

Optimal Unconstrained Self-Distillation in Ridge Regression: Strict Improvements, Precise Asymptotics, and One-Shot Tuning

Hien Dang[†]

hiendang@utexas.edu

Pratik Patil[†]

pratikpatil@utexas.edu

Alessandro Rinaldo[†]

alessandrorinaldo@utexas.edu

Abstract

Self-distillation (SD) is the process of retraining a student on a mixture of ground-truth labels and the teacher’s own predictions using the same architecture and training data. Although SD has been empirically shown to often improve generalization, its formal guarantees remain limited. We study SD for ridge regression in unconstrained setting in which the mixing weight ξ may be outside the unit interval. Conditioned on the training data and without any distributional assumptions, we prove that for any squared prediction risk (including out-of-distribution), the optimally mixed student strictly improves upon the ridge teacher for every regularization level $\lambda > 0$ at which the teacher ridge risk $R(\lambda)$ is nonstationary (i.e., $R'(\lambda) \neq 0$). We obtain a closed-form expression for the optimal mixing weight $\xi^*(\lambda)$ for any value of λ and show that it obeys the sign rule: $\text{sign}(\xi^*(\lambda)) = -\text{sign}(R'(\lambda))$. In particular, $\xi^*(\lambda)$ can be negative, which is the case in over-regularized regimes. To quantify the risk improvement due to SD, we derive exact deterministic equivalents for the optimal SD risk in the proportional asymptotics regime (where the sample and feature sizes n and p both diverge but their the aspect ratio p/n converges) under general anisotropic covariance and deterministic signals. Our asymptotic analysis extends standard second-order ridge deterministic equivalents to their fourth-order analogs using block linearization, which may be of independent interest. From a practical standpoint, we propose a consistent one-shot tuning method to estimate ξ^* without grid search, sample splitting, or refitting. Experiments on real-world datasets and pretrained neural network features support our theory and the one-shot tuning method.

1 Introduction

Knowledge distillation (KD), introduced by [Buciluă et al. \(2006\)](#); [Ba and Caruana \(2014\)](#); [Hinton et al. \(2015\)](#), is conventionally used for model compression, transferring knowledge from a large teacher to a smaller student. Recently, this paradigm has been adapted to the setting where teacher and student share the same architecture and training data, a process known as *self-distillation* (SD) ([Furlanello et al., 2018](#); [Zhang et al., 2021](#)). While it may seem counterintuitive that a model would improve by learning from its own predictions, extensive empirical evidence shows that SD can in fact boost generalization ([Chen et al., 2017, 2022](#); [Li et al., 2017](#); [Ahn et al., 2019](#); [Li et al., 2021](#); [Gou et al., 2021](#)). Despite these successes, it remains unclear whether and when such improvements can be guaranteed.

Formally, let f be a teacher trained on $\{(x_i, y_i)\}_{i=1}^n$ using a loss function ℓ . Self-distillation trains a student f_{sd} on the *same* data by minimizing a mixed objective that is an affine interpolation of the losses incurred with respect to the ground-truth labels y_i ’s and the teacher’s predictions $f(x_i)$ ’s. In

[†]Department of Statistics and Data Sciences, University of Texas, Austin, TX 78712, USA.

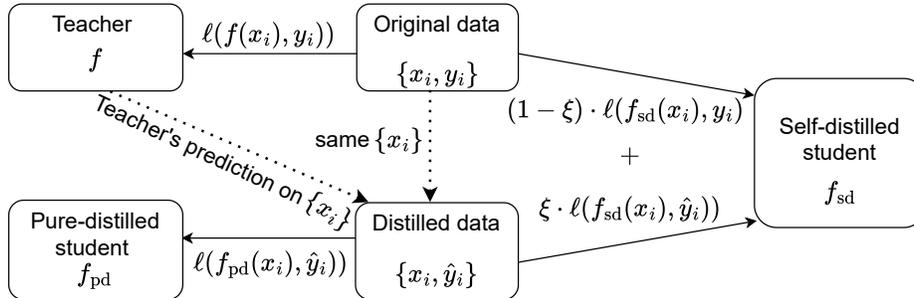


Figure 1: Visual illustration of the self-distillation process.

detail, the SD procedure seeks to find the student model f_{sd} minimizing

$$\frac{1}{n} \sum_{i=1}^n \left[(1 - \xi) \cdot \ell(y_i, f_{sd}(x_i)) + \xi \cdot \ell(f(x_i), f_{sd}(x_i)) \right], \quad (1)$$

where ξ is the mixing parameter (Lopez-Paz et al., 2015); see Figure 1. When $\xi = 1$, the student learns solely from the teacher’s predictions; we call this *pure-distillation* (PD) and denote the resulting predictor by f_{pd} .

The mixing parameter ξ balances the influence of ground-truth labels against teacher predictions. Standard distillation methods restrict ξ to lie in $[0, 1]$, interpreting the target loss as a convex combination. Recent work by Das and Sanghavi (2023) shows that this constraint can be suboptimal under high label noise, where the optimal mixing weight ξ^* may in fact be found to be greater than 1. Motivated by this, we adopt a fully unconstrained perspective and allow $\xi \in \mathbb{R}$, including negative values. Note that setting $\xi = 0$ recovers the teacher predictor, hence optimizing over ξ cannot perform worse than the teacher. With this in mind, we pose the key questions about SD:

- (Q1) When does the optimally mixed student f_{sd} trained using an optimal $\xi^* \in \mathbb{R}$ strictly outperform the teacher f , and how large can the gain be?
- (Q2) Can optimal SD from a suboptimal teacher achieve performance comparable to an optimally tuned teacher?
- (Q3) How can we efficiently tune the optimal $\xi^* \in \mathbb{R}$ without computationally expensive grid search?

We provide complete answers to all these questions for ridge regression, a model in which SD admits an explicit affine path (in the response) and the risk of both the teacher and the students can be characterized sharply, capturing the interplay between regularization and distillation.

1.1 Summary of Paper Contributions and Outline

Below we describe in detail the main contributions of the paper; see Figure 2 for a visual summary.

Structural nonasymptotic guarantees (Section 2). Addressing (Q1), we derive deterministic identities for self-distilled ridge that hold conditionally on the observed training data, without any distributional assumptions, and for *any* squared prediction risk (including the out-of-distribution risk). In particular, we show that for every $\lambda > 0$ such that the teacher ridge-path risk $\lambda \mapsto R(\lambda)$

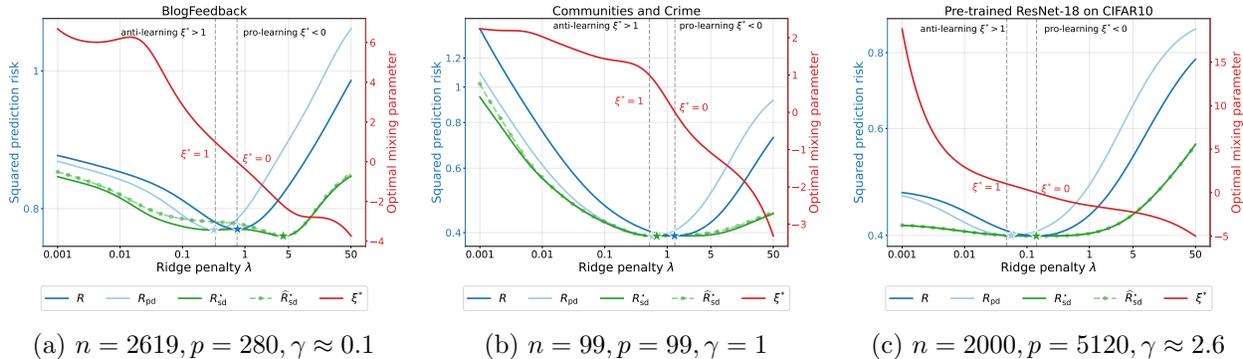


Figure 2: Strict improvement of SD risk with unconstrained mixing. Test squared prediction risk of ridge regression (R , in blue), pure-distilled ridge (R_{pd} , in light blue) and optimal self-distilled ridge (R_{sd}^* , in green) as functions of the ridge penalty λ . Results are shown on raw features from real-world datasets: BlogFeedback and Communities and Crime datasets, and on pretrained ResNet-18 features. The optimal mixing parameter $\xi^*(\lambda)$ is in red and the one-shot risk estimate $\hat{R}_{\text{sd}}^*(\lambda)$ computed from the training data is shown in green dashed line. Note that $\xi^*(\lambda)$ lies in $[0, 1]$ only for a narrow range of λ and can be strongly negative for large λ . We also observe that: (i) $R_{\text{sd}}^*(\lambda)$ is strictly smaller than $R(\lambda)$ at every λ that is not the stationary point of $R(\lambda)$, (ii) the sign of $\xi^*(\lambda)$ is opposite to the sign of $R'(\lambda)$, and (iii) the sign change of ξ^* happens at the stationary point of $R(\lambda)$. (Experiments with ξ restricted to $[0, 1]$ appear in Figure 16.)

is nonstationary (i.e., $R'(\lambda) \neq 0$), optimal mixing yields a *strict* improvement over the teacher (Theorem 2.2). Addressing (Q2), we provide a curvature-based sufficient condition under which the global minimum over λ of the SD risk $R_{\text{sd}}^*(\lambda)$, obtained using the optimal mixing $\xi^*(\lambda)$, is strictly smaller than the smallest ridge risk of the teacher (Proposition 2.3).

Precise proportional asymptotics (Section 3). Returning to (Q1) in the proportional regime – in which the sample and feature sizes $n, p \rightarrow \infty$ but their aspect ratio $p/n \rightarrow \gamma \in (0, \infty)$ – we derive exact deterministic equivalents for the optimal SD risk and mixing weight under general anisotropic covariance and deterministic signals (Theorem 3.1). These formulas quantify the SD gains in terms of γ , the signal-to-noise ratio, and the signal-covariance alignment. In particular, they characterize precisely when the optimal mixing weight ξ^* becomes *negative* (Corollary 3.2). Regarding (Q2), under an isotropic random signal we show that even for extremely under- or over-regularized teachers, optimal SD can closely approach the risk of the optimally tuned ridge predictor (Proposition 3.3).

One-shot tuning (Section 4). Addressing (Q3), we propose a consistent one-shot estimator of ξ^* based on generalized cross-validation (Theorem 4.1). The method avoids sample splitting and grid search over ξ (and hence does not require any refitting across candidate mixing weights), while remaining consistent in proportional asymptotics. We validate our theory and tuning methodology through extensive experiments on synthetic data, UCI regression benchmarks, and pretrained ResNet feature representations.

Extensions and variants (Section 5). We analyze several natural variants of the SD ridge framework, including multi-round distillation with a risk monotonicity property (Section D.1), self-distillation using teacher predictions on fresh features, which we show to be dominated by the

same- X setup in an isotropic setting (Section D.2), and extensions to generalized ridge and kernel ridge regression (Section D.3).

1.2 Related Works and Comparisons

Fixed-design analyses. Early theoretical work on distillation focused on when and how well a student can mimic a teacher (Phuong and Lampert, 2019; Ji and Zhu, 2020); for regression problems, this has typically been addressed using fixed-design. In kernel ridge regression, Mobahi et al. (2020) show that repeated SD shrinks the effective function class, resembling an increase in regularization; their emphasis is on the implicit regularization induced by SD rather than on guaranteeing performance improvements. In ridge regression, Das and Sanghavi (2023) analyze the bias–variance trade-off induced by the mixing weight ξ and give conditions under which the *globally* optimal SD risk (obtained by optimizing over λ) can go below the optimally tuned teacher ridge risk. Pareek et al. (2024) extend this line of work to multiple SD rounds, quantifying gains under an alignment condition when λ and ξ are tuned in each round. A key gap in these analyses is the question of pointwise improvement: whether SD can *strictly* improve a mis-regularized teacher at a fixed, suboptimal λ , which we show in Section 2. Such pointwise gains are important because finding global optimal λ is often challenging; for example, Stephenson et al. (2021) show that the leave-one-out cross-validation loss is generally neither convex nor even quasi-convex in λ .

A further distinction is the data model. Both the linear regression analysis in Das and Sanghavi (2023); Pareek et al. (2024) operate under fixed design and a well-specified linear response model. While analytically convenient, fixed-design analyses do not capture the contribution of test-feature randomness to prediction error, which is central in high-dimensional generalization (Hastie et al., 2022). In contrast, our structural results hold conditionally on the observed training data under random design and for *any* squared prediction risk, including out-of-distribution risk. We also work in proportional random-design asymptotics ($n, p \rightarrow \infty$ with $p/n \rightarrow \gamma$) in Section 3 and provide exact characterizations of generalization effects driven by the feature and signal structures.

Random-design analyses. Closer to our setting, Emrullah Ildiz et al. (2025) analyze KD under random design with anisotropic covariance, focusing on scaling laws and “weak-to-strong” generalization, a phenomenon where a strong student model is supervised by a weaker teacher (Burns et al., 2023). They provide nonasymptotic characterizations of pure-distilled ridge ($\xi = 1$ in our notation) and its scaling behavior. They primarily focus on the *ridgeless* case (minimum ℓ_2 -norm interpolator where $\lambda \rightarrow 0^+$). While interesting, such analysis obscures the interaction between explicit ridge penalty ($\lambda > 0$) and distillation weight ξ , which is central to our results. In particular, we identify over-regularized regimes where optimal SD requires a *negative* optimal mixing ξ^* to account for excessive shrinkage and thus strictly improve the teacher. In contrast, in pure-distillation ridge with $\lambda > 0$, the student can only improve upon the teacher for a range of small λ , and can fail to improve under over-regularization (or even being worse than the teacher) (Moniri and Hassani, 2025); see Figure 2.

A large body of recent work has focussed ridge regression risk and its variants under proportional asymptotics using tools from random matrix theory and statistical physics (e.g., Dobriban and Wang, 2018; Hastie et al., 2022; Patil et al., 2024 and references therein), shedding light on phenomena such as benign overfitting (Bartlett et al., 2020) and double descent (Belkin et al., 2019). We extend this literature by deriving exact deterministic equivalents for optimal SD risk under anisotropic feature covariance and deterministic signals. Because the teacher and student are trained on the *same* data, SD risks involve dependent resolvent-type quantities; unlike settings with independent refits

Table 1: **Settings overview in the self-distillation ridge/ridgeless literature.** Unless stated otherwise, risks in the proportional asymptotics setting are in-distribution squared prediction risks.

Paper	Focus		Nonasymptotic setting			Proportional asymptotics			Weights		Multi-rounds
	Type	Family	Design	Response	Risk	Features	Signal	Model	Range	Tuning	
DS23	SD	Ridge	Fixed	Linear	Estimation				\mathbb{R}		
PDO24	SD	Ridge	Fixed	Linear	Fixed-X prediction				\mathbb{R}	Split CV (three refits)	Yes
IGT+25	PD	Ridgeless	Random	Linear	In-distribution prediction				{1}		
MH25	PD	Ridge (generalized)				Isotropic	Isotropic	Linear	{1}		
GBS25	SD	Ridge (infinite rounds)				Anisotropic	Deterministic	Linear	[0, 1]		Infinite rounds only
Ours	SD	Ridge	Random	Any	Any squared prediction	Anisotropic	Deterministic	Any	\mathbb{R}	GCV (one-shot)	Yes

where one can directly combine known deterministic equivalents, our analysis requires higher-order deterministic equivalents obtained via block linearization (see Sections B.3 and B.6).

Unconstrained mixing and “anti-learning”. Standard KD typically restricts $\xi \in [0, 1]$. In the presence of label noise, Das and Sanghavi (2023) show that optimal mixing can satisfy $\xi^* > 1$, corresponding to a negative weight on the noisy ground truth (“anti-learning”). In binary classification, Javanmard et al. (2025) derive a nonlinear Bayes-optimal aggregation effect (using approximate message passing tools) that also effectively subtracts noisy labels during retraining. In our ridge setting, closed-form risks allow a clean three-regime picture: interpolation ($0 \leq \xi \leq 1$), extrapolation/anti-learning ($\xi > 1$), and over-regularization correction ($\xi < 0$), which we term “pro-learning”.

Synthetic-label retraining and model collapse. Beyond one-round SD, recent work has raised concerns about recursively training on model-generated labels leading to performance degradation, referred to as *model collapse* (Shumailov et al., 2024; Alemohammad et al., 2023; Dohmatob et al., 2024; Gerstgrasser et al., 2024). In contrast to the notion of “strong model collapse” in Dohmatob et al. (2025) (in which a model trained in the presence of synthetic data has worse asymptotic risk than a model trained solely on ground-truth data), optimal SD in our ridge setting exhibits no such degradation. Provided the nondegeneracy condition in Theorem 2.2 holds, the optimally mixed student strictly improves upon the teacher whenever $R'(\lambda) \neq 0$, with the two risks coinciding only at stationary points.

Closer to our setting, He et al. (2025); Garg et al. (2025) analyze infinite-round schemes that mix ground-truth and synthetic labels with a fixed weight w (analogous to our ξ) to prevent degradation under repeated synthetic training; the optimal mixing weight w^* in their settings lies in $[0, 1]$. In contrast, we characterize the risk-minimizing mixing for one-round SD and show that ξ^* can lie outside $[0, 1]$, including $\xi^* < 0$ in over-regularized regimes. To highlight the distinction, we run a synthetic experiment (Section E.2) comparing optimal one-round SD risk to 20-round SD with optimal fixed constrained weight in $[0, 1]$ for every round. As shown in Figure 15, unconstrained mixing can be crucial for improving the teacher’s performance, particularly in over-regularized regimes where negative weights becomes necessary. Finally, we also study a recursive multi-round variant with per-round unconstrained optimal mixing that yields a monotone (weakly) decreasing risk sequence (Section 5).

Tuning and risk estimation. A notable gap in existing theoretical results about SD is the choice of the optimal mixing hyperparameter ξ in practice. Most analyses assume oracle access to population risks or detailed spectral knowledge. In practice, cross-validation over ξ is computationally costly due to grid search and repeated refitting, and it can be statistically inefficient in high dimensions due to sample splitting (see, e.g., Rad and Maleki (2020)). Building on consistent risk estimation for high-dimensional ridge/ridgeless regression (see, e.g., Patil et al., 2021, 2022b; Wei et al., 2022; Han and Xu, 2023; Bellec et al., 2025; Koriyama et al., 2024 and references therein), we propose a one-shot generalized cross-validation estimator for SD ridge and prove its consistency in proportional asymptotics (Theorem 4.1). This enables data-efficient selection of unconstrained ξ without grid search over candidate mixing weights or hold-out sets.

Overall, relative to prior SD ridge analyses, we provide the following novel results: (i) pointwise strict improvement guarantees at any nonstationary $\lambda > 0$ together with a sign characterization of ξ^* (Theorem 2.2); (ii) exact proportional-asymptotic risk characterizations under anisotropic feature covariance and deterministic signals (Theorem 3.1); and (iii) a consistent one-shot GCV-based tuning method (Theorem 4.1). At the same time, our settings are more general than those considered in the literature so far: (a) a nonasymptotic setting with no distributional assumptions and any squared prediction risk; and (b) proportional asymptotics with general feature covariance and deterministic signals, without imposing a well-specified response model. A compact comparison of settings and assumptions across several related works is provided in Table 1.

2 Structural Nonasymptotic Results

This section described deterministic “structural” identities for self-distillation in ridge regression under any squared prediction risk (including the out-of-distribution risk). All the statements hold with virtually no distributional assumptions (other than the existence of second moments) and conditionally on the training data $\mathcal{D} = (X, y)$, where $X \in \mathbb{R}^{n \times p}$ is the design matrix containing p -dimensional observations $x_i \in \mathbb{R}^p$ as rows and $y \in \mathbb{R}^n$ is the response vector containing scalar ground-truth labels y_i for $i = 1, \dots, n$.

2.1 Self-Distillation with Ridge Regression

Given a regularization parameter $\lambda > 0$, define the *teacher* ridge predictor $x \in \mathbb{R}^p \mapsto f_\lambda(x) \in \mathbb{R}$ trained on \mathcal{D} as

$$x^\top \operatorname{argmin}_{\beta \in \mathbb{R}^p} \{ \|y - X\beta\|_2^2/n + \lambda \|\beta\|_2^2 \}. \quad (2)$$

Let $\hat{y}_\lambda = f_\lambda(X) = (f_\lambda(x_1), \dots, f_\lambda(x_n))^\top \in \mathbb{R}^n$ denote the teacher’s training predictions. The *pure-distilled* predictor $f_{\text{pd},\lambda}$ is ridge regression trained on the pseudo-labeled or “distilled” dataset (X, \hat{y}_λ) with the same penalty λ . For a mixing parameter $\xi \in \mathbb{R}$, the *self-distilled* predictor $x \in \mathbb{R}^p \mapsto f_{\text{sd},\lambda}(x) \in \mathbb{R}$ is the ridge fit obtained using the mixed-label loss:

$$x^\top \operatorname{argmin}_{\beta \in \mathbb{R}^p} \{ (1 - \xi) \|y - X\beta\|_2^2/n + \xi \|\hat{y}_\lambda - X\beta\|_2^2/n + \lambda \|\beta\|_2^2 \}. \quad (3)$$

A key simplification in ridge regression is that the ridge map is linear in the response vector. As a result, the SD predictor lies on a linear path between the teacher and the PD fit (see Section A):

$$f_{\text{sd},\lambda,\xi}(x) = (1 - \xi) f_\lambda(x) + \xi f_{\text{pd},\lambda}(x). \quad (4)$$

For a test point (x_0, y_0) (possibly out-of-distribution) of finite (data) conditional L_2 norm¹, we measure the performance of any (possibly data-dependent) predictor f by the conditional squared prediction risk

$$R(f) := \mathbb{E}[(y_0 - f(x_0))^2 \mid \mathcal{D}]. \quad (5)$$

For convenience, write $R(\lambda) := R(f_\lambda)$, $R_{\text{pd}}(\lambda) := R(f_{\text{pd},\lambda})$, and $R_{\text{sd}}(\lambda, \xi) := R(f_{\text{sd},\lambda,\xi})$. We note that these different types of risk are random, as they depend on training data.

2.2 Optimal SD Risk Decomposition

We first provide a useful decomposition of the optimal SD risk in terms of three key quantities, defined conditionally on \mathcal{D} ². First, we let

$$C(\lambda) := \mathbb{E}[(y_0 - f_\lambda(x_0))(y_0 - f_{\text{pd},\lambda}(x_0)) \mid \mathcal{D}], \quad (6)$$

be the conditional correlation between the residual errors of the teacher and PD predictors and

$$D(\lambda) = \mathbb{E}[(f_\lambda(x_0) - f_{\text{pd},\lambda}(x_0))^2 \mid \mathcal{D}] \geq 0. \quad (7)$$

be the conditional expected squared difference between the predictions of the teacher and PD student predictors (which is also the expected squared difference between their residuals).

The difference $D(\lambda)$ also admits an equivalent form $D(\lambda) = R(\lambda) + R_{\text{pd}}(\lambda) - 2C(\lambda)$ (see Section A). Let $\xi^*(\lambda) \in \operatorname{argmin}_{\xi \in \mathbb{R}} R_{\text{sd}}(\lambda, \xi)$, and define $R_{\text{sd}}^*(\lambda) := R_{\text{sd}}(\lambda, \xi^*)$, referred to as the optimal SD risk.

Proposition 2.1 (Optimal SD risk decomposition). Fix $\lambda > 0$ and assume $D(\lambda) > 0$. Then,

$$\xi^*(\lambda) = \frac{R(\lambda) - C(\lambda)}{D(\lambda)}, \quad R_{\text{sd}}^*(\lambda) = R(\lambda) - \frac{(R(\lambda) - C(\lambda))^2}{D(\lambda)}. \quad (8)$$

In particular, $R_{\text{sd}}^*(\lambda) \leq R(\lambda)$ for all λ , and $\xi^*(\lambda)$ may be negative.

The magnitude of the risk improvement $R(\lambda) - R_{\text{sd}}^*(\lambda)$ depends on (i) the numerator $R(\lambda) - C(\lambda)$, which also determines the sign of $\xi^*(\lambda)$, and (ii) the denominator $D(\lambda)$. From (7), note that $D(\lambda) = 0$ if and only if f_λ and $f_{\text{pd},\lambda}$ coincide almost surely on the test distribution, in which case $R_{\text{sd}}(\lambda, \xi) \equiv R(\lambda)$ for all ξ , so that each ξ is trivially optimal. Thus the nondegeneracy assumption $D(\lambda) > 0$ in Proposition 2.1 is innocuous.

2.3 Strict Pointwise Improvement and Sign of Optimal Mixing Weight

Building on Proposition 2.1, we next characterize when optimal SD is *strictly* better than the teacher, and determine the sign of the optimal mixing weight in terms of the derivative of the teacher risk. (It is worth mentioning that under finite conditional second moment conditions, the teacher risk $R(\lambda)$ is a smooth function of λ and hence the derivative is well-defined; see Lemma A.4.)

¹We do not need (x_0, y_0) to be independent of \mathcal{D} for results in this section. Thus they also apply when (x_0, y_0) is drawn from the empirical measure on \mathcal{D} , in which case $R(f)$ is simply the training error of f .

²Throughout this section, any identity, inequality, or stated property involving such quantities is implicitly assumed to hold with probability one with respect to the distribution of the training data.

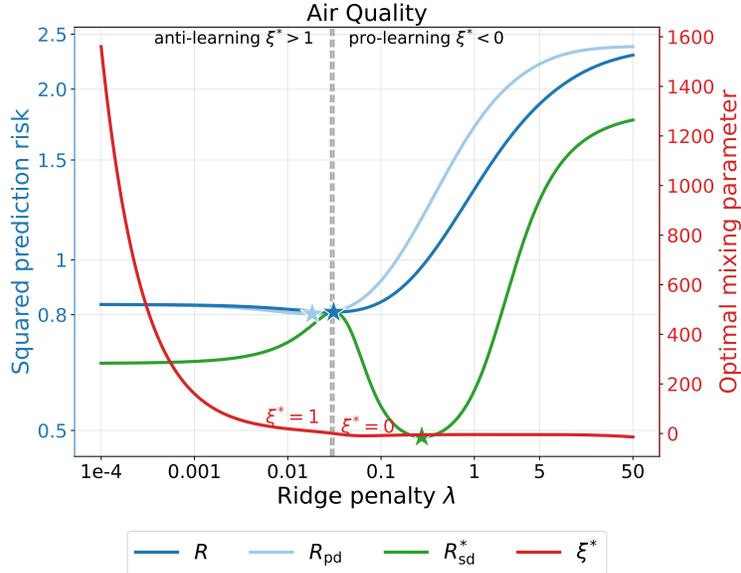


Figure 3: **Out-of-distribution SD risk improvement** Test prediction risk of ridge and optimal SD ridge on Air Quality dataset (see Section F for more details). SD yields strict improvements across λ and achieves a substantially smaller global minimum.

Theorem 2.2 (Strict improvement and sign rule). Fix $\lambda > 0$ and assume $D(\lambda) > 0$. Then

$$\xi^*(\lambda) = -\frac{\lambda}{2} \frac{R'(\lambda)}{D(\lambda)}, \quad R_{sd}^*(\lambda) = R(\lambda) - \frac{\lambda^2}{4} \frac{(R'(\lambda))^2}{D(\lambda)}. \quad (9)$$

In particular, if $R'(\lambda) \neq 0$, then $R_{sd}^*(\lambda) < R(\lambda)$ and $\text{sign}(\xi^*(\lambda)) = -\text{sign}(R'(\lambda))$.

Plainly, optimal self-distillation strictly improves the teacher at every *nonstationary* point of the teacher ridge risk (i.e., whenever $R'(\lambda) \neq 0$). Moreover, $\xi^*(\lambda)$ is positive in under-regularized regimes (where $R'(\lambda) < 0$) and negative in over-regularized regimes (where $R'(\lambda) > 0$), a behavior we refer to as “pro-learning” to contrast it with “anti-learning” (Das and Sanghavi, 2023). While prior work has noted improvements of optimal SD, to our knowledge, Theorem 2.2 is the first result establishing pointwise strict improvements for optimal self-distilled ridge, along with the sign rule for optimal mixing, at any nonstationary λ , for any squared prediction risk and without distributional assumptions. The strict improvement and sign rule are consistently observed in Figures 2 and 3.

2.4 Can Self-Distillation Beat Optimally Tuned Ridge?

Theorem 2.2 guarantees improvement at any suboptimal λ , but it also implies that at a ridge-optimal penalty λ^* , we must have $R_{sd}^*(\lambda^*) = R(\lambda^*)$. It is therefore natural to ask whether the *global* minimum of the SD risk curve can be strictly smaller than the teacher’s global minimum. We give a sufficient condition based on the local curvature of $R(\lambda)$.

Proposition 2.3 (Curvature test at the ridge-optimal λ). Let $\lambda^* \in \arg \min_{\lambda > 0} R(\lambda)$. If the follow-

Table 2: Curvature test at the ridge-optimal λ^* .

Dataset	Illustration	(10) holds?	Global gain?
BlogFeedback	Figure 2a	✓	✓
Communities and Crime	Figure 2b	✗	✗
CIFAR10	Figure 2c	✗	✗
Air Quality	Figure 3	✓	✓

ing curvature test at λ^* holds, i.e. if

$$D(\lambda^*) < \frac{\lambda^{*2}}{2} R''(\lambda^*), \quad (10)$$

then R_{sd}^* has negative curvature at λ^* , and, consequently,

$$\min_{\lambda>0} R_{\text{sd}}^*(\lambda) < \min_{\lambda>0} R(\lambda). \quad (11)$$

Proposition 2.3 formalizes a phenomenon visible in our experiments: the curves $R(\lambda)$ and $R_{\text{sd}}^*(\lambda)$ necessarily *touch* at ridge-optimal value λ^* of the regularization parameter, but R_{sd}^* may bend downward elsewhere and achieve a strictly smaller global minimum. The curvature test (10) can be viewed as a generalization of related curvature-based sufficient conditions by Das and Sanghavi (2023, Theorem 3.8), as it applies to arbitrary squared prediction risk and does not require that the response follows a well-specified linear model.

As mentioned above, the results in this section are concerned with the to out-of-distribution prediction risk, where (x_0, y_0) may be drawn from a distribution different from that of the training samples. In Figure 3, we showcase a real-data example on the Air Quality dataset with a distribution shift between the training and test sets due to time-dependent effects. We also numerically verify the curvature test in Table 2 for all of our real-data experiments and observe exact agreement in all cases.

While the results above guarantee a strict improvement at every nonstationary λ and require no distributional assumptions, they do not quantify the *magnitude* of the gain, i.e., how much smaller $R_{\text{sd}}^*(\lambda)$ can be than $R(\lambda)$. In the next section, we consider standard distributional assumptions in the proportional-asymptotics literature and derive deterministic equivalents for these risks and their improvements as explicit functions of the problem parameters.

3 Proportional Asymptotic Results

We now refine the structural identities from Section 2 by deriving deterministic limits for the optimal SD risk $R_{\text{sd}}^*(\lambda)$ and the optimal mixing weight $\xi^*(\lambda)$ in the proportional asymptotics regime in which $n, p \rightarrow \infty$ with $p/n \rightarrow \gamma \in (0, \infty)$. Our goals are to (i) obtain computable asymptotic expressions for $R_{\text{sd}}^*(\lambda)$ and its gain over the teacher risk $R(\lambda)$, and (ii) quantify how these quantities depend on the aspect ratio γ , the signal-to-noise ratio SNR, and signal-covariance alignment (to be defined below).

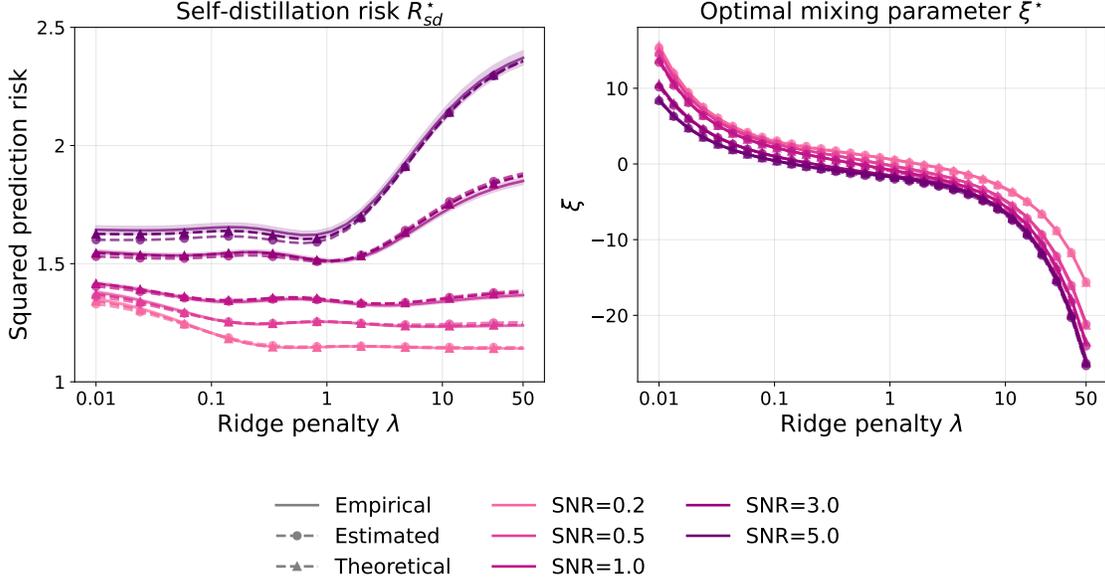


Figure 4: **Theoretical versus empirical risks.** Asymptotic SD risks and optimal mixing weights versus λ across multiple SNR values. Empirical curves are averaged over 30 simulations. Estimated curves are obtained using the proposed one-shot tuning method (Section 4), and the theoretical curves are obtained from Theorem 3.1. Setting: $n = 400$, $p = 200$, $\sigma^2 = 1$, $r^2 = \sigma^2 \text{SNR}$; Σ is AR1, β is a deterministic signal aligned with the top 10% eigenvectors of Σ (alignment factor 0.9). (See Section F for more details.)

3.1 Data Assumptions

We assume $\{(x_i, y_i)\}_{i=1}^n$ are drawn i.i.d. from a distribution $P_{x,y}$ satisfying the following:

Assumption A (Data distribution). The covariate vector $x \sim P_x$ admits the representation $x = \Sigma^{1/2}z$, where $\Sigma \in \mathbb{R}^{p \times p}$ is deterministic and positive definite with eigenvalues uniformly bounded away from 0 and ∞ , and $z \in \mathbb{R}^p$ has i.i.d. entries with mean 0, variance 1, and uniformly bounded $(4 + \mu)$ -th moment for some $\mu > 0$. The response $y \sim P_y$ has mean 0 and uniformly bounded $(4 + \nu)$ -th moment, for some $\nu > 0$.

The feature structure imposed in Assumption A is standard in random matrix analyses of high-dimensional regression; see, e.g., Bai and Silverstein (2010); Bartlett et al. (2021); Misiakiewicz and Montanari (2024). We do *not* assume a well-specified linear model for the response y . Instead, we parameterize $P_{x,y}$ by $(\Sigma, \beta, \sigma^2)$, where $\beta := \Sigma^{-1}\mathbb{E}[xy]$ denotes the parameter of the population linear L_2 projection of y onto x and $\sigma^2 := \text{Var}(y - x^\top \beta)$ denotes the corresponding residual variance. Finally, We denote the signal energy by $r^2 := \|\beta\|_2^2$ and $\text{SNR} := r^2/\sigma^2$.

Throughout this section, we focus on the in-distribution squared prediction risk in which the test point (x_0, y_0) in (5) is an independent copy drawn from the same distribution as the training samples. Extensions to out-of-distribution risks is possible building on results of Tripuraneni et al. (2021); Patil et al. (2024) but is technically involved and left for future work.

3.2 Asymptotics of Optimal Self-Distillation Risk and Mixing Weight

Recall from Section 2.2 that, for each fixed $\lambda > 0$, the optimal SD risk (8) can be expressed in terms of three random quantities: the teacher risk $R(\lambda)$, the PD risk $R_{\text{pd}}(\lambda)$, and the residual correlation term $C(\lambda)$. Under proportional asymptotics, these quantities concentrate and converge to deterministic limits, which we denote by $\mathcal{R}(\lambda)$, $\mathcal{R}_{\text{pd}}(\lambda)$, and $\mathcal{C}(\lambda)$, respectively, that only depend on $(\Sigma, \beta, \sigma^2)$, γ , and λ . To characterize these limits, we introduce some scalar parameters.

Write $\bar{\text{tr}}(A) := \text{tr}(A)/p$ for the normalized trace. For a fixed $\lambda > 0$, let $\kappa = \kappa(\lambda) > 0$ be the unique solution to the fixed-point equation

$$\kappa = \lambda + \gamma \kappa \bar{\text{tr}}(\Sigma(\Sigma + \kappa I_p)^{-1}). \quad (12)$$

Let $G := (\Sigma + \kappa I_p)^{-1}$ be the resolvent at κ and for any $k \in \{2, 3, 4\}$ let the trace and signal-covariance alignment functionals be

$$t_k := \gamma \bar{\text{tr}}(\Sigma^2 G^k), \quad q_k := \beta^\top G^k \Sigma \beta, \quad (13)$$

respectively. Next, setting $b := (1 - t_2)^{-1}$, define the variance-trace combinations as

$$u_2 := t_2 b, \quad u_3 := t_3 b^3, \quad u_4 := t_4 b^4 + 2t_3^2 b^5. \quad (14)$$

Finally, define the coefficients

$$a_2 := bE^2 + b^4 \kappa^2 \lambda^2 t_4 + b^5 \kappa^2 \lambda^2 t_3^2, \quad a_3 := 2b^2 \kappa \lambda E, \quad a_4 := b^3 \kappa^2 \lambda^2. \quad (15)$$

where $E := \kappa - b\lambda + b^2 \kappa \lambda t_3$. (The dependence on λ of here and throughout is upressed for readability.) We are now ready to state the main result of this section.

