
Generalized equivalences between subsampling and ridge regularization

Pratik Patil

Department of Statistics
University of California
Berkeley, CA 94720, USA
pratikpatil@berkeley.edu

Jin-Hong Du

Department of Statistics and Data Science
& Machine Learning Department
Carnegie Mellon University
Pittsburgh, PA 15213, USA
jinhongd@andrew.cmu.edu

Abstract

We establish precise structural and risk equivalences between subsampling and ridge regularization for ensemble ridge estimators. Specifically, we prove that linear and quadratic functionals of subsample ridge estimators, when fitted with different ridge regularization levels λ and subsample aspect ratios ψ , are asymptotically equivalent along specific paths in the (λ, ψ) -plane (where ψ is the ratio of the feature dimension to the subsample size). Our results only require bounded moment assumptions on feature and response distributions and allow for arbitrary joint distributions. Furthermore, we provide a data-dependent method to determine the equivalent paths of (λ, ψ) . An indirect implication of our equivalences is that optimally tuned ridge regression exhibits a monotonic prediction risk in the data aspect ratio. This resolves a recent open problem raised by Nakkiran et al. [1] for general data distributions under proportional asymptotics, assuming a mild regularity condition that maintains regression hardness through linearized signal-to-noise ratios.

1 Introduction

Ensemble methods, such as bagging [2, 3], are powerful tools that combine weak predictors to improve predictive stability and accuracy. This paper focuses on sampling-based ensembles, which exhibit an implicit regularization effect [4–6]. Specifically, we investigate subsample ridge ensembles, where the ridge predictors are fitted on independently subsampled datasets [7–9]. Recent work has demonstrated that a full ensemble (that is fitted on all possible subsampled datasets) of ridgeless [10] predictors achieves the same squared prediction risk as a ridge predictor fitted on full data [11–13].

To be precise, let ϕ be the limiting dataset aspect ratio p/n , where p is the feature dimension, and n is the number of observations. For a given ϕ , the limiting prediction risk of subsample ridge ensembles is parameterized by (λ, ψ) , where $\lambda \geq 0$ is the ridge regularization parameter and $\psi \geq \phi$ is the limiting subsample aspect ratio p/k , with k being the subsample size [11, 12]. Under isotropic features and a well-specified linear model, [11, 12] show that the squared prediction risk at $(0, \psi^*)$ (the optimal ridgeless ensemble) is the same as the risk at (λ^*, ϕ) (the optimal ridge), where ψ^* and λ^* are the optimal subsample aspect ratio and ridge penalty, respectively. Furthermore, this equivalence of prediction risk between subsampling and ridge regularization is extended in [13] to anisotropic linear models. As an application, [13] also demonstrates how generalized cross-validation for ridge regression can be naturally transferred to subsample ridge ensembles. In essence, these results suggest that subsampling a smaller number of observations has the same effect as adding an appropriately larger level of ridge penalty. These findings prompt two important open questions:

- (a) **The extent of equivalences.** The previous works all focus on equivalences of the squared prediction risk when the test distribution matches the train distribution. In real-world scenarios,

Table 1: Comparison with related work. The marker “ \checkmark° ” indicates a partial equivalence result connecting the *optimal* prediction risk of the ridge predictor and the full ridgeless ensemble predictor.

	Type of equivalence results			Type of data assumptions		
	pred. risk	gen. risk	estimator	response	feature	lim. spectrum
LeJeune et al. [11]	\checkmark°			linear	isotropic Gaussian	exists
Patil et al. [12]	\checkmark°			linear	isotropic RMT	exists
Du et al. [13]	\checkmark			linear	anisotropic RMT	exists
This work	\checkmark	\checkmark	\checkmark	arbitrary	anisotropic RMT	need not exist

however, there are often covariate and label shifts, making it crucial to consider prediction risks under such shifts. In addition, other risk measures, such as training error, estimation risk, coefficient errors, and more, are also of interest in various inferential tasks. A natural question is then whether similar equivalences hold for such “generalized” risk measures. At a basic level, the question boils down to whether any equivalences exist at the “estimator level”. Answering this question would establish a connection between the ensemble and the ridge estimators, facilitating the exchange of various downstream inferential statements between the two sets of estimators.

- (b) **The role of linear model.** All previous works assume a well-specified linear model between the responses and the features, which rarely holds in practical applications. A natural question is whether the equivalences hold for *arbitrary* joint distributions of the response and the features or whether they are merely an artifact of the linear model. Addressing this question broadens the applicability of such equivalences beyond simplistic models to real-world scenarios, where the relationship between the response and the features is typically intricate and unknown.

We provide answers to questions raised in both directions. We demonstrate that the equivalences hold for the generalized squared risks in the full ensemble. Further, these equivalences fundamentally occur at the estimator level for any arbitrary ensemble size. Importantly, these results are not limited to linear models of the data-generating process. Below we provide a summary of our main results.

1.1 Summary of contributions

- (1) **Risk equivalences.** We establish asymptotic equivalences of the full-ensemble ridge estimators at different ridge penalties λ and subsample ratios ψ along specific paths in the (λ, ψ) -plane for a variety of generalized risk functionals. This class of functionals includes commonly used risk measures (see Table 2), and our results hold for both independent and dependent coefficient matrices that define these risk functionals (see Theorem 1 and Proposition 2). In addition, we demonstrate that the equivalence path remains unchanged across all the functionals examined.
- (2) **Structural equivalences.** We establish structural equivalences in the form of linear functionals of the ensemble ridge estimators, which hold for arbitrary ensemble sizes (see Theorem 3). Our proofs for both structural and risk equivalences exploit properties of certain fixed equations that arise in our analysis, enabling us to explicitly characterize paths that yield equivalent estimators in the (λ, ψ) -plane (see Equation (5)). In addition, we provide an entirely data-driven construction of this path using certain companion Stieltjes transform relations of random matrices (see Proposition 4).
- (3) **Equivalence implications.** As an implication of our equivalences, we show that the prediction risk of an optimally-tuned ridge estimator is monotonically increasing in the data aspect ratio under mild regularity conditions (see Theorem 6). This is an indirect consequence of our general equivalence results that leverages the provable monotonicity of the subsample-optimized estimator. Under proportional asymptotics, our result settles a recent open question raised by Nakkiran et al. [1, Conjecture 1] concerning the monotonicity of optimal ridge regression under anisotropic features and general data models while maintaining a regularity condition that preserves the linearized signal-to-noise ratios across regression problems.
- (4) **Generality of equivalences.** Our main results apply to arbitrary responses with bounded $4 + \mu$ moments for some $\mu > 0$, as well as features with similar bounded moments and general

covariance structures. We demonstrate the practical implications of our findings on real-world datasets (see Section 6). On the technical side, we extend the tools developed in [14] to handle the dependency between the linear and the non-linear components and obtain model-free equivalences. Furthermore, we extend our analysis to include generalized ridge regression (see Corollary 5). Through experiments, we demonstrate the universality of the equivalences for random and kernel features and conjecture the associated data-driven prediction risk equivalence paths (see Section 6).

1.2 Related work

Ensemble learning. Ensemble methods yield strong predictors by combining weak predictors [15]. The most common strategy for building ensembles is based on subsampling observations, which includes bagging [2, 3], random forests [16], neural network ensembles [17–19], among the others. The effect of sampling-based ensembles has been studied in the statistical physics literature [20–22] as well as statistical learning literature [2, 3]. Recent works have attempted to explain the success of ensemble learning by suggesting that subsampling induces a regularizing effect [4, 6]. Under proportional asymptotics, the effect of ensemble random feature models has been investigated in [23–25]. There has been growing interest in connecting the effect of ensemble learning to explicit regularization: see [26] for related experimental evidence under the umbrella of “mini-patch” learning. Towards making these connections precise, some work has been done in the context of the ridge and ridgeless ensembles: [11–13] investigate the relationship between subsampling and regularization for ridge ensembles in the overparameterized regime. We will delve into these works in detail next.

Ensembles and ridge equivalences. In the study of ridge ensembles, LeJeune et al. [11] show that the optimal ridge predictor has the same prediction risk as the ridgeless ensemble under the Gaussian isotropic design, which has been extended in [12] for RMT features (see Assumption 2 for definition). Specifically, these works establish the prediction risk equivalences for a specific pair of (λ^*, ϕ) and $(0, \psi^*)$. The results are then extended to the entire range of (λ, ψ) using deterministic asymptotic risk formulas [13], assuming a well-specified linear model, RMT features, and the existence of a fixed limiting spectral distribution of the covariance matrix. Our work significantly broadens the scope of results connecting subsampling and ridge regularization. In particular, we allow for arbitrary joint distributions of the data, anisotropic features with only bounded moments, and do not assume the convergence of the spectrum of the population covariance matrix. Furthermore, we expand the applicability of equivalences to encompass both the estimators and various generalized squared risk functionals. See Table 1 for an explicit comparison and Section 3 for additional related comparisons.

Other ridge connections. Beyond the connection connections between implicit regularization induced by subsampling and explicit ridge regularization examined in this paper, ridge regression shares ties with various other forms of regularization. For example, ridge regression has been linked to dropout regularization [27, 28], variable splitting in random forests [2], noisy training [29, 30], random sketched regression [31–33], various forms of data augmentation [34], feature augmentation [9, 35], early stopping in gradient descent and its variants [36–39], among others. These connections highlight the pervasiveness and the significance of understanding various “facets” of ridge regularization. In that direction, our work contributes to expand the equivalences between subsampling and ridge regularization.

2 Notation and preliminaries

Ensemble estimators. Let $\mathcal{D}_n = \{(\mathbf{x}_i, y_i) : i \in [n]\}$ be a dataset containing i.i.d. samples in $\mathbb{R}^p \times \mathbb{R}$. Let $\mathbf{X} \in \mathbb{R}^{n \times p}$ and $\mathbf{y} \in \mathbb{R}^n$ be the feature matrix and response vector with \mathbf{x}_i^\top and y_i in i -th rows. For an index set $I \subseteq [n]$ of size k , let $\mathcal{D}_I = \{(\mathbf{x}_i, y_i) : i \in I\}$ be the associated subsampled dataset. Let $\mathbf{L}_I \in \mathbb{R}^{n \times n}$ be a diagonal matrix such that its i -th diagonal entry is 1 if $i \in I$ and 0 otherwise. Note that the feature matrix and response vector associated with \mathcal{D}_I respectively are $\mathbf{L}_I \mathbf{X}$ and $\mathbf{L}_I \mathbf{y}$. Given a ridge penalty $\lambda > 0$, the *ridge* estimator fitted on \mathcal{D}_I consisting of k samples is given by:

$$\hat{\beta}_k^\lambda(\mathcal{D}_I) = \operatorname{argmin}_{\beta \in \mathbb{R}^p} \frac{1}{k} \sum_{i \in I} (y_i - \mathbf{x}_i^\top \beta)^2 + \lambda \|\beta\|_2^2 = \left(\frac{1}{k} \mathbf{X}^\top \mathbf{L}_I \mathbf{X} + \lambda \mathbf{I}_p \right)^{-1} \frac{\mathbf{X}^\top \mathbf{L}_I \mathbf{y}}{k}. \quad (1)$$

The *ridgeless* estimator $\hat{\beta}_k^0(\mathcal{D}_I) = (\mathbf{X}^\top \mathbf{L}_I \mathbf{X} / k)^+ \mathbf{X}^\top \mathbf{L}_I \mathbf{y} / k$ is obtained by sending $\lambda \rightarrow 0^+$, where M^+ denotes the Moore-Penrose inverse of matrix M . To define the ensemble estimator, denote

Table 2: Summary of various generalized risks and their corresponding statistical learning problems. The asymptotic equivalences between subsampling and ridge regularization are analyzed in Theorem 1 and Proposition 2 for generalized risks (3) defined through functionals $L_{\mathbf{A}, \mathbf{b}}$ for various \mathbf{A} and \mathbf{b} .

Statistical learning problem	$L_{\mathbf{A}, \mathbf{b}}(\widehat{\beta} - \beta_0)$	\mathbf{A}	\mathbf{b}	$\text{nrow}(\mathbf{A})$
vector coefficient estimation	$\widehat{\beta} - \beta_0$	\mathbf{I}_p	0	p
projected coefficient estimation	$\mathbf{a}^\top(\widehat{\beta} - \beta_0)$	\mathbf{a}^\top	0	1
training error estimation	$\mathbf{X}\widehat{\beta} - \mathbf{y}$	\mathbf{X}	$-\mathbf{f}_{\text{NL}}$	n
in-sample prediction	$\mathbf{X}(\widehat{\beta} - \beta_0)$	\mathbf{X}	0	n
out-of-sample prediction	$\mathbf{x}_0^\top \widehat{\beta} - y_0$	\mathbf{x}_0^\top	$-\varepsilon_0$	1

the set of all k distinct elements from $[n]$ by $\mathcal{I}_k = \{\{i_1, i_2, \dots, i_k\} : 1 \leq i_1 < i_2 < \dots < i_k \leq n\}$. For $\lambda \geq 0$, the M -ensemble and the full-ensemble estimators are respectively defined as follows:

$$\widehat{\beta}_{k, M}^\lambda(\mathcal{D}_n; \{I_\ell\}_{\ell=1}^M) = \frac{1}{M} \sum_{\ell \in [M]} \widehat{\beta}_k^\lambda(\mathcal{D}_{I_\ell}), \quad \text{and} \quad \widehat{\beta}_{k, \infty}^\lambda(\mathcal{D}_n) = \mathbb{E}[\widehat{\beta}_k^\lambda(\mathcal{D}_I) \mid \mathcal{D}_n]. \quad (2)$$

Here $\{I_\ell\}_{\ell=1}^M$ and I are simple random samples from \mathcal{I}_k . In (2), the full-ensemble ridge estimator $\widehat{\beta}_{k, \infty}^\lambda(\mathcal{D}_n)$ is the average of predictors fitted on all possible subsampled datasets. It is shown in Lemma A.1 of [13] that $\widehat{\beta}_{k, \infty}^\lambda(\mathcal{D}_n)$ is also asymptotically equivalent to the limit of $\widehat{\beta}_{k, M}^\lambda(\mathcal{D}_n; \{I_\ell\}_{\ell=1}^M)$ as the ensemble size $M \rightarrow \infty$ (conditioning on the full dataset \mathcal{D}_n). This also justifies using ∞ in the subscript to denote the full-ensemble estimator. For brevity, we simply write the estimators as $\widehat{\beta}_{k, M}^\lambda$, $\widehat{\beta}_{k, \infty}^\lambda$, and drop the dependency on $\mathcal{D}_n, \{I_\ell\}_{\ell=1}^M$. We will show certain structural equivalences in terms of linear projections in the family of estimators $\widehat{\beta}_{k, M}^\lambda$ along certain paths in the $(\lambda, p/k)$ -plane for arbitrary ensemble size $M \in \mathbb{N} \cup \{\infty\}$ (in Theorem 3). Apart from connecting the estimators, we will also show equivalences of various notions of risks (in Theorem 1), which we describe next.

Generalized risks. Since the ensemble ridge estimators are linear estimators, we evaluate their performance relative to the oracle parameter: $\beta_0 = \mathbb{E}[\mathbf{x}\mathbf{x}^\top]^{-1}\mathbb{E}[\mathbf{x}y]$, which is the best (population) linear projection of y onto \mathbf{x} and minimizes the linear regression error (see, e.g., [40–42]). Note that we can decompose any response y into: $y = f_{\text{LI}}(\mathbf{x}) + f_{\text{NL}}(\mathbf{x})$, where $f_{\text{LI}}(\mathbf{x}) = \beta_0^\top \mathbf{x}$ is the oracle linear predictor, and $f_{\text{NL}}(\mathbf{x}) = y - f_{\text{LI}}(\mathbf{x})$ is the nonlinear component that is not explained by $f_{\text{LI}}(\mathbf{x})$. The best linear projection has the useful property that $f_{\text{LI}}(\mathbf{x})$ is (linearly) uncorrelated with $f_{\text{NL}}(\mathbf{x})$, although they are generally dependent. It is worth mentioning that this does not imply that y and \mathbf{x} follow a linear regression model. Indeed, our framework allows any nonlinear dependence structure between them and is model-free for the joint distribution of (\mathbf{x}, y) . Relative to β_0 , we measure the performance of an estimator $\widehat{\beta}$ via the generalized mean squared risk defined compactly as follows:

$$R(\widehat{\beta}; \mathbf{A}, \mathbf{b}, \beta_0) = \frac{1}{\text{nrow}(\mathbf{A})} \|L_{\mathbf{A}, \mathbf{b}}(\widehat{\beta} - \beta_0)\|_2^2, \quad (3)$$

where $L_{\mathbf{A}, \mathbf{b}}(\beta) = \mathbf{A}\beta + \mathbf{b}$ is a linear functional. Note that \mathbf{A} and \mathbf{b} can potentially depend on the data, and their dimensions can vary depending on the statistical learning problem at hand. This framework includes various important statistical learning problems, as summarized in Table 2. Observe that the framework also includes various notions of prediction risks. One can use the test error formulation in Table 2 to obtain the prediction risk at any test point (\mathbf{x}_0, y_0) , where $y_0 = \mathbf{x}_0^\top \beta_0 + \varepsilon_0$, and ε_0 may depend on \mathbf{x}_0 and have non-zero mean. This also permits the test point to be drawn from a distribution that differs from the training distribution. Specifically, when $\varepsilon_0 = f_{\text{NL}}(\mathbf{x}_0)$ but \mathbf{x}_0 has a distribution different from \mathbf{x} , we obtain the prediction error under *covariate shift*. Similarly, when $\varepsilon_0 \neq f_{\text{NL}}(\mathbf{x}_0)$ but \mathbf{x}_0 and \mathbf{x} have the same distribution, we get the prediction error under *label shift*.

Asymptotic equivalence. For our theoretical analysis, we consider the proportional asymptotics regime, where the ratio of the feature size p to the sample size n tends to a fixed limiting *data aspect ratio* $\phi \in (0, \infty)$. To concisely present our results, we will use the elegant framework of asymptotic equivalence [12, 43, 44]. Let \mathbf{A}_p and \mathbf{B}_p be sequences of (additively) conformable matrices of arbitrary dimensions (including vectors and scalars). We say that \mathbf{A}_p and \mathbf{B}_p are *asymptotically equivalent*, denoted as $\mathbf{A}_p \simeq \mathbf{B}_p$, if $\lim_{p \rightarrow \infty} |\text{tr}[\mathbf{C}_p(\mathbf{A}_p - \mathbf{B}_p)]| = 0$ almost surely for any sequence of random matrices \mathbf{C}_p with bounded trace norm that are (multiplicatively) conformable and independent of \mathbf{A}_p and \mathbf{B}_p . Note that for sequences of scalar random variables, the definition simply reduces to the typical almost sure convergence of sequences of random variables involved.

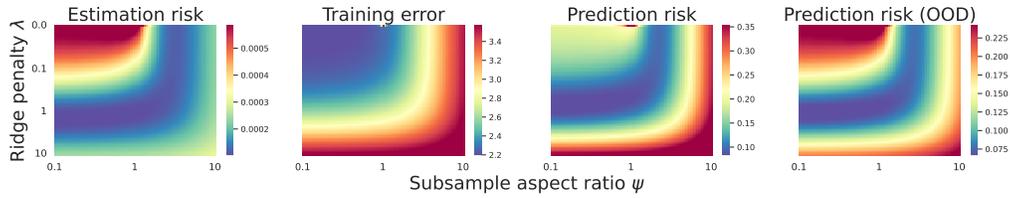


Figure 1: Heat map of the various generalized risks (estimation risk, training error, prediction risk, out-of-distribution (OOD) prediction risk) of full-ensemble ridge estimators (approximated with $M = 100$), for varying ridge penalties λ and subsample aspect ratios $\psi = p/k$ on the log-log scale. The data model is described in Appendix F.2 with $p = 500$, $n = 5000$, and $\phi = p/n = 0.1$. Observe that along the same path the values match well for all risks, in line with Theorem 1 and Proposition 2.

3 Generalized risk equivalences of ensemble estimators

We begin by examining the risk equivalences among different ensemble ridge estimators for various generalized risks defined as in (3). To prepare for our upcoming results, we impose two structural and moment assumptions on the distributions of the response variable and the feature vector.

Assumption 1 (Response variable distribution). Each response variable y_i for $i \in [n]$ has mean 0 and satisfies $\mathbb{E}[|y_i|^{4+\mu}] \leq M_\mu < \infty$ for some $\mu > 0$ and a constant M_μ .

Assumption 2 (Feature vector distribution). Each feature vector \mathbf{x}_i for $i \in [n]$ can be decomposed as $\mathbf{x}_i = \Sigma^{1/2} \mathbf{z}_i$, where $\mathbf{z}_i \in \mathbb{R}^p$ contains i.i.d. entries z_{ij} for $j \in [p]$ with mean 0, variance 1, and satisfy $\mathbb{E}[|z_{ij}|^{4+\nu}] \leq M_\nu < \infty$ for some $\nu > 0$ and a constant M_ν , and $\Sigma \in \mathbb{R}^{p \times p}$ is a deterministic and symmetric matrix with eigenvalues uniformly bounded between constants $r_{\min} > 0$ and $r_{\max} < \infty$.

Note that we do not impose any specific model assumptions on the response variable y in relation to the feature vector \mathbf{x} . We only assume bounded moments as stated in Assumption 1, making all of our subsequent results model-free. The zero-mean assumption for y is for simplicity since, in practice, centering can always be done by subtracting the sample mean. The bounded moment condition can also be satisfied if one imposes a stronger distributional assumption (e.g., sub-Gaussianity). Assumption 2 on the feature vector is common in the study of random matrix theory [45, 46] and the analysis of ridge and ridgeless regression [10, 47–49], which we refer to as *RMT features* for brevity.

Given a limiting data aspect ratio $\phi \in (0, \infty)$ and a limiting *subsample aspect ratio* $\bar{\psi} \in [\phi, \infty]$, our statement of equivalences between different ensemble estimators is defined through certain paths characterized by two endpoints $(0, \bar{\psi})$ and $(\bar{\lambda}, \phi)$. These endpoints correspond to the subsample ridgeless ensemble and the (non-ensemble) ridge predictor, respectively, with $\bar{\lambda}$ being the ridge penalty to be defined next. For that, let H_p be the empirical spectral distribution of Σ : $H_p(r) = p^{-1} \sum_{i=1}^p \mathbb{1}_{\{r_i \leq r\}}$, where r_i 's are the eigenvalues of Σ . Consider the following system of equations in $\bar{\lambda}$ and v :

$$\frac{1}{v} = \bar{\lambda} + \phi \int \frac{r}{1+vr} dH_p(r), \quad \text{and} \quad \frac{1}{v} = \bar{\psi} \int \frac{r}{1+vr} dH_p(r). \quad (4)$$

Existence and uniqueness of the solution $(\bar{\lambda}, v) \in [0, \infty]^2$ to the above equations are guaranteed by Corollary E.4. Now, define a path $\mathcal{P}(\bar{\lambda}; \phi, \bar{\psi})$ that passes through the endpoints $(0, \bar{\psi})$ and $(\bar{\lambda}, \phi)$:

$$\mathcal{P}(\bar{\lambda}; \phi, \bar{\psi}) = \{(1-\theta) \cdot (\bar{\lambda}, \phi) + \theta \cdot (0, \bar{\psi}) \mid \theta \in [0, 1]\}. \quad (5)$$

We are ready to state our results. We consider two cases for the generalized risk (3) depending on the relationships between (\mathbf{A}, \mathbf{b}) and \mathbf{X} . In the first case, when both \mathbf{A} and \mathbf{b} are independent of the data, Theorem 1 shows that the generalized risks are equivalent along the path (5).

Theorem 1 (Risk equivalences when $(\mathbf{A}, \mathbf{b}) \perp (\mathbf{X}, \mathbf{y})$). *Suppose Assumptions 1–2 hold. Let (\mathbf{A}, \mathbf{b}) be independent of (\mathbf{X}, \mathbf{y}) such that $\text{nrow}(\mathbf{A})^{-1/2} \|\mathbf{A}\|_{\text{op}}$ and $\|\mathbf{b}\|_2$ are almost surely bounded. Let $n, p \rightarrow \infty$ such that $p/n \rightarrow \phi \in (0, \infty)$. For any $\bar{\psi} \in [\phi, +\infty]$, let $\bar{\lambda}$ be as defined in (4). Then, for any pair of (λ_1, ψ_1) and (λ_2, ψ_2) on the path $\mathcal{P}(\bar{\lambda}; \phi, \bar{\psi})$ as defined in (5), the generalized risk functionals (3) of the full-ensemble estimator are asymptotically equivalent:*

$$R(\hat{\beta}_{\lfloor p/\psi_1 \rfloor, \infty}^{\lambda_1}; \mathbf{A}, \mathbf{b}, \beta_0) \simeq R(\hat{\beta}_{\lfloor p/\psi_2 \rfloor, \infty}^{\lambda_2}; \mathbf{A}, \mathbf{b}, \beta_0). \quad (6)$$

In other words, Theorem 1 establishes the equivalences of subsampling and ridge regularization in terms of generalized risk for many statistical learning problems, encompassing coefficient estimation,

coefficient confidence interval, and test error estimation. All of these problems fall in the category when (\mathbf{A}, \mathbf{b}) are independent of \mathbf{X} . However, there are other statistical learning problems not covered in Theorem 1, such as the in-sample prediction risk and the training error, which correspond to the case when $\mathbf{A} = \mathbf{X}$. Fortunately, the equivalences also apply to them, as summarized in Proposition 2.

