim Cn^{-\alpha}$ for some $C > 0$, and $\alpha \in (0, 2/5)$ is a consistent estimator of $\delta^*(x)$, i.e.

$\delta_n^*(x) \xrightarrow{p} \delta^*(x)$ for any $x \in \text{Supp}(F)$.

We provide the proof in the Appendix. It is instructive to note that the shrinkage rule $\delta_n^*(x)$ is exactly the same as the smoothed Stein shrinker in Ledoit and Wolf (2022) where the authors show that this same rule is the optimum when one considers estimating Σ_n under the Stein loss, i.e. $\mathbf{L}_n(\Sigma_n, \tilde{\Sigma}_n) = p^{-1} \text{tr}(\Sigma_n^{-1} \tilde{\Sigma}_n) - p^{-1} \log |\Sigma_n^{-1} \tilde{\Sigma}_n|$.

Example 1. Consider the case when the data matrix $W_n \stackrel{iid}{\sim} \mathcal{N}(0, 1)$ and $\Sigma_n = \mathbf{I}$. For this case, the distribution of the population eigenvalues is a point mass at 1, i.e. $H(x) = \mathbb{1}(x \geq 1)$, and the limiting shrinkage function is the famous Marcenko-Pastur law:

$$dF(x) = \frac{1}{2\pi cx} \sqrt{(c_+ - x)(x - c_-)} \cdot \mathbb{1}(c_- \leq x \leq c_+),$$

where $c_- = (1 - \sqrt{c})^2$ and $c_+ = (1 + \sqrt{c})^2$. It is well known that its Stieltjes transform $m_F(z) = (2cz)^{-1} (1 - c - z + \sqrt{(c_+ - z)(c_- - z)})$ for $z \in \mathbb{C}^+$, and hence

$$\check{m}_F(x) = \lim_{z \in \mathbb{C}^+ \rightarrow x} m_F(z) = \frac{1 - c - x + \sqrt{(c_+ - x)(c_- - x)}}{2cx}.$$

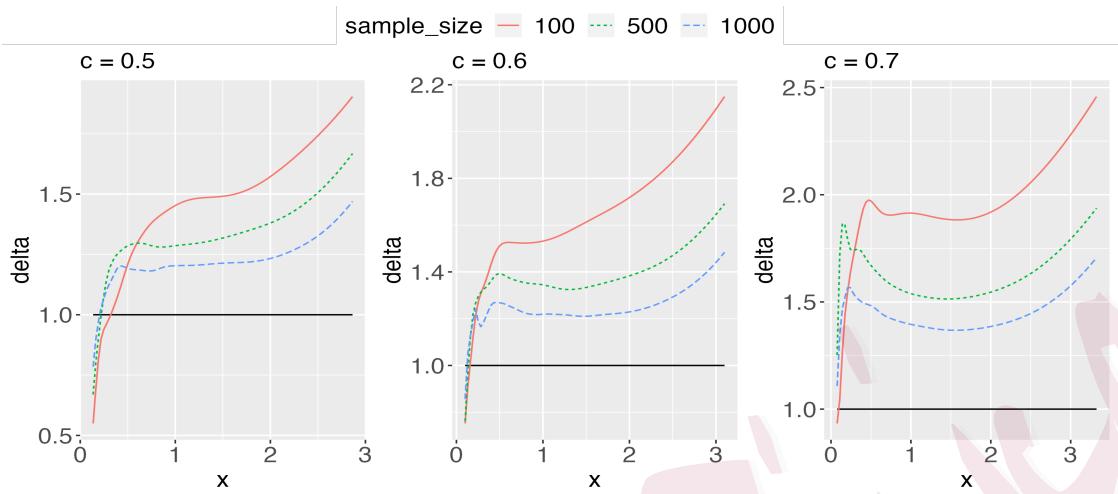


Figure 1: The population ($\delta^*(x)$) and sample shrinkage rule ($\delta(x)$) when the sample covariance matrix $S = n^{-1}XX^T$ has entries $X_{ij} \stackrel{iid}{\sim} N(0, 1)$. Three concentration ratios are considered $c = 0.5, 0.6, 0.7$. The black line corresponds to $\delta^*(x)$. The red, blue and chocolate lines correspond to $n = 100, 500, 1000$, respectively.

Thus, $\text{Re}[\check{m}_F(x)] = (1 - c - x)/2cx$ for $x \in \text{Supp}(F)$, which implies that $\delta^*(x) = 1$.

In Figure 1, we plot $\delta^*(\cdot)$ and its estimator $\delta(x)$ (a realization based on X_n) for three choices of the concentration ratio $c = 0.5, 0.6, 0.7$ and $n = 100, 500, 1000$. Here, the bandwidth parameter h is chosen to obtain smooth decision rules (large h). This parameter plays a critical role in that smaller values lead to fluctuating decision rules. Additionally, for each panel, h decreases with n to ensure the condition of Theorem 6 is satisfied. It can be seen from Figure 1, that the estimator δ^* converges to δ as n grows.

The shrinkage rule $\delta_n^*(x)$ has a bandwidth parameter h_n which approaches 0 as $n \rightarrow \infty$ according to Theorem 6. Ledoit and Wolf (2022) studied the dependence of the tuning parameter h_n on the concentration ratio $c = p/n$ through simulation experiments and suggested a default choice, which is $h_n = (c)^{0.7}(1/p)^{0.35}$. Although this choice works well in

simulations, it is not data-driven.

In the following subsection, we provide a data-dependent choice of h_n based on the estimation of the expected loss. While the results in this section till this point do not require any assumption about the distribution of Z_n other than the existence of its moments, in the following subsection we assume $Z_n \sim N(0, \Sigma_n)$.

3.1 (Nearly) Unbiased estimate of the risk

In this subsection, we develop a nearly unbiased estimate of the expected loss when $m = 1$, and find the tuning parameter h_n through minimizing the estimated expected loss. Suppose Assumptions 1 (a), 2, and 4 hold and in Assumption 3 elements of W_n are i.i.d $N(0, 1)$. Consider the estimator $\tilde{\Sigma}_n = U_n \tilde{\Delta}_n U_n^T$ where $\tilde{\Delta}_n$ is a diagonal matrix with diagonal elements $\delta_n^*(x)$ as defined in Theorem 5. The shrinkage rule $\delta_n^*(x)$ clearly depends on h_n . Hence, we write $\delta_n^*(x) = \delta_n^*(x; h_n)$ and the resulting estimator $\tilde{\Sigma}_n = \tilde{\Sigma}_n(h_n)$.

Our goal here is to propose a nearly unbiased estimate of the risk $\mathbb{E}_\Sigma[\mathbf{L}_{m,n}(\Sigma_n^{-1}, \tilde{\Sigma}_n^{-1}(h_n))]$ for any h_n and finite n when $m = 1$. Here, the expectation is computed with respect to the distribution of Z_n which depends on the unknown Σ_n . Since here we are concerned with finite n , we suppress the dependence on n and write $\Sigma_n = \Sigma$, $U_n = U$, $\tilde{\Delta}_n = \tilde{\Delta}$, $\delta_n^*(x) = \delta^*(x, h)$, $\lambda_{n,j} = \lambda_j$, $Z_n = Z$, $h_n = h$, $\tilde{\Sigma}_n = \tilde{\Sigma}(h)$ and $\mathbf{L}_{1,n} = \mathbf{L}(h) = \text{tr}[\{\Sigma^{-1} - \tilde{\Sigma}(h)^{-1}\}^2 S]$ where $S = n^{-1}Z^T Z$. The corresponding risk is $\mathbf{R}(\Sigma; h) = \mathbb{E}_\Sigma\{L(h)\} = n^{-1}\mathbb{E}_\Sigma\text{tr}[\{\Sigma^{-1} - \tilde{\Sigma}(h)^{-1}\}^2 (Z^T Z)]$. When $Z \sim N(0, \Sigma)$, we have $Z^T Z \sim W_p(n, \Sigma)$ where $W_p(n, \Sigma)$ is the Wishart distribution of dimension p , degrees of freedom n and parameter Σ .