Theorem 3.1 (Risk asymptotics). Under Assumption A, as $n, p \rightarrow \infty$ with $p/n \rightarrow \gamma \in (0, \infty)$, for each fixed $\lambda > 0$, we have

$$\xi^*(\lambda) \xrightarrow{P} \frac{\mathcal{R}(\lambda) - \mathcal{C}(\lambda)}{\mathcal{R}(\lambda) + \mathcal{R}_{\text{pd}}(\lambda) - 2\mathcal{C}(\lambda)}, \quad R_{\text{sd}}^*(\lambda) \xrightarrow{P} \mathcal{R}(\lambda) - \frac{(\mathcal{R}(\lambda) - \mathcal{C}(\lambda))^2}{\mathcal{R}(\lambda) + \mathcal{R}_{\text{pd}}(\lambda) - 2\mathcal{C}(\lambda)}, \quad (16)$$

where the individual component limits are as follows:

$$\begin{aligned} \mathcal{R}(\lambda) &:= \kappa^2 b q_2 + \sigma^2 u_2 + \sigma^2, \\ \mathcal{C}(\lambda) &:= 2\kappa^2 b q_2 - (\kappa b E q_2 + \kappa^2 b^2 \lambda q_3) + \sigma^2 (u_2 - \lambda u_3) + \sigma^2, \\ \mathcal{R}_{\text{pd}}(\lambda) &:= 4\kappa^2 b q_2 - 2(2\kappa b E q_2 + 2\kappa^2 b^2 \lambda q_3) + (a_2 q_2 + a_3 q_3 + a_4 q_4) + \sigma^2 (u_2 - 2\lambda u_3 + \lambda^2 u_4) + \sigma^2. \end{aligned}$$

To our knowledge, this is the first precise characterization of the optimal SD risk in the proportional asymptotics regime that holds under the general setting of anisotropic covariance and deterministic signal. The spectrum of Σ enters through resolvent traces such as t_2, t_3, t_4 , while the signal enters through the alignment functionals q_2, q_3, q_4 . As illustrated in Figure 4, the theoretical predictions closely match empirical risks even for moderate n and p .

In the isotropic-signal specialization when $\beta \sim \mathcal{N}(0, (r^2/p)I_p)$, the quadratic forms q_k simplify to trace functionals; see Section B.7. In this setting, the strict-improvement and sign behavior becomes particularly transparent, as illustrated in our next result.

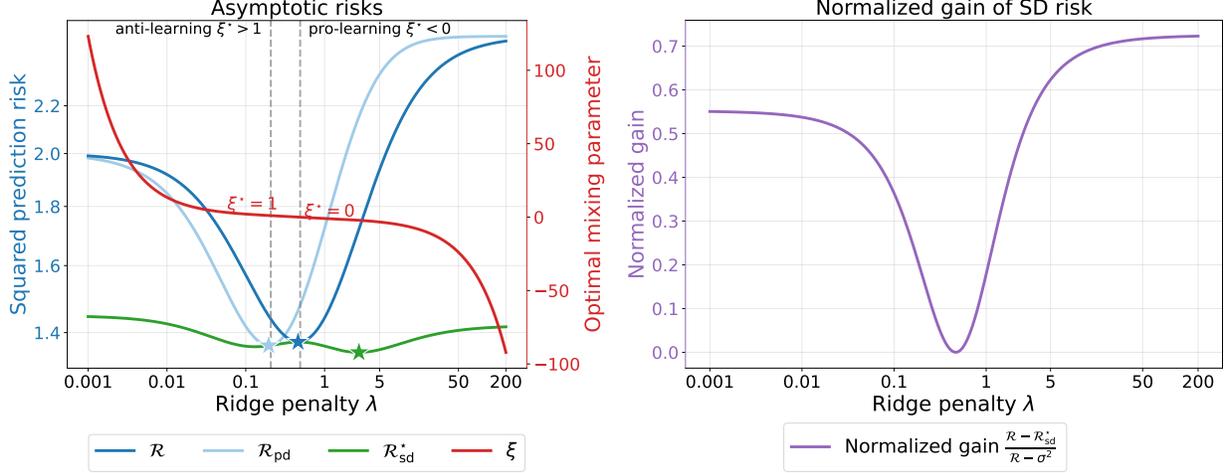


Figure 5: **Asymptotic gain over the teacher.** Deterministic limits of the risks and the gain $\mathcal{R}(\lambda) - \mathcal{R}_{\text{sd}}^*(\lambda)$. Same setting as Figure 4 with $r^2 = \sigma^2 = 1$.

Corollary 3.2. Under Assumption A with $\beta \sim \mathcal{N}(0, (r^2/p)I_p)$, we have

- (a) $\mathcal{R}_{\text{sd}}^*(\lambda) < \mathcal{R}(\lambda)$ for all $\lambda \neq \lambda^* := \gamma \frac{\sigma^2}{r^2}$.
- (b) $\xi^*(\lambda) < 0$ iff $\lambda > \lambda^*$, and $\xi^*(\lambda) > 0$ iff $\lambda < \lambda^*$, i.e., $\text{sign}(\xi^*(\lambda)) = \text{sign}(\lambda^* - \lambda)$ for all $\lambda \neq \lambda^*$.

Thus, in the isotropic-signal setting, optimal SD strictly improves upon ridge at every suboptimal value of λ , and it matches the (statistically optimal, for the isotropic design) ridge risk only at the ridge-optimal value λ^* . Moreover, the sign of the optimal mixing weight is determined solely by whether the ridge model is under-regularized ($\lambda < \lambda^*$) or over-regularized ($\lambda > \lambda^*$). Simulation studies show in Figure 5 are in close agreement with these findings: (i) $\mathcal{R}_{\text{sd}}^*(\lambda)$ lies below $\mathcal{R}(\lambda)$ for all nonstationary λ , (ii) $\xi^*(\lambda)$ flips sign at λ^* , and (iii) the gain is largest in strongly under- or over-regularized regimes. Additional illustrations for other covariance and signal geometries are provided in Section E.4.

3.3 Self-Distillation Risks with Extreme Regularization

Next, we study how close optimal SD can get to the best possible predictor when the teacher is *extremely* under- or over-regularized. We focus on the isotropic design and isotropic signal setting ($\Sigma = I_p$ and $\beta \sim \mathcal{N}(0, (r^2/p)I_p)$), where the ridge-optimal predictor is Bayes-optimal and its asymptotic risk \mathcal{R}^* is known in closed form (Dobriban and Wager, 2018; Hastie et al., 2022). We compare the limiting SD risk $\mathcal{R}_{\text{sd}}^*(\lambda)$ to \mathcal{R}^* as $\lambda \rightarrow 0$ and $\lambda \rightarrow \infty$.

Proposition 3.3 (Comparison with the optimal ridge). Assume $\Sigma = I_p$ and $\beta \sim \mathcal{N}(0, (r^2/p)I_p)$ and let $S^*(\text{SNR}, \gamma) := \frac{1}{2\gamma}(\text{SNR}(\gamma - 1) - \gamma + \sqrt{4\text{SNR}\gamma^2 + (\text{SNR}(\gamma - 1) - \gamma)^2})$. Then

$$\lim_{\lambda \rightarrow 0} \frac{\mathcal{R}_{\text{sd}}^*(\lambda) - \mathcal{R}^*}{\mathcal{R}^*} = \begin{cases} \frac{\text{SNR}(1 - \gamma)^2 + \gamma}{(\text{SNR}(1 - \gamma)^3 + \gamma(1 - \gamma^2))(S^* + 1)} - 1, & \gamma \in (0, 1), \\ \frac{\text{SNR}^2(\gamma - 1)^4 + \text{SNR}\gamma(2\gamma + 1)(\gamma - 1)^2 + \gamma^4}{(\text{SNR}\gamma(\gamma - 1)^3 + \gamma^2(\gamma^2 - 1))(S^* + 1)} - 1, & \gamma \in (1, \infty), \end{cases}$$

$$\lim_{\lambda \rightarrow \infty} \frac{\mathcal{R}_{\text{sd}}^*(\lambda) - \mathcal{R}^*}{\mathcal{R}^*} = \frac{\text{SNR}^2\gamma + \text{SNR}(2\gamma + 1) + \gamma}{(\text{SNR}(\gamma + 1) + \gamma)(S^* + 1)} - 1.$$

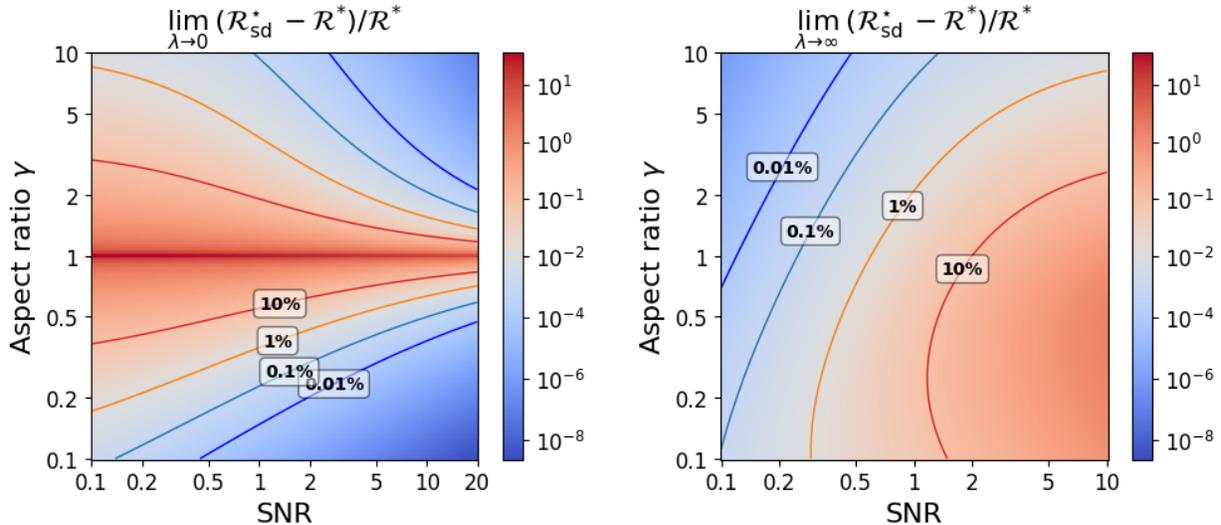


Figure 6: **Approximating the best predictor.** Percentage difference between $\mathcal{R}_{\text{sd}}^*(\text{SNR}, \gamma)$ and the Bayes-optimal ridge risk $\mathcal{R}^*(\text{SNR}, \gamma)$ in the isotropic design and signal setting.

Proposition 3.3 expresses the relative suboptimality of optimal SD under extreme regularization scenarios explicitly in terms of (γ, SNR) under the isotropic design and signal setting. The resulting percentage gaps are illustrated in Figure 6. Across a wide range of (γ, SNR) , optimal SD can be remarkably close to Bayes-optimal performance even when the teacher is extremely under- or over-regularized (e.g., within 0.01% for $(\text{SNR}, \gamma) = (2, 0.2)$ as $\lambda \rightarrow 0$). The heatmaps also suggest a qualitative dichotomy: for low SNR, over-regularization is broadly preferable across γ , while for high SNR, under-regularization is preferable except near the square design regime when $\gamma \approx 1$. For comparison, we also report the analogous extreme- λ gaps between the original ridge risk $\mathcal{R}(\lambda)$ and the ridge-optimal risk \mathcal{R}^* in Section E.5, along with experiments with other covariance structures.

4 One-Shot Tuning and Risk Estimation

The results in Section 3 quantify the benefits of optimal self-distillation. However, the deterministic equivalents in Theorem 3.1 depend on population quantities (e.g., the covariance spectrum and signal-covariance alignment), which are unknown in practice. For the practical question for tuning the mixing parameter ξ , a standard method involves grid search combined with split cross-validation (CV), but this is computationally expensive (requiring repeated retrainings over candidate ξ) and statistically inefficient in high dimensions due to sample splitting (see, e.g., Rad and Maleki (2020)). In this section, we propose a computationally efficient one-shot procedure that estimates $\xi^*(\lambda)$ from the training data without grid search or hold-out sets.

4.1 Risk Estimators via Generalized Cross-Validation

Our starting point is the closed-form identity for optimal SD given in Proposition 2.1, which expresses $\xi^*(\lambda)$ and $R_{\text{sd}}^*(\lambda)$ in terms of the teacher risk $R(\lambda)$, the PD risk $R_{\text{pd}}(\lambda)$, and the residual correlation term $C(\lambda)$. We construct estimators for these three terms using training data by means of generalized cross-validation (GCV) and its variants, along with plug-in estimation of $\xi^*(\lambda)$.

Let $\hat{y}_\lambda := f_\lambda(X)$ and $\hat{y}_{\text{pd},\lambda} := f_{\text{pd},\lambda}(X)$ be the fitted values of the ridge teacher and the pure-distilled predictors, respectively. Our estimators involve the notion of *effective degrees of freedom* from the theory of statistical optimism (Efron, 1983, 1986). The degrees of freedom $\text{df}(f)$ of a (possibly nonlinear) predictor f is measured by the trace of the operator $y \mapsto (\partial/\partial y)f(y)$ (Hastie and Tibshirani, 1990; Stein, 1981). In particular, for linear smoothers, this corresponds to the trace of the smoothing matrix (the so-called “hat” matrix).

Set $\text{df}_\lambda := \text{df}(f_\lambda)$ and $\text{df}_{\text{pd},\lambda} := \text{df}(f_{\text{pd},\lambda})$ and define the GCV-corrected residuals

$$\hat{r}_\lambda := \frac{y - \hat{y}_\lambda}{1 - \text{df}_\lambda/n}, \quad \hat{r}_{\text{pd},\lambda} := \frac{y - \hat{y}_{\text{pd},\lambda}}{1 - \text{df}_{\text{pd},\lambda}/n}. \quad (17)$$

We estimate the teacher and PD prediction risks, as well as their residual correlation term, by

$$\hat{R}(\lambda) := \frac{\|\hat{r}_\lambda\|_2^2}{n}, \quad \hat{R}_{\text{pd}}(\lambda) := \frac{\|\hat{r}_{\text{pd},\lambda}\|_2^2}{n}, \quad \hat{C}(\lambda) := \frac{\langle \hat{r}_\lambda, \hat{r}_{\text{pd},\lambda} \rangle}{n}, \quad (18)$$

respectively. Here $\hat{R}(\lambda)$ coincides with the standard ridge GCV estimator, while $\hat{R}_{\text{pd}}(\lambda)$ and $\hat{C}(\lambda)$ extend the same df correction principle to the PD and cross-term quantities appearing in Proposition 2.1.

4.2 One-Shot Estimators for Optimal Mixing Weight and Optimal SD Risk

Plugging the estimators (18) into the exact identities in (8) yields the one-shot estimators

$$\hat{\xi}^*(\lambda) := \frac{\hat{R}(\lambda) - \hat{C}(\lambda)}{\hat{R}(\lambda) + \hat{R}_{\text{pd}}(\lambda) - 2\hat{C}(\lambda)}, \quad \hat{R}_{\text{sd}}^*(\lambda) := \hat{R}(\lambda) - \frac{(\hat{R}(\lambda) - \hat{C}(\lambda))^2}{\hat{R}(\lambda) + \hat{R}_{\text{pd}}(\lambda) - 2\hat{C}(\lambda)}. \quad (19)$$

Note that the estimated denominator equals

$$\hat{D}(\lambda) := \hat{R}(\lambda) + \hat{R}_{\text{pd}}(\lambda) - 2\hat{C}(\lambda) = \frac{1}{n} \|\hat{r}_\lambda - \hat{r}_{\text{pd},\lambda}\|_2^2 \geq 0, \quad (20)$$

mirroring the nonnegativity of its population counterpart in (7). When $\hat{D}(\lambda)$ is extremely small in finite samples, one may stabilize (19) by adding a small ridge term to the denominator; in our experiments this was not necessary. Our next result shows that the one-shot estimates are consistent in the proportional regime.

Theorem 4.1 (Consistency of one-shot SD tuning). Under Assumption A, as $n, p \rightarrow \infty$ with $p/n \rightarrow \gamma \in (0, \infty)$, for each fixed $\lambda > 0$, we have

$$\hat{\xi}^*(\lambda) - \xi^*(\lambda) \xrightarrow{\text{P}} 0, \quad \hat{R}_{\text{sd}}^*(\lambda) - R_{\text{sd}}^*(\lambda) \xrightarrow{\text{P}} 0. \quad (21)$$

Compared with grid-search CV over ξ , the one-shot procedure has two key advantages: (i) CV requires retraining a student model for each candidate ξ , whereas (19) selects $\hat{\xi}^*(\lambda)$ in closed form from a single set of fitted quantities at the given λ . (ii) Split CV reduces the effective training sample size (e.g., to 4/5) and suffers from nonzero bias in high-dimensional settings where p is comparable to n ; in contrast, the one-shot estimators use all n samples without hold-out sets.

4.3 Real Data Experiments

To illustrate the utility of the one-shot tuning method, we apply (19) across a range of λ values on several real datasets. For regression, we consider UCI BlogFeedback and Communities and Crime. For classification, we apply ridge regression on pretrained neural network features (ResNet-18/34) for CIFAR10 and CIFAR100 (details in Section F).

Across these tasks, Figures 2 and 13 show that $\widehat{R}_{\text{sd}}^*(\lambda)$ closely tracks the test risk of the optimally distilled student, particularly in settings with small train and test discrepancy (e.g., CIFAR dataset benchmarks). Moreover, when the teacher is over-regularized, the one-shot estimate correctly selects negative $\widehat{\xi}^*(\lambda)$, so that the SD predictor corrects excessive shrinkage. This regime would be missed by restricting ξ to $[0, 1]$ (see Figure 16 for constrained SD risks as a comparison). Additional experiments and sample-size variations are provided in Section E.3, and for CIFAR10/CIFAR100 we also report the corresponding test accuracies in Figure 14.

5 Extensions and Variants

Our results so far focus on optimal *one-round* SD for *ordinary* ridge regression, where the student is refit on the *same* design matrix X as the teacher. In this section, we briefly outline several extensions that are naturally captured by the same “structural” viewpoint from Section 2.

5.1 Multiple Rounds of Self-Distillation

Fix $\lambda > 0$ and set $y^{(0)} := y$. A natural *recursive* multi-round scheme iteratively self-distills using the previous round’s ridge predictions. For each $k \geq 0$, let $f_\lambda^{(k)}$ denote the teacher ridge regression (2) trained on $(X, y^{(k)})$ with penalty λ , and write $\widehat{y}_\lambda^{(k)} := f_\lambda^{(k)}(X) \in \mathbb{R}^n$ for its fitted values. Let $f_{\text{pd},\lambda}^{(k)}$ denote the corresponding PD refit of round k , i.e., ridge trained on $(X, \widehat{y}_\lambda^{(k)})$ with the same penalty λ . Given a mixing weight $\xi_{k+1} \in \mathbb{R}$, define the round- $(k+1)$ SD predictor $f_{\text{sd},\lambda,\xi_{k+1}}^{(k+1)}$ as the ridge fit obtained from the same mixed-loss construction as in (3), but with base labels $y^{(k)}$ and teacher pseudo-labels $\widehat{y}_\lambda^{(k)}$. A simplification (see Section D.1.1) shows this is equivalent to fitting ridge regression on the mixed labels (see Figure 7 for an illustration):

$$y_{\lambda,\xi_{k+1}}^{(k+1)} := (1 - \xi_{k+1}) y^{(k)} + \xi_{k+1} \widehat{y}_\lambda^{(k)}. \quad (22)$$

Because ridge is linear in the response, the round- $(k+1)$ predictor $f_{\text{sd},\lambda,\xi_{k+1}}^{(k+1)}$ lies on the affine path:

$$f_{\text{sd},\lambda,\xi_{k+1}}^{(k+1)} = (1 - \xi_{k+1}) f_\lambda^{(k)} + \xi_{k+1} f_{\text{pd},\lambda}^{(k)}. \quad (23)$$

As in Section 2, we evaluate any predictor f by the conditional squared prediction risk $R(f) := \mathbb{E}[(y_0 - f(x_0))^2 \mid \mathcal{D}]$ for a (possibly out-of-distribution) test pair (x_0, y_0) with finite conditional second moments. Define the round- k risk: $R_k(\lambda) := R(f_\lambda^{(k)})$, round- k correlation: $C_k(\lambda) := \mathbb{E}[(y_0 - f_\lambda^{(k)}(x_0))(y_0 - f_{\text{pd},\lambda}^{(k)}(x_0)) \mid \mathcal{D}]$, and round- k discrepancy: $D_k(\lambda) := \mathbb{E}[(f_\lambda^{(k)}(x_0) - f_{\text{pd},\lambda}^{(k)}(x_0))^2 \mid \mathcal{D}] \geq 0$. For each round $k \geq 0$, let $\xi_{k+1}^* \in \arg \min_{\xi_{k+1} \in \mathbb{R}} R(f_{\text{sd},\lambda,\xi_{k+1}}^{(k+1)})$ denote the optimal (unconstrained) mixing weight. Write $y^{(k+1)} = y_{\lambda,\xi_{k+1}^*}^{(k+1)}$, $f_\lambda^{(k+1)} := f_{\text{sd},\lambda,\xi_{k+1}^*}^{(k+1)}$, and $R_{k+1}(\lambda) := R(f_\lambda^{(k+1)})$.

All the one-round structural identities from Section 2 apply at each round k by viewing $y^{(k)}$ as the “base” labels and $(f_\lambda^{(k)}, f_{\text{pd},\lambda}^{(k)})$ as the round- k (teacher, PD) pair. In particular, whenever $D_k(\lambda) > 0$,

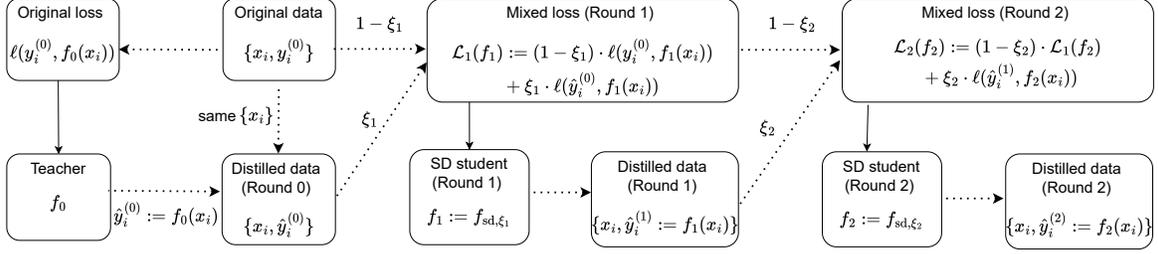


Figure 7: Visual illustration of the recursive multi-round self-distillation process for ridge regression.

the round-wise optimizer and improvement satisfy the same formulas as Proposition 2.1:

$$\xi_{k+1}^* = \frac{R_k(\lambda) - C_k(\lambda)}{D_k(\lambda)}, \quad R_{k+1}(\lambda) = R_k(\lambda) - \frac{(R_k(\lambda) - C_k(\lambda))^2}{D_k(\lambda)}.$$

This leads to the following risk monotonicity result for optimal multi-round SD:

Proposition 5.1 (Monotonicity of optimal recursive multi-round self-distillation). Fix $\lambda > 0$. The optimal SD risks are monotone (weakly) decreasing in the number of rounds k :

$$R_{k+1}(\lambda) \leq R_k(\lambda) \quad \text{for all } k \geq 0.$$

Assuming $D_k(\lambda) > 0$, the optimal SD mixing weights and risks admit the closed forms:

$$\xi_{k+1}^* = -\frac{\lambda}{2} \frac{R'_k(\lambda)}{D_k(\lambda)}, \quad R_{k+1}(\lambda) = R_k(\lambda) - \frac{\lambda^2}{4} \frac{(R'_k(\lambda))^2}{D_k(\lambda)}, \quad (24)$$

where $R'_k(\lambda)$ is the derivative along the ridge path for the round- k teacher with $y^{(k)}$ held fixed. Thus, if $R'_k(\lambda) \neq 0$, then $R_{k+1}(\lambda) < R_k(\lambda)$, and $\text{sign}(\xi_{k+1}^*(\lambda)) = -\text{sign}(R'_k(\lambda))$.

As with all results in Section 2, Proposition 5.1 is purely structural: it holds conditionally on $\mathcal{D} = (X, y)$ with probability one and for any squared prediction risk (including out-of-distribution).

Several repeated-distillation formulations instead remain *anchored* to the original (ground-truth) labels (Garg et al., 2025; Alemohammad et al., 2023). Specifically, given $\xi_{k+1} \in \mathbb{R}$, define the anchored round- $(k+1)$ SD predictor as the ridge fit obtained by mixing the loss (as in (3)) against the original labels $y^{(0)}$ and the current pseudo-labels $\hat{y}_\lambda^{(k)}$. As before, a simplification shows that this is equivalent to ordinary ridge regression trained on the anchored mixed labels:

$$y_{\lambda, \xi_{k+1}}^{(k+1), \text{anch}} := (1 - \xi_{k+1}) y^{(0)} + \xi_{k+1} \hat{y}_\lambda^{(k)}. \quad (25)$$

By linearity of ridge in the response, this anchored family is the affine path:

$$f_{\text{sd}, \lambda, \xi_{k+1}}^{(k+1), \text{anch}} = (1 - \xi_{k+1}) f_\lambda^{(0)} + \xi_{k+1} f_{\text{pd}, \lambda}^{(k)},$$

which in general need *not* contain the previous-round predictor $f_\lambda^{(k)}$. Consequently, the nesting argument behind Proposition 5.1 breaks down and monotonicity can fail; see Figure 8.

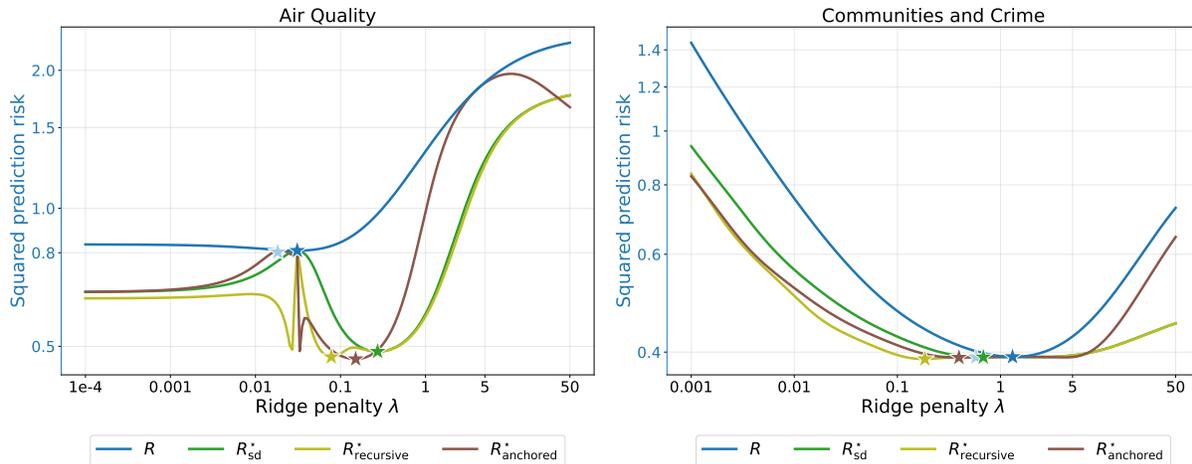


Figure 8: **Recursive versus anchored multi-round self-distillation.** Test risks of the teacher ridge (R , in blue, trained using $y^{(0)}$), one-round ($k = 1$) optimal self-distilled ridge (R_{sd}^* , in green, trained using $y^{(0)}$ and $\hat{y}^{(0)}$), two-round ($k = 2$) recursive ($R_{\text{recursive}}^*$, in olive, trained using $y^{(1)}$ and $\hat{y}^{(1)}$) and anchored (R_{anchored}^* , in brown, trained using $y^{(0)}$ and $\hat{y}^{(1)}$) self-distillation. Recursive mixing is monotone when ξ_k is optimized each round (Proposition 5.1); the two-round risk curve $R_{\text{recursive}}^*(\lambda)$ uniformly dominates the optimal one-round $R_{\text{sd}}^*(\lambda)$. Anchored mixing can be nonmonotone because the round- k family need not contain the round- $(k - 1)$ predictor; the two-round risk curve $R_{\text{anchored}}^*(\lambda)$ may be larger than optimal one-round $R_{\text{sd}}^*(\lambda)$.

5.2 Self-Distillation with Fresh Unlabeled Features

Our structural results in Section 2 rely on the student refit using the *same* design matrix X as the teacher. In particular, the same- X mixed-loss SD predictor (3) reduces to an affine path (Equation (4)) in ξ , so the risk $R_{\text{sd}}(\lambda, \xi)$ is a quadratic function of ξ and admits closed-form characterizations (Proposition 2.1 and Theorem 2.2).

A common pseudo-labeling variant instead refits using additional *fresh* unlabeled covariates; see Figure 9. Let $\tilde{X} \in \mathbb{R}^{m \times p}$ be independent of $\mathcal{D} = (X, y)$, and define the teacher pseudo-labels on \tilde{X} by $\tilde{y}_\lambda := f_\lambda(\tilde{X}) \in \mathbb{R}^m$, where f_λ is the ridge teacher trained on \mathcal{D} as in (2). In direct analogy with (3), one may define the fresh- X *mixed-loss* SD predictor $f_{\text{sd}, \lambda, \xi}^{\text{frmix}}(x)$ as

$$x^\top \underset{\beta \in \mathbb{R}^p}{\text{argmin}} \left\{ (1 - \xi) \|y - X\beta\|_2^2 / n + \xi \|\tilde{y}_\lambda - \tilde{X}\beta\|_2^2 / m + \lambda \|\beta\|_2^2 \right\}. \quad (26)$$

Unlike the same- X case, (26) does not reduce to ridge regression on a single mixed-label vector: because the two quadratic losses involve different Gram matrices, the map $\xi \mapsto f_{\text{sd}, \lambda, \xi}^{\text{frmix}}$ is generally *not* an affine path, and the risk $R_{\text{sd}}^{\text{frmix}}(\lambda, \xi) := R(f_{\text{sd}, \lambda, \xi}^{\text{frmix}})$ is no longer quadratic in ξ (see Lemma D.1).³ Consequently, the two-predictor identities for $\xi^*(\lambda)$ and $R_{\text{sd}}^*(\lambda)$ in Proposition 2.1 and the derivative identities in Theorem 2.2 do not apply directly in this fresh- X mixed-loss setting.

To retain an affine structure, one may instead define a fresh- X PD refit and then mix predictors explicitly. Let $f_{\text{pd}, \lambda}^{\text{fr}}$ denote ridge regression trained on the pseudo-labeled dataset $(\tilde{X}, \tilde{y}_\lambda)$ with the

³When $\xi \notin [0, 1]$, the objective in (26) need not be convex because it mixes two different quadratic forms. Throughout, we interpret (26) only for values of ξ such that the Hessian in β is positive definite, i.e., $(1 - \xi) X^\top X / n + \xi \tilde{X}^\top \tilde{X} / m + \lambda I_p \succ 0$, so the displayed argmin is well-defined.

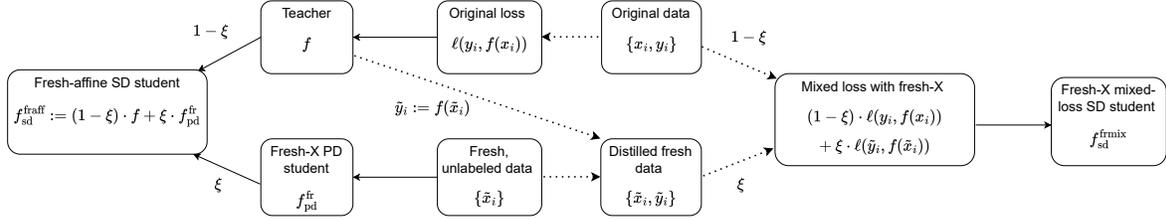


Figure 9: Visual illustration of self-distillation with fresh unlabeled features.

same penalty λ . Define the fresh- X *affine* SD family

$$f_{\text{sd},\lambda,\xi}^{\text{fraff}}(x) := (1 - \xi) f_{\lambda}(x) + \xi f_{\text{pd},\lambda}^{\text{fr}}(x). \quad (27)$$

Because (27) is an explicit two-predictor affine path, the corresponding risk $R_{\text{sd}}^{\text{fraff}}(\lambda, \xi) := R(f_{\text{sd},\lambda,\xi}^{\text{fraff}})$ is quadratic in ξ . Thus, the identities in Proposition 2.1 apply directly (with $f_{\text{pd},\lambda}$ replaced by $f_{\text{pd},\lambda}^{\text{fr}}$) to characterize the optimal mixing $\xi_{\text{sd}}^{*,\text{fraff}}(\lambda) \in \operatorname{argmin}_{\xi \in \mathbb{R}} R_{\text{sd}}^{\text{fraff}}(\lambda, \xi)$ and the optimal risk $R_{\text{sd}}^{*,\text{fraff}}(\lambda) := \min_{\xi \in \mathbb{R}} R_{\text{sd}}^{\text{fraff}}(\lambda, \xi)$. However, the same- X coupling behind the derivative-based characterization Theorem 2.2 is absent here, so pointwise strict improvements are no longer structurally guaranteed.

Empirically, when the refit sample size matches the teacher sample size (i.e., $m = n$), the same- X SD predictor tends to dominate both the fresh- X mixed-loss student (26) as well as the fresh- X affine student (27) across λ ; see Figure 10. Motivated by this behavior, we provide a prototype theoretical comparison in an isotropic in-distribution setting. Specializing the asymptotic analysis of Section 3 for same- X (Lemma D.2) and deriving a corresponding asymptotic analysis for affine fresh- X (Lemma D.3) under isotropic covariance and isotropic signal, we show that when $p/n \rightarrow \gamma$ and $p/m \rightarrow \gamma$, the same- X optimal SD risk uniformly dominates the fresh-affine optimal SD risk. A full analysis under general covariance and signal structure, and for the mixed-loss fresh- X student (26), is left for future work.

Theorem 5.2 (Same- X dominates fresh- X). Under Assumption A with $\Sigma = I_p$ and $\beta \sim \mathcal{N}(0, (r^2/p)I_p)$, as $m, n, p \rightarrow \infty$ with $p/m, p/n \rightarrow \gamma \in (0, \infty)$, for every fixed $\lambda > 0$, we have

$$\mathcal{R}_{\text{sd}}^{*,\text{same}}(\lambda) \leq \mathcal{R}_{\text{sd}}^{*,\text{fraff}}(\lambda),$$

where $\mathcal{R}_{\text{sd}}^{*,\text{same}}(\lambda)$ and $\mathcal{R}_{\text{sd}}^{*,\text{fraff}}(\lambda)$ denote the limiting (deterministic) optimal SD risks for the same- X SD predictor (4) and the affine fresh- X SD predictor (27), respectively

This result is somewhat counterintuitive as one might expect that additional unlabeled data would help self-distillation. It highlights that the gains characterized in the main paper are intimately tied to self-distillation on the same data. Understanding how the comparison changes when the amount of unlabeled data grows (e.g., $m \gg n$) is an interesting direction for future work.

5.3 Self-Distillation with Other Ridge Variants

The structural results of Section 2 extend beyond ordinary ridge regression to a class of (ridge) resolvent-based smoothers. At a high level, two properties drive the extension: (i) the teacher is a

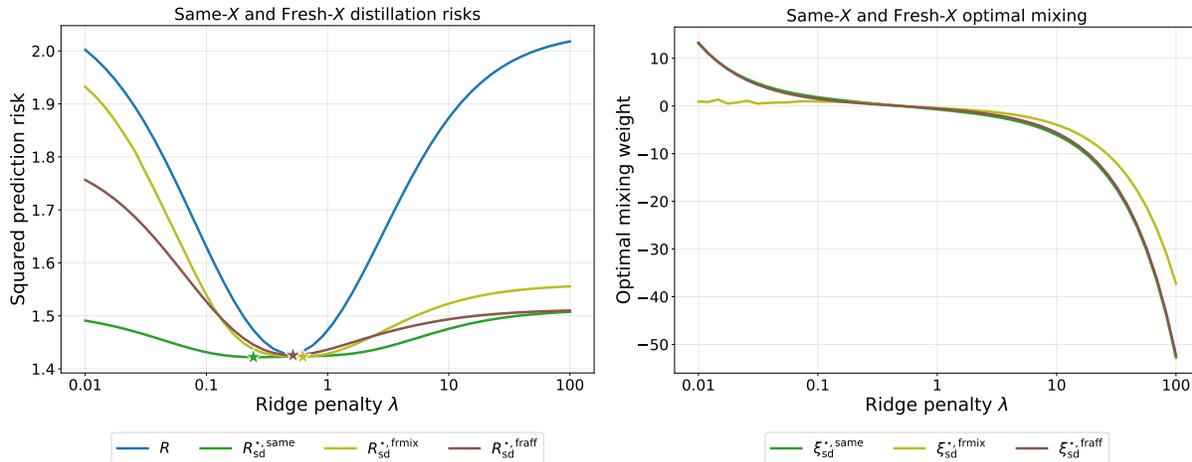


Figure 10: **Same- X versus fresh- X self-distillation risks.** Empirical test risks and optimal mixing weights (averaged over 30 simulations) in an isotropic setting with $n = 200$, $p = 100$, $r^2 = 1$, and $\sigma^2 = 1$. For each λ , we choose the fresh- X mixing weight ξ by grid search over 3,000 values in $[-100, 100]$, restricting to values for which the mixed-loss objective (26) is strictly convex, and picking the one with the lowest empirical test risk.

linear smoother in the labels, and (ii) the student refit applies the same *resolvent-based* smoothing family (at the same regularization level λ). Under (i), the SD family is again an affine path in the mixing weight ξ , so the two-predictor identities (e.g., Proposition 2.1) apply under any squared prediction risk. Under (ii), the teacher–PD gap admits a (surprising) derivative representation, resulting in strict improvement and sign rule properties analogous to Theorem 2.2. We state the generic result first, then highlight two representative examples.

