

Optimal Ridge Regularization for Out-of-Distribution Prediction

Pratik Patil[†]
pratikpatil@berkeley.edu

Jin-Hong Du^{‡§}
jinhongd@andrew.cmu.edu

Ryan J. Tibshirani[†]
ryantibs@berkeley.edu

Abstract

We study the behavior of optimal ridge regularization and optimal ridge risk for out-of-distribution prediction, where the test distribution deviates arbitrarily from the train distribution. We establish general conditions that determine the sign of the optimal regularization level under covariate and regression shifts. These conditions capture the alignment between the covariance and signal structures in the train and test data and reveal stark differences compared to the in-distribution setting. For example, a negative regularization level can be optimal under covariate shift or regression shift, even when the training features are isotropic or the design is underparameterized. Furthermore, we prove that the optimally tuned risk is monotonic in the data aspect ratio, even in the out-of-distribution setting and when optimizing over negative regularization levels. In general, our results do not make any modeling assumptions for the train or the test distributions, except for moment bounds, and allow for arbitrary shifts and the widest possible range of (negative) regularization levels.

1 Introduction

Regularization plays a crucial role in statistical modeling and is commonly incorporated into optimization-based models through a regularization term. Its effectiveness relies on properly scaling the regularization term, which is controlled by a penalty parameter that the data scientist needs to tune. Recent work in machine learning (precise references given shortly) has shed light on some rather surprising behavior exhibited by the optimal regularization level in overparameterized prediction models, which can be zero or even negative in certain problems with moderate signal-to-noise ratio and high dimensionality. This stands in contrast to the “typical” behavior from classical low-dimensional learning theory.

With this motivation, our paper focuses on two key questions for high-dimensional ridge regression:

- (Q1) What is the behavior of the *optimal ridge penalty*, as a function of parameters such as signal-to-noise ratio, data aspect ratio, feature correlations, and signal structure?
- (Q2) What is the behavior of the *optimally tuned ridge risk*, as a function of these same problem parameters?

To set the notation, let (x_i, y_i) for $i \in [n]$ be i.i.d. observations in $\mathbb{R}^p \times \mathbb{R}$ representing the feature vector and response. Denote the data matrix as $X = [x_1, \dots, x_n]^\top \in \mathbb{R}^{n \times p}$ and the associated response vector as $y = [y_1, \dots, y_n]^\top \in \mathbb{R}^n$. Given a ridge penalty $\lambda > 0$, recall the ridge regression fits:

$$\hat{\beta}^\lambda = \operatorname{argmin}_{b \in \mathbb{R}^p} \|y - Xb\|_2^2/n + \lambda \|b\|_2^2. \quad (1)$$

[†]Department of Statistics, University of California, Berkeley, CA 94720, USA.

[‡]Department of Statistics and Data Science, Carnegie Mellon University, Pittsburgh, PA 15213, USA.

[§]Machine Learning Department, Carnegie Mellon University, Pittsburgh, PA 15213, USA.

Ridge regression (Hoerl and Kennard, 1970a,b) has received considerable recent attention, particularly in settings involving overparameterization, such as double descent (see, e.g., Belkin et al., 2020; Hastie et al., 2022; Muthukumar et al., 2020, and references therein) and benign overfitting (Bartlett et al., 2020; Koehler et al., 2021; Mallinar et al., 2022). This interest in ridge regression, especially its “ridgeless” limit, where $\lambda \rightarrow 0^+$, owes to its peculiar double/multiple descent risk behavior in overparameterized regimes, which (on the surface) challenges the conventional understanding of the role of regularization (Hastie, 2020). By defining the ridge estimator as:

$$\widehat{\beta}^\lambda = (X^\top X/n + \lambda I_p)^\dagger X^\top y/n = X^\top (XX^\top/n + \lambda I_n)^\dagger y/n, \quad (2)$$

where A^\dagger denotes the Moore-Penrose pseudoinverse of A , we simultaneously accommodate the case of $\lambda > 0$, in which case (2) reduces to the unique ridge predictor obtained using (1), and the case of $\lambda = 0$, in which case (2) becomes the minimum ℓ_2 -norm interpolator among many solutions to problem (1) when $p \geq \text{rank}(X) = n$. Note that the above definition (2) is well defined even when $\lambda < 0$. For more background on the formulation of ridge regression with $\lambda < 0$, see Appendix B.

Partial answers to questions (Q1) and (Q2) are known for the *in-distribution* squared prediction risk, defined as:

$$R(\widehat{\beta}^\lambda) = \mathbb{E}_{x_0, y_0} [(y_0 - x_0^\top \widehat{\beta}^\lambda)^2 \mid X, y], \quad (3)$$

where (x_0, y_0) is a test point sampled independently from the *same* distribution $P_{x,y}$ as the training data. Note that the prediction risk (3) is conditional on the training data and is a function of (X, y) and the properties of $P_{x,y}$. Arguably, the cleanest answers to (Q1) and (Q2) are obtained under proportional asymptotics where the feature size p and the sample size n diverge proportionally to an *aspect ratio* $\phi \in (0, \infty)$. Under certain additional assumptions, the risk (3) then almost surely converges to a limit $\mathcal{R}(\lambda, \phi)$ that depends only on coarse properties of $P_{x,y}$. Analyzing the behavior of the optimal ridge penalty λ^* (which minimizes $\mathcal{R}(\lambda, \phi)$ over λ) and the optimal risk $\mathcal{R}(\lambda^*, \phi)$ then consequently allows us to answer questions (Q1) and (Q2), respectively.

For (Q1), consider a well-specified linear model $y = X\beta + \varepsilon$, where the noise $\varepsilon \sim (0_n, \sigma^2 I)$ is independent of X , and the signal is random with $\beta \sim (0_p, (\alpha^2/p)I)$. Remarkably, despite the lack of a closed-form expression for $\mathcal{R}(\lambda^*, \phi)$, Dobriban and Wager (2018) show that $\lambda^* = \phi/\text{SNR} > 0$, where $\text{SNR} = \alpha^2/\sigma^2$ is the signal-to-noise ratio. However, in real-data experiments, it has been observed that a negative ridge penalty can be optimal (Kobak et al., 2020). Motivated by this, Wu and Xu (2020); Richards et al. (2021) analyze the sign behavior of λ^* beyond random isotropic signals and establish sufficient conditions for when $\lambda^* < 0$ or $\lambda^* = 0$.

For (Q2), again remarkably, it follows from the results of Dobriban and Wager (2018) that for isotropic features and random isotropic signals, the risk $\mathcal{R}(\lambda^*, \phi)$ increases monotonically with the data aspect ratio ϕ . Recent work by Patil and Du (2023, Theorem 6) extends this result to anisotropic features and deterministic signals (with arbitrary response distributions of bounded moments), demonstrating that optimal ridge regression exhibits a monotonic risk profile and avoids double and multiple descents. In a certain sense, these results imply the sample-wise monotonicity of the optimal expected conditional risk (over the training data set).

We work to answer (Q1) and (Q2) more comprehensively across essentially all possible in-distribution (IND) ridge prediction problems. Furthermore, we will generalize this by also considering the *out-of-distribution* (OOD) setting, where the test point (x_0, y_0) in (3) has a different distribution P_{x_0, y_0} than the train distribution $P_{x,y} = P_x \cdot P_{y|x}$. Distribution shift occurs in many practical machine learning applications and has gained increasing attention in the learning community. We focus on two common types of distribution shifts:

- (i) *Covariate shift*: where $P_{x_0} \neq P_x$ but $P_{y_0|x_0} = P_{y|x}$.
- (ii) *Regression shift*: where $P_{y_0|x_0} \neq P_{y|x}$ but $P_{x_0} = P_x$.

Thus, we also study following generalizations of (Q1) and (Q2):

- (Q1') *How does distribution shift alter optimal regularization?* Specifically, under what types of shift is the optimal penalty λ^* in the OOD setting notably different (typically smaller) compared to the IND setting?
- (Q2') *How does distribution shift alter optimal risk behavior?* Specifically, is $\mathcal{R}(\lambda^*, \phi)$ still monotonic in ϕ when λ^* changes due to the distribution shift? Conversely, is optimal regularization *necessary* to ensure monotonic risk?

1.1 Summary of Results and Paper Outline

Extended risk characterization. In Section 2, we extend the scope of risk characterization for ridge regression for the out-of-distribution setting (Proposition 1) that: (i) does not assume any model for the train or the test distribution, apart from certain moment bounds on the train and test response distributions, (ii) does not assume that the spectrums of feature covariance or signal projections converge to fixed distributions, and (iii) allows for the widest possible range of (negative) regularization level λ_{\min} (see Definition 2).

Properties of optimal regularization. In Section 3, we characterize the conditions that determine the sign of λ^* under covariate shift (Theorem 4) and regression shift (Theorem 5). These conditions capture the general alignment between the signal and the covariance spectrum, which isolates the cases where the sign of λ^* under the OOD setting changes from the IND setting. Our work subsumes and extends previously known results on optimal ridge regularization to the best of our knowledge (see Table 1 for precise comparisons).

Properties of optimal risk. In Section 4, we show that the OOD risk of the optimally tuned ridge $\mathcal{R}(\lambda^*, \phi)$ is monotonic in the data aspect ratio ϕ and SNR (Theorem 7). Furthermore, we establish a partial converse (Theorem 8) that shows risk non-monotonicity under suboptimal regularization. To prove our results, we exploit the equivalences between subsampling and ridge regression from Patil and Du (2023) to the OOD setting and also allow for tuning over negative regularization (Theorem 9).

1.2 Related Work and Comparisons

Ridge risk characterization. The asymptotic risk of ridge regression has been extensively studied in the literature under proportional asymptotics when $p/n \rightarrow \phi \in (0, \infty)$, as $n, p \rightarrow \infty$ using tools from random matrix theory and statistical physics. For well-specified linear models, expressions of risk asymptotics for the IND setting are obtained by Dobriban and Wager (2018); Hastie et al. (2022), among others. Historically, heuristic derivations of these expressions have also been derived for Gaussian process regression by Sollich (2001). Additionally, several works have explored risk asymptotics and its implications in different variants of ridge regression (Wei et al., 2022; Mel and Ganguli, 2021; Loureiro et al., 2021; Jacot et al., 2020; Simon et al., 2021; Zhou et al., 2023; Bach, 2024; Pesce et al., 2023). The risk asymptotics for the OOD setting are obtained by Canatar et al. (2021), D'Amour et al. (2022, Section E.5), Patil et al. (2022, Section S.6.5), Tripuraneni et al. (2021). However, these works assume either random Gaussian features or a well-specified linear model or restrict to only the positive range of regularization. Our work extends this literature by allowing for general response models and the widest possible range of negative regularization.

Table 1: **Optimal regularization landscape in ridge regression.** Here, ‘ \circ ’ indicates either an isotropic feature or signal covariance, and ‘ \otimes ’ indicates anisotropic features or signal covariance. For the data aspect ratio ϕ , ‘all’ indicates $\phi \in (0, \infty)$, ‘under’ indicates $\phi \in (0, 1)$ for the underparameterized regime, and ‘over’ indicates $\phi \in (1, \infty)$ for the overparameterized regime. For the minimum penalty λ_{\min} , ‘neg’ and ‘neg*’ respectively indicate the naive (loose) and exact lower bound on the negative values (see Definition 2). For the optimal penalty λ^* , green and red contrast the cases where the sign changes. ‘Arb. Mod.’, ‘Arb. SNR.’, and ‘Arb. Spec.’ indicate allowing for arbitrary response model, signal-to-noise ratio, and feature covariance spectrum, respectively. (See Table 6 for the reference abbreviation key.)

| Σ | β | Σ_0 | β_0 | $\phi \leq 1$ | λ_{\min} | Arb. Mod. | Arb. SNR. | Arb. Spec. | Additional Specific Data Geometry Conditions | λ^* | Reference |
|---------------------|-----------|------------|-----------|---------------|------------------|--------------|--------------|--------------|---|-------------|----------------------------|
| In-distribution | | | | | | | | | | | |
| \otimes | \circ | Σ | β | all | zero | \times | \checkmark | \times | | + | [DW, Thm. 2.1] |
| \circ | \otimes | Σ | β | all | zero | \times | \checkmark | \times | | + | [HMRT, Cor. 5] |
| | | | | under | neg | \times | \checkmark | \times | | + | [WX, Prop. 6] |
| | | | | over | neg | \times | \times | \times | Strict misalignment of (Σ, β) | + | [WX, Thm. 4] |
| | | | | over | neg | \times | \times | \times | Strict alignment of (Σ, β) | - | [WX, Thm. 4, Prop. 7] |
| \otimes | \otimes | Σ | β | over | zero | \times | \times | \times | and/or special feature model | 0 | [RMR, Cor. 2] |
| | | | | under | neg* | \checkmark | \checkmark | \checkmark | | + | Theorem 2 (1) |
| | | | | over | neg* | \checkmark | \checkmark | \checkmark | General alignment of $(\Sigma, \beta, \sigma^2)$ | - | Theorem 2 (2) |
| Out-of-distribution | | | | | | | | | | | |
| \otimes | \circ | Σ_0 | β | all | neg* | \checkmark | \checkmark | \checkmark | | + | Proposition 3 |
| \otimes | \otimes | Σ_0 | β | under | neg* | \checkmark | \checkmark | \checkmark | | + | Theorem 4 (1) |
| \otimes | \otimes | I | β | over | neg* | \checkmark | \checkmark | \checkmark | | + | Theorem 4 (2) |
| \circ | \otimes | Σ_0 | β | over | neg* | \checkmark | \checkmark | \checkmark | General alignment of $(\Sigma_0, \beta, \sigma^2)$ | - | Theorem 4 (3) |
| | | | | under | neg* | \checkmark | \checkmark | \checkmark | General alignment of (Σ, β, β_0) | - | Theorem 5 (1), (39) |
| \otimes | \otimes | Σ | β_0 | under | neg* | \checkmark | \checkmark | \checkmark | General misalignment of (Σ, β, β_0) | + | Theorem 5 (1), (39) |
| | | | | over | neg* | \checkmark | \checkmark | \checkmark | General alignment of $(\Sigma, \beta, \beta_0, \sigma^2)$ | - | Theorem 5 (2) |

Behavior of optimal regularization. Under random signals with isotropic covariance, Dobriban and Wager (2018) show that the asymptotic risk over the positive range of regularization $\lambda > 0$ is minimized at $\lambda^* = \phi/\text{SNR}$. Here $\text{SNR} = \alpha^2/\sigma^2$, where σ^2 is the noise energy and α^2 is the signal energy. Remarkably, this result is invariant of the feature covariance. Similar results under Gaussian assumptions are derived by Dicker (2016), and Han and Xu (2023) extend the result to most signals in the unit ball with high probability. However, Kobak et al. (2020) demonstrate that optimal regularization can be negative for certain signal and covariance structures in real datasets. Motivated by these curious experiments, Wu and Xu (2020) provide sufficient conditions for optimal regularization of the Bayes risk under anisotropic feature covariance and random signal, assuming a limiting spectrum distribution of the covariance matrix Σ and alignment conditions between the eigenvalues of Σ and the projections of the signal β onto the eigenspace of Σ . Furthermore, Richards et al. (2021) consider strict alignment conditions for a special feature model but do not explicitly consider negative regularization. We refer to these conditions as *strict (mis)alignment* conditions in this paper.

Our paper extends the scope of the aforementioned results for the OOD setting with general response models and allows for the widest possible range of negative regularization. The main differences to Wu and Xu (2020) are the assumptions (linear models and limiting spectrum distribution) and the hypothesis of the theorem (aligned/misaligned). Their analysis only considers high SNR

regimes and analyzes the behavior of optimal regularization for bias when $\phi > 1$, although they also provide an upper bound for the noise level under special cases. We provide generic sufficient conditions. We call these *general (mis)alignment* conditions in the paper. Table 1 provides a detailed comparison summary.

Behavior of optimal risk. Under a random isotropic signal, Dobriban and Wager (2018) obtain the expression for limiting optimal risk, which can be shown to be monotonic in ϕ . Patil and Du (2023) extends these theoretical guarantees to features with an arbitrary covariance matrix and a general moment-bounded response. However, their analysis is limited to IND prediction risk and positive regularization. Recent works have also explored other aspects of the monotonicity of optimal risk; see, for example, Nakkiran et al. (2021); Simon et al. (2023); Yang and Suzuki (2023).

We extend the monotonicity result in Patil and Du (2023) to the OOD setting and allow for negative regularization. In addition, we also show the monotonicity of the optimal risk in SNR. The proof technique for risk monotonicity leverages the equivalences between subsampling and ridge regularization established by Du et al. (2023); Patil and Du (2023). However, the equivalences in these works only consider cases where the ridge penalty is non-negative. We extend their analyses to accommodate negative regularization, which requires extending the properties of the parameters that appear in certain fixed-point equations under negative regularization (see Appendix F).

2 Out-of-Distribution Risk Asymptotics

Before describing the properties of λ^* in Section 3 and the behavior of optimal risk $\mathcal{R}(\hat{\beta}^{\lambda^*})$ in Section 4, we provide the risk asymptotics in this section. For the reader’s convenience, we summarize all our notation in Appendix A. We state assumptions on the train and test distributions below.

2.1 Data Assumptions

We first define a general feature and response distribution, which we will use in our subsequent assumption shortly.

Definition 1 (General feature and response distribution). For $x \sim P_x$, it can be decomposed as $x = \Sigma^{1/2}z$, where $z \in \mathbb{R}^p$ contains i.i.d. entries with mean 0, variance 1, and $(4 + \mu)$ -th moment uniformly bounded for some $\mu > 0$. Here $\Sigma \in \mathbb{R}^{p \times p}$ is deterministic and symmetric with eigenvalues uniformly bounded away from 0 and $+\infty$. For $y \sim P_y$, it has mean 0 and $(4 + \nu)$ -th moment uniformly bounded for some $\nu > 0$. The L_2 linear projection parameter of y onto x is denoted by $\beta = \mathbb{E}[xx^\top]^{-1}\mathbb{E}[xy]$, and the variance of the conditional distribution $P_{y|x}$ is denoted by σ^2 . The joint distribution $P_{x,y}$ is parameterized by $(\Sigma, \beta, \sigma^2)$.

Definition 1 imposes weak moment assumptions on covariates and responses, which are commonly used in random matrix theory and overparameterized risk analysis (Hastie et al., 2022; Bartlett et al., 2021). These assumptions encode a wide class of distributions over \mathbb{R}^{p+1} . By decomposing $dP_y = dP_x \cdot dP_{y|x}$, we can express the response as $y = x^\top \beta + \varepsilon$, where ε is uncorrelated with x and $\mathbb{E}[\varepsilon^2] = \sigma^2$. Note that the L_2 projection parameter β minimizes the linear regression error (Györfi et al., 2006; Buja et al., 2019a,b)¹. Also note that this formulation does not impose any specific structure on the conditional distribution $P_{y|x}$ and does not imply that (x, y) follows a linear model,

¹Technically, our results can accommodate the conditional variance σ^2 depending on x with suitable regularity conditions, but for simplicity, we do not consider this variation in the current paper.

as ε is also a function of x . It is possible to further relax the assumption on the feature vector x to only require an appropriate convex concentration (that proves versions of the Marchenko-Pastur law) (Louart, 2022; Cheng and Montanari, 2022) or even certain infinitesimal asymptotic freeness between the population covariance matrix Σ and the sample covariance matrix $X^\top X/n$ (LeJeune et al., 2024; Patil and LeJeune, 2024). We do not consider such relaxations here.

Under Definition 1, the joint distribution $P_{x,y} = P_{x,y}(\Sigma, \beta, \sigma^2)$ is parameterized by $(\Sigma, \beta, \sigma^2)$. We next state assumptions on the train and test distributions in terms of these distributions, allowing for different sets of parameters.

Assumption A (Train and test distributions). Assume that $P_{x,y}$ and P_{x_0,y_0} are distributed according to Definition 1, parameterized by $(\Sigma, \beta, \sigma^2)$ and $(\Sigma_0, \beta_0, \sigma_0^2)$, respectively.

In this paper, we consider the following types of shifts:

- (i) *Covariate shift*: where $\Sigma \neq \Sigma_0$ but $(\beta, \sigma) = (\beta_0, \sigma_0)$.
- (ii) *Regression shift*: where $\Sigma = \Sigma_0$ but $(\beta, \sigma) \neq (\beta_0, \sigma_0)$.
- (iii) *Joint shift*: where $\Sigma \neq \Sigma_0$ and $(\beta, \sigma) \neq (\beta_0, \sigma_0)$.

Observe that this framework also encompasses various risk notions (even for the IND setting), including the *estimation risk*, which arises when $(\Sigma_0, \beta_0, \sigma_0^2) = (I, \beta, 0)$.

2.2 Out-of-Distribution Risk Asymptotics

In this section, we obtain the asymptotic risk of ridge regression in the OOD setting. Most of the papers on ridge regression consider the range of regularization $\lambda \geq 0$. Motivated by empirical findings in Kobak et al. (2020) that negative regularization can be optimal in real datasets, some recent works consider negative regularization; see, e.g., Wu and Xu (2020); Patil et al. (2021, 2022), using a naive lower bound of $-r_{\min}(1 - \sqrt{\phi})^2$ for λ . A tighter lower bound can be obtained from Theorem 3.1 of LeJeune et al. (2024), which provides an explicit characterization of the smallest nonzero eigenvalue of Wishart-type matrices. This bound is derived by explicitly identifying the analytic continuation to the real line of a unique solution to a certain fixed-point equation over the (upper) complex half-plane (Silverstein and Choi, 1995; Dobriban, 2015). The new bound can significantly outperform the previous naive lower bound (see Figure 1 of LeJeune et al. (2024)).

Definition 2 (Lower bound on negative regularization). Let $\mu_{\min} \in \mathbb{R}$ be the unique solution, satisfying $\mu_{\min} > -r_{\min}$, to the equation:

$$1 = \phi \bar{\text{tr}}[\Sigma^2(\Sigma + \mu_{\min}I)^{-2}], \quad (4)$$

and let $\lambda_{\min}(\phi)$ be given by:

$$\lambda_{\min}(\phi) = \mu_{\min} - \phi \bar{\text{tr}}[\Sigma(\Sigma + \mu_{\min}I)^{-1}]. \quad (5)$$

This enables feasible risk estimation over $\lambda \in (\lambda_{\min}(\phi), \infty)$. Here $\bar{\text{tr}}[A]$ denotes the average trace $\text{tr}[A]/p$ of a matrix $A \in \mathbb{R}^{p \times p}$. To reiterate, the difference between the bound (5) and the naive bound used in Wu and Xu (2020); Patil et al. (2021) can be significant, as seen in Figure 1 as well.

To characterize the asymptotic out-of-distribution (OOD) risk in Proposition 1, we first define the non-negative constants $\mu = \mu(\lambda, \phi)$ and $\tilde{v} = \tilde{v}(\lambda, \phi; \Sigma_0)$ as solutions of the following fixed-point

equations:

$$\mu = \lambda + \phi \overline{\text{tr}}[\mu \Sigma (\Sigma + \mu I)^{-1}], \quad \text{and} \quad \tilde{v} = \frac{\phi \overline{\text{tr}}[\Sigma_0 \Sigma (\Sigma + \mu I)^{-2}]}{1 - \phi \overline{\text{tr}}[\Sigma^2 (\Sigma + \mu I)^{-2}]}.$$

One can interpret μ as the level of implicit regularization “self-induced” by the data (Bartlett et al., 2021; Misiakiewicz and Montanari, 2023). Alternatively, it is also common to parameterize the equations using its inverse $v(\lambda, \phi) = \mu^{-1}(\lambda, \phi)$, which corresponds to the Stieltjes transform of the spectrum of the sample Gram matrix in the limit. With this notation in place, we can now extend the result in Eq. (11) of Patil and Du (2023) to the OOD setting as formalized below.

Proposition 1 (Deterministic equivalents for OOD risk). Under Assumption A, as $n, p \rightarrow \infty$ such that $p/n \rightarrow \phi \in (0, \infty)$ and $\lambda \in (\lambda_{\min}(\phi), \infty)$, the prediction risk $R(\hat{\beta}^\lambda)$ defined in (3) admits a deterministic equivalent $R(\hat{\beta}^\lambda) \simeq \mathcal{R}(\lambda, \phi)$, where the equivalent additively decomposes into:

$$\mathcal{R}(\lambda, \phi) := \mathcal{B}(\lambda, \phi) + \mathcal{V}(\lambda, \phi) + \mathcal{S}(\lambda, \phi) + \kappa^2, \quad (6)$$

with the following deterministic equivalents for the bias, variance, regression shift bias, and irreducible error:

$$\begin{aligned} \mathcal{B} &= \mu^2 \cdot \beta^\top (\Sigma + \mu I)^{-1} (\tilde{v} \Sigma + \Sigma_0) (\Sigma + \mu I)^{-1} \beta, \\ \mathcal{V} &= \sigma^2 \tilde{v}, \\ \mathcal{S} &= 2\mu \cdot \beta^\top (\Sigma + \mu I)^{-1} \Sigma_0 (\beta_0 - \beta), \\ \kappa^2 &= (\beta_0 - \beta)^\top \Sigma_0 (\beta_0 - \beta) + \sigma_0^2. \end{aligned}$$

Note that the deterministic equivalents presented in Proposition 1 depend not only on the regularization parameters (λ, ϕ) , but also on the problem parameters $(\Sigma, \beta, \sigma^2)$ and $(\Sigma_0, \beta_0, \sigma_0^2)$, which we have omitted for notational brevity. Since the risk depends additively on σ_0^2 , we focus mainly on the effect of (Σ_0, β_0) in our analysis. Extending this result to finite samples is possible by imposing additional distributional assumptions on the features and response. Techniques in Knowles and Yin (2017); Cheng and Montanari (2022); Louart (2022), among others, can be used to obtain non-asymptotic analogs of Proposition 1. In this paper, we will focus only on the deterministic equivalents, which capture the first-order information (akin to expectation) of interest for our goals.

3 Properties of Optimal Regularization

In this section, we focus on the optimal ridge penalty λ^* for the asymptotic out-of-distribution (OOD) risk, defined as²:

$$\lambda^* \in \underset{\lambda \geq \lambda_{\min}(\phi)}{\text{argmin}} \mathcal{R}(\lambda, \phi). \quad (7)$$

As discussed in Section 1.2, previous studies have explored the properties of λ^* for ridge regression summarized in Table 1. However, these studies predominantly focus on specific scenarios, such

²Over the extended reals, there is at least one solution to (7). In case there are multiple solutions λ^* to the problem (7), the subsequent guarantees stated in the paper hold for any solution λ^* .

as isotropic signals or features, and do not consider the full range of negative penalty values. Furthermore, their investigations are restricted mainly to the IND setting when $(\Sigma_0, \beta_0) = (\Sigma, \beta)$. We broaden the scope of these results, considering more general scenarios, including anisotropic signals, the full range of (negative) regularization, and both IND and OOD settings.

3.1 In-Distribution Optimal Regularization

We present our initial result for the IND setting, which encompasses and extends the scope of previous works. Based on Proposition 1, we can characterize the properties of the optimal ridge penalty λ^* defined in (7) as follows.

Theorem 2 (Optimal regularization sign for IND risk). Assume the setup of Proposition 1 with $(\Sigma_0, \beta_0) = (\Sigma, \beta)$.