Let $\underline{S} = Z^T Z$. Consider the spectral decomposition of $S = U \Lambda U^T$ which implies that the spectral decomposition of $\underline{S} = nS = U(n\Lambda)U^T = U\Lambda^*U^T$ with $\lambda_j^* = n\lambda_j$. Since

$\tilde{\Sigma}(h) = U\tilde{\Delta}(h)U^T$ with $\tilde{\Delta}(h) = \text{diag}(\delta^*(\lambda_1, h), \dots, \delta^*(\lambda_p, h))$, we have

$$\begin{aligned}\mathbf{R}(\Sigma; h) &= n^{-1}\mathbb{E}_\Sigma \left[\text{tr} \left(\tilde{\Sigma}(h)^{-2} \underline{S} \right) - 2\text{tr} \left(\Sigma^{-1} \underline{S} \tilde{\Sigma}(h)^{-1} \right) + \text{tr} \left(\Sigma^{-2} \underline{S} \right) \right] \\ &= n^{-1} \sum_{j=1}^p \mathbb{E}_\Sigma \left[\frac{\lambda_j^*}{\delta^*(\lambda_j, h)^2} \right] - 2n^{-1}\mathbb{E}_\Sigma \left[\text{tr} \left(\Sigma^{-1} \underline{S} \tilde{\Sigma}(h)^{-1} \right) \right] \\ &\quad + n^{-1}\mathbb{E}_\Sigma \left[\text{tr} \left(\Sigma^{-2} \underline{S} \right) \right].\end{aligned}\tag{3.6}$$

In the above display, the first term inside the expectation is computable for any observed S and given h , and it is trivially an unbiased estimate of its own expectation. We next focus on the third term $\mathbb{E}_\Sigma [\text{tr}(\Sigma^{-2} \underline{S})] = \mathbb{E}_\Sigma [\text{tr}(\Sigma^{-1} \underline{S} \Sigma^{-1})] = \mathbb{E}_\Sigma [\text{tr}(\underline{S}^*)]$ where $\underline{S}^* = \Sigma^{-1} \underline{S} \Sigma^{-1} \sim W_p(n, \Sigma^{-1})$. Let $\underline{\Sigma} = \Sigma^{-1}$. Then from properties of the Wishart distribution it is immediate that $\mathbb{E}_\Sigma [\text{tr}(\underline{S}^*)] = n \sum_{j=1}^p \underline{\Sigma}_{jj}$, where $\underline{\Sigma}_{jj}$ is the j -th diagonal element in $\underline{\Sigma}$. Since in the current setup $Z_i \sim N(0, \Sigma)$, $i = 1, \dots, n$, we have that $Z_{ij} \mid Z_{i,-j} \sim N(\alpha_j^T Z_{i,-j}, \underline{\Sigma}_{jj}^{-1})$ where $Z_{i,-j}$ is the vector obtained by excluding the j -th element from Z_i . Thus, an unbiased estimator of $\underline{\Sigma}_{jj}^{-1}$ is available from the error variance estimate obtained by linearly regressing the j -th column of Z on the rest of its columns. Inverting this estimator gives us a reasonably good estimator of $\underline{\Sigma}_{jj}$, denoted by $\hat{\underline{\Sigma}}_{jj}$, although strict unbiasedness can be maintained by employing the so-called sum-estimator (Glynn and Rhee, 2014). However, this sum-estimator might result in increased computational burden, and in our numerical experiments we observed the biased estimator to perform better than the unbiased estimator.

The key issue of estimating the risk is the second term in Equation (3.6), which is dealt with in the following theorem. The result is an application of Haff's identity (Haff, 1979) suited to our context. The proof is given in the Appendix.

Theorem 7. *Under Assumptions 1 (a), 2, 3, and when elements of W_n are distributed*

$N(0, 1)$ independently, we have

$$\mathbb{E}_{\Sigma} \left[\text{tr} \left(\Sigma^{-1} \underline{S} \tilde{\Sigma}(h)^{-1} \right) \right] = \mathbb{E}_{\Sigma} \left[\text{tr} \left\{ (n - p - 1) \tilde{\Sigma}(h)^{-1} + 2D_{\underline{S}}([\tilde{\Sigma}(h)^{-1}]^T \underline{S}) \right\} \right],$$

where $D_{\underline{S}}$ is a differential operator defined as $D_{\underline{S}} = \left\{ \frac{1}{2}(1 + d_{ij}) \frac{\partial}{\partial \underline{S}_{ij}} \right\}$ for $1 \leq i, j \leq p$ with $d_{ij} = 1$ if $i = j$ and 0 otherwise.

Therefore, we construct the following nearly unbiased estimate of the risk of the estimator $\tilde{\Sigma}(h)$:

$$\begin{aligned} \hat{R}(h) = & n^{-1} \sum_{j=1}^p \left[\frac{\lambda_j^*}{\delta^*(\lambda_j, h)^2} \right] - 2n^{-1} \left[\text{tr} \left\{ (n - p - 1) \tilde{\Sigma}(h)^{-1} \right\} \right] \\ & - 4n^{-1} \left[\text{tr} \left\{ D_{\underline{S}}([\tilde{\Sigma}(h)^{-1}]^T \underline{S}) \right\} \right] + \sum_{j=1}^p \hat{\Sigma}_{jj}. \end{aligned}$$

The proposed value \hat{h} for the tuning parameter h is the minimizer of $\hat{R}(h)$.

3.2 Shrinkage in the $p > n$ case

In this subsection, we extend our results for the relative savings loss to the $p > n$ case, where the sample covariance matrix is singular. The overall strategy remains the same. That is, we find an almost sure limit of the loss, find its minimizer, and then construct a consistent estimator of the minimizer. In this case, the limiting spectral distribution of the sample eigenvalues F is a mixture of a point mass at 0 and a continuous component, i.e. $F(x) = \{(c - 1)/c\} \mathbb{I}_{[0, \infty]}(x) + (1/c) \underline{F}(x)$ where $\underline{F}(x)$ is the continuous component whose support is bounded away from 0. The Steiltjes transform of F and \underline{F} are related as $\check{m}_{\underline{F}}(x) = c \check{m}_F(x) + (c - 1)/x$ for all $x \in \mathbb{R}$. From (Ledoit and Wolf, 2018, Lemma 14.1) we have $\Phi_n^{(-1)}(x)$ converges almost surely to $\Phi^{(-1)(x)}$ for all $x \in \mathbb{R} - \{0\}$. And, $\Phi^{(-1)}$ is continuously differentiable on $\mathbb{R} - \{0\}$ and can be expressed for all $x \in \mathbb{R}$ as

$\Phi^{(-1)}(x) = \int_{-\infty}^x \phi^{(-1)}(u) dF(u)$ where

$$\phi^{(-1)}(x) = \begin{cases} 0 & \text{if } x < 0 \\ \frac{c}{c-1} \check{m}_H(0) - \check{m}_{\underline{F}}(0) & \text{if } x = 0 \\ \frac{1 - c - 2cx \operatorname{Re}[\check{m}_F(x)]}{x} & \text{if } x > 0. \end{cases}$$

We then have the following limit of the relative savings loss, noting that in this case $dF(x) = (1/c)d\underline{F}(x)$ for $x > 0$.