Fix training inputs $X = (x_1^\top, \dots, x_n^\top)^\top$ and labels $y \in \mathbb{R}^n$. Let $\{f_\lambda\}_{\lambda>0}$ be a teacher predictor family such that for each λ there exists a vector-valued map $s_\lambda(\cdot) \in \mathbb{R}^n$ (depending on X, λ but not on y) satisfying $f_\lambda(x) = s_\lambda(x)^\top y$ and $\hat{y}_\lambda := f_\lambda(X) = S_\lambda y$ for some (smoothing) matrix $S_\lambda \in \mathbb{R}^{n \times n}$ depending on X, λ but not on y . Define the (same- λ) PD student refit by reapplying the same smoother to the teacher fitted values: $f_{\text{pd},\lambda}(x) := s_\lambda(x)^\top \hat{y}_\lambda$. Define the SD predictor by the same mixed-label construction as in (3) (now with teacher pseudo-labels \hat{y}_λ), which is equivalent to applying the smoother to the mixed labels $y(\xi) := (1 - \xi)y + \xi \hat{y}_\lambda$. Hence, the SD predictor can be expressed as (see Section D.3.1):

$$f_{\text{sd},\lambda,\xi}(x) = s_\lambda(x)^\top y(\xi) = (1 - \xi)f_\lambda(x) + \xi f_{\text{pd},\lambda}(x). \quad (28)$$

For any squared prediction risk $R(\cdot)$ as in (5), write $R(\lambda) := R(f_\lambda)$, $R_{\text{pd}}(\lambda) := R(f_{\text{pd},\lambda})$, and define:

$$C(\lambda) := \mathbb{E}[(y_0 - f_\lambda(x_0))(y_0 - f_{\text{pd},\lambda}(x_0)) \mid \mathcal{D}], \quad D(\lambda) := \mathbb{E}[(f_\lambda(x_0) - f_{\text{pd},\lambda}(x_0))^2 \mid \mathcal{D}].$$

(As in Section 2, it is easy to check that $D(\lambda) = R(\lambda) + R_{\text{pd}}(\lambda) - 2C(\lambda)$.) Assume in addition that the ridge-smoother family satisfies the derivative (tangent) identity

$$f_\lambda(x) - f_{\text{pd},\lambda}(x) = -\lambda \partial_\lambda f_\lambda(x) \quad \text{for all } x, \quad (29)$$

and that $\lambda \mapsto R(\lambda)$ is differentiable. Then the derivative-based characterization from Theorem 2.2 extends as follows.

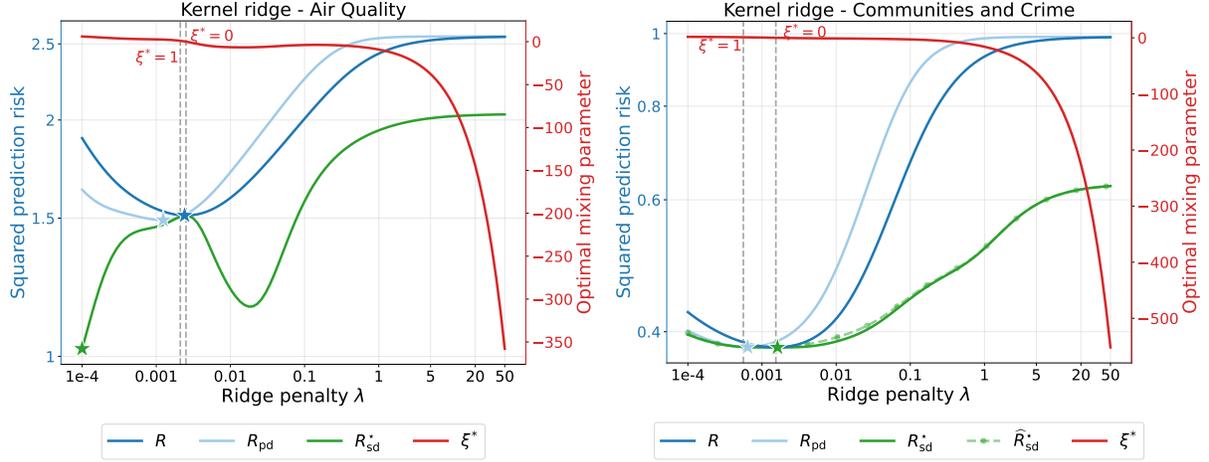


Figure 11: **Kernel ridge and kernel SD ridge regression.** Test risks of Gaussian kernel ridge and kernel SD ridge on two datasets.

Theorem 5.3 (Ridge-smoother strict improvement and sign rule). Fix $\lambda > 0$ and assume $D(\lambda) > 0$. Under (29), the optimal SD mixing weight and risk along the regularization path satisfy

$$\xi^*(\lambda) = -\frac{\lambda}{2} \frac{R'(\lambda)}{D(\lambda)}, \quad R_{\text{sd}}^*(\lambda) = R(\lambda) - \frac{\lambda^2}{4} \frac{(R'(\lambda))^2}{D(\lambda)}.$$

In particular, if $R'(\lambda) \neq 0$, then $R_{\text{sd}}^*(\lambda) < R(\lambda)$ and $\text{sign}(\xi^*(\lambda)) = -\text{sign}(R'(\lambda))$.

The derivative identity (29) surprisingly holds for many resolvent-based ridge estimators, including generalized or Tikhonov ridge (Lemma D.4) and kernel ridge regression (Lemma D.5):

- *Generalized ridge.* For a fixed penalty operator $\Omega \succ 0$ (independent of y), generalized ridge is a linear smoother with

$$s_\lambda^\Omega(x) = X(X^\top X + n\lambda\Omega)^{-1}x, \quad S_\lambda^\Omega = X(X^\top X + n\lambda\Omega)^{-1}X^\top.$$

This includes feature-wise shrinkage (diagonal Ω), graph or Laplacian regularization, spline-type penalties, and other quadratic Tikhonov penalties.

- *Kernel ridge.* For a PSD kernel with kernel matrix $K \in \mathbb{R}^{n \times n}$ and $k_x := (k(x, x_1), \dots, k(x, x_n))^\top$, kernel ridge is a linear smoother with

$$s_\lambda^{\text{kernel}}(x) = (K + n\lambda I_n)^{-1}k_x, \quad S_\lambda^{\text{kernel}} = K(K + n\lambda I_n)^{-1}.$$

This covers standard kernel ridge estimators (e.g., Gaussian and polynomial kernels) and their variants obtained by changing k . See Figure 11 for empirical illustrations with a Gaussian kernel.

Both examples can be shown to satisfy (29) (see Section D.3.3), hence inherit Theorem 5.3 under any squared prediction risk (including out-of-distribution).

6 Conclusion

In this paper, we asked whether self-distillation can provably deliver strict gains ((Q1)), when it can rival an optimally tuned teacher ((Q2)), and how to tune it efficiently ((Q3)). For ridge regression, and more broadly, for a class of resolvent-based ridge smoothers, we provide affirmative and sharp answers.

On the theory side, we give nonasymptotic structural identities that hold conditionally on the observed training data and for *any* squared prediction risk, including out-of-distribution risks. Under a mild nondegeneracy condition, we show that the optimally mixed student *strictly* improves the teacher at every nonstationary λ along the ridge path, and the optimal mixing weight obeys a simple sign rule: it has the *opposite* sign of the ridge-risk derivative, so it can be negative in over-regularized regimes. We then refine these identities in the proportional regime $p/n \rightarrow \gamma \in (0, \infty)$ by deriving deterministic equivalents for the optimal SD risk under anisotropic covariance and deterministic signals, which quantify how SD gains depend on (γ, SNR) and signal–covariance alignment. On the algorithmic side, we propose a consistent one-shot GCV-based tuning method that avoids grid search, repeated retraining, and data splitting, while matching predicted risk curves quite well in several real-data experiments.

Taken together, our results suggest a simple conceptual message: at least in ridge-type problems, optimal self-distillation acts as a cheap and tractable correction for a mis-regularized teacher using the teacher’s own fitted values to move *along* the regularization path in a direction determined by the risk slope. An important direction is to understand how much of this regularization path geometry persists beyond ridge regression, including classification and more general learners, and to develop equally sharp guarantees for fresh- X and other practically motivated distillation variants.

Acknowledgments

We are grateful to Sujay Sanghavi for introducing us to various surprises in self-distillation and for several helpful discussions. We also thank Adel Javanmard for an insightful seminar at the Institute for Foundations of Machine Learning (IFML) at UT Austin and the productive ensuing discussions. Computing support is in part provided by the Texas Advanced Computing Center (TACC).

References

- Sungsoo Ahn, Shell Xu Hu, Andreas Damianou, Neil D. Lawrence, and Zhenwen Dai. Variational information distillation for knowledge transfer. In *Conference on Computer Vision and Pattern Recognition*, 2019. URL <https://arxiv.org/abs/1904.05835>.
- Sina Alemohammad, Josue Casco-Rodriguez, Lorenzo Luzi, Ahmed Imtiaz Humayun, Hossein Babaei, Daniel LeJeune, Ali Siahkoochi, and Richard Baraniuk. Self-consuming generative models go mad. In *International Conference on Learning Representations*, 2023. URL <https://arxiv.org/abs/2307.01850>.
- Jimmy Ba and Rich Caruana. Do deep nets really need to be deep? In *Advances in Neural Information Processing Systems*, 2014. URL <https://arxiv.org/abs/1312.6184>.
- Z. Bai and Jack Silverstein. *Spectral Analysis of Large Dimensional Random Matrices*. Springer, 2010.

- 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. URL <http://dx.doi.org/10.1073/pnas.1907378117>.
- Peter L. Bartlett, Andrea Montanari, and Alexander Rakhlin. Deep learning: a statistical viewpoint. *Acta Numerica*, 30:87–201, 2021. URL <https://arxiv.org/abs/2103.09177>.
- Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, 2019. URL <http://dx.doi.org/10.1073/pnas.1903070116>.
- Pierre C. Bellec, Jin-Hong Du, Takuya Koriyama, Pratik Patil, and Kai Tan. Corrected generalized cross-validation for finite ensembles of penalized estimators. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 87(2):289–318, 2025. URL <https://pratikpatil.io/papers/cgcv.pdf>.
- Cristian Buciluă, Rich Caruana, and Alexandru Niculescu-Mizil. Model compression. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2006. URL <https://www.cs.cornell.edu/~caruana/compression.kdd06.pdf>.
- Collin Burns, Pavel Izmailov, Jan Hendrik Kirchner, Bowen Baker, Leo Gao, Leopold Aschenbrenner, Yining Chen, Adrien Ecoffet, Manas Joglekar, Jan Leike, Ilya Sutskever, and Jeff Wu. Weak-to-strong generalization: Eliciting strong capabilities with weak supervision. arXiv preprint arXiv:2312.09390, 2023. URL <https://arxiv.org/abs/2312.09390>.
- Krisztian Buza. BlogFeedback. UCI Machine Learning Repository, 2014.
- Defang Chen, Jian-Ping Mei, Hailin Zhang, Can Wang, Yan Feng, and Chun Chen. Knowledge distillation with the reused teacher classifier. In *Conference on Computer Vision and Pattern Recognition*, 2022. URL <https://arxiv.org/abs/2203.14001>.
- Guobin Chen, Wongun Choi, Xiang Yu, Tony Han, and Manmohan Chandraker. Learning efficient object detection models with knowledge distillation. In *Advances in Neural Information Processing Systems*, 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/e1e32e235eee1f970470a3a6658dfdd5-Paper.pdf.
- Rudrajit Das and Sujay Sanghavi. Understanding self-distillation in the presence of label noise. In *International Conference on Machine Learning*, 2023. URL <https://arxiv.org/abs/2301.13304>.
- Edgar Dobriban and Yue Sheng. Wonder: Weighted one-shot distributed ridge regression in high dimensions. *Journal of Machine Learning Research*, 21(66):1–52, 2020. URL <https://arxiv.org/abs/1903.09321>.
- Edgar Dobriban and Stefan Wager. High-dimensional asymptotics of prediction: Ridge regression and classification. *Annals of Statistics*, 46(1):247–279, 2018. URL <https://arxiv.org/abs/1507.03003>.
- Elvis Dohmatob, Yunzhen Feng, and Julia Kempe. Model collapse demystified: The case of regression. In *Advances in Neural Information Processing Systems*, 2024. URL <https://arxiv.org/abs/2402.07712>.

- Elvis Dohmatob, Yunzhen Feng, Arjun Subramonian, and Julia Kempe. Strong model collapse. In *International Conference on Learning Representations*, 2025. URL <https://arxiv.org/abs/2410.04840>.
- Bradley Efron. Estimating the error rate of a prediction rule: Improvement on cross-validation. *Journal of the American Statistical Association*, 78(382):316–331, 1983. URL <https://www.jstor.org/stable/2288636>.
- Bradley Efron. How biased is the apparent error rate of a prediction rule? *Journal of the American Statistical Association*, 81(394):461–470, 1986. URL <https://www.jstor.org/stable/2289236>.
- M Emrullah Ildiz, Halil Alperen Gozeten, Ege Onur Taga, Marco Mondelli, and Samet Oymak. High-dimensional analysis of knowledge distillation: Weak-to-strong generalization and scaling laws. In *International Conference on Learning Representations*, 2025. URL <https://arxiv.org/abs/2410.18837>.
- Tommaso Furlanello, Zachary Lipton, Michael Tschannen, Laurent Itti, and Anima Anandkumar. Born again neural networks. In *International Conference on Machine Learning*, 2018. URL <https://arxiv.org/abs/1805.04770>.
- Anvit Garg, Sohom Bhattacharya, and Pragya Sur. Preventing model collapse under over-parametrization: Optimal mixing ratios for interpolation learning and ridge regression. arXiv preprint arXiv:2509.22341, 2025. URL <https://arxiv.org/abs/2509.22341>.
- Matthias Gerstgrasser, Rylan Schaeffer, Apratim Dey, Rafael Rafailov, Henry Sleight, John Hughes, Tomasz Korbak, Rajashree Agrawal, Dhruv Pai, Andrey Gromov, et al. Is model collapse inevitable? breaking the curse of recursion by accumulating real and synthetic data. arXiv preprint arXiv:2404.01413, 2024. URL <https://arxiv.org/abs/2404.01413>.
- Jianping Gou, Baosheng Yu, Stephen J. Maybank, and Dacheng Tao. Knowledge distillation: A survey. *International Journal of Computer Vision*, 129(6):1789–1819, 2021. URL <http://dx.doi.org/10.1007/s11263-021-01453-z>.
- I. S. Gradshteyn and I. M. Ryzhik. *Table of Integrals, Series, and Products*. Elsevier/Academic Press, Amsterdam, seventh edition, 2007.
- Qiyang Han and Xiaocong Xu. The distribution of ridgeless least squares interpolators. *arXiv preprint arXiv:2307.02044*, 2023. URL <https://arxiv.org/abs/2307.02044>.
- Trevor Hastie and Robert Tibshirani. *Generalized Additive Models*. Chapman & Hall, 1990.
- Trevor Hastie, Andrea Montanari, Saharon Rosset, and Ryan J. Tibshirani. Surprises in high-dimensional ridgeless least squares interpolation. *Annals of Statistics*, 50(2):949, 2022. URL <https://arxiv.org/abs/1903.08560>.
- Hengzhi He, Shirong Xu, and Guang Cheng. Golden ratio weighting prevents model collapse. arXiv preprint arXiv:2502.18049, 2025. URL <https://arxiv.org/abs/2502.18049>.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Conference on Computer Vision and Pattern Recognition*, 2016. URL <https://arxiv.org/abs/1512.03385>.

- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. URL <https://arxiv.org/abs/1503.02531>.
- Adel Javanmard, Rudrajit Das, Alessandro Epasto, and Vahab Mirrokni. Self-boost via optimal retraining: An analysis via approximate message passing. In *Advances in Neural Information Processing Systems*, 2025. URL <https://arxiv.org/abs/2505.15195>.
- Guangda Ji and Zhanxing Zhu. Knowledge distillation in wide neural networks: Risk bound, data efficiency and imperfect teacher. In *Advances in Neural Information Processing Systems*, 2020. URL <https://arxiv.org/abs/2010.10090>.
- Antti Knowles and Jun Yin. Anisotropic local laws for random matrices. *Probability Theory and Related Fields*, 169(1):257–352, 2017. URL <https://arxiv.org/abs/1410.3516>.
- Takuya Koriyama, Pratik Patil, Jin-Hong Du, Kai Tan, and Pierre C. Bellec. Precise asymptotics of bagging regularized m-estimators. *arXiv preprint arXiv:2409.15252*, 2024. URL <https://pratikpatil.io/papers/subbagging-asymptotics.pdf>.
- Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. In *Advances in Neural Information Processing Systems*, 2021. URL <https://arxiv.org/abs/2107.07651>.
- Yuncheng Li, Jianchao Yang, Yale Song, Liangliang Cao, Jiebo Luo, and Li-Jia Li. Learning from noisy labels with distillation. In *International Conference on Computer Vision*, 2017. URL <https://arxiv.org/abs/1703.02391>.
- David Lopez-Paz, Léon Bottou, Bernhard Schölkopf, and Vladimir Vapnik. Unifying distillation and privileged information. *arXiv preprint arXiv:1511.03643*, 2015. URL <https://arxiv.org/abs/1511.03643>.
- Theodor Misiakiewicz and Andrea Montanari. Six lectures on linearized neural networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2024(10), 2024. URL <https://arxiv.org/abs/2308.13431>.
- Hossein Mobahi, Mehrdad Farajtabar, and Peter Bartlett. Self-distillation amplifies regularization in Hilbert space. In *Advances in Neural Information Processing Systems*, 2020. URL <https://arxiv.org/abs/2002.05715>.
- Behrad Moniri and Hamed Hassani. On the mechanisms of weak-to-strong generalization: A theoretical perspective. *arXiv preprint arXiv:2505.18346*, 2025. URL <https://arxiv.org/abs/2505.18346>.
- Divyansh Pareek, Simon S. Du, and Sewoong Oh. Understanding the gains from repeated self-distillation. In *Advances in Neural Information Processing Systems*, 2024. URL <https://arxiv.org/abs/2407.04600>.
- Pratik Patil and Jin-Hong Du. Generalized equivalences between subsampling and ridge regularization. In *Advances in Neural Information Processing Systems*, volume 36, pages 78926–78963, 2023. URL <https://pratikpatil.io/papers/generalized-equivalences.pdf>.

- Pratik Patil, Yuting Wei, Alessandro Rinaldo, and Ryan Tibshirani. Uniform consistency of cross-validation estimators for high-dimensional ridge regression. In *International Conference on Artificial Intelligence and Statistics*, 2021. URL <https://pratikpatil.io/papers/ridgecv-combined.pdf>.
- 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*, 2022a. URL <https://pratikpatil.io/papers/risk-monotonization.pdf>.
- Pratik Patil, Alessandro Rinaldo, and Ryan Tibshirani. Estimating functionals of the out-of-sample error distribution in high-dimensional ridge regression. In *International Conference on Artificial Intelligence and Statistics*, 2022b. URL <https://pratikpatil.io/papers/functionals-combined.pdf>.
- 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. URL <https://pratikpatil.io/papers/bagging.pdf>.
- Pratik Patil, Jin-Hong Du, and Ryan Tibshirani. Optimalridge regularization for out-of-distribution prediction. In *International Conference on Machine Learning*, 2024. URL <https://pratikpatil.io/papers/ridge-ood.pdf>.
- Mary Phuong and Christoph Lampert. Towards understanding knowledge distillation. In *International Conference on Machine Learning*, 2019. URL <https://arxiv.org/abs/2105.13093>.
- Kamiar Rahnama Rad and Arian Maleki. A scalable estimate of the out-of-sample prediction error via approximate leave-one-out cross-validation. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 82(4):965–996, 2020. URL <https://academic.oup.com/jrsssb/article/82/4/965/7056054>.
- Michael Redmond. Communities and Crime. UCI Machine Learning Repository, 2002.
- Francisco Rubio and Xavier Mestre. Spectral convergence for a general class of random matrices. *Statistics & Probability Letters*, 81(5):592–602, 2011. URL <https://www.sciencedirect.com/science/article/abs/pii/S0167715211000113>.
- Ilya Shumailov, Zakhar Shumaylov, Yiren Zhao, Nicolas Papernot, Ross Anderson, and Yarin Gal. AI models collapse when trained on recursively generated data. *Nature*, 631(8022):755–759, 2024. URL <https://www.nature.com/articles/s41586-024-07566-y>.
- Charles M. Stein. Estimation of the mean of a multivariate normal distribution. *Annals of Statistics*, pages 1135–1151, 1981. URL <https://projecteuclid.org/journals/annals-of-statistics/volume-9/issue-6/Estimation-of-the-Mean-of-a-Multivariate-Normal-Distribution/10.1214/aos/1176345632.full>.
- Will Stephenson, Zachary Frangella, Madeleine Udell, and Tamara Broderick. Can we globally optimize cross-validation loss? quasiconvexity in ridge regression. *Advances in Neural Information Processing Systems*, 2021.
- Nilesh Tripuraneni, Ben Adlam, and Jeffrey Pennington. Covariate shift in high-dimensional ran-

dom feature regression. *arXiv preprint arXiv:2111.08234*, 2021. URL <https://arxiv.org/abs/2111.08234>.

Saverio Vito. Air Quality. UCI Machine Learning Repository, 2008.

Alexander Wei, Wei Hu, and Jacob Steinhardt. More than a toy: Random matrix models predict how real-world neural representations generalize. In *International Conference on Machine Learning*, 2022. URL <https://arxiv.org/abs/2203.06176>.

Linfeng Zhang, Chenglong Bao, and Kaisheng Ma. Self-distillation: Towards efficient and compact neural networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(8):4388–4403, 2021. URL <https://ieeexplore.ieee.org/document/9381661>.

This supplement serves as a companion to the paper titled “Optimal Unconstrained Self-Distillation in Ridge Regression: Strict Improvements, Precise Asymptotics, and One-Shot Tuning”. We provide an outline of the supplement in Table 3 and a summary of general notation used throughout the paper in Table 4.

Organization

Table 3: Outline of the supplement.

Section	Subsection	Purpose
Proofs in Section 2		
Section A	Section A.1	Preliminaries
	Section A.2	Proof of Equation (4)
	Section A.3	Proof of Proposition 2.1
	Section A.4	Proof of Theorem 2.2
	Section A.5	Proof of Proposition 2.3
Proofs in Section 3		
Section B	Section B.1	Preliminaries
	Section B.2	Helper results (concentration components) for the proof of Theorem 3.1
	Section B.3	Helper results (deterministic equivalents) for the proof of Theorem 3.1
	Section B.4	Proof of Theorem 3.1
	Section B.5	Proofs of helper results (concentration components) for the proof of Theorem 3.1
	Section B.6	Proofs of helper results (deterministic equivalents) for the proof of Theorem 3.1
	Section B.7	Proof of Corollary 3.2
	Section B.8	Proof of Proposition 3.3
Proofs in Section 4		
Section C	Section C.1	Preliminaries
	Section C.2	Helper results (risk and correlation components consistency) for the proof of Theorem 4.1
	Section C.3	Proof of Theorem 4.1
	Section C.4	Technical lemmas
Proofs in Section 5		
Section D	Section D.1	Proof of Proposition 5.1 and other details in Section 5.1
	Section D.2	Proof of Theorem 5.2 and other details in Section 5.2
	Section D.3	Proof of Theorem 5.3 and other details in Section 5.3
Additional experiments		
Section E	Section E.1	Additional experiments on CIFAR datasets
	Section E.2	Illustrations for related works comparison in Section 1.2
	Section E.3	Additional illustrations in Section 2
	Section E.4	Additional illustrations in Section 3.2
	Section E.5	Additional illustrations in Section 3.3
Experimental details		
Section F	Section F.1	Real-world regression tasks
	Section F.2	Pretrained ResNet features on CIFAR datasets
	Section F.3	Synthetic asymptotic experiments

Notation

Table 4: Summary of general notation used throughout the paper.

Notation	Description
Typography	
Lowercase (e.g., x)	Scalars or vectors
Uppercase (e.g., X)	Matrices or linear operators
Calligraphic (e.g., \mathcal{D} , \mathcal{R})	Sets, σ -fields, events or certain limiting functions
Blackboard bold (e.g., \mathbb{R} , \mathbb{N})	Standard number systems
Analysis	
\mathbb{Z} , \mathbb{N} , \mathbb{R} , $\mathbb{R}_{\geq 0}$, $\overline{\mathbb{R}}$	Integers, positive integers, reals, nonnegative reals, extended reals
(a, b, c) , $\{a, b, c\}$	Ordered tuple and (unordered) set
$[n]$	Set $\{1, \dots, n\}$ for a positive integer n
$x \wedge y$, $x \vee y$	$\min\{x, y\}$ and $\max\{x, y\}$ for real numbers x, y
$\mathbb{1}_A$	Indicator random variable associated with event or set A
$\text{sign}(x)$	Sign of a real number x
C^∞	Function class of infinitely differentiable functions
Linear algebra	
$\text{tr}(A)$, $\overline{\text{tr}}(A)$, $\det(A)$	Trace, normalized 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
C^\dagger	Moore-Penrose inverse of a general rectangular matrix C
$\text{diag}(d_1, \dots, d_p)$	Diagonal matrix with diagonal entries d_1, \dots, d_p
$U^{1/2}$	Principal square root of a positive semidefinite matrix $U \succeq 0$
$f(W)$	Functional calculus for positive semidefinite matrix W (apply f to eigenvalues of W)
I , $\mathbf{1}$, $\mathbf{0}$	The identity matrix, the all-ones vector, the all-zeros vector
Inner products and norms	
$\langle u, v \rangle$	Euclidean inner product (or another inner product when specified)
$\ x\ _2$ (or simply $\ x\ $)	Euclidean norm of a vector x
$\ x\ _q$	ℓ_q norm of a vector ($q \geq 1$)
$\ x\ _A := \sqrt{x^\top A x}$	A -seminorm for $A \succeq 0$
$\ A\ _{\text{op}}$	Operator/spectral norm of a matrix A
$\ A\ _F$	Frobenius norm of a matrix A
$\ A\ _{\text{tr}}$	Trace/nuclear norm (sum of singular values)
$\ f\ _{L_q}$	L_q norm of a function f under the relevant measure ($q \geq 1$)
Probability	
$\mathbb{P}(\cdot)$, $\mathbb{E}[\cdot]$	Probability and expectation
$\mathbb{E}[\cdot \mid \mathcal{G}]$	Conditional expectation given σ -field \mathcal{G} (or given data, depending on context)
$\text{Var}(\cdot)$, $\text{Cov}(\cdot, \cdot)$	Variance and covariance
$X \sim \mathcal{N}(\mu, \Sigma)$	Gaussian random vector with mean μ and covariance Σ
$\stackrel{\text{d}}{=}$	Equality in distribution
Orders and asymptotics	
$X = \mathcal{O}_\alpha(Y)$, $X \lesssim_\alpha Y$	Deterministic upper bounds with constant possibly depending on parameter α
$o(\cdot)$, $\mathcal{O}(\cdot)$	Deterministic little- o and big- \mathcal{O}
$o_{\mathbb{P}}(\cdot)$, $\mathcal{O}_{\mathbb{P}}(\cdot)$	Probabilistic little- o and big- \mathcal{O}
\rightarrow , $\xrightarrow{\text{P}}$, $\xrightarrow{\text{a.s.}}$, $\xrightarrow{\text{d}}$	Convergence, in probability, almost surely, in distribution
\asymp	Asymptotic equivalence (see Definition 1 for more details)
C, C', c, c'	Generic positive constants (may change from line to line)

A Proofs in Section 2

A.1 Preliminaries

Define the empirical covariance matrix and empirical resolvent as

$$\widehat{\Sigma} := \frac{1}{n} X^\top X \in \mathbb{R}^{p \times p}, \quad \text{and} \quad Q_\lambda := (\widehat{\Sigma} + \lambda I_p)^{-1} \in \mathbb{R}^{p \times p}, \quad (30)$$

respectively. For any label vector $u \in \mathbb{R}^n$, define the ridge coefficient map and predictor as

$$\beta_\lambda(u) := Q_\lambda \frac{X^\top u}{n}, \quad \text{and} \quad f_{\lambda,u}(x) := x^\top \beta_\lambda(u), \quad (31)$$

respectively. In particular, the teacher predictor is $f_\lambda := f_{\lambda,y}$ with fitted values $\widehat{y}_\lambda := f_\lambda(X) = X\beta_\lambda(y) \in \mathbb{R}^n$. Note that $f_{\lambda,u}$ is linear in u .

The PD predictor $f_{\text{pd},\lambda}$ is the ridge predictor trained on (X, \widehat{y}_λ) with the same penalty λ , i.e., $f_{\text{pd},\lambda} := f_{\lambda, \widehat{y}_\lambda}$. Its coefficient vector admits the closed form

$$\beta_{\text{pd},\lambda} := \beta_\lambda(\widehat{y}_\lambda) = Q_\lambda \frac{X^\top (X\beta_\lambda(y))}{n} = Q_\lambda \widehat{\Sigma} \beta_\lambda(y). \quad (32)$$

Since $\widehat{\Sigma}$ and Q_λ commute (both are polynomials in $\widehat{\Sigma}$), we have $\widehat{\Sigma}Q_\lambda = I_p - \lambda Q_\lambda$, and therefore

$$\beta_\lambda(y) - \beta_{\text{pd},\lambda} = (I_p - Q_\lambda \widehat{\Sigma}) \beta_\lambda(y) = \lambda Q_\lambda \beta_\lambda(y). \quad (33)$$

A.2 Proof of Equation (4)

Lemma A.1 (SD ridge representation). Fix $\lambda > 0$ and $\xi \in \mathbb{R}$. Let $\widehat{y}_\lambda = f_\lambda(X)$ and define the mixed label vector

$$y^{(1)}(\xi) := (1 - \xi)y + \xi \widehat{y}_\lambda \in \mathbb{R}^n. \quad (34)$$

Then the SD predictor (3) coincides with ridge regression trained on $(X, y^{(1)}(\xi))$ with penalty λ , i.e., $f_{\text{sd},\lambda,\xi} = f_{\lambda, y^{(1)}(\xi)}$.

Proof. Expanding the SD objective (and dropping terms independent of β), we obtain that

$$\begin{aligned} & (1 - \xi) \frac{1}{n} \|y - X\beta\|_2^2 + \xi \frac{1}{n} \|\widehat{y}_\lambda - X\beta\|_2^2 + \lambda \|\beta\|_2^2 \\ &= \frac{1}{n} \|X\beta\|_2^2 - \frac{2}{n} \beta^\top X^\top ((1 - \xi)y + \xi \widehat{y}_\lambda) + \lambda \|\beta\|_2^2 + \text{const}, \end{aligned}$$

where const is a term that does not depend on β . This is exactly the ridge objective with response $y^{(1)}(\xi)$ as defined in (34). \square

Proof of Equation (4). By Lemma A.1, $f_{\text{sd},\lambda,\xi} = f_{\lambda, (1-\xi)y + \xi \widehat{y}_\lambda}$. Using linearity of $f_{\lambda,u}$ in u (cf. (31)),

$$f_{\text{sd},\lambda,\xi} = (1 - \xi)f_{\lambda,y} + \xi f_{\lambda, \widehat{y}_\lambda} = (1 - \xi)f_\lambda + \xi f_{\text{pd},\lambda},$$

which is (4). \square

A.3 Proof of Proposition 2.1

Throughout, $\lambda > 0$ is fixed and we abbreviate $f := f_\lambda$ and $g := f_{\text{pd},\lambda}$. Write

$$R := R(f), \quad R_{\text{pd}} := R(g), \quad C := \mathbb{E}[(y_0 - f(x_0))(y_0 - g(x_0)) \mid \mathcal{D}],$$

and define the nonnegative quantity y

$$D := R + R_{\text{pd}} - 2C = \mathbb{E}[(f(x_0) - g(x_0))^2 \mid \mathcal{D}] \geq 0.$$

Proof of Proposition 2.1. By (4), the SD predictor with mixing weight ξ is $f_{\text{sd}}(\xi) = (1 - \xi)f + \xi g$. Let the residuals be $r_f := y_0 - f(x_0)$ and $r_g := y_0 - g(x_0)$. Then

$$y_0 - f_{\text{sd}}(\xi)(x_0) = (1 - \xi)r_f + \xi r_g,$$

and thus

$$\begin{aligned} R_{\text{sd}}(\lambda, \xi) &= \mathbb{E}[(1 - \xi)r_f + \xi r_g]^2 \mid \mathcal{D}] \\ &= (1 - \xi)^2 \mathbb{E}[r_f^2 \mid \mathcal{D}] + \xi^2 \mathbb{E}[r_g^2 \mid \mathcal{D}] + 2\xi(1 - \xi) \mathbb{E}[r_f r_g \mid \mathcal{D}] \\ &= (1 - \xi)^2 R + \xi^2 R_{\text{pd}} + 2\xi(1 - \xi)C \\ &= R - 2\xi(R - C) + \xi^2 D. \end{aligned}$$

If $D > 0$, this is a strictly convex quadratic in ξ with a unique minimizer given by

$$0 = \partial_\xi R_{\text{sd}}(\lambda, \xi) = -2(R - C) + 2\xi D \implies \xi^* = \frac{R - C}{D}.$$

Substituting back yields

$$R_{\text{sd}}^* = R - \frac{(R - C)^2}{D}.$$

If $D = 0$, then $f(x_0) = g(x_0)$ a.s. under the test distribution, hence $R_{\text{sd}}(\lambda, \xi) \equiv R$ for all ξ and every ξ is optimal with $R_{\text{sd}}^* = R$. \square

A.4 Proof of Theorem 2.2

We first establish a derivative identity relating the PD predictor to the derivative of the teacher along the ridge path.

Lemma A.2 (Resolvent derivative rule). For a fixed $\lambda > 0$ and symmetric $A \succeq 0$, let $Q_\lambda := (A + \lambda I)^{-1}$. Then,

$$\partial_\lambda Q_\lambda = -Q_\lambda^2.$$

Proof. Differentiate $(A + \lambda I)Q_\lambda = I$ and left-multiply by Q_λ . \square

Lemma A.3 (Derivative identity for ridge versus pure-distillation). For any $\lambda > 0$ and $x \in \mathbb{R}^p$,

$$f_\lambda(x) - f_{\text{pd},\lambda}(x) = -\lambda \partial_\lambda f_\lambda(x). \quad (35)$$

Proof. From the definitions, $\beta_\lambda(y) = Q_\lambda X^\top y/n$ and $\beta_{\text{pd},\lambda} = Q_\lambda \widehat{\Sigma} \beta_\lambda(y)$; see (31) and (32). By Lemma A.2,

$$\partial_\lambda \beta_\lambda(y) = (\partial_\lambda Q_\lambda) \frac{X^\top y}{n} = -Q_\lambda^2 \frac{X^\top y}{n} = -Q_\lambda \beta_\lambda(y).$$

Thus $-\lambda \partial_\lambda \beta_\lambda(y) = \lambda Q_\lambda \beta_\lambda(y)$. Using (33), $\lambda Q_\lambda \beta_\lambda(y) = \beta_\lambda(y) - \beta_{\text{pd},\lambda}$. Multiplying by x^\top yields (35). \square

Next we relate $R(\lambda) - C(\lambda)$ to $R'(\lambda)$. Before we proceed, we justify that $R'(\lambda)$ is indeed well-defined under our setup.

Lemma A.4 (Smoothness of ridge predictions and risk). Fix the training data $\mathcal{D} = (X, y)$ and assume $\mathbb{E}[y_0^2 \mid \mathcal{D}] < \infty$ and $\mathbb{E}[\|x_0\|_2^2 \mid \mathcal{D}] < \infty$. Then for every $\lambda > 0$, the map $\lambda \mapsto f_\lambda(x_0)$ is C^∞ and the risk $R(\lambda) = \mathbb{E}[(y_0 - f_\lambda(x_0))^2 \mid \mathcal{D}]$ is C^∞ on $(0, \infty)$. Moreover,

$$R'(\lambda) = -2 \mathbb{E}[(y_0 - f_\lambda(x_0)) \partial_\lambda f_\lambda(x_0) \mid \mathcal{D}], \quad (36)$$

where $\partial_\lambda f_\lambda(x) = -x^\top Q_\lambda \beta_\lambda(y)$ with $Q_\lambda = (\widehat{\Sigma} + \lambda I)^{-1}$.