1. (*Underparameterized*) When $\phi < 1$, we have $\lambda^* \geq 0$.
2. (*Overparameterized*) When $\phi > 1$, if for all $v < 1/\mu(0, \phi)$, the following general alignment holds:

$$\frac{\overline{\text{tr}}[B\Sigma(v\Sigma + I)^{-2}] + \sigma^2}{\overline{\text{tr}}[B\Sigma(v\Sigma + I)^{-3}] + \sigma^2} > \frac{\overline{\text{tr}}[\Sigma(v\Sigma + I)^{-2}]}{\overline{\text{tr}}[\Sigma(v\Sigma + I)^{-3}]}, \quad (8)$$

where $B = \beta\beta^\top$, then we have $\lambda^* < 0$.

It is worth mentioning that although we state our results for general deterministic signals, our analysis can also incorporate random signals. In such cases, when β is random, one can simply replace B in the conclusion with its expectation $\mathbb{E}[B]$. Next, we highlight some special cases of Theorem 2 and compare them with previously known results.

When $\Sigma = I$ or $\mathbb{E}[B] = (\alpha^2/p)I$, it is easy to verify that the general alignment condition (8) does not hold (see Remark D.1). This corresponds to the special cases studied by Dicker (2016); Dobriban and Wager (2018), where $\lambda^* \geq 0$.

The general alignment condition (8) in Theorem 2 encompasses the strict alignment conditions in Wu and Xu (2020); Richards et al. (2021). Under strict alignment conditions, Wu and Xu (2020) demonstrate that the optimal ridge penalty is negative in the overparameterized and noiseless setting. They assume perfect alignment or misalignment between the signal distribution and the spectrum distribution of the covariance. This also includes the strong and weak features models considered in Richards et al. (2021). When the signal is strictly aligned with the spectrum of Σ and $\sigma^2 = 0$, it can be shown that (8) holds for all $\mu > 0$ (see Proposition D.1).

The general alignment condition (8) allows for a broader range of signal and covariance structures. For example, in scenarios where the signal is the average of the largest and smallest eigenvectors of Σ , the strict alignment condition does not hold. However, these scenarios can still satisfy the general alignment conditions, as we will illustrate shortly.

In the noiseless setting, when $\sigma^2 = 0$, the alignment condition (8) can be expressed succinctly as:

$$\frac{\partial h(\mu, \Sigma_\beta)}{\partial \mu} < \frac{h(\mu, \Sigma)}{\partial \mu}, \quad (9)$$

by defining the function $h(\cdot, \Sigma_\beta): \mu \mapsto \log \overline{\text{tr}}[B\Sigma(\Sigma + \mu I)^{-2}]$ and $\Sigma_\beta = \Sigma\beta\beta^\top$. At a high level, these




Figure 1: **Illustration of negative or positive optimal regularization under general alignment.** We plot the in-distribution risk of ridge regression against the penalty λ for varying data aspect ratios ϕ in the overparameterized regime. The left and right panels correspond to scenarios when SNR is high ($\sigma^2 = 0.01$) and low ($\sigma^2 = 1$), respectively. The data model has a covariance matrix $(\Sigma_{\text{ar}1})_{ij} := \rho_{\text{ar}1}^{|i-j|}$ with parameter $\rho_{\text{ar}1} = 0.5$, and a coefficient $\beta := \frac{1}{2}(w_{(1)} + w_{(p)})$, where $w_{(j)}$ is the j th eigenvector of $\Sigma_{\text{ar}1}$.

alignment conditions capture how aligned the signal vector is with the feature covariance matrix. When the alignment is strong, it indicates that the problem is effectively low-dimensional. In such cases, less regularization is needed if the signal energy in this effective direction is sufficiently large.

In the noisy setting, when $\sigma^2 \neq 0$, the condition (8) explicitly trades off the alignment of the signal and the noise level. While Wu and Xu (2020, Proposition 5) and Richards et al. (2021, Corollary 2) provide upper bounds on σ^2 for optimal negative regularization under restricted data models, (8) applies to a wider class of data models. In this sense, Theorem 2 extends the previous results on the occurrence of optimal negative regularization to a more general setting.

In Figure 1, we illustrate our theoretical result of Theorem 2. When SNR is high, we observe that the optimal ridge penalty can be negative in the overparameterized regime. In particular, for $\phi = 10$, the negative optimal ridge penalty exceeds the bound used in Wu and Xu (2020). It is worth noting that the scenario depicted in Figure 1 does not satisfy the strict alignment condition, but is incorporated by our characterization in Theorem 2.

3.2 Out-of-Distribution Optimal Regularization

We now investigate the behavior of the optimal ridge penalty under covariate shift or regression shift. In the IND setting, when $\Sigma_0 = \Sigma$ and the signals are isotropic ($\beta\beta^\top \simeq (\alpha^2/p)I$), previous works (Dobriban and Wager, 2018; Han and Xu, 2023) show that the optimal ridge penalty is $\lambda^* = \phi/\text{SNR} \geq 0$. Interestingly, even when allowing for negative regularization and covariate shift ($\Sigma_0 \neq \Sigma$), it is easy to check that the optimal λ remains positive in data models with isotropic signals for any $\phi \in (0, \infty)$.

Proposition 3 (Optimal regularization under covariate shift and random signal). When $\Sigma_0 \neq \Sigma$ and $\beta_0 = \beta$, assuming isotropic signals $\mathbb{E}[\beta\beta^\top] = (\alpha^2/p)I$, we have $\lambda^*(\phi) = \phi/\text{SNR} = \arg\min_\lambda \mathbb{E}_\beta[\mathcal{R}(\lambda, \phi)]$. Furthermore, for $\lambda_{\min} < 0$ such that $X^\top X/n \succ \lambda_{\min}I$, even the non-asymptotic OOD risk (3) is minimized at $\lambda_p^* = \phi_p/\text{SNR}$, where $\phi_p = p/n$.

We remark that Proposition 3 can also be seen as a result of a Bayes optimality argument. However, we provide a more direct argument in Appendix D.3. A result of this flavor for random features is also obtained by Tripuraneni et al. (2021, Proposition 6.1). An interesting observation from Proposition 3 is that the optimal ridge penalty does not depend on Σ_0 . This implies that when the signal is isotropic, one does not need to worry about the covariate shift when tuning the penalty. Therefore, generalized cross-validation for IND risks (Patil et al., 2021) can still yield optimal penalties even for OOD risks under isotropic signals!

However, it is important to note that this result holds specifically for random isotropic signals, which may not be realistic in practice. In scenarios where (near) random isotropic signals are not present, the question of data-dependent regularization tuning in the OOD setting is not as straightforward. We do not consider data-dependent tuning in the current paper and instead focus on the theoretical properties of the (oracle) optimal regularization (and the corresponding risk). Recently, Wang (2023) propose a method for data-dependent tuning of ridge regression under covariate shift by generating pseudo-labels with an undersmoothed ridge regression on the training data. While the method has adaptivity guarantees in the low-dimensional regime, its consistency in the proportional asymptotics regime is not clear yet, and is an interesting direction for future investigation.

Proposition 3 suggests that the optimal penalty λ^* remains invariant for OOD risks under isotropic signals. Although random isotropic signals make the theory more tractable, they generally are not realistic in practice. The following result examines the optimal ridge penalty under deterministic signals, with similar conclusions holding for random anisotropic signals.

Theorem 4 (Optimal regularization under covariate shift and deterministic signal). Assume the setting of Proposition 1 with $\Sigma_0 \neq \Sigma$, $\beta_0 = \beta$.

1. (*Underparameterized*) When $\phi < 1$, we have $\lambda^* \geq 0$.
2. (*Overparameterized*) When $\phi > 1$, if $\Sigma_0 = I$ (corresponding to the estimation risk), then we have $\lambda^* \geq 0$.
3. (*Overparameterized*) When $\phi > 1$, if $\Sigma = I$ and

$$\bar{\text{tr}}[\Sigma_0 B] > \bar{\text{tr}}[\Sigma_0] \left(\bar{\text{tr}}[B] + \frac{(1 + \mu(0, \phi))^3}{\mu(0, \phi)^3} \sigma^2 \right), \quad (10)$$

where $B = \beta\beta^\top$, then we have $\lambda^* < 0$.

When $\phi < 1$, Part (1) of Theorem 4 suggests that the optimal ridge penalty λ^* remains non-negative even under covariate shift. Similarly, when $\phi > 1$, Part (2) of Theorem 4 (for the estimation risk) also guarantees a non-negative λ^* . However, it is quite surprising that in the overparameterized regime, even with deterministic features, the optimal ridge penalty λ^* could be negative in Part (3). In particular, in noiseless setting ($\sigma^2 = 0$), the condition (10) reduces to the strict alignment

condition on (Σ_0, β) ; see Proposition D.1.

While Theorem 4 restricts $\Sigma = I$ to simplify the condition when $\phi > 1$, we provide the condition with general Σ in Equation (39). However, taking $\Sigma = I$ suffices to highlight this rather surprising sign-reversal phenomenon.

Theorem 5 (Optimal regularization under regression shift). Assume the setup of Proposition 1 with $\Sigma_0 = \Sigma$, $\beta_0 \neq \beta$.

1. (*Underparameterized*) When $\phi < 1$, if $\sigma^2 = o(1)$ and for all $\mu \geq 0$, the following general alignment holds:

$$\bar{\text{tr}}[B_0 \Sigma^2 (\Sigma + \mu I)^{-2}] > \bar{\text{tr}}[B \Sigma^2 (\Sigma + \mu I)^{-2}], \quad (11)$$

where $B = \beta \beta^\top$ and $B_0 = \beta_0 \beta_0^\top$, then we have $\lambda^* < 0$.

2. (*Overparameterized*) When $\phi > 1$, if the general alignment conditions (8) and (11) hold, then we have $\lambda^* < 0$.

Similarly to Part (3) of Theorem 4, Part (1) of Theorem 5 is rather surprising. As shown in Table 1, λ^* is always positive for the IND setting when $\phi < 1$. However, Theorem 5 suggests that λ^* can be negative, even for isotropic signals, when there is some alignment or misalignment between $\beta^\top \Sigma$ and $(\beta - \beta_0)^\top \Sigma$. In fact, when $\Sigma = \Sigma_0 = I$ and $\langle \beta, \beta_0 \rangle \geq \|\beta\|_2^2$, the alignment condition $\beta^\top \Sigma^2 (\Sigma + \mu I)^{-2} (\beta_0 - \beta) \geq 0$ in Theorem 5 always holds for all $\mu > 0$. It is worth noting that we assume $\sigma^2 = o(1)$ in Theorem 5 for simplicity, but a more general balance condition that holds for any $\sigma^2 > 0$ is provided in (39).

The numerical illustrations in Figure 2 demonstrate the results of Theorems 4 and 5. As shown, while the optimal ridge penalties λ^* for the IND prediction risks are positive, the OOD prediction risk can be negative and approach its lower limit. Similar observations also occur in real-world MNIST datasets (see Table 2 and Appendix G.1.2 for the experimental details). This phenomenon arises due to distribution shift, which effectively aligns β_0 and Σ_0 . In some cases, it may provide a possible explanation for the success of interpolators in practice, as the optimal ridge penalty λ^* can become negative under distribution shift, e.g., with random features regression (Tripuraneni et al., 2021).

Intuitively, negative regularization may be optimal in certain cases due to the implicit bias of the overparameterized ridge estimator, even when $\lambda = 0$. When the signal energy is sufficiently high, it can be beneficial to “subtract” some of this bias at the expense of increased variance. Negative regularization effectively reduces this bias and can, therefore, be the optimal choice. In a broader context, when there is implicit regularization, such as self-regularization resulting from the data structure, and we desire the overall regularization to be smaller than this inherent amount, negative “external” regularization can help counterbalance it.

4 Properties of Optimal Risk

Under the general covariance structure Σ , Theorem 6 in Patil and Du (2023) shows that optimal ridge regression exhibits a monotonic risk profile in the IND setting ($\Sigma_0 = \Sigma$) and effectively avoids the phenomena of double and multiple descents. This observation motivates a broader investigation




Figure 2: **Covariate and regression shift can lead to negative optimal regularization in both underparameterized and overparameterized regimes.** The plot shows the IND and OOD risks against λ in the high SNR setting ($\sigma^2 = 0.01$ and $\sigma_0^2 = 0$). The left panel shows the overparameterized regime ($\phi = 1.5$) where the optimal ridge penalty λ^* is negative under covariate shift, when $\Sigma = I$, $\Sigma_0 = \Sigma_{\text{ar1}}$, and $\beta = \beta_0 = \frac{1}{2}(w_{(1)} + w_{(p)})$. The right panel shows the underparameterized regime ($\phi = 0.5$) where the optimal ridge penalty λ^* is negative under regression shift, when $\Sigma = \Sigma_0 = \Sigma_{\text{ar1}}$, $\beta = \frac{1}{2}(w_{(1)} + w_{(p)})$, and $\beta_0 = 2\beta$.

Table 2: **Optimal ridge penalty for OOD risks on MNIST gradually becomes more negative with increasing distribution shift.** We gradually shift the test distribution by excluding samples with specific labels. For more details, please refer to Appendix G.1.2.

| Excl. labels: | \emptyset | {4} | {3, 4} | {2, 3, 4} | {1, 2, 3, 4} |
|---------------|-------------|------|--------|-----------|--------------|
| λ^* | 1.03 | 0.48 | 0.00 | -0.48 | 1.44 |

into the monotonic behavior of regularization optimization in the out-of-distribution setting. In this section, we investigate the monotonicity of the optimal OOD risk and converse for the suboptimal OOD risk.

4.1 Optimal Risk Monotonicity

To begin with, we examine the case of isotropic signals under covariate shift. A direct consequence of Proposition 3 is the monotonicity property of the risk in this scenario.

Proposition 6 (Optimal risk under isotropic signals). When $\Sigma_0 \neq \Sigma$ and $\beta = \beta_0$, assuming isotropic signals $\mathbb{E}[\beta\beta^\top] = (\alpha^2/p)I$ the optimal risk obtained at $\lambda^*(\phi) = \phi/\text{SNR}$ is given by:

$$\mathbb{E}_\beta[\mathcal{R}(\lambda^*, \phi)] = \alpha^2 \mu^* \overline{\text{tr}}[\Sigma_0(\Sigma + \mu^* I)^{-1}] + \sigma_0^2, \quad (12)$$

where $\mu^* = \mu(\lambda^*, \phi)$. Furthermore, the left side of (12) is strictly increasing in ϕ if $\text{SNR} \in (0, \infty)$ and σ_0^2 are fixed and strictly increasing in SNR if ϕ , σ^2 , and σ_0^2 are fixed.

Proposition 6 shows that the optimal OOD risk is a monotonic function of ϕ and SNR . This is intuitive because one would expect that having more data (smaller ϕ) or larger SNR would result in a lower prediction risk. In contrast, the ridge or ridgeless predictor computed on the full data does not exhibit this property (Hastie et al., 2022, Figure 2). It is worth noting that the optimal penalty λ^* is also monotonically increasing in the data aspect ratio ϕ when SNR is kept fixed.

However, under anisotropic signals, the optimal regularization penalty generally depends on the specific OOD risk being considered. In such cases, it is difficult to obtain analytical formulas for the optimal ridge penalty and the optimal risk. Nonetheless, we can still show that the optimal OOD risk monotonically increases in ϕ and SNR . We formalize this in the following result, which generalizes the aforementioned IND result in Patil and Du (2023, Theorem 6) to the OOD setting, allowing risk optimization over the possible range of negative regularization (with the lower limit as given in Definition 2).

Theorem 7 (Monotonicity of optimally tuned OOD risk). For $\lambda \geq \lambda_{\min}(\phi)$ where $\lambda_{\min}(\phi)$ is as in (5), for all $\epsilon > 0$ small enough, the risk of optimal ridge predictor satisfies:

$$\min_{\lambda \geq \lambda_{\min}(\phi) + \epsilon} R(\hat{\beta}^\lambda) \simeq \min_{\lambda \geq \lambda_{\min}(\phi)} \mathcal{R}(\lambda, \phi), \quad (13)$$

and right side of (13) is monotonically increasing in ϕ if SNR and σ_0^2 are fixed. In addition, when $\beta = \beta_0$ it is monotonically increasing in SNR if ϕ , σ^2 , and σ_0^2 are fixed.

We find the monotonicity of the optimal risk in ϕ remarkable because it even holds under *arbitrary* covariate and regression shift! The monotonicity in SNR under regression shift can be similarly analyzed but requires fixing more parameters. When considering the IND prediction risk, Patil and Du (2023, Theorem 6) demonstrate that the optimally tuned ridge over $\lambda \in (0, \infty)$ exhibits a monotonic risk profile in the data aspect ratio ϕ . It is somewhat surprising that optimizing over both $(0, \infty)$ and $(\lambda_{\min}(\phi), \infty)$ yields monotonic behavior, as numerically verified in Figure 3.

We also illustrate the optimal risk monotonicity behavior on MNIST in Figure 4. Further, the minimum risk over $(\lambda_{\min}(\phi), \infty)$ can be significantly lower than the minimum risk over $(0, \infty)$, particularly in the overparameterized regime. Although most software packages only consider positive ridge regularization for ridge tuning, Figure 3 suggests that allowing negative regularization can lead to significant improvements.

Finally, we consider the converse question: is optimal regularization *necessary* for achieving monotonic risk? When considering the excess IND prediction risk under isotropic design (i.e., $\Sigma = I$), the following theorem demonstrates that the prediction risk of ridge regression is generally non-monotonic in the data aspect ratio $\phi \in (0, \infty)$.




Figure 3: **Ridge regression optimized over $\lambda \geq \nu$ for different thresholds ν has monotonic risk profile.** We showcase the prediction risk of optimal ridge regression under the same data model as in Figure 1, with $\sigma^2 = 0.01$. The left panel shows the heatmap of the risks $\mathcal{R}(\lambda, \phi)$ of ridge regression for different ridge penalties λ and data aspect ratios ϕ . The lines indicate the optimized ridge risks $\min_{\lambda \geq \nu} \mathcal{R}(\lambda, \phi)$ at different thresholds ν . The right panel shows the optimized risk $\min_{\lambda \geq \nu} \mathcal{R}(\lambda, \phi)$ as a function of ϕ .




Figure 4: **Effect of distribution shift on the risk monotonicity behavior of optimal ridge on MNIST.** The figure illustrates the risk profile (against the training sample size) of optimal ridge regression on the MNIST dataset when subjected to different types of distribution shifts. We follow the same setup as for Table 2 (see Appendix G.1.2 for more details) and vary the number of training sample size n from 25 to 200, and inspect the OOD prediction risk of the optimal ridge predictor. Different colors represent different types of shift from less severe (Type 1) to more severe (Type 5). The y-axis represents the out-of-distribution prediction risk for the task of accurately predicting the digit value for unseen images. The figure shows a clear pattern where the optimal ridge exhibits a monotonically decreasing risk in the training sample size n .

Theorem 8 (Non-monotonicity of suboptimally tuned risk). When $(\Sigma_0, \beta_0) = (\Sigma, \beta)$ and $\Sigma = I$, the risk component equivalents defined in (6) have the following properties:

1. (*Bias component*) For all $\lambda > 0$, $\mathcal{B}(\lambda, \phi)$ is strictly increasing over $\phi \in (0, \lambda + 1)$ and strictly decreasing over $\phi \in (\lambda + 1, \infty)$.
2. (*Variance component*) For all $\lambda > 0$, $\mathcal{V}(\lambda, \phi)$ is strictly increasing over $\phi \in (0, \infty)$.
3. (*Risk*) When $\|\beta\|_2^2 > 0$, for all $\lambda > 0$ and $\epsilon > 0$, there exist $\sigma^2, \phi \in (0, \infty)$, such that $\partial \mathcal{R}(\lambda, \phi) / \partial \phi \leq -\epsilon$, i.e.,

$$\max_{\sigma^2, \phi \in (0, \infty)} \min_{\lambda \geq \lambda_{\min}(\phi)} \partial \mathcal{R}(\lambda, \phi) / \partial \phi \leq -\epsilon. \quad (14)$$

When σ^2 is small, Theorem 8 implies that the risk at any fixed λ is non-monotonic, even when $\Sigma = I$. It also explains the lack of risk monotonicity in ridgeless regression. This result extends known results on non-monotonic behavior of variance of ridgeless regression (Yang et al., 2020).

4.2 Connection to Subsampling and Ensembling

We now briefly discuss the connection between subsampling and ridge regularization (Patil and Du, 2023), used to prove the OOD risk monotonicity results in Section 4. To incorporate negative regularization and OOD risks, we extend these equivalences using tools from LeJeune et al. (2024).

Before presenting the extended equivalence results, we introduce several quantities related to the subsampled ridge ensembles. For an index set $\mathcal{I} \subseteq [n]$ of size k , let $L_{\mathcal{I}} \in \mathbb{R}^{n \times n}$ be a diagonal matrix with i -th diagonal 1 if $i \in \mathcal{I}$ and 0 otherwise. The feature matrix and the response vector associated with a subsampled dataset $\{(x_i, y_i) : i \in \mathcal{I}\}$ are $L_{\mathcal{I}}X$ and $L_{\mathcal{I}}y$, respectively. Given a ridge penalty λ , let $\hat{\beta}_k^\lambda(\mathcal{I})$ denote the ridge estimator fitted on the subsample $(L_{\mathcal{I}}X, L_{\mathcal{I}}y)$, consisting of k samples.

When we aggregate the estimators fitted on all subsampled datasets of size k , we obtain the so-called *full-ensemble* estimators $\hat{\beta}_{k, \infty}^\lambda$, which is almost surely $\mathbb{E}[\hat{\beta}_k^\lambda(\mathcal{I}) \mid X, y]$, if we draw \mathcal{I} independently from the set of index sets of size k . As $k, n, p \rightarrow \infty$ such that $p/n \rightarrow \phi \in (0, \infty)$ and $p/k \rightarrow \psi \in [\phi, \infty]$, Lemma E.1 implies the OOD risk equivalence: $R(\hat{\beta}_{k, \infty}^\lambda) \simeq \mathcal{R}(\lambda, \phi, \psi)$ as in (46). When $\psi = \phi$, the equivalent $\mathcal{R}(\lambda, \phi, \phi)$ reduces to $\mathcal{R}(\lambda, \phi)$ defined in (6) and analyzed in previous sections. We are now ready to present our main result in this section.

Theorem 9 (Optimal ensemble versus ridge regression under negative regularization). Let $\mathcal{R}^{**} := \min_{\psi \geq \phi, \lambda \geq \lambda_{\min}(\phi)} \mathcal{R}(\lambda, \phi, \psi)$. Then the following statements hold:

1. (*Underparameterized*) When $\phi < 1$ and $\beta_0 = \beta$, $\lambda^* \geq 0$,

$$\mathcal{R}^{**} = \min_{\lambda \geq 0} \mathcal{R}(\lambda, \phi, \phi) = \min_{\psi \geq \phi} \mathcal{R}(0; \phi, \psi).$$

2. (*Overparameterized*) When $\phi \geq 1$, $\lambda^* \geq \lambda_{\min}(\phi)$,

$$\mathcal{R}^{**} = \min_{\lambda \geq \lambda_{\min}(\phi)} \mathcal{R}(\lambda, \phi, \phi) = \min_{\psi \geq \phi} \mathcal{R}(\lambda_{\min}(\phi); \phi, \psi).$$

Theorem 9 establishes the OOD risk equivalences between ridge predictors and full-ensemble ridge-




Figure 5: **Negative regularization can help achieve optimal risk in both underparameterized and overparameterized regimes.** The heatmap illustrates the prediction risks for ridge regression as a function of the ridge penalty λ and subsample aspect ratio ψ in the full ensemble. We use the same data model as Figure 3 with $\sigma^2 = 0.01$. The left and right panels show the underparameterized ($\phi = 0.5$) and overparameterized regimes ($\phi = 2$), respectively. The red paths represent the optimal risks, while the blue and green stars indicate the optimal ridge predictor and the optimal full-ensemble ridge with the largest subsample aspect ratio.

less predictors. It shows that in the underparameterized regime, ridgeless ensembles are sufficient to achieve optimal IND risk, which is also supported by Du et al. (2023) when considering optimization over $\lambda \geq 0$. However, in the overparameterized regime, ridgeless ensembles alone are not enough to achieve optimal risk over $\lambda \geq \lambda_{\min}(\phi)$. Explicit negative regularization is required to obtain the optimal risk. In other words, the implicit regularization provided by subsampling is always positive, and under certain data geometries, using $\lambda < 0$ can improve predictive performance.

The proof of Theorem 9 establishes the risk equivalence between the ridge predictor and the full-ensemble ridgeless predictor. As shown in Figure 5, the risk profile of the full-ensemble ridge predictors demonstrates that negative regularization can help achieve optimal risk, particularly in the overparameterized regime when the subsample aspect ratio ψ is greater than 1.

5 Discussion

This paper investigates the optimal regularization and the optimal risk of ridge regression in the OOD setting. The analysis shows the differences between the IND and OOD settings. However, in both cases, the optimal risk is monotonic in the data aspect ratio. The main takeaway is that negative regularization can improve predictive performance in certain data geometries, even more so than the IND setting.

We close the paper by mentioning two future directions. The scope of this paper is limited to standard ridge regression. Under proportional asymptotics, similar conclusions are expected to hold for kernel ridge regression, as the gram matrix linearizes in the sense of asymptotic equivalence (Liang and Rakhlin, 2020; Bartlett et al., 2021; Sahraee-Ardakan et al., 2022; Misiakiewicz and Montanari, 2023). Beyond ridge variants, it is of interest to perform a similar analysis for the lasso.

Preliminary investigations in the literature suggest that, similar to optimal ridge regression, the optimal lasso also exhibits monotonic risk in the overparameterization ratio (Patil et al., 2022). For a comparative illustration, see Figures G13 and G14. It is also of interest to understand the monotonicity properties of the optimal risk for other convex penalized M-estimators in general (Thrampoulidis et al., 2018). Empirical investigations indicate that these estimators also exhibit monotonic risk in the data aspect ratio under optimal regularization (Patil and Du, 2023). Making these observations rigorous is a promising direction on which we hope to report in the near future.

Acknowledgments

We thank Edgar Dobriban, Avi Feller, Qiyang Han, Dmitry Kobak, Arun Kuchibhotla, Daniel LeJeune, Jaouad Mourtada, Alessandro Rinaldo, Kaizheng Wang, Alexander Wei, and Yuting Wei for helpful conversations surrounding this work. We thank the anonymous reviewers for their valuable feedback and suggestions that have improved the paper (in particular, the addition of Appendix B). We also thank the computing support provided by the ACCESS allocation MTH230020 for some of the experiments performed on the Bridges2 system at the Pittsburgh Supercomputing Center. PP and RJT were supported by ONR grant N00014-20-1-2787 during the course of this work.

References

- Francis Bach. High-dimensional analysis of double descent for linear regression with random projections. *SIAM Journal on Mathematics of Data Science*, 6(1):26–50, 2024.
- Peter L. Bartlett, Philip M. Long, Gábor Lugosi, and Alexander Tsigler. Benign overfitting in linear regression. *Proceedings of the National Academy of Sciences*, 117(48):30063–30070, 2020.
- Peter L. Bartlett, Andrea Montanari, and Alexander Rakhlin. Deep learning: A statistical viewpoint. *Acta Numerica*, 30:87–201, 2021.
- Mikhail Belkin, Daniel Hsu, and Ji Xu. Two models of double descent for weak features. *SIAM Journal on Mathematics of Data Science*, 2(4):1167–1180, 2020.
- Pierre Bellec, Jin-Hong Du, Takuya Koriyama, Pratik Patil, and Kai Tan. Corrected generalized cross-validation for finite ensembles of penalized estimators. *arXiv preprint arXiv:2310.01374*, 2023.
- Anders Björkström and Rolf Sundberg. A generalized view on continuum regression. *Scandinavian Journal of Statistics*, 26(1):17–30, 1999.
- 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, 2019a.
- Andreas Buja, Lawrence Brown, Arun Kumar Kuchibhotla, Richard Berk, Edward George, and Linda Zhao. Models as approximations II. *Statistical Science*, 34(4):545–565, 2019b.
- Abdulkadir Canatar, Blake Bordelon, and Cengiz Pehlevan. Out-of-distribution generalization in kernel regression. *Advances in Neural Information Processing Systems*, 2021.

