

Precise High-Dimensional Asymptotics for Quantifying Heterogeneous Transfers

Fan Yang

FYANGMATH@MAIL.TSINGHUA.EDU.CN

*Yau Mathematical Sciences Center, Tsinghua University
Beijing 100084, China*

Hongyang R. Zhang

HO.ZHANG@NORTHEASTERN.EDU

*Khoury College of Computer Sciences, Northeastern University
Boston, MA 02115, US*

Sen Wu

SENWU@CS.STANFORD.EDU

*Department of Computer Science, Stanford University
Stanford, CA 94305, US*

Christopher Ré

CHRISMRE@STANFORD.EDU

*Department of Computer Science, Stanford University
Stanford, CA 94305, US*

Weijie J. Su

SUW@WHARTON.UPENN.EDU

*Department of Statistics and Data Science, University of Pennsylvania
Philadelphia, PA 19104, US*

Editor: Mladen Kolar

Abstract

The problem of learning one task using samples from another task is central to transfer learning. In this paper, we focus on answering the following question: when does combining the samples from two related tasks perform better than learning with one target task alone? This question is motivated by an empirical phenomenon known as *negative transfer* often observed in transfer learning practice. While the transfer effect from one task to another depends on factors such as their sample sizes and the spectrum of their covariance matrices, precisely quantifying this dependence has remained a challenging problem. In order to compare a transfer learning estimator to single-task learning, one needs to compare the risks between the two estimators precisely. Further, the comparison depends on the distribution shifts between the two tasks. This paper applies recent developments of random matrix theory to tackle this challenge in a high-dimensional linear regression setting with two tasks. We provide precise high-dimensional asymptotics for the bias and variance of a classical hard parameter sharing (HPS) estimator in the proportional limit, when the sample sizes of both tasks increase proportionally with dimension at fixed ratios. The precise asymptotics apply to various types of distribution shifts, including covariate shifts, model shifts, and combinations of both. We illustrate these results in a random-effects model to mathematically prove a phase transition from positive to negative transfer as the number of source task samples increases. One insight from the analysis is that a rebalanced HPS estimator, which downsizes the source task when the model shift is high, achieves the minimax optimal rate. The finding regarding phase transition also applies to multiple tasks when feature covariates are shared across all tasks. Simulations validate the accuracy of the high-dimensional asymptotics for finite dimensions.

Keywords: transfer learning, negative transfer, covariate shift, high-dimensional asymptotics, random matrix theory

1. Introduction

Given samples from two tasks, does combining the samples from both tasks together yield a better estimator for the target task of interest? Concretely, suppose there are n_1 samples drawn from a p -dimensional distribution D_1 with real-valued labels in a *source* task. Suppose there are n_2 samples drawn from a p -dimensional distribution D_2 with real-valued labels in the *target* task. Does combining the samples from both the source and the target tasks together lead to an estimator whose performance would be better than the learning outcome from utilizing the n_2 samples alone?

This question is motivated by scenarios where one would like to use samples from an auxiliary task to augment the sample size of a target task. It is widely observed that when the two tasks are far apart, negative transfer can occur, where the transfer learning performance is worse than that of single-task learning (Pan and Yang, 2009; Wu et al., 2020). However, a rigorous proof of when and why negative transfer can happen (even within some simple statistical models) has remained elusive in the statistical learning literature.

Identifying negative transfer requires modeling the relationship between two tasks. Early work has approached such questions through structural learning (Ando and Zhang, 2005), where a shared, low-dimensional structure is extracted from several (possibly unlabeled) tasks via an alternating optimization framework. Another influential line of early work is the quantification of transfer learning with uniform convergence bounds (Ben-David et al., 2010a). Notice that negative transfer can happen if the label distribution of task one, conditioned on the covariates, differs significantly from task two. For brevity, we refer to settings where $D_1 \neq D_2$ as covariate shift and settings where label distributions differ (conditioned on the covariates) as model shift. However, analyzing these settings requires going beyond the in-distributional assumptions of supervised learning, which has led to recent results in transfer learning for the two-task linear regression setting (Lei et al., 2021; Dar and Baraniuk, 2022). Recent results by Dhifallah and Lu (2021) shed light on an intriguing phase transition that has been observed in empirical transfer learning. They study both hard transfer and soft transfer via a precise analysis of the transfer rates. In this paper, we propose to quantify transfer effects in high-dimensional linear regression by employing recently developed techniques from the random matrix theory literature in the proportional limit setting.

Concretely, we study a high-dimensional linear regression setting with two tasks under combinations of covariate and model shifts. We study a hard transfer estimation procedure in the high-dimensional limit. This procedure is rooted in the classical literature (Caruana, 1997), whose theoretical properties are not well-understood. We study functions involving two sample covariance matrices with arbitrarily different population covariances. We estimate the asymptotic limits of the bias and variance of hard parameter sharing and illustrate our estimates in a random-effects model. We achieve these results using tools from modern random matrix theory and free probability theory (Bai and Silverstein, 2010; Erdős and Yau, 2017; Nica and Speicher, 2006). In particular, building on recent developments in random matrix theory (Bao et al., 2017; Bloemendal et al., 2014; Knowles and Yin, 2016), we derive




Figure 1: We illustrate the performance of HPS vs. OLS (ordinary least squares) for different levels of model shifts, indicated by $\mu \approx \|\beta^{(1)} - \beta^{(2)}\|/\sqrt{2}$. We consider two linear regression tasks and vary the sample size of the source task, denoted as n_1 and μ , while holding the target task fixed. The gray line refers to the performance of OLS. Any point above the gray line represents negative transfer, while any point below represents positive transfer. In this simulation, we set $p = 100$, $n_2 = 300$, and $\sigma^2 = \frac{1}{4}$. We sample the covariates X from a p -dimensional isotropic Gaussian and plot the excess risk of HPS (see equation (2.6)). For further details about the setup, see Proposition 8, Section 4.

precise high-dimensional asymptotics involving combinations of two covariance matrices, along with *almost sharp convergence rates* to the limits.

1.1 Problem Setup

We first introduce our problem setup. For either task $i = 1$ or task $i = 2$, let $x_1^{(i)}, x_2^{(i)}, \dots, x_{n_i}^{(i)}$ denote the feature covariates and $y_1^{(i)}, y_2^{(i)}, \dots, y_{n_i}^{(i)}$ denote the prediction labels; Recall that n_1 and n_2 refer to the number of samples from task one and task two, respectively. We assume a linear model specified by an unknown parameter vector $\beta^{(i)} \in \mathbb{R}^p$:

$$y_j^{(i)} = (x_j^{(i)})^\top \beta^{(i)} + \varepsilon_j^{(i)}, \quad \text{for any } j = 1, \dots, n_i, \quad (1.1)$$

where $\varepsilon_j^{(i)} \in \mathbb{R}$ denotes a random noise variable with mean zero and variance σ^2 . We refer to the first task as the *source* task and the second task as the *target* task for ease of presentation. We assume $n_2 > p$ while n_1 can be either less than or greater than p .

Following the existing random matrix theory literature, we study random designs where the covariates are given by

$$x_j^{(i)} = (\Sigma^{(i)})^{\frac{1}{2}} z_j^{(i)}, \quad \text{for } i \in \{1, 2\} \text{ and for any } j = 1, \dots, n_i,$$

where $z_j^{(i)}$ consists of independent and identically distributed entries of mean zero and variance one, and $\Sigma^{(1)} \in \mathbb{R}^{p \times p}$ and $\Sigma^{(2)} \in \mathbb{R}^{p \times p}$ denote the population covariance matrices. Let $X^{(i)} = (x_1^{(i)}, \dots, x_{n_i}^{(i)})^\top \in \mathbb{R}^{n_i \times p}$ denote a matrix that corresponds to all the covariates of

task i , for $i \in \{1, 2\}$. Let $Y^{(1)} \in \mathbb{R}^{n_1}$, $Y^{(2)} \in \mathbb{R}^{n_2}$ be the vectors of labels. Then, we minimize the following objective, parameterized by a shared variable $B \in \mathbb{R}^p$:

$$\ell(B) = \left\| X^{(1)}B - Y^{(1)} \right\|^2 + \left\| X^{(2)}B - Y^{(2)} \right\|^2. \quad (1.2)$$

Notice that B provides a shared feature vector for both tasks. The results described in the next section will permit straightforward extensions to variants of objective (1.2), including adding weights to each task and a ridge penalty. Thus, we will describe our main results for unweighted and unregularized objectives without losing generality. One can see that when $n_2 > p$, there is a unique minimizer of f . We denote this estimator for the target task as $\hat{\beta}_2^{\text{HPS}}$, which is also called hard parameter sharing (HPS) or hard transfer in the multi-task learning literature (Caruana, 1997). We will analyze the bias and variance of the HPS estimator and connect the analysis to transfer effects.

Our motivation for studying this simple HPS estimator is that it provides a simple setting in which both positive and negative transfers can happen. In Figure 1, we illustrate that if the model shift is high (as indicated by $\mu = 0.45$), HPS always performs worse than OLS (ordinary least squares), regardless of the value of n_1 . If the model shift is moderate (as indicated by $\mu = 0.2, 0.3, 0.35$), the transfer effect is positive within a small range of n_1 until reaching a threshold and turns negative afterward. In summary, the HPS estimator provides an ideal scenario for understanding the effects of transfer learning under different types of distribution shifts between the source task and the target task.

1.2 Summary of Our Results

We give a summary of our results, which follow the assumptions of a proportional limit setting, with both n_1 and n_2 increasing to infinity in proportion as p goes to infinity. Our goal is to find the asymptotic limit of the bias and variance of the HPS estimator applied to the target task ($i = 2$). In particular, deriving the asymptotic limit of the bias and variance of the HPS estimator requires studying several functions involving two independent sample covariance matrices with different population covariance matrices. Below, we provide a simplified exposition and elaborate on the results later in Section 2.

First, the variance of the HPS estimator is equal to

$$\text{Tr} \left[\Sigma^{(2)} \left(X^{(1)\top} X^{(1)} + X^{(2)\top} X^{(2)} \right)^{-1} \right]. \quad (1.3)$$

The bias of the HPS estimator is more involved and is deferred until Lemma 2. The bias and variance both involve the inverse of the sum of the sample covariance matrices of both tasks, which exhibit a covariate shift between them.

When $n_2 > p$, the OLS estimator exists, and it is well-known that the limit of its excess risk is equal to $\frac{\sigma^2 p}{n_2 - p}$ (Bai and Silverstein, 2010). To the best of our knowledge, the asymptotic limit of equation (1.3) (when $n_1/p \rightarrow c_1, n_2/p \rightarrow c_2$ as $p \rightarrow \infty$) has not been identified in the literature. Our paper takes the first step to fill the gap. We summarize our main results as follows:

- First, we consider the covariate shift setting, where $\Sigma^{(1)}$ and $\Sigma^{(2)}$ are arbitrary and $\beta^{(1)} = \beta^{(2)}$. In Theorem 3, we describe the asymptotic limit of the variance formula as

a function of the singular values of the covariate shift matrix. This result generalizes classical results on the limit of the trace of the inverse of one sample covariance matrix to two covariance matrices under covariate shifts.

- Second, we consider the model shift setting, where $\beta^{(1)}$ and $\beta^{(2)}$ are arbitrary and $\Sigma^{(1)} = \Sigma^{(2)}$. In Theorem 6, we describe the asymptotic limit of the bias and variance of HPS as a function of the Euclidean distance between $\beta^{(1)}$ and $\beta^{(2)}$, and the sample sizes. We then illustrate the results in a random-effects model. Motivated by the phase transition, we show that *with a simple adjustment, HPS can match the minimax lower bound* (see Theorem 9).
- Third, we extend the above results in two aspects: i) When both covariate and model shifts are present; ii) When there are multiple tasks in the input. The detailed results, along with supporting simulations, are provided in Section 5. In particular, we demonstrate that the findings regarding complementary covariate shifts and phase transitions in transfer learning are applicable to more general settings.

Finally, we present a case study comparing the statistical behavior of a soft transfer estimator with HPS. The soft parameter sharing (SPS) estimator (which is also called soft transfer in Dhifallah and Lu (2021)) is defined by minimizing the following for $\lambda > 0$:

$$\ell(B, z) = \left\| X^{(1)}(B + z) - Y^{(1)} \right\|^2 + \left\| X^{(2)}B - Y^{(2)} \right\|^2 + \lambda \|z\|^2. \quad (1.4)$$

Let \hat{B}, \hat{z} denote the minimizer, and the SPS estimator for the target task is defined as $\hat{\beta}_2^{\text{SPS}}(\lambda) = \hat{B}$. The SPS estimator makes sense in the model shift setting where $\beta^{(1)}$ is far from $\beta^{(2)}$. We will provide the bias and variance of SPS under deterministic designs. Then, we provide an example to contrast the behavior of SPS with that of HPS. We find that the regularization effect of SPS is similar to downsizing the source task in HPS when the model shift is large.

We illustrate the above results in a random-effects model, where $\beta^{(i)}$ is equal to a shared model vector across all i plus an independent random effect for each task. The random-effects model provides a natural way to measure the heterogeneity of each task's underlying linear model and has been used to study distributed ridge regression (Dobriban and Sheng, 2020).

First, we demonstrate that covariate shifts can either enhance or hinder transfer learning performance compared to single-task learning. We describe an example in Proposition 4, showing that when $n_1 < n_2$, transferring from any covariate-shifted data source (i.e., $\Sigma^{(1)} \neq \Sigma^{(2)}$) achieves a lower excess risk of HPS than transferring from the data source with $\Sigma^{(1)} = \Sigma^{(2)}$. On the other hand, when $n_1 > n_2$, transferring from any covariate-shifted data source always incurs a higher excess risk of HPS than transferring from the data source with $\Sigma^{(1)} = \Sigma^{(2)}$.

Second, we identify three transfer regimes in the random-effects model. Let μ denote the Euclidean distance between $\beta^{(1)}$ and $\beta^{(2)}$. We show the following dichotomy between μ and n_1, n_2 :

- When $\mu^2 \frac{\text{Tr}[\Sigma^{(1)}]}{p} \leq \frac{\sigma^2 p}{2(n_2 - p)}$, the transfer effect is positive for any value of $n_1 > 0$: HPS always performs better than OLS.

- When $\frac{\sigma^2 p}{2(n_2-p)} < \mu^2 \frac{\text{Tr}[\Sigma^{(1)}]}{p} < \frac{\sigma^2 n_2}{2(n_2-p)}$, there exists a deterministic constant $\rho > 0$ such that the transfer effect is positive if and only if $\frac{n_1}{p} < \rho$. This corresponds to the crossing point between HPS and OLS in Figure 1.
- When $\mu^2 \frac{\text{Tr}[\Sigma^{(1)}]}{p} \geq \frac{\sigma^2 n_2}{2(n_2-p)}$, the transfer effect is negative for any value of $n_1 > 0$: HPS always performs worse than OLS.

Summary of contributions. This paper presents a rigorous study of the transfer effects of HPS estimators under various types of data set shifts between two tasks. We contribute to the transfer learning literature by using random matrix theory to analyze negative transfer of HPS in the *high-dimensional* asymptotic setting under covariate shifts. In particular, we derive the asymptotic limit of the HPS estimator under covariate shift (Theorem 3), which generalizes a classical result in the random matrix theory literature. For the model shift setting, our result (Theorem 6) can explain a phase transition (Proposition 8) that has been observed in empirical studies of transfer learning. We extend these results to show that the insights hold more generally beyond the specific examples. Simulation studies demonstrate the accuracy of the derived asymptotic in finite dimensions. We highlight numerous technically challenging open questions, and we hope our work can inspire further studies on applying random matrix theory to understand transfer learning.

1.3 Related Work

Our work expands the existing statistical learning literature by contributing a random matrix theory perspective to quantify transfer effects in high-dimensional linear regression with two tasks. Our perspective differs from the existing literature, which uses uniform convergence arguments (Baxter, 2000; Ben-David and Schuller, 2003; Maurer, 2006) and generalization bounds (Crammer et al., 2008; Ben-David et al., 2010a) to quantify the performance of transfer learning. The benefit of using precise asymptotics to study transfer learning performance compared to generalization bounds is that they can be used to compare the performances of different estimators, and their predictions remain accurate in finite dimensions. This is crucial for explaining phenomena such as phase transitions, as highlighted in the simulation.

There has been a long line of work that tries to formulate the notion of information or knowledge transfer in the statistical learning literature. In the context of the many Normal means models, lasso regularization can be used to leverage the shared support of multiple tasks (Kolar et al., 2011), and similar insights can also be fleshed out in multi-task learning with marginal regression (Kolar and Liu, 2012). The work of Dhifallah and Lu (2021) provides a precise analysis of both the hard transfer and soft transfer approaches in a student-teacher framework. They provide the convergence of the training loss in the asymptotic setting where $n_1/p, n_2/p$ converge to a constant. Then, they provide illustrative examples of the precise analysis in regression and classification models to explain the phase transition in transfer learning. One of the key differences between our results and their work is that we address covariate shifts in the high-dimensional setting, while also providing sharp convergence rates. This requires the use of random matrix theory techniques from more recent literature. Dar and Baraniuk (2022) derive generalization error rates for the two-task transfer learning in the important *overparameterized* regime (see also Ju et al. (2023) for further studies in the overparameterized transfer learning regime). They derive a generalization error formula for

the minimum ℓ_2 -norm solution of the target task, which can reproduce a double descent behavior and explain the negative transfer phenomenon (Dar et al., 2024). In this vein, the high-dimensional asymptotics we have provided under various distribution shifts may be viewed as complementary to this line of work, contributing to the mathematical formalization of transfer learning.

In the context of linear regression tasks, recent studies have investigated the optimal estimators for transfer learning under distribution shifts. Lei et al. (2021) give an in-depth study of minimax estimators for transfer learning from two linear regression tasks. They derive minimax optimal estimators that can be formulated as solving a convex program, and the approach can be applied under both covariate and model shifts. In the setting of multiple linear regression, Li et al. (2021) design an adaptive estimation procedure in a setting where many tasks with Gaussian covariates are present. Duan and Wang (2023) introduce an adaptive estimation procedure when many tasks are present in multi-task learning. There is another related line of work on multi-view regression, where the input variable (which is a real vector) can be partitioned into two different views, and it is assumed that either view of the input is sufficient to make accurate predictions (Kakade and Foster, 2007). Zhao et al. (2024) design a two-stage procedure for matrix estimation that uses a multi-task learning objective to capture shared and unique features, and then refines them to adjust for structural differences between the target and source matrices. Our paper complements these existing works, as we focus on the proportional limit setting to examine the performance of commonly used estimators and provide insights.

When covariates are sampled from Gaussian distributions, the precise asymptotic limit of the inverse of a sample covariance matrix can be derived directly from the properties of the Wishart distribution. In the high-dimensional setting, the eigenvalues of a Wishart matrix satisfy the well-known Marchenko–Pastur (MP) law, whose Stieltjes transform characterizes the variance limit. Furthermore, it is well known that the MP law holds universally regardless of the underlying data distribution of the covariates (see, e.g., Bai and Silverstein (2010)). Bloemendal et al. (2014) obtain a sharp convergence rate of the empirical spectral distribution (ESD) to the MP law for sample covariance matrices with isotropic population covariances. Knowles and Yin (2016) later extend this result to sample covariance matrices with arbitrary population covariances. These results are proved by establishing the optimal convergence estimates of the Stieltjes transforms of sample covariance matrices, also known as *local laws* in the random matrix theory literature. We refer interested readers to Erdős and Yau (2017) and the references therein for a detailed review. One technical contribution of this work is to extend these techniques to the two-task setting and prove an almost sharp local law for the sum of two sample covariance matrices with arbitrary covariate shifts. This local law allows us to derive the precise variance limit depending on the singular values of the covariate shift matrix. Finally, recent work has also studied the statistical behavior of interpolators in the high-dimensional setting (Sur and Candès, 2019), which would be another promising setting for exploring statistical transfer learning.

The asymptotic limit of equation (1.3) may also be derived using free probability theory. However, this approach is not fully justified when the covariates are sampled from non-Gaussian distributions with non-diagonal covariate shift matrices. Furthermore, our result provides almost sharp convergence rates to the asymptotic limit, while it is unclear how to obtain such rates using free probability techniques. The bias term involves asymmetric

matrices in terms of two sample covariance matrices, whose analysis is technically involved. Our techniques are inspired by free additions of random matrices (Nica and Speicher, 2006) and recent results (Bao et al., 2017). In particular, we provide the first precise bias limit in the model shift setting, assuming that $\Sigma^{(1)} = \Sigma^{(2)}$ or that $\Sigma^{(2)}$ is isotropic, and the covariates are sampled from Gaussian distributions. Showing the asymptotic bias limit under arbitrary covariate and model shifts is an interesting open problem for future work (See Remarks 7 and 13 for further discussions).

1.4 Organization

The rest of this paper is organized as follows. In Section 2, we state the data model and its underlying assumption. Then, we connect the transfer effect of the HPS estimator with its bias-variance decomposition. In Section 3, we present the high-dimensional asymptotic limits of the bias and variance of HPS under covariance shifts. In Section 4, we characterize the high-dimensional asymptotic limits of the bias and variance of HPS under model shifts. Section 5 extends our main findings from the above two settings to more general settings. Section 6 provides a preliminary study of the bias and variance of the SPS estimator. Finally, in Section 7, we state the conclusion of our paper. Appendix A to Appendix E provides detailed proofs of our results.

2. Preliminaries

This section describes the data model that we will work with, along with the underlying assumptions. Then, we describe the bias and variance of the HPS estimator. Finally, we connect the bias-variance decomposition to the transfer analysis in the rest of the paper.

2.1 Data Model and Assumptions

Recall that we have two tasks. For $i = 1, 2$, $X^{(i)} \in \mathbb{R}^{n_i \times p}$ corresponds to task i 's covariates and $Y^{(i)} \in \mathbb{R}^{n_i}$ corresponds to their labels. Moreover, let $\varepsilon^{(i)} \in \mathbb{R}^{n_i}$ be the vector notation corresponding to the additive noise of data set i . Then, equation (1.1) can be reformulated as $Y^{(i)} = X^{(i)}\beta^{(i)} + \varepsilon^{(i)}$. Assume that $\beta^{(1)}$ and $\beta^{(2)}$ are two arbitrary (deterministic or random) vectors that are independent of $X^{(i)}$ and $\varepsilon^{(i)}$. Throughout the paper, we make the following assumptions on $X^{(i)}$ and $\varepsilon^{(i)}$, which are standard in the random matrix theory literature (see, e.g., Tulino and Verdú (2004); Bai and Silverstein (2010)).

First, the row vectors of $X^{(i)}$ are i.i.d. centered random vectors with $p \times p$ population covariance matrix $\Sigma^{(i)}$ and an $n_i \times p$ random matrix $Z^{(i)} = [Z_{jk}^{(i)}]$ with independent entries of zero mean and unit variance:

$$X^{(i)} = Z^{(i)}(\Sigma^{(i)})^{1/2} \in \mathbb{R}^{n_i \times p}. \quad (2.1)$$

Let $\tau > 0$ be a small constant. Suppose the φ -th moment of each entry $Z_{jk}^{(i)}$ is bounded from above by $1/\tau$, for a constant $\varphi > 4$:

$$\mathbb{E} \left[|Z_{jk}^{(i)}|^\varphi \right] \leq \tau^{-1}. \quad (2.2)$$

The eigenvalues of $\Sigma^{(i)}$, denoted as $\sigma_1^{(i)}, \dots, \sigma_p^{(i)}$, are all bounded between τ and τ^{-1} :

$$\tau \leq \sigma_p^{(i)} \leq \dots \leq \sigma_2^{(i)} \leq \sigma_1^{(i)} \leq \tau^{-1}. \quad (2.3)$$

Second, $\varepsilon^{(i)} \in \mathbb{R}^{n_i}$ is a random vector with independent entries having mean zero, variance σ^2 , and bounded moments up to any order, i.e., for any fixed $k \in \mathbb{N}$, there exists a constant $C_k > 0$ such that

$$\mathbb{E} \left[|\varepsilon_j^{(i)}|^k \right] \leq C_k. \quad (2.4)$$

Third, the sample sizes are comparable to the dimension p . Denote by $\rho_1 = n_1/p$ and $\rho_2 = n_2/p$. Assume

$$0 < \rho_1 \leq p^{\tau-1}, \quad 1 + \tau \leq \rho_2 \leq p^{\tau-1}, \quad 0 < \rho_1/\rho_2 \leq \tau^{-1}. \quad (2.5)$$

The condition $\rho_2 \geq 1 + \tau$ ensures that the target task's sample covariance matrix is of full rank with high probability. The upper bound $\rho_i \leq p^{\tau-1}$ is a mild condition; otherwise, standard concentration results, such as the central limit theorem, already give accurate estimates in the linear model. The condition $\rho_1/\rho_2 \leq \tau^{-1}$ ensures the sample size imbalance between the two tasks is bounded by a factor that does not grow with p . In the transfer learning setting, one can think of ρ_1 and ρ_2 as fixed, positive constants that do not grow or diminish with p to model practical applications. To summarize, the underlying assumptions of the data model are as follows.

Assumption 1 *Let $\tau > 0$ be a small constant. Suppose $X^{(1)}$, $X^{(2)}$, $\varepsilon^{(1)}$, and $\varepsilon^{(2)}$ are mutually independent. Moreover, suppose the following holds for $i = 1$ and $i = 2$:*

1. $X^{(i)}$ takes the form of equation (2.1), where $Z^{(i)}$ is a random matrix with i.i.d. entries having zero mean, unit variance, and bounded moments as in equation (2.2), and $\Sigma^{(i)}$ is a deterministic positive definite symmetric matrix satisfying equation (2.3).
2. $\varepsilon^{(i)} \in \mathbb{R}^n$ is a random vector consisting of i.i.d. entries of zero mean, variance σ^2 , and bounded moments as in equation (2.4).
3. ρ_i satisfies equation (2.5).

2.2 Bias and Variance

The risk of $\hat{\beta}_i^{\text{HPS}}$ over an unseen sample (x, y) of the target task i is given by (under the mean squared loss):

$$\mathbb{E}_{x,y} \left[\left\| y - x^\top \hat{\beta}_i^{\text{HPS}} \right\|^2 \right] = \left\| \Sigma^{(i)\frac{1}{2}} \left(\hat{\beta}_i^{\text{HPS}} - \beta^{(i)} \right) \right\|^2 + \sigma^2.$$

The excess risk is the difference between the above risk and the expected risk of the population risk optimizer:

$$L(\hat{\beta}_i^{\text{HPS}}) := \left\| \Sigma^{(i)\frac{1}{2}} \left(\hat{\beta}_i^{\text{HPS}} - \beta^{(i)} \right) \right\|^2. \quad (2.6)$$

We shall treat task two as the target task and treat task one as the source task. The single-task learning estimator for the target task, which is the OLS estimator denoted by $\hat{\beta}_2^{\text{OLS}}$, is well-defined given Assumption 1, because the covariance matrix of the target task is invertible with high probability under the assumption.

Next, we present a bias-variance decomposition of the excess risk of HPS.¹ Denote the sum of both tasks' sample covariance matrix as:

$$\hat{\Sigma} = X^{(1)\top} X^{(1)} + X^{(2)\top} X^{(2)}. \quad (2.7)$$

The next result provides the bias and variance of the HPS estimator.

Lemma 2 *Under Assumption 1, for any small constant $c > 0$, with high probability over the randomness of the training samples $X^{(1)}, Y^{(1)}, X^{(2)}, Y^{(2)}$, the following estimates hold:*

$$L(\hat{\beta}_2^{\text{HPS}}) = \left(1 + O(p^{-\frac{1}{2}+c})\right) (L_{\text{bias}} + L_{\text{var}}), \quad (2.8)$$

where the bias and variance formulas are defined as

$$L_{\text{bias}} = \left\| \Sigma^{(2)\frac{1}{2}} \hat{\Sigma}^{-1} (X^{(1)\top} X^{(1)}) (\beta^{(1)} - \beta^{(2)}) \right\|^2, \quad (2.9)$$

$$L_{\text{var}} = \sigma^2 \text{Tr}[\Sigma^{(2)} \hat{\Sigma}^{-1}]. \quad (2.10)$$

Hereafter, an event Ξ is said to hold *with high probability* (w.h.p.) if $\mathbb{P}(\Xi) \rightarrow 1$ as $p \rightarrow \infty$. Moreover, the big-O notation $A = O(B)$ means that $|A| \leq C|B|$ for a constant $C > 0$ depending on the model parameters in Section 2.1 (i.e., τ , φ and C_k 's), but not on p , n_1 and n_2 . Hence, the event that equation (2.8) holds w.h.p. can be equivalently stated as

$$\lim_{p \rightarrow \infty} \mathbb{P} \left[\left| L(\hat{\beta}_2^{\text{HPS}}) - L_{\text{bias}} - L_{\text{var}} \right| \leq Cp^{-\frac{1}{2}+c} (L_{\text{bias}} + L_{\text{var}}) \right] = 1$$

for a constant $C > 0$. The proof of Lemma 2 can be found in Appendix A.3.

2.3 Using Bias and Variance to Analyze Transfer Effects

The bias and variance decomposition above allows us to reason about transfer effects by comparing the bias and variance of HPS with that of OLS. Notice that the bias of OLS is zero since $n_2 \geq (1 + \tau)p$ under Assumption 1. Hence, the bias of HPS is always larger than that of OLS. On the other hand, the variance of HPS is always lower than that of OLS. Let $X^{(1)} = 0$ in equation (2.10). By Woodbury matrix identity, the variance of OLS is always higher than formula (2.10):

$$\text{Tr} \left[\Sigma^{(2)} (X^{(2)\top} X^{(2)})^{-1} \right] \geq \text{Tr} \left[\Sigma^{(2)} \hat{\Sigma}^{-1} \right].$$

1. One way is to estimate the asymptotic limit of equation (2.6) directly. However, this would involve an estimation that needs to take into account the randomness of the covariates X along with the noise ε simultaneously. Instead, given that X and ε are independent, our approach is to first derive the bias-variance decomposition of the excess risk, which is equal to the expectation of $L(\hat{\beta}_i^{\text{HPS}})$ over the randomness of ε , and then compare the bias and variance of two different estimators.

Thus, the transfer effect of HPS is determined by the bias-variance decomposition: the bias always increases while the variance always decreases. Whether or not the transfer effect of HPS is positive depends on which effect dominates.

The above discussion highlights the need for precise bias-variance estimates to determine transfer effects. However, finding the exact limits is challenging because of data set shifts. Moreover, these shifts manifest in various forms. To make progress on this important yet challenging problem, we divide our study into various combinations of covariate and model shifts. Section 3 is devoted to the precise estimate of L_{var} for the HPS estimator under an arbitrary covariate shift but without model shift. Section 4 studies the effect of model shifts: Section 4.1 gives precise estimates of L_{bias} and L_{var} for the HPS estimator. Then, Section 4.2 discusses the adjusted HPS procedure. Finally, Section 4.3 states the minimax lower bound. Section 5 extends our main findings from the previous two sections to settings with both covariate and model shifts, as well as multiple tasks.

3. Covariate Shifts

This section presents a precise estimate of the variance L_{var} for the HPS estimator in equation (2.10) with an almost sharp convergence rate under two different population covariance matrices $\Sigma^{(1)}, \Sigma^{(2)}$. Then, we will provide two examples to illustrate the asymptotic limit, which provide varying levels of covariate shifts without model shifts. Our simulation study shows that the asymptotic estimates are quite accurate, even when the feature dimension p ranges from 50 to 100. Finally, we provide numerical comparisons of the excess risk of HPS with several other transfer learning estimators.

3.1 Main Results

Suppose the two tasks satisfy the same linear model ($\beta^{(1)} = \beta^{(2)}$) but have different population covariance matrices ($\Sigma^{(1)} \neq \Sigma^{(2)}$). Recall that the matrix $\hat{\Sigma}$ in equation (2.7) is a sum of two sample covariance matrices. Thus, the expectation of $\hat{\Sigma}$ is equal to a mixture of $\Sigma^{(1)}$ and $\Sigma^{(2)}$, with mixing proportions determined by the sample sizes n_1 and n_2 . Intuitively, the spectrum of $\hat{\Sigma}^{-1}$ not only depends on n_1, n_2 , but also depends on how well aligned $\Sigma^{(1)}$ and $\Sigma^{(2)}$ are. To capture this alignment, we introduce the covariate shift matrix $M := \Sigma^{(1)\frac{1}{2}}\Sigma^{(2)-\frac{1}{2}}$. Let $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_p$ be the singular values of M in descending order. Our first main result is the following theorem on the variance limit, which characterizes the exact dependence of L_{var} on the singular values of M and the sample sizes n_1 and n_2 .

Theorem 3 (Precise estimates under covariate shifts) *Under Assumption 1, for any small constant $c > 0$, with high probability over the randomness of $(X^{(1)}, X^{(2)})$, the following holds:*

$$\left| L_{\text{var}} - \frac{\sigma^2}{n_1 + n_2} \text{Tr} \left[\left(\alpha_1 M^\top M + \alpha_2 \text{Id}_{p \times p} \right)^{-1} \right] \right| \leq \sigma^2 p^{\frac{1}{2}} (n_1 + n_2)^{\frac{2}{p} - \frac{3}{2} + c}, \quad (3.1)$$

where α_1 and α_2 are the unique positive solutions of the following system of equations:

$$\alpha_1 + \alpha_2 = 1 - \frac{p}{n_1 + n_2}, \quad \alpha_1 + \frac{1}{n_1 + n_2} \sum_{i=1}^p \frac{\lambda_i^2 \alpha_1}{\lambda_i^2 \alpha_1 + \alpha_2} = \frac{n_1}{n_1 + n_2}. \quad (3.2)$$

Recalling that $\varphi > 4$, for a small enough constant c , equation (3.1) characterizes the limit of L_{var} with an error term that is smaller than the deterministic leading term by a factor $o(p^{-1/2})$.

Theorem 3 generalizes a classical result in multivariate statistics to the sum of two independent sample covariance matrices with arbitrary covariate shift: with high probability over the randomness of $X^{(2)}$,

$$\text{Tr} \left[\Sigma^{(2)} (X^{(2)\top} X^{(2)})^{-1} \right] = \frac{p}{n_2 - p} + O \left(p^{\frac{1}{2}} n^{\frac{2}{\varphi} - \frac{3}{2} + c} \right). \quad (3.3)$$

To see this, note that (3.3) is a special case of Theorem 3 with $n_1 = 0$. For this case, equation (3.2) implies that $\alpha_1 = 0$ and $\alpha_2 = (n_2 - p)/n_2$. Therefore,

$$\frac{1}{n_1 + n_2} \text{Tr} \left[\left(\alpha_1 M^\top M + \alpha_2 \text{Id}_{p \times p} \right)^{-1} \right] = \frac{p}{n_2 - p}.$$

The result (3.3) has a rich history in the literature of random matrix theory. When $Z^{(2)}$ is a Gaussian random matrix, the limit (without a convergence rate) follows from properties of the inverse Wishart distribution. Otherwise, one can use the Stieltjes transform method to derive the estimate (Lemma 3.11, Bai and Silverstein (2010)). The estimate (3.3) with sharp convergence rates can be found in Theorem 2.4 (Bloemendal et al., 2014).

Theorem 3 follows from a sharp *local law* on the resolvent of $\hat{\Sigma}$, which we will establish based on recently developed techniques in the random matrix theory literature (Knowles and Yin, 2016). We refer the reader to Appendix B.3 for more details.

As a remark, in the covariate shift setting, since $\beta^{(1)} = \beta^{(2)}$, we only need to study hard transfer, which does not add any model adjustment to the source task.

3.2 Illustrative Examples

The formulas in the above result are complex. To interpret their scaling, we provide two illustrative examples of Theorem 3, based on which we revisit the effect of covariate shift upon transfer. Notice that the answer to this question is not always clear in the transfer learning literature due to the intricacy of covariate shifts. Thus, we will leverage the precise asymptotics we have developed to provide a preliminary study of this question. Our goal is to provide several illustrative examples that are not meant to be exhaustive. Also note that in the case where both tasks share the same linear model ($\beta^{(1)} = \beta^{(2)}$), HPS always incurs a *lower* risk than OLS. Therefore, the transfer effect is strictly *positive* in the covariate shift setting, and our comparison will be between HPS estimators with different degrees of covariate shift.

In the first example, we show that the impact of covariate shift depends on the ratio between n_1 and n_2 . In this example, we demonstrate that having greater covariate shifts between source and target tasks can help improve estimation. Let p be an even integer, and let

$$g(M) = \frac{\sigma^2}{n_1 + n_2} \text{Tr} \left[\left(\alpha_1 M^\top M + \alpha_2 \text{Id}_{p \times p} \right)^{-1} \right].$$

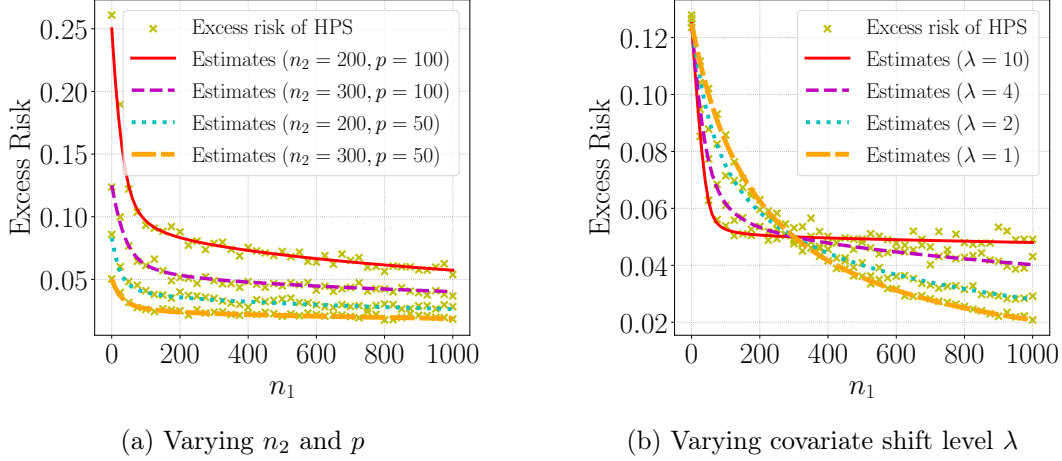


Figure 2: Illustration of Theorem 3 under various sample sizes and dimensions. Figure 2a shows that the variance estimate can accurately fit the excess risk of HPS. This simulation fixes $\lambda = 4$ and varies n_1, n_2, p . Figure 2b shows the dichotomy in Proposition 4. When $n_1 < n_2$, the lowest risk is achieved by transferring from a covariate-shifted data set. When $n_1 \geq n_2$, the lowest risk is achieved by transferring from a data set with $\Sigma^{(1)} = \Sigma^{(2)}$. This simulation fixes $p = 100, n_2 = 300$ and varies n_1, λ . Both simulations use $\sigma = 1/2$. The covariates are sampled independently from an isotropic Gaussian distribution.

We compare $g(M)$ for different M matrices within a set \mathcal{S} such that $\lambda_{p+1-i} = 1/\lambda_i$ for $i = 1, 2, \dots, p/2$. Therefore, every matrix in \mathcal{S} is normalized with $\det(M) = 1$.² In particular, $M = \text{Id}_{p \times p}$ represents two tasks with the same population covariance matrix. We show that when $n_1 < n_2$, the presence of covariate shifts can be beneficial for HPS. On the other hand, when $n_1 > n_2$, the presence of covariate shifts can hurt the performance of hard transfer.

Proposition 4 *Within the set $M \in \mathcal{S}$, the following dichotomy regarding $g(\text{Id}_{p \times p})$ and $g(M)$ holds: i) When $n_1 < n_2$, $g(M) \leq g(\text{Id}_{p \times p})$ for any $M \in \mathcal{S}$. ii) When $n_1 \geq n_2$, $g(\text{Id}_{p \times p}) \leq g(M)$ for any $M \in \mathcal{S}$.*

Figure 2 illustrates a special case where $\lambda_1 = \dots = \lambda_{p/2} = \lambda > 1$ and $\lambda_{p/2+1} = \dots = \lambda_p = 1/\lambda$. Thus, higher λ corresponds to a worse covariate shift. We plot the theoretical estimate using $g(M)$ and the excess risk using equation (2.10). Our theoretical estimate in Theorem 3 matches the empirical risk incredibly well. As a result, we indeed observe the dichotomy in Proposition 4. Furthermore, for larger λ , the excess risk of HPS decreases more slowly, indicating a worse “rate of transfer” from task one. As a remark, impossibility results for transfer learning under covariate shift have been observed for classification (Ben-David et al., 2010b). Our results are stated in the high-dimensional regression setting. A related

2. Note that the covariate shift intensifies when we scale M by a larger scalar $C > 1$. Moreover, intuitively, $g(CM)$ generally would decrease as the scalar C increases. Thus, to ensure a fair comparison between scenarios with and without covariate shift, we fix a normalization for the M matrix by requiring that $\det(M) = 1$.

result regarding the effect of covariate shift on transfer has also appeared in the work of Lei et al. (2021). They show that the minimax estimator depends on the singular values of the M matrix.

