

A Theory of Non-Linear Feature Learning with One Gradient Step in Two-Layer Neural Networks

Behrad Moniri*† Donghwan Lee*‡ Hamed Hassani† Edgar Dobriban§

Abstract

Feature learning is thought to be one of the fundamental reasons for the success of deep neural networks. It is rigorously known that in two-layer fully-connected neural networks under certain conditions, one step of gradient descent on the first layer can lead to feature learning; characterized by the appearance of a separated rank-one component—spike—in the spectrum of the feature matrix. However, with a constant gradient descent step size, this spike only carries information from the linear component of the target function and therefore learning non-linear components is impossible. We show that with a learning rate that grows with the sample size, such training in fact introduces multiple rank-one components, each corresponding to a specific polynomial feature. We further prove that the limiting large-dimensional and large sample training and test errors of the updated neural networks are fully characterized by these spikes. By precisely analyzing the improvement in the training and test errors, we demonstrate that these non-linear features can enhance learning.

1 Introduction

Learning non-linear features—or representations—from data is thought to be one of the fundamental reasons for the success of deep neural networks (see e.g., [BCV13, DKD16, YH21, SWL22, RBPB22], etc.). This has been observed in a wide range of domains, including computer vision and natural language processing. At the same time, the current theoretical understanding of feature learning is incomplete. In particular, among many theoretical approaches to study neural networks, much work has focused on two-layer fully-connected neural networks with a randomly generated, untrained first layer weights and a trained second layer—or *random features models* [RR07]. Despite their simplicity, random features models can capture various empirical properties of deep neural networks, and have been used to study generalization, overparametrization and “double descent”, adversarial robustness, transfer learning, estimation of out-of-distribution performance, and uncertainty quantification (see e.g., [MM22, HJ22, TAP21, LMH⁺23, BM23, CLKZ23, LD21, ALP22], etc.).

Nevertheless, feature learning is absent in random features models, because the first layer weights are assumed to be randomly generated, and then fixed. Although these models can represent non-linear functions of the data, in the commonly studied setting where the sample size, dimension,

*Equal Contribution.

†Department of Electrical and Systems Engineering, University of Pennsylvania.

‡Graduate Group in Applied Mathematics and Computational Science, University of Pennsylvania.

§Department of Statistics and Data Science, University of Pennsylvania.

{bemoniri, hassani}@seas.upenn.edu, dh7401@sas.upenn.edu, dobriban@wharton.upenn.edu.

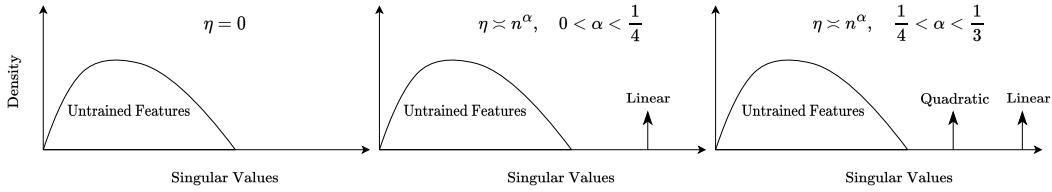


Figure 1: Spectrum of the updated feature matrix for different regimes of the gradient step size η . Spikes corresponding to monomial features are added to the spectrum of the initial matrix. The number of spikes depends on the range of α . See Theorems 3.2 and 3.3 for details.

and hidden layer size are proportional, under certain reasonable conditions they can only learn the *linear* component of the true model—or, teacher function—and other components of the teacher function effectively behave as Gaussian noise. Thus, in this setting, learning in a random features model is equivalent to learning in a *noisy linear model* with Gaussian features and Gaussian noise. This property is known as the *Gaussian equivalence property* (see e.g., [GLR⁺22](#), [ALP22](#), [AP20a](#), [MM22](#), [MS22](#), [HL23](#)). While other models such as the neural tangent kernel [JGH18](#), [DZPS19](#) can be more expressive, they also lack feature learning.

To bridge the gap between random features models and feature learning, several recent approaches have shown provable feature learning for neural networks under certain conditions; see Section 1.1 for details. In particular, the recent pioneering work of [BES⁺22](#) analyzed two-layer neural networks, trained with one gradient step on the first layer. They showed that when the step size is small, after one gradient step, the resulting two-layer neural network can learn linear features. However, it still behaves as a noisy linear model and does not capture non-linear components of a teacher function. Moreover, they showed that for a sufficiently large step size, under certain conditions, the one-step updated random features model can outperform linear and kernel predictors. However, the effects of a large gradient step size on the features is unknown. Moreover, what happens in the intermediate step size regime also remains unexplored. In this paper, we focus on the following key questions in this area:

What nonlinear features are learned by a two-layer neural network after one gradient update? How are these features reflected in the singular values and vectors of the feature matrix, and how does this depend on the scaling of the step size? What is the improvement in the training and test errors due to the nonlinear features learned?

Main Contributions. Toward answering the above questions, we make the following contributions:

- We study feature learning in two-layer neural networks. Specifically, we follow the training procedure introduced in [DLS22](#), [BES⁺22](#) where one step of gradient descent with step size η is applied to the first layer weights, and the second layer weights are found by solving ridge regression on the updated features. We consider a step size $\eta \asymp n^\alpha$, $\alpha \in (0, \frac{1}{2})$ that grows with the sample size n and examine how the learned features change with α (Section 2.1).
- In Section 3, we present a spectral analysis of the updated feature matrix. We first show that the spectrum of the feature matrix undergoes phase transitions depending on the range of α .

In particular, we find that if $\alpha \in (\frac{\ell-1}{2\ell}, \frac{\ell}{2\ell+2})$ for some $\ell \in \{1, 2, \dots\}$, then ℓ separated singular values—*spikes*—will be added to the spectrum of the initial feature matrix (Theorem 3.2). Figure 1 illustrates this finding.

- Building on perturbation theory for singular vectors, we argue that the left singular vectors (principal components) associated with the ℓ spikes are asymptotically aligned with polynomial features of different degrees (Theorem 3.3). In other words, the updated feature matrix will contain information about the degree- ℓ polynomial component of the target function.
- The Gaussian equivalence property [HL23, MM22], an essential tool to analyze random features models, fails after a gradient update with a large step size η . To overcome this difficulty, we establish equivalence theorems (Theorem 4.1 and 4.2) stating that the trained features \mathbf{F} can be replaced by sum of the untrained features \mathbf{F}_0 and ℓ spikes without changing the training and test errors. Then, by applying the Gaussian equivalence to the untrained component \mathbf{F}_0 , we provide a precise characterization of the training and test errors in the high-dimensional proportional regime (Theorem 4.3 and 4.5).
- From the derived results, we show that in the simple case where $\ell = 1$, the neural network does not learn non-linear functions. However, in the $\ell = 2$ regime, the neural network in fact learns quadratic components of the target function.

1.1 Related Works

Theory of shallow neural networks. Random features models [RR07] have been used to study various aspects of deep learning, such as generalization [MM22, ALP22, LD21, MP21], adversarial robustness [HJ22, BKM23], transfer learning [TAP21], out-of-distribution performance estimation [LMH⁺23], uncertainty quantification [CLKZ23], stability, and privacy [BM23]. This line of work builds upon nonlinear random matrix theory (see e.g., [PW17, LLC18, FW20, BP21], etc.) studying the spectrum of the feature matrix of two-layer neural networks at initialization.

Two-layer neural networks have been studied extensively in the mean-field regime (see e.g., [CB18, MMN18, MMM19, SS20, RVE22], etc.), and the neural tangent kernel (NTK) regime (see e.g., [JGH18, LXS⁺19, HY20], etc.). However, these results often require the neural net to have an extremely large width. In particular, in the NTK regime, this large width will result in features not evolving over the course of training and the model behaves similar to classic kernel methods. [GMMM21] show that for NTKs and other kernel methods, with a sample size linear in size of the input, non-linear functions cannot be learned. See also [Mis22, XHM⁺22, LY22]. Perturbative corrections to the large-width regime to capture feature learning have also been studied in the literature (see e.g., [Yai20, HN19, SNR23, NR21], etc.). See Section A for more discussion on related work in deep learning theory.

Feature learning. The problem of feature learning has been gaining a lot of attention recently. [DLS22] study the problem of learning polynomials with only a few relevant directions and show a sample complexity improvement over kernel methods. [CHS⁺23] extend these results and analyze multi-task feature learning in two-layer ReLU networks. [WES⁺22] empirically show that if learning rate is sufficiently large, an outlier in the spectrum of the weight and feature matrix emerges with the corresponding singular vector aligned to the structure of the training data. [NDL23, WNL23]

provide theoretical evidence that three-layer neural networks have provably richer feature learning capabilities than their two-layer counterparts.

Recently, [BES⁺22] show that in two-layer neural networks, when the dimension, sample size and hidden layer size are proportional, one gradient step with a constant step size on the first layer weights can lead to feature learning. However, non-linear components of a single-index target function are still not learned. They further show that with a sufficiently large step size, when the teacher function has a non-zero first Hermite coefficient, and under certain conditions, the updated neural networks can outperform linear and kernel methods. However, the precise effects of large gradient step sizes on learning nonlinear features, and their precise effects on the loss remain unexplored. [DKL⁺23] show that with a sample size proportional to the input dimension d , it is only possible to learn a single direction of multi-index teacher function using gradient updates on the first layer of a two-layer neural network. They also show that for single index models with information exponent (the index of the first non-zero Hermite coefficient) κ , there are hard directions whose learning requires a sample size of order $\Theta(d^\kappa)$. See also [BAGJ21].

High-dimensional asymptotics. We use tools developed in work on high-dimensional asymptotics, which dates back at least to the 1960s [Rau67, Dee70, Rau72]. Recently, these tools have been used in a wide range of areas such as wireless communications (e.g., [TV04, CD11], etc.), high-dimensional statistics (e.g., [RY04, Ser07, PA14, YBZ15, DW18], etc.), and machine learning (e.g., [GT90, Opp95, OK96, CL22, EVdB01], etc.). In particular, the spectrum of so-called information plus noise random matrices that arise in Gaussian equivalence results has been studied in [DS07, Péc19] and its spikes in [Cap14].

2 Preliminaries

Notation. We let $\mathbb{N} = \{1, 2, \dots\}$ be the set of positive integers. For a positive integer $d \geq 1$, we denote $[d] = \{1, \dots, d\}$. We use $O(\cdot)$ and $o(\cdot)$ for the standard big-O and little-o notation. For a matrix \mathbf{A} and a non-negative integer k , $\mathbf{A}^{\circ k} = \mathbf{A} \circ \mathbf{A} \circ \dots \circ \mathbf{A}$ is the matrix of the k -th powers of the elements of \mathbf{A} . For positive sequences $(A_n)_{n \geq 1}, (B_n)_{n \geq 1}$, we write $A_n = \Theta(B_n)$ or $A_n \asymp B_n$ or $A_n \equiv B_n$ if there is $C, C' > 0$ such that $CB_n \geq A_n \geq C'B_n$ for all n . We use $O_{\mathbb{P}}(\cdot)$, $o_{\mathbb{P}}(\cdot)$, and $\Theta_{\mathbb{P}}(\cdot)$ for the same notions holding in probability. The symbol \rightarrow_P denotes convergence in probability.

2.1 Problem Setting

In this paper, we study a supervised learning problem with training data $(\mathbf{x}_i, y_i) \in \mathbb{R}^d \times \mathbb{R}$, for $i \in [2n]$, where d is the feature dimension and $n \geq 2$ is the sample size. We assume that the data is generated according to

$$\mathbf{x}_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \mathbf{I}_d), \text{ and } y_i = f_{\star}(\mathbf{x}_i) + \varepsilon_i, \quad (1)$$

in which f_{\star} is the ground truth or *teacher function*, and $\varepsilon_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_{\varepsilon}^2)$ is additive noise.

We fit a model to the data in order to predict outcomes for unlabeled examples at test time; using a two-layer neural network. We let the width of the internal layer be $N \in \mathbb{N}$. For a weight matrix $\mathbf{W}_{\text{NN}} \in \mathbb{R}^{N \times d}$, an activation function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ applied element-wise, and the weights $\mathbf{a}_{\text{NN}} \in \mathbb{R}^N$ of a linear layer, we define the two-layer neural network as $f_{\mathbf{W}_{\text{NN}}, \mathbf{a}_{\text{NN}}}(\mathbf{x}) = \mathbf{a}_{\text{NN}}^{\top} \sigma(\mathbf{W}_{\text{NN}} \mathbf{x})$.

Following [DLS22, BES⁺22], for the convenience of the theoretical analysis, we split the training data into two parts: $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]^\top \in \mathbb{R}^{n \times d}$, $\mathbf{y} = (y_1, \dots, y_n)^\top \in \mathbb{R}^n$ and $\tilde{\mathbf{X}} = [\mathbf{x}_{n+1}, \dots, \mathbf{x}_{2n}]^\top \in \mathbb{R}^{n \times d}$, $\tilde{\mathbf{y}} = (y_{n+1}, \dots, y_{2n})^\top \in \mathbb{R}^n$.

We train the two layer neural network as follows. First, we initialize \mathbf{a}_{NN} with $\mathbf{a} = (a_1, \dots, a_N)^\top$ where $a_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1/N)$ and initialize \mathbf{W}_{NN} with $\mathbf{W}_0 = [\mathbf{w}_{0,1}, \dots, \mathbf{w}_{0,N}]^\top \in \mathbb{R}^{N \times d}$, $\mathbf{w}_{0,i} \stackrel{\text{i.i.d.}}{\sim} \text{Unif}(\mathbb{S}^{d-1})$ where \mathbb{S}^{d-1} is the unit sphere in \mathbb{R}^d and $\text{Unif}(\mathbb{S}^{d-1})$ is the uniform measure over it. Although we choose this initialization for a simpler analysis, many arguments can be shown to hold if we switch from the uniform distribution over the sphere to a Gaussian; for example, see Section O.5. Fixing \mathbf{a}_{NN} at initialization \mathbf{a} , we perform *one step of gradient descent* on \mathbf{W}_{NN} with respect to the squared loss computed on (\mathbf{X}, \mathbf{y}) . Recalling that \circ denotes element-wise multiplication, the negative gradient can be written as

$$\mathbf{G} := -\frac{\partial}{\partial \mathbf{W}_{\text{NN}}} \left[\frac{1}{2n} \left\| \mathbf{y} - \sigma(\mathbf{X} \mathbf{W}_{\text{NN}}^\top) \mathbf{a} \right\|_2^2 \right]_{\mathbf{W}_{\text{NN}}=\mathbf{W}_0} = \frac{1}{n} \left[(\mathbf{a} \mathbf{y}^\top - \mathbf{a} \mathbf{a}^\top \sigma(\mathbf{W}_0 \mathbf{X}^\top)) \circ \sigma'(\mathbf{W}_0 \mathbf{X}^\top) \right] \mathbf{X},$$

and the one-step update is $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_N]^\top = \mathbf{W}_0 + \eta \mathbf{G}$ for a *learning rate* or *step size* η .

After the update on \mathbf{W}_{NN} , we perform ridge regression on \mathbf{a} using $(\tilde{\mathbf{X}}, \tilde{\mathbf{y}})$. Let $\mathbf{F} = \sigma(\tilde{\mathbf{X}} \mathbf{W}^\top) \in \mathbb{R}^{n \times N}$ be the feature matrix after the one-step update. For a regularization parameter $\lambda > 0$, we set

$$\hat{\mathbf{a}} = \hat{\mathbf{a}}(\mathbf{F}) = \arg \min_{\mathbf{a} \in \mathbb{R}^N} \frac{1}{n} \|\tilde{\mathbf{y}} - \mathbf{F} \mathbf{a}\|_2^2 + \lambda \|\mathbf{a}\|_2^2 = \left(\mathbf{F}^\top \mathbf{F} + \lambda n \mathbf{I}_N \right)^{-1} \mathbf{F}^\top \tilde{\mathbf{y}}. \quad (2)$$

Then, for a test datapoint with features \mathbf{x} , we predict the outcome $\hat{y} = f_{\mathbf{W}, \hat{\mathbf{a}}}(\mathbf{x}) = \hat{\mathbf{a}}^\top \sigma(\mathbf{W} \mathbf{x})$.

2.2 Conditions

Our theoretical analysis applies under the following conditions:

Condition 2.1 (Asymptotic setting). *We assume that the sample size n , dimension d , and width of hidden layer N all tend to infinity with*

$$d/n \rightarrow \phi > 0, \quad \text{and} \quad d/N \rightarrow \psi > 0.$$

We further consider the following model for the teacher function, leading to a single-index model.

Condition 2.2. *We let $f_\star : \mathbb{R}^d \rightarrow \mathbb{R}$ be $f_\star(\mathbf{x}) = \sigma_\star(\mathbf{x}^\top \boldsymbol{\beta}_\star)$ for all \mathbf{x} , where $\boldsymbol{\beta}_\star \in \mathbb{R}^d$ is an unknown parameter with $\boldsymbol{\beta}_\star \sim \mathcal{N}(0, \frac{1}{d} \mathbf{I}_d)$ and $\sigma_\star : \mathbb{R} \rightarrow \mathbb{R}$ is a $\Theta(1)$ -Lipschitz teacher activation function.*

This condition is in line with prior work (see e.g., [BES⁺22, HL23, GLR⁺22], etc.). Recently, [DKL⁺23] showed that under Condition 2.1, two-layer neural networks trained with one gradient descent step can only learn a single-index approximation of a multi-index model. This shows that when studying a single-step update, 2.2 is not restrictive.

We let H_k , $k \geq 1$ be the (probabilist's) Hermite polynomials on \mathbb{R} defined by

$$H_k(x) = (-1)^k \exp(x^2/2) \frac{d^k}{dx^k} \exp(-x^2/2),$$

for any $x \in \mathbb{R}$. These polynomials form an orthogonal basis in the Hilbert space L^2 of measurable functions $f : \mathbb{R} \rightarrow \mathbb{R}$ such that $\int f^2(x) \exp(-x^2/2) dx < \infty$ with inner product $\langle f, g \rangle = \int f(x)g(x) \exp(-x^2/2) dx$. The first few Hermite polynomials are $H_0(x) = 1$, $H_1(x) = x$, and $H_2(x) = x^2 - 1$.

Condition 2.3. *The activation function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ has the following Hermite expansion in L^2 :*

$$\sigma(z) = \sum_{k=1}^{\infty} c_k H_k(z), \quad c_k = \frac{1}{k!} \mathbb{E}_{Z \sim \mathcal{N}(0,1)} [\sigma(Z) H_k(Z)].$$

The coefficients satisfy $c_1 \neq 0$ and $c_k^2 k! \leq C k^{-\frac{3}{2}-\omega}$ for some $C, \omega > 0$ and for all $k \geq 1$. Moreover, the first three derivatives of σ exist almost surely, and are bounded.

Note that in this paper, unlike [HL23], we do not require the activation function to be odd. The reason is that here, unlike [HL23], we do not analyze the problem for a general loss function and use a proof technique specialized for squared loss. We remark that the above condition requires $c_0 = 0$, i.e., that $\mathbb{E}\sigma(Z) = 0$ for $Z \sim \mathcal{N}(0, 1)$. This condition is in line with prior work in the area (e.g., [AP20a, BES⁺22], etc.), and could be removed at the expense of more complicated formulas and theoretical analysis. The smoothness assumption on σ is also in line with prior work in the area (see e.g., [HL23, BES⁺22], etc.). Note that the above condition is satisfied by many popular activation functions (after shifting) such as the ReLU $\sigma(x) = \max\{x, 0\} - \frac{1}{\sqrt{2\pi}}$, hyperbolic tangent $\sigma(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}}$, and sigmoid $\sigma(x) = \frac{1}{1+e^{-x}} - \frac{1}{2}$, for all x .

We also make similar assumptions on the teacher activation:

Condition 2.4. *The teacher activation $\sigma_{\star} : \mathbb{R} \rightarrow \mathbb{R}$ has the following Hermite expansion in L^2 :*

$$\sigma_{\star}(z) = \sum_{k=1}^{\infty} c_{\star,k} H_k(z), \quad c_{\star,k} = \frac{1}{k!} \mathbb{E}_Z [\sigma_{\star}(Z) H_k(Z)],$$

with $Z \sim \mathcal{N}(0, 1)$. Also, we define $c_{\star} = (\sum_{k=1}^{\infty} k! c_{\star,k}^2)^{\frac{1}{2}}$.

3 Analysis of the Feature Matrix

The first step in analyzing the spectrum of the feature matrix \mathbf{F} is to study the negative gradient \mathbf{G} . It is shown in [BES⁺22, Proposition 2] that in operator norm, the matrix \mathbf{G} can be approximated by the rank-one matrix $c_1 \mathbf{a} \boldsymbol{\beta}^{\top}$ with high probability, where the Hermite coefficient c_1 of the activation σ is defined in Condition 2.3, and $\boldsymbol{\beta} = \frac{1}{n} \mathbf{X}^{\top} \mathbf{y} \in \mathbb{R}^d$. Moreover, under Conditions 2.1–2.4, [BES⁺22] show that $\boldsymbol{\beta}$ can be understood as a noisy estimate of $\boldsymbol{\beta}_{\star}$, namely

$$\frac{|\boldsymbol{\beta}_{\star}^{\top} \boldsymbol{\beta}|}{\|\boldsymbol{\beta}_{\star}\|_2 \|\boldsymbol{\beta}\|_2} \xrightarrow{P} \frac{|c_{\star,1}|}{\sqrt{c_{\star,1}^2 + \phi(c_{\star}^2 + \sigma_{\varepsilon}^2)}}. \quad (3)$$

See also Lemma K.1. In particular, if the sample size used for the gradient update is very large; i.e., $\phi \rightarrow 0$, $\boldsymbol{\beta}$ will converge to being completely aligned to $\boldsymbol{\beta}_{\star}$.

Building on this result, we can prove the following rank-one approximation lemma. Note that the updated feature matrix can be written as $\mathbf{F} = \sigma(\tilde{\mathbf{X}}(\mathbf{W}_0 + \eta \mathbf{G})^{\top})$ and terms of the form $(\tilde{\mathbf{X}} \mathbf{G}^{\top})^{\circ k}$, $k \in \mathbb{N}$, will appear in polynomial and Taylor expansions of \mathbf{F} around \mathbf{F}_0 . In the following lemma, we show that for any fixed power k , these terms can be approximated by rank one terms.

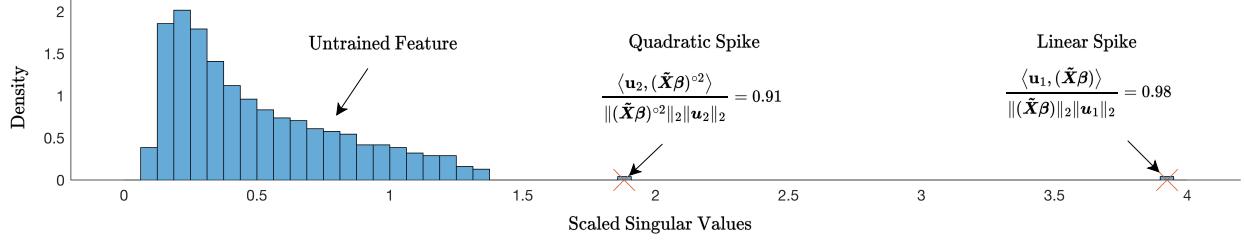


Figure 2: Histogram of the scaled singular values (divided by \sqrt{n}) of the feature matrix $\mathbf{F} = \sigma(\tilde{\mathbf{X}}\mathbf{W}^\top)$ after the update with step size $\eta = n^{0.29}$ ($\ell = 2$). In this regime, two isolated spikes appear in the spectrum as stated in Theorem 3.2. The top two left singular vectors \mathbf{u}_1 and \mathbf{u}_2 are aligned with $\tilde{\mathbf{X}}\beta$ and $(\tilde{\mathbf{X}}\beta)^{\circ 2}$, respectively. See Section 5 for the simulation details.

Lemma 3.1 (Rank-one approximation). *If Conditions 2.1–2.4 hold, then there exists $C > 0$ such that for c_1 from Condition 2.3, for any fixed $k \in \mathbb{N}$,*

$$\|(\tilde{\mathbf{X}}\mathbf{G}^\top)^{\circ k} - c_1^k (\tilde{\mathbf{X}}\beta)^{\circ k} (\mathbf{a}^{\circ k})^\top\|_{\text{op}} \leq C^k n^{-\frac{k}{2}} \log^{2k} n$$

with probability $1 - o(1)$.

Next, we will show that after the gradient step, the spectrum of the feature matrix \mathbf{F} will consist of a bulk of singular values that stick close together—given by the spectrum of the initial feature matrix $\mathbf{F}_0 = \sigma(\tilde{\mathbf{X}}\mathbf{W}_0^\top)$ —and ℓ separated spikes¹, where ℓ is an integer that depends on the step size used in the gradient update. Specifically, when the step size is $\eta \asymp n^\alpha$ with $\frac{\ell-1}{2\ell} < \alpha < \frac{\ell}{2\ell+2}$ for some $\ell \in \mathbb{N}$, the feature matrix \mathbf{F} can be approximated in operator norm by the untrained features $\mathbf{F}_0 = \sigma(\tilde{\mathbf{X}}\mathbf{W}_0^\top)$ plus ℓ rank-one terms, where the left singular vectors of the rank-one terms are aligned with the non-linear features $\tilde{\mathbf{X}} \mapsto (\tilde{\mathbf{X}}\beta)^{\circ k}$, for $k \in [\ell]$. See Figure 2.

Theorem 3.2 (Spectrum of feature matrix). *Let $\eta \asymp n^\alpha$ with $\frac{\ell-1}{2\ell} < \alpha < \frac{\ell}{2\ell+2}$ for some $\ell \in \mathbb{N}$. If Conditions 2.1–2.4 hold, then for c_k from Condition 2.3 and $\mathbf{F}_0 = \sigma(\tilde{\mathbf{X}}\mathbf{W}_0^\top)$,*

$$\mathbf{F} = \mathbf{F}_\ell + \Delta, \quad \text{with} \quad \mathbf{F}_\ell := \mathbf{F}_0 + \sum_{k=1}^{\ell} c_1^k c_k \eta^k (\tilde{\mathbf{X}}\beta)^{\circ k} (\mathbf{a}^{\circ k})^\top, \quad (4)$$

where $\|\Delta\|_{\text{op}} = o(\sqrt{n})$ with probability $1 - o(1)$.

To understand $(\tilde{\mathbf{X}}\beta)^{\circ k} (\mathbf{a}^{\circ k})^\top$, notice that for a datapoint with features $\tilde{\mathbf{x}}_i$, the activation of each neuron is proportional to the polynomial feature $(\tilde{\mathbf{x}}_i^\top \beta)^k$, with coefficients given by $\mathbf{a}^{\circ k}$ for the neurons. The spectrum of the initial feature matrix \mathbf{F}_0 is fully characterized in [PW17, BP21, BP22, LLC18, FW20], and its operator norm is known to be $\Theta_{\mathbb{P}}(\sqrt{n})$. Moreover, it follows from the proof that the operator norm of each of the terms $c_1^k c_k \eta^k (\tilde{\mathbf{X}}\beta)^{\circ k} (\mathbf{a}^{\circ k})^\top$, $k \in [\ell]$ is with high probability of order larger than \sqrt{n} . Thus, Theorem 3.2 identifies the spikes in the spectrum of the feature matrix.

Proof Idea. We approximate the feature matrix $\mathbf{F} = \sigma(\tilde{\mathbf{X}}(\mathbf{W}_0 + \eta\mathbf{G})^\top)$ by a polynomial using its Hermite expansion. Next, we use the binomial expansion and apply Lemma 3.1 to approximate

¹Using terminology from random matrix theory [BS10, YBZ15].

$(\tilde{\mathbf{X}}\mathbf{G}^\top)^{\circ k}$ by $c_1^k(\tilde{\mathbf{X}}\boldsymbol{\beta})^{\circ k}(\mathbf{a}^{\circ k})^\top$, for all k . Then, spike terms with $k \geq \ell + 1$ are negligible since we can show that their norm is $O_{\mathbb{P}}(n^{k\alpha+\frac{1}{2}-\frac{k-1}{2}}) = o_{\mathbb{P}}(\sqrt{n})$.

The special case where $\alpha = 0$ is discussed in [BES⁺22, Section 3], which focuses on the spectrum of the updated weight matrix $\mathbf{W} = \mathbf{W}_0 + \eta\mathbf{G}$. However, here we study the updated feature matrix $\mathbf{F} = \sigma(\tilde{\mathbf{X}}(\mathbf{W}_0 + \eta\mathbf{G})^\top)$ because that is more directly related to the learning problem—as we will discuss in the consequences for the training and test errors below.

In the following theorem, we argue that the subspace spanned by the non-linear features $\{\sigma(\tilde{\mathbf{X}}\mathbf{w}_i)\}_{i \in [N]}$ can be approximated by the subspace spanned by the monomials $\{(\tilde{\mathbf{X}}\boldsymbol{\beta})^{\circ k}\}_{k \in [\ell]}$. For two ℓ -dimensional subspaces $\mathcal{U}_1, \mathcal{U}_2 \subseteq \mathbb{R}^n$, with orthonormal bases $\mathbf{U}_1, \mathbf{U}_2 \in \mathbb{R}^{n \times \ell}$, recall the principal angle distance between $\mathcal{U}_1, \mathcal{U}_2$ defined by $d(\mathcal{U}_1, \mathcal{U}_2) = \min_{\mathbf{Q}} \|\mathbf{U}_1 - \mathbf{U}_2\mathbf{Q}\|_{\text{op}}$, where the minimum is over $\ell \times \ell$ orthogonal matrices [SS90]. This definition is invariant to the choice of $\mathbf{U}_1, \mathbf{U}_2$.

Theorem 3.3. *Let \mathcal{F}_ℓ be the ℓ -dimensional subspace of \mathbb{R}^n spanned by top- ℓ left singular vectors (principal components) of \mathbf{F} . Under the conditions of Theorem 3.2, we have*

$$d(\mathcal{F}_\ell, \text{span}\{(\tilde{\mathbf{X}}\boldsymbol{\beta})^{\circ k}\}_{k \in [\ell]}) \rightarrow_P 0.$$

This result shows that after one step of gradient descent with step size $\eta \asymp n^\alpha$ with $\frac{\ell-1}{2\ell} < \alpha < \frac{\ell}{2\ell+2}$, the subspace of the top- ℓ left singular vectors carries information from the polynomials $\{(\tilde{\mathbf{X}}\boldsymbol{\beta})^{\circ k}\}_{k \in [\ell]}$. Also, recall that by equation 3, the vector $\boldsymbol{\beta}$ is aligned with $\boldsymbol{\beta}_*$. Hence, it is shown that \mathcal{F}_ℓ carries information from the first ℓ polynomial components of the teacher function.

Proof Idea. We use Wedin’s theorem [Wed72] to characterize the distance between the left singular vector space of $\sum_{k=1}^{\ell} c_1^k c_k \eta^k (\tilde{\mathbf{X}}\boldsymbol{\beta})^{\circ k} (\mathbf{a}^{\circ k})^\top$ and that of \mathbf{F} . Here, we consider the matrix $\mathbf{F}_0 + \Delta$ as the perturbation term.

4 Learning Higher-Degree Polynomials

In the previous section, we studied the feature matrix \mathbf{F} and showed that when $\eta \asymp n^\alpha$ with $\frac{\ell-1}{2\ell} < \alpha < \frac{\ell}{2\ell+2}$, it can be approximated by $\mathbf{F}_0 = \sigma(\tilde{\mathbf{X}}\mathbf{W}_0^\top)$ plus ℓ rank-one or spike terms. We also saw that the left singular vectors of the spike terms are aligned with the non-linear functions $\tilde{\mathbf{X}} \mapsto (\tilde{\mathbf{X}}\boldsymbol{\beta})^{\circ k}$. Intuitively, this result suggests that after the gradient update, the trained weights are becoming aligned with the teacher model and we should expect the ridge regression estimator on the learned features to achieve better performance. In particular, when $\alpha > 0$, we expect the ridge regression estimator to capture the non-linear part of the teacher function. This is impossible for $\eta = O(1)$ [BES⁺22] or $\eta = 0$ [HL23, MM22].

In this section, we aim to make this intuition rigorous and show that the spikes in the feature matrix lead to a decrease in the error achieved by the estimator. Moreover, for large enough step sizes, the model can learn non-linear components of the teacher function. For this, we first need to prove *equivalence theorems* showing that instead of the true feature matrix \mathbf{F} , the approximations from Theorem 3.2 can be used to compute error terms (i.e., the effect of Δ on the error is negligible).

4.1 Equivalence Theorems

The Gaussian equivalence property ([GLR⁺22], [HL23], [MS22], etc.) implies that the training and test errors of a random features model are asymptotically the same as that of a noisy linear model. In other words, the limiting behavior of these quantities is unchanged if we replace the untrained feature matrix $\mathbf{F}_0 = \sigma(\tilde{\mathbf{X}}\mathbf{W}_0^\top)$ with $\mathbf{F}_0 = c_1\tilde{\mathbf{X}}\mathbf{W}_0^\top + c_{>1}\mathbf{Z}$, where $\mathbf{Z} \in \mathbb{R}^{n \times N}$ is an independent random matrix with i.i.d. $\mathcal{N}(0, 1)$ entries.

This property has been used extensively in work on random features models (see e.g., [AP20b], [AP20a], [TAP21], [MP21], etc.) as it provides a powerful tool to analyze non-linear random matrices. However, the Gaussian equivalence property fails when the weight matrix \mathbf{W} is updated with a large gradient descent step [BES⁺22], posing a significant challenge to the analysis.

In this section, we first prove that we can replace the trained features \mathbf{F} with their approximation \mathbf{F}_ℓ from Theorem 3.2 in terms of \mathbf{F}_0 and spikes, without changing the limiting training and test errors. Then, in the next sections we will see that the training and test errors can be derived by applying the Gaussian equivalence property to the untrained features \mathbf{F}_0 only.

Given a regularization parameter $\lambda > 0$, recalling the ridge estimator $\hat{\mathbf{a}}(\mathbf{F})$ from equation 2, we define the training loss

$$\mathcal{L}_{\text{tr}}(\mathbf{F}) = \frac{1}{n} \|\tilde{\mathbf{y}} - \mathbf{F}\hat{\mathbf{a}}(\mathbf{F})\|_2^2 + \lambda \|\hat{\mathbf{a}}(\mathbf{F})\|_2^2.$$

In the next theorem, we show that when $\eta \asymp n^\alpha$ with $\frac{\ell-1}{2\ell} < \alpha < \frac{\ell}{2\ell+2}$, the training loss $\mathcal{L}_{\text{tr}}(\mathbf{F})$ can be approximated with negligible error by $\mathcal{L}_{\text{tr}}(\mathbf{F}_\ell)$. In other words, the approximation of the feature matrix from Theorem 3.2 can be used to derive the asymptotics of the training loss.

Theorem 4.1 (Training loss equivalence). *Let $\eta \asymp n^\alpha$ with $\frac{\ell-1}{2\ell} < \alpha < \frac{\ell}{2\ell+2}$ for some $\ell \in \mathbb{N}$ and recall \mathbf{F}_ℓ from equation 4. If Conditions 2.1–2.4 hold, then for any fixed $\lambda > 0$, with probability $1 - o(1)$ we have*

$$\mathcal{L}_{\text{tr}}(\mathbf{F}) - \mathcal{L}_{\text{tr}}(\mathbf{F}_\ell) = o(1).$$

Similar equivalence results can also be proved for the test error, i.e., the average test loss. For any $\mathbf{a} \in \mathbb{R}^N$, we define the test error of \mathbf{a} as $\mathcal{L}_{\text{te}}(\mathbf{a}) = \mathbb{E}_{\mathbf{f}, y} (y - \mathbf{f}^\top \mathbf{a})^2$, in which the expectation is taken over (\mathbf{x}, y) where $\mathbf{f} = \sigma(\mathbf{W}\mathbf{x})$ with $\mathbf{x} \sim \mathcal{N}(0, \mathbf{I}_d)$ and $y = f_*(\mathbf{x}) + \varepsilon$ with $\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2)$. The next theorem shows that one can also use the approximation of the feature matrix from Theorem 3.2 to derive the asymptotics of the test error.

Theorem 4.2 (Test error equivalence). *Let $\eta \asymp n^\alpha$ with $\frac{\ell-1}{2\ell} < \alpha < \frac{\ell}{2\ell+2}$ for some $\ell \in \mathbb{N}$, and \mathbf{F}_ℓ be defined as in equation 4. If Conditions 2.1–2.4 hold, then for any $\lambda > 0$, if $\mathcal{L}_{\text{te}}(\hat{\mathbf{a}}(\mathbf{F})) \rightarrow_P \mathcal{L}_{\mathbf{F}}$ and $\mathcal{L}_{\text{te}}(\hat{\mathbf{a}}(\mathbf{F}_\ell)) \rightarrow_P \mathcal{L}_{\mathbf{F}_\ell}$, we have $\mathcal{L}_{\mathbf{F}} = \mathcal{L}_{\mathbf{F}_\ell}$.*

Proof Idea. To prove Theorem 4.2, we show first show that the norm of the trained second layer weight $\hat{\mathbf{a}}$, is $O_{\mathbb{P}}(1)$. Then, we use Theorem 3.2 to conclude the proof. To prove Theorem 4.2, we will use a *free-energy trick* [ASH19, [HL23], [HJ22]]. We first extend Theorem 4.1 and show that for any $\lambda, \zeta > 0$, the minima over \mathbf{a} of

$$\mathcal{R}_\zeta(\mathbf{a}, \bar{\mathbf{F}}) = \frac{1}{n} \|\tilde{\mathbf{y}} - \bar{\mathbf{F}}\mathbf{a}\|_2^2 + \lambda \|\mathbf{a}\|_2^2 + \zeta \mathcal{L}_{\text{te}}(\mathbf{a}),$$

for $\bar{\mathbf{F}} = \mathbf{F}$ and $\bar{\mathbf{F}} = \mathbf{F}_\ell$ are close. Then, we use this to argue that the limiting test errors are also close.

With Theorem 4.1 and 4.2 in hand, for $\eta \asymp n^\alpha$, we can use the approximation \mathbf{F}_ℓ —with the appropriate ℓ —of the feature matrix \mathbf{F} to analyze the training and test error.

4.2 Analysis of Training and Test Errors

In this section, we quantify the discrepancy between the training loss of the ridge estimator trained on the new—learned—feature matrix \mathbf{F} and the same ridge estimator trained on the untrained feature matrix \mathbf{F}_0 . We will do this for the step size $\eta \asymp n^\alpha$ with $\frac{\ell-1}{2\ell} < \alpha < \frac{\ell}{2\ell+2}$ for various $\ell \in \mathbb{N}$.

Our results depend on the limits of traces of the matrices $(\mathbf{F}_0 \mathbf{F}_0^\top + \lambda n \mathbf{I}_n)^{-1}$ and $\tilde{\mathbf{X}}^\top (\mathbf{F}_0 \mathbf{F}_0^\top + \lambda n \mathbf{I}_n)^{-1} \tilde{\mathbf{X}}$. These limits have been determined in [ALP22, AP20a], see also [PW17, Péc19], and depend on the values $m_1, m_2 > 0$, which are the unique solutions of the following system of coupled equations, for $\lambda > 0$:

$$\begin{aligned} \phi(m_1 - m_2)(c_{>1}^2 m_1 + c_1^2 m_2) + \Psi(m_1, m_2) &= 0, \\ \frac{\phi}{\psi}(c_1^2 m_1 m_2 + \phi(m_2 - m_1)) + \Psi(m_1, m_2) &= 0, \end{aligned} \quad (5)$$

where $\Psi(m_1, m_2) = c_1^2 m_1 m_2 (\lambda \psi m_1 / \phi - 1)$ and $c_{>1} = (\sum_{k=2}^{\infty} k! c_k^2)^{1/2}$. Here, m_1 is the limiting Stieltjes transform of the matrix $\frac{1}{n} \mathbf{F}_0 \mathbf{F}_0^\top$ and m_2 is an auxiliary transform. For instance, we leverage that under Condition 2.1, we have

$$\begin{aligned} \lim_{d, n, N \rightarrow \infty} \text{tr}(\tilde{\mathbf{X}}^\top (\mathbf{F}_0 \mathbf{F}_0^\top + \lambda n \mathbf{I}_n)^{-1} \tilde{\mathbf{X}}) / d &= \psi m_2 / \phi > 0, \\ \lim_{d, n, N \rightarrow \infty} \text{tr}((\mathbf{F}_0 \mathbf{F}_0^\top + \lambda n \mathbf{I}_n)^{-1}) &= \psi m_1 / \phi > 0. \end{aligned}$$

See Lemma K.4 and its proof for more details. For instance, as argued in [PW17, ALP22], these can be reduced to a quartic equation for m_1 and are convenient to solve numerically. However, the existence of these limits does not imply our results; on the contrary, the proofs of our results require extensive additional calculations and several novel ideas.

Theorem 4.3. *If Conditions 2.1–2.4 are satisfied, and we have $c_1, \dots, c_\ell \neq 0$, as well as $\eta \asymp n^\alpha$ with $\frac{\ell-1}{2\ell} < \alpha < \frac{\ell}{2\ell+2}$, then for the learned feature map \mathbf{F} and the untrained feature map \mathbf{F}_0 , we have*

$$\mathcal{L}_{\text{tr}}(\mathbf{F}_0) - \mathcal{L}_{\text{tr}}(\mathbf{F}) \rightarrow_P \Delta_\ell \geq 0,$$

where Δ_ℓ can be found in Section L.

The expression for Δ_ℓ is complex and given in Section L due to space limitations. For a better understanding of Theorem 4.3, we consider two specific cases, $\ell = 1$ and $\ell = 2$.

Corollary 4.4. *Under the assumptions of Theorem 4.3, for $\ell = 1$, we have $\mathcal{L}_{\text{tr}}(\mathbf{F}_0) - \mathcal{L}_{\text{tr}}(\mathbf{F}) \rightarrow_P \Delta_1$ with*

$$\Delta_1 := \frac{\psi \lambda c_{*,1}^4 m_2}{\phi [c_{*,1}^2 + \phi(c_*^2 + \sigma_\varepsilon^2)]} \geq 0. \quad (6)$$

For $\ell = 2$, we have $\mathcal{L}_{\text{tr}}(\mathbf{F}_0) - \mathcal{L}_{\text{tr}}(\mathbf{F}) \rightarrow_P \Delta_2$ with

$$\Delta_2 := \Delta_1 + \frac{4\psi \lambda c_{*,1}^4 c_{*,2}^2 m_1}{3\phi [\phi(c_*^2 + \sigma_\varepsilon^2) + c_{*,1}^2]^2} \geq 0. \quad (7)$$

The above result shows that after one gradient step with sufficiently large step size, the model can fit nonlinear components of the teacher function. This is impossible with a small step size. For example, when $\ell = 1$, the improvement in the loss is *increasing in the strength of the linear component* $c_{*,1}$, keeping the signal strength c_* fixed. This is not the case for the strength of the non-linear component $c_{*,>1}^2 = c_*^2 - c_{*,1}^2$.

When we further increase the step size to the $\ell = 2$ regime, the loss of the trained model will drop by an additional positive value, depending on the strength $c_{*,2}$ of the quadratic signal, which shows that the quadratic component of the target function is being fit. Also, note that if $c_{*,1} = 0$; i.e., if the *information exponent* (the index of the first non-zero Hermite coefficient) of σ_* is greater than one, the gradient step does not change the limiting loss. In this case, according to equation 3, the alignment between the learned direction β and the true direction β_* will converge to zero. It is known that learning single-index functions with information exponent greater than one requires a sample size of order larger than d [DKL+23].

The limiting value of the test error can be analyzed similarly.

Theorem 4.5. *Let Conditions 2.1–2.4 and the assumptions of Theorem 4.2 hold. If $c_1 \neq 0$, then for $\ell = 1$, we have $\mathcal{L}_{\text{te}}(\hat{\mathbf{a}}(\mathbf{F}_0)) - \mathcal{L}_{\text{te}}(\hat{\mathbf{a}}(\mathbf{F})) \rightarrow_P \Lambda_1$ with*

$$\Lambda_1 := \frac{c_{*,1}^4 \Gamma_1}{[c_{*,1}^2 + \phi(c_*^2 + \sigma_\varepsilon^2)]} \geq 0, \quad (8)$$

where Γ_1 does not depend on the target function.

If further $c_2 \neq 0$, then for $\ell = 2$, we have $\mathcal{L}_{\text{te}}(\hat{\mathbf{a}}(\mathbf{F}_0)) - \mathcal{L}_{\text{te}}(\hat{\mathbf{a}}(\mathbf{F})) \rightarrow_P \Lambda_2$ with

$$\Lambda_2 := \Lambda_1 + \frac{c_{*,1}^4 c_{*,2}^2 \Gamma_2}{[c_{*,1}^2 + \phi(c_*^2 + \sigma_\varepsilon^2)]^2} \geq 0, \quad (9)$$

where Γ_2 does not depend on the target function. The complete expressions for Γ_1 and Γ_2 can be found in equation 65 and equation 88, respectively.

Similar to the training loss result, when $\ell = 1$, the improvement in the test error is increasing in $c_{*,1}$; keeping the signal strength c_* fixed. Moreover, the improvement in the test error for $\ell = 2$ depends on the strength $c_{*,2}$ of the quadratic signal, showing that the nonlinear component is being learned.

Proof Idea. Using Theorem 4.1 and 4.2 in the expression for the training and test error, we replace the trained features matrix \mathbf{F} with its approximation \mathbf{F}_ℓ from Theorem 3.2. Then, by applying the Woodbury formula, we express the training and test errors in terms of $\bar{\mathbf{R}}_0 = (\mathbf{F}_0 \mathbf{F}_0^\top + \lambda n \mathbf{I}_n)^{-1}$ and the non-linear spikes from Theorem 3.2. Using the Gaussian equivalence property (Appendix J) for the untrained features \mathbf{F}_0 , we show that the interaction between the first ℓ Hermite components of $\tilde{\mathbf{y}}$ and the spike terms will result in non-vanishing terms corresponding to learning different components of the target function. Finally, we compute the limiting value of these terms in terms of m_1, m_2 and their derivatives using tools from random matrix theory.