- Chen Cheng and Andrea Montanari. Dimension free ridge regression. *arXiv preprint arXiv:2210.08571*, 2022.
- Alexander D’Amour, Katherine Heller, Dan Moldovan, Ben Adlam, Babak Alipanahi, Alex Beutel, Christina Chen, Jonathan Deaton, Jacob Eisenstein, Matthew D. Hoffman, et al. Underspecification presents challenges for credibility in modern machine learning. *The Journal of Machine Learning Research*, 23(1):10237–10297, 2022.
- James Davidson. *Stochastic Limit Theory: An Introduction for Econometricians*. Oxford University Press, 1994.
- Lee H. Dicker. Ridge regression and asymptotic minimax estimation over spheres of growing dimension. *Bernoulli*, 22(1):1–37, 2016.
- Edgar Dobriban. Efficient computation of limit spectra of sample covariance matrices. *Random Matrices: Theory and Applications*, 4(4), 2015.
- Edgar Dobriban and Stefan Wager. High-dimensional asymptotics of prediction: Ridge regression and classification. *The Annals of Statistics*, 46(1):247–279, 2018.
- Jin-Hong Du, Pratik Patil, and Arun Kumar Kuchibhotla. Subsample ridge ensembles: Equivalences and generalized cross-validation. In *International Conference on Machine Learning*, 2023.
- A. M. Fink and Max Jodeit Jr. On Chebyshev’s other inequality. *Lecture Notes-Monograph Series*, pages 115–120, 1984.
- László Györfi, Michael Kohler, Adam Krzyzak, and Harro Walk. *A Distribution-free Theory of Nonparametric Regression*. Springer, 2006.
- Qiyang Han and Xiacong Xu. The distribution of ridgeless least squares interpolators. *arXiv preprint arXiv:2307.02044*, 2023.
- Trevor Hastie. Ridge regularization: An essential concept in data science. *Technometrics*, 62(4):426–433, 2020.
- 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.
- Arthur E. Hoerl and Robert W. Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67, 1970a.
- Arthur E. Hoerl and Robert W. Kennard. Ridge regression: Applications to nonorthogonal problems. *Technometrics*, 12(1):69–82, 1970b.
- Tsushung A. Hua and Richard F. Gunst. Generalized ridge regression: a note on negative ridge parameters. *Communications in Statistics Theory and Methods*, 12(1):37–45, 1983.
- Arthur Jacot, Berfin Simsek, Francesco Spadaro, Clément Hongler, and Franck Gabriel. Kernel alignment risk estimator: Risk prediction from training data. *Advances in Neural Information Processing Systems*, 2020.

- Antti Knowles and Jun Yin. Anisotropic local laws for random matrices. *Probability Theory and Related Fields*, 169:257–352, 2017.
- Dmitry Kobak, Jonathan Lomond, and Benoit Sanchez. The optimal ridge penalty for real-world high-dimensional data can be zero or negative due to the implicit ridge regularization. *Journal of Machine Learning Research*, 21, 2020.
- Frederic Koehler, Lijia Zhou, Danica J. Sutherland, and Nathan Srebro. Uniform convergence of interpolators: Gaussian width, norm bounds and benign overfitting. *Advances in Neural Information Processing Systems*, 2021.
- Daniel LeJeune, Pratik Patil, Hamid Javadi, Richard G. Baraniuk, and Ryan J. Tibshirani. Asymptotics of the sketched pseudoinverse. *SIAM Journal on Mathematics of Data Science*, 2024.
- Tengyuan Liang and Alexander Rakhlin. Just interpolate: Kernel “ridgeless” regression can generalize. *The Annals of Statistics*, 2020.
- Cosme Louart. Sharp bounds for the concentration of the resolvent in convex concentration settings. *arXiv preprint arXiv:2201.00284*, 2022.
- Bruno Loureiro, Cedric Gerbelot, Hugo Cui, Sebastian Goldt, Florent Krzakala, Marc Mezard, and Lenka Zdeborova. Learning curves of generic features maps for realistic datasets with a teacher-student model. In *Advances in Neural Information Processing Systems*, 2021.
- Neil Mallinar, James B. Simon, Amirhesam Abedsoltan, Parthe Pandit, Mikhail Belkin, and Preetum Nakkiran. Benign, tempered, or catastrophic: A taxonomy of overfitting. *arXiv:2207.06569*, 2022.
- Gabriel Mel and Surya Ganguli. A theory of high dimensional regression with arbitrary correlations between input features and target functions: Sample complexity, multiple descent curves and a hierarchy of phase transitions. In *International Conference on Machine Learning*, 2021.
- Theodor Misiakiewicz and Andrea Montanari. Six lectures on linearized neural networks. *arXiv preprint arXiv:2308.13431*, 2023.
- Vidya Muthukumar, Kailas Vodrahalli, Vignesh Subramanian, and Anant Sahai. Harmless interpolation of noisy data in regression. *IEEE Journal on Selected Areas in Information Theory*, 1(1):67–83, 2020.
- Preetum Nakkiran, Prayaag Venkat, Sham M. Kakade, and Tengyu Ma. Optimal regularization can mitigate double descent. In *International Conference on Learning Representations*, 2021.
- Pratik Patil and Jin-Hong Du. Generalized equivalences between subsampling and ridge regularization. In *Conference on Neural Information Processing Systems*, 2023.
- Pratik Patil and Daniel LeJeune. Asymptotically free sketched ridge ensembles: Risks, cross-validation, and tuning. In *International Conference on Learning Representations*, 2024.
- Pratik Patil, Yuting Wei, Alessandro Rinaldo, and Ryan J. Tibshirani. Uniform consistency of cross-validation estimators for high-dimensional ridge regression. In *International Conference on Artificial Intelligence and Statistics*, 2021.

- Pratik Patil, Arun Kumar Kuchibhotla, Yuting Wei, and Alessandro Rinaldo. Mitigating multiple descents: A model-agnostic framework for risk monotonicization. *arXiv preprint arXiv:2205.12937*, 2022.
- Pratik Patil, Jin-Hong Du, and Arun Kumar Kuchibhotla. Bagging in overparameterized learning: Risk characterization and risk monotonicization. *Journal of Machine Learning Research*, 24(319): 1–113, 2023.
- Luca Pesce, Florent Krzakala, Bruno Loureiro, and Ludovic Stephan. Are Gaussian data all you need? Extents and limits of universality in high-dimensional generalized linear estimation. *arXiv:2302.08923*, 2023.
- Dominic Richards, Jaouad Mourtada, and Lorenzo Rosasco. Asymptotics of ridge (less) regression under general source condition. In *International Conference on Artificial Intelligence and Statistics*, 2021.
- Sheldon M. Ross. *Simulation*. Academic Press, 2022.
- 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.
- J. W. Silverstein and S. I. Choi. Analysis of the limiting spectral distribution of large dimensional random matrices. *Journal of Multivariate Analysis*, 54(2):295–309, 1995.
- James B. Simon, Madeline Dickens, Dhruva Karkada, and Michael Deweese. The eigenlearning framework: A conservation law perspective on kernel ridge regression and wide neural networks. *arxiv:2110.03922*, 2021.
- James B. Simon, Dhruva Karkada, Nikhil Ghosh, and Mikhail Belkin. More is better in modern machine learning: When infinite overparameterization is optimal and overfitting is obligatory, 2023.
- Peter Sollich. Gaussian process regression with mismatched models. *Advances in Neural Information Processing Systems*, 2001.
- Christos Thrampoulidis, Ehsan Abbasi, and Babak Hassibi. Precise error analysis of regularized M-estimators in high dimensions. *IEEE Transactions on Information Theory*, 64(8):5592–5628, 2018.
- Nilesh Tripuraneni, Ben Adlam, and Jeffrey Pennington. Covariate shift in high-dimensional random feature regression. *arXiv preprint arXiv:2111.08234*, 2021.
- Kaizheng Wang. Pseudo-labeling for kernel ridge regression under covariate shift. *arXiv preprint arXiv:2302.10160*, 2023.
- Alexander Wei, Wei Hu, and Jacob Steinhardt. More than a toy: Random matrix models predict how real-world neural representations generalize. *arXiv preprint arXiv:2203.06176*, 2022.
- Denny Wu and Ji Xu. On the optimal weighted ℓ_2 regularization in overparameterized linear regression. *Advances in Neural Information Processing Systems*, 33:10112–10123, 2020.

- Tian-Le Yang and Joe Suzuki. Dropout drops double descent. *arXiv preprint arXiv:2305.16179*, 2023.
- Zitong Yang, Yaodong Yu, Chong You, Jacob Steinhardt, and Yi Ma. Rethinking bias-variance trade-off for generalization of neural networks. In *International Conference on Machine Learning*, pages 10767–10777, 2020.
- Lijia Zhou, Frederic Koehler, Danica J. Sutherland, and Nathan Srebro. Optimistic rates: A unifying theory for interpolation learning and regularization in linear regression. *Journal of Data Science*, 1(1), 2023.

This supplement serves as a companion to the paper titled “Optimal Ridge Regularization for Out-of-Distribution Prediction”. In the first section, we provide an outline of the supplement in Table 3. We also include a summary of the general and specific notation used in the main paper and the supplement in Tables 4 and 5, respectively. In addition, in Table 6, we provide an explanation of and pointers to the abbreviations used in the references mentioned in Table 1.

A Organization and Notation

A.1 Organization

Table 3: Outline of the supplement.

| Section | Subsection | Purpose |
|------------------------------|--------------|--|
| Overview of supplement | | |
| Appendix A | Appendix A.1 | Organization |
| | Appendix A.2 | General notation |
| | Appendix A.3 | Specific notation |
| | Appendix A.4 | Reference key |
| Further details in Section 1 | | |
| Appendix B | Appendix B.1 | Background details on ridge regression with negative regularization |
| Proofs in Section 2 | | |
| Appendix C | Appendix C.1 | Proof of Proposition 1 (out-of-distribution risk asymptotics) |
| Proofs in Section 3 | | |
| Appendix D | Appendix D.1 | Proof of Theorem 2 (sign of optimal regularization for in-distribution) |
| | Appendix D.2 | Special cases of the general alignment condition in Theorem 2 |
| | Appendix D.3 | Proof of Proposition 3 (optimal regularization with isotropic signals) |
| | Appendix D.4 | Proof of Theorem 4 (sign of optimal regularization with covariate shift) |
| | Appendix D.5 | Proof of Theorem 5 (sign of optimal regularization with regression shift) |
| | Appendix D.6 | Helper lemmas (derivatives of components of risk deterministic approximation) |
| Proofs in Section 4 | | |
| Appendix E | Appendix E.1 | Proof of Proposition 6 (optimal risk under isotropic signals) |
| | Appendix E.2 | Proof of Theorem 7 (risk monotonicity with optimal regularization) |
| | Appendix E.3 | Proof of Theorem 8 (risk non-monotonicity with suboptimal regularization) |
| | Appendix E.4 | Proof of Theorem 9 (optimal subsample versus ridge with negative regularization) |
| | Appendix E.5 | Helper lemmas (optimal risk monotonicities, full ensemble OOD risk equivalences) |
| Some technical lemmas | | |
| Appendix F | Appendix F.1 | Properties of minimum limiting negative ridge penalty |
| | Appendix F.2 | Analytic properties of certain fixed-point solutions under negative regularization |
| | Appendix F.3 | Contour of fixed-point solutions under negative regularization |
| Additional experiments | | |
| Appendix G | Appendix G.1 | Additional illustrations for Section 3 |
| | Appendix G.2 | Additional illustrations for Section 5 |

A.2 General Notation

Table 4: A summary of general notation used in the paper and the supplement.

| Notation | Description |
|---|--|
| Fonts | |
| Non-bold lower case | Denotes vectors (e.g., a, b, c) |
| Non-bold upper case | Denotes matrices (e.g., A, B, C) |
| Calligraphic font | Denotes sets (e.g., $\mathcal{A}, \mathcal{B}, \mathcal{C}$) |
| Script font | Denotes certain limiting functions (e.g., $\mathcal{A}, \mathcal{B}, \mathcal{C}$) |
| Sets | |
| \mathbb{N} | Set of natural numbers |
| (a, b, c) | (Ordered) tuple of elements a, b, c |
| $\{a, b, c\}$ | Set of elements a, b, c |
| $[n]$ | Set $\{1, \dots, n\}$ for a natural number n |
| $\mathbb{R}, \mathbb{R}_{\geq 0}$ | Set of real and non-negative real numbers |
| $\mathbb{C}, \mathbb{C}^+, \mathbb{C}^-$ | Set of complex numbers and upper and lower complex half-planes |
| Operations | |
| $\text{tr}[A], \overline{\text{tr}}[A], \det(A)$ | Trace, average trace ($\text{tr}[A]/p$), and determinant of a square matrix $A \in \mathbb{R}^{p \times p}$ |
| B^{-1} | Inverse of an invertible square matrix $B \in \mathbb{R}^{p \times p}$ |
| C^\dagger | Moore-Penrose inverse of a general rectangular matrix $C \in \mathbb{R}^{n \times p}$ |
| $\text{rank}(C), \text{null}(C)$ | Rank and nullity of a general rectangular matrix $C \in \mathbb{R}^{n \times p}$ |
| $D^{1/2}$ | Principal square root of a positive semidefinite matrix D |
| $f(D)$ | Matrix obtained by applying a function $f : \mathbb{R} \rightarrow \mathbb{R}$ to a positive semidefinite matrix D |
| $I, \mathbf{1}, \mathbf{0}$ | The identity matrix, the all-ones vector, the all-zeros vector |
| Norms | |
| $\ u\ _q$ | The ℓ_q norm of a vector u for $q \geq 1$ |
| $\ u\ _A$ | The induced ℓ_2 norm of a vector u with respect to a positive semidefinite matrix A |
| $\ f\ _{L_q}$ | The L_q norm of a function f for $q \geq 1$ |
| $\ A\ _{\text{op}}$ | Operator (or spectral) norm of a real matrix A |
| $\ A\ _{\text{tr}}$ | Trace (or nuclear) norm of a real matrix A |
| $\ A\ _F$ | Frobenius norm of a real matrix A |
| Orders | |
| $X = \mathcal{O}_v(Y)$ | $ X \leq C_v Y$ for some constant C_v that may depend on the ambient parameter v |
| $X \lesssim_v Y$ | $X \leq C_v Y$ for some constant C_v that may depend on the ambient parameter v |
| $u \leq v$ | Lexicographic ordering for vectors u and v |
| $A \leq B$ | Loewner ordering for symmetric matrices A and B |
| Asymptotics | |
| $\mathcal{O}_{\mathbb{P}}, o_{\mathbb{P}}$ | Probabilistic big-O and little-o notation |
| $C \simeq D$ | Asymptotic equivalence of matrices C and D (see Section 2 for more details) |
| $\xrightarrow{d}, \xrightarrow{p}, \xrightarrow{\text{a.s.}}$ | Denotes convergence in distribution, probability, and almost sure convergence |

Some additional conventions used throughout the supplement:

- If a proof of a statement is separated from the statement, the statement is restated (while keeping the original numbering) along with the proof for convenience.
- If no subscript is present for the norm $\|u\|$ of a vector u , it is assumed to be the ℓ_2 norm of u .

A.3 Specific Notation

Table 5: A summary of specific notation used in the paper and the supplement.

| Notation | Description |
|--|--|
| In-distribution | |
| $P_x, P_{y x}, P_{x,y}$ | Distribution of the train features supported on \mathbb{R}^p , conditional distribution of the train response supported on \mathbb{R} , and the joint distribution |
| Σ | Covariance matrix in $\mathbb{R}^{p \times p}$ of the train feature distribution |
| $r_{\min} (r_{\max})$ | Minimum (maximum) eigenvalue of the train covariance matrix |
| σ^2, α^2 | Conditional variance and linearized signal energy of the train response distribution |
| SNR | In-distribution signal-to-noise ratio α^2/σ^2 |
| β | Coefficients of the (population) linear projection of y onto x |
| $\{(x_i, y_i)\}_{i=1}^n$ | Train samples of size n in $\mathbb{R}^p \times \mathbb{R}$ sampled i.i.d. from $P_{x,y}$ |
| X, y | Train feature matrix in $\mathbb{R}^{n \times p}$ and the corresponding response vector in \mathbb{R}^n |
| $\widehat{\beta}^\lambda$ | Ridge estimator fitted on train data (X, y) at regularization level λ |
| $\widehat{\beta}_k^\lambda(\mathcal{I})$ | Ridge estimator fitted on subsampled data $(L_{\mathcal{I}}X, L_{\mathcal{I}}y)$ at the regularization level λ |
| $\widehat{\beta}_{k,\infty}^\lambda$ | Full-ensemble ridge estimator at regularization level λ |
| Out-of-distribution | |
| $P_{x_0}, P_{y_0 x_0}, P_{x_0,y_0}$ | Distribution of the test features supported on \mathbb{R}^p , conditional distribution of the test response supported on \mathbb{R} , and the joint distribution |
| (x_0, y_0) | Test sample in $\mathbb{R}^p \times \mathbb{R}$ sampled from P_{x_0,y_0} independently of the train data (X, y) |
| Σ_0 | Covariance matrix in $\mathbb{R}^{p \times p}$ of the test feature distribution |
| σ_0^2 | Variance of the conditional test response distribution |
| β_0 | Coefficients of the L_2 linear projection of y_0 onto x_0 : $\mathbb{E}[x_0x_0^\top]^{-1}\mathbb{E}[x_0y_0]$ |
| $R(\widehat{\beta}^\lambda)$ | Squared test prediction risk of ridge estimator at level λ |
| Risk parameters | |
| ϕ, ψ | Limiting data and subsample aspect ratios |
| λ^* | Optimal ridge penalty |
| $\lambda_{\min}(\phi), v_{\min}$ | Minimum value of ridge regularization allowed at ϕ and the associated fixed-point parameter |
| $v_p(\lambda, \phi)$ | A fixed-point parameter |
| $\mu(\lambda, \phi)$ | Induced amount of implicit regularization at ridge regularization λ and aspect ratio ϕ |
| $\widetilde{v}_p(\lambda, \phi; \Sigma)$ | Function of a fixed-point parameter |
| $\mathcal{R}(\lambda, \phi), \mathcal{B}(\lambda, \phi), \mathcal{V}(\lambda, \phi), \mathcal{S}(\lambda, \phi)$ | Deterministic approximation to the squared risk, the bias, the variance, and the regression shift |
| κ^2 | Inflated out-of-distribution irreducible error |
| Σ_β | A certain matrix capturing the alignment of feature covariance and signal vector $\Sigma\beta\beta^\top$ |
| $\mathcal{R}(\lambda, \phi, \psi)$ | Deterministic approximation to the risk of the full-ensemble estimator |

A.4 Reference Abbreviation Key

Table 6: Links to references mentioned in Table 1.

| Abbreviation | Reference |
|--------------|--|
| HMRT | Hastie, Montanari, Rosset, and Tibshirani (2022) |
| DW | Dobriban and Wager (2018) |
| RMR | Richards, Mourtada, and Rosasco (2021) |
| WX | Wu and Xu (2020) |

B Further Details in Section 1

B.1 Background Details on Ridge Regression with Negative Regularization

Recall that the ridge regression estimator $\hat{\beta}_\lambda \in \mathbb{R}^p$, based on the training data X, y , is defined as the solution to the following regularized least squares problem:

$$\underset{\beta \in \mathbb{R}^p}{\text{minimize}} \underbrace{\frac{1}{n} \|y - X\beta\|_2^2}_{\mathcal{L}(\beta)} + \lambda \underbrace{\|\beta\|_2^2}_{\mathcal{R}(\beta)}. \quad (\text{P})$$

Here, λ is a regularization parameter. For ease of reference in the following, denote the least squares loss by \mathcal{L} and ℓ_2 regularization by \mathcal{R} . This is referred to as the primal form of ridge regression, denoted by (P). One can write a Lagrangian dual problem of (P) that optimizes weights $\theta \in \mathbb{R}^n$ over the data points as:

$$\underset{\theta \in \mathbb{R}^n}{\text{minimize}} \theta^\top (XX^\top/n + \lambda I_n)\theta/2 - \lambda \theta^\top y/n. \quad (\text{D})$$

This is referred to as the dual form of ridge regression, denoted by (D). Let us denote the solution of this problem by $\hat{\theta}_\lambda$. The estimator $\hat{\beta}$ is linked to the dual weights as $\lambda \hat{\beta}_\lambda = X^\top \hat{\theta}_\lambda$.

Now, there are three cases depending on whether λ is positive, zero, or negative.

- When $\lambda > 0$, regardless of whether $n \geq p$ or $p > n$, both problems (P) and (D) are strictly convex and lead to the following unique ridge estimator (expressed in primal and dual forms, respectively):

$$\hat{\beta}_\lambda = (X^\top X/n + \lambda I_p)^{-1} X^\top y/n = X^\top (XX^\top/n + \lambda I_n)^{-1} y/n. \quad (15)$$

- When $\lambda = 0$ (more precisely, defined through analytic continuation as $\lambda \rightarrow 0^+$) and $X^\top X$ is full rank (which is the case when $n \geq p$ and the design X has independent columns), the optimization problem (P) is still strictly convex and leads to the following unique solution:

$$\hat{\beta}_0 = (X^\top X/n)^{-1} X^\top y/n.$$

- When $\lambda = 0$ (again, defined using analytic continuation as $\lambda \rightarrow 0^+$) and XX^\top is full rank ($p > n$ and the design X has independent rows), the optimization problem (D) is still strictly convex and leads to the following unique solution:

$$\hat{\beta}_0 = X^\top (XX^\top/n)^{-1} y/n.$$

- More generally, when $n \geq p$ and $\lambda \geq -s_{\min}$, where s_{\min} is the smallest eigenvalue of $X^\top X/n$, the optimization problem (P) is still strictly convex and leads to the primal solution in (15).
- Similarly, when $p > n$ and $\lambda \geq -s_{\min}$, where s_{\min} is the smallest eigenvalue of XX^\top/n , the dual to the optimization problem (D) is strictly convex and leads to the dual solution in (15).

By defining the ridge estimator as:

$$\hat{\beta}_\lambda = (X^\top X/n + \lambda I_p)^\dagger X^\top y/n = X^\top (XX^\top/n + \lambda I_n)^\dagger y/n,$$

where A^\dagger denotes the Moore-Penrose pseudoinverse of the matrix A , we can handle all the cases mentioned above simultaneously. To gain intuition, we visually compare the ridge solutions with positive and negative regularization in Figures B6 and B8 for the under- and over-parameterized regimes, respectively. See [Hua and Gunst \(1983\)](#); [Björkström and Sundberg \(1999\)](#) also for some related discussion.

B.1.1 Visualizing ridge regression solutions for positive versus negative regularization (underparameterized regime)

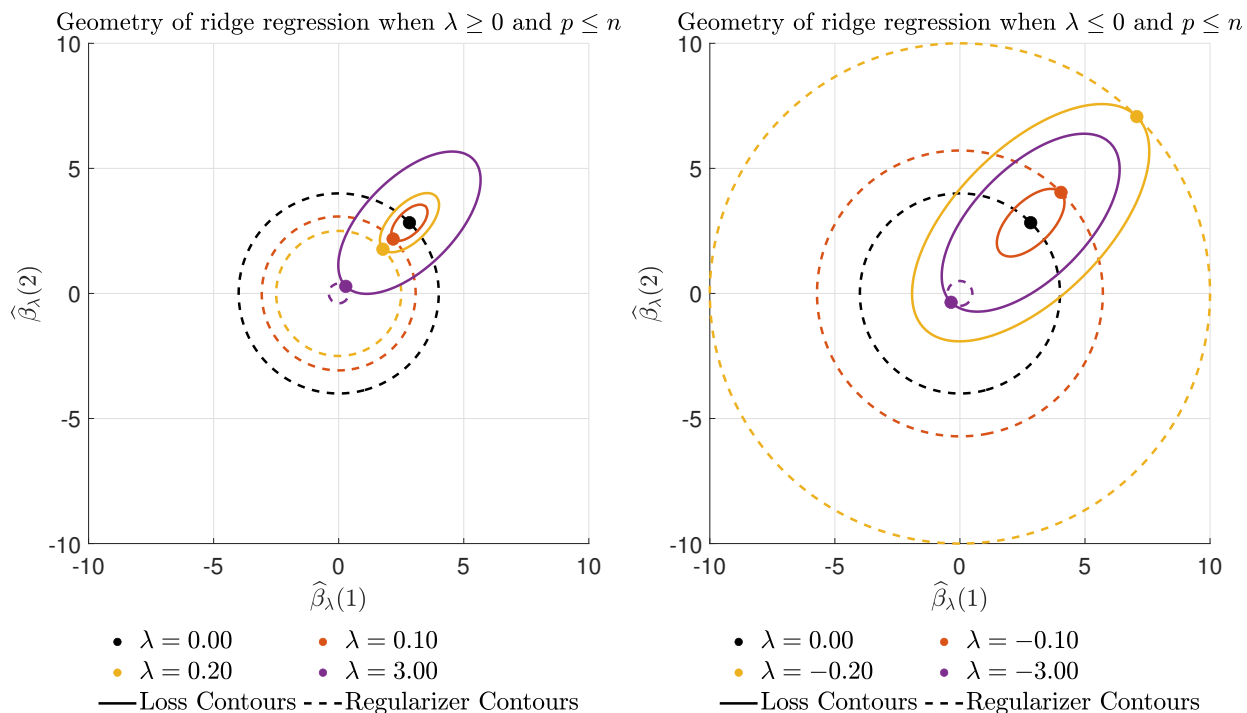


Figure B6: **Comparing the geometry of ridge solutions for positive versus negative regularization levels λ in the underparameterized regime when $p \leq n$.** We consider a two-dimensional underparameterized problem with a design matrix X having the smallest eigenvalue of $X^\top X/n$ as $s_{\min} = 4/3$ and the largest eigenvalue as $s_{\max} = 1/3$. We plot the contours of the ℓ_2 square loss \mathcal{L} and the ℓ_2 regularizer \mathcal{R} in the problem (P) for both positive and negative values of λ . For positive values of $\lambda \in [0, \infty]$, the loss contours touch the constraint contours from the *outside*, while for negative values of $\lambda \in (-\infty, -s_{\max}) \cup (-s_{\min}, 0)$, they touch from the *inside*. To better understand the trend, it helps to think of “tying” $+\infty$ to $-\infty$ on the real line, making it a “projective real line”. The values of λ are plotted in the following order: $\boxed{0.0} \rightarrow 0.1 \rightarrow 0.2 \rightarrow 3.0 \rightarrow -3.0 \rightarrow -0.2 \rightarrow -0.1 \rightarrow \boxed{0.0}$.




Figure B7: Illustration of the sequence of ridge regularization levels λ used in Figure B6 and the corresponding real projective line. (Note: The λ values are not necessarily drawn to scale.)