Theorem 8. *Under Assumptions 1 (b), 2, 3 and in the case when $c > 1$, $\mathbf{L}_{m,n} = \operatorname{tr}[(\Sigma_n^{-1} - \tilde{\Sigma}_n^{-1})^2 S_n^m]/p$ has the following almost sure limit:*

$$\mathbf{L}_m = \begin{cases} \int_{-\infty}^{\infty} x^m d\Phi^{(-2)}(x) - 2 \sum_{k=1}^K \int_{a_k}^{b_k} \frac{x^m}{\delta(x)} \phi^{(-1)}(x) dF(x) \\ + \sum_{k=1}^K \int_{a_k}^{b_k} \frac{x^m}{\delta^2(x)} dF(x), \quad m \geq 1 \\ \int_{-\infty}^{\infty} d\Phi^{(-2)}(x) - 2 \sum_{k=1}^K \int_{a_k}^{b_k} \frac{1}{\delta(x)} \phi^{(-1)}(x) dF(x) \\ + \sum_{k=1}^K \int_{a_k}^{b_k} \frac{1}{\delta^2(x)} dF(x) + \frac{c-1}{c} \left[\frac{1}{\delta^2(0)} - 2 \frac{\frac{c}{c-1} \check{m}_H(0) - \check{m}_{\underline{F}}(0)}{\delta(0)} \right], \quad m = 0. \end{cases}$$

The above limit immediately yields the following oracle decision rule, specifically for the case $m \geq 1$.

Corollary 2. *For $m \geq 1$, the limit \mathbf{L}_m in Theorem 8 is minimized at $\delta^*(x) = \frac{1}{\phi^{(-1)}(x)}$ for any $x > 0$.*

The next result gives us a consistent estimator of this optimal shrinkage rule.

Theorem 9. *Under the assumptions of Theorem 8, $\delta_n^*(x)$ for $x > 0$, defined as*

$$\delta_n^*(x) = \left[\left(\frac{p}{n} - 1 \right) x^{-1} + 2x^{-1} g_n^*(x^{-1}) \right]^{-1},$$

$$g_n^*(x) = \frac{1}{n} \sum_{k=p-n+1}^p \lambda_{n,k}^{-1} \frac{\lambda_{n,k}^{-1} - x}{(\lambda_{n,k}^{-1} - x)^2 + h_n^2 \lambda_{n,k}^{-2}},$$

with $h_n \sim Cn^{-\alpha}$ for some $C > 0$ and $\alpha \in (0, 2/5)$ is a consistent estimator of $\delta^*(x)$, i.e. $\delta_n^*(x) \xrightarrow{p} \delta^*(x)$ for any $x \in \text{Supp}(\underline{F})$.

The above two results imply that, when $m \geq 1$, the minimizer does not rely on the spectral distribution H and the shrinkage rule is only defined for the non-zero eigenvalues of the sample covariance matrix. This is due to the special nature of the relative savings loss for the case $m \geq 1$ as shown in Theorem 8. Indeed, when $p > n$, the sample covariance matrix S_n has n non-zero eigenvalues and $p - n$ zero eigenvalues. This implies for $m \geq 1$, and any $\tilde{\Sigma}_n = U_n \tilde{\Delta}_n U_n^T$, we have

$$\mathbf{L}_{m,n}(\Sigma_n^{-1}, \tilde{\Sigma}_n^{-1}) = \frac{1}{p} \sum_{j=1}^n (u_{n,j}^T \Sigma_n^{-2} u_{n,j}) \lambda_{n,j}^m - \frac{2}{p} \sum_{j=1}^n (u_{n,j}^T \Sigma_n^{-1} u_{n,j}) \frac{\lambda_{n,j}^m}{\delta_n(\lambda_{n,j})} + \frac{1}{p} \sum_{j=1}^n \frac{\lambda_{n,j}^m}{\delta_n^2(\lambda_{n,j})}.$$

In other words, the zero sample eigenvalues do not contribute to the loss. When $m = 0$, the oracle decision rule may involve the population eigenvalues according to Theorem 8 through the term $\check{m}_H(0)$, which can be estimated by inversion of the QuEst function (Ledoit and Wolf, 2018). In fact, Ledoit and Wolf (2022) and Ledoit and Wolf (2018) have considered this covariance estimation problem for $p > n$ under other loss functions, among which some loss functions such as Stein's loss also needs numerical inversion of the QuEst function.

For our purpose of estimating the parameter matrix B in Section 2, we are specifically interested in the case $m = 1$. Although decision rules for null eigenvalues do not affect the loss, for our primary objective of estimating B by $\widehat{B}(1 - \widehat{\Sigma}^{-1}Q)$, we do need $\widehat{\Sigma}^{-1}$, which would not be possible to compute with zero eigenvalues. For this reason, we consider the following shrinkage rule for zero eigenvalues:

$$\delta_n^*(x)^{-1} = \left(\frac{p}{n} - 1 \right) \frac{1}{n} \sum_{j=p-n+1}^p \lambda_{n,j}^{-1} \quad \text{for } x = 0, \quad (3.7)$$

which is the optimal shrinkage rule for estimation under Frobenius loss (Ledoit and Wolf, 2022, Section 5).

4 Coefficient matrix B estimation

We now again go back to our original problem of estimating the coefficient matrix B .

According to the results in Section 3, we propose to use

$$\bar{B} = \hat{B}(\mathbf{I} - Q^{1/2}\bar{\Sigma}_n Q^{1/2}) \quad (4.8)$$

with $\bar{\Sigma}_n = U_n \text{diag}(\delta_n^*(\lambda_{n,1}), \dots, \delta_n^*(\lambda_{n,p})) U_n^T$ to estimate B , where $\delta_n^*(x)$ is defined in Theorem 6 with $h_n = \hat{h}$ as defined in Section 3.2 if $p < n$ and is defined in Theorem 9 if $p > n$.

This estimator represents a global shrinkage rule on the data matrix \hat{B} .

In this section, we extend this global shrinkage estimator of $\beta^{(t)}$ to a local one with a mixture prior distribution. Scale mixtures of Gaussian (Polson and Scott, 2010) are the cornerstone of modern Bayesian sparse modeling techniques. The two component spike-slab prior is also a mixture of Gaussian random variables with appropriately small variance (Narisetty and He, 2014). These priors are obtained by mixing a Gaussian (typically with mean 0) with a scale parameter λ and some mixing distribution g over λ . Careful choices of g lead to flexible prior distributions with sufficiently heavier tails than a Normal distribution. Posterior means from these priors apply different levels of shrinkage depending on the magnitude of the signal. In other words, the shrinkage function is not global.

In particular, we consider the prior $\beta^{(t)} \sim \sum_{k=1}^K \pi_k N(0, Q^{1/2} \Omega_k Q^{1/2})$, so that the posterior mean is a weighted combination of linear shrinkage rules. Indeed, when $\beta^{(t)} \sim$

$\sum_{k=1}^K \pi_k N(0, Q^{1/2} \Omega_k Q^{1/2})$, we have

$$\begin{aligned} \mathbb{E}(\beta^{(t)} | \widehat{B}) &= (\mathbf{I} - \sum_{k=1}^K \pi_k^* C_k) \widehat{\beta}^{(t)}, \quad C_k = Q^{1/2} \Sigma_k^{-1} Q^{-1/2}, \Sigma_k = \mathbf{I} + \Omega_k \\ \pi_k^* &= \frac{\pi_k f(\widehat{\beta}^{(t)}; 0, Q^{1/2} \Sigma_k Q^{1/2})}{\sum_{k=1}^K \pi_k f(\widehat{\beta}^{(t)}; 0, Q^{1/2} \Sigma_k Q^{1/2})}, \end{aligned}$$

where $f(x; \mu, \Sigma)$ is the multivariate Gaussian density evaluated at x with mean μ and covariance Σ . Moreover, the marginal distribution of the observed data is a mixture of Gaussian as well: $\widehat{\beta}^{(t)} \sim \sum_{k=1}^K \pi_k^* N(0, Q^{1/2} \Sigma_k Q^{1/2})$. The posterior distribution of $\beta^{(t)}$ is also available in closed form: $\beta^{(t)} | \widehat{\beta}^{(t)} \sim \sum_{k=1}^K \pi_k^* N\{(\mathbf{I} - C_k) \widehat{\beta}^{(t)}, (\mathbf{I} - C_k) Q\}$. Note that more weight is given to component-specific posterior means that are supported by the data through the posterior probabilities.