Proposition 2 (Risk equivalences when $\mathbf{A} = \mathbf{X}$). *Under Assumptions 1–2, when $\mathbf{A} = \mathbf{X}$, the conclusion in Theorem 1 continue to hold for the cases of in-sample prediction risk ($\mathbf{b} = \mathbf{0}$) and training error ($\mathbf{b} = -\mathbf{f}_{\text{NL}}$).*

Both Theorem 1 and Proposition 2 provide specific second-order equivalences for the full-ensemble ridge estimators, in the sense that the quadratic functionals of the estimators associated with (\mathbf{A}, \mathbf{b}) are asymptotically the same. See Figure 1 for visual illustrations of both equivalences. These statements are presented separately because their proofs differ, with the proof for the dependent case being more intricate. We note that by combining the risk functionals in Table 2, it is possible to further extend the equivalences to other types of functionals not directly covered by statements above, using composition and continuous mapping. For example, it is not difficult to show that similar equivalences hold for the generalized cross-validation in the full ensembles [13]. This follows by combining the result for training error from Proposition 2 and for denominator concentration proved in Lemma 3.4 of [13].

Theorem 1 generalizes the existing equivalence results for the prediction risk [11–13] to include general risk functionals under milder assumptions. As summarized in Table 1, we allow for a general response model and feature covariance without assuming convergence of the spectral distributions H_p to a fixed distribution H . This flexibility is achieved by showing equivalences of the underlying resolvents in the estimators rather than deriving the limiting functionals as done in previous works. To extend results allowing for a general response, we generalize certain concentration results by leveraging tools from [14] to handle the dependency between the linear and non-linear components.

As alluded to earlier, it is worth noting that the generalized risk equivalences only hold in the full ensemble. However, by using the risk decomposition obtained from Lemma S.1.1 of [12] for finite ensembles, one can obtain the following relationship along the path in Theorem 3: $R(\hat{\beta}_{\lfloor p/\psi_1 \rfloor, M}^{\lambda_1}; \mathbf{A}, \mathbf{b}, \beta_0) - R(\hat{\beta}_{\lfloor p/\psi_2 \rfloor, M}^{\lambda_2}; \mathbf{A}, \mathbf{b}, \beta_0) \simeq \Delta/M$, for some Δ (independent of M) that is eventually almost surely bounded. In other words, the difference between any two estimators on the same path scales as $1/M$ when n is sufficiently large, as numerically verified in Figure F5.

4 Structural equivalences of ensemble estimators

The equivalences established in the previous section only hold for the full-ensemble estimators in a second-order sense. We can go a step further and ask if there exist any equivalences for the finite ensemble and if there are any equivalences at the estimator coordinate level between the estimators. Specifically, we aim to inspect whether each coordinate of the p -dimensional estimated coefficients asymptotically equals. This section establishes structural relationships between the estimators in a first-order sense that any bounded linear functionals of them are asymptotically the same.

Theorem 3 (Structural equivalences). *Suppose Assumptions 1–2 hold. Let $n, p \rightarrow \infty$ with $p/n \rightarrow \phi \in (0, \infty)$. For any $\bar{\psi} \in [\phi, +\infty]$, let $\bar{\lambda}$ be as in (4). Then, for any $M \in \mathbb{N} \cup \{\infty\}$ and any pair of (λ_1, ψ_1) and (λ_2, ψ_2) on the path (5), the M -ensemble estimators are asymptotically equivalent:*

$$\hat{\beta}_{\lfloor p/\psi_1 \rfloor, M}^{\lambda_1} \simeq \hat{\beta}_{\lfloor p/\psi_2 \rfloor, M}^{\lambda_2}. \quad (7)$$

Put another way, Theorem 3 implies that for any fixed ensemble size M , the M -ensemble ridgeless estimator at the limiting aspect ratio $\bar{\psi}$ is equivalent to ridge regression at the regularization level $\bar{\lambda}$. This equivalence is visually illustrated in Figure 2, where the data is simulated from an anisotropic and nonlinear model (see Appendix F.2 for more details). Furthermore, the contour path $\mathcal{P}(\bar{\lambda}; \phi, \bar{\psi})$ connecting the endpoints $(0, \bar{\psi})$ and $(\bar{\lambda}, \phi)$ is a straight line, although it may have varying slopes. Along any path, all M -ensemble estimators at the limiting aspect ratio ψ and ridge penalty λ are equivalent for all $(\lambda, \psi) \in \mathcal{P}(\bar{\lambda}; \phi, \bar{\psi})$. The equivalences here are for any arbitrary linear combinations of the estimators. In particular, this implies that the predicted values (or even any continuous function applied to them due to the continuous mapping theorem) of any test point will eventually be the same, almost surely, with respect to the training data. Finally, we remark that in the statement of Theorem 3, the equivalence is defined on extended reals in order to incorporate the ridgeless case when $\bar{\psi} = 1$.

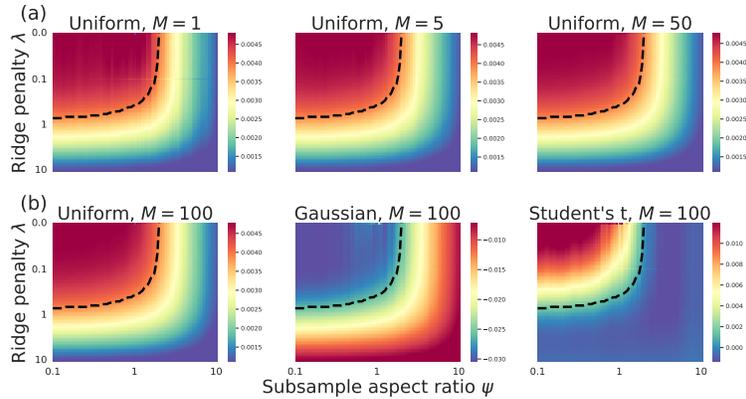


Figure 2: Heat map of the values of various linear functionals of ensemble ridge estimators $\mathbf{a}^\top \hat{\beta}_{k,M}^\lambda$, for varying ridge penalties λ , subsample aspect ratios $\psi = p/k$, and ensemble size M on the log-log scale. The data model is described in Appendix F.2 with $p = 1000$, $n = 10000$, and $\phi = p/n = 0.1$. (a) The values of the uniformly weighted projection for varying ensemble sizes ($M = 1, 5$, and 50). (b) The values of $M = 100$ for varying projection vectors (uniformly weighted, random Gaussian, and random student's t). The black dashed line is estimated based on Proposition 4 with $\bar{\psi} = 2$. Observe that the values along the data-dependent paths are indeed very similar, in line with Proposition 4.

The paths (5) in Theorem 3 are defined via the spectral distribution H_p , which requires knowledge of Σ . This is often difficult to obtain in practice from the observed data. Fortunately, we can provide an alternative characterization for the path (5) solely through the data, as summarized in Proposition 4.

Proposition 4 (Data-dependent equivalence path characterization). *Suppose Assumptions 1–2 hold. Define $\phi_n = p/n$. Let $k \leq n$ be the subsample size and denote by $\bar{\psi}_n = p/k$. For any $M \in \mathbb{N} \cup \{\infty\}$, let $\bar{\lambda}_n$ be the value that satisfies the following equation in ensemble ridgeless and ridge gram matrices:*

$$\frac{1}{M} \sum_{\ell=1}^M \frac{1}{k} \operatorname{tr} \left[\left(\frac{1}{k} \mathbf{L}_{I_\ell} \mathbf{X} \mathbf{X}^\top \mathbf{L}_{I_\ell} \right)^+ \right] = \frac{1}{n} \operatorname{tr} \left[\left(\frac{1}{n} \mathbf{X} \mathbf{X}^\top + \bar{\lambda}_n \mathbf{I}_n \right)^{-1} \right]. \quad (8)$$

Define the data-dependent path $\mathcal{P}_n = \mathcal{P}(\bar{\lambda}_n; \phi_n, \bar{\psi}_n)$. Then, the conclusion in Theorem 3 continues to hold if we replace the (population) path \mathcal{P} by the (data-dependent) path \mathcal{P}_n .

The term “data-dependent path” signifies that we can estimate the level of implicit regularization induced by subsampling solely based on the observed data by solving (8). We remark that (8) always has at least one solution for a given triplet of (n, k, p) . This is because the right-hand side of (8) is monotonically decreasing in λ , and the left-hand side always lies within the range of the right-hand side. The powerful implication of the characterization in Proposition 4 is that it enables practical computation of the path using real-world data. In Section 6, we will demonstrate this approach on various real-world datasets, allowing us to predict an equivalent amount of explicit regularization matching the implicit regularization due to subsampling.

One natural extension of the results is to consider the generalized ridge as a base estimator (1):

$$\hat{\beta}_k^{\lambda, \mathbf{G}}(\mathcal{D}_I) = \operatorname{argmin}_{\beta \in \mathbb{R}^p} \frac{1}{k} \sum_{i \in I} (y_i - \mathbf{x}_i^\top \beta)^2 + \lambda \|\mathbf{G}^{\frac{1}{2}} \beta\|_2^2, \quad (9)$$

where $\mathbf{G} \in \mathbb{R}^{p \times p}$ is a positive definite matrix. The equivalences still hold, albeit on different paths:

Corollary 5 (Equivalences for generalized ridge regression). *Suppose Assumptions 1–2 hold. Let $\mathbf{G} \in \mathbb{R}^{p \times p}$ be a deterministic and symmetric matrix with eigenvalues uniformly bounded between constants $g_{\min} > 0$ and $g_{\max} < \infty$. Let $n, p \rightarrow \infty$ such that $p/n \rightarrow \phi \in (0, \infty)$. For any $\bar{\psi} \in [\phi, +\infty]$, let $\bar{\lambda}$ be as defined in (4) with H_p replaced \tilde{H}_p , the empirical spectral distribution of $\mathbf{G}^{-1/2} \Sigma \mathbf{G}^{-1/2}$. Then, the conclusions in Theorems 1 and 3 continue to hold for the generalized ridge ensemble predictors.*

Another extension of our results is to incorporate subquadratic risk functionals in Theorem 1 to derive equivalences of the cumulative distribution functions of the predictive error distributions associated

with the estimators along the equivalence paths. One can use such equivalences for downstream inference questions. Finally, while our statements are asymptotic in nature, we expect a similar analysis as done in [50] can yield finite-sample statements. We leave these extensions for the future.

5 Implications of equivalences

The generalized risk equivalences establish a connection between the risks of different ridge ensembles, including the ridge predictors on full data and the full-ensemble ridgeless predictors. These connections enable us to transfer properties from one estimator to the other. For example, by examining the impact of subsampling on the estimator’s performance, we can better understand how to optimize the ridge penalty. We will next provide an example of this. Many common methods, such as ridgeless or lassoleless predictors, have been recently shown to exhibit non-monotonic behavior in the sample size or the limiting aspect ratio. An open problem raised by Nakkiran et al. [1, Conjecture 1] asks whether the prediction risk of the ridge regression with optimal ridge penalty λ^* is monotonically increasing in the data aspect ratio $\phi = p/n$. The non-monotonicity risk behavior implies that more data can hurt performance, and it’s important to investigate whether optimal ridge regression also suffers from this issue. Our equivalence results provide a surprising (even to us) indirect way to answer this question affirmatively under proportional asymptotics.

To analyze the monotonicity of the risk, we will need to impose two additional regularity assumptions to guarantee the convergence of the prediction risk. Let $\Sigma = \sum_{j=1}^p r_j \mathbf{w}_j \mathbf{w}_j^\top$ denote the eigenvalue decomposition, where (r_j, \mathbf{w}_j) ’s are pairs of associated eigenvalue and normalized eigenvector. The following assumptions ensure that the hardness of the underlying regression problems across different limiting data aspect ratios is comparable.

Assumption 3 (Spectral convergence). We assume there exists a deterministic distribution H such that the empirical spectral distribution of Σ , $H_p(r) = p^{-1} \sum_{i=1}^p \mathbb{1}_{\{r_i \leq r\}}$, weakly converges to H , almost surely (with respect to \mathbf{x}).

Assumption 4 (Limiting signal and noise energies). We assume there exists a deterministic distribution G such that the empirical distribution of β_0 ’s (squared) projection onto Σ ’s eigenspace, $G_p(r) = \|\beta_0\|_2^{-2} \sum_{i=1}^p (\beta_0^\top \mathbf{w}_i)^2 \mathbb{1}_{\{r_i \leq r\}}$, weakly converges to G . As $n, p \rightarrow \infty$ and $p/n \rightarrow \phi \in (0, \infty)$, the limiting linearized energy $\rho^2 = \lim \|\beta_0\|_2^2$ and the limiting nonlinearized energy $\sigma^2 = \lim \|f_{\text{NL}}\|_{L_2}^2$ are finite.

Assumption 3 is commonly used in random matrix theory and overparameterized learning [10, 12, 13] and ensures that the limiting spectral distribution of the sample covariance matrix converges to a fixed distribution. Under these assumptions, we show in Appendix C that there exists a deterministic function $\mathcal{R}(\lambda; \phi, \psi)$, such that $R(\beta_{k, \infty}^\lambda; \mathbf{A}, \mathbf{b}, \beta_0) \simeq \mathcal{R}(\lambda; \phi, \psi)$. Notably, the deterministic profile on the right-hand side is a monotonically increasing and a continuous function in the first parameter ϕ when limiting values of $\|\beta_0\|_2^2$ and $\|f_{\text{NL}}\|_{L_2}^2$ are fixed (i.e., when Assumption 4 holds). Moreover, the deterministic risk equivalent of the optimal ridge predictor matches that of the optimal full-ensemble ridgeless predictor; in other words, $\min_{\lambda \geq 0} \mathcal{R}(\lambda; \phi, \phi) = \min_{\psi \geq \phi} \mathcal{R}(0; \phi, \psi)$. The same monotonicity property of the two optimized functions leads to the following result.

Theorem 6 (Monotonicity of prediction risk with optimal ridge regularization). *Suppose the conditions of Theorem 1 and Assumptions 3–4 hold. Let $k, n, p \rightarrow \infty$ such that $p/n \rightarrow \phi \in (0, \infty)$ and $p/k \rightarrow \psi \in [\phi, \infty]$. Then, for $\mathbf{A} = \Sigma^{1/2}$ and $\mathbf{b} = \mathbf{0}$, the optimal risk of the ridgeless ensemble, $\min_{\psi \geq \phi} \mathcal{R}(0; \phi, \psi)$, is monotonically increasing in ϕ . Consequently, the optimal risk of the ridge predictor, $\min_{\lambda \geq 0} \mathcal{R}(\lambda; \phi, \phi)$, is also monotonically increasing in ϕ .*

Under Assumptions 3 and 4, the linearized signal-to-noise ratio (SNR) is maintained at the same level across varying distributions as ϕ changes. The key message of Theorem 6 is then that, for a sequence of problems with the same SNR (indicating the same level of regression hardness), the asymptotic prediction risk of optimized ridge risk gets monotonically worse as the aspect ratio of the problem increases. This is intuitive because a smaller ϕ corresponds to more samples than features. In this sense, Theorem 6 certifies that optimal ridge uses the available data effectively, avoiding sample-wise non-monotonicity. Moreover, as the null risk is finite, this also shows that the optimal ridge estimator mitigates the “double or multiple descents” behaviors in the generalization error [1, 5, 51].

Attempting to prove Theorem 6 directly is challenging due to the lack of a closed-form expression for the optimal risk of ridge regression in general. Moreover, the risk profile of ridge regression,

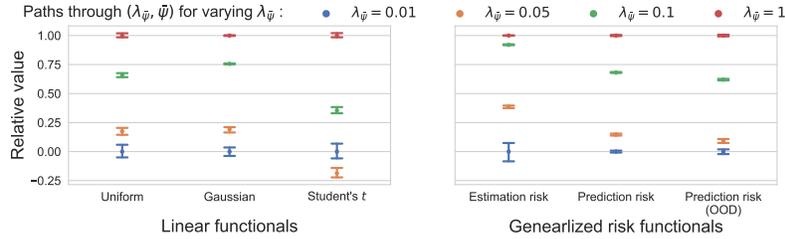


Figure 3: Comparison of various linear functionals (Theorem 3) and generalized risk functionals (Theorem 1) on different paths evaluated on CIFAR-10. The aspect ratios $\phi = p/n$ and $\bar{\psi} = 4\phi$ are estimated from the dataset and fixed. For different values of $\lambda_{\bar{\psi}}$ (different colors), we estimate $\bar{\lambda}$ from Proposition 4, which gives a path between $(\lambda_{\bar{\psi}}, \bar{\psi})$ and $(\bar{\lambda}, \phi)$. For each path, we uniformly sample 5 points and compute the functionals of the ridge ensemble using $M = 100$. The values of different paths are then normalized by subtracting the mean value in the first path and dividing by the mean difference of values on the last and first paths.

$\mathcal{R}(\lambda; \phi, \psi)$, does not exhibit any particular structure as a function of λ . On the other hand, the risk profile of the full-ensemble ridgeless, $\mathcal{R}(0; \phi, \psi)$, has a very nice structure in terms of ψ . This, coupled with our equivalence result, allows one to prove quite non-trivial behavior of optimal ridge. See Figure F6 for illustrations.

6 Discussion

Motivated by the recent subsampling and ridge equivalences for the prediction risk [11–13], this paper establishes generalized risk equivalences (Section 3) and structural equivalences (Section 4) within the family of ensemble ridge estimators. Our results precisely link the implicit regularization of subsampling to explicit ridge penalization via a path \mathcal{P} defined in (5), which connects two endpoints $(0, \bar{\psi})$ and $(\bar{\lambda}, \phi)$ regardless of the risk functionals. Furthermore, we provide a data-dependent method (Proposition 4) to estimate this path. Our results do not assume any specific relationship between the response variable y and the feature vector x . We next explore some extensions of these equivalences.

Real data. While our theoretical results assume RMT features (Assumption 2), we anticipate that the equivalences will very likely hold under more general features [50]. To verify this, we examine the equivalences in Proposition 4 with $M = 100$ on real-world datasets. Figure 3 depicts both the linear functionals in Theorem 3 and the generalized quadratic functionals in Theorem 1 computed along four paths, shown in different colors. We observe that the variation within each path is small, while different paths have different functional values. This suggests that the theoretical finding in the previous sections also hold quite well on real-world datasets. We investigate this further by varying $\bar{\psi}$ on the three image datasets, CIFAR-10, MNIST, and USPS. Figure F7 shows similar values and trends at the two points. These experiments demonstrate that the amount of explicit regularization induced by subsampling can indeed be accurately predicted based on the observed data using Proposition 4.

Random features. Closely related to two-layer neural networks [52], we can consider the random feature model, $f(x; \beta, \mathbf{F}) = \beta^\top \varphi(\mathbf{F}x)$, where $\mathbf{F} \in \mathbb{R}^{d \times p}$ is some randomly initialized weight matrix, and $\varphi: \mathbb{R} \rightarrow \mathbb{R}$ is a nonlinear activation function applied element-wise to $\mathbf{F}x$. As from “universality/invariance”, certain random feature models are asymptotically equivalent to a surrogate linear Gaussian model with a matching covariance matrix [53], we expect the theoretical results in Theorems 3–1 to likely hold, although the relationship (4) will now depend on the non-linear activation functions. Empirically, we indeed observe similar equivalence phenomena of the prediction risks with random features ridge regression using sigmoid, ReLU, and tanh activation functions, as shown in Figure 4. Analogous to Proposition 4, we conjecture a similar data-driven relationship between $(\bar{\lambda}_n, \phi_n)$ and $(0, \psi_n)$ for random features $\tilde{\mathbf{X}} = \varphi(\mathbf{X}\mathbf{F}^\top)$:

Conjecture 7 (Data-dependent equivalences for random features (informal)). *Suppose Assumptions 1–2 hold. Define $\phi_n = p/n$. Let $k \leq n$ be the subsample size and denote by $\psi_n = p/k$. Suppose φ satisfies certain regularity conditions. For any $M \in \mathbb{N} \cup \{\infty\}$, let $\bar{\lambda}_n$ be the value that satisfies*

$$\frac{1}{M} \sum_{\ell=1}^M \frac{1}{k} \operatorname{tr} \left[\left(\frac{1}{k} \varphi(\mathbf{L}_{I_\ell} \mathbf{X} \mathbf{F}^\top) \varphi(\mathbf{L}_{I_\ell} \mathbf{X} \mathbf{F}^\top)^\top \right)^+ \right] = \frac{1}{n} \operatorname{tr} \left[\left(\frac{1}{n} \varphi(\mathbf{X} \mathbf{F}^\top) \varphi(\mathbf{X} \mathbf{F}^\top)^\top + \bar{\lambda}_n \mathbf{I}_n \right)^{-1} \right].$$

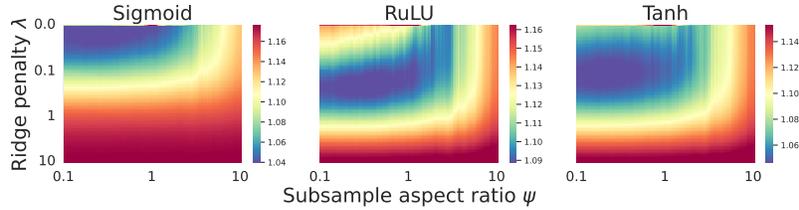


Figure 4: Heat map of prediction risks of full-ensemble ridge estimators (approximated with $M = 100$) using random features, for varying ridge penalties λ , subsample aspect ratios $\psi = p/d$ on the log-log scale. We consider random features $\varphi(\mathbf{F}\mathbf{x}_i) \in \mathbb{R}^d$, where φ is an activation function (sigmoid, ReLU, or tanh). The data model is described in Appendix F.6 with $p = 250$, $d = 500$, $n = 5000$, and $\phi = d/n = 0.1$. As in Theorem 1, we see clear risk equivalence paths across activations.

Define the data-dependent path $\mathcal{P}_n = \mathcal{P}(\bar{\lambda}_n; \phi_n, \bar{\psi}_n)$. Then, the conclusions in Theorem 1, Proposition 2, and Theorem 3 continue to hold on $(\varphi(\mathbf{X}\mathbf{F}^\top), \mathbf{y})$ with \mathcal{P} replaced by \mathcal{P}_n .

Kernel features. We can also consider kernel ridge regression. For a given feature map $\Phi : \mathbb{R}^p \rightarrow \mathbb{R}^d$, the kernel ridge estimator (in the primal form) is defined as:

$$\hat{\beta}_k^\lambda(\mathcal{D}_I) = \operatorname{argmin}_{\beta \in \mathbb{R}^p} \sum_{i \in I} (k^{-1/2}y_i - k^{-1/2}\Phi(\mathbf{x}_i)^\top \beta)^2 + \lambda \|\beta\|_2^2 + \frac{k}{p} \lambda \|\beta\|_2^2.$$

Leveraging the kernel trick, the preceding optimization problem translates to solving the following problem (in the dual domain): $\hat{\alpha}_k^\lambda(\mathcal{D}_I) = \operatorname{argmin}_{\alpha \in \mathbb{R}^k} \alpha^\top (\mathbf{K}_I + k\lambda \mathbf{I}_k) \alpha + \alpha^\top \mathbf{y}_I$, where $\mathbf{K}_I = \Phi_I \Phi_I^\top \in \mathbb{R}^{k \times k}$ is the kernel matrix and $\Phi_I = (\Phi(\mathbf{x}_i))_{i \in I} \in \mathbb{R}^{n \times d}$ is the feature matrix. The correspondence between the dual and primal solutions is simply given by: $\hat{\beta}_k^\lambda(\mathcal{D}_I) = \Phi_I^\top \hat{\alpha}_k^\lambda(\mathcal{D}_I)$.

Figure F8 illustrate results for kernel ridge regression using the same data-generating process as in the previous subsection. The figure shows the prediction risk of kernel ridge ensembles for polynomial, Gaussian, and Laplace kernels exhibit similar equivalence patterns, leading us to formulate an analogous conjecture for kernel ridge regression (which when $\Phi(\mathbf{x}) = \mathbf{x}$ gives back Proposition 4 with appropriate rescaling of features):

Conjecture 8 (Data-dependent equivalences for kernel features (informal)). *Suppose Assumptions 1–2 hold. Define $\phi_n = p/n$. Suppose the kernel K satisfies certain regularity conditions. Let $k \leq n$ be the subsample size and denote by $\bar{\psi}_n = p/k$. For any $M \in \mathbb{N} \cup \{\infty\}$, let $\bar{\lambda}_n$ be a solution to*

$$\frac{1}{M} \sum_{\ell=1}^M \operatorname{tr} [\mathbf{K}_{I_\ell}^+] = \operatorname{tr} \left[\left(\mathbf{K}_{[n]} + \frac{n}{p} \bar{\lambda}_n \mathbf{I}_n \right)^{-1} \right].$$

Define the data-dependent path $\mathcal{P}_n = \mathcal{P}(\bar{\lambda}_n; \phi_n, \bar{\psi}_n)$. Then, the conclusions in Theorem 1, Proposition 2, and Theorem 3 continue to hold for kernel ridge ensembles with \mathcal{P} replaced by \mathcal{P}_n .

The empirical evidence here strongly suggests that the relationship between subsampling and ridge regression, which we have proved for ridge regression, holds true at least when the gram matrix “linearizes” in the sense of asymptotic equivalence [14, 54, 55]. This intriguing observation opens up a whole host of avenues towards fully understanding the universality of this relationship and establishing precise connections for a broader range of models [17, 18, 20, 21]. We hope to share more on these compelling directions in the near future.

Tying other implicit regularizations. Finally, as noted in related work, the equivalences between ridge regularization and subsampling established in this paper naturally offer opportunities to interpret and understand the other implicit regularization effects such as dropout regularization [27, 28], early stopping in gradient descent variants [36, 38, 39]. Furthermore, these equivalences provide avenues to understand the combined effects of these forms of implicit regularization, such as the effects of both subsample aggregating and gradient descent in mini-batch gradient descent, and contrast them to explicit regularization. Whether the explicit regularization can always be cleanly expressed as ridge-like regularization or takes a more generic form is an intriguing question. Exploring this is exciting, and we hope our readers also share this excitement!