B.1.2 Visualizing ridge regression solutions for positive versus negative regularization (overparameterized regime)




Figure B8: **Comparing the geometry of ridge solutions for positive versus negative regularization levels λ in the overparameterized regime when $p > n$.** We consider a two-dimensional overparameterized problem with the design matrix X having the smallest *non-zero* eigenvalue of $X^\top X/n$ equal to $s_{\min}^+ = 1$. Similarly to Figure B6, we plot contours of the ℓ_2 squares loss \mathcal{L} and the ℓ_2 regularizer \mathcal{R} (from problem (P)) for both positive and negative values of λ . As before, note that for positive values of λ , the loss contours touch the constraint contours from the *outside*, while for negative values of λ , they touch from the *inside*. To reiterate, it helps to think of “tying” ∞ to $-\infty$ on the real line to better understand the trend. The values of λ are plotted in the following order: $\boxed{0.0} \rightarrow 0.1 \rightarrow 0.2 \rightarrow 3.0 \rightarrow -3.0 \rightarrow -0.2 \rightarrow -0.1 \rightarrow \boxed{0.0}$.




Figure B9: Illustration of the sequence of ridge regularization levels λ used in Figure B8 and the corresponding real projective line. (Note: Here s_{\min}^+ is the smallest *non-zero* eigenvalue of $X^\top X/n$.)

C Proofs in Section 2

C.1 Proof of Proposition 1

Proposition 1 (Deterministic equivalents for OOD risk). Under Assumption A, as $n, p \rightarrow \infty$ such that $p/n \rightarrow \phi \in (0, \infty)$ and $\lambda \in (\lambda_{\min}(\phi), \infty)$, the prediction risk $R(\hat{\beta}^\lambda)$ defined in (3) admits a deterministic equivalent $R(\hat{\beta}^\lambda) \simeq \mathcal{R}(\lambda, \phi)$, where the equivalent additively decomposes into:

$$\mathcal{R}(\lambda, \phi) := \mathcal{B}(\lambda, \phi) + \mathcal{V}(\lambda, \phi) + \mathcal{S}(\lambda, \phi) + \kappa^2, \quad (6)$$

with the following deterministic equivalents for the bias, variance, regression shift bias, and irreducible error:

$$\begin{aligned} \mathcal{B} &= \mu^2 \cdot \beta^\top (\Sigma + \mu I)^{-1} (\tilde{v}\Sigma + \Sigma_0) (\Sigma + \mu I)^{-1} \beta, \\ \mathcal{V} &= \sigma^2 \tilde{v}, \\ \mathcal{S} &= 2\mu \cdot \beta^\top (\Sigma + \mu I)^{-1} \Sigma_0 (\beta_0 - \beta), \\ \kappa^2 &= (\beta_0 - \beta)^\top \Sigma_0 (\beta_0 - \beta) + \sigma_0^2. \end{aligned}$$

Proof. The main workhorse of the proof is Lemma E.1. Recall that

$$\begin{aligned} R(\hat{\beta}^\lambda) &= (\hat{\beta}^\lambda - \beta_0)^\top \Sigma_0 (\hat{\beta}^\lambda - \beta_0) + \sigma_0^2 \\ &= (\hat{\beta}^\lambda - \beta)^\top \Sigma_0 (\hat{\beta}^\lambda - \beta) + 2(\hat{\beta}^\lambda - \beta)^\top \Sigma_0 (\beta - \beta_0) + \{(\beta_0 - \beta)^\top \Sigma_0 (\beta_0 - \beta) + \sigma_0^2\}. \end{aligned} \quad (16)$$

Note that the last quadratic term of (16) is just a constant. From Lemma E.1, we know that the first term of (16) has the following deterministic equivalent:

$$(\hat{\beta}^\lambda - \beta)^\top \Sigma_0 (\hat{\beta}^\lambda - \beta) \simeq \mathcal{Q}_p(\lambda, \phi, \phi), \quad (17)$$

where

$$\mathcal{Q}_p(\lambda, \phi, \psi) = \tilde{c}_p(\lambda, \phi, \psi, \Sigma_0) + \|f_{\text{NL}}\|_{L_2}^2 \tilde{v}_p(\lambda, \phi, \psi, \Sigma_0),$$

and the non-negative constants $\tilde{c}_p(\lambda, \phi, \psi, \Sigma_0)$ and $\tilde{v}_p(\lambda, \phi, \psi, \Sigma_0)$ are defined through the following equations:

$$\frac{1}{v_p(\lambda, \psi)} = \lambda + \psi \int \frac{r}{1 + v_p(\lambda, \psi)r} dH_p(r), \quad (18)$$

$$\tilde{v}_p(\lambda, \phi, \psi, \Sigma_0) = \frac{\phi \overline{\text{tr}}[\Sigma_0 \Sigma (v_p(\lambda, \psi) \Sigma + I)^{-2}]}{v_p(\lambda, \psi)^{-2} - \phi \int \frac{r^2}{(1 + v_p(\lambda, \psi)r)^2} dH_p(r)}, \quad (19)$$

$$\tilde{c}_p(\lambda, \phi, \psi, \Sigma_0) = \beta^\top (v_p(\lambda, \psi) \Sigma + I)^{-1} (\tilde{v}_p(\lambda, \phi, \psi, \Sigma_0) \Sigma + \Sigma_0) (v_p(\lambda, \psi) \Sigma + I)^{-1} \beta.$$

For the second term of (16), we also have

$$(\hat{\beta}^\lambda - \beta)^\top \Sigma_0 (\beta - \beta_0) = f_{\text{NL}}^\top \frac{X}{n} (\hat{\Sigma} + \lambda I)^{-1} \Sigma_0 (\beta - \beta_0) - \beta^\top \lambda (\hat{\Sigma} + \lambda I)^{-1} \Sigma_0 (\beta - \beta_0),$$

where f_{NL} is defined in the proof of Lemma E.1. The first term vanishes from Patil and Du (2023, Lemma D.2) because f_{NL} and X are uncorrelated, and the remaining part has operator norm tending to zero. It follows that

$$(\widehat{\beta}^\lambda - \beta)^\top \Sigma_0 (\beta - \beta_0) \simeq -\beta^\top (v_p(\lambda, \phi) \Sigma + I)^{-1} \Sigma_0 (\beta - \beta_0). \quad (20)$$

Combining (16), (17), and (20) yields

$$\begin{aligned} R(\widehat{\beta}^\lambda) &\simeq \mathcal{R}_p(\lambda, \phi, \phi) \\ &:= \mathcal{Q}_p(\lambda, \phi, \phi) - \beta^\top (v_p(\lambda, \phi) \Sigma + I)^{-1} \Sigma_0 (\beta - \beta_0) + \{(\beta_0 - \beta)^\top \Sigma_0 (\beta_0 - \beta) + \sigma_0^2\}, \end{aligned}$$

which completes the proof. \square

D Proofs in Section 3

Before presenting the proof of our main results, we introduce some notation to facilitate the upcoming presentation. Define

$$q_b(\Sigma_0, B) = \mu^2 \text{tr}[(\Sigma + \mu I)^{-1} \Sigma_0 (\Sigma + \mu I)^{-1} B] / p \quad (21)$$

$$q_v(\Sigma_0, B) = \phi \text{tr}[(\Sigma + \mu I)^{-1} \Sigma_0 (\Sigma + \mu I)^{-1} B] / p \quad (22)$$

$$l(\Sigma_0, B) = \text{tr}[(\Sigma + \mu I)^{-1} \Sigma_0 B] / p. \quad (23)$$

We denote the derivative with respect to μ by:

$$q'_b(\Sigma_0, B) := \frac{\partial q_b(\Sigma_0, B)}{\partial \mu},$$

$$q'_v(\Sigma_0, B) := \frac{\partial q_v(\Sigma_0, B)}{\partial \mu}.$$

We also let $B = \beta \beta^\top$ and $B_0 = \beta_0 \beta_0^\top$.

D.1 Proof of Theorem 2

Theorem 2 (Optimal regularization sign for IND risk). Assume the setup of Proposition 1 with $(\Sigma_0, \beta_0) = (\Sigma, \beta)$.

1. (*Underparameterized*) When $\phi < 1$, we have $\lambda^* \geq 0$.
2. (*Overparameterized*) When $\phi > 1$, if for all $v < 1/\mu(0, \phi)$, the following general alignment holds:

$$\frac{\overline{\text{tr}}[B \Sigma (v \Sigma + I)^{-2}] + \sigma^2}{\overline{\text{tr}}[B \Sigma (v \Sigma + I)^{-3}] + \sigma^2} > \frac{\overline{\text{tr}}[\Sigma (v \Sigma + I)^{-2}]}{\overline{\text{tr}}[\Sigma (v \Sigma + I)^{-3}]}, \quad (8)$$

where $B = \beta \beta^\top$, then we have $\lambda^* < 0$.

Proof. When $(\Sigma_0, \beta_0) = (\Sigma, \beta)$, the excess bias term $\mathcal{S}(\lambda, \phi) = 0$. Therefore, only the bias and variance terms contribute to the risk. We next split the proof into two cases.

Part (1) Underparameterized regime. When $\phi < 1$, from Lemma F.2, $\mu_p(0, \phi) = 0$, and thus, we have that

$$0 = \mathcal{B}(0, \phi) \leq \min_{\lambda \in [\lambda_{\min}(\phi), 0]} \mathcal{B}(\lambda, \phi). \quad (24)$$

On the other hand, when $\phi < 1$, because $\mathcal{V}'(\lambda, \phi) < 0$, the variance is strictly decreasing over $\lambda \in (\lambda_{\min}(\phi), 0)$. We have

$$\mathcal{V}(0, \phi) < \min_{\lambda \in (\lambda_{\min}(\phi), 0)} \mathcal{V}(\lambda, \phi) \quad (25)$$

It follows that

$$\min_{\lambda \in (\lambda_{\min}(\phi), 0]} \mathcal{R}(\lambda, \phi) \geq \mathcal{R}(0, \phi)$$

and

$$\min_{\lambda \geq \lambda_{\min}} \mathcal{R}(\lambda, \phi) = \min_{\lambda \geq 0} \mathcal{R}(\lambda, \phi).$$

Equivalently, we have $\lambda^* \geq 0$.

Part (2) Overparameterized regime. When $\phi > 1$, we begin by deriving the formula of the derivative of the bias term. When $(\Sigma_0, \beta_0) = (\Sigma, \beta)$, the bias term reduces to

$$\mathcal{B}(\lambda, \phi) = \frac{q_v(\Sigma, \Sigma)}{1 - q_v(\Sigma, \Sigma)} q_b(\Sigma, B) + q_b(\Sigma, B) = \frac{q_b(\Sigma, B)}{1 - q_v(\Sigma, \Sigma)}.$$

From Lemma D.2, its derivative satisfies that

$$\begin{aligned} (1 - q_v(\Sigma, \Sigma))^2 \mathcal{B}'(\lambda, \phi) &= (1 - q_v(\Sigma, \Sigma))^2 h'_2(\mu) \\ &= q'_b(\Sigma, B)(1 - q_v(\Sigma, \Sigma)) + q_b(\Sigma, B) q'_v(\Sigma, \Sigma) \\ &= \frac{\lambda}{\mu} q'_b(\Sigma, B) + \mu q'_b(\Sigma, B) q_v(\Sigma, I) + q_b(\Sigma, B) q'_v(\Sigma, \Sigma). \end{aligned} \quad (26)$$

where the last equality is from the identity

$$1 - q_v(\Sigma, \Sigma) = \frac{\lambda}{\mu} + \mu q_v(\Sigma, I) \quad (27)$$

The in-distribution excess term vanishes, i.e., $\mathcal{E}(\lambda, \phi) = 0$. Therefore, the derivative of the risk with respect to μ satisfies that

$$\begin{aligned} &(1 - q_v(\Sigma, \Sigma))^2 \mathcal{B}'(\lambda, \phi) \\ &= \frac{\lambda}{\mu} q'_b(\Sigma, B) + \mu q'_b(\Sigma, B) q_v(\Sigma, I) + q_b(\Sigma, B) q'_v(\Sigma, \Sigma) \\ &= \frac{\lambda}{\mu} q'_b(\Sigma, B) + \mu q'_b(\Sigma, B) q_v(\Sigma, I) + q_b(\Sigma, B) q'_v(\Sigma, \Sigma) \\ &= \frac{\lambda}{\mu} [\mu \bar{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-2}] - \mu^2 \bar{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-3}]] \\ &\quad + 2[\mu \bar{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-2}] - \mu^2 \bar{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-3}]] \cdot \mu \phi \bar{\text{tr}}[\Sigma(\Sigma + \mu I)^{-2}] \end{aligned} \quad (28)$$

$$\begin{aligned}
& - 2\mu \bar{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-2}] \cdot \mu\phi \bar{\text{tr}}[\Sigma^2(\Sigma + \mu I)^{-3}] \\
= & 2\lambda [\bar{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-3}(\Sigma + \mu I)] - \mu \bar{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-3}]] \\
& + 2[\mu \bar{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-2}] - \mu^2 \bar{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-3}]] \cdot \mu\phi \bar{\text{tr}}[\Sigma(\Sigma + \mu I)^{-2}] \\
& - 2\mu \bar{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-2}] \cdot \mu\phi \bar{\text{tr}}[\Sigma^2(\Sigma + \mu I)^{-3}] \\
= & 2\lambda \bar{\text{tr}}[\Sigma_\beta \Sigma(\Sigma + \mu I)^{-3}] \\
& + 2\mu^2 \phi \bar{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-2}] \cdot (\bar{\text{tr}}[\Sigma(\Sigma + \mu I)^{-2}] - \bar{\text{tr}}[\Sigma(\Sigma + \mu I)^{-3}]) \\
& - 2\mu^3 \phi \bar{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-3}] \cdot \bar{\text{tr}}[\Sigma^2(\Sigma + \mu I)^{-2}] \\
= & 2\lambda \bar{\text{tr}}[\Sigma_\beta \Sigma(\Sigma + \mu I)^{-3}] \\
& + 2\mu^3 \phi \bar{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-2}] \cdot \bar{\text{tr}}[\Sigma(\Sigma + \mu I)^{-3}] \\
& - 2\mu^3 \phi \bar{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-3}] \cdot \bar{\text{tr}}[\Sigma(\Sigma + \mu I)^{-2}]. \tag{29}
\end{aligned}$$

When $\lambda = 0$, it follows that

$$\begin{aligned}
& \mathcal{B}'(\lambda, \phi) \\
= & \frac{2\mu^3 \phi}{(1 - q_v(\Sigma, \Sigma))^2} (\bar{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-2}] \cdot \bar{\text{tr}}[\Sigma(\Sigma + \mu I)^{-3}] - \bar{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-3}] \cdot \bar{\text{tr}}[\Sigma(\Sigma + \mu I)^{-2}]),
\end{aligned}$$

and

$$\begin{aligned}
& \frac{\partial \mathcal{B}(\lambda, \phi)}{\partial \lambda} \\
= & \frac{2\mu^3 \phi}{(1 - q_v(\Sigma, \Sigma))^2} (\bar{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-2}] \cdot \bar{\text{tr}}[\Sigma(\Sigma + \mu I)^{-3}] - \bar{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-3}] \cdot \bar{\text{tr}}[\Sigma(\Sigma + \mu I)^{-2}]) \\
& \cdot \frac{\partial \mu(\lambda, \phi)}{\partial \lambda}, \tag{30}
\end{aligned}$$

where $\partial \mu(\lambda, \phi)/\partial \lambda > 0$ from Lemma F.2.

From Lemma F.2 and (30), since $\partial \mu(\lambda, \phi)/\partial \lambda > 0$, we further have that

$$\frac{\partial \mathcal{B}(\lambda, \phi)}{\partial \lambda} \propto \bar{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-2}] \cdot \bar{\text{tr}}[\Sigma(\Sigma + \mu I)^{-3}] - \bar{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-3}] \cdot \bar{\text{tr}}[\Sigma(\Sigma + \mu I)^{-2}],$$

which finishes the proof of the first conclusion.

To obtain the sign of the optimal ridge penalty, note that

$$\begin{aligned}
& (1 - q_v(\Sigma, \Sigma))^2 \mathcal{B}'(\lambda, \phi) \\
= & 2\lambda \bar{\text{tr}}[\Sigma_\beta \Sigma(\Sigma + \mu I)^{-3}] \\
& + 2\mu^3 \phi (\bar{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-2}] \cdot \bar{\text{tr}}[\Sigma(\Sigma + \mu I)^{-3}] - \bar{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-3}] \cdot \bar{\text{tr}}[\Sigma(\Sigma + \mu I)^{-2}]) \\
= & T_1 + T_2.
\end{aligned}$$

When $\phi > 1$, because $\mu \geq 0$, we have that $T_1 < 0$ when $\lambda < 0$ and $T_1 > 0$ when $\lambda > 0$. Also, under the assumption that $\bar{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-2}] \cdot \bar{\text{tr}}[\Sigma(\Sigma + \mu I)^{-3}] - \bar{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-3}] \cdot \bar{\text{tr}}[\Sigma(\Sigma + \mu I)^{-2}] > 0$, we have that $T_2 > 0$ for all $\lambda > \lambda_{\min}(\phi)$, with equality holds only when $\lambda = 0$. Thus, $\mathcal{B}(\lambda, \phi)$ is minimized at $\lambda < 0$.

For the variance term, from Lemma D.2, we have

$$\begin{aligned}
\mathcal{V}'(\lambda, \phi) &= \frac{\sigma^2 q'_v(\Sigma, \Sigma)}{(1 - q_v(\Sigma, \Sigma))^2} \\
&= -2\sigma^2 \phi \frac{\overline{\text{tr}}[\Sigma^2(\Sigma + \mu I)^{-3}]}{(1 - q_v(\Sigma, \Sigma))^2} \\
&= 2\phi \sigma^2 \frac{\mu \overline{\text{tr}}[\Sigma(\Sigma + \mu I)^{-3}] - \overline{\text{tr}}[\Sigma(\Sigma + \mu I)^{-2}]}{(1 - q_v(\Sigma, \Sigma))^2} \\
&:= \frac{T_3 + T_4}{(1 - q_v(\Sigma, \Sigma))^2} \\
&= \frac{-2\sigma^2 \phi \overline{\text{tr}}[\Sigma^2(\Sigma + \mu I)^{-3}]}{(1 - q_v(\Sigma, \Sigma))^2}. \tag{31}
\end{aligned}$$

Note that $-2\sigma^2 \phi \overline{\text{tr}}[\Sigma^2(\Sigma + \mu I)^{-3}] < 0$ and strictly increasing over $\lambda \geq 0$. Thus,

$$\begin{aligned}
&(1 - q_v(\Sigma, \Sigma))^2 \mathcal{R}'(\lambda, \phi) \\
&= (1 - q_v(\Sigma, \Sigma))^2 [\mathcal{B}'(\lambda, \phi) + \mathcal{V}'(\lambda, \phi)] \\
&= 2\lambda \overline{\text{tr}}[\Sigma_\beta \Sigma(\Sigma + \mu I)^{-3}] \\
&\quad + 2\mu^3 \phi (\overline{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-2}] \cdot \overline{\text{tr}}[\Sigma(\Sigma + \mu I)^{-3}] - \overline{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-3}] \cdot \overline{\text{tr}}[\Sigma(\Sigma + \mu I)^{-2}]) \\
&\quad + 2\phi \sigma^2 (\mu \overline{\text{tr}}[\Sigma(\Sigma + \mu I)^{-3}] - \overline{\text{tr}}[\Sigma(\Sigma + \mu I)^{-2}]) \\
&= 2\lambda \overline{\text{tr}}[\Sigma_\beta \Sigma(\Sigma + \mu I)^{-3}] \\
&\quad + 2\phi (\mu^3 \overline{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-2}] + \mu \sigma^2) \cdot \overline{\text{tr}}[\Sigma(\Sigma + \mu I)^{-3}] - 2\phi (\mu^3 \overline{\text{tr}}[\Sigma_\beta(\Sigma + \mu I)^{-3}] + \sigma^2) \cdot \overline{\text{tr}}[\Sigma(\Sigma + \mu I)^{-2}].
\end{aligned}$$

Under the condition that

$$\frac{\overline{\text{tr}}[\Sigma\{\mu^3(\Sigma + \mu I)^{-3}\}]}{\overline{\text{tr}}[\Sigma\{\mu^2(\Sigma + \mu I)^{-2}\}]} > \frac{\overline{\text{tr}}[\Sigma_\beta\{\mu^3(\Sigma + \mu I)^{-3}\}] + \sigma^2}{\overline{\text{tr}}[\Sigma_\beta\{\mu^2(\Sigma + \mu I)^{-2}\}] + \sigma^2},$$

it follows that for all $\lambda \geq 0$,

$$(1 - q_v(\Sigma, \Sigma))^2 \mathcal{R}'(\lambda, \phi) > 0.$$

This implies that $\mathcal{R}(\lambda, \phi)$ is minimized at $\lambda < 0$, which finishes the proof. \square

D.2 General Alignment Condition of Theorem 2 under Special Cases

Remark D.1 (Theorem 2 under isotropic features or signals). When $\phi > 1$ and $\Sigma = I$, the condition above is never satisfied because

$$\frac{\overline{\text{tr}}[\Sigma_\beta\{\mu^3(\Sigma + \mu I)^{-3}\}] + \sigma^2}{\overline{\text{tr}}[\Sigma_\beta\{\mu^2(\Sigma + \mu I)^{-2}\}] + \sigma^2} = \frac{\left(\frac{\mu}{1+\mu}\right)^3 \overline{\text{tr}}[B] + \sigma^2}{\left(\frac{\mu}{1+\mu}\right)^2 \overline{\text{tr}}[B] + \sigma^2} > \frac{\left(\frac{\mu}{1+\mu}\right)^3}{\left(\frac{\mu}{1+\mu}\right)^2} = \frac{\overline{\text{tr}}[\Sigma\{\mu^3(\Sigma + \mu I)^{-3}\}]}{\overline{\text{tr}}[\Sigma\{\mu^2(\Sigma + \mu I)^{-2}\}]}$$

Similarly, when β is isotropic, the condition above is never satisfied because

$$\frac{\overline{\text{tr}}[\Sigma_\beta\{\mu^3(\Sigma + \mu I)^{-3}\}] + \sigma^2}{\overline{\text{tr}}[\Sigma_\beta\{\mu^2(\Sigma + \mu I)^{-2}\}] + \sigma^2} = \frac{\overline{\text{tr}}[\Sigma\{\mu^3(\Sigma + \mu I)^{-3}\}] + \sigma^2}{\overline{\text{tr}}[\Sigma\{\mu^2(\Sigma + \mu I)^{-2}\}] + \sigma^2} > \frac{\overline{\text{tr}}[\Sigma\{\mu^3(\Sigma + \mu I)^{-3}\}]}{\overline{\text{tr}}[\Sigma\{\mu^2(\Sigma + \mu I)^{-2}\}]}$$

since $\overline{\text{tr}}[\Sigma\{\mu^3(\Sigma + \mu I)^{-3}\}] < \overline{\text{tr}}[\Sigma\{\mu^2(\Sigma + \mu I)^{-2}\}]$. To see this, note that $\mu > 0$ when $\phi > 1$, from Lemma F.2. Thus,

$$\begin{aligned}\overline{\text{tr}}[\Sigma\{\mu^3(\Sigma + \mu I)^{-3}\}] &= \overline{\text{tr}}[\Sigma\{\mu^2(\Sigma + \mu I)^{-2}\}\{\mu(\Sigma + \mu I)^{-1}\}] \\ &\leq \overline{\text{tr}}[\Sigma\{\mu^2(\Sigma + \mu I)^{-2}\}] \cdot \|\mu(\Sigma + \mu I)^{-1}\|_{\text{op}} \\ &\leq \overline{\text{tr}}[\Sigma\{\mu^2(\Sigma + \mu I)^{-2}\}].\end{aligned}$$

Proposition D.1 (Theorem 2 under strict alignment conditions). Assuming random signals, zero noise under the strict alignment conditions of Wu and Xu (2020), the general alignment condition (8) is satisfied.

Proof. Let h be a random variable following the empirical measure of the eigenvalues of $\Sigma = U\Lambda U$ (i.e., H_p) and g be a random variable following the empirical measure $\text{diag}(U\mathbb{E}[\beta^\top \beta]U)$. When the joint distribution of (h, g) exists, there exists a function f such that g has mass $r \mapsto f(r) dH_p(r)$. The strict alignment condition from Wu and Xu (2020) imposes f to be either strictly increasing or decreasing. This also implies that

$$\overline{\text{tr}}[\Sigma_\beta] = \overline{\text{tr}}[\Sigma B] \geq \overline{\text{tr}}[\Sigma] \overline{\text{tr}}[B]. \quad (32)$$

This holds because of the following Chebyshev's "other" inequality (Fink and Jodeit Jr., 1984).³

Fact D.1 (Positive/negative correlations of monotone functions; see, e.g., Appendix 9.9 of Ross (2022)). Let f and g be two real-valued functions of the same monotonicity. Let H be a probability measure on the real line. Then the following inequality holds:

$$\int f(r)g(r) dH(r) \geq \int f(r) dH(r) \cdot \int g(r) dH(r). \quad (33)$$

On the other hand, if f is non-decreasing and g is non-increasing, then the inequality in (33) is reversed.

Now we will verify that this condition indeed implies (8) when $\sigma^2 = 0$. Define $f(r) = \sum_{i=1}^p (\beta^\top w_i)^2 \mathbb{1}\{r = r_i\}$. $H_p(r) = p^{-1} \sum_{i=1}^p \mathbb{1}\{r = r_i\}$ and the transformed measure $\tilde{H}_p(r) = r(r + \mu)^{-2} p^{-1} \sum_{i=1}^p \mathbb{1}\{r = r_i\} / \int r(r + \mu)^{-2} dH_p(r)$.

Because $f(r)$ is increasing and $1/(r + \mu)$ is decreasing in r , it follows that

$$\int f(r) \frac{r}{r + \mu} d\tilde{H}_p(r) \leq \int f(r) d\tilde{H}_p(r) \int \frac{r}{r + \mu} d\tilde{H}_p(r).$$

Transforming back to $H_p(r)$ yields that

$$\int f(r) \frac{r}{(r + \mu)^3} dH_p(r) \cdot \int \frac{r}{(r + \mu)^2} dH_p(r) \leq \int f(r) \frac{r}{(r + \mu)^2} dH_p(r) \cdot \int \frac{r}{(r + \mu)^3} dH_p(r).$$

Equivalently,

$$\frac{\overline{\text{tr}}[\Sigma_\beta\{\mu^2(\Sigma + \mu I)^{-2}\}]}{\overline{\text{tr}}[\Sigma_\beta\{\mu^3(\Sigma + \mu I)^{-3}\}]} \geq \frac{\overline{\text{tr}}[\Sigma\{\mu^2(\Sigma + \mu I)^{-2}\}]}{\overline{\text{tr}}[\Sigma\{\mu^3(\Sigma + \mu I)^{-3}\}]},$$

with equality holds if and only if $f(r)$ is H_p -almost surely constant. This finishes the proof. \square

³This is the second (less well-known) inequality, and not the first (more well-known) tail inequality that may come to the reader's mind!