4.3 Staircase Property

Recently, [ABAB+21, AAM22] show that when learning Boolean functions, under certain conditions on the teacher function, a two-layer trained by SGD in the mean-field regime will learn the target

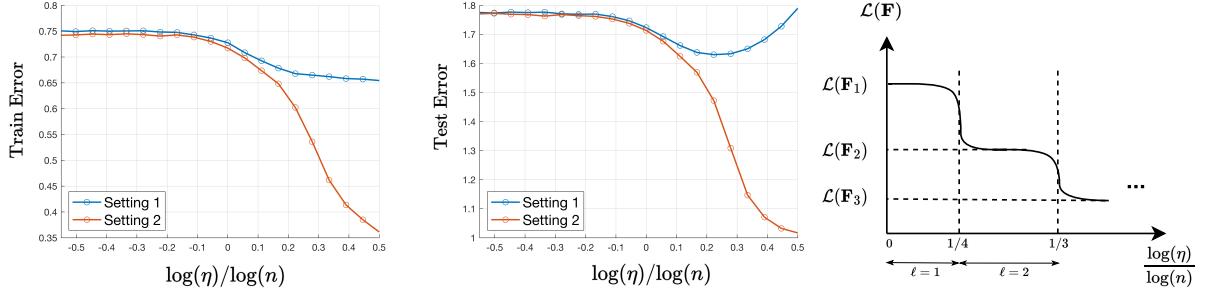


Figure 3: **(Left, Middle)** Training and test errors after one gradient as functions of $\log(\eta)/\log(n)$. **(Right)** A toy plot illustrating the theoretical training/test error curve as a function of $\log(\eta)/\log(n)$.

function incrementally; i.e., Fourier coefficients of higher order are sequentially learned over time. [\[BMZ23\]](#) study the problem of learning a single-index function using a wide two-layer neural network trained using gradient flow, and show that in a specific training setting where the stepsizes for the first layer are much smaller than those of the second layer, the decrease rate of training error is non-monotone; there are long plateaus where there is barely any progress, and there are intervals of rapid decrease.

In the one-step updated two layer neural network, we observe a similar phenomenon. Theorems [4.5](#) and Theorem [L.1](#) show that given $\ell \in \mathbb{N}$, the errors of the trained model is asymptotically constant for all $\eta = cn^\alpha$ with $\frac{\ell-1}{2\ell} < \alpha < \frac{\ell}{2\ell+2}$ and $c \in \mathbb{R}$. There are sharp jumps at the edges between regimes of α , whose size is precisely characterized above. This shows that a non-monotone rate of decrease in training and test error can also be seen after one step of gradient descent, as a function of step size. For an illustration of this phenomenon, see Figure 3 **(Right)**.

5 Numerical Simulations

To support and illustrate our theoretical results, we present some numerical simulations. We use the shifted ReLU activation $\sigma(x) = \max(x, 0) - 1/\sqrt{2\pi}$, $n = 1000$, $N = 500$, $d = 300$, and the regularization parameter $\lambda = 0.01$.

Singular Value Spectrum of \mathbf{F} . We let the teacher function be $f_\star(\mathbf{x}) = H_1(\beta_\star^\top \mathbf{x}) + H_2(\beta_\star^\top \mathbf{x})$, set the noise variance $\sigma_\varepsilon^2 = 0.5$, and the step size to $\eta = n^{0.29}$, so $\ell = 2$. We plot the histogram of singular values of the updated feature matrix \mathbf{F} . In Figure 2, we see two spikes corresponding to $\tilde{\mathbf{X}}\beta$, $(\tilde{\mathbf{X}}\beta)^{\circ 2}$ as suggested by Theorem [3.2](#) and [3.3](#). Since f_\star has a linear component H_1 and a quadratic component H_2 , these spikes will lead to feature learning.

Quadratic Feature Learning. To support the findings of Corollary [4.4](#) and Theorem [4.5](#) for $\ell = 2$, we consider the following two settings:

$$\begin{aligned} \text{Setting 1} : y &= H_1(\beta_\star^\top \mathbf{x}) + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, 1), \\ \text{Setting 2} : y &= H_1(\beta_\star^\top \mathbf{x}) + \frac{1}{\sqrt{2}} H_2(\beta_\star^\top \mathbf{x}). \end{aligned}$$

Note that $c_{*,1}$ and c_* + σ_ε^2 are same in these two settings. This ensures that the improvement due to learning the linear component is the same. We plot the training and test errors of the two-layer neural networks trained with the procedure described in Section 2.1 as functions of $\log(\eta)/\log(n)$. In Figure 3 (Left), we see that the errors decrease in the range $\log(\eta)/\log(n) \in (0, \frac{1}{4})$ as the model learns the linear component $H_1(\beta_\star^\top \mathbf{x})$. In the range $\log(\eta)/\log(n) \in (\frac{1}{4}, \frac{1}{3})$, the model starts to learn the quadratic feature. However since the quadratic feature is not present in Setting 1, the errors under the two settings diverge. Although the proofs reveal that the convergence rates of the training/test errors after one step can be slow, these results are consistent with Corollary 4.4 and Theorem 4.5.

6 Conclusion

We have studied feature learning in two-layer neural networks under one-step gradient descent with step size $\eta \asymp n^\alpha$, $\alpha \in (0, \frac{1}{2})$. We showed that if $\alpha \in ((\ell-1)/(2\ell), \ell/(2\ell+2))$, the gradient update will add ℓ separated singular values to the initial feature matrix spectrum corresponding to different nonlinear features. We then proved equivalence theorems and used them to derive a precise characterization of the training and test errors in the high-dimensional proportional limit. Using this, we showed that in certain examples, non-linear components of the teacher function are learned.

Future Work. In this paper, we only study the problem when $\eta \asymp n^\alpha$ with $\alpha \in ((\ell-1)/(2\ell), \ell/(2\ell+2))$. The boundary case where $\eta \asymp n^{(\ell-1)/(2\ell)}$ is an interesting problem and is left as future work.

Also, following prior work in the area (see e.g., [DLS22, BES⁺22, DKL⁺23, NDL23, WNL23], etc.), we use sample splitting in our two-step training procedure. Although this setting is natural for the analysis of pretrained models, it does not cover the case where feature learning and ridge regression use the same data. We leave this setting as a future direction.

In this paper, we focused on the case where $\alpha < 1/2$. The behaviour of the feature matrix can be significantly different when $\alpha = 1/2$. When $\alpha < 1/2$, as proved in this paper, the spectrum of the feature matrix will consist of a finite number of spikes added to the spectrum of the untrained feature matrix. However, when $\alpha = 1/2$, the behaviour can deviate from the spectrum of the untrained model in other ways. Note that according to Theorem 3.2, the number of spikes in the spectrum of \mathbf{F} will increase as we increase α from 0 to 1/2 and will diverge as we approach 1/2. The limiting empirical singular-value distribution of the feature matrix and the training and test errors of the network when $\alpha = 1/2$ is an open problem and we leave it as future work.

Acknowledgements

The work of Behrad Moniri is supported by The Institute for Learning-enabled Optimization at Scale (TILOS), under award number NSF-CCF-2112665. Donghwan Lee and Edgar Dobriban were supported in part by ARO W911NF-20-1-0080, ARO W911NF-23-1-0296, NSF 2031895, NSF DMS 2046874, ONR N00014-21-1-2843, and the Sloan Foundation. Hamed Hassani, Donghwan Lee and Behrad Moniri are also supported by EnCORE: Institute for Emerging CORE Methods in Data Science, under NSF award number 2217058.

References

- [AAM22] Emmanuel Abbe, Enric Boix Adsera, and Theodor Misiakiewicz. The merged-staircase property: a necessary and nearly sufficient condition for sgd learning of sparse functions on two-layer neural networks. In *Conference on Learning Theory*, pages 4782–4887. PMLR, 2022.
- [BABP⁺21] Emmanuel Abbe, Enric Boix-Adsera, Matthew S Brennan, Guy Bresler, and Dheeraj Nagaraj. The staircase property: How hierarchical structure can guide deep learning. In *Advances in Neural Information Processing Systems*, pages 26989–27002, 2021.
- [ALP22] Ben Adlam, Jake A Levinson, and Jeffrey Pennington. A random matrix perspective on mixtures of nonlinearities in high dimensions. In *International Conference on Artificial Intelligence and Statistics*, 2022.
- [AP20a] Ben Adlam and Jeffrey Pennington. The neural tangent kernel in high dimensions: Triple descent and a multi-scale theory of generalization. In *International Conference on Machine Learning*, 2020.
- [AP20b] Ben Adlam and Jeffrey Pennington. Understanding double descent requires a fine-grained bias-variance decomposition. In *Advances in Neural Information Processing Systems*, 2020.
- [AS68] Milton Abramowitz and Irene A Stegun. *Handbook of mathematical functions with formulas, graphs, and mathematical tables*, volume 55. US Government printing office, 1968.
- [ASH19] Ehsan Abbasi, Fariborz Salehi, and Babak Hassibi. Universality in learning from linear measurements. In *Advances in Neural Information Processing Systems*, 2019.
- [BAGJ21] Gerard Ben Arous, Reza Gheissari, and Aukosh Jagannath. Online stochastic gradient descent on non-convex losses from high-dimensional inference. *Journal of Machine Learning Research*, 22(106):1–51, 2021.
- [BCV13] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, 2013.
- [BES⁺22] Jimmy Ba, Murat A Erdogdu, Taiji Suzuki, Zhichao Wang, Denny Wu, and Greg Yang. High-dimensional asymptotics of feature learning: How one gradient step improves the representation. In *Advances in Neural Information Processing Systems*, 2022.
- [BKM23] Simone Bombari, Shayan Kiyani, and Marco Mondelli. Beyond the universal law of robustness: Sharper laws for random features and neural tangent kernels. In *International Conference on Machine Learning*, 2023.
- [BM23] Simone Bombari and Marco Mondelli. Stability, generalization and privacy: Precise analysis for random and NTK features. *arXiv preprint arXiv:2305.12100*, 2023.
- [BMP15] Marwa Banna, Florence Merlevède, and Magda Peligrad. On the limiting spectral distribution for a large class of symmetric random matrices with correlated entries. *Stochastic Processes and their Applications*, 125(7):2700–2726, 2015.
- [BMZ23] Raphaël Berthier, Andrea Montanari, and Kangjie Zhou. Learning time-scales in two-layers neural networks. *arXiv preprint arXiv:2303.00055*, 2023.
- [BNY20] Marwa Banna, Jamal Najim, and Jianfeng Yao. A CLT for linear spectral statistics of large random information-plus-noise matrices. *Stochastic Processes and their Applications*, 130(4):2250–2281, 2020.
- [BP21] Lucas Benigni and Sandrine Péché. Eigenvalue distribution of some nonlinear models of random matrices. *Electronic Journal of Probability*, 26:1–37, 2021.
- [BP22] Lucas Benigni and Sandrine Péché. Largest eigenvalues of the conjugate kernel of single-layered neural networks. *arXiv preprint arXiv:2201.04753*, 2022.

[BPH23] David Bosch, Ashkan Panahi, and Babak Hassibi. Precise asymptotic analysis of deep random feature models. In *Conference on Learning Theory*, 2023.

[BS10] Zhidong Bai and Jack W Silverstein. *Spectral Analysis of Large Dimensional Random Matrices*, volume 20. Springer, 2010.

[Cap14] Mireille Capitaine. Exact separation phenomenon for the eigenvalues of large information-plus-noise type matrices, and an application to spiked models. *Indiana University Mathematics Journal*, pages 1875–1910, 2014.

[CB18] Lenaic Chizat and Francis Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. In *Advances in neural information processing systems*, 2018.

[CCFM21] Yuxin Chen, Yuejie Chi, Jianqing Fan, and Cong Ma. Spectral methods for data science: A statistical perspective. *Foundations and Trends® in Machine Learning*, 14(5):566–806, 2021.

[CD11] Romain Couillet and Merouane Debbah. *Random Matrix Methods for Wireless Communications*. Cambridge University Press, 2011.

[CGKM22] Sitan Chen, Aravind Gollakota, Adam Klivans, and Raghu Meka. Hardness of noise-free learning for two-hidden-layer neural networks. In *Advances in Neural Information Processing Systems*, 2022.

[CHS⁺23] Liam Collins, Hamed Hassani, Mahdi Soltanolkotabi, Aryan Mokhtari, and Sanjay Shakkottai. Provable multi-task representation learning by two-layer ReLU neural networks. *arXiv preprint arXiv:2307.06887*, 2023.

[CKZ23] Hugo Cui, Florent Krzakala, and Lenka Zdeborova. Bayes-optimal learning of deep random networks of extensive-width. In *International Conference on Machine Learning*, 2023.

[CL22] Romain Couillet and Zhenyu Liao. *Random Matrix Methods for Machine Learning*. Cambridge University Press, 2022.

[CLKZ23] Lucas Clarté, Bruno Loureiro, Florent Krzakala, and Lenka Zdeborová. On double-descent in uncertainty quantification in overparametrized models. In *International Conference on Artificial Intelligence and Statistics*, 2023.

[CS13] Xiuyuan Cheng and Amit Singer. The spectrum of random inner-product kernel matrices. *Random Matrices: Theory and Applications*, 2(04):1350010, 2013.

[Dee70] AD Deev. Representation of statistics of discriminant analysis and asymptotic expansion when space dimensions are comparable with sample size. In *Sov. Math. Dokl.*, volume 11, pages 1547–1550, 1970.

[dGSB21] Stéphane d’Ascoli, Marylou Gabrié, Levent Sagun, and Giulio Biroli. On the interplay between data structure and loss function in classification problems. In *Advances in Neural Information Processing Systems*, 2021.

[DKD16] Jeff Donahue, Philipp Krähenbühl, and Trevor Darrell. Adversarial feature learning. In *International Conference on Learning Representations*, 2016.

[DKL⁺23] Yatin Dandi, Florent Krzakala, Bruno Loureiro, Luca Pesce, and Ludovic Stephan. Learning two-layer neural networks, one (giant) step at a time. *arXiv preprint arXiv:2305.18270*, 2023.

[DLS22] Alex Damian, Jason Lee, and Mahdi Soltanolkotabi. Neural networks can learn representations with gradient descent. In *Conference on Learning Theory*, 2022.

[DS07] R Brent Dozier and Jack W Silverstein. On the empirical distribution of eigenvalues of large dimensional information-plus-noise-type matrices. *Journal of Multivariate Analysis*, 98(4):678–694, 2007.

[DW18] Edgar Dobriban and Stefan Wager. High-dimensional asymptotics of prediction: Ridge regression and classification. *The Annals of Statistics*, 46(1):247–279, 2018.

[Dys62] Freeman J Dyson. A Brownian-motion model for the eigenvalues of a random matrix. *Journal of Mathematical Physics*, 3(6):1191–1198, 1962.

[DZPS19] Simon S Du, Xiyu Zhai, Barnabas Poczos, and Aarti Singh. Gradient descent provably optimizes over-parameterized neural networks. In *International Conference on Learning Representations*, 2019.

[EK10] Noureddine El Karoui. The spectrum of kernel random matrices. *The Annals of Statistics*, 38(1):1–50, 2010.

[EPR⁺10] László Erdős, Sandrine Péché, José A Ramírez, Benjamin Schlein, and Horng-Tzer Yau. Bulk universality for Wigner matrices. *Communications on Pure and Applied Mathematics*, 63(7):895–925, 2010.

[Erd19] Laszlo Erdős. The matrix Dyson equation and its applications for random matrices. *arXiv preprint arXiv:1903.10060*, 2019.

[EVdB01] Andreas Engel and Christian Van den Broeck. *Statistical mechanics of learning*. Cambridge University Press, 2001.

[EYY12] László Erdős, Horng-Tzer Yau, and Jun Yin. Bulk universality for generalized Wigner matrices. *Probability Theory and Related Fields*, 154(1):341–407, 2012.

[FM19] Zhou Fan and Andrea Montanari. The spectral norm of random inner-product kernel matrices. *Probability Theory and Related Fields*, 173(1):27–85, 2019.

[FW20] Zhou Fan and Zhichao Wang. Spectra of the conjugate kernel and neural tangent kernel for linear-width neural networks. In *Advances in Neural Information Processing Systems*, 2020.

[Gau61] Michel Gaudin. Sur la loi limite de l’espacement des valeurs propres d’une matrice aléatoire. *Nuclear Physics*, 25:447–458, 1961.

[GK19] Surbhi Goel and Adam R Klivans. Learning neural networks with two nonlinear layers in polynomial time. In *Conference on Learning Theory*, 2019.

[GLR⁺22] Sebastian Goldt, Bruno Loureiro, Galen Reeves, Florent Krzakala, Marc Mézard, and Lenka Zdeborová. The Gaussian equivalence of generative models for learning with shallow neural networks. In *Mathematical and Scientific Machine Learning*, pages 426–471, 2022.

[GMMM21] Behrooz Ghorbani, Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Linearized two-layers neural networks in high dimension. *The Annals of Statistics*, 49(2):1029–1054, 2021.

[GT90] Géza Györgyi and Naftali Tishby. Statistical theory of learning a rule. *Neural Networks and Spin Glasses*, pages 3–36, 1990.

[HJ22] Hamed Hassani and Adel Javanmard. The curse of overparametrization in adversarial training: Precise analysis of robust generalization for random features regression. *arXiv preprint arXiv:2201.05149*, 2022.

[HL23] Hong Hu and Yue M Lu. Universality laws for high-dimensional learning with random features. *IEEE Transactions on Information Theory*, 69(3), 2023.

[HN19] Boris Hanin and Mihai Nica. Finite depth and width corrections to the neural tangent kernel. In *International Conference on Learning Representations*, 2019.

[HY20] Jiaoyang Huang and Horng-Tzer Yau. Dynamics of deep neural networks and neural tangent hierarchy. In *International Conference on Machine Learning*, 2020.

[JGH18] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in Neural Information Processing Systems*, 2018.

[KC22] Arun Kumar Kuchibhotla and Abhishek Chakrabortty. Moving beyond sub-gaussianity in high-dimensional statistics: Applications in covariance estimation and linear regression. *Information and Inference: A Journal of the IMA*, 11(4):1389–1456, 2022.

[LD21] Licong Lin and Edgar Dobriban. What causes the test error? going beyond bias-variance via ANOVA. *Journal of Machine Learning Research*, 22:155–1, 2021.

[LGC⁺21] Bruno Loureiro, Cedric Gerbelot, Hugo Cui, Sebastian Goldt, Florent Krzakala, Marc Mezard, and Lenka Zdeborová. Learning curves of generic features maps for realistic datasets with a teacher-student model. In *Advances in Neural Information Processing Systems*, 2021.

[Lin22] Jarl Waldemar Lindeberg. Eine neue herleitung des exponentialgesetzes in der wahrscheinlichkeitsrechnung. *Mathematische Zeitschrift*, 15(1):211–225, 1922.

[LLC18] Cosme Louart, Zhenyu Liao, and Romain Couillet. A random matrix approach to neural networks. *The Annals of Applied Probability*, 28(2):1190–1248, 2018.

[LMH⁺23] Donghwan Lee, Behrad Moniri, Xinmeng Huang, Edgar Dobriban, and Hamed Hassani. Demystifying disagreement-on-the-line in high dimensions. In *International Conference on Machine Learning*, 2023.

[LXS⁺19] Jaehoon Lee, Lechao Xiao, Samuel Schoenholz, Yasaman Bahri, Roman Novak, Jascha Sohl-Dickstein, and Jeffrey Pennington. Wide neural networks of any depth evolve as linear models under gradient descent. In *Advances in Neural Information Processing Systems*, 2019.

[LY22] Yue M Lu and Horng-Tzer Yau. An equivalence principle for the spectrum of random inner-product kernel matrices. *arXiv preprint arXiv:2205.06308*, 2022.

[Meh04] Madan Lal Mehta. *Random matrices*. Elsevier, 2004.

[Mis22] Theodor Misiakiewicz. Spectrum of inner-product kernel matrices in the polynomial regime and multiple descent phenomenon in kernel ridge regression. *arXiv preprint arXiv:2204.10425*, 2022.

[MM22] Song Mei and Andrea Montanari. The generalization error of random features regression: Precise asymptotics and the double descent curve. *Communications on Pure and Applied Mathematics*, 75(4):667–766, 2022.

[MMM19] Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Mean-field theory of two-layers neural networks: dimension-free bounds and kernel limit. In *Conference on Learning Theory*, 2019.

[MMN18] Song Mei, Andrea Montanari, and Phan-Minh Nguyen. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33):E7665–E7671, 2018.

[MP21] Gabriel Mel and Jeffrey Pennington. Anisotropic random feature regression in high dimensions. In *International Conference on Learning Representations*, 2021.

[MRSY19] Andrea Montanari, Feng Ruan, Youngtak Sohn, and Jun Yan. The generalization error of max-margin linear classifiers: High-dimensional asymptotics in the overparametrized regime. *arXiv preprint arXiv:1911.01544*, 2019.

[MS22] Andrea Montanari and Basil N Saeed. Universality of empirical risk minimization. In *Conference on Learning Theory*, 2022.

[NDL23] Eshaan Nichani, Alex Damian, and Jason D Lee. Provable guarantees for nonlinear feature learning in three-layer neural networks. *arXiv preprint arXiv:2305.06986*, 2023.

[NR21] Gadi Naveh and Zohar Ringel. A self consistent theory of gaussian processes captures feature learning effects in finite cnns. In *Advances in Neural Information Processing Systems*, 2021.

[O'D14] Ryan O'Donnell. *Analysis of boolean functions*. Cambridge University Press, 2014.

[OK96] Manfred Opper and Wolfgang Kinzel. Statistical mechanics of generalization. In *Models of Neural Networks III*, pages 151–209. Springer, 1996.

[Opp95] Manfred Opper. Statistical mechanics of learning: Generalization. *The Handbook of Brain Theory and Neural Networks*, pages 922–925, 1995.

[PA14] Debashis Paul and Alexander Aue. Random matrix theory in statistics: A review. *Journal of Statistical Planning and Inference*, 150:1–29, 2014.

[P  c19] Sandrine P  ch  . A note on the Pennington-Worah distribution. *Electronic Communications in Probability*, 24:1–7, 2019.

[PW17] Jeffrey Pennington and Pratik Worah. Nonlinear random matrix theory for deep learning. In *Advances in Neural Information Processing Systems*, 2017.

[Rau67]   r  nas Raudys. On determining training sample size of linear classifier. *Computing Systems (in Russian)*, 28:79–87, 1967.

[Rau72]   r  nas Raudys. On the amount of a priori information in designing the classification algorithm. *Technical Cybernetics (in Russian)*, 4:168–174, 1972.

[RBPB22] Adityanarayanan Radhakrishnan, Daniel Beaglehole, Parthe Pandit, and Mikhail Belkin. Feature learning in neural networks and kernel machines that recursively learn features. *arXiv preprint arXiv:2212.13881*, 2022.

[RR07] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems*, 2007.

[RVE22] Grant Rotskoff and Eric Vanden-Eijnden. Trainability and accuracy of artificial neural networks: An interacting particle system approach. *Communications on Pure and Applied Mathematics*, 75(9):1889–1935, 2022.

[RY04]   r  nas Raudys and Dean M Young. Results in statistical discriminant analysis: A review of the former Soviet Union literature. *Journal of Multivariate Analysis*, 89(1):1–35, 2004.

[Sam23] Holger Sambale. Some notes on concentration for α -subexponential random variables. In *High Dimensional Probability IX: The Ethereal Volume*, pages 167–192. Springer, 2023.

[Ser07] Vadim Ivanovich Serdobolskii. *Multiparametric Statistics*. Elsevier, 2007.

[SNR23] Inbar Seroussi, Gadi Naveh, and Zohar Ringel. Separation of scales and a thermodynamic description of feature learning in some cnns. *Nature Communications*, 14(1):908, 2023.

[SS90] Gilbert W Stewart and Ji-guang Sun. *Matrix perturbation theory*. Elsevier, 1990.

[SS20] Justin Sirignano and Konstantinos Spiliopoulos. Mean field analysis of neural networks: A central limit theorem. *Stochastic Processes and their Applications*, 130(3):1820–1852, 2020.

[SWL22] Zhenmei Shi, Junyi Wei, and Yingyu Liang. A theoretical analysis on feature learning in neural networks: Emergence from inputs and advantage over fixed features. In *International Conference on Learning Representations*, 2022.

[TAP21] Nilesh Tripuraneni, Ben Adlam, and Jeffrey Pennington. Overparameterization improves robustness to covariate shift in high dimensions. In *Advances in Neural Information Processing Systems*, 2021.

[TV04] Antonio M Tulino and Sergio Verd  . Random matrix theory and wireless communications. *Communications and Information Theory*, 1(1):1–182, 2004.

[TV11] Terence Tao and Van Vu. Random matrices: universality of local eigenvalue statistics. *Acta mathematica*, 206(1):127–204, 2011.

[vdVW13] Aad van der Vaart and Jon Wellner. *Weak convergence and empirical processes: with applications to statistics*. Springer Science & Business Media, 2013.

[Ver12] Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. In Yonina C. Eldar and Gitta Kutyniok, editors, *Compressed Sensing: Theory and Applications*, page 210–268. Cambridge University Press, 2012.

[Ver18] Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*. Cambridge University Press, 2018.

[VGNA20] Mariia Vladimirova, Stéphane Girard, Hien Nguyen, and Julyan Arbel. Sub-weibull distributions: Generalizing sub-gaussian and sub-exponential properties to heavier tailed distributions. *Stat*, 9(1), 2020.

[Wed72] Per-Ake Wedin. Perturbation bounds in connection with singular value decomposition. *BIT Numerical Mathematics*, 12:99–111, 1972.

[WES⁺22] Zhichao Wang, Andrew Engel, Anand Sarwate, Ioana Dumitriu, and Tony Chiang. Spectral evolution and invariance in linear-width neural networks. *arXiv preprint arXiv:2211.06506*, 2022.

[Wig55] Eugene P Wigner. Characteristic vectors of bordered matrices with infinite dimensions. *Annals of Mathematics*, pages 548–564, 1955.

[WNL23] Zihao Wang, Eshaan Nichani, and Jason D Lee. Learning hierarchical polynomials with three-layer neural networks. *arXiv preprint arXiv:2311.13774*, 2023.

[XHM⁺22] Lechao Xiao, Hong Hu, Theodor Misiakiewicz, Yue Lu, and Jeffrey Pennington. Precise learning curves and higher-order scalings for dot-product kernel regression. In *Advances in Neural Information Processing Systems*, 2022.

[Yai20] Sho Yaida. Non-gaussian processes and neural networks at finite widths. In *Mathematical and Scientific Machine Learning*, pages 165–192. PMLR, 2020.

[YBZ15] Jianfeng Yao, Zhidong Bai, and Shurong Zheng. *Large Sample Covariance Matrices and High-Dimensional Data Analysis*. Cambridge University Press, 2015.

[YH21] Greg Yang and Edward J Hu. Feature learning in infinite-width neural networks. In *International Conference on Machine Learning*, 2021.

[ZC21] Huiming Zhang and Songxi Chen. Concentration inequalities for statistical inference. *Communications in Mathematical Research*, 37(1):1–85, 2021.

[ZWL22] Shi Zhenmei, Junyi Wei, and Yingyu Liang. A theoretical analysis on feature learning in neural networks: Emergence from inputs and advantage over fixed features. In *International Conference on Learning Representations*, 2022.

A Additional Related Work

[GK19] provide a polynomial time algorithm that learns neural networks with two non-linear layers. Our setting is different because we do not apply a non-linear activation after the second layer. [CGKM22] show that learning two-hidden-layer neural networks from noise-free Gaussian data requires superpolynomially many statistical queries. [ZWL22] show that neural networks trained by gradient descent can succeed on problems where the labels are determined by a set of class-relevant patterns and if these patterns are removed, no polynomial algorithm in the Statistical Query model can learn even weakly.

B Additional Notation and Terminology

In the appendix, we use the following additional notations. We let $\mathbb{N}_0 = \{0, 1, 2, \dots\}$ be the set of non-negative integers. For a set X and $x_1, x_2 \in X$, δ_{x_1, x_2} is the Kronecker delta, which equals unity if $x_1 = x_2$, and is zero otherwise. We use $\tilde{O}(\cdot)$ for the standard big-O notation up to logarithmic factors in n . For a positive integer k , $k!!$ is the product of all the positive integers up to n with the same parity as n . For two random quantities X, Y , $X \perp\!\!\!\perp Y$ denotes that X is independent of Y . By orderwise analysis, we mean bounding a term by the triangle inequality and the inequality $\|Ab\|_2 \leq \|A\|_{\text{op}}\|b\|_2$ for a conformable matrix-vector pair A, b , to reduce it to operator norms of matrices and Euclidean norms of vectors, and then use simple bounds for those quantities. Constants such as C, c' , etc., can change from line to line unless specified otherwise. For two random quantities A, B , $A =_d B$ denotes that A and B have the same distribution. Limits of random variables are understood in probability. For two matrices \mathbf{A}, \mathbf{B} with equal shape, we write $\mathbf{A} \circ \mathbf{B}$ to denote their entry-wise (Hadamard) product.

We denote $\mathbf{X}\boldsymbol{\beta} = \boldsymbol{\theta}$, $\tilde{\mathbf{X}}\boldsymbol{\beta} = \tilde{\boldsymbol{\theta}}$, $\mathbf{X}\boldsymbol{\beta}_* = \boldsymbol{\theta}_*$, and $\tilde{\mathbf{X}}\boldsymbol{\beta}_* = \tilde{\boldsymbol{\theta}}_*$. We also define $\bar{\mathbf{R}}_0 = (\mathbf{F}_0\mathbf{F}_0^\top + \lambda n\mathbf{I}_n)^{-1}$ and $\mathbf{R}_0 = (\mathbf{F}_0^\top\mathbf{F}_0 + \lambda n\mathbf{I}_N)^{-1}$.

C Basic Lemmas

Lemma C.1 (Orthogonality of Hermite polynomials). *Let (Z_1, Z_2) be jointly Gaussian with $\mathbb{E}[Z_1] = \mathbb{E}[Z_2] = 0$, $\mathbb{E}[Z_1^2] = \mathbb{E}[Z_2^2] = 1$, and $\mathbb{E}[Z_1 Z_2] = \rho$. Then for any $k_1, k_2 \in \mathbb{N}_0$,*

$$\mathbb{E}[H_{k_1}(Z_1)H_{k_2}(Z_2)] = k_1! \rho^{k_1} \delta_{k_1, k_2}.$$

In particular, if for some positive integer d , $\mathbf{Z} \sim \mathcal{N}(0, \mathbf{I}_d)$, and if $\mathbf{a}, \mathbf{b} \in \mathbb{S}^{d-1}$, then

$$\mathbb{E}[H_{k_1}(\mathbf{a}^\top \mathbf{Z})H_{k_2}(\mathbf{b}^\top \mathbf{Z})] = k_1! (\mathbf{a}^\top \mathbf{b})^{k_1} \delta_{k_1, k_2}.$$

Proof. See [O'D14, Chapter 11.2]. □

Lemma C.2 (Taylor expansion of Hermite polynomials). *For any $k \in \mathbb{N}_0$ and $x, y \in \mathbb{R}$,*

$$H_k(x + y) = \sum_{j=0}^k \binom{k}{j} x^j H_{k-j}(y).$$

Proof. Note that $\frac{d}{dx} H_k(x) = k H_{k-1}(x)$ [AS68, Equation 22.8.8] and thus $\frac{d^j}{dx^j} H_k(x) = \frac{k!}{(k-j)!} H_{k-j}(x)$. By Taylor expanding $H_k(x+y)$ at y , we find

$$H_k(x+y) = \sum_{j=0}^k \frac{x^j}{j!} \frac{d^j}{dy^j} H_k(y) = \sum_{j=0}^k \binom{k}{j} x^j H_{k-j}(y).$$

□

The following Lemma, proved in Section O.1, provides several bounds used in the proofs.

Lemma C.3. *Under Conditions 2.1–2.4, there exists $C > 0$ such that the following holds with probability $1 - o(1)$.*

- (a) $M_{\mathbf{a}} := \max_{1 \leq i \leq N} |a_i| \leq C n^{-\frac{1}{2}} \log^{\frac{1}{2}} n$,
- (b) $M_{\boldsymbol{\beta}} := \max_{1 \leq i \leq n} |\langle \tilde{\mathbf{x}}_i, \boldsymbol{\beta} \rangle| \leq C \log^{\frac{1}{2}} n$,
- (c) $M_{\mathbf{W}_0} := \sup_{k \geq 1} \|(\mathbf{W}_0 \mathbf{W}_0^\top)^{\circ k}\|_{\text{op}} \leq C$,
- (d) $\|\tilde{\mathbf{X}}\|_{\text{op}} \leq C \sqrt{n}$.

D Proof of equation 3

By Lemma K.1 with $\mathbf{v} = \boldsymbol{\beta}_*$ and $\mathbf{D} = \mathbf{I}_d$, we have

$$\boldsymbol{\beta}_*^\top \boldsymbol{\beta} \rightarrow_P c_{*,1} \|\boldsymbol{\beta}_*\|_2^2 = c_{*,1}, \quad \|\boldsymbol{\beta}\|_2^2 = \boldsymbol{\beta}^\top \boldsymbol{\beta} \rightarrow_P \phi(c_*^2 + \sigma_\varepsilon^2) + c_{*,1}^2 \boldsymbol{\beta}_*^\top \boldsymbol{\beta}_* = c_{*,1}^2 + \phi(c_*^2 + \sigma_\varepsilon^2).$$

By the continuous mapping theorem, we conclude

$$\frac{|\boldsymbol{\beta}_*^\top \boldsymbol{\beta}|}{\|\boldsymbol{\beta}_*\|_2 \|\boldsymbol{\beta}\|_2} \rightarrow_P \frac{|c_{*,1}|}{\sqrt{c_{*,1}^2 + \phi(c_*^2 + \sigma_\varepsilon^2)}}.$$

E Proof of Lemma 3.1

For $k = 1$, we have by [BES⁺22, Proposition 2]—substituting our c_1 for their μ_1 and using that $\boldsymbol{\beta} = \frac{1}{n} \mathbf{X}^\top \mathbf{y}$; as well as noting that by the discussion below [BES⁺22, Proposition 2], $\|\mathbf{G}\|_{\text{op}} = O_{\mathbb{P}}(1)$ —and by Lemma C.3 (d), that with probability $1 - o(1)$,

$$\|\tilde{\mathbf{X}} \mathbf{G}^\top - c_1 \tilde{\mathbf{X}} \boldsymbol{\beta} \mathbf{a}^\top\|_{\text{op}} = O(n^{-\frac{1}{2}} \log^2 n). \quad (10)$$

For $k \geq 2$, expanding $(\tilde{\mathbf{X}} \mathbf{G}^\top)^{\circ k} = (\tilde{\mathbf{X}} \mathbf{G}^\top - c_1 \tilde{\mathbf{X}} \boldsymbol{\beta} \mathbf{a}^\top + c_1 \tilde{\mathbf{X}} \boldsymbol{\beta} \mathbf{a}^\top)^{\circ k}$ using the binomial formula, we have

$$\begin{aligned} (\tilde{\mathbf{X}} \mathbf{G}^\top)^{\circ k} - c_1^k (\tilde{\mathbf{X}} \boldsymbol{\beta})^{\circ k} (\mathbf{a}^{\circ k})^\top &= \sum_{j=1}^k \binom{k}{j} (\tilde{\mathbf{X}} \mathbf{G}^\top - c_1 \tilde{\mathbf{X}} \boldsymbol{\beta} \mathbf{a}^\top)^{\circ j} \circ (c_1 \tilde{\mathbf{X}} \boldsymbol{\beta} \mathbf{a}^\top)^{\circ (k-j)} \\ &= \sum_{j=1}^k \binom{k}{j} c_1^{k-j} \text{diag}(\tilde{\mathbf{X}} \boldsymbol{\beta})^{k-j} (\tilde{\mathbf{X}} \mathbf{G}^\top - c_1 \tilde{\mathbf{X}} \boldsymbol{\beta} \mathbf{a}^\top)^{\circ j} \text{diag}(\mathbf{a})^{k-j}. \end{aligned}$$

Recalling $M_{\mathbf{a}}, M_{\boldsymbol{\beta}}$ from Lemma C.3, and using that

$$\|(\tilde{\mathbf{X}}\mathbf{G}^\top - c_1\tilde{\mathbf{X}}\boldsymbol{\beta}\mathbf{a}^\top)^{\circ j}\|_{\text{op}} \leq \|\tilde{\mathbf{X}}\mathbf{G}^\top - c_1\tilde{\mathbf{X}}\boldsymbol{\beta}\mathbf{a}^\top\|_{\text{op}}^j$$

(see e.g., [BS10, Corollary A.21]), we have

$$\begin{aligned} & \|\text{diag}(\tilde{\mathbf{X}}\boldsymbol{\beta})^{k-j}(\tilde{\mathbf{X}}\mathbf{G}^\top - c_1\tilde{\mathbf{X}}\boldsymbol{\beta}\mathbf{a}^\top)^{\circ j} \text{diag}(\mathbf{a})^{k-j}\|_{\text{op}} \\ & \leq \|\text{diag}(\tilde{\mathbf{X}}\boldsymbol{\beta})^{k-j}\|_{\text{op}} \|(\tilde{\mathbf{X}}\mathbf{G}^\top - c_1\tilde{\mathbf{X}}\boldsymbol{\beta}\mathbf{a}^\top)^{\circ j}\|_{\text{op}} \|\text{diag}(\mathbf{a})^{k-j}\|_{\text{op}} \\ & \leq (M_{\mathbf{a}}M_{\boldsymbol{\beta}})^{k-j} \|\tilde{\mathbf{X}}\mathbf{G}^\top - c_1\tilde{\mathbf{X}}\boldsymbol{\beta}\mathbf{a}^\top\|_{\text{op}}^j. \end{aligned}$$

Hence, by the triangle inequality,

$$\|(\tilde{\mathbf{X}}\mathbf{G}^\top)^{\circ k} - c_1^k(\tilde{\mathbf{X}}\boldsymbol{\beta})^{\circ k}(\mathbf{a}^{\circ k})^\top\|_{\text{op}} \leq \sum_{j=1}^k \binom{k}{j} (c_1 M_{\mathbf{a}} M_{\boldsymbol{\beta}})^{k-j} \|\tilde{\mathbf{X}}\mathbf{G}^\top - c_1\tilde{\mathbf{X}}\boldsymbol{\beta}\mathbf{a}^\top\|_{\text{op}}^j.$$

By Lemma C.3 (a), (b) and equation 10, there exists $C > 0$ such that for any $k \in \mathbb{N}$,

$$\begin{aligned} \sum_{j=1}^k \binom{k}{j} (c_1 M_{\mathbf{a}} M_{\boldsymbol{\beta}})^{k-j} \|\tilde{\mathbf{X}}\mathbf{G}^\top - c_1\tilde{\mathbf{X}}\boldsymbol{\beta}\mathbf{a}^\top\|_{\text{op}}^j & \leq (C/2)^k \sum_{j=1}^k \binom{k}{j} (n^{-\frac{1}{2}} \log n)^{k-j} (n^{-\frac{1}{2}} \log^2 n)^j \\ & \leq C^k n^{-\frac{k}{2}} \log^{2k} n \end{aligned}$$

with probability $1 - o(1)$.

F Proof of Theorem 3.2

We consider any fixed \mathbf{W}_0 such that the event $\Omega = \{\sup_{k \geq 1} \|(\mathbf{W}_0\mathbf{W}_0^\top)^{\circ k}\|_{\text{op}} \leq C\}$ from Lemma C.3 (c) holds. By Lemma C.1, each row of $H_j(\tilde{\mathbf{X}}\mathbf{W}_0^\top)$ has second moment matrix

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{N}(0, \mathbf{I}_d)} [H_j(\mathbf{W}_0\mathbf{x}) H_j(\mathbf{W}_0\mathbf{x})^\top] = j! (\mathbf{W}_0\mathbf{W}_0^\top)^{\circ j},$$

whose operator norm is $O(j!)$ on Ω . Thus by [Ver12, Theorem 5.48] and Markov's inequality, for any $j \in [L]$, for $t \geq (Cnj!)^{1/2}$, and with $M = \mathbb{E} \max_{i=1}^n \|H_j(\mathbf{W}_0\tilde{\mathbf{x}}_i)\|^2$, $\delta = C\sqrt{M \log \min(n, N)}$,

$$\begin{aligned} P(\|H_j(\tilde{\mathbf{X}}\mathbf{W}_0^\top)\|_{\text{op}} \geq t) & \leq P\left(\|H_j(\tilde{\mathbf{X}}\mathbf{W}_0^\top) H_j(\tilde{\mathbf{X}}\mathbf{W}_0^\top)^\top/n - j! (\mathbf{W}_0\mathbf{W}_0^\top)^{\circ j}\|_{\text{op}} \geq t^2/n - Cj!\right) \\ & \leq \frac{\mathbb{E}\|H_j(\tilde{\mathbf{X}}\mathbf{W}_0^\top) H_j(\tilde{\mathbf{X}}\mathbf{W}_0^\top)^\top/n - j! (\mathbf{W}_0\mathbf{W}_0^\top)^{\circ j}\|_{\text{op}}}{t^2/n - Cj!} \\ & \leq \frac{\delta \max((Cj!)^{1/2}, \delta)}{t^2/n - Cj!}. \end{aligned}$$

Next, we observe that since H_j is a j -th degree polynomial and the normal absolute moments increase with j , $M = \mathbb{E} \max_{i=1}^n \|H_j(\mathbf{W}_0\tilde{\mathbf{x}}_i)\|^2 \leq C_j \mathbb{E} \max_{i=1}^n \|(\mathbf{W}_0\tilde{\mathbf{x}}_i)^{\circ j}\|^2$. Now, note that for any vectors $\mathbf{x}_1, \mathbf{x}_2$, we have $\|\mathbf{x}_1 \circ \mathbf{x}_2\|^2 \leq \|\mathbf{x}_1\|^2 \|\mathbf{x}_2\|^2$ by simply expanding the norms. Thus, on the event Ω , one can verify that for all \mathbf{x} , $\|(\mathbf{W}_0\mathbf{x})^{\circ j}\|^2 \leq C'_j \|\mathbf{x}\|^{2j}$ for some $C'_j > 0$. Also, we have that $A_i = \|\tilde{\mathbf{x}}_i\|^{2j}/N$ for $i \in [n]$ are sub-Weibull random variables with tail parameter $1/(2j)$,

(see e.g., [VGNA20, ZC21]). Thus, by the maximal inequality for sub-Weibull random variables [KC22, Proposition A.6 and Remark A.1], it follows that for all $j \geq 1$, there is $C_j > 0$ such that $\mathbb{E} \max_{i=1}^n A_i \leq C_j (\log n)^{2j}$. Hence, $M \leq C_j'' N (\log n)^{2j}$.

Thus, choosing $t = C' \sqrt{n} j! (\log n)^j$ for sufficiently large C' leads to

$$\|H_j(\tilde{\mathbf{X}}\mathbf{W}_0^\top)\|_{\text{op}} = O\left(\sqrt{n} j! (\log n)^j\right) \quad (11)$$

with probability $1 - o(1)$.

Define, for all $z \in \mathbb{R}$, $\sigma_L(z) = \sum_{k=0}^L c_k H_k(z)$, where $L = \max\{\ell, \frac{\log n}{4(\ell+1) \log \log n}\}$. Each row of $(\sigma - \sigma_L)(\tilde{\mathbf{X}}\mathbf{W}_0^\top)$ has second moment matrix

$$\mathbb{E}_{\mathbf{x} \sim \mathcal{N}(0, \mathbf{I}_d)}[(\sigma - \sigma_L)(\mathbf{W}_0 \mathbf{x})(\sigma - \sigma_L)(\mathbf{W}_0 \mathbf{x})^\top] = \sum_{k=L+1}^{\infty} k! c_k^2 (\mathbf{W}_0 \mathbf{W}_0^\top)^{\circ k},$$

whose operator norm is $O(L^{-\frac{1}{2}-\omega})$ by Lemma C.3 (c) and Condition 2.3. Therefore,

$$\|(\sigma - \sigma_L)(\tilde{\mathbf{X}}\mathbf{W}_0^\top)\|_{\text{op}} = O(\sqrt{n \log n} L^{-\frac{1}{2}-\omega}) = o(\sqrt{n}) \quad (12)$$

with probability $1 - o(1)$. Since $\eta = o(\sqrt{n})$, the rows of \mathbf{W} have \mathbf{W} norm of $O_{\mathbb{P}}(1)$. Thus, we can repeat the same argument to show that with probability $1 - o(1)$, we have

$$\|(\sigma - \sigma_L)(\tilde{\mathbf{X}}\mathbf{W}^\top)\|_{\text{op}} = O(\sqrt{n \log n} L^{-\frac{1}{2}-\omega}) = o(\sqrt{n}). \quad (13)$$

Let $\mathbf{F}^{(L)} := \sigma_L(\tilde{\mathbf{X}}\mathbf{W}^\top)$ and $\mathbf{F}_0^{(L)} := \sigma_L(\tilde{\mathbf{X}}\mathbf{W}_0^\top)$. We can write

$$\mathbf{F}^{(L)} = \mathbf{F}_0^{(L)} + \sum_{k=1}^L c_k (H_k(\tilde{\mathbf{X}}\mathbf{W}^\top) - H_k(\tilde{\mathbf{X}}\mathbf{W}_0^\top)).$$

By Lemma C.2, using $\mathbf{W} = \mathbf{W}_0 + \eta \mathbf{G}$ so that $\tilde{\mathbf{X}}\mathbf{W}^\top = \tilde{\mathbf{X}}\mathbf{W}_0^\top + \eta \tilde{\mathbf{X}}\mathbf{G}^\top$, and using that $H_0(z) = 1$ for all $z \in \mathbb{R}$,

$$H_k(\tilde{\mathbf{X}}\mathbf{W}^\top) - H_k(\tilde{\mathbf{X}}\mathbf{W}_0^\top) = \eta^k (\tilde{\mathbf{X}}\mathbf{G}^\top)^{\circ k} + \sum_{j=1}^{k-1} \binom{k}{j} \eta^j H_{k-j}(\tilde{\mathbf{X}}\mathbf{W}_0^\top) \circ (\tilde{\mathbf{X}}\mathbf{G}^\top)^{\circ j}.$$

Therefore,

$$\begin{aligned} \mathbf{F}^{(L)} &= \mathbf{F}_0^{(L)} + \sum_{k=1}^{\ell} c_1^k c_k \eta^k (\tilde{\mathbf{X}}\boldsymbol{\beta})^{\circ k} (\mathbf{a}^{\circ k})^\top + \underbrace{\sum_{k=1}^L c_k \eta^k \left[(\tilde{\mathbf{X}}\mathbf{G}^\top)^{\circ k} - c_1^k (\tilde{\mathbf{X}}\boldsymbol{\beta})^{\circ k} (\mathbf{a}^{\circ k})^\top \right]}_{\Delta_1} \\ &\quad + \underbrace{\sum_{k=\ell+1}^L c_1^k c_k \eta^k (\tilde{\mathbf{X}}\boldsymbol{\beta})^{\circ k} (\mathbf{a}^{\circ k})^\top}_{\Delta_2} + \underbrace{\sum_{k=1}^L \sum_{j=1}^{k-1} c_k \binom{k}{j} \eta^j H_{k-j}(\tilde{\mathbf{X}}\mathbf{W}_0^\top) \circ \left[(\tilde{\mathbf{X}}\mathbf{G}^\top)^{\circ j} - c_1^j (\tilde{\mathbf{X}}\boldsymbol{\beta})^{\circ j} (\mathbf{a}^{\circ j})^\top \right]}_{\Delta_3} \\ &\quad + \underbrace{\sum_{k=1}^L \sum_{j=1}^{k-1} c_1^j c_k \binom{k}{j} \eta^j H_{k-j}(\tilde{\mathbf{X}}\mathbf{W}_0^\top) \circ \left[(\tilde{\mathbf{X}}\boldsymbol{\beta})^{\circ j} (\mathbf{a}^{\circ j})^\top \right]}_{\Delta_4}. \end{aligned}$$

We will show that each of $\|\Delta_1\|_{\text{op}}, \|\Delta_2\|_{\text{op}}, \|\Delta_3\|_{\text{op}}, \|\Delta_4\|_{\text{op}}$ is $o(\sqrt{n})$ with probability $1 - o(1)$.