As in Section 2, we consider the case when the prior parameters Ω_k are unknown. In order to obtain the posterior mean $\mathbb{E}(\beta^{(t)} | \widehat{B})$, we need estimates of Σ_k . Recall that $\widehat{\beta}_*^{(t)} = Q^{-1/2} \widehat{\beta}^{(t)}$. To estimate Σ_k , we introduce latent indicators $Z \in \{0, 1\}^{K \times 1}$ such that $\widehat{\beta}_*^{(t)} | Z_k = 1 \sim N(0, \Sigma_k)$. Thus, given these indicators, estimating Σ_k can be achieved by resorting to covariance matrix estimates obtained in Section 3. Since the indicators are unknown, we use a Markov chain Monte Carlo (MCMC) sampler to sample $Z | \Sigma_k$, and then estimate Σ_k conditional on the sampled indicators Z - the standard data augmentation technique to fit Bayesian mixture models (Hobert, 2011). The resulting algorithm is summarized in Algorithm S.9 of the supplement. Regarding the choice of K in practice, we suggest dividing the data into training and validation subsets, and determining the optimal K by comparing the prediction errors of the proposed method on the validation set for various candidate values of K .

In addition, we provide a discussion of the risk of the proposed global and local shrinkage

rules in Section S.2 of the supplement, which quantifies the potential risk reduction for the proposed estimators compared to the maximum likelihood estimator.

5 Simulation experiments

In this section, we conduct a series of extensive simulation experiments, and compare the proposed method with existing methods in the literature. We focus on the performance of the linear shrinkage estimator from Section 4 in estimating the parameter matrix B . R codes to implement our method are available at <https://github.com/antik015/EBayes-Integration>.

To benchmark the performance of the proposed method, we compare it to the following list of methods - 1) Unified Test for MOlecular SignaTures (UTMOST) developed by Hu et al. (2019) for cross-tissue TWAS models where the estimator is obtained by optimizing an error sum of squares for all tissues with added ℓ_1 -penalties on the columns of the mean matrix, and ℓ_2 -penalties on the rows of the mean matrix, 2) Iterated stable autoencoder (ISA) developed by Josse and Wager (2016) where the estimator is obtained by creating an autoencoder of the covariate matrix using a careful bootstrap scheme, and 3) Multivariate Adaptive Shrinkage (MASH) developed by Kim et al. (2024) where continuous shrinkage priors are approximated by finite mixtures of Gaussian distributions, and variational methods for scalable computation are combined to estimate the mean matrix. The first two methods make a working structural assumption on the parameter to be estimated, in that it is either sparse or low-rank. The last method is under the empirical Bayes framework.

We consider $N = 200$ points within each source of data (i.e., tissue in the context of TWAS), and $n = 40$ or 50 as the number of data sources. Within each source, we vary the number of predictors to be $p = 10, 20$, or 30 when $n = 40$, and $p = 20, 30$, or 40 when

$n = 50$. We provide additional simulations for high-dimensional settings in Section S.4 of the supplement.

Each row of the matrix of covariates X for all combinations was generated independently from $N(0, \Sigma_X)$ where $\Sigma_X = (1 - \rho)I_p + \rho\mathbf{J}_p$ where \mathbf{J}_p is a $p \times p$ matrix with all entries equal to 1. We consider $\rho = 0, 0.5$, or 0.8 in our experiments. The true parameter matrix $B_0 \in \mathbb{R}^{p \times n}$ with columns $\beta_0^{(t)}$, $t = 1, \dots, n$ were generated using four specific settings - a) low-rank (LR) where $B_0 = F_0 G_0^T$, and $F_0 \in \mathbb{R}^{p \times r}$, $G_0 \in \mathbb{R}^{n \times r}$ with $r = 8$, b) approximate sparse (AS) where each entry of B_0 is generated independently from $N(0, \tau_0^2)$ with $\tau_0 = 0.2$, so that $\mathbb{E}[\|B_0\|_F^2] = 0.2^2 np$, c) Horseshoe (HS) where elements of $\beta_0^{(t)}$ are generated from the Horseshoe prior (Carvalho et al., 2010) by first generating $\tau \sim \text{Unif}[0, 1]$, then $\lambda_j \sim C^+(0, 1)$, and $\beta_{0j}^{(t)} \mid \tau, \lambda_j \sim N(0, \lambda_j^2 \tau^2)$ where $C^+(0, 1)$ is the standard Half-Cauchy distribution with pdf $f(x) \propto (1 + x^2)^{-1}$, d) mixture (Mix) where $\beta_0^{(t)} \sim \pi_1 N(0, A_1) + \pi_2 N(0, A_2)$ with $A_1 = 0.1^2 I_p$, and $A_2 = 10^2 I_p$. Each of the design settings considered here focuses on specific cases. For example, the low-rank setting is favorable to methods that are designed to recover parameters with this structure, the Horseshoe setting is designed to have a mix of signal strengths within each column of the true coefficient matrix. On the other hand, in the approximate sparse setting all elements of B_0 are very small and methods that penalize the size of the coefficient matrix are expected to perform well, such as UTMOST. Finally, the mixture setting generates coefficients that have strong signals for some data sources and weak signals for remaining data sources, a typical scenario in TWAS.

We then generate the response within each source as $y^{(t)} = X\beta_0^{(t)} + \epsilon^{(t)}$ where $\epsilon^{(t)} \sim N(0, 1)$. For every method, we compute the mean square error (MSE) computed as $(pn)^{-1} \|\tilde{B} - B_0\|_F^2$ and predictive mean squared error (PE) computed as $(nN_t)^{-1} \|X_t \tilde{B} - X_t B_0\|_F^2$, av-