D.3 Proof of Proposition 3

Proposition 3 (Optimal regularization under covariate shift and random signal). When $\Sigma_0 \neq \Sigma$ and $\beta_0 = \beta$, assuming isotropic signals $\mathbb{E}[\beta\beta^\top] = (\alpha^2/p)I$, we have $\lambda^*(\phi) = \phi/\text{SNR} = \text{argmin}_\lambda \mathbb{E}_\beta[\mathcal{R}(\lambda, \phi)]$. Furthermore, for $\lambda_{\min} < 0$ such that $X^\top X/n \succ \lambda_{\min}I$, even the non-asymptotic OOD risk (3) is minimized at $\lambda_p^* = \phi_p/\text{SNR}$, where $\phi_p = p/n$.

Proof. With slight abuse of notations, we use \mathcal{R} , \mathcal{B} , \mathcal{V} , and \mathcal{S} to denote $\mathbb{E}_\beta[\mathcal{R}]$, $\mathbb{E}_\beta[\mathcal{B}]$, $\mathbb{E}_\beta[\mathcal{V}]$, and $\mathbb{E}_\beta[\mathcal{S}]$ as well. We split the proof into two parts.

Part (1) Asymptotic risk. When $\beta_0 = \beta$, the extra bias term is zero, that is, $\mathcal{S}(\lambda, \phi) = 0$. So, only the bias and variance components contribute to the risk.

We begin by deriving the formula for the derivative of the bias. For isotropic signals, from Lemma D.2 and (27), we have

$$\begin{aligned}
& (1 - q_v(\Sigma, \Sigma))^2 \mathcal{B}'(\lambda, \phi) \\
&= \alpha^2 [q'_v(\Sigma_0, \Sigma) q_b(\Sigma, I) (1 - q_v(\Sigma, \Sigma)) \\
&\quad + q_v(\Sigma_0, \Sigma) q'_b(\Sigma, I) (1 - q_v(\Sigma, \Sigma)) + q_v(\Sigma_0, \Sigma) q_b(\Sigma, I) q'_v(\Sigma, \Sigma) \\
&\quad + q'_b(\Sigma_0, I) (1 - q_v(\Sigma, \Sigma))^2] \\
&= \alpha^2 \left[q'_v(\Sigma_0, \Sigma) q_b(\Sigma, I) \left(\frac{\lambda}{\mu} + \mu q_v(\Sigma, I) \right) \right. \\
&\quad + q_v(\Sigma_0, \Sigma) q'_b(\Sigma, I) \left(\frac{\lambda}{\mu} + \mu q_v(\Sigma, I) \right) + q_v(\Sigma_0, \Sigma) q_b(\Sigma, I) q'_v(\Sigma, \Sigma) \\
&\quad \left. + q'_b(\Sigma_0, I) \left(\frac{\lambda}{\mu} + \mu q_v(\Sigma, I) \right)^2 \right] \\
&= \frac{\alpha^2 \lambda}{\mu} \left(q'_v(\Sigma_0, \Sigma) q_b(\Sigma, I) + q_v(\Sigma_0, \Sigma) q'_b(\Sigma, I) + \frac{\lambda}{\mu} q'_b(\Sigma_0, I) + 2\mu q'_b(\Sigma_0, I) q_v(\Sigma, I) \right) \\
&\quad + \alpha^2 \left[\mu q'_v(\Sigma_0, \Sigma) q_b(\Sigma, I) q_v(\Sigma, I) \right. \\
&\quad + \mu q_v(\Sigma_0, \Sigma) q'_b(\Sigma, I) q_v(\Sigma, I) + q_v(\Sigma_0, \Sigma) q_b(\Sigma, I) q'_v(\Sigma, \Sigma) \\
&\quad \left. + \mu^2 q'_b(\Sigma_0, I) q_v(\Sigma, I)^2 \right].
\end{aligned}$$

Notice that

$$\begin{aligned}
q_b(\Sigma_0, I) &= \mu^2 \overline{\text{tr}}[\Sigma_0(\Sigma + \mu I)^{-2}] \\
&= \frac{\mu^2}{\phi} q_v(\Sigma_0, I) \\
q'_b(\Sigma_0, I) &= 2\mu \overline{\text{tr}}[\Sigma_0(\Sigma + \mu I)^{-2}] - 2\mu^2 \overline{\text{tr}}[\Sigma_0(\Sigma + \mu I)^{-3}] \\
&= 2\mu \overline{\text{tr}}[\Sigma_0 \Sigma (\Sigma + \mu I)^{-3}] \\
&= -\frac{\mu}{\phi} q'_v(\Sigma_0, \Sigma),
\end{aligned}$$

thus, it follows that

$$q'_v(\Sigma_0, \Sigma) q_b(\Sigma, I) + \mu q'_b(\Sigma_0, I) q_v(\Sigma, I) = q'_v(\Sigma_0, \Sigma) \left[q_b(\Sigma, I) - \frac{\mu^2}{\phi} q_v(\Sigma, I) \right] = 0$$

$$\mu q'_b(\Sigma, I)q_v(\Sigma, I) + q_b(\Sigma, I)q'_v(\Sigma, \Sigma) = \frac{\mu^2}{\phi} [-q'_v(\Sigma, \Sigma)q_v(\Sigma, I) + q'_v(\Sigma, \Sigma)q_v(\Sigma, I)] = 0.$$

Therefore, we have

$$\begin{aligned} & (1 - q_v(\Sigma, \Sigma))^2 \mathcal{B}'(\lambda, \phi) \\ &= \frac{\alpha^2 \lambda}{\mu} \left(q_v(\Sigma_0, \Sigma)q'_b(\Sigma, I) + \frac{\lambda}{\mu} q'_b(\Sigma_0, I) + \mu q'_b(\Sigma_0, I)q_v(\Sigma, I) \right) \\ &= -\frac{\alpha^2 \lambda}{\phi} q_v(\Sigma_0, \Sigma)q'_v(\Sigma, \Sigma) - \frac{\alpha^2 \lambda^2}{\mu \phi} q'_v(\Sigma_0, \Sigma) - \frac{\alpha^2 \lambda \mu}{\phi} q'_v(\Sigma_0, \Sigma)q_v(\Sigma, I) \\ &= -\frac{\alpha^2 \lambda}{\phi} [q_v(\Sigma_0, \Sigma)q'_v(\Sigma, \Sigma) + \mu q'_v(\Sigma_0, \Sigma)q_v(\Sigma, I)] - \frac{\alpha^2 \lambda^2}{\mu \phi} q'_v(\Sigma_0, \Sigma). \end{aligned}$$

From (27), we further have that

$$\begin{aligned} \mathcal{B}'(\lambda, \phi) &= -\frac{\alpha^2 \lambda}{\phi(1 - q_v(\Sigma, \Sigma))^2} [q_v(\Sigma_0, \Sigma)q'_v(\Sigma, \Sigma) - q'_v(\Sigma_0, \Sigma)q_v(\Sigma, \Sigma) + q'_v(\Sigma_0, \Sigma)] \\ &= -\frac{\alpha^2 \lambda}{\phi} \frac{\partial}{\partial \mu} \frac{q_v(\Sigma_0, \Sigma)}{1 - q_v(\Sigma, \Sigma)}. \end{aligned}$$

From Lemma D.2, we know that

$$\frac{\partial}{\partial \mu} \frac{q_v(\Sigma_0, \Sigma)}{1 - q_v(\Sigma, \Sigma)} < 0$$

for all $\lambda \in (-\lambda_{\min}(\phi), +\infty)$. Furthermore, we have

$$\begin{aligned} \mathcal{R}'(\lambda, \phi) &= \mathcal{B}'(\lambda, \phi) + \mathcal{V}'(\lambda, \phi) \\ &= -\frac{\alpha^2 \lambda}{\phi} \frac{\partial}{\partial \mu} \frac{q_v(\Sigma_0, \Sigma)}{1 - q_v(\Sigma, \Sigma)} + \sigma^2 \frac{\partial}{\partial \mu} \frac{q_v(\Sigma_0, \Sigma)}{1 - q_v(\Sigma, \Sigma)}. \end{aligned}$$

Setting the above to zero gives $\lambda^* = \phi \sigma^2 / \alpha^2$. Also note that $\mathcal{R}'(\lambda, \phi)$ is negative when $\lambda < \lambda^*$ and positive when $\lambda > \lambda^*$, as well as $\partial \mu / \partial \lambda > 0$ from Lemma F.2 for all $\lambda \geq \lambda_{\min}(\phi)$. Thus, λ^* gives the optimal risk.

Part (2) Finite-sample risk. Recall the finite-sample risk is given by (3). Under Assumption A and isotropic signals, it follows that

$$R(\hat{\beta}^\lambda) = \lambda^2 \alpha^2 \bar{\text{tr}}[(\hat{\Sigma} + \lambda I)^{-2} \Sigma_0] + \sigma^2 \phi \bar{\text{tr}}[(\hat{\Sigma} + \lambda I)^{-2} \hat{\Sigma} \Sigma_0].$$

Taking the derivative with respect to λ yields that

$$\begin{aligned} \frac{\partial R(\hat{\beta}^\lambda)}{\partial \lambda} &= 2\lambda \alpha^2 \bar{\text{tr}}[(\hat{\Sigma} + \lambda I)^{-2} \Sigma_0] - 2\lambda^2 \alpha^2 \bar{\text{tr}}[(\hat{\Sigma} + \lambda I)^{-3} \Sigma_0] - 2\sigma^2 \phi \bar{\text{tr}}[(\hat{\Sigma} + \lambda I)^{-3} \hat{\Sigma} \Sigma_0] \\ &= 2\lambda \alpha^2 \bar{\text{tr}}[(\hat{\Sigma} + \lambda I)^{-3} \hat{\Sigma} \Sigma_0] - 2\sigma^2 \phi \bar{\text{tr}}[(\hat{\Sigma} + \lambda I)^{-3} \hat{\Sigma} \Sigma_0] \\ &= 2(\lambda \alpha^2 - \sigma^2 \phi) \bar{\text{tr}}[(\hat{\Sigma} + \lambda I)^{-3} \hat{\Sigma} \Sigma_0]. \end{aligned}$$

Setting the above to zero gives $\lambda_p^* = \phi_n / \text{SNR}$. □

D.4 Proof of Theorem 4

Theorem 4 (Optimal regularization under covariate shift and deterministic signal). Assume the setting of Proposition 1 with $\Sigma_0 \neq \Sigma$, $\beta_0 = \beta$.

1. (*Underparameterized*) When $\phi < 1$, we have $\lambda^* \geq 0$.
2. (*Overparameterized*) When $\phi > 1$, if $\Sigma_0 = I$ (corresponding to the estimation risk), then we have $\lambda^* \geq 0$.
3. (*Overparameterized*) When $\phi > 1$, if $\Sigma = I$ and

$$\bar{\text{tr}}[\Sigma_0 B] > \bar{\text{tr}}[\Sigma_0] \left(\bar{\text{tr}}[B] + \frac{(1 + \mu(0, \phi))^3}{\mu(0, \phi)^3} \sigma^2 \right), \quad (10)$$

where $B = \beta\beta^\top$, then we have $\lambda^* < 0$.

Proof. When $\beta_0 = \beta$, the extra bias term is zero, that is, $\mathcal{S}(\lambda, \phi) = 0$. So, only bias and variance contribute to the risk. We next split the proof into two cases.

Part (1) Underparameterized regime. When $\phi < 1$, from Lemma F.3 we have $\mu(0, \phi) = 0$ and the bias defined in Proposition 1 becomes zero, i.e., $\mathcal{B}(0, \phi) = 0$. From the fixed-point equation (53), we also have $\lim_{\lambda \rightarrow 0^+} \lambda/\mu(\lambda, \phi) = 0$. From Lemma D.2 and (27), we have

$$\begin{aligned} & (1 - q_v(\Sigma, \Sigma))^2 \mathcal{B}'(\lambda, \phi) \\ &= q'_v(\Sigma_0, \Sigma) q_b(\Sigma, B) (1 - q_v(\Sigma, \Sigma)) \\ & \quad + q_v(\Sigma_0, \Sigma) q'_b(\Sigma, B) (1 - q_v(\Sigma, \Sigma)) + q_v(\Sigma_0, \Sigma) q_b(\Sigma, B) q'_v(\Sigma, \Sigma) \\ & \quad + q'_b(\Sigma_0, B) (1 - q_v(\Sigma, \Sigma))^2 \\ &= q'_v(\Sigma_0, \Sigma) q_b(\Sigma, B) \left(\frac{\lambda}{\mu} + \mu q_v(\Sigma, I) \right) \\ & \quad + q_v(\Sigma_0, \Sigma) q'_b(\Sigma, B) \left(\frac{\lambda}{\mu} + \mu q_v(\Sigma, I) \right) + q_v(\Sigma_0, \Sigma) q_b(\Sigma, B) q'_v(\Sigma, \Sigma) \\ & \quad + q'_b(\Sigma_0, B) \left(\frac{\lambda}{\mu} + \mu q_v(\Sigma, I) \right)^2 \\ &= \frac{\lambda}{\mu} \left(q'_v(\Sigma_0, \Sigma) q_b(\Sigma, B) + q_v(\Sigma_0, \Sigma) q'_b(\Sigma, B) + \frac{\lambda}{\mu} q'_b(\Sigma_0, B) + 2\mu q'_b(\Sigma_0, B) q_v(\Sigma, I) \right) \\ & \quad + \left[\mu q'_v(\Sigma_0, \Sigma) q_b(\Sigma, B) q_v(\Sigma, I) \right. \\ & \quad + \mu q_v(\Sigma_0, \Sigma) q'_b(\Sigma, B) q_v(\Sigma, I) + q_v(\Sigma_0, \Sigma) q_b(\Sigma, B) q'_v(\Sigma, \Sigma) \\ & \quad \left. + \mu^2 q'_b(\Sigma_0, B) q_v(\Sigma, I)^2 \right]. \end{aligned}$$

Thus, the derivative of the bias term becomes zero, i.e., $\mathcal{B}'(0, \phi) = 0$. Note that the variance term is strictly decreasing over $(\lambda_{\min}(\phi), +\infty)$ from Lemma D.2. Similarly to part (1) of the proof of Theorem 2, it follows that $\lambda^* \geq 0$.

Part (2) Overparameterized regime and $\Sigma = I$. When $\Sigma = I$, the above derivative in Part (1) becomes

$$(1 - q_v(I, I))^2 \mathcal{B}'(\lambda, \phi)$$

$$\begin{aligned}
&= \frac{\lambda}{\mu} \left(q'_v(\Sigma_0, I)q_b(I, B) + q_v(\Sigma_0, I)q'_b(I, B) + \frac{\lambda}{\mu}q'_b(\Sigma_0, B) + 2\mu q'_b(\Sigma_0, B)q_v(I, I) \right) \\
&\quad + \left[\mu q'_v(\Sigma_0, I)q_b(I, B)q_v(I, I) \right. \\
&\quad + \mu q_v(\Sigma_0, I)q'_b(I, B)q_v(I, I) + q_v(\Sigma_0, I)q_b(I, B)q'_v(I, I) \\
&\quad \left. + \mu^2 q'_b(\Sigma_0, B)q_v(I, I)^2 \right].
\end{aligned}$$

Note from (21) and (22), we have that

$$\begin{aligned}
q_b(\Sigma_0, B) &= \frac{\mu^2}{(1+\mu)^2} \bar{\text{tr}}[\Sigma_0 B] \\
q'_b(\Sigma_0, B) &= \left(-\frac{2\mu^2}{(1+\mu)^3} + \frac{2\mu}{(1+\mu)^2} \right) \bar{\text{tr}}[\Sigma_0 B] = \frac{2\mu}{(1+\mu)^3} \bar{\text{tr}}[\Sigma_0 B] = \frac{2}{\mu(1+\mu)} q_b(\Sigma_0, B), \\
q_v(I, I) &= \frac{\phi}{(1+\mu)^2} \\
q'_v(I, I) &= -\frac{2\phi}{(1+\mu)^3} = -\frac{2}{1+\mu} q_v(I, I) \\
q_v(\Sigma_0, I) &= \frac{\phi}{(1+\mu)^2} \bar{\text{tr}}[\Sigma_0] \\
q'_v(\Sigma_0, I) &= -\frac{2\phi}{(1+\mu)^3} \bar{\text{tr}}[\Sigma_0] = -\frac{2}{1+\mu} q_v(\Sigma_0, I).
\end{aligned}$$

We further have

$$\begin{aligned}
&(1 - q_v(I, I))^2 \mathcal{B}'(\lambda, \phi) \\
&= \frac{\lambda}{\mu} \left(-\frac{2}{1+\mu} q_v(\Sigma_0, I)q_b(I, B) + \frac{2}{\mu(1+\mu)} q_v(\Sigma_0, I)q_b(I, B) + \frac{2\lambda}{\mu^2(1+\mu)} q_b(\Sigma_0, B) + \frac{4}{1+\mu} q_b(\Sigma_0, B)q_v(I, I) \right) \\
&\quad + \left[-\frac{2\mu}{1+\mu} q_v(\Sigma_0, I)q_b(I, B)q_v(I, I) \right. \\
&\quad + \frac{2}{1+\mu} q_v(\Sigma_0, I)q_b(I, B)q_v(I, I) - \frac{2}{1+\mu} q_v(\Sigma_0, I)q_b(I, B)q_v(I, I) \\
&\quad \left. + \frac{2\mu}{1+\mu} q_b(\Sigma_0, B)q_v(I, I)^2 \right] \\
&= \frac{\lambda}{\mu} \left(\frac{2\phi\mu(1-\mu)}{(1+\mu)^5} \bar{\text{tr}}[\Sigma_0] \bar{\text{tr}}[B] + \frac{2\lambda}{(1+\mu)^3} \bar{\text{tr}}[\Sigma_0 B] + \frac{4\phi\mu^2}{(1+\mu)^5} \bar{\text{tr}}[\Sigma_0 B] \right) \\
&\quad - \frac{2\mu}{1+\mu} q_v(\Sigma_0, I)q_b(I, B)q_v(I, I) \\
&\quad + \frac{2\mu}{1+\mu} q_b(\Sigma_0, B)q_v(I, I)^2 \\
&= \frac{\lambda}{\mu} \left(\frac{2\phi\mu(1-\mu)}{(1+\mu)^5} \bar{\text{tr}}[\Sigma_0] \bar{\text{tr}}[B] + \frac{2\lambda}{(1+\mu)^3} \bar{\text{tr}}[\Sigma_0 B] + \frac{4\phi\mu^2}{(1+\mu)^5} \bar{\text{tr}}[\Sigma_0 B] \right) \\
&\quad + \frac{2\phi^2\mu^3}{(1+\mu)^7} (\bar{\text{tr}}[\Sigma_0 B] - \bar{\text{tr}}[\Sigma_0] \bar{\text{tr}}[B]) \\
&= \frac{\lambda}{\mu} \left(\frac{2\phi\mu(1-\mu)}{(1+\mu)^5} \bar{\text{tr}}[\Sigma_0] \bar{\text{tr}}[B] + \frac{2[\mu(1+\mu)^2 - \phi\mu(1-\mu)]}{(1+\mu)^5} \bar{\text{tr}}[\Sigma_0 B] \right) \\
&\quad + \frac{2\phi^2\mu^3}{(1+\mu)^7} (\bar{\text{tr}}[\Sigma_0 B] - \bar{\text{tr}}[\Sigma_0] \bar{\text{tr}}[B])
\end{aligned}$$

$$\begin{aligned}
&= \frac{\lambda}{\mu} \left(\frac{2\phi\mu(1-\mu)}{(1+\mu)^5} (\bar{\text{tr}}[\Sigma_0] \bar{\text{tr}}[B] - \bar{\text{tr}}[\Sigma_0 B]) + \frac{2\mu}{(1+\mu)^3} \bar{\text{tr}}[\Sigma_0 B] \right) \\
&\quad + \frac{2\phi^2\mu^3}{(1+\mu)^7} (\bar{\text{tr}}[\Sigma_0 B] - \bar{\text{tr}}[\Sigma_0] \bar{\text{tr}}[B]).
\end{aligned}$$

For the variance term, from Lemma D.2, we have

$$\begin{aligned}
(1 - q_v(I, I))^2 \mathcal{V}'(\lambda, \phi) &= \sigma^2 [q'_v(\Sigma_0, I)(1 - q_v(I, I)) + q_v(\Sigma_0, I)q'_v(I, I)] \\
&= \sigma^2 \left[q'_v(\Sigma_0, I) \left(\frac{\lambda}{\mu} + \mu q_v(I, I) \right) + q_v(\Sigma_0, I)q'_v(I, I) \right] \\
&= -\sigma^2 \frac{2\lambda}{\mu(1+\mu)} q_v(\Sigma_0, I) - \sigma^2 \frac{2}{1+\mu} [\mu q_v(\Sigma_0, I)q_v(I, I) + q_v(\Sigma_0, I)q_v(I, I)] \\
&= -\sigma^2 \frac{2\lambda\phi}{\mu(1+\mu)^3} \bar{\text{tr}}[\Sigma_0] - \sigma^2 \frac{2\phi^2}{(1+\mu)^4} \bar{\text{tr}}[\Sigma_0].
\end{aligned}$$

From the fixed-point equation (53) we have $\mu(1+\mu) - \phi\mu = \lambda(1+\mu)$ and thus,

$$\begin{aligned}
(1 - q_v(I, I))^2 \mathcal{V}'(\lambda, \phi) &= -\sigma^2 \frac{2\phi(\mu(1+\mu) - \phi\mu)}{\mu(1+\mu)^4} \bar{\text{tr}}[\Sigma_0] - \sigma^2 \frac{2\phi^2}{(1+\mu)^4} \bar{\text{tr}}[\Sigma_0] \\
&= -\sigma^2 \frac{2\phi}{(1+\mu)^3} \bar{\text{tr}}[\Sigma_0],
\end{aligned}$$

which is strictly increasing in $\lambda \geq 0$.

Then we have

$$\begin{aligned}
(1 - q_v(I, I))^2 \mathcal{R}'(\lambda, \phi) &= (1 - q_v(\Sigma, \Sigma))^2 [\mathcal{B}'(\lambda, \phi) + \mathcal{V}'(\lambda, \phi)] \\
&\geq (1 - q_v(\Sigma, \Sigma))^2|_{\lambda=0} [\mathcal{B}'(0; \phi) + \mathcal{V}'(0; \phi)] \\
&= 2\lambda \left(\frac{\phi(1-\mu)}{(1+\mu)^5} (\bar{\text{tr}}[\Sigma_0] \bar{\text{tr}}[B] - \bar{\text{tr}}[\Sigma_0 B]) + \frac{1}{(1+\mu)^3} \bar{\text{tr}}[\Sigma_0 B] \right) \\
&\quad + \frac{2\phi^2\mu^3}{(1+\mu)^7} (\bar{\text{tr}}[\Sigma_0 B] - \bar{\text{tr}}[\Sigma_0] \bar{\text{tr}}[B]) - \sigma^2 \frac{2\phi}{(1+\mu)^3} \bar{\text{tr}}[\Sigma_0] \\
&= \frac{2\lambda\phi}{(1+\mu)^5} \left(\frac{(1+\mu)^2 - \phi}{\phi} \bar{\text{tr}}[\Sigma_0 B] + \bar{\text{tr}}[\Sigma_0] \bar{\text{tr}}[B] \right) \tag{34}
\end{aligned}$$

$$+ \left(\frac{2\lambda\phi\mu}{(1+\mu)^5} - \frac{2\lambda\phi\mu^2}{(1+\mu)^6} \right) (\bar{\text{tr}}[\Sigma_0 B] - \bar{\text{tr}}[\Sigma_0] \bar{\text{tr}}[B]) \tag{35}$$

$$+ \frac{2\phi\mu^3}{(1+\mu)^6} \left(\bar{\text{tr}}[\Sigma_0 B] - \bar{\text{tr}}[\Sigma_0] \left(\bar{\text{tr}}[B] + \frac{(1+\mu)^3}{\mu^3} \sigma^2 \right) \right). \tag{36}$$

Note that when $\phi > 1$, $\mu(\lambda, \phi) > 0$ from Lemma F.2. Under the alignment condition, we have

$$\bar{\text{tr}}[\Sigma_0 B] > \bar{\text{tr}}[\Sigma_0] \left(\bar{\text{tr}}[B] + \frac{(1+\mu(0, \phi))^3}{\mu(0, \phi)^3} \sigma^2 \right) \geq \bar{\text{tr}}[\Sigma_0] \bar{\text{tr}}[B].$$

Thus, for $\lambda \geq 0$, the second term (35) is non-negative, and the third term (36) is strictly positive. When $\lambda \geq 0$, from the fixed-point equation (53) we have $\phi = (1+\mu) - \lambda(1+\mu)/\mu \leq 1+\mu$ and $(1+\mu)^2 - \phi \geq \phi(1+\mu) - \phi = \phi\mu$. Then, we know that the first term (34) is non-negative. Therefore, it follows that for all $\lambda \geq 0$,

$$(1 - q_v(\Sigma, \Sigma))^2 \mathcal{R}'(\lambda, \phi) > 0.$$

This implies that $\mathcal{R}(\lambda, \phi)$ is minimized at $\lambda < 0$.