In the second example, we consider the case where $n_1 \geq n_2$. A typical scenario in transfer learning is that the source data set is much larger than the target data set. We show that when n_1 is larger than n_2 times a sufficiently large constant, $M = \text{Id}_{p \times p}$ indeed minimizes $g(M)$ within a bounded set of matrices whose determinants are equal to one.

Proposition 5 *Let $c \in (0, 1)$ be a fixed constant. Let \mathcal{S}' be a set of matrices such that for any $M \in \mathcal{S}'$: (i) $\det(M) = 1$; (ii) the eigenvalues $\lambda_1^2, \dots, \lambda_p^2$ of $M^\top M$ are bounded between c and $1/c$. If $n_1 - p \geq c^{-1}n_2$, then g achieves the global minimum at $\text{Id}_{p \times p}$, i.e.,*

$$g(\text{Id}_{p \times p}) \leq g(M), \text{ for any } M \in \mathcal{S}'. \quad (3.4)$$

We remark that the condition that the eigenvalues $\lambda_1^2, \dots, \lambda_p^2$ are bounded would be necessary for equation (3.4) to hold. For example, choose $M = M_x$ with $1 < \lambda_1^2 = x$ and $\lambda_2^2 = \dots = \lambda_p^2 = x^{-(p-1)^{-1}}$. Let $(\alpha_1(x), \alpha_2(x))$ be the solution to the system of equations

$$\alpha_1 + \alpha_2 = 1 - \frac{p}{n_1 + n_2}, \alpha_1 + \frac{1}{n_1 + n_2} \left(\frac{(p-1)x^{-(p-1)^{-1}}\alpha_1}{x^{-(p-1)^{-1}}\alpha_1 + \alpha_2} + \frac{x\alpha_1}{x\alpha_1 + \alpha_2} \right) = \frac{n_1}{n_1 + n_2}.$$

From this equation, we obtain that $\alpha_1(x) \rightarrow \frac{n_1-1}{n_1+n_2}$ and $\alpha_2(x) \rightarrow \frac{n_2+1-p}{n_1+n_2}$ as $x \rightarrow \infty$. Thus,

$$\begin{aligned} \lim_{x \rightarrow \infty} g(M_x) &= \lim_{x \rightarrow \infty} \frac{\sigma^2}{n_1 + n_2} \left(\frac{p-1}{x^{-(p-1)^{-1}}\alpha_1(x) + \alpha_2(x)} + \frac{1}{x\alpha_1(x) + \alpha_2(x)} \right) \\ &= \frac{\sigma^2(p-1)}{n_2 + 1 - p} \geq \frac{\sigma^2 p}{n_1 + n_2 - p} = g(\text{Id}_{p \times p}). \end{aligned}$$

Similar arguments, by comparing $g(\text{Id}_{p \times p})$ with $g(M)$ for certain M on the boundary of \mathcal{S}' , show that $(n_1 - p)/n_2$ should have a proper lower bound for $g(\text{Id}_{p \times p})$ to be the global minimum. We will use an induction argument to show that $(n_1 - p)/n_2 \geq c^{-1}$ is a sufficient condition. The proof of Propositions 4 and 5 can be found in Appendix B.1 and Appendix B.2, respectively.

Note that in both Propositions 4 and 5, our goal is to illustrate the exact asymptotics of Theorem 3, since the formula is intricate. Relatedly, Dhifallah and Lu (2021) have also conducted a precise analysis of the hard transfer and soft transfer estimation. These precise estimates seem necessary when examining the effects of transfer learning and studying negative transfer, as we need to compare the performance of joint learning with that of single-task learning. The main intuition behind the above examples is that when the source and target tasks have complementary spectra, combining both tasks can help transfer by increasing the smallest eigenvalue in the combined covariance matrix $\hat{\Sigma}$. Later in Section 5.1, we will extend these examples to a setting involving both covariate and model shifts (see Figures 4a and 4b for the illustration).

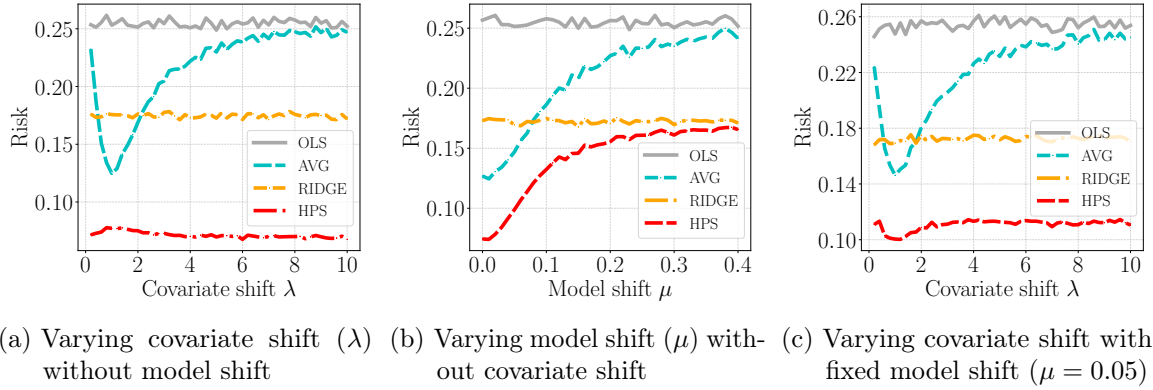


Figure 3: Comparing the excess risks of several estimators, showing that HPS achieves the lowest excess risk compared to OLS, AVG, and RIDGE in various settings with a combination of covariate and model shifts.

3.3 Numerical Comparisons

We show that the HPS estimator provides strong empirical performance compared to several transfer learning estimators. For this comparison, we consider the OLS estimator and the ridge estimator (RIDGE), given by:

$$\hat{\beta}_2^{\text{RIDGE}}(k) = (X^{(2)\top} X^{(2)} + k \cdot \text{Id}_{p \times p})^{-1} X^{(2)\top} Y^{(2)}.$$

We also consider an averaging estimator (AVG), which takes a convex combination of their OLS estimators:

$$\hat{\beta}_2^{\text{AVG}}(b) = b \cdot \hat{\beta}_1^{\text{OLS}} + (1 - b) \cdot \hat{\beta}_2^{\text{OLS}}.$$

The parameters b and k are optimized using a validation set independent of the training set. For HPS, a weight parameter is added for each task, along with a ridge penalty. Both of these are optimized using the validation set. We remark that all the high-dimensional asymptotic limits for HPS can be extended to this setting.

For the comparison in Figure 3a, we sample the entries of $Z^{(1)}$ and $Z^{(2)}$ from Gaussian distributions and consider the same covariate shift as in Figure 2, that is, $\lambda_1 = \dots = \lambda_{\lfloor \frac{p}{2} \rfloor} = \lambda > 1$ and $\lambda_{\lfloor \frac{p}{2} \rfloor + 1} = \dots = \lambda_p = \lambda^{-1}$. Then, we compare the excess risk of equation (2.6) for different estimators under varying levels of λ . For this simulation, we set $p = 50, n_1 = n_2 = 100$, and $\sigma = \frac{1}{2}$. Figure 3a shows that HPS can consistently outperform OLS, RIDGE, and AVG in this simulation.

In Figures 3b and 3c, we show that similar results continue to hold in more general settings with a combination of covariate and model shifts. In particular, we generate covariate-shifted features and different linear models for the two tasks. We take the average of 100 random seeds because of high variances due to small sample sizes. We find that HPS achieves strong empirical performance under various settings of covariate and model shifts. The experiment code for reproducing these simulation results can be found at https://github.com/VirtuosoResearch/Transfer_learning_random_matrix_simulations.

4. Model Shifts

This section presents precise estimates of the bias and variance of HPS under model shifts, assuming no covariate shifts are present. With these results, we can connect the bias-variance decomposition to transfer effects by comparing the bias and variance of transfer learning estimators with those of single-task learning. In particular, we will present a detailed analysis in a random-effects model, which uncovers a phase transition in transfer learning. This motivates us to design an adjusted HPS procedure that downsizes the source task when the model shift is too high. We show that the performance of this procedure can match a minimax lower bound for estimating one task from the samples of two linear regression tasks.

4.1 Precise Estimates

In this subsection, we study the impact of model shifts on the performance of the HPS estimator when there is no covariate shift. Particularly, both tasks have the same population covariance matrix ($\Sigma^{(1)} = \Sigma^{(2)}$) but follow different linear models ($\beta^{(1)} \neq \beta^{(2)}$). The following result states the exact asymptotic limit of the excess risk of HPS in this case.

Theorem 6 (Precise estimates under model shifts) *Let $\Sigma^{(1)} = \Sigma^{(2)}$ be the same. Under Assumption 1, suppose $Z^{(1)}$ and $Z^{(2)}$ are both Gaussian random matrices. Then, for any small constant $c > 0$, with high probability over the randomness of $Z^{(1)}, Z^{(2)}$, the following estimates hold:*

$$L_{\text{var}} = \sigma^2 L_1 + O\left(\sigma^2 \sqrt{p}(n_1 + n_2)^{-\frac{3}{2}+c}\right), \quad (4.1)$$

$$L_{\text{bias}} = \left\| \Sigma^{(1)\frac{1}{2}} (\beta^{(1)} - \beta^{(2)}) \right\|^2 L_2 + O\left(\frac{p^{-\frac{1}{2}+c} n_1^2}{(n_1 + n_2)^2} \left\| \beta^{(1)} - \beta^{(2)} \right\|^2\right). \quad (4.2)$$

Above, L_1 and L_2 are defined as

$$L_1 = \frac{p}{n_1 + n_2 - p}, \quad L_2 = \frac{n_1^2(n_1 + n_2 - p) + pn_1n_2}{(n_1 + n_2)^2(n_1 + n_2 - p)}.$$

Combining equations (6) with (2) results in an exact estimate for the excess risk of HPS under model shifts. The variance estimate (4.1) is a special case of Theorem 3 with $M = \text{Id}_{p \times p}$ and $\varphi \rightarrow \infty$ (since Gaussian random variables have bounded moments up to any order).

The bias estimate (4.2) requires the assumption that both $Z^{(1)}$ and $Z^{(2)}$ are Gaussian. We briefly describe the proof ideas. Let $\mathbf{v} = (\Sigma^{(1)})^{1/2} (\beta^{(1)} - \beta^{(2)})$. Then, the bias formula (2.9) can be written as:

$$L_{\text{bias}} = \mathbf{v}^\top Z^{(1)\top} Z^{(1)} \left(Z^{(1)\top} Z^{(1)} + Z^{(2)\top} Z^{(2)} \right)^{-2} Z^{(1)\top} Z^{(1)} \mathbf{v}.$$

Since both $Z^{(1)}$ and $Z^{(2)}$ are Gaussian random matrices, the distributions of $Z^{(1)\top} Z^{(1)}$ and $Z^{(2)\top} Z^{(2)}$ are rotation-invariant. Thus, the following (approximate) identity holds up to a

small error:

$$\begin{aligned} L_{\text{bias}} &\approx \frac{\|\mathbf{v}\|^2}{p} \text{Tr} \left[\left(Z^{(1)\top} Z^{(1)} \right)^2 \left(Z^{(1)\top} Z^{(1)} + Z^{(2)\top} Z^{(2)} \right)^{-2} \right] \\ &= \frac{\|\mathbf{v}\|^2}{p} \left. \frac{d}{dx} \right|_{x=0} \text{Tr} \left[\left(Z^{(1)\top} Z^{(1)} + x \left(Z^{(1)\top} Z^{(1)} \right)^2 + Z^{(2)\top} Z^{(2)} \right)^{-1} \right]. \end{aligned} \quad (4.3)$$

Due to the rotation invariance, the asymptotic limit of equation (4.3) is determined by the free addition of $Z^{(1)\top} Z^{(1)} + x(Z^{(1)\top} Z^{(1)})^2$ and $Z^{(2)\top} Z^{(2)}$. Building on free probability techniques (see, e.g., Nica and Speicher (2006); Bao et al. (2017)), an explicit formula for this free addition can be derived for any x around zero. This observation allows us to derive equation (4.2) by taking the derivative with respect to x as in equation (4.3). More details including the complete proof of Theorem 6 are presented in Appendix C.1.

Remark 7 *We conjecture that the bias limit (4.2) is still the exact asymptotic form even if $Z^{(1)}$ and $Z^{(2)}$ are non-Gaussian random matrices. One approach to show this is using the local laws for polynomials of random matrices. This requires checking certain technical regularity conditions, which are left as an open question for future work.*

4.2 Examples and Adjustments

We will consider a random-effects model (Dobriban and Sheng, 2020). Each $\beta^{(i)}$ consists of two components, in this case, one shared by all tasks and one that is task-specific. Let β_0 be the shared component and γ_i be the i -th task-specific component. For any i , the i -th model vector is equal to $\beta^{(i)} = \beta_0 + \gamma_i$. The entries of the task-specific component γ_i are drawn independently from a Gaussian distribution with mean zero and variance $p^{-1}\mu^2$, for a parameter $\mu > 0$. In expectation, the Euclidean distance between the two model vectors is equal to $2\mu^2$.

Based on Theorem 6, we present a precise analysis of information transfer in this random-effects model using the precise limits. The proof of the following proposition will be presented in Appendix C.2.

Proposition 8 (Phase transition in the random-effects model) *Under the assumptions of Theorem 6, suppose the random-effect model applies and $n_2 \geq 3p$. For any small constant $c > 0$, the following statements hold with high probability over the randomness of training samples and model vectors.*

1. If $\mu^2 \frac{\text{Tr}[\Sigma^{(1)}]}{p} \leq \frac{\sigma^2 p}{2(n_2 - p)}$, then the transfer effect is always positive:

$$L(\hat{\beta}_2^{\text{HPS}}) \leq \left(1 + O(p^{-1/2+c})\right) \cdot L(\hat{\beta}_2^{\text{OLS}}). \quad (4.4)$$

2. If $\frac{\sigma^2 p}{2(n_2 - p)} < \mu^2 \frac{\text{Tr}[\Sigma^{(1)}]}{p} < \frac{\sigma^2 n_2}{2(n_2 - p)}$, then there exists a deterministic value $n_0 > 0$ (which may not be an integer) such that if $n_1 \leq n_0$, then equation (4.4) holds; otherwise if $n_1 > n_0$, then

$$L(\hat{\beta}_2^{\text{OLS}}) \leq \left(1 + O(p^{-1/2+c})\right) \cdot L(\hat{\beta}_2^{\text{HPS}}). \quad (4.5)$$

3. If $\mu^2 \frac{\text{Tr}[\Sigma^{(1)}]}{p} \geq \frac{\sigma^2 n_2}{2(n_2 - p)}$, then equation (4.5) holds for any $n_1 > p$.

Figure 1 illustrates phases 2 and 3 of Proposition 8 for multiple values of μ . We plot our estimate vs. the excess risk using equation (2.7). This simulation shows that our estimate accurately matches the bias and variance. For this simulation, we set $\Sigma^{(1)} = \Sigma^{(2)} = \text{Id}_{p \times p}$, and fix $p = 100$, $n_2 = 300$, $\sigma = 1/2$ while varying n_1 and μ . Notice that the above result requires $n_2 \geq 3p$; when $p < n_2 < 3p$, the threshold conditions can be derived with a similar analysis.

Figure 1 suggests that a natural training procedure is to downsample the sample size of the source task until the transfer learning performance starts to drop. In the setting of Figure 1, this procedure will terminate at the optimal value of n_1 . We can rigorously show this *dichotomy* in the setting of Proposition 8 (see Appendix C.2 for further details):

1. If $\mu^2 \frac{\text{Tr}[\Sigma^{(1)}]}{p} \geq \frac{\sigma^2 n_2}{2(n_2 - p)}$, then HPS gets strictly worse as n_1 increases in $[0, \infty)$.
2. If $\mu^2 \frac{\text{Tr}[\Sigma^{(1)}]}{p} < \frac{\sigma^2 n_2}{2(n_2 - p)}$, then there exists $n_0 > 0$ such that HPS first improves as n_1 increases in $[0, n_0]$ but turns worse as soon as $n_1 > n_0$.

The above claim demonstrates that a simple adjustment to the HPS estimator can be made to reduce negative transfers. Essentially, we will choose between HPS and OLS, whichever performs better for the target task. In other words, if μ is too large, we will throw away the source task data. Next, we will compare the performance of this adjusted HPS estimation with a minimax lower bound, which shows that our estimate for the adjusted HPS estimator matches the lower bound.

4.3 A Matching Lower Bound

Lastly, we complement our study in the model shift setting with a minimax lower bound. Suppose we are trying to estimate an unknown parameter denoted as $\beta^{(2)}$. We are given n_2 samples drawn from a linear parametric model, following $\beta^{(2)}$, with isotropic Gaussian covariates $X^{(2)}$, contaminated by Gaussian noise with variance σ^2 . Then, we are given n_1 samples from another linear parametric model, following $\beta^{(1)}$, again with isotropic Gaussian covariates $X^{(1)}$, contaminated by Gaussian noise with variance σ^2 . The parameter vectors belong to the set

$$\Theta(\mu) = \{\beta^{(1)} \in \mathbb{R}^p, \beta^{(2)} \in \mathbb{R}^p : \|\beta^{(1)} - \beta^{(2)}\| \leq \mu, \|\beta^{(1)}\| \leq 1, \|\beta^{(2)}\| \leq 1\}.$$

Note that our proof can be readily extended to cases with anisotropic Gaussian covariates as in equation (2.1), and the length constraints on $\beta^{(1)}$ and $\beta^{(2)}$ can be replaced with any other constant.³

Let $\hat{\beta}$ be any estimation procedure that, given the above $n_1 + n_2$ samples, produces an estimate of the unknown vector $\beta^{(2)}$. We prove the following minimax rate on the estimation error of $\hat{\beta}$.

3. In particular, to extend the proof to the case with covariate shifts, one can replace the bound in equation (C.25), Appendix C.3, with different values of c for each respective $\Sigma^{(1)}$ and $\Sigma^{(2)}$. The rest of the covering arguments will still go through. Then, one would set $\mu = 0$ in the specification of $\Theta(\mu)$.

Theorem 9 *In the setting described above, let $\hat{\beta}$ be a fixed estimation procedure. Assume that $n_1 \geq (1 + \tau)p$ and $n_2 \geq (1 + \tau)p$ for a constant $\tau > 0$. For any $(\beta^{(1)}, \beta^{(2)})$ within the set $\Theta(\mu)$, we have that*

$$\inf_{\hat{\beta}} \sup_{\Theta(\mu)} \mathbb{E} \left[\frac{1}{n_2} \left\| \tilde{X}^{(2)} \left(\hat{\beta} - \beta^{(2)} \right) \right\|^2 \right] \geq c \left(\min \left(\frac{n_1^2 \mu^2}{(n_1 + n_2)^2}, \frac{\sigma^2 p}{n_2} \right) + \frac{\sigma^2 p}{n_1 + n_2} \right), \quad (4.6)$$

where the expectation is over the randomness of $X^{(1)}, X^{(2)}, Y^{(1)}, Y^{(2)}$ and an independently drawn $\tilde{X}^{(2)}$ that follows the same distribution as $X^{(2)}$. c is a fixed constant that does not grow with p .

The lower bound in the right-hand side of equation (4.6) involves two parts:

- For the first part, if μ is large, the source samples are not helpful, so the rate $\frac{\sigma^2 p}{n_2}$ is the OLS rate using only the target task samples. If μ is small and n_1 is large, then the rate μ^2 appears if we use the OLS estimator for the source task.
- For the second part, $\frac{\sigma^2 p}{n_1 + n_2}$ is the rate in the case without model drift, i.e., $\beta^{(1)} = \beta^{(2)}$.

We now compare the minimax lower bound in equation (4.6) to the rates in Theorem 6 and show that for an adjusted HPS estimation, the two bounds match up to each other up to constant factors. Let $\|\beta^{(1)} - \beta^{(2)}\| = \mu$. In the setting of Theorem 6, one can see that the bias in equation (4.1) and the variance in equation (4.2) are given by the following orders:

$$L_{\text{bias}} = \Theta \left(\frac{n_1^2 \mu^2}{(n_1 + n_2)^2} \right), \quad L_{\text{var}} = \Theta \left(\frac{\sigma^2 p}{n_1 + n_2} \right). \quad (4.7)$$

These estimates hold more generally for the bias-variance decomposition in Lemma 2 (beyond the setting of Theorem 6); See the derivation in Appendix A.3 (equations (A.26) and (A.27)). In particular, we can derive these estimates using concentration estimates on the eigenvalues of random matrices.

If $\mu^2 \leq \frac{\sigma^2 p (n_1 + n_2)^2}{n_1^2 n_2}$, we can see that $L_{\text{bias}} + L_{\text{var}}$ already matches the lower bound in equation (4.6). Then, by invoking Lemma 2, we conclude that the excess risk of HPS matches the minimax lower bound.

Otherwise, if $\mu^2 > \frac{\sigma^2 p (n_1 + n_2)^2}{n_1^2 n_2}$, the adjusted HPS procedure—which throws away all the source task data—gives an excess risk at most $\frac{\sigma^2 p}{n_2 - p}$ (plus small concentration errors). On the other hand, the lower bound in equation (4.6) becomes $O(\sigma^2 p (\frac{1}{n_1} + \frac{1}{n_1 + n_2}))$. Since n_1 and n_2 are within a constant away of each other, the rate of the adjusted HPS still matches the lower bound in the case that μ is larger than the above threshold. This concludes our study of the model shift setting. The proof of this minimax lower bound is based on a covering argument (see, e.g., Chapter 12, Zhang (2023)), and the details can be found in Appendix C.3.

5. Extensions

This section presents two extensions to reinforce our findings from previous sections. First, we consider a setting that involves both covariate shifts and model shifts, and our goal is to

show that the insights obtained in Section 3 can still be extended to this setting. Second, we consider a setting involving multiple regression tasks with shared feature covariates, but following different linear models. The goal is to leverage these shared structures to help estimate the target task. See, e.g., Ando and Zhang (2005); Wang et al. (2016) for some prior works that study similar settings. In particular, we can still uncover a phase transition similar to our findings from Section 4.

5.1 A Setting with Both Covariate and Model Shifts

We first consider a setting with both covariate and model shifts. Define a function that depends on the source task's covariance:

$$g_0(x) = \frac{1}{n_1} \text{Tr} \left[\Sigma^{(1)} (\text{Id}_{p \times p} + x \cdot \Sigma^{(1)})^{-1} \right] - \frac{1}{x}. \quad (5.1)$$

Let $g'_0(x)$ be the derivative of $g_0(x)$. It can be shown that there is a unique positive solution to the equation

$$(1+x)g_0(x) = -n_1^{-1}(n_1 + n_2 - p). \quad (5.2)$$

Let y_0 denote this solution. The following theorem provides the exact asymptotic variance and bias limits when the population covariance of the target task is isotropic.

Theorem 10 (Precise estimates under covariate and model shifts) *Under Assumption 1, suppose that $\Sigma^{(2)} = \text{Id}_{p \times p}$ and $Z^{(1)}$ and $Z^{(2)}$ are both Gaussian random matrices. Suppose further that the random-effects model applies. Let*

$$f_1 = \frac{n_1}{p} y_0 + \frac{n_1 - p}{p \cdot g_0(y_0)}, \quad f_2 = \frac{n_1}{p \cdot g'_0(y_0)} - \frac{n_1 - p}{p \cdot (g_0(y_0))^2}, \quad \text{and } f_3 = -g_0(y_0). \quad (5.3)$$

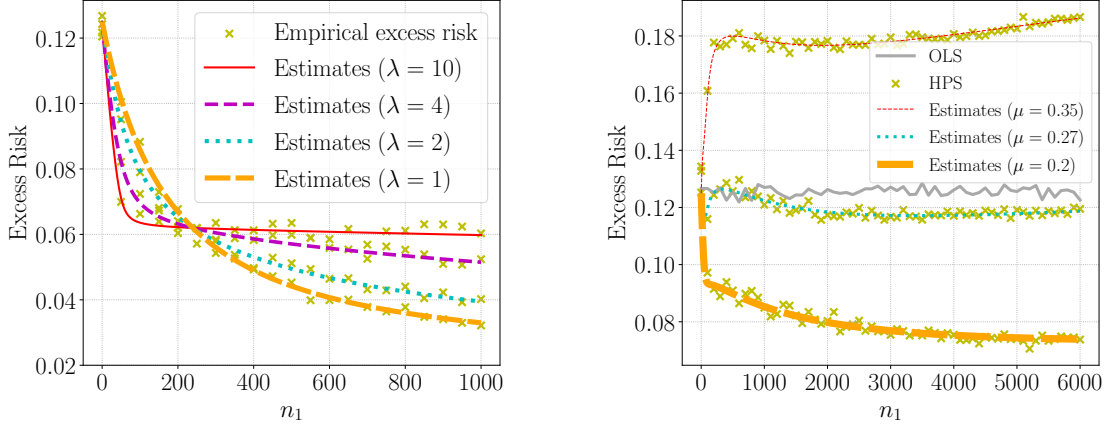
Then, for any small constant $c > 0$, with high probability over the randomness of training samples and model vectors, the following estimates hold:

$$L_{\text{var}} = \left(1 + O(p^{-\frac{1}{2}+c})\right) \sigma^2 p n_1^{-1} f_1, \\ L_{\text{bias}} = \left(1 + O(p^{-\frac{1}{2}+c})\right) 2\mu^2 \frac{1 - 2f_1 f_3 + f_2 f_3^2}{1 - n_2^{-1} p f_2 f_3^2}.$$

When the population covariance of the source task is also isotropic (i.e., $\Sigma^{(1)} = \text{Id}_{p \times p}$), we can verify that the above result is consistent with the variance and bias limits in Theorem 6. The proof of Theorem 10 uses similar techniques as Theorem 6 (see Appendix C.4).

5.1.1 THE ANISOTROPIC CASE

Theorem 10 states exact estimates when the target task's population covariance matrix is isotropic. In the general anisotropic case, it is possible to estimate the bias formula (2.9) by approximating the source task's sample covariance matrix with its expectation. It turns out that this renders the analysis of the bias formula similar to that of the variance formula in Theorem 3. However, this approximation results in an error term that decreases



(a) Varying covariate shift with fixed model shift (b) Varying model shift with fixed covariate shift

Figure 4: Showing that our findings from Section 3 continue to hold under the presence of both covariate and model shifts in the data.

to zero (relative to the estimated term) as n_1/p increases to infinity; see the right-hand side of equation (C.30). Thus, the accuracy of the estimate increases as n_1/p increases. This approximation is motivated by transfer learning scenarios where the source task is significantly larger than the target task. The statement of this result for the anisotropic case (Theorem 44), together with a complete proof, is described in Appendix C.5.

Remark 11 *In order to obtain an exact asymptotic limit of L_{bias} under general forms of $\Sigma^{(2)}$, one needs to study the singular values and singular vectors of an asymmetric random matrix*

$$(X^{(1)\top} X^{(1)})^{-1} X^{(2)\top} X^{(2)} + \text{Id}_{p \times p}. \quad (5.4)$$

The eigenvalues of this matrix have been studied in the name of Fisher matrices (Zheng et al., 2017). Notice, however, that its singular values are different from its eigenvalues because of asymmetry. This asymmetry makes it difficult to characterize the asymptotic behavior of its resolvent. We leave deriving a precise estimate of the anisotropic case to future work.

5.2 Numerical Comparisons

We complement our theoretical analysis of HPS with empirical evaluations. Figure 4 illustrates the result for a setting where half of the eigenvalues of $\Sigma^{(1)}$ are equal to $\lambda > 1$ and the other half are equal to $1/\lambda$. Our theoretical estimates consistently match the empirical risks. Figure 4a shows a similar dichotomy as in Figure 2b. We fix model shift $\mu = 0.1$ while varying covariate shift λ for each curve. Figure 4b illustrates different transfer effects as in Figure 1. We fix the level of covariate shift $\lambda = 4$ while varying μ for each curve. Both simulations use $p = 100$, $n_2 = 300$, and $\sigma = 1/2$.

5.3 Multitask Learning with Shared Features and Different Models

Next, we extend our findings from Section 4 to multiple source tasks involving model shifts. Suppose there are t tasks whose feature covariates are all equal to $X \in \mathbb{R}^{n \times p}$. The label vector of the i -th task follows a linear model with an unknown p dimensional vector $\beta^{(i)}$, for $i = 1, 2, \dots, t$:

$$Y^{(i)} = X\beta^{(i)} + \varepsilon^{(i)}. \quad (5.5)$$

This setting may also be referred to as multiple label regression, and variants of this setting have been studied in the prior literature (Kolar et al., 2011; Wang et al., 2016). Similar to the two-task case, one of the tasks is viewed as the primary target task of interest, while the others are used as source tasks that aid in learning.

We assume that $X = Z\Sigma^{\frac{1}{2}} \in \mathbb{R}^{n \times p}$ is a random matrix satisfying the same assumption as $X^{(2)}$ in Assumption 1, and $\varepsilon^{(i)} \in \mathbb{R}^n$, $i = 1, 2, \dots, t$, are independent random vectors, each of which is independent of X and satisfies the same assumption as $\varepsilon^{(2)}$ in Assumption 1. Furthermore, each $\beta^{(i)} \in \mathbb{R}^p$ is a (random or deterministic) vector independent of any other $\beta^{(j)}$ for $j \neq i$, the matrix X , and $\varepsilon^{(j)}$ for all $j = 1, \dots, t$. The sample size n satisfies $1 + \tau \leq \frac{n}{p} \leq p^{\tau-1}$.

We combine the samples from all the tasks with a shared feature matrix B and a task-specific prediction vector for each task. Specifically, let $B \in \mathbb{R}^{p \times r}$ denote the shared feature matrix and let $A = [A_1, A_2, \dots, A_t] \in \mathbb{R}^{r \times t}$ denote the combined prediction variables. We note that, unlike the two-task setting concerning covariate shifts, both A and B are part of the model in the multi-task setting. The reason is twofold. First, this paper focuses on the under-parameterized setting. Thus, in the two-task setting, one can show that the optimal rank of B would be one, i.e., $r = 1$, in which case B becomes a scalar. Second, to tackle more than two tasks, we need to incorporate more of the structural information from the other tasks into the model (Ando and Zhang, 2005), which is encoded as the low-rank structures of A and B . By contrast, in the two-task setting, we can focus on the setting where we have a shared parameter vector for both tasks.

We can now write down the loss objective as follows:

$$\ell(A, B) = \sum_{j=1}^t \left\| XBA_j - Y^{(j)} \right\|^2, \quad (5.6)$$

In particular, we focus on the case where $r < t$. Otherwise, if $r \geq t$, the problem reduces to single-task learning (for reference, see Proposition 1 from Wu et al. (2020)).

Let (\hat{A}, \hat{B}) denote the global minimizer of $\ell(A, B)$. In particular, we compute the global minimizer of $\ell(A, B)$ by first solving \hat{B} as a function of \hat{A} —this turns out to be $(X^\top X)^{-1}X^\top Y$ multiplied by $A^\top(AA^\top)^+$, with $(AA^\top)^+$ denoting the pseudo-inverse of AA^\top . Next, we could find the optimal \hat{A} by taking the rank- r SVD of $Y^\top X(X^\top X)^{-1}X^\top Y$ to get the leading r left singular vectors as $U_r \in \mathbb{R}^{t \times r}$; See Appendix D for the derivation. Then, we can obtain $\hat{B}\hat{A}$ as $(X^\top X)^{-1}X^\top YU_rU_r^\top$.

We now define the HPS estimator for task i as $\hat{\beta}_i^{\text{HPS}} = \hat{B}\hat{A}_i$, where \hat{A}_i denotes the i -th column of \hat{A} . The excess risk of $\hat{\beta}_i^{\text{HPS}}$ is equal to:

$$L_i(\hat{\beta}_i^{\text{HPS}}) = \left\| \Sigma^{\frac{1}{2}} \left(\hat{\beta}_i^{\text{HPS}} - \beta^{(i)} \right) \right\|^2. \quad (5.7)$$

5.3.1 RESULTS

We demonstrate that in the aforementioned multi-task setting, hard transfer identifies the optimal rank- r approximation, which is utilized to share information across multiple tasks. To describe the result, we introduce several notations. Let $B^\star := [\beta^{(1)}, \beta^{(2)}, \dots, \beta^{(t)}] \in \mathbb{R}^{p \times t}$ be the matrix of concatenated model vectors. Let $A^\star A^{\star\top}$ be the best approximation of $B^{\star\top} \Sigma B^\star$ in the set of rank- r subspaces in \mathbb{R}^t :

$$A^\star := \underset{U \in \mathbb{R}^{t \times r}: U^\top U = \text{Id}_{r \times r}}{\text{argmax}} \quad \langle UU^\top, B^{\star\top} \Sigma B^\star \rangle, \quad (5.8)$$

where $\langle \cdot, \cdot \rangle$ denotes the Frobenius inner product between two matrices. Let $a_i^\star \in \mathbb{R}^r$ be the i -th column vector of $A^\star A^{\star\top}$. For any matrix X , let $\|X\|_2$ be its spectral norm and $\|X\|_F$ be its Frobenius norm. A precise estimate of the excess risk of HPS is stated below.

Theorem 12 *Suppose the setting described above holds. Let $r < t$ be a positive integer. Suppose the r -th largest eigenvalue λ_r of $B^{\star\top} \Sigma B^\star$ is strictly larger than its $(r+1)$ -th largest eigenvalue λ_{r+1} . Then, for $i = 1, \dots, t$ and any small constant $c > 0$, the following estimate holds with high probability over the randomness of the training samples:*

$$\left| L_i(\hat{\beta}_i^{\text{HPS}}) - L_i(B^\star a_i^\star) - \frac{\sigma^2 p \|a_i^\star\|^2}{n-p} \right| \leq \left(\|B^{\star\top} \Sigma B^\star\|_2 + \sigma^2 \right) \sqrt{\frac{\frac{\|B^{\star\top} \Sigma B^\star\|_2}{n^{\frac{1}{2}-\frac{2}{\varphi}-c}} + \frac{\sigma^2}{n^{\frac{1}{2}-c}}}{\lambda_r - \lambda_{r+1}}}. \quad (5.9)$$

In equation (5.9), $L_i(B^\star a_i^\star)$ is the asymptotic limit of the bias of HPS, while $\frac{\sigma^2 p}{n-p} \|a_i^\star\|^2$ is the asymptotic limit of the variance of HPS. The proof of Theorem 12 relies on a characterization of the global minimizer of $\ell(A, B)$ by rank- r SVD. The detailed proof can be found in Appendix D.1.

Remark 13 *The data model in equation (5.5) studied in this section has also been considered by several prior works in which there is a shared set of nonzero coordinates in the model vectors $\beta^{(1)}, \dots, \beta^{(t)}$ among all tasks (Kolar et al., 2011; Wang et al., 2016). The goal is to leverage this shared sparsity to recover the support set of nonzero coordinates. Duan and Wang (2023) examined an adaptive multi-task learning setting that does not require X to be shared among all tasks. They introduced an adaptive estimation procedure that can achieve minimax optimality. Notice that in both the work of Duan and Wang (2023) and the setting of Theorem 12, only model shifts are assumed to be present, and instead, covariate shifts are not imposed across different tasks. It is an interesting question to study a combination of covariate and model shifts in the multi-task learning setting. More recently, Li et al. (2023) introduce a task modeling framework to capture task relationships through influence functions. Li et al. (2024) design a boosting procedure on top of this influence estimation framework. It is an interesting question to theoretically formalize the benefit of boosting for multitask learning when tasks have strong negative interference. One promising direction is to revisit methods from active clustering and connect that task structures (Voevodski et al., 2012).*

5.3.2 ILLUSTRATIVE EXAMPLES

We illustrate the above result with the random-effects model. The model vector of every task is equal to a shared vector β_0 plus a task-specific component γ_i , for $i = 1, 2, \dots, t$:

$$\beta^{(i)} = \beta_0 + \gamma_i. \quad (5.10)$$

The entries of γ_i are drawn independently from a Gaussian distribution with mean zero and variance $p^{-1}\mu^2$. We study two natural questions:

1. What is the optimal rank r of the shared feature matrix B ?
2. When does HPS transfer positively to a particular task, depending on the sample size n and the model shift parameter μ ?

For the first question, since the setting is symmetric in the t tasks, we analyze the averaged limiting risk

$$g_r(n, \mu) := \frac{1}{t} \left\| \Sigma^{1/2} B^* (A^* A^{*\top} - \text{Id}_{t \times t}) \right\|_F^2 + \frac{\sigma^2 p}{n-p} \frac{r}{t},$$

which accurately approximates $t^{-1} \sum_{i=1}^t L_i(\hat{\beta}_i^{\text{HPS}})$ by Theorem 12. To see this, we notice that the average of the bias and variance limits in equation (5.9) is

$$\frac{1}{t} \sum_{i=1}^t \left(L_i(B^* a_i^*) + \frac{\sigma^2 p}{n-p} \|a_i^*\|^2 \right) = \frac{1}{t} \left\| \Sigma^{1/2} B^* (A^* A^{*\top} - \text{Id}_{t \times t}) \right\|_F^2 + \frac{\sigma^2 p}{n-p} \frac{r}{t}.$$

Above, we have used the matrix notation A^* to rewrite the bias component. Moreover, we apply the identity $\sum_{i=1}^t \|a_i^*\|^2 = \text{Tr}[(A^* A^{*\top})^2] = r$ to the variance component, since $A^{*\top} A^* = \text{Id}_{r \times r}$ following Definition (5.8). Then, the optimal rank of B can be identified by sweeping through $g_r(n, \mu)$ for different values of r .

For the second question, recall that by equation (3.3), the excess risk of the OLS estimator is given by $\frac{\sigma^2 p}{n-p}$. Thus, by comparing $g_r(n, \mu)$ to $\frac{\sigma^2 p}{n-p}$, the exact threshold between positive and negative transfer can be identified. This is stated as follows.

Proposition 14 (Phase transition in multi-task learning) *Consider the multi-task setting stated above. Assume further that the random-effects model in equation (5.10) holds. Then, for any small constant $c > 0$, the following claims hold with high probability over the randomness of the training data samples and the underlying linear models:*

1. When $\mu^2 > \left(1 + p^{-\frac{1}{2}+c}\right) \frac{\sigma^2 p^2}{(n-p)\text{Tr}[\Sigma]}$, then $g_r(n, \mu) \geq \frac{\sigma^2 p}{n-p}$ for any $1 \leq r < t$.
2. When $\mu^2 < \left(1 - p^{-\frac{1}{2}+c}\right) \frac{\sigma^2 p^2}{(n-p)\text{Tr}[\Sigma]}$, then $g_r(n, \mu)$ is minimized when $r = 1$. Furthermore, when $r = 1$, we have that $g_1(n, \mu) \leq \frac{\sigma^2 p}{n-p}$.

Figure 5 illustrates the above result with ten tasks of dimension $p = 100$ and noise variance $\sigma = \frac{1}{2}$. We plot the theoretical estimate using $g_r(n, \mu)$ and the empirical value of bias plus variance of $\hat{\beta}_t^{\text{HPS}}$ (equation (D.16) in Section D). Figure 5a shows that the performance of HPS is optimized when $r = 1$. Figure 5b shows different transfer effects by varying n and μ in the multi-task setting. The results under different values of μ also match the conditions in Proposition 14. The proof of Proposition 14 can be found in Appendix D.2.

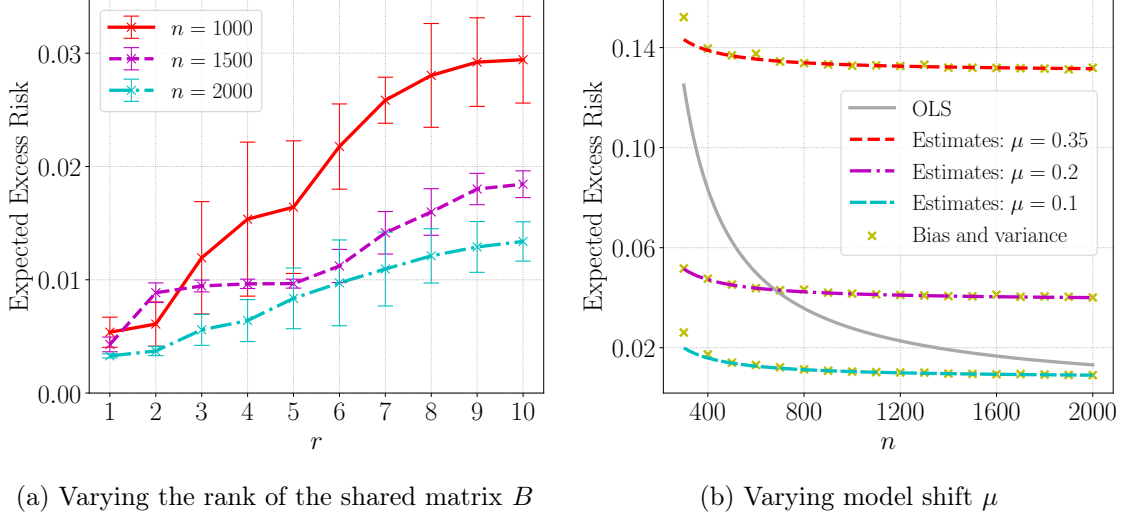


Figure 5: Illustration of transfer effects in the random-effects model with multiple source tasks. Figure 5a shows that for rank r from 1 to 10, the lowest excess risk of task t is achieved at $r = 1$ (averaged over three random seeds). Figure 5b fixes $r = 1$ and varies μ, n . Similar to Figure 1, the transfer effect of HPS can be either positive or negative depending on μ, n .