Acknowledgments and Disclosure of Funding

We are indebted to Ryan Tibshirani, Alessandro Rinaldo, Arun Kuchibhotla, Yuting Wei, Matey Neykov, Edgar Dobriban, Daniel LeJeune, Shamindra Shrotriya, Yuchen Wu for many insightful conversations. We are also grateful to anonymous reviewers of the precursor [13] for encouraging us to work on several follow-up directions that make up parts of this successor. On a more personal note: The lack of proof of risk monotonicity of optimally-tuned ridge in general had been “monotone bugging” (in a good way) the first author, ever since plotting Figure 8 (and other similar plots) in [51]. The resolution in Theorem 6 is of great personal significance to him and he is heartily appreciative of all the collaborators on related threads, especially the second author, for being a dedicated climbing partner on this technical “ascent”. At the risk of over analogizing, the metaphor of letting a nut soak in water instead of forcibly cracking it open with a hammer, from Grothendieck², may perhaps describe the proof of Theorem 6. Establishing subsampling and ridge equivalences first and using it to demonstrate ridge monotonicity, in retrospect, seems a much less arduous route than a direct brute-force attempt (with calculus, say). But we will let the readers decide!

This work used the Bridges2 system at the Pittsburgh Supercomputing Center (PSC) through allocations MTH230020 from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) program. The code for reproducing the results of this paper can be found at <https://jaydu1.github.io/overparameterized-ensembling/equiv>.

References

- [1] Preetum Nakkiran, Prayaag Venkat, Sham M. Kakade, and Tengyu Ma. Optimal regularization can mitigate double descent. In *International Conference on Learning Representations*, 2021.
- [2] Leo Breiman. Bagging predictors. *Machine Learning*, 24(2):123–140, 1996.
- [3] Peter Bühlmann and Bin Yu. Analyzing bagging. *The Annals of Statistics*, 30(4):927–961, 2002.
- [4] Abraham J. Wyner, Matthew Olson, Justin Bleich, and David Mease. Explaining the success of adaboost and random forests as interpolating classifiers. *The Journal of Machine Learning Research*, 18(1):1558–1590, 2017.
- [5] Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, 2019.
- [6] Lucas Mentch and Siyu Zhou. Randomization as regularization: A degrees of freedom explanation for random forest success. *The Journal of Machine Learning Research*, 21(1):6918–6953, 2020.
- [7] Arthur E. Hoerl and Robert W. Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67, 1970.
- [8] Arthur E. Hoerl and Robert W. Kennard. Ridge regression: Applications to nonorthogonal problems. *Technometrics*, 12(1):69–82, 1970.
- [9] Trevor Hastie. Ridge regularization: An essential concept in data science. *Technometrics*, 62(4):426–433, 2020.
- [10] Trevor Hastie, Andrea Montanari, Saharon Rosset, and Ryan J. Tibshirani. Surprises in high-dimensional ridgeless least squares interpolation. *The Annals of Statistics*, 50(2):949–986, 2022.
- [11] Daniel LeJeune, Hamid Javadi, and Richard Baraniuk. The implicit regularization of ordinary least squares ensembles. In *International Conference on Artificial Intelligence and Statistics*, 2020.

²<http://math.stanford.edu/~vakil/216blog/FOAGmay2523public.pdf>, page 9.

- [12] Pratik Patil, Jin-Hong Du, and Arun Kumar Kuchibhotla. Bagging in overparameterized learning: Risk characterization and risk monotonicity. *arXiv preprint arXiv:2210.11445*, 2022.
- [13] Jin-Hong Du, Pratik Patil, and Arun Kumar Kuchibhotla. Subsample ridge ensembles: Equivalences and generalized cross-validation. In *International Conference on Machine Learning*, 2023.
- [14] Peter L. Bartlett, Andrea Montanari, and Alexander Rakhlin. Deep learning: A statistical viewpoint. *Acta numerica*, 30:87–201, 2021.
- [15] Trevor Hastie, Robert Tibshirani, and Jerome H. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, 2009. Second edition.
- [16] Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- [17] L.K. Hansen and P. Salamon. Neural network ensembles. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(10):993–1001, 1990.
- [18] Anders Krogh and Jesper Vedelsby. Neural network ensembles, cross validation, and active learning. *Advances in neural information processing systems*, 1994.
- [19] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems*, 2017.
- [20] Michael Perrone. Putting it all together: Methods for combining neural networks. In *Advances in Neural Information Processing Systems*, 1993.
- [21] Anders Krogh and Peter Sollich. Learning with ensembles: How overfitting can be useful. In *Advances in Neural Information Processing Systems*, 1995.
- [22] Anders Krogh and Peter Sollich. Statistical mechanics of ensemble learning. *Physical Review E*, 55(6):811, 1997.
- [23] Stéphane d’Ascoli, Maria Refinetti, Giulio Biroli, and Florent Krzakala. Double trouble in double descent: Bias and variance (s) in the lazy regime. In *International Conference on Machine Learning*, pages 2280–2290. PMLR, 2020.
- [24] Ben Adlam and Jeffrey Pennington. Understanding double descent requires a fine-grained bias-variance decomposition. In *Advances in Neural Information Processing Systems*, 2020.
- [25] Bruno Loureiro, Cedric Gerbelot, Maria Refinetti, Gabriele Sicuro, and Florent Krzakala. Fluctuations, bias, variance & ensemble of learners: Exact asymptotics for convex losses in high-dimension. In *International Conference on Machine Learning*, 2022.
- [26] Tianyi Yao, Daniel LeJeune, Hamid Javadi, Richard G. Baraniuk, and Genevera I. Allen. Minipatch learning as implicit ridge-like regularization. In *International Conference on Big Data and Smart Computing*, 2021.
- [27] Stefan Wager, Sida Wang, and Percy S. Liang. Dropout training as adaptive regularization. In *Advances in Neural Information Processing Systems*, 2013.
- [28] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- [29] Andrew R. Webb. Functional approximation by feed-forward networks: A least-squares approach to generalization. *IEEE Transactions on Neural Networks*, 5(3):363–371, 1994.
- [30] Chris M. Bishop. Training with noise is equivalent to Tikhonov regularization. *Neural computation*, 7(1):108–116, 1995.

- [31] Gian-Andrea Thanei, Christina Heinze, and Nicolai Meinshausen. Random projections for large-scale regression. *Big and Complex Data Analysis: Methodologies and Applications*, pages 51–68, 2017.
- [32] Alessandro Rudi, Raffaello Camoriano, and Lorenzo Rosasco. Less is more: Nyström computational regularization. In *Advances in Neural Information Processing Systems*, volume 28, 2015.
- [33] Daniel LeJeune, Pratik Patil, Hamid Javadi, Richard G Baraniuk, and Ryan J. Tibshirani. Asymptotics of the sketched pseudoinverse. *arXiv preprint arXiv:2211.03751*, 2022.
- [34] Chi-Heng Lin, Chiraag Kaushik, Eva L. Dyer, and Vidya Muthukumar. The good, the bad and the ugly sides of data augmentation: An implicit spectral regularization perspective. *arXiv preprint arXiv:2210.05021*, 2022.
- [35] Marina Skurichina and Robert P. W. Duin. Regularization by adding redundant features. In *Advances in Pattern Recognition: Joint IAPR International Workshops SSPR'98 and SPR'98 Sydney, Australia, August 11–13, 1998 Proceedings*, pages 564–572. Springer, 1998.
- [36] Gergely Neu and Lorenzo Rosasco. Iterate averaging as regularization for stochastic gradient descent. In *Conference On Learning Theory*, pages 3222–3242. PMLR, 2018.
- [37] Arun Suggala, Adarsh Prasad, and Pradeep K. Ravikumar. Connecting optimization and regularization paths. In *Advances in Neural Information Processing Systems*, 2018.
- [38] Alnur Ali, Edgar Dobriban, and Ryan J. Tibshirani. The implicit regularization of stochastic gradient flow for least squares. In *International conference on machine learning*, 2020.
- [39] Alnur Ali, J. Zico Kolter, and Ryan J. Tibshirani. A continuous-time view of early stopping for least squares regression. In *International Conference on Artificial Intelligence and Statistics*, 2019.
- [40] A. Buja, R. Berk, L. Brown, E. George, E. Pitkin, M. Traskin, K. Zhang, and L. Zhao. A conspiracy of random predictors and model violations against classical inference in regression. *arXiv preprint arXiv:1404.1578*, 2014.
- [41] Andreas Buja, Lawrence Brown, Richard Berk, Edward George, Emil Pitkin, Mikhail Traskin, Kai Zhang, and Linda Zhao. Models as approximations I. *Statistical Science*, 34(4):523–544, 2019.
- [42] Andreas Buja, Lawrence Brown, Arun Kumar Kuchibhotla, Richard Berk, Edward George, and Linda Zhao. Models as approximations II. *Statistical Science*, 34(4):545–565, 2019.
- [43] Edgar Dobriban and Yue Sheng. Wonder: Weighted one-shot distributed ridge regression in high dimensions. *Journal of Machine Learning Research*, 21(66):1–52, 2020.
- [44] Edgar Dobriban and Yue Sheng. Distributed linear regression by averaging. *The Annals of Statistics*, 49(2):918–943, 2021.
- [45] Zhidong Bai and Jack W. Silverstein. *Spectral Analysis of Large Dimensional Random Matrices*. Springer, 2010. Second edition.
- [46] Noureddine El Karoui. The spectrum of kernel random matrices. *The Annals of Statistics*, 38(1):1–50, 2010.
- [47] Noureddine El Karoui and Holger Kösters. Geometric sensitivity of random matrix results: Consequences for shrinkage estimators of covariance and related statistical methods. *arXiv preprint arXiv:1105.1404*, 2011.
- [48] Noureddine El Karoui. Asymptotic behavior of unregularized and ridge-regularized high-dimensional robust regression estimators: rigorous results. *arXiv preprint arXiv:1311.2445*, 2013.
- [49] Edgar Dobriban and Stefan Wager. High-dimensional asymptotics of prediction: Ridge regression and classification. *The Annals of Statistics*, 46(1):247–279, 2018.

- [50] Chen Cheng and Andrea Montanari. Dimension free ridge regression. *arXiv preprint arXiv:2210.08571*, 2022.
- [51] Pratik Patil, Arun Kumar Kuchibhotla, Yuting Wei, and Alessandro Rinaldo. Mitigating multiple descents: A model-agnostic framework for risk monotonicity. *arXiv preprint arXiv:2205.12937*, 2022.
- [52] 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.
- [53] Hong Hu and Yue M. Lu. Universality laws for high-dimensional learning with random features. *IEEE Transactions on Information Theory*, 69(3):1932–1964, 2023.
- [54] Tengyuan Liang and Alexander Rakhlin. Just interpolate: Kernel “ridgeless” regression can generalize. *The Annals of Statistics*, 2020.
- [55] Mojtaba Sahraee-Ardakan, Melikasadat Emami, Parthe Pandit, Sundeep Rangan, and Alyson K Fletcher. Kernel methods and multi-layer perceptrons learn linear models in high dimensions. *arXiv preprint arXiv:2201.08082*, 2022.
- [56] Antti Knowles and Jun Yin. Anisotropic local laws for random matrices. *Probability Theory and Related Fields*, 169:257–352, 2017.
- [57] László Erdős and Horng-Tzer Yau. *A Dynamical Approach to Random Matrix Theory*. American Mathematical Society, 2017.
- [58] Francisco Rubio and Xavier Mestre. Spectral convergence for a general class of random matrices. *Statistics & probability letters*, 81(5):592–602, 2011.
- [59] TorchVision. Torchvision: PyTorch’s computer vision library. <https://github.com/pytorch/vision>, 2016.
- [60] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

Supplementary material for “Generalized equivalences between subsampling and ridge regularization”

This document serves as a supplement to the paper “Generalized equivalences between subsampling and ridge regularization.” The initial (unnumbered) section of the supplement provides a summary of the notation used in both the paper and supplement, followed by an organization for the supplement.

Notation. Below we provide an overview of the notation used throughout.

General notation:

Notation	Description
Non-bold	Denotes univariates (scalars, functions, distributions, etc.) (e.g., λ, f, H).
Bold lower case	Denotes vectors (e.g., $\mathbf{x}, \boldsymbol{\beta}$).
Bold upper case	Denotes matrices (e.g., \mathbf{X}).
Calligraphic font	Denotes sets (e.g., \mathcal{D}).
Script font	Denotes certain limiting functions (e.g., \mathcal{R} in (27)).
\mathbb{N}	Set of positive integers.
\mathbb{R}, \mathbb{R}_+	Set of real numbers and positive real numbers.
\mathbb{C}, \mathbb{C}^+	Set of complex numbers and complex numbers with positive imaginary part.
$[n]$	Set $\{1, \dots, n\}$ for natural number n .
$(x)_+$	The positive part of real number x .
$\lfloor x \rfloor, \lceil x \rceil$	The floor and ceiling of real number x .
$\ f\ _{L_2}$	The L_2 norm of function f , $\ f\ _{L_2} = \mathbb{E}_{\mathbf{x}}[f^2(\mathbf{x})]$.
$\ \boldsymbol{\beta}\ _2$	The ℓ_2 norm of vector $\boldsymbol{\beta}$.
$\langle \mathbf{v}, \mathbf{w} \rangle$	The inner product of vectors \mathbf{v} and \mathbf{w} .
$\mathbf{X}^\top, \mathbf{X}^+$	The transpose and Moore-Penrose inverse of matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$.
$\text{tr}[\mathbf{A}], \mathbf{A}^{-1}$	The trace and inverse (if invertible) of a square matrix $\mathbf{A} \in \mathbb{R}^{p \times p}$.
$\boldsymbol{\Sigma}^{1/2}$	The principal square root of positive semi-definite matrix $\boldsymbol{\Sigma}$.
\mathbf{I}_p or \mathbf{I}	The $p \times p$ identity matrix.
$\ \mathbf{X}\ _{\text{op}}$	The operator norm (spectral norm) of real matrix \mathbf{X} .
$\ \mathbf{X}\ _{\text{tr}}$	The trace norm (nuclear norm) $\text{tr}[(\mathbf{X}^\top \mathbf{X})^{1/2}]$ of real matrix \mathbf{X} .
$f(\mathbf{A})$	The matrix $\mathbf{V}f(\mathbf{R})\mathbf{V}^{-1}$ where $\mathbf{A} = \mathbf{V}\mathbf{R}\mathbf{V}^{-1}$ is eigenvalue decomposition, and $f(\mathbf{R})$ is f applied (elementwise) to the diagonals of \mathbf{R} .
$\mathbf{A} \preceq \mathbf{B}$	The Loewner ordering for symmetric matrices \mathbf{A} and \mathbf{B} .
$\xrightarrow{p}, \xrightarrow{\text{a.s.}}, \xrightarrow{d}$	Almost sure convergence, convergence in probability, and weak convergence.

Specific notation:

Notation	Description
(\mathbf{x}, y)	The population random vector in $\mathbb{R}^p \times \mathbb{R}$.
\mathcal{D}_n	A dataset containing i.i.d. samples of (\mathbf{x}, y) : $\mathcal{D}_n = \{(\mathbf{x}_i, y_i) : i \in [n]\}$.
\mathbf{X}, \mathbf{y}	The feature matrix in $\mathbb{R}^{n \times p}$ and the response vector in \mathbb{R}^n .
\mathcal{D}_I	A subsampled dataset $\mathcal{D}_I = \{(\mathbf{x}_i, y_i) : i \in I\}$.
$\hat{\boldsymbol{\beta}}_k^\lambda(\mathcal{D}_I)$	A ridge estimator fitted on \mathcal{D}_I with $ I = k$ observations and ridge penalty λ .
$\hat{\boldsymbol{\beta}}_{k,M}^\lambda(\mathcal{D}_n; \{I_\ell\}_{\ell=1}^M)$	An M -ensemble ridge estimator fitted on $\{\mathcal{D}_{I_\ell}\}_{\ell=1}^M$ with subsample size k and ridge penalty λ .
$\boldsymbol{\beta}_0$	The best (population) linear projection coefficient of y onto \mathbf{x} : $\mathbb{E}[\mathbf{x}\mathbf{x}^\top]^{-1}\mathbb{E}[\mathbf{x}y]$.
$f_{\text{LI}}(\mathbf{x})$	The best (population) linear projection of y onto \mathbf{x} : $\mathbf{x}^\top \boldsymbol{\beta}_0$.
$f_{\text{NL}}(\mathbf{x})$	The component of y that is not explained by \mathbf{x} : $y - \mathbf{x}^\top \boldsymbol{\beta}_0$.
$\mathbf{f}_{\text{LI}}, \mathbf{f}_{\text{NL}}$	$\mathbf{f}_{\text{LI}} = \mathbf{X}\boldsymbol{\beta}_0$, $\mathbf{f}_{\text{NL}} = [f_{\text{NL}}(\mathbf{x}_i)]_{i \in [n]}$.

A note on indexing of sequences:

In the subsequent sections, we will prove the results for n, k, p being a sequence of integers $\{n_m\}_{m=1}^\infty, \{k_m\}_{m=1}^\infty, \{p_m\}_{m=1}^\infty$. Alternatively, one can also view k and p as sequences k_n and p_n that are indexed by n . For notational brevity, we drop the subscript when it is clear from the context.

Organization. Below we outline the structure of the rest of the supplement. The technical lemmas refer to the main ingredients that we use to prove results in the corresponding sections.

Table 3: Outline of the supplement.

Section	Description	Technical lemmas
Appendix A	Proof of Theorem 1, Proposition 2 (from Section 3)	Lemma A.1, Lemma A.2, Lemma A.3
Appendix B	Proof of Theorem 3, Proposition 4, Corollary 5 (from Section 4)	Lemma B.1, Lemma B.2
Appendix C	Proof of Theorem 6 (from Section 5)	Lemma C.1
Appendix D	New asymptotic equivalents, concentration results and other useful lemmas (used in Appendices A–C)	Lemma D.1, Lemma D.2, Lemma D.3, Lemma D.5
Appendix E	Asymptotic equivalents: background and known results (used in Appendices A–D)	
Appendix F	Additional details for all the experiments	

A Proofs of results in Section 3

A.1 Proof of Theorem 1

Given an observation (\mathbf{x}, y) , recall the decomposition $y = f_{\text{LI}}(\mathbf{x}) + f_{\text{NL}}(\mathbf{x})$ explained in Section 2. For n i.i.d. samples from the same distribution as (\mathbf{x}, y) , we define analogously the vector decomposition:

$$\mathbf{y} = \mathbf{f}_{\text{LI}} + \mathbf{f}_{\text{NL}}, \quad (10)$$

where $\mathbf{f}_{\text{LI}} = \mathbf{X}\beta_0$ and $\mathbf{f}_{\text{NL}} = [f_{\text{NL}}(\mathbf{x}_i)]_{i \in [n]}$. Let $n_{\mathbf{A}} = \text{nrow}(\mathbf{A})$. Note that

$$R(\hat{\beta}_{k,\infty}^\lambda; \mathbf{A}, \mathbf{b}, \beta_0) = R(\hat{\beta}_{k,\infty}^\lambda; \mathbf{A}, \mathbf{0}, \beta_0) + 2n_{\mathbf{A}}^{-1} \mathbf{b}^\top \mathbf{A}(\hat{\beta}_{k,\infty}^\lambda - \beta_0) + n_{\mathbf{A}}^{-1} \|\mathbf{b}\|_2^2.$$

By Theorem 3, the cross term vanishes, i.e., $n_{\mathbf{A}}^{-1} \mathbf{b}^\top \mathbf{A}(\hat{\beta}_{k,\infty}^\lambda - \beta_0) \xrightarrow{\text{a.s.}} 0$. We then have

$$|R(\hat{\beta}_{k_1,\infty}^{\lambda_1}; \mathbf{A}, \mathbf{b}, \beta_0) - R(\hat{\beta}_{k_2,\infty}^{\lambda_2}; \mathbf{A}, \mathbf{b}, \beta_0)| \xrightarrow{\text{a.s.}} |R(\hat{\beta}_{k_1,\infty}^{\lambda_1}; \mathbf{A}, \mathbf{0}, \beta_0) - R(\hat{\beta}_{k_2,\infty}^{\lambda_2}; \mathbf{A}, \mathbf{0}, \beta_0)|.$$

It suffices to analyze $R(\hat{\beta}_{k,\infty}^\lambda; \mathbf{A}, \mathbf{0}, \beta_0)$.

For simplicity, we treat \mathbf{A} as the normalized matrix $n_{\mathbf{A}}^{-1/2} \mathbf{A}$ to avoid the notation of $\text{nrow}(\mathbf{A})$. Note that $\mathbf{y} = \mathbf{X}\beta_0 + \mathbf{f}_{\text{NL}} + \boldsymbol{\varepsilon}$, where β_0 is the best linear projection of \mathbf{y} on \mathbf{X} , \mathbf{f}_{NL} is the nonlinear residual and $\boldsymbol{\varepsilon}$ is the independent noise. Let $\mathbf{A}_k^\lambda = \mathbb{E}_{I \sim \mathcal{I}_k}[(\mathbf{X}^\top \mathbf{L}_I \mathbf{X} / k + \lambda \mathbf{I}_p)^{-1} \mathbf{X}^\top \mathbf{L}_I / k]$ and $\mathbf{B}_k^\lambda = \mathbf{I}_p - \mathbf{A}_k^\lambda \mathbf{X} = \mathbb{E}_{I \sim \mathcal{I}_k}[\lambda(\mathbf{X}^\top \mathbf{L}_I \mathbf{X} / k + \lambda \mathbf{I}_p)^{-1}]$. We begin by decomposing the generalized risk for arbitrary (λ, k) into different terms:

$$\begin{aligned} R(\hat{\beta}_{k,\infty}^\lambda; \mathbf{A}, \mathbf{0}, \beta_0) &= \|\mathbf{A}(\hat{\beta}_{k,\infty}^\lambda - \beta_0)\|_2^2 \\ &= \|\mathbf{A}(\mathbf{A}_k^\lambda \mathbf{y} - \beta_0)\|_2^2 \\ &= \|\mathbf{A} \mathbf{A}_k^\lambda (\mathbf{X}\beta_0 + \mathbf{f}_{\text{NL}}) - \mathbf{A}\beta_0\|_2^2 \\ &= \beta_0^\top (\mathbf{A}_k^\lambda \mathbf{X} - \mathbf{I}_p)^\top \mathbf{A}^\top \mathbf{A} (\mathbf{A}_k^\lambda \mathbf{X} - \mathbf{I}_p) \beta_0 + \mathbf{f}_{\text{NL}}^\top (\mathbf{A}_k^\lambda)^\top \mathbf{A}^\top \mathbf{A} \mathbf{A}_k^\lambda \mathbf{f}_{\text{NL}} \\ &\quad + 2 \mathbf{f}_{\text{NL}}^\top (\mathbf{A}_k^\lambda)^\top \mathbf{A}^\top \mathbf{A} (\mathbf{A}_k^\lambda \mathbf{X} - \mathbf{I}_p) \beta_0 \\ &= \beta_0^\top \mathbf{B}_k^\lambda \mathbf{A}^\top \mathbf{A} \mathbf{B}_k^\lambda \beta_0 + \mathbf{f}_{\text{NL}}^\top (\mathbf{A}_k^\lambda)^\top \mathbf{A}^\top \mathbf{A} \mathbf{A}_k^\lambda \mathbf{f}_{\text{NL}} - 2 \mathbf{f}_{\text{NL}}^\top (\mathbf{A}_k^\lambda)^\top \mathbf{A}^\top \mathbf{A} \mathbf{B}_k^\lambda \beta_0. \end{aligned}$$

When $\lambda > 0$, from Lemma A.3, we know that the cross terms $2\mathbf{f}_{\text{NL}}^\top (\mathbf{A}_k^\lambda)^\top \mathbf{A}^\top \mathbf{A} \mathbf{B}_k^\lambda \boldsymbol{\beta}_0$ vanishes as p tends to infinity. Further, by Lemma A.1 and Lemma A.2, it follows that when $\mathbf{A} \perp\!\!\!\perp (\mathbf{X}, \mathbf{y})$, as $p/n \rightarrow \phi$ and $p/k \rightarrow \psi$,

$$|R(\widehat{\boldsymbol{\beta}}_{k,\infty}^\lambda; \mathbf{A}, \mathbf{0}, \boldsymbol{\beta}_0) - R_p(\lambda; \phi, \psi)| \xrightarrow{\text{a.s.}} 0,$$

where the function R_p is defined as

$$R_p(\lambda; \phi, \psi) = \tilde{c}_p(-\lambda; \phi, \psi, \mathbf{A}^\top \mathbf{A}) + \|\mathbf{f}_{\text{NL}}\|_{L_2}^2 \tilde{v}_p(-\lambda; \phi, \psi, \mathbf{A}^\top \mathbf{A}), \quad (11)$$

and the nonnegative constants $\tilde{c}_p(-\lambda; \phi, \psi, \mathbf{A}^\top \mathbf{A})$ and $\tilde{v}_p(-\lambda; \phi, \psi, \mathbf{A}^\top \mathbf{A})$ are as defined in Lemma A.1.