By Lemma 3.1,

$$\|\Delta_1\|_{\text{op}} \leq \sum_{k=1}^L c_k C^k \eta^k n^{-\frac{k}{2}} \log^{2k} n = \tilde{O}(\eta/\sqrt{n}) = o(\sqrt{n})$$

with probability $1 - o(1)$.

By Lemma C.3 (a) and (b), using that $\alpha < \frac{\ell}{2\ell+2}$,

$$\|\Delta_2\|_{\text{op}} \leq \sum_{k=\ell+1}^L c_1^k c_k \eta^k \|(\tilde{\mathbf{X}}\beta)^{\circ k}\|_2 \|\mathbf{a}^{\circ k}\|_2 \leq \sum_{k=\ell+1}^L c_1^k c_k \eta^k n M_{\mathbf{a}}^k M_{\beta}^k = \tilde{O}(n(\eta/\sqrt{n})^{\ell+1}) = o(\sqrt{n})$$

with probability $1 - o(1)$.

By [BS10, Corollary A.21], equation 11, and Lemma 3.1,

$$\|\Delta_3\|_{\text{op}} \leq \sum_{k=1}^L \sum_{j=1}^{k-1} c_k C^j \binom{k}{j} \sqrt{(k-j)!} \eta^j n^{-\frac{j}{2} + \frac{1}{2}} \log^{k+j} n = \tilde{O}(\eta) = o(\sqrt{n}).$$

Finally, since

$$\begin{aligned} \|H_{k-j}(\tilde{\mathbf{X}}\mathbf{W}_0^\top) \circ (\tilde{\mathbf{X}}\beta)^{\circ j} (\mathbf{a}^{\circ j})^\top\|_{\text{op}} &\leq (M_{\mathbf{a}} M_{\beta})^j \sqrt{(k-j)!} n^{\frac{1}{2}} \log^{k-j} n \\ &\leq C^j \sqrt{(k-j)!} n^{-\frac{j}{2} + \frac{1}{2}} \log^k n, \end{aligned}$$

we also have

$$\|\Delta_4\|_{\text{op}} \leq \sum_{k=1}^L \sum_{j=1}^{k-1} c_k C^j \binom{k}{j} \sqrt{(k-j)!} \eta^j n^{-\frac{j}{2} + \frac{1}{2}} \log^k n = \tilde{O}(\eta) = o(\sqrt{n}).$$

This proves that with probability $1 - o(1)$, we have $\mathbf{F}^{(L)} = \mathbf{F}_0^{(L)} + \sum_{k=1}^{\ell} c_1^k c_k \eta^k (\tilde{\mathbf{X}}\beta)^{\circ k} (\mathbf{a}^{\circ k})^\top + \Delta$, with $\|\Delta\|_{\text{op}} = o(1)$. This, alongside equation 12 and equation 13, concludes the proof.

G Proof of Theorem 3.3

By Theorem 3.2, letting $\mathbf{E} = \mathbf{F}_0 + \Delta$, we have $\|\mathbf{E}\|_{\text{op}} = O_{\mathbb{P}}(\sqrt{n})$. Note that $\sum_{k=1}^{\ell} c_1^k c_k \eta^k (\tilde{\mathbf{X}}\beta)^{\circ k} (\mathbf{a}^{\circ k})^\top$ has rank ℓ almost surely and its left singular vector space is $\text{span}\{(\tilde{\mathbf{X}}\beta)^{\circ k}\}_{k \in [\ell]}$. Also, the subspace spanned by the top- ℓ left singular vectors of \mathbf{F} is \mathcal{F}_ℓ . By Wedin's theorem [Wed72], [CCFM21, Theorem 2.9], and as $\alpha > \frac{\ell-1}{2\ell}$, we have

$$d(\mathcal{F}_\ell, \text{span}\{(\tilde{\mathbf{X}}\beta)^{\circ k}\}_{k \in [\ell]}) = O_{\mathbb{P}} \left(\frac{\|\mathbf{E}\|_{\text{op}}}{\eta^\ell n^{\frac{1}{2} - \frac{\ell-1}{2}} - \|\mathbf{E}\|_{\text{op}}} \right) = O_{\mathbb{P}}(n^{\frac{\ell-1}{2} - \alpha\ell}) = o_{\mathbb{P}}(1),$$

which concludes the proof.

H Proof of Theorem 4.1

By the definition of $\hat{\mathbf{a}}(\mathbf{F})$, we have

$$\begin{aligned} \max \left\{ \frac{1}{n} \|\tilde{\mathbf{y}} - \mathbf{F}\hat{\mathbf{a}}(\mathbf{F})\|_2^2, \lambda \|\hat{\mathbf{a}}(\mathbf{F})\|_2^2 \right\} &\leq \frac{1}{n} \|\tilde{\mathbf{y}} - \mathbf{F}\hat{\mathbf{a}}(\mathbf{F})\|_2^2 + \lambda \|\hat{\mathbf{a}}(\mathbf{F})\|_2^2 \\ &\leq \frac{1}{n} \|\tilde{\mathbf{y}} - \mathbf{F} \cdot \mathbf{0}\|_2^2 + \lambda \|\mathbf{0}\|_2^2 = \frac{1}{n} \|\tilde{\mathbf{y}}\|_2^2 = O_{\mathbb{P}}(1). \end{aligned}$$

Thus,

$$\|\hat{\mathbf{a}}(\mathbf{F})\|_2 = O_{\mathbb{P}}(1), \quad \|\tilde{\mathbf{y}} - \mathbf{F}\hat{\mathbf{a}}(\mathbf{F})\|_2 = O_{\mathbb{P}}(\sqrt{n}). \quad (14)$$

A similar argument gives

$$\|\hat{\mathbf{a}}(\mathbf{F}_\ell)\|_2 = O_{\mathbb{P}}(1), \quad \|\tilde{\mathbf{y}} - \mathbf{F}_\ell\hat{\mathbf{a}}(\mathbf{F}_\ell)\|_2 = O_{\mathbb{P}}(\sqrt{n}). \quad (15)$$

Also, by the triangle inequality, and using equation 15 and Theorem 3.2, which states $\|\mathbf{F}_\ell - \mathbf{F}\|_{\text{op}} = o_{\mathbb{P}}(\sqrt{n})$, we have

$$\begin{aligned} \|\tilde{\mathbf{y}} - \mathbf{F}\hat{\mathbf{a}}(\mathbf{F}_\ell)\|_2 &\leq \|\tilde{\mathbf{y}} - \mathbf{F}_\ell\hat{\mathbf{a}}(\mathbf{F}_\ell)\|_2 + \|(\mathbf{F}_\ell - \mathbf{F})\hat{\mathbf{a}}(\mathbf{F}_\ell)\|_2 \\ &\leq \|\tilde{\mathbf{y}} - \mathbf{F}_\ell\hat{\mathbf{a}}(\mathbf{F}_\ell)\|_2 + \|\mathbf{F}_\ell - \mathbf{F}\|_{\text{op}} \|\hat{\mathbf{a}}(\mathbf{F}_\ell)\|_2 = O_{\mathbb{P}}(\sqrt{n}). \end{aligned} \quad (16)$$

Similarly, we can prove that

$$\|\tilde{\mathbf{y}} - \mathbf{F}_\ell\hat{\mathbf{a}}(\mathbf{F})\|_2 = O_{\mathbb{P}}(\sqrt{n}). \quad (17)$$

For $\mathbf{a} = \hat{\mathbf{a}}(\mathbf{F})$ or $\mathbf{a} = \hat{\mathbf{a}}(\mathbf{F}_\ell)$,

$$\begin{aligned} \frac{1}{n} (\|\tilde{\mathbf{y}} - \mathbf{F}\mathbf{a}\|_2^2 - \|\tilde{\mathbf{y}} - \mathbf{F}_\ell\mathbf{a}\|_2^2) &= \frac{1}{n} \langle (\mathbf{F}_\ell - \mathbf{F})\mathbf{a}, \tilde{\mathbf{y}} - \mathbf{F}\mathbf{a} + \tilde{\mathbf{y}} - \mathbf{F}_\ell\mathbf{a} \rangle \\ &\leq \frac{1}{n} \|\mathbf{F}_\ell - \mathbf{F}\|_{\text{op}} \|\mathbf{a}\|_2 (\|\tilde{\mathbf{y}} - \mathbf{F}\mathbf{a}\|_2 + \|\tilde{\mathbf{y}} - \mathbf{F}_\ell\mathbf{a}\|_2) = o_{\mathbb{P}}(1) \end{aligned}$$

by equation 14, equation 15, equation 16, equation 17, and Theorem 3.2. Therefore, using the definition of $\hat{\mathbf{a}}(\mathbf{F}_\ell)$,

$$\begin{aligned} \frac{1}{n} \|\tilde{\mathbf{y}} - \mathbf{F}_\ell\hat{\mathbf{a}}(\mathbf{F}_\ell)\|_2^2 + \lambda \|\hat{\mathbf{a}}(\mathbf{F}_\ell)\|_2^2 &\leq \frac{1}{n} \|\tilde{\mathbf{y}} - \mathbf{F}_\ell\hat{\mathbf{a}}(\mathbf{F})\|_2^2 + \lambda \|\hat{\mathbf{a}}(\mathbf{F})\|_2^2 \\ &= \frac{1}{n} \|\tilde{\mathbf{y}} - \mathbf{F}\hat{\mathbf{a}}(\mathbf{F})\|_2^2 + \lambda \|\hat{\mathbf{a}}(\mathbf{F})\|_2^2 + o_{\mathbb{P}}(1) \end{aligned}$$

and using the definition of $\hat{\mathbf{a}}(\mathbf{F})$,

$$\begin{aligned} \frac{1}{n} \|\tilde{\mathbf{y}} - \mathbf{F}\hat{\mathbf{a}}(\mathbf{F})\|_2^2 + \lambda \|\hat{\mathbf{a}}(\mathbf{F})\|_2^2 &\leq \frac{1}{n} \|\tilde{\mathbf{y}} - \mathbf{F}\hat{\mathbf{a}}(\mathbf{F}_\ell)\|_2^2 + \lambda \|\hat{\mathbf{a}}(\mathbf{F}_\ell)\|_2^2 \\ &= \frac{1}{n} \|\tilde{\mathbf{y}} - \mathbf{F}_\ell\hat{\mathbf{a}}(\mathbf{F}_\ell)\|_2^2 + \lambda \|\hat{\mathbf{a}}(\mathbf{F}_\ell)\|_2^2 + o_{\mathbb{P}}(1). \end{aligned}$$

These together prove the theorem.

I Proof of Theorem 4.2

First, we will prove a general lemma regarding the equivalence of an augmented training loss. We will later use this result to prove the equivalence of the test error.

Lemma I.1. *Let $\eta \asymp n^\alpha$ with $\frac{\ell-1}{2\ell} < \alpha < \frac{\ell}{2\ell+2}$ for some $\ell \in \mathbb{N}$ and \mathbf{F}_ℓ be defined as in equation 4. For the test error \mathcal{L}_{te} from Section 4.1, define*

$$\mathcal{R}_\zeta(\mathbf{a}, \mathbf{F}) = \frac{1}{n} \|\tilde{\mathbf{y}} - \mathbf{F}\mathbf{a}\|_2^2 + \lambda \|\mathbf{a}\|_2^2 + \zeta \mathcal{L}_{\text{te}}(\mathbf{a}).$$

Then, for any $\lambda > 0$, $\zeta > 0$, we have

$$\left| \min_{\mathbf{a}} \mathcal{R}_\zeta(\mathbf{a}, \mathbf{F}_\ell) - \min_{\mathbf{a}} \mathcal{R}_\zeta(\mathbf{a}, \mathbf{F}) \right| = o(1), \quad (18)$$

with probability $1 - o(1)$.

Proof of Lemma I.1. Letting $\hat{\mathbf{a}}_\zeta(\mathbf{F}) = \arg \min_{\mathbf{a}} \mathcal{R}_\zeta(\mathbf{a}, \mathbf{F})$, we can write

$$\begin{aligned} \max \left\{ \frac{1}{n} \|\tilde{\mathbf{y}} - \mathbf{F}\hat{\mathbf{a}}_\zeta(\mathbf{F})\|_2^2, \lambda \|\hat{\mathbf{a}}_\zeta(\mathbf{F})\|_2^2, \zeta \mathcal{L}_{\text{te}}(\hat{\mathbf{a}}_\zeta(\mathbf{F})) \right\} &\leq \frac{1}{n} \|\tilde{\mathbf{y}} - \mathbf{F}\hat{\mathbf{a}}_\zeta(\mathbf{F})\|_2^2 + \lambda \|\hat{\mathbf{a}}_\zeta(\mathbf{F})\|_2^2 + \zeta \mathcal{L}_{\text{te}}(\hat{\mathbf{a}}_\zeta(\mathbf{F})) \\ &\leq \frac{1}{n} \|\tilde{\mathbf{y}} - \mathbf{F} \cdot \mathbf{0}\|_2^2 + \lambda \|\mathbf{0}\|_2^2 + \zeta \mathcal{L}_{\text{te}}(\mathbf{0}) = O_{\mathbb{P}}(1). \end{aligned}$$

Thus,

$$\mathcal{L}_{\text{te}}(\hat{\mathbf{a}}_\zeta(\mathbf{F})) = O_{\mathbb{P}}(1), \quad \|\hat{\mathbf{a}}_\zeta(\mathbf{F})\|_2 = O_{\mathbb{P}}(1), \quad \|\tilde{\mathbf{y}} - \mathbf{F}\hat{\mathbf{a}}_\zeta(\mathbf{F})\|_2 = O_{\mathbb{P}}(\sqrt{n}). \quad (19)$$

A similar argument gives

$$\mathcal{L}_{\text{te}}(\hat{\mathbf{a}}_\zeta(\mathbf{F}_\ell)) = O_{\mathbb{P}}(1), \quad \|\hat{\mathbf{a}}_\zeta(\mathbf{F}_\ell)\|_2 = O_{\mathbb{P}}(1), \quad \|\tilde{\mathbf{y}} - \mathbf{F}_\ell \hat{\mathbf{a}}_\zeta(\mathbf{F}_\ell)\|_2 = O_{\mathbb{P}}(\sqrt{n}). \quad (20)$$

Also by the triangle inequality, equation 20 and Theorem 3.2, which states $\|\mathbf{F}_\ell - \mathbf{F}\|_{\text{op}} = o_{\mathbb{P}}(\sqrt{n})$,

$$\begin{aligned} \|\tilde{\mathbf{y}} - \mathbf{F}\hat{\mathbf{a}}_\zeta(\mathbf{F}_\ell)\|_2 &\leq \|\tilde{\mathbf{y}} - \mathbf{F}_\ell \hat{\mathbf{a}}_\zeta(\mathbf{F}_\ell)\|_2 + \|(\mathbf{F}_\ell - \mathbf{F})\hat{\mathbf{a}}_\zeta(\mathbf{F}_\ell)\|_2 \\ &\leq \|\tilde{\mathbf{y}} - \mathbf{F}_\ell \hat{\mathbf{a}}_\zeta(\mathbf{F}_\ell)\|_2 + \|\mathbf{F}_\ell - \mathbf{F}\|_{\text{op}} \|\hat{\mathbf{a}}_\zeta(\mathbf{F}_\ell)\|_2 = O_{\mathbb{P}}(\sqrt{n}). \end{aligned} \quad (21)$$

Similarly, we can show that

$$\|\tilde{\mathbf{y}} - \mathbf{F}_\ell \hat{\mathbf{a}}_\zeta(\mathbf{F})\|_2 = O_{\mathbb{P}}(\sqrt{n}). \quad (22)$$

For $\mathbf{a} = \hat{\mathbf{a}}_\zeta(\mathbf{F})$ or $\mathbf{a} = \hat{\mathbf{a}}_\zeta(\mathbf{F}_\ell)$,

$$\begin{aligned} \frac{1}{n} (\|\tilde{\mathbf{y}} - \mathbf{F}\mathbf{a}\|_2^2 - \|\tilde{\mathbf{y}} - \mathbf{F}_\ell \mathbf{a}\|_2^2) &= \frac{1}{n} \langle (\mathbf{F}_\ell - \mathbf{F})\mathbf{a}, \tilde{\mathbf{y}} - \mathbf{F}\mathbf{a} + \tilde{\mathbf{y}} - \mathbf{F}_\ell \mathbf{a} \rangle \\ &\leq \frac{1}{n} \|\mathbf{F}_\ell - \mathbf{F}\|_{\text{op}} \|\mathbf{a}\|_2 (\|\tilde{\mathbf{y}} - \mathbf{F}\mathbf{a}\|_2 + \|\tilde{\mathbf{y}} - \mathbf{F}_\ell \mathbf{a}\|_2) = o_{\mathbb{P}}(1) \end{aligned}$$

by equation 19, equation 20, equation 21, equation 22, and Theorem 3.2. Therefore, using the definition of $\hat{\mathbf{a}}_\zeta(\mathbf{F}_\ell)$,

$$\begin{aligned} \frac{1}{n} \|\tilde{\mathbf{y}} - \mathbf{F}_\ell \hat{\mathbf{a}}_\zeta(\mathbf{F}_\ell)\|_2^2 + \lambda \|\hat{\mathbf{a}}_\zeta(\mathbf{F}_\ell)\|_2^2 + \zeta \mathcal{L}_{\text{te}}(\hat{\mathbf{a}}_\zeta(\mathbf{F}_\ell)) \\ \leq \frac{1}{n} \|\tilde{\mathbf{y}} - \mathbf{F}_\ell \hat{\mathbf{a}}_\zeta(\mathbf{F})\|_2^2 + \lambda \|\hat{\mathbf{a}}_\zeta(\mathbf{F})\|_2^2 + \zeta \mathcal{L}_{\text{te}}(\hat{\mathbf{a}}_\zeta(\mathbf{F})) \\ = \frac{1}{n} \|\tilde{\mathbf{y}} - \mathbf{F} \hat{\mathbf{a}}_\zeta(\mathbf{F})\|_2^2 + \lambda \|\hat{\mathbf{a}}_\zeta(\mathbf{F})\|_2^2 + \zeta \mathcal{L}_{\text{te}}(\hat{\mathbf{a}}_\zeta(\mathbf{F})) + o_{\mathbb{P}}(1), \end{aligned}$$

and using the definition of $\hat{\mathbf{a}}_\zeta(\mathbf{F})$,

$$\begin{aligned} & \frac{1}{n} \|\tilde{\mathbf{y}} - \mathbf{F}\hat{\mathbf{a}}_\zeta(\mathbf{F})\|_2^2 + \lambda \|\hat{\mathbf{a}}_\zeta(\mathbf{F})\|_2^2 + \zeta \mathcal{L}_{\text{te}}(\hat{\mathbf{a}}_\zeta(\mathbf{F})) \\ & \leq \frac{1}{n} \|\tilde{\mathbf{y}} - \mathbf{F}\hat{\mathbf{a}}_\zeta(\mathbf{F}_\ell)\|_2^2 + \lambda \|\hat{\mathbf{a}}_\zeta(\mathbf{F}_\ell)\|_2^2 + \zeta \mathcal{L}_{\text{te}}(\hat{\mathbf{a}}_\zeta(\mathbf{F}_\ell)) \\ & = \frac{1}{n} \|\tilde{\mathbf{y}} - \mathbf{F}_\ell \hat{\mathbf{a}}_\zeta(\mathbf{F}_\ell)\|_2^2 + \lambda \|\hat{\mathbf{a}}_\zeta(\mathbf{F}_\ell)\|_2^2 + \zeta \mathcal{L}_{\text{te}}(\hat{\mathbf{a}}_\zeta(\mathbf{F}_\ell)) + o_{\mathbb{P}}(1). \end{aligned}$$

Putting these together, we have

$$|\min_{\mathbf{a}} \mathcal{R}_\zeta(\mathbf{a}, \mathbf{F}_\ell) - \min_{\mathbf{a}} \mathcal{R}_\zeta(\mathbf{a}, \mathbf{F})| = o_{\mathbb{P}}(1), \quad (23)$$

which concludes the proof. \square

Now, we use this lemma to prove the equivalence of the test error.

Proof of Theorem 4.2. We will argue by contradiction. Assume that $\mathcal{L}_{\mathbf{F}} \neq \mathcal{L}_{\mathbf{F}_\ell}$ and let $\mathcal{L} = \frac{1}{2}(\mathcal{L}_{\mathbf{F}} + \mathcal{L}_{\mathbf{F}_\ell})$. Now, consider the following two optimization problems:

$$\mathcal{L}_1 = \min_{\mathcal{L}_{\text{te}}(\mathbf{a}) \leq \mathcal{L}} \frac{1}{n} \|\tilde{\mathbf{y}} - \mathbf{F}\mathbf{a}\|_2^2 + \lambda \|\mathbf{a}\|_2^2, \quad \mathcal{L}_2 = \min_{\mathcal{L}_{\text{te}}(\mathbf{a}) \leq \mathcal{L}} \frac{1}{n} \|\tilde{\mathbf{y}} - \mathbf{F}_\ell \mathbf{a}\|_2^2 + \lambda \|\mathbf{a}\|_2^2.$$

Without loss of generality, assume that $\mathcal{L}_{\mathbf{F}} < \mathcal{L}_{\mathbf{F}_\ell}$. The solution of the first optimization problem will still converge to $\mathcal{L}_{\text{tr}}(\mathbf{F})$ because $\mathcal{L}_{\mathbf{F}} < \mathcal{L}$. However, the solution of the second optimization problem will converge to a value greater than $\mathcal{L}_{\text{tr}}(\mathbf{F}_\ell)$, because $\mathcal{L}_{\mathbf{F}_\ell} > \mathcal{L}$ and the objective is λ -strongly convex. Note that by Theorem 4.1, we asymptotically have $\mathcal{L}_{\text{tr}}(\mathbf{F}_\ell) = \mathcal{L}_{\text{tr}}(\mathbf{F})$. Thus \mathcal{L}_1 and \mathcal{L}_2 converge to different quantities as $n \rightarrow \infty$. However, using the minimax theorem and since the objectives are λ -strongly convex, we can write

$$\begin{aligned} \mathcal{L}_1 &= \max_{\zeta > 0} -\zeta \mathcal{L} + \min_{\mathbf{a}} \left[\frac{1}{n} \|\tilde{\mathbf{y}} - \mathbf{F}\mathbf{a}\|_2^2 + \lambda \|\mathbf{a}\|_2^2 + \zeta \mathcal{L}_{\text{te}}(\mathbf{a}) \right], \\ \mathcal{L}_2 &= \max_{\zeta > 0} -\zeta \mathcal{L} + \min_{\mathbf{a}} \left[\frac{1}{n} \|\tilde{\mathbf{y}} - \mathbf{F}_\ell \mathbf{a}\|_2^2 + \lambda \|\mathbf{a}\|_2^2 + \zeta \mathcal{L}_{\text{te}}(\mathbf{a}) \right]. \end{aligned}$$

According to Lemma I.1, the two minima above converge to the same value for any fixed ζ . Note that, as functions of ζ , both maxima are concave as they are minima of linear functions of ζ . Hence, by using the concave version of [ASH19, Lemma 1], we have that \mathcal{L}_1 and \mathcal{L}_2 converge to the same value, which is a contradiction. \square

J Gaussian Equivalence Property

Gaussian equivalence results for non-linear random matrices were introduced in [EK10, CS13, FM19]. They have been repeatedly used in recent studies of random feature models [MM22, MRSY19, AP20a, AP20b, TAP21, GLR⁺22, MP21, dGSB21, LGC⁺21, LMH⁺23, HJ22, HL23, MS22]. Also, there has been progress on proving the Gaussian equivalence property for a multi-layer network with only the final layer trained [BPH23, CKZ23].

In more distantly related work in random matrix theory literature, the phenomenon that eigenvalue statistics in the bulk spectrum of a random matrix do not depend on the specific law of the matrix entries is referred to as “bulk universality” [Wig55, Gau61, Meh04, Dys62, EPR⁺10, EYY12, EK10, TV11].

[Erd19] shows that local spectral laws of correlated random Hermitian matrices can be fully determined by their first and second moments, through the matrix Dyson equation. Also, [BMP15, BNY20] show that spectral distributions of correlated symmetric random matrices can be characterized by Gaussian matrices with matching correlation structures.

In our case, we apply the Gaussian equivalence property to the following quantities for $p, q \in \mathbb{N}_0$ and $\beta_1, \beta_2 \in \{\beta, \beta_\star\}$: $H_p(\tilde{\mathbf{X}}\beta_1)^\top \tilde{\mathbf{R}}_0 H_q(\tilde{\mathbf{X}}\beta_2)$.

J.1 Proof Sketch of Gaussian Equivalence Property

In this section, we describe the proof idea of the Gaussian equivalence property. We use the Lindeberg exchange method [Lin22] in which we replace each column $\mathbf{g}_i = \sigma(\tilde{\mathbf{X}}\mathbf{w}_{0,i})$ of \mathbf{F}_0 with its Gaussian equivalent $\tilde{\mathbf{g}}_i = c_1 \tilde{\mathbf{X}}\mathbf{w}_{0,i} + c_{>1} \mathbf{z}_i$, $\mathbf{z}_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \mathbf{I}_n)$. Hereafter, we condition on all random variables except \mathbf{W}_0 . Then, $H_p(\tilde{\mathbf{X}}\beta_1)$ and $H_q(\tilde{\mathbf{X}}\beta_2)$ become deterministic vectors with $O(1)$ entries. We write $\mathbf{v} = H_p(\tilde{\mathbf{X}}\beta_1)$, $\mathbf{w} = H_q(\tilde{\mathbf{X}}\beta_2)$, and for all $i \in [N]$, $\mathbf{M}_i = \sum_{j=1}^{i-1} \tilde{\mathbf{g}}_j \mathbf{e}_j^\top + \sum_{j=i+1}^N \mathbf{g}_j \mathbf{e}_j^\top$.

Let $\tilde{\mathbf{F}}_0 = [\tilde{\mathbf{g}}_1 \cdots \tilde{\mathbf{g}}_N] \in \mathbb{R}^{n \times N}$ be the Gaussian equivalent of \mathbf{F}_0 and let $\tilde{\mathbf{R}}_0 = (\tilde{\mathbf{F}}_0 \tilde{\mathbf{F}}_0^\top + \lambda n \mathbf{I}_n)^{-1}$. By the triangle inequality,

$$\begin{aligned} & |\mathbf{v}^\top \mathbb{E} \mathbf{R}_0 \mathbf{w} - \mathbf{v}^\top \mathbb{E} \tilde{\mathbf{R}}_0 \mathbf{w}| \\ & \leq \sum_{i=1}^N |\mathbf{v}^\top \mathbb{E}(\mathbf{M}_i \mathbf{M}_i^\top + \mathbf{g}_i \mathbf{g}_i^\top + \lambda n \mathbf{I}_n)^{-1} \mathbf{w} - \mathbf{v}^\top \mathbb{E}(\mathbf{M}_i \mathbf{M}_i^\top + \tilde{\mathbf{g}}_i \tilde{\mathbf{g}}_i^\top + \lambda n \mathbf{I}_n)^{-1} \mathbf{w}|. \end{aligned} \quad (24)$$

Defining $\mathbf{S}_i = (\mathbf{M}_i \mathbf{M}_i^\top + \lambda n \mathbf{I}_n)^{-1}$, we have by the Sherman-Morrison formula that

$$\begin{aligned} (\mathbf{M}_i \mathbf{M}_i^\top + \mathbf{g}_i \mathbf{g}_i^\top + \lambda n \mathbf{I}_n)^{-1} &= \mathbf{S}_i - \frac{\mathbf{S}_i \mathbf{g}_i \mathbf{g}_i^\top \mathbf{S}_i}{1 + \mathbf{g}_i^\top \mathbf{S}_i \mathbf{g}_i}, \quad \text{and} \\ (\mathbf{M}_i \mathbf{M}_i^\top + \tilde{\mathbf{g}}_i \tilde{\mathbf{g}}_i^\top + \lambda n \mathbf{I}_n)^{-1} &= \mathbf{S}_i - \frac{\mathbf{S}_i \tilde{\mathbf{g}}_i \tilde{\mathbf{g}}_i^\top \mathbf{S}_i}{1 + \tilde{\mathbf{g}}_i^\top \mathbf{S}_i \tilde{\mathbf{g}}_i}. \end{aligned}$$

Thus,

$$\begin{aligned} & \mathbf{v}^\top \mathbb{E}(\mathbf{M}_i \mathbf{M}_i^\top + \mathbf{g}_i \mathbf{g}_i^\top + \lambda n \mathbf{I}_n)^{-1} \mathbf{w} - \mathbf{v}^\top \mathbb{E}(\mathbf{M}_i \mathbf{M}_i^\top + \tilde{\mathbf{g}}_i \tilde{\mathbf{g}}_i^\top + \lambda n \mathbf{I}_n)^{-1} \mathbf{w} \\ &= \mathbb{E} \frac{\mathbf{v}^\top \mathbf{S}_i \tilde{\mathbf{g}}_i \tilde{\mathbf{g}}_i^\top \mathbf{S}_i \mathbf{w}}{1 + \tilde{\mathbf{g}}_i^\top \mathbf{S}_i \tilde{\mathbf{g}}_i} - \mathbb{E} \frac{\mathbf{v}^\top \mathbf{S}_i \mathbf{g}_i \mathbf{g}_i^\top \mathbf{S}_i \mathbf{w}}{1 + \mathbf{g}_i^\top \mathbf{S}_i \mathbf{g}_i}. \end{aligned}$$

Let $\Sigma_g = \mathbb{E} \mathbf{g}_i \mathbf{g}_i^\top$ and $\Sigma_{\tilde{g}} = \mathbb{E} \tilde{\mathbf{g}}_i \tilde{\mathbf{g}}_i^\top$. By the Hanson-Wright concentration inequality, $\mathbf{g}_i^\top \mathbf{S}_i \mathbf{g}_i = \text{tr}(\mathbf{S}_i \Sigma_g) + O_{\mathbb{P}}(1/\sqrt{n})$ and $\tilde{\mathbf{g}}_i^\top \mathbf{S}_i \tilde{\mathbf{g}}_i = \text{tr}(\mathbf{S}_i \Sigma_{\tilde{g}}) + O_{\mathbb{P}}(1/\sqrt{n})$. Hence,

$$\begin{aligned} \mathbb{E} \frac{\mathbf{v}^\top \mathbf{S}_i \mathbf{g}_i \mathbf{g}_i^\top \mathbf{S}_i \mathbf{w}}{1 + \mathbf{g}_i^\top \mathbf{S}_i \mathbf{g}_i} &= \frac{\text{tr}(\mathbf{S}_i \mathbf{w} \mathbf{v}^\top \mathbf{S}_i \Sigma_g)}{1 + \text{tr}(\mathbf{S}_i \Sigma_g)} + O_{\mathbb{P}}(n^{-\frac{3}{2}}), \\ \mathbb{E} \frac{\mathbf{v}^\top \mathbf{S}_i \tilde{\mathbf{g}}_i \tilde{\mathbf{g}}_i^\top \mathbf{S}_i \mathbf{w}}{1 + \tilde{\mathbf{g}}_i^\top \mathbf{S}_i \tilde{\mathbf{g}}_i} &= \frac{\text{tr}(\mathbf{S}_i \mathbf{w} \mathbf{v}^\top \mathbf{S}_i \Sigma_{\tilde{g}})}{1 + \text{tr}(\mathbf{S}_i \Sigma_{\tilde{g}})} + O_{\mathbb{P}}(n^{-\frac{3}{2}}). \end{aligned}$$

Plugging into equation 24,

$$\begin{aligned} |\mathbf{v}^\top \bar{\mathbf{E}} \mathbf{R}_0 \mathbf{w} - \mathbf{v}^\top \mathbb{E} \tilde{\mathbf{R}}_0 \mathbf{w}| &\leq \sum_{i=1}^N \left| \frac{\text{tr}(\mathbf{S}_i \mathbf{w} \mathbf{v}^\top \mathbf{S}_i \Sigma_g)}{1 + \text{tr}(\mathbf{S}_i \Sigma_g)} - \frac{\text{tr}(\mathbf{S}_i \mathbf{w} \mathbf{v}^\top \mathbf{S}_i \Sigma_{\tilde{g}})}{1 + \text{tr}(\mathbf{S}_i \Sigma_{\tilde{g}})} \right| + O_{\mathbb{P}}(n^{-\frac{1}{2}}) \\ &\leq \sum_{i=1}^N \left(\frac{|\text{tr}(\mathbf{S}_i \mathbf{w} \mathbf{v}^\top \mathbf{S}_i (\Sigma_g - \Sigma_{\tilde{g}}))|}{1 + \text{tr}(\mathbf{S}_i \Sigma_g)} + \frac{|\text{tr}(\mathbf{S}_i \mathbf{w} \mathbf{v}^\top \mathbf{S}_i \Sigma_{\tilde{g}}) \text{tr}(\mathbf{S}_i (\Sigma_g - \Sigma_{\tilde{g}}))|}{(1 + \text{tr}(\mathbf{S}_i \Sigma_g))(1 + \text{tr}(\mathbf{S}_i \Sigma_{\tilde{g}}))} \right) + O_{\mathbb{P}}(n^{-\frac{1}{2}}). \end{aligned}$$

Now, we have

$$\begin{aligned} |\text{tr}(\mathbf{S}_i \mathbf{w} \mathbf{v}^\top \mathbf{S}_i (\Sigma_g - \Sigma_{\tilde{g}}))| &\leq \|\mathbf{S}_i \mathbf{v}\|_2 \|\mathbf{S}_i \mathbf{w}\|_2 \|\Sigma_g - \Sigma_{\tilde{g}}\|_{\text{op}} = O_{\mathbb{P}}(\|\Sigma_g - \Sigma_{\tilde{g}}\|_{\text{op}}/n), \\ |\text{tr}(\mathbf{S}_i \mathbf{w} \mathbf{v}^\top \mathbf{S}_i \Sigma_{\tilde{g}})| &\leq \|\mathbf{S}_i \mathbf{v}\|_2 \|\mathbf{S}_i \mathbf{w}\|_2 \|\Sigma_{\tilde{g}}\|_{\text{op}} = O_{\mathbb{P}}(1/n), \\ |\text{tr}(\mathbf{S}_i (\Sigma_g - \Sigma_{\tilde{g}}))| &\leq \|\mathbf{S}_i\|_F \|\Sigma_g - \Sigma_{\tilde{g}}\|_F \leq \sqrt{n} \|\mathbf{S}_i\|_F \|\Sigma_g - \Sigma_{\tilde{g}}\|_{\text{op}} = O_{\mathbb{P}}(\|\Sigma_g - \Sigma_{\tilde{g}}\|_{\text{op}}), \end{aligned}$$

where the first and second inequalities follow from the definition of the operator norm, and the last one follows from the Cauchy-Schwarz inequality.

By [EK10, Theorem 2.1], $\|\Sigma_g - \Sigma_{\tilde{g}}\|_{\text{op}} \rightarrow_P 0$. Therefore,

$$|\mathbf{v}^\top \bar{\mathbf{E}} \mathbf{R}_0 \mathbf{w} - \mathbf{v}^\top \mathbb{E} \tilde{\mathbf{R}}_0 \mathbf{w}| \leq O_{\mathbb{P}}(\|\Sigma_g - \Sigma_{\tilde{g}}\|_{\text{op}}) + O_{\mathbb{P}}(n^{-\frac{1}{2}}) \rightarrow_P 0.$$

K Proofs of Results from Section 4.2

Here, we will prove the results in Section 4.2. First, we will provide several lemmas, which will be used in our proofs. The first lemma allows us to approximate linear and quadratic forms of β in terms of β_* ; the quadratic form result is from [BES⁺22]. Its proof is in Section O.2.

Lemma K.1. *For any $d \in \mathbb{N}$, let $\mathbf{v} \in \mathbb{R}^d$ and $\mathbf{D} \in \mathbb{R}^{d \times d}$ be vectors and matrices, fixed or independent of $\mathbf{X}, \beta_*, \varepsilon_1, \dots, \varepsilon_n$, and satisfy $\|\mathbf{v}\|_2, \|\mathbf{D}\|_{\text{op}} \leq C$ almost surely, uniformly for some constant $C > 0$. Under Condition 2.1, we have*

$$\left| \mathbf{v}^\top \beta - c_{*,1} \mathbf{v}^\top \beta_* \right| \rightarrow 0, \quad \left| \beta^\top \mathbf{D} \beta - \frac{1}{n} (c_*^2 + \sigma_\varepsilon^2) \text{tr} \mathbf{D} - c_{*,1}^2 \beta_*^\top \mathbf{D} \beta_* \right| \rightarrow 0$$

in probability as $d \rightarrow \infty$.

We will use the expression derived for the training loss in the following lemma; see Section O.3 for the proof.

Lemma K.2. *The training loss $\mathcal{L}_{\text{tr}}(\mathbf{F})$ can be written as $\mathcal{L}_{\text{tr}}(\mathbf{F}) = \lambda \tilde{\mathbf{y}}^\top (\mathbf{F} \mathbf{F}^\top + \lambda n \mathbf{I}_n)^{-1} \tilde{\mathbf{y}}$.*

The following lemma will be used in proving concentration of certain quadratic forms appearing in the proofs; see Section O.4 for the proof.

Lemma K.3. *Let $g : \mathbb{R} \rightarrow \mathbb{R}$ be a polynomial, $\mathbf{D} \in \mathbb{R}^{n \times n}$ be a matrix with $\|\mathbf{D}\|_{\text{op}} = O_{\mathbb{P}}(1/n)$, and $\mathbf{Z} \in \mathbb{R}^n$ be a vector of i.i.d. Gaussian random variables with bounded variance independent of \mathbf{D} . We have*

$$\left| g(\mathbf{Z})^\top \mathbf{D} g(\mathbf{Z}) - \mathbb{E}[g(\mathbf{Z})^\top \mathbf{D} g(\mathbf{Z})] \right| \rightarrow_P 0,$$

in which g is applied elementwise.

The limiting values of two key quadratic forms appearing in the proof are derived in the following lemma, whose proof is deferred to Section [O.5](#).

Lemma K.4. *Let m_1 and m_2 be the solutions to the system of fixed point equations from equation [5](#). Then, the following holds:*

$$(a) \ \boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} = \psi(c_\star^2 + \sigma_\varepsilon^2)m_2 + \frac{\psi}{\phi}c_{\star,1}^2m_2 + o_{\mathbb{P}}(1) = \Theta_{\mathbb{P}}(1).$$

$$(b) \ \mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} - \|\mathbf{a}\|_2^2 = -\lambda \frac{\psi^2}{\phi^2}m_1 + \frac{\psi}{\phi} - 1 + o_{\mathbb{P}}(1) = \Theta_{\mathbb{P}}(1).$$

In particular, $\psi(c_\star^2 + \sigma_\varepsilon^2)m_2 + \frac{\psi}{\phi}c_{\star,1}^2m_2 \neq 0$ and $-\lambda \frac{\psi^2}{\phi^2}m_1 + \frac{\psi}{\phi} - 1 \neq 0$.

The following lemmas will be used in the computations. We defer the proofs of these lemmas to Sections [O.6](#), [O.7](#), [O.8](#), and [O.9](#) respectively.

Lemma K.5. *For any $p, q \in \mathbb{N}_0$, $p \neq q$ and any vector $\mathbf{u} \in \mathbb{R}^n$, with $\|\mathbf{u}\|_2 = 1$ independent of $\bar{\mathbf{R}}_0$, we have $H_q(\tilde{\mathbf{X}}\mathbf{u})^\top \bar{\mathbf{R}}_0 H_p(\tilde{\mathbf{X}}\mathbf{u}) = o_{\mathbb{P}}(1)$.*

Lemma K.6. *For any $p \in \mathbb{N}$, we have*

$$(a) \ \sqrt{N}H_p(\tilde{\mathbf{X}}\boldsymbol{\beta}_\star) \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}^{\circ 2} = o_{\mathbb{P}}(1),$$

$$(b) \ \sqrt{N}H_p(\tilde{\mathbf{X}}\boldsymbol{\beta}) \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}^{\circ 2} = o_{\mathbb{P}}(1).$$

Lemma K.7. *For $s \in \{1, 2\}$, $p \in \mathbb{N}$, and $p \neq s$, we have $H_p(\tilde{\mathbf{X}}\boldsymbol{\beta}_\star)^\top \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}}\boldsymbol{\beta})^{\circ s} = o_{\mathbb{P}}(1)$. Further,*

$$\lim_{n, N, d \rightarrow \infty} H_2(\tilde{\mathbf{X}}\boldsymbol{\beta}_\star)^\top \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}}\boldsymbol{\beta})^{\circ 2} = 2c_{\star,1}^2 \frac{\psi m_1}{\phi}.$$

Lemma K.8. *We have*

$$\lim_{n, N, d \rightarrow \infty} (\tilde{\mathbf{X}}\boldsymbol{\beta})^{\circ 2\top} \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}}\boldsymbol{\beta})^{\circ 2} = \frac{3\psi m_1}{\phi} [\phi(c_\star^2 + \sigma_\varepsilon^2) + c_{\star,1}^2]^2.$$

Now, we will first provide a proof of Theorem [4.3](#) in the case of $\ell = 1$ and $\ell = 2$ for a better insight into the proof techniques. We will then prove the general form in Section [L](#).

K.1 Proof for $\ell = 1$

In the $\ell = 1$ regime, due to Theorem [4.1](#), we can replace \mathbf{F} by \mathbf{F}_1 (defined in equation [4](#)) to compute the training loss. Hence, from now on we let $\mathbf{F} = \mathbf{F}_1$. We can write $\mathbf{F}\mathbf{F}^\top = \mathbf{F}_0\mathbf{F}_0^\top + \mathbf{U}\mathbf{K}\mathbf{U}^\top$ where $\mathbf{U} = [\mathbf{F}_0 \mathbf{a} \mid \tilde{\mathbf{X}}\boldsymbol{\beta}]$ and

$$\mathbf{K} = \begin{bmatrix} 0 & c_1^2 \eta \\ c_1^2 \eta & c_1^4 \eta^2 \|\mathbf{a}\|_2^2 \end{bmatrix}.$$

Based on Lemma [K.2](#), the training loss depends on $\bar{\mathbf{R}} = (\mathbf{F}\mathbf{F}^\top + \lambda n \mathbf{I}_n)^{-1}$. Using the Woodbury formula, this matrix can be written in terms of $\bar{\mathbf{R}}_0 = (\mathbf{F}_0\mathbf{F}_0^\top + \lambda n \mathbf{I}_n)^{-1}$ as

$$\bar{\mathbf{R}} = \bar{\mathbf{R}}_0 - \bar{\mathbf{R}}_0 \mathbf{U} (\mathbf{K}^{-1} + \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{U})^{-1} \mathbf{U}^\top \bar{\mathbf{R}}_0. \quad (25)$$

Defining $\mathbf{T} = (\mathbf{K}^{-1} + \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{U})^{-1} \in \mathbb{R}^{2 \times 2}$ and substituting $\bar{\mathbf{R}} = \bar{\mathbf{R}}_0 - \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0$ in the formula for training loss in Lemma K.2, we find

$$\mathcal{L}_{\text{tr}}(\mathbf{F}_0) - \mathcal{L}_{\text{tr}}(\mathbf{F}) = \lambda \tilde{\mathbf{y}}^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{y}}. \quad (26)$$

Using equation 26 and $\mathbf{U} = [\mathbf{F}_0 \mathbf{a} \mid \tilde{\mathbf{X}} \boldsymbol{\beta}]$, the loss difference can be written as

$$\begin{aligned} \mathcal{L}_{\text{tr}}(\mathbf{F}_0) - \mathcal{L}_{\text{tr}}(\mathbf{F}) \\ = \lambda \left[T_{11}(\tilde{\mathbf{y}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a})^2 + (T_{12} + T_{21}) \tilde{\mathbf{y}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} \cdot \mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{y}} + T_{22}(\tilde{\mathbf{y}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta})^2 \right], \end{aligned} \quad (27)$$

in which T_{ij} are the elements of the matrix \mathbf{T} . Using

$$\mathbf{T} = \frac{\begin{bmatrix} \boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} & -\frac{1}{c_1^2 \eta} - \mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} \\ -\frac{1}{c_1^2 \eta} - \boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} & \mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} - \|\mathbf{a}\|_2^2 \end{bmatrix}}{\left(\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} \right) \left(\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} - \|\mathbf{a}\|_2^2 \right) - \left(\frac{1}{c_1^2 \eta} + \mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} \right)^2}, \quad (28)$$

we will compute the limit of each term appearing in equation 27 separately:

Term 1. The first term can be written as

$$\delta_1 = \lambda T_{11}(\tilde{\mathbf{y}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a})^2 = \frac{\lambda (\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}) (\tilde{\mathbf{y}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a})^2}{\left(\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} \right) \left(\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} - \|\mathbf{a}\|_2^2 \right) - \left(\frac{1}{c_1^2 \eta} + \mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} \right)^2}.$$

Based on Lemma K.4, we know that $\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}$ and $\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} - \|\mathbf{a}\|_2^2$ are $\Theta_{\mathbb{P}}(1)$. Also, it can easily be seen that

$$\|\mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}\|_2 \leq \|\mathbf{F}_0\|_{\text{op}} \|\bar{\mathbf{R}}_0\|_{\text{op}} \|\tilde{\mathbf{X}} \boldsymbol{\beta}\|_2 = O_{\mathbb{P}}(1).$$

Hence,

$$\left(\frac{1}{c_1^2 \eta} + \mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} \right)^2 = o_{\mathbb{P}}(1) \quad (29)$$

because $\mathbf{a} \perp \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}$. Also, using that $\bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{F}_0^\top = (\mathbf{F}_0 \mathbf{F}_0^\top + \lambda n \mathbf{I}_n)^{-1} \mathbf{F}_0 \mathbf{F}_0^\top = \mathbf{I} - \lambda n \bar{\mathbf{R}}_0$,

$$\begin{aligned} (\tilde{\mathbf{y}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a})^2 &= \tilde{\mathbf{y}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbb{E}_{\mathbf{a}} [\mathbf{a} \mathbf{a}^\top] \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{y}} + o_{\mathbb{P}}(1) \\ &= \frac{1}{N} \tilde{\mathbf{y}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{y}} + o_{\mathbb{P}}(1) = \frac{1}{N} \tilde{\mathbf{y}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{y}} - \frac{\lambda n}{N} \tilde{\mathbf{y}}^\top \bar{\mathbf{R}}_0^2 \tilde{\mathbf{y}} + o_{\mathbb{P}}(1) = o_{\mathbb{P}}(1), \end{aligned}$$

where in the last inequality, we used that $\tilde{\mathbf{y}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{y}} \leq \frac{1}{\lambda n} \|\tilde{\mathbf{y}}\|_2^2 = O_{\mathbb{P}}(1)$ and $\tilde{\mathbf{y}}^\top \bar{\mathbf{R}}_0^2 \tilde{\mathbf{y}} \leq \frac{1}{(\lambda n)^2} \|\tilde{\mathbf{y}}\|_2^2 = o_{\mathbb{P}}(1)$. Putting everything together, it follows that $\delta_1 = o_{\mathbb{P}}(1)$ in probability.