6. Comparison to Soft Transfer

In this section, we provide a case study of the statistical behavior of soft transfer estimation. Soft transfer works by penalizing the distance between the linear models for the target task and the source task. In general, soft transfer offers a more flexible regularization approach for adjusting the distance. Thus, one can expect the statistical behavior of soft transfer to be different from that of HPS. In this section, we provide a case study in several settings to illustrate this difference.

6.1 Bias-Variance under Deterministic Designs

We first derive the expectation of the excess risk $L(\hat{\beta}_2^{\text{SPS}}(\lambda))$ over the randomness of $\varepsilon^{(1)}$ and $\varepsilon^{(2)}$ to get the bias and variance of SPS. For simplicity, we denote $\hat{\Sigma}^{(1)} = X^{(1)\top} X^{(1)}$, $\hat{\Sigma}^{(2)} = X^{(2)\top} X^{(2)}$, and $\hat{\Sigma} = \hat{\Sigma}^{(1)} + \hat{\Sigma}^{(2)}$ as in (2.7). Let $\hat{\Sigma}_\lambda$ (which depends on the regularization parameter $\lambda > 0^4$) be

$$\hat{\Sigma}_\lambda := \lambda^{-1} \hat{\Sigma}^{(1)} \hat{\Sigma}^{(2)} + \hat{\Sigma}. \quad (6.1)$$

The bias and variance of the SPS estimator can be derived from straightforward algebraic calculations.

4. The case that $\lambda = 0$ is not interesting since there would be no information transfer in this case.

Lemma 15 *In expectation over the randomness of $\varepsilon^{(1)}$ and $\varepsilon^{(2)}$, we have the following equation for any $\lambda > 0$:*

$$\mathbb{E}_{\varepsilon^{(1)}, \varepsilon^{(2)}} \left[L(\hat{\beta}_2^{\text{SPS}}(\lambda)) \right] = L_{\text{bias}}(\lambda) + L_{\text{var}}(\lambda), \quad (6.2)$$

where the bias and variance formulas are defined as

$$L_{\text{bias}}(\lambda) = \left\| \Sigma^{(2)^{\frac{1}{2}}} (\hat{\Sigma}_\lambda)^{-1} \hat{\Sigma}^{(1)} (\beta^{(1)} - \beta^{(2)}) \right\|^2, \quad (6.3)$$

$$L_{\text{var}}(\lambda) = \sigma^2 \text{Tr} \left[\left((\hat{\Sigma}_\lambda)^{-1} \right)^\top \Sigma^{(2)} (\hat{\Sigma}_\lambda)^{-1} \left(\hat{\Sigma}^{(1)} + \left(\frac{\hat{\Sigma}^{(1)}}{\lambda} + \text{Id} \right) \hat{\Sigma}^{(2)} \left(\frac{\hat{\Sigma}^{(1)}}{\lambda} + \text{Id} \right) \right) \right]. \quad (6.4)$$

As a sanity check, one can see that when λ tends to infinity, the above reduces to the bias and variance of SPS, since $\hat{\Sigma}_\lambda$ tends to $\hat{\Sigma}$ and $\hat{\Sigma}^{(1)}/\lambda$ tends to zero.

6.2 An Illustrative Example

Above, we have considered $\hat{\Sigma}^{(1)}$ and $\hat{\Sigma}^{(2)}$ to be deterministic covariance matrices by conditioning on $X^{(1)}$ and $X^{(2)}$. Next, we are given a special example to compare the behavior of SPS with HPS, assuming that $(Z^{(2)})^\top Z^{(2)}/n_2 = \text{Id}_{p \times p}$ while task one still follows a random design. Given any $\lambda > 0$, denote

$$\tilde{\Sigma}(\lambda) = \frac{n_1}{\lambda} \Sigma^{(2)} + \frac{n_1}{n_2} \text{Id}_{p \times p}.$$

Let a_0 be the unique positive solution to the equation

$$\frac{1}{a_0} = 1 + \frac{1}{n_1} \text{Tr} \left[\frac{\tilde{\Sigma}(\lambda)}{\text{Id}_{p \times p} + a_0 \tilde{\Sigma}(\lambda)} \right].$$

Then, let

$$x_0 = \frac{1}{n_1} \text{Tr} \left[\left(\text{Id}_{p \times p} + a_0 \tilde{\Sigma}(\lambda) \right)^{-2} \right], \quad y_0 = a_0^2 x_0 \left(1 - \frac{p}{n_1} - 2a_0 + x_0 \right)^{-1}. \quad (6.5)$$

Lastly, define two matrices as follows:

$$M_1 = \frac{n_1 a_0}{n_2} \left(\text{Id}_{p \times p} + a_0 \tilde{\Sigma}(\lambda) \right)^{-1}, \quad M_2 = \left(\frac{n_1}{n_2} \right)^2 (a_0^2 - y_0) \left(\text{Id}_{p \times p} + a_0 \tilde{\Sigma}(\lambda) \right)^{-2}. \quad (6.6)$$

We state the limits of bias and variance of SPS as follows.

Proposition 16 *Suppose Assumption 1 holds. Additionally, assume that $\Sigma^{(1)} = \Sigma^{(2)}$ and $Z^{(2)\top} Z^{(2)}/n_2 = \text{Id}_{p \times p}$. Then, for any small constant $c > 0$, with high probability over the randomness of the training samples, the following estimates hold:*

$$L_{\text{bias}} = (\beta^{(1)} - \beta^{(2)})^\top \Sigma^{(2)^{\frac{1}{2}}} M_2 \Sigma^{(2)^{\frac{1}{2}}} (\beta^{(1)} - \beta^{(2)}) + O \left(n^{\frac{1}{\varphi} - \frac{1}{4} + c} \|\beta^{(1)} - \beta^{(2)}\|^2 \right), \quad (6.7)$$

$$L_{\text{var}} = \frac{\sigma^2 p}{n_2} - \frac{\sigma^2}{n_2} \text{Tr}[M_1] - \frac{\sigma^2}{\lambda} \text{Tr}[\Sigma^{(2)} M_2] + O \left(\left(\frac{p \sigma^2}{n_2} + \frac{p \sigma^2}{\lambda} \right) n^{\frac{1}{\varphi} - \frac{1}{4} + c} \right), \quad (6.8)$$

where we denote $n = n_1 + n_2$.




Figure 6: We find that the regularization effect of SPS can help mitigate negative transfer, which is more prominent when μ is higher (as indicated by the two lines where $\mu = 0.3$ and $\mu = 0.25$). We leave a thorough understanding of this phenomenon, including soft transfer, for future work.

Recalling that $\varphi > 4$, for a small enough constant c , equations (6.7) and (6.8) characterize the limits of L_{var} and L_{bias} with an error term that is smaller than the deterministic leading term by an asymptotically vanishing factor. The proof of Lemma 15 and Proposition 16 can be found in Appendix E.

It is an interesting question to derive the asymptotic bias and variance limits for the SPS estimator under random designs for both $\hat{\Sigma}^{(1)}$ and $\hat{\Sigma}^{(2)}$, with arbitrary covariate and model shifts. This requires developing new tools beyond the current random matrix theory literature. Similar to the issues in Remark 11, one needs to deal with an even more complicated asymmetric random matrix than equation (5.4), which is left for future work:

$$(\hat{\Sigma}^{(1)})^{-1} \hat{\Sigma}_{\lambda} = \lambda^{-1} \hat{\Sigma}^{(2)} + (\hat{\Sigma}^{(1)})^{-1} \hat{\Sigma}^{(2)} + \text{Id}_{p \times p}.$$

6.3 Numerical Comparisons

We conduct a numerical comparison of SPS, OLS, and HPS. In Figure 6, we use the same data-generating process as Figure 1, except that we now take $Z^{(2)\top} Z^{(2)} = n_2 \text{Id}_{p \times p}$. We find that the regularization effect of SPS can generally help reduce negative transfer, especially when the model shift is large. Additionally, our estimates in Proposition 16 match the bias and variance values in finite dimensions. Here, we set $p = 100$.

6.4 Discussion

Our analysis of SPS suggests that it behaves similarly to HPS when the model shift is low or modest, and SPS can generally help reduce negative transfer when the model shift is high. In this sense, the regularization of SPS behaves similarly to the data set shrinkage applied to HPS. More general forms of soft transfer can be found in transfer learning with foundation models, such as fine-tuning and instruction tuning, where overfitting often occurs when a large model is fine-tuned on a small target data set. For such scenarios, applying the regularization in soft transfer can effectively mitigate overfitting. It would be interesting to further study the regularization effect of soft transfer in large models, which would also

require developing tools for analyzing over-parameterized models (and perhaps a separate data model). These are left for future work.

7. Conclusion

This paper presents several high-dimensional asymptotic results for analyzing transfer learning with two linear regression tasks: one is the source task, and the other is the target task. The main technical ingredients involve estimating the high-dimensional asymptotic limits of several functions involving two independent sample covariance matrices with different population covariance matrices. In particular, our main result involves:

- A generalization of a well-known result from the random matrix theory literature on the trace of the inverse of one sample covariance matrix to the sum of two sample covariance matrices under covariate shifts (Theorem 3), along with an almost sharp convergence rate to this limit.
- A phase transition of transfer learning under model shifts between two high-dimensional linear regression tasks, depending on the level of model shifts and the sample sizes of source and target tasks (Theorem 6). With a simple adjustment, HPS is nearly optimal, matching the minimax lower bound up to constant factors.
- An extension of the findings to a setting involving both covariate shift and model shift, and a multi-task learning setting where the feature covariates are shared among all tasks.

Numerical simulations are provided to show the accuracy of our estimates in finite dimensions.

Our work establishes a rigorous connection between transfer learning and random matrix theory, highlighting several technically challenging questions that are worth further study. The theoretical modeling of transfer learning is still a relatively new research area. We hope our work can inspire future studies that apply random matrix theory to transfer learning.

Acknowledgments

The first two authors, Fan Yang and Hongyang Zhang, contributed equally and are listed in alphabetical order.

Thanks to Tony Cai, Edgar Dobriban, Kaizheng Wang, Yaqi Duan, Wei Hu, and Hongji Wei for helpful discussions at various stages of this project. We are grateful to the action editor and the anonymous referees for their helpful comments, which helped improve this paper.

Fan Yang is supported in part by the National Key R&D Program of China (No. 2023YFA1010400), and is also affiliated with Beijing Institute of Mathematical Sciences and Applications. Hongyang Zhang is supported in part by NSF award IIS-2412008. Fan Yang and Hongyang Zhang were partly supported by the Wharton Dean's Fund for Postdoctoral Research during the early stage of this project at UPenn. Weijie Su is supported in part by NSF through DMS-1847415 and the Wharton Dean's Fund for Postdoctoral Research.

Appendix A. Basic Tools

Here is a roadmap of the appendix materials. In Appendix A.1 and A.2, we describe some basic tools used to establish the proofs. Then, in Appendix A.3, we prove the bias-variance decomposition of the HPS estimation from Section 2. In Appendix B, we prove the results stated in Section 3. In Appendix C, we give the complete proofs for the estimates under model shifts from Section 4 and the extension in Section 5.1. We state the proofs for the multi-task case in Appendix D. Finally, Appendix E provides proofs for the SPS estimation.

A.1 Notations

In this supplement, we will use the following notations. The fundamental large parameter in our proof is p . All quantities that are not explicitly constant may depend on p , and we usually omit p from our notations. Given any matrix X , let $\|X\|$ or $\|X\|_2$ denote its operator norm (or equivalently, the largest singular value); let $\|X\|_F$ denote its Frobenius norm; let $\lambda_{\max}(X)$ and $\lambda_{\min}(X)$ denote its largest and smallest singular values, respectively; let $\lambda_1(X) \geq \lambda_2(X) \geq \dots$ denote the singular values of X in descending order; let X^+ denote the Moore-Penrose pseudoinverse of X . As a special case, the operator or Frobenius norm of a vector v is also its Euclidean norm, i.e., $\|v\| = \|v\|_2 = \|v\|_F$. We will often write an identity matrix as 1 if this does not cause any confusion in the specific context.

We say an event Ξ holds with high probability if $\mathbb{P}(\Xi) \rightarrow 1$ as $p \rightarrow \infty$.

We will use the big-O notation $g(p) = O(f(p))$ if there exists a constant $C > 0$ such that $|g(p)| \leq C f(p)$ for large enough p . Moreover, for any $g(p) \geq 0$ and $f(p) \geq 0$, we will use the notations $g(p) \lesssim f(p)$ if $g(p) = O(f(p))$, and $g(p) \sim f(p)$ if $g(p) \lesssim f(p)$ and $f(p) \lesssim g(p)$.

A.2 Concentration Estimates

In this section, we collect some useful tools that will be used in the proof. First, it is convenient to introduce the following notation.

Definition 17 (Overwhelming probability) *We say an event Ξ holds with overwhelming probability (w.o.p.) if for any constant $D > 0$, $\mathbb{P}(\Xi) \geq 1 - p^{-D}$ for large enough p . Moreover, we say Ξ holds with overwhelming probability in an event Ω if for any constant $D > 0$, $\mathbb{P}(\Omega \setminus \Xi) \leq p^{-D}$ for large enough p .*

The following notion of stochastic domination is commonly used to study random matrices.

Definition 18 (Stochastic domination) *Let $\xi \equiv \xi^{(p)}$ and $\zeta \equiv \zeta^{(p)}$ be two p -dependent random variables. We say that ξ is stochastically dominated by ζ , denoted by $\xi \prec \zeta$ or $\xi = O_{\prec}(\zeta)$, if for any small constant $c > 0$ and large constant $D > 0$, there exists a $p_0(c, D) \in \mathbb{N}$ such that for all $p > p_0(c, D)$,*

$$\mathbb{P}(|\xi| > p^c |\zeta|) \leq p^{-D}.$$

In other words, $\xi \prec \zeta$ if $|\xi| \leq p^c |\zeta|$ with overwhelming probability for any small constant $c > 0$. If $\xi(u)$ and $\zeta(u)$ are functions of u supported in a set \mathcal{U} , then we say $\xi(u)$ is stochastically dominated by $\zeta(u)$ uniformly in \mathcal{U} if

$$\mathbb{P}(\cup_{u \in \mathcal{U}} \{|\xi(u)| > p^c |\zeta(u)|\}) \leq p^{-D}$$

for large enough p . Given any event Ω , we say $\xi \prec \zeta$ on Ω if $\mathbf{1}_\Omega \xi \prec \zeta$.

Remark 19 We make two simple remarks. First, since we allow for an p^c factor in stochastic domination, we can ignore log factors without loss of generality since $(\log p)^C \prec 1$ for any constant $C > 0$. Second, given a random variable ξ with unit variance and finite moments up to any order as in (2.4), we have that $|\xi| \prec 1$. This is because, by Markov's inequality, there is

$$\mathbb{P}(|\xi| \geq p^c) \leq p^{-kc} \mathbb{E}|\xi|^k \leq p^{-D},$$

as long as k is taken to be larger than D/c .

The following lemma collects several basic properties of stochastic domination that will be used tacitly in the proof. Roughly speaking, it says that the stochastic domination " \prec " can be treated as the conventional less-than sign " $<$ " in some sense.

Lemma 20 (Lemma 3.2 in Bloemendal et al. (2014)) *Let ξ and ζ be two families of nonnegative random variables depending on some parameters $u \in \mathcal{U}$ and $v \in \mathcal{V}$. Let $C > 0$ be an arbitrary constant.*

- (i) **Sum.** *Suppose that $\xi(u, v) \prec \zeta(u, v)$ uniformly in $u \in \mathcal{U}$ and $v \in \mathcal{V}$. If $|\mathcal{V}| \leq p^C$, then $\sum_{v \in \mathcal{V}} \xi(u, v) \prec \sum_{v \in \mathcal{V}} \zeta(u, v)$ uniformly in u .*
- (ii) **Product.** *If $\xi_1(u) \prec \zeta_1(u)$ and $\xi_2(u) \prec \zeta_2(u)$ uniformly in $u \in \mathcal{U}$, then $\xi_1(u)\xi_2(u) \prec \zeta_1(u)\zeta_2(u)$ uniformly in $u \in \mathcal{U}$.*
- (iii) **Expectation.** *Suppose that $\Psi(u) \geq p^{-C}$ is a family of deterministic parameters, and $\xi(u)$ satisfies $\mathbb{E}\xi(u)^2 \leq p^C$. If $\xi(u) \prec \Psi(u)$ uniformly in u , then we also have $\mathbb{E}\xi(u) \prec \Psi(u)$ uniformly in u .*

Let $Q > 0$ be a (p -dependent) deterministic parameter. We say a random matrix $Z \in \mathbb{R}^{n \times p}$ satisfies the *bounded support condition* with Q (or Z has *bounded support* Q) if

$$\max_{1 \leq i \leq n, 1 \leq j \leq p} |Z_{ij}| \prec Q. \quad (\text{A.1})$$

As shown in Remark 19, if the entries of Z have finite moments up to any order, then Z has bounded support $Q = 1$. More generally, if every entry of Z has a finite φ -th moment as in (2.2) and $n \geq p$, then using Markov's inequality and a simple union bound we get that

$$\begin{aligned} \mathbb{P}\left(\max_{1 \leq i \leq n, 1 \leq j \leq p} |Z_{ij}| \geq (\log n)n^{\frac{2}{\varphi}}\right) &\leq \sum_{i=1}^n \sum_{j=1}^p \mathbb{P}\left(|Z_{ij}| \geq (\log n)n^{\frac{2}{\varphi}}\right) \\ &\lesssim \sum_{i=1}^n \sum_{j=1}^p \left[(\log n)n^{\frac{2}{\varphi}}\right]^{-\varphi} \leq (\log n)^{-\varphi}. \end{aligned} \quad (\text{A.2})$$

In other words, Z has bounded support $Q = n^{2/\varphi}$ with high probability.

The following lemma, which can also be found from Lemma 3.8 in Erdős et al. (2013), gives sharp concentration bounds for linear and quadratic forms of random variables with bounded support.

Lemma 21 *Let the sequence $(x_i), (y_j)$ be families of centered and independent random variables, and $(A_i), (B_{ij})$ be families of deterministic complex numbers. Suppose the entries x_i and y_j have variance at most 1, and satisfy the bounded support condition (A.1) for a deterministic parameter $Q \geq 1$. Then, we have the following estimates:*

$$\left| \sum_{i=1}^n A_i x_i \right| \prec Q \max_{i=1}^n |A_i| + \left(\sum_{i=1}^n |A_i|^2 \right)^{1/2}, \quad (\text{A.3})$$

$$\left| \sum_{i,j=1}^n x_i B_{ij} y_j \right| \prec Q^2 B_d + Q n^{1/2} B_o + \left(\sum_{1 \leq i \neq j \leq n} |B_{ij}|^2 \right)^{1/2}, \quad (\text{A.4})$$

$$\left| \sum_{i=1}^n (|x_i|^2 - \mathbb{E}|x_i|^2) B_{ii} \right| \prec Q n^{1/2} B_d, \quad (\text{A.5})$$

$$\left| \sum_{1 \leq i \neq j \leq n} x_i B_{ij} x_j \right| \prec Q n^{1/2} B_o + \left(\sum_{1 \leq i \neq j \leq n} |B_{ij}|^2 \right)^{1/2}, \quad (\text{A.6})$$

where we denote $B_d := \max_i |B_{ii}|$ and $B_o := \max_{i \neq j} |B_{ij}|$. Moreover, if x_i and y_j have finite moments up to any order, then we have the following stronger estimates:

$$\left| \sum_{i=1}^n A_i x_i \right| \prec \left(\sum_{i=1}^n |A_i|^2 \right)^{1/2}, \quad (\text{A.7})$$

$$\left| \sum_{i,j=1}^n x_i B_{ij} y_j \right| \prec \left(\sum_{i,j=1}^n |B_{ij}|^2 \right)^{1/2}, \quad (\text{A.8})$$

$$\left| \sum_{i=1}^n (|x_i|^2 - \mathbb{E}|x_i|^2) B_{ii} \right| \prec \left(\sum_{i=1}^n |B_{ii}|^2 \right)^{1/2}, \quad (\text{A.9})$$

$$\left| \sum_{1 \leq i \neq j \leq n} x_i B_{ij} x_j \right| \prec \left(\sum_{1 \leq i \neq j \leq n} |B_{ij}|^2 \right)^{1/2}. \quad (\text{A.10})$$

It is well-known that the empirical spectral distributions of $Z^{(1)\top} Z^{(1)}$ and $Z^{(2)\top} Z^{(2)}$ satisfy the Marchenko-Pastur (MP) law (Marčenko and Pastur, 1967). Moreover, their eigenvalues are all inside the support of the MP law with high probability. In the proof, we will need a slightly stronger result that holds with overwhelming probability, as given by the following lemma.

Lemma 22 *Suppose $Z \in \mathbb{R}^{n \times p}$ is an $n \times p$ random matrix satisfying the same assumptions as $Z^{(2)}$ in Assumption 1. Suppose $1 + \tau \leq n/p \leq p^{\tau-1}$, and Z satisfies the bounded support condition (A.1) for a deterministic parameter Q such that $1 \leq Q \leq n^{1/2-c_Q}$ for a constant $c_Q > 0$. Then, we have that*

$$(\sqrt{n} - \sqrt{p})^2 - O_{\prec}(\sqrt{n} \cdot Q) \leq \lambda_p(Z^\top Z) \leq \lambda_1(Z^\top Z) \leq (\sqrt{n} - \sqrt{p})^2 + O_{\prec}(\sqrt{n} \cdot Q). \quad (\text{A.11})$$

Proof When Q is of order 1, this lemma follows from Theorem 2.10 of Bloemendal et al. (2014). The result for the general case with $1 \leq Q \leq n^{1/2-c_Q}$ follows from Lemma 3.11 of Ding and Yang (2018). \blacksquare

Using a standard cut-off argument, we can extend Lemma 21 and Lemma 22 to random matrices whose entries satisfy only certain moment assumptions but not necessarily the bounded support condition.

Corollary 23 *Suppose $Z \in \mathbb{R}^{n \times p}$ is an $n \times p$ random matrix satisfying the same assumptions as $Z^{(2)}$ in Assumption 1. Suppose $1 + \tau \leq n/p \leq p^{\tau-1}$. Then, equation (A.11) holds on a high probability event with $Q = n^{2/\varphi}$, where φ refers to the constant in equation (2.2).*

Proof We introduce a truncated matrix \tilde{Z} with entries

$$\tilde{Z}_{ij} := \mathbf{1}(|Z_{ij}| \leq Q \log n) \cdot Z_{ij}. \quad (\text{A.12})$$

From equation (A.2), we get that

$$\mathbb{P}(\tilde{Z} = Z) = 1 - \mathbb{P}\left(\max_{i,j} |Z_{ij}| > Q \log n\right) = 1 - O((\log n)^{-\varphi}). \quad (\text{A.13})$$

By definition, we have

$$\begin{aligned} \mathbb{E} \tilde{Z}_{ij} &= -\mathbb{E}[\mathbf{1}(|Z_{ij}| > Q \log n) Z_{ij}], \\ \mathbb{E} |\tilde{Z}_{ij}|^2 &= 1 - \mathbb{E}[\mathbf{1}(|Z_{ij}| > Q \log n) |Z_{ij}|^2]. \end{aligned} \quad (\text{A.14})$$

Using the tail probability expectation formula, we can check that

$$\begin{aligned} \mathbb{E}[\mathbf{1}(|Z_{ij}| > Q \log n) Z_{ij}] &= \int_0^\infty \mathbb{P}(\mathbf{1}(|Z_{ij}| > Q \log n) Z_{ij} > s) ds \\ &= \int_0^{Q \log n} \mathbb{P}(|Z_{ij}| > Q \log n) ds + \int_{Q \log n}^\infty \mathbb{P}(|Z_{ij}| > s) ds \\ &\lesssim \int_0^{Q \log n} (Q \log n)^{-\varphi} ds + \int_{Q \log n}^\infty s^{-\varphi} ds \leq n^{-2(\varphi-1)/\varphi}, \end{aligned}$$

where in the third step we use the finite φ -th moment condition (2.2) for Z_{ij} and Markov's inequality. Similarly, we can obtain that

$$\begin{aligned} \mathbb{E}[\mathbf{1}(|Z_{ij}| > Q \log n) |Z_{ij}|^2] &= 2 \int_0^\infty s \mathbb{P}(\mathbf{1}(|Z_{ij}| > Q \log n) |Z_{ij}| > s) ds \\ &= 2 \int_0^{Q \log n} s \mathbb{P}(|Z_{ij}| > Q \log n) ds + 2 \int_{Q \log n}^\infty s \mathbb{P}(|Z_{ij}| > s) ds \\ &\lesssim \int_0^{Q \log n} s (Q \log n)^{-\varphi} ds + \int_{Q \log n}^\infty s^{-\varphi+1} ds \leq n^{-2(\varphi-2)/\varphi}. \end{aligned}$$

Plugging the above two estimates into (A.14) and using $\varphi > 4$, we get that

$$|\mathbb{E} \tilde{Z}_{ij}| = O(n^{-3/2}), \quad \mathbb{E} |\tilde{Z}_{ij}|^2 = 1 + O(n^{-1}). \quad (\text{A.15})$$

From the first estimate in equation (A.15), we also get a bound on the operator norm:

$$\|\mathbb{E}[\tilde{Z}]\| = O(n^{-1/2}). \quad (\text{A.16})$$

Then, we centralize and rescale \tilde{Z} as

$$\hat{Z} := (\tilde{Z} - \mathbb{E}\tilde{Z})/\sqrt{\mathbb{E}|\tilde{Z}_{11}|^2}. \quad (\text{A.17})$$

Now, \hat{Z} is a matrix satisfying the assumptions of Lemma 22 with bounded support Q . Thus, we get that

$$(\sqrt{n} - \sqrt{p})^2 - O_{\prec}(\sqrt{n} \cdot Q) \leq \lambda_p(\hat{Z}^\top \hat{Z}) \leq \lambda_1(\hat{Z}^\top \hat{Z}) \leq (\sqrt{n} - \sqrt{p})^2 + O_{\prec}(\sqrt{n} \cdot Q).$$

Combining this estimate with lines (A.15) and (A.16), we can readily get that equation (A.11) holds for the eigenvalues of $\tilde{Z}^\top \tilde{Z}$, which concludes the proof by equation (A.13). \blacksquare

Corollary 24 *Suppose $Z \in \mathbb{R}^{n \times p}$ is an $n \times p$ random matrix satisfying the same assumptions as $Z^{(2)}$ in Assumption 1. Then, there exists a high probability event on which the following estimate holds for any deterministic vector $v \in \mathbb{R}^p$:*

$$\left| \|Zv\|^2 - n\|v\|^2 \right| \prec n^{1/2}Q\|v\|^2, \quad \text{for } Q = n^{2/\varphi}. \quad (\text{A.18})$$

Proof Similar to the proof of Corollary 23, we truncate Z as in (A.12) and define \hat{Z} by (A.17). By (A.15) and (A.16), we see that to conclude (A.18), it suffices to show that

$$\left| \|\hat{Z}v\|^2 - n\|v\|^2 \right| \prec n^{1/2}Q\|v\|^2. \quad (\text{A.19})$$

To prove equation (A.19), we first notice that $\hat{Z}v \in \mathbb{R}^n$ is a random vector with i.i.d. entries of mean zero and variance $\|v\|^2$. Furthermore, with equation (A.3) we get that

$$|(\hat{Z}v)_i| \prec Q \max_{j=1}^p |v_j| + \|v\| \leq 2Q\|v\|, \quad i = 1, 2, \dots, n.$$

Hence, $(\hat{Z}v)/\|v\|$ consists of i.i.d. random entries with zero mean, unit variance, and bounded support Q . Then, applying equation (A.5), we get that

$$\left| \|\hat{Z}v\|^2 - n\|v\|^2 \right| = \|v\|^2 \left| \sum_{i=1}^n \left(\frac{|(\hat{Z}v)_i|^2}{\|v\|^2} - \mathbb{E} \frac{|(\hat{Z}v)_i|^2}{\|v\|^2} \right) \right| \prec Qn^{1/2}\|v\|^2.$$

This concludes the proof. \blacksquare

From Lemma 21, we can obtain the following concentration estimates for zero-mean noise vectors (which may also be found in standard texts such as Wainwright (2019)).

Corollary 25 *Suppose $\varepsilon^{(1)}, \dots, \varepsilon^{(t)} \in \mathbb{R}^n$ are independent random vectors satisfying Assumption 1 (ii). Then, we have that for any deterministic vector $v \in \mathbb{R}^n$,*

$$|v^\top \varepsilon^{(i)}| \prec \sigma\|v\|, \quad i = 1, \dots, t, \quad (\text{A.20})$$

and for any deterministic matrix $B \in \mathbb{R}^{n \times n}$,

$$\left| (\varepsilon^{(i)})^\top B \varepsilon^{(j)} - \delta_{i,j} \cdot \sigma^2 \text{Tr}[B] \right| \prec \sigma^2 \|B\|_F, \quad \text{for any } i, j = 1, \dots, t, \quad (\text{A.21})$$

where $\delta_{i,j} = 1$ if $i = j$, else $\delta_{i,j} = 0$ if $i \neq j$.

Proof Note that $\varepsilon^{(i)}/\sigma$ is a random vector with independent entries of zero mean, unit variance, and bounded moments up to any order. Then, the result in equation (A.20) implies that equation (A.7) holds. Using equation (A.8), we obtain that for $i \neq j$,

$$\left| (\varepsilon^{(i)})^\top B \varepsilon^{(j)} \right| = \left| \sum_{k,l=1}^n \varepsilon_k^{(i)} \varepsilon_l^{(j)} B_{kl} \right| \prec \sigma^2 \left(\sum_{k,l=1}^n |B_{kl}|^2 \right)^{1/2} = \sigma^2 \|B\|_F.$$

Using the two estimates (A.9) and (A.10), we obtain that

$$\begin{aligned} \left| (\varepsilon^{(i)})^\top B \varepsilon^{(i)} - \sigma^2 \text{Tr}(B) \right| &\leq \left| \sum_k \left(|\varepsilon_k^{(i)}|^2 - \mathbb{E}|\varepsilon_k^{(i)}|^2 \right) B_{kk} \right| + \left| \sum_{k \neq l} \varepsilon_k^{(i)} \varepsilon_l^{(i)} B_{kl} \right| \\ &\prec \sigma^2 \left(\sum_{k=1}^p |B_{kk}|^2 \right)^{\frac{1}{2}} + \sigma^2 \left(\sum_{k \neq l} |B_{kl}|^2 \right)^{\frac{1}{2}} \leq \sqrt{2} \sigma^2 \|B\|_F. \end{aligned}$$

Hence, we have proved that the estimate (A.21) hold. \blacksquare

A.3 Proof of Lemma 2

Proof For the optimization objective $\ell(B)$ in equation (1.2), using the local optimality condition $\nabla \ell(B) = 0$, we can solve that

$$\hat{B} = \hat{\Sigma}^{-1} \left(X^{(1)\top} Y^{(1)} + X^{(2)\top} Y^{(2)} \right),$$

where $\hat{\Sigma}$ is the sum of the sample covariance matrices (see definition in equation (2.7)). Then, the HPS estimator for task two is equal to $\hat{\beta}_2^{\text{HPS}} = \hat{B}$. Plugging it into the excess risk $L(\hat{\beta}_2^{\text{HPS}})$, we find that it is equal to

$$\left\| \Sigma^{(1)/2} \hat{\Sigma}^{-1} X^{(1)\top} X^{(1)} (\beta^{(1)} - \beta^{(2)}) + \Sigma^{(2)/2} \hat{\Sigma}^{-1} \left(X^{(1)\top} \varepsilon^{(1)} + X^{(2)\top} \varepsilon^{(2)} \right) \right\|^2,$$

which can be expanded as $L_{\text{bias}} + 2h_1 + 2h_2 + h_3 + h_4 + 2h_5$, where

$$\begin{aligned} h_1 &:= \varepsilon^{(1)\top} X^{(1)} \hat{\Sigma}^{-1} \Sigma^{(2)} \hat{\Sigma}^{-1} X^{(1)\top} X^{(1)} (\beta^{(1)} - \beta^{(2)}), \\ h_2 &:= \varepsilon^{(2)\top} X^{(2)} \hat{\Sigma}^{-1} \Sigma^{(2)} \hat{\Sigma}^{-1} X^{(1)\top} X^{(1)} (\beta^{(1)} - \beta^{(2)}), \\ h_3 &:= \varepsilon^{(1)\top} X^{(1)} \hat{\Sigma}^{-1} \Sigma^{(2)} \hat{\Sigma}^{-1} X^{(1)\top} \varepsilon^{(1)}, \\ h_4 &:= \varepsilon^{(2)\top} X^{(2)} \hat{\Sigma}^{-1} \Sigma^{(2)} \hat{\Sigma}^{-1} X^{(2)\top} \varepsilon^{(2)}, \\ h_5 &:= \varepsilon^{(1)\top} X^{(1)} \hat{\Sigma}^{-1} \Sigma^{(2)} \hat{\Sigma}^{-1} X^{(2)\top} \varepsilon^{(2)}. \end{aligned}$$

We estimate these terms one by one using Lemma 22 and Corollary 25.

First, under assumption (2.3), we can obtain the following estimates using Corollary 23. There exists a high probability event on which all of the following estimates hold

$$\left\| X^{(1)} \right\|_2 \lesssim \max(\sqrt{n_1}, \sqrt{p}), \quad \left\| X^{(2)} \right\|_2 \lesssim \sqrt{n_2}, \quad \left\| \hat{\Sigma}^{-1} \right\|_2 \lesssim (n_1 + n_2)^{-1}. \quad (\text{A.22})$$

As for the last bound, we can derive it in more detail as follows (recall that λ_1, λ_p refer to the largest and smallest singular values, respectively): using estimate (A.11) and the assumptions $\rho_2 \geq 1 + \tau$ and $\rho_2 \geq \tau \rho_1$ in equation (2.5), we obtain that with high probability,

$$\lambda_1(\hat{\Sigma}) \leq \lambda_1(X^{(1)\top} X^{(1)}) + \lambda_1(X^{(2)\top} X^{(2)}) \lesssim n_1 + n_2, \quad (\text{A.23})$$

$$\lambda_p(\hat{\Sigma}) \geq \lambda_p(X^{(2)\top} X^{(2)}) \gtrsim n_2, \quad (\text{A.24})$$

which is again at the order of $n_1 + n_2$ since they are within a constant factor of each other. With equation (A.24), we conclude the last estimate in equation (A.22):

$$\|\hat{\Sigma}^{-1}\|_2 = \frac{1}{\lambda_p(\hat{\Sigma})} \lesssim \frac{1}{n_1 + n_2}. \quad (\text{A.25})$$

In addition, based on the estimates (A.23)–(A.25), we can obtain that L_{bias} and L_{var} (cf. equations (2.9) and (2.10)) are of the following orders (up to constants that do not grow with p from both above and below):

$$L_{\text{bias}} \sim \frac{1}{(n_1 + n_2)^2} \left\| X^{(1)\top} X^{(1)} (\beta^{(1)} - \beta^{(2)}) \right\|^2, \quad L_{\text{var}} \sim \frac{p\sigma^2}{n_1 + n_2}. \quad (\text{A.26})$$

We first look at L_{var} . Notice that

$$\begin{aligned} L_{\text{var}} &= \sigma^2 \text{Tr}[\Sigma^{(2)} \hat{\Sigma}^{-1}] \leq \sigma^2 p \left\| \Sigma^{(2)} \hat{\Sigma}^{-1} \right\|_2 \\ &\leq \sigma^2 p \left\| \Sigma^{(2)} \right\|_2 \cdot \left\| \hat{\Sigma}^{-1} \right\|_2 \\ &\leq \sigma^2 p \tau^{-1} \cdot \left\| \hat{\Sigma}^{-1} \right\|_2 \quad (\text{by Assumption (2.3) on } \Sigma^{(2)}) \\ &\lesssim \frac{\sigma^2 p}{n_1 + n_2}. \quad (\text{by estimate (A.25)}) \end{aligned}$$

Likewise, using the assumption that $\lambda_p(\Sigma^{(2)}) \geq \tau$, one can show that

$$L_{\text{var}} \geq \sigma^2 p \tau \cdot \lambda_p(\hat{\Sigma}^{-1}) \gtrsim \frac{\sigma^2 p}{n_1 + n_2},$$

where we have used that $\lambda_p(\hat{\Sigma}^{-1}) = \lambda_1(\hat{\Sigma})^{-1} \gtrsim (n_1 + n_2)^{-1}$ by estimate (A.23). Next, we look at L_{bias} . We first examine its upper estimate:

$$\begin{aligned} L_{\text{bias}} &= \left\| \Sigma^{(2)\frac{1}{2}} \hat{\Sigma}^{-1} (X^{(1)\top} X^{(1)}) (\beta^{(1)} - \beta^{(2)}) \right\|^2 \\ &\leq \left\| \Sigma^{(2)} \right\|_2 \left\| \hat{\Sigma}^{-1} \right\|_2^2 \left\| X^{(1)\top} X^{(1)} (\beta^{(1)} - \beta^{(2)}) \right\|^2 \\ &\lesssim \frac{1}{(n_1 + n_2)^2} \left\| X^{(1)\top} X^{(1)} (\beta^{(1)} - \beta^{(2)}) \right\|^2. \quad (\text{by Eq. (2.3) and estimate (A.25)}) \end{aligned}$$

Then, we can derive the lower estimate with similar arguments as follows:

$$\begin{aligned} L_{\text{bias}} &\geq \left(\lambda_p(\Sigma^{(2)\frac{1}{2}} \hat{\Sigma}^{-1}) \right)^2 \cdot \left\| X^{(1)\top} X^{(1)} (\beta^{(1)} - \beta^{(2)}) \right\|^2 \\ &\gtrsim \frac{1}{(n_1 + n_2)^2} \left\| X^{(1)\top} X^{(1)} (\beta^{(1)} - \beta^{(2)}) \right\|^2. \quad (\text{by Eq. (2.3) and estimate (A.23)}) \end{aligned}$$

In the first step, we use the fact that $\|Xv\| \geq (\lambda_{\min}(X))^2 \|v\|^2$ for any matrix X and vector v . In the second step, we use that $u^\top \hat{\Sigma}^{-1} \Sigma^{(2)} \hat{\Sigma}^{-1} u \geq \lambda_{\min}(\Sigma^{(2)}) \cdot u^\top \hat{\Sigma}^{-2} u$ for any vector u , so $\lambda_p(\Sigma^{(2)\frac{1}{2}} \hat{\Sigma}^{-1}) \gtrsim (n_1 + n_2)^{-1}$. Taken together, we have proved the estimates in (A.26).

Note that if we further assume that $\rho_1 \geq 1 + \tau$ as in Theorem 9, then by Assumption (2.3) on $\Sigma^{(1)}$ and the estimate (A.11), we have that with high probability, $n_1 \lesssim \lambda_p(X^{(1)\top} X^{(1)}) \leq \lambda_1(X^{(1)\top} X^{(1)}) \lesssim n_1$. Consequently, the estimation of L_{bias} in equation (A.26) simplifies to

$$L_{\text{bias}} \sim \frac{n_1^2}{(n_1 + n_2)^2} \left\| \beta^{(1)} - \beta^{(2)} \right\|^2. \quad (\text{A.27})$$

This equation has been used to establish the estimate (4.7).

The rest of the proof is always restricted to the high probability event Ξ , on which the estimates (A.22)–(A.26) all hold true. Using equation (A.20) in Corollary (25), we can bound h_1 on Ξ as

$$\begin{aligned} |h_1| &< \sigma \left\| X^{(1)} \hat{\Sigma}^{-1} \Sigma^{(2)} \hat{\Sigma}^{-1} X^{(1)\top} X^{(1)} (\beta^{(1)} - \beta^{(2)}) \right\| \\ &\leq \sigma \left\| X^{(1)\top} X^{(1)} (\beta^{(1)} - \beta^{(2)}) \right\| \cdot \left\| X^{(1)} \hat{\Sigma}^{-1} \Sigma^{(2)} \hat{\Sigma}^{-1} \right\|_2 \\ &\lesssim \sigma \left\| X^{(1)\top} X^{(1)} (\beta^{(1)} - \beta^{(2)}) \right\| \cdot \frac{\max(\sqrt{n_1}, \sqrt{p})}{(n_1 + n_2)^2} \quad (\text{by estimate (A.22)}) \\ &\leq \frac{p^{-\frac{1}{4}} \left\| X^{(1)\top} X^{(1)} (\beta^{(1)} - \beta^{(2)}) \right\|}{n_1 + n_2} \cdot \frac{p^{\frac{1}{4}} \sigma}{(n_1 + n_2)^{\frac{1}{2}}} \quad (\text{since } \max(\sqrt{n_1}, \sqrt{p}) \leq \sqrt{n_1 + n_2}) \\ &\leq \frac{p^{-\frac{1}{2}} \left\| X^{(1)\top} X^{(1)} (\beta^{(1)} - \beta^{(2)}) \right\|^2}{(n_1 + n_2)^2} + \frac{\sigma^2 \sqrt{p}}{n_1 + n_2} \quad (\text{by the AM-GM inequality}) \\ &\lesssim \frac{L_{\text{bias}} + L_{\text{var}}}{\sqrt{p}}. \quad (\text{by estimate (A.26)}) \end{aligned}$$

With a similar argument, we can show that $|h_2| < \frac{L_{\text{bias}} + L_{\text{var}}}{\sqrt{p}}$.