When $\lambda = 0$, as in the proof of Theorem 3, we again show that $P_{n,\lambda} - Q_{n,\lambda}$ is equicontinuous over $\Lambda = [0, \lambda_{\max}]$ for any $\lambda_{\max} \in (0, \infty)$ fixed, where we define

$$P_{n,\lambda} = R(\widehat{\boldsymbol{\beta}}_{k,\infty}^\lambda; \mathbf{A}, \mathbf{0}, \boldsymbol{\beta}_0), \quad \text{and} \quad Q_{n,\lambda} = R_p(\lambda; \phi, \psi).$$

When $\psi \neq 1$, it can be verified that $|P_{n,\lambda}|$, $|\partial P_{n,\lambda}/\partial \lambda|$, $|Q_{n,\lambda}|$, $|\partial Q_{n,\lambda}/\partial \lambda|$ are bounded almost surely. Thus, by the Moore-Osgood theorem, it follows that, with probability one,

$$\lim_{n \rightarrow \infty} |R(\widehat{\boldsymbol{\beta}}_{k,\infty}^0; \mathbf{A}, \mathbf{0}, \boldsymbol{\beta}_0) - R_p(0; \phi, \psi)| = \lim_{\lambda \rightarrow 0^+} \lim_{n \rightarrow \infty} |R(\widehat{\boldsymbol{\beta}}_{k,\infty}^\lambda; \mathbf{A}, \mathbf{0}, \boldsymbol{\beta}_0) - R_p(\lambda; \phi, \psi)| = 0.$$

Note that $\tilde{c}_p(-\lambda; \phi, \psi, \mathbf{A}^\top \mathbf{A})$ and $\tilde{v}_p(-\lambda; \phi, \psi, \mathbf{A}^\top \mathbf{A})$ are functions of the fixed-point solution $v_p(-\lambda; \psi)$. By Lemma C.1 we have that for $\bar{\psi} \in [\phi, \infty]$ there exists a segment \mathcal{P} such that for all $(\lambda_1, \psi_1), (\lambda_2, \psi_2) \in \mathcal{P}$, it holds that $v(-\lambda_1; \psi_1) = v(-\lambda_2; \psi_2)$ as $p/k_1 \rightarrow \psi_1$ and $p/k_2 \rightarrow \psi_2$. This implies that

$$R_p(\lambda_1; \phi, \psi_1) = R_p(\lambda_2; \phi, \psi_2).$$

Thus, by triangle inequality, we have

$$|R(\widehat{\boldsymbol{\beta}}_{k_1,\infty}^{\lambda_1}; \mathbf{A}, \mathbf{0}, \boldsymbol{\beta}_0) - R(\widehat{\boldsymbol{\beta}}_{k_2,\infty}^{\lambda_2}; \mathbf{A}, \mathbf{0}, \boldsymbol{\beta}_0)| \xrightarrow{\text{a.s.}} 0,$$

which completes the proof.

A.2 Proof of Proposition 2

When $\mathbf{A} = \mathbf{X}$ and $\mathbf{b} = \mathbf{f}_{\text{NL}}$, the generalized risk reduces to

$$\begin{aligned} R(\widehat{\boldsymbol{\beta}}_{k,\infty}^\lambda; \mathbf{A}, \mathbf{b}, \boldsymbol{\beta}_0) &= \frac{1}{n} \|\mathbf{X}(\widehat{\boldsymbol{\beta}}_{k,\infty}^\lambda - \boldsymbol{\beta}_0) - \mathbf{f}_{\text{NL}}\|_2^2 \\ &= \frac{1}{n} \|\mathbf{X}\widehat{\boldsymbol{\beta}}_{k,\infty}^\lambda - (\mathbf{X}\boldsymbol{\beta}_0 + \mathbf{f}_{\text{NL}})\|_2^2 \\ &= \frac{1}{n} \|\mathbf{X}\mathbf{A}_k^\lambda(\mathbf{X}\boldsymbol{\beta}_0 + \mathbf{f}_{\text{NL}}) - (\mathbf{X}\boldsymbol{\beta}_0 + \mathbf{f}_{\text{NL}})\|_2^2 \\ &= \boldsymbol{\beta}_0^\top \mathbf{B}_k^\lambda \widehat{\boldsymbol{\Sigma}} \mathbf{B}_k^\lambda \boldsymbol{\beta}_0 + \frac{1}{n} \mathbf{f}_{\text{NL}}^\top (\mathbf{I}_p - \mathbf{X}\mathbf{A}_k^\lambda)^\top (\mathbf{I}_p - \mathbf{X}\mathbf{A}_k^\lambda) \mathbf{f}_{\text{NL}} \\ &\quad + \frac{2}{n} \mathbf{f}_{\text{NL}}^\top (\mathbf{I}_p - \mathbf{X}\mathbf{A}_k^\lambda)^\top \mathbf{X}\mathbf{B}_k^\lambda \boldsymbol{\beta}_0. \end{aligned}$$

The proof then follows analogously as in Theorem 1 by involving Lemma A.1 and Lemma A.2 with $\mathbf{A} = n^{-1/2} \mathbf{X}$.

When $\mathbf{A} = \mathbf{X}$ and $\mathbf{b} = \mathbf{0}$, we have

$$R(\widehat{\boldsymbol{\beta}}_{k,\infty}^\lambda; \mathbf{A}, \mathbf{b}, \boldsymbol{\beta}_0) = \boldsymbol{\beta}_0^\top \mathbf{B}_k^\lambda \widehat{\boldsymbol{\Sigma}} \mathbf{B}_k^\lambda \boldsymbol{\beta}_0 + \mathbf{f}_{\text{NL}}^\top (\mathbf{A}_k^\lambda)^\top \widehat{\boldsymbol{\Sigma}} \mathbf{A}_k^\lambda \mathbf{f}_{\text{NL}} - 2\mathbf{f}_{\text{NL}}^\top (\mathbf{A}_k^\lambda)^\top \widehat{\boldsymbol{\Sigma}} \mathbf{B}_k^\lambda \boldsymbol{\beta}_0.$$

Invoking Lemma A.2 for $\mathbf{f}_{\text{NL}}^\top (\mathbf{A}_k^\lambda)^\top \widehat{\boldsymbol{\Sigma}} \mathbf{A}_k^\lambda \mathbf{f}_{\text{NL}}$ instead of $\mathbf{f}_{\text{NL}}^\top (\mathbf{I}_p - \mathbf{X}\mathbf{A}_k^\lambda)^\top (\mathbf{I}_p - \mathbf{X}\mathbf{A}_k^\lambda) \mathbf{f}_{\text{NL}}$, the proof follows analogously.

A.3 Technical lemmas

Lemma A.1 (Bias of generalized risk). *Suppose the same assumptions in Theorem 1 hold with $\lambda > 0$. Let $\mathbf{B}_k^\lambda = \mathbf{I}_p - \mathbf{A}_k^\lambda \mathbf{X} = \mathbb{E}_{I \sim \mathcal{I}_k} [\lambda (\mathbf{X}^\top \mathbf{L}_I \mathbf{X} / k + \lambda \mathbf{I}_p)^{-1}]$. For any \mathbf{A} with $\limsup \|\mathbf{A}\|_{\text{op}}$ bounded almost surely, the following statements hold:*

(1) If $\mathbf{A} \perp (\mathbf{X}, \mathbf{y})$, then

$$\beta_0^\top \mathbf{B}_k^\lambda \mathbf{A}^\top \mathbf{A} \mathbf{B}_k^\lambda \beta_0 - \tilde{c}_p(-\lambda; \phi, \psi, \mathbf{A}^\top \mathbf{A}) \xrightarrow{\text{a.s.}} 0.$$

(2) If $\mathbf{A} = n^{-1/2} \mathbf{X}$, then

$$\beta_0^\top \mathbf{B}_k^\lambda \mathbf{A}^\top \mathbf{A} \mathbf{B}_k^\lambda \beta_0 - \left(1 - \phi \int \frac{v_p(-\lambda; \psi)r}{1 + v_p(-\lambda; \psi)r} dH_p(r)\right)^2 \tilde{c}_p(-\lambda; \phi, \psi, \Sigma) \xrightarrow{\text{a.s.}} 0.$$

Here the nonnegative constants $v_p(-\lambda; \psi)$, $\tilde{v}_p(-\lambda; \phi, \psi, \mathbf{A}^\top \mathbf{A})$ and $\tilde{c}_p(-\lambda; \phi, \psi, \mathbf{A}^\top \mathbf{A})$ are defined through the following equations:

$$\begin{aligned} \frac{1}{v_p(-\lambda; \psi)} &= \lambda + \psi \int \frac{r}{1 + v_p(-\lambda; \psi)r} dH_p(r), \\ \tilde{v}_p(-\lambda; \phi, \psi, \mathbf{A}^\top \mathbf{A}) &= \frac{\phi \text{tr}[\mathbf{A}^\top \mathbf{A} \Sigma (v_p(-\lambda; \psi) \Sigma + \mathbf{I}_p)^{-2}] / p}{v_p(-\lambda; \psi)^{-2} - \phi \int \frac{r^2}{(1 + v_p(-\lambda; \psi)r)^2} dH_p(r)}, \\ \tilde{c}_p(-\lambda; \phi, \psi, \mathbf{A}^\top \mathbf{A}) &= \beta_0^\top (v_p(-\lambda; \psi) \Sigma + \mathbf{I}_p)^{-1} (\tilde{v}_p(-\lambda; \phi, \psi, \mathbf{A}^\top \mathbf{A}) \Sigma + \mathbf{A}^\top \mathbf{A}) (v_p(-\lambda; \psi) \Sigma + \mathbf{I}_p)^{-1} \beta_0. \end{aligned}$$

Proof. Note that β_0 is independent of

$$\mathbf{B}_k^\lambda \mathbf{A}^\top \mathbf{A} \mathbf{B}_k^\lambda = \lambda^2 \mathbb{E}_{I_1, I_2 \sim \mathcal{I}_k} [\mathbf{M}_1 \mathbf{A}^\top \mathbf{A} \mathbf{M}_2],$$

where $\widehat{\Sigma}_{1 \cap 2} = \mathbf{X}^\top \mathbf{L}_{I_1 \cap I_2} \mathbf{X} / |I_1 \cap I_2|$ and $\mathbf{M}_j = (\mathbf{X}^\top \mathbf{L}_{I_j} \mathbf{X} / k + \lambda \mathbf{I}_p)^{-1}$ for $j = 1, 2$. We analyze the deterministic equivalents of the latter for the two cases.

(1) $\mathbf{A} \perp (\mathbf{X}, \mathbf{y})$.

From Lemma D.1 (2), we know that when \mathbf{A} is independent to (\mathbf{X}, \mathbf{y}) , it follows that

$$\begin{aligned} &\lambda^2 \mathbb{E}_{I_1, I_2 \sim \mathcal{I}_k} [\mathbf{M}_1 \mathbf{A}^\top \mathbf{A} \mathbf{M}_2] \\ &\simeq (v_p(-\lambda; \psi) \Sigma + \mathbf{I}_p)^{-1} (\tilde{v}_p(-\lambda; \phi, \psi, \mathbf{A}^\top \mathbf{A}) \Sigma + \mathbf{A}^\top \mathbf{A}) (v_p(-\lambda; \psi) \Sigma + \mathbf{I}_p)^{-1}. \end{aligned}$$

Then, by the trace property of deterministic equivalents in Lemma E.3 (4), we have

$$\begin{aligned} &\beta_0^\top \mathbf{B}_k^\lambda \mathbf{A}^\top \mathbf{A} \mathbf{B}_k^\lambda \beta_0 \\ &\stackrel{\text{a.s.}}{=} \beta_0^\top (v_p(-\lambda; \psi) \Sigma + \mathbf{I}_p)^{-1} (\tilde{v}_p(-\lambda; \phi, \psi, \mathbf{A}^\top \mathbf{A}) \Sigma + \mathbf{A}^\top \mathbf{A}) (v_p(-\lambda; \psi) \Sigma + \mathbf{I}_p)^{-1} \beta_0 \\ &\stackrel{\text{a.s.}}{=} \tilde{c}_p(-\lambda; \phi, \psi, \mathbf{A}^\top \mathbf{A}). \end{aligned}$$

(2) $\mathbf{A} = n^{-1/2} \mathbf{X}$.

From Lemma D.1 (4), it follows that

$$\begin{aligned} &\lambda^2 \mathbb{E}_{I_1, I_2 \sim \mathcal{I}_k} [\mathbf{M}_1 \widehat{\Sigma} \mathbf{M}_2] \simeq \left(1 - \phi \int \frac{v_p(-\lambda; \psi)r}{1 + v_p(-\lambda; \psi)r} dH_p(r)\right) \\ &\quad \cdot (1 + \tilde{v}_p(-\lambda; \phi, \psi)) (v_p(-\lambda; \psi) \Sigma + \mathbf{I}_p)^{-2} \Sigma. \end{aligned}$$

Then, by the trace property of deterministic equivalents in Lemma E.3 (4), we have

$$\begin{aligned} &\beta_0^\top \mathbf{B}_k^\lambda \widehat{\Sigma} \mathbf{B}_k^\lambda \beta_0 \stackrel{\text{a.s.}}{=} \left(1 - \phi \int \frac{v_p(-\lambda; \psi)r}{1 + v_p(-\lambda; \psi)r} dH_p(r)\right) (1 + \tilde{v}_p(-\lambda; \phi, \psi)) \cdot \\ &\quad \beta_0^\top (v_p(-\lambda; \psi) \Sigma + \mathbf{I}_p)^{-1} \Sigma (v_p(-\lambda; \psi) \Sigma + \mathbf{I}_p)^{-1} \beta_0 \\ &\stackrel{\text{a.s.}}{=} \left(1 - \phi \int \frac{v_p(-\lambda; \psi)r}{1 + v_p(-\lambda; \psi)r} dH_p(r)\right)^2 \tilde{c}_p(-\lambda; \phi, \psi, \Sigma). \end{aligned}$$

□

Lemma A.2 (Variance term of generalized risk). *Suppose the same assumptions in Theorem 1 hold with $\lambda > 0$. Let $\mathbf{A}_k^\lambda = \mathbb{E}_{I \sim \mathcal{I}_k}[(\mathbf{X}^\top \mathbf{L}_I \mathbf{X}/k + \lambda \mathbf{I}_p)^{-1} \mathbf{X}^\top \mathbf{L}_I/k]$ and $\mathbf{B}_k^\lambda = \mathbf{I}_p - \mathbf{A}_k^\lambda \mathbf{X} = \mathbb{E}_{I \sim \mathcal{I}_k}[\lambda(\mathbf{X}^\top \mathbf{L}_I \mathbf{X}/k + \lambda \mathbf{I}_p)^{-1}]$. For any \mathbf{A} with $\limsup \|\mathbf{A}\|_{\text{op}}$ bounded almost surely, the following statements hold:*

(1) *If $\mathbf{A} \perp (\mathbf{X}, \mathbf{y})$, then*

$$\mathbf{f}_{\text{NL}}^\top (\mathbf{A}_k^\lambda)^\top \mathbf{A}^\top \mathbf{A} \mathbf{A}_k^\lambda \mathbf{f}_{\text{NL}} - \|\mathbf{f}_{\text{NL}}\|_{L^2}^2 (1 + \tilde{v}_p(-\lambda; \phi, \psi, \boldsymbol{\Sigma})) \xrightarrow{\text{a.s.}} 0.$$

(2) *If $\mathbf{A} = n^{-1/2} \mathbf{X}$, then*

$$\begin{aligned} & \mathbf{f}_{\text{NL}}^\top (\mathbf{A}_k^\lambda)^\top \mathbf{A}^\top \mathbf{A} \mathbf{A}_k^\lambda \mathbf{f}_{\text{NL}} \\ & \xrightarrow{\text{a.s.}} \|\mathbf{f}_{\text{NL}}\|_{L^2}^2 \left(1 - \phi \int \frac{v_p(-\lambda; \psi)r}{1 + v_p(-\lambda; \psi)r} dH_p(r)\right)^2 \cdot (1 + \tilde{v}_p(-\lambda; \phi, \psi, \boldsymbol{\Sigma})) \\ & \quad + \|\mathbf{f}_{\text{NL}}\|_{L^2}^2 \left(2\phi \int \frac{v_p(-\lambda; \psi)r}{1 + v_p(-\lambda; \psi)r} dH_p(r) - 1\right). \\ & \quad \frac{1}{n} \mathbf{f}_{\text{NL}}^\top (\mathbf{X} \mathbf{A}_k^\lambda - \mathbf{I}_n)^\top (\mathbf{X} \mathbf{A}_k^\lambda - \mathbf{I}_n) \mathbf{f}_{\text{NL}} \\ & \xrightarrow{\text{a.s.}} \|\mathbf{f}_{\text{NL}}\|_{L^2}^2 \left(1 - \phi \int \frac{v_p(-\lambda; \psi)r}{1 + v_p(-\lambda; \psi)r} dH_p(r)\right)^2 \cdot \sigma^2 (1 + \tilde{v}_p(-\lambda; \phi, \psi, \boldsymbol{\Sigma})), \end{aligned}$$

Here the nonnegative constants $v_p(-\lambda; \psi)$ and $\tilde{v}_p(-\lambda; \phi, \psi, \mathbf{A}^\top \mathbf{A})$ are defined through the following equations:

$$\begin{aligned} \frac{1}{v_p(-\lambda; \psi)} &= \lambda + \psi \int \frac{r}{1 + v_p(-\lambda; \psi)r} dH_p(r), \\ \tilde{v}_p(-\lambda; \phi, \psi, \mathbf{A}^\top \mathbf{A}) &= \frac{\phi \text{tr}[\mathbf{A}^\top \mathbf{A} \boldsymbol{\Sigma} (v_p(-\lambda; \psi) \boldsymbol{\Sigma} + \mathbf{I}_p)^{-2}]/p}{v_p(-\lambda; \psi)^{-2} - \phi \int \frac{r^2}{(1 + v_p(-\lambda; \psi)r)^2} dH_p(r)}. \end{aligned}$$

Proof. The two cases are treated separately below.

(1) $\mathbf{A} \perp (\mathbf{X}, \mathbf{y})$.

We split the proof into three different parts.

Part (1) Intersection concentration. Let $\mathbf{f}_{\text{NL}} = \mathbf{L}_{I_1 \cap I_2} \mathbf{f}_{\text{NL}} + \mathbf{L}_{I_1 \setminus I_2} \mathbf{f}_{\text{NL}} + \mathbf{L}_{I_2 \setminus I_1} \mathbf{f}_{\text{NL}} =: \mathbf{f}_0 + \mathbf{f}_1 + \mathbf{f}_2$. Note that

$$\mathbf{f}_{\text{NL}}^\top \mathbf{A}_k^{\lambda_1 \top} \mathbf{A}^\top \mathbf{A} \mathbf{A}_k^{\lambda_1} \mathbf{f}_{\text{NL}} = \mathbb{E}_{I_1, I_2 \sim \mathcal{I}_k} \left[(\mathbf{f}_0 + \mathbf{f}_1)^\top \frac{\mathbf{X}}{k} \mathbf{M}_1 \mathbf{A}^\top \mathbf{A} \mathbf{M}_2 \frac{\mathbf{X}^\top}{k} (\mathbf{f}_0 + \mathbf{f}_2) \right].$$

where $\mathbf{M}_j = (\mathbf{X}^\top \mathbf{L}_{I_j} \mathbf{X}/k + \lambda \mathbf{I}_p)^{-1}$. By conditional independence and Lemma S.8.5 of [51], the cross terms vanish

$$\mathbf{f}_1^\top \frac{\mathbf{X}}{k} \mathbf{M}_1 \mathbf{A}^\top \mathbf{A} \mathbf{M}_2 \frac{\mathbf{X}^\top}{k} \mathbf{L}_2 \mathbf{f}_{\text{NL}} \xrightarrow{\text{a.s.}} 0, \quad \text{and} \quad \mathbf{f}_{\text{NL}}^\top \frac{\mathbf{L}_1 \mathbf{X}}{k} \mathbf{M}_1 \mathbf{A}^\top \mathbf{A} \mathbf{M}_2 \frac{\mathbf{X}^\top}{k} \mathbf{f}_2 \xrightarrow{\text{a.s.}} 0. \tag{12}$$

It remains to analyze the quadratic term of \mathbf{f}_0 :

$$\mathbf{f}_0^\top \frac{\mathbf{L}_{I_1 \cap I_2} \mathbf{X}}{k} \mathbf{M}_1 \mathbf{A}^\top \mathbf{A} \mathbf{M}_2 \frac{\mathbf{X}^\top \mathbf{L}_{I_1 \cap I_2}}{k} \mathbf{f}_0.$$

By conditioning on $\mathbf{L}_{I_1 \cap I_2} \mathbf{X}$, from Lemma S.7.10 (1) of [12], we have

$$\mathbf{M}_j \simeq \mathbf{M}^{\text{det}} := \frac{k}{i_0} (\widehat{\boldsymbol{\Sigma}}_0 + \lambda \mathbf{I}_p + \lambda \mathbf{C})^{-1},$$

where $\widehat{\Sigma}_0 = \mathbf{X}_0^\top \mathbf{X}_0 / i_0$, $\mathbf{X}_0 = \mathbf{L}_{I_1 \cap I_2} \mathbf{X}$, $\mathbf{C} = (k - i_0) / i_0 \cdot (v(-\lambda; \gamma_1, \Sigma_{C_1}) \Sigma + \mathbf{I}_p)$, $C_1 = i_0 (\lambda(k - i_0))^{-1} (\widehat{\Sigma}_0 + \lambda \mathbf{I}_p)$ and $i_0 = |I_1 \cap I_2|$. Then we have

$$\begin{aligned} \frac{\mathbf{X}_0}{k} \mathbf{M}_1 \mathbf{A}^\top \mathbf{A} \mathbf{M}_2 \frac{\mathbf{X}_0^\top}{k} &\simeq \frac{k^2}{i_0^2} \frac{\mathbf{X}_0}{k} \mathbf{M}^{\det} \mathbf{A}^\top \mathbf{A} \mathbf{M}^{\det} \frac{\mathbf{X}_0^\top}{k} \\ &= \frac{k^2}{i_0^2} \frac{\mathbf{X}'_0}{k} \mathbf{M}' \mathbf{A}' \mathbf{M}' \frac{\mathbf{X}'_0{}^\top}{k} \end{aligned} \quad (13)$$

where $\mathbf{X}'_0 = \mathbf{X}_0 (\mathbf{I}_p + \mathbf{C})^{-1/2}$, $\mathbf{A}' = (\mathbf{I}_p + \mathbf{C})^{-1/2} \mathbf{A}^\top \mathbf{A} (\mathbf{I}_p + \mathbf{C})^{-1/2}$, and $\mathbf{M}' = ((\mathbf{I}_p + \mathbf{C})^{-1/2} \widehat{\Sigma}_0 (\mathbf{I}_p + \mathbf{C})^{-1/2} + \lambda \mathbf{I}_p)^{-1}$.

Part (2) Diagonal concentration. We next use a similar strategy as in the proof of Lemma A.16 in [14] by showing that the off-diagonal summation vanishes. Let $\Sigma' = (\mathbf{I}_p + \mathbf{C})^{-1/2} \Sigma (\mathbf{I}_p + \mathbf{C})^{-1/2}$. Since $\mathbf{X}_0 = \mathbf{Z}_0 \Sigma^{1/2}$, we have $\mathbf{X}'_0 = \mathbf{Z}_0 \Sigma'^{1/2}$. Then, the quadratic form becomes:

$$\begin{aligned} &\frac{1}{k^2} \mathbf{f}_0^\top \mathbf{X}'_0 \mathbf{M}' \mathbf{A}' \mathbf{M}' \mathbf{X}'_0{}^\top \mathbf{f}_0 \\ &= \frac{i_0}{k^2} \mathbf{f}_0^\top \left(\frac{\mathbf{Z}_0 \Sigma' \mathbf{Z}_0^\top}{i_0} + \lambda \mathbf{I}_n \right)^{-1} \frac{\mathbf{Z}_0}{\sqrt{i_0}} \Sigma'^{1/2} \mathbf{A}' \Sigma'^{1/2} \frac{\mathbf{Z}_0^\top}{\sqrt{i_0}} \left(\frac{\mathbf{Z}_0 \Sigma' \mathbf{Z}_0^\top}{i_0} + \lambda \mathbf{I}_n \right)^{-1} \mathbf{f}_0 \\ &=: \frac{i_0}{k^2} \mathbf{f}_0^\top \mathbf{B}_1^{-1} \mathbf{B}_2 \mathbf{B}_1^{-1} \mathbf{f}_0. \end{aligned} \quad (14)$$

Note that from Lemma D.4, we have for any $t > 0$,

$$\mathbf{B}_1^{-1} \mathbf{B}_2 \mathbf{B}_1^{-1} = \frac{1}{t} (\mathbf{B}_1^{-1} - (\mathbf{B}_1 + t \mathbf{B}_2)^{-1}) + t \mathbf{B}_1^{-1} \mathbf{B}_2 (\mathbf{B}_1 + t \mathbf{B}_2)^{-1} \mathbf{B}_2 \mathbf{B}_1^{-1}.$$

Let $\mathbf{U} \in \mathbb{R}^{n \times n}$ with $U_{ij} = [\mathbf{f}_0]_i [\mathbf{f}_0]_j \mathbb{1}\{i \neq j\}$. We then have

$$\begin{aligned} &\left| \sum_{1 \leq i \neq j \leq n} [\mathbf{B}_1^{-1} \mathbf{B}_2 \mathbf{B}_1^{-1}]_{ij} [\mathbf{f}_0]_i [\mathbf{f}_0]_j \right| \\ &= |\langle \mathbf{B}_1^{-1} \mathbf{B}_2 \mathbf{B}_1^{-1}, \mathbf{U} \rangle| \\ &\leq \frac{1}{t} |\langle \mathbf{B}_1^{-1}, \mathbf{U} \rangle| - \frac{1}{t} |\langle (\mathbf{B}_1 + t \mathbf{B}_2)^{-1}, \mathbf{U} \rangle| + t \|\mathbf{B}_1^{-1}\|_{\text{op}}^2 \|\mathbf{B}_2\|_{\text{op}}^2 \|(\mathbf{B}_1 + t \mathbf{B}_2)^{-1}\|_{\text{op}} \|\mathbf{U}\|_{\text{tr}}. \end{aligned} \quad (15)$$

For the first two terms, Lemma D.3 implies that

$$\frac{1}{k} |\langle \mathbf{B}_1^{-1}, \mathbf{U} \rangle| \xrightarrow{\text{a.s.}} 0, \quad \text{and} \quad \frac{1}{k} |\langle (\mathbf{B}_1 + t \mathbf{B}_2)^{-1}, \mathbf{U} \rangle| \xrightarrow{\text{a.s.}} 0.$$