Proof. Fix $\lambda_0 > 0$ and consider $\lambda \in [\lambda_0/2, 3\lambda_0/2]$. Then $\|Q_\lambda\|_{\text{op}} \leq 2/\lambda_0$, and with $b := X^\top y/n$ we have $\beta_\lambda = Q_\lambda b$ and $\partial_\lambda \beta_\lambda = -Q_\lambda \beta_\lambda$. Hence there exist constants (depending on \mathcal{D} and λ_0) such that uniformly over this interval,

$$|f_\lambda(x_0)| \leq c_1 \|x_0\|_2, \quad |\partial_\lambda f_\lambda(x_0)| \leq c_2 \|x_0\|_2.$$

Therefore,

$$|\partial_\lambda (y_0 - f_\lambda(x_0))^2| = 2|y_0 - f_\lambda(x_0)| |\partial_\lambda f_\lambda(x_0)| \leq c_3 (|y_0| \|x_0\|_2 + \|x_0\|_2^2),$$

whose conditional expectation is finite by Cauchy–Schwarz and the assumed second-moment bounds. Dominated convergence justifies differentiating under $\mathbb{E}[\cdot \mid \mathcal{D}]$, yielding (36). Higher derivatives follow similarly since $f_\lambda(x_0)$ is a rational (hence analytic) function of λ for $\lambda > 0$. \square

Lemma A.5 (Derivative identity for $R(\lambda) - C(\lambda)$). For any $\lambda > 0$,

$$R(\lambda) - C(\lambda) = -\frac{\lambda}{2} R'(\lambda), \quad (37)$$

where $R(\lambda) = R(f_\lambda)$ and $C(\lambda) = \mathbb{E}[(y_0 - f_\lambda(x_0))(y_0 - f_{\text{pd},\lambda}(x_0)) \mid \mathcal{D}]$.

Proof. Write $r_\lambda := y_0 - f_\lambda(x_0)$. First note that

$$R(\lambda) - C(\lambda) = \mathbb{E}[r_\lambda^2 - r_\lambda(y_0 - f_{\text{pd},\lambda}(x_0)) \mid \mathcal{D}] = \mathbb{E}[r_\lambda (f_{\text{pd},\lambda}(x_0) - f_\lambda(x_0)) \mid \mathcal{D}].$$

Using Lemma A.3, $f_{\text{pd},\lambda}(x_0) - f_\lambda(x_0) = \lambda \partial_\lambda f_\lambda(x_0)$, hence

$$R(\lambda) - C(\lambda) = \lambda \mathbb{E}[r_\lambda \partial_\lambda f_\lambda(x_0) \mid \mathcal{D}]. \quad (38)$$

On the other hand, from Lemma A.4, we have

$$R'(\lambda) = -2 \mathbb{E}[r_\lambda \partial_\lambda f_\lambda(x_0) \mid \mathcal{D}].$$

Combining with (38) yields (37). \square

Proof of Theorem 2.2. Let $D(\lambda) := R(\lambda) + R_{\text{pd}}(\lambda) - 2C(\lambda)$. When $D(\lambda) > 0$, Proposition 2.1 gives

$$\xi^*(\lambda) = \frac{R(\lambda) - C(\lambda)}{D(\lambda)}, \quad R_{\text{sd}}^*(\lambda) = R(\lambda) - \frac{(R(\lambda) - C(\lambda))^2}{D(\lambda)}.$$

Apply Lemma A.5 to substitute $R(\lambda) - C(\lambda) = -(\lambda/2)R'(\lambda)$, obtaining

$$\xi^*(\lambda) = -\frac{\lambda R'(\lambda)}{2 D(\lambda)}, \quad R_{\text{sd}}^*(\lambda) = R(\lambda) - \frac{\lambda^2 (R'(\lambda))^2}{4 D(\lambda)}.$$

If $R'(\lambda) \neq 0$, then the subtracted term is strictly positive (since $D(\lambda) > 0$), hence $R_{\text{sd}}^*(\lambda) < R(\lambda)$. Finally, $\text{sign}(\xi^*(\lambda)) = -\text{sign}(R'(\lambda))$ because $\lambda > 0$ and $D(\lambda) > 0$. \square

A.5 Proof of Proposition 2.3

Proof of Proposition 2.3. Define

$$g(\lambda) := R(\lambda) - C(\lambda), \quad \text{and} \quad D(\lambda) := R(\lambda) + R_{\text{pd}}(\lambda) - 2C(\lambda).$$

For $D(\lambda) > 0$, Proposition 2.1 gives

$$R_{\text{sd}}^*(\lambda) = R(\lambda) - \frac{g(\lambda)^2}{D(\lambda)}. \quad (39)$$

By Lemma A.5, $g(\lambda) = -(\lambda/2)R'(\lambda)$, hence at λ^* ,

$$g(\lambda^*) = 0, \quad g'(\lambda^*) = -\frac{\lambda^*}{2}R''(\lambda^*). \quad (40)$$

Differentiate (39) twice and evaluate at λ^* . All terms involving $g(\lambda^*)$ vanish, yielding

$$R_{\text{sd}}^{*''}(\lambda^*) = R''(\lambda^*) - \frac{2(g'(\lambda^*))^2}{D(\lambda^*)}. \quad (41)$$

Substituting (40) into (41) gives

$$R_{\text{sd}}^{*''}(\lambda^*) = R''(\lambda^*) - \frac{\lambda^{*2} (R''(\lambda^*))^2}{2 D(\lambda^*)}.$$

Therefore, if $D(\lambda^*) < \frac{\lambda^{*2}}{2}R''(\lambda^*)$, then $R_{\text{sd}}^{*''}(\lambda^*) < 0$, so λ^* cannot be a local minimizer of R_{sd}^* . Since $g(\lambda^*) = 0$, (39) gives $R_{\text{sd}}^*(\lambda^*) = R(\lambda^*) = \min_{\lambda \geq 0} R(\lambda)$. Hence there exists $\lambda > 0$ such that $R_{\text{sd}}^*(\lambda) < R_{\text{sd}}^*(\lambda^*) = \min_{\lambda \geq 0} R(\lambda)$, which implies $\min_{\lambda > 0} R_{\text{sd}}^*(\lambda) < \min_{\lambda \geq 0} R(\lambda)$. \square

B Proofs in Section 3

B.1 Preliminaries

Recall the random-design model in Assumption A. The design matrix $X \in \mathbb{R}^{n \times p}$ has rows $x_i^\top = z_i^\top \Sigma^{1/2}$, where $z_i \in \mathbb{R}^p$ has i.i.d. entries with mean 0, variance 1, and uniformly bounded $(4 + \mu)$ -th moment for some $\mu > 0$, and $\Sigma \in \mathbb{R}^{p \times p}$ is deterministic positive definite with spectrum uniformly bounded away from 0 and ∞ . The response y has mean 0 and uniformly bounded $(4 + \nu)$ -th moment.

Define the population L_2 -projection coefficient, residual, and residual variance by

$$\beta := \mathbb{E}[xx^\top]^{-1}\mathbb{E}[xy] = \Sigma^{-1}\mathbb{E}[xy], \quad \varepsilon := y - x^\top\beta, \quad \sigma^2 := \mathbb{E}[\varepsilon^2].$$

Then $\mathbb{E}[x\varepsilon] = 0$ (equivalently $\mathbb{E}[z\varepsilon] = 0$) and $\mathbb{E}[\varepsilon] = 0$. Throughout, we consider the in-distribution prediction setting, i.e., the test pair (x_0, y_0) is an independent copy of (x, y) , so $y_0 = x_0^\top\beta + \varepsilon_0$ with $\mathbb{E}[x_0\varepsilon_0] = 0$ and $\mathbb{E}[\varepsilon_0^2] = \sigma^2$.

Write the sample covariance and ridge resolvent as

$$\widehat{\Sigma} := \frac{1}{n}X^\top X, \quad Q_\lambda := (\widehat{\Sigma} + \lambda I_p)^{-1}, \quad \lambda > 0,$$

and set $\overline{\text{tr}}(A) := \text{tr}(A)/p$.

Define the ridge estimator and its pure-distilled transform by

$$\beta_\lambda := Q_\lambda \frac{X^\top y}{n}, \quad M_\lambda := \widehat{\Sigma}(\widehat{\Sigma} + \lambda I_p)^{-1} = \widehat{\Sigma}Q_\lambda = I_p - \lambda Q_\lambda, \quad \beta_{\text{pd},\lambda} := M_\lambda \beta_\lambda.$$

We use the Σ -inner product and norm

$$\langle u, v \rangle_\Sigma := u^\top \Sigma v, \quad \|u\|_\Sigma^2 := \langle u, u \rangle_\Sigma.$$

For any estimator $\widehat{\beta} = \widehat{\beta}(X, y)$, the in-distribution prediction MSE decomposes as

$$\mathbb{E}[(y_0 - x_0^\top \widehat{\beta})^2 \mid X, y] = \|\widehat{\beta} - \beta\|_\Sigma^2 + \sigma^2,$$

and similarly, for two estimators $\widehat{\beta}, \widetilde{\beta}$,

$$\mathbb{E}[(y_0 - x_0^\top \widehat{\beta})(y_0 - x_0^\top \widetilde{\beta}) \mid X, y] = \langle \widehat{\beta} - \beta, \widetilde{\beta} - \beta \rangle_\Sigma + \sigma^2.$$

Thus, adding σ^2 to the teacher/PD risks and correlation term does not affect the optimal mixing weight (constants cancel), and it only adds σ^2 to the optimal mixed prediction MSE.

Accordingly, for a fixed training dataset (X, y) , we define the *excess* scalars

$$\overline{R}(\lambda) := \|\beta_\lambda - \beta\|_\Sigma^2, \quad \overline{R}_{\text{pd}}(\lambda) := \|\beta_{\text{pd},\lambda} - \beta\|_\Sigma^2, \quad \overline{C}(\lambda) := \langle \beta_\lambda - \beta, \beta_{\text{pd},\lambda} - \beta \rangle_\Sigma,$$

and note that the full in-distribution quantities in the main paper are $\overline{R}(\lambda) + \sigma^2$, $\overline{R}_{\text{pd}}(\lambda) + \sigma^2$, and $\overline{C}(\lambda) + \sigma^2$.

Recall from Section 2 that the self-distilled estimator is

$$\beta_{\text{sd},\lambda}(\xi) := (1 - \xi)\beta_\lambda + \xi\beta_{\text{pd},\lambda},$$

and its excess risk is

$$\|\beta_{\text{sd},\lambda}(\xi) - \beta\|_\Sigma^2 = (1 - \xi)^2 \overline{R}(\lambda) + \xi^2 \overline{R}_{\text{pd}}(\lambda) + 2\xi(1 - \xi)\overline{C}(\lambda).$$

Therefore the optimal mixing weight and optimal excess SD risk are given by

$$\xi^*(\lambda) = \frac{\overline{R}(\lambda) - \overline{C}(\lambda)}{\overline{R}(\lambda) + \overline{R}_{\text{pd}}(\lambda) - 2\overline{C}(\lambda)}, \quad R_{\text{sd}}^*(\lambda) = \sigma^2 + \overline{R}(\lambda) - \frac{(\overline{R}(\lambda) - \overline{C}(\lambda))^2}{\overline{R}(\lambda) + \overline{R}_{\text{pd}}(\lambda) - 2\overline{C}(\lambda)}, \quad (42)$$

whenever the denominator is positive. Thus, to prove Theorem 3.1, it suffices to show that $\overline{R}(\lambda), \overline{C}(\lambda), \overline{R}_{\text{pd}}(\lambda)$ converge in probability to deterministic limits $\mathcal{R}(\lambda) - \sigma^2, \mathcal{C}(\lambda) - \sigma^2, \mathcal{R}_{\text{pd}}(\lambda) - \sigma^2$ (i.e., the excess parts of the main-paper limits), and then apply the continuous mapping theorem to (42) (and finally add σ^2 back to match the full risks stated in Theorem 3.1).

B.2 Helpers Results (Concentration Components) for the Proof of Theorem 3.1

Define the residual vector $\varepsilon := (\varepsilon_1, \dots, \varepsilon_n)^\top$ with $\varepsilon_i := y_i - x_i^\top \beta$, so that $y = X\beta + \varepsilon$ holds identically.

Lemma B.1 (Exact decomposition of coefficient errors). For any $\lambda > 0$,

$$\begin{aligned}\beta_\lambda - \beta &= -\lambda Q_\lambda \beta + Q_\lambda X^\top \varepsilon / n, \\ \beta_{\text{pd},\lambda} - \beta &= (-2\lambda Q_\lambda + \lambda^2 Q_\lambda^2) \beta + (I_p - \lambda Q_\lambda) Q_\lambda X^\top \varepsilon / n.\end{aligned}$$

Proof. Using $y = X\beta + \varepsilon$ and $\widehat{\Sigma} = X^\top X / n$,

$$\beta_\lambda = Q_\lambda (\widehat{\Sigma} \beta + X^\top \varepsilon / n) = (I_p - \lambda Q_\lambda) \beta + Q_\lambda X^\top \varepsilon / n,$$

so $\beta_\lambda - \beta = -\lambda Q_\lambda \beta + Q_\lambda X^\top \varepsilon / n$. Then $\beta_{\text{pd},\lambda} = (I_p - \lambda Q_\lambda) \beta_\lambda$ and expanding yields the second identity. \square

Define

$$Q_2(\lambda) := \lambda^2 Q_\lambda \Sigma Q_\lambda, \quad Q_3(\lambda) := \lambda^3 (Q_\lambda \Sigma Q_\lambda^2 + Q_\lambda^2 \Sigma Q_\lambda), \quad Q_4(\lambda) := \lambda^4 Q_\lambda^2 \Sigma Q_\lambda^2, \quad (43)$$

and

$$U_k(\lambda) := \frac{1}{n} \text{tr}(\Sigma \widehat{\Sigma} Q_\lambda^k), \quad k \in \{2, 3, 4\}. \quad (44)$$

Lemma B.2 (Expansion of $\bar{R}(\lambda), \bar{C}(\lambda), R_{\text{pd}}(\lambda)$). Let $\lambda > 0$ be fixed. We have

$$\bar{R}(\lambda) = \beta^\top Q_2(\lambda) \beta + \frac{1}{n^2} \varepsilon^\top X Q_\lambda \Sigma Q_\lambda X^\top \varepsilon - \frac{2\lambda}{n} \beta^\top Q_\lambda \Sigma Q_\lambda X^\top \varepsilon, \quad (45)$$

$$\bar{C}(\lambda) = \beta^\top \left(2Q_2(\lambda) - \frac{1}{2} Q_3(\lambda) \right) \beta + \frac{1}{n^2} \varepsilon^\top X Q_\lambda \Sigma (I_p - \lambda Q_\lambda) Q_\lambda X^\top \varepsilon + \text{Lin}_C(\lambda), \quad (46)$$

$$\begin{aligned}\bar{R}_{\text{pd}}(\lambda) &= \beta^\top \left(4Q_2(\lambda) - 2Q_3(\lambda) + Q_4(\lambda) \right) \beta + \frac{1}{n^2} \varepsilon^\top X Q_\lambda (I_p - \lambda Q_\lambda) \Sigma (I_p - \lambda Q_\lambda) Q_\lambda X^\top \varepsilon \\ &\quad + \text{Lin}_{\text{pd}}(\lambda),\end{aligned} \quad (47)$$

where $\text{Lin}_C(\lambda)$ and $\text{Lin}_{\text{pd}}(\lambda)$ are terms linear in ε . Moreover, the quadratic ε -terms in (46)–(47) can be rewritten as

$$\begin{aligned}\frac{1}{n^2} \varepsilon^\top X Q_\lambda \Sigma (I_p - \lambda Q_\lambda) Q_\lambda X^\top \varepsilon &= \frac{1}{n^2} \varepsilon^\top X Q_\lambda \Sigma Q_\lambda X^\top \varepsilon - \lambda \cdot \frac{1}{n^2} \varepsilon^\top X Q_\lambda \Sigma Q_\lambda^2 X^\top \varepsilon, \quad (48) \\ \frac{1}{n^2} \varepsilon^\top X Q_\lambda (I_p - \lambda Q_\lambda) \Sigma (I_p - \lambda Q_\lambda) Q_\lambda X^\top \varepsilon &= \frac{1}{n^2} \varepsilon^\top X Q_\lambda \Sigma Q_\lambda X^\top \varepsilon - 2\lambda \cdot \frac{1}{n^2} \varepsilon^\top X Q_\lambda \Sigma Q_\lambda^2 X^\top \varepsilon \\ &\quad + \lambda^2 \cdot \frac{1}{n^2} \varepsilon^\top X Q_\lambda^2 \Sigma Q_\lambda^2 X^\top \varepsilon.\end{aligned} \quad (49)$$

Proof. Plug the decompositions from Lemma B.1 into the definitions of $R = \|e\|_\Sigma^2$, $C = \langle e, \tilde{e} \rangle_\Sigma$, $R_{\text{pd}} = \|\tilde{e}\|_\Sigma^2$ and expand. The bias-only terms give the stated Q_2, Q_3, Q_4 combinations (using symmetry to replace scalars of the form $\beta^\top Q_\lambda \Sigma Q_\lambda^2 \beta$ by $\frac{1}{2} \beta^\top (Q_\lambda \Sigma Q_\lambda^2 + Q_\lambda^2 \Sigma Q_\lambda) \beta$). The terms quadratic in ε are the displayed quadratic forms. The linear terms $\text{Lin}_C, \text{Lin}_{\text{pd}}$ collect the bias- ε cross-terms. Finally, (48)–(49) follow by expanding $(I_p - \lambda Q_\lambda)$ and using symmetry of Q_λ and Σ . \square

In the misspecified setting, we cannot condition on X and drop the bias–noise cross-terms. Instead we use leave-one-out concentration for ridge-type linear and quadratic forms.

Lemma B.3 (Response-noise concentration for ridge-type bilinear/quadratic forms). Assume Assumption A. Fix $\lambda > 0$ and let $\varepsilon_i := y_i - x_i^\top \beta$, so that $\mathbb{E}[z_i \varepsilon_i] = 0$ and $\mathbb{E}[\varepsilon_i^2] = \sigma^2$. Then, as $n, p \rightarrow \infty$ with $p/n \rightarrow \gamma$:

1. All terms linear in ε appearing in (45)–(47) satisfy $\text{Lin}_\bullet(\lambda) = o_{\mathbb{P}}(1)$.
2. The quadratic ε -forms satisfy

$$\begin{aligned}\frac{1}{n^2} \varepsilon^\top X Q_\lambda \Sigma Q_\lambda X^\top \varepsilon &= \sigma^2 U_2(\lambda) + o_{\mathbb{P}}(1), \\ \frac{1}{n^2} \varepsilon^\top X Q_\lambda \Sigma Q_\lambda^2 X^\top \varepsilon &= \sigma^2 U_3(\lambda) + o_{\mathbb{P}}(1), \\ \frac{1}{n^2} \varepsilon^\top X Q_\lambda^2 \Sigma Q_\lambda^2 X^\top \varepsilon &= \sigma^2 U_4(\lambda) + o_{\mathbb{P}}(1),\end{aligned}$$

with $U_k(\lambda)$ defined in (44).

Proof outline. This follows from leave-one-out expansions in Lemmas D.2 (linear forms) and D.3 (off-diagonal quadratic forms) of Patil and Du (2023). At a high level: (i) each linear term is a normalized sum $\frac{1}{n} \sum_{i=1}^n \alpha_i \varepsilon_i$ where α_i is a leave-one-out coefficient built from Q_λ and x_i ; Lemma D.2 gives $o_{\mathbb{P}}(1)$ under $\mathbb{E}[z_i \varepsilon_i] = 0$ and bounded moments; (ii) each quadratic form is $\frac{1}{n^2} \sum_{i,j} \varepsilon_i \varepsilon_j K_{ij}$ for a ridge-type kernel K . Lemma D.3 controls the off-diagonal sum $\sum_{i \neq j}$ (giving $o_{\mathbb{P}}(1)$), while the diagonal sum reduces to $\sigma^2 \frac{1}{n} \text{tr}(\cdot)$ up to $o_{\mathbb{P}}(1)$ by the same leave-one-out decoupling. A detailed proof is given in Section B.5. \square

Combining Lemmas B.2 and B.3 yields

$$\bar{R}(\lambda) = \beta^\top Q_2(\lambda) \beta + \sigma^2 U_2(\lambda) + o_{\mathbb{P}}(1), \quad (50)$$

$$\bar{C}(\lambda) = \beta^\top \left(2Q_2(\lambda) - \frac{1}{2} Q_3(\lambda) \right) \beta + \sigma^2 (U_2(\lambda) - \lambda U_3(\lambda)) + o_{\mathbb{P}}(1), \quad (51)$$

$$\bar{R}_{\text{pd}}(\lambda) = \beta^\top \left(4Q_2(\lambda) - 2Q_3(\lambda) + Q_4(\lambda) \right) \beta + \sigma^2 (U_2(\lambda) - 2\lambda U_3(\lambda) + \lambda^2 U_4(\lambda)) + o_{\mathbb{P}}(1). \quad (52)$$

B.3 Helper Results (Deterministic Equivalents) for the Proof of Theorem 3.1

Fix $\lambda > 0$ and define $\kappa = \kappa(\lambda) > 0$ as the unique solution of

$$\kappa = \lambda + \gamma \kappa \bar{\text{tr}}(\Sigma(\Sigma + \kappa I_p)^{-1}). \quad (53)$$

Let $G := (\Sigma + \kappa I_p)^{-1}$ and define

$$t_k := \gamma \bar{\text{tr}}(\Sigma^2 G^k), \quad k \in \{2, 3, 4\}, \quad b := \frac{1}{1 - t_2},$$

and the signal–covariance quadratic forms

$$q_k := \beta^\top G^k \Sigma \beta, \quad k \in \{2, 3, 4\}.$$

Finally define

$$E := \kappa - b\lambda + b^2 \kappa \lambda t_3,$$

$$a_2 := bE^2 + b^4\kappa^2\lambda^2t_4 + b^5\kappa^2\lambda^2t_3^2, \quad a_3 := 2b^2\kappa\lambda E, \quad a_4 := b^3\kappa^2\lambda^2,$$

and

$$u_2 := t_2b, \quad u_3 := t_3b^3, \quad u_4 := t_4b^4 + 2t_3^2b^5.$$

We use standard anisotropic deterministic-equivalent (local-law) results for sample-covariance resolvents under bounded-spectrum and bounded-moment assumptions; see, e.g., [Knowles and Yin \(2017\)](#); [Dobriban and Wager \(2018\)](#); [Dobriban and Sheng \(2020\)](#); [Patil et al. \(2022a, 2023\)](#). In particular, $Q_\lambda = (\widehat{\Sigma} + \lambda I_p)^{-1}$ admits deterministic equivalents uniformly over λ in compact subsets of $(0, \infty)$, and these equivalents can be differentiated with respect to scalar parameters (see [Section B.6.1](#) for more details; see also [Appendix E of Patil and Du \(2023\)](#) for a general background on asymptotic equivalents and a summary of various calculus rules for asymptotic equivalents).

Lemma B.4 (Deterministic equivalents for Q_2, Q_3, Q_4). For each fixed $\lambda > 0$, we have

$$\begin{aligned} Q_2(\lambda) &\asymp \kappa^2b G^2\Sigma, \\ Q_3(\lambda) &\asymp 2\kappa bE G^2\Sigma + 2\kappa^2b^2\lambda G^3\Sigma, \\ Q_4(\lambda) &\asymp a_2 G^2\Sigma + a_3 G^3\Sigma + a_4 G^4\Sigma. \end{aligned}$$

Lemma B.5 (Deterministic limits for U_2, U_3, U_4). For each fixed $\lambda > 0$, we have

$$U_2(\lambda) \xrightarrow{\mathbb{P}} u_2, \quad U_3(\lambda) \xrightarrow{\mathbb{P}} u_3, \quad U_4(\lambda) \xrightarrow{\mathbb{P}} u_4.$$

Proof outlines for Lemmas B.4 and B.5. These follow from: (i) the anisotropic resolvent deterministic equivalent for Q_λ together with the fixed point [\(53\)](#); (ii) differentiation of the two-point deterministic equivalent for $R_{\lambda_1}\Sigma R_{\lambda_2}$ to generate $R\Sigma R^2$, $R^2\Sigma R$, and $R^2\Sigma R^2$; and (iii) the identities $U_3 = -\frac{1}{2}U_2'$ and $U_4 = \frac{1}{6}U_2''$ together with deterministic-equivalent calculus. Detailed proofs are provided below in [Section B.6](#). \square

B.4 Proof of [Theorem 3.1](#)

Proof of [Theorem 3.1](#). Fix $\lambda > 0$. By [\(50\)](#)–[\(52\)](#),

$$\begin{aligned} \bar{R}(\lambda) &= \beta^\top Q_2(\lambda)\beta + \sigma^2 U_2(\lambda) + o_{\mathbb{P}}(1), \\ \bar{C}(\lambda) &= \beta^\top \left(2Q_2(\lambda) - \frac{1}{2}Q_3(\lambda) \right) \beta + \sigma^2 (U_2(\lambda) - \lambda U_3(\lambda)) + o_{\mathbb{P}}(1), \\ \bar{R}_{\text{pd}}(\lambda) &= \beta^\top \left(4Q_2(\lambda) - 2Q_3(\lambda) + Q_4(\lambda) \right) \beta + \sigma^2 (U_2(\lambda) - 2\lambda U_3(\lambda) + \lambda^2 U_4(\lambda)) + o_{\mathbb{P}}(1). \end{aligned}$$

We now apply [Lemma B.4](#) in bilinear forms with the deterministic vector β (assumed to have $\|\beta\|_2 = O(1)$) so these bilinear forms are $O(1)$:

$$\begin{aligned} \beta^\top Q_2(\lambda)\beta &\rightarrow \kappa^2b q_2, \\ \beta^\top \left(2Q_2(\lambda) - \frac{1}{2}Q_3(\lambda) \right) \beta &\rightarrow 2\kappa^2b q_2 - (\kappa bE q_2 + \kappa^2b^2\lambda q_3), \\ \beta^\top \left(4Q_2(\lambda) - 2Q_3(\lambda) + Q_4(\lambda) \right) \beta &\rightarrow 4\kappa^2b q_2 - 2(2\kappa bE q_2 + 2\kappa^2b^2\lambda q_3) + (a_2q_2 + a_3q_3 + a_4q_4). \end{aligned}$$

Also by [Lemma B.5](#),

$$U_2(\lambda) \rightarrow u_2, \quad U_3(\lambda) \rightarrow u_3, \quad U_4(\lambda) \rightarrow u_4.$$

Substituting yields the excess-risk limits:

$$\bar{R}(\lambda) \rightarrow \kappa^2 b q_2 + \sigma^2 u_2,$$

$$\bar{C}(\lambda) \rightarrow 2\kappa^2 b q_2 - (\kappa b E q_2 + \kappa^2 b^2 \lambda q_3) + \sigma^2 (u_2 - \lambda u_3),$$

$$\bar{R}_{\text{pd}}(\lambda) \rightarrow 4\kappa^2 b q_2 - 2(2\kappa b E q_2 + 2\kappa^2 b^2 \lambda q_3) + (a_2 q_2 + a_3 q_3 + a_4 q_4) + \sigma^2 (u_2 - 2\lambda u_3 + \lambda^2 u_4).$$

Finally, recall from the discussion in Section B.1 that the *full* in-distribution quantities in the main paper are obtained by adding σ^2 to each of $\bar{R}(\lambda), \bar{C}(\lambda), \bar{R}_{\text{pd}}(\lambda)$. Therefore the full limits match exactly the statement of Theorem 3.1. \square

B.5 Proof of Lemma B.3

Proof of Lemma B.3. Fix $\lambda > 0$. Let $\varepsilon_i := y_i - x_i^\top \beta$ and $\varepsilon := (\varepsilon_i)_{i=1}^n$. By construction, $\mathbb{E}[\varepsilon_i] = 0$ and $\mathbb{E}[z_i \varepsilon_i] = 0$, and $\mathbb{E}[|\varepsilon_i|^{4+\nu}] < \infty$.

Write $X = Z\Sigma^{1/2}$, $\hat{\Sigma} = X^\top X/n$, and

$$Q_\lambda := (\hat{\Sigma} + \lambda I_p)^{-1}.$$

Define the dual (Gram) resolvent

$$\bar{Q}_\lambda := \left(\frac{1}{n} X X^\top + \lambda I_n \right)^{-1} = \left(\frac{1}{n} Z \Sigma Z^\top + \lambda I_n \right)^{-1}.$$

A standard push-through identity gives, for every integer $m \geq 1$,

$$Q_\lambda^m \frac{X^\top}{n} = \frac{X^\top}{n} \bar{Q}_\lambda^m, \quad \text{and hence} \quad \Sigma^{1/2} Q_\lambda^m \frac{X^\top}{n} = \frac{\Sigma Z^\top}{n} \bar{Q}_\lambda^m. \quad (54)$$

Define also

$$B_2 := \frac{1}{n} X \Sigma X^\top = \frac{1}{n} Z \Sigma^2 Z^\top.$$

All linear terms in Lemma B.2 can be reduced (using (54)) to forms

$$\frac{1}{n} a^\top \Sigma Z^\top \bar{Q}_\lambda^m \varepsilon, \quad m \in \{1, 2\}, \quad (55)$$

where $a \in \mathbb{R}^p$ is deterministic (or independent of Z) with $\|a\|_2 = O(1)$. Likewise, the quadratic noise terms reduce to

$$\frac{1}{n} \varepsilon^\top \bar{Q}_\lambda^m B_2 \bar{Q}_\lambda^\ell \varepsilon, \quad (m, \ell) \in \{(1, 1), (1, 2), (2, 2)\}. \quad (56)$$

Linear forms. For $m = 1$, (55) is exactly Lemma D.2 of Patil and Du (2023) (apply it with $D = \Sigma$ and $g(z_i) = \varepsilon_i$), yielding convergence in probability to 0. For $m = 2$, use the exact identity (Lemma D.4 of Patil and Du (2023))

$$\bar{Q}_\lambda^2 = \frac{1}{t} (\bar{Q}_\lambda - \bar{Q}_{\lambda+t}) + t \bar{Q}_\lambda \bar{Q}_{\lambda+t} \bar{Q}_\lambda, \quad t > 0.$$

Plugging into (55), the difference-quotient term reduces to a difference of two $m = 1$ linear forms at λ and $\lambda + t$, which vanish in probability by Lemma D.2. The remainder is controlled by operator norms: $\|\bar{Q}_\lambda\|_{\text{op}} \leq \lambda^{-1}$, $\|\bar{Q}_{\lambda+t}\|_{\text{op}} \leq (\lambda + t)^{-1}$, $\|\Sigma Z^\top\|_{\text{op}} = \mathcal{O}_{\mathbb{P}}(\sqrt{n})$, and $\|\varepsilon\|_2 = \mathcal{O}_{\mathbb{P}}(\sqrt{n})$. Choosing

a deterministic $t = t_n \rightarrow 0$ (e.g. $t_n = n^{-1/4}$) makes the remainder $o_{\mathbb{P}}(1)$. Hence all linear terms are $o_{\mathbb{P}}(1)$.

Quadratic forms. Fix (m, ℓ) and set $M_{m, \ell} := \bar{Q}_\lambda^m B_2 \bar{Q}_\lambda^\ell$. Write

$$\varepsilon^\top M_{m, \ell} \varepsilon = \sum_i (M_{m, \ell})_{ii} \varepsilon_i^2 + \sum_{i \neq j} (M_{m, \ell})_{ij} \varepsilon_i \varepsilon_j.$$

The off-diagonal sum divided by n converges to 0 in probability by Lemma D.3 of [Patil and Du \(2023\)](#) (after representing $M_{m, \ell}$ via resolvent identities as a finite linear combination of ridge-type resolvent kernels). The diagonal sum equals $\sigma^2 \text{tr}(M_{m, \ell}) + o_{\mathbb{P}}(n)$ by the same leave-one-out decoupling (applied to $\varepsilon_i^2 - \sigma^2$). Therefore,

$$\frac{1}{n} \varepsilon^\top M_{m, \ell} \varepsilon = \frac{\sigma^2}{n} \text{tr}(M_{m, \ell}) + o_{\mathbb{P}}(1).$$

Finally, use cyclicity of trace and commutativity of Q_λ with $\widehat{\Sigma}$ to identify these traces with U_2, U_3, U_4 (as in the proof outline following Lemma B.3). Substituting back into Lemma B.2 proves Lemma B.3. \square

B.6 Proofs of Lemmas B.4 and B.5

B.6.1 Background

We use an anisotropic (bilinear-form) notion of deterministic equivalent.

Definition 1 (Deterministic equivalent). Let $A = A_p$ be a (possibly random) $p \times p$ matrix and $\bar{A} = \bar{A}_p$ a deterministic $p \times p$ matrix. We write $A \asymp \bar{A}$ if for every pair of deterministic vectors $u = u_p, v = v_p$ with $\|u\|_2, \|v\|_2 = O(1)$,

$$u^\top (A - \bar{A}) v \xrightarrow{\mathbb{P}} 0.$$

If $A(\theta), \bar{A}(\theta)$ depend on a parameter θ in an open set Θ , we write $A(\theta) \asymp \bar{A}(\theta)$ *uniformly on compact subsets of Θ* if the convergence above holds uniformly over θ in any compact $K \subset \Theta$.

This notion implies trace convergence whenever operator norms are uniformly bounded. In particular, if $A \asymp \bar{A}$ and $\|A\|_{\text{op}}, \|\bar{A}\|_{\text{op}} = O_{\mathbb{P}}(1)$, then $\overline{\text{tr}}(A) - \overline{\text{tr}}(\bar{A}) \xrightarrow{\mathbb{P}} 0$.

We will also use that uniform deterministic equivalents can be differentiated.

Lemma B.6 (Differentiate a deterministic equivalent). Let $A(\theta)$ be random and $\bar{A}(\theta)$ deterministic, both entrywise differentiable in θ in a neighborhood of θ_0 . Assume $A(\theta) \asymp \bar{A}(\theta)$ uniformly in θ in that neighborhood. Then $A'(\theta_0) \asymp \bar{A}'(\theta_0)$.

Recall $\widehat{\Sigma} = X^\top X/n$ and $Q_\lambda = (\widehat{\Sigma} + \lambda I_p)^{-1}$. A standard anisotropic local law for sample-covariance resolvents gives the following asymptotic equivalence; see, e.g., [Rubio and Mestre \(2011\)](#); [Knowles and Yin \(2017\)](#); [Patil and Du \(2023\)](#):

Lemma B.7 (Scaled resolvent deterministic equivalent). Under Assumption A, for each fixed $\lambda > 0$,

$$\lambda Q_\lambda \asymp \kappa(\lambda) (\Sigma + \kappa(\lambda) I_p)^{-1} = \kappa G, \tag{57}$$

where $\kappa = \kappa(\lambda) > 0$ is the unique solution to

$$\kappa = \lambda + \gamma \kappa \overline{\text{tr}}(\Sigma(\Sigma + \kappa I_p)^{-1}), \tag{58}$$

and $G = (\Sigma + \kappa I_p)^{-1}$. The equivalence holds uniformly for λ in compact subsets of $(0, \infty)$.

Define, as in the main paper,

$$t_k := \gamma \overline{\text{tr}}(\Sigma^2 G^k), \quad k \in \{2, 3, 4\}, \quad b := \frac{1}{1 - t_2}.$$

Lemma B.8 (Derivative of the fixed-point solution). The map $\lambda \mapsto \kappa(\lambda)$ is differentiable for $\lambda > 0$ and

$$\kappa'(\lambda) = b(\lambda) = \frac{1}{1 - t_2(\lambda)}.$$

B.6.2 A Two-Point Deterministic Equivalent via Block Linearization

For $\lambda_1, \lambda_2 > 0$, define $R_i = (\widehat{\Sigma} + \lambda_i I_p)^{-1}$, $\kappa_i = \kappa(\lambda_i)$, $G_i = (\Sigma + \kappa_i I_p)^{-1}$, and

$$t_{12} := \gamma \overline{\text{tr}}(\Sigma^2 G_1 G_2) = \gamma \overline{\text{tr}}(\Sigma G_1 \Sigma G_2). \quad (59)$$

Lemma B.9 (Two-point bias-resolvent deterministic equivalent). For each fixed $\lambda_1, \lambda_2 > 0$,

$$R_1 \Sigma R_2 \asymp \frac{\kappa_1 \kappa_2}{\lambda_1 \lambda_2} \cdot \frac{1}{1 - t_{12}} G_1 \Sigma G_2. \quad (60)$$

Proof. We use a block-resolvent generating function. For a scalar coupling $J \in \mathbb{R}$, define

$$\mathbb{H}(J) := \begin{pmatrix} \widehat{\Sigma} + \lambda_1 I_p & J \Sigma \\ J \Sigma & \widehat{\Sigma} + \lambda_2 I_p \end{pmatrix}, \quad \mathbb{G}(J) := \mathbb{H}(J)^{-1}.$$

Differentiate $\mathbb{H}(J)\mathbb{G}(J) = I$:

$$\mathbb{G}'(J) = -\mathbb{G}(J) \mathbb{H}'(J) \mathbb{G}(J), \quad \mathbb{H}'(0) = \begin{pmatrix} 0 & \Sigma \\ \Sigma & 0 \end{pmatrix}.$$

At $J = 0$, $\mathbb{G}(0) = \text{diag}(R_1, R_2)$, hence the $(1, 2)$ block satisfies the exact identity

$$(\mathbb{G}'(0))_{12} = -R_1 \Sigma R_2. \quad (61)$$

By anisotropic local laws for such deterministic 2×2 linearizations (uniformly for J in a neighborhood of 0) and Lemma B.6, $(\mathbb{G}'(0))_{12}$ admits a deterministic equivalent obtained by differentiating the deterministic equivalent of $\mathbb{G}(J)$ at $J = 0$. On the diagonal blocks, Lemma B.7 gives $R_i \asymp (\kappa_i / \lambda_i) G_i$. The linearized equation for the off-diagonal block reduces (by commutativity of G_1, G_2 with Σ) to the scalar equivalence

$$(\mathbb{G}'(0))_{12} \asymp -\frac{\kappa_1 \kappa_2}{\lambda_1 \lambda_2} \rho_{12} G_1 \Sigma G_2,$$

together with the scalar self-consistency relation $\rho_{12} = 1 + t_{12} \rho_{12}$, where t_{12} is given by (59). Solving yields $\rho_{12} = 1/(1 - t_{12})$, and combining with (61) gives (60). \square

B.6.3 Proof of Lemma B.4

Recall Q_2, Q_3, Q_4 from (43) and define $\mathcal{B}_k := G^k \Sigma$ for $k \in \{2, 3, 4\}$.