Next, using equation (A.21) in Corollary 25, we can estimate h_3 as follows

$$\begin{aligned} \left| h_3 - \sigma^2 \text{Tr} \left[X^{(1)} \hat{\Sigma}^{-1} \Sigma^{(2)} \hat{\Sigma}^{-1} X^{(1)\top} \right] \right| &< \sigma^2 \left\| X^{(1)} \hat{\Sigma}^{-1} \Sigma^{(2)} \hat{\Sigma}^{-1} X^{(1)\top} \right\|_F \\ &\lesssim \frac{\sigma^2}{(n_1 + n_2)^2} \left\| X^{(1)} X^{(1)\top} \right\|_F \quad (\text{by equations (2.3) and (A.22)}) \\ &\lesssim \frac{\sigma^2 p^{\frac{1}{2}} \max(n_1, p)}{(n_1 + n_2)^2} \quad (\text{by estimate (A.22)}) \\ &\lesssim \frac{L_{\text{var}}}{\sqrt{p}}, \quad (\text{by estimate (A.26)}) \end{aligned}$$

where we use the following estimate in the third line

$$\left\| X^{(1)} X^{(1)\top} \right\|_F \leq \sqrt{p} \left\| X^{(1)} X^{(1)\top} \right\|_2 \lesssim \sqrt{p} \max(n_1, p).$$

With a similar argument, we can show that

$$\left| h_4 - \sigma^2 \text{Tr} \left[X^{(2)} \hat{\Sigma}^{-1} \Sigma^{(2)} \hat{\Sigma}^{-1} X^{(2)\top} \right] \right| \prec \frac{L_{\text{var}}}{\sqrt{p}},$$

and $|h_5| \prec L_{\text{var}}/\sqrt{p}$.

Finally, combining the above estimates on h_1, h_2, h_3, h_4, h_5 , and using that

$$\sigma^2 \text{Tr} \left[X^{(1)} \hat{\Sigma}^{-1} \Sigma^{(2)} \hat{\Sigma}^{-1} X^{(1)\top} \right] + \sigma^2 \text{Tr} \left[X^{(2)} \hat{\Sigma}^{-1} \Sigma^{(2)} \hat{\Sigma}^{-1} X^{(2)\top} \right] = \sigma^2 \text{Tr} \left[\Sigma^{(2)} \hat{\Sigma}^{-1} \right] = L_{\text{var}},$$

we conclude that conditioned on event Ξ , the following estimate holds true:

$$\left| L(\hat{\beta}_2^{\text{HPS}}) - L_{\text{bias}} - L_{\text{var}} \right| \prec \frac{L_{\text{bias}} + L_{\text{var}}}{\sqrt{p}}. \quad (\text{A.28})$$

This concludes the proof of equation (2.8). \blacksquare

Appendix B. Proofs for the Covariate Shift Setting

Before presenting the proof of Theorem 3, we first use its conclusion to complete the proofs of Propositions 4 and 5.

B.1 Proof of Proposition 4

Proof By definition, we can expand $g(M)$ as

$$g(M) = \frac{\sigma^2}{n_1 + n_2} \sum_{i=1}^{\frac{p}{2}} \left(\frac{1}{\lambda_i^2 \alpha_1 + \alpha_2} + \frac{1}{\lambda_i^{-2} \alpha_1 + \alpha_2} \right).$$

When $M = \text{Id}_{p \times p}$, we have $g(\text{Id}_{p \times p}) = \frac{\sigma^2 p}{n_1 + n_2 - p}$, which can also be written as $\frac{\sigma^2}{n_1 + n_2} \sum_{i=1}^{\frac{p}{2}} \frac{2}{\alpha_1 + \alpha_2}$, since $\alpha_1 + \alpha_2 = 1 - \frac{p}{n_1 + n_2}$ by equation (3.2). Then, we subtract $g(\text{Id}_{p \times p})$ from $g(M)$ to get

$$g(M) - g(\text{Id}_{p \times p}) = \frac{\sigma^2}{n_1 + n_2 - p} \sum_{i=1}^{p/2} \frac{(\lambda_i^2 - 1)^2 \alpha_1 (\alpha_1 - \alpha_2)}{(\alpha_1 + \lambda_i^2 \alpha_2)(\lambda_i^2 \alpha_1 + \alpha_2)}.$$

We claim that $\alpha_1 > \alpha_2$ if and only if $n_1 > n_2$, thus proving the dichotomy. When $\alpha_1 > \alpha_2$, the first part of equation (3.2) implies $\alpha_1 > \frac{1}{2}(1 - \frac{p}{n_1 + n_2})$. The second part of equation (3.2) gives

$$\begin{aligned} \frac{n_1}{n_1 + n_2} &= \alpha_1 + \frac{1}{n_1 + n_2} \sum_{i=1}^p \frac{\lambda_i^2 \alpha_1}{\lambda_i^2 \alpha_1 + \alpha_2} \\ &> \frac{1}{2} \left(1 - \frac{p}{n_1 + n_2} \right) + \frac{1}{n_1 + n_2} \sum_{i=1}^{p/2} \left(\frac{\lambda_i^2}{\lambda_i^2 + 1} + \frac{\lambda_i^{-2}}{\lambda_i^{-2} + 1} \right) \\ &= \frac{1}{2} \left(1 - \frac{p}{n_1 + n_2} \right) + \frac{p}{2(n_1 + n_2)} = \frac{1}{2}, \end{aligned}$$

implying $n_1 > n_2$. When $\alpha_1 < \alpha_2$, one can show $n_1 < n_2$ using similar steps. This completes the proof. \blacksquare

B.2 Proof of Proposition 5

Proof From the second part of equation (3.2), we get

$$\frac{n_1}{n_1 + n_2} = \alpha_1 + \frac{1}{n_1 + n_2} \sum_{i=1}^p \frac{\lambda_i^2 \alpha_1}{\lambda_i^2 \alpha_1 + \alpha_2} < \alpha_1 + \frac{p}{n_1 + n_2},$$

which gives that $\alpha_1 > \frac{n_1 - p}{n_1 + n_2}$. Combined with the first part of equation (3.2), it yields that

$$\alpha_2 = 1 - \frac{p}{n_1 + n_2} - \alpha_1 < \frac{n_2}{n_1 + n_2}.$$

Hence, we have $\alpha_1/\alpha_2 > (n_1 - p)/n_2 \geq c^{-1}$.

Now, abbreviating $\mathbf{x}_p = (x_1, \dots, x_p)$ with $x_i = \lambda_i^2$, the trace of $(\alpha_1 M^\top M + \alpha_2)^{-1}$ can be written as

$$f_p(\mathbf{x}_p) := \text{Tr} \left[\frac{1}{\alpha_1 M^\top M + \alpha_2} \right] = \sum_{i=1}^p \frac{1}{\alpha_1 x_i + \alpha_2}.$$

To conclude the proof, it suffices to prove the following claim: for any $k \in \{1, \dots, p\}$, $c \leq b \leq 1$, and α_1, α_2 such that $\alpha_1/\alpha_2 \geq c^{-1}$, $f_k(\mathbf{x}_k)$ is minimized at $\mathbf{x}_k = (b, \dots, b) \in \mathbb{R}^k$ subject to the constraints $c \leq x_i \leq c^{-1}$ and $\prod_{i=1}^k x_i = b^k$. Note that in contrast to the original setting, α_1 and α_2 are now taken as fixed parameters that do not depend on \mathbf{x}_k , which will greatly simplify the optimization problem.

We prove the above claim by induction. First, the claim is trivial when $k = 1$. For $k = 2$, a similar argument as in the proof of Proposition 4 indeed gives that f_2 is minimized at $x_1 = x_2 = b$ as long as $b\alpha_1 \geq \alpha_2$. Now, suppose the claim holds for f_k for some $2 \leq k \leq p-1$. Consider $\mathbf{x}_{k+1} = (x_1, \dots, x_{k+1})$ with $c \leq x_i \leq c^{-1}$ and $\prod_{i=1}^{k+1} x_i = b^{k+1}$. Renaming the indices if necessary, we may assume that x_{k+1} is the largest entry, i.e., $x_{k+1} \geq \max_{i=1}^k x_i$. Then, we can write $x_{k+1} = ba$ for some $1 \leq a \leq \min\{(bc)^{-1}, b^k c^{-k}\}$, where the condition $a \leq b^k c^{-k}$ comes from the constraints $\prod_{i=1}^{k+1} x_i = b^{k+1}$ and $\prod_{i=1}^k x_i \geq c^k$. The remaining variables satisfy $c \leq x_i \leq c^{-1}$ and $\prod_{i=1}^k x_i = b^k/a$. Applying the induction hypothesis for f_k , we get that

$$\sum_{i=1}^k \frac{1}{\alpha_1 x_i + \alpha_2} \geq \frac{k}{\alpha_1 b a^{-1/k} + \alpha_2}.$$

Hence, for $\mathbf{x}_{k+1} = (x_1, \dots, x_k, ba)$, we have

$$f_{k+1}(\mathbf{x}_{k+1}) \geq g_k(a) := \frac{k}{\alpha_1 b a^{-1/k} + \alpha_2} + \frac{1}{\alpha_1 b a + \alpha_2}.$$

To conclude the proof, we only need to show that $g_k(a)$ achieves a minimum when $a = 1$, i.e.,

$$g_k(a) \geq \frac{k+1}{\alpha_1 b + \alpha_2},$$

which is equivalent to

$$\frac{k(a^{1/k} - 1)}{w + a^{1/k}} \geq \frac{a - 1}{wa + 1}, \quad (\text{B.1})$$

where we introduce the simplified notation $w = b\alpha_1/\alpha_2 > bc^{-1} \geq 1$. We rewrite the inequality (B.1) as $h_k(a) \geq 0$ with

$$h_k(a) := k(a^{\frac{1}{k}} - 1)(wa + 1) - (w + a^{\frac{1}{k}})(a - 1).$$

Its derivative satisfies

$$h'_k(a) = w(k + 1)(a^{\frac{1}{k}} - 1) - \frac{k + 1}{k} a^{\frac{1}{k}}(1 - a^{-1}) \geq 0,$$

where, in the second step, we apply that $k(1 - a^{-\frac{1}{k}}) \geq 1 - a^{-1}$. Hence, we have $h_k(a) \geq h_k(1) = 0$ for $a \geq 1$, i.e., the inequality of (B.1) indeed holds, which concludes the proof. ■

B.3 Proof of Theorem 3

For the rest of this section, we present the proof of Theorem 3, which is one of the main results of this paper. The central quantity of interest is the matrix $\hat{\Sigma}^{-1}$ in equation (2.7). Assume that $M = \Sigma^{(1)1/2} \Sigma^{(2)-1/2}$ has a singular value decomposition

$$M = U\Lambda V^\top, \quad \text{where } \Lambda := \text{diag}(\lambda_1, \dots, \lambda_p).$$

Then, we can write equation (2.10) as

$$L_{\text{var}} = \sigma^2 \text{Tr}[\Sigma^{(2)} \hat{\Sigma}^{-1}] = \frac{\sigma^2}{n} \text{Tr}[W^{-1}], \quad (\text{B.2})$$

where we denote $n := n_1 + n_2$ and

$$W := n^{-1} \left(\Lambda U^\top Z^{(1)\top} Z^{(1)} U \Lambda + V^\top Z^{(2)\top} Z^{(2)} V \right).$$

The *resolvent* or *Green's function* of W is defined as $(W - z\text{Id}_{p \times p})^{-1}$ for $z \in \mathbb{C}$.

In this section, we will prove a local convergence of this resolvent with a *sharp convergence rate*, which is conventionally referred to as “the local law” (Bloemendal et al., 2014; Erdős and Yau, 2017; Knowles and Yin, 2016). In particular, the key technical challenge is to establish an (almost) sharp convergence rate of the bias and variance equations to their asymptotic limits.

B.3.1 RESOLVENT AND LOCAL LAW

We can write W as $W = FF^\top$ for a $p \times n$ matrix

$$F := n^{-1/2} [\Lambda U^\top Z^{(1)\top}, V^\top Z^{(2)\top}]. \quad (\text{B.3})$$

We introduce a convenient self-adjoint linearization trick to study W . It has been proved to be useful in studying random matrices of Gram type (Knowles and Yin, 2016).

Definition 26 (Self-adjoint linearization and resolvent) We define the following $(p+n) \times (p+n)$ symmetric block matrix

$$H := \begin{pmatrix} 0 & F \\ F^\top & 0 \end{pmatrix}, \quad (\text{B.4})$$

and its resolvent as

$$G(z) \equiv G(Z^{(1)}, Z^{(2)}, z) := \left[H - \begin{pmatrix} z\text{Id}_{p \times p} & 0 \\ 0 & \text{Id}_{(n_1+n_2) \times (n_1+n_2)} \end{pmatrix} \right]^{-1}, \quad z \in \mathbb{C},$$

as long as the inverse exists. Furthermore, we define the following (weighted) partial traces

$$\begin{aligned} m(z) &:= \frac{1}{p} \sum_{i \in \mathcal{I}_0} G_{ii}(z), & m_0(z) &:= \frac{1}{p} \sum_{i \in \mathcal{I}_0} \lambda_i^2 G_{ii}(z), \\ m_1(z) &:= \frac{1}{n_1} \sum_{\mu \in \mathcal{I}_1} G_{\mu\mu}(z), & m_2(z) &:= \frac{1}{n_2} \sum_{\nu \in \mathcal{I}_2} G_{\nu\nu}(z), \end{aligned} \quad (\text{B.5})$$

where \mathcal{I}_i , $i = 0, 1, 2$, are index sets defined as

$$\mathcal{I}_0 := \llbracket 1, p \rrbracket, \quad \mathcal{I}_1 := \llbracket p+1, p+n_1 \rrbracket, \quad \mathcal{I}_2 := \llbracket p+n_1+1, p+n_1+n_2 \rrbracket.$$

We will consistently use latin letters $i, j \in \mathcal{I}_0$ and greek letters $\mu, \nu \in \mathcal{I}_1 \cup \mathcal{I}_2$. Correspondingly, the indices $Z^{(1)}$ and $Z^{(2)}$ are labelled as

$$Z^{(1)} = \left[Z_{\mu i}^{(1)} : i \in \mathcal{I}_0, \mu \in \mathcal{I}_1 \right], \quad Z^{(2)} = \left[Z_{\nu i}^{(2)} : i \in \mathcal{I}_0, \nu \in \mathcal{I}_2 \right].$$

Moreover, we define the set of all indices $\mathcal{I} := \mathcal{I}_0 \cup \mathcal{I}_1 \cup \mathcal{I}_2$, and label the indices in \mathcal{I} as \mathbf{a} , \mathbf{b} , \mathbf{c} and so on.

Using the Schur complement formula for the inverse of a block matrix, we get that

$$G(z) = \begin{pmatrix} (W - z\text{Id})^{-1} & (W - z\text{Id})^{-1}F \\ F^\top(W - z\text{Id})^{-1} & z(F^\top F - z\text{Id})^{-1} \end{pmatrix}. \quad (\text{B.6})$$

In particular, the upper left block of G is exactly the resolvent of W we are interested in. Compared with $(W - z\text{Id})^{-1}$, it turns out that $G(z)$ is more convenient to deal with because H is a linear function of $Z^{(1)}$ and $Z^{(2)}$. This is why we have chosen to work with $G(z)$.

We define the matrix limit of $G(z)$ as

$$\mathfrak{G}(z) := \begin{pmatrix} [\alpha_1(z)\Lambda^2 + (\alpha_2(z) - z)]^{-1} & 0 & 0 \\ 0 & -\frac{n}{n_1}\alpha_1(z)\text{Id}_{n_1 \times n_1} & 0 \\ 0 & 0 & -\frac{n}{n_2}\alpha_2(z)\text{Id}_{n_2 \times n_2} \end{pmatrix}, \quad (\text{B.7})$$

where $(\alpha_1(z), \alpha_2(z))$ is the unique solution to the following system of self-consistent equations:

$$\begin{aligned} \alpha_1(z) + \alpha_2(z) &= 1 - \frac{1}{n} \sum_{i=1}^p \frac{\lambda_i^2 \alpha_1(z) + \alpha_2(z)}{\lambda_i^2 \alpha_1(z) + \alpha_2(z) - z}, \\ \alpha_1(z) + \frac{1}{n} \sum_{i=1}^p \frac{\lambda_i^2 \alpha_1(z)}{\lambda_i^2 \alpha_1(z) + \alpha_2(z) - z} &= \frac{n_1}{n}, \end{aligned} \quad (\text{B.8})$$

such that $\text{Im } \alpha_1(z) \leq 0$ and $\text{Im } \alpha_2(z) \leq 0$ whenever $\text{Im } z > 0$. The existence and uniqueness of solutions to the above system will be proved in Lemma 28.

We now state the main result, Theorem 27, of this section, which shows that for z in a small neighborhood around 0, $G(z)$ converges to the limit $\mathfrak{G}(z)$ when p goes to infinity. Moreover, it also gives an almost sharp convergence rate of $G(z)$. Such an estimate is conventionally called an *anisotropic local law* (Knowles and Yin, 2016). We define a domain of the spectral parameter z as

$$\mathbf{D} := \{z = E + i\eta \in \mathbb{C}_+ : |z| \leq (\log n)^{-1}\}.$$

Theorem 27 *Suppose that $Z^{(1)}$, $Z^{(2)}$, ρ_1 and ρ_2 satisfy Assumption 1. Suppose that $Z^{(1)}$ and $Z^{(2)}$ satisfy the bounded support condition (A.1) with $Q = n^{2/\varphi}$. Suppose that the singular values of M satisfy that*

$$\lambda_p \leq \dots \leq \lambda_2 \leq \lambda_1 \leq \tau^{-1}. \quad (\text{B.9})$$

Then, the following local laws hold.

(1) **Averaged local law:** *We have that*

$$\sup_{z \in \mathbf{D}} \left| p^{-1} \sum_{i \in \mathcal{I}_0} [G_{ii}(z) - \mathfrak{G}_{ii}(z)] \right| \prec (np)^{-1/2} Q, \quad (\text{B.10})$$

$$\sup_{z \in \mathbf{D}} \left| p^{-1} \sum_{i \in \mathcal{I}_0} \lambda_i^2 [G_{ii}(z) - \mathfrak{G}_{ii}(z)] \right| \prec (np)^{-1/2} Q. \quad (\text{B.11})$$

(2) **Anisotropic local law:** *For any deterministic unit vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^{p+n}$, we have that*

$$\sup_{z \in \mathbf{D}} \left| \mathbf{u}^\top [G(z) - \mathfrak{G}(z)] \mathbf{v} \right| \prec n^{-1/2} Q. \quad (\text{B.12})$$

With Theorem 27, we can complete the proof of Theorem 3 with a standard cutoff argument.

Proof Similar to equation (A.12), we introduce the truncated matrices $\tilde{Z}^{(1)}$ and $\tilde{Z}^{(2)}$ with entries

$$\tilde{Z}_{\mu i}^{(1)} := \mathbf{1} \left(|Z_{\mu i}^{(1)}| \leq Q \log n \right) \cdot Z_{\mu i}^{(1)}, \quad \tilde{Z}_{\nu i}^{(2)} := \mathbf{1} \left(|Z_{\nu i}^{(2)}| \leq Q \log n \right) \cdot Z_{\nu i}^{(2)}, \quad (\text{B.13})$$

for $Q = n^{2/\varphi}$. From equation (A.2), we get that

$$\mathbb{P}(\tilde{Z}^{(1)} = Z^{(1)}, \tilde{Z}^{(2)} = Z^{(2)}) = 1 - O((\log n)^{-\varphi}).$$

By equations (A.15) and (A.16), we have that

$$\begin{aligned} \mathbb{E}|\tilde{Z}_{\mu i}^{(1)}| &= O(n^{-3/2}), \quad \mathbb{E}|\tilde{Z}_{\nu i}^{(2)}| = O(n^{-3/2}), \\ \mathbb{E}|\tilde{Z}_{\mu i}^{(1)}|^2 &= 1 + O(n^{-1}), \quad \mathbb{E}|\tilde{Z}_{\nu i}^{(2)}|^2 = 1 + O(n^{-1}), \end{aligned} \quad (\text{B.14})$$

and

$$\|\mathbb{E}\tilde{Z}^{(1)}\| = O(n^{-1/2}), \quad \|\mathbb{E}\tilde{Z}^{(2)}\| = O(n^{-1/2}). \quad (\text{B.15})$$

Then, we centralize and rescale $\tilde{Z}^{(1)}$ and $\tilde{Z}^{(2)}$ as

$$\hat{Z}^{(1)} := (\mathbb{E}|\tilde{Z}_{\mu i}^{(1)}|^2)^{-1/2}(\tilde{Z}^{(1)} - \mathbb{E}\tilde{Z}^{(1)}), \quad \hat{Z}^{(2)} := (\mathbb{E}|\tilde{Z}_{\nu i}^{(2)}|^2)^{-1/2}(\tilde{Z}^{(2)} - \mathbb{E}\tilde{Z}^{(2)}),$$

for arbitrary $\mu \in \mathcal{I}_1$, $\nu \in \mathcal{I}_2$ and $i \in \mathcal{I}_0$. Now, $\hat{Z}^{(1)}$ and $\hat{Z}^{(2)}$ satisfy the assumptions of Theorem 27. Hence, $G(\hat{Z}^{(1)}, \hat{Z}^{(2)}, z)$ satisfies equation (B.10), where $G(\hat{Z}^{(1)}, \hat{Z}^{(2)}, z)$ is defined in the same way as $G(z)$ but with $(Z^{(1)}, Z^{(2)})$ replaced by $(\hat{Z}^{(1)}, \hat{Z}^{(2)})$. By equations (B.14) and (B.15), we have that for $k = 1, 2$,

$$\|\hat{Z}^{(k)} - \tilde{Z}^{(k)}\| \lesssim n^{-1}\|\hat{Z}^{(k)}\| + \|\mathbb{E}\tilde{Z}^{(k)}\| \prec n^{-1/2},$$

where we use Lemma 22 to bound the operator norm of $\hat{Z}^{(k)}$. Together with the estimate (B.36) below, this bound implies that

$$\max_{i \in \mathcal{I}_0} \left| G_{ii}(\hat{Z}^{(1)}, \hat{Z}^{(2)}, z) - G_{ii}(\tilde{Z}^{(1)}, \tilde{Z}^{(2)}, z) \right| \prec n^{-1/2} \sum_{k=1}^2 \|\hat{Z}^{(k)} - \tilde{Z}^{(k)}\| \prec n^{-1}.$$

Combining this estimate with the local law (B.10) for $G(\hat{Z}^{(1)}, \hat{Z}^{(2)}, z)$, we obtain that (B.10) also holds for $G(z) \equiv G(Z^{(1)}, Z^{(2)}, z)$ on the event $\Xi_1 := \{\tilde{Z}^{(1)} = Z^{(1)}, \tilde{Z}^{(2)} = Z^{(2)}\}$.

Now, we are ready to prove equation (3.1). By estimate (B.2), we have that

$$L_{\text{var}}(a) = \frac{\sigma^2}{n} \sum_{i \in \mathcal{I}_0} G_{ii}(0).$$

We also notice that the equations in estimate (B.8) reduce to the equations in equation (3.2) when $z = 0$, which gives that $\alpha_1 = \alpha_1(0)$ and $\alpha_2 = \alpha_2(0)$. Hence, the partial trace of \mathfrak{G} in (B.7) over \mathcal{I}_0 is equal to

$$\sum_{i \in \mathcal{I}_0} \mathfrak{G}_{ii}(0) = \text{Tr} \left[\frac{1}{\alpha_1 \Lambda^2 + \alpha_2} \right] = \text{Tr} \left[\frac{1}{\alpha_1 M^\top M + \alpha_2} \right].$$

Now, applying estimate (B.10) to $G(z)$, we conclude that on the event Ξ_1 ,

$$\left| L_{\text{var}} - \frac{\sigma^2}{n_1 + n_2} \text{Tr} \left[\frac{1}{\alpha_1 M^\top M + \alpha_2} \right] \right| \prec \frac{n^{2/\varphi}}{(np)^{1/2}} \cdot \frac{p\sigma^2}{n_1 + n_2}.$$

This concludes the proof of Theorem 3. ■

B.3.2 SELF-CONSISTENT EQUATIONS

Proof The rest of this section is devoted to the proof of Theorem 27. In this subsection, we show that the self-consistent equation (B.8) has a unique solution $(\alpha_1(z), \alpha_2(z))$ for any $z \in \mathbf{D}$. Otherwise, Theorem 27 will be a vacuous result. For simplicity of notation, we define the following ratios

$$\gamma_n := \frac{p}{n}, \quad r_1 := \frac{n_1}{n}, \quad r_2 := \frac{n_2}{n}.$$

When $z = 0$, estimate (B.8) reduces to the system of equations in (3.2), from which we can derive an equation of $\alpha_1 \equiv \alpha_1(0)$ only:

$$f(\alpha_1) = r_1, \quad \text{with} \quad f(\alpha_1) := \alpha_1 + \frac{1}{n} \sum_{i=1}^p \frac{\lambda_i^2 \alpha_1}{\lambda_i^2 \alpha_1 + (1 - \gamma_n - \alpha_1)}. \quad (\text{B.16})$$

We can calculate that

$$f'(\alpha_1) = 1 + \frac{1}{n} \sum_{i=1}^p \frac{\lambda_i^2 (1 - \gamma_n)}{[\lambda_i^2 \alpha_1 + (1 - \gamma_n - \alpha_1)]^2} > 0.$$

Hence, f is strictly increasing on $[0, 1 - \gamma_n]$. Moreover, we have $f(0) = 0 < r_1$, $f(1 - \gamma_n) = 1 > r_1$, and $f(r_1) > r_1$ if $r_1 \leq 1 - \gamma_n$. Hence, there exists a unique solution α_1 to (B.16) satisfying $0 < \alpha_1 < \min\{1 - \gamma_n, r_1\}$. Furthermore, using that $f'(x) = O(1)$ for any fixed $x \in (0, 1 - \gamma_n)$, it is not hard to check that

$$r_1 \tau \leq \alpha_1 \leq \min\{1 - \gamma_n, r_1\} \quad (\text{B.17})$$

for a small constant $\tau > 0$. From equation (3.2), we can also derive a equation of $\alpha_2 \equiv \alpha_2(0)$ only. With a similar argument as above, we get that

$$r_2 \tau \leq \alpha_2 \leq \min\{1 - \gamma_n, r_2\} \quad (\text{B.18})$$

for a small constant $\tau > 0$.

Next, we prove the existence and uniqueness of the solution to the self-consistent equation (B.8) for a general $z \in \mathbf{D}$. For the proof of Theorem 27, it is more convenient to use the following rescaled functions of $\alpha_1(z)$ and $\alpha_2(z)$:

$$m_{1c}(z) := -r_1^{-1} \alpha_1(z), \quad m_{2c}(z) := -r_2^{-1} \alpha_2(z), \quad (\text{B.19})$$

which, as we will see later, are the classical values (i.e., asymptotic limits) of $m_1(z)$ and $m_2(z)$, respectively. Moreover, it is not hard to check that equation (B.8) is equivalent to the following system of self-consistent equations of $m_{1c}(z), m_{2c}(z)$:

$$\begin{aligned} \frac{1}{m_{1c}} &= \frac{\gamma_n}{p} \sum_{i=1}^p \frac{\lambda_i^2}{z + \lambda_i^2 r_1 m_{1c} + r_2 m_{2c}} - 1, \\ \frac{1}{m_{2c}} &= \frac{\gamma_n}{p} \sum_{i=1}^p \frac{1}{z + \lambda_i^2 r_1 m_{1c} + r_2 m_{2c}} - 1. \end{aligned} \quad (\text{B.20})$$

When $z = 0$, using equations (3.2), (B.17), and (B.18), we get that

$$\tau \leq -m_{1c}(0) \leq 1, \quad \tau \leq -m_{2c}(0) \leq 1, \quad -r_1 m_{1c}(0) - r_2 m_{2c}(0) = 1 - \gamma_n. \quad (\text{B.21})$$

Now, we claim the following lemma, which gives the existence and uniqueness of the solutions $m_{1c}(z)$ and $m_{2c}(z)$ to the system of equations (B.20).

Lemma 28 *There exist constants $c_0, C_0 > 0$ depending only on τ in Assumption 1 and equation (B.21) such that the following statements hold. There exists a unique solution $m_{1c}(z), m_{2c}(z)$ to equation (B.20) under the conditions*

$$|z| \leq c_0, \quad |m_{1c}(z) - m_{1c}(0)| + |m_{2c}(z) - m_{2c}(0)| \leq c_0. \quad (\text{B.22})$$

Moreover, the solution satisfies

$$|m_{1c}(z) - m_{1c}(0)| + |m_{2c}(z) - m_{2c}(0)| \leq C_0 |z|. \quad (\text{B.23})$$

Proof The proof is based on a standard application of the contraction principle. First, it is easy to check that the system of equations in equation (B.20) is equivalent to

$$r_1 m_{1c} = -(1 - \gamma_n) - r_2 m_{2c} - z (m_{2c}^{-1} + 1), \quad g_z(m_{2c}(z)) = 1, \quad (\text{B.24})$$

where

$$g_z(m_{2c}) := -m_{2c} + \frac{\gamma_n}{p} \sum_{i=1}^p \frac{m_{2c}}{z - \lambda_i^2(1 - \gamma_n) + (1 - \lambda_i^2)r_2 m_{2c} - \lambda_i^2 z (m_{2c}^{-1} + 1)}.$$

We first show that there exists a unique solution $m_{2c}(z)$ to the equation $g_z(m_{2c}(z)) = 1$ under the conditions in equation (B.22). We abbreviate $\delta(z) := m_{2c}(z) - m_{2c}(0)$. From equation (B.24), we obtain that

$$0 = [g_z(m_{2c}(z)) - g_0(m_{2c}(0)) - g'_z(m_{2c}(0))\delta(z)] + g'_z(m_{2c}(0))\delta(z),$$

which gives that

$$\delta(z) = -\frac{g_z(m_{2c}(0)) - g_0(m_{2c}(0))}{g'_z(m_{2c}(0))} - \frac{g_z(m_{2c}(0) + \delta(z)) - g_z(m_{2c}(0)) - g'_z(m_{2c}(0))\delta(z)}{g'_z(m_{2c}(0))}.$$

Inspired by this equation, we define iteratively a sequence $\delta^{(k)}(z) \in \mathbb{C}$ such that $\delta^{(0)} = 0$ and

$$\delta^{(k+1)} = -\frac{g_z(m_{2c}(0)) - g_0(m_{2c}(0))}{g'_z(m_{2c}(0))} - \frac{g_z(m_{2c}(0) + \delta^{(k)}) - g_z(m_{2c}(0)) - g'_z(m_{2c}(0))\delta^{(k)}}{g'_z(m_{2c}(0))}. \quad (\text{B.25})$$

Then, equation (B.25) defines a mapping $h_z : \mathbb{C} \rightarrow \mathbb{C}$, which maps $\delta^{(k)}$ to $\delta^{(k+1)} = h_z(\delta^{(k)})$.

Through a straightforward calculation, we get that

$$g'_z(m_{2c}(0)) = -1 - \frac{\gamma_n}{p} \sum_{i=1}^p \frac{\lambda_i^2(1 - \gamma_n) - z [1 - \lambda_i^2 (2m_{2c}^{-1}(0) + 1)]}{[z - \lambda_i^2(1 - \gamma_n) + (1 - \lambda_i^2)r_2 m_{2c}(0) - \lambda_i^2 z (m_{2c}^{-1}(0) + 1)]^2}.$$

Using equations (B.9) and (B.21), it is easy to check that

$$|z - \lambda_i^2(1 - \gamma_n) + (1 - \lambda_i^2)r_2 m_{2c}(0) - \lambda_i^2 z (m_{2c}^{-1}(0) + 1)| \geq c_\tau - c_\tau^{-1}|z|,$$

for a constant $c_\tau > 0$ depending only on τ . Hence, as long as we choose $c_0 \leq c_\tau^2/2$, we have

$$|z - \lambda_i^2(1 - \gamma_n) + (1 - \lambda_i^2)r_2 m_{2c}(0) - \lambda_i^2 z (m_{2c}^{-1}(0) + 1)| \geq c_\tau/2.$$

With this estimate, we can check that there exist constants $\tilde{c}_\tau, \tilde{C}_\tau > 0$ depending only on τ such that the following estimates hold: for all z, δ_1, δ_2 such that $\min\{|z|, |\delta_1|, |\delta_2|\} \leq \tilde{c}_\tau$,

$$\left| \frac{1}{g'_z(m_{2c}(0))} \right| \leq \tilde{C}_\tau, \quad \left| \frac{g_z(m_{2c}(0)) - g_0(m_{2c}(0))}{g'_z(m_{2c}(0))} \right| \leq \tilde{C}_\tau |z|, \quad (\text{B.26})$$

and

$$\left| \frac{g_z(m_{2c}(0) + \delta_1) - g_z(m_{2c}(0) + \delta_2) - g'_z(m_{2c}(0))(\delta_1 - \delta_2)}{g'_z(m_{2c}(0))} \right| \leq \tilde{C}_\tau |\delta_1 - \delta_2|^2. \quad (\text{B.27})$$

Using equations (B.26) and (B.27), we find that there exists a sufficiently small constant $c_1 > 0$ depending only on \tilde{C}_τ such that $h_z : B_d \rightarrow B_d$ is a self-mapping on the ball $B_d := \{\delta \in \mathbb{C} : |\delta| \leq d\}$, as long as $d \leq c_1$ and $|z| \leq c_1$. Now, it suffices to prove that h_z restricted to B_d is a contraction, which then implies that $\delta := \lim_{k \rightarrow \infty} \delta^{(k)}$ exists and $m_{2c}(0) + \delta(z)$ is a unique solution to the equation $g_z(m_{2c}(z)) = 1$ subject to $|\delta| \leq d$.

From the iteration relation (B.25), using equation (B.27) we readily get that

$$\delta^{(k+1)} - \delta^{(k)} = h_z(\delta^{(k)}) - h_z(\delta^{(k-1)}) \leq \tilde{C}_\tau |\delta^{(k)} - \delta^{(k-1)}|^2.$$

Hence, h is indeed a contraction mapping on B_d as long as d is chosen sufficiently small such that $2d\tilde{C}_\tau \leq 1/2$. This proves both the existence and uniqueness of the solution $m_{2c}(z) = m_{2c}(0) + \delta(z)$ if we choose c_0 in equation (B.22) to be $c_0 = \min\{c_\tau^2/2, c_1, (4\tilde{C}_\tau)^{-1}\}$. After obtaining $m_{2c}(z)$, we then solve $m_{1c}(z)$ using the first equation in (B.24).

It remains to prove the estimate (B.23). Using equation (B.26) and $\delta^{(0)} = 0$, we can obtain from equation (B.25) that $|\delta^{(1)}(z)| \leq \tilde{C}_\tau |z|$. Then, by the contraction mapping, we have the bound

$$|\delta| \leq \sum_{k=0}^{\infty} |\delta^{(k+1)} - \delta^{(k)}| \leq 2\tilde{C}_\tau |z| \Rightarrow |m_{2c}(z) - m_{2c}(0)| \leq 2\tilde{C}_\tau |z|.$$

Finally, using the first equation in (B.24), we also obtain the bound $|m_{1c}(z) - m_{1c}(0)| \lesssim |z|$. ■

As a byproduct of the above contraction mapping argument, we also obtain the following stability result that will be used in the proof of Theorem 27. Roughly speaking, it states that if two analytic functions $m_1(z)$ and $m_2(z)$ satisfy the self-consistent equation (B.20) approximately up to some small errors, then $m_1(z)$ and $m_2(z)$ will be close to the solutions $m_{1c}(z)$ and $m_{2c}(z)$.

Lemma 29 *There exist constants $c_0, C_0 > 0$ depending only on τ in Assumption 1 and (B.21) such that the system of self-consistent equations (B.20) is stable in the following sense. Suppose $|z| \leq c_0$, and $m_1(z)$ and $m_2(z)$ are analytic functions of z such that*

$$|m_1(z) - m_{1c}(0)| + |m_2(z) - m_{2c}(0)| \leq c_0. \quad (\text{B.28})$$

Moreover, assume that m_1, m_2 satisfies the system of equations

$$\begin{aligned} \frac{1}{m_1} + 1 - \frac{\gamma_n}{p} \sum_{i=1}^p \frac{\lambda_i^2}{z + \lambda_i^2 r_1 m_1 + r_2 m_2} &= \mathcal{E}_1, \\ \frac{1}{m_2} + 1 - \frac{\gamma_n}{p} \sum_{i=1}^p \frac{1}{z + \lambda_i^2 r_1 m_1 + r_2 m_2} &= \mathcal{E}_2, \end{aligned} \quad (\text{B.29})$$

for some (deterministic or random) errors such that $|\mathcal{E}_1| + |\mathcal{E}_2| \leq (\log n)^{-1/2}$. Then, we have

$$|m_1(z) - m_{1c}(z)| + |m_2(z) - m_{2c}(z)| \leq C_0 (|\mathcal{E}_1| + |\mathcal{E}_2|).$$

Proof Under condition (B.28), we can obtain equation (B.24) approximately:

$$r_1 m_1 = -(1 - \gamma_n) - r_2 m_2 - z (m_2^{-1} + 1) + \tilde{\mathcal{E}}_1(z), \quad g_z(m_2(z)) = 1 + \tilde{\mathcal{E}}_2(z), \quad (\text{B.30})$$

where the errors satisfy that $|\tilde{\mathcal{E}}_1(z)| + |\tilde{\mathcal{E}}_2(z)| \lesssim |\mathcal{E}_1| + |\mathcal{E}_2|$. Then, we subtract equation (B.24) from equation (B.30), and consider the contraction principle for the function $\delta(z) := m_2(z) - m_{2c}(z)$. The rest of the proof is exactly the same as the one for Lemma 28, so we omit the details. \blacksquare

B.3.3 MULTIVARIATE GAUSSIAN CASE

One main difficulty for the proof is that the entries of $Z^{(1)}U\Lambda$ and $Z^{(2)}V$ are not independent. However, if the entries of $Z^{(1)}$ and $Z^{(2)}$ are i.i.d. Gaussian, then by the rotational invariance of the multivariate Gaussian distribution, we have that

$$Z^{(1)}U\Lambda \stackrel{d}{=} Z^{(1)}\Lambda, \quad Z^{(2)}V \stackrel{d}{=} Z^{(2)}.$$

In this case, the problem is reduced to proving the local laws for $G(z)$ with $U = \text{Id}_{p \times p}$ and $V = \text{Id}_{p \times p}$, such that the entries of $Z^{(1)}\Lambda$ and $Z^{(2)}$ are independent. In this case, we use the standard resolvent method, as in e.g., Bloemendal et al. (2014), to prove the following proposition. Note that if the entries of $Z^{(1)}$ and $Z^{(2)}$ are i.i.d. Gaussian, then $Z^{(1)}$ and $Z^{(2)}$ have bounded support $Q = 1$ by Remark 19.

Proposition 30 *Under the setting of Theorem 27, assume further that the entries of $Z^{(1)}$ and $Z^{(2)}$ are i.i.d. Gaussian random variables, and $U = V = \text{Id}_{p \times p}$. Then, the estimates (B.10), (B.11), and (B.12) hold with $Q = 1$.*

The proof of Proposition 30 is based on the following entrywise local law.

Lemma 31 (Entrywise local law) *Under the setting of Proposition 30, the averaged local laws (B.10) and (B.11) and the following entrywise local law hold with $Q = 1$:*

$$\sup_{z \in \mathbf{D}} \max_{a, b \in \mathcal{I}} |G_{ab}(z) - \mathfrak{G}_{ab}(z)| \prec n^{-1/2}. \quad (\text{B.31})$$

With Lemma 31, we can complete the proof of Proposition 30.

Proof With estimate (B.31), we can use the polynomial method in Section 5 of Bloemendal et al. (2014) to get the anisotropic local law (B.12) with $Q = 1$. The proof is exactly the same, except for some minor differences in notations. Hence, we omit the details. \blacksquare

The rest of this subsection is devoted to the proof of Lemma 31, where the resolvent G in Definition 26 becomes

$$G(z) = \begin{pmatrix} -z \text{Id}_{p \times p} & n^{-1/2} \Lambda Z^{(1)\top} & n^{-1/2} Z^{(2)\top} \\ n^{-1/2} Z^{(1)} \Lambda & -\text{Id}_{n_1 \times n_1} & 0 \\ n^{-1/2} Z^{(2)} & 0 & -\text{Id}_{n_2 \times n_2} \end{pmatrix}^{-1}. \quad (\text{B.32})$$

To deal with the matrix inverse, we introduce resolvent minors.