Term 2 and Term 3. The second and third terms can be written as

$$\begin{aligned} \delta_2 = \delta_3 &= \lambda T_{12} \tilde{\mathbf{y}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} \mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{y}} \\ &= \frac{\lambda \left(-\frac{1}{c_1^2 \eta} - \mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} \right) (\tilde{\mathbf{y}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} \mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{y}})}{\left(\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} \right) \left(\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} - \|\mathbf{a}\|_2^2 \right) - \left(\frac{1}{c_1^2 \eta} + \mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} \right)^2}. \end{aligned}$$

Recall from the above argument that the denominator is $\Theta_{\mathbb{P}}(1)$ and that $\frac{1}{c_1^2 \eta} + \mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} = o_{\mathbb{P}}(1)$. Also, $\tilde{\mathbf{y}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} \mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{y}} \leq \frac{1}{(\lambda n)^2} \|\tilde{\mathbf{y}}\|_2^2 \|\tilde{\mathbf{X}} \boldsymbol{\beta}\|_2 \|\mathbf{a}\|_2 \|\mathbf{F}_0\|_{\text{op}} = O_{\mathbb{P}}(1)$. Therefore, we find $\delta_2 = \delta_3 = o_{\mathbb{P}}(1)$.

Term 4. This term can be written as

$$\begin{aligned}\delta_4 &= \lambda T_{22}(\tilde{\mathbf{y}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta})^2 \\ &= \frac{\lambda (\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} - \|\mathbf{a}\|_2^2) (\tilde{\mathbf{y}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta})^2}{\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} (\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} - \|\mathbf{a}\|_2^2) - \left(\frac{1}{c_1^2 \eta} + \mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}\right)^2} = \lambda \frac{(\tilde{\mathbf{y}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta})^2}{\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}} + o_{\mathbb{P}}(1),\end{aligned}$$

since $\frac{1}{c_1^2 \eta} + \mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} = o_{\mathbb{P}}(1)$ and $\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} - \|\mathbf{a}\|_2^2 = \Theta_{\mathbb{P}}(1) \neq 0$ by Lemma K.4. By equation 1 and Condition 2.4, we can write $\tilde{\mathbf{y}} = \sum_{p=1}^{\infty} c_{\star,p} H_p(\tilde{\mathbf{X}} \boldsymbol{\beta}_\star) + \boldsymbol{\varepsilon}$, where $\boldsymbol{\varepsilon} \in \mathbb{R}^n$ is additive Gaussian noise. Note that

$$\tilde{\mathbf{y}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} = c_{\star,1} \boldsymbol{\beta}_\star^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} + o_{\mathbb{P}}(1)$$

by Lemma K.7 and since $\|\bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}\|_2 = O_{\mathbb{P}}(1/\sqrt{n})$ and $\boldsymbol{\varepsilon} \perp \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}$.

Further by Lemma K.1,

$$c_{\star,1} \boldsymbol{\beta}_\star^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} = c_{\star,1}^2 \boldsymbol{\beta}_\star^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}_\star + o_{\mathbb{P}}(1) \rightarrow_P c_{\star,1}^2 \frac{\psi}{\phi} m_2, \quad (30)$$

where the final limit follows from the proof of Lemma K.4. By summing up the fours terms computed above and using Lemma K.4, we find

$$\mathcal{L}_{\text{tr}}(\mathbf{F}_0) - \mathcal{L}_{\text{tr}}(\mathbf{F}) \rightarrow_P \Delta_1 = \frac{\psi \lambda c_{\star,1}^4 m_2}{\phi [c_{\star,1}^2 + \phi(c_\star^2 + \sigma_\varepsilon^2)]} > 0, \quad (31)$$

which concludes the proof for $\ell = 1$.

K.2 Proof for $\ell = 2$

In the $\ell = 2$ regime, based on Theorem 4.1, we can replace \mathbf{F} with \mathbf{F}_2 (defined in equation 4) to compute the training loss. Hence, from now on we let $\mathbf{F} = \mathbf{F}_2$. We can write $\mathbf{F} \mathbf{F}^\top = \mathbf{F}_0 \mathbf{F}_0^\top + \mathbf{U} \mathbf{K} \mathbf{U}^\top$ where $\mathbf{U} = [\mathbf{F}_0 \mathbf{a} \mid \mathbf{F}_0 \mathbf{a}^{\circ 2} \sqrt{N} \mid \tilde{\mathbf{X}} \boldsymbol{\beta} \mid (\tilde{\mathbf{X}} \boldsymbol{\beta})^{\circ 2}]$ and

$$\mathbf{K} = \begin{bmatrix} 0 & 0 & c_1^2 \eta & 0 \\ 0 & 0 & 0 & c_1^2 c_2 \eta^2 / \sqrt{N} \\ c_1^2 \eta & 0 & c_1^4 \eta^2 \|\mathbf{a}\|_2^2 & c_1^4 c_2 \eta^3 \langle \mathbf{a}, \mathbf{a}^{\circ 2} \rangle \\ 0 & c_1^2 c_2 \eta^2 / \sqrt{N} & c_1^4 c_2 \eta^3 \langle \mathbf{a}^{\circ 2}, \mathbf{a} \rangle & c_1^4 c_2^2 \eta^4 \langle \mathbf{a}^{\circ 2}, \mathbf{a}^{\circ 2} \rangle \end{bmatrix}.$$

Recalling $\bar{\mathbf{R}} = (\mathbf{F} \mathbf{F}^\top + \lambda n \mathbf{I}_n)^{-1}$ and $\bar{\mathbf{R}}_0 = (\mathbf{F}_0 \mathbf{F}_0^\top + \lambda n \mathbf{I}_n)^{-1}$, we still have equation 25. Defining $\mathbf{T} = (\mathbf{K}^{-1} + \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{U})^{-1} \in \mathbb{R}^{4 \times 4}$, we have the following analogue to equation 26:

$$\mathcal{L}_{\text{tr}}(\mathbf{F}_0) - \mathcal{L}_{\text{tr}}(\mathbf{F}) = \lambda \tilde{\mathbf{y}}^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{y}}. \quad (32)$$

Denoting in what follows $\mathbf{Q} = \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0$, the inverse \mathbf{T}^{-1} can be written as follows:

$$\begin{bmatrix}
\mathbf{a}^\top \mathbf{Q} \mathbf{a} - \|\mathbf{a}\|_2^2 & N^{\frac{1}{2}} \mathbf{a}^\top (\mathbf{Q} - \mathbf{I}) \mathbf{a}^{\circ 2} & \mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}} + \frac{1}{c_1^2 \eta} & \mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ 2} \\
N^{\frac{1}{2}} \mathbf{a}^\top (\mathbf{Q} - \mathbf{I}) \mathbf{a}^{\circ 2} & N \mathbf{a}^{\circ 2 \top} \mathbf{Q} \mathbf{a}^{\circ 2} - N \|\mathbf{a}^{\circ 2}\|_2^2 & N^{\frac{1}{2}} \mathbf{a}^{\circ 2 \top} \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}} & N^{\frac{1}{2}} \mathbf{a}^{\circ 2 \top} \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ 2} + \frac{N^{\frac{1}{2}}}{c_1^2 c_2 \eta^2} \\
\tilde{\boldsymbol{\theta}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} + \frac{1}{c_1^2 \eta} & N^{\frac{1}{2}} \tilde{\boldsymbol{\theta}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}^{\circ 2} & \tilde{\boldsymbol{\theta}}^\top \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}} & \tilde{\boldsymbol{\theta}}^\top \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ 2} \\
\tilde{\boldsymbol{\theta}}^{\circ 2 \top} \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} & N^{\frac{1}{2}} \tilde{\boldsymbol{\theta}}^{\circ 2 \top} \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}^{\circ 2} + \frac{N^{\frac{1}{2}}}{c_1^2 c_2 \eta^2} & \tilde{\boldsymbol{\theta}}^{\circ 2 \top} \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}} & \tilde{\boldsymbol{\theta}}^{\circ 2 \top} \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ 2}
\end{bmatrix}. \quad (33)$$

K.2.1 Analysis of Terms in \mathbf{T}^{-1} and \mathbf{T}

In the following section, we will first analyze the elements of \mathbf{T}^{-1} :

(1,1): The term $\mathbf{a}^\top \mathbf{Q} \mathbf{a} - \|\mathbf{a}\|_2^2$ has already been analyzed in Lemma K.4 and is $\Theta_{\mathbb{P}}(1)$.

(1,2) and (2,1): Recalling $\mathbf{Q} = \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0$ and $\mathbf{R}_0 = (\mathbf{F}_0^\top \mathbf{F}_0 + \lambda n \mathbf{I}_N)^{-1}$, we can write

$$\begin{aligned}
[\mathbf{T}^{-1}]_{1,2} &= [\mathbf{T}^{-1}]_{2,1} = \sqrt{N} \mathbf{a}^\top \mathbf{Q} \mathbf{a}^{\circ 2} - \sqrt{N} \langle \mathbf{a}, \mathbf{a}^{\circ 2} \rangle \\
&= -\lambda n \sqrt{N} \mathbf{a}^\top \mathbf{R}_0 \mathbf{a}^{\circ 2} = -\lambda n \sqrt{N} \mathbf{a}^\top \mathbf{R}_0 (\mathbf{a}^{\circ 2} - 1/N \mathbf{1}_N + 1/N \mathbf{1}_N).
\end{aligned}$$

Introducing $\tilde{\mathbf{a}} = \sqrt{N} \mathbf{a} \sim \mathcal{N}(0, \mathbf{I}_N)$, and as $H_2(x) = x^2 - 1$ for all x , we find

$$[\mathbf{T}^{-1}]_{1,2} = [\mathbf{T}^{-1}]_{2,1} = -\frac{\lambda n}{N} \tilde{\mathbf{a}}^\top \mathbf{R}_0 H_2(\tilde{\mathbf{a}}) - \frac{\lambda n}{\sqrt{N}} \mathbf{a}^\top \mathbf{R}_0 \mathbf{1}_N.$$

The second term converges to zero as $n \rightarrow \infty$ because $\mathbf{a} \sim \mathcal{N}(0, \frac{1}{N} \mathbf{I}_N)$ is independent of \mathbf{R}_0 , and $\left\| \frac{n}{\sqrt{N}} \mathbf{R}_0 \mathbf{1}_N \right\|_2 = O_{\mathbb{P}}(1)$. Moreover, the first term also converges to zero; indeed,

$$\tilde{\mathbf{a}}^\top \mathbf{R}_0 H_2(\tilde{\mathbf{a}}) = \left(\frac{\tilde{\mathbf{a}}^\top + H_2(\tilde{\mathbf{a}})}{2} \right)^\top \mathbf{R}_0 \left(\frac{\tilde{\mathbf{a}}^\top + H_2(\tilde{\mathbf{a}})}{2} \right) - \left(\frac{\tilde{\mathbf{a}}^\top - H_2(\tilde{\mathbf{a}})}{2} \right)^\top \mathbf{R}_0 \left(\frac{\tilde{\mathbf{a}}^\top - H_2(\tilde{\mathbf{a}})}{2} \right).$$

Lemma K.3 can be used with $\mathbf{D} = \mathbf{R}$ to prove the concentration of both term around their expectation. Note that the expectation of $\tilde{\mathbf{a}}^\top \mathbf{R}_0 H_2(\tilde{\mathbf{a}})$ is zero because of the orthogonality property of Hermite polynomials and the independence of $\tilde{\mathbf{a}}$ and \mathbf{R}_0 . Putting everything together, we conclude that $[\mathbf{T}^{-1}]_{1,2} = [\mathbf{T}^{-1}]_{2,1} = o_{\mathbb{P}}(1)$.

(1,3) and (3,1): Recalling that $\tilde{\boldsymbol{\theta}} = \tilde{\mathbf{X}} \boldsymbol{\beta}$, it follows from equation 29 that this term is $o_{\mathbb{P}}(1)$.

(1,4) and (4,1): To bound $\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ 2}$, note that

$$\|\mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ 2}\|_{\text{op}} \leq \|\mathbf{F}_0\|_{\text{op}} \|\bar{\mathbf{R}}_0\|_{\text{op}} \|\tilde{\boldsymbol{\theta}}^{\circ 2}\|_2 = O_{\mathbb{P}}(1).$$

Hence, because $\mathbf{a} \sim \mathcal{N}(0, \frac{1}{N} \mathbf{I}_N)$ is independent of $\mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ 2}$, we have

$$[\mathbf{T}^{-1}]_{1,4} = [\mathbf{T}^{-1}]_{4,1} = \mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ 2} = o_{\mathbb{P}}(1).$$

(2,2): This term is $O_{\mathbb{P}}(1)$, because $\mathbf{a} \sim \mathcal{N}(0, \frac{1}{N} \mathbf{I}_N)$, so

$$\begin{aligned} [\mathbf{T}^{-1}]_{2,2} &= N \mathbf{a}^{\circ 2 \top} \mathbf{Q} \mathbf{a}^{\circ 2} - N \|\mathbf{a}^{\circ 2}\|_2^2 = -\lambda N n \mathbf{a}^{\circ 2 \top} (\mathbf{F}_0^\top \mathbf{F}_0 + \lambda n \mathbf{I}_N)^{-1} \mathbf{a}^{\circ 2} \\ &\leq \lambda N n \|\mathbf{a}^{\circ 2}\|_2^2 \cdot \|(\mathbf{F}_0^\top \mathbf{F}_0 + \lambda n \mathbf{I}_N)^{-1}\|_{\text{op}} = O_{\mathbb{P}}(1). \end{aligned}$$

(2,3) and (3,2): To bound $\sqrt{N} \mathbf{a}^{\circ 2 \top} \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}$, note that

$$\|\sqrt{N} \mathbf{a}^{\circ 2 \top} \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}}\|_2 \leq \|\sqrt{N} \mathbf{a}^{\circ 2}\|_2 \|\mathbf{F}_0\|_{\text{op}} \|\bar{\mathbf{R}}_0\|_{\text{op}} \|\tilde{\mathbf{X}}\|_{\text{op}} \leq C \cdot \sqrt{N} \cdot \frac{1}{n} \cdot \sqrt{N} = O_{\mathbb{P}}(1).$$

Also, by Lemma [K.1](#), we have

$$[\mathbf{T}^{-1}]_{2,3} = [\mathbf{T}^{-1}]_{3,2} = \sqrt{N} \mathbf{a}^{\circ 2 \top} \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} = c_{\star,1} \sqrt{N} \mathbf{a}^{\circ 2 \top} \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}_\star + o_{\mathbb{P}}(1),$$

which converges to zero, because $\boldsymbol{\beta}_\star \sim \mathcal{N}(0, \frac{1}{d} \mathbf{I}_d)$ and is independent of $\sqrt{N} \mathbf{a}^{\circ 2 \top} \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}}$, which has bounded norm in probability.

(2,4) and (4,2): First note that in the regime where $\ell = 2$, we have $\frac{\sqrt{N}}{\eta^2} \rightarrow 0$. Hence, we can write

$$[\mathbf{T}^{-1}]_{2,4} = \sqrt{N} (\tilde{\mathbf{X}} \boldsymbol{\beta})^{\circ 2 \top} \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}^{\circ 2} + o_{\mathbb{P}}(1) = \sqrt{N} H_2(\tilde{\mathbf{X}} \boldsymbol{\beta})^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}^{\circ 2} + \sqrt{N} \mathbf{1}_n^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}^{\circ 2} + o_{\mathbb{P}}(1). \quad (34)$$

By Lemma [K.6](#), the first term converges in probability to zero. Moreover, $\mathbf{a} \sim \mathcal{N}(0, \frac{1}{N} \mathbf{I}_N)$ is independent of $\bar{\mathbf{R}}_0 \mathbf{F}_0$, and $\|\mathbf{1}_n^\top \bar{\mathbf{R}}_0 \mathbf{F}_0\|_2 = O_{\mathbb{P}}(1)$. Thus, we have that $\sqrt{N} \mathbf{1}_n^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 (\mathbf{a}^{\circ 2} - 1/N \mathbf{1}_N) \rightarrow_P 0$. Hence, we find

$$[\mathbf{T}^{-1}]_{2,4} = \sqrt{N} \mathbf{1}_n^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{1}_N / N + o_{\mathbb{P}}(1).$$

Based on the Gaussian equivalence from Appendix [J](#), we can replace \mathbf{F}_0 with $\mathbf{F}_0 = c_1 \tilde{\mathbf{X}} \mathbf{W}_0^\top + c_{>1} \mathbf{Z}$, where $\mathbf{Z} \in \mathbb{R}^{n \times d}$ is an independent random matrix with $\mathcal{N}(0, 1)$ entries, without changing the limit. Now, the linearized \mathbf{F}_0 is left-orthogonally invariant, hence \mathbf{F}_0 has the same distribution as $\mathbf{O} \mathbf{F}_0$, where \mathbf{O} is uniformly distributed over the Haar measure of d -dimensional orthogonal matrices, independently of all other randomness. Hence, $N^{-1/2} \mathbf{1}_n^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{1}_N =_d N^{-1/2} \mathbf{1}_n^\top \mathbf{O} \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{1}_N$. Now, $\mathbf{O}^\top \mathbf{1}_n =_d \sqrt{n} \mathbf{z} / \|\mathbf{z}\|_2$, where $\mathbf{z} \sim \mathcal{N}(0, \mathbf{I}_n)$. Moreover $\|\mathbf{z}\|_2 = \sqrt{n}(1 + o_{\mathbb{P}}(1))$, hence replacing $\mathbf{O}^\top \mathbf{1}_n$ with \mathbf{z}^\top introduces negligible error. Hence, $[\mathbf{T}^{-1}]_{2,4} =_d N^{-1/2} \mathbf{z}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{1}_N + o_{\mathbb{P}}(1)$. Now, $\mathbf{z}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{1}_N \sim \mathcal{N}(0, \|\bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{1}_N\|_2^2)$, and $\|\bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{1}_N\|_2 = O_{\mathbb{P}}(1)$, thus $[\mathbf{T}^{-1}]_{2,4} \rightarrow_P 0$.

(3,3): We have $\|\tilde{\boldsymbol{\theta}}\|_2 = O_{\mathbb{P}}(\sqrt{N})$ and $\|\bar{\mathbf{R}}_0\|_{\text{op}} = O_{\mathbb{P}}(1/n)$. Thus, $[\mathbf{T}^{-1}]_{3,3} = O_{\mathbb{P}}(1)$.

(3,4) and (4,3): First, note that defining $\tilde{\boldsymbol{\beta}} = \frac{\boldsymbol{\beta}}{\|\boldsymbol{\beta}\|_2}$, and as $H_2(x) = x^2 - 1$ for all x , we can write

$$\begin{aligned} [\mathbf{T}^{-1}]_{3,4} &= [\mathbf{T}^{-1}]_{4,3} = \tilde{\boldsymbol{\theta}}^\top \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ 2} = \|\boldsymbol{\beta}\|_2^3 \left((\tilde{\mathbf{X}} \tilde{\boldsymbol{\beta}})^\top \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}} \tilde{\boldsymbol{\beta}})^{\circ 2} \right) \\ &= \|\boldsymbol{\beta}\|_2^3 \left((\tilde{\mathbf{X}} \tilde{\boldsymbol{\beta}})^\top \bar{\mathbf{R}}_0 H_2(\tilde{\mathbf{X}} \tilde{\boldsymbol{\beta}}) \right) + \|\boldsymbol{\beta}\|_2^2 \left(\tilde{\boldsymbol{\theta}}^\top \bar{\mathbf{R}}_0 \mathbf{1}_N \right). \end{aligned}$$

Now, by Lemma [K.1](#), we have $\tilde{\boldsymbol{\theta}}^\top \bar{\mathbf{R}}_0 \mathbf{1}_N = c_{\star,1} \tilde{\boldsymbol{\theta}}_\star^\top \bar{\mathbf{R}}_0 \mathbf{1}_N + o_{\mathbb{P}}(1)$. Now, note that $\|\tilde{\mathbf{X}} \bar{\mathbf{R}}_0 \mathbf{1}_N\|_2 = O_{\mathbb{P}}(1)$ and $\boldsymbol{\beta}_\star \sim \mathcal{N}(0, \frac{1}{d} \mathbf{I}_d)$ is independent of $\tilde{\mathbf{X}} \bar{\mathbf{R}}_0 \mathbf{1}_N$, which implies that the second term converges to zero. By using Lemma [K.5](#) for $\mathbf{u} = \tilde{\boldsymbol{\beta}}$, the first term also converges to zero. Putting these together, we have $[\mathbf{T}^{-1}]_{3,4} = [\mathbf{T}^{-1}]_{4,3} = o_{\mathbb{P}}(1)$.

(4,4): We have $\|\tilde{\boldsymbol{\theta}}^{\circ 2}\|_2 = O_{\mathbb{P}}(\sqrt{N})$ and $\|\bar{\mathbf{R}}_0\|_{\text{op}} = O_{\mathbb{P}}(1/n)$. Thus, $[\mathbf{T}^{-1}]_{4,4} = O_{\mathbb{P}}(1)$.

Now, putting everything together, the matrix \mathbf{T}^{-1} can be written as

$$\mathbf{T}^{-1} = \begin{bmatrix} [\mathbf{T}^{-1}]_{1,1} & 0 & 0 & 0 \\ 0 & [\mathbf{T}^{-1}]_{2,2} & 0 & 0 \\ 0 & 0 & [\mathbf{T}^{-1}]_{3,3} & 0 \\ 0 & 0 & 0 & [\mathbf{T}^{-1}]_{4,4} \end{bmatrix} + \Delta_1,$$

where the all elements of Δ_1 are $o_{\mathbb{P}}(1)$. Thus the matrix \mathbf{T} is equal to

$$\mathbf{T} = \begin{bmatrix} \frac{1}{[\mathbf{T}^{-1}]_{1,1}} & 0 & 0 & 0 \\ 0 & \frac{1}{[\mathbf{T}^{-1}]_{2,2}} & 0 & 0 \\ 0 & 0 & \frac{1}{[\mathbf{T}^{-1}]_{3,3}} & 0 \\ 0 & 0 & 0 & \frac{1}{[\mathbf{T}^{-1}]_{4,4}} \end{bmatrix} + \Delta_2, \quad (35)$$

where the all elements of Δ_2 are $o_{\mathbb{P}}(1)$.

K.2.2 Computing the training loss

Having computed the limit of the matrix \mathbf{T}^{-1} and \mathbf{T} , we are now ready to put everything together and compute the limiting train loss. One can write the outcome vector $\tilde{\mathbf{y}}$ as $\tilde{\mathbf{y}} = \sigma_{\star}(\tilde{\mathbf{X}}\boldsymbol{\beta}_{\star}) + \boldsymbol{\varepsilon}$, where $\boldsymbol{\varepsilon} \in \mathbb{R}^n$ is the noise term. Thus, using equation 32, we find

$$\begin{aligned} \mathcal{L}_{\text{tr}}(\mathbf{F}_0) - \mathcal{L}_{\text{tr}}(\mathbf{F}) &= \lambda \sigma_{\star}(\tilde{\mathbf{X}}\boldsymbol{\beta}_{\star})^{\top} \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^{\top} \bar{\mathbf{R}}_0 \sigma_{\star}(\tilde{\mathbf{X}}\boldsymbol{\beta}_{\star}) \\ &\quad + 2\lambda \sigma_{\star}(\tilde{\mathbf{X}}\boldsymbol{\beta}_{\star})^{\top} \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^{\top} \bar{\mathbf{R}}_0 \boldsymbol{\varepsilon} + \lambda \boldsymbol{\varepsilon}^{\top} \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^{\top} \bar{\mathbf{R}}_0 \boldsymbol{\varepsilon}. \end{aligned} \quad (36)$$

We will first argue the second and third term will go to zero in probability. To do this, we note that $\|\mathbf{T}\|_{\text{op}} = O_{\mathbb{P}}(1)$ and also $\|\mathbf{U}^{\top} \bar{\mathbf{R}}_0\|_2 \leq \|\mathbf{U}\|_{\text{op}} \|\bar{\mathbf{R}}_0\|_{\text{op}} = O_{\mathbb{P}}(1/\sqrt{n})$. We have $\boldsymbol{\varepsilon} \sim \mathcal{N}(0, \sigma_{\varepsilon}^2 \mathbf{I}_n)$ and it is independent of $\bar{\mathbf{R}}_0$, \mathbf{U} , \mathbf{T} , $\tilde{\mathbf{X}}$, and $\boldsymbol{\beta}_{\star}$. Also note that $\|\sigma_{\star}(\tilde{\mathbf{X}}\boldsymbol{\beta}_{\star})^{\top} \bar{\mathbf{R}}_0 \mathbf{U}\|_2 = O_{\mathbb{P}}(1)$. Thus, the second and third term in equation 36 go to zero and we have

$$\mathcal{L}_{\text{tr}}(\mathbf{F}_0) - \mathcal{L}_{\text{tr}}(\mathbf{F}) = \lambda \sigma_{\star}(\tilde{\mathbf{X}}\boldsymbol{\beta}_{\star})^{\top} \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^{\top} \bar{\mathbf{R}}_0 \sigma_{\star}(\tilde{\mathbf{X}}\boldsymbol{\beta}_{\star}) + o_{\mathbb{P}}(1).$$

If we expand $\sigma_{\star}(\tilde{\mathbf{X}}\beta_{\star})$ in the Hermite basis as $\sigma_{\star}(\tilde{\mathbf{X}}\beta_{\star}) = \sum_{p=1}^{\infty} c_{\star,p} H_p(\tilde{\theta}_{\star})$, we can write

$$\mathcal{L}_{\text{tr}}(\mathbf{F}_0) - \mathcal{L}_{\text{tr}}(\mathbf{F}) = \lambda \sum_{p,q=1}^{\infty} c_{\star,p} c_{\star,q} H_p(\tilde{\theta}_{\star})^{\top} \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^{\top} \bar{\mathbf{R}}_0 H_q(\tilde{\theta}_{\star}) + o_{\mathbb{P}}(1).$$

We define $\Delta_{p,q} = H_p(\tilde{\theta}_{\star})^{\top} \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^{\top} \bar{\mathbf{R}}_0 H_q(\tilde{\theta}_{\star}) = \delta_1^{p,q} + \delta_2^{p,q} + \delta_3^{p,q} + \delta_4^{p,q}$ in which, with $T_{i,j}$ being the (i,j) -th elements of the matrix \mathbf{T} ,

$$\begin{aligned} \delta_1^{p,q} &= T_{1,1} H_p(\tilde{\theta}_{\star})^{\top} \bar{\mathbf{R}}_0 (\mathbf{F}_0 \mathbf{a}) (\mathbf{F}_0 \mathbf{a})^{\top} \bar{\mathbf{R}}_0 H_q(\tilde{\theta}_{\star}) \\ &\quad + T_{1,2} H_p(\tilde{\theta}_{\star})^{\top} \bar{\mathbf{R}}_0 (\mathbf{F}_0 \mathbf{a}) (\sqrt{N} \mathbf{F}_0 \mathbf{a}^{\circ 2})^{\top} \bar{\mathbf{R}}_0 H_q(\tilde{\theta}_{\star}) \\ &\quad + T_{1,3} H_p(\tilde{\theta}_{\star})^{\top} \bar{\mathbf{R}}_0 (\mathbf{F}_0 \mathbf{a}) \tilde{\theta}^{\top} \bar{\mathbf{R}}_0 H_q(\tilde{\theta}_{\star}) \\ &\quad + T_{1,4} H_p(\tilde{\theta}_{\star})^{\top} \bar{\mathbf{R}}_0 (\mathbf{F}_0 \mathbf{a}) \tilde{\theta}^{\circ 2 \top} \bar{\mathbf{R}}_0 H_q(\tilde{\theta}_{\star}), \end{aligned} \quad (37)$$

$$\begin{aligned} \delta_2^{p,q} &= T_{2,1} H_p(\tilde{\theta}_{\star})^{\top} \bar{\mathbf{R}}_0 (\sqrt{N} \mathbf{F}_0 \mathbf{a}^{\circ 2}) (\mathbf{F}_0 \mathbf{a})^{\top} \bar{\mathbf{R}}_0 H_q(\tilde{\theta}_{\star}) \\ &\quad + T_{2,2} H_p(\tilde{\theta}_{\star})^{\top} \bar{\mathbf{R}}_0 (\sqrt{N} \mathbf{F}_0 \mathbf{a}^{\circ 2}) (\sqrt{N} \mathbf{F}_0 \mathbf{a}^{\circ 2})^{\top} \bar{\mathbf{R}}_0 H_q(\tilde{\theta}_{\star}) \\ &\quad + T_{2,3} H_p(\tilde{\theta}_{\star})^{\top} \bar{\mathbf{R}}_0 (\sqrt{N} \mathbf{F}_0 \mathbf{a}^{\circ 2}) \tilde{\theta}^{\top} \bar{\mathbf{R}}_0 H_q(\tilde{\theta}_{\star}) \\ &\quad + T_{2,4} H_p(\tilde{\theta}_{\star})^{\top} \bar{\mathbf{R}}_0 (\sqrt{N} \mathbf{F}_0 \mathbf{a}^{\circ 2}) \tilde{\theta}^{\circ 2 \top} \bar{\mathbf{R}}_0 H_q(\tilde{\theta}_{\star}), \end{aligned} \quad (38)$$

$$\begin{aligned} \delta_3^{p,q} &= T_{3,1} H_p(\tilde{\theta}_{\star})^{\top} \bar{\mathbf{R}}_0 \tilde{\theta} (\mathbf{F}_0 \mathbf{a})^{\top} \bar{\mathbf{R}}_0 H_q(\tilde{\theta}_{\star}) \\ &\quad + T_{3,2} H_p(\tilde{\theta}_{\star})^{\top} \bar{\mathbf{R}}_0 \tilde{\theta} (\sqrt{N} \mathbf{F}_0 \mathbf{a}^{\circ 2})^{\top} \bar{\mathbf{R}}_0 H_q(\tilde{\theta}_{\star}) \\ &\quad + T_{3,3} H_p(\tilde{\theta}_{\star})^{\top} \bar{\mathbf{R}}_0 \tilde{\theta} \tilde{\theta}^{\top} \bar{\mathbf{R}}_0 H_q(\tilde{\theta}_{\star}) \\ &\quad + T_{3,4} H_p(\tilde{\theta}_{\star})^{\top} \bar{\mathbf{R}}_0 \tilde{\theta} \tilde{\theta}^{\circ 2 \top} \bar{\mathbf{R}}_0 H_q(\tilde{\theta}_{\star}), \end{aligned} \quad (39)$$

and

$$\begin{aligned} \delta_4^{p,q} &= T_{4,1} H_p(\tilde{\theta}_{\star})^{\top} \bar{\mathbf{R}}_0 \tilde{\theta}^{\circ 2} (\mathbf{F}_0 \mathbf{a})^{\top} \bar{\mathbf{R}}_0 H_q(\tilde{\theta}_{\star}) \\ &\quad + T_{4,2} H_p(\tilde{\theta}_{\star})^{\top} \bar{\mathbf{R}}_0 \tilde{\theta}^{\circ 2} (\sqrt{N} \mathbf{F}_0 \mathbf{a}^{\circ 2})^{\top} \bar{\mathbf{R}}_0 H_q(\tilde{\theta}_{\star}) \\ &\quad + T_{4,3} H_p(\tilde{\theta}_{\star})^{\top} \bar{\mathbf{R}}_0 \tilde{\theta}^{\circ 2} \tilde{\theta}^{\top} \bar{\mathbf{R}}_0 H_q(\tilde{\theta}_{\star}) \\ &\quad + T_{4,4} H_p(\tilde{\theta}_{\star})^{\top} \bar{\mathbf{R}}_0 \tilde{\theta}^{\circ 2} \tilde{\theta}^{\circ 2 \top} \bar{\mathbf{R}}_0 H_q(\tilde{\theta}_{\star}). \end{aligned} \quad (40)$$

We will now look at each $\delta_i^{p,q}$ for $i \in \{1, 2, 3, 4\}$.

Term $\delta_1^{p,q}$: To prove that the term in equation 37 are asymptotically negligible, note that $\mathbf{a} \sim \mathcal{N}(0, \frac{1}{N} \mathbf{I}_N)$ is independent of $H_p(\tilde{\theta}_{\star}) \bar{\mathbf{R}}_0 \mathbf{F}_0$ and we have $\|H_p(\tilde{\theta}_{\star}) \bar{\mathbf{R}}_0 \mathbf{F}_0\|_2 = O_{\mathbb{P}}(1)$. Thus, $H_p(\tilde{\theta}_{\star}) \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} = o_{\mathbb{P}}(1)$ and all other terms multiplying this are $O_{\mathbb{P}}(1)$. This implies that for any $p, q \in \mathbb{N}$, we have $\delta_1^{p,q} = o_{\mathbb{P}}(1)$.

Term $\delta_2^{p,q}$: All four terms in equation 38 converge to zero. To prove this, we will use the Lemma K.6. In equation 38, all terms multiplied by $\sqrt{N} H_p(\tilde{\theta}_{\star}) \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}^{\circ 2}$ are $O_{\mathbb{P}}(1)$. Thus, $\delta_2^{p,q} = o_{\mathbb{P}}(1)$ for any $p, q \in \mathbb{N}$.

Term $\delta_3^{p,q}$: The first term in equation 39 converges to zero in probability due to an argument similar to the arguments used for $\delta_1^{p,q}$; and the same holds for the second term in equation 39, by arguing similarly as for $\delta_2^{p,q}$. We have shown that $T_{3,4} = o_{\mathbb{P}}(1)$, and by a norm argument, we can see that $H_p(\tilde{\boldsymbol{\theta}}_*)^\top \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}$ and $\tilde{\boldsymbol{\theta}}^{\circ 2\top} \bar{\mathbf{R}}_0 H_q(\tilde{\boldsymbol{\theta}}_*)$ are $O_{\mathbb{P}}(1)$. Hence,

$$\delta_3^{p,q} = T_{3,3} (H_p(\tilde{\boldsymbol{\theta}}_*)^\top \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}) (\tilde{\boldsymbol{\theta}}^{\top} \bar{\mathbf{R}}_0 H_q(\tilde{\boldsymbol{\theta}}_*)) + o_{\mathbb{P}}(1).$$

Term $\delta_4^{p,q}$: The first two terms in equation 40 converge to zero by the same reasoning used for $\delta_1^{p,q}$ and $\delta_2^{p,q}$, respectively. The third term can also be shown to converge to zero by recalling that $T_{4,3} = o_{\mathbb{P}}(1)$. Hence, we can write

$$\delta_4^{p,q} = T_{4,4} (H_p(\tilde{\boldsymbol{\theta}}_*)^\top \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ 2}) (\tilde{\boldsymbol{\theta}}^{\circ 2\top} \bar{\mathbf{R}}_0 H_q(\tilde{\boldsymbol{\theta}}_*)) + o_{\mathbb{P}}(1).$$

Putting everything together, we find

$$\begin{aligned} L_{\text{tr}}(\mathbf{F}_0) - L_{\text{tr}}(\mathbf{F}) &= \lambda T_{3,3} \sum_{p,q=1}^{\infty} c_{\star,p} c_{\star,q} (H_p(\tilde{\boldsymbol{\theta}}_*)^\top \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}) (\tilde{\boldsymbol{\theta}}^{\top} \bar{\mathbf{R}}_0 H_q(\tilde{\boldsymbol{\theta}}_*)) \\ &\quad + \lambda T_{4,4} \sum_{p,q=1}^{\infty} c_{\star,p} c_{\star,q} (H_p(\tilde{\boldsymbol{\theta}}_*)^\top \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ 2}) (\tilde{\boldsymbol{\theta}}^{\circ 2\top} \bar{\mathbf{R}}_0 H_q(\tilde{\boldsymbol{\theta}}_*)) + o_{\mathbb{P}}(1). \end{aligned}$$

Using Lemma K.7, we know that in the sums above, the terms corresponding to $(p, q) = (1, 1)$ and $(p, q) = (2, 2)$ are the only non-negligible terms in the first and second sum respectively.

Hence, as $T_{3,3} = 1/(\tilde{\boldsymbol{\theta}}^{\top} \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}) + o_{\mathbb{P}}(1)$ and $T_{4,4} = 1/(\tilde{\boldsymbol{\theta}}^{\circ 2\top} \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ 2}) + o_{\mathbb{P}}(1)$, from Lemmas K.1, K.4, K.7 and K.8, we can write,

$$\begin{aligned} \mathcal{L}_{\text{tr}}(\mathbf{F}) - \mathcal{L}_{\text{tr}}(\mathbf{F}_0) &= \lambda T_{3,3} c_{\star,1}^2 (\tilde{\boldsymbol{\theta}}^{\top} \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}})^2 + \lambda T_{4,4} c_{\star,2}^2 (H_2(\tilde{\boldsymbol{\theta}}_*)^\top \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ 2})^2 + o_{\mathbb{P}}(1) \\ &= \lambda \frac{c_{\star,1}^2 (\tilde{\boldsymbol{\theta}}^{\top} \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}})^2}{\tilde{\boldsymbol{\theta}}^{\top} \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}} + \lambda c_{\star,2}^2 \frac{(H_2(\tilde{\boldsymbol{\theta}}_*)^\top \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ 2})^2}{\tilde{\boldsymbol{\theta}}^{\circ 2\top} \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ 2}} + o_{\mathbb{P}}(1) \\ &\xrightarrow{P} \Delta_2 = \frac{\psi \lambda c_{\star,1}^4 m_2}{\phi [c_{\star,1}^2 + \phi(c_{\star}^2 + \sigma_{\varepsilon}^2)]} + \frac{4\psi \lambda c_{\star,1}^4 c_{\star,2}^2 m_1}{3\phi [\phi(c_{\star}^2 + \sigma_{\varepsilon}^2) + c_{\star,1}^2]^2}, \end{aligned}$$

proving the theorem for $\ell = 2$.

L Asymptotics of the Training Loss for General ℓ

We define the values $\xi_{i,j}$ for all $i, j \in \{0, 1, \dots\}$ such that for any $p \in \mathbb{N}$ and $x \in \mathbb{R}$, we have $x^p = \sum_{i=0}^p \xi_{p,i} H_i(x)$.

Theorem L.1. *Let $\ell \in \mathbb{N}$. If Conditions 2.1-2.4 hold, while we also have $c_1, \dots, c_{\ell} \neq 0$, and $\eta \asymp n^{\alpha}$ with $\frac{\ell-1}{2\ell} < \alpha < \frac{\ell}{2\ell+2}$, then for the learned feature map \mathbf{F} and the untrained feature map \mathbf{F}_0 , we have $\mathcal{L}_{\text{tr}}(\mathbf{F}_0) - \mathcal{L}_{\text{tr}}(\mathbf{F}) \xrightarrow{P} \Delta_{\ell} > 0$, where*

$$\Delta_{\ell} = \lambda \sum_{p=1}^{\ell} \sum_{q=1}^{\ell} c_{\star,p} c_{\star,q} r_p r_q \sum_{i=1}^{\ell} \sum_{j=1}^{\ell} \boldsymbol{\Omega}_{i,j} (\phi(c_{\star}^2 + \sigma_{\varepsilon}^2) + c_{\star,1}^2)^{(i+j)/2} \xi_{i,p} \xi_{j,q} + o_{\mathbb{P}}(1),$$

in which Ω is an invertible matrix with

$$[\Omega^{-1}]_{i,j} = (c_{\star,1}^2 + \phi(c_{\star}^2 + \sigma_{\varepsilon}^2))^{(i+j)/2} \frac{\psi}{\phi} \left[m_2 \xi_{i,1} \xi_{j,1} + m_1 \sum_{k=0, k \neq 1}^{\min(i,j)} k! \xi_{i,k} \xi_{j,k} \right], \quad \forall i, j \in [\ell],$$

and for $p \in \mathbb{N}$,

$$r_p = \begin{cases} \frac{p! \psi m_1}{\phi} \left(\frac{c_{\star,1}}{\sqrt{\phi(c_{\star}^2 + \sigma_{\varepsilon}^2) + c_{\star,1}^2}} \right)^p & p \neq 1 \\ \frac{\psi m_2}{\phi} \frac{c_{\star,1}}{\sqrt{\phi(c_{\star}^2 + \sigma_{\varepsilon}^2) + c_{\star,1}^2}} & p = 1 \end{cases}$$

Proof of Theorem L.1. In the regime where $\eta \asymp n^{\alpha}$ with $\frac{\ell-1}{2\ell} < \alpha < \frac{\ell}{2\ell+2}$, according to the equivalence theorem 4.1, we can replace \mathbf{F} with \mathbf{F}_{ℓ} when computing the limiting training loss. To compute the limiting training loss difference according to lemma K.2, we study the matrix $\bar{\mathbf{R}} = (\mathbf{F}\mathbf{F}^{\top} + \lambda n \mathbf{I}_n)^{-1}$. Due to equation 4, we can write

$$\begin{aligned} \mathbf{F}\mathbf{F}^{\top} &= \mathbf{F}_0\mathbf{F}_0^{\top} + \sum_{k=1}^{\ell} c_1^k c_k \eta^k \tilde{\boldsymbol{\theta}}^{\circ k} (\mathbf{F}_0 \mathbf{a}^{\circ k})^{\top} \\ &\quad + \sum_{k=1}^{\ell} c_1^k c_k \eta^k (\mathbf{F}_0 \mathbf{a}^{\circ k}) \tilde{\boldsymbol{\theta}}^{\circ k \top} + \sum_{j=1}^{\ell} \sum_{i=1}^{\ell} c_1^{i+j} c_i c_j \eta^{i+j} (\mathbf{a}^{\circ i})^{\top} (\mathbf{a}^{\circ j}) \tilde{\boldsymbol{\theta}}^{\circ i} \tilde{\boldsymbol{\theta}}^{\circ j \top}. \end{aligned}$$

Defining the matrix \mathbf{U} as

$$\mathbf{U} = \left[\underbrace{\mathbf{F}_0 \mathbf{a} \mid \cdots \mid N^{(\ell-1)/2} \mathbf{F}_0 \mathbf{a}^{\circ \ell}}_{\ell \text{ columns}} \mid \underbrace{\tilde{\boldsymbol{\theta}} \mid \cdots \mid \tilde{\boldsymbol{\theta}}^{\circ \ell}}_{\ell \text{ columns}} \right] \in \mathbb{R}^{n \times 2\ell},$$

we can write

$$\mathbf{F}\mathbf{F}^{\top} = \mathbf{F}_0\mathbf{F}_0^{\top} + \mathbf{U}\mathbf{K}\mathbf{U}^{\top}, \text{ in which } \mathbf{K} = \begin{bmatrix} \mathbf{0}_{\ell \times \ell} & \mathbf{K}_o \\ \mathbf{K}_o & \tilde{\mathbf{K}} \end{bmatrix} \in \mathbb{R}^{2\ell \times 2\ell},$$

where $\mathbf{K}_o = \text{diag} \left(\frac{c_1 c_1 \eta}{N^0}, \dots, \frac{c_1^\ell c_\ell \eta^\ell}{N^{(\ell-1)/2}} \right) \in \mathbb{R}^{\ell \times \ell}$, and $\tilde{\mathbf{K}} \in \mathbb{R}^{\ell \times \ell}$ with $[\tilde{\mathbf{K}}]_{i,j} = c_1^{i+j} c_i c_j \eta^{i+j} \langle \mathbf{a}^{\circ i}, \mathbf{a}^{\circ j} \rangle$, for all $i, j \in [\ell]$.

Using the Woodbury formula, the matrix $\bar{\mathbf{R}}$ can be written in terms of $\bar{\mathbf{R}}_0 = (\mathbf{F}_0\mathbf{F}_0^{\top} + \lambda n \mathbf{I}_n)^{-1}$ and $\mathbf{T} = (\mathbf{K}^{-1} + \mathbf{U}^{\top} \bar{\mathbf{R}}_0 \mathbf{U})^{-1} \in \mathbb{R}^{2\ell \times 2\ell}$ as $\bar{\mathbf{R}} = \bar{\mathbf{R}}_0 - \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^{\top} \bar{\mathbf{R}}_0$. Now

$$\mathbf{K}^{-1} = \begin{bmatrix} \hat{\mathbf{K}} & \mathbf{K}_o^{-1} \\ \mathbf{K}_o^{-1} & \mathbf{0}_{\ell \times \ell} \end{bmatrix}, \text{ where } \mathbf{K}_o^{-1} = \text{diag} \left(\frac{N^0}{c_1 c_1 \eta}, \dots, \frac{N^{\frac{\ell-1}{2}}}{c_1^\ell c_\ell \eta^\ell} \right),$$

and $[\hat{\mathbf{K}}]_{i,j} = -N^{(i-1)/2} N^{(j-1)/2} \langle \mathbf{a}^{\circ i}, \mathbf{a}^{\circ j} \rangle$, for all $i, j \in [\ell]$. We define $\mathbf{M}_1, \mathbf{M}_2, \mathbf{M}_o \in \mathbb{R}^{\ell \times \ell}$ as the following blocks of \mathbf{T}^{-1} :

$$\mathbf{T}^{-1} = \begin{bmatrix} \mathbf{M}_1 & \mathbf{M}_o \\ \mathbf{M}_o & \mathbf{M}_2 \end{bmatrix}.$$

Hence, we have

$$\begin{cases} [\mathbf{M}_1]_{i,j} = N^{(i-1)/2} N^{(j-1)/2} \mathbf{a}^{\circ i \top} (\mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 - \mathbf{I}) \mathbf{a}^{\circ j}, \\ [\mathbf{M}_o]_{i,j} = N^{(i-1)/2} \mathbf{a}^{\circ i \top} \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ j} + o_{\mathbb{P}}(1), \\ [\mathbf{M}_2]_{i,j} = \tilde{\boldsymbol{\theta}}^{\circ i \top} \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ j}. \end{cases}$$

We can expand the monomials in terms of the Hermite polynomials, for scalars $\xi_{i,k}$, $k \in [i]$, as follows:

$$(N^{1/2} \mathbf{a})^{\circ i} = \sum_{k=0}^i \xi_{i,k} H_k(N^{1/2} \mathbf{a}), \quad \text{and} \quad (\tilde{\mathbf{X}} \boldsymbol{\beta})^{\circ i} = \|\boldsymbol{\beta}\|_2^i \sum_{k=0}^i \xi_{i,k} H_k(\tilde{\mathbf{X}} \boldsymbol{\beta} / \|\boldsymbol{\beta}\|_2).$$

Using these, we will analyze each matrix $\mathbf{M}_1, \mathbf{M}_2, \mathbf{M}_o$ separately.