Proof of Lemma B.4. Fix $\lambda > 0$ and let $\kappa = \kappa(\lambda)$, $G = (\Sigma + \kappa I_p)^{-1}$.

Equivalent for Q_2 . Apply Lemma B.9 with $\lambda_1 = \lambda_2 = \lambda$. Then $G_1 = G_2 = G$, $\kappa_1 = \kappa_2 = \kappa$, and $t_{12} = t_2$, so

$$Q_\lambda \Sigma Q_\lambda \asymp \frac{\kappa^2}{\lambda^2} \cdot \frac{1}{1-t_2} G \Sigma G = \frac{\kappa^2 b}{\lambda^2} G^2 \Sigma.$$

Multiplying by λ^2 gives

$$Q_2(\lambda) = \lambda^2 Q_\lambda \Sigma Q_\lambda \asymp \kappa^2 b \mathcal{B}_2.$$

Equivalent for Q_3 . Let $F(\lambda_1, \lambda_2) := R_{\lambda_1} \Sigma R_{\lambda_2}$. The exact identities

$$\partial_{\lambda_2} F(\lambda_1, \lambda_2) = -R_{\lambda_1} \Sigma R_{\lambda_2}^2, \quad \partial_{\lambda_1} F(\lambda_1, \lambda_2) = -R_{\lambda_1}^2 \Sigma R_{\lambda_2}$$

imply

$$Q_3(\lambda) = \lambda^3 (Q_\lambda \Sigma Q_\lambda^2 + Q_\lambda^2 \Sigma Q_\lambda) = -\lambda^3 (\partial_{\lambda_1} + \partial_{\lambda_2}) F(\lambda_1, \lambda_2) \Big|_{\lambda_1 = \lambda_2 = \lambda}.$$

By Lemma B.9, $F \asymp \bar{F} := \alpha H$ where

$$\alpha(\lambda_1, \lambda_2) = \frac{\kappa_1 \kappa_2}{\lambda_1 \lambda_2} \cdot \frac{1}{1-t_{12}}, \quad H(\lambda_1, \lambda_2) = G_1 \Sigma G_2.$$

By Lemma B.6, $\partial_{\lambda_i} F \asymp \partial_{\lambda_i} \bar{F}$. A direct product-rule computation on the diagonal $\lambda_1 = \lambda_2 = \lambda$, using $\kappa'(\lambda) = b$ (Lemma B.8) and the trace derivative $\partial_{\lambda_2} t_{12} \Big|_{\text{diag}} = -b t_3$, yields

$$-\lambda^3 (\partial_{\lambda_1} + \partial_{\lambda_2}) \bar{F} \Big|_{\text{diag}} = 2\kappa b E \mathcal{B}_2 + 2\kappa^2 b^2 \lambda \mathcal{B}_3,$$

where $E := \kappa - b\lambda + b^2 \kappa \lambda t_3$. Therefore

$$Q_3(\lambda) \asymp 2\kappa b E \mathcal{B}_2 + 2\kappa^2 b^2 \lambda \mathcal{B}_3.$$

Equivalent for Q_4 . The exact mixed-derivative identity

$$\partial_{\lambda_1} \partial_{\lambda_2} F(\lambda_1, \lambda_2) = R_{\lambda_1}^2 \Sigma R_{\lambda_2}^2$$

implies

$$Q_4(\lambda) = \lambda^4 Q_\lambda^2 \Sigma Q_\lambda^2 = \lambda^4 \partial_{\lambda_1} \partial_{\lambda_2} F(\lambda_1, \lambda_2) \Big|_{\lambda_1 = \lambda_2 = \lambda}.$$

Again by Lemma B.6, we may replace F by $\bar{F} = \alpha H$ and differentiate. On the diagonal, the derivatives satisfy

$$H = \mathcal{B}_2, \quad \partial_{\lambda_1} H = \partial_{\lambda_2} H = -b \mathcal{B}_3, \quad \partial_{\lambda_1} \partial_{\lambda_2} H = b^2 \mathcal{B}_4,$$

and the overlap factor $b_{12} = (1 - t_{12})^{-1}$ satisfies

$$\partial_{\lambda_1} b_{12} \Big|_{\text{diag}} = \partial_{\lambda_2} b_{12} \Big|_{\text{diag}} = -b^3 t_3, \quad \partial_{\lambda_1} \partial_{\lambda_2} b_{12} \Big|_{\text{diag}} = b^4 t_4 + 2b^5 t_3^2,$$

with $t_4 = \gamma \bar{\text{tr}}(\Sigma^2 G^4)$. Carrying out the product-rule expansion of $\partial_{\lambda_1} \partial_{\lambda_2} \bar{F}$ and collecting terms in $\mathcal{B}_2, \mathcal{B}_3, \mathcal{B}_4$ yields

$$\lambda^4 \partial_{\lambda_1} \partial_{\lambda_2} \bar{F} \Big|_{\text{diag}} = a_2 \mathcal{B}_2 + a_3 \mathcal{B}_3 + a_4 \mathcal{B}_4,$$

where

$$a_4 = b^3 \kappa^2 \lambda^2, \quad a_3 = 2b^2 \kappa \lambda E, \quad a_2 = bE^2 + b^4 \kappa^2 \lambda^2 t_4 + b^5 \kappa^2 \lambda^2 t_3^2.$$

Therefore $Q_4(\lambda) \asymp a_2 \mathcal{B}_2 + a_3 \mathcal{B}_3 + a_4 \mathcal{B}_4$, completing the proof. \square

B.6.4 Proof of Lemma B.5

Proof of Lemma B.5. Fix $\lambda > 0$. Recall $U_k(\lambda) = \frac{1}{n} \text{tr}(\Sigma \hat{\Sigma} Q_\lambda^k)$ and $p/n \rightarrow \gamma$.

Limit for U_2 . Use $\hat{\Sigma} Q_\lambda = I_p - \lambda Q_\lambda$ to get the exact identity

$$\hat{\Sigma} Q_\lambda^2 = (I_p - \lambda Q_\lambda) Q_\lambda = Q_\lambda - \lambda Q_\lambda^2,$$

hence

$$U_2(\lambda) = \frac{1}{n} \text{tr}(\Sigma(Q_\lambda - \lambda Q_\lambda^2)) = \gamma \bar{\text{tr}}(\Sigma Q_\lambda) - \lambda \gamma \bar{\text{tr}}(\Sigma Q_\lambda^2). \quad (62)$$

Define $s(\lambda) := \gamma \bar{\text{tr}}(\Sigma Q_\lambda)$. Since $Q'_\lambda = -Q_\lambda^2$,

$$s'(\lambda) = \gamma \bar{\text{tr}}(\Sigma Q'_\lambda) = -\gamma \bar{\text{tr}}(\Sigma Q_\lambda^2),$$

so (62) becomes

$$U_2(\lambda) = s(\lambda) + \lambda s'(\lambda). \quad (63)$$

By Lemma B.7 and trace convergence,

$$s(\lambda) = \gamma \bar{\text{tr}}(\Sigma Q_\lambda) \rightarrow \gamma \bar{\text{tr}}\left(\Sigma \cdot \frac{\kappa}{\lambda} G\right) = \frac{\gamma \kappa}{\lambda} \bar{\text{tr}}(\Sigma G).$$

Using the fixed-point equation (58), $\gamma \kappa \bar{\text{tr}}(\Sigma G) = \kappa - \lambda$, hence

$$s(\lambda) \rightarrow \frac{\kappa - \lambda}{\lambda} = \frac{\kappa}{\lambda} - 1.$$

By uniformity in λ locally and Lemma B.6,

$$s'(\lambda) \rightarrow \frac{\kappa'}{\lambda} - \frac{\kappa}{\lambda^2}.$$

Substitute into (63):

$$U_2(\lambda) \rightarrow \left(\frac{\kappa}{\lambda} - 1\right) + \lambda \left(\frac{\kappa'}{\lambda} - \frac{\kappa}{\lambda^2}\right) = \kappa'(\lambda) - 1 = b - 1 = \frac{t_2}{1 - t_2} = t_2 b =: u_2.$$

Limits for U_3 and U_4 . Since $Q'_\lambda = -Q_\lambda^2$, we have the exact identities

$$\frac{d}{d\lambda} Q_\lambda^2 = -2Q_\lambda^3, \quad \frac{d}{d\lambda} Q_\lambda^3 = -3Q_\lambda^4,$$

and therefore

$$U_2'(\lambda) = -2U_3(\lambda), \quad U_3'(\lambda) = -3U_4(\lambda),$$

equivalently

$$U_3(\lambda) = -\frac{1}{2}U_2'(\lambda), \quad U_4(\lambda) = \frac{1}{6}U_2''(\lambda). \quad (64)$$

From the U_2 step, $U_2(\lambda) \rightarrow u_2(\lambda) = t_2 b$. By uniformity and Lemma B.6, we may differentiate to obtain $U_2' \rightarrow u_2'$ and $U_2'' \rightarrow u_2''$. Thus by (64),

$$U_3(\lambda) \rightarrow -\frac{1}{2}u_2'(\lambda), \quad U_4(\lambda) \rightarrow \frac{1}{6}u_2''(\lambda).$$

Finally, compute these derivatives. Since $t_2'(\kappa) = -2t_3$ and $\kappa' = b$, we have $t_2'(\lambda) = -2t_3 b$ and hence

$$u_2(\lambda) = \frac{t_2}{1-t_2} \Rightarrow u_2'(\lambda) = \frac{t_2'(\lambda)}{(1-t_2)^2} = (-2t_3 b) b^2 = -2t_3 b^3,$$

so $-\frac{1}{2}u_2' = t_3 b^3 =: u_3$. Similarly, using $t_3'(\kappa) = -3t_4$ and $b'(\lambda) = -2t_3 b^3$, we obtain

$$u_2''(\lambda) = \frac{d}{d\lambda}(-2t_3 b^3) = -2(t_3'(\lambda)b^3 + t_3 \cdot 3b^2 b') = 6t_4 b^4 + 12t_3^2 b^5,$$

hence $\frac{1}{6}u_2'' = t_4 b^4 + 2t_3^2 b^5 =: u_4$. This proves $U_3 \rightarrow u_3$ and $U_4 \rightarrow u_4$. \square

B.7 Proof of Corollary 3.2

Proof of Corollary 3.2. Fix $\lambda > 0$. Recall from Theorem 3.1 that the optimal mixing weight and optimal SD risk satisfy

$$\xi^*(\lambda) \xrightarrow{P} \frac{\mathcal{R}(\lambda) - \mathcal{C}(\lambda)}{\mathcal{R}(\lambda) + \mathcal{R}_{\text{pd}}(\lambda) - 2\mathcal{C}(\lambda)}, \quad \mathcal{R}_{\text{sd}}^*(\lambda) \xrightarrow{P} \mathcal{R}_{\text{sd}}^*(\lambda) := \mathcal{R}(\lambda) - \frac{(\mathcal{R}(\lambda) - \mathcal{C}(\lambda))^2}{\mathcal{R}(\lambda) + \mathcal{R}_{\text{pd}}(\lambda) - 2\mathcal{C}(\lambda)}.$$

In particular, writing $\mathcal{D}(\lambda) := \mathcal{R}(\lambda) + \mathcal{R}_{\text{pd}}(\lambda) - 2\mathcal{C}(\lambda) \geq 0$, we have $\mathcal{R}_{\text{sd}}^*(\lambda) < \mathcal{R}(\lambda)$ whenever $\mathcal{R}(\lambda) - \mathcal{C}(\lambda) \neq 0$ and $\mathcal{D}(\lambda) > 0$, and moreover $\text{sign}(\xi^*(\lambda)) = \text{sign}(\mathcal{R}(\lambda) - \mathcal{C}(\lambda))$ whenever $\mathcal{D}(\lambda) > 0$. Thus it remains to characterize the sign of $\mathcal{R}(\lambda) - \mathcal{C}(\lambda)$ under the isotropic signal prior.

As before, let $\widehat{\Sigma} := X^\top X/n$ and $Q_\lambda := (\widehat{\Sigma} + \lambda I_p)^{-1}$. Under the isotropic prior $\beta \sim \mathcal{N}(0, (r^2/p)I_p)$, for any random matrix M (measurable with respect to X), $\mathbb{E}[\beta^\top M \beta \mid X] = \text{tr}(M \mathbb{E}[\beta \beta^\top]) = r^2 \text{tr}(M)/p$. Likewise, for a noise vector ε independent of X with $\mathbb{E}[\varepsilon] = 0$ and $\mathbb{E}[\varepsilon \varepsilon^\top] = \sigma^2 I_n$, for any random matrix M , $\mathbb{E}[\varepsilon^\top M \varepsilon \mid X] = \sigma^2 \text{tr}(M)$. Applying these X -conditional limits on the conditional risk and correlation expansions (50)–(52) yields (with $o_{\mathbb{P}}(1)$ remainders absorbed since they are mean-zero after conditioning on X):

$$\bar{R}_X(\lambda) := \mathbb{E}[\bar{R}(\lambda) \mid X] = \frac{r^2}{p} \text{tr}(\Sigma \lambda^2 Q_\lambda^2) + \frac{\sigma^2}{n} \text{tr}(\Sigma Q_\lambda (I_p - \lambda Q_\lambda)), \quad (65)$$

$$\bar{R}_{\text{pd},X}(\lambda) := \mathbb{E}[\bar{R}_{\text{pd}}(\lambda) \mid X] = \frac{r^2}{p} \text{tr}(\Sigma \lambda^2 Q_\lambda^2 (2I_p - \lambda Q_\lambda)^2) + \frac{\sigma^2}{n} \text{tr}(\Sigma Q_\lambda (I_p - \lambda Q_\lambda)^3), \quad (66)$$

$$\bar{C}_X(\lambda) := \mathbb{E}[\bar{C}(\lambda) \mid X] = \frac{r^2}{p} \text{tr}(\Sigma \lambda^2 Q_\lambda^2 (2I_p - \lambda Q_\lambda)) + \frac{\sigma^2}{n} \text{tr}(\Sigma Q_\lambda (I_p - \lambda Q_\lambda)^2). \quad (67)$$

Here $\bar{R}, \bar{R}_{\text{pd}}, \bar{C}$ denote the corresponding excess components as above, i.e., $R(\lambda) = \sigma^2 + \bar{R}(\lambda)$, $R_{\text{pd}}(\lambda) = \sigma^2 + \bar{R}_{\text{pd}}(\lambda)$, and $C(\lambda) = \sigma^2 + \bar{C}(\lambda)$.

Define the nonasymptotic critical value $\lambda_{n,p}^* := \frac{p}{n} \frac{\sigma^2}{r^2}$ so that $\lambda_{n,p}^* \rightarrow \lambda^* := \gamma \frac{\sigma^2}{r^2}$. A direct simplification of (65) and (67) gives

$$\bar{R}_X(\lambda) - \bar{C}_X(\lambda) = \lambda \left(\frac{\sigma^2}{n} - \frac{r^2}{p} \lambda \right) \left(\text{tr}(\Sigma Q_\lambda^2) - \lambda \text{tr}(\Sigma Q_\lambda^3) \right)$$

$$\begin{aligned}
&= \lambda(\lambda_{n,p}^* - \lambda) \frac{r^2}{p} \operatorname{tr}(\Sigma Q_\lambda^2 (I_p - \lambda Q_\lambda)) \\
&= \lambda(\lambda_{n,p}^* - \lambda) \frac{r^2}{p} \operatorname{tr}(\Sigma Q_\lambda^2 \widehat{\Sigma} Q_\lambda),
\end{aligned}$$

where in the last step we used $I_p - \lambda Q_\lambda = \widehat{\Sigma} Q_\lambda$, which follows from $(\widehat{\Sigma} + \lambda I_p) Q_\lambda = I_p$. The trace factor is strictly positive for $\lambda > 0$ with probability one because $\Sigma \succ 0$ and $Q_\lambda \succ 0$ for $\lambda > 0$ and hence the trace vanishes if and only if $\widehat{\Sigma} = 0$ (which under Assumption A is a measure zero event). Therefore, for each fixed $\lambda > 0$, $\operatorname{sign}(\bar{R}_X(\lambda) - \bar{C}_X(\lambda)) = \operatorname{sign}(\lambda_{n,p}^* - \lambda)$ with probability one.

Now, by the same concentration arguments of Lemma B.3 used to prove Theorem 3.1, the random quantities $\bar{R}(\lambda) - \bar{C}(\lambda)$ concentrate around their conditional means, and $\bar{R}(\lambda) - \bar{C}(\lambda) - (\bar{R}_X(\lambda) - \bar{C}_X(\lambda)) \xrightarrow{P} 0$. Moreover, $\lambda_{n,p}^* \rightarrow \lambda^*$ and $\bar{R}_X(\lambda) - \bar{C}_X(\lambda)$ converges to $\mathcal{R}(\lambda) - \mathcal{C}(\lambda)$ (since σ^2 cancels in the difference). Hence, $\mathcal{R}(\lambda) - \mathcal{C}(\lambda) = 0$ if and only if $\lambda = \lambda^*$, and $\operatorname{sign}(\mathcal{R}(\lambda) - \mathcal{C}(\lambda)) = \operatorname{sign}(\lambda^* - \lambda)$. Finally, for $\lambda > 0$, note that the limiting discrepancy $\mathcal{D}(\lambda) = \mathcal{R}(\lambda) + \mathcal{R}_{\text{pd}}(\lambda) - 2\mathcal{C}(\lambda)$ is strictly positive with probability one in this case, as it is the limit of $\mathbb{E}[(f_\lambda(x_0) - f_{\text{pd},\lambda}(x_0))^2 \mid \mathcal{D}]$, which is nondegenerate in this case (see also the proof of Theorem 5.2 for a more direct argument). \square

B.8 Proof of Proposition 3.3

We provide two proofs. The first evaluates the needed trace functionals by diagonalizing $\widehat{\Sigma}$ and invoking the Marchenko–Pastur law (including negative moments) in the limits $\lambda \rightarrow 0$ and $\lambda \rightarrow \infty$. The second is a specialization of the general deterministic equivalent in Theorem 3.1 to $\Sigma = I_p$.

B.8.1 First proof

For the first approach, we characterize separately the extreme- λ limits of the ridge risk $\mathcal{R}(\lambda)$, the ridge-optimal risk \mathcal{R}^* , and the optimal SD risk $\mathcal{R}_{\text{sd}}^*(\lambda)$, and then take ratios as in Proposition 3.3.

We first recall some known results from the literature; see, e.g., Dobriban and Wager (2018).

Lemma B.10. Assume $\Sigma = I_p$ and $\beta \sim \mathcal{N}(0, (r^2/p)I_p)$, and consider proportional asymptotics $p, n \rightarrow \infty$ with $p/n \rightarrow \gamma$. Then the asymptotic ridge prediction risk $\mathcal{R}(\lambda)$ satisfies:

$$\lim_{\lambda \rightarrow 0} \mathcal{R}(\lambda) = \begin{cases} \frac{\sigma^2 \gamma}{1-\gamma} + \sigma^2, & \gamma \in (0, 1), \\ \frac{r^2(\gamma-1)}{\gamma} + \frac{\sigma^2}{\gamma-1} + \sigma^2, & \gamma \in (1, \infty), \end{cases} \quad (68)$$

$$\lim_{\lambda \rightarrow \infty} \mathcal{R}(\lambda) = r^2 + \sigma^2, \quad \forall \gamma \in (0, \infty). \quad (69)$$

Lemma B.11. Assume $\Sigma = I_p$. Then the (proportional) asymptotic ridge-optimal risk is:

$$\mathcal{R}^* = \sigma^2 + \frac{1}{2\gamma} \left(-\gamma\sigma^2 + r^2(\gamma-1) + \sqrt{4\gamma^2 r^2 \sigma^2 + (\gamma\sigma^2 - r^2(\gamma-1))^2} \right).$$

Next we characterize the extreme- λ limits of $\mathcal{R}_{\text{sd}}^*(\lambda)$.

Lemma B.12. Assume $\Sigma = I_p$ and $\beta \sim \mathcal{N}(0, (r^2/p)I_p)$, and consider proportional asymptotics $p, n \rightarrow \infty$ with $p/n \rightarrow \gamma$. Then the asymptotic optimal SD prediction risk $\mathcal{R}_{\text{sd}}^*(\lambda)$ satisfies

$$\lim_{\lambda \rightarrow 0} \mathcal{R}_{\text{sd}}^*(\lambda) = \begin{cases} \frac{r^2 \sigma^2 \gamma (1-\gamma)^2 + \sigma^4 \gamma^3}{r^2 (1-\gamma)^3 + \sigma^2 \gamma (1-\gamma^2)} + \sigma^2, & \gamma \in (0, 1), \\ \frac{r^4 (\gamma-1)^4 + r^2 \sigma^2 \gamma (\gamma+2) (\gamma-1)^2 + \sigma^4 \gamma^2}{r^2 \gamma (\gamma-1)^3 + \sigma^2 \gamma^2 (\gamma^2-1)} + \sigma^2, & \gamma \in (1, \infty), \end{cases} \quad (70)$$

$$\lim_{\lambda \rightarrow \infty} \mathcal{R}_{\text{sd}}^*(\lambda) = \frac{r^4\gamma + r^2\sigma^2\gamma}{r^2(\gamma + 1) + \sigma^2\gamma} + \sigma^2, \quad \gamma \in (0, \infty). \quad (71)$$

Proof of Lemma B.12. As before, let $Q_\lambda := (\widehat{\Sigma} + \lambda I_p)^{-1}$ with $\widehat{\Sigma} := X^\top X/n$. Recall from (42) the closed-form expression of the optimal SD risk, in terms of X -conditional quantities:

$$\mathcal{R}_{\text{sd}}^*(\lambda) = \sigma^2 + \lim_{n,p \rightarrow \infty} \left(\frac{\overline{R}_X(\lambda)\overline{R}_{\text{pd},X}(\lambda) - \overline{C}_X(\lambda)^2}{\overline{R}_X(\lambda) + \overline{R}_{\text{pd},X}(\lambda) - 2\overline{C}_X(\lambda)} \right).$$

Recall the expressions for these terms from equations (65)–(67). Based on the assumption $\Sigma = I$, we have

$$\begin{aligned} & \overline{R}_X(\lambda)\overline{R}_{\text{pd},X}(\lambda) - \overline{C}_X(\lambda)^2 \\ &= \frac{r^4}{p^2} \left[\text{tr}(\lambda^2 Q_\lambda^2) \text{tr}(\lambda^2 Q_\lambda^2 (2I - \lambda Q_\lambda)^2) - \text{tr}(\lambda^2 Q_\lambda^2 (2I - \lambda Q_\lambda))^2 \right] \\ &+ \frac{r^2\sigma^2}{pn\lambda} \left[\text{tr}(\lambda^2 Q_\lambda^2) \text{tr}(\lambda Q_\lambda (I - \lambda Q_\lambda)^3) + \text{tr}(\lambda Q_\lambda (I - \lambda Q_\lambda)) \text{tr}(\lambda^2 Q_\lambda^2 (2I - \lambda Q_\lambda)^2) \right. \\ &\quad \left. - 2 \text{tr}(\lambda^2 Q_\lambda^2 (2I - \lambda Q_\lambda)) \text{tr}(\lambda Q_\lambda (I - \lambda Q_\lambda)^2) \right] \\ &+ \frac{\sigma^4}{n^2\lambda^2} \left[\text{tr}(\lambda Q_\lambda (I - \lambda Q_\lambda)) \text{tr}(\lambda Q_\lambda (I - \lambda Q_\lambda)^3) - \text{tr}(\lambda Q_\lambda (I - \lambda Q_\lambda)^2)^2 \right], \\ &:= \frac{r^4}{p^2} A_1 + \frac{r^2\sigma^2}{pn\lambda} A_2 + \frac{\sigma^4}{n^2\lambda^2} A_3, \\ &\overline{R}_X(\lambda) + \overline{R}_{\text{pd},X}(\lambda) - 2\overline{C}_X(\lambda) = \frac{r^2}{p} \text{tr}(\lambda^2 Q_\lambda^2 (I - \lambda Q_\lambda)^2) + \frac{\sigma^2}{n\lambda} \text{tr}(\lambda^3 Q_\lambda^3 (I - \lambda Q_\lambda)). \end{aligned}$$

Denote the eigenvalues of the sample covariance matrix $\widehat{\Sigma}$ as $\{s_i\}_{i=1}^p$. Then we can rewrite the traces in the above as:

$$\begin{aligned} A_1 &= \text{tr}(\lambda^2 Q_\lambda^2) \text{tr}(\lambda^2 Q_\lambda^2 (2I - \lambda Q_\lambda)^2) - \text{tr}(\lambda^2 Q_\lambda^2 (2I - \lambda Q_\lambda))^2 \\ &= \left(\sum_{i=1}^p \frac{\lambda^2}{(s_i + \lambda)^2} \right) \left(\sum_{i=1}^p \frac{\lambda^2}{(s_i + \lambda)^2} \left(2 - \frac{\lambda}{s_i + \lambda} \right)^2 \right) - \left(\sum_{i=1}^p \frac{\lambda^2}{(s_i + \lambda)^2} \left(2 - \frac{\lambda}{s_i + \lambda} \right) \right)^2 \\ &\stackrel{(1)}{=} \frac{1}{2} \sum_{i,j=1}^p \left(\frac{\lambda}{s_i + \lambda} \frac{\lambda}{s_j + \lambda} \left(2 - \frac{\lambda}{s_j + \lambda} \right) - \frac{\lambda}{s_j + \lambda} \frac{\lambda}{s_i + \lambda} \left(2 - \frac{\lambda}{s_i + \lambda} \right) \right)^2 \\ &= \sum_{i,j=1}^p \frac{1}{2} \frac{\lambda^2}{(s_i + \lambda)^2 (s_j + \lambda)^2} \left(\frac{\lambda}{s_i + \lambda} - \frac{\lambda}{s_j + \lambda} \right)^2 \\ &= \frac{1}{2} \sum_{i,j=1}^p \frac{\lambda^6 (s_i - s_j)^2}{(s_i + \lambda)^4 (s_j + \lambda)^4}, \\ A_2 &= \left(\sum_{i=1}^p \frac{\lambda^2}{(s_i + \lambda)^2} \right) \left(\sum_{i=1}^p \frac{\lambda}{s_i + \lambda} \frac{s_i^3}{(s_i + \lambda)^3} \right) + \left(\sum_{i=1}^p \frac{\lambda}{s_i + \lambda} \frac{s_i}{s_i + \lambda} \right) \left(\sum_{i=1}^p \frac{\lambda^2}{(s_i + \lambda)^2} \left(2 - \frac{\lambda}{s_i + \lambda} \right)^2 \right) \\ &\quad - 2 \left(\sum_{i=1}^p \frac{\lambda^2}{(s_i + \lambda)^2} \left(2 - \frac{\lambda}{s_i + \lambda} \right) \right) \left(\sum_{i=1}^p \frac{\lambda}{s_i + \lambda} \frac{s_i^2}{(s_i + \lambda)^2} \right) \end{aligned}$$

$$\begin{aligned}
& \stackrel{(2)}{=} \sum_{i,j=1}^p \left(\frac{\lambda}{s_i + \lambda} \sqrt{\frac{\lambda}{s_j + \lambda}} \frac{s_j^{3/2}}{(s_j + \lambda)^{3/2}} - \frac{\lambda}{s_i + \lambda} \left(2 - \frac{\lambda}{s_i + \lambda} \right) \sqrt{\frac{\lambda}{s_j + \lambda}} \frac{s_j}{s_j + \lambda} \right)^2 \\
& = \sum_{i,j=1}^p \frac{\lambda^2}{(s_i + \lambda)^2} \frac{\lambda}{s_j + \lambda} \frac{s_j}{s_j + \lambda} \left(\frac{s_j}{s_j + \lambda} - 2 + \frac{\lambda}{s_i + \lambda} \right)^2 \\
& = \sum_{i,j=1}^p \frac{\lambda^3 s_j}{(s_i + \lambda)^2 (s_j + \lambda)^2} \left(\frac{-\lambda}{s_j + \lambda} + \frac{\lambda}{s_i + \lambda} - 1 \right)^2, \\
A_3 & = \left(\sum_{i=1}^p \frac{\lambda}{s_i + \lambda} \frac{s_i}{s_i + \lambda} \right) \left(\sum_{i=1}^p \frac{\lambda}{s_i + \lambda} \frac{s_i^3}{(s_i + \lambda)^3} \right) - \left(\sum_{i=1}^p \frac{\lambda}{s_i + \lambda} \frac{s_i^2}{(s_i + \lambda)^2} \right)^2 \\
& \stackrel{(1)}{=} \frac{1}{2} \sum_{i,j=1}^p \left(\sqrt{\frac{\lambda}{s_i + \lambda}} \frac{s_i}{s_i + \lambda} \sqrt{\frac{\lambda}{s_j + \lambda}} \frac{s_j^{3/2}}{(s_j + \lambda)^{3/2}} - \sqrt{\frac{\lambda}{s_j + \lambda}} \frac{s_j}{s_j + \lambda} \sqrt{\frac{\lambda}{s_i + \lambda}} \frac{s_i^{3/2}}{(s_i + \lambda)^{3/2}} \right)^2 \\
& = \frac{1}{2} \sum_{i,j=1}^p \frac{\lambda^2 s_i s_j}{(s_i + \lambda)^2 (s_j + \lambda)^2} \left(\frac{s_i}{s_i + \lambda} - \frac{s_j}{s_j + \lambda} \right)^2 \\
& = \frac{1}{2} \sum_{i,j=1}^p \frac{\lambda^4 s_i s_j (s_i - s_j)^2}{(s_i + \lambda)^4 (s_j + \lambda)^4},
\end{aligned}$$

and

$$\bar{R}_X(\lambda) + \bar{R}_{\text{pd},X}(\lambda) - 2\bar{C}_X(\lambda) = \frac{r^2}{p} \sum_{i=1}^p \frac{\lambda^2 s_i^2}{(s_i + \lambda)^4} + \frac{\sigma^2}{n} \sum_{i=1}^p \frac{\lambda^2 s_i}{(s_i + \lambda)^4}, \quad (72)$$

where (1) and (2) are from the following equalities:

$$\begin{aligned}
& \left(\sum_{i=1}^p x_i^2 \right) \left(\sum_{i=1}^p y_i^2 \right) - \left(\sum_{i=1}^p x_i y_i \right)^2 = \frac{1}{2} \sum_{i,j=1}^p \left(x_i y_j - x_j y_i \right)^2, \\
& \left(\sum_{i=1}^p x_i^2 \right) \left(\sum_{i=1}^p y_i^2 \right) + \left(\sum_{i=1}^p z_i^2 \right) \left(\sum_{i=1}^p t_i^2 \right) - 2 \left(\sum_{i=1}^p x_i t_i \right) \left(\sum_{i=1}^p y_i z_i \right) = \sum_{i,j=1}^p \left(x_i y_j - z_j t_i \right)^2.
\end{aligned}$$

Case 1: $\lambda \rightarrow \infty$ and $\gamma \in (0, \infty)$. For this case, we have

$$\begin{aligned}
\lambda^2 \frac{r^4}{p^2} A_1 & = \frac{r^4}{2p^2} \sum_{i,j=1}^p \frac{\lambda^8 (s_i - s_j)^2}{(s_i + \lambda)^4 (s_j + \lambda)^4} = \frac{r^4}{2p^2} \sum_{i,j=1}^p \frac{(s_i - s_j)^2}{(s_i/\lambda + 1)^4 (s_j/\lambda + 1)^4} \\
& \xrightarrow{\lambda \rightarrow \infty} \frac{r^4}{2p^2} \sum_{i,j=1}^p (s_i - s_j)^2 \xrightarrow{n,p \rightarrow \infty} r^4 \gamma,
\end{aligned}$$

where we used the fact that $\sum_{i,j=1}^p (s_i - s_j)^2/p^2 = 2 \sum_i s_i^2/p - 2(\sum_i s_i/p)^2 \rightarrow 2(1 + \gamma) - 2 = 2\gamma$ based on the Marchenko–Pastur law. Similarly,

$$\lambda^2 \frac{r^2 \sigma^2}{pn\lambda} A_2 = \frac{r^2 \sigma^2}{pn} \sum_{i,j=1}^p \frac{\lambda^4 s_j}{(s_i + \lambda)^2 (s_j + \lambda)^2} \left(\frac{-\lambda}{s_j + \lambda} + \frac{\lambda}{s_i + \lambda} - 1 \right)^2$$

$$\begin{aligned}
&= \frac{r^2 \sigma^2}{pn} \sum_{i,j=1}^p \frac{s_j}{(s_i/\lambda + 1)^2 (s_j/\lambda + 1)^2} \left(\frac{-1}{s_j/\lambda + 1} + \frac{1}{s_i/\lambda + 1} - 1 \right)^2 \\
&\xrightarrow{\lambda \rightarrow \infty} \frac{r^2 \sigma^2}{pn} \sum_{i,j=1}^p s_j = \frac{r^2 \sigma^2}{n} \sum_{j=1}^p s_j \xrightarrow{n,p \rightarrow \infty} r^2 \sigma^2 \gamma, \\
\lambda^2 \frac{\sigma^4}{n^2 \lambda^2} A_3 &= \frac{\sigma^2}{2n^2} \sum_{i,j=1}^p \frac{\lambda^4 s_i s_j (s_i - s_j)^2}{(s_i + \lambda)^4 (s_j + \lambda)^4} \xrightarrow{\lambda \rightarrow \infty} 0, \\
\lambda^2 (\bar{R}_X(\lambda) + \bar{R}_{\text{pd},X}(\lambda) - 2\bar{C}_X(\lambda)) &= \frac{r^2}{p} \sum_{i=1}^p \frac{\lambda^4 s_i^2}{(s_i + \lambda)^4} + \frac{\sigma^2}{n} \sum_{i=1}^p \frac{\lambda^4 s_i}{(s_i + \lambda)^4} \\
&\xrightarrow{\lambda \rightarrow \infty} \frac{r^2}{p} \sum_{i=1}^p s_i^2 + \frac{\sigma^2}{n} \sum_{i=1}^p s_i \xrightarrow{n,p \rightarrow \infty} \frac{r^2}{p} (1 + \gamma) + \sigma^2 \gamma.
\end{aligned}$$

From the above limits, we have

$$\lim_{\lambda \rightarrow \infty} \mathcal{R}_{\text{sd}}^*(\lambda) = \sigma^2 + \lim_{\lambda \rightarrow \infty} \lim_{n,p \rightarrow \infty} \left(\frac{\lambda^2 (\bar{R}_X(\lambda) \bar{R}_{\text{pd},X}(\lambda) - \bar{C}_X(\lambda)^2)}{\lambda^2 (\bar{R}_X(\lambda) + \bar{R}_{\text{pd},X}(\lambda) - 2\bar{C}_X(\lambda))} \right) = \sigma^2 + \frac{r^4 \gamma + r^2 \sigma^2 \gamma}{r^2 (1 + \gamma) + \sigma^2 \gamma}.$$

Case 2: $\lambda \rightarrow 0$ and $\gamma < 1$. First, we state the following results for negative moments of Marchenko–Pastur law. Assume $X \sim \text{MP}(\gamma)$ with $\gamma < 1$, then from Lemma B.13, we have

$$\mathbb{E}[X^{-1}] = \frac{1}{1 - \gamma}, \quad \mathbb{E}[X^{-2}] = \frac{1}{(1 - \gamma)^3}, \quad \mathbb{E}[X^{-3}] = \frac{1 + \gamma}{(1 - \gamma)^5}.$$

Back to this case, we have

$$\begin{aligned}
\lambda^{-2} \frac{r^4}{p^2} A_1 &= \frac{r^4}{2p^2} \sum_{i,j=1}^p \frac{\lambda^4 (s_i - s_j)^2}{(s_i + \lambda)^4 (s_j + \lambda)^4} \xrightarrow{\lambda \rightarrow 0} 0, \\
\lambda^{-2} \frac{r^2 \sigma^2}{pn \lambda} A_2 &= \frac{r^2 \sigma^2}{pn} \sum_{i,j=1}^p \frac{s_j}{(s_i + \lambda)^2 (s_j + \lambda)^2} \left(\frac{-\lambda}{s_j + \lambda} + \frac{\lambda}{s_i + \lambda} - 1 \right)^2 \\
&\xrightarrow{\lambda \rightarrow 0} \frac{r^2 \sigma^2}{pn} \sum_{i,j=1}^p \frac{1}{s_i^2 s_j} = \frac{r^2 \sigma^2}{pn} \left(\sum_{i=1}^p \frac{1}{s_i^2} \right) \left(\sum_{i=1}^p \frac{1}{s_i} \right) \\
&= \frac{r^2 \sigma^2 p}{n} \left(\frac{1}{p} \sum_{i=1}^p \frac{1}{s_i^2} \right) \left(\frac{1}{p} \sum_{i=1}^p \frac{1}{s_i} \right) \xrightarrow{n,p \rightarrow \infty} \frac{r^2 \sigma^2 \gamma}{(1 - \gamma)^4}, \\
\lambda^{-2} \frac{\sigma^4}{n^2 \lambda^2} A_3 &= \frac{\sigma^4}{2n^2} \sum_{i,j=1}^p \frac{s_i s_j (s_i - s_j)^2}{(s_i + \lambda)^4 (s_j + \lambda)^4} \xrightarrow{\lambda \rightarrow 0} \frac{\sigma^4}{2n^2} \sum_{i,j=1}^p \frac{(s_i - s_j)^2}{s_i^3 s_j^3} \\
&= \frac{\sigma^4}{n^2} \left(\left(\sum_{i=1}^p \frac{1}{s_i} \right) \left(\sum_{i=1}^p \frac{1}{s_i^3} \right) - \left(\sum_{i=1}^p \frac{1}{s_i^2} \right)^2 \right) \\
&\xrightarrow{n,p \rightarrow \infty} \sigma^4 \gamma^2 \left(\frac{1 + \gamma}{(1 - \gamma)^6} - \frac{1}{(1 - \gamma)^6} \right) = \sigma^4 \gamma^2 \frac{\gamma}{(1 - \gamma)^6}, \\
\lambda^{-2} (\bar{R}_X(\lambda) + \bar{R}_{\text{pd},X}(\lambda) - 2\bar{C}_X(\lambda)) &= \frac{r^2}{p} \sum_{i=1}^p \frac{s_i^2}{(s_i + \lambda)^4} + \frac{\sigma^2}{n} \sum_{i=1}^p \frac{s_i}{(s_i + \lambda)^4}
\end{aligned}$$

$$\xrightarrow{\lambda \rightarrow 0} \frac{r^2}{p} \sum_{i=1}^p \frac{1}{s_i^2} + \frac{\sigma^2}{n} \sum_{i=1}^p \frac{1}{s_i^3} \xrightarrow{n, p \rightarrow \infty} \frac{r^2}{(1-\gamma)^3} + \frac{\sigma^2 \gamma (1+\gamma)}{(1-\gamma)^5}$$