Definition 32 (Resolvent minors) *Given a $(p+n) \times (p+n)$ matrix \mathcal{A} and $\mathbf{c} \in \mathcal{I}$, the minor of \mathcal{A} after removing the \mathbf{c} -th row and column is a $(p+n-1) \times (p+n-1)$ matrix denoted by $\mathcal{A}^{(\mathbf{c})} := [\mathcal{A}_{\mathbf{ab}} : \mathbf{a}, \mathbf{b} \in \mathcal{I} \setminus \{\mathbf{c}\}]$. We keep the names of indices when defining $\mathcal{A}^{(\mathbf{c})}$, i.e., $\mathcal{A}_{\mathbf{ab}}^{(\mathbf{c})} = \mathcal{A}_{\mathbf{ab}}$ for $\mathbf{a}, \mathbf{b} \neq \mathbf{c}$. Correspondingly, we define the resolvent minor of $G(z)$ by*

$$G^{(\mathbf{c})}(z) := \left[\begin{pmatrix} -z\text{Id}_{p \times p} & n^{-1/2}\Lambda Z^{(1)\top} & n^{-1/2}Z^{(2)\top} \\ n^{-1/2}Z^{(1)}\Lambda & -\text{Id}_{n_1 \times n_1} & 0 \\ n^{-1/2}Z^{(2)} & 0 & -\text{Id}_{n_2 \times n_2} \end{pmatrix}^{(\mathbf{c})} \right]^{-1}.$$

We define the partial traces $m^{(\mathbf{c})}(z)$, $m_0^{(\mathbf{c})}(z)$, $m_1^{(\mathbf{c})}(z)$ and $m_2^{(\mathbf{c})}(z)$ by replacing $G(z)$ with $G^{(\mathbf{c})}(z)$ in equation (B.5). For convenience, we will adopt the convention that $G_{\mathbf{ab}}^{(\mathbf{c})} = 0$ if $\mathbf{a} = \mathbf{c}$ or $\mathbf{b} = \mathbf{c}$.

The following resolvent identities are important tools for our proof. All of them can be proved directly using Schur's complement formula, cf. Lemma 4.4 of Knowles and Yin (2016). Recall that the matrix F is defined in equation (B.3). We do not assume U and V are identity matrices for Lemma 33 and Lemma 34 below.

Lemma 33 *We have the following resolvent identities.*

(i) *For $i \in \mathcal{I}_0$ and $\mu \in \mathcal{I}_1 \cup \mathcal{I}_2$, we have*

$$\frac{1}{G_{ii}} = -z - \left(F G^{(i)} F^\top \right)_{ii}, \quad \frac{1}{G_{\mu\mu}} = -1 - \left(F^\top G^{(\mu)} F \right)_{\mu\mu}. \quad (\text{B.33})$$

(ii) *For $i \in \mathcal{I}_0$, $\mu \in \mathcal{I}_1 \cup \mathcal{I}_2$, $\mathbf{a} \in \mathcal{I} \setminus \{i\}$, and $\mathbf{b} \in \mathcal{I} \setminus \{\mu\}$, we have*

$$G_{i\mathbf{a}} = -G_{ii} \left(F G^{(i)} \right)_{i\mathbf{a}}, \quad G_{\mu\mathbf{b}} = -G_{\mu\mu} \left(F^\top G^{(\mu)} \right)_{\mu\mathbf{b}}. \quad (\text{B.34})$$

(iii) *For $\mathbf{c} \in \mathcal{I}$ and $\mathbf{a}, \mathbf{b} \in \mathcal{I} \setminus \{\mathbf{c}\}$, we have*

$$G_{\mathbf{ab}}^{(\mathbf{c})} = G_{\mathbf{ab}} - \frac{G_{\mathbf{ac}} G_{\mathbf{cb}}}{G_{\mathbf{cc}}}. \quad (\text{B.35})$$

We claim the following a priori estimate on the resolvent $G(z)$ for $z \in \mathbf{D}$.

Lemma 34 *Under the setting of Theorem 27, there exists a constant $C > 0$ such that with overwhelming probability the following estimates hold uniformly in $z, z' \in \mathbf{D}$:*

$$\|G(z)\| \leq C, \text{ and} \quad (\text{B.36})$$

$$\|G(z) - G(z')\| \leq C|z - z'|. \quad (\text{B.37})$$

Proof Our proof is a simple application of the spectral decomposition of G . Let

$$F = \sum_{k=1}^p \sqrt{\mu_k} \xi_k \zeta_k^\top, \quad \mu_1 \geq \mu_2 \geq \dots \geq \mu_p \geq 0 = \mu_{p+1} = \dots = \mu_n,$$

be a singular value decomposition of F , where $\{\xi_k\}_{k=1}^p$ are the left-singular vectors and $\{\zeta_k\}_{k=1}^n$ are the right-singular vectors. Then, using equation (B.6), we get that for $i, j \in \mathcal{I}_1$ and $\mu, \nu \in \mathcal{I}_1 \cup \mathcal{I}_2$,

$$G_{ij} = \sum_{k=1}^p \frac{\xi_k(i) \xi_k^\top(j)}{\mu_k - z}, \quad G_{\mu\nu} = z \sum_{k=1}^n \frac{\zeta_k(\mu) \zeta_k^\top(\nu)}{\mu_k - z}, \quad (\text{B.38})$$

$$G_{i\mu} = G_{\mu i} = \sum_{k=1}^p \frac{\sqrt{\mu_k} \xi_k(i) \zeta_k^\top(\mu)}{\mu_k - z}. \quad (\text{B.39})$$

Using the fact $n^{-1} V^\top Z^{(2)\top} Z^{(2)} V \preceq F F^\top$ and Lemma 22, we obtain that

$$\mu_p \geq \lambda_p \left(n^{-1} Z^{(2)\top} Z^{(2)} \right) \geq c_\tau \quad \text{w.o.p.,}$$

for a constant $c_\tau > 0$ depending only on τ . This further implies that

$$\inf_{z \in \mathbf{D}} \min_{1 \leq k \leq p} |\mu_k - z| \geq c_\tau - (\log n)^{-1}.$$

Combining this bound with estimates (B.38) and (B.39), we can conclude that estimates (B.36) and (B.37) are both true. \blacksquare

Now, we are ready to give the proof of Lemma 31.

Proof In the setting of Lemma 31, we can write equation (B.3) as

$$F \stackrel{d}{=} n^{-1/2} [\Lambda Z^{(1)\top}, Z^{(2)\top}].$$

We use the resolvent in equation (B.32) throughout the following proof. Our proof is divided into four steps.

Step 1: Large deviation estimates. In this step, we prove some large deviation estimates on the off-diagonal G entries and the following \mathcal{Z} variables. In analogy to Section 3 of Erdős et al. (2013) and Section 5 of Knowles and Yin (2016), we introduce the \mathcal{Z} variables

$$\mathcal{Z}_{\mathbf{a}} := (1 - \mathbb{E}_{\mathbf{a}})[(G_{\mathbf{a}\mathbf{a}})^{-1}], \quad \mathbf{a} \in \mathcal{I},$$

where $\mathbb{E}_{\mathbf{a}}[\cdot] := \mathbb{E}[\cdot \mid H^{(\mathbf{a})}]$ denotes the partial expectation over the entries in the \mathbf{a} -th row and column of H . Using equation (B.33), we get that for $i \in \mathcal{I}_0$,

$$\begin{aligned} \mathcal{Z}_i &= \frac{\lambda_i^2}{n} \sum_{\mu, \nu \in \mathcal{I}_1} G_{\mu\nu}^{(i)} \left(\delta_{\mu\nu} - Z_{\mu i}^{(1)} Z_{\nu i}^{(1)} \right) + \frac{1}{n} \sum_{\mu, \nu \in \mathcal{I}_2} G_{\mu\nu}^{(i)} \left(\delta_{\mu\nu} - Z_{\mu i}^{(2)} Z_{\nu i}^{(2)} \right) \\ &\quad - 2 \frac{\lambda_i}{n} \sum_{\mu \in \mathcal{I}_1, \nu \in \mathcal{I}_2} Z_{\mu i}^{(1)} Z_{\nu i}^{(2)} G_{\mu\nu}^{(i)}, \end{aligned} \quad (\text{B.40})$$

and for $\mu \in \mathcal{I}_1$ and $\nu \in \mathcal{I}_2$,

$$\mathcal{Z}_\mu = \frac{1}{n} \sum_{i,j \in \mathcal{I}_0} \lambda_i \lambda_j G_{ij}^{(\mu)} \left(\delta_{ij} - Z_{\mu i}^{(1)} Z_{\mu j}^{(1)} \right), \quad \mathcal{Z}_\nu = \frac{1}{n} \sum_{i,j \in \mathcal{I}_0} G_{ij}^{(\nu)} \left(\delta_{ij} - Z_{\nu i}^{(2)} Z_{\nu j}^{(2)} \right). \quad (\text{B.41})$$

Moreover, we introduce the random error

$$\Lambda_o := \max_{\mathfrak{a} \neq \mathfrak{b}} |G_{\mathfrak{a}\mathfrak{a}}^{-1} G_{\mathfrak{a}\mathfrak{b}}|, \quad (\text{B.42})$$

which controls the size of off-diagonal entries.

Lemma 35 *Under the assumptions of Proposition 30, the following estimate holds uniformly in $z \in \mathbf{D}$:*

$$\Lambda_o + \max_{\mathfrak{a} \in \mathcal{I}} |\mathcal{Z}_{\mathfrak{a}}| \prec n^{-1/2}. \quad (\text{B.43})$$

Proof Note that for $\mathfrak{a} \in \mathcal{I}$, $H^{(\mathfrak{a})}$ and $G^{(\mathfrak{a})}$ also satisfy the assumptions of Lemma 34. Hence, the estimates (B.36) and (B.37) hold for $G^{(\mathfrak{a})}$ with overwhelming probability. For any $i \in \mathcal{I}_0$, since $G^{(i)}$ is independent of the entries in the i -th row and column of H , we can apply equations (A.8), (A.9), and (A.10) to equation (B.40) to obtain that

$$\begin{aligned} |\mathcal{Z}_i| &\lesssim \frac{1}{n} \sum_{k=1}^2 \left| \sum_{\mu, \nu \in \mathcal{I}_k} G_{\mu\nu}^{(i)} \left(\delta_{\mu\nu} - Z_{\mu i}^{(k)} Z_{\nu i}^{(k)} \right) \right| + \frac{1}{n} \left| \sum_{\mu \in \mathcal{I}_1, \nu \in \mathcal{I}_2} Z_{\mu i}^{(1)} Z_{\nu i}^{(2)} G_{\mu\nu}^{(i)} \right| \\ &\prec \frac{1}{n} \left(\sum_{\mu, \nu \in \mathcal{I}_1 \cup \mathcal{I}_2} |G_{\mu\nu}^{(i)}|^2 \right)^{1/2} \prec n^{-1/2}. \end{aligned}$$

Here, in the last step we used equation (B.36) to get that for $\mu \in \mathcal{I}_1 \cup \mathcal{I}_2$,

$$\sum_{\nu \in \mathcal{I}_1 \cup \mathcal{I}_2} |G_{\mu\nu}^{(i)}|^2 \leq \sum_{\mathfrak{a} \in \mathcal{I}} |G_{\mu\mathfrak{a}}^{(i)}|^2 = \left(G^{(i)} G^{(i)*} \right)_{\mu\mu} = O(1), \quad (\text{B.44})$$

with overwhelming probability, where $G^{(i)*}$ denotes the conjugate transpose of $G^{(i)}$. Similarly, applying equations (A.8), (A.9), and (A.10) to \mathcal{Z}_μ and \mathcal{Z}_ν in equation (B.41) and using equation (B.36), we can obtain the same bound. This gives that $\max_{\mathfrak{a} \in \mathcal{I}} |\mathcal{Z}_{\mathfrak{a}}| \prec n^{-1/2}$.

Next, we prove the off-diagonal estimate on Λ_o . For $i \in \mathcal{I}_1$ and $\mathfrak{a} \in \mathcal{I} \setminus \{i\}$, using equations (B.34), (A.8) and (B.36), we can obtain that

$$\begin{aligned} |G_{ii}^{-1} G_{i\mathfrak{a}}| &\lesssim n^{-1/2} \left| \sum_{\mu \in \mathcal{I}_1} Z_{\mu i}^{(1)} G_{\mu\mathfrak{a}}^{(i)} \right| + n^{-1/2} \left| \sum_{\mu \in \mathcal{I}_2} Z_{\mu i}^{(2)} G_{\mu\mathfrak{a}}^{(i)} \right| \\ &\prec n^{-1/2} \left(\sum_{\mu \in \mathcal{I}_1 \cup \mathcal{I}_2} |G_{\mu\mathfrak{a}}^{(i)}|^2 \right)^{1/2} \prec n^{-1/2}. \end{aligned}$$

Using exactly the same argument, we can get a similar estimate on $|G_{\mu\mu}^{-1} G_{\mu\mathfrak{b}}|$ for $\mu \in \mathcal{I}_1 \cup \mathcal{I}_2$ and $\mathfrak{b} \in \mathcal{I} \setminus \{\mu\}$. This gives that $\Lambda_o \prec n^{-1/2}$. \blacksquare

Note that combining equation (B.43) with the bound in equation (B.36), we immediately conclude equation (B.31) for off-diagonal resolvent entries $G_{\mathbf{a}\mathbf{b}}$ with $\mathbf{a} \neq \mathbf{b}$.

Step 2: Self-consistent equations. In this step, we show that $(m_1(z), m_2(z))$ satisfies the system of approximate self-consistent equations in equation (B.29) for some small errors \mathcal{E}_1 and \mathcal{E}_2 . Later in Step 3, we will apply Lemma 29 to show that $(m_1(z), m_2(z))$ is close to $(m_{1c}(z), m_{2c}(z))$.

By equation (B.23), the following estimates hold for $z \in \mathbf{D}$:

$$|m_{1c}(z) - m_{1c}(0)| \lesssim (\log n)^{-1}, \quad |m_{2c}(z) - m_{2c}(0)| \lesssim (\log n)^{-1}.$$

Combining them with the estimates in (B.21), we obtain that uniformly in $z \in \mathbf{D}$,

$$|m_{1c}(z)| \sim |m_{2c}(z)| \sim 1, \quad |z + \lambda_i^2 r_1 m_{1c}(z) + r_2 m_{2c}(z)| \sim 1. \quad (\text{B.45})$$

Moreover, using equation (B.20), we get that uniformly in $z \in \mathbf{D}$,

$$|1 + \gamma_n m_c(z)| = |m_{2c}^{-1}(z)| \sim 1, \quad |1 + \gamma_n m_{0c}(z)| = |m_{1c}^{-1}(z)| \sim 1, \quad (\text{B.46})$$

where we introduce the notations

$$m_c(z) := -\frac{1}{p} \sum_{i=1}^p \frac{1}{z + \lambda_i^2 r_1 m_{1c}(z) + r_2 m_{2c}(z)}, \quad (\text{B.47})$$

$$m_{0c}(z) := -\frac{1}{p} \sum_{i=1}^p \frac{\lambda_i^2}{z + \lambda_i^2 r_1 m_{1c}(z) + r_2 m_{2c}(z)}. \quad (\text{B.48})$$

Later, we will see that $m_c(z)$ and $m_{0c}(z)$ are respectively the asymptotic limits of $m(z)$ and $m_0(z)$ defined in equation (B.5). Applying equation (B.45) to equation (B.7) and using equation (B.19), we get that

$$|\mathfrak{G}_{\mathbf{a}\mathbf{a}}(z)| \sim 1 \quad \text{uniformly in } z \in \mathbf{D} \quad \text{and} \quad \mathbf{a} \in \mathcal{I}.$$

Now, we define the following z -dependent event

$$\Xi(z) := \left\{ |m_1(z) - m_{1c}(0)| + |m_2(z) - m_{2c}(0)| \leq (\log n)^{-1/2} \right\}. \quad (\text{B.49})$$

With equation (B.45), we get that on $\Xi(z)$,

$$|m_1(z)| \sim |m_2(z)| \sim 1, \quad |z + \lambda_i^2 r_1 m_1(z) + r_2 m_2(z)| \sim 1. \quad (\text{B.50})$$

We claim the following key lemma, which shows that $(m_1(z), m_2(z))$ satisfies equation (B.29) on $\Xi(z)$ for some small (random) errors \mathcal{E}_1 and \mathcal{E}_2 .

Lemma 36 *Under the setting of Lemma 31, the following estimates hold uniformly in $z \in \mathbf{D}$:*

$$\mathbf{1}(\Xi) \left| \frac{1}{m_1} + 1 - \frac{\gamma_n}{p} \sum_{i=1}^p \frac{\lambda_i^2}{z + \lambda_i^2 r_1 m_1 + r_2 m_2} \right| \prec n^{-1} + n^{-\frac{1}{2}} \Theta + |[\mathcal{Z}]_0| + |[\mathcal{Z}]_1|, \quad (\text{B.51})$$

and

$$\mathbf{1}(\Xi) \left| \frac{1}{m_2} + 1 - \frac{\gamma_n}{p} \sum_{i=1}^p \frac{1}{z + \lambda_i^2 r_1 m_1 + r_2 m_2} \right| \prec n^{-1} + n^{-\frac{1}{2}} \Theta + |[\mathcal{Z}]| + |[\mathcal{Z}]_2|, \quad (\text{B.52})$$

where we introduce the notations

$$\Theta := |m_1(z) - m_{1c}(z)| + |m_2(z) - m_{2c}(z)|, \quad (\text{B.53})$$

and

$$\begin{aligned} [\mathcal{Z}] &:= \frac{1}{p} \sum_{i \in \mathcal{I}_0} \frac{\mathcal{Z}_i}{(z + \lambda_i^2 r_1 m_{1c} + r_2 m_{2c})^2}, \\ [\mathcal{Z}]_0 &:= \frac{1}{p} \sum_{i \in \mathcal{I}_0} \frac{\lambda_i^2 \mathcal{Z}_i}{(z + \lambda_i^2 r_1 m_{1c} + r_2 m_{2c})^2}, \\ [\mathcal{Z}]_1 &:= \frac{1}{n_1} \sum_{\mu \in \mathcal{I}_1} \mathcal{Z}_\mu, \quad [\mathcal{Z}]_2 := \frac{1}{n_2} \sum_{\nu \in \mathcal{I}_2} \mathcal{Z}_\nu. \end{aligned} \quad (\text{B.54})$$

Proof Using equations (B.33), (B.40) and (B.41), we can write that

$$\frac{1}{G_{ii}} = -z - \frac{\lambda_i^2}{n} \sum_{\mu \in \mathcal{I}_1} G_{\mu\mu}^{(i)} - \frac{1}{n} \sum_{\mu \in \mathcal{I}_2} G_{\mu\mu}^{(i)} + \mathcal{Z}_i = -z - \lambda_i^2 r_1 m_1 - r_2 m_2 + \mathcal{E}_i, \forall i \in \mathcal{I}_0, \quad (\text{B.55})$$

$$\frac{1}{G_{\mu\mu}} = -1 - \frac{1}{n} \sum_{i \in \mathcal{I}_0} \lambda_i^2 G_{ii}^{(\mu)} + \mathcal{Z}_\mu = -1 - \gamma_n m_0 + \mathcal{E}_\mu, \forall \mu \in \mathcal{I}_1, \quad (\text{B.56})$$

$$\frac{1}{G_{\nu\nu}} = -1 - \frac{1}{n} \sum_{i \in \mathcal{I}_0} G_{ii}^{(\nu)} + \mathcal{Z}_\nu = -1 - \gamma_n m + \mathcal{E}_\nu, \forall \nu \in \mathcal{I}_2, \quad (\text{B.57})$$

where we denote (recall the notations in equation (B.5) and Definition 32)

$$\begin{aligned} \mathcal{E}_i &:= \mathcal{Z}_i + \lambda_i^2 r_1 (m_1 - m_1^{(i)}) + r_2 (m_2 - m_2^{(i)}), \text{ and} \\ \mathcal{E}_\mu &:= \mathcal{Z}_\mu + \gamma_n (m_0 - m_0^{(\mu)}), \quad \mathcal{E}_\nu := \mathcal{Z}_\nu + \gamma_n (m - m^{(\nu)}). \end{aligned}$$

Using equations (B.35), (B.42) and (B.43), we can bound that

$$|m_1 - m_1^{(i)}| \leq \frac{1}{n_1} \sum_{\mu \in \mathcal{I}_1} \left| \frac{G_{\mu i} G_{i \mu}}{G_{ii}} \right| \leq |\Lambda_0|^2 |G_{ii}| \prec n^{-1}. \quad (\text{B.58})$$

Similarly, we also have that

$$|m_2 - m_2^{(i)}| \prec n^{-1}, \quad |m_0 - m_0^{(\mu)}| \prec n^{-1}, \quad |m - m^{(\nu)}| \prec n^{-1}, \quad (\text{B.59})$$

for $i \in \mathcal{I}_0$, $\mu \in \mathcal{I}_1$ and $\nu \in \mathcal{I}_2$. Combining the above estimates with equation (B.43), we obtain that

$$\max_{i \in \mathcal{I}_0} |\mathcal{E}_i| + \max_{\mu \in \mathcal{I}_1 \cup \mathcal{I}_2} |\mathcal{E}_\mu| \prec n^{-1/2}. \quad (\text{B.60})$$

From equation (B.55), we obtain that on Ξ ,

$$\begin{aligned} G_{ii} &= -\frac{1}{z + \lambda_i^2 r_1 m_1 + r_2 m_2} - \frac{\mathcal{E}_i}{(z + \lambda_i^2 r_1 m_1 + r_2 m_2)^2} + O_{\prec}(n^{-1}) \\ &= -\frac{1}{z + \lambda_i^2 r_1 m_1 + r_2 m_2} - \frac{\mathcal{Z}_i}{(z + \lambda_i^2 r_1 m_{1c} + r_2 m_{2c})^2} + O_{\prec}\left(n^{-1} + n^{-\frac{1}{2}}\Theta\right), \end{aligned} \quad (\text{B.61})$$

where in the first step we use estimates (B.60) and (B.50) on Ξ , and in the second step we use estimates (B.53), (B.58) and (B.60). Plugging equation (B.61) into the definitions of m and m_0 in equation (B.5) and using equation (B.54), we get that on Ξ ,

$$m = -\frac{1}{p} \sum_{i \in \mathcal{I}_0} \frac{1}{z + \lambda_i^2 r_1 m_1 + r_2 m_2} - [\mathcal{Z}] + O_{\prec}\left(n^{-1} + n^{-1/2}\Theta\right), \quad (\text{B.62})$$

$$m_0 = -\frac{1}{p} \sum_{i \in \mathcal{I}_0} \frac{\lambda_i^2}{z + \lambda_i^2 r_1 m_1 + r_2 m_2} - [\mathcal{Z}]_0 + O_{\prec}\left(n^{-1} + n^{-1/2}\Theta\right). \quad (\text{B.63})$$

Comparing these two equations with estimates (B.47) and (B.48), we obtain that

$$|m(z) - m_c(z)| + |m_0(z) - m_{0c}(z)| \lesssim (\log n)^{-1/2}, \quad \text{w.o.p. on } \Xi.$$

Together with equation (B.46), this estimate implies that

$$|1 + \gamma_n m(z)| \sim 1, \quad |1 + \gamma_n m_0(z)| \sim 1, \quad \text{w.o.p. on } \Xi. \quad (\text{B.64})$$

With a similar argument as above, from equations (B.56) and (B.57) we obtain that on Ξ ,

$$G_{\mu\mu} = -\frac{1}{1 + \gamma_n m_0} - \frac{\mathcal{Z}_{\mu}}{(1 + \gamma_n m_0)^2} + O_{\prec}\left(n^{-1} + n^{-\frac{1}{2}}\Theta\right), \quad \mu \in \mathcal{I}_1, \quad (\text{B.65})$$

$$G_{\nu\nu} = -\frac{1}{1 + \gamma_n m} - \frac{\mathcal{Z}_{\nu}}{(1 + \gamma_n m)^2} + O_{\prec}\left(n^{-1} + n^{-\frac{1}{2}}\Theta\right), \quad \nu \in \mathcal{I}_2, \quad (\text{B.66})$$

where we used estimates (B.53), (B.59), (B.60) and (B.64) in the derivation. Taking averages of equations (B.65) and (B.66) over $\mu \in \mathcal{I}_1$ and $\nu \in \mathcal{I}_2$, we get that on Ξ ,

$$m_1 = -\frac{1}{1 + \gamma_n m_0} - \frac{[\mathcal{Z}]_1}{(1 + \gamma_n m_0)^2} + O_{\prec}\left(n^{-1} + n^{-\frac{1}{2}}\Theta\right), \quad (\text{B.67})$$

$$m_2 = -\frac{1}{1 + \gamma_n m} - \frac{[\mathcal{Z}]_2}{(1 + \gamma_n m)^2} + O_{\prec}\left(n^{-1} + n^{-\frac{1}{2}}\Theta\right), \quad (\text{B.68})$$

which further implies that on Ξ ,

$$\frac{1}{m_1} + 1 + \gamma_n m_0 \prec n^{-1} + n^{-\frac{1}{2}}\Theta + |[\mathcal{Z}]_1|, \quad (\text{B.69})$$

$$\frac{1}{m_2} + 1 + \gamma_n m \prec n^{-1} + n^{-\frac{1}{2}}\Theta + |[\mathcal{Z}]_2|. \quad (\text{B.70})$$

Finally, plugging estimates (B.62) and (B.63) into equations (B.69) and (B.70), we get that equations (B.51) and (B.52) are true. \blacksquare

Step 3: Entrywise local law. In this step, we show that the event $\Xi(z)$ in (B.49) actually holds with overwhelming probability for all $z \in \mathbf{D}$. Once we have proved this fact, applying Lemma 29 to equations (B.51) and (B.52) gives that $(m_1(z), m_2(z))$ is close to $(m_{1c}(z), m_{2c}(z))$ up to an error of order $O_{\prec}(n^{-1/2})$, with which we can conclude the entrywise local law (B.31).

First, we claim that it suffices to show

$$|m_1(0) - m_{1c}(0)| + |m_2(0) - m_{2c}(0)| \prec n^{-\frac{1}{2}}. \quad (\text{B.71})$$

In fact, by equation (B.37) we have that uniformly in $z \in \mathbf{D}$,

$$|m_1(z) - m_1(0)| + |m_2(z) - m_2(0)| \lesssim (\log n)^{-1} \quad \text{w.o.p.}$$

Thus, if equation (B.71) holds, using the triangle inequality, we obtain from the above estimate that

$$\sup_{z \in \mathbf{D}} (|m_1(z) - m_{1c}(0)| + |m_2(z) - m_{2c}(0)|) \lesssim (\log n)^{-1} \quad \text{w.o.p.} \quad (\text{B.72})$$

The equation (B.72) shows that the event $\Xi(z)$ holds w.o.p., and it also verifies the condition (B.28) of Lemma 29. Now, applying Lemma 29 to equations (B.51) and (B.52), we obtain

$$\begin{aligned} \Theta(z) &= |m_1(z) - m_{1c}(z)| + |m_2(z) - m_{2c}(z)| \\ &\prec n^{-1} + n^{-\frac{1}{2}}\Theta(z) + |[\mathcal{Z}]| + |[\mathcal{Z}]_0| + |[\mathcal{Z}]_1| + |[\mathcal{Z}]_2|, \end{aligned}$$

which implies that

$$\Theta(z) \prec n^{-1} + |[\mathcal{Z}]| + |[\mathcal{Z}]_0| + |[\mathcal{Z}]_1| + |[\mathcal{Z}]_2| \prec n^{-\frac{1}{2}}, \quad (\text{B.73})$$

uniformly in all $z \in \mathbf{D}$, where we used equation (B.43) in the second step. On the other hand, with equations (B.65) to (B.68), we obtain that

$$\max_{\mu \in \mathcal{I}_1} |G_{\mu\mu}(z) - m_1(z)| + \max_{\nu \in \mathcal{I}_2} |G_{\nu\nu}(z) - m_2(z)| \prec n^{-\frac{1}{2}}.$$

Combining this estimate with equation (B.73), we get that

$$\max_{\mu \in \mathcal{I}_1} |G_{\mu\mu}(z) - m_{1c}(z)| + \max_{\nu \in \mathcal{I}_2} |G_{\nu\nu}(z) - m_{2c}(z)| \prec n^{-\frac{1}{2}}. \quad (\text{B.74})$$

Next, plugging equation (B.73) into equation (B.61) and recalling equations (B.7) and (B.19), we obtain that

$$\max_{i \in \mathcal{I}_1} |G_{ii}(z) - \mathfrak{G}_{ii}(z)| \prec n^{-\frac{1}{2}}.$$

Together with equation (B.74), it gives the diagonal estimate

$$\max_{\mathfrak{a} \in \mathcal{I}} |G_{\mathfrak{a}\mathfrak{a}}(z) - \Pi_{\mathfrak{a}\mathfrak{a}}(z)| \prec n^{-\frac{1}{2}}. \quad (\text{B.75})$$

Combining equation (B.75) with the off-diagonal estimate on Λ_o in equation (B.43), we conclude the entrywise local law (B.31).

It remains to prove equation (B.71). Using equations (B.36) and (B.38), w.o.p.,

$$1 \gtrsim m(0) = \frac{1}{p} \sum_{i \in \mathcal{I}_0} G_{ii}(0) = \frac{1}{p} \sum_{i \in \mathcal{I}_0} \sum_{k=1}^p \frac{|\xi_k(i)|^2}{\mu_k} \geq \mu_1^{-1} \gtrsim 1,$$

where we used Lemma 22 to bound μ_1 . Similarly, we can get that $m_0(0) > 0$ and $m_0(0) \sim 1$. Hence, we have the estimates

$$1 + \gamma_n m(0) \sim 1, \quad 1 + \gamma_n m_0(0) \sim 1. \quad (\text{B.76})$$

Combining these estimates with equations (B.56), (B.57) and (B.60), we obtain that equations (B.67) and (B.68) hold at $z = 0$ without requiring the event $\Xi(0)$ to hold. This further gives that w.o.p.,

$$|\lambda_i^2 r_1 m_1(0) + r_2 m_2(0)| = \left| \frac{\lambda_i^2 r_1}{1 + \gamma_n m_0(0)} + \frac{r_2}{1 + \gamma_n m(0)} + O_{\prec}(n^{-1/2}) \right| \sim 1.$$

Then, combining this estimate with (B.55) and (B.60), we obtain that equations (B.62) and (B.63) also hold at $z = 0$ without requiring the event $\Xi(0)$ to hold. Finally, plugging equations (B.62) and (B.63) into equations (B.69) and (B.70), we conclude that equations (B.51) and (B.52) hold at $z = 0$, that is,

$$\begin{aligned} \left| \frac{1}{m_1(0)} + 1 - \frac{1}{n} \sum_{i=1}^p \frac{\lambda_i^2}{\lambda_i^2 r_1 m_1(0) + r_2 m_2(0)} \right| &< n^{-\frac{1}{2}}, \\ \left| \frac{1}{m_2(0)} + 1 - \frac{1}{n} \sum_{i=1}^p \frac{1}{\lambda_i^2 r_1 m_1(0) + r_2 m_2(0)} \right| &< n^{-\frac{1}{2}}. \end{aligned} \quad (\text{B.77})$$

Denoting $y_1 = -m_1(0)$ and $y_2 = -m_2(0)$, by estimates (B.67) and (B.68) at $z = 0$, we have that

$$y_1 = \frac{1}{1 + \gamma_n m_0(0)} + O_{\prec}(n^{-1/2}), \quad y_2 = \frac{1}{1 + \gamma_n m(0)} + O_{\prec}(n^{-1/2}).$$

Together with equation (B.76), it implies that

$$c \leq y_1 \leq 1, \quad c \leq y_2 \leq 1, \quad \text{w.o.p.}, \quad (\text{B.78})$$

for a small constant $c > 0$. Also with equation (B.77), we can check that $(r_1 y_1, r_2 y_2)$ satisfies approximately the same system of equations as equations (3.2) and (B.16):

$$r_1 y_1 + r_2 y_2 = 1 - \gamma_n + O_{\prec}(n^{-1/2}), \quad f(r_1 y_1) = r_1 + O_{\prec}(n^{-1/2}). \quad (\text{B.79})$$

The first equation of (B.79) and equation (B.78) together imply that $r_1 y_1 \in [0, 1 - \gamma_n]$ with overwhelming probability. We also know that $r_1 y_1 = \alpha_1$ is a solution to the second equation of (B.79). Moreover, we have checked that the function f is strictly increasing and has bounded derivative on $[0, r_1^{-1}(1 - \gamma_n)]$.

With basic calculus, we can get that

$$|m_1(0) - m_{1c}(0)| = |y_1 - r_1^{-1}\alpha_1| \prec n^{-1/2}.$$

Plugging it into the first equation of (B.79), we get

$$|m_2(0) - m_{2c}(0)| = |y_2 - r_2^{-1}\alpha_2| \prec n^{-1/2}.$$

The above two estimates conclude that equation (B.71) is true. This concludes the proof of Lemma B.31. \blacksquare

Step 4: Averaged local law. Finally, we prove the averaged local laws (B.10) and (B.11). For this purpose, we need to use the following *fluctuation averaging* lemma, whose proof is omitted as it is a standard result in the literature.

Lemma 37 (Fluctuation averaging) *Under the setting of Proposition 30, suppose the entrywise local law (B.31) holds uniformly in $z \in \mathbf{D}$. Then, we have that*

$$|[\mathcal{Z}]| + |[\mathcal{Z}]_0| + |[\mathcal{Z}]_1| + |[\mathcal{Z}]_2| \prec (np)^{-1/2}, \quad (\text{B.80})$$

uniformly in $z \in \mathbf{D}$.

Proof The proof can be found in the works of Erdős et al. (2013). \blacksquare

Now, plugging equations (B.73) and (B.80) into equations (B.51) and (B.52) and applying Lemma 29, we get

$$|m_1(z) - m_{1c}(z)| + |m_2(z) - m_{2c}(z)| \prec (np)^{-1/2}. \quad (\text{B.81})$$

Then, subtracting equation (B.47) from equation (B.62) and using equations (B.80) and (B.81), we obtain that

$$|m(z) - m_c(z)| \prec (np)^{-1/2}.$$

This is the averaged local law (B.10) with $Q = 1$. The proof of law (B.11) is similar. \blacksquare

B.3.4 ANISOTROPIC LOCAL LAW

Proof In this section, we prove the anisotropic local law in Theorem 27 by extending from Gaussian random matrices to generally distributed random matrices. With Proposition 30, it suffices to prove that for $Z^{(1)}$ and $Z^{(2)}$ satisfying the assumptions of Theorem 27, we have

$$\left| \mathbf{u}^\top (G(Z, z) - G(Z^{\text{Gauss}}, z)) \mathbf{v} \right| \prec n^{-1/2} Q,$$

for any deterministic unit vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^{p+n}$ and $z \in \mathbf{D}$, where we abbreviate that

$$Z := \begin{pmatrix} Z^{(1)} \\ Z^{(2)} \end{pmatrix} \quad \text{and} \quad Z^{\text{Gauss}} := \begin{pmatrix} (Z^{(1)})^{\text{Gauss}} \\ (Z^{(2)})^{\text{Gauss}} \end{pmatrix}.$$

Here, $(Z^{(1)})^{\text{Gauss}}$ and $(Z^{(2)})^{\text{Gauss}}$ are Gaussian random matrices satisfying the assumptions of Proposition 30. We will prove the above statement using a continuous comparison argument developed in Knowles and Yin (2016). Since the arguments are similar to the ones in Sections 7 and 8 of Knowles and Yin (2016), we will not write down all the details.

We define a continuous sequence of interpolating matrices between Z^{Gauss} and Z .

Definition 38 (Interpolation) We denote $Z^0 := Z^{\text{Gauss}}$ and $Z^1 := Z$. Let $\rho_{\mu i}^0$ and $\rho_{\mu i}^1$ be respectively the laws of $Z_{\mu i}^0$ and $Z_{\mu i}^1$ for $i \in \mathcal{I}_0$ and $\mu \in \mathcal{I}_1 \cup \mathcal{I}_2$. For any $\theta \in [0, 1]$, we define the interpolated law $\rho_{\mu i}^\theta := (1 - \theta)\rho_{\mu i}^0 + \theta\rho_{\mu i}^1$. We work on the probability space consisting of triples (Z^0, Z^θ, Z^1) of independent $n \times p$ random matrices, where the matrix $Z^\theta = (Z_{\mu i}^\theta)$ has law

$$\prod_{i \in \mathcal{I}_0} \prod_{\mu \in \mathcal{I}_1 \cup \mathcal{I}_2} \rho_{\mu i}^\theta(dZ_{\mu i}^\theta). \quad (\text{B.82})$$

For $\lambda \in \mathbb{R}$, $i \in \mathcal{I}_0$ and $\mu \in \mathcal{I}_1 \cup \mathcal{I}_2$, we define the matrix $Z_{(\mu i)}^{\theta, \lambda}$ through

$$\left(Z_{(\mu i)}^{\theta, \lambda} \right)_{\nu j} := \begin{cases} Z_{\mu i}^\theta, & \text{if } (j, \nu) \neq (i, \mu) \\ \lambda, & \text{if } (j, \nu) = (i, \mu) \end{cases},$$

that is, it is obtained by replacing the (μ, i) -th entry of Z^θ with λ . We abbreviate that

$$G^\theta(z) := G(Z^\theta, z), \quad G_{(\mu i)}^{\theta, \lambda}(z) := G(Z_{(\mu i)}^{\theta, \lambda}, z).$$

We will prove the anisotropic local law (B.12) through interpolating matrices Z^θ between Z^0 and Z^1 . We have seen that equation (B.12) holds for $G(Z^0, z)$ by Proposition 30. Using equation (B.82) and basic calculus, we readily get the following interpolation formula: for any differentiable function $F : \mathbb{R}^{n \times p} \rightarrow \mathbb{C}$,

$$\frac{d}{d\theta} \mathbb{E} F(Z^\theta) = \sum_{i \in \mathcal{I}_0} \sum_{\mu \in \mathcal{I}_1 \cup \mathcal{I}_2} \left[\mathbb{E} F \left(Z_{(\mu i)}^{\theta, Z_{\mu i}^1} \right) - \mathbb{E} F \left(Z_{(\mu i)}^{\theta, Z_{\mu i}^0} \right) \right], \quad (\text{B.83})$$

provided all expectations exist. We will apply equation (B.83) to the function $F(Z) := F_{\mathbf{uv}}^s(Z, z)$ for any fixed $s \in 2\mathbb{N}$, where

$$F_{\mathbf{uv}}(Z, z) := \left| \mathbf{u}^\top [G(Z, z) - \mathfrak{G}(z)] \mathbf{v} \right|. \quad (\text{B.84})$$

The main part of the proof is to show the following self-consistent estimate for the right-hand side of equation (B.83): for any fixed $s \in 2\mathbb{N}$, constant $c > 0$, and $\theta \in [0, 1]$,

$$\left| \sum_{i \in \mathcal{I}_0} \sum_{\mu \in \mathcal{I}_1 \cup \mathcal{I}_2} \left[\mathbb{E} F_{\mathbf{uv}}^s \left(Z_{(\mu i)}^{\theta, Z_{\mu i}^1}, z \right) - \mathbb{E} F_{\mathbf{uv}}^s \left(Z_{(\mu i)}^{\theta, Z_{\mu i}^0}, z \right) \right] \right| \lesssim (n^c q)^s + \mathbb{E} F_{\mathbf{uv}}^s(Z^\theta, z), \quad (\text{B.85})$$

where we abbreviate that $q := n^{-1/2}Q$. If equation (B.85) holds, then combining equation (B.83) with Grönwall's inequality, we obtain that for any fixed $s \in 2\mathbb{N}$ and constant $c > 0$,

$$\mathbb{E} \left| \mathbf{u}^\top [G(Z^1, z) - \Pi(z)] \mathbf{v} \right|^s \lesssim (n^c q)^s. \quad (\text{B.86})$$

Then, applying Markov's inequality and noticing that c can be arbitrarily small, we conclude that equation (B.12) is true.