Analysis of \mathbf{M}_1 . It is easily seen that the elements of this matrix are $O_{\mathbb{P}}(1)$.

Analysis of \mathbf{M}_2 . To analyze these terms, we need the following lemma, whose proof is deferred to Section [O.10](#).

Lemma L.2. *For any $i, j \in \mathbb{N}_0$, we have*

$$(\tilde{\mathbf{X}} \boldsymbol{\beta})^{\circ i \top} \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}} \boldsymbol{\beta})^{\circ j} \rightarrow_P (c_{\star,1}^2 + \phi(c_{\star}^2 + \sigma_{\varepsilon}^2))^{(i+j)/2} \left[\xi_{i,1} \xi_{j,1} \frac{\psi m_2}{\phi} + \frac{\psi m_1}{\phi} \sum_{k=0, k \neq 1}^{\min(i,j)} k! \xi_{i,k} \xi_{j,k} \right].$$

Defining the matrix $\bar{\mathbf{M}}_2 \in \mathbb{R}^{\ell \times \ell}$ with entries

$$[\bar{\mathbf{M}}_2]_{i,j} = (c_{\star,1}^2 + \phi(c_{\star}^2 + \sigma_{\varepsilon}^2))^{(i+j)/2} \left[\xi_{i,1} \xi_{j,1} \frac{\psi m_2}{\phi} + \frac{\psi m_1}{\phi} \sum_{k=0, k \neq 1}^{\min(i,j)} k! \xi_{i,k} \xi_{j,k} \right],$$

for all $i, j \in [\ell]$, we have $[\mathbf{M}_2]_{i,j} \rightarrow_P [\bar{\mathbf{M}}_2]_{i,j}$. Note that we can write

$$\bar{\mathbf{M}}_2 = \frac{\psi}{\phi} \mathbf{B} \mathbf{Z} \mathbf{M} \mathbf{Z}^\top \mathbf{B} + \frac{\psi m_1}{\phi} \mathbf{e} \mathbf{e}^\top,$$

where we define $b = (c_{\star,1}^2 + \phi(c_{\star}^2 + \sigma_{\varepsilon}^2))^{1/2}$, $\mathbf{B} = \text{diag}(b^1, \dots, b^\ell) \in \mathbb{R}^{\ell \times \ell}$, $\mathbf{e} = \mathbf{B}[\xi_{1,0}, \dots, \xi_{\ell,0}]^\top$,

$$\mathbf{M} = \begin{bmatrix} 1! m_2 & 0 & \cdots & 0 \\ 0 & 2! m_1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \ell! m_1 \end{bmatrix} \in \mathbb{R}^{\ell \times \ell}, \quad \text{and} \quad \mathbf{Z} = \begin{bmatrix} \xi_{1,1} & \cdots & \xi_{1,\ell} \\ \vdots & \ddots & \vdots \\ \xi_{\ell,1} & \cdots & \xi_{\ell,\ell} \end{bmatrix} \in \mathbb{R}^{\ell \times \ell}.$$

Recalling that for all $i, j \in \{0, 1, \dots\}$, $\xi_{i,j}$ are such that such that for any $p \in \mathbb{N}$ and $x \in \mathbb{R}$, we have $x^p = \sum_{i=0}^p \xi_{p,i} H_i(x)$, it follows that the matrix \mathbf{Z} is lower-triangular with unit diagonal; hence invertible. Thus, since \mathbf{B}, \mathbf{M} are diagonal with positive entries, the matrix $\mathbf{B} \mathbf{Z} \mathbf{M} \mathbf{Z}^\top \mathbf{B}$ is positive definite. This implies that $\bar{\mathbf{M}}_2$ is invertible. We will denote $\boldsymbol{\Omega} = \bar{\mathbf{M}}_2^{-1}$.

Analysis of \mathbf{M}_o . We analyze $[\mathbf{M}_o]_{i,j}$ by writing $N^{(i-1)/2} \mathbf{a}^{\circ i}$ in the Hermite basis, finding

$$[\mathbf{M}_o]_{i,j} = \sum_{k=0}^i \frac{\xi_{i,k}}{\sqrt{N}} H_k(N^{1/2} \mathbf{a})^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ j} + o_{\mathbb{P}}(1).$$

The terms with $k > 0$ are all $o_{\mathbb{P}}(1)$ because $\frac{H_k(N^{1/2} \mathbf{a})}{\sqrt{N}}$ is a norm $O_{\mathbb{P}}(1)$ vector with mean zero, independent from the vector $\mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ j}$ with norm $O_{\mathbb{P}}(1)$. Thus, $[\mathbf{M}_o]_{i,j} = o_{\mathbb{P}}(1)$. The term with $k = 0$ can also be shown to be $o_{\mathbb{P}}(1)$ by using that the linearized \mathbf{F}_0 is left-orthogonally invariant, via an argument identical to the one used to analyze equation 34.

Hence, putting these together, the matrix \mathbf{T} can be written as

$$\mathbf{T} = \begin{bmatrix} \mathbf{M}_1^{-1} & \mathbf{0}_{\ell \times \ell} \\ \mathbf{0}_{\ell \times \ell} & \bar{\mathbf{M}}_2^{-1} \end{bmatrix} + o_{\mathbb{P}}(1).$$

Using lemma K.2, we can write the training loss difference as $\mathcal{L}_{\text{tr}}(\mathbf{F}_0) - \mathcal{L}_{\text{tr}}(\mathbf{F}) = \lambda \mathbf{y}^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{y}$. Plugging in the teacher function f_\star , we find

$$\begin{aligned} \mathcal{L}_{\text{tr}}(\mathbf{F}_0) - \mathcal{L}_{\text{tr}}(\mathbf{F}) &= \sum_{p,q} \lambda c_{\star,p} c_{\star,q} H_p(\tilde{\boldsymbol{\theta}}_\star)^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 H_q(\tilde{\boldsymbol{\theta}}_\star) \\ &\quad + 2\lambda \sum_p \left(c_{\star,p} H_p(\tilde{\boldsymbol{\theta}}_\star)^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \boldsymbol{\varepsilon} \right) + \lambda \boldsymbol{\varepsilon}^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \boldsymbol{\varepsilon}. \end{aligned}$$

Note that the second term can be shown to be $o_{\mathbb{P}}(1)$ because $\boldsymbol{\varepsilon} \sim \mathcal{N}(0, \sigma_\varepsilon^2 \mathbf{I}_n)$ and it is independent from $H_p(\tilde{\boldsymbol{\theta}}_\star)^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0$, and $\|H_p(\tilde{\boldsymbol{\theta}}_\star)^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0\|_{\text{op}} = O_{\mathbb{P}}(1/\sqrt{N})$ with a simple orderwise analysis. The third can also be shown to be $o_{\mathbb{P}}(1)$ by noting that $\boldsymbol{\varepsilon}$ is independent from $\bar{\mathbf{R}}_0 \mathbf{U}$, $\|\bar{\mathbf{R}}_0 \mathbf{U}\|_{\text{op}} = O_{\mathbb{P}}(1/\sqrt{n})$ and that the elements of \mathbf{T} are $O_{\mathbb{P}}(1)$.

To analyze the first term, we define $\delta_{p,q} = H_p(\tilde{\boldsymbol{\theta}}_\star)^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 H_q(\tilde{\boldsymbol{\theta}}_\star)$ for all non-negative integers p, q . To analyze such terms, we first expand $\mathbf{U} \mathbf{T} \mathbf{U}^\top$ as

$$\mathbf{U} \mathbf{T} \mathbf{U}^\top = \sum_{i=1}^{\ell} \sum_{j=1}^{\ell} N^{(i+j)/2-1} [\mathbf{M}_1^{-1}]_{i,j} (\mathbf{F}_0 \mathbf{a}^{\circ i}) (\mathbf{F}_0 \mathbf{a}^{\circ j})^\top + \sum_{i=1}^{\ell} \sum_{j=1}^{\ell} [\bar{\mathbf{M}}_2^{-1}]_{i,j} \tilde{\boldsymbol{\theta}}^{\circ i} \tilde{\boldsymbol{\theta}}^{\circ j \top}.$$

Thus, for any $p, q \in \mathbb{N}_0$, the terms $\delta_{p,q}$ can be written as

$$\begin{aligned} \delta_{p,q} &= \sum_{i=1}^{\ell} \sum_{j=1}^{\ell} N^{(i+j)/2-1} [\mathbf{M}_1^{-1}]_{i,j} H_p(\tilde{\boldsymbol{\theta}}_\star)^\top \bar{\mathbf{R}}_0 (\mathbf{F}_0 \mathbf{a}^{\circ i}) (\mathbf{F}_0 \mathbf{a}^{\circ j})^\top \bar{\mathbf{R}}_0 H_q(\tilde{\boldsymbol{\theta}}_\star) \\ &\quad + \sum_{i=1}^{\ell} \sum_{j=1}^{\ell} [\bar{\mathbf{M}}_2^{-1}]_{i,j} H_p(\tilde{\boldsymbol{\theta}}_\star)^\top \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ i} \tilde{\boldsymbol{\theta}}^{\circ j \top} \bar{\mathbf{R}}_0 H_q(\tilde{\boldsymbol{\theta}}_\star). \end{aligned}$$

By an argument identical to the argument for the terms in \mathbf{M}_o , the first sum goes to zero in probability. Denoting $\boldsymbol{\beta}/\|\boldsymbol{\beta}\|_2 := \tilde{\boldsymbol{\beta}}$, we can expand $(\tilde{\mathbf{X}} \boldsymbol{\beta})^{\circ i} = \|\boldsymbol{\beta}\|_2^i \sum_{k=0}^i \xi_{i,k} H_k(\tilde{\mathbf{X}} \boldsymbol{\beta}/\|\boldsymbol{\beta}\|_2)$. To analyze $\delta_{p,q}$, we need the following result, whose proof is deferred to Section O.11.

Lemma L.3. For any $p, q \in \mathbb{N}_0$, we have

$$H_p(\tilde{\mathbf{X}}\beta_\star)^\top \bar{\mathbf{R}}_0 H_q(\tilde{\mathbf{X}}\tilde{\beta}) \rightarrow_P \begin{cases} \frac{p! \psi m_1}{\phi} \left(\frac{c_{\star,1}}{\sqrt{\phi(c_\star^2 + \sigma_\varepsilon^2) + c_{\star,1}^2}} \right)^p & p = q \neq 1 \\ \frac{\psi m_2}{\phi} \frac{c_{\star,1}}{\sqrt{\phi(c_\star^2 + \sigma_\varepsilon^2) + c_{\star,1}^2}} & p = q = 1 \\ 0 & p \neq q. \end{cases}$$

We can now use Lemma L.3 and that $\|\beta\|_2 \rightarrow_P (\phi(c_\star^2 + \sigma_\varepsilon^2) + c_{\star,1}^2)^{1/2}$ to write

$$\begin{aligned} \delta_{p,q} &= \sum_{i=1}^{\ell} \sum_{j=1}^{\ell} [\bar{\mathbf{M}}_2^{-1}]_{i,j} \|\beta\|_2^{i+j} \xi_{i,p} \xi_{j,q} H_p(\tilde{\theta}_\star)^\top \bar{\mathbf{R}}_0 H_p(\tilde{\mathbf{X}}\tilde{\beta}) \cdot H_q(\tilde{\mathbf{X}}\tilde{\beta})^\top \bar{\mathbf{R}}_0 H_q(\tilde{\theta}_\star) + o_{\mathbb{P}}(1) \\ &= \sum_{i=1}^{\ell} \sum_{j=1}^{\ell} [\Omega]_{i,j} (\phi(c_\star^2 + \sigma_\varepsilon^2) + c_{\star,1}^2)^{(i+j)/2} \xi_{i,p} \xi_{j,q} r_p r_q + o_{\mathbb{P}}(1), \end{aligned}$$

for $p, q \in [\ell]$, which concludes the proof. \square

M Infinite sample limit

In the infinite sample limit, where $n \gg N, d$, we have $\phi \rightarrow 0$. In this extreme case, the expressions for m_1, m_2 will further simplify as $m_1, m_2 \rightarrow \phi/\lambda\psi$. Note that in this limit, we have $\mathcal{L}_{\text{tr}}(\mathbf{F}_0) \rightarrow \sigma_\varepsilon^2 + c_\star^2$ (see e.g., MM22, Section 6]. Using Corollary 4.4, we see that for example when $\ell = 2$, we have $\mathcal{L}(\mathbf{F}) \rightarrow \sigma_\varepsilon^2 + \frac{2c_{\star,2}^2}{3} + c_{\star,2}^2$. In particular, the term corresponding to the linear component of the teacher function in $\mathcal{L}(\mathbf{F}_0)$ cancels out with the corresponding term in Δ_2 .

N Proof of Theorem 4.5

Let $(\mathbf{x}_{\text{te}}, y_{\text{te}})$ follow the model from (1). Recall that the test error can be written as

$$\mathcal{L}_{\text{te}}(\hat{\mathbf{a}}(\mathbf{F})) = \mathbb{E}_{\mathbf{x}_{\text{te}}, y_{\text{te}}} (y_{\text{te}} - \hat{\mathbf{a}}^\top \sigma(\mathbf{W}\mathbf{x}_{\text{te}}))^2 = \mathbb{E}_{\mathbf{x}_{\text{te}}, y_{\text{te}}} (y_{\text{te}}^2) + \hat{\mathbf{a}}^\top \Sigma_f \hat{\mathbf{a}} - 2\hat{\mathbf{a}}^\top \mu_f, \quad (41)$$

where $\Sigma_f = \mathbb{E}_{\mathbf{x}_{\text{te}}} [\sigma(\mathbf{W}\mathbf{x}_{\text{te}})\sigma(\mathbf{W}\mathbf{x}_{\text{te}})^\top]$ and $\mu_f = \mathbb{E}_{\mathbf{x}_{\text{te}}, y_{\text{te}}} [y_{\text{te}}\sigma(\mathbf{W}\mathbf{x}_{\text{te}})]$. First, we will show that in the definition of Σ_f and μ_f , we can replace the test feature $\sigma(\mathbf{W}\mathbf{x}_{\text{te}})$ with the spiked approximation from Theorem 3.2

$$f_\ell = \sigma(\mathbf{W}_0 \mathbf{x}_{\text{te}}) + \sum_{k=1}^{\ell} c_1^k c_k \eta^k (\beta^\top \mathbf{x}_{\text{te}})^k \mathbf{a}^{\circ k}, \quad (42)$$

without changing the test error. To do this, consider an independent test set $\{\mathbf{x}_{\text{te},i}, y_{\text{te},i}\}_{i=1}^{n_{\text{te}}}$ with n_{te} test samples following the data generation distribution in equation 1 and we define $\mathbf{X} = [\mathbf{x}_{\text{te},1}, \dots, \mathbf{x}_{\text{te},n}]^\top \in \mathbb{R}^{n_{\text{te}} \times d}$. Let $\mathbf{F}_{\text{te}} = \sigma(\mathbf{X}_{\text{te}} \mathbf{W}^\top) \in \mathbb{R}^{n_{\text{te}} \times N}$ be the test feature matrix. We can write the test error as

$$\mathcal{L}_{\text{te}}(\hat{\mathbf{a}}(\mathbf{F})) = \lim_{n_{\text{te}} \rightarrow \infty} \frac{1}{n_{\text{te}}} \|\mathbf{y}_{\text{te}} - \mathbf{F}_{\text{te}} \hat{\mathbf{a}}(\mathbf{F})\|_2^2.$$

Now, consider the spiked approximation of the test feature matrix $\mathbf{F}_{\text{te},\ell} \in \mathbb{R}^{n_{\text{te}}, N}$ where each row of $\mathbf{F}_{\text{te},\ell}$ follows the approximation in equation 42. Using Theorem 3.2, we have $\|\mathbf{F}_{\text{te}} - \mathbf{F}_{\text{te},\ell}\|_{\text{op}} = o_{\mathbb{P}}(\sqrt{n_{\text{te}}})$. Thus,

$$\begin{aligned} \left| \frac{1}{n_{\text{te}}} \|\mathbf{y}_{\text{te}} - \mathbf{F}_{\text{te}} \hat{\mathbf{a}}(\mathbf{F})\|_2^2 - \frac{1}{n_{\text{te}}} \|\mathbf{y}_{\text{te}} - \mathbf{F}_{\text{te},\ell} \hat{\mathbf{a}}(\mathbf{F})\|_2^2 \right| &\leq \frac{1}{n_{\text{te}}} \left| \left\langle (\mathbf{F}_{\text{te}} - \mathbf{F}_{\text{te},\ell}) \hat{\mathbf{a}}(\mathbf{F}), 2\mathbf{y}_{\text{te}} - \mathbf{F}_{\text{te}} - \mathbf{F}_{\text{te},\ell} \right\rangle \right| \\ &\leq \frac{C}{\sqrt{n_{\text{te}}}} \|\mathbf{F}_{\text{te}} - \mathbf{F}_{\text{te},\ell}\|_{\text{op}} \cdot \|\hat{\mathbf{a}}(\mathbf{F})\|_2 \rightarrow_{\mathbb{P}} 0, \end{aligned}$$

where the last line is due to the fact that under the assumption of Theorem 4.2, we have $\|\mathbf{y}_{\text{te}} - \mathbf{F}_{\text{te}} \hat{\mathbf{a}}(\mathbf{F})\|_2^2 = O_{\mathbb{P}}(\sqrt{n_{\text{te}}})$, and that $\|\hat{\mathbf{a}}(\mathbf{F})\|_2 = O_{\mathbb{P}}(1)$ from the proof of Theorem 4.1. Thus in the test error, we can replace the test features with their spiked approximation without changing the limiting test error. With this, we can write

$$\begin{aligned} \Sigma_f &= \Sigma_f^0 + \mathbb{E}_{\mathbf{x}_{\text{te}}} \left[\sum_{k=1}^{\ell} c_1^k c_k \eta^k (\mathbf{x}_{\text{te}}^\top \boldsymbol{\beta})^k \left(\mathbf{a}^{\circ k} \sigma(\mathbf{W}_0 \mathbf{x}_{\text{te}})^\top + \sigma(\mathbf{W}_0 \mathbf{x}_{\text{te}}) \mathbf{a}^{\circ k \top} \right) \right] \\ &\quad + \mathbb{E}_{\mathbf{x}_{\text{te}}} \left[\sum_{i=1}^{\ell} \sum_{j=1}^{\ell} c_1^{i+j} \eta^{i+j} (\mathbf{x}_{\text{te}}^\top \boldsymbol{\beta})^{i+j} c_i c_j (\mathbf{a}^{\circ i} \mathbf{a}^{\circ j \top}) \right] \\ &= \Sigma_f^0 + \mathbb{E}_{\mathbf{x}_{\text{te}}} \left[\sum_{k=1}^{\ell} c_1^k c_k \eta^k \left(\mathbf{a}^{\circ k} \boldsymbol{\nu}_k^\top + \boldsymbol{\nu}_k \mathbf{a}^{\circ k \top} \right) \right] + \mathbb{E}_{\mathbf{x}_{\text{te}}} \left[\sum_{i=1}^{\ell} \sum_{j=1}^{\ell} c_1^{i+j} \eta^{i+j} c_i c_j \aleph_{i+j} (\mathbf{a}^{\circ i} \mathbf{a}^{\circ j \top}) \right], \quad (43) \end{aligned}$$

where $\Sigma_f^0 = \mathbb{E}_{\mathbf{x}_{\text{te}}} [\sigma(\mathbf{W}_0 \mathbf{x}_{\text{te}}) \sigma(\mathbf{W}_0 \mathbf{x}_{\text{te}})^\top]$, $\boldsymbol{\nu}_k = \mathbb{E}_{\mathbf{x}_{\text{te}}} [(\mathbf{x}_{\text{te}}^\top \boldsymbol{\beta})^k \sigma(\mathbf{W}_0 \mathbf{x}_{\text{te}})]$, and $\aleph_k = \mathbb{E}_{\mathbf{x}_{\text{te}}} (\mathbf{x}_{\text{te}}^\top \boldsymbol{\beta})^k$ for all $k \in [\ell]$. Also,

$$\boldsymbol{\mu}_f = \mathbb{E}_{\mathbf{x}_{\text{te}}, y_{\text{te}}} [y_{\text{te}} \sigma(\mathbf{W} \mathbf{x}_{\text{te}})] = \boldsymbol{\mu}_f^0 + \sum_{k=1}^{\ell} c_1^k c_k \tau_k \eta^k \mathbf{a}^{\circ k}, \quad (44)$$

where $\boldsymbol{\mu}_f^0 = \mathbb{E}_{\mathbf{x}_{\text{te}}, y_{\text{te}}} [y_{\text{te}} \sigma(\mathbf{W}_0 \mathbf{x}_{\text{te}})]$, and $\tau_k = \mathbb{E}_{\mathbf{x}_{\text{te}}} [y_{\text{te}} (\mathbf{x}_{\text{te}}^\top \boldsymbol{\beta})^k]$.

N.1 Proof for $\ell = 1$

Without loss of generality, assume that $c_1 = 1$. First, note that

$$\begin{aligned} \boldsymbol{\nu}_1 &= \mathbb{E}_{\mathbf{x}_{\text{te}}} [(\mathbf{x}_{\text{te}}^\top \boldsymbol{\beta}) \sigma(\mathbf{W}_0 \mathbf{x}_{\text{te}})] = \mathbf{W}_0 \boldsymbol{\beta}, \\ \aleph_1 &= \mathbb{E}_{\mathbf{x}_{\text{te}}} (\mathbf{x}_{\text{te}}^\top \boldsymbol{\beta}) = 0, \quad \text{and} \quad \aleph_2 = \mathbb{E}_{\mathbf{x}_{\text{te}}} (\mathbf{x}_{\text{te}}^\top \boldsymbol{\beta})^2 = \|\boldsymbol{\beta}\|_2^2 \rightarrow_P c_{\star,1}^2 + \phi(c_{\star}^2 + \sigma_{\varepsilon}^2), \\ \tau_1 &= \mathbb{E} [y_{\text{te}} (\mathbf{x}_{\text{te}}^\top \boldsymbol{\beta})] = c_{\star,1} \boldsymbol{\beta}^\top \boldsymbol{\beta} \rightarrow_P c_{\star,1}^2, \end{aligned}$$

where the convergence follows from equation 3 and the computations are in its proof.

Thus, using equation 43, we have

$$\boldsymbol{\mu}_f = \boldsymbol{\mu}_f^0 + c_{\star,1}^2 \eta \mathbf{a}, \quad \Sigma_f = \Sigma_f^0 + \eta \left(\mathbf{a} (\mathbf{W}_0 \boldsymbol{\beta})^\top + \mathbf{W}_0 \boldsymbol{\beta} \mathbf{a}^\top \right) + \eta^2 (c_{\star,1}^2 + \phi(c_{\star}^2 + \sigma_{\varepsilon}^2)) \mathbf{a} \mathbf{a}^\top.$$

From Section K.1, we have $\bar{\mathbf{R}} = \bar{\mathbf{R}}_0 - \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0$, where \mathbf{T} is defined in equation 28 and $\mathbf{U} = [\mathbf{F}_0 \mathbf{a} \mid \tilde{\mathbf{X}} \boldsymbol{\beta}]$. In Section K.1 it was shown that

$$\mathbf{T} = \frac{\begin{bmatrix} \boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} & -\frac{1}{\eta} - \mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} \\ -\frac{1}{\eta} - \boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} & \mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} - \|\mathbf{a}\|_2^2 \end{bmatrix}}{\left(\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} \right) (\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} - \|\mathbf{a}\|_2^2) - \left(\frac{1}{\eta} + \mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} \right)^2}, \quad (45)$$

where $T_{11}, T_{22} = \Theta_{\mathbb{P}}(1)$, and $T_{21}, T_{12} = O_{\mathbb{P}}(1/\eta)$.

Now, we are ready to study the terms $E_1 = \mathbb{E}_{\mathbf{x}_{\text{te}}, y_{\text{te}}}(y_{\text{te}}^2)$, $E_2 = \hat{\mathbf{a}}^\top \boldsymbol{\Sigma}_f \hat{\mathbf{a}}$, and $E_3 = -2\hat{\mathbf{a}}^\top \boldsymbol{\mu}_f$ that appear in the test error in equation 41.

N.1.1 Analysis of E_1 .

This term can be readily computed as

$$E_1 = \mathbb{E}(y_{\text{te}}^2) = \sigma_\varepsilon^2 + \mathbb{E}_{\mathbf{x}_{\text{te}}} \left(\sigma_\star(\boldsymbol{\beta}_\star^\top \mathbf{x}_{\text{te}}) \right)^2 = \sigma_\varepsilon^2 + c_\star^2. \quad (46)$$

N.1.2 Analysis of E_2 .

Recall that $\hat{\mathbf{a}} = (\mathbf{F}^\top \mathbf{F} + \lambda n \mathbf{I}_N)^{-1} \mathbf{F}^\top \tilde{\mathbf{y}}$, $\bar{\mathbf{R}}_0 = (\mathbf{F}_0^\top \mathbf{F}_0 + \lambda n \mathbf{I}_n)^{-1}$, and $\bar{\mathbf{R}} = (\mathbf{F} \mathbf{F}^\top + \lambda n \mathbf{I}_n)^{-1}$. Using these, we can write

$$\begin{aligned} E_2 &= \hat{\mathbf{a}}^\top \boldsymbol{\Sigma}_f \hat{\mathbf{a}} = \tilde{\mathbf{y}}^\top \mathbf{F} (\mathbf{F}^\top \mathbf{F} + \lambda n \mathbf{I}_N)^{-1} \boldsymbol{\Sigma}_f (\mathbf{F}^\top \mathbf{F} + \lambda n \mathbf{I}_N)^{-1} \mathbf{F}^\top \tilde{\mathbf{y}} \\ &= \tilde{\mathbf{y}}^\top (\mathbf{F} \mathbf{F}^\top + \lambda n \mathbf{I}_n)^{-1} \mathbf{F} \boldsymbol{\Sigma}_f \mathbf{F}^\top (\mathbf{F} \mathbf{F}^\top + \lambda n \mathbf{I}_n)^{-1} \tilde{\mathbf{y}} = \tilde{\mathbf{y}}^\top \bar{\mathbf{R}} \mathbf{F} \boldsymbol{\Sigma}_f \mathbf{F}^\top \bar{\mathbf{R}} \tilde{\mathbf{y}}. \end{aligned} \quad (47)$$

We have $\tilde{\mathbf{y}} = \mathbf{f}_\star + \boldsymbol{\varepsilon}$, thus

$$E_2 = \mathbf{f}_\star^\top \bar{\mathbf{R}} \mathbf{F} \boldsymbol{\Sigma}_f \mathbf{F}^\top \bar{\mathbf{R}} \mathbf{f}_\star + 2\boldsymbol{\varepsilon}^\top \bar{\mathbf{R}} \mathbf{F} \boldsymbol{\Sigma}_f \mathbf{F}^\top \bar{\mathbf{R}} \mathbf{f}_\star + \boldsymbol{\varepsilon}^\top \bar{\mathbf{R}} \mathbf{F} \boldsymbol{\Sigma}_f \mathbf{F}^\top \bar{\mathbf{R}} \boldsymbol{\varepsilon}. \quad (48)$$

We will now analyze the terms in equation 48. First note that $\boldsymbol{\varepsilon}^\top \bar{\mathbf{R}} \mathbf{F} \boldsymbol{\Sigma}_f \mathbf{F}^\top \bar{\mathbf{R}} \mathbf{f}_\star = o_{\mathbb{P}}(1)$ using a simple order-wise argument. To analyze the third term in equation 48, we write

$$\mathbf{F}^\top \bar{\mathbf{R}} \boldsymbol{\varepsilon} = \left(\mathbf{F}_0 + \eta(\tilde{\mathbf{X}} \boldsymbol{\beta}) \mathbf{a}^\top \right)^\top \bar{\mathbf{R}} \boldsymbol{\varepsilon} = \mathbf{F}_0^\top \bar{\mathbf{R}} \boldsymbol{\varepsilon} + \eta(\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}} \boldsymbol{\varepsilon}) \mathbf{a}.$$

Using a simple order-wise analysis, we have $\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}} \boldsymbol{\varepsilon} = O_{\mathbb{P}}(1/\sqrt{n})$. Thus, the third term can be written as

$$\boldsymbol{\varepsilon}^\top \bar{\mathbf{R}} \mathbf{F} \boldsymbol{\Sigma}_f \mathbf{F}^\top \bar{\mathbf{R}} \boldsymbol{\varepsilon} = \underbrace{\boldsymbol{\varepsilon}^\top \bar{\mathbf{R}} \mathbf{F}_0 \boldsymbol{\Sigma}_f \mathbf{F}_0^\top \bar{\mathbf{R}} \boldsymbol{\varepsilon}}_{q_1} + \underbrace{2\eta(\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}} \boldsymbol{\varepsilon}) \boldsymbol{\varepsilon}^\top \bar{\mathbf{R}} \mathbf{F}_0 \boldsymbol{\Sigma}_f \mathbf{a}}_{q_2} + \underbrace{\eta^2 (\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}} \boldsymbol{\varepsilon})^2 \mathbf{a}^\top \boldsymbol{\Sigma}_f \mathbf{a}}_{q_3}.$$

The term q_1 can be computed as

$$\begin{aligned} q_1 &= \boldsymbol{\varepsilon}^\top \bar{\mathbf{R}} \mathbf{F}_0 \boldsymbol{\Sigma}_f \mathbf{F}_0^\top \bar{\mathbf{R}} \boldsymbol{\varepsilon} = \boldsymbol{\varepsilon}^\top \bar{\mathbf{R}} \mathbf{F}_0 \left(\boldsymbol{\Sigma}_f^0 + \eta \left(\mathbf{a}(\mathbf{W}_0 \boldsymbol{\beta})^\top + \mathbf{W}_0 \boldsymbol{\beta} \mathbf{a}^\top \right) + \eta^2 \|\boldsymbol{\beta}\|_2^2 \mathbf{a} \mathbf{a}^\top \right) \mathbf{F}_0^\top \bar{\mathbf{R}} \boldsymbol{\varepsilon} \\ &= \boldsymbol{\varepsilon}^\top \bar{\mathbf{R}} \mathbf{F}_0 \boldsymbol{\Sigma}_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}} \boldsymbol{\varepsilon} + o_{\mathbb{P}}(1) = \boldsymbol{\varepsilon}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \boldsymbol{\Sigma}_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \boldsymbol{\varepsilon} + o_{\mathbb{P}}(1), \end{aligned}$$

where in the last line we have used that $\bar{\mathbf{R}} = \bar{\mathbf{R}}_0 - \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0$ and an order-wise analysis for various terms. To analyze q_2 , note that

$$\boldsymbol{\varepsilon}^\top \bar{\mathbf{R}} \mathbf{F}_0 \boldsymbol{\Sigma}_f \mathbf{a} = \boldsymbol{\varepsilon}^\top \bar{\mathbf{R}} \mathbf{F}_0 \left(\boldsymbol{\Sigma}_f^0 + \eta \left(\mathbf{a} (\mathbf{W}_0 \boldsymbol{\beta})^\top + \mathbf{W}_0 \boldsymbol{\beta} \mathbf{a}^\top \right) + \eta^2 \|\boldsymbol{\beta}\|_2^2 \mathbf{a} \mathbf{a}^\top \right) \mathbf{a} = o_{\mathbb{P}}(1),$$

using a simple order-wise analysis, thus $q_2 = o_{\mathbb{P}}(1)$. Similarly, for q_3 , we have

$$q_3 = \eta^2 (\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}} \boldsymbol{\varepsilon})^2 \mathbf{a}^\top \left(\boldsymbol{\Sigma}_f^0 + \eta \left(\mathbf{a} (\mathbf{W}_0 \boldsymbol{\beta})^\top + \mathbf{W}_0 \boldsymbol{\beta} \mathbf{a}^\top \right) + \eta^2 \|\boldsymbol{\beta}\|_2^2 \mathbf{a} \mathbf{a}^\top \right) \mathbf{a} = o_{\mathbb{P}}(1).$$

Hence, summing everything up

$$E_2 = \boldsymbol{\varepsilon}^\top \bar{\mathbf{R}} \mathbf{F}_0 \boldsymbol{\Sigma}_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \boldsymbol{\varepsilon} + \mathbf{f}_\star^\top \bar{\mathbf{R}} \mathbf{F} \boldsymbol{\Sigma}_f \mathbf{F}^\top \bar{\mathbf{R}} \mathbf{f}_\star + o_{\mathbb{P}}(1). \quad (49)$$

Next, we will study the term $\mathbf{f}_\star^\top \bar{\mathbf{R}} \mathbf{F} \boldsymbol{\Sigma}_f \mathbf{F}^\top \bar{\mathbf{R}} \mathbf{f}_\star$. We can write

$$\begin{aligned} \hat{\mathbf{a}} &= \mathbf{F}^\top \bar{\mathbf{R}} \mathbf{f}_\star = \left(\mathbf{F}_0 + \eta (\tilde{\mathbf{X}} \boldsymbol{\beta}) \mathbf{a}^\top \right)^\top \left(\bar{\mathbf{R}}_0 - \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \right) \mathbf{f}_\star \\ &= \underbrace{\mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star}_{p_1} - \underbrace{\mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star}_{p_2} + \underbrace{\eta \left[(\tilde{\mathbf{X}} \boldsymbol{\beta})^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star - (\tilde{\mathbf{X}} \boldsymbol{\beta})^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star \right] \mathbf{a}}_{p_3}. \end{aligned} \quad (50)$$

Thus, defining $E_2^{ij} = \mathbf{p}_i^\top \boldsymbol{\Sigma}_f \mathbf{p}_j$ for $i, j \in [3]$, we have $E_2 = \sum_{i,j=1}^3 E_2^{ij}$. In the following sections, we will analyze each term in this sum separately.

Preliminary Computations. Before starting the computation, we define K as

$$K := \boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star - \boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star. \quad (51)$$

Recalling that $T_{12} = T_{21} = O(1/\eta)$ and $\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} = O(1/\sqrt{n})$, the variable K can be simplified as follows:

$$K = \boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star - T_{22} (\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}) (\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star) + O_{\mathbb{P}} \left(\frac{1}{\eta \sqrt{n}} \right).$$

Also, from the definition of the matrix \mathbf{T} in equation 45, we have

$$T_{22} = \frac{1}{\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}} + \frac{1}{(\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta})^2 (\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} - \|\mathbf{a}\|_2^2)} \cdot \frac{1}{\eta^2} + o_{\mathbb{P}}(1/\eta^2).$$

Hence, putting everything together, and using Lemma K.4, we can write

$$K = - \frac{\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star}{(\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}) (\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} - \|\mathbf{a}\|_2^2)} \cdot \frac{1}{\eta^2} + o_{\mathbb{P}}(1/\eta^2). \quad (52)$$

Next, we will study the limit of ηT_{12} . For this, we can use equation 45 to write

$$T_{12} = \frac{-1}{(\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}) (\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} - \|\mathbf{a}\|_2^2)} \cdot \frac{1}{\eta} + O_{\mathbb{P}}(1/\eta^3). \quad (53)$$

Analysis of E_2^{11} . Noting that $\eta = o(n^{1/4})$ and $\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star = O_{\mathbb{P}}(1/\sqrt{n})$, we can simplify this term as follows:

$$\begin{aligned} E_2^{11} &= \mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star \\ &= \mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \left(\Sigma_f^0 + \eta \left(\mathbf{a}(\mathbf{W}_0 \boldsymbol{\beta})^\top + \mathbf{W}_0 \boldsymbol{\beta} \mathbf{a}^\top \right) + \eta^2 \|\boldsymbol{\beta}\|_2^2 \mathbf{a} \mathbf{a}^\top \right) \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star \\ &= \mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star + o_{\mathbb{P}}(1). \end{aligned} \quad (54)$$

Analysis of E_2^{12} and E_2^{21} . By expanding Σ_f and $\mathbf{U} \mathbf{T} \mathbf{U}^\top$, we have

$$\begin{aligned} E_2^{21} = E_2^{12} &= -\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star \\ &= -\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \left(\Sigma_f^0 + \eta \left(\mathbf{a}(\mathbf{W}_0 \boldsymbol{\beta})^\top + \mathbf{W}_0 \boldsymbol{\beta} \mathbf{a}^\top \right) + \eta^2 \|\boldsymbol{\beta}\|_2^2 \mathbf{a} \mathbf{a}^\top \right) \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star \\ &= E_2^{21(1)} + E_2^{21(2)} + E_2^{21(3)} + E_2^{21(4)}, \end{aligned}$$

in which

$$\begin{aligned} E_2^{21(1)} &= -\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star, \\ E_2^{21(2)} &= -\eta (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}) (\boldsymbol{\beta}^\top \mathbf{W}_0^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star), \\ E_2^{21(3)} &= -\eta (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{W}_0 \boldsymbol{\beta}) (\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star), \\ E_2^{21(4)} &= -\eta^2 (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}) (\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star). \end{aligned}$$

These terms can be simplified as follows. By expanding $\mathbf{U} \mathbf{T} \mathbf{U}^\top$, we have

$$\begin{aligned} E_2^{21(1)} &= -\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star \\ &= -\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \times \\ &\quad \left[T_{11}(\mathbf{F}_0 \mathbf{a})(\mathbf{F}_0 \mathbf{a})^\top + T_{12}(\mathbf{F}_0 \mathbf{a})(\tilde{\mathbf{X}} \boldsymbol{\beta})^\top + T_{21}(\tilde{\mathbf{X}} \boldsymbol{\beta})(\mathbf{F}_0 \mathbf{a})^\top + T_{22}(\tilde{\mathbf{X}} \boldsymbol{\beta})(\tilde{\mathbf{X}} \boldsymbol{\beta})^\top \right] \bar{\mathbf{R}}_0 \mathbf{f}_\star \\ &= -T_{22} (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}) (\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star) + o_{\mathbb{P}}(1), \end{aligned}$$

where the last line uses that $\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} = o_{\mathbb{P}}(1)$ and $\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} = o_{\mathbb{P}}(1)$. Next, noting that $\eta (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}) = o_{\mathbb{P}}(1)$ and $\boldsymbol{\beta}^\top \mathbf{W}_0^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star = \Theta_{\mathbb{P}}(1)$, we can write

$$E_2^{21(2)} = -\eta (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}) (\boldsymbol{\beta}^\top \mathbf{W}_0^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star) = o_{\mathbb{P}}(1).$$

Using a similar argument, we have

$$\begin{aligned} E_2^{21(3)} &= -\eta (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{W}_0 \boldsymbol{\beta}) (\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star) \\ &= -\eta T_{12} (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{W}_0 \boldsymbol{\beta}) (\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}) (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}) + o_{\mathbb{P}}(1). \end{aligned}$$

Finally, again note that $\eta^2 (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a})^2 = o_{\mathbb{P}}(1)$. Thus,

$$E_2^{21(4)} = -\eta^2 (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}) (\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star) = o_{\mathbb{P}}(1).$$

Putting all together, we arrive at

$$\begin{aligned} E_2^{21} = E_2^{12} &= -T_{22} (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}) (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}) \\ &\quad - \eta T_{12} (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{W}_0 \boldsymbol{\beta}) (\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}) (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}) + o_{\mathbb{P}}(1). \end{aligned} \quad (55)$$

Analysis of E_2^{22} . Once again, by expanding Σ_f and $\mathbf{U}\mathbf{T}\mathbf{U}^\top$, we have

$$\begin{aligned} E_2^{22} &= \mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star \\ &= \mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \left[T_{11}(\mathbf{F}_0 \mathbf{a})(\mathbf{F}_0 \mathbf{a})^\top + T_{12}(\mathbf{F}_0 \mathbf{a})(\tilde{\mathbf{X}} \boldsymbol{\beta})^\top + T_{21}(\tilde{\mathbf{X}} \boldsymbol{\beta})(\mathbf{F}_0 \mathbf{a})^\top + T_{22}(\tilde{\mathbf{X}} \boldsymbol{\beta})(\tilde{\mathbf{X}} \boldsymbol{\beta})^\top \right] \bar{\mathbf{R}}_0 \mathbf{F}_0 \\ &\quad \times \left[\Sigma_f^0 + \eta \left(\mathbf{a}(\mathbf{W}_0 \boldsymbol{\beta})^\top + \mathbf{W}_0 \boldsymbol{\beta} \mathbf{a}^\top \right) + \eta^2 \|\boldsymbol{\beta}\|_2^2 \mathbf{a} \mathbf{a}^\top \right] \\ &\quad \times \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \left[T_{11}(\mathbf{F}_0 \mathbf{a})(\mathbf{F}_0 \mathbf{a})^\top + T_{12}(\mathbf{F}_0 \mathbf{a})(\tilde{\mathbf{X}} \boldsymbol{\beta})^\top + T_{21}(\tilde{\mathbf{X}} \boldsymbol{\beta})(\mathbf{F}_0 \mathbf{a})^\top + T_{22}(\tilde{\mathbf{X}} \boldsymbol{\beta})(\tilde{\mathbf{X}} \boldsymbol{\beta})^\top \right] \bar{\mathbf{R}}_0 \mathbf{f}_\star. \end{aligned}$$

Note that $\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} = O_{\mathbb{P}}(1/\sqrt{n})$ and $\eta^2 = o(\sqrt{n})$. Hence,

$$\begin{aligned} E_2^{22} &= (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta})^2 \left[T_{21}(\mathbf{F}_0 \mathbf{a})^\top + T_{22}(\tilde{\mathbf{X}} \boldsymbol{\beta})^\top \right] \\ &\quad \times \bar{\mathbf{R}}_0 \mathbf{F}_0 \left[\Sigma_f^0 + \eta \left(\mathbf{a}(\mathbf{W}_0 \boldsymbol{\beta})^\top + \mathbf{W}_0 \boldsymbol{\beta} \mathbf{a}^\top \right) + \eta^2 \|\boldsymbol{\beta}\|_2^2 \mathbf{a} \mathbf{a}^\top \right] \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \\ &\quad \times \left[T_{12}(\mathbf{F}_0 \mathbf{a}) + T_{22}(\tilde{\mathbf{X}} \boldsymbol{\beta}) \right] + o_{\mathbb{P}}(1) \end{aligned} \quad (56)$$

$$\begin{aligned} E_2^{22} &= (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta})^2 \left[T_{22}^2 \left(\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} \right) \right. \\ &\quad \left. + 2\eta T_{12} T_{22} (\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}) (\boldsymbol{\beta}^\top \mathbf{W}_0^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}) + \|\boldsymbol{\beta}\|_2^2 \eta^2 T_{12}^2 (\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a})^2 \right] + o_{\mathbb{P}}(1). \end{aligned} \quad (57)$$

Analysis of E_2^{13} and E_2^{31} . Recalling the definition of K in equation 51, this term can be written as

$$\begin{aligned} E_2^{13} &= E_2^{31} = \eta K \mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f \mathbf{a} \\ &= \eta K \mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \left[\Sigma_f^0 + \eta \left(\mathbf{a}(\mathbf{W}_0 \boldsymbol{\beta})^\top + \mathbf{W}_0 \boldsymbol{\beta} \mathbf{a}^\top \right) + \eta^2 \|\boldsymbol{\beta}\|_2^2 \mathbf{a} \mathbf{a}^\top \right] \mathbf{a} \\ &= \eta^2 K (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{W}_0 \boldsymbol{\beta}) + o_{\mathbb{P}}(1), \end{aligned} \quad (58)$$

where the last line uses equation 52 and that $\eta \mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} = o_{\mathbb{P}}(1)$.