Part (3) Overparameterized regime and $\Sigma_0 = I$. When $\Sigma_0 = I$, the above derivative in Part (1) becomes

$$\begin{aligned} & (1 - q_v(\Sigma, \Sigma))^2 \mathcal{B}'(\lambda, \phi) \\ &= \frac{\lambda}{\mu} \left(q'_v(I, \Sigma) q_b(\Sigma, B) + q_v(I, \Sigma) q'_b(\Sigma, B) + \frac{\lambda}{\mu} q'_b(I, B) + 2\mu q'_b(I, B) q_v(\Sigma, I) \right) \\ & \quad + [\mu q'_v(I, \Sigma) q_b(\Sigma, B) q_v(\Sigma, I) + \mu q_v(I, \Sigma) q'_b(\Sigma, B) q_v(\Sigma, I)] \end{aligned} \quad (37)$$

$$+ q_v(I, \Sigma) q_b(\Sigma, B) q'_v(\Sigma, \Sigma) + \mu^2 q'_b(I, B) q_v(\Sigma, I)^2] \quad (38)$$

Recalling the definitions of $q_b(\cdot, \cdot)$ and $q_v(\cdot, \cdot)$ from (21) and (22), observe that

$$\begin{aligned} q_b(I, B) &= \mu^2 \bar{\text{tr}}[(\Sigma + \mu I)^{-2} B] \\ q'_b(I, B) &= 2\mu \bar{\text{tr}}[(\Sigma + \mu I)^{-2} B] - 2\mu^2 \bar{\text{tr}}[(\Sigma + \mu I)^{-3} B] = \frac{2}{\mu} q_b(I, B) - 2\mu^2 \bar{\text{tr}}[(\Sigma + \mu I)^{-3} B] \\ &= 2\mu \bar{\text{tr}}[(\Sigma + \mu I)^{-3} \Sigma B], \\ q_v(I, \Sigma) &= \phi \bar{\text{tr}}[(\Sigma + \mu I)^{-2} \Sigma] \\ q'_v(I, \Sigma) &= -2\phi \bar{\text{tr}}[(\Sigma + \mu I)^{-3} \Sigma]. \end{aligned}$$

Next, we work on terms (37) and (38) that do not involve a factor of λ/μ .

$$\begin{aligned} & \mu q'_v(I, \Sigma) q_b(\Sigma, B) q_v(\Sigma, I) \\ & \quad + \mu q_v(I, \Sigma) q'_b(\Sigma, B) q_v(\Sigma, I) + q_v(I, \Sigma) q_b(\Sigma, B) q'_v(\Sigma, \Sigma) \\ & \quad + \mu^2 q'_b(I, B) q_v(\Sigma, I)^2 \\ &= -2\mu\phi \bar{\text{tr}}[(\Sigma + \mu I)^{-3} \Sigma] q_b(\Sigma, B) q_v(\Sigma, I) \\ & \quad + \mu \left(\frac{2}{\mu} q_b(\Sigma, B) - 2\mu^2 \bar{\text{tr}}[(\Sigma + \mu I)^{-3} \Sigma B] \right) q_v(I, \Sigma) q_v(\Sigma, I) \\ & \quad - 2\phi \bar{\text{tr}}[(\Sigma + \mu I)^{-3} \Sigma^2] q_v(I, \Sigma) q_b(\Sigma, B) \\ & \quad + \mu^2 \left(\frac{2}{\mu} q_b(I, B) - 2\mu^2 \bar{\text{tr}}[(\Sigma + \mu I)^{-3} B] \right) q_v(\Sigma, I)^2 \\ &= 2\mu\phi \left(-\bar{\text{tr}}[(\Sigma + \mu I)^{-3} \Sigma] + \frac{1}{\mu} \bar{\text{tr}}[(\Sigma + \mu I)^{-2} \Sigma] - \bar{\text{tr}}[(\Sigma + \mu I)^{-3} \Sigma^2] \right) q_b(\Sigma, B) q_v(\Sigma, \Sigma) \\ & \quad + 2\mu^2 \left(\frac{1}{\mu} q_b(I, B) - \mu^2 \bar{\text{tr}}[(\Sigma + \mu I)^{-3} B] - \mu \bar{\text{tr}}[(\Sigma + \mu I)^{-3} \Sigma B] \right) q_v(\Sigma, I)^2 \\ &= 0. \end{aligned}$$

This implies that $\mathcal{B}'(0, \phi) = 0$ and

$$\begin{aligned} & (1 - q_v(\Sigma, \Sigma))^2 \mathcal{B}'(\lambda, \phi) \\ &= \frac{\lambda}{\mu} (q'_v(I, \Sigma) q_b(\Sigma, B) + q_v(I, \Sigma) q'_b(\Sigma, B) + 2\mu q'_b(I, B) q_v(\Sigma, I)) + \frac{\lambda^2}{\mu^2} q'_b(I, B) \\ &= \frac{\lambda}{\mu} (q'_v(I, \Sigma) q_b(\Sigma, B) + q_v(I, \Sigma) q'_b(\Sigma, B) + 2\mu q'_b(I, B) q_v(\Sigma, I)) + \frac{\lambda^2}{\mu^2} q'_b(I, B) \\ &= \frac{\lambda}{\mu} \left(\frac{2}{\mu} (-\mu\phi \bar{\text{tr}}[(\Sigma + \mu I)^{-3} \Sigma] + q_v(\Sigma, I)) q_b(\Sigma, B) - 2\mu^2 \bar{\text{tr}}[(\Sigma + \mu I)^{-3} \Sigma B] q_v(\Sigma, I) \right) \end{aligned}$$

$$\begin{aligned}
& + 4\mu^2 \bar{\text{tr}}[(\Sigma + \mu I)^{-3}\Sigma B]q_v(\Sigma, I) \Big) + \frac{\lambda^2}{\mu^2}q'_b(I, B) \\
& = \frac{\lambda}{\mu} \left(\frac{2}{\mu}\phi \bar{\text{tr}}[(\Sigma + \mu I)^{-3}\Sigma^2]q_b(\Sigma, B) + 2\mu^2 \bar{\text{tr}}[(\Sigma + \mu I)^{-3}\Sigma B]q_v(\Sigma, I) \right) + \frac{\lambda^2}{\mu^2}q'_b(I, B) \\
& = 2\phi\lambda \left(\bar{\text{tr}}[(\Sigma + \mu I)^{-3}\Sigma^2] \bar{\text{tr}}[(\Sigma + \mu I)^{-2}\Sigma B] + \mu \bar{\text{tr}}[(\Sigma + \mu I)^{-3}\Sigma B] \bar{\text{tr}}[(\Sigma + \mu I)^{-2}\Sigma] \right) \\
& \quad + 2\lambda (1 - \phi \bar{\text{tr}}[(\Sigma + \mu I)^{-1}\Sigma]) \bar{\text{tr}}[(\Sigma + \mu I)^{-3}\Sigma B] \\
& = 2\phi\lambda \left(\bar{\text{tr}}[(\Sigma + \mu I)^{-3}\Sigma^2] \bar{\text{tr}}[(\Sigma + \mu I)^{-2}\Sigma B] - \bar{\text{tr}}[(\Sigma + \mu I)^{-3}\Sigma B] \bar{\text{tr}}[(\Sigma + \mu I)^{-2}\Sigma^2] \right) \\
& \quad + 2\lambda \bar{\text{tr}}[(\Sigma + \mu I)^{-3}\Sigma B] \\
& = 2\phi\lambda \bar{\text{tr}}[(\Sigma + \mu I)^{-3}\Sigma^2] \bar{\text{tr}}[(\Sigma + \mu I)^{-2}\Sigma B] \\
& \quad + 2\lambda \bar{\text{tr}}[(\Sigma + \mu I)^{-3}\Sigma B](1 - \phi \bar{\text{tr}}[(\Sigma + \mu I)^{-2}\Sigma^2]).
\end{aligned}$$

From Lemma F.3 (3), we know that $1 - \phi \bar{\text{tr}}[(\Sigma + \mu I)^{-2}\Sigma^2] \geq 0$. When $\phi > 1$, it then follows that \mathcal{B}' is strictly negative for all $\lambda \in [\lambda_{\min}(\phi), 0)$ and strictly positive for all $\lambda > 0$ because $\mu(\lambda, \phi) > 0$.

For the variance term, from Lemma D.2, we have

$$\begin{aligned}
& (1 - q_v(\Sigma, \Sigma))^2 \mathcal{V}'(\lambda, \phi) \\
& = \sigma^2 [q'_v(I, \Sigma)(1 - q_v(\Sigma, \Sigma)) + q_v(I, \Sigma)q'_v(\Sigma, \Sigma)] \\
& = \sigma^2 \left[q'_v(I, \Sigma) \left(\frac{\lambda}{\mu} + \mu q_v(\Sigma, I) \right) + q_v(I, \Sigma)q'_v(\Sigma, \Sigma) \right] \\
& = -2\phi\sigma^2 \left[\bar{\text{tr}}[(\Sigma + \mu I)^{-3}\Sigma] \left(\frac{\lambda}{\mu} + \mu q_v(\Sigma, I) \right) + \bar{\text{tr}}[(\Sigma + \mu I)^{-3}\Sigma^2]q_v(I, \Sigma) \right] \\
& = -2\phi\sigma^2 \left[\frac{\lambda}{\mu} \bar{\text{tr}}[(\Sigma + \mu I)^{-3}\Sigma] + \bar{\text{tr}}[(\Sigma + \mu I)^{-2}\Sigma]q_v(\Sigma, I) \right] \\
& = -2\phi\sigma^2 \left[(1 - \phi \bar{\text{tr}}[(\Sigma + \mu I)^{-1}\Sigma]) \bar{\text{tr}}[(\Sigma + \mu I)^{-3}\Sigma] + \bar{\text{tr}}[(\Sigma + \mu I)^{-2}\Sigma]q_v(\Sigma, I) \right] \\
& = -2\phi\sigma^2 \left[\bar{\text{tr}}[(\Sigma + \mu I)^{-3}\Sigma] + \phi(\bar{\text{tr}}[(\Sigma + \mu I)^{-2}\Sigma]^2 - \bar{\text{tr}}[(\Sigma + \mu I)^{-1}\Sigma] \bar{\text{tr}}[(\Sigma + \mu I)^{-3}\Sigma]) \right],
\end{aligned}$$

which is strictly negative for all $\lambda \geq \lambda_{\min}(\phi)$.

Combining the above two derivatives, we conclude that $\lambda^* > 0$ in this case. \square

D.5 Proof of Theorem 5

Theorem 5 (Optimal regularization under regression shift). Assume the setup of Proposition 1 with $\Sigma_0 = \Sigma$, $\beta_0 \neq \beta$.

1. (*Underparameterized*) When $\phi < 1$, if $\sigma^2 = o(1)$ and for all $\mu \geq 0$, the following general alignment holds:

$$\bar{\text{tr}}[B_0\Sigma^2(\Sigma + \mu I)^{-2}] > \bar{\text{tr}}[B\Sigma^2(\Sigma + \mu I)^{-2}], \quad (11)$$

where $B = \beta\beta^\top$ and $B_0 = \beta_0\beta_0^\top$, then we have $\lambda^* < 0$.

2. (*Overparameterized*) When $\phi > 1$, if the general alignment conditions (8) and (11) hold, then we have $\lambda^* < 0$.

Proof. We split the proof into two parts.

Part (1) Underparameterized regime. From the proof of Theorem 2, we know that when $\Sigma_0 = \Sigma$, the bias term satisfies that

$$\min_{\lambda \in [\lambda_{\min}(\phi), 0]} \mathcal{B}(\lambda, \phi) \geq \mathcal{B}(0, \phi) = 0,$$

from (24). From Lemma D.2, the excess bias term has the following derivative:

$$\mathcal{S}'(\lambda, \phi) = -2\beta^\top \Sigma(\Sigma + \mu I)^{-2} \Sigma(\beta_0 - \beta) = -2\beta^\top \Sigma(\Sigma + \mu I)^{-2} \Sigma(\beta - \beta_0),$$

which is zero when $\lambda = 0$ because $\mu(0, \phi) = 0$ when $\phi < 1$, from Lemma F.2. Under the condition that $\beta^\top \Sigma(\Sigma + \mu I)^{-2} \Sigma(\beta - \beta_0) < 0$ for all $\mu \geq 0$, we have that $\mathcal{S}(\lambda, \phi)$ is strictly increasing in $\lambda \geq 0$.

From (31), we have

$$\mathcal{V}'(\lambda, \phi) = \frac{-2\sigma^2 \phi \bar{\text{tr}}[\Sigma^2(\Sigma + \mu I)^{-3}]}{(1 - q_v(\Sigma, \Sigma))^2}.$$

Also, since $\mathcal{B} \geq 0$ with equality holds if $\lambda = 0$. Then we have, if

$$\mathcal{S}'(\lambda, \phi) + \mathcal{V}'(\lambda, \phi) = -2\beta^\top \Sigma(\Sigma + \mu I)^{-2} \Sigma(\beta - \beta_0) + \frac{-2\sigma^2 \phi \bar{\text{tr}}[\Sigma^2(\Sigma + \mu I)^{-3}]}{(1 - q_v(\Sigma, \Sigma))^2} > 0 \quad (39)$$

for all $\mu \geq \mu(0, \phi)$, then $\lambda^* < 0$. Note that when $\lambda \rightarrow +\infty$, $\mu \rightarrow +\infty$ and the denominator of the second term tends to one, so we have $\mathcal{V}'(\lambda, \phi) \asymp \mu^{-3}$. On the other hand, the first term scales as $\mathcal{S}'(\lambda, \phi) \asymp \mu^{-2}$. Eventually, the first term dominates. Thus, the condition (39) could hold when $\mathcal{S}'(\lambda, \phi)$ is positive and large enough.

Especially, when $\sigma^2 = o(1)$ and the assumed alignment condition is met, it follows that $\mathcal{R}(\lambda, \phi) = \mathcal{B}(\lambda, \phi) + \mathcal{V}(\lambda, \phi) + \mathcal{S}(\lambda, \phi)$ is minimized over $\lambda < 0$.

Part (2) Overparameterized regime. From the proof of Theorem 2, we know that when (8), $\mathcal{B}(\lambda, \phi) + \mathcal{V}(\lambda, \phi)$ is minimized at $\lambda < 0$. Under the condition that $\beta^\top \Sigma(\Sigma + \mu I)^{-2} \Sigma(\beta - \beta_0) \leq 0$ for all $\mu > 0$, we have $\mathcal{S}'(\lambda, \phi) \geq 0$ over $\lambda \in (\lambda_{\min}(\phi), +\infty)$. This implies that $\mathcal{S}(\lambda, \phi)$ is increasing over $\lambda \in (\lambda_{\min}(\phi), +\infty)$. Combining the two results, we further see that the risk $\mathcal{R}(\lambda, \phi)$ is minimized at $\lambda < 0$. \square

D.6 Helper Lemmas

Lemma D.2 (Out-of-distribution risk derivatives). *Under the same conditions as in Proposition 1, we have*

$$\frac{\partial \mathcal{R}(\lambda, \phi)}{\partial \lambda} = (\mathcal{B}'(\lambda, \phi) + \mathcal{V}'(\lambda, \phi) + \mathcal{S}'(\lambda, \phi)) \frac{\partial \mu}{\partial \lambda}, \quad (40)$$

where

$$\begin{aligned} \mathcal{B}'(\lambda, \phi) &:= \frac{\partial \mathcal{B}(\lambda, \phi)}{\partial \mu} = \frac{1}{(1 - q_v(\Sigma, \Sigma))^2} \left[q'_v(\Sigma_0, \Sigma) q_b(\Sigma, B) (1 - q_v(\Sigma, \Sigma)) \right. \\ &\quad \left. + q_v(\Sigma_0, \Sigma) q'_b(\Sigma, B) (1 - q_v(\Sigma, \Sigma)) + q_v(\Sigma_0, \Sigma) q_b(\Sigma, B) q'_v(\Sigma, \Sigma) \right. \\ &\quad \left. + q'_b(\Sigma_0, B) (1 - q_v(\Sigma, \Sigma))^2 \right] \\ \mathcal{V}'(\lambda, \phi) &:= \frac{\partial \mathcal{V}(\lambda, \phi)}{\partial \mu} = \sigma^2 \frac{q'_v(\Sigma_0, \Sigma) - q'_v(\Sigma_0, \Sigma) q_v(\Sigma, \Sigma) + q'_v(\Sigma, \Sigma) q_v(\Sigma_0, \Sigma)}{(1 - q_v(\Sigma, \Sigma))^2} \end{aligned}$$

$$\mathcal{S}'(\lambda, \phi) := \frac{\partial \mathcal{S}(\lambda, \phi)}{\partial \mu} = -\frac{2}{\mu^2} q_b(\Sigma_0, (B - B_0)\Sigma).$$

and $\mu = 1/v_p(\lambda, \phi)$. Furthermore, we have $\mathcal{V}'(\lambda, \phi) < 0$ for all $\lambda \in (\lambda_{\min}(\phi), +\infty)$.

Proof. We split the proof into different parts.

Part (1) Bias term. Recall the expression for bias for out-of-distribution squared risk:

$$\mathcal{B}(\lambda, \phi) = \frac{q_v(\Sigma_0, \Sigma)}{1 - q_v(\Sigma, \Sigma)} q_b(\Sigma, B) + q_b(\Sigma_0, B)$$

Define

$$h_1(\mu) = q_v(\Sigma_0, \Sigma), \quad h_2(\mu) = \frac{q_b(\Sigma, B)}{1 - q_v(\Sigma, \Sigma)}, \quad h_3(\mu) = q_b(\Sigma_0, B). \quad (41)$$

Then we have

$$\mathcal{B}(\lambda, \phi) = h_1(\mu)h_2(\mu) + h_3(\mu),$$

and

$$\begin{aligned} \mathcal{B}'(\lambda, \phi) &= h_1'(\mu)h_2(\mu) + h_1(\mu)h_2'(\mu) + h_3'(\mu) \\ &= q_v'(\Sigma_0, \Sigma) \frac{q_b(\Sigma, B)}{1 - q_v(\Sigma, \Sigma)} + q_v(\Sigma_0, \Sigma) \frac{q_b'(\Sigma, B)(1 - q_v(\Sigma, \Sigma)) + q_b(\Sigma, B)q_v'(\Sigma, \Sigma)}{(1 - q_v(\Sigma, \Sigma))^2} + q_b'(\Sigma_0, B) \\ &= \frac{1}{(1 - q_v(\Sigma, \Sigma))^2} \left[q_v'(\Sigma_0, \Sigma)q_b(\Sigma, B)(1 - q_v(\Sigma, \Sigma)) \right. \\ &\quad \left. + q_v(\Sigma_0, \Sigma)q_b'(\Sigma, B)(1 - q_v(\Sigma, \Sigma)) + q_v(\Sigma_0, \Sigma)q_b(\Sigma, B)q_v'(\Sigma, \Sigma) \right. \\ &\quad \left. + q_b'(\Sigma_0, B)(1 - q_v(\Sigma, \Sigma))^2 \right]. \end{aligned}$$

Part (2) Variance term. Recall that the variance term is given by:

$$\mathcal{V}(\lambda, \phi) = \sigma^2 \frac{q_v(\Sigma_0, \Sigma)}{1 - q_v(\Sigma, \Sigma)}.$$

The derivative in μ is:

$$\mathcal{V}' = \sigma^2 \frac{1}{(1 - q_v(\Sigma, \Sigma))^2} \{q_v'(\Sigma_0, \Sigma)(1 - q_v(\Sigma, \Sigma)) + q_v(\Sigma_0, \Sigma)q_v'(\Sigma, \Sigma)\}.$$

Note that

$$\begin{aligned} q_v'(\Sigma, \Sigma) &= -2\phi \bar{\text{tr}}[\Sigma^2(\Sigma + \mu I)^{-2}] < 0 \\ q_v'(\Sigma_0, \Sigma) &= -2\phi \bar{\text{tr}}[\Sigma_0 \Sigma(\Sigma + \mu I)^{-2}] = -2\phi \bar{\text{tr}}[(\Sigma + \mu I)^{-1} \Sigma^{1/2} \Sigma_0 \Sigma^{1/2} (\Sigma + \mu I)^{-1}] < 0. \end{aligned}$$

From Lemma F.3, we also have $q_v(\Sigma, \Sigma) > 0$ and $1 - q_v(\Sigma, \Sigma) > 0$ when $\lambda \in (\lambda_{\min}(\phi), +\infty)$. Therefore, it holds that

$$\mathcal{V}'(\lambda, \phi) < 0.$$

Part (3) Extra bias term. Recall that the extra bias term is given by

$$\mathcal{S}(\lambda, \phi) = -2\beta^\top (v\Sigma + I)^{-1} \Sigma_0 (\beta - \beta_0) = -2\mu\beta^\top (\Sigma + \mu I)^{-1} \Sigma_0 (\beta - \beta_0) = l(\Sigma_0, (B - B_0)\Sigma).$$

Then, the derivative is given by:

$$\mathcal{S}'(\lambda, \phi) = -2\beta^\top \Sigma (\Sigma + \mu I)^{-2} \Sigma_0 (\beta - \beta_0) = -\frac{2}{\mu^2} q_b(\Sigma_0, (B - B_0)\Sigma).$$

□

E Proofs in Section 4

E.1 Proof of Proposition 6

Proposition 6 (Optimal risk under isotropic signals). When $\Sigma_0 \neq \Sigma$ and $\beta = \beta_0$, assuming isotropic signals $\mathbb{E}[\beta\beta^\top] = (\alpha^2/p)I$ the optimal risk obtained at $\lambda^*(\phi) = \phi/\text{SNR}$ is given by:

$$\mathbb{E}_\beta[\mathcal{R}(\lambda^*, \phi)] = \alpha^2 \mu^* \bar{\text{tr}}[\Sigma_0(\Sigma + \mu^* I)^{-1}] + \sigma_0^2, \quad (12)$$

where $\mu^* = \mu(\lambda^*, \phi)$. Furthermore, the left side of (12) is strictly increasing in ϕ if $\text{SNR} \in (0, \infty)$ and σ_0^2 are fixed and strictly increasing in SNR if ϕ , σ^2 , and σ_0^2 are fixed.

Proof. When $\beta = \beta_0$ and $\beta\beta^\top \simeq \alpha^2 I$, from (27), we have that

$$\begin{aligned} \mathcal{B}(\lambda, \phi) &= \alpha^2 \frac{q_v(\Sigma_0, \Sigma)}{1 - q_v(\Sigma, \Sigma)} q_b(\Sigma, I) + \alpha^2 q_b(\Sigma_0, I) \\ &= \alpha^2 \frac{\mu^2}{\phi} \left(\frac{q_v(\Sigma_0, \Sigma)}{1 - q_v(\Sigma, \Sigma)} q_v(\Sigma, I) + q_v(\Sigma_0, I) \right) \\ &= \alpha^2 \frac{\mu}{\phi} \left(\frac{q_v(\Sigma_0, \Sigma)}{1 - q_v(\Sigma, \Sigma)} \left(1 - q_v(\Sigma, \Sigma) - \frac{\lambda}{\mu} \right) + \mu q_v(\Sigma_0, I) \right) \\ &= \alpha^2 \frac{\mu}{\phi} [q_v(\Sigma_0, \Sigma) + \mu q_v(\Sigma_0, I)] - \alpha^2 \frac{\lambda}{\phi} \frac{q_v(\Sigma_0, \Sigma)}{1 - q_v(\Sigma, \Sigma)} \\ &= \alpha^2 \mu \bar{\text{tr}}[\Sigma_0(\Sigma + \mu I)^{-1}] - \alpha^2 \frac{\lambda}{\phi} \frac{q_v(\Sigma_0, \Sigma)}{1 - q_v(\Sigma, \Sigma)}. \end{aligned}$$

Therefore, from Proposition 1 and Theorem 5, the optimal risk is given by

$$\begin{aligned} R(\hat{\beta}^{\lambda^*}) &\simeq \mathcal{B}_p(\lambda^*, \phi) + \mathcal{V}_p(\lambda^*, \phi) + \sigma_0^2 \\ &= \alpha^2 \mu^* \bar{\text{tr}}[\Sigma_0(\Sigma + \mu^* I)^{-1}] + \left(\sigma^2 - \frac{\alpha^2 \lambda^*}{\phi} \right) \frac{q_v(\Sigma_0, \Sigma)}{1 - q_v(\Sigma, \Sigma)} + \sigma_0^2 \\ &= \alpha^2 \mu^* \bar{\text{tr}}[\Sigma_0(\Sigma + \mu^* I)^{-1}] + \sigma_0^2 \\ &= \alpha^2 \bar{\text{tr}}[\Sigma_0(v^* \Sigma + I)^{-1}] + \sigma_0^2, \end{aligned}$$

where $\mu^* = \mu(\lambda^*, \phi)$ and $v^* = v(\lambda^*, \phi)$.

Note that when σ^2 and σ_0^2 are fixed, $\mathcal{R}(\lambda^*, \phi) = \phi \sigma^2 \cdot (\mu^*/\lambda^*) \cdot \bar{\text{tr}}[\Sigma_0(\Sigma + \mu^* I)^{-1}] + \sigma_0^2$ is simply a function of λ^* . From Lemma F.2 (4), we have that μ^*/λ^* is strictly decreasing in λ^* . Also note that $\bar{\text{tr}}[\Sigma_0(\Sigma + \mu^* I)^{-1}]$ is strictly decreasing in λ^* . Thus, we know that $\mathcal{R}(\lambda^*, \phi)$ is strictly decreasing in λ^* when ϕ, σ, σ_0 are fixed. Because λ^* is strictly decreasing in SNR , we further find that $\mathcal{R}(\lambda^*, \phi)$ is strictly increasing in SNR . □

E.2 Proof of Theorem 7

Theorem 7 (Monotonicity of optimally tuned OOD risk). For $\lambda \geq \lambda_{\min}(\phi)$ where $\lambda_{\min}(\phi)$ is as in (5), for all $\epsilon > 0$ small enough, the risk of optimal ridge predictor satisfies:

$$\min_{\lambda \geq \lambda_{\min}(\phi) + \epsilon} R(\widehat{\beta}^\lambda) \simeq \min_{\lambda \geq \lambda_{\min}(\phi)} \mathcal{R}(\lambda, \phi), \quad (13)$$

and right side of (13) is monotonically increasing in ϕ if SNR and σ_0^2 are fixed. In addition, when $\beta = \beta_0$ it is monotonically increasing in SNR if ϕ , σ^2 , and σ_0^2 are fixed.

Proof. We split the proof into different parts.

Part (1) Risk characterization and equivalence. From the proof of Theorem 9, we have

$$R(\widehat{\beta}^\lambda) \simeq \mathcal{R}_p(\lambda, \phi, \phi)$$

where \mathcal{R}_p is defined in (50). Furthermore, for any $\bar{\psi} \in [\phi, +\infty]$, there exists a unique $\bar{\lambda} \geq \lambda_{\min}(\phi)$ defined through (44). 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 (55), we have that:

$$\mathcal{R}_p(\lambda_1; \phi, \psi_1) = \mathcal{R}_p(\lambda_2; \phi, \psi_2).$$

Part (2) Risk monotonicity. From Lemma F.3 (3), we know that the denominator of \tilde{v}_p defined in (19) is non-negative:

$$v_p(\lambda, \psi)^{-2} - \phi \int r^2 (1 + v_p(\lambda, \psi)r)^{-2} dH_p(r) \geq v_p(\lambda, \psi)^{-2} - \psi \int r^2 (1 + v_p(\lambda, \psi)r)^{-2} dH_p(r) \geq 0$$

for $\lambda \geq \lambda_{\min}(\phi)$. Therefore, $\mathcal{R}_p(\lambda, \phi, \psi)$ is increasing in ϕ for any fixed (λ, ψ) . Furthermore, since $\mathcal{R}_p(\lambda, \phi, \psi)$ is a continuous function of ϕ and $v(\lambda; \psi)$, it follows that for $0 < \phi_1 \leq \phi_2 < \infty$,

$$\min_{\psi \geq \phi_1} \mathcal{R}_p(\lambda, \phi_1, \psi) \leq \min_{\psi \geq \phi_2} \mathcal{R}_p(\lambda, \phi_1, \psi) \leq \min_{\psi \geq \phi_2} \mathcal{R}_p(\lambda, \phi_2, \psi),$$

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

Part (3) Optimal subsampling and regularization. Similar to the proof of Part (3) in Lemma E.1, we have

$$\min_{\psi \geq \phi} \mathcal{R}_p(0; \phi, \psi) \simeq \min_{\lambda \geq \lambda_{\min}(\phi) + \epsilon} R(\widehat{\beta}^\lambda).$$

From Part (3), we know that the former is continuous and monotonically increasing in ϕ , which finishes the proof.