In order to prove equation (B.85), we compare $Z_{(\mu i)}^{\theta, Z_{\mu i}^{\alpha}}$ and $Z_{(\mu i)}^{\theta, Z_{\mu i}^1}$ via a common $Z_{(\mu i)}^{\theta, 0}$, that is, we will prove that for any constant $c > 0$, $\alpha \in \{0, 1\}$, and $\theta \in [0, 1]$,

$$\left| \sum_{i \in \mathcal{I}_0} \sum_{\mu \in \mathcal{I}_1 \cup \mathcal{I}_2} \left[\mathbb{E} F_{\mathbf{uv}}^s \left(Z_{(\mu i)}^{\theta, Z_{\mu i}^{\alpha}}, z \right) - \mathbb{E} F_{\mathbf{uv}}^s \left(Z_{(\mu i)}^{\theta, 0}, z \right) \right] - \mathcal{A} \right| \lesssim (n^c q)^s + \mathbb{E} F_{\mathbf{uv}}^s(Z^{\theta}, z), \quad (\text{B.87})$$

for a quantity \mathcal{A} that does not depend on $\alpha \in \{0, 1\}$. Underlying the proof of equation (B.87) is an expansion approach which we describe now. We define the $\mathcal{I} \times \mathcal{I}$ matrix $\Delta_{(\mu i)}^{\lambda}$ as

$$\Delta_{(\mu i)}^{\lambda} := \lambda \begin{pmatrix} 0 & \mathbf{u}_i^{(\mu)} \mathbf{e}_{\mu}^{\top} \\ \mathbf{e}_{\mu} \mathbf{u}_i^{(\mu)\top} & 0 \end{pmatrix},$$

where $\mathbf{u}_i^{(\mu)} := \Lambda U^{\top} \mathbf{e}_i$ if $\mu \in \mathcal{I}_1$, $\mathbf{u}_i^{(\mu)} := V^{\top} \mathbf{e}_i$ if $\mu \in \mathcal{I}_2$, and \mathbf{e}_i and \mathbf{e}_{μ} denote the standard basis vectors along the i -th and μ -th directions. Then, by the definition of H in equation (B.4), applying a Taylor expansion, we get that for any $\lambda, \lambda' \in \mathbb{R}$ and $K \in \mathbb{N}$,

$$G_{(\mu i)}^{\theta, \lambda'} = G_{(\mu i)}^{\theta, \lambda} + n^{-\frac{k}{2}} \sum_{k=1}^K G_{(\mu i)}^{\theta, \lambda} \left(\Delta_{(\mu i)}^{\lambda - \lambda'} G_{(\mu i)}^{\theta, \lambda} \right)^k + n^{-\frac{K+1}{2}} G_{(\mu i)}^{\theta, \lambda'} \left(\Delta_{(\mu i)}^{\lambda - \lambda'} G_{(\mu i)}^{\theta, \lambda} \right)^{K+1}. \quad (\text{B.88})$$

Using this expansion and the bound in equation (B.36), it is easy to prove the following estimate: if y is a random variable satisfying $|y| \prec Q$ (note that the entries of any interpolating matrix Z^{θ} satisfy this bound), then w.o.p.,

$$\max_{i \in \mathcal{I}_0} \max_{\mu \in \mathcal{I}_1 \cup \mathcal{I}_2} G_{(\mu i)}^{\theta, y} = O(1). \quad (\text{B.89})$$

For simplicity of notations, in the following proof, we denote

$$f_{(\mu i)}(\lambda) := F_{\mathbf{uv}}^s \left(Z_{(\mu i)}^{\theta, \lambda} \right) = \left| \mathbf{u}^{\top} \left(G \left(Z_{(\mu i)}^{\theta, \lambda}, z \right) - \mathfrak{G}(z) \right) \mathbf{v} \right|^s.$$

We use $f_{(\mu i)}^{(r)}$ to denote the r -th order derivative of $f_{(\mu i)}$. By equation (B.89), it is easy to see that for any fixed $r \in \mathbb{N}$, $f_{(\mu i)}^{(r)}(y) = O(1)$ w.o.p. for any random variable y satisfying $|y| \prec Q$. Then, the Taylor expansion of $f_{(\mu i)}$ gives

$$f_{(\mu i)}(y) = \frac{1}{n^{r/2}} \sum_{r=0}^{s+4} \frac{y^r}{r!} f_{(\mu i)}^{(r)}(0) + O_{\prec}(q^{s+4}).$$

Therefore, we have that for $\alpha \in \{0, 1\}$,

$$\begin{aligned} & \mathbb{E} F_{\mathbf{uv}}^s \left(Z_{(\mu i)}^{\theta, Z_{\mu i}^{\alpha}} \right) - \mathbb{E} F_{\mathbf{uv}}^s \left(Z_{(\mu i)}^{\theta, 0} \right) = \mathbb{E} \left[f_{(\mu i)} \left(Z_{(\mu i)}^{\alpha} \right) - f_{(\mu i)}(0) \right] \\ & = \mathbb{E} f_{(\mu i)}(0) + \frac{1}{2n} \mathbb{E} f_{(\mu i)}^{(2)}(0) + \sum_{r=3}^{s+4} \frac{n^{-r/2}}{r!} \mathbb{E} f_{(\mu i)}^{(r)}(0) \cdot \mathbb{E} \left(Z_{(\mu i)}^{\alpha} \right)^r + O_{\prec}(q^{s+4}). \end{aligned} \quad (\text{B.90})$$

To illustrate the idea in a more concise way, we assume the extra condition

$$\mathbb{E}(Z_{\mu i}^1)^3 = 0, \quad 1 \leq \mu \leq n, \quad 1 \leq i \leq p. \quad (\text{B.91})$$

Hence, the $r = 3$ term in the Taylor expansion (B.90) vanishes. However, this condition is not necessary as we will explain at the end of the proof.

Recall that the entries of $Z_{\mu i}^1$ have finite fourth moments by equation (2.2). Together with the bounded support condition, it gives that

$$|\mathbb{E}(Z_{\mu i}^a)^r| \prec Q^{r-4}, \quad r \geq 4.$$

We take

$$\mathcal{A} = \mathbb{E} f_{(\mu i)}(0) + \frac{1}{2n} \mathbb{E} f_{(\mu i)}^{(2)}(0).$$

Thus, to show equation (B.87) under equation (B.91), we only need to prove that for $r = 4, \dots, s+4$,

$$n^{-2} q^{r-4} \sum_{i \in \mathcal{I}_0} \sum_{\mu \in \mathcal{I}_1 \cup \mathcal{I}_2} \left| \mathbb{E} f_{(\mu i)}^{(r)}(0) \right| \lesssim (n^c q)^s + \mathbb{E} F_{\mathbf{uv}}^s(Z^\theta, z). \quad (\text{B.92})$$

In order to get a self-consistent estimate in terms of the matrix Z^θ only on the right-hand side of equation (B.92), we want to replace $Z_{(\mu i)}^{\theta, 0}$ in $f_{(\mu i)}(0) = F_{\mathbf{uv}}^s(Z_{(\mu i)}^{\theta, 0})$ with $Z^\theta \equiv Z_{(\mu i)}^{\theta, Z_{\mu i}^\theta}$.

Lemma 39 *Suppose that*

$$n^{-2} q^{r-4} \sum_{i \in \mathcal{I}_0} \sum_{\mu \in \mathcal{I}_1 \cup \mathcal{I}_2} \left| \mathbb{E} f_{(\mu i)}^{(r)}(Z_{\mu i}^\theta) \right| \lesssim (n^c q)^s + \mathbb{E} F_{\mathbf{uv}}^s(Z^\theta) \quad (\text{B.93})$$

holds for $r = 4, \dots, s+4$. Then, equation (B.92) holds for $r = 4, \dots, s+4$.

Proof The proof is the same as that for Lemma 7.16 of Knowles and Yin (2016). \blacksquare

What remains now is the proof of equation (B.93). For simplicity of notations, we will abbreviate $Z^\theta \equiv Z$ in the following proof. For any $k \in \mathbb{N}$, we denote

$$A_{\mu i}(k) := \left(\frac{\partial}{\partial Z_{\mu i}} \right)^k \mathbf{u}^\top (G - \mathfrak{G}) \mathbf{v}. \quad (\text{B.94})$$

The derivative on the right-hand side can be calculated using the expansion (B.88). In particular, it is easy to check the following bound

$$|A_{\mu i}(k)| \prec \begin{cases} (\mathcal{R}_i^{(\mu)})^2 + \mathcal{R}_\mu^2, & \text{if } k \geq 2, \\ \mathcal{R}_i^{(\mu)} \mathcal{R}_\mu, & \text{if } k = 1, \end{cases} \quad (\text{B.95})$$

where for $i \in \mathcal{I}_1$ and $\mu \in \mathcal{I}_1 \cup \mathcal{I}_2$, we denote

$$\mathcal{R}_i^{(\mu)} := |\mathbf{u}^\top G \mathbf{u}_i^{(\mu)}| + |\mathbf{v}^\top G \mathbf{u}_i^{(\mu)}|, \quad \mathcal{R}_\mu := |\mathbf{u}^\top G \mathbf{e}_\mu| + |\mathbf{v}^\top G \mathbf{e}_\mu|. \quad (\text{B.96})$$

Now, we can calculate the derivative

$$f_{(\mu i)}^{(r)}(Z_{\mu i}) = \left(\frac{\partial}{\partial Z_{\mu i}} \right)^r F_{\mathbf{uv}}^s(Z) = \sum_{k_1 + \dots + k_s = r} \prod_{t=1}^{s/2} \left(A_{\mu i}(k_t) \overline{A_{\mu i}(k_{t+s/2})} \right).$$

Then, to prove equation (B.93), it suffices to show that

$$n^{-2} q^{r-4} \sum_{i \in \mathcal{I}_0} \sum_{\mu \in \mathcal{I}_1 \cup \mathcal{I}_2} \left| \mathbb{E} \prod_{t=1}^{s/2} A_{\mu i}(k_t) \overline{A_{\mu i}(k_{t+s/2})} \right| \lesssim (n^c q)^s + \mathbb{E} F_{\mathbf{uv}}^s(Z, z), \quad (\text{B.97})$$

for $r = 4, \dots, s+4$ and any $(k_1, \dots, k_s) \in \mathbb{N}^s$ satisfying $k_1 + \dots + k_s = r$. Treating zero k_t 's separately (note that $A_{\mu i}(0) = G_{\mathbf{uv}} - \mathfrak{G}_{\mathbf{uv}}$ by definition), we find that it suffices to prove

$$n^{-2} q^{r-4} \sum_{i \in \mathcal{I}_0} \sum_{\mu \in \mathcal{I}_1 \cup \mathcal{I}_2} \mathbb{E} |A_{\mu i}(0)|^{s-l} \prod_{t=1}^l |A_{\mu i}(k_t)| \lesssim (n^c q)^s + \mathbb{E} F_{\mathbf{uv}}^s(Z, z), \quad (\text{B.98})$$

for $r = 4, \dots, s+4$ and $l = 1, \dots, s$. Here, without loss of generality, we have assumed that $k_t = 0$ for $l+1 \leq t \leq s$, $k_t \geq 1$ for $1 \leq t \leq l$, and $\sum_{t=1}^l k_t = r$.

There is at least one non-zero k_t in all cases, while in the case with $r \leq 2l-2$, there exist at least two k_t 's with $k_t = 1$ by the pigeonhole principle. Therefore, with equation (B.95), we get that

$$\prod_{t=1}^l |A_{\mu i}(k_t)| \prec \mathbf{1}(r \geq 2l-1) \left[(\mathcal{R}_i^{(\mu)})^2 + \mathcal{R}_\mu^2 \right] + \mathbf{1}(r \leq 2l-2) (\mathcal{R}_i^{(\mu)})^2 \mathcal{R}_\mu^2. \quad (\text{B.99})$$

Using equation (B.36) and a similar argument as in equation (B.44), we can obtain that

$$\sum_{i \in \mathcal{I}_0} (\mathcal{R}_i^{(\mu)})^2 = O(1), \quad \sum_{\mu \in \mathcal{I}_1 \cup \mathcal{I}_2} \mathcal{R}_\mu^2 = O(1), \quad (\text{B.100})$$

with overwhelming probability. Using equations (B.99) and (B.100), and $n^{-1/2} \leq n^{-1/2} Q = q$, we get that

$$\begin{aligned} & n^{-2} q^{r-4} \sum_{i \in \mathcal{I}_0} \sum_{\mu \in \mathcal{I}_1 \cup \mathcal{I}_2} |A_{\mu i}(0)|^{s-l} \prod_{t=1}^l |A_{\mu i}(k_t)| \\ & \prec q^{r-4} F_{\mathbf{uv}}^{s-l}(Z) \left[\mathbf{1}(r \geq 2l-1) n^{-1} + \mathbf{1}(r \leq 2l-2) n^{-2} \right] \\ & \leq F_{\mathbf{uv}}^{s-l}(Z) \left[\mathbf{1}(r \geq 2l-1) q^{r-2} + \mathbf{1}(r \leq 2l-2) q^r \right]. \end{aligned} \quad (\text{B.101})$$

If $r \leq 2l-2$, then we have $q^r \leq q^l$ by the trivial inequality $r \geq l$. On the other hand, if $r \geq 4$ and $r \geq 2l-1$, then $r \geq l+2$ and we have $q^{r-2} \leq q^l$. Thus, with equation (B.101), we conclude that

$$\begin{aligned} n^{-2} q^{r-4} \sum_{i \in \mathcal{I}_0} \sum_{\mu \in \mathcal{I}_1 \cup \mathcal{I}_2} \mathbb{E} |A_{\mu i}(0)|^{s-l} \prod_{t=1}^l |A_{\mu i}(k_t)| & \prec \mathbb{E} F_{\mathbf{uv}}^{s-l}(Z) q^l \leq [\mathbb{E} F_{\mathbf{uv}}^s(Z)]^{\frac{s-l}{s}} q^l \\ & \lesssim \mathbb{E} F_{\mathbf{uv}}^s(Z) + q^s, \end{aligned}$$

where we used Hölder's inequality in the second step and Young's inequality in the last step. This gives equation (B.98), which concludes the proof of equation (B.93), and hence of equation (B.87), and hence of equation (B.85), which concludes equation (B.86) and completes the proof of the anisotropic local law of equation (B.12) under the condition of equation (B.91).

Finally, if the condition (B.91) does not hold, then there is also an $r = 3$ term in the Taylor expansion (B.90):

$$\frac{1}{6n^{3/2}} \mathbb{E} f_{(\mu i)}^{(3)}(0) \cdot \mathbb{E} (Z_{i\mu}^\alpha)^3.$$

But the sum over i and μ in equation (B.87) gives a factor n^2 , which cannot be canceled by the $n^{-3/2}$ factor in the above equation. In fact, $\mathbb{E} f_{(\mu i)}^{(3)}(0)$ will provide an extra $n^{-1/2}$ factor to compensate the remaining $n^{1/2}$ factor. This follows from an improved self-consistent comparison argument for sample covariance matrices in Section 8 of Knowles and Yin (2016). The argument for our setting is almost the same except for some notational differences, so we omit the details. This completes the proof of equation (B.12) without the condition (B.91). ■

B.3.5 AVERAGED LOCAL LAW

Proof Finally, in this subsection, we focus on the proof of the averaged local law (B.10) in the setting of Theorem 27, while the proof of the law (B.11) is exactly the same. Our proof is similar to that for the anisotropic local law (B.12) in the previous subsection, and we only explain the main differences. In analogy to equation (B.84)), we define

$$\tilde{F}(Z, z) := \left| \frac{1}{p} \sum_{i \in \mathcal{I}_0} [G_{ii}(Z, z) - \mathfrak{G}_{ii}(z)] \right|.$$

Under the notations of Definition 38, we have proved that $\tilde{F}(Z^0, z) \prec (np)^{-1/2}$ in Lemma 31. To illustrate the idea, we again assume that the condition (B.91) holds in the following proof. Using the arguments in Section B.3.4, analogous to equation (B.93) we only need to prove that for $Z = Z^\theta$, $q = n^{-1/2}Q$, any small constant $c > 0$, and fixed $s \in 2\mathbb{N}$,

$$n^{-2}q^{r-4} \sum_{i \in \mathcal{I}_0} \sum_{\mu \in \mathcal{I}_1 \cup \mathcal{I}_2} \left| \mathbb{E} \left(\frac{\partial}{\partial Z_{\mu i}} \right)^r \tilde{F}^s(Z) \right| \lesssim (p^{-1/2+c}q)^s + \mathbb{E} \tilde{F}^s(Z), \quad (\text{B.102})$$

for $r = 4, \dots, K$, where $K \in \mathbb{N}$ is a large enough constant. Similar to equation (B.94), let

$$A_{j,\mu i}(k) := \left(\frac{\partial}{\partial Z_{\mu i}} \right)^k (G_{jj} - \mathfrak{G}_{jj}).$$

Analogous to estimate (B.97), it suffices to prove that

$$\begin{aligned} & n^{-2}q^{r-4} \sum_{i \in \mathcal{I}_0} \sum_{\mu \in \mathcal{I}_1 \cup \mathcal{I}_2} \left| \mathbb{E} \prod_{t=1}^{s/2} \left(\frac{1}{p} \sum_{j \in \mathcal{I}_0} A_{j,\mu i}(k_t) \right) \left(\frac{1}{p} \sum_{j \in \mathcal{I}_0} \overline{A_{j,\mu i}(k_{t+s/2})} \right) \right| \\ & \lesssim (p^{-1/2+c}q)^s + \mathbb{E} \tilde{F}^s(Z), \end{aligned}$$

for $r = 4, \dots, K$ and any $(k_1, \dots, k_s) \in \mathbb{N}^s$ satisfying $k_1 + \dots + k_s = r$. Without loss of generality, we assume that $k_t = 0$ for $l+1 \leq t \leq s$, $k_t \geq 1$ for $1 \leq t \leq l$, and $\sum_{t=1}^l k_t = r$. Then, it suffices to prove that

$$n^{-2}q^{r-4} \sum_{i \in \mathcal{I}_0} \sum_{\mu \in \mathcal{I}_1 \cup \mathcal{I}_2} \mathbb{E} \tilde{F}^{s-l}(Z) \prod_{t=1}^l \left| \frac{1}{p} \sum_{j \in \mathcal{I}_0} A_{j,\mu i}(k_t) \right| \lesssim (p^{-1/2+c}q)^s + \mathbb{E} \tilde{F}^s(Z),$$

for $r = 4, \dots, K$ and any $1 \leq l \leq s$.

Using equation (B.36) and a similar argument as in equation (B.44), we obtain that for $1 \leq t \leq l$,

$$\left| \sum_{j \in \mathcal{I}_0} A_{j,\mu i}(k_t) \right| \lesssim 1 \quad (\text{B.103})$$

with overwhelming probability. Moreover, taking $\mathbf{u} = \mathbf{v} = \mathbf{e}_j$ in equation (B.96), we define

$$\mathcal{R}_{j,i}^{(\mu)} := |\mathbf{e}_j^\top G \mathbf{u}_i^{(\mu)}|, \quad \mathcal{R}_{j,\mu} := |\mathbf{e}_j^\top G \mathbf{e}_\mu|.$$

Similar to equation (B.100), we have that

$$\sum_{i \in \mathcal{I}_0} (\mathcal{R}_{j,i}^{(\mu)})^2 = O(1), \quad \sum_{\mu \in \mathcal{I}_1 \cup \mathcal{I}_2} \mathcal{R}_{j,\mu}^2 = O(1), \quad \text{w.o.p.} \quad (\text{B.104})$$

Using equation (B.103) and applying equation (B.95) to $A_{j,\mu i}(k_t)$, we obtain that

$$\prod_{t=1}^l \left| \frac{1}{p} \sum_{j \in \mathcal{I}_0} A_{j,\mu i}(k_t) \right| \prec \mathbf{1}(r \geq 2l-1)p^{-l} + \mathbf{1}(r \leq 2l-2)p^{-l} \sum_{j_1, j_2 \in \mathcal{I}_0} \mathcal{R}_{j_1,i}^{(\mu)} \mathcal{R}_{j_2,i}^{(\mu)} \mathcal{R}_{j_1,\mu} \mathcal{R}_{j_2,\mu},$$

where we use an argument that is similar to the one for estimate (B.99). Summing this equation over $i \in \mathcal{I}_0$, $\mu \in \mathcal{I}_1 \cup \mathcal{I}_2$ and using estimate (B.104), we get that

$$\begin{aligned} & n^{-2}q^{r-4} \sum_{i \in \mathcal{I}_0} \sum_{\mu \in \mathcal{I}_1 \cup \mathcal{I}_2} \mathbb{E} \tilde{F}^{s-l}(Z) \prod_{t=1}^l \left| \frac{1}{p} \sum_{j \in \mathcal{I}_0} A_{j,\mu i}(k_t) \right| \\ & \prec q^{r-4} \mathbb{E} \tilde{F}^{s-l}(Z) \left[\mathbf{1}(r \geq 2l-1)p^{-(l-1)}n^{-1} + \mathbf{1}(r \leq 2l-2)n^{-2}p^{-(l-2)} \right]. \end{aligned}$$

We consider the following cases (recall that $r \geq 4$ and $r \geq l$).

- If $r \geq 2l-1$ and $l \geq 2$, then we have $r \geq l+2$ and $l-1 \geq l/2$, which gives

$$p^{-(l-1)}q^{r-4}n^{-1} \leq p^{-l/2}q^{r-2} \leq (p^{-1/2}q)^l.$$

- If $r \geq 2l-1$ and $l = 1$, then we have

$$p^{-(l-1)}q^{r-4}n^{-1} \leq n^{-1} \leq (p^{-1/2}q)^l.$$

- If $r \leq 2l-2$ and $l \geq 4$, then we have $r \geq l$ and $l-2 \geq l/2$, which gives

$$p^{-(l-2)}q^{r-4}n^{-2} \leq p^{-l/2}q^r \leq (p^{-1/2}q)^l.$$

- If $r \leq 2l - 2$ and $l < 4$, then we have $r = 4$ and $l = 3$, which gives

$$p^{-(l-2)} q^{r-4} n^{-2} \leq p^{-1} n^{-2} \leq (p^{-1/2} q)^l.$$

Combining the above cases, we get that

$$n^{-2} q^{r-4} \sum_{i \in \mathcal{I}_0} \sum_{\mu \in \mathcal{I}_1 \cup \mathcal{I}_2} \mathbb{E} \tilde{F}^{s-l}(Z) \prod_{t=1}^l \left| \frac{1}{p} \sum_{j \in \mathcal{I}_0} A_{j, \mu i}(k_t) \right| \prec \mathbb{E} \tilde{F}^{s-l}(X) (p^{-1/2} q)^l.$$

Applying Holder's inequality and Young's inequality, we can conclude that equation (B.102) holds, which completes the proof of the averaged local law (B.10) under the condition (B.91).

Finally, even if the condition (B.91) does not hold, using the self-consistent comparison argument in Section 9 of Knowles and Yin (2016), we can still prove equation (B.10). Since (almost) the same argument also works in our setting, we omit the details. \blacksquare

Appendix C. Proofs for the Model Shift Setting

In this section, we present the proofs of Theorems 6, 9, and 10, and Proposition 8. At the end of this section, we will also derive an approximate estimate of the bias under arbitrary covariate and model shifts.

C.1 Proof of Theorem 6

The variance estimate (4.1) can be derived from Theorem 3 or equation (3.3). Hence, we focus on the proof of the bias estimate (4.2). In the setting of Theorem 6, we can write that

$$\begin{aligned} L_{\text{bias}} &= \mathbf{v}^\top Z^{(1)\top} Z^{(1)} \left(Z^{(1)\top} Z^{(1)} + Z^{(2)\top} Z^{(2)} \right)^{-2} Z^{(1)\top} Z^{(1)} \mathbf{v} \\ &= \mathbf{v}^\top \mathcal{Q}^{(1)} \left(\mathcal{Q}^{(1)} + \frac{n_2}{n_1} \mathcal{Q}^{(2)} \right)^{-2} \mathcal{Q}^{(1)} \mathbf{v}, \end{aligned}$$

where we use the simplified notations

$$\mathbf{v} := \Sigma^{(1)1/2} (\beta^{(1)} - \beta^{(2)}), \quad \mathcal{Q}^{(1)} := \frac{1}{n_1} Z^{(1)\top} Z^{(1)}, \quad \mathcal{Q}^{(2)} := \frac{1}{n_2} Z^{(2)\top} Z^{(2)}.$$

Note that $\mathcal{Q}^{(1)}$ and $\mathcal{Q}^{(2)}$ are both Wishart matrices. Using the rotational invariance of the laws of $\mathcal{Q}^{(1)}$ and $\mathcal{Q}^{(2)}$, we can simplify L_{bias} as follows.

Lemma 40 *In the setting of Theorem 6, we have that*

$$L_{\text{bias}} = \left[1 + O_{\prec}(p^{-1/2}) \right] \frac{\|\mathbf{v}\|^2}{p} \text{Tr} \left[\left(\mathcal{Q}^{(1)} + \frac{n_2}{n_1} \mathcal{Q}^{(2)} \right)^{-2} (\mathcal{Q}^{(1)})^2 \right]. \quad (\text{C.1})$$

Proof It is easy to see that the law of the matrix $\mathcal{A} := \mathcal{Q}^{(1)}(\mathcal{Q}^{(1)} + \frac{n_2}{n_1}\mathcal{Q}^{(2)})^{-2}\mathcal{Q}^{(1)}$ is rotationally invariant. Thus, we have that

$$\mathbf{v}^\top \mathcal{A} \mathbf{v} \stackrel{d}{=} \|\mathbf{v}\|^2 \frac{\mathbf{g}^\top}{\|\mathbf{g}\|} \mathcal{A} \frac{\mathbf{g}}{\|\mathbf{g}\|}, \quad (\text{C.2})$$

where “ $\stackrel{d}{=}$ ” means “equal in distribution”, and $\mathbf{g} = (g_1, \dots, g_p)$ is a random vector that is independent of \mathcal{A} and has i.i.d. Gaussian entries of mean zero and variance one. By Lemma 22, we have that w.o.p.,

$$\|\mathcal{A}\|_F \leq p^{1/2} \|\mathcal{A}\| \lesssim p^{-1/2} \text{Tr}[\mathcal{A}].$$

Using equation (A.21), we get that

$$\|\mathbf{g}\| = p + O_{\prec}(p^{1/2}), \quad |\mathbf{g}^\top \mathcal{A} \mathbf{g} - \text{Tr} \mathcal{A}| \prec \|\mathcal{A}\|_F \prec p^{-1/2} \text{Tr} \mathcal{A}. \quad (\text{C.3})$$

Plugging equation (C.3) into equation (C.2), we conclude that equation (C.1) is true. \blacksquare

Now, based on estimate (C.1), we can write that

$$L_{\text{bias}} = - \left[1 + O_{\prec}(p^{-1/2}) \right] \|\mathbf{v}\|^2 \cdot \frac{dh_\alpha(t)}{dt} \Big|_{t=0}, \quad \text{where } \alpha := n_2/n_1 \text{ and} \quad (\text{C.4})$$

$$h_\alpha(t) := \frac{1}{p} \text{Tr} \left[\frac{1}{\mathcal{Q}^{(1)} + t(\mathcal{Q}^{(1)})^2 + \alpha \mathcal{Q}^{(2)}} \right].$$

Hence, to obtain the asymptotic limit of L_{bias} , we need to calculate the values of $h_\alpha(t)$ for t around 0. Without loss of generality, in the following proof, we mainly focus on the case that $0 < \alpha \leq 1$. That is, the ratio between n_2 and n_1 does not grow with p . Later, we will explain how to extend the proof to the $\alpha > 1$ case.

We calculate $h_\alpha(t)$ using the Stieltjes transform method in random matrix theory and free additive convolution (or free addition) in free probability theory. We briefly describe the basic concepts that are needed for the proof and refer the interested readers to classical texts such as Bai and Silverstein (2010) for a more thorough introduction. The Stieltjes transform of a probability measure μ supported on \mathbb{R} is a complex function defined as

$$m_\mu(z) := \int_0^\infty \frac{d\mu(x)}{x - z}, \quad \text{for } z \in \mathbb{C} \setminus \text{supp}(\mu). \quad (\text{C.5})$$

Given any $p \times p$ symmetric matrix M , let $\mu_M := p^{-1} \sum_i \delta_{\lambda_i(M)}$ denote the empirical spectral distribution (ESD) of M , where $\lambda_i(M)$ denotes the i -th eigenvalue of M and $\delta_{\lambda_i(M)}$ is the point mass measure at $\lambda_i(M)$. Then, it is easy to see that the Stieltjes transform of μ_M is

$$m_{\mu_M}(z) := \frac{1}{p} \sum_{i=1}^p \frac{1}{\lambda_i(M) - z} = \frac{1}{p} \text{Tr} [(M - z \text{Id}_{p \times p})^{-1}].$$

Given two $p \times p$ matrices A_p and B_p , suppose their ESDs μ_{A_p} and μ_{B_p} converge weakly to probability measures μ_A and μ_B , respectively. Let U_p be a sequence of $p \times p$ Haar distributed orthogonal matrices. Then, it is known in free probability theory that the ESD of $A_p + U_p B_p U_p^\top$ converges to the *free addition* of μ_A and μ_B , denoted by $\mu_A \boxplus \mu_B$.

It is well-known that the ESDs of Wishart matrices $\mathcal{Q}^{(1)}$ and $\mathcal{Q}^{(2)}$ converge weakly to the famous Marchenko-Pastur (MP) law (Marčenko and Pastur, 1967): $\mu_{\mathcal{Q}^{(1)}} \Rightarrow \mu^{(1)}$ and $\mu_{\mathcal{Q}^{(2)}} \Rightarrow \mu^{(2)}$, where $\mu^{(1)}$ and $\mu^{(2)}$ have densities

$$\rho^{(i)}(x) = \frac{1}{2\pi\xi_i x} \sqrt{(\lambda_+^{(i)} - x)(x - \lambda_-^{(i)})} \mathbf{1}_{x \in [\lambda_-^{(i)}, \lambda_+^{(i)}]}, \quad i = 1, 2,$$

where we denote $\xi_i = p/n_i$ and the spectrum edges are given by $\lambda_{\pm}^{(i)} := (1 \pm \sqrt{\xi_i})^2$. Moreover, the Stieltjes transforms of $\mu^{(1)}$ and $\mu^{(2)}$ satisfy the self-consistent equations

$$z\xi_i m_{\mu^{(i)}}^2 - (1 - \xi_i - z)m_{\mu^{(i)}} + 1 = 0, \quad i = 1, 2.$$

With this equation, we can check that $g_i(m_{\mu^{(i)}}(z)) = z$, where the function g_i is defined by

$$g_i(m) = \frac{1}{1 + \xi_i m} - \frac{1}{m}, \quad i = 1, 2. \quad (\text{C.6})$$

The sharp convergence rates of $\mu_{\mathcal{Q}^{(1)}}$ and $\mu_{\mathcal{Q}^{(2)}}$ have also been obtained in Theorem 3.3 of Pillai and Yin (2014), that is,

$$d_K(\mu_{\mathcal{Q}^{(i)}}, \mu^{(i)}) \prec p^{-1}, \quad i = 1, 2, \quad (\text{C.7})$$

where d_K denotes the Kolmogorov distance between two probability measures:

$$d_K(\mu_{\mathcal{Q}^{(i)}}, \mu^{(i)}) := \sup_{x \in \mathbb{R}} \left| \mu_{\mathcal{Q}^{(i)}}((-\infty, x]) - \mu^{(i)}((-\infty, x]) \right|.$$

For any fixed $\alpha, t \geq 0$, the ESDs of $\alpha\mathcal{Q}^{(2)}$ and $\mathcal{Q}^{(1)} + t(\mathcal{Q}^{(1)})^2$ converge weakly to two measures $\mu_{\alpha}^{(2)}$ and $\mu_t^{(1)}$ defined through

$$\mu_{\alpha}^{(2)}((-\infty, x]) = \int \mathbf{1}_{\alpha y \in (-\infty, x]} d\mu^{(2)}(y), \quad \mu_t^{(1)}((-\infty, x]) = \int \mathbf{1}_{y+ty^2 \in (-\infty, x]} d\mu^{(1)}(y).$$

Hence, their Stieltjes transforms are given by

$$m_{\mu_{\alpha}^{(2)}}(z) = \frac{1}{\alpha} m_{\mu^{(2)}}\left(\frac{z}{\alpha}\right), \quad m_{\mu_t^{(1)}}(z) = \int \frac{d\mu^{(1)}(x)}{x + tx^2 - z}. \quad (\text{C.8})$$

Note that the eigenmatrices of $\mathcal{Q}^{(1)} + t(\mathcal{Q}^{(1)})^2$ and $\alpha\mathcal{Q}^{(2)}$ are independent Haar-distributed orthogonal matrices. Hence, the ESD of $\mathcal{Q}^{(1)} + t(\mathcal{Q}^{(1)})^2 + \alpha\mathcal{Q}^{(2)}$ converges weakly to the free addition $\mu_t^{(1)} \boxplus \mu_{\alpha}^{(2)}$. In particular, we will use the following almost sharp estimate on the difference between the Stieltjes transforms of $\mu_{\mathcal{Q}^{(1)}+t(\mathcal{Q}^{(1)})^2+\alpha\mathcal{Q}^{(2)}}$ and $\mu_t^{(1)} \boxplus \mu_{\alpha}^{(2)}$.

Lemma 41 *In the setting of Theorem 6, suppose $\alpha, t \in [0, C]$ for a constant $C > 0$. Then, we have that*

$$\begin{aligned} & \left| \frac{1}{p} \text{Tr} \left[\frac{1}{\mathcal{Q}^{(1)} + t(\mathcal{Q}^{(1)})^2 + \alpha\mathcal{Q}^{(2)}} \right] - m_{\mu_t^{(1)} \boxplus \mu_{\alpha}^{(2)}}(z=0) \right| \\ & \prec \frac{1}{p} + d_K(\mu_{\mathcal{Q}^{(1)}+t(\mathcal{Q}^{(1)})^2}, \mu_t^{(1)}) + d_K(\mu_{\alpha\mathcal{Q}^{(2)}}, \mu_{\alpha}^{(2)}). \end{aligned} \quad (\text{C.9})$$

Proof This lemma is a consequence of Theorem 2.5 of Bao et al. (2017). In fact, equation (C.9) is proved for $z = E + i\eta$ with $E \in \text{supp}(\mu_t^{(1)} \boxplus \mu_\alpha^{(2)})$ and $\eta > 0$ in a follow-up work of Bao et al. (2017), but the proof there can be repeated almost verbatim in our setting with $z = 0$. ■

With the above lemma, in order to calculate the right-hand side of (C.4), we need to calculate $\partial_t m_{\mu_t^{(1)} \boxplus \mu_\alpha^{(2)}}(z = 0)$ at $t = 0$. This is given by the following lemma.

Lemma 42 *We have*

$$\left. \frac{d}{dt} m_{\mu_t^{(1)} \boxplus \mu_\alpha^{(2)}}(0) \right|_{t=0} = - \frac{1 - 2f_1(\alpha)f_3(\alpha) + f_2(\alpha)f_3(\alpha)^2}{1 - \xi_2 f_2(\alpha)f_3(\alpha)^2}, \quad (\text{C.10})$$

where the functions f_1 , f_2 and f_3 are defined as

$$\begin{aligned} f_1(\alpha) &:= m_{\mu_0^{(1)} \boxplus \mu_\alpha^{(2)}}(0) \\ &= \frac{2}{\alpha(1 - \xi_2) + (1 - \xi_1) + \sqrt{[\alpha(1 - \xi_2) + (1 - \xi_1)]^2 + 4\alpha(\xi_1 + \xi_2 - \xi_1\xi_2)}}, \end{aligned} \quad (\text{C.11})$$

$$f_2(\alpha) := \left(\frac{1}{f_1(\alpha)^2} - \frac{\xi_1}{(1 + \xi_1 f_1(\alpha))^2} \right)^{-1}, \quad (\text{C.12})$$

$$f_3(\alpha) := \frac{\alpha}{1 + \alpha\xi_2 f_1(\alpha)}. \quad (\text{C.13})$$

Proof We calculate the Stieltjes transform of the free addition $\mu_t^{(1)} \boxplus \mu_\alpha^{(2)}$ using the following lemma. The following lemma is known in the literature.

Lemma 43 *Given two probability measures, μ_1 and μ_2 on \mathbb{R} , there exist unique analytic functions $\omega_1, \omega_2 : \mathbb{C}^+ \rightarrow \mathbb{C}^+$, where $\mathbb{C}^+ := \{z \in \mathbb{C} : \text{Im } z > 0\}$ is the upper half complex plane, such that the following equations hold: for any $z \in \mathbb{C}^+$,*

$$m_{\mu_1}(\omega_2(z)) = m_{\mu_2}(\omega_1(z)), \quad \omega_1(z) + \omega_2(z) - z = -\frac{1}{m_{\mu_1}(\omega_2(z))}. \quad (\text{C.14})$$

Moreover, $m_{\mu_1}(\omega_2(z))$ is the Stieltjes transform of $\mu_1 \boxplus \mu_2$, that is,

$$m_{\mu_1 \boxplus \mu_2}(z) = m_{\mu_1}(\omega_2(z)).$$

We now solve equation (C.14) for $\mu_1 = \mu_t^{(1)}$ and $\mu_2 = \mu_\alpha^{(2)}$ when $z \rightarrow 0$:

$$m_{\mu_t^{(1)}}(\omega_2(\alpha, t)) = m_{\mu_\alpha^{(2)}}(\omega_1(\alpha, t)), \quad \omega_1(\alpha, t) + \omega_2(\alpha, t) = -\frac{1}{m_{\mu_t^{(1)}}(\omega_2(\alpha, t))}, \quad (\text{C.15})$$

where, for simplicity, we omit the argument $z = 0$ from $\omega_1(z = 0, \alpha, t)$ and $\omega_2(z = 0, \alpha, t)$. Using the definition of $m_{\mu_\alpha^{(2)}}$ in equation (C.8), we can check that

$$\alpha g_2(\alpha m_{\mu_\alpha^{(2)}}(z)) = z, \quad (\text{C.16})$$

where g_2 is defined in equation (C.6). Applying equation (C.16) to the first equation of (C.15), we get that

$$\omega_1 = \frac{\alpha}{1 + \xi_2 \alpha m_{\mu_t^{(1)}}(\omega_2)} - \frac{1}{m_{\mu_t^{(1)}}(\omega_2)}.$$

Plugging this equation into the second equation of (C.15), we get that

$$\frac{\alpha}{1 + \xi_2 \alpha m_{\mu_t^{(1)}}(\omega_2)} + \omega_2 = 0 \Leftrightarrow \alpha + \omega_2 \left[1 + \alpha \xi_2 m_{\mu_t^{(1)}}(\omega_2) \right] = 0. \quad (\text{C.17})$$

This gives a self-consistent equation of $m_{\mu_t^{(1)} \boxplus \mu_\alpha^{(2)}}(z=0) = m_{\mu_t^{(1)}}(\omega_2(\alpha, t))$.

Now, we define the following quantities at $t=0$:

$$f_1(\alpha) := m_{\mu_0^{(1)}}(\omega_2(\alpha, 0)) = m_{\mu_0^{(1)} \boxplus \mu_\alpha^{(2)}}(z=0),$$

$$f_2(\alpha) := \left. \frac{dm_{\mu_0^{(1)}}(z)}{dz} \right|_{z=\omega_2(\alpha, 0)} = \int \frac{d\mu^{(1)}(x)}{[x - \omega_2(\alpha, 0)]^2}, \quad f_3(\alpha) := -\omega_2(\alpha, 0).$$

First, from equation (C.17) we can obtain that equation (C.13) is true. Using the fact that g_1 in equation (C.6) is the inverse function of $m_{\mu_0^{(1)}}$, we can write equation (C.17) into an equation of f_1 only when $t=0$:

$$\alpha + \left(\frac{1}{1 + \xi_1 f_1} - \frac{1}{f_1} \right) (1 + \alpha \xi_2 f_1) = 0.$$

This equation can be reduced to a quadratic equation:

$$\alpha (\xi_1 + \xi_2 - \xi_1 \xi_2) f_1^2 + [\alpha(1 - \xi_2) + (1 - \xi_1)] f_1 - 1 = 0. \quad (\text{C.18})$$

By definition, f_1 is the Stieltjes transform of $\mu_0^{(1)} \boxplus \mu_\alpha^{(2)}$ at $z=0$. Since $\mu_0^{(1)} \boxplus \mu_\alpha^{(2)}$ is supported on $(0, \infty)$, f_1 is positive by equation (C.5). Then, it is not hard to see that the only positive solution of equation (C.18) is given by equation (C.11). Finally, calculating the derivative of $m_{\mu_0^{(1)}}$ using its inverse function, we obtain that $f_2(\alpha) = [g_1'(f_1)]^{-1}$, which implies that equation (C.12) holds.