For the last term, note that $\|\mathbf{B}_2\|_{\text{op}} \leq \|\mathbf{A}\|_{\text{op}}^2 \|\widehat{\Sigma}_0\|_{\text{op}}$, where $\|\mathbf{A}\|_{\text{op}}$ is almost surely bounded as assumed, and $\|\widehat{\Sigma}_0\|_{\text{op}} \leq r_{\max} (1 + \sqrt{\psi^2 / \phi})^2$ almost surely as $k, n, p \rightarrow \infty$ and $p/n \rightarrow \phi \in (0, \infty)$, $p/k \rightarrow \psi \in [\phi, \infty]$ (see, e.g., [45]). Also, $\|\mathbf{U}\|_* \leq 2 \|\mathbf{f}_{\text{NL}}\|_2^2 \xrightarrow{\text{a.s.}} 2 \|f_{\text{NL}}\|_{L_2}^2 < \infty$ from the strong law of large numbers, Lemma D.5, and $\|\mathbf{B}_1\|_{\text{op}} \leq \lambda^{-1}$. Thus, the last term is almost surely bounded. It then follows that $k^{-1} |\langle \mathbf{B}_1^{-1} \mathbf{B}_2 \mathbf{B}_1^{-1}, \mathbf{U} \rangle| \xrightarrow{\text{a.s.}} 0$. Therefore,

$$\left| \frac{1}{k} \mathbf{f}_0^\top \mathbf{B}_1^{-1} \mathbf{B}_2 \mathbf{B}_1^{-1} \mathbf{f}_0 - \frac{1}{k} \sum_{i=1}^n [\mathbf{B}_1^{-1} \mathbf{B}_2 \mathbf{B}_1^{-1}]_{ii} [\mathbf{f}_0]_i^2 \right| \xrightarrow{\text{a.s.}} 0.$$

Part (3) Trace concentration. From the results in [56], it holds that

$$\max_{1 \leq i \leq n} \left| [\mathbf{B}_1^{-1} \mathbf{B}_2 \mathbf{B}_1^{-1}]_{ii} - \frac{1}{n} \text{tr}[\mathbf{B}_1^{-1} \mathbf{B}_2 \mathbf{B}_1^{-1}] \right| \xrightarrow{\text{a.s.}} 0.$$

Further, $n^{-1} \sum_{i=1}^n \mathbf{f}_i^2 \xrightarrow{\text{a.s.}} \|f_{\text{NL}}\|_{L_2}^2$ by strong law of large number and Lemma D.5. Therefore, we have

$$\frac{1}{k} |\mathbf{f}^\top \mathbf{B}_1^{-1} \mathbf{B}_2 \mathbf{B}_1^{-1} \mathbf{f} - \text{tr}[\mathbf{B}_1^{-1} \mathbf{B}_2 \mathbf{B}_1^{-1}] \|f_{\text{NL}}\|_{L_2}^2| \xrightarrow{\text{a.s.}} 0. \quad (16)$$

Combining (12), (13), and (16) yields that

$$|\mathbf{f}^\top (\mathbf{A}_k^\lambda)^\top \mathbf{A}^\top \mathbf{A} \mathbf{A}_k^\lambda \mathbf{f} - \text{tr}[(\mathbf{A}_k^\lambda)^\top \mathbf{A}^\top \mathbf{A} \mathbf{A}_k^\lambda]| \cdot \|f_{\text{NL}}\|_{L^2}^2 \xrightarrow{\text{a.s.}} 0.$$

By the trace property $\text{tr}[(\mathbf{A}_k^\lambda)^\top \mathbf{A}^\top \mathbf{A} \mathbf{A}_k^\lambda] = \text{tr}[\mathbf{A}_k^\lambda (\mathbf{A}_k^\lambda)^\top \mathbf{A}^\top \mathbf{A}]$, it remains to derive the deterministic equivalents of $\mathbf{A}_k^\lambda (\mathbf{A}_k^\lambda)^\top \mathbf{A}^\top \mathbf{A}$. Note that

$$\mathbf{A}_k^\lambda (\mathbf{A}_k^\lambda)^\top \mathbf{A}^\top \mathbf{A} = \frac{|I_1 \cap I_2|}{k^2} \mathbb{E}_{I_1, I_2 \sim \mathcal{I}_k} [\mathbf{M}_1 \widehat{\Sigma}_{1 \cap 2} \mathbf{M}_2] \mathbf{A}^\top \mathbf{A}$$

where $\widehat{\Sigma}_{1 \cap 2} = \mathbf{X}^\top \mathbf{L}_{I_1 \cap I_2} \mathbf{X} / |I_1 \cap I_2|$ and $\mathbf{M}_j = (\mathbf{X}^\top \mathbf{L}_{I_j} \mathbf{X} / k + \lambda \mathbf{I}_p)^{-1}$ for $j = 1, 2$. The above quantity is well-defined almost surely because $|I_1 \cap I_2|$ converges to some positive quantity almost surely. Next, we analyze the trace term for the two cases.

From Lemma D.1 (3) we know that when \mathbf{A} is independent to (\mathbf{X}, \mathbf{y}) , it follows that

$$\mathbb{E}_{I_1, I_2 \sim \mathcal{I}_k} [\mathbf{M}_1 \widehat{\Sigma}_{1 \cap 2} \mathbf{M}_2] \mathbf{A}^\top \mathbf{A} \simeq \phi^{-1} \tilde{v}_v(-\lambda; \phi, \psi) (v_p(-\lambda; \psi) \Sigma + \mathbf{I}_p)^{-2} \Sigma \mathbf{A}^\top \mathbf{A}.$$

We now have

$$\begin{aligned} \mathbf{f}_{\text{NL}}^\top (\mathbf{A}_k^\lambda)^\top \mathbf{A}^\top \mathbf{A} \mathbf{A}_k^\lambda \mathbf{f}_{\text{NL}} &\stackrel{\text{a.s.}}{=} \|f_{\text{NL}}\|_{L^2}^2 \frac{p}{k} \cdot \frac{|I_1 \cap I_2|}{k} \cdot \frac{1}{p} \text{tr}[\mathbb{E}_{I_1, I_2 \sim \mathcal{I}_k} [\mathbf{M}_1 \widehat{\Sigma}_{1 \cap 2} \mathbf{M}_2] \mathbf{A}^\top \mathbf{A}] \\ &\stackrel{\text{a.s.}}{=} \|f_{\text{NL}}\|_{L^2}^2 \psi \cdot \frac{\phi}{\psi} \cdot \frac{1}{\phi} \tilde{v}_v(-\lambda; \phi, \psi, \mathbf{A}^\top \mathbf{A}) \\ &= \|f_{\text{NL}}\|_{L^2}^2 \tilde{v}_v(-\lambda; \phi, \psi, \mathbf{A}^\top \mathbf{A}), \end{aligned}$$

where the second convergence is from Lemma S.8.3 of [12] and the trace property in Lemma E.3 (4).

(2) $\mathbf{A} = n^{-1/2} \mathbf{X}$.

Instead of working on (13), we use the following decomposition:

$$\begin{aligned} \mathbf{M}_1 \widehat{\Sigma} \mathbf{M}_2 &= \mathbf{M}_1 \widehat{\Sigma}_1 \mathbf{M}_2 + \widehat{\Sigma}_2 \mathbf{M}_2 - \mathbf{M}_1 \widehat{\Sigma}_{1 \cap 2} \mathbf{M}_2 + \mathbf{M}_1 \widehat{\Sigma}_{(1 \cup 2)^c} \mathbf{M}_2 \\ &= \sum_{j=1}^2 \mathbf{M}_j - 2\lambda \mathbf{M}_1 \mathbf{M}_2 - \mathbf{M}_1 \widehat{\Sigma}_{1 \cap 2} \mathbf{M}_2 + \mathbf{M}_1 \widehat{\Sigma}_{(1 \cup 2)^c} \mathbf{M}_2 \end{aligned}$$

Then, repeating Part (2) and (3) as above, it suffices to derive the deterministic equivalents of $(\mathbf{A}_k^\lambda)^\top \widehat{\Sigma} \mathbf{A}_k^\lambda$. We decompose this term into:

$$(\mathbf{A}_k^\lambda)^\top \mathbf{A}^\top \mathbf{A} \mathbf{A}_k^\lambda = \frac{1}{n} (\mathbf{X} \mathbf{A}_k^\lambda - \mathbf{I}_n)^\top (\mathbf{X} \mathbf{A}_k^\lambda - \mathbf{I}_n) + \frac{1}{n} (\mathbf{X} \mathbf{A}_k^\lambda + \mathbf{X}^\top (\mathbf{A}_k^\lambda)^\top) - \frac{1}{n} \mathbf{I}_n. \quad (17)$$

For the last two terms of the right hand side, following the similar argument as in Part (1), it holds that

$$\frac{2}{n} \text{tr}[\mathbf{X} \mathbf{A}_k^\lambda] - 2\phi \int \frac{v_p(-\lambda; \psi)r}{1 + v_p(-\lambda; \psi)r} dH_p(r) \xrightarrow{\text{a.s.}} 0, \quad \text{and} \quad \frac{1}{n} \text{tr}[\mathbf{I}_n] = 1. \quad (18)$$

For the first term of the right hand side, notice that it is the variance term of the mean squared error computed on all samples, which is the variance term of the numerator of the generalized cross-validation (GCV) estimator in [13]. It has been shown in Proposition 3.6 of [13] that in the full ensemble, the GCV estimator is consistent with the prediction risk, which is also true for the variance term:

$$\frac{1}{n} \text{tr}[(\mathbf{X} \mathbf{A}_k^\lambda - \mathbf{I}_n)^\top (\mathbf{X} \mathbf{A}_k^\lambda - \mathbf{I}_n)] \quad (19)$$

$$\xrightarrow{\text{a.s.}} \left(1 - \phi \int \frac{v_p(-\lambda; \psi)r}{1 + v_p(-\lambda; \psi)r} dH_p(r)\right)^2 \cdot \sigma^2 (1 + \tilde{v}_v(-\lambda; \phi, \psi, \Sigma)), \quad (20)$$

Finally, combining (14)-(16) with the above finishes the proof for the first convergence result of $\mathbf{A} = n^{-1/2} \mathbf{X}$.

Analogously, from (20), we also have

$$\begin{aligned} & \frac{1}{n} \mathbf{f}_{\text{NL}}^\top (\mathbf{X} \mathbf{A}_k^\lambda - \mathbf{I}_n)^\top (\mathbf{X} \mathbf{A}_k^\lambda - \mathbf{I}_n) \mathbf{f}_{\text{NL}} \\ & \xrightarrow{\text{a.s.}} \| \mathbf{f}_{\text{NL}} \|_{L^2}^2 \left(1 - \phi \int \frac{v_p(-\lambda; \psi) r}{1 + v_p(-\lambda; \psi) r} dH_p(r) \right)^2 \cdot \sigma^2 (1 + \tilde{v}_p(-\lambda; \phi, \psi, \boldsymbol{\Sigma})), \end{aligned}$$

which finishes the proof for the second convergence result of $\mathbf{A} = n^{-1/2} \mathbf{X}$.

□

Lemma A.3 (Cross terms of generalized risk). *Suppose the same assumptions in Theorem 1 hold with $\lambda > 0$. Let $\mathbf{A}_k^\lambda = \mathbb{E}_{I \sim \mathcal{I}_k} [(\mathbf{X}^\top \mathbf{L}_I \mathbf{X} / k + \lambda \mathbf{I}_p)^{-1} \mathbf{X}^\top \mathbf{L}_I / k]$ and $\mathbf{B}_k^\lambda = \mathbf{I}_p - \mathbf{A}_k^\lambda \mathbf{X} = \mathbb{E}_{I \sim \mathcal{I}_k} [\lambda (\mathbf{X}^\top \mathbf{L}_I \mathbf{X} / k + \lambda \mathbf{I}_p)^{-1}]$. For any \mathbf{A} with $\limsup \| \mathbf{A} \|_{\text{op}}$ bounded almost surely, it holds that*

$$\mathbf{f}_{\text{NL}}^\top (\mathbf{A}_k^\lambda)^\top \mathbf{A}^\top \mathbf{A} \mathbf{B}_k^\lambda \boldsymbol{\beta}_0 \xrightarrow{\text{a.s.}} 0.$$

Proof. Note that

$$\begin{aligned} n \| (\mathbf{A}_k^\lambda)^\top \mathbf{A}^\top \mathbf{A} \mathbf{B}_k^\lambda \boldsymbol{\beta}_0 \|_2^2 & \leq \frac{n}{k} \| \mathbf{A}_k^\lambda \|_{\text{op}}^2 \| \mathbf{A} \|_{\text{op}}^4 \| \mathbf{B}_k^\lambda \|_{\text{op}}^2 \| \boldsymbol{\beta}_0 \|_2^2 \\ & \leq \frac{n}{k} \mathbb{E}_{I \sim \mathcal{I}_k} [\| \widehat{\boldsymbol{\Sigma}}_I \|_{\text{op}}^{\frac{1}{2}} \| \lambda (\mathbf{X}^\top \mathbf{L}_I \mathbf{X} / k + \lambda \mathbf{I}_p)^{-1} \|_{\text{op}}]^2 \cdot \| \mathbf{B}_k^\lambda \|_{\text{op}}^2 \| \mathbf{A} \|_{\text{op}}^4 \| \boldsymbol{\beta}_0 \|_2^2 \\ & \leq \frac{n}{k} \mathbb{E}_{I \sim \mathcal{I}_k} [\| \widehat{\boldsymbol{\Sigma}}_I \|_{\text{op}}^{\frac{1}{2}}]^2 \cdot \| \mathbf{A} \|_{\text{op}}^4 \| \boldsymbol{\beta}_0 \|_2^2. \end{aligned}$$

Let s_j^2 be the singular value of $\widehat{\boldsymbol{\Sigma}}_j$. From the results in [45], we have $\limsup \| \widehat{\boldsymbol{\Sigma}}_I \|_{\text{op}} \leq \limsup \max_{1 \leq i \leq p} s_i^2 \leq r_{\max} (1 + \sqrt{\psi})^2$ almost surely as $k, p \rightarrow \infty$ and $p/k \rightarrow \psi \in (0, \infty]$. On the other hand, $n/k \xrightarrow{\text{a.s.}} \psi/\phi$ and $\| \mathbf{A} \|_2$ is uniformly bounded as assumed. Thus, we have $\limsup n \| (\mathbf{A}_k^\lambda)^\top \mathbf{A}^\top \mathbf{A} \mathbf{B}_k^\lambda \boldsymbol{\beta}_0 \|_2^2$ is bounded almost surely.

Similar to the proof of Theorem 1, we decompose \mathbf{f}_{NL} as $\mathbf{f}_{\text{NL}} = \mathbf{L}_{I_1 \cap I_2} \mathbf{f}_{\text{NL}} + \mathbf{L}_{I_1 \setminus I_2} \mathbf{f}_{\text{NL}} + \mathbf{L}_{I_2 \setminus I_1} \mathbf{f}_{\text{NL}} =: \mathbf{f}_0 + \mathbf{f}_1 + \mathbf{f}_2$. We then have

$$\begin{aligned} & \mathbf{f}_{\text{NL}}^\top \frac{\mathbf{L}_{I_1 \cap I_2} \mathbf{X}}{k} \mathbf{M}_1 \mathbf{A}^\top \mathbf{A} \mathbf{M}_2 \frac{\mathbf{X}^\top \mathbf{L}_{I_1 \cap I_2} \mathbf{X}}{k} \boldsymbol{\beta}_0 \\ & \xrightarrow{\text{a.s.}} \mathbf{f}_0^\top \frac{\mathbf{L}_{I_1 \cap I_2} \mathbf{X}}{k} \mathbf{M}_1 \mathbf{A}^\top \mathbf{A} \mathbf{M}_2 \frac{\mathbf{X}^\top \mathbf{L}_{I_1 \cap I_2} \mathbf{X}}{k} \boldsymbol{\beta}_0 \\ & \xrightarrow{\text{a.s.}} \frac{i_0}{k} \mathbf{f}_0^\top \mathbf{L}_{I_1 \cap I_2} \frac{k^2}{i_0^2} \frac{\mathbf{X}'_0}{k} \mathbf{M}' \mathbf{A}' \mathbf{M}' \widehat{\boldsymbol{\Sigma}}'_0 \boldsymbol{\beta}_0 \\ & = \frac{i_0}{k} \mathbf{f}_0^\top \mathbf{L}_{I_1 \cap I_2} \frac{k^2}{i_0^2} \frac{\mathbf{X}'_0}{k} \mathbf{M}' \mathbf{A}' \boldsymbol{\beta}_0 - \lambda \frac{i_0}{k} \mathbf{f}_0^\top \mathbf{L}_{I_1 \cap I_2} \frac{k^2}{i_0^2} \frac{\mathbf{X}'_0}{k} \mathbf{M}' \mathbf{A}' \mathbf{M}' \boldsymbol{\beta}_0, \end{aligned}$$

where the first convergence is from Lemma S.8.5 of [51] and the second convergence is from (13), with $\mathbf{X}'_0 = \mathbf{X}_0 (\mathbf{I}_p + \mathbf{C})^{-1/2}$, $\mathbf{A}' = (\mathbf{I}_p + \mathbf{C})^{-1/2} \mathbf{A}^\top \mathbf{A} (\mathbf{I}_p + \mathbf{C})^{-1/2}$, and $\mathbf{M}' = (\widehat{\boldsymbol{\Sigma}}'_0 + \lambda \mathbf{I}_p)^{-1}$ and $\widehat{\boldsymbol{\Sigma}}'_0 = (\mathbf{I}_p + \mathbf{C})^{-1/2} \widehat{\boldsymbol{\Sigma}}_0 (\mathbf{I}_p + \mathbf{C})^{-1/2}$. From Lemma D.2, the first term in the above display vanishes. Analogous to (15), the second term also vanishes by splitting and applying Lemma D.2. Thus, we have for $I_1, I_2 \sim \mathcal{I}_k$,

$$\mathbf{f}_{\text{NL}}^\top \frac{\mathbf{L}_{I_1 \cap I_2} \mathbf{X}}{k} \mathbf{M}_1 \mathbf{A}^\top \mathbf{A} \mathbf{M}_2 \frac{\mathbf{X}^\top \mathbf{L}_{I_1 \cap I_2} \mathbf{X}}{k} \boldsymbol{\beta}_0 \xrightarrow{\text{a.s.}} 0.$$

Finally, by Lemma G.5 (2) of [13], the conclusion follows. □

B Proofs of results in Section 4

B.1 Proof of Theorem 3

We consider two cases.

(1) $\lambda_1, \lambda_2 > 0$.

We begin to prove for the full ensemble when $M = \infty$. For $j = 1, 2$ and $I_j \sim \mathcal{I}_{k_j}$, define $\widehat{\Sigma}_j = \mathbf{X}^\top \mathbf{L}_{I_j} \mathbf{X} / n$ and $\mathbf{M}_j = (\mathbf{X}^\top \mathbf{L}_{I_j} \mathbf{X} / k_j + \lambda_j \mathbf{I}_p)^\dagger$. Recall that $\mathbf{y} = \mathbf{X} \boldsymbol{\beta}_0 + \mathbf{f}_{\text{NL}}$ and

$$\begin{aligned} \widehat{\boldsymbol{\beta}}_{k_j, \infty}^{\lambda_j} &= \mathbb{E}_{I_j \sim \mathcal{I}_{k_j}} \left[\left(\frac{1}{k_j} \mathbf{X}^\top \mathbf{L}_{I_j} \mathbf{X} + \lambda_j \mathbf{I}_p \right)^{-1} \frac{\mathbf{X}^\top \mathbf{L}_{I_j} \mathbf{y}}{k_j} \right] \\ &= \mathbb{E}_{I_j \sim \mathcal{I}_{k_j}} [\mathbf{M}_j \widehat{\Sigma}_j] \boldsymbol{\beta}_0 + \mathbb{E}_{I_j \sim \mathcal{I}_{k_j}} [\mathbf{M}_j \mathbf{X}^\top \mathbf{L}_{I_j} / k_j] \mathbf{f}_{\text{NL}}. \end{aligned}$$

Then, for $\mathbf{a} \in \mathbb{R}^p$ with bounded l_2 norm, we have

$$\begin{aligned} \mathbf{a}^\top (\widehat{\boldsymbol{\beta}}_{k_1, \infty}^{\lambda_1} - \widehat{\boldsymbol{\beta}}_{k_2, \infty}^{\lambda_2}) &= \mathbf{a}^\top \underbrace{(\mathbb{E}_{I_1 \sim \mathcal{I}_{k_1}} [\mathbf{M}_1 \widehat{\Sigma}_1] - \mathbb{E}_{I_2 \sim \mathcal{I}_{k_2}} [\mathbf{M}_2 \widehat{\Sigma}_2])}_{T_1} \boldsymbol{\beta}_0 \\ &\quad + \mathbf{a}^\top \underbrace{(\mathbb{E}_{I_1 \sim \mathcal{I}_{k_1}} [\mathbf{M}_1 \mathbf{X}^\top \mathbf{L}_{I_1}] / k_1 - \mathbb{E}_{I_2 \sim \mathcal{I}_{k_2}} [\mathbf{M}_2 \mathbf{X}^\top \mathbf{L}_{I_2}] / k_2)}_{T_2} \mathbf{f}_{\text{NL}}. \end{aligned} \tag{21}$$

Next, we analyze the three terms separately.

For the first term, from Lemma D.1 (1) we have for $\lambda_j > 0$,

$$\mathbb{E}_{I_j \sim \mathcal{I}_{k_j}} [\mathbf{M}_j \widehat{\Sigma}_j] = \mathbf{I}_p - \mathbb{E}_{I_j \sim \mathcal{I}_{k_j}} [\lambda_j \mathbf{M}_j] \simeq \mathbf{I}_p - (v(-\lambda_j; \psi_j) \boldsymbol{\Sigma} + \mathbf{I}_p)^{-1},$$

where $v(-\lambda_j; \psi_j)$ is as defined in (40). By Lemma C.1 we have that for $\bar{\psi} \in [\phi, \infty]$ there exists a segment \mathcal{P} such that for all $(\lambda_1, \psi_1), (\lambda_2, \psi_2) \in \mathcal{P}$, it holds that $v(-\lambda_1; \psi_1) = v(-\lambda_1; \psi_2)$ as $p/k_1 \rightarrow \psi_1$ and $p/k_2 \rightarrow \psi_2$. By the definition of deterministic equivalents (Definition E.1), it follows that

$$\begin{aligned} T_1 &= \text{tr}[\boldsymbol{\beta}_0 \mathbf{a}^\top \mathbb{E}_{I_1 \sim \mathcal{I}_{k_1}} [\mathbf{M}_1 \widehat{\Sigma}_1]] - \text{tr}[\boldsymbol{\beta}_0 \mathbf{a}^\top \mathbb{E}_{I_2 \sim \mathcal{I}_{k_2}} [\mathbf{M}_2 \widehat{\Sigma}_2]] \\ &\xrightarrow[p \rightarrow \infty]{\text{a.s.}} \text{tr}[\boldsymbol{\beta}_0 \mathbf{a}^\top (\mathbf{I}_p - (v(-\lambda_1; \psi_1) \boldsymbol{\Sigma} + \mathbf{I}_p)^{-1})] - \text{tr}[\boldsymbol{\beta}_0 \mathbf{a}^\top (\mathbf{I}_p - (v(-\lambda_2; \psi_2) \boldsymbol{\Sigma} + \mathbf{I}_p)^{-1})] \\ &= 0. \end{aligned}$$

For the second term, notice that by Lemma D.2,

$$\frac{1}{k_j} \mathbf{a}^\top \mathbf{M}_j \mathbf{X}^\top \mathbf{L}_{I_j} \mathbf{f}_{\text{NL}} = \mathbf{a}^\top \frac{\mathbf{Z} \boldsymbol{\Sigma}^{\frac{1}{2}}}{k_j} \left(\frac{\mathbf{Z} \boldsymbol{\Sigma} \mathbf{Z}^\top}{k_j} + \lambda_j \mathbf{I}_p \right)^{-1} \mathbf{f}_{\text{NL}} \xrightarrow{\text{a.s.}} 0.$$

From Lemma G.5 (2) of [13], it follows that

$$\mathbf{a}^\top \mathbb{E}_{I_j \sim \mathcal{I}_{k_j}} [\mathbf{M}_j \mathbf{X}^\top \mathbf{L}_{I_j}] / k_j \mathbf{f}_{\text{NL}} \xrightarrow{\text{a.s.}} 0, \quad j = 1, 2,$$

and $T_2 \xrightarrow{\text{a.s.}} 0$.

Combining the above results and applying triangle inequality on (21) yields that

$$|\mathbf{a}^\top (\widehat{\boldsymbol{\beta}}_{k_1, \infty}^{\lambda_1} - \widehat{\boldsymbol{\beta}}_{k_2, \infty}^{\lambda_2})| \leq |T_1| + |T_2| \xrightarrow{\text{a.s.}} 0,$$

which completes the proof when $M = \infty$. When $M \in \mathbb{N}$, replacing the above expectation by the average over M simple random samples $I_1, I_2 \sim \mathcal{I}_k$ completes the proof.

(2) $\lambda_1 \lambda_2 = 0$.