Part (4) Monotonicity in signal-to-noise ratio. When $\beta_0 = \beta$, the extra bias term \mathcal{S} is zero. When σ^2 is fixed, note that \mathcal{B} and κ are strictly increasing in α^2 while \mathcal{V} does not depend on α^2 . Thus, we know that \mathcal{R}_p is strictly increasing in SNR. Consequently, $\min_{\psi \geq \phi} \mathcal{R}_p(0; \phi, \psi)$ is strictly increasing in SNR. \square

E.3 Proof of Theorem 8

Theorem 8 (Non-monotonicity of suboptimally tuned risk). When $(\Sigma_0, \beta_0) = (\Sigma, \beta)$ and $\Sigma = I$, the risk component equivalents defined in (6) have the following properties:

1. (*Bias component*) For all $\lambda > 0$, $\mathcal{B}(\lambda, \phi)$ is strictly increasing over $\phi \in (0, \lambda + 1)$ and strictly decreasing over $\phi \in (\lambda + 1, \infty)$.
2. (*Variance component*) For all $\lambda > 0$, $\mathcal{V}(\lambda, \phi)$ is strictly increasing over $\phi \in (0, \infty)$.
3. (*Risk*) When $\|\beta\|_2^2 > 0$, for all $\lambda > 0$ and $\epsilon > 0$, there exist $\sigma^2, \phi \in (0, \infty)$, such that $\partial \mathcal{R}(\lambda, \phi) / \partial \phi \leq -\epsilon$, i.e.,

$$\max_{\sigma^2, \phi \in (0, \infty)} \min_{\lambda \geq \lambda_{\min}(\phi)} \partial \mathcal{R}(\lambda, \phi) / \partial \phi \leq -\epsilon. \quad (14)$$

Proof. For isotropic features $\Sigma = I$, analogous to the proof of Theorem 7, the excess prediction risk is given by

$$R(\widehat{\beta}^\lambda) \simeq R_p(\lambda, \phi, \phi)$$

where $R_p(\lambda, \phi, \psi)$ is defined in (6). When $\Sigma = I$, the non-negative constants $\tilde{c}_p(\lambda, \psi)$ and $\tilde{v}_p(\lambda, \phi, \psi)$ are defined through the following equations:

$$\begin{aligned} v_p(\lambda, \phi) &= \frac{\sqrt{(\phi + \lambda - 1)^2 + 4\lambda} - (\phi + \lambda - 1)}{2\lambda}, \\ \tilde{v}_p(\lambda, \phi, \psi) &= \frac{\frac{1}{\phi(1 + v_p(\lambda, \psi))^2}}{v_p(\lambda, \psi)^{-2} - \frac{1}{\phi(1 + v_p(\lambda, \psi))^2}}, \\ \tilde{c}_p(\lambda, \psi) &= (v_p(\lambda, \psi)\Sigma + I)^{-2}\alpha^2. \end{aligned}$$

From Theorem 1 of Patil and Du (2023), for all $(\lambda, \phi) \in (0, \infty)^2$, there exists $\psi = \psi(\lambda, \phi)$ such that the prediction risk (45) of the full-ensemble estimator are asymptotically equivalent:

$$R_p(\lambda, \phi, \phi) \simeq R_p(0; \phi, \psi).$$

Note that the left-hand side is simply the risk of the ridge predictor with ridge penalty λ . Furthermore, from (44), it holds that

$$\lambda + \phi \int \frac{r}{1 + v_p(0; \psi)r} dH(r) = \psi \int \frac{r}{1 + v_p(0; \psi)r} dH(r).$$

Taking the derivative with respect to ϕ on both sides yields that

$$\frac{\partial \psi}{\partial \phi} = 1 - \frac{\lambda v_p(0; \psi)}{1 - \lambda v_p(0; \psi)} \tilde{v}_p(0; \phi, \psi). \quad (42)$$

We consider three cases below.

(1) $\alpha^2 = 0$ and $\sigma^2 > 0$. In this case, the excess risk equals the variance component:

$$\begin{aligned} R_p'(\lambda, \phi, \phi) &= \sigma^2 \tilde{v}_p(\lambda, \phi, \phi) \\ &= \sigma^2 \left(\frac{1}{1 - \phi \int \left(\frac{v_p(\lambda, \phi)r}{1 + v_p(\lambda, \phi)r} \right)^2 dH_p(r)} - 1 \right). \end{aligned}$$

Let $f(\phi) := \phi \int (v_p(\lambda, \phi)r / (1 + v_p(\lambda, \phi)r))^2 dH_p(r)$. Then, the monotonicity of the above display in ϕ is the same as that of f in ϕ . Note that

$$f(\phi) = \phi \left(\frac{\sqrt{(\phi + \lambda - 1)^2 + 4\lambda} - (\phi + \lambda - 1)}{\sqrt{(\phi + \lambda - 1)^2 + 4\lambda} - (\phi - \lambda - 1)} \right)^2. \quad (43)$$

Taking the derivative with respect to ϕ , we have

$$f'(\phi) = \frac{(-\phi + \lambda + 1) \left(\sqrt{(\phi + \lambda - 1)^2 + 4\lambda} - \phi - \lambda + 1 \right)^2}{\sqrt{(\phi + \lambda - 1)^2 + 4\lambda} \left(\sqrt{(\phi + \lambda - 1)^2 + 4\lambda} + \lambda - \phi + 1 \right)^2}.$$

Since $f'(\phi) > 0$ when $\phi \in (0, \lambda + 1)$ and $f'(\phi) < 0$ when $\phi \in (\lambda + 1, +\infty)$, we know that R_p is strictly increasing over $\phi \in (0, \lambda + 1)$ and strictly decreasing over $\phi \in (\lambda + 1, +\infty)$. Thus, the monotonicity of $R_p(\lambda, \phi, \phi)$ follows.

(2) $\alpha^2 > 0$ and $\sigma^2 = 0$. In this case, the excess risk equals the bias component:

$$\begin{aligned} R_p(\lambda, \phi, \phi) &= \tilde{c}_p(-\lambda; \phi) (\tilde{v}_p(\lambda, \phi, \phi) + 1) \\ &= \alpha^2 \frac{1}{1 - \phi \frac{1}{(1 + v_p(\lambda, \phi))^2}} \\ &= \frac{\alpha^2}{(1 - \phi)v_p(\lambda, \phi)^2 + 2v_p(\lambda, \phi) + 1}. \end{aligned}$$

Let $g(\phi) = (1 - \phi)v_p(\lambda, \phi)^2 + 2v_p(\lambda, \phi) + 1$. Taking the derivative with respect to ϕ yields

$$\begin{aligned} g'(\phi) &= \frac{\left(\sqrt{(\phi + \lambda - 1)^2 + 4\lambda} - \lambda - \phi + 1 \right)}{4\lambda^2 \sqrt{(\lambda + \phi - 1)^2 + 4\lambda}} \\ &\quad \cdot \left(\lambda + 3(\phi - 1)\sqrt{\lambda^2 + 2\lambda(\phi + 1) + \phi^2 - 2\phi + 1} - \lambda^2 - 4\lambda(\phi + 1) - 3(\phi - 1)^2 \right) \\ &=: \frac{\left(\sqrt{(\phi + \lambda - 1)^2 + 4\lambda} - \lambda - \phi + 1 \right)}{4\lambda^2 \sqrt{(\lambda + \phi - 1)^2 + 4\lambda}} h(\phi). \end{aligned}$$

By simple calculations, one can show that $h'(\phi) < 0$ and $h(\phi) \leq h(0) < -2\lambda$ when $\lambda > 0$. Therefore, we have $g'(\phi) < 0$ and R_p is strictly increasing over $\phi \in (0, \infty)$.

(3) **General cases when $\alpha^2 > 0$.** Note that

$$\begin{aligned} R_p(\lambda, \phi, \phi) &= \sigma^2 \tilde{v}_p(\lambda, \phi, \psi) + \tilde{c}_p(-\lambda; \phi)(\tilde{v}_p(\lambda, \phi, \phi) + 1) \\ &=: f_1(\phi) + f_2(\phi), \end{aligned}$$

where f_1 first increases and then decreases in ϕ , and f_2 is a strictly increasing function. Note that only f_1 depends on σ^2 . Because for any $\lambda > 0$ and $\phi \in (\lambda + 1, \infty)$, $f_1'(\phi) < 0$ and its scale is proportional to σ^2 , we have that for all $\epsilon > 0$, there exists $\sigma^2 > 0$ such that $-f_1'(\phi) > f_2'(\phi) + \epsilon$. This implies that

$$\max_{\sigma^2, \phi \in (0, \infty)} \min_{\lambda \geq 0} \frac{\partial \mathcal{R}(\lambda, \phi)}{\partial \phi} \leq -\epsilon,$$

which completes the proof. \square

E.4 Proof of Theorem 9

Theorem 9 (Optimal ensemble versus ridge regression under negative regularization). Let $\mathcal{R}^{**} := \min_{\psi \geq \phi, \lambda \geq \lambda_{\min}(\phi)} \mathcal{R}(\lambda, \phi, \psi)$. Then the following statements hold:

1. (*Underparameterized*) When $\phi < 1$ and $\beta_0 = \beta$, $\lambda^* \geq 0$,

$$\mathcal{R}^{**} = \min_{\lambda \geq 0} \mathcal{R}(\lambda, \phi, \phi) = \min_{\psi \geq \phi} \mathcal{R}(0; \phi, \psi).$$

2. (*Overparameterized*) When $\phi \geq 1$, $\lambda^* \geq \lambda_{\min}(\phi)$,

$$\mathcal{R}^{**} = \min_{\lambda \geq \lambda_{\min}(\phi)} \mathcal{R}(\lambda, \phi, \phi) = \min_{\psi \geq \phi} \mathcal{R}(\lambda_{\min}(\phi); \phi, \psi).$$

Proof. We split the proof into different parts.

Part (1) Risk characterization. From Lemma E.2, we have $R(\hat{\beta}_{k, \infty}^\lambda) \simeq \mathcal{R}_p(\lambda, \phi, \psi)$.

Part (2) Risk equivalence. From Lemma F.4, for any $\bar{\psi} \in [\phi, +\infty]$, there exists a unique $\bar{\lambda} \geq \lambda_{\min}(\phi)$ such that

$$\frac{1}{v} = \bar{\lambda} + \phi \int \frac{r}{1 + vr} dH_p(r), \quad \text{and} \quad \frac{1}{v} = \bar{\psi} \int \frac{r}{1 + vr} dH_p(r). \quad (44)$$

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 (55), we have that:

$$\mathcal{R}_p(\lambda_1; \phi, \psi_1) = \mathcal{R}_p(\lambda_2; \phi, \psi_2).$$

Part (3) Optimal risk. From (49), we have that for any $\lambda \geq \lambda_{\min}(\phi)$, there exists $\psi \geq 1$ such that $|\mathcal{R}_p(\lambda_{\min}(\phi); \phi, \psi) - \mathcal{R}_p(\lambda, \phi, \phi)| \xrightarrow{\text{a.s.}} 0$. From Lemma F.1 and Lemma F.3, $1/v_p(-\lambda; \psi) \in [-r_{\min}, \infty]$ and

$$\lim_{\psi \rightarrow \phi} \mathcal{R}_p(\lambda_{\min}(\phi); \phi, \psi) = \lim_{\lambda \rightarrow \lambda_{\min}(\phi)^+} \mathcal{R}_p(\lambda, \phi, \phi) = +\infty.$$

Similar to the proof of Lemma 28 from Bellec et al. (2023), one can show that the sequence of functions $\{\mathcal{R}_p(\lambda_{\min}(\phi); \phi, \psi(\lambda)) - \mathcal{R}_p(\lambda, \phi, \phi)\}_{p \in \mathbb{N}}$ is uniformly equicontinuous on $\lambda \in \Lambda =$

$[\lambda_{\min}(\phi) + \epsilon, \infty]$ almost surely for some small $\epsilon > 0$ such that $\mathcal{R}_p(\lambda, \phi, \psi)$ is no larger than the null risk $\mathcal{R}_p(+\infty; \phi, \phi)$ when $\lambda \in [\lambda_{\min}(\phi) + \epsilon, +\infty]$. From Theorem 21.8 of [Davidson \(1994\)](#), it further follows that the sequences converge to zero uniformly over Λ almost surely. This implies that

$$\begin{aligned}
0 &= \limsup_p \max_{\lambda \geq \lambda_{\min}(\phi) + \epsilon} [\mathcal{R}_p(\lambda_{\min}(\phi); \phi, \psi(\lambda)) - \mathcal{R}_p(\lambda, \phi, \phi)] \\
&\geq \min_{\psi \geq \phi} \limsup_p \max_{\lambda \geq \lambda_{\min}(\phi) + \epsilon} [\mathcal{R}_p(\lambda_{\min}(\phi); \phi, \psi) - \mathcal{R}_p(\lambda, \phi, \phi)] \\
&\geq \limsup_p \min_{\psi \geq \phi} \max_{\lambda \geq \lambda_{\min}(\phi) + \epsilon} [\mathcal{R}_p(\lambda_{\min}(\phi); \phi, \psi) - \mathcal{R}_p(\lambda, \phi, \phi)] \\
&= \limsup_p \left[\min_{\psi \geq \phi} \mathcal{R}_p(\lambda_{\min}(\phi); \phi, \psi) - \min_{\lambda \geq \lambda_{\min}(\phi) + \epsilon} \mathcal{R}_p(\lambda, \phi, \phi) \right].
\end{aligned}$$

Conversely, since for any $\psi \geq \psi(\lambda_{\min}(\phi) + \epsilon)$, there exists $\lambda \geq \lambda_{\min}(\phi) + \epsilon$ such that $|\mathcal{R}_p(\lambda_{\min}(\phi); \phi, \psi) - \mathcal{R}_p(\lambda(\psi); \phi, \phi)| \xrightarrow{\text{a.s.}} 0$. Similarly, we can show that $\{\mathcal{R}_p(\lambda_{\min}(\phi); \phi, \psi) - \mathcal{R}_p(\lambda(\psi); \phi, \phi)\}_{p \in \mathbb{N}}$ is uniformly equicontinuous on $\psi \in \Psi = [\psi(\lambda_{\min}(\phi) + \epsilon), \infty]$ almost surely. This also implies that

$$\begin{aligned}
0 &= \liminf_p \min_{\psi \geq \psi(\lambda_{\min}(\phi) + \epsilon)} [\mathcal{R}_p(\lambda_{\min}(\phi); \phi, \psi) - \mathcal{R}_p(\lambda(\psi); \phi, \phi)] \\
&\leq \max_{\lambda \geq \lambda_{\min}(\phi)} \liminf_p \min_{\psi \geq \psi(\lambda_{\min}(\phi) + \epsilon)} [\mathcal{R}_p(\lambda_{\min}(\phi); \phi, \psi) - \mathcal{R}_p(\lambda, \phi, \phi)] \\
&\leq \liminf_p \max_{\lambda \geq \lambda_{\min}(\phi)} \min_{\psi \geq \psi(\lambda_{\min}(\phi) + \epsilon)} [\mathcal{R}_p(\lambda_{\min}(\phi); \phi, \psi) - \mathcal{R}_p(\lambda, \phi, \phi)] \\
&= \liminf_p \left[\min_{\psi \geq \psi(\lambda_{\min}(\phi) + \epsilon)} \mathcal{R}_p(\lambda_{\min}(\phi); \phi, \psi) - \min_{\lambda \geq \lambda_{\min}(\phi)} \mathcal{R}_p(\lambda, \phi, \phi) \right].
\end{aligned}$$

Combining the previous two inequalities implies that

$$\min_{\lambda \geq \lambda_{\min}(\phi) + \epsilon} \mathcal{R}_p(\lambda, \phi, \phi) \simeq \min_{\psi \geq \psi(\lambda_{\min}(\phi) + \epsilon)} \mathcal{R}_p(\lambda_{\min}(\phi); \phi, \psi).$$

Since

$$\begin{aligned}
\min_{\lambda \geq \lambda_{\min}(\phi) + \epsilon} \mathcal{R}_p(\lambda, \phi, \phi) &= \min_{\lambda \geq \lambda_{\min}(\phi)} \mathcal{R}_p(\lambda, \phi, \phi) \\
\min_{\psi \geq \psi(\lambda_{\min}(\phi) + \epsilon)} \mathcal{R}_p(\lambda_{\min}(\phi); \phi, \psi) &= \min_{\psi \geq \phi} \mathcal{R}_p(\lambda_{\min}(\phi); \phi, \psi),
\end{aligned}$$

we further have

$$\min_{\lambda \geq \lambda_{\min}(\phi)} \mathcal{R}_p(\lambda, \phi, \phi) = \min_{\psi \geq \phi} \mathcal{R}_p(\lambda_{\min}(\phi); \phi, \psi),$$

which holds for $\phi \in (0, \infty)$. This finishes the proof of the second conclusion.

Part (4) Optimal risk when $\phi < 1$. When $\beta_0 = \beta$, the excess bias term $\mathcal{S} \equiv 0$. From Lemma [D.2](#), we have that the bias component $\tilde{c}_p(\lambda, \phi, \phi)$ of the risk equivalent is minimized at $\lambda = 0$ when $\phi < 1$. Since $\tilde{v}_p(\lambda, \phi, \phi)$ is a strictly increasing function in $v_p(\lambda, \phi)$ and $v_p(\lambda, \phi)$ is a strictly decreasing function in λ , we see that $\tilde{v}_p(\lambda, \phi, \phi)$ is a strictly decreasing function in λ . Thus, we have that

$$\min_{\lambda \in [\lambda_{\min}, 0]} \mathcal{R}_p(\lambda, \phi, \phi) \geq \mathcal{R}_p(0; \phi, \phi).$$

This implies that

$$\min_{\lambda \geq \lambda_{\min}} \mathcal{R}_p(\lambda, \phi, \phi) \geq \min_{\lambda \geq 0} \mathcal{R}_p(0; \phi, \phi),$$

which finishes the proof of the first conclusion. \square

E.5 Helper Lemmas

Lemma E.1 (Monotonicity of generalized prediction risk with optimal ridge regularization). *Suppose Assumption A hold. Define the generalized mean squared risk for a estimator $\hat{\beta}$ as:*

$$R(\hat{\beta}; A, b, \beta_0) = \|L_{A,b}(\hat{\beta} - \beta_0)\|_2^2, \quad (45)$$

where $L_{A,b}(\beta) = A\beta + b$ is a linear functional and (A, b) is independent of (X, y) such that $\|A\|_{\text{op}}$ and $\|b\|_2$ are almost surely bounded. Then, as $k, n, p \rightarrow \infty$ such that $p/n \rightarrow \phi \in (0, \infty)$ and $p/k \rightarrow \psi \in [\phi, \infty]$, there exists a sequence of random variables $\{\mathcal{Q}_p(\lambda, \phi, \psi)\}_{p=1}^\infty$ that is asymptotically equivalent to the risk of the full-ensemble ridge predictor,

$$R(\hat{\beta}_{k,\infty}^\lambda; A, b, \beta_0) \simeq \mathcal{Q}_p(\lambda, \phi, \psi) := \tilde{c}_p(\lambda, \phi, \psi, A^\top A) + \|f_{\text{NL}}\|_{L_2}^2 \tilde{v}_p(\lambda, \phi, \psi, A^\top A), \quad (46)$$

where the non-negative constants $\tilde{c}_p(\lambda, \phi, \psi, A^\top A)$ and $\tilde{v}_p(\lambda, \phi, \psi, A^\top A)$ are defined through the following equations:

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

For the ridge predictor when $k = n$ and $\lambda_{\min}(\phi)$ defined in (5), the optimal risk equivalence $\min_{\lambda \geq \lambda_{\min}(\phi)} \mathcal{Q}_p(\lambda, \phi, \phi)$ is monotonically increasing in ϕ .

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

$$y = f_{\text{LI}} + f_{\text{NL}}, \quad (47)$$

where $f_{\text{LI}} = X\beta_0$ and $f_{\text{NL}} = [f_{\text{NL}}(x_i)]_{i \in [n]}$.

Note that

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

By Theorem 3 of Patil and Du (2023), the cross term vanishes, that is, $b^\top A(\hat{\beta}_{k,\infty}^\lambda - \beta_0) \xrightarrow{\text{a.s.}} 0$. We then have

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

Thus, it suffices to analyze $R(\hat{\beta}_{k,\infty}^\lambda; A, 0, \beta_0)$. To simplify the notation, we define

$$R_p(\lambda, \phi, \psi) := R(\hat{\beta}_{[p/\psi],\infty}^{\lambda_1}; A, 0_p, \beta_0) \quad (48)$$

to indicate the dependency solely on p and (λ, ϕ, ψ) . We split the proof into different parts.

Part (1) Risk characterization. Under Assumption A, from Equation (11) of Patil and Du (2023), we have that for $\lambda \geq 0$,

$$R(\widehat{\beta}_{k,\infty}^\lambda; A, b, \beta_0) \simeq \mathcal{Q}_p(\lambda, \phi, \psi)$$

where \mathcal{Q}_p is defined in (46). Note that for $\lambda \in [\lambda_{\min}, 0)$, the fixed-point solution $v_p(\lambda, \psi)$ satisfies the same fixed-point equation as the one for $\lambda \geq 0$. Since the above deterministic equivalent depends on λ only through $v_p(\lambda, \psi)$, it also applies to $\lambda \in [\lambda_{\min}, 0)$.

Part (2) Risk equivalence. From Lemma F.4, we have that, for any $\bar{\psi} \in [\phi, +\infty]$, there exists $\bar{\lambda}$ uniquely defined through (44). 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 (55), the generalized risk functionals (45) of the full-ensemble estimator are asymptotically equivalent:

$$R_p(\lambda_1; \phi, \psi_1) \simeq \mathcal{Q}_p(\lambda_1; \phi, \psi_1) = \mathcal{Q}_p(\lambda_2; \phi, \psi_2) \simeq R_p(\lambda_2; \phi, \psi_2). \quad (49)$$

Part (3) Risk monotonicity. Note that from Lemma F.10 (3) and Lemma F.11 (3) Du et al. (2023), $v_p(\lambda, \psi)^{-2} - \phi \int r^2(1 + v_p(\lambda, \psi)r)^{-2} dH_p(r)$ is non-negative. Then we have

$$\tilde{v}_p(\lambda, \phi, \psi, A^\top A) = \frac{\text{tr}[A^\top A \Sigma(v_p(\lambda, \psi)\Sigma + I)^{-2}]}{\phi^{-1}v_p(\lambda, \psi)^{-2} - \int \frac{r^2}{(1 + v_p(\lambda, \psi)r)^2} dH_p(r)}$$

is increasing in $\phi \in (0, \psi]$ for any fixed ψ . Thus, $\mathcal{Q}_p(\lambda, \phi, \psi)$ is increasing in ϕ for any fixed (λ, ψ) . Furthermore, since $\mathcal{Q}_p(\lambda, \phi, \psi)$ is a continuous function of ϕ and $v(\lambda; \psi)$, it follows that for $0 < \phi_1 \leq \phi_2 < \infty$,

$$\min_{\psi \geq \phi_1} \mathcal{Q}_p(\lambda, \phi_1, \psi) \leq \min_{\psi \geq \phi_2} \mathcal{Q}_p(\lambda, \phi_1, \psi) \leq \min_{\psi \geq \phi_2} \mathcal{Q}_p(\lambda, \phi_2, \psi),$$

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

Part (4) Optimal subsampling and regularization. From (49), we have that for any $\lambda \geq \lambda_{\min}(\phi)$, there exists $\psi \geq 1$ such that $|\mathcal{Q}_p(\lambda_{\min}(\phi); \phi, \psi) - R_p(\lambda, \phi, \phi)| \xrightarrow{\text{a.s.}} 0$. From Lemma F.1 and Lemma F.3, $1/v_p(-\lambda; \psi) \in [-r_{\min}, \infty]$ and $\lim_{\psi \rightarrow \phi} \mathcal{Q}_p(\lambda_{\min}(\phi); \phi, \psi) = +\infty$. Similarly to the proof of Lemma 28 from Bellec et al. (2023), one can show that the sequence of functions $\{\mathcal{Q}_p(\lambda_{\min}(\phi); \phi, \psi(\lambda)) - R_p(\lambda, \phi, \phi)\}_{p \in \mathbb{N}}$ is uniformly equicontinuous on $\lambda \in \Lambda = [\lambda_{\min}(\phi) + \epsilon, \infty]$ almost surely for some small $\epsilon > 0$ such that $|\mathcal{Q}_p(\lambda_{\min}(\phi); \phi, \psi(\lambda))|$ is not greater than the null risk $\mathcal{Q}_p(+\infty; \phi, \phi)$ when $\lambda \in [\lambda_{\min}(\phi) + \epsilon, +\infty]$. From Theorem 21.8 of Davidson (1994), it further follows that the sequences converge to zero uniformly over Λ almost surely. This implies that

$$\begin{aligned} 0 &= \limsup_p \max_{\lambda \geq \lambda_{\min}(\phi) + \epsilon} [\mathcal{Q}_p(\lambda_{\min}(\phi); \phi, \psi(\lambda)) - R_p(\lambda, \phi, \phi)] \\ &\geq \min_{\psi \geq \phi} \limsup_p \max_{\lambda \geq \lambda_{\min}(\phi) + \epsilon} [\mathcal{Q}_p(\lambda_{\min}(\phi); \phi, \psi) - R_p(\lambda, \phi, \phi)] \\ &\geq \limsup_p \min_{\psi \geq \phi} \max_{\lambda \geq \lambda_{\min}(\phi) + \epsilon} [\mathcal{Q}_p(\lambda_{\min}(\phi); \phi, \psi) - R_p(\lambda, \phi, \phi)] \\ &= \limsup_p \left[\min_{\psi \geq \phi} \mathcal{Q}_p(\lambda_{\min}(\phi); \phi, \psi) - \min_{\lambda \geq \lambda_{\min}(\phi) + \epsilon} R_p(\lambda, \phi, \phi) \right]. \end{aligned}$$

Conversely, since for any $\psi \geq \psi(\lambda_{\min}(\phi) + \epsilon)$, there exists $\lambda \geq \lambda_{\min}(\phi) + \epsilon$ such that $|\mathcal{Q}_p(\lambda_{\min}(\phi); \phi, \psi) - R_p(\lambda(\psi); \phi, \phi)| \xrightarrow{\text{a.s.}} 0$. Similarly, we can show that $\{\mathcal{Q}_p(\lambda_{\min}(\phi); \phi, \psi) - R_p(\lambda(\psi); \phi, \phi)\}_{p \in \mathbb{N}}$ is uniformly equicontinuous on $\psi \in \Psi = [\psi(\lambda_{\min}(\phi) + \epsilon), \infty]$ almost surely. This also implies that

$$\begin{aligned} 0 &= \liminf_p \min_{\psi \geq \psi(\lambda_{\min}(\phi) + \epsilon)} [\mathcal{Q}_p(\lambda_{\min}(\phi); \phi, \psi) - R_p(\lambda(\psi); \phi, \phi)] \\ &\leq \max_{\lambda \geq \lambda_{\min}(\phi)} \liminf_p \min_{\psi \geq \psi(\lambda_{\min}(\phi) + \epsilon)} [\mathcal{Q}_p(\lambda_{\min}(\phi); \phi, \psi) - R_p(\lambda, \phi, \phi)] \\ &\leq \liminf_p \max_{\lambda \geq \lambda_{\min}(\phi)} \min_{\psi \geq \psi(\lambda_{\min}(\phi) + \epsilon)} [\mathcal{Q}_p(\lambda_{\min}(\phi); \phi, \psi) - R_p(\lambda, \phi, \phi)] \\ &= \liminf_p \left[\min_{\psi \geq \psi(\lambda_{\min}(\phi) + \epsilon)} \mathcal{Q}_p(\lambda_{\min}(\phi); \phi, \psi) - \min_{\lambda \geq \lambda_{\min}(\phi)} R_p(\lambda, \phi, \phi) \right]. \end{aligned}$$

Combining the previous two inequalities implies that

$$\min_{\lambda \geq \lambda_{\min}(\phi) + \epsilon} R_p(\lambda, \phi, \phi) \simeq \min_{\psi \geq \psi(\lambda_{\min}(\phi) + \epsilon)} \mathcal{Q}_p(\lambda_{\min}(\phi); \phi, \psi).$$