From the above limits, we have

$$\lim_{\lambda \rightarrow 0} \mathcal{R}_{\text{sd}}^*(\lambda) = \sigma^2 + \lim_{\lambda \rightarrow 0} \lim_{n, p \rightarrow \infty} \left(\frac{\lambda^{-2} (\bar{R}_X(\lambda) \bar{R}_{\text{pd}, X}(\lambda) - \bar{C}_X(\lambda)^2)}{\lambda^{-2} (\bar{R}_X(\lambda) + \bar{R}_{\text{pd}, X}(\lambda) - 2\bar{C}_X(\lambda))} \right) = \sigma^2 + \frac{r^2 \sigma^2 \gamma (1-\gamma)^2 + \sigma^4 \gamma^3}{r^2 (1-\gamma)^3 + \sigma^2 \gamma (1-\gamma^2)}.$$

Case 3: $\lambda \rightarrow 0$ and $\gamma > 1$. For this case, the sample covariance is at most rank $n < p$, so we have $s_i = 0$ for $i > n$. Thus

$$\begin{aligned} \lambda^{-2} \frac{r^4}{p^2} A_1 &= \frac{r^4}{2p^2} \sum_{i,j=1}^p \frac{\lambda^4 (s_i - s_j)^2}{(s_i + \lambda)^4 (s_j + \lambda)^4} \\ &= \frac{r^4}{2p^2} \left(\sum_{i=1}^n \sum_{j=1}^n \frac{\lambda^4 (s_i - s_j)^2}{(s_i + \lambda)^4 (s_j + \lambda)^4} + 2 \sum_{i=1}^n \sum_{j=n+1}^p \frac{\lambda^4 (s_i - s_j)^2}{(s_i + \lambda)^4 (s_j + \lambda)^4} \right) \\ &= \frac{r^4}{2p^2} \left(\sum_{i=1}^n \sum_{j=1}^n \frac{\lambda^4 (s_i - s_j)^2}{(s_i + \lambda)^4 (s_j + \lambda)^4} + 2(p-n) \sum_{i=1}^n \frac{\lambda^4 s_i^2}{(s_i + \lambda)^4 \lambda^4} \right) \\ &\xrightarrow{\lambda \rightarrow 0} \frac{r^4}{2p^2} \left(0 + 2(p-n) \sum_{i=1}^n \frac{1}{s_i^2} \right) = r^4 \frac{p-n}{p} \left(\frac{1}{p} \sum_{i=1}^n \frac{1}{s_i^2} \right) \\ &\xrightarrow{n, p \rightarrow \infty} r^4 \left(1 - \frac{1}{\gamma} \right) \frac{1}{(\gamma-1)^3} = \frac{r^4}{\gamma(\gamma-1)^2}, \\ \lambda^{-2} \frac{r^2 \sigma^2}{pn\lambda} A_2 &= \frac{r^2 \sigma^2}{pn} \sum_{i,j=1}^p \frac{s_j}{(s_i + \lambda)^2 (s_j + \lambda)^2} \left(\frac{-\lambda}{s_j + \lambda} + \frac{\lambda}{s_i + \lambda} - 1 \right)^2 \\ &= \frac{r^2 \sigma^2}{pn} \left(\sum_{i=1}^n \sum_{j=1}^n \frac{s_j}{(s_i + \lambda)^2 (s_j + \lambda)^2} \left(\frac{-\lambda}{s_j + \lambda} + \frac{\lambda}{s_i + \lambda} - 1 \right)^2 + \sum_{i=n+1}^p \sum_{j=1}^n \frac{s_j}{\lambda^2 (s_j + \lambda)^2 (s_j + \lambda)^2} \right) \\ &\xrightarrow{\lambda \rightarrow 0} \frac{r^2 \sigma^2}{pn} \left(\sum_{i=1}^n \sum_{j=1}^n \frac{1}{s_i^2 s_j} + (p-n) \sum_{j=1}^n \frac{1}{s_j^3} \right) \\ &= \frac{r^2 \sigma^2 p}{n} \left[\left(\frac{1}{p} \sum_{i=1}^n \frac{1}{s_i^2} \right) \left(\frac{1}{p} \sum_{i=1}^n \frac{1}{s_i} \right) + \frac{p-n}{p} \left(\frac{1}{p} \sum_{i=1}^n \frac{1}{s_i^3} \right) \right] \\ &\xrightarrow{n, p \rightarrow \infty} r^2 \sigma^2 \gamma \left[\frac{1}{\gamma(\gamma-1)^4} + \left(1 - \frac{1}{\gamma} \right) \frac{\gamma+1}{(\gamma-1)^5} \right] = \frac{r^2 \sigma^2 (\gamma+2)}{(\gamma-1)^4}, \\ \lambda^{-2} \frac{\sigma^4}{n^2 \lambda^2} A_3 &= \frac{\sigma^4}{2n^2} \sum_{i,j=1}^p \frac{s_i s_j (s_i - s_j)^2}{(s_i + \lambda)^4 (s_j + \lambda)^4} = \frac{\sigma^4}{2n^2} \sum_{i,j=1}^n \frac{s_i s_j (s_i - s_j)^2}{(s_i + \lambda)^4 (s_j + \lambda)^4} \\ &\xrightarrow{\lambda \rightarrow 0} \frac{\sigma^4}{2n^2} \sum_{i,j=1}^n \frac{(s_i - s_j)^2}{s_i^3 s_j^3} \\ &= \frac{\sigma^4 p^2}{n^2} \left(\left(\frac{1}{p} \sum_{i=1}^n \frac{1}{s_i} \right) \left(\frac{1}{p} \sum_{i=1}^n \frac{1}{s_i^3} \right) - \left(\frac{1}{p} \sum_{i=1}^n \frac{1}{s_i^2} \right)^2 \right) \end{aligned}$$

$$\xrightarrow{n,p \rightarrow \infty} \sigma^4 \gamma^2 \left(\frac{\gamma + 1}{\gamma(\gamma - 1)^6} - \frac{1}{(\gamma - 1)^6} \right) = \frac{\sigma^4 \gamma}{(\gamma - 1)^6},$$

where we used the following facts about the MP law's negative moments when $\gamma > 1$:

$$\frac{1}{p} \sum_{i=1}^p \frac{1}{s_i} \rightarrow \frac{1}{\gamma(\gamma - 1)}, \quad \frac{1}{p} \sum_{i=1}^p \frac{1}{s_i^2} \rightarrow \frac{1}{(\gamma - 1)^3}, \quad \frac{1}{p} \sum_{i=1}^p \frac{1}{s_i^3} \rightarrow \frac{\gamma + 1}{(\gamma - 1)^5}.$$

Similarly, we have

$$\begin{aligned} \lambda^{-2}(\bar{R}_X(\lambda) + \bar{R}_{\text{pd},X}(\lambda) - 2\bar{C}_X(\lambda)) &= \frac{r^2}{p} \sum_{i=1}^p \frac{s_i^2}{(s_i + \lambda)^4} + \frac{\sigma^2}{n} \sum_{i=1}^p \frac{s_i}{(s_i + \lambda)^4} \\ &= \frac{r^2}{p} \sum_{i=1}^n \frac{s_i^2}{(s_i + \lambda)^4} + \frac{\sigma^2}{n} \sum_{i=1}^n \frac{s_i}{(s_i + \lambda)^4} \\ &\xrightarrow{\lambda \rightarrow 0} \frac{r^2}{p} \sum_{i=1}^n \frac{1}{s_i^2} + \frac{\sigma^2}{n} \sum_{i=1}^n \frac{1}{s_i^3} \\ &\xrightarrow{n,p \rightarrow \infty} \frac{r^2}{(\gamma - 1)^3} + \frac{\sigma^2 \gamma(\gamma + 1)}{(\gamma - 1)^5}. \end{aligned}$$

Thus, for $\gamma > 1$, we have

$$\lim_{\lambda \rightarrow 0} \mathcal{R}_{\text{sd}}^*(\lambda) = \sigma^2 + \frac{r^4(\gamma - 1)^4 + r^2 \sigma^2 \gamma(\gamma + 2)(\gamma - 1)^2 + \sigma^4 \gamma^2}{r^2 \gamma(\gamma - 1)^3 + \sigma^2 \gamma^2(\gamma^2 - 1)}.$$

Finally, the ratios in Proposition 3.3 are obtained by plugging in the asymptotic limits of the corresponding risks using Lemmas B.10–B.12. \square

Lemma B.13. Assume the random variable $X \sim \text{MP}(\gamma)$ with $\gamma < 1$. Then we have

$$\mathbb{E}[X^{-1}] = \frac{1}{1 - \gamma}, \quad \mathbb{E}[X^{-2}] = \frac{1}{(1 - \gamma)^3}, \quad \mathbb{E}[X^{-3}] = \frac{1 + \gamma}{(1 - \gamma)^5}.$$

Proof. Denote $b = (1 + \sqrt{\gamma})^2$ and $a = (1 - \sqrt{\gamma})^2$. Let $R := \sqrt{(b - x)(x - a)}$, from Equation 2.265 of Gradshteyn and Ryzhik (2007), for $m \geq 2$ we have

$$\int_a^b \frac{\sqrt{R}}{x^m} dx = \frac{(2m - 5)(b + a)}{2(m - 1)ba} \int_a^b \frac{\sqrt{R}}{x^{m-1}} dx - \frac{m - 4}{(m - 1)ba} \int_a^b \frac{\sqrt{R}}{x^{m-2}} dx. \quad (73)$$

Recall the density of the MP law $f(x) = \frac{\sqrt{R}}{2\pi\gamma x}$ for $x \in (a, b)$. For $k \geq 1$, using the above identity, we have

$$\begin{aligned} \mathbb{E}(X^{-k}) &= \int_a^b \frac{1}{x^k} f(x) dx = \frac{1}{2\pi\gamma} \int_a^b \frac{\sqrt{R}}{x^{k+1}} dx \\ &= \frac{(2k - 3)(1 + \gamma)}{k(1 - \gamma)^2} \mathbb{E}(X^{-k+1}) - \frac{k - 3}{k(1 - \gamma)^2} \mathbb{E}(X^{-k+2}). \end{aligned}$$

Plug-in $k = 1, 2, 3$ and noting that $\mathbb{E}(X) = 1$, we obtain the desired results. \square

B.8.2 Second proof

For the second approach, we will specialize the general deterministic equivalent in Theorem 3.1 to the isotropic case.

Proof of Proposition 3.3. We specialize Theorem 3.1 to $\Sigma = I_p$. In this case, we have $G = (I_p + \kappa I_p)^{-1} = (1 + \kappa)^{-1} I_p$, and thus

$$t_k = \gamma \overline{\text{tr}}(I_p^2 G^k) = \frac{\gamma}{(1 + \kappa)^k}, \quad q_k = \beta^\top G^k I_p \beta = \frac{\|\beta\|_2^2}{(1 + \kappa)^k}.$$

With $\beta \sim \mathcal{N}(0, (r^2/p)I_p)$, $\|\beta\|_2^2 \rightarrow r^2$ in probability, so $q_k \rightarrow r^2/(1 + \kappa)^k$ in probability. The fixed-point equation becomes

$$\kappa = \lambda + \frac{\gamma \kappa}{1 + \kappa}, \quad (74)$$

whose unique nonnegative solution is

$$\kappa(\lambda) = \frac{1}{2} \left((\lambda + \gamma - 1) + \sqrt{(\lambda + \gamma - 1)^2 + 4\lambda} \right). \quad (75)$$

In particular,

$$\kappa(\lambda) \xrightarrow{\lambda \rightarrow \infty} \infty, \quad \kappa(\lambda) \xrightarrow{\lambda \rightarrow 0} (\gamma - 1)_+.$$

Let $\overline{\mathcal{R}}_{\text{sd}}^*(\lambda) := \mathcal{R}_{\text{sd}}^*(\lambda) - \sigma^2$ denote the excess prediction risk. Substituting the isotropic specializations into the general formula and simplifying yields the explicit rational form (as a function of $\kappa = \kappa(\lambda)$):

$$\overline{\mathcal{R}}_{\text{sd}}^*(\lambda) = \frac{\gamma \left(r^4 \kappa^4 + r^2 \sigma^2 (1 + \kappa)^4 + \gamma^2 \sigma^2 (r^2 + \sigma^2) - 2\gamma r^2 \sigma^2 (2\kappa + 1) \right)}{\left((1 + \kappa)^2 - \gamma \right) \left(r^2 \left((1 + \gamma)(1 + \kappa)^2 - 4\gamma(1 + \kappa) + \gamma(1 + \gamma) \right) + \sigma^2 \gamma \left((1 + \kappa)^2 + \gamma \right) \right)}. \quad (76)$$

Case 1: $\lambda \rightarrow \infty$. Using $\kappa(\lambda) \rightarrow \infty$, the leading terms in (76) give

$$\lim_{\lambda \rightarrow \infty} \overline{\mathcal{R}}_{\text{sd}}^*(\lambda) = \frac{\gamma r^2 (r^2 + \sigma^2)}{r^2 (\gamma + 1) + \gamma \sigma^2},$$

and adding back σ^2 and writing in terms of SNR yields the stated $\lambda \rightarrow \infty$ limit.

Case 2: $\lambda \rightarrow 0$. If $\gamma \in (0, 1)$ then $\kappa(\lambda) \rightarrow 0$; substituting $\kappa = 0$ in (76) yields the stated $\gamma < 1$ limit. If $\gamma \in (1, \infty)$ then $\kappa(\lambda) \rightarrow \gamma - 1$; substituting $\kappa = \gamma - 1$ yields the stated $\gamma > 1$ limit. Adding back σ^2 and expressing in terms of SNR completes the proof. \square

C Proofs in Section 4

C.1 Preliminaries

Fix $\lambda > 0$. Let $X \in \mathbb{R}^{n \times p}$ and $y \in \mathbb{R}^n$, and define the sample covariance $\widehat{\Sigma} := X^\top X/n$ and ridge resolvent $Q_\lambda := (\widehat{\Sigma} + \lambda I_p)^{-1}$. Define the ridge hat (smoother) matrix

$$H_\lambda := X Q_\lambda X^\top / n \in \mathbb{R}^{n \times n}.$$

The ridge teacher coefficient and fitted values are

$$\beta_\lambda := Q_\lambda X^\top y/n, \quad \hat{y}_\lambda := X\beta_\lambda = H_\lambda y.$$

Define the pure-distilled ridge coefficient (ridge refit on pseudo-labels \hat{y}_λ)

$$\beta_{\text{pd},\lambda} := Q_\lambda X^\top \hat{y}_\lambda/n, \quad \hat{y}_{\text{pd},\lambda} := X\beta_{\text{pd},\lambda}.$$

Since $\hat{y}_\lambda = H_\lambda y$ and $H_\lambda = XQ_\lambda X^\top/n$, a direct calculation yields the identity used repeatedly below:

$$\hat{y}_{\text{pd},\lambda} = H_\lambda^2 y. \quad (77)$$

Thus both the teacher and PD predictors are linear smoothers in y , with smoothing matrices H_λ and H_λ^2 , respectively. Consequently, their degrees of freedom equal the traces of these matrices:

$$\text{df}_\lambda = \text{tr}(H_\lambda), \quad \text{df}_{\text{pd},\lambda} = \text{tr}(H_\lambda^2).$$

Recall the GCV residuals and risk/correlation estimators defined in Equations (17) and (18):

$$\hat{r}_\lambda := \frac{y - \hat{y}_\lambda}{1 - \text{df}_\lambda/n}, \quad \hat{r}_{\text{pd},\lambda} := \frac{y - \hat{y}_{\text{pd},\lambda}}{1 - \text{df}_{\text{pd},\lambda}/n},$$

and

$$\hat{R}(\lambda) := \frac{\|\hat{r}_\lambda\|_2^2}{n}, \quad \hat{R}_{\text{pd}}(\lambda) := \frac{\|\hat{r}_{\text{pd},\lambda}\|_2^2}{n}, \quad \hat{C}(\lambda) := \frac{\langle \hat{r}_\lambda, \hat{r}_{\text{pd},\lambda} \rangle}{n}.$$

Let (x_0, y_0) be an independent test pair distributed as in Assumption A and independent of the training data. Define the out-of-sample errors

$$e_\lambda := y_0 - x_0^\top \beta_\lambda, \quad e_{\text{pd},\lambda} := y_0 - x_0^\top \beta_{\text{pd},\lambda},$$

and the corresponding conditional (oracle) quantities

$$R(\lambda) := \mathbb{E}[e_\lambda^2 | X, y], \quad R_{\text{pd}}(\lambda) := \mathbb{E}[e_{\text{pd},\lambda}^2 | X, y], \quad C(\lambda) := \mathbb{E}[e_\lambda e_{\text{pd},\lambda} | X, y].$$

Finally define

$$D(\lambda) := R(\lambda) + R_{\text{pd}}(\lambda) - 2C(\lambda) = \mathbb{E}[(e_\lambda - e_{\text{pd},\lambda})^2 | X, y] \geq 0. \quad (78)$$

We will use the following ridge functional estimation result (Theorem 4 of Patil et al. (2022b)):

Lemma C.1 (Ridge functional estimation via GCV/LOOCV). Assume the conditions of Assumption A and $p/n \rightarrow \gamma$. Fix $\lambda > 0$. Let $t : \mathbb{R} \rightarrow \mathbb{R}$ be continuous and satisfy a quadratic growth condition. Define $T_\lambda := \mathbb{E}[t(e_\lambda) | X, y]$. Let T_λ^{gcv} and T_λ^{loo} be the plug-in estimators based on the usual ridge GCV and LOOCV residuals. Then

$$T_\lambda^{\text{gcv}} - T_\lambda \xrightarrow{\text{P}} 0, \quad T_\lambda^{\text{loo}} - T_\lambda \xrightarrow{\text{P}} 0.$$

C.2 Helper Results (Risk and Correlation Components Consistency) for the Proof of Theorem 4.1

Lemma C.2 (Consistency of the ridge-risk estimator). Under Assumption A and $p/n \rightarrow \gamma$, for each fixed $\lambda > 0$,

$$\widehat{R}(\lambda) - R(\lambda) \xrightarrow{P} 0.$$

Proof. This simply follows from Lemma C.1 with $t(z) = z^2$. Then $T_\lambda = \mathbb{E}[e_\lambda^2 | X, y] = R(\lambda)$, and the ridge GCV plug-in estimator is precisely $\widehat{R}(\lambda) = \|\widehat{r}_\lambda\|_2^2/n$. \square

Lemma C.3 (Consistency of the pure-distilled risk and correlation estimators). Under Assumption A and $p/n \rightarrow \gamma$, for each fixed $\lambda > 0$,

$$\widehat{R}_{\text{pd}}(\lambda) - R_{\text{pd}}(\lambda) \xrightarrow{P} 0, \quad \widehat{C}(\lambda) - C(\lambda) \xrightarrow{P} 0.$$

Proof. We follow the same three-part argument used by Patil et al. (2022b, Proof of Theorem 3): (i) relate the target risks/correlation to leave-one-out (LOO) errors, (ii) show a stability approximation that replaces true LOO errors by diagonal-corrected residuals computed from the full-sample smoother, and (iii) replace the diagonal correction by the GCV trace correction using diagonal–trace equivalence.

Part (i): LOO concentration around $R_{\text{pd}}(\lambda)$ and $C(\lambda)$. Let $\beta_\lambda^{(-i)}$ and $\beta_{\text{pd},\lambda}^{(-i)}$ be the teacher and PD coefficients trained without sample i , and define the corresponding LOO prediction errors

$$e_{i,\lambda}^{\text{loo}} := y_i - x_i^\top \beta_\lambda^{(-i)}, \quad e_{i,\text{pd},\lambda}^{\text{loo}} := y_i - x_i^\top \beta_{\text{pd},\lambda}^{(-i)}.$$

Since (x_i, y_i) is independent of $\sigma\{(x_j, y_j) : j \neq i\}$ and identically distributed to (x_0, y_0) , we have

$$\mathbb{E}[(e_{i,\text{pd},\lambda}^{\text{loo}})^2 | X^{(-i)}, y^{(-i)}] = R_{\text{pd}}^{(-i)}(\lambda), \quad \mathbb{E}[e_{i,\lambda}^{\text{loo}} e_{i,\text{pd},\lambda}^{\text{loo}} | X^{(-i)}, y^{(-i)}] = C^{(-i)}(\lambda),$$

where $R_{\text{pd}}^{(-i)}(\lambda)$ and $C^{(-i)}(\lambda)$ denote the analogues of $R_{\text{pd}}(\lambda)$ and $C(\lambda)$ for the $(n-1)$ -sample-trained predictors.

Using the same martingale-difference decomposition and Burkholder inequality arguments as in Patil et al. (2022b, Supplement S.1.2), together with the moment bounds in Assumption A, one obtains

$$\left| \frac{1}{n} \sum_{i=1}^n (e_{i,\text{pd},\lambda}^{\text{loo}})^2 - \frac{1}{n} \sum_{i=1}^n R_{\text{pd}}^{(-i)}(\lambda) \right| \xrightarrow{P} 0, \quad \left| \frac{1}{n} \sum_{i=1}^n e_{i,\lambda}^{\text{loo}} e_{i,\text{pd},\lambda}^{\text{loo}} - \frac{1}{n} \sum_{i=1}^n C^{(-i)}(\lambda) \right| \xrightarrow{P} 0.$$

Moreover, by LOO stability of ridge and the PD stability in Lemma C.5 (see Part (ii) below),

$$\frac{1}{n} \sum_{i=1}^n (R_{\text{pd}}^{(-i)}(\lambda) - R_{\text{pd}}(\lambda)) \xrightarrow{P} 0, \quad \frac{1}{n} \sum_{i=1}^n (C^{(-i)}(\lambda) - C(\lambda)) \xrightarrow{P} 0,$$

because (conditionally on the training data) the test risks are Lipschitz in the coefficient vector: for instance,

$$|R_{\text{pd}}^{(-i)}(\lambda) - R_{\text{pd}}(\lambda)| = |\mathbb{E}[(x_0^\top (\beta_{\text{pd},\lambda}^{(-i)} - \beta_{\text{pd},\lambda}))^2 | X, y]| \leq \|\Sigma\|_{\text{op}} \|\beta_{\text{pd},\lambda}^{(-i)} - \beta_{\text{pd},\lambda}\|_2^2.$$

Combining yields

$$\frac{1}{n} \sum_{i=1}^n (e_{i,\text{pd},\lambda}^{\text{loo}})^2 \xrightarrow{\text{P}} R_{\text{pd}}(\lambda), \quad \frac{1}{n} \sum_{i=1}^n e_{i,\lambda}^{\text{loo}} e_{i,\text{pd},\lambda}^{\text{loo}} \xrightarrow{\text{P}} C(\lambda). \quad (79)$$

Part (ii): Equivalence of diagonal-corrected residuals and LOO errors. Define the diagonal-corrected residuals based on the full-sample smoothing matrices:

$$\bar{r}_{\lambda,i} := \frac{y_i - (H_{\lambda}y)_i}{1 - (H_{\lambda})_{ii}}, \quad \bar{r}_{\text{pd},i} := \frac{y_i - (H_{\lambda}^2y)_i}{1 - (H_{\lambda}^2)_{ii}}.$$

For ridge, the shortcut formula is exact: $e_{i,\lambda}^{\text{loo}} = \bar{r}_{\lambda,i}$ for all i . For the PD two-stage procedure, leaving out one sample changes both the teacher and student fits; however, by the stability statement in Lemma C.5 (together with ridge stability), these perturbations are negligible on average. In particular, one obtains

$$\frac{1}{n} \sum_{i=1}^n (e_{i,\text{pd},\lambda}^{\text{loo}} - \bar{r}_{\text{pd},i})^2 \xrightarrow{\text{P}} 0. \quad (80)$$

Combining (79) and (80) yields

$$\frac{1}{n} \sum_{i=1}^n \bar{r}_{\text{pd},i}^2 \xrightarrow{\text{P}} R_{\text{pd}}(\lambda), \quad \frac{1}{n} \sum_{i=1}^n \bar{r}_{\lambda,i} \bar{r}_{\text{pd},i} \xrightarrow{\text{P}} C(\lambda), \quad (81)$$

where for the cross term we use $\bar{r}_{\lambda,i} = e_{i,\lambda}^{\text{loo}}$ exactly and Cauchy–Schwarz.

Part (iii): Equivalence of trace-corrected residuals to diagonal-corrected residuals. By Lemma C.6,

$$\max_{1 \leq i \leq n} \left| \frac{1}{1 - \text{tr}(H_{\lambda}^2)/n} - \frac{1}{1 - (H_{\lambda}^2)_{ii}} \right| \xrightarrow{\text{P}} 0,$$

and the analogous ridge diagonal–trace equivalence (see Patil et al. (2022b, Supplement S.1.3)) gives the same statement with H_{λ} in place of H_{λ}^2 . Since $\hat{r}_{\text{pd},\lambda,i} = (y_i - (H_{\lambda}^2y)_i)/(1 - \text{tr}(H_{\lambda}^2)/n)$ and $\hat{r}_{\lambda,i} = (y_i - (H_{\lambda}y)_i)/(1 - \text{tr}(H_{\lambda})/n)$, it follows that

$$\frac{1}{n} \sum_{i=1}^n (\hat{r}_{\text{pd},\lambda,i} - \bar{r}_{\text{pd},i})^2 \xrightarrow{\text{P}} 0, \quad \frac{1}{n} \sum_{i=1}^n (\hat{r}_{\lambda,i} - \bar{r}_{\lambda,i})^2 \xrightarrow{\text{P}} 0.$$

Combining with (81) and using Cauchy–Schwarz yields

$$\frac{1}{n} \sum_{i=1}^n \hat{r}_{\text{pd},\lambda,i}^2 \xrightarrow{\text{P}} R_{\text{pd}}(\lambda), \quad \frac{1}{n} \sum_{i=1}^n \hat{r}_{\lambda,i} \hat{r}_{\text{pd},\lambda,i} \xrightarrow{\text{P}} C(\lambda),$$

which is exactly $\hat{R}_{\text{pd}}(\lambda) \xrightarrow{\text{P}} R_{\text{pd}}(\lambda)$ and $\hat{C}(\lambda) \xrightarrow{\text{P}} C(\lambda)$. □

C.3 Proof of Theorem 4.1

Proof of Theorem 4.1. By Lemmas C.2 and C.3,

$$(\hat{R}(\lambda), \hat{R}_{\text{pd}}(\lambda), \hat{C}(\lambda)) \xrightarrow{\text{P}} (R(\lambda), R_{\text{pd}}(\lambda), C(\lambda)).$$

In particular,

$$\hat{D}(\lambda) := \hat{R}(\lambda) + \hat{R}_{\text{pd}}(\lambda) - 2\hat{C}(\lambda) \xrightarrow{\text{P}} D(\lambda) := R(\lambda) + R_{\text{pd}}(\lambda) - 2C(\lambda).$$

Under the proportional regime, Theorem 3.1 implies that $D(\lambda)$ converges in probability to a deterministic limit $\mathcal{D}(\lambda)$ (given by $\mathcal{R}(\lambda) + \mathcal{R}_{\text{pd}}(\lambda) - 2\mathcal{C}(\lambda)$). In all nondegenerate settings, $\mathcal{D}(\lambda) > 0$ (equivalently, the teacher and PD predictions do not coincide asymptotically), hence $D(\lambda)$ and $\widehat{D}(\lambda)$ are bounded away from 0 in probability. Therefore, the plug-in maps

$$(a, b, c) \mapsto \frac{a - c}{a + b - 2c}, \quad (a, b, c) \mapsto a - \frac{(a - c)^2}{a + b - 2c}$$

are continuous with high probability in a neighborhood of $(R(\lambda), R_{\text{pd}}(\lambda), C(\lambda))$, and the continuous mapping theorem yields

$$\widehat{\xi}^*(\lambda) - \xi^*(\lambda) \xrightarrow{\mathbb{P}} 0, \quad \widehat{R}_{\text{sd}}^*(\lambda) - R_{\text{sd}}^*(\lambda) \xrightarrow{\mathbb{P}} 0.$$

(If $\mathcal{D}(\lambda) = 0$, then the oracle SD objective is asymptotically flat in ξ ; in this case $R_{\text{sd}}^*(\lambda) = R(\lambda)$ and the risk-consistency conclusion remains valid, while $\xi^*(\lambda)$ is not identifiable.) \square

C.4 Technical Lemmas

We establish two technical ingredients for the pure-distilled smoother H_λ^2 : (i) leave-one-out stability and (ii) diagonal-trace equivalence.

C.4.1 Leave-One-Out Stability

For $i \in \{1, \dots, n\}$, let $(X^{(-i)}, y^{(-i)})$ denote the dataset with observation i removed and $\widehat{\Sigma}^{(-i)} := (X^{(-i)})^\top X^{(-i)} / (n-1)$. Let $Q_\lambda^{(-i)} := (\widehat{\Sigma}^{(-i)} + \lambda I_p)^{-1}$ and $\beta_\lambda^{(-i)} := Q_\lambda^{(-i)} (X^{(-i)})^\top y^{(-i)} / (n-1)$. Define the PD coefficient trained without i by

$$\beta_{\text{pd}, \lambda}^{(-i)} := Q_\lambda^{(-i)} (X^{(-i)})^\top \widehat{y}_\lambda^{(-i)} / (n-1), \quad \widehat{y}_\lambda^{(-i)} := X^{(-i)} \beta_\lambda^{(-i)}.$$

Lemma C.4 (Operator-Lipschitz property of $M_\lambda(\cdot)$). For $\lambda > 0$, define $M_\lambda(A) := A(A + \lambda I)^{-1} = I - \lambda(A + \lambda I)^{-1}$. Then for any symmetric $A, B \succeq 0$,

$$\|M_\lambda(A) - M_\lambda(B)\|_{\text{op}} \leq \frac{1}{\lambda} \|A - B\|_{\text{op}}.$$

Proof. Using $M_\lambda(A) = I - \lambda(A + \lambda I)^{-1}$ and the resolvent identity,

$$M_\lambda(A) - M_\lambda(B) = -\lambda[(A + \lambda I)^{-1} - (B + \lambda I)^{-1}] = -\lambda(A + \lambda I)^{-1}(B - A)(B + \lambda I)^{-1}.$$

Taking operator norms and using $\|(A + \lambda I)^{-1}\|_{\text{op}} \leq 1/\lambda$ and $\|(B + \lambda I)^{-1}\|_{\text{op}} \leq 1/\lambda$ yields the claim. \square

Lemma C.5 (Average stability of the pure-distilled coefficient). Under Assumption A and $p/n \rightarrow \gamma$, for each fixed $\lambda > 0$,

$$\frac{1}{n} \sum_{i=1}^n \|\beta_{\text{pd}, \lambda} - \beta_{\text{pd}, \lambda}^{(-i)}\|_2^2 \xrightarrow{\mathbb{P}} 0.$$

Proof. Write $\beta_{\text{pd}, \lambda} = M_\lambda(\widehat{\Sigma})\beta_\lambda$ and $\beta_{\text{pd}, \lambda}^{(-i)} = M_\lambda(\widehat{\Sigma}^{(-i)})\beta_\lambda^{(-i)}$. Then

$$\beta_{\text{pd}, \lambda} - \beta_{\text{pd}, \lambda}^{(-i)} = M_\lambda(\widehat{\Sigma})(\beta_\lambda - \beta_\lambda^{(-i)}) + (M_\lambda(\widehat{\Sigma}) - M_\lambda(\widehat{\Sigma}^{(-i)}))\beta_\lambda^{(-i)}.$$

Using $(a + b)^2 \leq 2a^2 + 2b^2$ and $\|Au\|_2 \leq \|A\|_{\text{op}}\|u\|_2$, we obtain

$$\frac{1}{n} \sum_{i=1}^n \|\beta_{\text{pd},\lambda} - \beta_{\text{pd},\lambda}^{(-i)}\|_2^2 \leq \frac{2}{n} \sum_{i=1}^n \|M_\lambda(\widehat{\Sigma})\|_{\text{op}}^2 \|\beta_\lambda - \beta_\lambda^{(-i)}\|_2^2 + \frac{2}{n} \sum_{i=1}^n \|(M_\lambda(\widehat{\Sigma}) - M_\lambda(\widehat{\Sigma}^{(-i)}))\beta_\lambda^{(-i)}\|_2^2. \quad (82)$$

Since $\|M_\lambda(\cdot)\|_{\text{op}} \leq 1$, the first term in (82) is bounded by $\frac{2}{n} \sum_i \|\beta_\lambda - \beta_\lambda^{(-i)}\|_2^2$, which converges to 0 in probability by the ridge stability result of [Patil et al. \(2022b\)](#) (used there to control LOOCV/GCV).