Analysis of E_2^{23} and E_2^{32} . Again, recalling the definition of K in equation 51, we have

$$\begin{aligned} E_2^{23} &= E_2^{32} = -\eta K \mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f \mathbf{a} \\ &= -\eta K \mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \left[T_{11}(\mathbf{F}_0 \mathbf{a})(\mathbf{F}_0 \mathbf{a})^\top + T_{12}(\mathbf{F}_0 \mathbf{a})(\tilde{\mathbf{X}} \boldsymbol{\beta})^\top + T_{21}(\tilde{\mathbf{X}} \boldsymbol{\beta})(\mathbf{F}_0 \mathbf{a})^\top + T_{22}(\tilde{\mathbf{X}} \boldsymbol{\beta})(\tilde{\mathbf{X}} \boldsymbol{\beta})^\top \right] \\ &\quad \times \left[\bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f^0 \mathbf{a} + \eta \left(\boldsymbol{\beta}^\top \mathbf{W}_0^\top \mathbf{a} \right) \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} + \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{W}_0 \boldsymbol{\beta} \right] + \eta^2 \|\boldsymbol{\beta}\|_2^2 \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} \\ &= E_2^{23(1)} + E_2^{23(2)} + E_2^{23(3)} + E_2^{23(4)}, \end{aligned}$$

in which each term can be written as follows:

$$\begin{aligned} E_2^{23(1)} &= -\eta K T_{11}(\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}) \left[\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f^0 \mathbf{a} + \eta (\mathbf{a}^\top \mathbf{W}_0 \boldsymbol{\beta}) (\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}) \right. \\ &\quad \left. + \eta (\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{W}_0 \boldsymbol{\beta}) + \eta^2 \|\boldsymbol{\beta}\|_2^2 (\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}) \right] = o_{\mathbb{P}}(1), \end{aligned}$$

where we have used that $K = \Theta_{\mathbb{P}}(1/\eta^2)$, $\mathbf{a}^\top \mathbf{W}_0 \boldsymbol{\beta} = o_{\mathbb{P}}(1)$, $\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{W}_0 \boldsymbol{\beta} = o_{\mathbb{P}}(1)$, and $\eta \mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{a} = o_{\mathbb{P}}(1)$. Also,

$$\begin{aligned} E_2^{23(2)} &= -\eta K T_{12}(\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}) \left[\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f^0 \mathbf{a} + \eta (\mathbf{a}^\top \mathbf{W}_0 \boldsymbol{\beta}) (\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}) \right. \\ &\quad \left. + \eta (\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{W}_0 \boldsymbol{\beta}) + \eta^2 \|\boldsymbol{\beta}\|_2^2 (\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}) \right] = o_{\mathbb{P}}(1), \end{aligned}$$

with a very similar argument to that for $E_2^{23(1)}$. Next,

$$\begin{aligned} E_2^{23(3)} &= -\eta K T_{21} (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}) \left[\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \boldsymbol{\Sigma}_f^0 \mathbf{a} + \eta (\mathbf{a}^\top \mathbf{W}_0 \boldsymbol{\beta}) (\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}) \right. \\ &\quad \left. + \eta (\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{W}_0 \boldsymbol{\beta}) + \eta^2 \|\boldsymbol{\beta}\|_2^2 (\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}) \right] \\ &= -\eta^3 K T_{21} \|\boldsymbol{\beta}\|_2^2 (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}) (\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}) + o_{\mathbb{P}}(1), \end{aligned}$$

in which we have used the fact that $T_{21} = \Theta_{\mathbb{P}}(1/\eta)$, and $K = \Theta(1/\eta^2)$. Finally, with a similar argument

$$\begin{aligned} E_2^{23(4)} &= -\eta K T_{22} (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}) \left[\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \boldsymbol{\Sigma}_f^0 \mathbf{a} + \eta (\mathbf{a}^\top \mathbf{W}_0 \boldsymbol{\beta}) (\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}) \right. \\ &\quad \left. + \eta (\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{W}_0 \boldsymbol{\beta}) + \eta^2 \|\boldsymbol{\beta}\|_2^2 (\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}) \right] \\ &= -\eta^2 K T_{22} (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}) (\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{W}_0 \boldsymbol{\beta}) + o_{\mathbb{P}}(1). \end{aligned}$$

Putting everything together, we have

$$\begin{aligned} E_2^{23} &= -\eta^3 K T_{21} \|\boldsymbol{\beta}\|_2^2 (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}) (\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}) \\ &\quad - \eta^2 K T_{22} (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}) (\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{W}_0 \boldsymbol{\beta}) + o_{\mathbb{P}}(1). \end{aligned} \quad (59)$$

Analysis of E_2^{33} . This term can be analyzed by expanding $\boldsymbol{\Sigma}_f$ as follows:

$$\begin{aligned} E_2^{33} &= \eta^2 K^2 \mathbf{a}^\top \left[\boldsymbol{\Sigma}_f^0 + \eta \left(\mathbf{a}(\mathbf{W}_0 \boldsymbol{\beta})^\top + \mathbf{W}_0 \boldsymbol{\beta} \mathbf{a}^\top \right) + \eta^2 \|\boldsymbol{\beta}\|_2^2 \mathbf{a} \mathbf{a}^\top \right] \mathbf{a} \\ &= \eta^2 K^2 \mathbf{a}^\top \boldsymbol{\Sigma}_f^0 \mathbf{a} + 2\eta^3 K^2 (\mathbf{a}^\top \mathbf{W}_0 \boldsymbol{\beta}) + \eta^4 K^2 \|\boldsymbol{\beta}\|_2^2 = \eta^4 K^2 \|\boldsymbol{\beta}\|_2^2 + o_{\mathbb{P}}(1), \end{aligned} \quad (60)$$

where we have used that $K = \Theta_{\mathbb{P}}(1/\eta^2)$ and $\eta \mathbf{a}^\top \mathbf{W}_0 \boldsymbol{\beta} = o_{\mathbb{P}}(1)$.

Putting Everything Together. Now, we can put together the results from previous sections to derive the limiting value of E_2 . First, we will explicitly derive the limit of each component. To do so, recall that using equation 3 and Lemma K.4, we have

$$\begin{aligned} \|\boldsymbol{\beta}\|_2^2 &\rightarrow_P [\phi(c_\star^2 + \sigma_\varepsilon^2) + c_{\star,1}^2], \quad \mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} \rightarrow_P c_{\star,1}^2 \psi m_2 / \phi \\ \boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} &\rightarrow_P [\phi(c_\star^2 + \sigma_\varepsilon^2) + c_{\star,1}^2] \psi m_2 / \phi, \quad \text{and} \quad \mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} \rightarrow_P \psi / \phi - \lambda \psi^2 m_1 / \phi^2. \end{aligned}$$

Also, using Lemma K.1, we have

$$\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \boldsymbol{\Sigma}_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} \rightarrow_P c_{\star,1}^2 M, \quad \text{and} \quad \boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \boldsymbol{\Sigma}_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta} \rightarrow_P [\phi(c_\star^2 + \sigma_\varepsilon^2) + c_{\star,1}^2] M,$$

in which $M := \lim_{n,N,d \rightarrow \infty} \boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \boldsymbol{\Sigma}_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}$. This limit has been computed in [AP20a]. Using the diagram in the proof of Lemma K.4 that shows how the notations of [AP20a] match ours, we find that E_{32} in [AP20a, S148] equals our M . Thus, we find

$$M = 1 - \frac{2m_2}{m_1} - \frac{m_2'}{m_1^2}, \quad (61)$$

where m_2' is the derivative of m_2 with respect to λ . Also, again using K.1, we have

$$\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{W}_0 \boldsymbol{\beta} \rightarrow_P c_{\star,1}^2 \bar{M}, \quad \text{and} \quad \boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{W}_0 \boldsymbol{\beta} \rightarrow_P [\phi(c_\star^2 + \sigma_\varepsilon^2) + c_{\star,1}^2] \bar{M},$$

in which $\bar{M} := \lim_{n,N,d \rightarrow \infty} \beta_\star^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{W}_0 \beta_\star$. This limit has been computed in [BES⁺22]. Specifically, using T_9^0 in (C.16), and noting that their Φ_0 translates to \mathbf{F}_0/\sqrt{N} , their \mathbf{W}_0 translates to \mathbf{W}_0^\top , and their $\tilde{\lambda}$ translates to $\frac{\lambda n}{N}$ in our notation, using [BES⁺22, Proposition 29] we find

$$\bar{M} = 1 - \frac{m_2}{m_1}. \quad (62)$$

Now, we can use equation 55, equation 56, equation 58, equation 59, and equation 60, respectively, to derive the following expressions:

$$\begin{aligned} E_2^{21} = E_2^{12} &= \frac{c_{\star,1}^4}{\phi(c_\star^2 + \sigma_\varepsilon^2) + c_{\star,1}^2} \left[-M + \bar{M} \left(\frac{\psi/\phi - \lambda\psi^2 m_1/\phi^2}{\psi/\phi - \lambda\psi^2 m_1/\phi^2 - 1} \right) \right], \\ E_2^{22} &= \frac{c_{\star,1}^4}{\phi(c_\star^2 + \sigma_\varepsilon^2) + c_{\star,1}^2} \left[M - 2\bar{M} \left(\frac{\psi/\phi - \lambda\psi^2 m_1/\phi^2}{\psi/\phi - \lambda\psi^2 m_1/\phi^2 - 1} \right) + \left(\frac{\psi/\phi - \lambda\psi^2 m_1/\phi^2}{\psi/\phi - \lambda\psi^2 m_1/\phi^2 - 1} \right)^2 \right], \\ E_2^{13} = E_2^{31} &= -\frac{c_{\star,1}^4}{\phi(c_\star^2 + \sigma_\varepsilon^2) + c_{\star,1}^2} \left[\frac{\bar{M}}{\psi/\phi + \lambda\psi^2/\phi^2 - 1} \right], \\ E_2^{23} = E_2^{32} &= \frac{c_{\star,1}^4}{\phi(c_\star^2 + \sigma_\varepsilon^2) + c_{\star,1}^2} \left[-\frac{\psi/\phi - \lambda\psi^2 m_1/\phi^2}{(\psi/\phi - \lambda\psi^2 m_1/\phi^2 - 1)^2} + \frac{\bar{M}}{\psi/\phi - \lambda\psi^2 m_1/\phi^2 - 1} \right], \\ E_2^{33} &= \frac{c_{\star,1}^4}{\phi(c_\star^2 + \sigma_\varepsilon^2) + c_{\star,1}^2} \left[\frac{1}{(\psi/\phi - \lambda\psi^2 m_1/\phi^2 - 1)^2} \right]. \end{aligned}$$

Thus, summing these terms up, we conclude that

$$E_2 - \tilde{\mathbf{y}}^\top \bar{\mathbf{R}}_0 \mathbf{F} \Sigma_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{y}} \rightarrow_P \frac{c_{\star,1}^4 (1 - M)}{\phi(c_\star^2 + \sigma_\varepsilon^2) + c_{\star,1}^2}, \quad (63)$$

wrapping up the derivation of the limiting value of E_2 .

N.1.3 Analysis of E_3 .

To analyze this term, first note that

$$E_3 = -2\hat{\mathbf{a}}^\top \boldsymbol{\mu}_f = -2\tilde{\mathbf{y}}^\top \bar{\mathbf{R}} \mathbf{F} \boldsymbol{\mu}_f = -2\mathbf{f}_\star^\top \bar{\mathbf{R}} \mathbf{F} \boldsymbol{\mu}_f - 2\varepsilon^\top \mathbf{F} \bar{\mathbf{R}} \boldsymbol{\mu}_f.$$

With a simple order-wise analysis, the second term can be shown to be $o_{\mathbb{P}}(1)$. To analyze the first term, we will again use the decomposition from equation 50. We can write $-2\mathbf{f}_\star^\top \bar{\mathbf{R}} \mathbf{F} \boldsymbol{\mu}_f = E_3^{(1)} + E_3^{(2)} + E_3^{(3)}$, in which $E_3^{(i)} = -2\mathbf{p}_i^\top \boldsymbol{\mu}_f$. We will analyze these terms separately. For the first term, we have

$$E_3^{(1)} = -2\mathbf{p}_1^\top \boldsymbol{\mu}_f = -2\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 (c_{\star,1} \mathbf{W}_0 \boldsymbol{\beta}_\star + c_{\star,1}^2 \eta \mathbf{a}) = -2c_{\star,1} \mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{W}_0 \boldsymbol{\beta}_\star + o_{\mathbb{P}}(1),$$

where we have used that $\eta \mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} = o_{\mathbb{P}}(1)$. Similarly, the second term can be written as

$$\begin{aligned}
E_3^{(2)} &= -2 \mathbf{p}_2^\top \boldsymbol{\mu}_f = 2 \mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 (c_{\star,1} \mathbf{W}_0 \boldsymbol{\beta}_\star + c_{\star,1}^2 \eta \mathbf{a}) \\
&= 2 \mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \left[T_{11}(\mathbf{F}_0 \mathbf{a})(\mathbf{F}_0 \mathbf{a})^\top + T_{12}(\mathbf{F}_0 \mathbf{a})(\tilde{\mathbf{X}} \boldsymbol{\beta})^\top + T_{21}(\tilde{\mathbf{X}} \boldsymbol{\beta})(\mathbf{F}_0 \mathbf{a})^\top + T_{22}(\tilde{\mathbf{X}} \boldsymbol{\beta})(\tilde{\mathbf{X}} \boldsymbol{\beta})^\top \right] \\
&\quad \times \bar{\mathbf{R}}_0 \mathbf{F}_0 (c_{\star,1} \mathbf{W}_0 \boldsymbol{\beta}_\star + c_{\star,1}^2 \eta \mathbf{a}) \\
&= 2 (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}) \left[T_{21}(\mathbf{F}_0 \mathbf{a})^\top + T_{22}(\tilde{\mathbf{X}} \boldsymbol{\beta})^\top \right] \bar{\mathbf{R}}_0 \mathbf{F}_0 (c_{\star,1} \mathbf{W}_0 \boldsymbol{\beta}_\star + c_{\star,1}^2 \eta \mathbf{a}) + o_{\mathbb{P}}(1) \\
&= 2 (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}) \left[c_{\star,1}^2 \eta T_{21} \mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} + c_{\star,1} T_{22} \boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{W}_0 \boldsymbol{\beta}_\star \right] + o_{\mathbb{P}}(1).
\end{aligned}$$

Also, for the third term we can write

$$E_3^{(3)} = -2 \mathbf{p}_3^\top \boldsymbol{\mu}_f = -2 \eta K \mathbf{a}^\top (c_{\star,1} \mathbf{W}_0 \boldsymbol{\beta}_\star + c_{\star,1}^2 \eta \mathbf{a}) = -2 c_{\star,1}^2 \eta^2 K + o_{\mathbb{P}}(1).$$

Summing up, the limiting value for E_3 is

$$E_3 + 2 c_{\star,1} \tilde{\mathbf{y}}^\top \mathbf{F}_0 \bar{\mathbf{R}}_0 \mathbf{W}_0 \boldsymbol{\beta}_\star + \rightarrow_P \frac{c_{\star,1}^4}{\phi(c_\star^2 + \sigma_\varepsilon^2) + c_{\star,1}^2} [2\bar{M} - 2]. \quad (64)$$

N.1.4 The Final Result

Putting equations 46, 63, 64 together, we have

$$\begin{aligned}
\mathcal{L}_{\text{te}}(\hat{\mathbf{a}}(\mathbf{F})) - \left(\sigma_\varepsilon^2 + c_\star^2 + \tilde{\mathbf{y}}^\top \bar{\mathbf{R}}_0 \mathbf{F} \boldsymbol{\Sigma}_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{y}} - 2 c_{\star,1} \tilde{\mathbf{y}}^\top \mathbf{F}_0 \bar{\mathbf{R}}_0 \mathbf{W}_0 \boldsymbol{\beta}_\star \right) \\
\rightarrow_P \frac{c_{\star,1}^4}{\phi(c_\star^2 + \sigma_\varepsilon^2) + c_{\star,1}^2} [2\bar{M} - 1 - M].
\end{aligned}$$

The test error of the untrained random features model can be written as

$$\mathcal{L}_{\text{te}}(\hat{\mathbf{a}}(\mathbf{F}_0)) = \sigma_\varepsilon^2 + c_\star^2 + \tilde{\mathbf{y}}^\top \bar{\mathbf{R}}_0 \mathbf{F} \boldsymbol{\Sigma}_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{y}} - 2 c_{\star,1} \tilde{\mathbf{y}}^\top \mathbf{F}_0 \bar{\mathbf{R}}_0 \mathbf{W}_0 \boldsymbol{\beta}_\star.$$

Hence, the improvement over the untrained random features model in terms of test error is equal to

$$\mathcal{L}_{\text{te}}(\hat{\mathbf{a}}(\mathbf{F}_0)) - \mathcal{L}_{\text{te}}(\hat{\mathbf{a}}(\mathbf{F})) \rightarrow_P \frac{c_{\star,1}^4}{\phi(c_\star^2 + \sigma_\varepsilon^2) + c_{\star,1}^2} [1 + M - 2\bar{M}].$$

Hence, using equation 62 and equation 61, we find

$$\mathcal{L}_{\text{te}}(\hat{\mathbf{a}}(\mathbf{F}_0)) - \mathcal{L}_{\text{te}}(\hat{\mathbf{a}}(\mathbf{F})) \rightarrow_P \frac{-c_{\star,1}^4}{(\phi(c_\star^2 + \sigma_\varepsilon^2) + c_{\star,1}^2) m_1^2} \frac{\partial m_2}{\partial \lambda}, \quad (65)$$

where $\frac{\partial m_1}{\partial \lambda} \leq 0$ using that $\text{tr}(\tilde{\mathbf{X}}^\top (\mathbf{F}_0 \mathbf{F}_0^\top + \lambda n \mathbf{I}_n)^{-1} \tilde{\mathbf{X}})/d = \psi m_2/\phi$. This concludes the proof for $\ell = 1$.

N.2 Proof for $\ell = 2$

First, similar to the proof for $\ell = 1$, we have

$$\begin{aligned}\nu_2 &= \mathbb{E}_{\mathbf{x}_{\text{te}}}[(\mathbf{x}_{\text{te}}^\top \boldsymbol{\beta})^2 \sigma(\mathbf{W}_0 \mathbf{x}_{\text{te}})] = 2c_2(\mathbf{W}_0 \boldsymbol{\beta})^{\circ 2}, \\ \aleph_3 &= \mathbb{E}_{\mathbf{x}_{\text{te}}}(\mathbf{x}_{\text{te}}^\top \boldsymbol{\beta})^3 = 0, \text{ and } \aleph_4 = \mathbb{E}_{\mathbf{x}_{\text{te}}}(\mathbf{x}_{\text{te}}^\top \boldsymbol{\beta})^4 = 3\|\boldsymbol{\beta}\|_2^2 \rightarrow_P 3[c_{\star,1}^2 + \phi(c_{\star}^2 + \sigma_\varepsilon^2)], \\ \tau_2 &= \mathbb{E} \left[y_{\text{te}}(\mathbf{x}_{\text{te}}^\top \boldsymbol{\beta})^2 \right] = 2c_{\star,2}(\boldsymbol{\beta}^\top \boldsymbol{\beta})^2 \rightarrow_P 2c_{\star,1}^2 c_{\star,2},\end{aligned}$$

in which we have used Lemma C.1. Thus, using equation 43 and equation 44, we have

$$\begin{aligned}\boldsymbol{\Sigma}_f &= \boldsymbol{\Sigma}_f^0 + \eta \left(\mathbf{a}(\mathbf{W}_0 \boldsymbol{\beta})^\top + (\mathbf{W}_0 \boldsymbol{\beta}) \mathbf{a}^\top \right) + \frac{2c_2^2 \eta^2}{\sqrt{N}} \left((\sqrt{N} \mathbf{a}^{\circ 2})(\mathbf{W}_0 \boldsymbol{\beta})^{\circ 2\top} + (\mathbf{W}_0 \boldsymbol{\beta})^{\circ 2}(\sqrt{N} \mathbf{a}^{\circ 2\top}) \right) \\ &\quad + \eta^2 \|\boldsymbol{\beta}\|_2^2 \mathbf{a} \mathbf{a}^\top + \frac{3c_2^2 \eta^4}{N} \|\boldsymbol{\beta}\|_2^4 (\sqrt{N} \mathbf{a}^{\circ 2})(\sqrt{N} \mathbf{a}^{\circ 2\top}), \quad \text{and} \\ \boldsymbol{\mu}_f &= \boldsymbol{\mu}_f^0 + c_{\star,1}^2 \eta \mathbf{a} + 2c_2 c_{\star,1}^2 c_{\star,2} \eta^2 \mathbf{a}^{\circ 2}.\end{aligned}$$

Also, from Section K.2, we have $\bar{\mathbf{R}} = \bar{\mathbf{R}}_0 - \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0$, where the matrix \mathbf{T}^{-1} is defined in equation 33 and $\mathbf{U} = [\mathbf{F}_0 \mathbf{a} \mid \mathbf{F}_0 \mathbf{a}^{\circ 2} \sqrt{N} \mid \tilde{\mathbf{X}} \boldsymbol{\beta} \mid (\tilde{\mathbf{X}} \boldsymbol{\beta})^{\circ 2}]$. Using the analysis in Section K.2.1, we note that

$$\mathbf{T}^{-1} = \begin{bmatrix} \mathbf{a}^\top (\mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 - \mathbf{I}) \mathbf{a} & 0 & \frac{1}{c_1^2 \eta} & 0 \\ 0 & N \mathbf{a}^{\circ 2\top} (\mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 - \mathbf{I}) \mathbf{a}^{\circ 2} & 0 & \frac{N^{\frac{1}{2}}}{c_1^2 c_2 \eta^2} \\ \frac{1}{c_1^2 \eta} & 0 & (\tilde{\mathbf{X}} \boldsymbol{\beta})^\top \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}} \boldsymbol{\beta}) & 0 \\ 0 & \frac{N^{\frac{1}{2}}}{c_1^2 c_2 \eta^2} & 0 & (\tilde{\mathbf{X}} \boldsymbol{\beta})^{\circ 2\top} \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}} \boldsymbol{\beta})^{\circ 2} \end{bmatrix} + \boldsymbol{\Delta},$$

in which the elements of $\boldsymbol{\Delta}$ are $O_{\mathbb{P}}(1/\sqrt{n})$. Hence, the matrix $\mathbf{T} = [T_{i,j}]$ has entries

$$\begin{aligned}T_{11} &= \frac{1}{\mathbf{a}^\top (\mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 - \mathbf{I}) \mathbf{a}} + \frac{1}{\eta^2} \cdot \frac{1}{(\mathbf{a}^\top (\mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 - \mathbf{I}) \mathbf{a})^2 ((\tilde{\mathbf{X}} \boldsymbol{\beta})^\top \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}} \boldsymbol{\beta}))} + o_{\mathbb{P}}(1/\eta^2), \\ T_{13} = T_{31} &= -\frac{1}{\eta} \cdot \frac{1}{(\mathbf{a}^\top (\mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 - \mathbf{I}) \mathbf{a}) ((\tilde{\mathbf{X}} \boldsymbol{\beta})^\top \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}} \boldsymbol{\beta}))} + o_{\mathbb{P}}(1/\eta), \\ T_{22} &= \frac{1}{N \mathbf{a}^{\circ 2\top} (\mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 - \mathbf{I}) \mathbf{a}^{\circ 2}} + \frac{N}{c_2^2 \eta^4} \cdot \frac{1}{(N \mathbf{a}^{\circ 2\top} (\mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 - \mathbf{I}) \mathbf{a}^{\circ 2})^2 ((\tilde{\mathbf{X}} \boldsymbol{\beta})^{\circ 2\top} \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}} \boldsymbol{\beta})^{\circ 2})} + o_{\mathbb{P}}(N/\eta^4), \\ T_{24} = T_{42} &= -\frac{\sqrt{N}}{c_2 \eta^2} \cdot \frac{1}{(N \mathbf{a}^{\circ 2\top} (\mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 - \mathbf{I}) \mathbf{a}^{\circ 2}) ((\tilde{\mathbf{X}} \boldsymbol{\beta})^{\circ 2\top} \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}} \boldsymbol{\beta})^{\circ 2})} + o_{\mathbb{P}}(1/\eta^2), \\ T_{33} &= \frac{1}{(\tilde{\mathbf{X}} \boldsymbol{\beta})^\top \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}} \boldsymbol{\beta})} + \frac{1}{\eta^2} \cdot \frac{1}{(\mathbf{a}^\top (\mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 - \mathbf{I}) \mathbf{a}) ((\tilde{\mathbf{X}} \boldsymbol{\beta})^\top \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}} \boldsymbol{\beta}))^2} + o_{\mathbb{P}}(1/\eta^2), \\ T_{44} &= \frac{1}{(\tilde{\mathbf{X}} \boldsymbol{\beta})^{\circ 2\top} \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}} \boldsymbol{\beta})^{\circ 2}} + \frac{N}{c_2^2 \eta^4} \cdot \frac{1}{(N \mathbf{a}^{\circ 2\top} (\mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 - \mathbf{I}) \mathbf{a}^{\circ 2\top}) ((\tilde{\mathbf{X}} \boldsymbol{\beta})^{\circ 2\top} \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}} \boldsymbol{\beta})^{\circ 2})^2} + o_{\mathbb{P}}(N/\eta^4),\end{aligned}$$

and its other elements are $O_{\mathbb{P}}(1/\sqrt{n})$.

Next, we will study the terms $E_1 = \mathbb{E}_{\mathbf{x}_{\text{te}}, y_{\text{te}}}(y_{\text{te}}^2)$, $E_2 = \hat{\mathbf{a}}^\top \boldsymbol{\Sigma}_f \hat{\mathbf{a}}$, and $E_3 = -2\hat{\mathbf{a}}^\top \boldsymbol{\mu}_f$ that appear in the decomposition of the test error in equation 41.

N.2.1 Analysis of E_1

Similar to the $\ell = 1$ case, we have

$$E_1 = \mathbb{E}(y_{\text{te}}^2) = \sigma_\varepsilon^2 + \mathbb{E}_{\mathbf{x}_{\text{te}}} \left(\sigma_\star(\boldsymbol{\beta}_\star^\top \mathbf{x}_{\text{te}}) \right)^2 = \sigma_\varepsilon^2 + c_\star^2. \quad (66)$$

N.2.2 Analysis of E_2

Recall that using equation 48, we have

$$E_2 = \mathbf{f}_\star^\top \bar{\mathbf{R}} \mathbf{F} \boldsymbol{\Sigma}_f \mathbf{F}^\top \bar{\mathbf{R}} \mathbf{f}_\star + 2\varepsilon^\top \bar{\mathbf{R}} \mathbf{F} \boldsymbol{\Sigma}_f \mathbf{F}^\top \bar{\mathbf{R}} \tilde{\mathbf{f}}_\star + \varepsilon^\top \bar{\mathbf{R}} \mathbf{F} \boldsymbol{\Sigma}_f \mathbf{F}^\top \bar{\mathbf{R}} \varepsilon.$$

First, note that $\varepsilon^\top \bar{\mathbf{R}} \mathbf{F} \boldsymbol{\Sigma}_f \mathbf{F}^\top \bar{\mathbf{R}} \mathbf{f}_\star = o_{\mathbb{P}}(1)$ using a simple order-wise argument. Using an argument similar to the one for $\ell = 1$, the third term can be written as $\varepsilon^\top \bar{\mathbf{R}} \mathbf{F} \boldsymbol{\Sigma}_f \mathbf{F}^\top \bar{\mathbf{R}} \varepsilon = \varepsilon^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \boldsymbol{\Sigma}_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \varepsilon + o_{\mathbb{P}}(1)$.

Next, we will study the term $\mathbf{f}_\star^\top \bar{\mathbf{R}} \mathbf{F} \boldsymbol{\Sigma}_f \mathbf{F}^\top \bar{\mathbf{R}} \mathbf{f}_\star$. We can write

$$\begin{aligned} \mathbf{F}^\top \bar{\mathbf{R}} \mathbf{f}_\star &= \left(\mathbf{F}_0 + \eta(\tilde{\mathbf{X}}\boldsymbol{\beta}) \mathbf{a}^\top + c_2 \eta^2 (\tilde{\mathbf{X}}\boldsymbol{\beta})^{\circ 2} \mathbf{a}^{\circ 2\top} \right)^\top \left(\bar{\mathbf{R}}_0 - \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \right) \mathbf{f}_\star \\ &= \underbrace{\mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star}_{p_1} - \underbrace{\mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star}_{p_2} + \underbrace{\eta K_1 \mathbf{a}}_{p_3} + \underbrace{c_2 \eta^2 K_2 \mathbf{a}^{\circ 2}}_{p_4}, \end{aligned} \quad (67)$$

in which K_1 and K_2 are defined as

$$\begin{aligned} K_1 &= \left[(\tilde{\mathbf{X}}\boldsymbol{\beta})^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star - (\tilde{\mathbf{X}}\boldsymbol{\beta})^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star \right], \\ K_2 &= \left[(\tilde{\mathbf{X}}\boldsymbol{\beta})^{\circ 2\top} \bar{\mathbf{R}}_0 \mathbf{f}_\star - (\tilde{\mathbf{X}}\boldsymbol{\beta})^{\circ 2\top} \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star \right]. \end{aligned}$$

Using this notation, we have $\mathbf{f}_\star^\top \bar{\mathbf{R}} \mathbf{F} \boldsymbol{\Sigma}_f \mathbf{F}^\top \bar{\mathbf{R}} \mathbf{f}_\star = \sum_{i,j \in [4]} E_2^{ij}$, where $E_2^{ij} = \mathbf{p}_i^\top \boldsymbol{\Sigma}_f \mathbf{p}_j^\top$. In the following sections, we will compute each term separately.

Preliminary Computations. First, we will analyze K_1 and K_2 . Recall that

$$\begin{aligned} \mathbf{U} \mathbf{T} \mathbf{U}^\top &= T_{11}(\mathbf{F}_0 \mathbf{a})(\mathbf{F}_0 \mathbf{a})^\top + T_{13}(\mathbf{F}_0 \mathbf{a})(\tilde{\mathbf{X}}\boldsymbol{\beta})^\top \\ &\quad + NT_{22}(\mathbf{F}_0 \mathbf{a}^{\circ 2})(\mathbf{F}_0 \mathbf{a}^{\circ 2})^\top + \sqrt{N}T_{24}(\mathbf{F}_0 \mathbf{a}^{\circ 2})(\tilde{\mathbf{X}}\boldsymbol{\beta})^{\circ 2\top} \\ &\quad + T_{31}(\tilde{\mathbf{X}}\boldsymbol{\beta})(\mathbf{F}_0 \mathbf{a})^\top + T_{33}(\tilde{\mathbf{X}}\boldsymbol{\beta})(\tilde{\mathbf{X}}\boldsymbol{\beta})^\top + \sqrt{N}T_{42}(\tilde{\mathbf{X}}\boldsymbol{\beta})^{\circ 2}(\mathbf{F}_0 \mathbf{a}^{\circ 2\top}) \\ &\quad + T_{44}(\tilde{\mathbf{X}}\boldsymbol{\beta})^{\circ 2}(\tilde{\mathbf{X}}\boldsymbol{\beta})^{\circ 2\top}. \end{aligned}$$

Thus, we have

$$(\tilde{\mathbf{X}}\boldsymbol{\beta})^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star = T_{33} \left[(\tilde{\mathbf{X}}\boldsymbol{\beta})^\top \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}}\boldsymbol{\beta}) \right] \cdot \left[(\tilde{\mathbf{X}}\boldsymbol{\beta})^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star \right] + O_{\mathbb{P}}(1/\sqrt{n}),$$

which gives

$$K_1 = -\frac{1}{\eta^2} \cdot \frac{\boldsymbol{\beta}^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0^\top \mathbf{f}_\star}{(\mathbf{a}^\top (\mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 - \mathbf{I}) \mathbf{a}) ((\tilde{\mathbf{X}}\boldsymbol{\beta})^\top \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}}\boldsymbol{\beta}))} + o_{\mathbb{P}}(1/\eta^2). \quad (68)$$

Similarly, for K_2 we have

$$(\tilde{\mathbf{X}}\beta)^{\circ 2\top} \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star = T_{44} \left[(\tilde{\mathbf{X}}\beta)^{\circ 2\top} \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}}\beta)^{\circ 2} \right] \cdot \left[(\tilde{\mathbf{X}}\beta)^{\circ 2\top} \bar{\mathbf{R}}_0 \mathbf{f}_\star \right] + o_{\mathbb{P}}(1/\sqrt{n}),$$

which gives

$$K_2 = -\frac{N}{c_2^2 \eta^4} \cdot \frac{(\tilde{\mathbf{X}}\beta)^{\circ 2\top} \bar{\mathbf{R}}_0^\top \mathbf{f}_\star}{(N \mathbf{a}^{\circ 2\top} (\mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 - \mathbf{I}) \mathbf{a}^{\circ 2\top}) ((\tilde{\mathbf{X}}\beta)^{\circ 2\top} \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}}\beta)^{\circ 2})} + o_{\mathbb{P}}(1/\eta^4). \quad (69)$$

Analysis of E_2^{11} . For this term, by expanding Σ_f , we can write

$$E_2^{11} = \mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star = \mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star + o_{\mathbb{P}}(1). \quad (70)$$

This holds because $\eta \mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} = o_{\mathbb{P}}(1)$, and $\frac{\eta^2}{\sqrt{N}} \mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 (\sqrt{N} \mathbf{a}^{\circ 2}) = o_{\mathbb{P}}(1)$.

Analysis of E_2^{12} and E_2^{21} . We have $E_2^{21} = E_2^{12} = -\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star$. Using the expression for \mathbf{T} , we can write

$$\mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star = \begin{bmatrix} T_{11} \mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star + T_{13} (\tilde{\mathbf{X}}\beta)^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star \\ T_{22} \sqrt{N} \mathbf{a}^{\circ 2\top} \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star + T_{24} (\tilde{\mathbf{X}}\beta)^{\circ 2\top} \bar{\mathbf{R}}_0 \mathbf{f}_\star \\ T_{31} \mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star + T_{33} (\tilde{\mathbf{X}}\beta)^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star \\ T_{42} \sqrt{N} \mathbf{a}^{\circ 2\top} \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star + T_{44} (\tilde{\mathbf{X}}\beta)^{\circ 2\top} \bar{\mathbf{R}}_0 \mathbf{f}_\star \end{bmatrix}. \quad (71)$$

On the other hand, by expanding \mathbf{U} and Σ_f , we similarly have

$$(\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{U}) = \begin{bmatrix} \eta (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{W}_0 \beta) (\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}) \\ \frac{2c_2^2 \eta^2}{\sqrt{N}} (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 (\mathbf{W}_0 \beta)^{\circ 2}) ((\sqrt{N} \mathbf{a}^{\circ 2})^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 (\sqrt{N} \mathbf{a}^{\circ 2})) \\ \mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}}\beta \\ \mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}}\beta)^{\circ 2} \end{bmatrix}^\top. \quad (72)$$

This gives

$$\begin{aligned} E_2^{21} = E_2^{12} = & -T_{13} \eta (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{W}_0 \beta) (\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}) ((\tilde{\mathbf{X}}\beta)^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star) \\ & - \frac{2c_2^2 \eta^2 T_{24}}{\sqrt{N}} (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 (\mathbf{W}_0 \beta)^{\circ 2}) ((\sqrt{N} \mathbf{a}^{\circ 2})^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 (\sqrt{N} \mathbf{a}^{\circ 2})) ((\tilde{\mathbf{X}}\beta)^{\circ 2\top} \bar{\mathbf{R}}_0 \mathbf{f}_\star) \\ & - T_{33} (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}}\beta) ((\tilde{\mathbf{X}}\beta)^\top \bar{\mathbf{R}}_0 \mathbf{f}_\star) \\ & - T_{44} (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}}\beta)^{\circ 2}) ((\tilde{\mathbf{X}}\beta)^{\circ 2\top} \bar{\mathbf{R}}_0 \mathbf{f}_\star) + o_{\mathbb{P}}(1). \end{aligned} \quad (73)$$

Analysis of E_2^{13} and E_2^{31} . Recalling equation 68, by expanding Σ_f we have

$$E_2^{13} = E_2^{31} = \eta K_1 \mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f \mathbf{a} = \eta^2 K_1 (\mathbf{f}_\star^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{W}_0 \beta) + o_{\mathbb{P}}(1), \quad (74)$$

in which we have used that $K_1 = O_{\mathbb{P}}(1/\eta^2)$.

Analysis of E_2^{14} and E_2^{41} . Recalling equation 69, by expanding Σ_f we have

$$E_2^{14} = E_2^{41} = c_2 \eta^2 K_2 \mathbf{f}_*^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f \mathbf{a}^{\circ 2} = o_{\mathbb{P}}(1), \quad (75)$$

in which we have used that $\eta^4 K_2 / N = O_{\mathbb{P}}(1)$.

Analysis of E_2^{22} . This term is equal to $E_2^{22} = \mathbf{f}_*^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_*$. Using equation 71, we can write

$$\begin{aligned} \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_* &= [T_{11} \mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{f}_* + T_{13} (\tilde{\mathbf{X}}\beta)^\top \bar{\mathbf{R}}_0 \mathbf{f}_*] \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} \\ &\quad + [T_{22} \sqrt{N} \mathbf{a}^{\circ 2\top} \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{f}_* + T_{24} (\tilde{\mathbf{X}}\beta)^{\circ 2\top} \bar{\mathbf{R}}_0 \mathbf{f}_*] \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 (\sqrt{N} \mathbf{a}^{\circ 2}) \\ &\quad + [T_{31} \mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{f}_* + T_{33} (\tilde{\mathbf{X}}\beta)^\top \bar{\mathbf{R}}_0 \mathbf{f}_*] \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}}\beta \\ &\quad + [T_{42} \sqrt{N} \mathbf{a}^{\circ 2\top} \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{f}_* + T_{44} (\tilde{\mathbf{X}}\beta)^{\circ 2\top} \bar{\mathbf{R}}_0 \mathbf{f}_*] \mathbf{F}_0^\top \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}}\beta)^{\circ 2}. \end{aligned} \quad (76)$$

Further, by expanding Σ_f , we have

$$E_2^{22} = E_2^{22(1)} + E_2^{22(2)} + E_2^{22(3)} + E_2^{22(4)} + o_{\mathbb{P}}(1), \quad (77)$$

in which

$$\begin{aligned} E_2^{22(1)} &= [T_{33} (\tilde{\mathbf{X}}\beta)^\top \bar{\mathbf{R}}_0 \mathbf{f}_*]^2 (\beta^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}}\beta) \\ &\quad + [T_{44} (\tilde{\mathbf{X}}\beta)^{\circ 2\top} \bar{\mathbf{R}}_0 \mathbf{f}_*]^2 ((\tilde{\mathbf{X}}\beta)^{\circ 2\top} \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}}\beta)^{\circ 2}), \\ E_2^{22(2)} &= 2\eta T_{13} T_{33} \cdot ((\tilde{\mathbf{X}}\beta)^\top \bar{\mathbf{R}}_0 \mathbf{f}_*)^2 (\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}) (\beta^\top \mathbf{W}_0^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}}\beta), \\ E_2^{22(3)} &= \frac{4c_2^2 \eta^2}{\sqrt{N}} T_{24} T_{44} ((\tilde{\mathbf{X}}\beta)^{\circ 2\top} \bar{\mathbf{R}}_0 \mathbf{f}_*)^2 \left((\sqrt{N} \mathbf{a}^{\circ 2})^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 (\sqrt{N} \mathbf{a}^{\circ 2}) \right) \left((\mathbf{W}_0 \beta)^{\circ 2} \mathbf{F}_0^\top \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}}\beta)^{\circ 2} \right), \\ E_2^{22(4)} &= \eta^2 T_{13}^2 \|\beta\|_2^2 ((\tilde{\mathbf{X}}\beta)^\top \bar{\mathbf{R}}_0 \mathbf{f}_*)^2 (\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a})^2 \\ &\quad + \frac{3c_2^2 \eta^4}{N} T_{24}^2 \|\beta\|_2^4 \left((\tilde{\mathbf{X}}\beta)^{\circ 2\top} \bar{\mathbf{R}}_0 \mathbf{f}_* \right)^2 \left((\sqrt{N} \mathbf{a}^{\circ 2})^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 (\sqrt{N} \mathbf{a}^{\circ 2}) \right)^2. \end{aligned}$$

Analysis of E_2^{23} and E_2^{32} . We have $E_2^{23} = E_2^{32} = -\eta K_1 \mathbf{a}^\top \Sigma_f \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_*$. Recall that the vector $\mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_*$ has been computed in equation 76. With this, we have

$$\begin{aligned} E_2^{23} = E_2^{32} &= -\eta^2 K_1 T_{33} (\beta^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{f}_*) (\beta^\top \mathbf{W}_0^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}}\beta) \\ &\quad - \eta^3 K_1 T_{13} \|\beta\|_2^2 (\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}) (\beta^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \mathbf{f}_*) + o_{\mathbb{P}}(1), \end{aligned} \quad (78)$$

in which we used that $\eta^2 K_1 = O_{\mathbb{P}}(1)$ and $\eta T_{13} = O_{\mathbb{P}}(1)$.

Analysis of E_2^{24} and E_2^{42} . This term can be written as

$$E_{24} = E_{42} = -c_2 \eta^2 K_2 \mathbf{a}^{\circ 2\top} \Sigma_f \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_*.$$

The vector $\mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^\top \bar{\mathbf{R}}_0 \mathbf{f}_*$ has been computed in equation 76. By expanding Σ_f , we have

$$\begin{aligned} E_2^{24} = E_2^{42} &= -\frac{2c_2^3 \eta^4 K_2 T_{44}}{N} \cdot \left((\mathbf{W}_0 \beta)^{\circ 2\top} \mathbf{F}_0^\top \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}}\beta)^{\circ 2} \right) \cdot \left((\tilde{\mathbf{X}}\beta)^{\circ 2\top} \bar{\mathbf{R}}_0 \mathbf{f}_* \right) \\ &\quad - \frac{3c_2^3 \eta^6 K_2 T_{24} \|\beta\|_2^4}{N^{3/2}} \left((\sqrt{N} \mathbf{a}^{\circ 2})^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 (\sqrt{N} \mathbf{a}^{\circ 2}) \right) \cdot \left((\tilde{\mathbf{X}}\beta)^{\circ 2\top} \bar{\mathbf{R}}_0 \mathbf{f}_* \right) + o_{\mathbb{P}}(1). \end{aligned}$$

Now, noting that $K_2 = O_{\mathbb{P}}(N/\eta^4)$ from equation 69 and $\|(\mathbf{W}_0\beta)^{\circ 2}\|_2 = O_{\mathbb{P}}(1/\sqrt{N})$, we find that the first term is $o_{\mathbb{P}}(1)$ and we have

$$E_2^{24} = E_2^{42} = -\frac{3c_2^3\eta^6K_2T_{24}\|\beta\|_2^4}{N^{3/2}} \left((\sqrt{N}\mathbf{a}^{\circ 2})^\top \bar{\mathbf{F}}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 (\sqrt{N}\mathbf{a}^{\circ 2}) \right) \cdot \left((\tilde{\mathbf{X}}\beta)^{\circ 2\top} \bar{\mathbf{R}}_0 \mathbf{f}_\star \right) + o_{\mathbb{P}}(1). \quad (79)$$

Analysis of E_2^{33} . Similar to the $\ell = 1$ case, this term can be written as

$$E_2^{33} = \eta^2 K_1^2 \mathbf{a}^\top \Sigma_f \mathbf{a} = \eta^4 K_1^2 \|\beta\|_2^2 + o_{\mathbb{P}}(1), \quad (80)$$

noting that $K_1 = O_{\mathbb{P}}(1/\eta^2)$.

Analysis of E_2^{34} and E_2^{43} . By expanding Σ_f , we readily arrive at

$$E_2^{34} = E_2^{43} = \eta^3 K_1 K_2 c_2 \mathbf{a}^\top \Sigma_f \mathbf{a}^{\circ 2} = o_{\mathbb{P}}(1). \quad (81)$$

Analysis of E_2^{44} . This term can be written as $E_2^{44} = c_2^2 \eta^4 K_2^2 \mathbf{a}^{\circ 2\top} \Sigma_f \mathbf{a}^{\circ 2}$. By expanding Σ_f and noting that $K_2 = O_{\mathbb{P}}(N/\eta^4)$, we can write

$$E_2^{44} = \frac{3c_2^4 \eta^8 K_2^2}{N^2} + o_{\mathbb{P}}(1). \quad (82)$$

Putting Everything Together. Now, we can use the results derived above to compute the limiting value of E_2 . Recall that from Lemma K.7 and Lemma K.8 we have

$$H_2(\tilde{\mathbf{X}}\beta_\star)^\top \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}}\beta)^{\circ 2} \rightarrow_P 2c_{\star,1}^2 \frac{\psi m_1}{\phi}, \quad \text{and} \quad (\tilde{\mathbf{X}}\beta)^{\circ 2\top} \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}}\beta)^{\circ 2} \rightarrow_P \frac{3\psi m_1}{\phi} [\phi(c_\star^2 + \sigma_\varepsilon^2) + c_{\star,1}^2]^2.$$

Also using an argument similar to the argument in the proof of Lemma K.7 and Lemma K.8, we have

$$\begin{aligned} H_2(\tilde{\mathbf{X}}\beta_\star)^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}}\beta)^{\circ 2} &\rightarrow_p 2c_{\star,1}^2 \hat{M}, \\ (\tilde{\mathbf{X}}\beta)^{\circ 2\top} \bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}}\beta)^{\circ 2} &\rightarrow 3\hat{M}[\phi(c_\star^2 + \sigma_\varepsilon^2) + c_{\star,1}^2]^2, \end{aligned}$$

in which $\hat{M} = \lim_{n,N,d \rightarrow \infty} \frac{1}{n} \text{tr} [\bar{\mathbf{R}}_0 \mathbf{F}_0 \Sigma_f^0 \mathbf{F}_0^\top \bar{\mathbf{R}}_0]$. This term has been computed in [AP20a]. Using the diagram in the proof of Lemma K.4 that shows how the notations of [AP20a] match ours, we find that we can use (S142) with $\sigma_\varepsilon = 0$ in [AP20a], to find

$$\hat{M} = -\frac{m'_1}{m_1^2} - 1, \quad (83)$$

where m'_1 is the derivative of m_1 with respect to λ .

For brevity, we will define $\tilde{A} := \lim_{n,N,d \rightarrow \infty} \left[(\sqrt{N}\mathbf{a}^{\circ 2\top}) \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 (\sqrt{N}\mathbf{a}^{\circ 2}) \right]$. With this, equations 73, 74, 75, 77, 78, 79, 80, 81, and 82 give

$$E_2^{12} = E_2^{21} \rightarrow_P \frac{c_{\star,1}^4}{\phi(c_\star^2 + \sigma_\varepsilon^2) + c_{\star,1}^2} \left[\bar{M} \left(\frac{\psi/\phi - \lambda\psi^2 m_1/\phi^2}{\psi/\phi - \lambda\psi^2 m_1/\phi^2 - 1} \right) - M \right] - \frac{4c_{\star,1}^4 c_{\star,2}^2 \hat{M}}{3[\phi(c_\star^2 + \sigma_\varepsilon^2) + c_{\star,1}^2]^2},$$

$$E_2^{13} = E_2^{31} \rightarrow_P -\frac{c_{\star,1}^4}{\phi(c_\star^2 + \sigma_\varepsilon^2) + c_{\star,1}^2} \cdot \frac{\bar{M}}{\psi/\phi - \lambda\psi^2 m_1/\phi^2 - 1}, \quad E_2^{14} = E_2^{41} \rightarrow_P 0,$$

as well as

$$\begin{aligned}
E_2^{22} &\rightarrow_P \frac{c_{\star,1}^4}{\phi(c_{\star}^2 + \sigma_{\varepsilon}^2) + c_{\star,1}^2} \left[M - 2\bar{M} \left(\frac{\psi/\phi - \lambda\psi^2 m_1/\phi^2}{\psi/\phi - \lambda\psi^2 m_1/\phi^2 - 1} \right) + \left(\frac{\psi/\phi - \lambda\psi^2 m_1/\phi^2}{\psi/\phi - \lambda\psi^2 m_1/\phi^2 - 1} \right)^2 \right] \\
&\quad + \frac{c_{\star,1}^4 c_{\star,2}^2}{[\phi(c_{\star}^2 + \sigma_{\varepsilon}^2) + c_{\star,1}^2]^2} \left[4\hat{M}/3 + \frac{4\tilde{A}^2}{3(\tilde{A}-1)^2} \right], \\
E_2^{23} = E_2^{32} &\rightarrow_P \frac{c_{\star,1}^4}{\phi(c_{\star}^2 + \sigma_{\varepsilon}^2) + c_{\star,1}^2} \left[\frac{\bar{M}}{\psi/\phi - \lambda\psi^2 m_1/\phi^2 - 1} - \frac{\psi/\phi - \lambda\psi^2 m_1/\phi^2}{(\psi/\phi - \lambda\psi^2 m_1/\phi^2 - 1)^2} \right],
\end{aligned}$$

and also

$$\begin{aligned}
E_2^{24} = E_2^{42} &\rightarrow_P -\frac{c_{\star,1}^4 c_{\star,2}^2}{[\phi(c_{\star}^2 + \sigma_{\varepsilon}^2) + c_{\star,1}^2]^2} \left[\frac{4\tilde{A}}{3(\tilde{A}-1)^2} \right], \\
E_2^{33} &\rightarrow_P \frac{c_{\star,1}^4}{\phi(c_{\star}^2 + \sigma_{\varepsilon}^2) + c_{\star,1}^2} \left[\frac{1}{(\psi/\phi - \lambda\psi^2 m_1/\phi^2 - 1)^2} \right], \\
E_2^{34} = E_2^{43} &\rightarrow_P 0, \quad E_2^{44} = E_2^{44} \rightarrow_P \frac{c_{\star,1}^4 c_{\star,2}^2}{[\phi(c_{\star}^2 + \sigma_{\varepsilon}^2) + c_{\star,1}^2]^2} \left[\frac{4}{3(\tilde{A}-1)^2} \right],
\end{aligned}$$

respectively. Putting these together, the component E_2 can be written as

$$E_2 - \tilde{\mathbf{y}}^{\top} \bar{\mathbf{R}}_0 \mathbf{F} \Sigma_f^0 \mathbf{F}_0^{\top} \bar{\mathbf{R}}_0 \tilde{\mathbf{y}} \rightarrow_P \frac{c_{\star,1}^4(1-M)}{\phi(c_{\star}^2 + \sigma_{\varepsilon}^2) + c_{\star,1}^2} + \frac{4c_{\star,1}^4 c_{\star,2}^2(1-\hat{M})}{3[\phi(c_{\star}^2 + \sigma_{\varepsilon}^2) + c_{\star,1}^2]^2}. \quad (84)$$

N.2.3 Analysis of E_3 .