To conclude the proof, we still need to calculate $\partial_t m_{\mu_t^{(1)}}(\omega_2(\alpha, t))|_{t=0}$. Taking the derivative of equation (C.17) with respect to t at $t=0$, we get that

$$\partial_t \omega_2(\alpha, 0) \cdot [1 + \alpha \xi_2 f_1(\alpha)] - \alpha \xi_2 f_3(\alpha) \cdot \partial_t m_{\mu_t^{(1)}}(\omega_2(\alpha, t)) \Big|_{t=0} = 0. \quad (\text{C.19})$$

Using equation (C.8), we can calculate that

$$\partial_t m_{\mu_t^{(1)}}(\omega_2(\alpha, t)) = \partial_t \int \frac{d\mu^{(1)}(x)}{x + tx^2 - \omega_2(\alpha, t)} = - \int \frac{[x^2 - \partial_t \omega(\alpha, t)] d\mu^{(1)}(x)}{[x + tx^2 - \omega_2(\alpha, t)]^2}.$$

Taking $t=0$ in the above equation, we get that

$$\begin{aligned} & \partial_t m_{\mu_t^{(1)}}(\omega_2(\alpha, t)) \Big|_{t=0} \\ &= \partial_t \omega(\alpha, 0) \cdot f_2(\alpha) - \int \frac{[(x - \omega_2(\alpha, 0))^2 + 2\omega_2(\alpha, 0)(x - \omega_2(\alpha, 0)) + \omega_2(\alpha, 0)^2] d\mu^{(1)}(x)}{[x - \omega_2(\alpha, 0)]^2} \\ &= \partial_t \omega(\alpha, 0) \cdot f_2(\alpha) - 1 + 2f_1(\alpha)f_3(\alpha) - f_2(\alpha)f_3(\alpha)^2. \end{aligned}$$

We can solve from this equation that

$$\partial_t \omega(\alpha, 0) = \frac{1}{f_2(\alpha)} \left[\partial_t m_{\mu_t^{(1)}}(\omega_2(\alpha, t)) \Big|_{t=0} + 1 - 2f_1(\alpha)f_3(\alpha) + f_2(\alpha)f_3(\alpha)^2 \right].$$

Inserting it into equation (C.19), we can solve that

$$\partial_t m_{\mu_t^{(1)}}(\omega_2(\alpha, t)) \Big|_{t=0} = -\frac{1 - 2f_1(\alpha)f_3(\alpha) + f_2(\alpha)f_3(\alpha)^2}{1 - \frac{\alpha\xi_2 f_2(\alpha)f_3(\alpha)}{1 + \alpha\xi_2 f_1(\alpha)}}.$$

Using $(1 + \alpha\xi_2 f_1(\alpha))^{-1} = \alpha^{-1} f_3(\alpha)$ by equation (C.17), we conclude Lemma 42. \blacksquare

Now, we are ready to give the proof of Theorem 6.

Proof The estimate (4.1) is a special case of Theorem 3 with $M = \text{Id}_{p \times p}$. For a sanity check, we show that it is consistent with the result obtained from the free addition technique. We can write the variance term (2.10) as

$$L_{\text{var}} = \sigma^2 \text{Tr} \left[\frac{1}{n_1 \mathcal{Q}^{(1)} + n_2 \mathcal{Q}^{(2)}} \right] = \frac{p\sigma^2}{n_1} \cdot \frac{1}{p} \text{Tr} \left[\frac{1}{\mathcal{Q}^{(1)} + \alpha \mathcal{Q}^{(2)}} \right]. \quad (\text{C.20})$$

Combining Lemma 41 with estimate (C.7), we get that

$$\left| \frac{1}{p} \text{Tr} \left[\frac{1}{\mathcal{Q}^{(1)} + \alpha \mathcal{Q}^{(2)}} \right] - m_{\mu_0^{(1)} \boxplus \mu_\alpha^{(2)}}(0) \right| \prec \frac{1}{p}.$$

Recall that $m_{\mu_0^{(1)} \boxplus \mu_\alpha^{(2)}}(0) = f_1(\alpha)$ by estimate (C.11). Plugging into $\alpha = n_2/n_1$, we indeed get that $\frac{p}{n_1} f(\alpha) = L_1$.

Next, for the bias term, we first consider the case $\alpha = n_2/n_1 \leq 1$. Recall that the bias limit is given by estimate (C.4). Taking $t = p^{-1/2}$, we can use Lemma 22 to check that w.o.p.,

$$\left| \frac{dh_\alpha(t)}{dt} \Big|_{t=0} - \frac{h_\alpha(t) - h_\alpha(0)}{t} \right| \lesssim t.$$

Similarly, we have w.o.p.,

$$\left| \frac{dm_{\mu_t^{(1)} \boxplus \mu_\alpha^{(2)}}(0)}{dt} \Big|_{t=0} - \frac{m_{\mu_t^{(1)} \boxplus \mu_\alpha^{(2)}}(0) - m_{\mu_0^{(1)} \boxplus \mu_\alpha^{(2)}}(0)}{t} \right| \lesssim t.$$

On the other hand, using Lemma 41 and estimate (C.7), we get that

$$\left| h_\alpha(0) - m_{\mu_0^{(1)} \boxplus \mu_\alpha^{(2)}}(0) \right| + \left| h_\alpha(t) - m_{\mu_t^{(1)} \boxplus \mu_\alpha^{(2)}}(0) \right| \prec \frac{1}{p}.$$

Combining the above three estimates, we obtain that

$$\left| \frac{dh_\alpha(t)}{dt} \Big|_{t=0} - \frac{dm_{\mu_t^{(1)} \boxplus \mu_\alpha^{(2)}}(0)}{dt} \Big|_{t=0} \right| \prec t + \frac{|h_\alpha(0) - m_{\mu_0^{(1)} \boxplus \mu_\alpha^{(2)}}(0)| + |h_\alpha(t) - m_{\mu_t^{(1)} \boxplus \mu_\alpha^{(2)}}(0)|}{t} \prec p^{-1/2}. \quad (\text{C.21})$$

Plugging this estimate into estimate (C.4), after a straightforward calculation estimate using equation (C.10), we can get the estimate (4.2) when $\alpha \leq 1$.

Finally, we consider the $\alpha > 1$ (recall that $\alpha = \frac{n_2}{n_1}$) case for the bias term. We rewrite equations (C.1) and (C.4) as

$$\begin{aligned} L_{\text{bias}} &= \left(1 + O_{\prec}(p^{-\frac{1}{2}})\right) \cdot \frac{\|\mathbf{v}\|^2}{\alpha^2 p} \text{Tr} \left[\left(\alpha^{-1} \mathcal{Q}^{(1)} + \mathcal{Q}^{(2)} \right)^{-2} (\mathcal{Q}^{(1)})^2 \right] \\ &= - \left(1 + O_{\prec}(p^{-\frac{1}{2}})\right) \|\mathbf{v}\|^2 \cdot \tilde{\alpha}^2 \left. \frac{d\tilde{h}_{\tilde{\alpha}}(t)}{dt} \right|_{t=0}, \end{aligned}$$

where $\tilde{\alpha} := \alpha^{-1} = n_1/n_2$ and $\tilde{h}_{\tilde{\alpha}}(t)$ is defined by

$$\tilde{h}_{\tilde{\alpha}}(t) := \frac{1}{p} \text{Tr} \left[\frac{1}{\tilde{\alpha} \mathcal{Q}^{(1)} + t(\mathcal{Q}^{(1)})^2 + \mathcal{Q}^{(2)}} \right].$$

We can estimate $\left. \frac{d\tilde{h}_{\tilde{\alpha}}(t)}{dt} \right|_{t=0}$ in the same way as $\left. \frac{dh_{\alpha}(t)}{dt} \right|_{t=0}$ for $\tilde{\alpha} < 1$, by establishing analogues of Lemma 41 and Lemma 42. Through these calculations, we can find that $\tilde{\alpha}^2 \left. \frac{d\tilde{h}_{\tilde{\alpha}}(t)}{dt} \right|_{t=0}$ converges to the same asymptotic limit as $\left. \frac{dh_{\alpha}(t)}{dt} \right|_{t=0}$ as expected. However, the error term becomes

$$O_{\prec}(\tilde{\alpha}^2 p^{-\frac{1}{2}}) = O_{\prec} \left(\frac{p^{-\frac{1}{2}} n_1^2}{(n_1 + n_2)^2} \right),$$

where we use that $\tilde{\alpha} < \frac{2n_1}{n_1+n_2}$ when $n_1 < n_2$. This concludes equation (4.2) for the case when $\alpha > 1$. \blacksquare

C.2 Proof of Proposition 8

Proof Let $a = p^{-1} \text{Tr} [\Sigma^{(1)}]$. Since $\|(\Sigma^{(1)})^{1/2} (\beta^{(1)} - \beta^{(2)})\|^2 = (2 + O(p^{-1/2+c}))\mu^2 a$ with high probability, the limit of $L(\hat{\beta}_2^{\text{HPS}})$ is equal to

$$\ell(n_1) := \sigma^2 L_1 + 2\mu^2 a L_2 = \frac{\sigma^2 p}{n_1 + n_2 - p} + 2\mu^2 a \cdot \frac{n_1^2(n_1 + n_2 - p) + p n_1 n_2}{(n_1 + n_2)^2(n_1 + n_2 - p)}. \quad (\text{C.22})$$

By Theorem 6, we have that with high probability,

$$L(\hat{\beta}_2^{\text{HPS}}) = \left(1 + O(p^{-1/2+c})\right) \cdot \ell(n_1). \quad (\text{C.23})$$

By equation (3.3), with high probability over the randomness of $X^{(2)}$, the excess risk of the OLS estimator is

$$L(\hat{\beta}_2^{\text{OLS}}) = \sigma^2 \cdot \text{Tr} \left[\Sigma^{(2)} (X^{(2)\top} X^{(2)})^{-1} \right] = (1 + O(p^{-1+c})) \cdot \frac{\sigma^2 p}{n_2 - p}. \quad (\text{C.24})$$

Thus, whether or not $L(\hat{\beta}_2^{\text{HPS}}) \leq L(\hat{\beta}_2^{\text{OLS}})$ reduces to comparing $\ell(n_1)$ and $\frac{\sigma^2 p}{n_2 - p}$. Let $h(n_1)$ be their difference:

$$h(n_1) = 2\mu^2 a \cdot \frac{n_1^2(n_1 + n_2 - p) + pn_1 n_2}{(n_1 + n_2)^2(n_1 + n_2 - p)} - \frac{\sigma^2 p n_1}{(n_1 + n_2 - p)(n_2 - p)}.$$

The sign of $h(n_1)$ is the same as the sign of the following second-order polynomial in n_1 :

$$\begin{aligned} \tilde{h}(n_1) &= 2\mu^2 a(n_2 - p)(n_1(n_1 + n_2 - p) + pn_2) - \sigma^2 p(n_1 + n_2)^2 \\ &= (2\mu^2 a(n_2 - p) - \sigma^2 p)n_1^2 + (2\mu^2 a(n_2 - p)^2 - 2\sigma^2 p n_2)n_1 + (2\mu^2 a(n_2 - p)pn_2 - \sigma^2 p n_2^2). \end{aligned}$$

Let C_0, C_1, C_2 be the coefficients of n_1^0, n_1^1, n_1^2 terms in $\tilde{h}(n_1)$. We prove each claim as follows.

1. If $\mu^2 a \leq \frac{\sigma^2 p}{2(n_2 - p)}$, then C_0, C_1, C_2 are all non-positive, which gives $\tilde{h}(n_1) \leq 0$.
2. If $\frac{\sigma^2 p}{2(n_2 - p)} < \mu^2 a < \frac{\sigma^2 n_2}{2(n_2 - p)}$, then $C_2 > 0$ and $C_0 < 0$. Thus, $\tilde{h}(n_1)$ has a positive root and a negative root. Let the positive root of $\tilde{h}(n_1)$ be $n_0 \in \mathbb{R}$. Then, $\tilde{h}(n_1) \leq 0$ if $n_1 \leq n_0$, and $\tilde{h}(n_1) \geq 0$ otherwise.
3. If $\mu^2 a \geq \frac{\sigma^2 n_2}{2(n_2 - p)}$, then C_0 and C_2 are both non-negative. Furthermore, we can check that $C_1 \geq 0$ using the assumption $n_2 \geq 3p$. Hence, we have $\tilde{h}(n_1) \geq 0$ for all n_1 .

Combining these three cases with equations (C.23) and (C.24) concludes the proof. \blacksquare

C.2.1 PROOF OF THE DICHOTOMY

Proof The derivative of $\ell(n_1)$ is equal to

$$\ell'(n_1) = -\frac{p\sigma^2}{(n-p)^2} + \frac{2\mu^2 a n_2}{(n-p)^2} \left(2 \left(1 - \frac{n_2}{n} \right) \left(1 - \frac{2p}{n} \right) \left(1 - \frac{p}{n} \right) + \frac{p(n_2 - p)}{n^2} \right),$$

where we denote $n = n_1 + n_2$. Hence, the sign of $\ell'(n_1)$ is given by the sign of the following function:

$$f(n) := 2 \left(1 - \frac{n_2}{n} \right) \left(1 - \frac{2p}{n} \right) \left(1 - \frac{p}{n} \right) + \frac{p(n_2 - p)}{n^2} - \frac{p\sigma^2}{2\mu^2 a n_2}.$$

Under the assumption $n \geq n_2 \geq 3p$, we see that

$$f'(n) > \frac{2n_2}{n^2} \left(1 - \frac{2p}{n} \right) \left(1 - \frac{p}{n} \right) - \frac{2p(n_2 - p)}{n^3} \geq 0,$$

i.e., $f(n)$ is strictly increasing with respect to n . Hence, if $f(n_2) \geq 0$, then case 1 holds; otherwise, if $f(n_2) < 0$, then case 2 holds. This concludes the proof. \blacksquare

C.3 Proof of Theorem 9

Proof Since $X^{(1)}$ and $X^{(2)}$ are both isotropic Gaussian matrices, with high probability, they satisfy the following normalization conditions: there exists a constant $c > 0$, such that for $i = 1, 2$ and any vector $\theta \in \mathbb{R}^p$,

$$c \|\theta\| \leq \frac{1}{\sqrt{n_i}} \|X^{(i)} \theta\| \leq \frac{1}{c} \|\theta\|, \quad (\text{C.25})$$

where we use $X_j^{(i)}$ to denote the j -th column of the matrix $X^{(i)}$. For reference, see, e.g., Section 3.2 of Raskutti et al. (2011). These two estimates also follow directly from Lemma 22. For the rest of the proof, we will condition on $X^{(1)}$ and $X^{(2)}$, which are regarded as fixed design matrices satisfying the normalization condition (C.25). For convenience, denote $n = n_1 + n_2$ and $r_1 = n_1/n$. Our proof is divided into two cases.

- i) $\frac{n_1^2 \mu^2}{(n_1 + n_2)^2} \leq \frac{\sigma^2 p}{n_1 + n_2}$. Notice that $\frac{\sigma^2 p}{n_1 + n_2}$ is the minimax lower bound for fixed design linear regression when $\mu = 0$, i.e., all of the $n_1 + n_2$ samples are drawn from the same linear regression model. Thus, by invoking some standard arguments (see e.g., Example 15.14 of Wainwright (2019)), we can show equation (4.6).
- ii) $\frac{n_1^2 \mu^2}{(n_1 + n_2)^2} > \frac{\sigma^2 p}{n_1 + n_2}$. Then, we need to show that the minimax rate is at least of order $\min(\frac{n_1^2 \mu^2}{(n_1 + n_2)^2}, \frac{\sigma^2 p}{n_2})$. This can be done by constructing an $O(r_1 \mu)$ -cover of the subset of parameter vectors.

We next prove each case separately following the arguments in Example 15.14 of Wainwright (2019).

For case i), we design a covering set, denoted as $\mathcal{S} = \{\tilde{\theta}^1, \tilde{\theta}^2, \dots, \tilde{\theta}^M\} \cup \{\theta^1, \theta^2, \dots, \theta^M\}$, such that the followings hold:

- $\{\theta^1, \theta^2, \dots, \theta^M\}$ is a 2δ -packing of the subset of p -dimensional vectors of Euclidean length at most 4δ .
- $\|\tilde{\theta}^i - \theta^i\| = \mu$ for $i = 1, 2, \dots, M$. We achieve this by adding a fixed shift a of length μ to θ^i , i.e., $\tilde{\theta}^i = \theta^i + a$.
- The size of \mathcal{S} can be as large as 2^{p+1} .

For any $j = 1, \dots, M$, let \mathbb{P}_j denote the law of the label vector $y = \begin{pmatrix} Y^{(1)} \\ Y^{(2)} \end{pmatrix}$ of the two tasks when the true regression vectors are $\beta^{(1)} = \tilde{\theta}^j$ and $\beta^{(2)} = \theta^j$. By definition of the linear model, conditioning on $X^{(1)}$ and $X^{(2)}$, \mathbb{P}_j follows a multivariate normal distribution $\mathcal{N}\left(\begin{pmatrix} X^{(1)} \tilde{\theta}^j \\ X^{(2)} \theta^j \end{pmatrix}, \sigma^2 \text{Id}_{n \times n}\right)$. Similarly, for any $k \neq j$, let \mathbb{P}_k denote the law of the label vector when the true regression vectors are $\tilde{\theta}^k$ and θ^k . By standard facts on the KL divergence

between two normal distributions, we have

$$\begin{aligned} D_{KL}(\mathbb{P}_j \| \mathbb{P}_k) &= \frac{1}{2\sigma^2} \left(\left\| X^{(1)}(\tilde{\theta}^j - \tilde{\theta}^k) \right\|^2 + \left\| X^{(2)}(\theta^j - \theta^k) \right\|^2 \right) \\ &\leq \frac{1}{2\sigma^2} \left(c^{-2}n_1 \|\tilde{\theta}^j - \tilde{\theta}^k\|^2 + c^{-2}n_2 \|\theta^j - \theta^k\|^2 \right) \\ &\leq \frac{n}{2c^2\sigma^2} \|\theta^j - \theta^k\|^2 \leq \frac{32n\delta^2}{c^2\sigma^2}, \end{aligned}$$

where we use equation (C.25) in the second step. We need the above bound to satisfy the following condition (see equation (15.35b), Section 15.3.3 of Wainwright (2019)):

$$\frac{1}{2} \log |\mathcal{S}| \geq \frac{32n\delta^2}{c^2\sigma^2} + \log 2, \quad (\text{C.26})$$

which can be satisfied by setting $\delta^2 = \frac{c^2\sigma^2 p}{100n}$ for large enough p . Thus, we conclude that

$$\begin{aligned} \inf_{\hat{\beta}} \sup_{\Theta(\mu)} \mathbb{E} \left[\frac{1}{n_2} \left\| \tilde{X}^{(2)}(\hat{\beta} - \beta^{(2)}) \right\|^2 \right] &= \inf_{\hat{\beta}} \sup_{\Theta(\mu)} \mathbb{E} \left[\left\| \hat{\beta} - \beta^{(2)} \right\|^2 \right] \\ &\geq \frac{\delta^2}{2} \geq \frac{c^2\sigma^2 p}{200n} \geq \frac{c^2}{400} \left(\frac{n_1^2 \mu^2}{(n_1 + n_2)^2} + \frac{\sigma^2 p}{n_1 + n_2} \right). \end{aligned} \quad (\text{C.27})$$

Next, we consider case ii), which is further divided into two cases. If $r_1^2 \mu^2 \geq \frac{\sigma^2 p}{n_2}$, then the source task samples are not helpful in improving the rate. In this case, we just need to follow a similar calculation as above, but use the label vector $y = Y^{(2)}$ instead with the KL divergence $\frac{1}{2\sigma^2} \|X^{(2)}(\theta^i - \theta^k)\|^2$. Then, the lower bound in equation (C.27) becomes

$$\frac{c^2}{200} \frac{\sigma^2 p}{n_2} \geq \frac{c^2}{400} \left(\frac{\sigma^2 p}{n_2} + \frac{\sigma^2 p}{n_1 + n_2} \right).$$

On the other hand, if $\frac{\sigma^2 p}{n_1 + n_2} < r_1^2 \mu^2 < \frac{\sigma^2 p}{n_2}$, we again construct a 2δ -packing of the subset of p -dimensional vectors of Euclidean length at most 4δ , but we choose the label vector as $y = Y^{(2)}$. Following a similar argument as above, we reach the conclusion that

$$D_{KL}(\mathbb{P}_j \| \mathbb{P}_k) \leq \frac{32n_2\delta^2}{c^2\sigma^2}.$$

By setting $\delta^2 = \frac{c^2 r_1^2 \mu^2}{100}$, we know that the above bound is less than $\frac{1}{2}(p-1) \log 2$. Thus, the condition for the KL bound (C.26) is satisfied. Hence, the lower bound in equation (C.27) becomes

$$\frac{\delta^2}{2} = \frac{c^2 r_1^2 \mu^2}{200} \geq \frac{c^2}{400} \left(\frac{n_1^2 \mu^2}{(n_1 + n_2)^2} + \frac{\sigma^2 p}{n_1 + n_2} \right).$$

Combining all three cases together and renaming the constant c , we complete the proof of estimate (4.6). \blacksquare

C.4 Proof of Theorem 10

Theorem 10 can be proved in a similar way to Theorem 6 by using free additions.

Proof In the current setting with $\Sigma^{(2)} = \text{Id}$, we denote $\alpha = n_2/n_1$ and

$$\mathcal{Q}^{(1)} := \frac{1}{n_1} X^{(1)\top} X^{(1)} = \frac{1}{n_1} \Sigma^{(1)1/2} Z^{(1)\top} Z^{(1)} \Sigma^{(1)1/2}, \quad \mathcal{Q}^{(2)} := \frac{1}{n_2} Z^{(2)\top} Z^{(2)}.$$

In terms of these notations, L_{var} still satisfies equation (C.20), and L_{bias} satisfies a similar equality as estimate (C.1):

$$\begin{aligned} L_{\text{bias}} &= \left(1 + O_{\prec}(p^{-1/2})\right) \cdot \frac{2\mu^2}{p} \text{Tr} \left[\left(\mathcal{Q}^{(1)} + \alpha \mathcal{Q}^{(2)} \right)^{-2} (\mathcal{Q}^{(1)})^2 \right] \\ &= - \left(1 + O_{\prec}(p^{-1/2})\right) \cdot 2\mu^2 \left. \frac{dh_{\alpha}(t)}{dt} \right|_{t=0}, \end{aligned} \tag{C.28}$$

where we use that under the random-effect model, $\beta^{(1)} - \beta^{(2)}$ has i.i.d. Gaussian entries of mean zero and variance $2\mu^2/p$.

The ESD of $\mathcal{Q}^{(1)}$ converges weakly to a deformed MP law $\mu^{(1)}$ determined by the eigenvalues of $\Sigma^{(1)}$ (Marčenko and Pastur, 1967) as described below. Let $\mu^{(0)}$ be the asymptotic ESD of $\mathcal{Q}^{(0)} := X^{(1)} X^{(1)\top} / n_1$. Since $\mathcal{Q}^{(0)}$ has the same non-zero eigenvalues as $\mathcal{Q}^{(1)}$ and $n_1 - p$ more zero eigenvalues, we have that

$$\mu^{(0)} = \frac{p}{n_1} \mu^{(1)} + \frac{n_1 - p}{n_1} \delta_0.$$

Hence, the Stieltjes transforms of $\mu^{(0)}$ and $\mu^{(1)}$ satisfy the following relation:

$$m_{\mu^{(1)}} = \frac{n_1}{p} m_{\mu^{(0)}} + \frac{n_1 - p}{pz}.$$

Moreover, similar to (E.3), $m_{\mu^{(0)}}$ satisfies the deformed MP equation

$$\frac{1}{m_{\mu^{(0)}}(z)} = -z + \frac{1}{n_1} \sum_{i=1}^p \frac{\sigma_i^{(1)}}{1 + \sigma_i^{(1)} m_{\mu^{(0)}}(z)},$$

where $\sigma_i^{(1)}$, $i = 1, \dots, p$, are the eigenvalues of $\Sigma^{(1)}$. Thus, we have that $g_0(m_{\mu^{(0)}}(z)) = z$, where recall that the function g_0 is defined in (5.1). The sharp convergence rate of $\mu_{\mathcal{Q}^{(0)}}$ follows from Theorem 3.12 of Knowles and Yin (2016):

$$d_K(\mu_{\mathcal{Q}^{(0)}}, \mu^{(0)}) \prec p^{-1}.$$

Similar to the proof of Theorem 6 in Appendix C.1, we need to calculate $m_{\mu_t^{(1)} \boxplus \mu_{\alpha}^{(2)}}(0)$ and its derivative $\partial_t m_{\mu_t^{(1)} \boxplus \mu_{\alpha}^{(2)}}(0)$ at $t = 0$. First, we follow exactly the same argument as in the proof of Lemma 42 and derive the self-consistent equation (C.17) for $\omega_2(\alpha, t)$. Then, we again introduce the following functions:

$$f_1(\alpha) := m_{\mu_0^{(1)}}(\omega_2(\alpha, 0)) = \frac{n_1}{p} m_{\mu^{(0)}}(\omega_2(\alpha, 0)) + \frac{n_1 - p}{p\omega_2(\alpha, 0)},$$

$$f_2(\alpha) := \left. \frac{dm_{\mu_0^{(1)}}(z)}{dz} \right|_{z=\omega_2(\alpha,0)}, \quad f_3(\alpha) := -\omega_2(\alpha,0).$$

We rewrite equation (C.17) at $t = 0$ and $\alpha = n_2/n_1$ as

$$\frac{n_2}{n_1} \left(1 + \frac{\xi_2}{\xi_1} - \xi_2 \right) + \omega_2 \cdot \left[1 + m_{\mu^{(0)}}(\omega_2(\alpha,0)) \right] = 0. \quad (\text{C.29})$$

Using the fact that g_0 is the inverse function of $m_{\mu^{(0)}}$, we can rewrite (C.29) into an equation of $y_0 := m_{\mu^{(0)}}(\omega_2(\alpha,0))$ at $t = 0$:

$$\frac{n_1 + n_2 - p}{n_1} + g_0(y_0) \cdot (1 + y_0) = 0,$$

which is the equation (5.2). We can express $f_1(\alpha)$ and $f_3(\alpha)$ as in equation (5.3). Regarding f_2 , we have that

$$f_2(\alpha) = \frac{n_1}{p} m'_{\mu^{(0)}}(\omega_2(\alpha,0)) - \frac{n_1 - p}{p[\omega_2(\alpha,0)]^2} = \frac{n_1}{p \cdot g'_0(y_0)} - \frac{n_1 - p}{p \cdot [g_0(y_0)]^2}.$$

Finally, following the argument below line (C.19), we can obtain line (C.10).

From these calculations, we obtain $m_{\mu_0^{(1)} \boxplus \mu_\alpha^{(2)}}(0) = f_1(\alpha)$ and $\partial_t m_{\mu_t^{(1)} \boxplus \mu_\alpha^{(2)}}(0)|_{t=0}$ given by equation (C.10). Plugging them into equations (C.20) and (C.28), we conclude the proof. ■

C.5 An Approximate Estimate under Covariate and Model Shifts

With the local convergence of the resolvent, we can also derive an approximate estimate of the bias under arbitrary covariate and model shifts, which may be of independent interest. The accuracy of this estimate increases as n_1/p increases.

Theorem 44 (Anisotropic covariance) *Under Assumption 1, if $n_1 \gg p$, then for any small constant $c > 0$, with high probability over the randomness of the training samples, the following holds:*

$$\begin{aligned} & \left| L_{\text{bias}} - \left\| \Pi^{1/2} \Sigma^{(1)1/2} (\beta^{(1)} - \beta^{(2)}) \right\|^2 \right| \\ &= O \left(\left(\sqrt{\frac{p}{n_1}} + n_1^{2/\varphi - 1/2 + c} \right) \frac{\lambda_1^2}{\lambda_p^2} \right) \left\| \Sigma^{(1)1/2} (\beta^{(1)} - \beta^{(2)}) \right\|^2. \end{aligned} \quad (\text{C.30})$$

Above, λ_1 and λ_p are the largest and smallest singular values of M , respectively. Π is a $p \times p$ matrix defined as

$$\Pi = \frac{n_1^2}{(n_1 + n_2)^2} M \frac{\alpha_3 M^\top M + \alpha_4 + 1}{(\alpha_1 M^\top M + \alpha_2)^2} M^\top,$$

where α_1 and α_2 are defined in equation (3.2), and α_3 and α_4 are the unique positive solutions of the following system of equations:

$$\begin{aligned}\alpha_3 + \alpha_4 &= \frac{1}{n_1 + n_2} \sum_{i=1}^p \frac{1}{\lambda_i^2 \alpha_1 + \alpha_2}, \\ \alpha_3 + \frac{1}{n_1 + n_2} \sum_{i=1}^p \frac{\lambda_i^2 (\alpha_2 \alpha_3 - \alpha_1 \alpha_4)}{(\lambda_i^2 \alpha_1 + \alpha_2)^2} &= \frac{1}{n_1 + n_2} \sum_{i=1}^p \frac{\lambda_i^2 \alpha_1}{(\lambda_i^2 \alpha_1 + \alpha_2)^2}.\end{aligned}\tag{C.31}$$

Proof We first show an estimate on

$$\tilde{L}_{\text{bias}} := n_1^2 \left\| \Sigma^{(2)1/2} \hat{\Sigma}^{-1} \Sigma^{(1)} \left(\beta^{(1)} - \beta^{(2)} \right) \right\|^2.$$

We claim that for any small constant $c > 0$, there exists a high probability event on which the following estimate holds:

$$\left| \tilde{L}_{\text{bias}} - \left\| \Pi^{1/2} \Sigma^{(1)1/2} (\beta^{(1)} - \beta^{(2)}) \right\|^2 \right| \prec n^{-1/2} Q \left\| \Sigma^{(1)1/2} (\beta^{(1)} - \beta^{(2)}) \right\|^2.\tag{C.32}$$

Define the vector $\mathbf{v} := V^\top \Sigma^{(2)-1/2} \Sigma^{(1)} (\beta^{(1)} - \beta^{(2)}) \in \mathbb{R}^p$, and its embedding in \mathbb{R}^{p+n} , $\mathbf{w} = (\mathbf{v}^\top, \mathbf{0}_n^\top)^\top$, where $\mathbf{0}_n$ is an n -dimensional zero vector. Then, we can write that

$$\tilde{L}_{\text{bias}} = \mathbf{w}^\top \frac{n_1^2}{(\Lambda U^\top Z^{(1)\top} Z^{(1)} U \Lambda + V^\top Z^{(2)\top} Z^{(2)} V)^2} \mathbf{w} = \frac{n_1^2}{n^2} \mathbf{w}^\top G'(0) \mathbf{w},$$

where $G'(0)$ denotes the derivative of $G(z)$ with respect to z at $z = 0$. Again, we introduce the truncated matrices $\tilde{Z}^{(1)}$ and $\tilde{Z}^{(2)}$ in equation (B.13). With a similar argument as in the above proof of Theorem 3, we can show that (B.12) holds for $\mathbf{w}^\top G(z) \mathbf{w}$ on the event $\Xi_1 = \{\tilde{Z}^{(1)} = Z^{(1)}, \tilde{Z}^{(2)} = Z^{(2)}\}$. Combining equation (B.12) with Cauchy's integral formula, we get that on Ξ_1 ,

$$\begin{aligned}\mathbf{w}^\top G'(0) \mathbf{w} &= \frac{1}{2\pi i} \oint_{\mathcal{C}} \frac{\mathbf{w}^\top G(z) \mathbf{w}}{z^2} dz = \frac{1}{2\pi i} \oint_{\mathcal{C}} \frac{\mathbf{w}^\top \mathfrak{G}(z) \mathbf{w}}{z^2} dz + O_{\prec}(n^{-\frac{1}{2}} Q \|\mathbf{w}\|^2) \\ &= \mathbf{w}^\top \mathfrak{G}'(0) \mathbf{w} + O_{\prec}(n^{-\frac{1}{2}} Q \|\mathbf{w}\|^2),\end{aligned}\tag{C.33}$$

where \mathcal{C} is the contour $\{z \in \mathbb{C} : |z| = (\log n)^{-1}\}$. With equation (B.7), we can calculate the derivative $\mathbf{w}^\top \mathfrak{G}'(0) \mathbf{w}$ as

$$\mathbf{w}^\top \mathfrak{G}'(0) \mathbf{w} = \mathbf{v}^\top \frac{\alpha_3 \Lambda^2 + (1 + \alpha_4) \text{Id}_{p \times p}}{(\alpha_1 \Lambda^2 + \alpha_2 \text{Id}_{p \times p})^2} \mathbf{v},$$

where

$$\alpha_3 := - \left. \frac{d\alpha_1(z)}{dz} \right|_{z=0}, \quad \alpha_4 := - \left. \frac{d\alpha_2(z)}{dz} \right|_{z=0}.$$

Taking derivatives of the system of equations in equation (B.8) with respect to z at $z = 0$, we can derive equation (C.31). Together with equation (C.33), this concludes that equation (C.32) is true.

Now, with equation (C.32), to conclude equation (C.30) it remains to bound

$$L_{\text{bias}} - \tilde{L}_{\text{bias}} = 2n_1(\beta^{(1)} - \beta^{(2)})^\top \Sigma^{(1)1/2} \Delta \left[\Sigma^{(1)1/2} \hat{\Sigma}^{-1} \Sigma^{(2)} \hat{\Sigma}^{-1} \Sigma^{(1)1/2} \right] \Sigma^{(1)1/2} (\beta^{(1)} - \beta^{(2)}) \\ + \left\| \Sigma^{(2)1/2} \hat{\Sigma}^{-1} \Sigma^{(1)1/2} \Delta \Sigma^{(1)1/2} (\beta^{(1)} - \beta^{(2)}) \right\|^2,$$

where we abbreviate $\Delta = Z^{(1)\top} Z^{(1)} - n_1 \text{Id}_{p \times p}$. From this equation, we get that

$$\left| L_{\text{bias}} - \tilde{L}_{\text{bias}} \right| \leq \left((n_1 + \|\Delta\|)^2 - n_1^2 \right) \left\| \Sigma^{(1)1/2} \hat{\Sigma}^{-1} \Sigma^{(2)} \hat{\Sigma}^{-1} \Sigma^{(1)1/2} \right\| \cdot \left\| \Sigma^{(1)1/2} (\beta^{(1)} - \beta^{(2)}) \right\|^2.$$

By Corollary 23, the following estimates holds with high probability for any constant $c > 0$:

$$(n_1 + \|\Delta\|)^2 - n_1^2 \leq n_1^2 \left[\left(1 + \sqrt{\frac{p}{n_1}} \right)^4 - 1 + n_1^{-1/2+2/\varphi+c} \right], \\ \left\| \Sigma^{(2)1/2} \hat{\Sigma}^{-1} \Sigma^{(1)1/2} \right\| \leq \frac{1 + n_1^{-1/2+2/\varphi+c}}{\lambda_p^2(\sqrt{n_1} - \sqrt{p})^2 + (\sqrt{n_2} - \sqrt{p})^2}.$$

Combining the above three estimates and $\|M\|_2 \leq \lambda_1$, we obtain that with high probability,

$$\left| L_{\text{bias}} - \tilde{L}_{\text{bias}} \right| \leq \left(\left(1 + \sqrt{\frac{p}{n_1}} \right)^4 - 1 + n_1^{-1/2+2/\varphi+c} \right) \frac{\lambda_1^2 n_1^2 \left\| \Sigma^{(1)1/2} (\beta^{(1)} - \beta^{(2)}) \right\|^2}{(\lambda_p^2(\sqrt{n_1} - \sqrt{p})^2 + (\sqrt{n_2} - \sqrt{p})^2)^2}.$$

Combining this estimate with equation (C.32), we conclude the proof of Theorem 44. \blacksquare

Appendix D. Proofs for the Multiple Tasks Setting

We first derive the HPS estimator. For the optimization objective $\ell(A, B)$ in equation (5.6), using the local optimality condition $\frac{\partial f}{\partial B} = 0$, we can obtain \hat{B} as a function of A :

$$\hat{B}(A) = (X^\top X)^{-1} X^\top Y A^\top (A A^\top)^+,$$

where $Y := [Y^{(1)}, Y^{(2)}, \dots, Y^{(t)}]$ and $(A A^\top)^+$ denotes the pseudoinverse of $A A^\top$. Plugging $\hat{B}(A)$ into equation (5.6), we obtain the following objective that depends only on A (in matrix notation):

$$g(A) = \left\| X(X^\top X)^{-1} X^\top Y A^\top (A A^\top)^+ A - Y \right\|_F^2. \quad (\text{D.1})$$

Note that $A^\top (A A^\top)^+ A$ is a projection onto the subspace spanned by the rows of A . For simplicity, we write it into the form

$$A^\top (A A^\top)^+ A = U_A U_A^\top,$$

where $U_A \in \mathbb{R}^{t \times r}$ is a $t \times r$ partial orthonormal matrix (i.e., $U_A^\top U_A = \text{Id}_{r \times r}$). Hence, we also denote the function $g(A)$ by $g(U_A)$.

Now, to find the optimal solution for U_A , we write $X(X^\top X)^{-1}X^\top = P_X$ (which is an n by n projection matrix of rank p), and expand $g(A)$ as:

$$\begin{aligned} g(A) &= \left\| P_X Y U_A U_A^\top - Y \right\|_F^2 = \left\| P_X Y U_A U_A^\top \right\|_F^2 + \|Y\|_F^2 - 2\langle P_X Y U_A U_A^\top, Y \rangle \\ &= \|Y\|_F^2 - \langle P_X Y U_A U_A^\top, Y \rangle = \|Y\|_F^2 - \langle U_A U_A^\top, Y^\top P_X Y \rangle, \end{aligned}$$

where we use the fact that $P_X^2 = P_X$ since it is a projection matrix and $U_A^\top U_A = \text{Id}_{r \times r}$. Thus, to minimize $g(A)$, we just need to maximize the inner product between $Y^\top P_X Y$ and $U_A U_A^\top$ —this can be achieved by finding the best rank- r SVD of $Y^\top P_X Y$ and taking the top- r singular vectors, which form the partial orthonormal matrix $U_r = U_{\hat{A}}$. Lastly, we can insert U_r as A back to $\hat{B}A$ to obtain $\hat{B}\hat{A}$ as the HPS estimators for all tasks from 1 to t .

D.1 Proof of Theorem 12

First, we can use the concentration estimates in Appendix A to simplify the expression of $g(U_A)$. In the following proof, we always let $Q = n^{2/\varphi}$.

Lemma 45 *In the setting of Theorem 12, for any small constant $c > 0$ and large constant $C > 0$, there exists a high probability event Ξ , on which the following estimate holds:*

$$|g(U_A) - h(U_A)| \leq Q n^{\frac{1}{2}+c} \left\| \Sigma^{\frac{1}{2}} B^* (U_A U_A^\top - \text{Id}_{t \times t}) \right\|_F^2 + \sigma^2 n^{\frac{1}{2}+c} + p^{-C} \left\| \Sigma^{\frac{1}{2}} B^* \right\|_F^2, \quad (\text{D.2})$$

uniformly in all rank- r partial orthonormal matrices $U_A \in \mathbb{R}^{t \times r}$. Here, h is defined by

$$h(U_A) := n \left\| \Sigma^{1/2} B^* (U_A U_A^\top - \text{Id}_{t \times t}) \right\|_F^2 + \sigma^2 (nt - pr). \quad (\text{D.3})$$

Proof With Corollary 23 and Corollary 24, we can choose a high probability event Ξ_1 on which the estimate (A.18) holds and

$$C^{-1}n \leq \lambda_p(Z^\top Z) \leq \lambda_1(Z^\top Z) \leq Cn \quad (\text{D.4})$$

for a large constant $C > 0$. Throughout the following proof, we assume that event Ξ_1 holds.

To facilitate the analysis, we introduce the following matrix notations. Denote

$$\mathcal{E} := [\varepsilon^{(1)}, \varepsilon^{(2)}, \dots, \varepsilon^{(t)}], \quad \text{and} \quad \mathcal{W} := X(X^\top X)^{-1}X^\top \mathcal{E} U_A U_A^\top.$$

For any $j = 1, 2, \dots, t$, denote

$$H_j := B^* (U_A U_A^\top - \text{Id}_{t \times t}) e_j, \quad E_j := (\mathcal{W} - \mathcal{E}) e_j, \quad (\text{D.5})$$

where e_j is the standard basis unit vector along the j -th direction. Then, plugging $Y = XB^* + \mathcal{E}$ into equation (D.1), we can write the function $g(U_A)$ as

$$g(U_A) = \sum_{j=1}^t \|X H_j + E_j\|^2.$$

We will divide $g(A)$ into three parts.

The first part is

$$\sum_{j=1}^t \|XH_j\|^2 = \left\| XB^*(U_A U_A^\top - \text{Id}_{t \times t}) \right\|_F^2.$$

Applying estimate (A.18) to $XH_j = Z\Sigma^{\frac{1}{2}}H_j$, we obtain that on Ξ_1 ,

$$\begin{aligned} \|XH_j\|^2 &= n \|\Sigma^{\frac{1}{2}}H_j\|^2 \cdot \left(1 + O_{\prec}(n^{-\frac{1}{2}}Q)\right) \\ &= n \left\| \Sigma^{\frac{1}{2}}B^*(U_A U_A^\top - \text{Id}_{t \times t})e_j \right\|^2 \cdot \left(1 + O_{\prec}(n^{-\frac{1}{2}}Q)\right), \end{aligned}$$

which gives that

$$\left| \sum_{j=1}^t \|XH_j\|^2 - n \left\| \Sigma^{1/2}B^*(U_A U_A^\top - \text{Id}_{t \times t}) \right\|_F^2 \right| \prec Qn^{1/2} \left\| \Sigma^{1/2}B^*(U_A U_A^\top - \text{Id}_{t \times t}) \right\|_F^2.$$

The second part is the cross term

$$2 \sum_{j=1}^t \langle XH_j, E_j \rangle = 2 \langle XB^*(U_A U_A^\top - \text{Id}_{t \times t}), \mathcal{W} - \mathcal{E} \rangle = -2 \langle XB^*(U_A U_A^\top - \text{Id}_{t \times t}), \mathcal{E} \rangle.$$

Using estimate (A.20), we obtain that

$$\begin{aligned} |\langle XB^*(U_A U_A^\top - \text{Id}_{t \times t}), \mathcal{E} \rangle| &\prec \sigma \left\| XB^*(U_A U_A^\top - \text{Id}_{t \times t}) \right\|_F \\ &\lesssim \sigma n^{1/2} \left\| \Sigma^{1/2}B^*(U_A U_A^\top - \text{Id}_{t \times t}) \right\|_F \\ &\leq \sigma^2 n^{1/2} + n^{1/2} \left\| \Sigma^{1/2}B^*(U_A U_A^\top - \text{Id}_{t \times t}) \right\|_F^2, \end{aligned}$$

where we use $X = Z\Sigma^{\frac{1}{2}}$ and equation (D.4) in the second step, and use the AM-GM inequality in the third step.