For the second term, use the exact identity $M_\lambda(A) = I - \lambda(A + \lambda I)^{-1}$ to write

$$M_\lambda(\widehat{\Sigma}) - M_\lambda(\widehat{\Sigma}^{(-i)}) = -\lambda(Q_\lambda - Q_\lambda^{(-i)}).$$

By the resolvent identity,

$$Q_\lambda - Q_\lambda^{(-i)} = Q_\lambda(\widehat{\Sigma}^{(-i)} - \widehat{\Sigma})Q_\lambda^{(-i)}.$$

Moreover, $\widehat{\Sigma} = \frac{n-1}{n}\widehat{\Sigma}^{(-i)} + \frac{1}{n}x_i x_i^\top$, so $\widehat{\Sigma}^{(-i)} - \widehat{\Sigma} = \frac{1}{n}(\widehat{\Sigma}^{(-i)} - x_i x_i^\top)$. Define $v_i := Q_\lambda^{(-i)}\beta_\lambda^{(-i)}$. Then

$$(M_\lambda(\widehat{\Sigma}) - M_\lambda(\widehat{\Sigma}^{(-i)}))\beta_\lambda^{(-i)} = \frac{\lambda}{n} Q_\lambda(x_i x_i^\top - \widehat{\Sigma}^{(-i)})v_i.$$

Since $\|Q_\lambda\|_{\text{op}} \leq 1/\lambda$, it suffices to control $n^{-1} \sum_i \|(x_i x_i^\top - \widehat{\Sigma}^{(-i)})v_i\|_2^2/n^2$. Condition on $\sigma\{(x_j, y_j) : j \neq i\}$, under which v_i is deterministic and independent of x_i . Using $\|x_i x_i^\top v_i\|_2^2 = \|x_i\|_2^2 (x_i^\top v_i)^2$ and bounded-spectrum Σ from Assumption A,

$$\mathbb{E}\left[\|x_i x_i^\top v_i\|_2^2 \mid X^{(-i)}, y^{(-i)}\right] = \mathbb{E}[\|x_i\|_2^2 (x_i^\top v_i)^2 \mid X^{(-i)}, y^{(-i)}] \lesssim p \|v_i\|_2^2,$$

while $\|\widehat{\Sigma}^{(-i)}v_i\|_2^2 \lesssim \|v_i\|_2^2$ since $\|\widehat{\Sigma}^{(-i)}\|_{\text{op}} = O_p(1)$ for proportional random design. Therefore,

$$\mathbb{E}\left[\|(x_i x_i^\top - \widehat{\Sigma}^{(-i)})v_i\|_2^2 \mid X^{(-i)}, y^{(-i)}\right] \lesssim p \|v_i\|_2^2,$$

and hence

$$\mathbb{E}\left[\|(M_\lambda(\widehat{\Sigma}) - M_\lambda(\widehat{\Sigma}^{(-i)}))\beta_\lambda^{(-i)}\|_2^2\right] \lesssim \frac{p}{n^2} \mathbb{E}[\|v_i\|_2^2].$$

Finally, $\|v_i\|_2 = \|Q_\lambda^{(-i)}\beta_\lambda^{(-i)}\|_2 \leq \|Q_\lambda^{(-i)}\|_{\text{op}}\|\beta_\lambda^{(-i)}\|_2 \leq \lambda^{-1}\|\beta_\lambda^{(-i)}\|_2$, and for fixed $\lambda > 0$ we have $\|\beta_\lambda^{(-i)}\|_2 = O_p(1)$ uniformly in i under Assumption A. Since $p/n \rightarrow \gamma$, we have $p/n^2 = O(1/n)$, and averaging over i yields

$$\frac{1}{n} \sum_{i=1}^n \|(M_\lambda(\widehat{\Sigma}) - M_\lambda(\widehat{\Sigma}^{(-i)}))\beta_\lambda^{(-i)}\|_2^2 \xrightarrow{p} 0.$$

Together with the first term in (82), this proves the claim. \square

C.4.2 Diagonal-Trace Equivalence

Let $S := XX^\top/n \in \mathbb{R}^{n \times n}$ and $G_\lambda := (S + \lambda I_n)^{-1}$. Recall the standard identity

$$H_\lambda = S(S + \lambda I_n)^{-1} = I_n - \lambda G_\lambda,$$

so

$$H_\lambda^2 = (I_n - \lambda G_\lambda)^2 = I_n - 2\lambda G_\lambda + \lambda^2 G_\lambda^2. \quad (83)$$

Lemma C.6 (Diagonal-trace equivalence for H_λ^2). Fix $\lambda > 0$. Under Assumption A and $p/n \rightarrow \gamma$,

$$\max_{1 \leq i \leq n} \left| (H_\lambda^2)_{ii} - \frac{1}{n} \operatorname{tr}(H_\lambda^2) \right| \xrightarrow{p} 0,$$

and hence

$$\max_{1 \leq i \leq n} \left| \frac{1}{1 - \operatorname{tr}(H_\lambda^2)/n} - \frac{1}{1 - (H_\lambda^2)_{ii}} \right| \xrightarrow{p} 0.$$

Proof. By (83), for each i ,

$$(H_\lambda^2)_{ii} - \frac{1}{n} \operatorname{tr}(H_\lambda^2) = -2\lambda \left((G_\lambda)_{ii} - \frac{1}{n} \operatorname{tr}(G_\lambda) \right) + \lambda^2 \left((G_\lambda^2)_{ii} - \frac{1}{n} \operatorname{tr}(G_\lambda^2) \right).$$

Thus it suffices to prove

$$\max_i \left| (G_\lambda)_{ii} - \frac{1}{n} \operatorname{tr}(G_\lambda) \right| \xrightarrow{p} 0, \quad \max_i \left| (G_\lambda^2)_{ii} - \frac{1}{n} \operatorname{tr}(G_\lambda^2) \right| \xrightarrow{p} 0.$$

1. *Control for G_λ :* Since $1 - (H_\lambda)_{ii} = \lambda(G_\lambda)_{ii}$ and $1 - \operatorname{tr}(H_\lambda)/n = \lambda \operatorname{tr}(G_\lambda)/n$, the diagonal-trace equivalence for ridge (see Patil et al. (2022b, Supplement S.1.3)) implies $\max_i |(G_\lambda)_{ii} - \operatorname{tr}(G_\lambda)/n| \xrightarrow{p} 0$.
2. *Control for G_λ^2 :* Fix any $t > 0$ and define $G_{\lambda+t} := (S + (\lambda + t)I_n)^{-1}$. Since G_λ and $G_{\lambda+t}$ are both functions of S , they commute, and we have the exact identity

$$G_\lambda^2 = \frac{G_\lambda - G_{\lambda+t}}{t} + t G_{\lambda+t} G_\lambda^2. \quad (84)$$

Taking diagonal-trace differences and maxima gives

$$\max_i \left| (G_\lambda^2)_{ii} - \frac{1}{n} \operatorname{tr}(G_\lambda^2) \right| \leq \frac{1}{t} \max_i \left| (G_\lambda)_{ii} - \frac{1}{n} \operatorname{tr}(G_\lambda) \right| + \frac{1}{t} \max_i \left| (G_{\lambda+t})_{ii} - \frac{1}{n} \operatorname{tr}(G_{\lambda+t}) \right| + 2t \|G_{\lambda+t} G_\lambda^2\|_{\text{op}}.$$

Since $\|G_\lambda\|_{\text{op}} \leq 1/\lambda$ and $\|G_{\lambda+t}\|_{\text{op}} \leq 1/(\lambda + t)$, the last term is bounded by $2t/\lambda^3$. Letting $n \rightarrow \infty$ and using Step 1 at λ and at $\lambda + t$ yields

$$\limsup_{n \rightarrow \infty} \max_i \left| (G_\lambda^2)_{ii} - \frac{1}{n} \operatorname{tr}(G_\lambda^2) \right| \leq \frac{2t}{\lambda^3} \quad \text{in probability.}$$

Since $t > 0$ was arbitrary, sending $t \rightarrow 0$ gives the desired diagonal-trace control for G_λ^2 .

This proves the first claim. For the second claim, note that $0 \preceq H_\lambda^2 \preceq I_n$, so $1 - (H_\lambda^2)_{ii} > 0$ and $1 - \operatorname{tr}(H_\lambda^2)/n > 0$. Moreover, for fixed $\lambda > 0$ these denominators are bounded away from 0 with high probability. Hence

$$\left| \frac{1}{1 - \operatorname{tr}(H_\lambda^2)/n} - \frac{1}{1 - (H_\lambda^2)_{ii}} \right| = \frac{|(H_\lambda^2)_{ii} - \operatorname{tr}(H_\lambda^2)/n|}{(1 - \operatorname{tr}(H_\lambda^2)/n)(1 - (H_\lambda^2)_{ii})},$$

and the second claim follows from the first. \square

D Proofs in Section 5

D.1 Proofs and Details in Section 5.1

D.1.1 Proof of Equation (23)

Proof of Equation (23). Fix $k \geq 0$ and $\xi \in \mathbb{R}$. Recall that the round- $(k+1)$ SD predictor $f_{\text{sd},\lambda,\xi}^{(k+1)}$ is defined as:

$$f_{\text{sd},\lambda,\xi}^{(k+1)}(x) := x^\top \operatorname{argmin}_{\beta \in \mathbb{R}^p} \left\{ (1-\xi) \|y^{(k)} - X\beta\|_2^2/n + \xi \|\widehat{y}_\lambda^{(k)} - X\beta\|_2^2/n + \lambda \|\beta\|_2^2 \right\}. \quad (85)$$

Expand the objective in (85) (dropping constants independent of β):

$$\begin{aligned} & (1-\xi) \|y^{(k)} - X\beta\|_2^2/n + \xi \|\widehat{y}_\lambda^{(k)} - X\beta\|_2^2/n + \lambda \|\beta\|_2^2 \\ & \equiv \|(1-\xi)y^{(k)} + \xi \widehat{y}_\lambda^{(k)} - X\beta\|_2^2/n + \lambda \|\beta\|_2^2 = \|y_{\lambda,\xi}^{(k+1)} - X\beta\|_2^2/n + \lambda \|\beta\|_2^2, \end{aligned}$$

where $y_{\lambda,\xi}^{(k+1)} = (1-\xi)y^{(k)} + \xi \widehat{y}_\lambda^{(k)}$ is exactly the mixed-label vector in (22). Thus $f_{\text{sd},\lambda,\xi}^{(k+1)}$ coincides with ridge regression trained on $(X, y_{\lambda,\xi}^{(k+1)})$ with penalty λ . Since ridge is linear in its response vector, we have

$$f_{\text{sd},\lambda,\xi}^{(k+1)} = (1-\xi) f_\lambda^{(k)} + \xi f_{\text{pd},\lambda}^{(k)},$$

because $f_\lambda^{(k)}$ is ridge on response $y^{(k)}$ and $f_{\text{pd},\lambda}^{(k)}$ is ridge on response $\widehat{y}_\lambda^{(k)} = f_\lambda^{(k)}(X)$. This is (23). \square

D.1.2 Proof of Proposition 5.1

Proof of Proposition 5.1. Fix $k \geq 0$. By Equation (23), the family $\{f_{\text{sd},\lambda,\xi}^{(k+1)} : \xi \in \mathbb{R}\}$ is the affine path between $f_\lambda^{(k)}$ and $f_{\text{pd},\lambda}^{(k)}$.

Monotonicity. Since $\xi = 0$ is feasible and $f_{\text{sd},\lambda,0}^{(k+1)} = f_\lambda^{(k)}$, we have

$$R_{k+1}(\lambda) = \min_{\xi \in \mathbb{R}} R(f_{\text{sd},\lambda,\xi}^{(k+1)}) \leq R(f_{\text{sd},\lambda,0}^{(k+1)}) = R(f_\lambda^{(k)}) = R_k(\lambda).$$

Closed forms. When $D_k(\lambda) > 0$, applying Proposition 2.1 to the affine path between $f_\lambda^{(k)}$ and $f_{\text{pd},\lambda}^{(k)}$ yields the displayed round-wise closed forms (with (R, C, D) replaced by (R_k, C_k, D_k)). Likewise, (24) follows by applying Theorem 2.2 at round k , again with $(f_\lambda, f_{\text{pd},\lambda}, R, C, D)$ replaced by $(f_\lambda^{(k)}, f_{\text{pd},\lambda}^{(k)}, R_k, C_k, D_k)$, and interpreting $R'_k(\lambda)$ as the derivative with respect to the ridge penalty while holding the base labels $y^{(k)}$ fixed. \square

D.2 Proofs and Details in Section 5.2

D.2.1 Fresh-X Mixed-Loss Student Representation

Recall the fresh- X mixed-loss SD predictor $f_{\text{sd},\lambda,\xi}^{\text{frmix}}$ from (26). Writing $f_{\text{sd},\lambda,\xi}^{\text{frmix}}(x) = x^\top \beta_{\text{sd},\lambda}^{\text{frmix}}(\xi)$, we derive an explicit closed form for $\beta_{\text{sd},\lambda}^{\text{frmix}}(\xi)$ and show that, unlike the same- X setting, the map $\xi \mapsto f_{\text{sd},\lambda,\xi}^{\text{frmix}}$ is generally not an affine path.

Let $\widehat{\Sigma} := X^\top X/n$, $\widetilde{\Sigma} := \widetilde{X}^\top \widetilde{X}/m$, $\widehat{v} := X^\top y/n$, and define the ridge resolvents:

$$Q_\lambda := (\widehat{\Sigma} + \lambda I_p)^{-1}, \quad \widetilde{Q}_\lambda := (\widetilde{\Sigma} + \lambda I_p)^{-1}.$$

The ridge teacher coefficient is $\beta_\lambda = Q_\lambda \widehat{v}$. The fresh- X PD refit trained on $(\widetilde{X}, \widetilde{y}_\lambda)$ with $\widetilde{y}_\lambda = \widetilde{X} \beta_\lambda$ has coefficient

$$\beta_{\text{pd},\lambda}^{\text{fr}} := (\widetilde{\Sigma} + \lambda I_p)^{-1} \frac{1}{m} \widetilde{X}^\top \widetilde{y}_\lambda = (\widetilde{\Sigma} + \lambda I_p)^{-1} \widetilde{\Sigma} \beta_\lambda = \widetilde{Q}_\lambda \widetilde{\Sigma} \beta_\lambda.$$

(Here $\widetilde{Q}_\lambda \widetilde{\Sigma} = \widetilde{\Sigma} \widetilde{Q}_\lambda$ since both are functions of $\widetilde{\Sigma}$.)

Lemma D.1 (Fresh- X mixed-loss student representation). Fix $\lambda > 0$ and let $\xi \in \mathbb{R}$ be such that the mixed-loss objective (26) is strictly convex in β , equivalently $A_\xi := (1 - \xi) \widehat{\Sigma} + \xi \widetilde{\Sigma} + \lambda I_p \succ 0$. (In particular, this holds for all $\xi \in [0, 1]$.) Then the coefficient of the fresh- X mixed-loss student satisfies

$$\beta_{\text{sd},\lambda}^{\text{frmix}}(\xi) = A_\xi^{-1} ((1 - \xi) \widehat{v} + \xi \widetilde{\Sigma} \beta_\lambda) = ((1 - \xi) Q_\lambda^{-1} + \xi \widetilde{Q}_\lambda^{-1})^{-1} ((1 - \xi) \widehat{v} + \xi \widetilde{\Sigma} \beta_\lambda). \quad (86)$$

Moreover, for $\xi \notin \{0, 1\}$, writing $S := Q_\lambda \widetilde{Q}_\lambda^{-1}$ (so $S^{-1} = \widetilde{Q}_\lambda Q_\lambda^{-1}$), we have the matrix-weighted decomposition

$$\beta_{\text{sd},\lambda}^{\text{frmix}}(\xi) = \left(I + \frac{\xi}{1-\xi} S \right)^{-1} \beta_\lambda + \left(I + \frac{1-\xi}{\xi} S^{-1} \right)^{-1} \beta_{\text{pd},\lambda}^{\text{fr}}, \quad (87)$$

with the boundary cases $\xi \in \{0, 1\}$ obtained by continuity.

Proof. Expanding (26) and setting the gradient with respect to β to zero yields the normal equations

$$((1 - \xi) \widehat{\Sigma} + \xi \widetilde{\Sigma} + \lambda I_p) \beta = (1 - \xi) \widehat{v} + \xi \widetilde{\Sigma} \beta_\lambda,$$

which gives the first identity in (86); the second follows from $Q_\lambda^{-1} = \widehat{\Sigma} + \lambda I_p$ and $\widetilde{Q}_\lambda^{-1} = \widetilde{\Sigma} + \lambda I_p$. For (87), substitute $\widehat{v} = Q_\lambda^{-1} \beta_\lambda$ and $\widetilde{\Sigma} \beta_\lambda = \widetilde{Q}_\lambda^{-1} \beta_{\text{pd},\lambda}^{\text{fr}}$ into (86), then split the two terms and factor Q_λ^{-1} (respectively $\widetilde{Q}_\lambda^{-1}$) from the left. \square

The upshot of Lemma D.1 is that, unless $S = I_p$ (equivalently $Q_\lambda = \widetilde{Q}_\lambda$), the map $\xi \mapsto \beta_{\text{sd},\lambda}^{\text{frmix}}(\xi)$ is not a scalar affine combination $(1 - \xi) \beta_\lambda + \xi \beta_{\text{pd},\lambda}^{\text{fr}}$. In particular, $\xi \mapsto f_{\text{sd},\lambda,\xi}^{\text{frmix}}$ is generally not an affine path, unlike the same- X case in (4).

D.2.2 Same- X SD under Isotropic Setting

Here we specialize the general deterministic equivalents from Theorem 3.1 to the isotropic setting $\Sigma = I_p$ with isotropic signal $\beta \sim \mathcal{N}(0, (r^2/p) I_p)$.

Define the companion Stieltjes transform $v = v(\lambda) > 0$ as the unique solution of

$$\frac{1}{v} = \lambda + \frac{\gamma}{1+v}, \quad \lambda > 0, \quad (88)$$

and set $\kappa(\lambda) := 1/v(\lambda)$. In the isotropic case, these admit the explicit closed forms

$$\kappa(\lambda) = \frac{(\lambda + \gamma - 1) + \sqrt{(\lambda + \gamma - 1)^2 + 4\lambda}}{2}, \quad v(\lambda) = \frac{\sqrt{(\lambda + \gamma - 1)^2 + 4\lambda} - (\lambda + \gamma - 1)}{2\lambda}.$$

Since $\Sigma = I_p$, we have $G = (\Sigma + \kappa I_p)^{-1} = (1 + \kappa)^{-1} I_p$. Hence for $k \in \{2, 3, 4\}$,

$$t_k := \gamma \frac{\text{tr}(\Sigma^2 G^k)}{p} = \gamma \left(\frac{1}{1 + \kappa} \right)^k = \gamma \left(\frac{v}{1 + v} \right)^k, \quad (89)$$

$$q_k := \mathbb{E}[\beta^\top G^k \Sigma \beta] = r^2 \left(\frac{1}{1+\kappa} \right)^k = r^2 \left(\frac{v}{1+v} \right)^k.$$

(Equivalently, $q_k = \lim_{p \rightarrow \infty} \beta^\top G^k \Sigma \beta$ in probability, by isotropy of β .)

Define

$$b := (1 - t_2)^{-1}, \quad E := \kappa - b\lambda + b^2 \kappa \lambda t_3,$$

the variance-trace limits

$$u_2 := t_2 b, \quad u_3 := t_3 b^3, \quad u_4 := t_4 b^4 + 2t_3^2 b^5,$$

and

$$a_2 := bE^2 + b^4 \kappa^2 \lambda^2 t_4 + b^5 \kappa^2 \lambda^2 t_3^2, \quad a_3 := 2b^2 \kappa \lambda E, \quad a_4 := b^3 \kappa^2 \lambda^2.$$

Lemma D.2 (Isotropic limits for same- X SD). Fix $\lambda > 0$. Under Assumption A with $\Sigma = I_p$ and $\beta \sim \mathcal{N}(0, (r^2/p)I_p)$, as $n, p \rightarrow \infty$ with $p/n \rightarrow \gamma \in (0, \infty)$, the same- X optimal SD risk satisfies

$$R_{\text{sd}}^{\star, \text{same}}(\lambda) \xrightarrow{P} \mathcal{R}_{\text{sd}}^{\star, \text{same}}(\lambda) := \mathcal{R}(\lambda) - \frac{(\mathcal{R}(\lambda) - \mathcal{C}^{\text{same}}(\lambda))^2}{\mathcal{R}(\lambda) + \mathcal{R}_{\text{pd}}^{\text{same}}(\lambda) - 2\mathcal{C}^{\text{same}}(\lambda)}.$$

Here $\mathcal{R}(\lambda)$, $\mathcal{R}_{\text{pd}}^{\text{same}}(\lambda)$, and $\mathcal{C}^{\text{same}}(\lambda)$ are the deterministic limits of the (one-round) teacher risk $R(\lambda)$, same- X PD risk $R_{\text{pd}}^{\text{same}}(\lambda)$, and residual correlation $\mathcal{C}^{\text{same}}(\lambda)$, respectively (in the same- X setting of Section 2, these coincide with $R(\lambda)$, $R_{\text{pd}}(\lambda)$, and $C(\lambda)$), given by:

$$\begin{aligned} \mathcal{R}(\lambda) &= \sigma^2 + \kappa^2 b q_2 + \sigma^2 u_2, \\ \mathcal{C}^{\text{same}}(\lambda) &= \sigma^2 + 2\kappa^2 b q_2 - (\kappa b E q_2 + \kappa^2 b^2 \lambda q_3) + \sigma^2 (u_2 - \lambda u_3), \\ \mathcal{R}_{\text{pd}}^{\text{same}}(\lambda) &= \sigma^2 + 4\kappa^2 b q_2 - 2(2\kappa b E q_2 + 2\kappa^2 b^2 \lambda q_3) + (a_2 q_2 + a_3 q_3 + a_4 q_4) + \sigma^2 (u_2 - 2\lambda u_3 + \lambda^2 u_4). \end{aligned}$$

Proof. This follows by specializing Theorem 3.1 to $\Sigma = I_p$ and isotropic β . \square

For an illustration of the empirical versus theoretical risks in Lemma D.2, see the second panel of Figure 12.

D.2.3 Fresh- X (affine) SD under Isotropic Setting

Throughout, we work in the balanced proportional regime $m, n, p \rightarrow \infty$ with $p/m \rightarrow \gamma$ and $p/n \rightarrow \gamma$, and under the isotropic assumptions $\Sigma = I_p$ and $\beta \sim \mathcal{N}(0, (r^2/p)I_p)$ from Section D.2. Recall the companion Stieltjes transform $v(\lambda)$ and $\kappa(\lambda) = 1/v(\lambda)$ from Section D.2.2, as well as $b(\lambda) = \kappa'(\lambda) = (1 - t_2(\lambda))^{-1}$.

Let $\tilde{\Sigma} := \tilde{X}^\top \tilde{X}/m$ and define the (fresh) shrinkage matrix $M_\lambda := \tilde{\Sigma}(\tilde{\Sigma} + \lambda I_p)^{-1}$. Let

$$s(\lambda) := \lim_{m, p \rightarrow \infty} \frac{1}{p} \text{tr}(M_\lambda), \quad s_2(\lambda) := \lim_{m, p \rightarrow \infty} \frac{1}{p} \text{tr}(M_\lambda^2).$$

Writing $\tilde{S} := \tilde{X} \tilde{X}^\top / m$ and $\tilde{G}_\lambda := (\tilde{S} + \lambda I_m)^{-1}$, we have $\tilde{S}(\tilde{S} + \lambda I_m)^{-1} = I_m - \lambda \tilde{G}_\lambda$, and therefore

$$\frac{1}{p} \text{tr}(M_\lambda) = \frac{m}{p} \cdot \frac{1}{m} \text{tr}(\tilde{S}(\tilde{S} + \lambda I_m)^{-1}) = \frac{1}{\gamma} \left(1 - \lambda \frac{1}{m} \text{tr}(\tilde{G}_\lambda) \right).$$

Since $p/m \rightarrow \gamma$, we have $\frac{1}{m} \text{tr}(\tilde{G}_\lambda) \rightarrow v(\lambda)$, where $v(\lambda)$ solves (88). Hence

$$s(\lambda) = \frac{1 - \lambda v(\lambda)}{\gamma} = \frac{1 - \lambda/\kappa(\lambda)}{\gamma}. \quad (90)$$

Similarly, using $M_\lambda = I_p - \lambda(\tilde{\Sigma} + \lambda I_p)^{-1}$ and the identity $\frac{1}{m} \text{tr}(\tilde{G}_\lambda^2) = -v'(\lambda)$, we obtain

$$s_2(\lambda) = \frac{1 - 2\lambda v(\lambda) - \lambda^2 v'(\lambda)}{\gamma} = \frac{1 - 2\lambda v(\lambda) + \lambda^2 b(\lambda) v(\lambda)^2}{\gamma} = \frac{1 - 2\lambda/\kappa(\lambda) + \lambda^2 b(\lambda)/\kappa(\lambda)^2}{\gamma}. \quad (91)$$

Lemma D.3 (Isotropic limits for fresh- X (affine) SD). Fix $\lambda > 0$. Under Assumption A with $\Sigma = I_p$ and $\beta \sim \mathcal{N}(0, (r^2/p)I_p)$, as $m, n, p \rightarrow \infty$ with $p/m, p/n \rightarrow \gamma \in (0, \infty)$, the optimal fresh- X SD risk satisfies

$$R_{\text{sd}}^{\star, \text{fraff}}(\lambda) \xrightarrow{P} \mathcal{R}_{\text{sd}}^{\star, \text{fraff}}(\lambda) := \mathcal{R}(\lambda) - \frac{(\mathcal{R}(\lambda) - \mathcal{C}^{\text{fr}}(\lambda))^2}{\mathcal{R}(\lambda) + \mathcal{R}_{\text{pd}}^{\text{fr}}(\lambda) - 2\mathcal{C}^{\text{fr}}(\lambda)}.$$

Here $\mathcal{R}(\lambda)$ is the same teacher risk limit from Lemma D.2, and $\mathcal{R}_{\text{pd}}^{\text{fr}}(\lambda)$ and $\mathcal{C}^{\text{fr}}(\lambda)$ are the deterministic limits of the fresh- X PD risk $R_{\text{pd}}^{\text{fr}}(\lambda)$ and the corresponding residual correlation $C^{\text{fr}}(\lambda)$, given by:

$$\begin{aligned} \mathcal{C}^{\text{fr}}(\lambda) &= \sigma^2 + s(\lambda)(\mathcal{R}(\lambda) - \sigma^2) + r^2(1 - s(\lambda))^2, \\ \mathcal{R}_{\text{pd}}^{\text{fr}}(\lambda) &= \sigma^2 + s_2(\lambda)(\mathcal{R}(\lambda) - \sigma^2) + r^2(1 - 2s(\lambda))^2 + (2s(\lambda) - 1)s_2(\lambda). \end{aligned}$$

Proof. Let $\beta_{\text{pd}, \lambda}^{\text{fr}} := M_\lambda \beta_\lambda$ be the coefficient of the fresh- X PD refit (trained on $(\tilde{X}, \tilde{y}_\lambda)$ with $\tilde{y}_\lambda = \tilde{X} \beta_\lambda$), and write $\Delta_\lambda := \beta_\lambda - \beta$, and $\Delta_{\text{pd}, \lambda}^{\text{fr}} := \beta_{\text{pd}, \lambda}^{\text{fr}} - \beta$. In the isotropic in-distribution setting, the teacher and PD risks and their residual correlation satisfy the identities:

$$R(\lambda) = \sigma^2 + \|\Delta_\lambda\|_2^2, \quad R_{\text{pd}}^{\text{fr}}(\lambda) = \sigma^2 + \|\Delta_{\text{pd}, \lambda}^{\text{fr}}\|_2^2, \quad C^{\text{fr}}(\lambda) = \sigma^2 + \langle \Delta_\lambda, \Delta_{\text{pd}, \lambda}^{\text{fr}} \rangle.$$

Since $\Delta_{\text{pd}, \lambda}^{\text{fr}} = M_\lambda \beta_\lambda - \beta = M_\lambda \Delta_\lambda + (M_\lambda - I_p)\beta$, we have

$$\langle \Delta_\lambda, \Delta_{\text{pd}, \lambda}^{\text{fr}} \rangle = \Delta_\lambda^\top M_\lambda \Delta_\lambda + \Delta_\lambda^\top (M_\lambda - I_p) \beta, \quad (92)$$

and

$$\|\Delta_{\text{pd}, \lambda}^{\text{fr}}\|_2^2 = \Delta_\lambda^\top M_\lambda^2 \Delta_\lambda + \beta^\top (M_\lambda - I_p)^2 \beta + 2 \Delta_\lambda^\top M_\lambda (M_\lambda - I_p) \beta. \quad (93)$$

Thus, we need asymptotics for linear and quadratic forms involving Δ_λ , which we obtain below.

Linear and quadratic forms of Δ_λ . Write the ridge teacher error as

$$\Delta_\lambda = -\lambda(\hat{\Sigma} + \lambda I_p)^{-1} \beta + (\hat{\Sigma} + \lambda I_p)^{-1} \frac{X^\top \varepsilon}{n} =: \Delta_\lambda^{\text{bias}} + \Delta_\lambda^{\text{var}},$$

where $\hat{\Sigma} = X^\top X/n$ and $\varepsilon := y - X\beta$. As in the proof of Theorem 3.1, the cross term $\langle \Delta_\lambda^{\text{bias}}, \Delta_\lambda^{\text{var}} \rangle$ is $o_p(1)$, and $\|\Delta_\lambda\|_2^2 = \|\Delta_\lambda^{\text{bias}}\|_2^2 + \|\Delta_\lambda^{\text{var}}\|_2^2 + o_p(1)$. Moreover, by Lemma D.2, $\|\Delta_\lambda\|_2^2 \xrightarrow{P} \mathcal{R}(\lambda) - \sigma^2$.

Because \tilde{X} is independent of (X, β, ε) , the matrix M_λ is independent of (Δ_λ, β) and is orthogonally invariant, with $\|M_\lambda\|_{\text{op}} \leq 1$. By standard concentration for quadratic forms of an independent orthogonally invariant matrix, similar to Lemma B.3 in the proof of Theorem 3.1, we have

$$\Delta_\lambda^\top M_\lambda \Delta_\lambda = \left(\frac{1}{p} \text{tr}(M_\lambda) \right) \|\Delta_\lambda\|_2^2 + o_p(1), \quad (94)$$

$$\Delta_\lambda^\top M_\lambda^2 \Delta_\lambda = \left(\frac{1}{p} \text{tr}(M_\lambda^2) \right) \|\Delta_\lambda\|_2^2 + o_p(1), \quad (95)$$

$$\Delta_\lambda^\top M_\lambda \beta = \left(\frac{1}{p} \text{tr}(M_\lambda) \right) \langle \Delta_\lambda, \beta \rangle + o_p(1), \quad (96)$$

$$\Delta_\lambda^\top M_\lambda^2 \beta = \left(\frac{1}{p} \text{tr}(M_\lambda^2) \right) \langle \Delta_\lambda, \beta \rangle + o_p(1), \quad (97)$$

and similarly,

$$\beta^\top (M_\lambda - I_p)^2 \beta = \left(\frac{1}{p} \text{tr}((M_\lambda - I_p)^2) \right) \|\beta\|_2^2 + o_p(1). \quad (98)$$

Now, note that

$$\langle \Delta_\lambda, \beta \rangle = -\lambda \beta^\top (\widehat{\Sigma} + \lambda I_p)^{-1} \beta + \beta^\top (\widehat{\Sigma} + \lambda I_p)^{-1} \frac{X^\top \varepsilon}{n}.$$

The second term is mean-zero conditional on (X, β) and is $o_p(1)$ by Lemma B.3. For the first term, isotropy of β and standard quadratic-form concentration give

$$\beta^\top (\widehat{\Sigma} + \lambda I_p)^{-1} \beta = \frac{\|\beta\|_2^2}{p} \text{tr}(\widehat{\Sigma} + \lambda I_p)^{-1} + o_p(1).$$

Using $\|\beta\|_2^2 \xrightarrow{P} r^2$ and $\widehat{\Sigma}(\widehat{\Sigma} + \lambda I_p)^{-1} = I_p - \lambda(\widehat{\Sigma} + \lambda I_p)^{-1}$, we obtain

$$\lambda \cdot \frac{1}{p} \text{tr}(\widehat{\Sigma} + \lambda I_p)^{-1} = 1 - \frac{1}{p} \text{tr}(\widehat{\Sigma}(\widehat{\Sigma} + \lambda I_p)^{-1}).$$

In the isotropic proportional regime with $p/n \rightarrow \gamma$, the term $\frac{1}{p} \text{tr}(\widehat{\Sigma}(\widehat{\Sigma} + \lambda I_p)^{-1})$ converges to the same deterministic limit as $\frac{1}{p} \text{tr}(M_\lambda)$ (since $\widehat{\Sigma}$ and $\widetilde{\Sigma}$ are independent Wishart matrices with the same aspect ratio), namely $s(\lambda)$ from (90). Therefore,

$$\langle \Delta_\lambda, \beta \rangle \xrightarrow{P} -r^2(1 - s(\lambda)). \quad (99)$$

We are now ready to obtain the asymptotics for the fresh- X PD risk and residual correlation.

Residual correlation asymptotics. Combining (92) with (94) and (96) gives

$$\langle \Delta_\lambda, \Delta_{\text{pd},\lambda}^{\text{fr}} \rangle = \left(\frac{1}{p} \text{tr}(M_\lambda) \right) \|\Delta_\lambda\|_2^2 + \left(\frac{1}{p} \text{tr}(M_\lambda) - 1 \right) \langle \Delta_\lambda, \beta \rangle + o_p(1).$$

Using $\frac{1}{p} \text{tr}(M_\lambda) \rightarrow s(\lambda)$, $\|\Delta_\lambda\|_2^2 \xrightarrow{P} \mathcal{R}(\lambda) - \sigma^2$, and (99), we obtain

$$\langle \Delta_\lambda, \Delta_{\text{pd},\lambda}^{\text{fr}} \rangle \xrightarrow{P} s(\lambda)(\mathcal{R}(\lambda) - \sigma^2) + r^2(1 - s(\lambda))^2.$$

Therefore,

$$\mathcal{C}^{\text{fr}}(\lambda) = \sigma^2 + \langle \Delta_\lambda, \Delta_{\text{pd},\lambda}^{\text{fr}} \rangle \xrightarrow{P} \sigma^2 + s(\lambda)(\mathcal{R}(\lambda) - \sigma^2) + r^2(1 - s(\lambda))^2 = \mathcal{C}^{\text{fr}}(\lambda).$$

PD risk asymptotics. Similarly, combining (93) with (95), (98), and (97) yields:

$$\|\Delta_{\text{pd},\lambda}^{\text{fr}}\|_2^2 = \left(\frac{1}{p} \text{tr}(M_\lambda^2) \right) \|\Delta_\lambda\|_2^2 + \left(\frac{1}{p} \text{tr}((M_\lambda - I_p)^2) \right) \|\beta\|_2^2 + 2 \left(\frac{1}{p} \text{tr}(M_\lambda^2) - \frac{1}{p} \text{tr}(M_\lambda) \right) \langle \Delta_\lambda, \beta \rangle + o_p(1).$$

Using $\frac{1}{p} \text{tr}(M_\lambda^2) \rightarrow s_2(\lambda)$, $\frac{1}{p} \text{tr}((M_\lambda - I_p)^2) \rightarrow 1 - 2s(\lambda) + s_2(\lambda)$, $\|\beta\|_2^2 \xrightarrow{P} r^2$, $\|\Delta_\lambda\|_2^2 \xrightarrow{P} \mathcal{R}(\lambda) - \sigma^2$, and (99), we get

$$\|\Delta_{\text{pd},\lambda}^{\text{fr}}\|_2^2 \xrightarrow{P} s_2(\lambda)(\mathcal{R}(\lambda) - \sigma^2) + r^2(1 - 2s(\lambda)^2 + (2s(\lambda) - 1)s_2(\lambda)).$$

Hence

$$R_{\text{pd}}^{\text{fr}}(\lambda) = \sigma^2 + \|\Delta_{\text{pd},\lambda}^{\text{fr}}\|_2^2 \xrightarrow{P} \sigma^2 + s_2(\lambda)(\mathcal{R}(\lambda) - \sigma^2) + r^2(1 - 2s(\lambda)^2 + (2s(\lambda) - 1)s_2(\lambda)) = \mathcal{R}_{\text{pd}}^{\text{fr}}(\lambda).$$

Finally, because $f_{\text{sd},\lambda,\xi}^{\text{fraff}} = (1 - \xi)f_\lambda + \xi f_{\text{pd},\lambda}^{\text{fr}}$ is a two-predictor affine path, combining the oracle formula in Proposition 2.1 with the component convergences proved above finishes the proof. \square

For an illustration of the empirical versus theoretical risks in Lemma D.3, see the second panel of Figure 12.

D.2.4 Proof of Theorem 5.2

Proof of Theorem 5.2. As before, let $v = v(\lambda) > 0$ denote the (unique) solution of (88), and define the signal-to-noise ratio $\text{SNR} := r^2/\sigma^2 > 0$. Set $D_0 := (1 + v)^2 - \gamma v^2 = (1 + v)^2(1 - t_2)$, where $t_2 = \gamma(v/(1 + v))^2$ as in (89). Note that $D_0 > 0$. This is because differentiating (88) shows that

$$v'(\lambda) = -\frac{v(\lambda)^2}{1 - \gamma(v(\lambda)/(1 + v(\lambda)))^2} = -\frac{v(\lambda)^2}{1 - t_2(\lambda)} < 0,$$

so $1 - t_2(\lambda) > 0$ and hence $D_0 > 0$ for all $\lambda > 0$.

Using the explicit isotropic expressions from Lemmas D.2 and D.3 and simplifying, one obtains the factorization

$$\mathcal{R}_{\text{sd}}^{\text{*,fraff}}(\lambda) - \mathcal{R}_{\text{sd}}^{\text{*,same}}(\lambda) = \sigma^2 \frac{\gamma v \Delta(v)^2 \Xi(v)}{(1 + v)^3 D_0 \Pi_1(v) \Pi_2(v)}, \quad (100)$$

where

$$\Delta(v) := \gamma \text{SNR} v + \gamma v^2 + \gamma v - \text{SNR} v - \text{SNR}, \quad (101)$$

$$\Pi_1(v) := -\gamma \text{SNR} v + \gamma \text{SNR} + \gamma v + \gamma + \text{SNR} v + \text{SNR}, \quad (102)$$

$$\begin{aligned} \Pi_2(v) &:= \gamma^2 \text{SNR} v^2 + \gamma^2 v^2 - 2\gamma \text{SNR} v^2 - 2\gamma \text{SNR} v + \gamma \text{SNR} \\ &\quad + \gamma v^2 + 2\gamma v + \gamma + \text{SNR} v^2 + 2\text{SNR} v + \text{SNR}, \end{aligned} \quad (103)$$

$$\Xi(v) := \gamma^2 v^3 + \text{SNR} \left(\frac{(D_0 - 1)^2}{v} + D_0 v + v + 2 \right). \quad (104)$$

All factors on the right-hand side of (100) are nonnegative: $\sigma^2 > 0$, $\gamma > 0$, $v > 0$, and $\Delta(v)^2 \geq 0$. Moreover, $D_0 > 0$ as shown above, and $\Xi(v) > 0$ since it is a sum of strictly positive terms.

It remains to note that $\Pi_1(v)$ and $\Pi_2(v)$ are strictly positive. Indeed, in both the same- X and fresh- X settings, the two-predictor oracle formula Proposition 2.1 involves the discrepancy denominator

$$\mathcal{D}^\bullet(\lambda) := \mathcal{R}(\lambda) + \mathcal{R}_{\text{pd}}^\bullet(\lambda) - 2\mathcal{C}^\bullet(\lambda) \quad \bullet \in \{\text{same}, \text{fr}\},$$

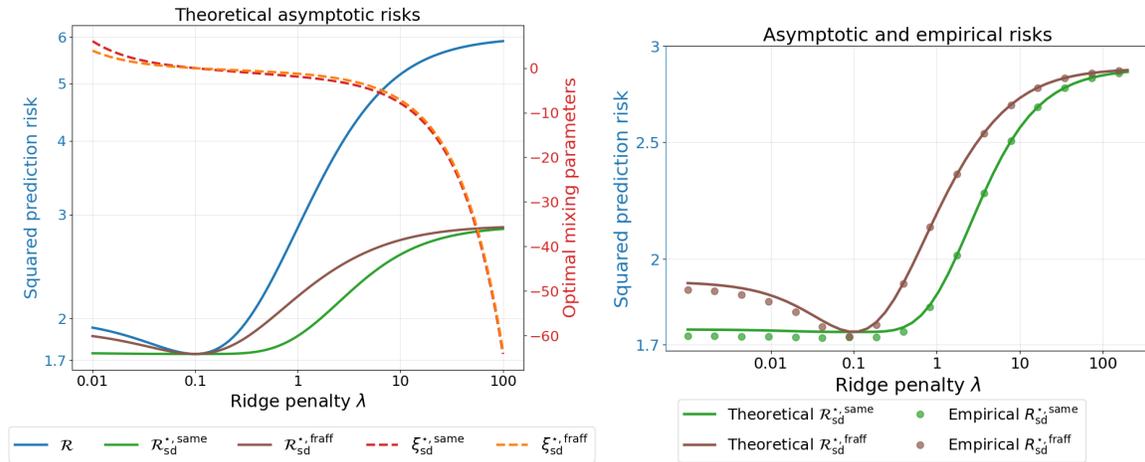


Figure 12: **Same- X versus fresh- X self-distillation risks.** Theoretical asymptotic (Theorem 5.2) and empirical risks (averaged over 20 simulations) and optimal mixing weights for an isotropic setting with $n = 400$, $p = 200$, $r^2 = 5$, and $\sigma^2 = 1$.

which is strictly positive for every $\lambda > 0$ (the teacher and the corresponding PD refit do not coincide in this isotropic model). A direct identification shows that $\Pi_1(v)$ and $\Pi_2(v)$ are proportional to $\mathcal{D}^{\text{fr}}(\lambda)$ and $\mathcal{D}^{\text{same}}(\lambda)$, respectively, with positive proportionality constants; hence $\Pi_1(v), \Pi_2(v) > 0$.