To analyze this component, first note that

$$E_3 = -2\hat{\mathbf{a}}^{\top} \boldsymbol{\mu}_f = -2\tilde{\mathbf{y}}^{\top} \bar{\mathbf{R}} \mathbf{F} \boldsymbol{\mu}_f = -2\mathbf{f}_{\star}^{\top} \bar{\mathbf{R}} \mathbf{F} \boldsymbol{\mu}_f - 2\varepsilon^{\top} \bar{\mathbf{R}} \mathbf{F} \boldsymbol{\mu}_f.$$

Using a simple order-wise analysis, it can be shown that the second term is $o_{\mathbb{P}}(1)$. Now, recalling equation 67 we write $E_3 = \sum_{i=1}^4 E_3^{(i)}$, where $E_3^{(i)} = -2\mathbf{p}_i^{\top} \boldsymbol{\mu}_f$. By expanding $\boldsymbol{\mu}_f$ and recalling that $\boldsymbol{\mu}_f^0 = c_{\star,1} \mathbf{W}_0 \boldsymbol{\beta}$, we have

$$E_3^{(1)} = -2\mathbf{f}_{\star}^{\top} \bar{\mathbf{R}}_0 \mathbf{F}_0 \boldsymbol{\mu}_f^0 + o_{\mathbb{P}}(1). \quad (85)$$

Next, recall that the matrix $\mathbf{F}_0^{\top} \bar{\mathbf{R}}_0 \mathbf{U} \mathbf{T} \mathbf{U}^{\top} \bar{\mathbf{R}}_0 \mathbf{f}_{\star}$ is analyzed in equation 76. Using this, and by expanding $\boldsymbol{\mu}_f$, we get

$$\begin{aligned}
E_3^{(2)} &= 2(\mathbf{f}_{\star}^{\top} \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}) \left[T_{33}(\boldsymbol{\mu}_f^{0\top} \mathbf{F}_0^{\top} \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \boldsymbol{\beta}) + c_{\star,1}^2 T_{13} \eta(\mathbf{a}^{\top} \mathbf{F}_0^{\top} \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}) \right] \\
&\quad + \frac{4c_2 c_{\star,1}^2 c_{\star,2} T_{24} \eta^2}{\sqrt{N}} \left[(\sqrt{N} \mathbf{a}^{\circ 2})^{\top} \mathbf{F}_0^{\top} \bar{\mathbf{R}}_0 \mathbf{F}_0 (\sqrt{N} \mathbf{a}^{\circ 2}) \right] \left((\tilde{\mathbf{X}} \boldsymbol{\beta})^{\circ 2\top} \bar{\mathbf{R}}_0 \mathbf{f}_{\star} \right) + o_{\mathbb{P}}(1). \quad (86)
\end{aligned}$$

Similarly, by expanding μ_f , we arrive at

$$E_3^{(3)} = -2c_{\star,1}^2\eta^2K_1 + o_{\mathbb{P}}(1), \quad \text{and} \quad E_3^{(4)} = -\frac{4c_2^2c_{\star,1}^2c_{\star,2}\eta^4K_2}{N} + o_{\mathbb{P}}(1).$$

Similar to the computation for E_2 , we can derive the limiting values of the components in E_3 as

$$\begin{aligned} E_3^{(2)} &\rightarrow_P \frac{c_{\star,1}^4}{\phi(c_{\star}^2 + \sigma_{\varepsilon}^2) + c_{\star,1}^2} \left(2\bar{M} - \frac{2\tilde{A}}{\tilde{A} - 1} \right) - \frac{c_{\star,1}^4c_{\star,2}^2}{[\phi(c_{\star}^2 + \sigma_{\varepsilon}^2) + c_{\star,1}^2]^2} \left(\frac{8\tilde{A}}{3(\tilde{A} - 1)} \right), \\ E_3^{(3)} &\rightarrow_P \frac{c_{\star,1}^4}{\phi(c_{\star}^2 + \sigma_{\varepsilon}^2) + c_{\star,1}^2} \left(\frac{2}{\tilde{A} - 1} \right), \quad E_3^{(4)} \rightarrow_P \frac{c_{\star,1}^4c_{\star,2}^2}{[\phi(c_{\star}^2 + \sigma_{\varepsilon}^2) + c_{\star,1}^2]^2} \left(\frac{8}{3(\tilde{A} - 1)} \right). \end{aligned}$$

Putting these together, we have

$$E_3 + 2f_{\star}^{\top} \bar{\mathbf{R}}_0 \mathbf{F}_0 \mu_f^0 \rightarrow_P \frac{c_{\star,1}^4(2\bar{M} - 2)}{\phi(c_{\star}^2 + \sigma_{\varepsilon}^2) + c_{\star,1}^2} - \frac{8c_{\star,1}^4c_{\star,2}^2}{3[\phi(c_{\star}^2 + \sigma_{\varepsilon}^2) + c_{\star,1}^2]^2}. \quad (87)$$

N.2.4 The Final Result

Putting equations 66, 84, and 87, we have

$$\mathcal{L}_{\text{te}}(\hat{\mathbf{a}}(\mathbf{F}_0)) - \mathcal{L}_{\text{te}}(\hat{\mathbf{a}}(\mathbf{F})) \rightarrow_P \frac{c_{\star,1}^4(1 + M - 2\bar{M})}{\phi(c_{\star}^2 + \sigma_{\varepsilon}^2) + c_{\star,1}^2} + \frac{4c_{\star,1}^4c_{\star,2}^2(1 + \hat{M})}{3[\phi(c_{\star}^2 + \sigma_{\varepsilon}^2) + c_{\star,1}^2]^2},$$

where $\mathcal{L}_{\text{te}}(\hat{\mathbf{a}}(\mathbf{F}_0))$ is the test error of the untrained random feature model. Further, using equation 61, equation 62, and equation 83, we get

$$\mathcal{L}_{\text{te}}(\hat{\mathbf{a}}(\mathbf{F}_0)) - \mathcal{L}_{\text{te}}(\hat{\mathbf{a}}(\mathbf{F})) \rightarrow_P - \frac{c_{\star,1}^4}{(\phi(c_{\star}^2 + \sigma_{\varepsilon}^2) + c_{\star,1}^2)m_1^2} \frac{\partial m_2}{\partial \lambda} - \frac{4c_{\star,1}^4c_{\star,2}^2}{3[\phi(c_{\star}^2 + \sigma_{\varepsilon}^2) + c_{\star,1}^2]^2m_1^2} \frac{\partial m_1}{\partial \lambda}. \quad (88)$$

Note that $\frac{\partial m_1}{\partial \lambda}, \frac{\partial m_2}{\partial \lambda} \leq 0$, concluding the proof.

O Proofs of Supplementary Lemmas

O.1 Proof of Lemma C.3

Recalling $a_i \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1/N)$ and $\langle \tilde{\mathbf{x}}_i, \beta \rangle | \beta \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \|\beta\|_2^2)$, claims (a) and (b) follow from standard Gaussian maximal inequalities [vdVW13, Section 2.2] and from $\|\beta\|_2^2 = O_{\mathbb{P}}(1)$; the latter follows by writing $\beta = n^{-1}\mathbf{X}^{\top}(\sigma_{\star}(\mathbf{X}\beta_{\star}) + \varepsilon)$, where $\varepsilon = (\varepsilon_1, \dots, \varepsilon_n)^{\top}$ and using our distributional assumptions on \mathbf{X}, ε , as well as Condition 2.2.

By [Ver12, Theorem 5.39] and [BS10, Corollary A.21], we have $\|\mathbf{W}_0 \mathbf{W}_0^{\top}\|_{\text{op}}, \|(\mathbf{W}_0 \mathbf{W}_0^{\top})^{\circ 2}\|_{\text{op}} = O_{\mathbb{P}}(1)$. Also, by [Ver18, Theorem 3.4.6] and Gaussian maximal inequalities [vdVW13, Section 2.2], we have $\max_{1 \leq i \neq j \leq N} \langle \mathbf{w}_{0,i}, \mathbf{w}_{0,j} \rangle = O_{\mathbb{P}}(n^{-\frac{1}{2}} \log^{\frac{1}{2}} n)$. For $k \geq 3$,

$$\begin{aligned} \|(\mathbf{W}_0 \mathbf{W}_0^{\top})^{\circ k}\|_{\text{op}} &\leq \|(\mathbf{W}_0 \mathbf{W}_0^{\top})^{\circ k} - \mathbf{I}_N\|_{\text{op}} + 1 \leq \|(\mathbf{W}_0 \mathbf{W}_0^{\top})^{\circ k} - \mathbf{I}_N\|_{\text{F}} + 1 \\ &\leq \left(\sum_{1 \leq i \neq j \leq N} \langle \mathbf{w}_{0,i}, \mathbf{w}_{0,j} \rangle^{2k} \right)^{\frac{1}{2}} + 1 = o_{\mathbb{P}}(1) + 1. \end{aligned}$$

Therefore,

$$M_{W_0} \leq \max \left\{ \|\mathbf{W}_0 \mathbf{W}_0^\top\|_{\text{op}}, \|(\mathbf{W}_0 \mathbf{W}_0^\top)^{\circ 2}\|_{\text{op}}, \sup_{k \geq 3} \|(\mathbf{W}_0 \mathbf{W}_0^\top)^{\circ k}\|_{\text{op}} \right\} = O_{\mathbb{P}}(1).$$

Claim (d) is standard, see e.g. [Ver18, Theorem 4.4.5].

O.2 Proof of Lemma K.1

We can write

$$\begin{aligned} \mathbf{v}^\top (\boldsymbol{\beta} - c_{\star,1} \boldsymbol{\beta}_\star) &= n^{-1} \mathbf{v}^\top (\mathbf{X}^\top (\sigma_\star(\mathbf{X} \boldsymbol{\beta}_\star) + \boldsymbol{\varepsilon})) - c_{\star,1} \boldsymbol{\beta}_\star \\ &= n^{-1} \sum_{i=1}^n (\mathbf{v}^\top \mathbf{x}_i \sigma_\star(\mathbf{x}_i^\top \boldsymbol{\beta}_\star) - c_{\star,1} \mathbf{v}^\top \boldsymbol{\beta}_\star) + n^{-1} \mathbf{v}^\top \boldsymbol{\varepsilon}. \end{aligned}$$

Now $n^{-1} \mathbf{v}^\top \boldsymbol{\varepsilon} \sim \mathcal{N}(0, \sigma_\varepsilon^2 \|\mathbf{v}\|_2^2)/n \rightarrow_P 0$. Moreover, by Condition 2.4, we can write $\sigma_\star(\mathbf{x}_i^\top \boldsymbol{\beta}_\star) = c_{\star,0} + c_{\star,1} \mathbf{x}_i^\top \boldsymbol{\beta}_\star + (P_{>1} \sigma_\star)(\mathbf{x}_i^\top \boldsymbol{\beta}_\star)$, where conditional on $\boldsymbol{\beta}_\star$, $(P_{>1} \sigma_\star)(\mathbf{x}_i^\top \boldsymbol{\beta}_\star)$ is orthogonal in L^2 to the constant function and to $\mathbf{x}_i^\top \boldsymbol{\beta}_\star$. Hence the first sum above equals

$$n^{-1} c_{\star,0} \mathbf{v}^\top \sum_{i=1}^n \mathbf{x}_i + n^{-1} c_{\star,1} \mathbf{v}^\top \left(\sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i^\top - \mathbf{I} \right) \boldsymbol{\beta}_\star + n^{-1} \sum_{i=1}^n \mathbf{v}^\top \mathbf{x}_i (P_{>1} \sigma_\star)(\mathbf{x}_i^\top \boldsymbol{\beta}_\star).$$

For the first term, $n^{-1} c_{\star,0} \mathbf{v}^\top \sum_{i=1}^n \mathbf{x}_i \sim n^{-1} c_{\star,0} \cdot \mathcal{N}(0, n \|\mathbf{v}\|_2^2) \rightarrow_P 0$. The second term is $c_{\star,1}$ times a sample mean of i.i.d. random variables of the form $\mathbf{v}^\top (\mathbf{x}_i \mathbf{x}_i^\top - 1) \boldsymbol{\beta}_\star$, which have zero mean by the Gaussianity of \mathbf{x}_i , and for which all moments are finite. Hence, by the weak law of large numbers, this term converges to zero in probability.

Similarly, the third term is a sample mean of i.i.d. random variables of the form $\mathbf{v}^\top \mathbf{x}_i (P_{>1} \sigma_\star)(\mathbf{x}_i^\top \boldsymbol{\beta}_\star)$, which have zero mean by the Gaussianity of \mathbf{x}_i and Lemma C.1, and whose second moments are finite since σ_\star is Lipschitz. Hence, by the weak law of large numbers, this term also converges to zero in probability. This finishes the proof of the first claim.

Next, the second statement follows from [BES⁺22, Lemma 18]. While that work has slightly different assumptions on the teacher function f_\star , it is straightforward to check that their proof goes through unchanged under our assumptions. Specifically, their proof requires that $\mathbf{x} \mapsto f_\star(\mathbf{x}) = \sigma_\star(\mathbf{x}^\top \boldsymbol{\beta}_\star)$ is $O(1)$ -Lipschitz, which holds in our case because σ_\star is $O(1)$ -Lipschitz, and $\|\boldsymbol{\beta}_\star\|_2 = O_{\mathbb{P}}(1)$.

O.3 Proof of Lemma K.2

By plugging in $\hat{\mathbf{a}}$ into the training loss, we find

$$\begin{aligned} \mathcal{L}_{\text{tr}}(\mathbf{F}) &= \frac{1}{n} \|\tilde{\mathbf{y}} - \mathbf{F} \hat{\mathbf{a}}\|_2^2 + \lambda \|\hat{\mathbf{a}}\|_2^2 = \frac{1}{n} \|\tilde{\mathbf{y}}\|_2^2 - \frac{2}{n} \tilde{\mathbf{y}}^\top \mathbf{F} \hat{\mathbf{a}} + \frac{1}{n} \hat{\mathbf{a}}^\top (\mathbf{F}^\top \mathbf{F} + \lambda n \mathbf{I}_N) \hat{\mathbf{a}} \\ &= \frac{1}{n} \|\tilde{\mathbf{y}}\|_2^2 - \frac{1}{n} \tilde{\mathbf{y}}^\top \mathbf{F} \hat{\mathbf{a}} = \frac{1}{n} \|\tilde{\mathbf{y}}\|_2^2 - \frac{1}{n} \tilde{\mathbf{y}}^\top \mathbf{F} (\mathbf{F}^\top \mathbf{F} + \lambda n \mathbf{I}_N)^{-1} \mathbf{F}^\top \tilde{\mathbf{y}} \\ &= \frac{1}{n} \|\tilde{\mathbf{y}}\|_2^2 - \frac{1}{n} \tilde{\mathbf{y}}^\top \mathbf{F} \mathbf{F}^\top (\mathbf{F} \mathbf{F}^\top + \lambda n \mathbf{I}_n)^{-1} \tilde{\mathbf{y}} \\ &= \frac{1}{n} \|\tilde{\mathbf{y}}\|_2^2 - \frac{1}{n} \tilde{\mathbf{y}}^\top (\mathbf{F} \mathbf{F}^\top + \lambda n \mathbf{I}_n) (\mathbf{F} \mathbf{F}^\top + \lambda n \mathbf{I}_n)^{-1} \tilde{\mathbf{y}} + \lambda \tilde{\mathbf{y}}^\top (\mathbf{F} \mathbf{F}^\top + \lambda n \mathbf{I}_n)^{-1} \tilde{\mathbf{y}} \\ &= \lambda \tilde{\mathbf{y}}^\top (\mathbf{F} \mathbf{F}^\top + \lambda n \mathbf{I}_n)^{-1} \tilde{\mathbf{y}}, \end{aligned}$$

which proves the lemma.

O.4 Proof of Lemma K.3

To prove the concentration of this term around its mean, we will use the generalized Hanson-Wright inequality [Sam23, Theorem 2.1] for α -subexponential random variables. Note that, by definition, if Z is a Gaussian random variable, $H_p(Z)$ is $2/p$ -subexponential (see the definition in equation (1.1) of [Sam23] and for these variables the Orlicz norm of order $2/p$ is bounded (see equation (1.3) of [Sam23]). Also note that $\|\mathbf{D}\|_{\text{Fr}} \leq \sqrt{n}\|\mathbf{D}\|_{\text{op}} = O_{\mathbb{P}}(1/\sqrt{n})$. Thus, using [Sam23, Theorem 2.1] and setting $t = \frac{\log(n)}{\sqrt{n}}$, we find

$$\mathbb{P}\left(\left|g(\mathbf{Z})^\top \mathbf{D} g(\mathbf{Z}) - \mathbb{E}[g(\mathbf{Z})^\top \mathbf{D} g(\mathbf{Z})]\right| \geq \frac{\log n}{\sqrt{n}}\right) \leq 2 \exp\left(-C \min\left\{\log^2(n), (\sqrt{n} \log n)^{1/p}\right\}\right),$$

where $C > 0$ is some constant. This concludes the proof.

O.5 Proof of Lemma K.4

First, we show that switching from $\mathbf{w}_{0,i} \stackrel{i.i.d.}{\sim} \text{Unif}(\mathbb{S}^{d-1})$ to $\hat{\mathbf{w}}_{0,i} \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \frac{1}{d}\mathbf{I}_d)$ will not change the limit of the terms $\frac{1}{d}\mathbb{E}\text{tr}(\tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}})$ and $\mathbb{E}\text{tr}(\bar{\mathbf{R}}_0)$ which will appear later in the proof. First, we define $\hat{\mathbf{W}}_0 = [\hat{\mathbf{w}}_{0,1}, \dots, \hat{\mathbf{w}}_{0,N}]^\top$,

$$\begin{aligned} \mathbf{D} &= \text{diag}\left(\frac{1}{\|\hat{\mathbf{w}}_{0,1}\|_2}, \dots, \frac{1}{\|\hat{\mathbf{w}}_{0,N}\|_2}\right), \quad \mathbf{W}_0 =_d \mathbf{D}\hat{\mathbf{W}}_0, \quad \hat{\mathbf{F}}_0 = \sigma(\tilde{\mathbf{X}}\hat{\mathbf{W}}_0^\top), \\ \text{and} \quad \hat{\bar{\mathbf{R}}}_0 &= (\hat{\mathbf{F}}_0 \hat{\mathbf{F}}_0^\top + \lambda n \mathbf{I}_n)^{-1}. \end{aligned}$$

Then,

$$\begin{aligned} \left| \text{tr} \left[\bar{\mathbf{R}}_0 - \hat{\bar{\mathbf{R}}}_0 \right] \right| &= \left| \text{tr} \left[(\mathbf{F}_0 \mathbf{F}_0^\top + \lambda n \mathbf{I}_n)^{-1} - (\hat{\mathbf{F}}_0 \hat{\mathbf{F}}_0^\top + \lambda n \mathbf{I}_n)^{-1} \right] \right| \\ &= \left| \text{tr} \left[(\mathbf{F}_0 \mathbf{F}_0^\top + \lambda n \mathbf{I}_n)^{-1} (\mathbf{F}_0 \mathbf{F}_0^\top - \hat{\mathbf{F}}_0 \hat{\mathbf{F}}_0^\top) (\hat{\mathbf{F}}_0 \hat{\mathbf{F}}_0^\top + \lambda n \mathbf{I}_n)^{-1} \right] \right| \\ &\leq \text{tr}(\mathbf{F}_0 \mathbf{F}_0^\top + \lambda n \mathbf{I}_n)^{-1} \|(\hat{\mathbf{F}}_0 \hat{\mathbf{F}}_0^\top + \lambda n \mathbf{I}_n)^{-1}\|_{\text{op}} \|\mathbf{F}_0 \mathbf{F}_0 - \hat{\mathbf{F}}_0 \hat{\mathbf{F}}_0\|_{\text{op}} \\ &\leq \frac{C}{n} \|\mathbf{F}_0 \mathbf{F}_0 - \hat{\mathbf{F}}_0 \hat{\mathbf{F}}_0\|_{\text{op}}. \end{aligned}$$

Now, using the Gaussian equivalence from Appendix J, we can replace \mathbf{F}_0 and $\hat{\mathbf{F}}_0$ with $\mathbf{F}_0 = c_1 \tilde{\mathbf{X}} \mathbf{W}_0^\top + c_{>1} \mathbf{Z}$ and $\hat{\mathbf{F}}_0 = c_1 \tilde{\mathbf{X}} \hat{\mathbf{W}}_0^\top + c_{>1} \mathbf{Z}$, respectively, without changing the limit. With this, we have

$$\mathbf{F}_0 \mathbf{F}_0^\top - \hat{\mathbf{F}}_0 \hat{\mathbf{F}}_0^\top = c_1^2 \tilde{\mathbf{X}} (\mathbf{W}_0 \mathbf{W}_0^\top - \hat{\mathbf{W}}_0 \hat{\mathbf{W}}_0^\top) \tilde{\mathbf{X}}^\top + c_1 c_{>1} \tilde{\mathbf{X}} (\mathbf{W}_0 - \hat{\mathbf{W}}_0)^\top \mathbf{Z}^\top + c_1 c_{>1} \mathbf{Z} (\mathbf{W}_0 - \hat{\mathbf{W}}_0) \tilde{\mathbf{X}}^\top.$$

Now,

$$\|\mathbf{W}_0 \mathbf{W}_0^\top - \hat{\mathbf{W}}_0 \hat{\mathbf{W}}_0^\top\|_{\text{op}} \leq \|\mathbf{I}_N - \mathbf{D}\|_{\text{op}} \|\mathbf{W}_0 \mathbf{W}_0^\top\|_{\text{op}} (\|\mathbf{D}\|_{\text{op}} + 1).$$

Note that $\|\mathbf{W}_0 \mathbf{W}_0^\top\|_{\text{op}} = O_{\mathbb{P}}(1)$, $\|\mathbf{D}\|_{\text{op}} = O_{\mathbb{P}}(1)$, and $\|\mathbf{I}_N - \mathbf{D}\|_{\text{op}} = o_{\mathbb{P}}(1)$. Thus $\|\mathbf{W}_0 \mathbf{W}_0^\top - \hat{\mathbf{W}}_0 \hat{\mathbf{W}}_0^\top\|_{\text{op}} = o_{\mathbb{P}}(1)$. Also, similarly, $\|\mathbf{W}_0 - \hat{\mathbf{W}}_0\|_{\text{op}} = o_{\mathbb{P}}(1)$. Hence, noting that $\|\tilde{\mathbf{X}}\|_{\text{op}}$ and $\|\mathbf{Z}\|_{\text{op}}$

are both $O_{\mathbb{P}}(\sqrt{N})$, we have $\frac{1}{n} \|\mathbf{F}_0 \mathbf{F}_0^\top - \hat{\mathbf{F}}_0 \hat{\mathbf{F}}_0^\top\|_{\text{op}} \rightarrow_P 0$. This implies that $|\text{tr}[\bar{\mathbf{R}}_0 - \hat{\bar{\mathbf{R}}}_0]| = o_{\mathbb{P}}(1)$. Also,

$$\left| \frac{1}{d} \text{tr} [\tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}}] - \frac{1}{d} \text{tr} [\tilde{\mathbf{X}}^\top \hat{\bar{\mathbf{R}}}_0 \tilde{\mathbf{X}}] \right| \leq |\text{tr}[\bar{\mathbf{R}}_0 - \hat{\bar{\mathbf{R}}}_0]| \frac{\|\tilde{\mathbf{X}} \tilde{\mathbf{X}}^\top\|_{\text{op}}}{d} \rightarrow_P 0.$$

Finally, we can prove the required claims as follows:

- (a) Since $\beta_\star \sim \mathcal{N}(0, \frac{1}{d} \mathbf{I}_d)$, we have $\beta_\star^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}} \beta_\star = \frac{1}{d} \mathbb{E} \text{tr}(\tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}}) + o_{\mathbb{P}}(1)$, by the Hanson-Wright inequality. Note that by the argument above, we can assume that $\hat{w}_{0,i} \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \frac{1}{d} \mathbf{I}_d)$ without changing the limiting trace. Further, from [AP20a, Proposition 1], see also [ALP22], we have $\frac{1}{d} \mathbb{E} \text{tr}(\tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0 \tilde{\mathbf{X}}) \rightarrow \frac{\psi}{\phi} m_2$; see the discussion at the end of this proof for the detailed explanation. Now, we arrive at the conclusion by applying Lemma K.1.
- (b) Since $\mathbf{a} \sim \mathcal{N}(0, \frac{1}{N} \mathbf{I}_N)$, we have $\mathbf{a}^\top \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a} - \|\mathbf{a}\|_2^2 = \frac{1}{N} \text{tr}(\mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0) - 1 + o_{\mathbb{P}}(1)$ by the Hanson-Wright inequality. Moreover,

$$\begin{aligned} \mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0 &= \mathbf{F}_0^\top \mathbf{F}_0 (\mathbf{F}_0^\top \mathbf{F}_0 + \lambda n \mathbf{I}_N)^{-1} \\ &= (\mathbf{F}_0^\top \mathbf{F}_0 + \lambda n \mathbf{I}_N - \lambda n \mathbf{I}_N) (\mathbf{F}_0^\top \mathbf{F}_0 + \lambda n \mathbf{I}_N)^{-1} = \mathbf{I}_N - \lambda n (\mathbf{F}_0^\top \mathbf{F}_0 + \lambda n \mathbf{I}_N)^{-1}. \end{aligned}$$

Hence, $\frac{1}{N} \text{tr}(\mathbf{F}_0^\top \bar{\mathbf{R}}_0 \mathbf{F}_0) - 1 = -\frac{\lambda n}{N} \text{tr}(\mathbf{F}_0^\top \mathbf{F}_0 + \lambda n \mathbf{I}_N)^{-1}$. From the argument above, we can assume that $\hat{w}_{0,i} \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \frac{1}{d} \mathbf{I}_d)$ without changing the limiting trace. It follows from [AP20a, Proposition 1] that $\mathbb{E} \text{tr} \bar{\mathbf{R}}_0 \rightarrow \frac{\psi}{\phi} m_1$; again see the discussion at the end of this proof for the detailed explanation. Note that $\lim \mathbb{E} \text{tr}(\mathbf{F}_0^\top \mathbf{F}_0 + \lambda n \mathbf{I}_N)^{-1}$ is the limiting companion Stieltjes transform of $\mathbf{F}_0 \mathbf{F}_0^\top$. Hence, $\bar{m}_1 = \lim \mathbb{E} \text{tr}(\mathbf{F}_0^\top \mathbf{F}_0 + \lambda n \mathbf{I}_N)^{-1}$ is the limiting companion Stieltjes transform of m_1 which is given by

$$\bar{m}_1 = \frac{\psi}{\phi} m_1 - \left(1 - \frac{\phi}{\psi}\right) \frac{1}{\lambda}. \quad (89)$$

This concludes the proof.

For the reader's convenience, we provide the following diagram that shows how the notations of [AP20a] (left) match (\Leftrightarrow) ours (right):

$$\begin{aligned} n_0 &\Leftrightarrow d, & n_1 &\Leftrightarrow N, & m &\Leftrightarrow n, & \phi, \psi &\Leftrightarrow \phi, \psi, \\ \mathbf{X}^\top \in \mathbb{R}^{m \times n_0} &\Leftrightarrow \tilde{\mathbf{X}} \in \mathbb{R}^{n \times d}, & \mathbf{F}^\top \in \mathbb{R}^{m \times n_1} &\Leftrightarrow \mathbf{F}_0 \in \mathbb{R}^{n \times N}, & \sigma_{W_2} &= 0, \\ \frac{1}{n_1} \mathbf{K}(\lambda m / n_1)^{-1} &= \frac{1}{n_1} \mathbf{F}^\top \mathbf{F} + \lambda \mathbf{I}_m \Leftrightarrow \bar{\mathbf{R}}_0^{-1} = \mathbf{F}_0 \mathbf{F}_0^\top + \lambda n \mathbf{I}_n, & \zeta &\Leftrightarrow c_1^2, & \eta &\Leftrightarrow c_1^2 + c_{>1}^2, \\ \tau_1 = \frac{1}{m} \mathbb{E} \text{tr} \mathbf{K}^{-1} &\Leftrightarrow m_1 = \frac{N}{n} \mathbb{E} \text{tr} \bar{\mathbf{R}}_0, & \tau_2 = \frac{1}{mn_0} \mathbb{E} \text{tr} \mathbf{X}^\top \mathbf{X} \mathbf{K}^{-1} &\Leftrightarrow m_2 = \frac{N}{nd} \mathbb{E} \text{tr} \tilde{\mathbf{X}} \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0. \end{aligned}$$

O.6 Proof of Lemma K.5

Define $\hat{\mathbf{X}} = \tilde{\mathbf{X}} - \tilde{\mathbf{X}} \mathbf{u} \mathbf{u}^\top$, which implies $\hat{\mathbf{X}} \perp \tilde{\mathbf{X}} \mathbf{u}$ due to the Gaussianity of \mathbf{X} . Based on the Gaussian equivalence from Appendix J, we can replace \mathbf{F}_0 with $c_1 \tilde{\mathbf{X}} \mathbf{W}_0^\top + c_{>1} \mathbf{Z}$, where $\mathbf{Z} \in \mathbb{R}^{n \times d}$ is an independent random matrix with $\mathcal{N}(0, 1)$ entries, without changing the conclusion. Hence, from now on, we write $\mathbf{F}_0 = c_1 \tilde{\mathbf{X}} \mathbf{W}_0^\top + c_{>1} \mathbf{Z}$. Further, we define

$$\hat{\mathbf{F}}_0 = c_1 \hat{\mathbf{X}} \mathbf{W}_0^\top + c_{>1} \mathbf{Z}. \quad (90)$$

Thus, by the definition of $\hat{\mathbf{X}}$, $\hat{\mathbf{F}}_0 = \mathbf{F}_0 - c_1 \tilde{\mathbf{X}}\mathbf{u}(\mathbf{W}_0\mathbf{u})^\top$. As a consequence, we also have $\mathbf{F}_0\mathbf{F}_0^\top = \hat{\mathbf{F}}_0\hat{\mathbf{F}}_0^\top + \mathbf{V}\mathbf{D}\mathbf{V}^\top$, where $\mathbf{V} = [\hat{\mathbf{F}}_0\mathbf{W}_0\mathbf{u} \quad \tilde{\mathbf{X}}\mathbf{u}] \in \mathbb{R}^{n \times 2}$ and

$$\mathbf{D} = \begin{bmatrix} 0 & c_1 \\ c_1 & c_1^2 \|\mathbf{W}_0\mathbf{u}\|_2^2 \end{bmatrix}.$$

Noting that D is invertible, and using the Woodbury formula, with $\hat{\mathbf{R}}_0 = (\hat{\mathbf{F}}_0\hat{\mathbf{F}}_0^\top + \lambda n \mathbf{I}_n)^{-1}$, we find

$$\bar{\mathbf{R}}_0 = \hat{\mathbf{R}}_0 - \hat{\mathbf{R}}_0\mathbf{V}(\mathbf{D}^{-1} + \mathbf{V}^\top\hat{\mathbf{R}}_0\mathbf{V})^{-1}\mathbf{V}^\top\hat{\mathbf{R}}_0.$$
 (91)

Now, we can write

$$H_q(\tilde{\mathbf{X}}\mathbf{u})^\top\bar{\mathbf{R}}_0H_p(\tilde{\mathbf{X}}\mathbf{u}) = H_q(\tilde{\mathbf{X}}\mathbf{u})^\top\hat{\mathbf{R}}_0H_p(\tilde{\mathbf{X}}\mathbf{u}) - H_q(\tilde{\mathbf{X}}\mathbf{u})^\top\hat{\mathbf{R}}_0\mathbf{V}(\mathbf{D}^{-1} + \mathbf{V}^\top\hat{\mathbf{R}}_0\mathbf{V})^{-1}\mathbf{V}^\top\hat{\mathbf{R}}_0H_p(\tilde{\mathbf{X}}\mathbf{u}).$$

Next, we can analyze each term in the above sum separately.

The first term on the right hand side converges to zero by using Lemma K.3 to prove the concentration of this term around its mean and noting that the mean is zero using the orthogonality property of Hermite polynomials (Lemma C.1).

To analyze the second term, we first study the matrix $\mathbf{K} = (\mathbf{D}^{-1} + \mathbf{V}^\top\hat{\mathbf{R}}_0\mathbf{V})^{-1}$, writing

$$\mathbf{K}^{-1} = (\mathbf{D}^{-1} + \mathbf{V}^\top\hat{\mathbf{R}}_0\mathbf{V}) = \begin{bmatrix} \mathbf{u}^\top\mathbf{W}_0^\top\hat{\mathbf{F}}_0^\top\hat{\mathbf{R}}_0\hat{\mathbf{F}}_0\mathbf{W}_0\mathbf{u} - \|\mathbf{W}_0\mathbf{u}\|_2^2 & \mathbf{u}^\top\mathbf{W}_0^\top\hat{\mathbf{F}}_0^\top\hat{\mathbf{R}}_0\tilde{\mathbf{X}}\mathbf{u} - \frac{1}{c_1} \\ \mathbf{u}^\top\tilde{\mathbf{X}}^\top\hat{\mathbf{R}}_0\hat{\mathbf{F}}_0\mathbf{W}_0\mathbf{u} - \frac{1}{c_1} & \mathbf{u}^\top\tilde{\mathbf{X}}^\top\hat{\mathbf{R}}_0\tilde{\mathbf{X}}\mathbf{u} \end{bmatrix}.$$

It can readily verified that all elements in this matrix are $O_{\mathbb{P}}(1)$ by checking the order of the operator and Euclidean norms. Next, we analyze the terms in the expression

$$\begin{aligned} H_q(\tilde{\mathbf{X}}\mathbf{u})^\top\hat{\mathbf{R}}_0\mathbf{V}\mathbf{K}\mathbf{V}^\top\hat{\mathbf{R}}_0H_p(\tilde{\mathbf{X}}\mathbf{u}) &= [\mathbf{K}]_{1,1}H_q(\tilde{\mathbf{X}}\mathbf{u})^\top\hat{\mathbf{R}}_0(\hat{\mathbf{F}}_0\mathbf{W}_0\mathbf{u})(\hat{\mathbf{F}}_0\mathbf{W}_0\mathbf{u})^\top\hat{\mathbf{R}}_0H_p(\tilde{\mathbf{X}}\mathbf{u}) \\ &\quad + [\mathbf{K}]_{1,2}H_q(\tilde{\mathbf{X}}\mathbf{u})^\top\hat{\mathbf{R}}_0(\hat{\mathbf{F}}_0\mathbf{W}_0\mathbf{u})(\tilde{\mathbf{X}}\mathbf{u})^\top\hat{\mathbf{R}}_0H_p(\tilde{\mathbf{X}}\mathbf{u}) \\ &\quad + [\mathbf{K}]_{2,1}H_q(\tilde{\mathbf{X}}\mathbf{u})^\top\hat{\mathbf{R}}_0(\tilde{\mathbf{X}}\mathbf{u})(\hat{\mathbf{F}}_0\mathbf{W}_0\mathbf{u})^\top\hat{\mathbf{R}}_0H_p(\tilde{\mathbf{X}}\mathbf{u}) \\ &\quad + [\mathbf{K}]_{2,2}H_q(\tilde{\mathbf{X}}\mathbf{u})^\top\hat{\mathbf{R}}_0(\tilde{\mathbf{X}}\mathbf{u})(\tilde{\mathbf{X}}\mathbf{u})^\top\hat{\mathbf{R}}_0H_p(\tilde{\mathbf{X}}\mathbf{u}). \end{aligned}$$

Without loss of generality, we can assume that $p \neq 1$.

- **First Term.** Note that $H_q(\tilde{\mathbf{X}}\mathbf{u})^\top$ and $H_p(\tilde{\mathbf{X}}\mathbf{u})$ are orthogonal in L^2 by the properties of the Hermite polynomials, and conditional on \mathbf{u} , they are independent of $\hat{\mathbf{R}}_0(\hat{\mathbf{F}}_0\mathbf{W}_0\mathbf{u})(\hat{\mathbf{F}}_0\mathbf{W}_0\mathbf{u})^\top\hat{\mathbf{R}}_0$. Moreover, $\|\hat{\mathbf{R}}_0(\hat{\mathbf{F}}_0\mathbf{W}_0\mathbf{u})(\hat{\mathbf{F}}_0\mathbf{W}_0\mathbf{u})^\top\hat{\mathbf{R}}_0\|_{\text{op}} = O_{\mathbb{P}}(1/n)$. Thus, by using Lemma K.3, this term converges to zero.
- **Second Term.** Similar to the argument above, we can show that $(\tilde{\mathbf{X}}\mathbf{u})^\top\hat{\mathbf{R}}_0H_p(\tilde{\mathbf{X}}\mathbf{u})$ converges to zero. Also, by analyzing the operator norms, we have $H_q(\tilde{\mathbf{X}}\mathbf{u})^\top\hat{\mathbf{R}}_0(\hat{\mathbf{F}}_0\mathbf{W}_0\mathbf{u}) = O(1)$. This implies that the second term converges to zero.
- **Third Term.** First, note that by a simple order-wise analysis, $H_q(\tilde{\mathbf{X}}\mathbf{u})^\top\hat{\mathbf{R}}_0(\tilde{\mathbf{X}}\mathbf{u}) = O_{\mathbb{P}}(1)$. Now, we have $H_p(\tilde{\mathbf{X}}\mathbf{u})$ is independent of $(\hat{\mathbf{F}}_0\mathbf{W}_0\mathbf{u})^\top\hat{\mathbf{R}}_0$ and $\|(\hat{\mathbf{F}}_0\mathbf{W}_0\mathbf{u})^\top\hat{\mathbf{R}}_0\|_2 = O_{\mathbb{P}}(1/\sqrt{n})$. The term $(\hat{\mathbf{F}}_0\mathbf{W}_0\mathbf{u})^\top\hat{\mathbf{R}}_0H_p(\tilde{\mathbf{X}}\mathbf{u})$ converges to zero in probability by noting that $H_p(\tilde{\mathbf{X}}\mathbf{u})$ is mean zero for $p \neq 0$. For the $p = 0$ case, we can use an orthogonality invariance argument identical to the one used to analyze equation 34.
- **Fourth Term.** This term also converges to zero because $(\tilde{\mathbf{X}}\mathbf{u})^\top\hat{\mathbf{R}}_0H_p(\tilde{\mathbf{X}}\mathbf{u})$ converges to zero, as argued above.

Putting everything together, the proof is completed.

O.7 Proof of Lemma K.6

We will prove part (a) first. To do this, we will first handle the cases where $p = 0$ and $p = 1$.

For $p = 0$, we have $\sqrt{N}H_0(\tilde{\boldsymbol{\theta}}_*)\bar{\mathbf{R}}_0\mathbf{F}_0\mathbf{a}^{\circ 2} = \sqrt{N}$. This is identical to the second term in equation 34 and it is shown to be $O_{\mathbb{P}}(1)$.

For $p = 1$, we need to analyze $\sqrt{N}H_1(\tilde{\boldsymbol{\theta}}_*)\bar{\mathbf{R}}_0\mathbf{F}_0\mathbf{a}^{\circ 2} = \sqrt{N}\boldsymbol{\beta}_*^\top \tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0\mathbf{F}_0\mathbf{a}^{\circ 2}$. Note that $\boldsymbol{\beta}_* \sim \mathcal{N}(0, \frac{1}{d}\mathbf{I}_d)$ is independent of $\sqrt{N}\tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0\mathbf{F}_0\mathbf{a}^{\circ 2}$ and

$$\|\sqrt{N}\tilde{\mathbf{X}}^\top \bar{\mathbf{R}}_0\mathbf{F}_0\mathbf{a}^{\circ 2}\|_2 \leq \sqrt{N}\|\tilde{\mathbf{X}}\|_{\text{op}} \cdot \|\bar{\mathbf{R}}_0\|_{\text{op}} \cdot \|\mathbf{F}_0\|_{\text{op}} \cdot \|\mathbf{a}^{\circ 2}\|_2 = O_{\mathbb{P}}(1).$$

Thus, we can conclude that $\sqrt{N}H_1(\tilde{\boldsymbol{\theta}}_*)\bar{\mathbf{R}}_0\mathbf{F}_0\mathbf{a}^{\circ 2} \rightarrow 0$ in probability.

To analyze the case where $p > 1$, we first define $\hat{\mathbf{X}} = \tilde{\mathbf{X}} - \tilde{\boldsymbol{\theta}}_*\boldsymbol{\beta}_*^\top$. By construction, we have $\hat{\mathbf{X}} \perp \tilde{\boldsymbol{\theta}}_*$. As in the proof of Lemma K.5, Based on the Gaussian equivalence from Appendix J, we can replace \mathbf{F}_0 with $c_1\tilde{\mathbf{X}}\mathbf{W}_0^\top + c_{>1}\mathbf{Z}$ in our computations without changing the limiting result, where $\mathbf{Z} \in \mathbb{R}^{n \times d}$ is an independent random matrix with $\mathcal{N}(0, 1)$ entries. Thus, from now on, we denote $\mathbf{F}_0 = c_1\tilde{\mathbf{X}}\mathbf{W}_0^\top + c_{>1}\mathbf{Z}$. We define $\hat{\mathbf{F}}_0$ as in equation 90. Thus, $\hat{\mathbf{F}}_0 = \mathbf{F}_0 - c_1\tilde{\boldsymbol{\theta}}_*(\mathbf{W}_0\boldsymbol{\beta}_*)^\top$. As a consequence, we can write $\mathbf{F}_0\mathbf{F}_0^\top = \hat{\mathbf{F}}_0\hat{\mathbf{F}}_0^\top + \mathbf{V}\mathbf{D}\mathbf{V}^\top$, where $\mathbf{V} = [\hat{\mathbf{F}}_0\mathbf{W}_0\boldsymbol{\beta}_* \quad \tilde{\boldsymbol{\theta}}_*] \in \mathbb{R}^{n \times 2}$ and

$$\mathbf{D} = \begin{bmatrix} 0 & c_1 \\ c_1 & c_1^2\|\mathbf{W}_0\boldsymbol{\beta}_*\|_2^2 \end{bmatrix}.$$

Using the Woodbury formula, we find that equation 91 still holds. Now, we can write

$$\begin{aligned} & \sqrt{N}H_p(\tilde{\boldsymbol{\theta}}_*)^\top \bar{\mathbf{R}}_0\mathbf{F}_0\mathbf{a}^{\circ 2} \\ &= \sqrt{N}H_p(\tilde{\boldsymbol{\theta}}_*)^\top \hat{\mathbf{R}}_0\mathbf{F}_0\mathbf{a}^{\circ 2} - \sqrt{N}H_p(\tilde{\boldsymbol{\theta}}_*)^\top \hat{\mathbf{R}}_0\mathbf{V}(\mathbf{D}^{-1} + \mathbf{V}^\top \hat{\mathbf{R}}_0\mathbf{V})^{-1}\mathbf{V}^\top \hat{\mathbf{R}}_0\mathbf{F}_0\mathbf{a}^{\circ 2} \\ &= \sqrt{N}H_p(\tilde{\boldsymbol{\theta}}_*)^\top \hat{\mathbf{R}}_0(\hat{\mathbf{F}}_0 + c_1\tilde{\boldsymbol{\theta}}_*(\mathbf{W}_0\boldsymbol{\beta}_*)^\top)\mathbf{a}^{\circ 2} \\ &\quad - \sqrt{N}H_p(\tilde{\boldsymbol{\theta}}_*)^\top \hat{\mathbf{R}}_0\mathbf{V}(\mathbf{D}^{-1} + \mathbf{V}^\top \hat{\mathbf{R}}_0\mathbf{V})^{-1}\mathbf{V}^\top \hat{\mathbf{R}}_0(\hat{\mathbf{F}}_0 + c_1\tilde{\boldsymbol{\theta}}_*(\mathbf{W}_0\boldsymbol{\beta}_*)^\top)\mathbf{a}^{\circ 2}. \end{aligned} \tag{92}$$

Now, we can analyze each term in the above sum separately.

Term 1. Note that by a simple orderwise analysis,

$$\|\sqrt{N}\hat{\mathbf{R}}_0\hat{\mathbf{F}}_0\mathbf{a}^{\circ 2}\|_{\text{op}} \leq \sqrt{N}\|\hat{\mathbf{R}}_0\|_{\text{op}}\|\hat{\mathbf{F}}_0\|_{\text{op}}\|\mathbf{a}^{\circ 2}\|_2 = O(1/\sqrt{N}).$$

We have $\|H_p(\tilde{\boldsymbol{\theta}}_*)\|_2 = O_{\mathbb{P}}(\sqrt{N})$, $\mathbb{E}[H_p(\tilde{\boldsymbol{\theta}}_*)] = 0$, and $H_p(\tilde{\boldsymbol{\theta}}_*)$ has independent entries. Also $H_p(\tilde{\boldsymbol{\theta}}_*) \perp \hat{\mathbf{R}}_0\hat{\mathbf{F}}_0\mathbf{a}^{\circ 2}$. Thus, $\sqrt{N}H_p(\tilde{\boldsymbol{\theta}}_*)^\top \hat{\mathbf{R}}_0\hat{\mathbf{F}}_0\mathbf{a}^{\circ 2} \rightarrow_P 0$.