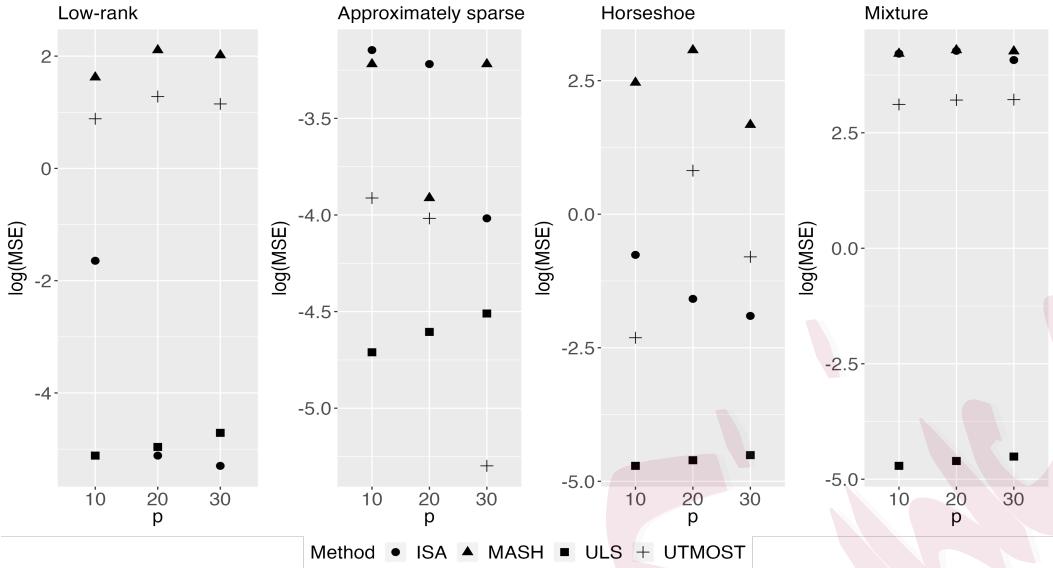
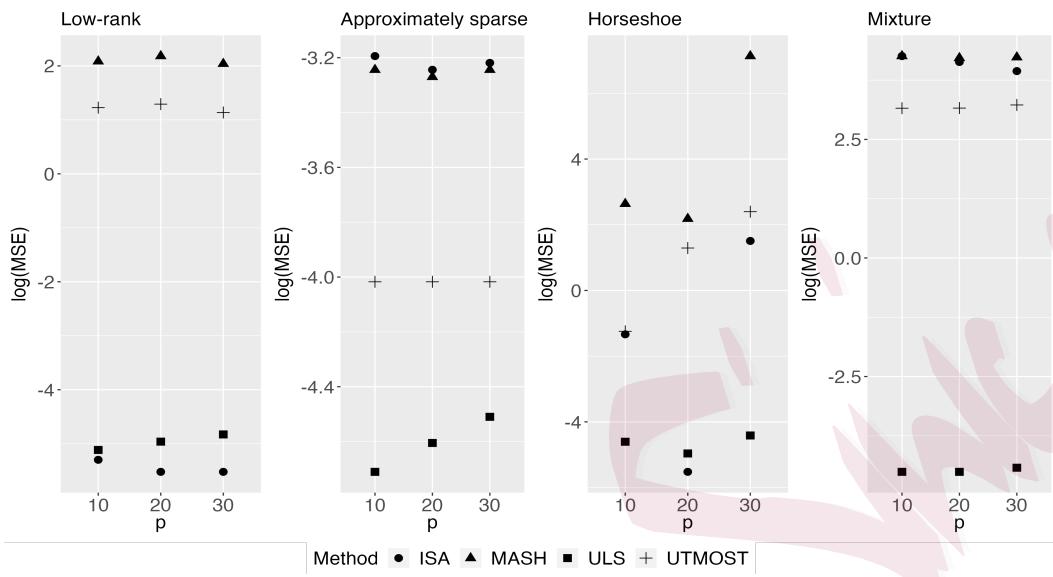
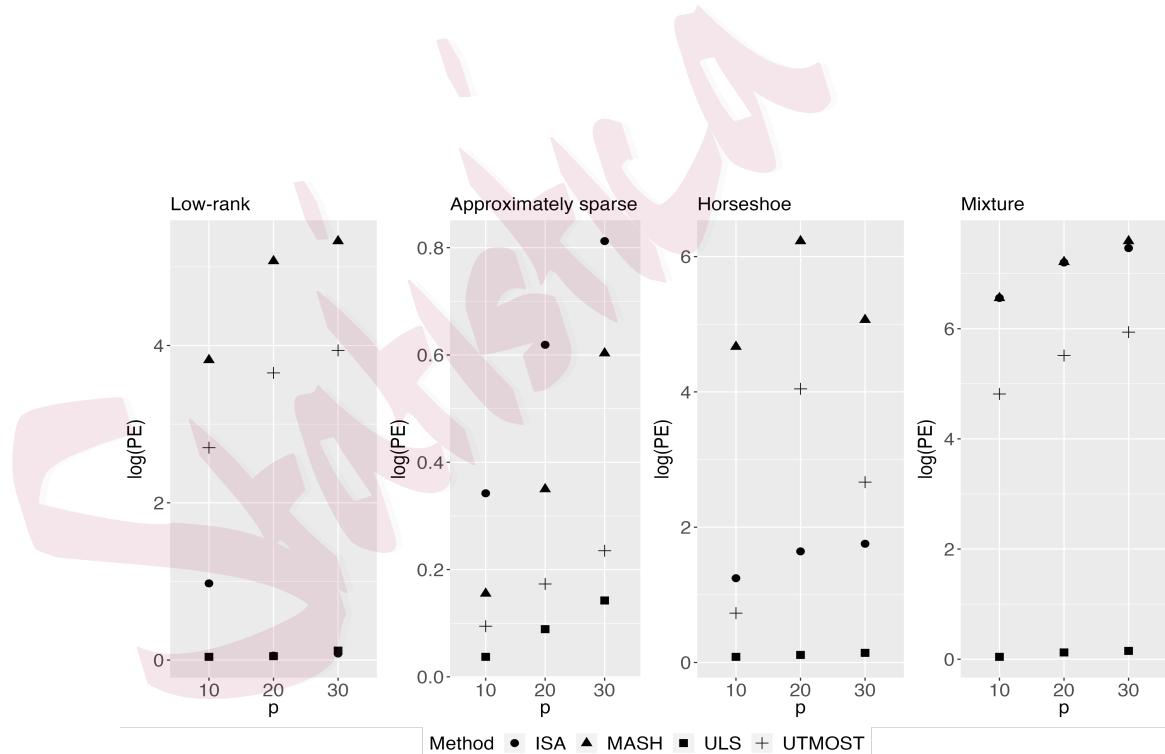


Figure 2: Logarithm of MSE when $n = 40$ and $\rho = 0.5$ under the four experimental settings.

eraged over 20 independent replications within each setting, where \tilde{B} is an estimator for B_0 . Here X_t is a $N_t \times p$ matrix of testing data on the covariates; we set $N_t = 20$. For the proposed local linear shrinkage estimator, we consider three choices of $K = 2, 3, 4$ abbreviated as LLS-2, LLS-3, LLS-4. The proposed global linear shrinkage estimator defined in Equation (4.8) is abbreviated as ULS for unbiased linear shrinkage. We also consider the proposed global linear shrinkage estimator with the default smoothing parameter suggested by Ledoit and Wolf (2022) as the h_n , which is abbreviated as LS.

The detailed results for all the methods are reported in Tables S.6–S.9, which can be found in the Appendix. In particular, we compare the proposed ULS method with existing methods in Figures 2–5. Naturally, when the true parameter matrix has a low-rank structure, the best performing method is ISA, although errors for ULS are very close as shown in the figures. In other cases, ULS is the best performing method, especially when it comes to MSE. As shown in Tables S.6–S.9, other proposed estimators also maintains

Figure 3: Logarithm of MSE when $n = 50$ and $\rho = 0.5$ under the four experimental settings.Figure 4: Logarithm of PE when $n = 40$ and $\rho = 0.5$ under the four experimental settings.

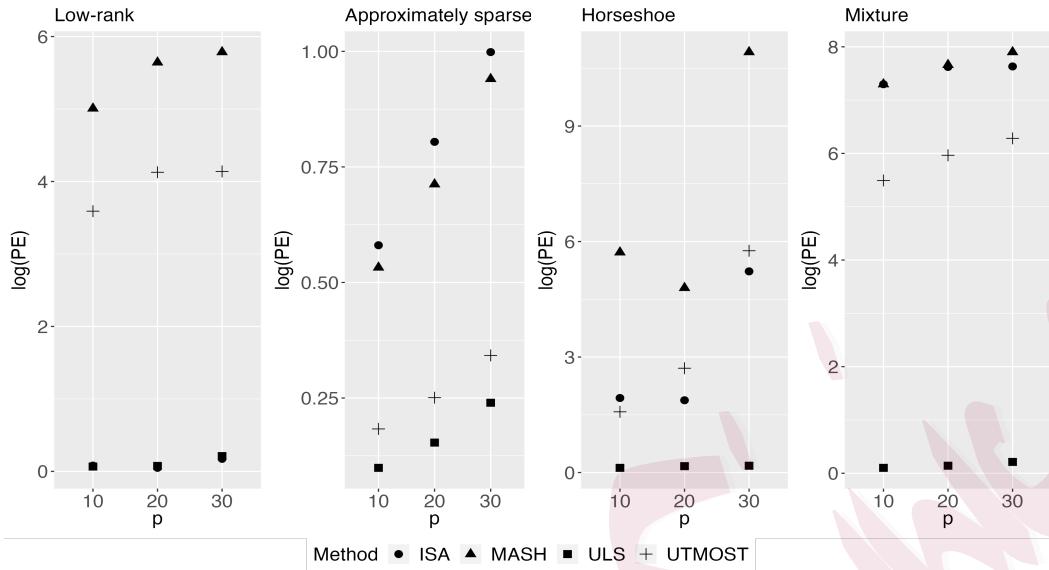


Figure 5: Logarithm of PE when $n = 50$ and $\rho = 0.5$ under the four experimental settings.

good performance across different settings. When the true parameter is generated from a two-component mixture, errors of LLS with $K = 3, 4$ are very close to errors of $K = 2$. In addition, we expect the level of error of the proposed method to increase as p grows, since Figure S.1 in the supplement shows that the value of the loss function $\mathbf{L}_{1,n}$ increases as $c = \lim_{n \rightarrow \infty} p/n$ increases for a fixed n .