When $\lambda_1 = \lambda_2 = 0$, it is trivially true as $k_1 = k_2$. Otherwise, without loss of generality, we assume $\lambda_1 = 0$ and $\lambda_2 > 0$. We use the same decomposition (21) and first analyze T_1 . From part one we have that for $\lambda > 0$, $P_{n, \lambda} - Q_{n, \lambda} \xrightarrow{\text{a.s.}} 0$ where

$$\begin{aligned} P_{n, \lambda} &:= \text{tr}[\boldsymbol{\beta}_0 \mathbf{a}^\top \mathbb{E}_{I_1 \sim \mathcal{I}_{k_1}} [(\mathbf{X}^\top \mathbf{L}_{I_1} \mathbf{X} / k_1 + \lambda \mathbf{I}_p)^\dagger \widehat{\Sigma}_1]], \\ Q_{n, \lambda} &:= \text{tr}[\boldsymbol{\beta}_0 \mathbf{a}^\top (\mathbf{I}_p - (v(-\lambda; \psi_1) \boldsymbol{\Sigma} + \mathbf{I}_p)^{-1})]. \end{aligned}$$

We next show that

$$\lim_{n \rightarrow \infty} \lim_{\lambda \rightarrow 0^+} P_{n, \lambda} = \lim_{\lambda \rightarrow 0^+} \lim_{n \rightarrow \infty} P_{n, \lambda} = 0, \tag{22}$$

by proving that the function $P_{n,\lambda}$ is equicontinuous family in λ over $\Lambda = [0, \lambda_{\max}]$ for any $\lambda_{\max} \in (0, \infty)$ fixed. Note that for all $\lambda \in \Lambda$,

$$|P_{n,\lambda}| \leq \mathbb{E}_{I_1 \sim \mathcal{I}_{k_1}} [\|(\mathbf{X}^\top \mathbf{L}_{I_1} \mathbf{X}/k_1 + \lambda \mathbf{I}_p)^+ \widehat{\boldsymbol{\Sigma}}_1\|_{\text{op}}] \|\boldsymbol{\beta}_0 \mathbf{a}^\top\|_{\text{tr}} \leq \|\boldsymbol{\beta}_0\|_2^2 \|\mathbf{a}\|_2^2$$

and its derivative

$$\begin{aligned} \left| \frac{\partial}{\partial \lambda} P_{n,\lambda} \right| &= \left| \mathbb{E}_{I_1 \sim \mathcal{I}_{k_1}} [\text{tr}[(\mathbf{X}^\top \mathbf{L}_{I_1} \mathbf{X}/k_1 + \lambda \mathbf{I}_p)^{-2} \widehat{\boldsymbol{\Sigma}}_1 \boldsymbol{\beta}_0 \mathbf{a}^\top]] \right| \\ &\leq \mathbb{E}_{I_1 \sim \mathcal{I}_{k_1}} [\|(\mathbf{X}^\top \mathbf{L}_{I_1} \mathbf{X}/k_1 + \lambda \mathbf{I}_p)^{-2} \widehat{\boldsymbol{\Sigma}}_1\|_{\text{op}}] \|\boldsymbol{\beta}_0 \mathbf{a}^\top\|_{\text{tr}} \\ &\leq \|\boldsymbol{\beta}_0\|_2^2 \|\mathbf{a}\|_2^2 \end{aligned}$$

are uniformly bounded in λ almost surely. In the above inequalities, the operator norm is bounded because $\|(\mathbf{X}^\top \mathbf{L}_{I_1} \mathbf{X}/k_1 + \lambda \mathbf{I}_p)^{-1} \widehat{\boldsymbol{\Sigma}}_1\|_{\text{op}} \leq s_i/(s_i + \lambda) \leq 1$ and $\|(\mathbf{X}^\top \mathbf{L}_{I_1} \mathbf{X}/k_1 + \lambda \mathbf{I}_p)^{-2} \widehat{\boldsymbol{\Sigma}}_1\|_{\text{op}} \leq s_i/(s_i + \lambda)^2 \leq 1/(s_i + \lambda) \leq 1$, by noting that $\limsup \|\widehat{\boldsymbol{\Sigma}}_j\|_{\text{op}} \leq \limsup \max_{1 \leq i \leq p} s_i^2 \leq r_{\max}(1 + \sqrt{\psi_1})^2$ and $\liminf \|\widehat{\boldsymbol{\Sigma}}_j\|_{\text{op}} \geq \liminf \min_{1 \leq i \leq p} s_i^2 \geq r_{\min}(1 - \sqrt{\psi_1})^2$ almost surely as $k_1, p \rightarrow \infty$ and $p/k_1 \rightarrow \psi_1 \in (0, \infty) \setminus \{1\}$ [45]. On the other hand, since $Q_{n,\lambda}$ is a continuous function of $v(-\lambda; \psi_1)$, we have

$$\begin{aligned} |Q_{n,\lambda}| &\leq \|v(-\lambda; \psi_1) \boldsymbol{\Sigma} (v(-\lambda; \psi_1) \boldsymbol{\Sigma} + \mathbf{I}_p)^{-1}\|_{\text{op}} \|\boldsymbol{\beta}_0 \mathbf{a}^\top\|_{\text{tr}} \leq \|\boldsymbol{\beta}_0\|_2^2 \|\mathbf{a}\|_2^2 \\ \left| \frac{\partial}{\partial \lambda} Q_{n,\lambda} \right| &= \left| \text{tr}[\boldsymbol{\beta}_0 \mathbf{a}^\top (v(-\lambda; \psi_1) \boldsymbol{\Sigma} + \mathbf{I}_p)^{-2} \boldsymbol{\Sigma}] \frac{\partial v(-\lambda; \psi_1)}{\partial \lambda} \right| \\ &\leq \|\boldsymbol{\beta}_0 \mathbf{a}^\top\|_{\text{op}} \left| \frac{\partial v(-\lambda; \psi_1)}{\partial \lambda} \right| \int \frac{r}{(1 + v(-\lambda; \psi_1)r)^2} dH_p(r). \end{aligned}$$

When $\psi_1 > 1$, $\partial v(-\lambda; \psi_1)/\partial \lambda$ is bounded over Λ [13, Lemma E.10 and Lemma E.11] and thus $|\partial Q_{n,\lambda}/\partial \lambda|$ is also bounded almost surely. When $\psi_1 < 1$, from Lemma E.12 of [13] we have

$$\begin{aligned} \left| \frac{\partial v(-\lambda; \psi_1)}{\partial \lambda} \right| \int \frac{r}{(1 + v(-\lambda; \psi_1)r)^2} dH_p(r) &= -\frac{\partial v(-\lambda; \psi_1)}{\partial \lambda} \int \frac{r}{(1 + v(-\lambda; \psi_1)r)^2} dH_p(r) \\ &= \frac{\int \frac{r}{(1 + v(-\lambda; \psi_1)r)^2} dH_p(r)}{\frac{1}{v(-\lambda; \psi_1)^2} - \psi_1 \int \frac{r^2}{(1 + v(-\lambda; \psi_1)r)^2} dH_p(r)} \\ &\leq \frac{1}{1 - \psi_1}, \end{aligned}$$

since $v(-\lambda; \psi_1) \leq v(0; \psi_1) = +\infty$. Therefore, $|\partial Q_{n,\lambda}/\partial \lambda|$ is uniformly bounded over Λ for $\psi_1 \in (0, \infty) \setminus \{1\}$. Thus, by the Moore-Osgood theorem, the convergence is uniform in λ and (22) follows. Finally, since T_2 is bounded analogously as in Part (1), the equivalence holds when $\lambda_1 = 0$.

B.2 Proof of Proposition 4

From Lemma B.2, it follows that

$$v(-\lambda; \phi_n) \simeq \frac{1}{n} \text{tr} \left[\left(\frac{1}{n} \mathbf{X} \mathbf{X}^\top + \lambda \mathbf{I}_p \right)^{-1} \right]. \quad (23)$$

By the continuity of function $\phi \mapsto v(-\lambda; \phi)$ from Lemma F.10 and F.11 of [13], we have $v(-\lambda; \phi_n) \simeq v(-\lambda; \phi)$ as $\phi_n = p/n \rightarrow \phi$. From Lemma C.1, there exists $\bar{\lambda}_n$ such that

$$v(0; \psi_n) = v(-\bar{\lambda}_n; \phi_n),$$

as $\psi_n \rightarrow \bar{\psi}$ and $\phi_n \rightarrow \phi$. Involving (23) on the both sides yields that

$$\frac{1}{n} \text{tr} \left[\left(\frac{1}{n} \mathbf{X} \mathbf{X}^\top + \bar{\lambda}_n \mathbf{I}_p \right)^{-1} \right] \simeq \frac{1}{k} \text{tr} \left[\left(\frac{1}{k} \mathbf{X} \mathbf{L}_{I_1} \mathbf{X}^\top \right)^+ \right] \simeq \frac{1}{Mk} \sum_{\ell=1}^M \text{tr} \left[\left(\frac{1}{k} \mathbf{X} \mathbf{L}_{I_\ell} \mathbf{X}^\top \right)^+ \right]$$

where $\{I_1, \dots, I_M\}$ is a simple random sample from \mathcal{I}_k .

B.3 Proof of Corollary 5

We will use a structural equivalence (much more direct than the first-order equivalence considered in the paper) between generalized ridge predictor and the isotropic ridge predictor to prove the result.

The generalized ridge predictor (9), trained on the subsampled dataset \mathcal{D}_I , can be expressed as:

$$\widehat{\beta}_k^{\lambda, \mathbf{G}}(\mathcal{D}_I) = \left(\frac{1}{k} \mathbf{X}^\top \mathbf{L}_I \mathbf{X} + \lambda \mathbf{G} \right)^{-1} \frac{\mathbf{X}^\top \mathbf{L}_I \mathbf{y}}{k}.$$

Observe that we can equivalently manipulate the generalized ridge estimator into:

$$\widehat{\beta}_k^{\lambda, \mathbf{G}}(\mathcal{D}_I) = \mathbf{G}^{-1/2} \left(\frac{1}{k} \mathbf{G}^{-1/2} \mathbf{X}^\top \mathbf{L}_I \mathbf{X} \mathbf{G}^{-1/2} + \lambda \mathbf{I}_p \right)^{-1} \mathbf{G}^{-1/2} \frac{\mathbf{X}^\top \mathbf{L}_I \mathbf{y}}{k}.$$

Recalling from Assumption 2 that $\mathbf{X} = \mathbf{Z} \boldsymbol{\Sigma}^{1/2}$, where $\mathbf{Z} \in \mathbb{R}^{n \times p}$ contains \mathbf{z}_i^\top in the i -th row for $i \in [n]$, we obtain

$$\widehat{\beta}_k^{\lambda, \mathbf{G}}(\mathcal{D}_I) = \mathbf{G}^{-1/2} \left(\frac{1}{k} \mathbf{G}^{-1/2} \boldsymbol{\Sigma}^{1/2} \mathbf{Z}^\top \mathbf{L}_I \mathbf{Z} \boldsymbol{\Sigma}^{1/2} \mathbf{G}^{-1/2} + \lambda \mathbf{I}_p \right)^{-1} \frac{\mathbf{G}^{-1/2} \boldsymbol{\Sigma}^{1/2} \mathbf{Z}^\top \mathbf{L}_I \mathbf{y}}{k}. \quad (24)$$

We now define $\boldsymbol{\Sigma}_{\mathbf{G}} = \mathbf{G}^{-1/2} \boldsymbol{\Sigma} \mathbf{G}^{-1/2}$ and consider a transformed feature matrix $\mathbf{X}_{\mathbf{G}} = \mathbf{Z} \boldsymbol{\Sigma}_{\mathbf{G}}^{1/2}$. Denote the transformed dataset corresponding to the new feature matrix by $\mathcal{D}^{\mathbf{G}}$ (where we keep the response vector as is). Using (24), we can write relate the generalized ridge estimator fitted on the dataset \mathcal{D} to the standard ridge estimator fitted on the data $\mathcal{D}^{\mathbf{G}}$ (both at the same scalar regularization level λ) by:

$$\widehat{\beta}_k^{\lambda, \mathbf{G}}(\mathcal{D}_I) = \mathbf{G}^{-1/2} \widehat{\beta}_k^{\lambda}(\mathcal{D}_I^{\mathbf{G}}).$$

Observe that the transformed dataset $\mathcal{D}^{\mathbf{G}}$ satisfies Assumption 2 as the eigenvalues of \mathbf{G} are bounded away from 0 and ∞ . Further, note that both Theorem 1 and Theorem 3 are invariant to linear transformations of the estimator. Thus, we conclude that the results of both the theorems continue to hold. The equivalence path now use the modified $\boldsymbol{\Sigma}_{\mathbf{G}}$, which in turn changes the spectral distribution H_p to that of $\boldsymbol{\Sigma}_{\mathbf{G}}$, in defining the end points via (4), given by:

$$\widetilde{H}_p = \frac{1}{p} \sum_{i=1}^p \mathbf{1}_{\{\widetilde{r}_i \leq r\}},$$

where \widetilde{r}_i for $i \in [p]$ are eigenvalues of $\boldsymbol{\Sigma}_{\mathbf{G}}$. This completes the proof.

B.4 Technical lemmas

Lemma B.1 (Relationship between \widehat{m} and \widehat{v}). For $\mathbf{X} \in \mathbb{R}^{n \times p}$ and $\phi_n = p/n$, define

$$\widehat{m}(z; \phi_n) = \frac{1}{p} \operatorname{tr} \left[\left(\frac{1}{n} \mathbf{X}^\top \mathbf{X} - z \mathbf{I}_p \right)^{-1} \right], \quad \text{and} \quad \widehat{v}(z; \phi_n) = \frac{1}{n} \operatorname{tr} \left[\left(\frac{1}{n} \mathbf{X} \mathbf{X}^\top - z \mathbf{I}_p \right)^{-1} \right].$$

It holds that

$$\phi_n z \widehat{m}(z; \phi_n) + \phi_n - 1 = z \widehat{v}(z; \phi_n). \quad (25)$$

Proof. Let r be the rank of the matrix \mathbf{X} . Denote by s_i , $i = 1, \dots, r$, the non-zero eigenvalues of the matrix $\mathbf{X}^\top \mathbf{X}$. Note that these are the same non-zero eigenvalues of the matrix $\mathbf{X} \mathbf{X}^\top$. Define function S such that for $z \neq 0$,

$$S(z) = \sum_{i=1}^r \frac{1}{s_i - z}.$$

For $z \neq 0$, write out \widehat{v} and \widehat{m} in terms of $S(z)$ as:

$$\widehat{v}(z; \phi_n) = \frac{1}{n} \sum_{i=1}^r \frac{1}{s_i - z} - \frac{1}{n} \frac{n-r}{z} = \frac{1}{n} S(z) - \frac{1}{n} \frac{n-r}{z}$$

$$\widehat{m}(z; \phi_n) = \frac{1}{p} \sum_{i=1}^r \frac{1}{s_i - z} - \frac{1}{p} \frac{p-r}{z} = \frac{1}{p} S(z) - \frac{1}{p} \frac{p-r}{z}.$$

We now expand the left-hand side of (25):

$$\begin{aligned} \phi_n z \widehat{m}(z; \phi_n) + \phi_n - 1 &= \frac{p}{n} z \left(\frac{1}{p} S(z) - \frac{1}{p} \frac{p-r}{z} \right) + \frac{p}{n} - 1 \\ &= \frac{1}{n} z S(z) - \frac{p-r}{n} + \frac{p}{n} - 1 \\ &= \frac{1}{n} z S(z) - \frac{n-r}{n} + \frac{n-r}{n} - \frac{p-r}{n} + \frac{p}{n} - 1 \\ &= z \widehat{v}(z; \phi_n) + \frac{n-p}{n} + \frac{p}{n} - 1 \\ &= z \widehat{v}(z; \phi_n), \end{aligned}$$

which finishes the proof. \square

Lemma B.2 (Equivalence of \widehat{v} and v). *Suppose Assumptions 1–2 hold. Then it holds that*

$$v(-\lambda; \phi_n) \simeq \frac{1}{n} \operatorname{tr} \left[\left(\frac{1}{n} \mathbf{X} \mathbf{X}^\top + \lambda I_p \right)^{-1} \right].$$

Proof. From Lemma B.1, we have

$$z \widehat{v}(z; \phi_n) + (1 - \phi_n) = \phi_n z \widehat{m}(z; \phi_n).$$

Substituting $z = -\lambda$ yields that

$$\lambda \frac{1}{n} \operatorname{tr} \left[\left(\frac{1}{n} \mathbf{X} \mathbf{X}^\top + \lambda I_p \right)^{-1} \right] + (\phi_n - 1) = \phi_n \lambda \frac{1}{p} \operatorname{tr} \left[\left(\frac{1}{n} \mathbf{X}^\top \mathbf{X} + \lambda I_p \right)^{-1} \right]. \quad (26)$$

From Corollary E.4, we have

$$\begin{aligned} \frac{1}{v(-\lambda; \phi_n)} &= \lambda + \phi \operatorname{tr}[\Sigma(v(-\lambda; \phi_n)\Sigma + I_p)^{-1}]/p, \\ \frac{1}{p} \operatorname{tr}[(v(-\lambda; \phi_n)\Sigma + I_p)^{-1}] &\simeq \frac{1}{p} \lambda \operatorname{tr} \left[\left(\frac{1}{n} \mathbf{X}^\top \mathbf{X} + \lambda I_p \right)^{-1} \right]. \end{aligned}$$

This implies

$$\begin{aligned} 1 &= \lambda v(-\lambda; \phi_n) + \phi_n \operatorname{tr}[v(-\lambda; \phi_n)\Sigma(v(-\lambda; \phi_n)\Sigma + I_p)^{-1}]/p \\ &= \lambda v(-\lambda; \phi_n) + \phi_n - \phi_n \operatorname{tr}[(v(-\lambda; \phi_n)\Sigma + I_p)^{-1}]/p \\ &\simeq \lambda v(-\lambda; \phi_n) + \phi_n - \phi_n \lambda \operatorname{tr} \left[\left(\frac{1}{n} \mathbf{X}^\top \mathbf{X} + \lambda I_p \right)^{-1} \right] / p \\ &= 1 + \lambda v(-\lambda; \phi_n) - \frac{1}{n} \lambda \operatorname{tr} \left[\left(\frac{1}{n} \mathbf{X} \mathbf{X}^\top + \lambda I_p \right)^{-1} \right], \end{aligned}$$

where the last equality follows from (26). This concludes the proof. \square

C Proof of results in Section 5

C.1 Proof of Theorem 6

By Assumptions 3–4, for $\mathbf{A} = \Sigma^{1/2}$ and $\mathbf{b} = \mathbf{0}$, $R_p(\lambda; \phi, \psi)$, as defined in (11), converges to

$$R_p(\lambda; \phi, \psi) \xrightarrow{\text{a.s.}} \rho^2 \widetilde{c}(-\lambda; \phi, \psi) + \sigma^2 \widetilde{v}(-\lambda; \phi, \psi) =: \mathcal{R}(\lambda; \phi, \psi) \quad (27)$$

where the nonnegative constants $\tilde{c}(-\lambda; \phi, \psi)$ and $\tilde{v}(-\lambda; \phi, \psi)$ are defined through the following equations:

$$\begin{aligned}\tilde{v}(-\lambda; \phi, \psi) &= \frac{\phi \int \frac{r^2}{(1 + v(-\lambda; \psi)r)^2} dH(r)}{v(-\lambda; \psi)^{-2} - \phi \int \frac{r^2}{(1 + v(-\lambda; \psi)r)^2} dH(r)}, \\ \tilde{c}(-\lambda; \phi, \psi) &= (\tilde{v}(-\lambda; \phi, \psi) + 1) \int \frac{r}{1 + v(-\lambda; \psi)r} dG(r), \\ \frac{1}{v(-\lambda; \psi)} &= \lambda + \psi \int \frac{r}{1 + v(-\lambda; \psi)r} dH(r).\end{aligned}$$

From the proof of Theorem 1, we also have that $\mathcal{R}(\lambda; \phi, \psi) \simeq R_p(\lambda; \phi, \psi) \simeq R(\widehat{\beta}_{k, \infty}^\lambda; \Sigma, \mathbf{0}, \beta_0)$.

Since $\mathcal{R}(0; \phi, \psi)$ is a continuous function of ϕ and $v(0; \psi)$ and is increasing in ϕ for any fixed ψ , it follows that for $0 < \phi_1 \leq \phi_2 < \infty$,

$$\min_{\psi \geq \phi_1} \mathcal{R}(0; \phi_1, \psi) \leq \min_{\psi \geq \phi_2} \mathcal{R}(0; \phi_1, \psi) \leq \min_{\psi \geq \phi_2} \mathcal{R}(0; \phi_2, \psi),$$

where the first inequality follows because $\{\psi : \psi \geq \phi_1\} \supseteq \{\psi : \psi \geq \phi_2\}$, and the second inequality follows because $\mathcal{R}(0; \phi, \psi)$ is increasing in ϕ for a fixed ψ . Thus, $\min_{\psi \geq \phi} \mathcal{R}(0; \phi, \psi)$ is a continuous and monotonically increasing function in ϕ .

Finally, note that from Lemma C.1, for any λ , there exists ψ such that $v(0; \phi, \psi) = v(-\lambda; \phi, \phi)$. This implies that $\mathcal{R}(0; \phi, \psi) = \mathcal{R}(\lambda; \phi, \phi)$. Then we have $\min_{\psi \geq \phi} \mathcal{R}(0; \phi, \psi) \leq \min_{\lambda \geq 0} \mathcal{R}(\lambda; \phi, \phi)$. Conversely, since for any ψ , there exists λ such that $v(0; \phi, \psi) = v(-\lambda; \phi, \phi)$, we also have $\min_{\psi \geq \phi} \mathcal{R}(0; \phi, \psi) \geq \min_{\lambda \geq 0} \mathcal{R}(\lambda; \phi, \phi)$. Combining the two inequalities yields the conclusion.

C.2 Technical lemmas

In this section, we gather results on certain analytic properties of the fixed-point solution $v(-\lambda; \phi)$ defined in (40).

Lemma C.1 (Contour of fixed-point solutions). *As $n, p \rightarrow \infty$ such that $p/n \rightarrow \phi \in (0, \infty)$, for any $\bar{\psi} \in [\phi, +\infty]$, there exists a unique value $\bar{\lambda} \geq 0$ (or conversely for $\bar{\lambda} \in [0, \infty]$, there exists a unique value $\bar{\psi} \in [\phi \vee 1, \infty]$) such that for all (λ, ψ) on the path*

$$\mathcal{P} = \{(1 - \theta) \cdot (\bar{\lambda}, \phi) + \theta \cdot (0, \bar{\psi}) \mid \theta \in [0, 1]\},$$

it holds that

$$v_p(-\lambda; \psi) = v(-\bar{\lambda}; \phi) = v(0; \bar{\psi}).$$

where $v_p(-\lambda; \psi)$ is as defined in (40).

Proof. Since when $\phi < 1$, $v(0; \psi) = +\infty$ for all $\psi \in [\phi, 1]$, we can restrict ourselves to $\psi \in [\phi \vee 1, +\infty]$. From Lemma E.11 (1) of [13], the function $\psi \mapsto v(0; \psi)$ is strictly decreasing over $\psi \in [\phi \vee 1, \infty]$ with range

$$v(0; \phi \vee 1) = \begin{cases} v(0; \phi), & \phi \in (1, \infty) \\ \lim_{\psi \rightarrow 1^+} v(0; \psi) = +\infty, & \phi \in (0, 1] \end{cases}, \quad v(0; +\infty) := \lim_{\psi \rightarrow +\infty} v(0; \psi) = 0.$$

From Lemma E.12 (3) of [13], the function $\lambda \mapsto v(-\lambda; \phi)$ is strictly decreasing over $\lambda \in [0, \infty]$ with range

$$v(0; \phi) = \begin{cases} v(0; \phi), & \phi \in (1, \infty) \\ \lim_{\lambda \rightarrow 0^+} v(-\lambda; \phi) = +\infty, & \phi \in (0, 1] \end{cases}, \quad v(-\infty; \phi) := \lim_{\lambda \rightarrow +\infty} v(-\lambda; \phi) = 0.$$

Note that $v(0; \phi \vee 1) = v(0; \phi)$. For $\bar{\psi} \in [\phi \vee 1, \infty]$, by the intermediate value theorem, there exists unique $\bar{\lambda} \in [0, \infty]$ such that $v(-\bar{\lambda}; \phi) = v(0; \bar{\psi})$. Furthermore, when $\bar{\psi} \leq 1$, $\bar{\lambda} = 0$ is also the unique

value such that $v(0; \bar{\psi}) = v(-\bar{\lambda}; \phi)$. Conversely, for $\bar{\lambda} \in [0, \infty]$, there also exists $\bar{\psi} \in [\phi \vee 1, \infty]$ such that $v(-\bar{\lambda}; \phi) = v(0; \bar{\psi})$.