The last part is

$$\sum_{j=1}^t \|E_j\|^2 = \|\mathcal{W} - \mathcal{E}\|_F^2 = \|\mathcal{E}\|_F^2 - \langle \mathcal{W}, \mathcal{E} \rangle,$$

where in the second step we use $\|\mathcal{W}\|_F^2 = \langle \mathcal{W}, \mathcal{E} \rangle$ by algebraic calculation. Using estimate (A.21), we get that

$$\left| \|\varepsilon^{(i)}\|^2 - \sigma^2 n \right| \prec \sigma^2 \sqrt{n}, \quad (\text{D.6})$$

and

$$\begin{aligned} &\left| \varepsilon^{(i)\top} X(X^\top X)^{-1} X^\top \varepsilon^{(j)} - \delta_{ij} \cdot p\sigma^2 \right| \\ &= \left| \varepsilon^{(i)\top} X(X^\top X)^{-1} X^\top \varepsilon^{(j)} - \delta_{ij} \cdot \sigma^2 \text{Tr}[X(X^\top X)^{-1} X^\top] \right| \\ &\prec \sigma^2 \left\| X(X^\top X)^{-1} X^\top \right\|_F = \sigma^2 \sqrt{p}. \end{aligned} \quad (\text{D.7})$$

Summing up equation (D.6) over i , we obtain that

$$\left| \|\mathcal{E}\|_F^2 - \sigma^2 nt \right| \leq \sum_{i=1}^t \left| \|\varepsilon^{(i)}\|^2 - \sigma^2 n \right| \prec \sigma^2 \sqrt{n}. \quad (\text{D.8})$$

Using line (D.7), we can estimate the inner product between \mathcal{W} and \mathcal{E} as

$$\begin{aligned} |\langle \mathcal{W}, \mathcal{E} \rangle - \sigma^2 pr| &= \left| \text{Tr} \left[\left(\mathcal{E}^\top X (X^\top X)^{-1} X^\top \mathcal{E} - p\sigma^2 \cdot \text{Id}_{t \times t} \right) U_A U_A^\top \right] \right| \\ &\leq r \left\| \mathcal{E}^\top X (X^\top X)^{-1} X^\top \mathcal{E} - p\sigma^2 \cdot \text{Id}_{t \times t} \right\| \prec \sigma^2 \sqrt{n}. \end{aligned} \quad (\text{D.9})$$

Combining equations (D.8) and (D.9), we obtain that

$$\sum_{j=1}^t \|E_j\|^2 = \sigma^2(nt - pr) + O_{\prec}(\sigma^2 \sqrt{n}).$$

Combining the concentration estimates for all three parts, we obtain that on event Ξ_1 ,

$$|g(U_A) - h(U_A)| \prec Q\sqrt{n} \left\| \Sigma^{1/2} B^* (U_A U_A^\top - \text{Id}_{t \times t}) \right\|_F^2 + \sigma^2 \sqrt{n}$$

for any fixed U_A . Then, using a standard ε -net argument with $\varepsilon = p^{-C}$, we can conclude that equation (D.2) is true uniformly in all rank- r partial orthonormal matrices U_A . \blacksquare

From equation (D.3), we see that the global minimizer of $h(U_A)$ is the best rank- r approximation A^* of $B^{*\top} \Sigma B^*$ defined in equation (5.8). On the other hand, let $U_{\hat{A}}$ be a global minimizer of $g(U_A)$. We have the following characterization of $U_{\hat{A}} U_{\hat{A}}^\top$, which is a consequence of Lemma 45.

Lemma 46 *In the setting of Theorem 12, we have that*

$$\left\| U_{\hat{A}} U_{\hat{A}}^\top - A^* A^{*\top} \right\|_F^2 \lesssim n^{-1/2+c} \frac{Q \|B^{*\top} \Sigma B^*\| + \sigma^2}{\lambda_r - \lambda_{r+1}},$$

on the high probability event Ξ in Lemma 45.

Proof Using triangle inequality, we upper bound $h(U_{\hat{A}}) - h(A^*)$ as

$$\begin{aligned} h(U_{\hat{A}}) - h(A^*) &\leq (g(U_{\hat{A}}) - g(A^*)) + |g(A^*) - h(A^*)| + |g(U_{\hat{A}}) - h(U_{\hat{A}})| \\ &\leq |g(A^*) - h(A^*)| + |g(U_{\hat{A}}) - h(U_{\hat{A}})| \lesssim n^{1/2+c} \left(Q \|\Sigma^{1/2} B^*\|_F^2 + \sigma^2 \right), \end{aligned} \quad (\text{D.10})$$

on the event Ξ . Here, in the second step, we used the fact that $U_{\hat{A}}$ is the global minimizer of $g(\cdot)$, so that $g(U_{\hat{A}}) \leq g(A^*)$, and in the third step we used equation (D.2), $\|U_A U_A^\top - \text{Id}_{t \times t}\| \leq 1$ and $p^{-C} \leq Q n^{1/2+c}$ for large enough $C > 0$. Using the definition of $h(U_A)$ in equation (D.3), we can check that

$$h(U_{\hat{A}}) - h(A^*) = n \cdot \text{Tr} \left[B^{*\top} \Sigma B^* (A^* A^{*\top} - U_{\hat{A}} U_{\hat{A}}^\top) \right].$$

For $1 \leq i \leq t$, let λ_i be the i -th largest eigenvalue of $B^\star{}^\top \Sigma B^\star$ and v_i be the corresponding eigenvector. Then, we have $A^\star A^{\star\top} = \sum_{i=1}^r v_i v_i^\top$, and

$$\begin{aligned} h(U_{\hat{A}}) - h(A^\star) &= n \sum_{i=1}^r \lambda_i - n \sum_{i=1}^t \lambda_i \|U_{\hat{A}}^\top v_i\|^2 \\ &= n \sum_{i=1}^r \lambda_i \left(1 - \|U_{\hat{A}}^\top v_i\|^2\right) - n \sum_{i=r+1}^t \lambda_i \|U_{\hat{A}}^\top v_i\|^2 \\ &\geq n(\lambda_r - \lambda_{r+1}) \sum_{i=r+1}^t \|U_{\hat{A}}^\top v_i\|^2, \end{aligned} \quad (\text{D.11})$$

where in the last step we use that

$$\sum_{i=1}^r \left(1 - \|U_{\hat{A}}^\top v_i\|^2\right) = r - \sum_{i=1}^r \|U_{\hat{A}}^\top v_i\|^2 = \sum_{i=r+1}^t \|U_{\hat{A}}^\top v_i\|^2.$$

From equation (D.10) and equation (D.11), we obtain that

$$\sum_{i=r+1}^t \|U_{\hat{A}}^\top v_i\|^2 \lesssim n^{-1/2+c} \frac{Q \|\Sigma^{1/2} B^\star\|_F^2 + \sigma^2}{\lambda_r - \lambda_{r+1}}. \quad (\text{D.12})$$

On the other hand, we have

$$\left\|A^\star A^{\star\top} - U_{\hat{A}} U_{\hat{A}}^\top\right\|_F^2 = 2r - 2 \left\langle A^\star A^{\star\top}, U_{\hat{A}} U_{\hat{A}}^\top \right\rangle = 2 \sum_{i=r+1}^t \|U_{\hat{A}}^\top v_i\|^2.$$

Combined with equation (D.12) and $\|\Sigma^{1/2} B^\star\|_F^2 \lesssim \|B^\star{}^\top \Sigma B^\star\|$, we conclude the proof. \blacksquare

The last piece of the proof of Theorem 12 is the following concentration estimate on the prediction loss of $\hat{\beta}_i^{\text{HPS}}(A)$.

Lemma 47 *In the setting of Theorem 12, denote*

$$\hat{\beta}_i^{\text{HPS}}(U_A) := \hat{B}(A)A = (X^\top X)^{-1} X^\top Y U_A U_A^\top,$$

and $a_i := U_A U_A^\top e_i$. Then, for any small constant $c > 0$ and large constant $C > 0$, there exists a high probability event Ξ , on which the following estimate holds:

$$\begin{aligned} &\left| L_i(\hat{\beta}_i^{\text{HPS}}(U_A)) - L_i(B^\star a_i) - \sigma^2 \|a_i\|^2 \cdot \text{Tr} \left[\Sigma (X^\top X)^{-1} \right] \right| \\ &\leq n^{-1/2+c} (L_i(B^\star a_i) + \sigma^2 \|a_i\|^2) + p^{-C} \left\| \Sigma^{1/2} B^\star \right\|_F^2, \end{aligned} \quad (\text{D.13})$$

uniformly in all rank- r partial orthonormal matrices $U_A \in \mathbb{R}^{t \times r}$, where L_i is defined in (5.7).

Proof The following proof is restricted to a high probability event Ξ_1 on which equation (D.4) holds. For any fixed U_A , the prediction loss of $\hat{\beta}_i^{\text{HPS}}(U_A)$ is

$$L_i(\hat{\beta}_i^{\text{HPS}}(U_A)) = \left\| \Sigma^{1/2} \left((X^\top X)^{-1} X^\top Y a_i - \beta^{(i)} \right) \right\|^2 = \left\| \Sigma^{1/2} (H_i + R_i) \right\|^2,$$

where we denote $R_i = (X^\top X)^{-1} X^\top \mathcal{E} a_i$ and H_i is defined in equation (D.5). The rest of the proof is similar to that of Lemma 45. We divide the prediction loss into three parts.

Part 1: The first part is the bias term $\|\Sigma^{1/2} H_i\|^2 = L_i(B^* a_i)$.

Part 2: The second part is the cross term $2\langle \Sigma^{1/2} H_i, \Sigma^{1/2} R_i \rangle$. We can bound it as

$$\begin{aligned} \left| \langle \Sigma^{1/2} H_i, \Sigma^{1/2} R_i \rangle \right| &= \left| \langle X(X^\top X)^{-1} \Sigma H_i, \mathcal{E} a_i \rangle \right| \leq \sum_{j=1}^t |a_i(j)| \cdot \left| \langle X(X^\top X)^{-1} \Sigma H_i, \varepsilon^{(j)} \rangle \right| \\ &\prec \sum_{j=1}^t |a_i(j)| \cdot \sigma \left\| X(X^\top X)^{-1} \Sigma H_i \right\| \lesssim \frac{\|a_i\| \sigma}{n^{1/2}} \left\| \Sigma^{1/2} H_i \right\| \\ &\leq n^{-1/2} \sigma^2 \|a_i\|^2 + n^{-1/2} L_i(B^* a_i). \end{aligned}$$

Here, in the second step, we use $a_i(j)$ to denote the j -th coordinate of a_i ; in the third step, we use equation (A.20); in the fourth step, we use equations (D.4), (2.3) and $\sum_j |a_i(j)| \leq \sqrt{t} \|a_i\|$ by Cauchy-Schwarz inequality; in the last step, we use AM-GM inequality and $\|\Sigma^{1/2} H_i\|^2 = L_i(B^* a_i)$.

Part 3: The last part is

$$\begin{aligned} \|\Sigma^{1/2} R_i\|^2 &= \left\| \sum_{j=1}^t a_i(j) \Sigma^{1/2} (X^\top X)^{-1} X^\top \varepsilon^{(j)} \right\|^2 \\ &= \sum_{1 \leq j, k \leq t} a_i(j) a_i(k) \varepsilon^{(j)\top} X (X^\top X)^{-1} \Sigma (X^\top X)^{-1} X^\top \varepsilon^{(k)}. \end{aligned} \quad (\text{D.14})$$

Using estimate (A.21) and $\text{Tr}[X(X^\top X)^{-1} \Sigma (X^\top X)^{-1} X^\top] = \text{Tr}[\Sigma (X^\top X)^{-1}]$, we obtain that

$$\begin{aligned} &\left| \varepsilon^{(j)\top} X (X^\top X)^{-1} \Sigma (X^\top X)^{-1} X^\top \varepsilon^{(k)} - \delta_{jk} \cdot \sigma^2 \text{Tr}[\Sigma (X^\top X)^{-1}] \right| \\ &\prec \sigma^2 \left\| X (X^\top X)^{-1} \Sigma (X^\top X)^{-1} X^\top \right\|_F \lesssim \sigma^2 p^{1/2} n^{-1} \leq \sigma^2 n^{-1/2}, \end{aligned} \quad (\text{D.15})$$

where we use estimate (D.4) in the second step. Plugging estimate (D.15) into estimate (D.14), we obtain that

$$\begin{aligned} \left| \left\| \Sigma^{1/2} R_i \right\|^2 - \sigma^2 \|a_i\|^2 \cdot \text{Tr}[\Sigma (X^\top X)^{-1}] \right| &\prec \sigma^2 n^{-1/2} \sum_{1 \leq j, k \leq t} |a_i(j)| |a_i(k)| \\ &\lesssim n^{-1/2} \sigma^2 \|a_i\|^2. \end{aligned}$$

Combining the concentration error estimates for all three parts, we obtain that on event Ξ_1 ,

$$\left| L_i(\hat{\beta}_i^{\text{HPS}}) - L_i(B^* a_i) - \sigma^2 \|a_i\|^2 \text{Tr}[\Sigma (X^\top X)^{-1}] \right| \prec n^{-\frac{1}{2}} (L_i(B^* a_i) + \sigma^2 \|a_i\|^2), \quad (\text{D.16})$$

for any fixed U_A . Then, using a standard ε -net argument with $\varepsilon = p^{-C}$, we can conclude equation (D.13) holds uniformly in all rank- r partial orthonormal matrices U_A . \blacksquare

Combining Lemma 46 and Lemma 47, we are ready to prove Theorem 12.

Proof Using Lemma 47 and applying Theorem 3 (with $n_1 = 0$, $n_2 = n$ and $\Sigma^{(2)} = \Sigma$) to $\text{Tr} [\Sigma(X^\top X)^{-1}]$, we get that for $\hat{a}_i := U_{\hat{A}} U_{\hat{A}}^\top e_i$,

$$\begin{aligned} & \left| L_i(\hat{\beta}_i^{\text{HPS}}(U_{\hat{A}})) - L_i(B^\star \hat{a}_i) - \frac{p\sigma^2}{n-p} \|\hat{a}_i\|^2 \right| \\ & \leq n^{-\frac{1}{2}+c} L_i(B^\star \hat{a}_i) + p^{-\frac{1}{2}} n^{\frac{2}{\varphi}-\frac{1}{2}+c} \frac{p\sigma^2 \|\hat{a}_i\|^2}{n} + p^{-C} \|\Sigma^{1/2} B^\star\|_F^2 \end{aligned} \quad (\text{D.17})$$

with high probability. From this equation, using $\varphi > 4$, we obtain that

$$\begin{aligned} & \left| L_i(\hat{\beta}_i^{\text{HPS}}(U_{\hat{A}})) - L_i(B^\star a_i^\star) - \frac{p\sigma^2}{n-p} \|a_i^\star\|^2 \right| \\ & \leq n^{-1/2+c} \left(\left\| \Sigma^{1/2} B^\star \right\|_2^2 + \sigma^2 \|a_i^\star\|^2 \right) + \left(\left\| \Sigma^{1/2} B^\star \right\|_2^2 + \sigma^2 \right) \|\hat{a}_i - a_i^\star\|, \end{aligned} \quad (\text{D.18})$$

where we also use that $L_i(B^\star a_i^\star) \leq \left\| \Sigma^{1/2} B^\star \right\|_2^2$. On the other hand, using Lemma 46, we can bound that with high probability,

$$\|\hat{a}_i - a_i^\star\|^2 \lesssim n^{-1/2+c} \frac{Q \left\| B^{\star\top} \Sigma B^\star \right\|_2 + \sigma^2}{\lambda_r - \lambda_{r+1}}.$$

Plugging it into equation (D.18), we conclude that equation (5.9) is true. \blacksquare

D.2 Proof of Proposition 14

Proof First, as a corollary of equation (5.9), by averaging over $i = 1, \dots, t$, with high probability over the randomness of the training examples, the averaged risk satisfies

$$\begin{aligned} & \left| \frac{1}{t} \sum_{i=1}^t L_i(\hat{\beta}_i^{\text{HPS}}) - \frac{1}{t} \left\| \Sigma^{1/2} B^\star (A^\star A^{\star\top} - \text{Id}_{t \times t}) \right\|_F^2 - \frac{\sigma^2 p}{n-p} \frac{r}{t} \right| \\ & \leq \left\| B^{\star\top} \Sigma B^\star \right\|_2 n^{\frac{2}{\varphi}-\frac{1}{2}+c} + \sigma^2 n^{-\frac{1}{2}+c}. \end{aligned} \quad (\text{D.19})$$

However, note that the error bound in equation (D.19) for the averaged risk is tighter than that in equation (5.9) for each individual risk. The proof of estimate (D.19) is postponed until we complete the proof of Proposition 14.

By Definition (5.8), the bias term is equal to the sum of the smallest $(t-r)$ eigenvalues of $B^{\star\top} \Sigma B^\star$:

$$\left\| \Sigma^{1/2} B^\star (A^\star A^{\star\top} - \text{Id}_{t \times t}) \right\|_F^2 = \text{Tr} \left[B^{\star\top} \Sigma B^\star (\text{Id}_{p \times p} - A^\star A^{\star\top}) \right].$$

By the concentration estimates of random vectors in Lemma 21, with high probability, the (i, j) -th entry of $B^{\star\top} \Sigma B^{\star}$ is equal to

$$\beta^{(i)\top} \Sigma \beta^{(j)} = \beta_0^\top \Sigma \beta_0 + \delta_{ij} \frac{\mu^2}{p} \text{Tr}[\Sigma] + O\left(p^{-1/2+c'} \|\beta_0\|^2 + p^{-1/2+c'} \mu^2\right),$$

for any constant $c' > 0$. Thus, $B^{\star\top} \Sigma B^{\star}$ is (approximately) equal to a rank-1 matrix with spectral norm $t \cdot \beta_0^\top \Sigma \beta_0$ plus a scalar matrix $\frac{\mu^2}{p} \text{Tr}[\Sigma] \text{Id}_{t \times t}$. Thus, we have that

$$\begin{aligned} \lambda_1 &= \left(1 + O(p^{-1/2+c'})\right) \left(t \cdot \beta_0^\top \Sigma \beta_0 + \frac{\mu^2}{p} \text{Tr}[\Sigma]\right), \\ \lambda_i &= \left(1 + O(p^{-1/2+c'})\right) \frac{\mu^2}{p} \text{Tr}[\Sigma], \quad \forall i = 2, \dots, t. \end{aligned}$$

This implies that the sum of the smallest $(t - r)$ eigenvalues of $B^{\star\top} \Sigma B^{\star}$ is equal to

$$\left(1 + O(p^{-1/2+c'})\right) \cdot (t - r) \frac{\mu^2 \text{Tr}[\Sigma]}{p},$$

with high probability. Thus, we conclude that with high probability,

$$\begin{aligned} g_r(n, \mu) - \frac{\sigma^2 p}{n - p} &= \left(1 + O(p^{-1/2+c'})\right) \cdot \left(1 - \frac{r}{t}\right) \frac{\mu^2}{p} \text{Tr}[\Sigma] + \frac{r}{t} \cdot \frac{\sigma^2 p}{n - p} - \frac{\sigma^2 p}{n - p} \\ &= \left(1 - \frac{r}{t}\right) \left[\left(1 + O(p^{-1/2+c'})\right) \cdot \frac{\mu^2 \text{Tr}[\Sigma]}{p} - \frac{\sigma^2 p}{n - p} \right]. \end{aligned} \quad (\text{D.20})$$

Now, we are ready to finish the proof. Let c' be a positive constant smaller than c .

1. If $\mu^2 > (1 + p^{-\frac{1}{2}+c}) \frac{\sigma^2 p^2}{(n-p) \text{Tr}[\Sigma]}$, the coefficient of $(1 - \frac{r}{t})$ in equation (D.20) is nonnegative and Claim 1 follows.
2. If $\mu^2 < (1 - p^{-\frac{1}{2}+c}) \frac{\sigma^2 p^2}{(n-p) \text{Tr}[\Sigma]}$, the coefficient of $(1 - \frac{r}{t})$ in equation (D.20) is negative. Thus, equation (D.20) is minimized when $r = 1$, and Claim 2 also follows.

This completes the proof of Proposition 14. ■

It remains to prove the estimate in equation (D.19).

Proof Summing up equation (D.17) over i , we obtain that

$$\begin{aligned} &\left| \sum_i L_i(\hat{\beta}_i^{\text{HPS}}(U_{\hat{A}})) - \left\| \Sigma^{\frac{1}{2}} B^{\star}(U_{\hat{A}} U_{\hat{A}}^\top - \text{Id}_{t \times t}) \right\|_F^2 - \frac{\sigma^2 p \cdot r}{n - p} \right| \\ &\leq n^{-\frac{1}{2}+c} \left(\left\| \Sigma^{\frac{1}{2}} B^{\star}(U_{\hat{A}} U_{\hat{A}}^\top - \text{Id}_{t \times t}) \right\|_F^2 + \sigma^2 \right) + p^{-C} \left\| \Sigma^{\frac{1}{2}} B^{\star} \right\|_F^2, \end{aligned} \quad (\text{D.21})$$

with high probability, where we use that $\sum_{i=1}^t \|\hat{a}_i\|^2 = r$ and

$$\sum_{i=1}^t L_i(B^{\star} \hat{a}_i) = \left\| \Sigma^{\frac{1}{2}} B^{\star}(U_{\hat{A}} U_{\hat{A}}^\top - \text{Id}_{t \times t}) \right\|_F^2.$$

With the result from equation (D.3), we can get

$$\begin{aligned} \left\| \Sigma^{\frac{1}{2}} B^* (U_{\hat{A}} U_{\hat{A}}^\top - \text{Id}_{t \times t}) \right\|_F^2 &= \frac{h(U_{\hat{A}}) - \sigma^2(nt - pr)}{n} \\ &\geq \frac{h(A^*) - \sigma^2(nt - pr)}{n} = \left\| \Sigma^{\frac{1}{2}} B^* (A^* A^{*\top} - \text{Id}_{t \times t}) \right\|_F^2, \end{aligned}$$

where in the second step we use the fact that A^* is the global minimizer of h . On the other hand, using equation (D.10), we can get that with high probability

$$\frac{h(U_{\hat{A}}) - \sigma^2(nt - pr)}{n} \leq \frac{h(A^*) - \sigma^2(nt - pr)}{n} + O\left(n^{-\frac{1}{2}+c} \left(Q \|\Sigma^{\frac{1}{2}} B^*\|_F^2 + \sigma^2\right)\right).$$

Combining the above two estimates, we get that with high probability,

$$\left\| \Sigma^{\frac{1}{2}} B^* (U_{\hat{A}} U_{\hat{A}}^\top - \text{Id}_{t \times t}) \right\|_F^2 = \left\| \Sigma^{\frac{1}{2}} B^* (A^* A^{*\top} - \text{Id}_{t \times t}) \right\|_F^2 + O\left(n^{-\frac{1}{2}+c} \left(Q \|\Sigma^{\frac{1}{2}} B^*\|_F^2 + \sigma^2\right)\right).$$

Plugging this estimate into equation (D.21), we conclude that equation (D.19) is true. \blacksquare

Appendix E. Proofs for SPS Estimation

We first derive the bias-variance decomposition for the SPS estimator in Lemma 15.

E.1 Proof of Lemma 15

Proof Minimizing equation (1.4) leads to the solution for z as:

$$\hat{z} = (X^{(1)\top} X^{(1)} + \lambda)^{-1} X^{(1)\top} (Y^{(1)} - X^{(1)} B).$$

Inserting the expression of \hat{z} into $\ell(B, z)$ leads us to the solution of B , which is the SPS estimator given as follows (recall that $\hat{\Sigma}_\lambda$ has been defined in equation (6.1)):

$$\begin{aligned} \hat{\beta}_2^{\text{SPS}} &= \hat{\Sigma}_\lambda^{-1} \left(\lambda^{-1} X^{(1)\top} X^{(1)} X^{(2)\top} Y^{(2)} + X^{(1)\top} Y^{(1)} + X^{(2)\top} Y^{(2)} \right) \\ &= \beta^{(2)} + \hat{\Sigma}_\lambda^{-1} \left(X^{(1)\top} X^{(1)} (\beta^{(1)} - \beta^{(2)}) + X^{(1)\top} \varepsilon^{(1)} + (\lambda^{-1} X^{(1)\top} X^{(1)} + 1) X^{(2)\top} \varepsilon^{(2)} \right). \end{aligned}$$

Hence, the excess risk of $\hat{\beta}_2^{\text{SPS}}$ is given by

$$L(\hat{\beta}_2^{\text{SPS}}) = \left\| \Sigma^{(2)\frac{1}{2}} \hat{\Sigma}_\lambda^{-1} \left(\hat{\Sigma}^{(1)} (\beta^{(1)} - \beta^{(2)}) + X^{(1)\top} \varepsilon^{(1)} + (\lambda^{-1} \hat{\Sigma}^{(1)} + 1) X^{(2)\top} \varepsilon^{(2)} \right) \right\|^2.$$

We can now take the expectation of $L(\hat{\beta}_2^{\text{SPS}})$ with respect to $\varepsilon^{(1)}$ and $\varepsilon^{(2)}$. Equations (6.3) and (6.4) follow after rearranging the terms in the expectation, and we can conclude that equation (6.2) is true. \blacksquare

E.2 Proof of Proposition 16

As part of the proof, we are going to need the following result.

Proposition 48 *Under the setting of Proposition 16, suppose that $Z^{(1)}$ satisfies the bounded support condition (A.1) with $Q = n^{\frac{2}{\varphi}}$. Then, the following anisotropic local law holds: for any deterministic unit vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^p$, we have that*

$$\sup_{t \in [-1/2, 1/2]} \left| \mathbf{u}^\top [G_t(-1) - \mathfrak{G}_t(-1)] \mathbf{v} \right| \prec n^{-1/2} Q, \quad (\text{E.1})$$

Proof For each fixed t , the anisotropic local law was established in the work of Ding and Yang (2018) (and its follow-up). Then, using a standard ε -net argument in t , we can extend the result uniformly to all $t \in [-1/2, 1/2]$ and conclude that equation (E.1) is true. \blacksquare

Now we are ready to present the proof of Proposition 16.

Proof As in Lemma 2, we can first show the following estimate that is similar to equation (A.28):

$$\left| L(\hat{\beta}_2^{\text{HPS}}) - L_{\text{bias}}(\lambda) - L_{\text{var}}(\lambda) \right| \prec \frac{L_{\text{bias}}(\lambda) + L_{\text{var}}(\lambda)}{\sqrt{p}}.$$

The proof essentially relies on the estimates (A.22) and (A.26) for the SPS estimation, employing the same reasoning as outlined in Appendix A.3.

Next, for $t \in \mathbb{R}$, we denote

$$\mathbf{v} := \Sigma^{(2)\frac{1}{2}} \left(\beta^{(1)} - \beta^{(2)} \right), \quad H_t := n_2 \left(Z^{(1)\top} Z^{(1)} \right)^{-1} + \frac{n_2}{\lambda} \Sigma^{(2)} + (1+t) \text{Id}_{p \times p}.$$

Then, under the setting of Theorem 16, we can write $L_{\text{bias}}(\lambda)$ and $L_{\text{var}}(\lambda)$ in equations (6.3) and (6.4) as

$$L_{\text{bias}}(\lambda) = \mathbf{v}^\top H_0^{-2} \mathbf{v}, \quad L_{\text{var}}(\lambda) = \frac{p\sigma^2}{n_2} - \frac{\sigma^2}{n_2} \text{Tr} [H_0^{-1}] - \frac{\sigma^2}{\lambda} \text{Tr} [\Sigma^{(2)} H_0^{-2}]. \quad (\text{E.2})$$

Note that $H_0^{-2} = -\partial_t H_t^{-1}|_{t=0}$. Hence, to evaluate equation (E.2), we need to derive precise estimates on $\text{Tr} [H_t^{-1}]$ and $\text{Tr} [\Sigma^{(2)} H_t^{-1}]$ for t around 0. For this purpose, we rewrite H_t^{-1} as

$$H_t^{-1} = \frac{n_1}{n_2} \tilde{\Sigma}_t^{-\frac{1}{2}} \left(\text{Id}_{p \times p} - \frac{1}{n_1^{-1} \tilde{\Sigma}_t^{\frac{1}{2}} Z^{(1)\top} Z^{(1)} \tilde{\Sigma}_t^{\frac{1}{2}} + 1} \right) \tilde{\Sigma}_t^{-\frac{1}{2}},$$

where $\tilde{\Sigma}_t := \frac{n_1}{\lambda} \Sigma^{(2)} + (1+t) \frac{n_1}{n_2} \text{Id}_{p \times p}$. Now, define the resolvent $G_t(z)$ and its averaged trace as

$$G_t(z) := \left(\frac{1}{n_1} \tilde{\Sigma}_t^{\frac{1}{2}} Z^{(1)\top} Z^{(1)} \tilde{\Sigma}_t^{\frac{1}{2}} - z \text{Id}_{p \times p} \right)^{-1}, \quad m_t(z) := \frac{1}{n_1} \text{Tr} [G_t(z)].$$

It is well-known that $m_t(z)$ converges to the Stieltjes transform of the MP law (Marčenko and Pastur, 1967), defined as the unique solution $a_t(z)$ to the deformed MP equation

$$\frac{1}{a_t(z)} = -z + \frac{1}{n_1} \text{Tr} \left[\frac{\tilde{\Sigma}_t}{\text{Id}_{p \times p} + a_t(z) \tilde{\Sigma}_t} \right] \quad (\text{E.3})$$

such that $\text{Im } a_t(z) \geq 0$ whenever $\text{Im } z \geq 0$. Furthermore, $G_t(z)$ is close to a deterministic matrix

$$\mathfrak{G}_t(z) := -\frac{1}{z(1 + a_t(z)\widetilde{\Sigma}_t)},$$

in the sense of Proposition 48, in which we take $z = -1$.

Now, we define

$$\Pi_t := \frac{n_1}{n_2} \widetilde{\Sigma}_t^{-\frac{1}{2}} [\text{Id}_{p \times p} - \mathfrak{G}_t(-1)] \widetilde{\Sigma}_t^{-\frac{1}{2}} = \frac{n_1}{n_2} \frac{a_t(-1)}{\text{Id}_{p \times p} + a_t(-1)\widetilde{\Sigma}_t}. \quad (\text{E.4})$$

By Proposition 48, H_t^{-1} is well-approximated by Π_t in the sense of anisotropic local law. Furthermore, using the argument in the proof of Theorem 6 (i.e., the argument above equation (C.21)), we obtain that for any deterministic unit vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^p$,

$$\left| \mathbf{u}^\top (H_0^{-2} + \partial_t \Pi_t|_{t=0}) \mathbf{v} \right| \prec n^{-\frac{1}{4}} Q^{\frac{1}{2}}. \quad (\text{E.5})$$

With estimates (E.2), (E.1) and (E.5), to conclude estimates (6.7) and (6.8), it remains to calculate Π_0 and $\partial_t \Pi_t|_{t=0}$.

First, when $t = 0$, Π_0 is equal to the matrix M_1 defined in equation (6.6) with $a_0(-1) = a_0$ and $\widetilde{\Sigma}_0 = \widetilde{\Sigma}$. Second, taking the derivative of equation (E.3) with respect to t at $t = 0$, we obtain that

$$\frac{\partial_t a_t(-1)|_{t=0}}{a_0^2} = -\frac{n_1}{n_2} x_0 + \partial_t a_t(-1)|_{t=0} \frac{1}{n_1} \text{Tr} \left[\left(\frac{\widetilde{\Sigma}_0}{\text{Id}_{p \times p} + a_0 \widetilde{\Sigma}_0} \right)^2 \right], \quad (\text{E.6})$$

where x_0 is defined in equation (6.5). Using equation (E.3), we obtain that

$$\begin{aligned} \frac{1}{n_1} \text{Tr} \left[\left(\frac{\widetilde{\Sigma}_0}{\text{Id}_{p \times p} + a_0 \widetilde{\Sigma}_0} \right)^2 \right] &= \frac{1}{a_0^2} \left(-\frac{p}{n_1} + \frac{2a_0}{n_1} \text{Tr} \frac{\widetilde{\Sigma}_0}{1 + a_0 \widetilde{\Sigma}_0} + x_0 \right) \\ &= \frac{1}{a_0^2} \left[-\frac{p}{n_1} + 2a_0 \left(\frac{1}{a_0} - 1 \right) + x_0 \right] = \frac{1}{a_0^2} \left(2 - \frac{p}{n_1} - 2a_0 + x_0 \right). \end{aligned}$$

Plugging it into estimate (E.6), we can solve that

$$\partial_t a_t(-1)|_{t=0} = \frac{n_1}{n_2} y_0, \quad (\text{E.7})$$

where y_0 is defined in equation (6.5). With estimate (E.7), we can calculate that

$$\partial_t \Pi_t|_{t=0} = \left(\frac{n_1}{n_2} \right)^2 \frac{y_0 - a_0^2}{(\text{Id}_{p \times p} + a_0 \widetilde{\Sigma}_0)^2} = -M_2, \quad (\text{E.8})$$

for M_2 defined in equation (6.6).

Finally, we combine the above ingredients to complete the proof. Using the truncation argument in the proof of Theorem 3 (i.e., the argument below equation (B.13)), it suffices to assume that $Z^{(1)}$ satisfies the bounded support condition (A.1) with $Q = n^{\frac{2}{\varphi}}$. Then, applying equations (E.1), (E.4), (E.5) and (E.8), we can evaluate equation (E.2) and conclude that both equations (6.7) and (6.8) are true. The proof of Proposition 16 is finished. \blacksquare

References

- Rie Kubota Ando and Tong Zhang. A framework for learning predictive structures from multiple tasks and unlabeled data. *Journal of Machine Learning Research*, 6(61):1817–1853, 2005.
- Zhidong Bai and Jack W. Silverstein. *Spectral Analysis of Large Dimensional Random Matrices*. Springer Series in Statistics. Springer, New York, 2nd edition, 2010.
- Zhigang Bao, László Erdős, and Kevin Schnelli. Local law of addition of random matrices on optimal scale. *Communications in Mathematical Physics*, 349(3):947–990, 2017.
- Jonathan Baxter. A model of inductive bias learning. *Journal of Artificial Intelligence Research*, 12:149–198, 2000.
- Shai Ben-David and Reba Schuller. Exploiting task relatedness for multiple task learning. In *Learning Theory and Kernel Machines*, pages 567–580. Springer, 2003.
- Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine Learning*, 79(1-2):151–175, 2010a.
- Shai Ben-David, Tyler Lu, Teresa Luu, and Dávid Pál. Impossibility theorems for domain adaptation. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 129–136, 2010b.
- Alex Bloemendal, László Erdős, Antti Knowles, Horng-Tzer Yau, and Jun Yin. Isotropic local laws for sample covariance and generalized Wigner matrices. *Electronic Journal of Probability*, 19(33):1–53, 2014.
- Rich Caruana. Multitask learning. *Machine Learning*, 28(1):41–75, 1997.
- Koby Crammer, Michael Kearns, and Jennifer Wortman Vaughan. Learning from multiple sources. *Journal of Machine Learning Research*, 9(8), 2008.
- Yehuda Dar and Richard G Baraniuk. Double double descent: On generalization errors in transfer learning between linear regression tasks. *SIAM Journal on Mathematics of Data Science*, 4(4):1447–1472, 2022.
- Yehuda Dar, Daniel LeJeune, and Richard G Baraniuk. The common intuition to transfer learning can win or lose: Case studies for linear regression. *SIAM Journal on Mathematics of Data Science*, 6(2):454–480, 2024.
- Oussama Dhifallah and Yue M Lu. Phase transitions in transfer learning for high-dimensional perceptrons. *Entropy*, 23(4):400, 2021.
- Xiucui Ding and Fan Yang. A necessary and sufficient condition for edge universality at the largest singular values of covariance matrices. *The Annals of Applied Probability*, 28(3):1679–1738, 2018.

- Edgar Dobriban and Yue Sheng. Wonder: Weighted one-shot distributed ridge regression in high dimensions. *Journal of Machine Learning Research*, 21(66):1–52, 2020.
- Yaqi Duan and Kaizheng Wang. Adaptive and robust multi-task learning. *The Annals of Statistics*, 51(5):2015–2039, 2023.
- László Erdős and Horng-Tzer Yau. *A Dynamical Approach to Random Matrix Theory*, volume 28 of *Courant Lecture Notes in Mathematics*. American Mathematical Soc., 2017.
- László Erdős, Antti Knowles, Horng-Tzer Yau, and Jun Yin. Spectral statistics of Erdős-Rényi graphs I: Local semicircle law. *The Annals of Probability*, 41(3B):2279–2375, 2013.
- Peizhong Ju, Sen Lin, Mark S Squillante, Yingbin Liang, and Ness B Shroff. Generalization performance of transfer learning: Overparameterized and underparameterized regimes. *arXiv preprint arXiv:2306.04901*, 2023.
- Sham M Kakade and Dean P Foster. Multi-view regression via canonical correlation analysis. In *International Conference on Computational Learning Theory*, pages 82–96. Springer, 2007.
- Antti Knowles and Jun Yin. Anisotropic local laws for random matrices. *Probability Theory and Related Fields*, pages 1–96, 2016.
- Mladen Kolar and Han Liu. Marginal regression for multitask learning. In *Artificial Intelligence and Statistics*, pages 647–655. PMLR, 2012.
- Mladen Kolar, John Lafferty, and Larry Wasserman. Union support recovery in multi-task learning. *Journal of Machine Learning Research*, 12(7), 2011.
- Qi Lei, Wei Hu, and Jason Lee. Near-optimal linear regression under distribution shift. In *International Conference on Machine Learning*, pages 6164–6174. PMLR, 2021.
- Dongyue Li, Huy L. Nguyen, and Hongyang R. Zhang. Identification of negative transfers in multitask learning using surrogate models. *Transactions on Machine Learning Research*, 2023.
- Dongyue Li, Aneesh Sharma, and Hongyang R. Zhang. Scalable multitask learning using gradient-based estimation of task affinity. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 1542–1553, 2024.
- Sai Li, T. Tony Cai, and Hongzhe Li. Transfer learning for high-dimensional linear regression: Prediction, estimation and minimax optimality. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(1):149–173, 2021.
- Vladimir A. Marčenko and Leonid A. Pastur. Distribution of eigenvalues for some sets of random matrices. *Mathematics of the USSR-Sbornik*, 1:457, 1967.
- Andreas Maurer. Bounds for linear multi-task learning. *Journal of Machine Learning Research*, 7(5):117–139, 2006.

- Alexandru Nica and Roland Speicher. *Lectures on the Combinatorics of Free Probability*, volume 13. Cambridge University Press, 2006.
- Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359, 2009.
- Natesh S. Pillai and Jun Yin. Universality of covariance matrices. *The Annals of Applied Probability*, 24:935–1001, 2014.
- Garvesh Raskutti, Martin J. Wainwright, and Bin Yu. Minimax rates of estimation for high-dimensional linear regression over ℓ_q -balls. *IEEE Transactions on Information Theory*, 57(10):6976–6994, 2011.
- Pragya Sur and Emmanuel J. Candès. A modern maximum-likelihood theory for high-dimensional logistic regression. *Proceedings of the National Academy of Sciences*, 116(29):14516–14525, 2019.
- Antonia M. Tulino and Sergio Verdú. *Random Matrix Theory and Wireless Communications*, volume 1 of *Foundations and Trends in Communications and Information Theory*. Now Publishers Inc, 2004.
- Konstantin Voevodski, Maria-Florina Balcan, Heiko Röglin, Shang-Hua Teng, and Yu Xia. Active clustering of biological sequences. *Journal of Machine Learning Research*, 13(1):203–225, 2012.
- Martin J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*, volume 48. Cambridge University Press, 2019.
- Jialei Wang, Mladen Kolar, and Nathan Srebro. Distributed multi-task learning. In *Artificial Intelligence and Statistics*, pages 751–760. PMLR, 2016.
- Sen Wu, Hongyang R. Zhang, and Christopher Ré. Understanding and improving information transfer in multi-task learning. In *International Conference on Learning Representations*, 2020.
- Tong Zhang. *Mathematical Analysis of Machine Learning Algorithms*. Cambridge University Press, 2023.
- Boxin Zhao, Cong Ma, and Mladen Kolar. Trans-glasso: A transfer learning approach to precision matrix estimation. *arXiv preprint arXiv:2411.15624*, 2024.
- Shurong Zheng, Zhidong Bai, and Jianfeng Yao. CLT for eigenvalue statistics of large-dimensional general Fisher matrices with applications. *Bernoulli*, 23(2):1130–1178, 2017.