We observe a lot of variation of errors for the Horseshoe case. The Horseshoe prior is a heavy-tailed prior (polynomially decreasing tails) with an infinite spike at 0. As a result, samples from this distribution can be very small and very large. The MASH/ISA/UTMOST algorithm is particularly sensitive to these kind of parameter settings and hence the large variations. In particular, the variation comes from the fact that in some of the replications, these algorithms did not converge.

Overall, the results clearly point towards the versatility of the proposed shrinkage method, and highlights the benefits of the approach in that when the underlying struc-

ture of the parameter is unknown, it is perhaps better to shrink. This approach may not always provide the best performance. However, one can still achieve reasonable performance in a variety of settings.

6 Real data application

In this section, we apply the proposed method to the GTEx data and compare it with the existing OLS, UTMOST, ISA, and MASH methods in predicting tissue-specific expressions using cis-SNPs and estimating the corresponding effects.

We aim to study the relationship between the cis-SNPs and gene expression across various tissues and genes. The GTEx project, initiated in 2010 as part of the NIH Common Fund, provides a comprehensive public resource database to study the association between genetic variation and gene expression. This project has collected genotype and gene expression data from 838 participants in 49 tissue types, which are extracted from tissue samples by the Laboratory, Data Analysis and Coordinating Center (LDACC) (Lonsdale et al., 2013). However, we do not have tissue samples of all the 49 tissue types from each participant. In this paper, we focus only on 32 tissues, each of which contains at least 200 samples.

We obtain gene expression values from “GTEx_Analysis_v8_eQTL_expression_matrices.tar” at <https://www.gtexportal.org/home/downloads/adult-gtex/qlt>. These values have been fully processed and normalized by the GTEx project (Consortium, 2020; Lonsdale et al., 2013). The genotype data in “GTEx_Analysis_2017-06-05_v8_WholeGenomeSeq_838_Indiv_Analysis_Freeze.SHAPEIT2phased.vcf.g” are also used and processed as outlined below. The SNPs with minor allele frequencies less than 5% are excluded. The remaining

SNPs are filtered for linkage disequilibrium (LD) using PLINK 1.9 with a window size of 50 SNPs, a step size of 5 SNPs, and an R^2 threshold of 0.2. Furthermore, we identify cis-eQTLs for each gene based on the approach in Wang et al. (2016). As a result, 3,201 genes have at least two and at most twenty associated cis-SNPs.

We adopt the proposed LS method, along with the OLS, UTMOST, ISA, and MASH methods, to predict gene expression based on the corresponding cis-SNPs for each gene. Here we only use the LS method since other proposed methods perform similarly as the LS based on the simulations in Section 5. To assess the prediction performance of each method, we conduct a 10-fold cross-validation analysis for each gene. Specifically, for each tissue, we randomly split all the observed samples into 10 equally sized folds, and name them Folds 1–10. For each $i = 1, \dots, 10$, we treat Fold i in all the tissues as a testing set, and the remaining folds in all the tissues together as a training set. We then use the average prediction mean squared error (PMSE) across all the folds and genes for the evaluation of prediction accuracy.

We provide the average PMSE across all the 3,201 genes for each method in the first row of Table 1, which shows that the proposed method, UTMOST, and MASH perform slightly better than the OLS method. In the second and third rows of Table 1, we provide the average PMSEs across 2,736 and 2,019 genes which have at most 10 and 5 cis-SNPs, respectively, showing that the proposed method and the UTMOST perform the best among all the methods. The MASH performs worse than these two methods when there are fewer cis-SNPs.

To further compare all the methods, we plot heat maps of estimated cis-SNP effects by all the methods for the ZNF138 gene, which is shown in Figure 6. In sub-figures of Figure 6,

each row of the heat map represents a tissue and each column represents a cis-SNP. We can observe that the MASH and UTMOST methods tend to shrink SNP effects to zero, while the proposed method allows non-zero small effects. Moreover, the MASH method is likely to fuse effects from different tissues to the same value, while the proposed method enable variability of effects across tissues. In fact, there could exist considerable differences among the cis-SNP effects in different tissue types (Fu et al., 2012). Furthermore, gene expression could be complex and controlled by many SNPs with small effects (Heap et al., 2009; Gresle et al., 2020; Lloyd-Jones et al., 2017; Boyle et al., 2017).

In addition, we have applied the proposed LS method to the Yeast Cell Cycle dataset and compared that with existing methods. The results are provided in Section S.5 of the supplement. The proposed method works well in the multi-response regression for the Yeast Cell Cycle dataset.

Table 1: Average prediction mean squared errors (PMSEs) for different methods, with standard deviation (SD) in the parentheses. “PMSE” represents the average PMSE across all the 3,201 genes. “PMSE₁₀” and “PMSE₅” represent average PMSEs for genes with at most 10 and 5 cis-SNPs.

Methods	Proposed (LS)	OLS	UTMOST	ISA	MASH
PMSE	0.964 (0.018)	0.968 (0.017)	0.962 (0.016)	0.982 (0.023)	0.964 (0.024)
PMSE ₁₀	0.966 (0.011)	0.970 (0.011)	0.966 (0.010)	0.987 (0.015)	0.970 (0.015)
PMSE ₅	0.969 (0.007)	0.972 (0.007)	0.969 (0.006)	0.988 (0.010)	0.974 (0.011)