Based on the definition of fixed-point solutions, it follows that

$$\frac{1}{v(-\bar{\lambda}; \phi)} = \bar{\lambda} + \phi \int \frac{r}{1 + v(-\bar{\lambda}; \phi)r} dH_p(r) = \bar{\psi} \int \frac{r}{1 + v(0; \bar{\psi})r} dH_p(r) = \frac{1}{v(0; \bar{\psi})}.$$

Then, for any $(\lambda, \psi) = (1 - \theta)(\bar{\lambda}, \phi) + \theta(0, \bar{\psi})$ on the path \mathcal{P} , we have

$$\begin{aligned} \frac{1}{v(-\bar{\lambda}; \phi)} &= (1 - \theta) \frac{1}{v(-\bar{\lambda}; \phi)} + \theta \frac{1}{v(0; \bar{\psi})} \\ &= (1 - \theta)\bar{\lambda} + (1 - \theta)\phi \int \frac{r}{1 + v(-\bar{\lambda}; \phi)r} dH_p(r) + \theta\bar{\psi} \int \frac{r}{1 + v(0; \bar{\psi})r} dH_p(r) \\ &= \lambda + \psi \int \frac{r}{1 + v(-\bar{\lambda}; \phi)r} dH_p(r). \end{aligned}$$

Because $v_p(-\lambda; \psi)$ is the unique solution to the fixed-point equation:

$$\frac{1}{v_p(-\lambda; \psi)} = \lambda + \psi \int \frac{r}{1 + v_p(-\lambda; \psi)r} dH_p(r),$$

it then follows that $v_p(-\lambda; \psi) = v(-\bar{\lambda}; \phi) = v(0; \bar{\psi})$. □

D Asymptotic equivalents, concentration results, and other useful lemmas

D.1 Full-ensemble resolvents

Lemma D.1 (Full-ensemble resolvents). *Let $\widehat{\Sigma} = \mathbf{X}^\top \mathbf{X}/n$, $\widehat{\Sigma}_j = \mathbf{X}^\top \mathbf{L}_{I_j} \mathbf{X}/k$ where $I_j \sim \mathcal{I}_k$. Let $\Sigma_{1 \cap 2} = \mathbf{X}^\top \mathbf{L}_{I_1 \cap I_2} \mathbf{X}/|I_1 \cap I_2|$ and $\mathbf{C} \in \mathbb{R}^{p \times p}$ with bounded operator norm almost surely. As $n, p, k \rightarrow \infty$, $p/n \rightarrow \phi$, $p/k \rightarrow \psi$, the following asymptotic equivalences hold:*

(1) *Basic ridge resolvent:*

$$\lambda \mathbb{E}_{I_1 \sim \mathcal{I}_k} [(\widehat{\Sigma}_1 + \lambda \mathbf{I}_p)^{-1}] \simeq (v_p(-\lambda; \psi) \Sigma + \mathbf{I}_p)^{-1}.$$

(2) *Bias resolvent with $\mathbf{C} \perp \mathbf{X}$:*

$$\begin{aligned} &\mathbb{E}_{I_1, I_2 \sim \mathcal{I}_k} [\lambda^2 (\widehat{\Sigma}_1 + \lambda \mathbf{I}_p)^{-1} \mathbf{C} (\widehat{\Sigma}_2 + \lambda \mathbf{I}_p)^{-1}] \\ &\simeq (v_p(-\lambda; \psi) \Sigma + \mathbf{I}_p)^{-1} (\tilde{v}_p(-\lambda; \phi, \psi, \mathbf{C}) \Sigma + \mathbf{C}) (v_p(-\lambda; \psi) \Sigma + \mathbf{I}_p)^{-1}. \end{aligned}$$

(3) *Variance resolvent with $\mathbf{C} \perp \mathbf{X}$:*

$$\begin{aligned} &\mathbb{E}_{I_1, I_2 \sim \mathcal{I}_k} [(\widehat{\Sigma}_1 + \lambda \mathbf{I}_p)^{-1} \widehat{\Sigma}_{1 \cap 2} (\widehat{\Sigma}_2 + \lambda \mathbf{I}_p)^{-1} \mathbf{C}] \\ &\simeq \phi^{-1} \tilde{v}_v(-\lambda; \phi, \psi) (v_p(-\lambda; \psi) \Sigma + \mathbf{I}_p)^{-2} \Sigma \mathbf{C}. \end{aligned}$$

(4) *Bias resolvent with $\mathbf{C} = \widehat{\Sigma}$:*

$$\begin{aligned} &\mathbb{E}_{I_1, I_2 \sim \mathcal{I}_k} [\lambda^2 (\widehat{\Sigma}_1 + \lambda \mathbf{I}_p)^{-1} \widehat{\Sigma} (\widehat{\Sigma}_2 + \lambda \mathbf{I}_p)^{-1}] \\ &\simeq \tilde{d}(-\lambda; \phi, \psi) (1 + \tilde{v}_p(-\lambda; \phi, \psi)) (v_p(-\lambda; \psi) \Sigma + \mathbf{I}_p)^{-2} \Sigma. \end{aligned}$$

Here $v(-\lambda; \phi)$ is as defined in (41), and we let $\tilde{v}_p(-\lambda; \phi, \psi) = \tilde{v}_p(-\lambda; \phi, \psi, \Sigma)$,

$$\tilde{v}_p(-\lambda; \phi, \psi, \mathbf{C}) = \frac{\lim_{k, n, p} \phi \operatorname{tr}[\mathbf{C} \Sigma (v_p(-\lambda; \psi) \Sigma + \mathbf{I}_p)^{-2}]/p}{v_p(-\lambda; \psi)^{-2} - \phi \int \frac{r^2}{(1 + v_p(-\lambda; \psi)r)^2} dH_p(r)},$$

$$\tilde{v}_v(-\lambda; \phi, \psi) = \frac{1}{v_p(-\lambda; \psi)^{-2} - \phi \int \frac{r^2}{(1 + v_p(-\lambda; \psi)r)^2} dH_p(r)},$$

$$\tilde{d}(-\lambda; \phi, \psi) = \left(1 - \phi \int \frac{v_p(-\lambda; \psi)r}{1 + v_p(-\lambda; \psi)r} dH_p(r)\right)^2.$$

The empirical distribution H_n of eigenvalues can be replaced by the limiting distribution H whenever it exists.

Proof. The proofs for different parts is separated below.

(1) Basic ridge resolvent.

From Definition E.2, we know that $\lambda(\widehat{\Sigma}_j + \lambda \mathbf{I}_p)^{-1} \simeq (v_p(-\lambda; \psi)\Sigma + \mathbf{I}_p)^{-1}$. By the definition of deterministic equivalents in Definition E.1, for any $\mathbf{A} \in \mathbb{R}^{p \times p}$ that has bounded trace norm and is independent to \mathbf{X} , we have

$$\text{tr}[\mathbf{A}(\lambda(\widehat{\Sigma}_j + \lambda \mathbf{I}_p)^{-1} - (v_p(-\lambda; \psi)\Sigma + \mathbf{I}_p)^{-1})] \xrightarrow{\text{a.s.}} 0.$$

From Lemma G.5 (2) of [13], it follows that

$$\begin{aligned} & \text{tr}[\mathbf{A}(\lambda \mathbb{E}_{I_1 \sim \mathcal{I}_k}[(\widehat{\Sigma}_j + \lambda \mathbf{I}_p)^{-1}] - (v_p(-\lambda; \psi)\Sigma + \mathbf{I}_p)^{-1})] \\ &= \mathbb{E}_{I_1 \sim \mathcal{I}_k}[\text{tr}[\mathbf{A}(\lambda(\widehat{\Sigma}_j + \lambda \mathbf{I}_p)^{-1} - (v_p(-\lambda; \psi)\Sigma + \mathbf{I}_p)^{-1})]] \xrightarrow{\text{a.s.}} 0, \end{aligned} \quad (28)$$

which implies that

$$\lambda \mathbb{E}_{I_1 \sim \mathcal{I}_k}[(\widehat{\Sigma}_1 + \lambda \mathbf{I}_p)^{-1}] \simeq (v_p(-\lambda; \psi)\Sigma + \mathbf{I}_p)^{-1}.$$

(2) Bias resolvent with $\mathbf{C} \perp\!\!\!\perp \mathbf{X}$.

Denote $\mathbf{M}_j = (\widehat{\Sigma}_j + \lambda \mathbf{I}_p)^{-1}$ for $j = 1, 2$. From Part (c) of the proof for Lemma S.2.4 in [12], it follows that for $I_1, I_2 \sim \mathcal{I}_k$,

$$\lambda^2 \mathbf{M}_1 \mathbf{C} \mathbf{M}_2 \simeq (v_p(-\lambda; \psi)\Sigma + \mathbf{I}_p)^{-1} (\tilde{v}_p(-\lambda; \phi, \psi, \mathbf{C})\Sigma + \mathbf{C}) (v_p(-\lambda; \psi)\Sigma + \mathbf{I}_p)^{-1}.$$

By the same argument as in (28), the conclusion follows.

(3) Variance resolvent with $\mathbf{C} \perp\!\!\!\perp \mathbf{X}$.

From Lemma E.8 (3) of [13], it follows that for $I_1, I_2 \sim \mathcal{I}_k$,

$$\mathbf{M}_1 \widehat{\Sigma}_{I_1 \cap I_2} \mathbf{M}_2 \mathbf{C} \simeq \phi^{-1} \tilde{v}_v(-\lambda; \phi, \psi) (v_p(-\lambda; \psi)\Sigma + \mathbf{I}_p)^{-2} \Sigma \mathbf{C}.$$

By the same argument as in (28), the conclusion follows.

(4) Bias resolvent with $\mathbf{C} = \widehat{\Sigma}$.

We begin by decomposing the object inside expectation into two terms:

$$\lambda^2 \mathbf{M}_1 \widehat{\Sigma} \mathbf{M}_2 = \frac{|I_1 \cup I_2|}{n} \lambda^2 \mathbf{M}_1 \widehat{\Sigma}_{I_1 \cup I_2} \mathbf{M}_2 + \frac{n - |I_1 \cup I_2|}{n} \lambda^2 \mathbf{M}_1 \widehat{\Sigma}_{(I_1 \cup I_2)^c} \mathbf{M}_2 \quad (29)$$

where $\widehat{\Sigma}_{I_1 \cup I_2} = \mathbf{X}^\top \mathbf{L}_{I_1 \cup I_2} \mathbf{X} / |I_1 \cup I_2|$ and $\widehat{\Sigma}_{(I_1 \cup I_2)^c} = \mathbf{X}^\top (\mathbf{I} - \mathbf{L}_{I_1 \cup I_2}) \mathbf{X} / (n - |I_1 \cup I_2|)$. Next, we analyze each of them.

For the first term, from Lemma D.6 (3) of [13] we have

$$\begin{aligned} \mathbf{M}_1 \widehat{\Sigma}_{I_1 \cup I_2} \mathbf{M}_2 &\simeq v_p(-\lambda; \psi)^2 (1 + \tilde{v}_p(-\lambda; \phi, \psi)) \\ &\quad \left(\frac{2(\psi - \phi)}{2\psi - \phi} \frac{1}{\lambda v_p(-\lambda; \psi)} + \frac{\phi}{2\psi - \phi} \right) (v_p(-\lambda; \psi)\Sigma + \mathbf{I}_p)^{-2} \Sigma \end{aligned} \quad (30)$$

where $\tilde{v}_p(-\lambda; \phi, \psi) = \tilde{v}_p(-\lambda; \phi, \psi, \Sigma)$.

For the second term, from Lemma E.8 (1) of [13] and the product rule of calculus in Lemma E.3 (3), we have

$$\lambda^2 \mathbf{M}_1 \widehat{\Sigma}_{(I_1 \cup I_2)^c} \mathbf{M}_2 \simeq \lambda^2 \mathbf{M}_1 \Sigma \mathbf{M}_2.$$

From Part (2) and Lemma E.3 (1), it follows that

$$\lambda^2 \mathbf{M}_1 \widehat{\Sigma}_{(1 \cup 2)^c} \mathbf{M}_2 \simeq \lambda^2 \mathbf{M}_1 \Sigma \mathbf{M}_2 \simeq (1 + \tilde{v}_p(-\lambda; \phi, \psi)) (v_p(-\lambda; \psi) \Sigma + \mathbf{I}_p)^{-2} \Sigma. \quad (31)$$

Finally, for the coefficients of the two terms, from Lemma S.8.3 of [12], we have that

$$\frac{|I_1 \cup I_2|}{n} = \frac{|I_1| + |I_2| - |I_1 \cap I_2|}{n} \xrightarrow{\text{a.s.}} \frac{\phi(2\psi - \phi)}{\psi^2}, \quad \frac{n - |I_1 \cup I_2|}{n} \xrightarrow{\text{a.s.}} \frac{(\psi - \phi)^2}{\psi^2}. \quad (32)$$

Combining (29)-(32) yields that

$$\begin{aligned} & \lambda^2 \mathbf{M}_1 \widehat{\Sigma} \mathbf{M}_2 \\ & \simeq \frac{1}{\psi^2} (2(\psi - \phi)\phi\lambda v_p(-\lambda; \psi) + \phi^2 \lambda^2 v_p(-\lambda; \psi)^2 + (\psi - \phi)^2) \\ & \quad (1 + \tilde{v}_p(-\lambda; \phi, \psi))(v_p(-\lambda; \psi) \Sigma + \mathbf{I}_p)^{-2} \Sigma \\ & = \frac{(\phi\lambda v_p(-\lambda; \psi) + \psi - \phi)^2}{\psi^2} (1 + \tilde{v}_p(-\lambda; \phi, \psi))(v_p(-\lambda; \psi) \Sigma + \mathbf{I}_p)^{-2} \Sigma \\ & = \left(1 - \phi \int \frac{v_p(-\lambda; \psi)r}{1 + v_p(-\lambda; \psi)r} dH_p(r)\right)^2 (1 + \tilde{v}_p(-\lambda; \phi, \psi))(v_p(-\lambda; \psi) \Sigma + \mathbf{I}_p)^{-2} \Sigma \end{aligned}$$

where the last equality follows by substituting

$$\lambda = v_p(-\lambda; \psi)^{-1} - \psi \int r(1 + v_p(-\lambda; \psi)r)^{-1} dH_p(r)$$

based on the fixed-point equation. □

D.2 Convergence of random linear and quadratic forms

Lemma D.2 (Concentration of linear form with independent components and varying coefficients). *Let $\mathbf{z}_i \in \mathbb{R}^p$ for $i = 1, \dots, n$ be a sequence of random vectors with i.i.d. entries z_{ij} , $j = 1, \dots, p$ such that for each i, j , $\mathbb{E}[z_{ij}] = 0$, $\mathbb{E}[z_{ij}^2] = 1$, $\mathbb{E}[|z_{ij}|^{4+\alpha}] \leq M_\alpha$ for some $\alpha > 0$ and constant $M_\alpha < \infty$. Let $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_n]^\top \in \mathbb{R}^{n \times p}$ be the random matrix formed by concatenating \mathbf{z}_i 's. Let $g : \mathbb{R}^p \rightarrow \mathbb{R}$ be any measurable function such that $\mathbb{E}[g(\mathbf{z}_i)] = 0$ and $\mathbb{E}[\mathbf{z}_i g(\mathbf{z}_i)] = 0$ for $i = 1, \dots, n$. Let $\mathbf{a}_p \in \mathbb{R}^p$ be a sequence of random vectors independent of \mathbf{z}_p such that $\limsup_p \|\mathbf{a}_p\|^2/p \leq M_0$ almost surely for a constant $M_0 < \infty$. Let \mathbf{D} be a positive semidefinite matrix such that $\limsup \|\mathbf{D}\|_{\text{op}} \leq M_0$ almost surely as $p \rightarrow \infty$ for some constant $M_0 < \infty$. Then, as $n, p \rightarrow \infty$ such that $p/n \rightarrow \phi \in (0, \infty)$, we have*

$$\left| \frac{1}{n} \sum_{i=1}^n \left[\mathbf{a}^\top \left(\frac{\mathbf{D}^{\frac{1}{2}} \mathbf{Z}^\top \mathbf{Z} \mathbf{D}^{\frac{1}{2}}}{n} + \lambda \mathbf{I}_p \right)^{-1} \mathbf{D}^{\frac{1}{2}} \mathbf{Z}^\top \right]_i g(\mathbf{z}_i) \right| \xrightarrow{\text{a.s.}} 0. \quad (33)$$

Proof. We will use the standard leave-one-out trick to break the dependence between the i -th component multiplier in the summation in (33) and $g(\mathbf{z}_i)$ for each $i = 1, \dots, n$. To that end, using the Woodbury matrix identity, first observe that

$$\begin{aligned} & \left(\frac{\mathbf{D}^{\frac{1}{2}} \mathbf{Z}^\top \mathbf{Z} \mathbf{D}^{\frac{1}{2}}}{n} + \lambda \mathbf{I}_p \right)^{-1} \\ & = \left(\frac{\mathbf{D}^{\frac{1}{2}} \mathbf{Z}_{-i}^\top \mathbf{Z}_{-i} \mathbf{D}^{\frac{1}{2}}}{n} + \frac{\mathbf{D}^{\frac{1}{2}} \mathbf{z}_i \mathbf{z}_i^\top \mathbf{D}^{\frac{1}{2}}}{n} + \lambda \mathbf{I}_p \right)^{-1} \\ & = \left(\frac{\mathbf{D}^{\frac{1}{2}} \mathbf{Z}_{-i}^\top \mathbf{Z}_{-i} \mathbf{D}^{\frac{1}{2}}}{n} + \lambda \mathbf{I}_p \right)^{-1} \end{aligned}$$

uncorrelatedness of \mathbf{z} and $g(\mathbf{z})$ along with standard concentration of a quadratic form with respect to an independent matrix. Similar strategy has been used in [14] when analyzing kernel ridge regression under proportional asymptotics.

Denote $\mathbf{Y} = (\mathbf{Z}\mathbf{D}\mathbf{Z}^\top/n + \lambda\mathbf{I}_n)^{-1}$. For each pair of $i, j \in [n]$ and $i \neq j$, we let $\mathbf{Z}_{-(ij)} \in \mathbb{R}^{(n-2) \times p}$ be the matrix comprising the $n - 2$ rows of \mathbf{Z} excluding the i th and j th rows, and $\mathbf{U} \in \mathbb{R}^{p \times 2}$ be the matrix with columns $\mathbf{U}_{ij}\mathbf{e}_i = \mathbf{z}_i, \mathbf{U}\mathbf{e}_j = \mathbf{z}_j$. We finally define the matrices $\mathbf{R}_{-(ij)} \in \mathbb{R}^{p \times p}$ as follows:

$$\mathbf{R}_{-(ij)} = \lambda\mathbf{D}^{1/2} \left(\mathbf{D}^{1/2}\mathbf{Z}_{-(ij)}^\top\mathbf{Z}_{-(ij)}\mathbf{D}^{1/2}/n + \lambda\mathbf{I}_p \right)^{-1} \mathbf{D}^{1/2},$$

and let $\tilde{\mathbf{Y}}_{ij} = \{Y_{m,\ell}\}_{m,\ell \in \{i,j\}}$ be the submatrix of \mathbf{Y} . Then, using the block matrix inversion formula (see, e.g., Appendix A.1.4 of [45]) and the Woodbury matrix identity, one can show that

$$\begin{aligned} \tilde{\mathbf{Y}}_{ij} &= (\mathbf{U}_{ij}^\top\mathbf{R}_{-(ij)}\mathbf{U}_{ij}/n + \lambda\mathbf{I}_2)^{-1}, \\ Y_{ij} &= -\frac{\langle \mathbf{z}_i, \mathbf{R}_{-(ij)}\mathbf{z}_j \rangle / n}{d_{ij}}, \end{aligned}$$

where d_{ij} is given by:

$$d_{ij} = (\lambda + \langle \mathbf{z}_i, \mathbf{R}_{-(ij)}\mathbf{z}_j \rangle / n) (\lambda + \langle \mathbf{z}_j, \mathbf{R}_{-(ij)}\mathbf{z}_i \rangle / n) - \langle \mathbf{z}_i, \mathbf{R}_{-(ij)}\mathbf{z}_j \rangle^2 / n^2.$$

Since $\mathbf{z}_i, \mathbf{z}_j$, and $\mathbf{R}_{-(ij)}$ are mutually independent, by Lemma S.8.5 of [51], we have $\langle \mathbf{z}_i, \mathbf{R}_{-(ij)}\mathbf{z}_j \rangle / n \xrightarrow{\text{a.s.}} 0$ and the denominator concentrates on λ^2 . By Lemma G.5 (1) of [13], it follows that $\max_{1 \leq i \neq j \leq n} d_{ij} \xrightarrow{\text{a.s.}} \lambda^2$. Thus, there exists $N_0 \in \mathbb{N}$ such that for all $n \geq N_0$, $\max_{1 \leq i \neq j \leq n} d_{ij} \geq \lambda^2/2$ almost surely. Then, it follows that for $n \geq N_0$,

$$\begin{aligned} \left| \frac{1}{p} \sum_{1 \leq i \neq j \leq n} Y_{ij}g(\mathbf{z}_i)g(\mathbf{z}_j) \right| &\leq \frac{2}{\lambda^2} \frac{1}{n^2} \sum_{1 \leq i \neq j \leq n} |\langle \mathbf{z}_i, \mathbf{R}_{-(ij)}\mathbf{z}_j \rangle g(\mathbf{z}_i)g(\mathbf{z}_j)| \\ &= \frac{2}{\lambda^2} \frac{1}{n^2} \sum_{1 \leq i \neq j \leq n} (\mathbf{z}_i g(\mathbf{z}_i))^\top \mathbf{R}_{-(ij)} (\mathbf{z}_j g(\mathbf{z}_j)) \\ &\xrightarrow{\text{a.s.}} 0, \end{aligned}$$

where the last convergence is from Lemma S.8.5 of [51] and Lemma G.5 (2) of [13]. \square

Lemma D.4. For any two conforming matrices \mathbf{A} and \mathbf{B} and any $t \neq 0$, we have

$$\mathbf{A}^{-1}\mathbf{B}\mathbf{A}^{-1} = (\mathbf{A}^{-1} - (\mathbf{A} + t\mathbf{B})^{-1})/t + t\mathbf{A}^{-1}\mathbf{B}(\mathbf{A} + t\mathbf{B})^{-1}\mathbf{B}\mathbf{A}^{-1}.$$

Proof. We recall the Woodbury matrix identity:

$$(\mathbf{A} + \mathbf{UCV})^{-1} = \mathbf{A}^{-1} - \mathbf{A}^{-1}\mathbf{U}(\mathbf{C}^{-1} + \mathbf{VA}^{-1}\mathbf{U})^{-1}\mathbf{VA}^{-1}.$$

This holds for any conforming matrices $\mathbf{A}, \mathbf{U}, \mathbf{C}$, and \mathbf{V} . We will need to apply the Woodbury matrix identity twice below.

1. Applying the Woodbury identity for the first time with $\mathbf{A} = \mathbf{A}, \mathbf{U} = t, \mathbf{C} = \mathbf{B}$, and $\mathbf{V} = \mathbf{1}$, we get

$$(\mathbf{A} + t\mathbf{B})^{-1} = \mathbf{A}^{-1} - t\mathbf{A}^{-1}(\mathbf{B}^{-1} + t\mathbf{A}^{-1})^{-1}\mathbf{A}^{-1}.$$

Rearranging, this yields

$$(\mathbf{A}^{-1} - (\mathbf{A} + t\mathbf{B})^{-1})/t = \mathbf{A}^{-1}(\mathbf{B}^{-1} + t\mathbf{A}^{-1})^{-1}\mathbf{A}^{-1} = \mathbf{A}^{-1}(t\mathbf{A}^{-1} + \mathbf{B}^{-1})^{-1}\mathbf{A}^{-1}. \quad (38)$$

2. Applying the Woodbury identity the second time with $\mathbf{A} = t\mathbf{B}, \mathbf{U} = \mathbf{1}, \mathbf{C} = \mathbf{A}$, and $\mathbf{V} = \mathbf{1}$, we get

$$(\mathbf{A} + t\mathbf{B})^{-1} = t^{-1}\mathbf{B}^{-1} - t^{-1}\mathbf{B}^{-1}(\mathbf{A}^{-1} + t^{-1}\mathbf{B}^{-1})^{-1}t^{-1}\mathbf{B}^{-1}.$$

Multiplying by $t\mathbf{B}$ from the left yields

$$t\mathbf{B}(\mathbf{A} + t\mathbf{B})^{-1} = \mathbf{I} - (\mathbf{A}^{-1} + t^{-1}\mathbf{B}^{-1})^{-1}t^{-1}\mathbf{B}^{-1} = \mathbf{I} - (t\mathbf{A}^{-1} + \mathbf{B}^{-1})^{-1}\mathbf{B}^{-1}.$$

Now, multiplying by \mathbf{B} on the right, notice that

$$t\mathbf{B}(\mathbf{A} + t\mathbf{B})^{-1}\mathbf{B} = \mathbf{B} - (t\mathbf{A}^{-1} + \mathbf{B}^{-1})^{-1}. \quad (39)$$

From (38), observe that

$$\begin{aligned} \mathbf{A}^{-1} \mathbf{B} \mathbf{A}^{-1} - (\mathbf{A}^{-1} - (\mathbf{A} + t\mathbf{B})^{-1})/t &= \mathbf{A}^{-1} \mathbf{B} \mathbf{A}^{-1} - \mathbf{A}^{-1} (t\mathbf{A}^{-1} + \mathbf{B}^{-1})^{-1} \mathbf{A}^{-1} \\ &= \mathbf{A}^{-1} (\mathbf{B} - (t\mathbf{A}^{-1} + \mathbf{B}^{-1})^{-1}) \mathbf{A}^{-1}. \end{aligned}$$

Using (39), we then have

$$\mathbf{A}^{-1} \mathbf{B} \mathbf{A}^{-1} - (\mathbf{A}^{-1} - (\mathbf{A} + t\mathbf{B})^{-1})/t = t\mathbf{A}^{-1} \mathbf{B} (\mathbf{A} + t\mathbf{B})^{-1} \mathbf{B} \mathbf{A}^{-1}.$$

Rearranging, we arrive at the desired equality:

$$\mathbf{A}^{-1} \mathbf{B} \mathbf{A}^{-1} = (\mathbf{A}^{-1} - (\mathbf{A} + t\mathbf{B})^{-1})/t + t\mathbf{A}^{-1} \mathbf{B} (\mathbf{A} + t\mathbf{B})^{-1} \mathbf{B} \mathbf{A}^{-1}.$$

□

D.3 Concentration of energy of linear and nonlinear components

Lemma D.5. *Under Assumptions 1–2, the best linear estimator β_0 and the nonlinear component f_{NL} as defined in (10) satisfies that $\|\beta_0\|_2$ and $\|f_{\text{NL}}\|_{L_{4+\delta}}$ are bounded almost surely.*

Proof. By Jensen's inequality and Assumption 1, we have $\mathbb{E}[y^2] \leq \mathbb{E}[y^{4+\delta}]^{2/(4+\delta)} \leq C_0^{2/(4+\delta)}$ for some constant $C_0 > 0$.

By the orthogonality, we have that $\mathbb{E}[y^2] = \mathbb{E}[(\mathbf{x}^\top \beta_0)^2] + \|f_{\text{NL}}\|_{L_2}^2$, which implies that $\mathbb{E}[(\mathbf{x}^\top \beta_0)^2]$ and $\|f_{\text{NL}}\|_{L_2}^2$ is bounded. Since by Assumption 2, the eigenvalues of Σ is lower bounded by $r_{\min} > 0$, we have that $\mathbb{E}[(\mathbf{x}^\top \beta_0)^2] = \beta_0^\top \Sigma \beta_0 \geq r_{\min} \|\beta_0\|_2^2$ and thus $\|\beta_0\|_2$ is bounded.