We now need to analyze $\sqrt{N}H_p(\tilde{\boldsymbol{\theta}}_*)^\top \hat{\mathbf{R}}_0\tilde{\boldsymbol{\theta}}_*\boldsymbol{\beta}_*^\top \mathbf{W}_0^\top \mathbf{a}^{\circ 2}$. Note that $H_p(\tilde{\boldsymbol{\theta}}_*)^\top \hat{\mathbf{R}}_0\tilde{\boldsymbol{\theta}}_* = O_{\mathbb{P}}(1)$ by a simple order analysis of the norms. We also have $\sqrt{N}\boldsymbol{\beta}_*^\top \mathbf{W}_0^\top \mathbf{a}^{\circ 2} \rightarrow_P 0$, because $\boldsymbol{\beta}_* \sim \mathcal{N}(0, \frac{1}{d}\mathbf{I}_d)$ is independent of the norm bounded vector $\sqrt{N}\mathbf{W}_0^\top \mathbf{a}^{\circ 2}$.

Term 2. To analyze the second term, we first study the matrix $\mathbf{K} = (\mathbf{D}^{-1} + \mathbf{V}^\top \hat{\mathbf{R}}_0\mathbf{V})^{-1}$:

$$\mathbf{K}^{-1} = (\mathbf{D}^{-1} + \mathbf{V}^\top \hat{\mathbf{R}}_0\mathbf{V}) = \begin{bmatrix} \boldsymbol{\beta}_*^\top \mathbf{W}_0^\top \hat{\mathbf{F}}_0^\top \hat{\mathbf{R}}_0 \hat{\mathbf{F}}_0 \mathbf{W}_0 \boldsymbol{\beta}_* - \|\mathbf{W}_0 \boldsymbol{\beta}_*\|_2^2 & \boldsymbol{\beta}_*^\top \mathbf{W}_0^\top \hat{\mathbf{F}}_0^\top \hat{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}_* - \frac{1}{c_1} \\ \boldsymbol{\beta}_*^\top \tilde{\mathbf{X}}^\top \hat{\mathbf{R}}_0 \hat{\mathbf{F}}_0 \mathbf{W}_0 \boldsymbol{\beta}_* - \frac{1}{c_1} & \boldsymbol{\beta}_*^\top \tilde{\mathbf{X}}^\top \hat{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}_* \end{bmatrix}.$$

By orderwise analysis, all elements in this matrix converge to deterministic $O_{\mathbb{P}}(1)$ values in probability. We write the second term in equation 92 as follows:

$$\begin{aligned}
& \sqrt{N} H_p(\tilde{\theta}_*)^\top \hat{\mathbf{R}}_0 \mathbf{V} \mathbf{K} \mathbf{V}^\top \hat{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}^{\circ 2} \\
&= [\mathbf{K}]_{1,1} H_p(\tilde{\theta}_*)^\top \hat{\mathbf{R}}_0 (\hat{\mathbf{F}}_0 \mathbf{W}_0 \beta_*) (\hat{\mathbf{F}}_0 \mathbf{W}_0 \beta_*)^\top \hat{\mathbf{R}}_0 \mathbf{F}_0 (\sqrt{N} \mathbf{a}^{\circ 2}) \\
&+ [\mathbf{K}]_{1,2} H_p(\tilde{\theta}_*)^\top \hat{\mathbf{R}}_0 (\hat{\mathbf{F}}_0 \mathbf{W}_0 \beta_*) \tilde{\theta}_*^\top \hat{\mathbf{R}}_0 \mathbf{F}_0 (\sqrt{N} \mathbf{a}^{\circ 2}) \\
&+ [\mathbf{K}]_{2,1} H_p(\tilde{\theta}_*)^\top \hat{\mathbf{R}}_0 (\tilde{\theta}_*) (\hat{\mathbf{F}}_0 \mathbf{W}_0 \beta_*)^\top \hat{\mathbf{R}}_0 \mathbf{F}_0 (\sqrt{N} \mathbf{a}^{\circ 2}) \\
&+ [\mathbf{K}]_{2,2} H_p(\tilde{\theta}_*)^\top \hat{\mathbf{R}}_0 (\tilde{\theta}_*) \tilde{\theta}_*^\top \hat{\mathbf{R}}_0 \mathbf{F}_0 (\sqrt{N} \mathbf{a}^{\circ 2}).
\end{aligned}$$

In the sum above, we will show that each term converges to zero.

- First term: By orderwise analysis, we have $\|\hat{\mathbf{R}}_0 (\hat{\mathbf{F}}_0 \mathbf{W}_0 \beta_*)\|_{\text{op}} = O_{\mathbb{P}}(1/\sqrt{N})$. Further, $H_p(\tilde{\theta}_*)$ is independent of it (only considering the randomness in \mathbf{X}) with mean zero and $\|H_p(\tilde{\theta}_*)\|_2 = O_{\mathbb{P}}(\sqrt{N})$. This implies that

$$H_p(\tilde{\theta}_*)^\top \hat{\mathbf{R}}_0 (\hat{\mathbf{F}}_0 \mathbf{W}_0 \beta_*) \rightarrow_P 0. \quad (93)$$

We can use a simple order argument to show that $\sqrt{N} (\hat{\mathbf{F}}_0 \mathbf{W}_0 \beta_*)^\top \hat{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}^{\circ 2} = O_{\mathbb{P}}(1)$. Thus, the first term converges to zero.

- Second term: For this term, we use that $H_p(\tilde{\theta}_*)^\top \hat{\mathbf{R}}_0 (\hat{\mathbf{F}}_0 \mathbf{W}_0 \beta_*) \rightarrow_P 0$. We can also use an orderwise analysis to prove that $\sqrt{N} (\tilde{\theta}_*)^\top \hat{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}^{\circ 2} = O_{\mathbb{P}}(1)$. This proves that the second term also converges to zero.
- Third term: By a simple orderwise analysis, we have $\sqrt{N} (\hat{\mathbf{F}}_0 \mathbf{W}_0 \beta_*)^\top \hat{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}^{\circ 2} = O_{\mathbb{P}}(1)$. To show that the third term converges to zero, it is enough to show that $H_p(\tilde{\theta}_*)^\top \hat{\mathbf{R}}_0 (\tilde{\theta}_*) \rightarrow_P 0$, which is true for $p \neq 1$ by using Lemma K.3 and the orthogonality property of Hermite polynomials (Lemma C.1).
- Fourth term: By a simple orderwise analysis, we have $\sqrt{N} \tilde{\theta}_*^\top \hat{\mathbf{R}}_0 \mathbf{F}_0 \mathbf{a}^{\circ 2} = O_{\mathbb{P}}(1)$. Again, to show that the fourth term converges to zero, it is enough to show that $H_p(\tilde{\theta}_*)^\top \hat{\mathbf{R}}_0 (\tilde{\theta}_*) \rightarrow_P 0$, which is true for $p \neq 1$ as argued above.

Putting everything together, part (a) follows. The proof for part (b) is identical and omitted.

O.8 Proof of Lemma K.7

We will study the cases where $s = 1$ and $s = 2$ separately. For $s = 1$, we can use Lemma K.1 to show that $H_p(\tilde{\theta}_*) \bar{\mathbf{R}}_0 \tilde{\theta} = c_{*,1} H_p(\tilde{\theta}_*) \bar{\mathbf{R}}_0 \tilde{\theta}_* + o_{\mathbb{P}}(1)$. Also, by Lemma K.5, we have $H_p(\tilde{\theta}_*) \bar{\mathbf{R}}_0(\tilde{\theta}_*) = o(1)$ in probability if $p \neq 1$, which proves the lemma.

For the case $s = 2$, we define $\tilde{\beta} = \beta / \|\beta\|_2$ and write

$$H_p(\tilde{\theta}_*) \bar{\mathbf{R}}_0(\tilde{\theta})^{\circ 2} = \|\beta\|_2^2 H_p(\tilde{\theta}_*) \bar{\mathbf{R}}_0(\tilde{\mathbf{X}} \tilde{\beta})^{\circ 2} = \|\beta\|_2^2 H_p(\tilde{\theta}_*) \bar{\mathbf{R}}_0 H_2(\tilde{\mathbf{X}} \tilde{\beta}) + o_{\mathbb{P}}(1).$$

Now, we define $\beta_{\perp} = \frac{\beta_{*} - \langle \beta_{*}, \tilde{\beta} \rangle \tilde{\beta}}{\|\beta_{*} - \langle \beta_{*}, \tilde{\beta} \rangle \tilde{\beta}\|_2}$, and set $\hat{\mathbf{X}} = \tilde{\mathbf{X}} - \tilde{\mathbf{X}} \tilde{\beta} \tilde{\beta}^\top - \tilde{\mathbf{X}} \beta_{\perp} \beta_{\perp}^\top$. By construction, we have $\hat{\mathbf{X}} \perp \tilde{\mathbf{X}} \tilde{\beta}, \tilde{\theta}_*$. Based on the Gaussian equivalence from Appendix J, we can again replace \mathbf{F}_0 with

$\mathbf{F}_0 = c_1 \tilde{\mathbf{X}} \mathbf{W}_0^\top + c_{>1} \mathbf{Z}$, where $\mathbf{Z} \in \mathbb{R}^{n \times d}$ is an independent random matrix with $\mathcal{N}(0, 1)$ entries. Again, we define $\hat{\mathbf{F}}_0$ as in equation 90. Thus, $\hat{\mathbf{F}}_0 = \mathbf{F}_0 - c_1 \tilde{\mathbf{X}} \tilde{\boldsymbol{\beta}} (\mathbf{W}_0 \tilde{\boldsymbol{\beta}})^\top - c_1 \tilde{\mathbf{X}} \boldsymbol{\beta}_\perp (\mathbf{W}_0 \boldsymbol{\beta}_\perp)^\top$. As a consequence, we also have $\mathbf{F}_0 \mathbf{F}_0^\top = \hat{\mathbf{F}}_0 \hat{\mathbf{F}}_0^\top + \mathbf{V} \mathbf{D} \mathbf{V}^\top$, where $\mathbf{V} = [\tilde{\mathbf{X}} \tilde{\boldsymbol{\beta}} \quad \tilde{\mathbf{X}} \boldsymbol{\beta}_\perp \quad \hat{\mathbf{F}}_0 \mathbf{W}_0 \tilde{\boldsymbol{\beta}} \quad \hat{\mathbf{F}}_0 \mathbf{W}_0 \boldsymbol{\beta}_\perp] \in \mathbb{R}^{n \times 4}$ and

$$\mathbf{D} = \begin{bmatrix} c_1^2 \langle \mathbf{W}_0 \tilde{\boldsymbol{\beta}}, \mathbf{W}_0 \tilde{\boldsymbol{\beta}} \rangle & c_1^2 \langle \mathbf{W}_0 \tilde{\boldsymbol{\beta}}, \mathbf{W}_0 \boldsymbol{\beta}_\perp \rangle & c_1 & 0 \\ c_1^2 \langle \mathbf{W}_0 \tilde{\boldsymbol{\beta}}, \mathbf{W}_0 \boldsymbol{\beta}_\perp \rangle & c_1^2 \langle \mathbf{W}_0 \boldsymbol{\beta}_\perp, \mathbf{W}_0 \boldsymbol{\beta}_\perp \rangle & 0 & c_1 \\ c_1 & 0 & 0 & 0 \\ 0 & c_1 & 0 & 0 \end{bmatrix}.$$

Using the Woodbury formula, we find that equation 91 still holds. We can write

$$\begin{aligned} H_p(\tilde{\boldsymbol{\theta}}_*)^\top \bar{\mathbf{R}}_0 H_2(\tilde{\mathbf{X}} \tilde{\boldsymbol{\beta}}) &= H_p(\tilde{\boldsymbol{\theta}}_*)^\top \hat{\mathbf{R}}_0 H_2(\tilde{\mathbf{X}} \tilde{\boldsymbol{\beta}}) \\ &\quad - H_p(\tilde{\boldsymbol{\theta}}_*)^\top \hat{\mathbf{R}}_0 \mathbf{V} (\mathbf{D}^{-1} + \mathbf{V}^\top \hat{\mathbf{R}}_0 \mathbf{V})^{-1} \mathbf{V}^\top \hat{\mathbf{R}}_0 H_2(\tilde{\mathbf{X}} \tilde{\boldsymbol{\beta}}). \end{aligned} \quad (94)$$

The first term converges to zero for any $p \neq 2$, analogously to the argument in Section K.2.1 for the term (1,2).

To prove that the second term will also converge to zero, we first observe that the elements of $\mathbf{K} = (\mathbf{D}^{-1} + \mathbf{V}^\top \hat{\mathbf{R}}_0 \mathbf{V})^{-1}$ are all $O_{\mathbb{P}}(1)$. The second term will involve quantities of the form $[\mathbf{K}]_{i,j} H_p(\tilde{\boldsymbol{\theta}}_*)^\top \hat{\mathbf{R}}_0 \mathbf{v}_i \mathbf{v}_j^\top \hat{\mathbf{R}}_0 H_2(\tilde{\mathbf{X}} \tilde{\boldsymbol{\beta}})$, where \mathbf{v}_i , for $i \in \{1, 2, 3, 4\}$, is the i -th column of the matrix $\mathbf{V} = [\tilde{\mathbf{X}} \tilde{\boldsymbol{\beta}} \quad \tilde{\mathbf{X}} \boldsymbol{\beta}_\perp \quad \hat{\mathbf{F}}_0 \mathbf{W}_0 \tilde{\boldsymbol{\beta}} \quad \hat{\mathbf{F}}_0 \mathbf{W}_0 \boldsymbol{\beta}_\perp]$. We can argue that all these terms converge to zero, as follows:

- The terms where $j = 1$ converge to zero because $(\tilde{\mathbf{X}} \tilde{\boldsymbol{\beta}})^\top \hat{\mathbf{R}}_0 H_2(\tilde{\mathbf{X}} \tilde{\boldsymbol{\beta}})$ converges to zero analogously to the argument in Section K.2.1 for the term (1,2). The same argument applies to the terms where $j = 2$, via the convergence of $(\tilde{\mathbf{X}} \boldsymbol{\beta}_\perp)^\top \hat{\mathbf{R}}_0 H_2(\tilde{\mathbf{X}} \tilde{\boldsymbol{\beta}})$ to zero.
- For $j = 3, 4$, since $H_2(\tilde{\mathbf{X}} \tilde{\boldsymbol{\beta}})$ is independent of $\hat{\mathbf{R}}_0 [\hat{\mathbf{F}}_0 \mathbf{W}_0 \tilde{\boldsymbol{\beta}} \quad \hat{\mathbf{F}}_0 \mathbf{W}_0 \boldsymbol{\beta}_\perp]$, and has zero-mean i.i.d. entries, it also follows that these entries converge to zero in probability.

Finally we study $H_2(\tilde{\boldsymbol{\theta}}_*)^\top \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ 2}$, by analyzing the terms in equation 94 for $p = 2$.

For $H_2(\tilde{\boldsymbol{\theta}}_*)^\top \hat{\mathbf{R}}_0 H_2(\tilde{\mathbf{X}} \tilde{\boldsymbol{\beta}})$, since $H_2(\tilde{\boldsymbol{\theta}}_*)$, $H_2(\tilde{\mathbf{X}} \tilde{\boldsymbol{\beta}})$ are independent of $\hat{\mathbf{R}}_0$, it follows from Lemma K.3, as in the analysis of term (1, 2) in Section K.2, that $H_2(\tilde{\boldsymbol{\theta}}_*)^\top \hat{\mathbf{R}}_0 H_2(\tilde{\mathbf{X}} \tilde{\boldsymbol{\beta}}) - \mathbb{E} \hat{\mathbf{R}}_0 \cdot \mathbb{E} H_2(\tilde{\boldsymbol{\theta}}_*)^\top H_2(\tilde{\mathbf{X}} \tilde{\boldsymbol{\beta}}) \rightarrow_P 0$. Now notice that $\hat{\mathbf{F}}_0$ is left-orthogonally invariant in distribution, and thus $\hat{\mathbf{R}}_0 =_d \mathbf{O} \hat{\mathbf{R}}_0 \mathbf{O}^\top$, where \mathbf{O} is uniformly distributed over the Haar measure of n -dimensional orthogonal matrices, independently of all other randomness. Hence, $\mathbb{E} \hat{\mathbf{R}}_0 = \mathbb{E} \text{tr} \hat{\mathbf{R}}_0 \mathbf{I}_n / n$. Moreover, from the Woodbury formula in equation 25,

$$|\text{tr} \bar{\mathbf{R}}_0 - \text{tr} \hat{\mathbf{R}}_0| \leq |\text{tr} \hat{\mathbf{R}}_0 \mathbf{V} (\mathbf{D}^{-1} + \mathbf{V}^\top \hat{\mathbf{R}}_0 \mathbf{V})^{-1} \mathbf{V}^\top \hat{\mathbf{R}}_0| \leq |\text{tr} (\mathbf{D}^{-1} + \mathbf{V}^\top \hat{\mathbf{R}}_0 \mathbf{V})^{-1} \mathbf{V}^\top \mathbf{V}| \cdot \|\hat{\mathbf{R}}_0\|_{\text{op}}^2.$$

From our previous analysis and as the entries of $\mathbf{V}^\top \mathbf{V}$ are $O_{\mathbb{P}}(n)$, it follows that the first term is $O_{\mathbb{P}}(n)$; whereas $\|\hat{\mathbf{R}}_0\|_{\text{op}}^2 = O(1/n^2)$. Hence, $|\text{tr} \bar{\mathbf{R}}_0 - \text{tr} \hat{\mathbf{R}}_0| \rightarrow_P 0$, and thus by the bounded convergence theorem $|\mathbb{E} \text{tr} \bar{\mathbf{R}}_0 - \mathbb{E} \text{tr} \hat{\mathbf{R}}_0| \rightarrow_P 0$. Moreover, we have already argued in the proof of Lemma K.4 that $\mathbb{E} \text{tr} \bar{\mathbf{R}}_0 \rightarrow \psi m_1 / \phi$.

Further, by Lemmas C.1 and K.1,

$$\begin{aligned} \mathbb{E} H_2(\tilde{\boldsymbol{\theta}}_*)^\top H_2(\tilde{\mathbf{X}} \tilde{\boldsymbol{\beta}}) &= n \cdot \mathbb{E} H_2(\tilde{\mathbf{x}}_1^\top \boldsymbol{\beta}_*) H_2(\tilde{\mathbf{x}}_1^\top \tilde{\boldsymbol{\beta}}) \\ &= 2n \mathbb{E} (\boldsymbol{\beta}_*^\top \tilde{\boldsymbol{\beta}})^2 = 2n \mathbb{E} \frac{(\boldsymbol{\beta}_*^\top \boldsymbol{\beta})^2}{\|\boldsymbol{\beta}\|^2} = 2n \frac{c_{*,1}^2}{\phi(c_*^2 + \sigma_\varepsilon^2) + c_{*,1}^2} + o_{\mathbb{P}}(1). \end{aligned}$$

This shows that

$$H_2(\tilde{\boldsymbol{\theta}}_*)^\top \hat{\mathbf{R}}_0 H_2(\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}}) \rightarrow_P 2 \frac{\psi m_1}{\phi} \frac{c_{*,1}^2}{\phi(c_{*1}^2 + \sigma_\varepsilon^2) + c_{*,1}^2}.$$

Next, we consider $H_2(\tilde{\boldsymbol{\theta}}_*)^\top \hat{\mathbf{R}}_0 \mathbf{V}$ with $\mathbf{V} = [\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}} \quad \tilde{\mathbf{X}}\boldsymbol{\beta}_\perp \quad \hat{\mathbf{F}}_0 \mathbf{W}_0 \tilde{\boldsymbol{\beta}} \quad \hat{\mathbf{F}}_0 \mathbf{W}_0 \boldsymbol{\beta}_\perp]$. For the first two entries of the vector $H_2(\tilde{\boldsymbol{\theta}}_*)^\top \hat{\mathbf{R}}_0 \mathbf{V}$, an analysis very similar to the one above for $H_2(\tilde{\boldsymbol{\theta}}_*)^\top \hat{\mathbf{R}}_0 H_2(\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}})$ shows that they converge to zero in probability. For the last two entries, since $H_2(\tilde{\boldsymbol{\theta}}_*)$ is independent of $\hat{\mathbf{R}}_0[\hat{\mathbf{F}}_0 \mathbf{W}_0 \tilde{\boldsymbol{\beta}} \quad \hat{\mathbf{F}}_0 \mathbf{W}_0 \boldsymbol{\beta}_\perp]$, and has zero-mean i.i.d. entries, it also follows that these entries converge to zero in probability. Moreover, the limiting entries of $(\mathbf{D}^{-1} + \mathbf{V}^\top \hat{\mathbf{R}}_0 \mathbf{V})^{-1}$ have been shown to be bounded in our above analysis. Hence, the second term converges to zero in probability.

Now, note that $\tilde{\boldsymbol{\beta}} = \boldsymbol{\beta}/\|\boldsymbol{\beta}\|_2$. From Lemma K.1, $\|\boldsymbol{\beta}\|^2 \rightarrow_P \phi(c_{*1}^2 + \sigma_\varepsilon^2) + c_{*,1}^2$. Hence,

$$H_2(\tilde{\boldsymbol{\theta}}_*)^\top \hat{\mathbf{R}}_0 H_2(\tilde{\mathbf{X}}\boldsymbol{\beta}) = 2 \frac{\psi m_1}{\phi} \frac{c_{*,1}^2 \|\boldsymbol{\beta}\|_2^2}{\phi(c_{*1}^2 + \sigma_\varepsilon^2) + c_{*,1}^2} + o_P(1) \rightarrow_P \frac{2c_{*,1}^2 \psi m_1}{\phi},$$

which concludes the proof.

O.9 Proof of Lemma K.8

As in the proof of Lemma K.6, we define $\hat{\mathbf{X}} = \mathbf{X} - \tilde{\boldsymbol{\theta}}\boldsymbol{\beta}^\top$. By construction, we have $\hat{\mathbf{X}} \perp \tilde{\boldsymbol{\theta}}$. As in the proof of Lemma K.5, based on the Gaussian equivalence from Appendix J, we can replace \mathbf{F}_0 with $c_1 \mathbf{X} \mathbf{W}_0^\top + c_{>1} \mathbf{Z}$ in our computations without changing the limiting result, where $\mathbf{Z} \in \mathbb{R}^{n \times d}$ is an independent random matrix with $\mathcal{N}(0, 1)$ entries. Thus, from now on, we denote $\mathbf{F}_0 = c_1 \mathbf{X} \mathbf{W}_0^\top + c_{>1} \mathbf{Z}$. We define $\hat{\mathbf{F}}_0$ as in equation 90; thus, $\hat{\mathbf{F}}_0 = \mathbf{F}_0 - c_1 \tilde{\boldsymbol{\theta}}(\mathbf{W}_0 \boldsymbol{\beta})^\top$. As a consequence, we can write $\mathbf{F}_0 \mathbf{F}_0^\top = \hat{\mathbf{F}}_0 \hat{\mathbf{F}}_0^\top + \mathbf{V} \mathbf{D} \mathbf{V}^\top$, where $\mathbf{V} = [\hat{\mathbf{F}}_0 \mathbf{W}_0 \boldsymbol{\beta} \quad \tilde{\boldsymbol{\theta}}] \in \mathbb{R}^{n \times 2}$ and

$$\mathbf{D} = \begin{bmatrix} 0 & c_1 \\ c_1 & c_1^2 \|\mathbf{W}_0 \boldsymbol{\beta}\|_2^2 \end{bmatrix}.$$

Using the Woodbury formula, we find that equation 91 still holds. Now, we can write

$$\tilde{\boldsymbol{\theta}}^{\circ 2\top} \bar{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ 2} = \tilde{\boldsymbol{\theta}}^{\circ 2\top} \hat{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ 2} - \tilde{\boldsymbol{\theta}}^{\circ 2\top} \hat{\mathbf{R}}_0 \mathbf{V} (\mathbf{D}^{-1} + \mathbf{V}^\top \hat{\mathbf{R}}_0 \mathbf{V})^{-1} \mathbf{V}^\top \hat{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ 2}. \quad (95)$$

We can analyze each term in the above sum separately.

By Lemma K.3, $\tilde{\boldsymbol{\theta}}^{\circ 2\top} \hat{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ 2} - \mathbb{E} \tilde{\boldsymbol{\theta}}^{\circ 2\top} \hat{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ 2} \rightarrow_P 0$. Further, conditional on $\boldsymbol{\beta}$, $\mathbb{E} \tilde{\boldsymbol{\theta}}^{\circ 2\top} \hat{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ 2} = 3\|\boldsymbol{\beta}\|_2^4 \mathbb{E} \text{tr} \hat{\mathbf{R}}_0$; and as in the proof of Lemma K.7, $\mathbb{E} \text{tr} \hat{\mathbf{R}}_0 - \mathbb{E} \text{tr} \bar{\mathbf{R}}_0 \rightarrow 0$. Moreover, we have already argued in the proof of Lemma K.4 that $\mathbb{E} \text{tr} \bar{\mathbf{R}}_0 \rightarrow \psi m_1/\phi$. In addition, from Lemma K.1, $\|\boldsymbol{\beta}\|^2 \rightarrow_P \phi(c_{*1}^2 + \sigma_\varepsilon^2) + c_{*,1}^2$. Hence,

$$\tilde{\boldsymbol{\theta}}^{\circ 2\top} \hat{\mathbf{R}}_0 \tilde{\boldsymbol{\theta}}^{\circ 2} \rightarrow_P 3\psi m_1 [\phi(c_{*1}^2 + \sigma_\varepsilon^2) + c_{*,1}^2]^2 / \phi.$$

To analyze the second term in equation 95, we first study $\tilde{\boldsymbol{\theta}}^{\circ 2\top} \hat{\mathbf{R}}_0 \hat{\mathbf{F}}_0 \mathbf{W}_0 \boldsymbol{\beta}$. By an argument similar to the ones above, we can show that it concentrates around $\mathbf{1}_n^\top \hat{\mathbf{R}}_0 \hat{\mathbf{F}}_0 \mathbf{W}_0 \boldsymbol{\beta} = \mathbf{1}_n^\top \hat{\mathbf{F}}_0 \hat{\mathbf{R}}_0 \mathbf{W}_0 \boldsymbol{\beta}$. Since $\hat{\mathbf{F}}_0$ is left-orthogonally invariant, $\mathbf{1}_n^\top \hat{\mathbf{F}}_0 \hat{\mathbf{R}}_0 \mathbf{W}_0 \boldsymbol{\beta} = \mathbf{1}_n^\top \mathbf{O} \hat{\mathbf{F}}_0 \hat{\mathbf{R}}_0 \mathbf{W}_0 \boldsymbol{\beta}$, where \mathbf{O} is uniformly distributed over the Haar measure of n -dimensional orthogonal matrices, independently of all other randomness. Then, it follows as in the analysis of term (1,2) from Section K.2 that $\mathbf{1}_n^\top \mathbf{O} \hat{\mathbf{F}}_0 \hat{\mathbf{R}}_0 \mathbf{W}_0 \boldsymbol{\beta} \rightarrow_P 0$; and hence $\tilde{\boldsymbol{\theta}}^{\circ 2\top} \hat{\mathbf{R}}_0 \hat{\mathbf{F}}_0 \mathbf{W}_0 \boldsymbol{\beta} \rightarrow_P 0$.

Moreover, the limiting entries of $(\mathbf{D}^{-1} + \mathbf{V}^\top \hat{\mathbf{R}}_0 \mathbf{V})^{-1}$ can be shown to be bounded by a simple orderwise analysis. Hence, the second term in equation 95 is $o_P(1)$.

O.10 Proof of Lemma L.2

Denoting $\tilde{\beta} = \beta / \|\beta\|_2$, we have

$$\begin{aligned}
(\tilde{\mathbf{X}}\beta)^{\circ i \top} \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}}\beta)^{\circ j} &= \|\beta\|_2^{i+j} (\tilde{\mathbf{X}}\tilde{\beta})^{\circ i \top} \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}}\tilde{\beta})^{\circ j} = \|\beta\|_2^{i+j} \sum_{k_1=0}^i \sum_{k_2=0}^j \xi_{i,k_1} \xi_{j,k_2} H_{k_1} (\tilde{\mathbf{X}}\tilde{\beta})^\top \bar{\mathbf{R}}_0 H_{k_2} (\tilde{\mathbf{X}}\tilde{\beta}) \\
&= \|\beta\|_2^{i+j} \sum_{k=0}^{\min(i,j)} \xi_{j,k} \xi_{i,k} H_k (\tilde{\mathbf{X}}\tilde{\beta})^\top \bar{\mathbf{R}}_0 H_k (\tilde{\mathbf{X}}\tilde{\beta}) + o_{\mathbb{P}}(1) \\
&= \|\beta\|_2^{i+j} \left[\xi_{i,1} \xi_{j,1} (\tilde{\mathbf{X}}\tilde{\beta})^\top \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}}\tilde{\beta}) + \sum_{k=0, k \neq 1}^{\min(i,j)} \xi_{i,k} \xi_{j,k} H_k (\tilde{\mathbf{X}}\tilde{\beta})^\top \bar{\mathbf{R}}_0 H_k (\tilde{\mathbf{X}}\tilde{\beta}) \right] + o_{\mathbb{P}}(1).
\end{aligned}$$

The third line follows from Lemma K.5. Now, we claim that for any $p \in \{0, 2, 3, \dots\}$, we have $H_p(\tilde{\mathbf{X}}\beta / \|\beta\|_2)^\top \bar{\mathbf{R}}_0 H_p(\tilde{\mathbf{X}}\beta / \|\beta\|_2) \rightarrow_P p! \psi m_1 / \phi$. Using this claim, the facts that $\|\beta\|_2^2 \rightarrow_P c_{\star,1}^2 + \phi(c_{\star}^2 + \sigma_{\varepsilon}^2)$, and $\text{tr}(\tilde{\mathbf{X}}^\top (\mathbf{F}_0 \mathbf{F}_0^\top + \lambda n \mathbf{I}_n)^{-1} \tilde{\mathbf{X}}) / d \rightarrow_P \psi m_2 / \phi$, we can conclude

$$(\tilde{\mathbf{X}}\beta)^{\circ i \top} \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}}\beta)^{\circ j} \rightarrow_P (c_{\star,1}^2 + \phi(c_{\star}^2 + \sigma_{\varepsilon}^2))^{(i+j)/2} \left[\xi_{i,1} \xi_{j,1} \frac{\psi m_2}{\phi} + \frac{\psi m_1}{\phi} \sum_{k=0, k \neq 1}^{\min(i,j)} k! \xi_{i,k} \xi_{j,k} \right].$$

Now, it remains to prove the claim that for any $p \in \{0, 2, 3, \dots\}$, we have

$$H_p(\tilde{\mathbf{X}}\beta / \|\beta\|_2)^\top \bar{\mathbf{R}}_0 H_p(\tilde{\mathbf{X}}\beta / \|\beta\|_2) \rightarrow_P p! \psi m_1 / \phi.$$

As in the proof of Lemma K.8, we define $\hat{\mathbf{X}} = \tilde{\mathbf{X}} - \tilde{\mathbf{X}}\tilde{\beta}\tilde{\beta}^\top$. By construction, we have $\hat{\mathbf{X}} \perp \tilde{\mathbf{X}}\tilde{\beta}$. As in the proof of Lemma K.5, based on the Gaussian equivalence from Appendix J, we can replace \mathbf{F}_0 with $c_1 \tilde{\mathbf{X}} \mathbf{W}_0^\top + c_{>1} \mathbf{Z}$ in our computations without changing the limiting result, where $\mathbf{Z} \in \mathbb{R}^{n \times d}$ is an independent random matrix with $\mathcal{N}(0, 1)$ entries. Thus, from now on, we denote $\mathbf{F}_0 = c_1 \tilde{\mathbf{X}} \mathbf{W}_0^\top + c_{>1} \mathbf{Z}$. We define $\hat{\mathbf{F}}_0$ as in equation 90; thus, $\hat{\mathbf{F}}_0 = \mathbf{F}_0 - c_1 \tilde{\mathbf{X}}\tilde{\beta}(\mathbf{W}_0\tilde{\beta})^\top$. As a consequence, we can write $\mathbf{F}_0 \mathbf{F}_0^\top = \hat{\mathbf{F}}_0 \hat{\mathbf{F}}_0^\top + \mathbf{V} \mathbf{D} \mathbf{V}^\top$, where $\mathbf{V} = [\hat{\mathbf{F}}_0 \mathbf{W}_0\tilde{\beta} \quad \tilde{\mathbf{X}}\tilde{\beta}] \in \mathbb{R}^{n \times 2}$ and

$$\mathbf{D} = \begin{bmatrix} 0 & c_1 \\ c_1 & c_1^2 \|\mathbf{W}_0\tilde{\beta}\|_2^2 \end{bmatrix}.$$

Using the Woodbury formula, we find that equation 91 still holds. Now, we can write

$$\begin{aligned}
H_p(\tilde{\mathbf{X}}\tilde{\beta})^\top \bar{\mathbf{R}}_0 H_p(\tilde{\mathbf{X}}\tilde{\beta}) &= H_p(\tilde{\mathbf{X}}\tilde{\beta})^\top \hat{\mathbf{R}}_0 H_p(\tilde{\mathbf{X}}\tilde{\beta}) \\
&\quad - H_p(\tilde{\mathbf{X}}\tilde{\beta})^\top \hat{\mathbf{R}}_0 \mathbf{V} (\mathbf{D}^{-1} + \mathbf{V}^\top \hat{\mathbf{R}}_0 \mathbf{V})^{-1} \mathbf{V}^\top \hat{\mathbf{R}}_0 H_p(\tilde{\mathbf{X}}\tilde{\beta}).
\end{aligned} \tag{96}$$

We can analyze each term in the above sum separately.

By Lemma K.3, $H_p(\tilde{\mathbf{X}}\tilde{\beta})^\top \hat{\mathbf{R}}_0 H_p(\tilde{\mathbf{X}}\tilde{\beta}) - \mathbb{E} H_p(\tilde{\mathbf{X}}\tilde{\beta})^\top \hat{\mathbf{R}}_0 H_p(\tilde{\mathbf{X}}\tilde{\beta}) \rightarrow_P 0$. Further, conditional on $\tilde{\beta}$, and using C.1, we have

$$\mathbb{E} H_p(\tilde{\mathbf{X}}\tilde{\beta})^\top \hat{\mathbf{R}}_0 H_p(\tilde{\mathbf{X}}\tilde{\beta}) = \mathbb{E} \text{tr} [\hat{\mathbf{R}}_0 H_p(\tilde{\mathbf{X}}\tilde{\beta}) H_p(\tilde{\mathbf{X}}\tilde{\beta})^\top] = p! \mathbb{E} \text{tr} [\hat{\mathbf{R}}_0],$$

and as in the proof of Lemma K.7, $\mathbb{E} \text{tr} \hat{\mathbf{R}}_0 - \mathbb{E} \text{tr} \bar{\mathbf{R}}_0 \rightarrow 0$. Moreover, we have already argued in the proof of Lemma K.4 that $\mathbb{E} \text{tr} \bar{\mathbf{R}}_0 \rightarrow \psi m_1 / \phi$. Hence, $H_p(\tilde{\mathbf{X}}\tilde{\beta})^\top \hat{\mathbf{R}}_0 H_p(\tilde{\mathbf{X}}\tilde{\beta}) \rightarrow_P p! \psi m_1 / \phi$. To

analyze the second term in equation 96, we first study $H_p(\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}})^\top \hat{\mathbf{R}}_0 \hat{\mathbf{F}}_0 \mathbf{W}_0 \boldsymbol{\beta}$. Conditional on $\tilde{\boldsymbol{\beta}}$, $H_p(\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}})$ is a vector with independent mean-zero, bounded variance entries, independent of the vector $\hat{\mathbf{R}}_0 \hat{\mathbf{F}}_0 \mathbf{W}_0 \tilde{\boldsymbol{\beta}}$ that has norm $O(1/\sqrt{n})$. Hence, we conclude that this term goes to zero. Next, note that $H_p(\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}})^\top \hat{\mathbf{R}}_0(\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}}) \rightarrow_P 0$ using Lemma K.3 and Lemma C.1. Moreover, the limiting entries of $(\mathbf{D}^{-1} + \mathbf{V}^\top \hat{\mathbf{R}}_0 \mathbf{V})^{-1}$ can be shown to be bounded by a simple orderwise analysis. Hence, the second term in equation 96 is $o_{\mathbb{P}}(1)$. This concludes the proof.

O.11 Proof of Lemma L.3

We define $\boldsymbol{\beta}_\perp = \frac{\boldsymbol{\beta}_* - \langle \boldsymbol{\beta}_*, \tilde{\boldsymbol{\beta}} \rangle \tilde{\boldsymbol{\beta}}}{\|\boldsymbol{\beta}_* - \langle \boldsymbol{\beta}_*, \tilde{\boldsymbol{\beta}} \rangle \tilde{\boldsymbol{\beta}}\|_2}$, and set $\hat{\mathbf{X}} = \tilde{\mathbf{X}} - \tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}}\tilde{\boldsymbol{\beta}}^\top - \tilde{\mathbf{X}}\boldsymbol{\beta}_\perp\boldsymbol{\beta}_\perp^\top$. By construction, we have $\hat{\mathbf{X}} \perp \tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}}, \tilde{\boldsymbol{\theta}}_*$. Based on the Gaussian equivalence from Appendix J, we can again replace \mathbf{F}_0 with $\mathbf{F}_0 = c_1 \tilde{\mathbf{X}}\mathbf{W}_0^\top + c_{>1} \mathbf{Z}$, where $\mathbf{Z} \in \mathbb{R}^{n \times d}$ is an independent random matrix with $\mathcal{N}(0, 1)$ entries. Again, we define $\hat{\mathbf{F}}_0$ as in equation 90. Thus, $\hat{\mathbf{F}}_0 = \mathbf{F}_0 - c_1 \tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}}(\mathbf{W}_0\tilde{\boldsymbol{\beta}})^\top - c_1 \tilde{\mathbf{X}}\boldsymbol{\beta}_\perp(\mathbf{W}_0\boldsymbol{\beta}_\perp)^\top$. As a consequence, we also have $\mathbf{F}_0\mathbf{F}_0^\top = \hat{\mathbf{F}}_0\hat{\mathbf{F}}_0^\top + \mathbf{V}\mathbf{D}\mathbf{V}^\top$, where $\mathbf{V} = [\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}} \quad \tilde{\mathbf{X}}\boldsymbol{\beta}_\perp \quad \hat{\mathbf{F}}_0\mathbf{W}_0\tilde{\boldsymbol{\beta}} \quad \hat{\mathbf{F}}_0\mathbf{W}_0\boldsymbol{\beta}_\perp] \in \mathbb{R}^{n \times 4}$ and

$$\mathbf{D} = \begin{bmatrix} c_1^2 \langle \mathbf{W}_0\tilde{\boldsymbol{\beta}}, \mathbf{W}_0\tilde{\boldsymbol{\beta}} \rangle & c_1^2 \langle \mathbf{W}_0\tilde{\boldsymbol{\beta}}, \mathbf{W}_0\boldsymbol{\beta}_\perp \rangle & c_1 & 0 \\ c_1^2 \langle \mathbf{W}_0\tilde{\boldsymbol{\beta}}, \mathbf{W}_0\boldsymbol{\beta}_\perp \rangle & c_1^2 \langle \mathbf{W}_0\boldsymbol{\beta}_\perp, \mathbf{W}_0\boldsymbol{\beta}_\perp \rangle & 0 & c_1 \\ c_1 & 0 & 0 & 0 \\ 0 & c_1 & 0 & 0 \end{bmatrix}.$$

Using the Woodbury formula, we find that equation 91 still holds. We can write

$$\begin{aligned} H_p(\tilde{\boldsymbol{\theta}}_*)^\top \bar{\mathbf{R}}_0 H_q(\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}}) \\ = H_p(\tilde{\boldsymbol{\theta}}_*)^\top \hat{\mathbf{R}}_0 H_q(\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}}) - H_p(\tilde{\boldsymbol{\theta}}_*)^\top \hat{\mathbf{R}}_0 \mathbf{V} (\mathbf{D}^{-1} + \mathbf{V}^\top \hat{\mathbf{R}}_0 \mathbf{V})^{-1} \mathbf{V}^\top \hat{\mathbf{R}}_0 H_q(\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}}). \end{aligned} \quad (97)$$

$p \neq q$ case: The first term converges to zero for any $p \neq q$, analogously to the argument in Section K.2.1 for the terms (1,2) and (2,4). In particular, for $p = 0$, we can use orthogonal invariance as in the analysis of the term (2,4). To prove that the second term will also converge to zero when $p \neq q$, we first observe that the elements of $\mathbf{K} = (\mathbf{D}^{-1} + \mathbf{V}^\top \hat{\mathbf{R}}_0 \mathbf{V})^{-1}$ are all $O(1)$. The second term will involve quantities of the form $[\mathbf{K}]_{i,j} H_p(\tilde{\boldsymbol{\theta}}_*)^\top \hat{\mathbf{R}}_0 \mathbf{v}_i \mathbf{v}_j^\top \hat{\mathbf{R}}_0 H_q(\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}})$, where \mathbf{v}_i , for $i \in \{1, 2, 3, 4\}$, is the i -th column of the matrix $\mathbf{V} = [\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}} \quad \tilde{\mathbf{X}}\boldsymbol{\beta}_\perp \quad \hat{\mathbf{F}}_0\mathbf{W}_0\tilde{\boldsymbol{\beta}} \quad \hat{\mathbf{F}}_0\mathbf{W}_0\boldsymbol{\beta}_\perp]$. We can argue that all these terms converge to zero, as follows. Because $p \neq q$, without loss of generality, assume that $q \neq 1$.

- The terms where $j = 1$ converge to zero because $(\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}})^\top \hat{\mathbf{R}}_0 H_q(\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}})$ converges to zero using the concentration argument from Lemma K.3 and the orthogonality of Hermite polynomials from Lemma C.1. The same argument applies to the terms where $j = 2$, via the convergence of $(\tilde{\mathbf{X}}\boldsymbol{\beta}_\perp)^\top \hat{\mathbf{R}}_0 H_q(\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}})$ to zero.
- For $j = 3, 4$, and for $q > 0$, since $H_q(\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}})$ is independent of $\hat{\mathbf{R}}_0[\hat{\mathbf{F}}_0\mathbf{W}_0\tilde{\boldsymbol{\beta}} \quad \hat{\mathbf{F}}_0\mathbf{W}_0\boldsymbol{\beta}_\perp]$, and has zero-mean i.i.d. entries, it also follows that these entries converge to zero in probability. For $q = 0$, we can again use orthogonal invariance as in the analysis of the term (2,4).

The case when $p = q \neq 1$: Finally we study $H_p(\tilde{\boldsymbol{\theta}}_*)^\top \bar{\mathbf{R}}_0 H_p(\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}})$, by analyzing the terms in equation 94.

For $H_p(\tilde{\boldsymbol{\theta}}_*)^\top \hat{\mathbf{R}}_0 H_p(\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}})$, since $H_p(\tilde{\boldsymbol{\theta}}_*)$, $H_p(\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}})$ are independent of $\hat{\mathbf{R}}_0$, it follows from Lemma K.3, as in the analysis of term (1, 2) in the Section K.2, that

$$H_p(\tilde{\boldsymbol{\theta}}_*)^\top \hat{\mathbf{R}}_0 H_p(\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}}) - \mathbb{E}\hat{\mathbf{R}}_0 \cdot \mathbb{E}H_p(\tilde{\boldsymbol{\theta}}_*)^\top H_p(\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}}) \rightarrow_P 0.$$

Now notice that $\hat{\mathbf{F}}_0$ is left-orthogonally invariant in distribution, and thus $\hat{\mathbf{R}}_0 =_d \mathbf{O}\hat{\mathbf{R}}_0\mathbf{O}^\top$, where \mathbf{O} is uniformly distributed over the Haar measure of n -dimensional orthogonal matrices, independently of all other randomness. Hence, $\mathbb{E}\hat{\mathbf{R}}_0 = \mathbb{E}\text{tr } \hat{\mathbf{R}}_0 \mathbf{I}_n/n$. Also, similar to the proof of Lemma K.7, we have $|\text{tr } \bar{\mathbf{R}}_0 - \text{tr } \hat{\mathbf{R}}_0| = o_{\mathbb{P}}(1)$. Moreover, we have already argued in the proof of Lemma K.4 that $\mathbb{E}\text{tr } \bar{\mathbf{R}}_0 \rightarrow \psi m_1/\phi$. Further, by Lemmas C.1 and K.1,

$$H_p(\tilde{\boldsymbol{\theta}}_*)^\top \hat{\mathbf{R}}_0 H_p(\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}}) \rightarrow_P p! \frac{\psi m_1}{\phi} \left(\frac{c_{*,1}}{\sqrt{\phi(c_*^2 + \sigma_\varepsilon^2) + c_{*,1}^2}} \right)^p.$$

Next, we consider $H_2(\tilde{\boldsymbol{\theta}}_*)^\top \hat{\mathbf{R}}_0 \mathbf{V}$ with $\mathbf{V} = [\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}} \quad \tilde{\mathbf{X}}\boldsymbol{\beta}_\perp \quad \hat{\mathbf{F}}_0 \mathbf{W}_0 \tilde{\boldsymbol{\beta}} \quad \hat{\mathbf{F}}_0 \mathbf{W}_0 \boldsymbol{\beta}_\perp]$. For the first two entries of the vector $H_p(\tilde{\boldsymbol{\theta}}_*)^\top \hat{\mathbf{R}}_0 \mathbf{V}$, an analysis very similar to the one above for $H_p(\tilde{\boldsymbol{\theta}}_*)^\top \hat{\mathbf{R}}_0 H_p(\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}})$ shows that they converge to zero in probability. For the last two entries, since $H_p(\tilde{\boldsymbol{\theta}}_*)$ is independent of $\hat{\mathbf{R}}_0[\hat{\mathbf{F}}_0 \mathbf{W}_0 \tilde{\boldsymbol{\beta}} \quad \hat{\mathbf{F}}_0 \mathbf{W}_0 \boldsymbol{\beta}_\perp]$, and has zero-mean i.i.d. entries, it also follows that these entries converge to zero in probability. Moreover, the limiting entries of $(\mathbf{D}^{-1} + \mathbf{V}^\top \hat{\mathbf{R}}_0 \mathbf{V})^{-1}$ have been shown to be bounded in our above analysis. Hence, the second term converges to zero in probability.

The case when $p = q = 1$: In this case, we have

$$(\tilde{\mathbf{X}}\boldsymbol{\beta}_*)^\top \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}}) = \frac{(\tilde{\mathbf{X}}\boldsymbol{\beta}_*)^\top \bar{\mathbf{R}}_0 (\tilde{\mathbf{X}}\boldsymbol{\beta})}{\|\boldsymbol{\beta}\|_2} = \frac{c_{*,1} \frac{\psi m_2}{\phi}}{\sqrt{\phi(c_*^2 + \sigma_\varepsilon^2) + c_{*,1}^2}} + o_{\mathbb{P}}(1),$$

using Lemma K.1 and by arguments similar to the ones in the proof of Lemma K.4.

Putting everything together concludes the proof.