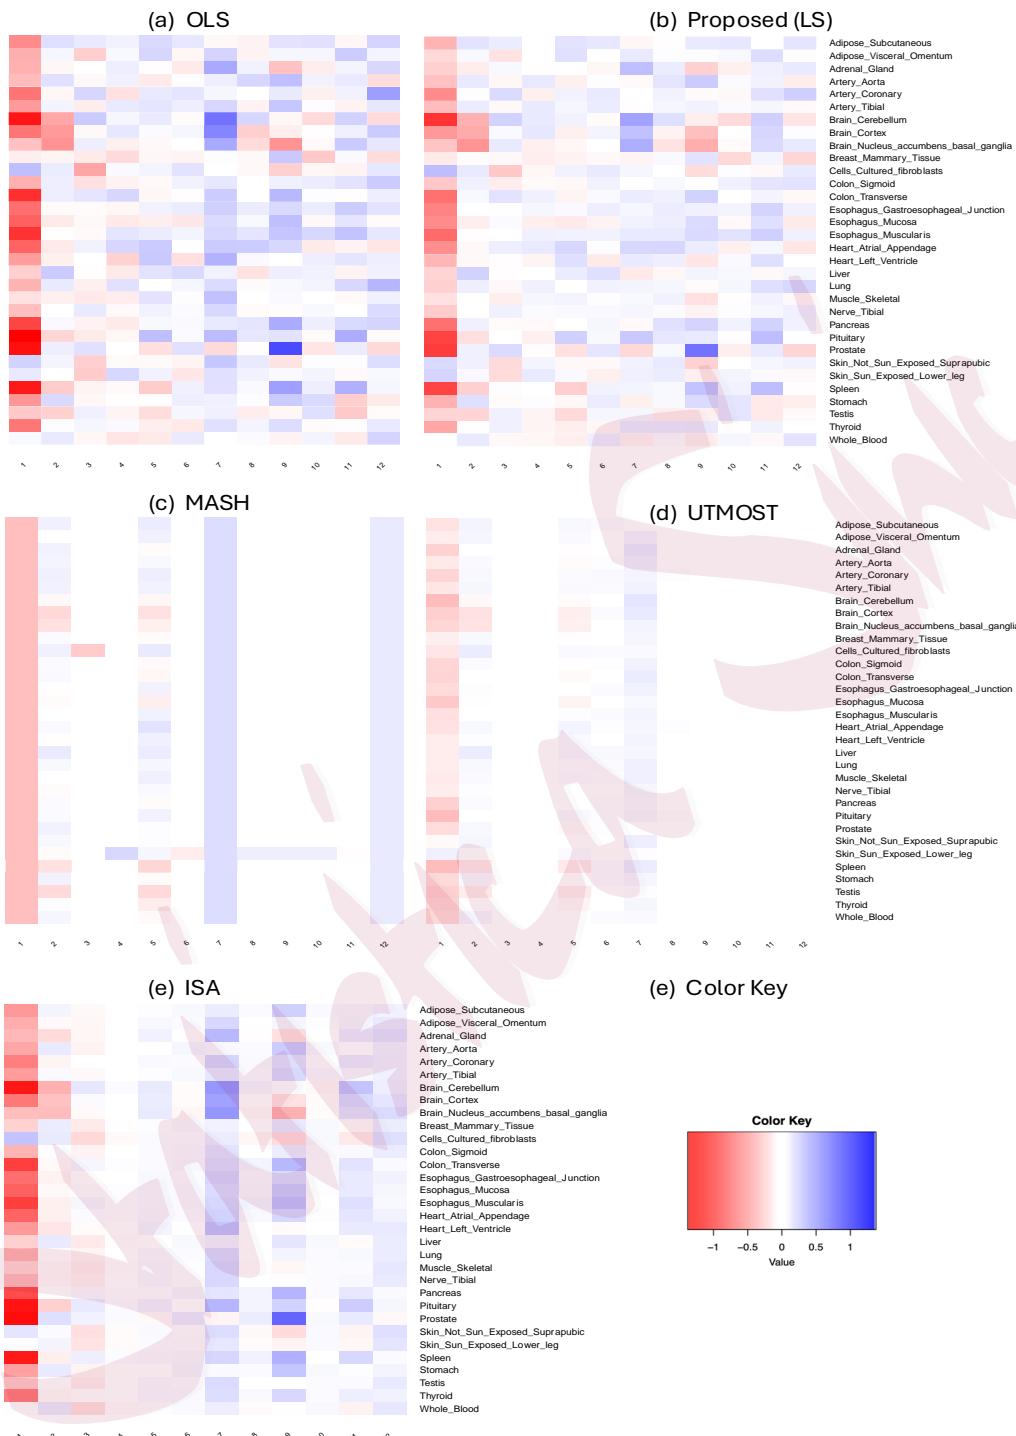


Figure 6: Heatmaps of estimated coefficients of 12 cis-SNPs for the ZNF138 gene by different methods.

7 Discussion

We developed a formal data-integration method for multiple-source linear regressions with applications in TWAS studies. Existing literature views this problem through the lens of multi-response regression, although proposed solutions often make assumptions on the parameter structure that are not possible to verify in practice. We take an empirical Bayes approach here which naturally leads to shrinkage estimators, and are shown through extensive numerical experiments to have excellent performance under several parameter settings. Our work also builds upon the close connection between estimating the regression coefficient matrix and estimating covariance matrices.

Addressing uncertainty in the estimates is a natural next step. This could be either handled directly or by computational methods such as bootstrap. Another interesting future direction is to consider nonparametric regressions with more predictive power within each source. To integrate predictions from different sources, the proposed empirical Bayes idea could still be useful. In this paper, we only considered the low-dimensional cases where the number of subjects is larger than the number of predictors. In fact, high-dimensional cases are also important and could be another future direction, where we may leverage the asymptotic distribution of the de-biased lasso (Javanmard and Montanari, 2014) and incorporate more genotype biomarkers. Moreover, for large genetic datasets, summary statistics, such as correlations and standard errors, are easier to share than individual-level data due to privacy issues. Thus, it would be worthwhile to explore extending the proposed method to a summary statistics-based approach in the future.

Furthermore, in TWAS applications, certain domain knowledge can potentially be incorporated to further improve prediction and interpretability. For example, one may inte-

grate information of known gene co-expression networks (Langfelder and Horvath, 2008) or tissue similarity (Hu et al., 2019) (e.g., from GTEx tissue correlation structure). Specifically, suppose that $\Sigma_T^{n \times n}$ is a positive definite symmetric matrix with $\Sigma_T(i, j)$ encoding the similarity/distance between tissue i and tissue j . Such a matrix is assumed to be constructed by the user. Then an observation model for the OLS estimates can be formed as $\widehat{B} \sim \text{MN}(B, \Sigma_T, Q)$ with the prior $B \sim \text{MN}(0, \Sigma_T, Q^{1/2}\Omega Q^{1/2})$. This scenario can be easily incorporated within our method since $\mathbb{E}(B \mid \widehat{B}) = \widehat{B} - \Sigma_T^{-1}\widehat{B}Q^{1/2}(\mathbf{I} + \Omega)^{-1}Q^{-1/2}$, which requires estimating the inverse of the marginal covariance matrix $(\mathbf{I} + \Omega)^{-1}$ under a rescaled relative savings loss (the rescaling is through Σ_T^{-1}).

Supplementary Material

Additional simulation results, proofs, and algorithms are provided in the supplementary material.

Acknowledgments

This work was supported in part by the National Science Foundation under Grant DMS 2210860. We are also grateful to Yunlong Liu of Department of Medical and Molecular Genetics in Indiana University School of Medicine, for his helpful discussions on real data.

References

Anderson, T. W. (1951). Estimating linear restrictions on regression coefficients for multivariate normal distributions. *The Annals of Mathematical Statistics*, 327–351.

Bai, R. and M. Ghosh (2018). High-dimensional multivariate posterior consistency under global-local shrinkage priors. *Journal of Multivariate Analysis* 167, 157–170.

Bai, Z.-D. and J. W. Silverstein (1998). No eigenvalues outside the support of the limiting spectral distribution of large-dimensional sample covariance matrices. *The Annals of Probability* 26(1), 316–345.

Banerjee, T., G. Mukherjee, and D. Paul (2021). Improved shrinkage prediction under a spiked covariance structure. *Journal of Machine Learning Research* 22(180), 1–40.

Boukehil, D., D. Fourdrinier, F. Mezoued, and W. E. Strawderman (2021). Estimation of the inverse scatter matrix for a scale mixture of wishart matrices under Efron–Morris type losses. *Journal of Statistical Planning and Inference* 215, 368–387.

Boyle, E. A., Y. I. Li, and J. K. Pritchard (2017). An expanded view of complex traits: from polygenic to omnigenic. *Cell* 169(7), 1177–1186.

Bunea, F., Y. She, and M. H. Wegkamp (2011). Optimal selection of reduced rank estimators of high-dimensional matrices. *The Annals of Statistics* 39(2), 1282–1309.

Bunea, F., Y. She, M. H. Wegkamp, et al. (2012). Joint variable and rank selection for parsimonious estimation of high-dimensional matrices. *The Annals of Statistics* 40(5), 2359–2388.

Carvalho, C., N. Polson, and J. Scott (2010). The horseshoe estimator for sparse signals. *Biometrika* 97(2), 465–480.

Chakraborty, A., A. Bhattacharya, and B. K. Mallick (2020). Bayesian sparse multiple regression for simultaneous rank reduction and variable selection. *Biometrika* 107(1), 205–221.

Chen, L. and J. Z. Huang (2012). Sparse reduced-rank regression for simultaneous dimension reduction and variable selection. *Journal of the American Statistical Association* 107(500), 1533–1545.

Consortium, G. (2020). The GTEx consortium atlas of genetic regulatory effects across human tissues. *Science* 369(6509), 1318–1330.

Efron, B. and C. Morris (1972). Empirical Bayes on vector observations: An extension of Stein's method. *Biometrika* 59(2), 335–347.

Efron, B. and C. Morris (1976). Multivariate empirical Bayes and estimation of covariance matrices. *The Annals of Statistics*, 22–32.