Since $f_{\text{NL}} = y - \mathbf{x}^\top \beta_0$, by triangle inequality we have that

$$\begin{aligned} \|f_{\text{NL}}\|_{L_{4+\delta}} &\leq \|y\|_{L_{4+\delta}} + \|\mathbf{x}^\top \beta_0\|_{L_{4+\delta}} \\ &\leq \|y\|_{L_{4+\delta}} + C\|\beta_0\|_2 \end{aligned}$$

for some constant $C > 0$. The second inequality of the above display is from Lemma 7.8 of [57]. Since $\|y\|_{L_{4+\delta}}$ and $\|\beta_0\|_2$ are bounded, it follows that $\|f_{\text{NL}}\|_{L_{4+\delta}}$ is also bounded. □

E Asymptotic equivalents: background and known results

Preliminary background. In several proofs, we employ the concept of asymptotic equivalence of sequences of random matrices; see [12, 44, 49, 51]. This section provides a brief overview of the associated terminology and calculus principles.

Definition E.1 (Asymptotic equivalence). *Consider sequences $\{\mathbf{A}_p\}_{p \geq 1}$ and $\{\mathbf{B}_p\}_{p \geq 1}$ of (random or deterministic) matrices of growing dimensions. We say that \mathbf{A}_p and \mathbf{B}_p are equivalent and write $\mathbf{A}_p \simeq \mathbf{B}_p$ if $\lim_{p \rightarrow \infty} |\text{tr}[\mathbf{C}_p(\mathbf{A}_p - \mathbf{B}_p)]| = 0$ almost surely for every sequence of random matrices \mathbf{C}_p independent to \mathbf{A}_p and \mathbf{B}_p , with bounded trace norm such that $\limsup_{p \rightarrow \infty} \|\mathbf{C}_p\|_{\text{tr}} < \infty$ almost surely.*

The notion of asymptotic equivalence of two sequences of random matrices above can be further extended to incorporate conditioning on another sequence of random matrices; see [12] for more details.

Definition E.2 (Conditional asymptotic equivalence). *Consider sequences $\{\mathbf{A}_p\}_{p \geq 1}$, $\{\mathbf{B}_p\}_{p \geq 1}$ and $\{\mathbf{D}_p\}_{p \geq 1}$ of (random or deterministic) matrices of growing dimensions. We say that \mathbf{A}_p and \mathbf{B}_p are equivalent given \mathbf{D}_p and write $\mathbf{A}_p \simeq \mathbf{B}_p \mid \mathbf{D}_p$ if $\lim_{p \rightarrow \infty} |\text{tr}[\mathbf{C}_p(\mathbf{A}_p - \mathbf{B}_p)]| = 0$ almost surely conditional on $\{\mathbf{D}_p\}_{p \geq 1}$, i.e.,*

$$\mathbb{P} \left(\lim_{p \rightarrow \infty} |\text{tr}[\mathbf{C}_p(\mathbf{A}_p - \mathbf{B}_p)]| = 0 \mid \{\mathbf{D}_p\}_{p \geq 1} \right) = 1,$$

for any sequence of random matrices \mathbf{C}_p independent to \mathbf{A}_p and \mathbf{B}_p conditional on \mathbf{D}_p , with bounded trace norm such that $\limsup_{p \rightarrow \infty} \|\mathbf{C}_p\|_{\text{tr}} < \infty$ as $p \rightarrow \infty$.

Below we summarize the calculus rules for conditional asymptotic equivalence in Definition E.2. These are adapted from Lemma S.7.4 and S.7.6 of [12].

Lemma E.3 (Calculus of asymptotic equivalents). *Let $\mathbf{A}_p, \mathbf{B}_p, \mathbf{C}_p$ and \mathbf{D}_p be sequences of random matrices. The calculus of asymptotic equivalents satisfies the following properties:*

- (1) *Equivalence: The relation \simeq is an equivalence relation.*
- (2) *Sum: If $\mathbf{A}_p \simeq \mathbf{B}_p \mid \mathbf{E}_p$ and $\mathbf{C}_p \simeq \mathbf{D}_p \mid \mathbf{E}_p$, then $\mathbf{A}_p + \mathbf{C}_p \simeq \mathbf{B}_p + \mathbf{D}_p \mid \mathbf{E}_p$.*
- (3) *Product: If \mathbf{A}_p has bounded operator norms such that $\limsup_{p \rightarrow \infty} \|\mathbf{A}_p\|_{\text{op}} < \infty$, \mathbf{A}_p is conditional independent to \mathbf{B}_p and \mathbf{C}_p given \mathbf{E}_p for $p \geq 1$, and $\mathbf{B}_p \simeq \mathbf{C}_p \mid \mathbf{E}_p$, then $\mathbf{A}_p \mathbf{B}_p \simeq \mathbf{A}_p \mathbf{C}_p \mid \mathbf{E}_p$.*
- (4) *Trace: If $\mathbf{A}_p \simeq \mathbf{B}_p \mid \mathbf{E}_p$, then $\text{tr}[\mathbf{A}_p]/p - \text{tr}[\mathbf{B}_p]/p \rightarrow 0$ almost surely when conditioning on \mathbf{E}_p .*
- (5) *Differentiation: Suppose $f(z, \mathbf{A}_p) \simeq g(z, \mathbf{B}_p) \mid \mathbf{E}_p$ where the entries of f and g are analytic functions in $z \in S$ and S is an open connected subset of \mathbb{C} . Suppose for any sequence \mathbf{C}_p of deterministic matrices with bounded trace norm we have $|\text{tr}[\mathbf{C}_p(f(z, \mathbf{A}_p) - g(z, \mathbf{B}_p))]| \leq M$ for every p and $z \in S$. Then, we have $f'(z, \mathbf{A}_p) \simeq g'(z, \mathbf{B}_p) \mid \mathbf{E}_p$ for every $z \in S$, where the derivatives are taken entrywise with respect to z .*
- (6) *Unconditioning: If $\mathbf{A}_p \simeq \mathbf{B}_p \mid \mathbf{E}_p$, then $\mathbf{A}_p \simeq \mathbf{B}_p$.*
- (7) *Substitution: Let $v : \mathbb{R}^{p \times p} \rightarrow \mathbb{R}$ and $f(v(\mathbf{C}), \mathbf{C}) : \mathbb{R}^{p \times p} \rightarrow \mathbb{R}^{p \times p}$ be a matrix function for matrix $\mathbf{C} \in \mathbb{R}^{p \times p}$ and $p \in \mathbb{N}$, that is continuous in the first argument with respect to operator norm. If $v(\mathbf{C}) \stackrel{\text{a.s.}}{=} v(\mathbf{D})$ such that \mathbf{C} is independent to \mathbf{D} , then $f(v(\mathbf{C}), \mathbf{C}) \simeq f(v(\mathbf{D}), \mathbf{C}) \mid \mathbf{C}$.*

Standard ridge resolvents and various extensions. In this section, we collect various asymptotic equivalents. The following corollary is a simple consequence of Theorem 1 in [58]. It provides deterministic equivalent for the scaled ridge resolvent.

Corollary E.4 (Deterministic equivalent for scaled ridge resolvent). *Suppose $\mathbf{x}_i \in \mathbb{R}^p$, $1 \leq i \leq n$, are i.i.d. random vectors such that each $\mathbf{x}_i = \mathbf{z}_i \Sigma^{1/2}$, where \mathbf{z}_i is a random vector consisting of i.i.d. entries z_{ij} , $1 \leq j \leq p$, satisfying $\mathbb{E}[z_{ij}] = 0$, $\mathbb{E}[z_{ij}^2] = 1$, and $\mathbb{E}[|z_{ij}|^{8+\alpha}] \leq M_\alpha$ for some constants $\alpha > 0$ and $M_\alpha < \infty$, and $\Sigma \in \mathbb{R}^{p \times p}$ is a positive semidefinite matrix satisfying $0 \preceq \Sigma \preceq r_{\max} \mathbf{I}_p$ for some constant $r_{\max} < \infty$ (independent of p). Let $\mathbf{X} \in \mathbb{R}^{n \times p}$ the concatenated matrix with \mathbf{x}_i^\top , $1 \leq i \leq n$, as rows, and let $\widehat{\Sigma} \in \mathbb{R}^{p \times p}$ denote the random matrix $\mathbf{X}^\top \mathbf{X}/n$. Let $\gamma = p/n$. Then, for $z \in \mathbb{C}^+$, as $n, p \rightarrow \infty$ such that $0 < \liminf \gamma \leq \limsup \gamma < \infty$, and $\lambda > 0$, we have*

$$\lambda(\widehat{\Sigma} + \lambda \mathbf{I}_p)^{-1} \simeq (v(-\lambda; \gamma) \Sigma + \mathbf{I}_p)^{-1},$$

where $v(-\lambda; \gamma) > 0$ is the unique solution to the fixed-point equation

$$\frac{1}{v(-\lambda; \gamma)} = \lambda + \gamma \int \frac{r}{1 + v(-\lambda; \gamma)r} dH_p(r). \quad (40)$$

Here H_n is the empirical distribution (supported on $\mathbb{R}_{\geq 0}$) of the eigenvalues of Σ .

Side remark: The parameter $v(-\lambda; \gamma)$ in Corollary E.4 is also the companion Stieltjes transform of the spectral distribution of the sample covariance matrix $\widehat{\Sigma}$. It is also the Stieltjes transform of the spectral distribution of the gram matrix $\mathbf{X} \mathbf{X}^\top / n$. This is essentially the characterization we use in proving Proposition 4.

The following lemma uses Corollary E.4 along with calculus of deterministic equivalents (from Lemma E.3), and provides deterministic equivalents for resolvents needed to obtain asymptotic bias and variance of standard ridge regression. It is adapted from Lemma S.6.10 of [51].

Lemma E.5 (Deterministic equivalents for ridge resolvents associated with generalized bias and variance). *Suppose $\mathbf{x}_i \in \mathbb{R}^p$, $1 \leq i \leq n$, are i.i.d. random vectors with each $\mathbf{x}_i = \mathbf{z}_i \Sigma^{1/2}$, where $\mathbf{z}_i \in \mathbb{R}^p$ is a random vector that contains i.i.d. random variables z_{ij} , $1 \leq j \leq p$, each with $\mathbb{E}[z_{ij}] = 0$, $\mathbb{E}[z_{ij}^2] = 1$, and $\mathbb{E}[|z_{ij}|^{4+\alpha}] \leq M_\alpha$ for some constants $\alpha > 0$ and $M_\alpha < \infty$, and $\Sigma \in \mathbb{R}^{p \times p}$ is a positive semidefinite matrix with $r_{\min} \mathbf{I}_p \preceq \Sigma \preceq r_{\max} \mathbf{I}_p$ for some constants $r_{\min} > 0$ and $r_{\max} < \infty$ (independent of p). Let $\mathbf{X} \in \mathbb{R}^{n \times p}$ be the concatenated random matrix with \mathbf{x}_i , $1 \leq i \leq n$, as its rows, and define $\widehat{\Sigma} = \mathbf{X}^\top \mathbf{X}/n \in \mathbb{R}^{p \times p}$. Let $\gamma = p/n$. Then, for $\lambda > 0$, as $n, p \rightarrow \infty$ with $0 < \liminf \gamma \leq \limsup \gamma < \infty$, the following statements hold:*

(1) Bias of ridge regression:

$$\lambda^2(\widehat{\Sigma} + \lambda \mathbf{I}_p)^{-1} \mathbf{A}(\widehat{\Sigma} + \lambda \mathbf{I}_p)^{-1} \\ \simeq (v(-\lambda; \gamma, \Sigma) \Sigma + \mathbf{I}_p)^{-1} (\tilde{v}_b(-\lambda; \gamma, \Sigma, \mathbf{A}) \Sigma + \mathbf{A}) (v(-\lambda; \gamma, \Sigma) \Sigma + \mathbf{I}_p)^{-1}.$$

(2) Variance of ridge regression:

$$(\widehat{\Sigma} + \lambda \mathbf{I}_p)^{-2} \widehat{\Sigma} \mathbf{A} \simeq \tilde{v}_v(-\lambda; \gamma, \Sigma) (v(-\lambda; \gamma, \Sigma) \Sigma + \mathbf{I}_p)^{-2} \Sigma \mathbf{A}.$$

Here $v(-\lambda; \gamma, \Sigma) > 0$ is the unique solution to the fixed-point equation

$$\frac{1}{v(-\lambda; \gamma, \Sigma)} = \lambda + \int \frac{\gamma r}{1 + v(-\lambda; \gamma, \Sigma) r} dH_n(r; \Sigma), \quad (41)$$

and $\tilde{v}_b(-\lambda; \gamma, \Sigma)$ and $\tilde{v}_v(-\lambda; \gamma, \Sigma)$ are defined through $v(-\lambda; \gamma, \Sigma)$ by the following equations:

$$\tilde{v}_b(-\lambda; \gamma, \Sigma, \mathbf{A}) = \frac{\gamma \operatorname{tr}[\mathbf{A} \Sigma (v(-\lambda; \gamma, \Sigma) \Sigma + \mathbf{I}_p)^{-2}] / p}{v(-\lambda; \gamma, \Sigma)^{-2} - \int \gamma r^2 (1 + v(-\lambda; \gamma, \Sigma) r)^{-2} dH_n(r; \Sigma)}, \quad (42)$$

$$\tilde{v}_v(-\lambda; \gamma, \Sigma)^{-1} = v(-\lambda; \gamma, \Sigma)^{-2} - \int \gamma r^2 (1 + v(-\lambda; \gamma, \Sigma) r)^{-2} dH_n(r; \Sigma), \quad (43)$$

where $H_n(\cdot; \Sigma)$ is the empirical distribution (supported on $[r_{\min}, r_{\max}]$) of the eigenvalues of Σ .

Though Lemma E.5 states the dependency explicitly, we will simply write $H_p(r)$, $v(-\lambda; \gamma)$, $\tilde{v}_b(-\lambda; \gamma, \mathbf{A})$, and $\tilde{v}_v(-\lambda; \gamma)$ to denote $H_n(r; \Sigma)$, $v(-\lambda; \gamma, \Sigma)$, $\tilde{v}_b(-\lambda; \gamma, \Sigma, \mathbf{A})$, and $\tilde{v}_v(-\lambda; \gamma, \Sigma)$, respectively, for simplifying various notations when it is clear from the context. When $\mathbf{A} = \Sigma$, we simply write $\tilde{v}_b(-\lambda; \gamma) = \tilde{v}_b(-\lambda; \gamma, \mathbf{A})$.

F Experiment details

F.1 Reproducibility and compute details

The source code for generating all of our figures is included with the supplementary material. The source code also includes details about the computational resources used to run the code and other timing details.

F.2 Simulation details

The covariance matrix of an auto-regressive process of order 1 (AR(1)) is given by Σ_{ar1} , where $(\Sigma_{\text{ar1}})_{ij} = \rho_{\text{ar1}}^{|i-j|}$ for some parameter $\rho_{\text{ar1}} \in (0, 1)$. Define $\beta_0 := \frac{1}{5} \sum_{j=1}^5 \mathbf{w}_{(j)}$ where $\mathbf{w}_{(j)}$ is the eigenvector of Σ_{ar1} associated with the top j th eigenvalue $r_{(j)}$. We generated data (\mathbf{x}_i, y_i) for $i = 1, \dots, n$ from a nonlinear model:

$$y_i = \mathbf{x}_i^\top \beta_0 + \frac{1}{p} (\|\mathbf{x}_i\|_2^2 - \operatorname{tr}[\Sigma_{\text{ar1}}]) + \varepsilon_i, \quad \mathbf{x}_i = \Sigma_{\text{ar1}}^{\frac{1}{2}} \mathbf{z}_i, \quad z_{ij} \stackrel{iid}{\sim} \frac{t_5}{\sigma_5}, \quad \varepsilon_i \sim \frac{t_5}{\sigma_5}, \quad (\text{M-AR1})$$

where $\sigma_5 = \sqrt{5/3}$ is the standard deviation of t_5 distribution. The benefit of using the above nonlinear model is that we can clearly separate the linear and the nonlinear components and compute the quantities of interest because β_0 happens to be the best linear projection.

For the simulations, we set $\rho_{\text{ar1}} = 0.5$. For finite ensembles, the risks are averaged across 50 simulations.

F.3 Risk comparisons along equivalence paths for finite ensembles

Figure F5 verifies the following linear (in M) relationship along the path in Theorem 3 for finite ensembles: $R(\hat{\beta}_{\lfloor p/\psi_1 \rfloor, M}^{\lambda_1}; \mathbf{A}, \mathbf{b}, \beta_0) - R(\hat{\beta}_{\lfloor p/\psi_2 \rfloor, M}^{\lambda_2}; \mathbf{A}, \mathbf{b}, \beta_0) \simeq \Delta/M$, for some Δ (independent of M), which is eventually almost surely bounded.

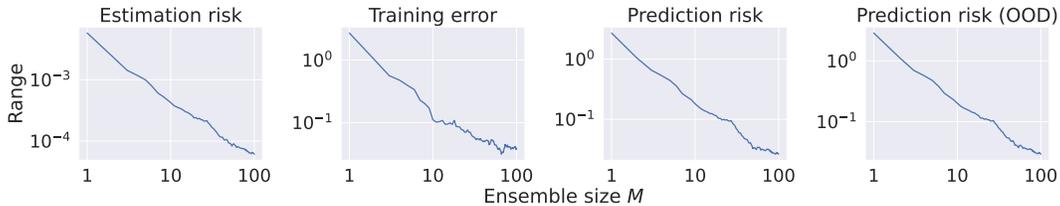


Figure F5: The range (the difference between the maximum and the maximum along the path) of the finite-sample generalized risk of M -ensemble ridge estimators on the path through $(\lambda, \psi) = (0, 2)$, for varying ensemble size M under the same setting as in Figure 1. The generalized risks include the estimation risk, the training error, the prediction risk, and the out-of-distribution (OOD) prediction risk.

F.4 Illustration of risk monotonicity of optimal ridge

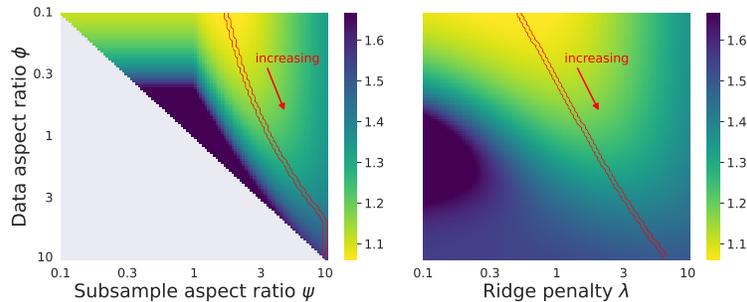


Figure F6: Illustration of risk monotonicity of optimal subsampled ridgeless regression versus optimal regularized ridge regression. The underlying model has a non-isotropic covariance as described in Appendix F.2 with a linearized signal-to-noise ratio of 0.6 and $\|\mathbf{f}_{\text{NL}}\|_{L^2} = 1$. This results in a null risk of 1.6, which is kept the same across all regression problems. The left panel shows the limiting risk of the full-ensemble ridgeless regression at various data and subsample aspect ratios (ϕ, ψ) . The right panel shows the limiting risk of the ridge predictor (on the full data) at various data aspect ratios and regularization penalties (ϕ, λ) . Optimal risks for each data aspect ratio are highlighted using slender red lines in both panels. Observe that the optimal risk in both cases is increasing as a function of ϕ . In the left panel, in addition, for every subsample aspect ratio ψ , the risk of subsampled ridgeless is increasing in the data aspect ratio ϕ . This in turn implies the risk monotonicity claim for the optimal ridge regression in Theorem 6 through a slick triangle inequality argument (see the proof in Appendix C for more details).

F.5 Real-world datasets

We conduct experiments on real-world datasets to examine the equivalence in a more general setting. We utilized three image datasets for our experimental analysis: CIFAR 10, MNIST, and USPS [59]. For the CIFAR 10 dataset, we subset the images labeled as “dog” and “cat”. For other datasets, we subset the images labeled “3” and “8”. Then we treat them as binary labels $y \in \{0, 1\}$ and use the flattened image as our feature vector \mathbf{x} . The training sample sizes, the feature dimensions, and the test sample sizes (n, p, n_{te}) are (10000, 3072, 2000), (12873, 784, 2145), and (22358, 3072, 7981) for the three datasets, respectively.

For the first experiment in Figure 3, we fix $\phi = p/n$ and $\bar{\psi} = 4\phi$ using the training set of CIFAR-10. For $\lambda_{\bar{\psi}} \in \{0.01, 0.05, 0.1, 1\}$, we compute the data-dependent value of $\bar{\lambda}$ based on Proposition 4. Each value of $\lambda_{\bar{\psi}}$ gives a path between $(\lambda_{\bar{\psi}}, \bar{\psi})$ and $(\bar{\lambda}, \phi)$.

For the second experiment in Figure F7, by varying the values of subsample aspect ratios $\bar{\psi}$, we compare the random Gaussian linear projection of the estimators and the prediction risk at the two endpoints $(\bar{\lambda}, \phi)$ and $(0, \bar{\psi})$, on CIFAR-10, MNIST, and USPS datasets.

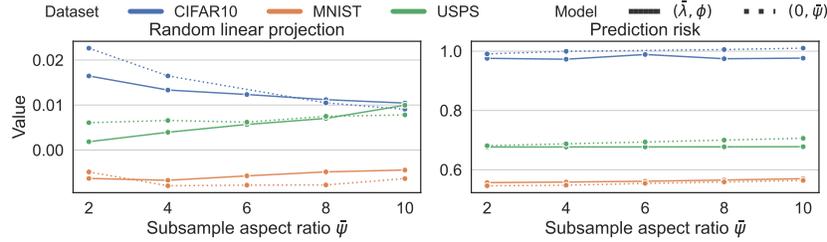


Figure F7: The functional equivalences of subsampling and ridge regularization on different datasets. For each dataset with an aspect ratio $\phi = p/n$ and each value of $\bar{\psi}$, the corresponding ridge penalty $\bar{\lambda}$ is estimated by Proposition 4. The two models with ridge penalty and subsample aspect ratio $(\bar{\lambda}, \phi)$ and $(0, \bar{\psi})$ are compared. Note that the model corresponding to $(\bar{\lambda}, \phi)$ is ridge regression at ridge penalty $\bar{\lambda}$ without subsampling, and the model corresponding to $(0, \bar{\psi})$ is full-ensemble ridgeless at a subsample aspect ratio $\bar{\psi}$.

F.6 Random features model

The data (\mathbf{x}_i, y_i) is generate from the nonlinear model

$$y_i = \mathbf{x}_i^\top \beta_0 + \frac{1}{p} (\|\mathbf{x}_i\|_2^2 - p) + \varepsilon_i,$$

where $x_{ij} \stackrel{iid}{\sim} \mathcal{N}(0, 1)$, and $\varepsilon_i \sim \mathcal{N}(0, 1)$. The prediction risk of the ridge ensemble ($M = 100$) is computed based on the random feature $\varphi(\mathbf{F}\mathbf{x}_i)$ and the response y_i , where $\mathbf{F} \in \mathbb{R}^{d \times p}$ is the random weight matrix with $F_{ij} \stackrel{iid}{\sim} \mathcal{N}(0, p^{-1})$. Here, φ is a nonlinear activation function (sigmoid, ReLU, or tanh) that is applied entry-wise to $\mathbf{F}\mathbf{x}_i$. For the experiment, we set $p = 250$, $d = 500$, and $\phi = d/n = 0.1$.

F.7 Kernel ridge regression

For a given feature map $\Phi : \mathbb{R}^p \rightarrow \mathbb{R}^d$, the kernel ridge estimator is defined as:

$$\hat{\beta}_k^\lambda(\mathcal{D}_I) = \operatorname{argmin}_{\beta \in \mathbb{R}^p} \sum_{i \in I} (k^{-1/2} y_i - k^{-1/2} \Phi(\mathbf{x}_i)^\top \beta)^2 + \lambda \|\beta\|_2^2 + \frac{k}{p} \lambda \|\beta\|_2^2.$$

By using the kernel trick, the above optimization problem is equivalent to solving:

$$\hat{\alpha}_k^\lambda(\mathcal{D}_I) = \operatorname{argmin}_{\alpha \in \mathbb{R}^k} \alpha^\top (\mathbf{K}_I + k\lambda \mathbf{I}_k) \alpha + \alpha^\top \mathbf{y}_I,$$

where $\mathbf{K}_I = \Phi_I \Phi_I^\top \in \mathbb{R}^{k \times k}$ is the kernel matrix and $\Phi_I = (\Phi(\mathbf{x}_i))_{i \in I} \in \mathbb{R}^{n \times d}$ is the feature matrix. Simple calculation shows that $\hat{\beta}_k^\lambda(\mathcal{D}_I) = \Phi_I^\top \hat{\alpha}_k^\lambda(\mathcal{D}_I)$.

Using the same data-generating process as in the previous subsection, the results on kernel ridge regression are illustrated in Figure F8.

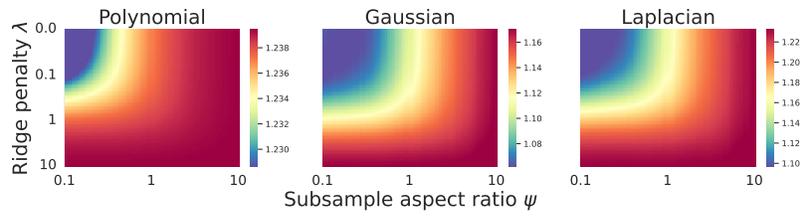


Figure F8: Heat map of the empirical distribution prediction risk of full-ensemble kernel ridge estimators, for varying ridge penalties λ and subsample aspect ratio $\psi = p/d$ on the log-log scale. The prediction risk of the ridge ensemble ($M = 100$) is computed using polynomial, Gaussian, and Laplacian kernel, using the default parameters as in Python package `scikit-learn` v1.2.2 [60] without the intercept terms. We scale λ by k/p so that we can obtain non-null estimators for the polynomial inner-product kernel (of degree 3).