Since

$$\min_{\psi \geq \psi(\lambda_{\min}(\phi) + \epsilon)} \mathcal{Q}_p(\lambda_{\min}(\phi); \phi, \psi) = \min_{\psi \geq \psi(\lambda_{\min}(\phi))} \mathcal{Q}_p(\lambda_{\min}(\phi); \phi, \psi) = \min_{\psi \geq \phi} \mathcal{Q}_p(\lambda_{\min}(\phi); \phi, \psi),$$

we further have

$$\min_{\lambda \geq \lambda_{\min}(\phi) + \epsilon} R_p(\lambda, \phi, \phi) \simeq \min_{\psi \geq \phi} \mathcal{Q}_p(\lambda_{\min}(\phi); \phi, \psi).$$

From the second part, we know that the latter is continuous and monotonically increasing in ϕ , which finishes the proof. \square

Lemma E.2 (Out-of-distribution full-ensemble risk asymptotics). *Under Assumption A, as $k, n, p \rightarrow \infty$ such that $p/n \rightarrow \phi \in (0, \infty)$ and $p/k \rightarrow \psi \in [\phi, \infty]$, for $\lambda \geq \lambda_{\min}(\phi)$, it holds that*

$$\begin{aligned} R(\widehat{\beta}_{k, \infty}^\lambda) &\simeq \mathcal{R}_p(\lambda, \phi, \psi) \\ &:= \mathcal{Q}_p(\lambda, \phi, \psi) - 2\beta^\top (v_p(\lambda, \phi)\Sigma + I)^{-1} \Sigma_0 (\beta - \beta_0) + [(\beta_0 - \beta)^\top \Sigma_0 (\beta_0 - \beta) + \sigma_0^2]. \end{aligned} \quad (50)$$

where

$$\mathcal{Q}_p(\lambda, \phi, \psi) := \tilde{c}_p(\lambda, \phi, \psi, \Sigma_0) + \|f_{\text{NL}}\|_{L_2}^2 \tilde{v}_p(\lambda, \phi, \psi, \Sigma_0),$$

where the non-negative constants $\tilde{c}_p(\lambda, \phi, \psi, \Sigma_0)$ and $\tilde{v}_p(\lambda, \phi, \psi, \Sigma_0)$ are defined through the following equations:

$$\begin{aligned} \frac{1}{v_p(\lambda, \psi)} &= \lambda + \psi \int \frac{r}{1 + v_p(\lambda, \psi)r} dH_p(r), \\ \tilde{v}_p(\lambda, \phi, \psi, \Sigma_0) &= \frac{\phi \overline{\text{tr}}[\Sigma_0 \Sigma (v_p(\lambda, \psi)\Sigma + I)^{-2}]}{v_p(\lambda, \psi)^{-2} - \phi \int \frac{r^2}{(1 + v_p(\lambda, \psi)r)^2} dH_p(r)}, \\ \tilde{c}_p(\lambda, \phi, \psi, \Sigma_0) &= \beta^\top (v_p(\lambda, \psi)\Sigma + I)^{-1} (\tilde{v}_p(\lambda, \phi, \psi, \Sigma_0)\Sigma + \Sigma_0) (v_p(\lambda, \psi)\Sigma + I)^{-1} \beta. \end{aligned}$$

Proof. Similar to the proof of Proposition 1, we have the decomposition

$$\begin{aligned}
R(\widehat{\beta}_{k,\infty}^\lambda) &= (\widehat{\beta}_{k,\infty}^\lambda - \beta_0)^\top \Sigma_0 (\widehat{\beta}_{k,\infty}^\lambda - \beta_0) + \sigma_0^2 \\
&= (\widehat{\beta}_{k,\infty}^\lambda - \beta)^\top \Sigma_0 (\widehat{\beta}_{k,\infty}^\lambda - \beta) + 2(\widehat{\beta}_{k,\infty}^\lambda - \beta)^\top \Sigma_0 (\beta - \beta_0) + [(\beta_0 - \beta)^\top \Sigma_0 (\beta_0 - \beta) + \sigma_0^2] \\
&\simeq \mathcal{Q}_p(\lambda, \phi, \psi) + 2\mathbb{E}_{I \sim \mathcal{I}_k} [(\widehat{\beta}(I) - \beta)^\top \Sigma_0 (\beta - \beta_0) \mid \mathcal{D}_n] + [(\beta_0 - \beta)^\top \Sigma_0 (\beta_0 - \beta) + \sigma_0^2] \\
&\simeq \mathcal{Q}_p(\lambda, \phi, \psi) - 2\beta^\top (v_p(\lambda, \psi)\Sigma + I)^{-1} \Sigma_0 (\beta - \beta_0) + [(\beta_0 - \beta)^\top \Sigma_0 (\beta_0 - \beta) + \sigma_0^2],
\end{aligned}$$

where the first asymptotic equivalent is from Lemma E.1 and the second is from Patil and Du (2023, Lemma D.1 (1)). This finishes the proof. \square

F Technical Lemmas

F.1 Fixed-Point Equations for Minimum Ridge Penalty

Recall that under Assumption A, the minimum ridge penalty $\lambda_{\min} = \lambda_{\min}(\phi)$ can be determined by the following equations:

$$\begin{aligned}
1 &= \phi \int \left(\frac{v_0 r}{1 + v_0 r} \right)^2 dP(r), \quad -r_{\min} < v_0^{-1} \leq 0, \\
\frac{1}{v_0} &= \lambda_{\min} + \phi \int \frac{r}{1 + v_0 r} dP(r),
\end{aligned}$$

or equivalently

$$\begin{aligned}
1 &= \phi \int \left(\frac{r}{\mu_0 + r} \right)^2 dP(r), \quad -r_{\min} < \mu_0 \leq 0, \\
\mu_0 &= \lambda_{\min} + \phi \int \frac{\mu_0 r}{\mu_0 + r} dP(r).
\end{aligned}$$

We next analyze the properties of μ_0 in ϕ .

Lemma F.1 (Continuity properties with the minimum regularization parameter). *Let $a > 0$ and $b < \infty$ be real numbers. Let P be a probability measure supported on $[a, b]$. Consider the function $v_0(\cdot) : \phi \mapsto \mu_0(\phi)$, over $(0, \infty)$, where $-a \leq \mu_0(\phi) \leq 0$ is the unique solution to the following fixed-point equation:*

$$1 = \phi \int \left(\frac{r}{\mu_0(\phi) + r} \right)^2 dP(r), \quad (51)$$

Then the following properties hold:

- (1) *The range of $\mu_0(\phi)$ is $[-a, \infty)$.*
- (2) *$\mu_0(\phi)$ is continuous and strictly increasing over $\phi \in (0, \infty)$. In addition, $\mu_0(\phi)$ has a root at $\phi = 1$.*
- (3) *$\lambda_{\min}(\phi) = \mu_0(\phi)(1 - \phi \int r/(\mu_0(\phi) + r) dP(r))$ is non-positive over $\phi \in (0, \infty)$ with zero obtained when $\phi = 1$. Furthermore, it is strictly increasing over $\phi \in (0, 1)$ and strictly decreasing over $\phi \in (1, \infty)$.*

Proof. The existence of the solution $v_0(\phi) = 1/\mu_0(\phi)$ to the fixed-point equation

$$\frac{1}{\phi} = \int \left(\frac{v_0(\phi)r}{1 + v_0(\phi)r} \right)^2 dP(r)$$

follows from Theorem 3.1 of [LeJeune et al. \(2024\)](#). Next, we split the proof into different parts.

Part (1). Define $h(x) = \int (xr)^2/(1+xr)^2 dP(r)$.

When $\phi < 1$, $h(v_0(\phi)) = 1/\phi > 1$, which implies that $1/v_0(\phi) \in [-a, 0)$. When $\phi = +\infty$, $1/v_0(\phi) = -a$ if $P(a) > 0$.

When $\phi \geq 1$, $h(v_0(\phi)) = 1/\phi \leq 1$, which implies that $v_0(\phi) \in (0, \infty]$, with infinity obtained when $\phi = 1$, or equivalently, $1/v_0(\phi) \in [0, \infty)$.

Part (2). Since $g(t) = h(t^{-1})^{-1}$ is positive, strictly increasing and continuous over $t \in [-a, \infty)$, by the continuous inverse theorem, we have that $1/v_0(\phi) = g^{-1}(\phi)$ is strictly increasing and continuous over $\phi \in (0, \infty)$. From Part (1), we also have $1/v_0(1) = 0$, which finishes the proof.

Part (3). Consider the function $f(x) = 1 - \phi \int r/(x+r) dP(r)$. When $\phi < 1$, we know that $f(\mu_0(\phi)) > 0$ because from [\(51\)](#),

$$1 = \phi \int \left(\frac{r}{\mu_0(\phi) + r} \right)^2 dP(r) > \phi \int \frac{r}{\mu_0(\phi) + r} dP(r),$$

which holds because $\mu_0(\phi) < 0$ from Part (2). Analogously, when $\phi > 1$, $f(\mu_0(\phi)) < 0$. Therefore, $\lambda_{\min}(\phi) = \mu_0(\phi)f(\mu_0(\phi)) \leq 0$ with equality obtained when $\phi = 1$.

Because $f(x)$ is strictly decreasing in both ϕ and x , combining the sign properties, we further find that $\lambda_{\min}(\phi) = \mu_0(\phi)f(\mu_0(\phi))$ is strictly increasing over $\phi \in (0, 1)$ and strictly decreasing over $\phi \in (1, \infty)$. \square

F.2 Properties of Fixed-Point Equations under Negative Regularization

Our analysis involves $v(\lambda, \phi)$ as a unique solution to the fixed-point equation in

$$\frac{1}{v(\lambda, \phi)} = \lambda + \phi \int \frac{r}{1 + v(\lambda, \phi)r} dH_p(r). \quad (52)$$

Define $\mu(\lambda, \phi) = 1/v(\lambda, \phi)$. Equivalently, we have that $\mu(\lambda, \phi)$ is a unique solution to the following fixed-point equation:

$$\mu(\lambda, \phi) = \lambda + \phi \int \frac{\mu(\lambda, \phi)r}{\mu(\lambda, \phi) + r} dH_p(r). \quad (53)$$

The analytic properties of the function $\lambda \mapsto v(\lambda, \phi)$ on $(\lambda_{\min}(\phi), +\infty)$ for $\phi \in (1, \infty)$ are detailed in [Lemma F.2](#).

Lemma F.2 (Analytic properties in the regularization parameter). *Let $0 < a \leq b < \infty$ be real numbers. Let P be a probability measure supported on $[a, b]$. Let $\phi > 0$ be a real number. For $\lambda > \lambda_{\min}(\phi)$, let $v(\lambda, \phi) > 0$ denote the solution to the fixed-point equation*

$$\mu(\lambda, \phi) = \lambda + \phi \int \frac{\mu(\lambda, \phi)r}{r + \mu(\lambda, \phi)} dP(r).$$

Then the following properties hold:

- (1) (Monotonicity) For $\phi \in (0, \infty)$, the function $\lambda \mapsto \mu(\lambda, \phi)$ is strictly increasing in $\lambda \in (\lambda_{\min}(\phi), \infty)$.
- (2) (Range) For $\phi \in (0, 1]$, $\lim_{\lambda \rightarrow \lambda_{\min}(\phi)^-} \mu(\lambda, \phi) = -\infty$ and $\mu(0, \phi) = 0$. For $\phi \in (1, \infty)$, $\lim_{\lambda \rightarrow \lambda_{\min}(\phi)^-} \mu(\lambda, \phi) \in (0, \infty)$. For $\phi \in (0, \infty)$, $\lim_{\lambda \rightarrow +\infty} \mu(\lambda, \phi) = +\infty$.
- (3) (Differentiability) For $\phi \in (0, \infty)$, the function $\lambda \mapsto \mu(\lambda, \phi)$ is differentiable over Λ .
- (4) The function $\lambda \mapsto \lambda/\mu_p(\lambda, \phi)$ is strictly increasing in λ with $\lim_{\lambda \rightarrow 0} \lambda/\mu_p(\lambda, \phi) = 0$ and $\lim_{\lambda \rightarrow \infty} \lambda/\mu_p(\lambda, \phi) = 1$.

Proof. Note that

$$\lambda = \mu(\lambda, \phi) - \phi \int \frac{\mu(\lambda, \phi)r}{r + \mu(\lambda, \phi)} dP(r).$$

Define a function f by

$$f(x) = x - \phi \int \frac{xr}{x + r} dP(r).$$

Observe that $\mu(\lambda, \phi) = f^{-1}(\lambda)$. The claim of differentiability of the function $\lambda \mapsto \mu(\lambda, \phi)$ follows from the differentiability and strict monotonicity of f , similar to [Patil et al. \(2022, Lemma S.6.14\)](#).

For the last property, from the definition of the fixed-point equation, we have

$$1 = \frac{\lambda}{\mu_p(\lambda, \phi)} + \phi \int \frac{r}{r + \mu(\lambda, \phi)} dP(r).$$

Because $\phi \int \frac{r}{r + \mu(\lambda, \phi)} dP(r)$ is strictly decreasing in λ , we have that $\lambda/\mu_p(\lambda, \phi)$ is strictly increasing in λ . Because $\lim_{\lambda \rightarrow +\infty} \mu_p(\lambda, \phi) = +\infty$, we know that

$$\lim_{\lambda \rightarrow \infty} \phi \int \frac{r}{r + \mu(\lambda, \phi)} dP(r) = 0.$$

and thus,

$$\lim_{\lambda \rightarrow \infty} \frac{\lambda}{\mu_p(\lambda, \phi)} = 1.$$

On the other hand, because $\lim_{\lambda \rightarrow 0} \mu_p(\lambda, \phi) = 0$ for $\phi \in (0, 1]$ and $\lim_{\lambda \rightarrow 0} \mu_p(\lambda, \phi) < \infty$, it directly implies that

$$\lim_{\lambda \rightarrow 0} \frac{\lambda}{\mu_p(\lambda, \phi)} = 0.$$

Consequently, the conclusion follows. □

Lemma F.3 (Continuity properties in the aspect ratio for ridge regression, adapted from [Patil et al. \(2023\)](#)). *Let $a > 0$, $b < \infty$ and $\lambda \in \mathbb{R}$ be real numbers. Let P be a probability measure supported on $[a, b]$. For $\lambda \in \mathbb{R}$, let $\Phi(\lambda) = \{\phi \in (0, \infty) \mid \lambda_{\min}(\phi) < \lambda\}$. Consider the function $\mu(\lambda, \cdot) : \phi \mapsto \mu(\lambda, \phi)$, over $\phi \in \Phi(\lambda)$ such that $\lambda_{\min}(\phi) \leq \lambda$, where $\mu(\lambda, \phi) \geq -a$ is the unique solution to the following fixed-point equation:*

$$\mu(\lambda, \phi) = \lambda + \phi \int \frac{\mu(\lambda, \phi)r}{\mu(\lambda, \phi) + r} dP(r). \tag{54}$$

Then the following properties hold:

- (1) The range of the function $\mu(-\lambda; \cdot)$ is a subset of $[\lambda, \infty]$ when $\lambda \geq 0$ and $[-a, \infty]$ when $\lambda < 0$.
- (2) The function $\mu(-\lambda; \cdot)$ is continuous and strictly increasing over $\Phi(\lambda)$. Furthermore, $\lim_{\phi \rightarrow \infty} \mu(\lambda, \phi) = +\infty$, $\lim_{\phi \rightarrow 0^+} \mu(\lambda, \phi) = \lambda$ when $\lambda \geq 0$, and $\lim_{\phi \rightarrow 0^+} \mu(\lambda, \phi) = -\infty$ when $\lambda < 0$.
- (3) The function $\tilde{v}_v(-\lambda; \cdot) : \phi \mapsto \tilde{v}_v(\lambda, \phi)$, where

$$\tilde{v}_v(\lambda, \phi) = \left(\mu(\lambda, \phi)^2 - \phi \int (\mu(\lambda, \phi)r)^2 (\mu(\lambda, \phi) + r)^{-2} dP(r) \right)^{-1},$$

is positive and continuous over $\Phi(\lambda)$.

- (4) The function $\tilde{v}_b(-\lambda; \cdot) : \phi \mapsto \tilde{v}_b(\lambda, \phi)$, where

$$\tilde{v}_b(\lambda, \phi) = \tilde{v}_v(\lambda, \phi) \int \phi (\mu(\lambda, \phi)r)^2 (\mu(\lambda, \phi) + r)^{-2} dP(r),$$

is positive and continuous over $\Phi(\lambda)$.

Proof. Properties (1)-(4) follow similarly to [Patil et al. \(2022, Lemma S.6.15\)](#). \square

F.3 Contours of Fixed-Point Solutions under Negative Regularization

Lemma F.4 (Contours of fixed-point solutions). *As $n, p \rightarrow \infty$ such that $p/n \rightarrow \phi \in (0, \infty)$, let $\lambda_{\min} : \phi \mapsto \lambda_{\min}(\phi)$ as defined in (5). For any $\bar{\psi} \in [\phi, +\infty]$, there exists a unique value $\bar{\lambda} \geq \lambda_{\min}(\bar{\psi})$ (or conversely for $\bar{\lambda} \in [\lambda_{\min}(\phi), \infty]$, there exists a unique value $\bar{\psi} \in [\phi, \infty]$) such that for all (λ, ψ) on the path (as in Figure F10)*

$$\mathcal{P} = \{(1 - \theta) \cdot (\bar{\lambda}, \phi) + \theta \cdot (\lambda_{\min}(\bar{\psi}), \bar{\psi}) \mid \theta \in [0, 1]\}, \quad (55)$$

it holds that

$$\mu(\lambda, \psi) = \mu(\bar{\lambda}, \phi) = \mu(\lambda_{\min}(\bar{\psi}), \bar{\psi}),$$

where $\mu(\lambda, \psi)$ is as defined in (53).

Proof. Note that $\mu(\lambda_{\min}(\phi), \phi) = \mu_0(\phi)$ which is the solution to the fixed-point equation (51). From Lemma F.1 (2), we see that the function $\psi \mapsto \mu(\lambda_{\min}(\psi), \psi)$ is strictly increasing over $\psi \in [\phi, \infty]$ with the range

$$\lim_{\psi \rightarrow \phi} \mu(\lambda_{\min}(\psi), \psi) = \mu_0(\phi), \quad \lim_{\psi \rightarrow +\infty} \mu(\lambda_{\min}(\psi), \psi) = +\infty.$$

From Lemma F.3, the function $\lambda \mapsto \mu(\lambda, \phi)$ is strictly increasing over $\lambda \in [\lambda_{\min}, \infty]$ with range

$$\lim_{\lambda \rightarrow \lambda_{\min}} \mu(\lambda, \phi) = \mu_0(\phi), \quad \lim_{\lambda \rightarrow +\infty} \mu(\lambda, \phi) = +\infty.$$

For $\bar{\psi} \in [\phi, \infty]$, by the intermediate value theorem, there exists a unique $\bar{\lambda} \in [\lambda_{\min}(\phi), \infty]$ such that $v(\bar{\lambda}, \phi) = v(-\lambda_{\min}(\bar{\psi}); \bar{\psi})$. Conversely, for $\bar{\lambda} \in [\lambda_{\min}(\phi), \infty]$, there also exists a unique $\bar{\psi} \in [\phi, \infty]$ such that $v(\bar{\lambda}, \phi) = v(-\lambda_{\min}(\bar{\psi}); \bar{\psi})$.




Figure F10: Heatmap of $\mu_p(\lambda, \psi)$ for isotropic covariance matrix $\Sigma = I$ in the symmetric log-log scale (where the logarithmic scale is applied symmetrically to both positive and negative values on the x -axis and y -axis). The 5 black dashed lines indicate 5 different equivalence paths. The boundary of negative ridge penalties is given by $-(1 - \sqrt{\psi})^2$.

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

$$\mu(\bar{\lambda}, \phi) = \bar{\lambda} + \phi \int \frac{\mu(\bar{\lambda}, \phi)r}{\mu(\bar{\lambda}, \phi) + r} dH_p(r) = \lambda_{\min}(\phi) + \bar{\psi} \int \frac{\mu(-\lambda_{\min}; \bar{\psi})r}{\mu(-\lambda_{\min}; \bar{\psi}) + r} dH_p(r) = \mu(-\lambda_{\min}; \bar{\psi}).$$

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

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

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

$$\mu(\lambda, \psi) = \lambda + \psi \int \frac{\mu(\lambda, \psi)r}{\mu(\lambda, \psi) + r} dH_p(r),$$

it then follows that $\mu(\lambda, \psi) = \mu(\bar{\lambda}, \phi) = \mu(\lambda_{\min}(\bar{\psi}), \bar{\psi})$. This completes the proof. \square

G Experimental Details and Additional Numerical Illustrations

The source code for reproducing the results of this paper can be found at the following location: <https://github.com/jaydu1/ood-ridge>.

G.1 Additional Illustrations for Section 3

G.1.1 Numerical verification of (8) under generic alignment scenarios

Under Assumption A, suppose the data model has a covariance matrix Σ to be $(\Sigma_{\text{ar1}})_{ij} := \rho_{\text{ar1}}^{|i-j|}$ with parameter $\rho_{\text{ar1}} \in (0, 1)$, and a coefficient $\beta := \frac{1}{2}(w_{(1)} + w_{(p)})$, where $w_{(j)}$ is the j th eigenvector of Σ_{ar1} .

$$\bar{\text{tr}}[\Sigma B] = \frac{1}{4p}(w_{(1)} + w_{(p)})^\top \Sigma (w_{(1)} + w_{(p)}) = \frac{1}{4p}w_{(1)}^\top \Sigma w_{(1)} + \frac{1}{4p}w_{(p)}^\top \Sigma w_{(p)} = \frac{1}{4}(r_{(1)} + r_{(p)}).$$

On the other hand, we also have

$$\bar{\text{tr}}[\Sigma] \bar{\text{tr}}[B] = \sum_{j=1}^p \frac{r_{(j)}}{p} \cdot \frac{\|\beta\|_2^2}{p} = \frac{1}{2p^2} \sum_{j=1}^p r_{(j)}$$

One can numerically verify that when $p = 500$ and $\rho_{\text{ar1}} = 0.5$,

$$r_{(1)} + r_{(p)} \approx 3.33 > 2 = \frac{2}{p} \sum_{j=1}^p r_{(j)}$$

which contradicts the implication (32) of the strict alignment condition.




Figure G11: Numerical evaluation of $b_2/b_3 - s_2/s_3$ under the same data model in Figure 1, where $s_k = \bar{\text{tr}}[\Sigma\{\mu^k(\Sigma + \mu I)^{-k}\}]$ and $b_k = \bar{\text{tr}}[\Sigma_\beta\{\mu^k(\Sigma + \mu I)^{-k}\}]$ for $k = 2, 3$.

On the other hand, from Figure G11, we see that the general alignment condition

$$\frac{\bar{\text{tr}}[\Sigma_\beta\{\mu^2(\Sigma + \mu I)^{-2}\}]}{\bar{\text{tr}}[\Sigma_\beta\{\mu^3(\Sigma + \mu I)^{-3}\}]} = \frac{\bar{\text{tr}}[\Sigma\{\mu^2(\Sigma + \mu I)^{-2}\}]}{\bar{\text{tr}}[\Sigma\{\mu^3(\Sigma + \mu I)^{-3}\}]}$$

holds for various data aspect ratios in the noiseless setting.

G.1.2 Real Data Illustration

Following the approach of Kobak et al. (2020), we consider a similar setup using random Fourier features on MNIST.

Feature generation. The pixel values are normalized to the range $[-1, 1]$. The features are then mapped from the original dimension of 28×28 to 1000 random Fourier features. This is achieved by multiplying the features with a 784×500 random matrix W , where the elements are independently drawn from a normal distribution with mean 0 and standard deviation 0.2. The real and imaginary parts of $\exp(-iXW)$ are taken as separate features, where $X \in \mathbb{R}^{n \times 784}$.

Training details. For training, we randomly select $n = 64$ images, and for testing, we hold out 10,000 images. The response variable y represents the digit value ranging from 0 to 9. Our model includes an intercept term, which is not subject to penalization. To estimate the expected out-of-distribution risk, we average the risks across 100 random samples from the training set.

Distribution shift. To generate distribution shift, we gradually exclude samples with given labels in the test set:

- Type 1: \emptyset
- Type 2: $\{4\}$
- Type 3: $\{3, 4\}$
- Type 4: $\{2, 3, 4\}$
- Type 5: $\{1, 2, 3, 4\}$

In other words, Type 1 represents no covariate shift, while Type 5 represents the case with potentially the most severe covariate shift. The results are summarized in Figure G12. We observe a clear pattern where the optimal ridge penalty shifts toward negative values, suggesting that Theorem 4 may occur in real-world datasets. However, for Type 5, the optimal ridge penalty becomes positive again. This could be due to the removal of Class 1, which reduces the degree of alignment. Also, observe that the minimum risks increase from Type 1 to 5.




Figure G12: **Effect of distribution shift on the optimal ridge penalty on MNIST.** (a) The left panel illustrates the risk profile (against the regularization penalty) of ridge regression on the MNIST dataset when subjected to different types of distribution shifts. Different colors represent different types of shift from less severe (Type 1) to more severe (Type 5). The y-axis represents the out-of-distribution prediction risk for the task of accurately predicting the digit value for unseen images. The figure shows a clear pattern where the optimal ridge penalty shifts towards negative values, in the spirit of Theorem 5. The only exception seems to be Type 5 for which the optimal ridge penalty becomes positive again. It is likely due to the removal of Class 1, which reduces the degree of alignment. (b) The right panel shows the relative OOD prediction risk computed only on the excluded class, compared to the IND prediction risk of the ridge predictor fitted only on the training data of the same class. The reason we compare the relative prediction risk is to compensate for the differences in the conditional variances in the different classes. Observe that Class 1 (Type 5) has the lowest relative prediction risk, which partially explains the increase in optimal regularization of Type 5 in the left panel.

G.2 Additional Illustrations for Section 5 (Ridge versus Lasso Monotonicities)

G.2.1 Subsampled ridgeless versus full ridge




Figure G13: **Ridge optimal risk monotonicity and suboptimal risk non-monotonicity.** Optimal ridge has a monotonic risk in the data aspect ratio, but the risk for fixed ridge penalty λ may not be monotonic in the data aspect ratio. The data model has $\Sigma = \Sigma_{\text{ar1}}$ with $\rho_{\text{ar1}} = 0.25$ (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 $\rho_{\text{ar1}} \in (0, 1)$), β being the leading eigenvector of Σ and $\sigma = 0.5$. The leftmost panel shows the limiting risk of the full-ensemble ridgeless regression at various data and subsample aspect ratios (ϕ, ψ) . The middle panel shows the limiting risk of the ridge predictor (on the full data) at various data aspect ratios and regularization penalties (ϕ, λ) . We highlight the optimal risks at a given data aspect ratio for the leftmost and middle panels using slender red lines. Observe that the optimal risk in both cases is increasing as a function of ϕ .

G.2.2 Subsampled lassoless versus full lasso




Figure G14: **Lasso optimal risk monotonicity and suboptimal risk non-monotonicity.** Similar to ridge regression, optimal lasso has a monotonic risk in the over-parameterization ratio, but the risk for fixed lasso penalty λ may not be monotonic in the data aspect ratio. The data model has $\Sigma = I$ and $\beta_j \stackrel{\text{i.i.d.}}{\sim} \epsilon P_{1/\sqrt{\phi\epsilon}} + (1 - \epsilon)P_0$ where the sparsity level is $\epsilon = 0.01$ and $\sigma^2 = 1$, such that $\text{SNR} = 1$. As for ridge regression in Figure G13, optimal risks for each data aspect ratio are highlighted using slender red lines in the left and middle panels. Similarly to the ridge curves in Figure G13, observe that the optimal risk in both cases is increasing as a function of ϕ .