Fu, J., M. G. Wolfs, P. Deelen, H.-J. Westra, R. S. Fehrman, G. J. Te Meerman, W. A. Buurman, S. S. Rensen, H. J. Groen, R. K. Weersma, et al. (2012). Unraveling the regulatory mechanisms underlying tissue-dependent genetic variation of gene expression. *PLoS Genetics* 8(1), e1002431.

Geweke, J. (1991). Efficient simulation from the multivariate normal and student-t distributions subject to linear constraints and the evaluation of constraint probabilities. In *Computing science and statistics: Proceedings of the 23rd symposium on the interface*, pp. 571–578. Fairfax, Virginia: Interface Foundation of North America, Inc.

Glynn, P. W. and C.-h. Rhee (2014). Exact estimation for Markov chain equilibrium expectations. *Journal of Applied Probability* 51(A), 377–389.

Gresle, M. M., M. A. Jordan, J. Stankovich, T. Spelman, L. J. Johnson, L. Laverick, A. Hamlett, L. D. Smith, V. G. Jokubaitis, J. Baker, et al. (2020). Multiple sclerosis risk variants regulate gene expression in innate and adaptive immune cells. *Life Science Alliance* 3(7).

Haff, L. (1979). Estimation of the inverse covariance matrix: Random mixtures of the inverse wishart matrix and the identity. *The Annals of Statistics*, 1264–1276.

Heap, G. A., G. Trynka, R. C. Jansen, M. Bruinenberg, M. A. Swertz, L. C. Dinesen, K. A. Hunt, C. Wijmenga, D. A. Vanheel, and L. Franke (2009). Complex nature of SNP genotype effects on gene expression in primary human leucocytes. *BMC Medical Genomics* 2, 1–13.

Hobert, J. P. (2011). The data augmentation algorithm: Theory and methodology. *Handbook of Markov Chain Monte Carlo*, 253–293.

Hu, Y., M. Li, Q. Lu, H. Weng, J. Wang, S. M. Zekavat, Z. Yu, B. Li, J. Gu, S. Muchnik, et al. (2019). A statistical framework for cross-tissue transcriptome-wide association analysis. *Nature Genetics* 51(3), 568–576.

Izenman, A. J. (1975). Reduced-rank regression for the multivariate linear model. *Journal of Multivariate Analysis* 5(2), 248–264.

Javanmard, A. and A. Montanari (2014). Confidence intervals and hypothesis testing for high-dimensional regression. *The Journal of Machine Learning Research* 15(1), 2869–2909.

Josse, J. and S. Wager (2016). Bootstrap-based regularization for low-rank matrix estimation. *Journal of Machine Learning Research* 17(124), 1–29.

Kim, Y., W. Wang, P. Carbonetto, and M. Stephens (2024). A flexible empirical Bayes approach to multiple linear regression and connections with penalized regression. *Journal of Machine Learning Research* 25(185), 1–59.

Kubokawa, T. and M. S. Srivastava (2008). Estimation of the precision matrix of a singular wishart distribution and its application in high-dimensional data. *Journal of multivariate Analysis* 99(9), 1906–1928.

Langfelder, P. and S. Horvath (2008). Wgcna: an r package for weighted correlation network analysis. *BMC bioinformatics* 9(1), 559.

Ledoit, O. and S. Péché (2011). Eigenvectors of some large sample covariance matrix ensembles. *Probability Theory and Related Fields* 151(1-2), 233–264.

Ledoit, O. and M. Wolf (2004). A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis* 88(2), 365–411.

Ledoit, O. and M. Wolf (2012). Nonlinear shrinkage estimation of large-dimensional covariance matrices. *The Annals of Statistics* 40(2), 1024 – 1060.

Ledoit, O. and M. Wolf (2018). Optimal estimation of a large-dimensional covariance matrix under Stein’s loss. *Bernoulli* 24(4B), 3791 – 3832.

Ledoit, O. and M. Wolf (2022). Quadratic shrinkage for large covariance matrices. *Bernoulli* 28(3), 1519–1547.

Lloyd-Jones, L. R., A. Holloway, A. McRae, J. Yang, K. Small, J. Zhao, B. Zeng, A. Bakshi, A. Metspalu, M. Dermitzakis, et al. (2017). The genetic architecture of gene expression in peripheral blood. *The American Journal of Human Genetics* 100(2), 228–237.

Lonsdale, J., J. Thomas, M. Salvatore, R. Phillips, E. Lo, S. Shad, R. Hasz, G. Walters, F. Garcia, N. Young, et al. (2013). The genotype-tissue expression (GTEx) project. *Nature Genetics* 45(6), 580–585.

Mai, J., M. Lu, Q. Gao, J. Zeng, and J. Xiao (2023). Transcriptome-wide association studies: recent advances in methods, applications and available databases. *Communications Biology* 6(1), 899.

Matsuda, T. and F. Komaki (2015). Singular value shrinkage priors for Bayesian prediction. *Biometrika* 102(4), 843–854.

Matsuda, T. and F. Komaki (2019). Empirical Bayes matrix completion. *Computational Statistics & Data Analysis* 137, 195–210.

Matsuda, T. and W. E. Strawderman (2022). Estimation under matrix quadratic loss and matrix superharmonicity. *Biometrika* 109(2), 503–519.

Narisetty, N. N. and X. He (2014). Bayesian variable selection with shrinking and diffusing priors. *The Annals of Statistics* 42(2), 789 – 817.

Park, T. and G. Casella (2008). The Bayesian lasso. *Journal of the American Statistical Association* 103(482), 681–686.

Paul, D. (2007). Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statistica Sinica*, 1617–1642.

Polson, N. G. and J. G. Scott (2010). Shrink globally, act locally: sparse Bayesian regularization and prediction. *Bayesian Statistics* 9, 501–538.

Silverstein, J. W. (1995). Strong convergence of the empirical distribution of eigenvalues of large dimensional random matrices. *Journal of Multivariate Analysis* 55(2), 331–339.

Silverstein, J. W. and Z. Bai (1995). On the empirical distribution of eigenvalues of a class of large dimensional random matrices. *Journal of Multivariate Analysis* 54(2), 175–192.

Silverstein, J. W. and S.-I. Choi (1995). Analysis of the limiting spectral distribution of large dimensional random matrices. *Journal of Multivariate Analysis* 54(2), 295–309.

Stein, C. M. (1981). Estimation of the mean of a multivariate normal distribution. *The Annals of Statistics*, 1135–1151.

Velu, R. and G. C. Reinsel (2013). *Multivariate reduced-rank regression: theory and applications*, Volume 136. Springer Science & Business Media.

Wainberg, M., N. Sinnott-Armstrong, N. Mancuso, A. N. Barbeira, D. A. Knowles, D. Golan, R. Ermel, A. Rusalepp, T. Quertermous, K. Hao, et al. (2019). Opportunities and challenges for transcriptome-wide association studies. *Nature Genetics* 51(4), 592–599.

Wang, J., E. R. Gamazon, B. L. Pierce, B. E. Stranger, H. K. Im, R. D. Gibbons, N. J. Cox, D. L. Nicolae, and L. S. Chen (2016). Imputing gene expression in uncollected tissues within and beyond GTEx. *The American Journal of Human Genetics* 98(4), 697–708.

Wang, Y. and S. D. Zhao (2021). Linear shrinkage for predicting responses in large-scale multivariate linear regression. *arXiv preprint arXiv:2104.08970*.

Xue, F. and H. Li (2022). An empirical bayes regression for multi-tissue eqtl data analysis. *arXiv preprint arXiv:2211.13889*.

Yuan, M., A. Ekici, Z. Lu, and R. Monteiro (2007). Dimension reduction and coefficient estimation in multivariate linear regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 69(3), 329–346.

Department of Statistics, Purdue University

E-mail: antik015@purdue.edu

ORCID: <https://orcid.org/0000-0001-8255-3060>

Department of Statistics, Purdue University

E-mail: feixue@purdue.edu

ORCID: <https://orcid.org/0000-0003-1135-4894>