Therefore the right-hand side of (100) is nonnegative for all $\lambda > 0$, showing the desired domination. \square

For a visual illustration of the risk asymptotics of the same- X versus (affine) fresh- X optimal SD risks, see the first panel of Figure 12. Consistent with Theorem 5.2, the same- X optimal SD risk uniformly dominates the fresh- X optimal SD risk across λ .

D.3 Proofs and Details in Section 5.3

D.3.1 Proof of Equation (28)

Proof of Equation (28). Recall that $f_\lambda(x) = s_\lambda(x)^\top y$ and $\hat{y}_\lambda = S_\lambda y$, and that the SD refit applies the same smoother to the mixed labels $y^{(\xi)} := (1 - \xi)y + \xi\hat{y}_\lambda$. Then, for every x ,

$$f_{\text{sd},\lambda,\xi}(x) = s_\lambda(x)^\top y^{(\xi)} = (1 - \xi)s_\lambda(x)^\top y + \xi s_\lambda(x)^\top \hat{y}_\lambda = (1 - \xi)f_\lambda(x) + \xi f_{\text{pd},\lambda}(x),$$

where $f_{\text{pd},\lambda}(x) := s_\lambda(x)^\top \hat{y}_\lambda$ by definition. \square

D.3.2 Proof of Theorem 5.3

Proof of Theorem 5.3. Fix $\lambda > 0$. By (28), the SD family is the affine path $f_{\text{sd},\lambda,\xi} = (1 - \xi)f_\lambda + \xi f_{\text{pd},\lambda}$. Let $e_\lambda := y_0 - f_\lambda(x_0)$ and $e_{\text{pd}} := y_0 - f_{\text{pd},\lambda}(x_0)$. Expanding the square gives, for every $\xi \in \mathbb{R}$,

$$\begin{aligned} R_{\text{sd}}(\lambda, \xi) &= \mathbb{E}\left[\left((1 - \xi)e_\lambda + \xi e_{\text{pd}}\right)^2 \middle| \mathcal{D}\right] \\ &= (1 - \xi)^2 R(\lambda) + \xi^2 R_{\text{pd}}(\lambda) + 2\xi(1 - \xi)C(\lambda) \\ &= R(\lambda) - 2\xi(R(\lambda) - C(\lambda)) + \xi^2 D(\lambda), \end{aligned}$$

where $D(\lambda) = R(\lambda) + R_{\text{pd}}(\lambda) - 2C(\lambda)$. Under $D(\lambda) > 0$, this is a strictly convex quadratic in ξ , hence the unique minimizer is

$$\xi^*(\lambda) = \frac{R(\lambda) - C(\lambda)}{D(\lambda)}, \quad R_{\text{sd}}^*(\lambda) = R(\lambda) - \frac{(R(\lambda) - C(\lambda))^2}{D(\lambda)}.$$

Next, assume (29) holds. Since $\lambda \mapsto R(\lambda)$ is differentiable, we may differentiate inside the conditional expectation to get

$$R'(\lambda) = \partial_\lambda \mathbb{E}[e_\lambda^2 \mid \mathcal{D}] = \mathbb{E}[2e_\lambda \partial_\lambda e_\lambda \mid \mathcal{D}] = -2 \mathbb{E}[e_\lambda \partial_\lambda f_\lambda(x_0) \mid \mathcal{D}].$$

On the other hand,

$$\begin{aligned} R(\lambda) - C(\lambda) &= \mathbb{E}[e_\lambda(e_\lambda - e_{\text{pd}}) \mid \mathcal{D}] = \mathbb{E}[e_\lambda (f_{\text{pd},\lambda}(x_0) - f_\lambda(x_0)) \mid \mathcal{D}] \\ &= -\mathbb{E}[e_\lambda (f_\lambda(x_0) - f_{\text{pd},\lambda}(x_0)) \mid \mathcal{D}] = -\mathbb{E}[e_\lambda (-\lambda \partial_\lambda f_\lambda(x_0)) \mid \mathcal{D}] \\ &= \lambda \mathbb{E}[e_\lambda \partial_\lambda f_\lambda(x_0) \mid \mathcal{D}] = -\frac{\lambda}{2} R'(\lambda). \end{aligned}$$

Substituting $R(\lambda) - C(\lambda) = -(\lambda/2)R'(\lambda)$ into the closed forms above yields

$$\xi^*(\lambda) = -\frac{\lambda}{2} \frac{R'(\lambda)}{D(\lambda)}, \quad R_{\text{sd}}^*(\lambda) = R(\lambda) - \frac{\lambda^2}{4} \frac{(R'(\lambda))^2}{D(\lambda)}.$$

Finally, if $R'(\lambda) \neq 0$ then $R_{\text{sd}}^*(\lambda) < R(\lambda)$ and $\text{sign}(\xi^*(\lambda)) = -\text{sign}(R'(\lambda))$. \square

D.3.3 Verifying Derivative Property (29) for Common Ridge Variants

Lemma D.4 (Generalized ridge satisfies (29)). Fix $\Omega \succ 0$ and $\lambda > 0$, and let

$$f_\lambda^\Omega(x) := x^\top (X^\top X + n\lambda\Omega)^{-1} X^\top y.$$

Let $f_{\text{pd},\lambda}^\Omega$ denote the PD refit obtained by training generalized ridge at the same (X, Ω, λ) on pseudo-labels $\hat{y}_\lambda^\Omega = f_\lambda^\Omega(X)$. Then, for all x ,

$$f_\lambda^\Omega(x) - f_{\text{pd},\lambda}^\Omega(x) = -\lambda \partial_\lambda f_\lambda^\Omega(x).$$

Proof. Let $A_\lambda := X^\top X + n\lambda\Omega$. Then $f_\lambda^\Omega(x) = x^\top A_\lambda^{-1} X^\top y$ and $f_{\text{pd},\lambda}^\Omega(x) = x^\top A_\lambda^{-1} X^\top X A_\lambda^{-1} X^\top y$. Hence

$$f_\lambda^\Omega(x) - f_{\text{pd},\lambda}^\Omega(x) = x^\top A_\lambda^{-1} (A_\lambda - X^\top X) A_\lambda^{-1} X^\top y = n\lambda x^\top A_\lambda^{-1} \Omega A_\lambda^{-1} X^\top y.$$

Moreover, $\partial_\lambda A_\lambda^{-1} = -A_\lambda^{-1} (\partial_\lambda A_\lambda) A_\lambda^{-1} = -n A_\lambda^{-1} \Omega A_\lambda^{-1}$, so

$$\partial_\lambda f_\lambda^\Omega(x) = x^\top (\partial_\lambda A_\lambda^{-1}) X^\top y = -n x^\top A_\lambda^{-1} \Omega A_\lambda^{-1} X^\top y,$$

and multiplying by $-\lambda$ gives the claim. \square

Lemma D.5 (Kernel ridge satisfies (29)). Fix $\lambda > 0$ and a PSD kernel with kernel matrix $K \in \mathbb{R}^{n \times n}$ and $k_x := (k(x, x_1), \dots, k(x, x_n))^\top$. Let

$$f_\lambda^{\text{kern}}(x) := k_x^\top (K + n\lambda I_n)^{-1} y,$$

and let $f_{\text{pd},\lambda}^{\text{kern}}$ denote the PD refit trained at the same λ on pseudo-labels $\hat{y}_\lambda^{\text{kern}} = f_\lambda^{\text{kern}}(X)$. Then, for all x ,

$$f_\lambda^{\text{kern}}(x) - f_{\text{pd},\lambda}^{\text{kern}}(x) = -\lambda \partial_\lambda f_\lambda^{\text{kern}}(x).$$

Proof. Let $B_\lambda := K + n\lambda I_n$. Then $f_\lambda^{\text{kern}}(x) = k_x^\top B_\lambda^{-1}y$ and $f_{\text{pd},\lambda}^{\text{kern}}(x) = k_x^\top B_\lambda^{-1}KB_\lambda^{-1}y$. Thus

$$f_\lambda^{\text{kern}}(x) - f_{\text{pd},\lambda}^{\text{kern}}(x) = k_x^\top B_\lambda^{-1}(B_\lambda - K)B_\lambda^{-1}y = n\lambda k_x^\top B_\lambda^{-2}y.$$

Since $\partial_\lambda B_\lambda^{-1} = -B_\lambda^{-1}(\partial_\lambda B_\lambda)B_\lambda^{-1} = -nB_\lambda^{-2}$, we have

$$\partial_\lambda f_\lambda^{\text{kern}}(x) = k_x^\top (\partial_\lambda B_\lambda^{-1})y = -n k_x^\top B_\lambda^{-2}y,$$

and multiplying by $-\lambda$ yields the claim. □

E Additional Numerical Illustrations

The source code for reproducing the results of this paper can be found at the following location: https://github.com/hhd357/optimal_self_distillation_ridge

E.1 Additional Experiments on CIFAR Datasets

E.1.1 CIFAR100 with Pretrained ResNet-34 Features

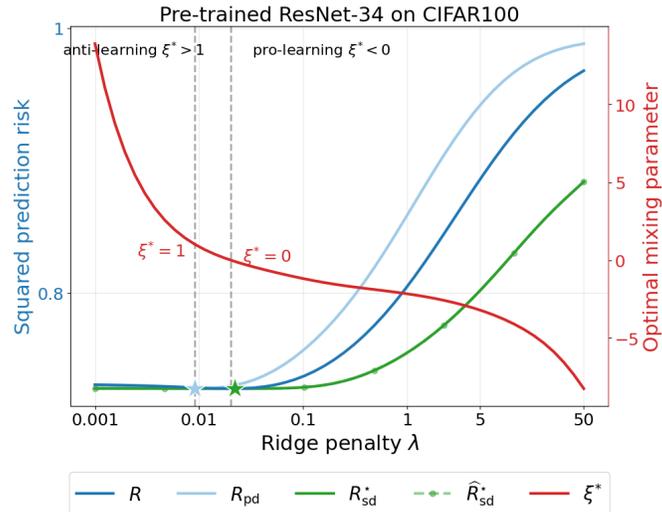


Figure 13: Squared prediction risks and optimal mixing parameter for ridge, pure-distilled and self-distilled ridge on pretrained ResNet-34 features on CIFAR100 dataset.

E.1.2 CIFAR10 and CIFAR100 Classification Accuracies

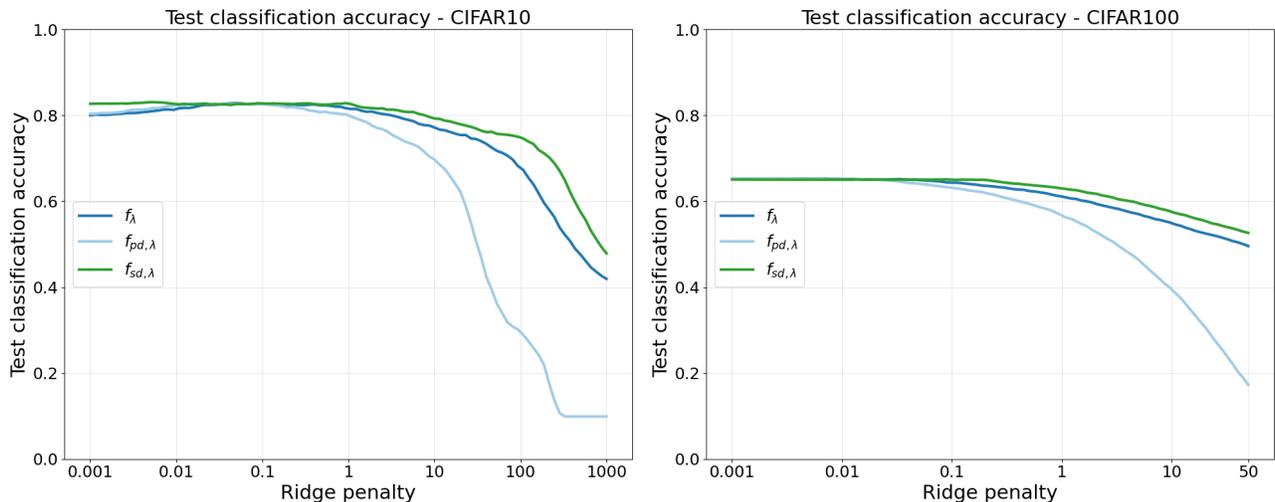


Figure 14: Test classification accuracy of the ridge and self-distilled ridges on CIFAR10 and CIFAR100 experiments. Although the optimal mixing is chosen to minimize the squared prediction risk, the strict improvement property still somewhat holds for test accuracy as well.

E.2 Illustration for Related Works Comparison in Section 1.2

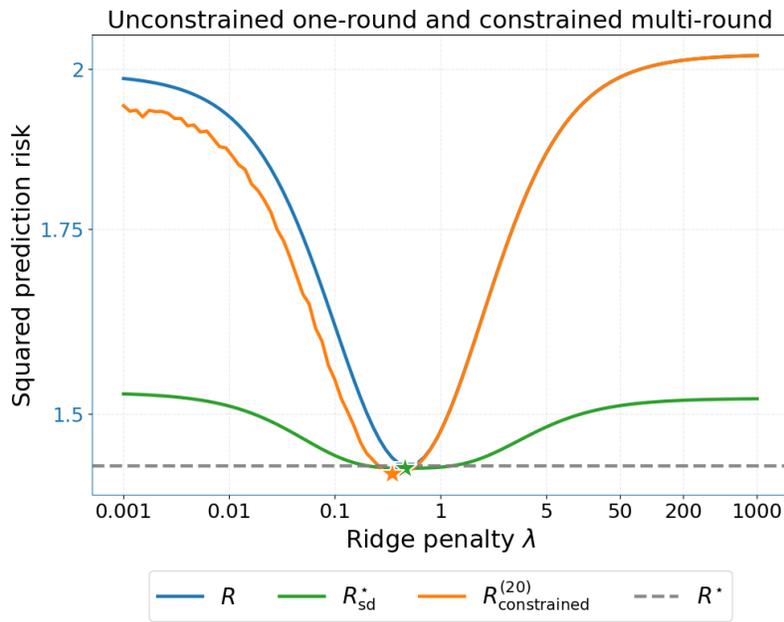


Figure 15: **One-round unconstrained versus multi-round constrained mixing weight.** Squared prediction risk empirical curves averaged over 20 simulations, $n = 400, p = 200$ with isotropic design and isotropic signal. We choose the optimal ξ over a grid of 200 values in $[0, 1]$ and pick the one with the lowest risk at the 20-th round.

E.3 Additional Illustrations Real-World Regression Tasks

E.3.1 Comparison with Constrained Self-Distillation

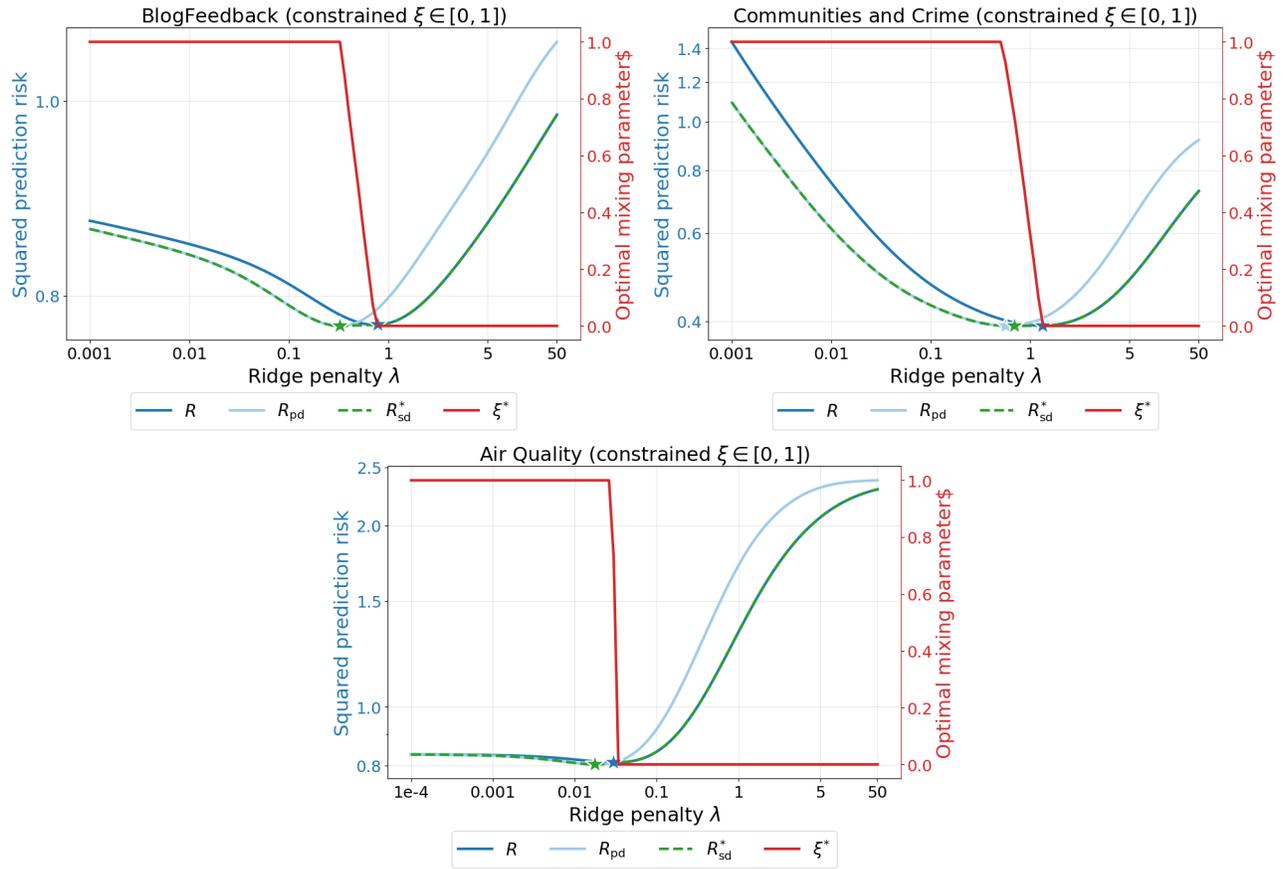


Figure 16: Test-set squared prediction risk and optimal mixing parameter, being constrained in $[0, 1]$. We choose the optimal ξ over a grid of 100 values in $[0, 1]$ and pick the one with the lowest SD risk. The linestyle for the SD risk is changed only for this figure since it mostly overlaps with either the original ridge or the pure-distilled ridge curves. When $\xi = 1$, the SD risk exactly matches the PD risk and when $\xi = 0$, it exactly matches the teacher ridge risk.

E.3.2 Varying Train-to-Test Split Ratios and Gains

BlogFeedback dataset:

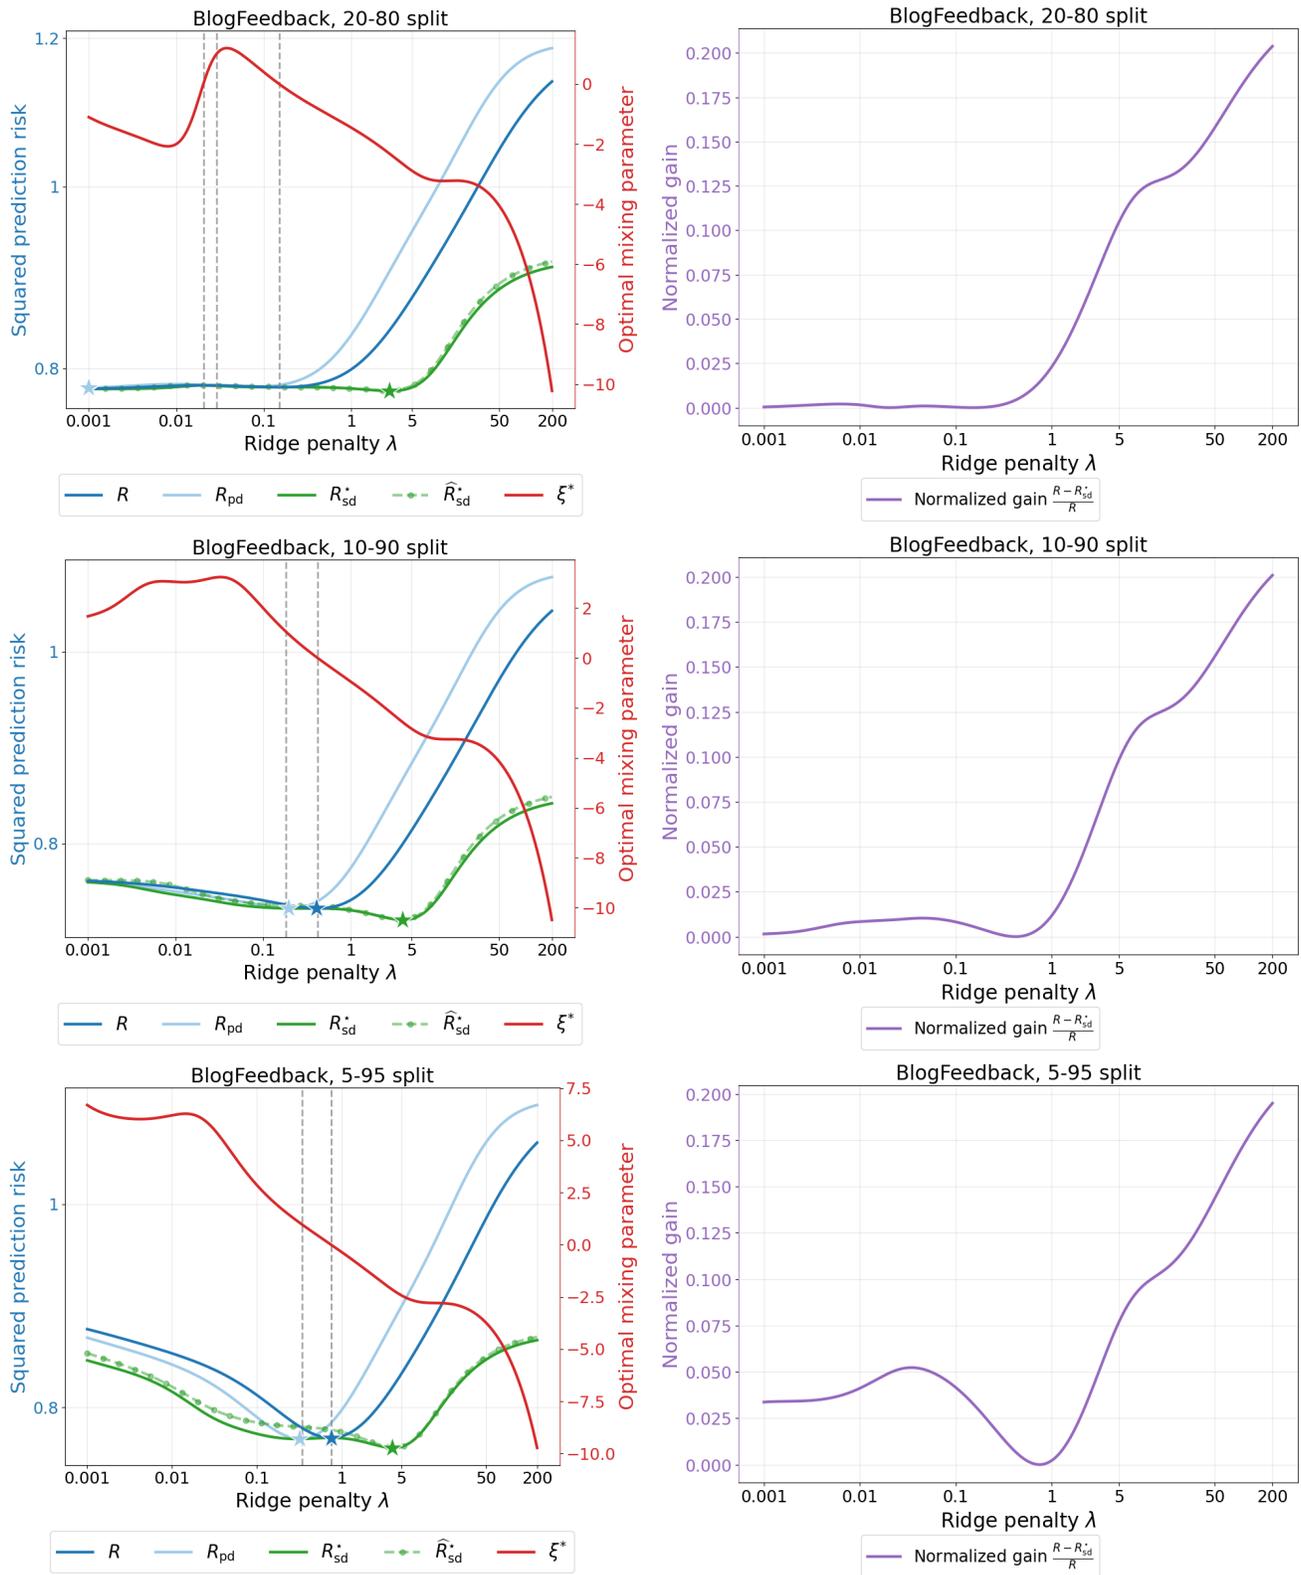


Figure 17: Squared prediction risks and gain curve on BlogFeedback dataset with different split ratios.

Communities and Crime dataset:

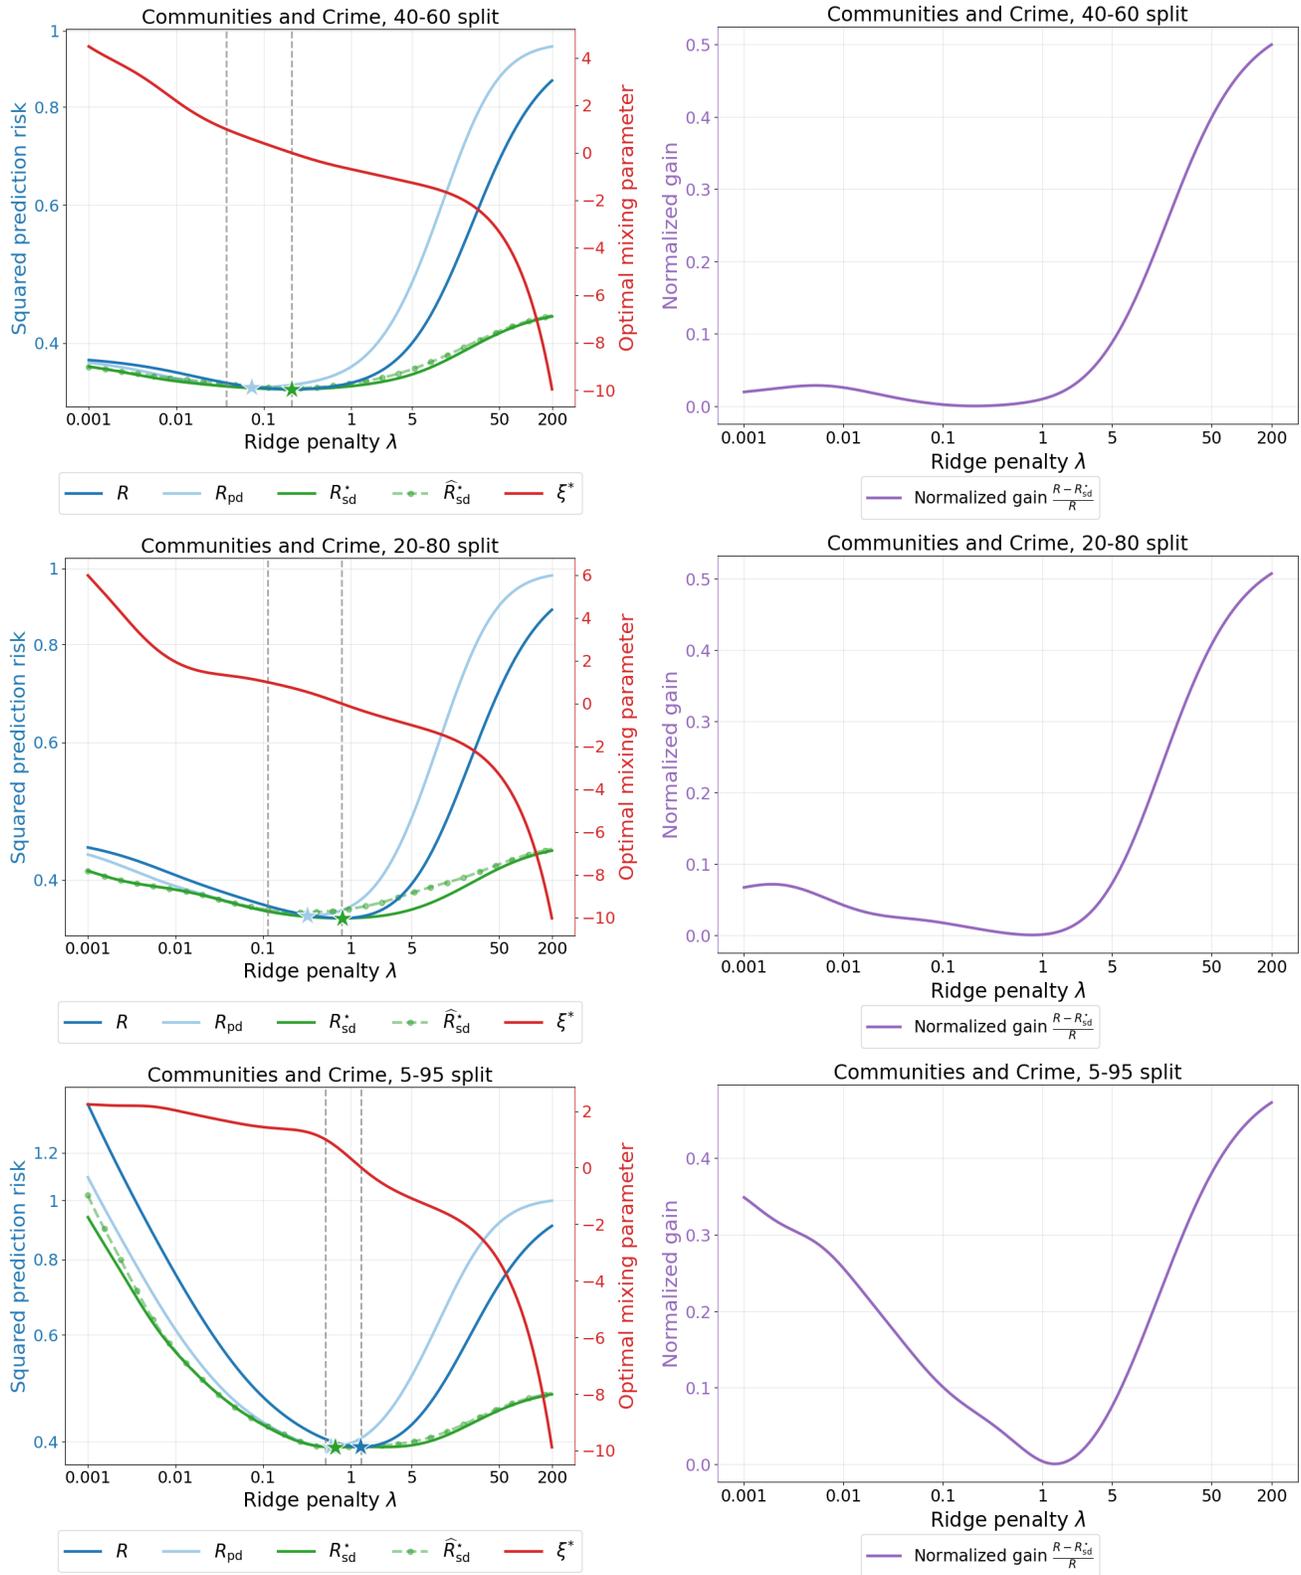


Figure 18: Squared prediction risks and gain curve on Communities and Crime dataset with different split ratios.

Air Quality dataset:

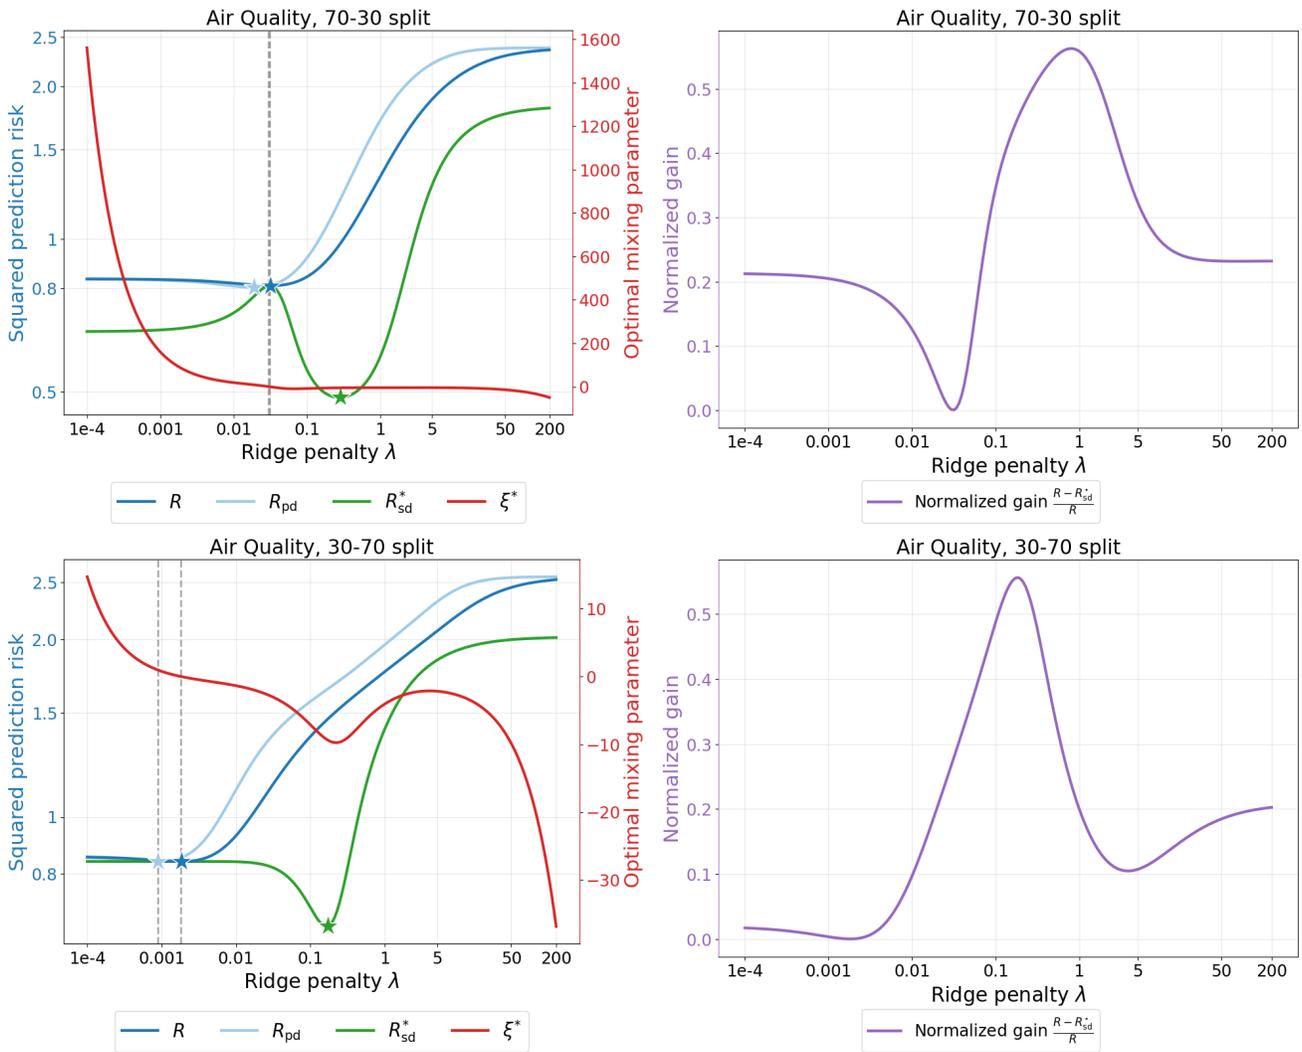


Figure 19: Squared prediction risks and gain curve on Air Quality dataset with different split ratios.

E.4 Additional Illustrations on Proportional Asymptotic Risks

E.4.1 Varying Signal-to-Noise Ratios

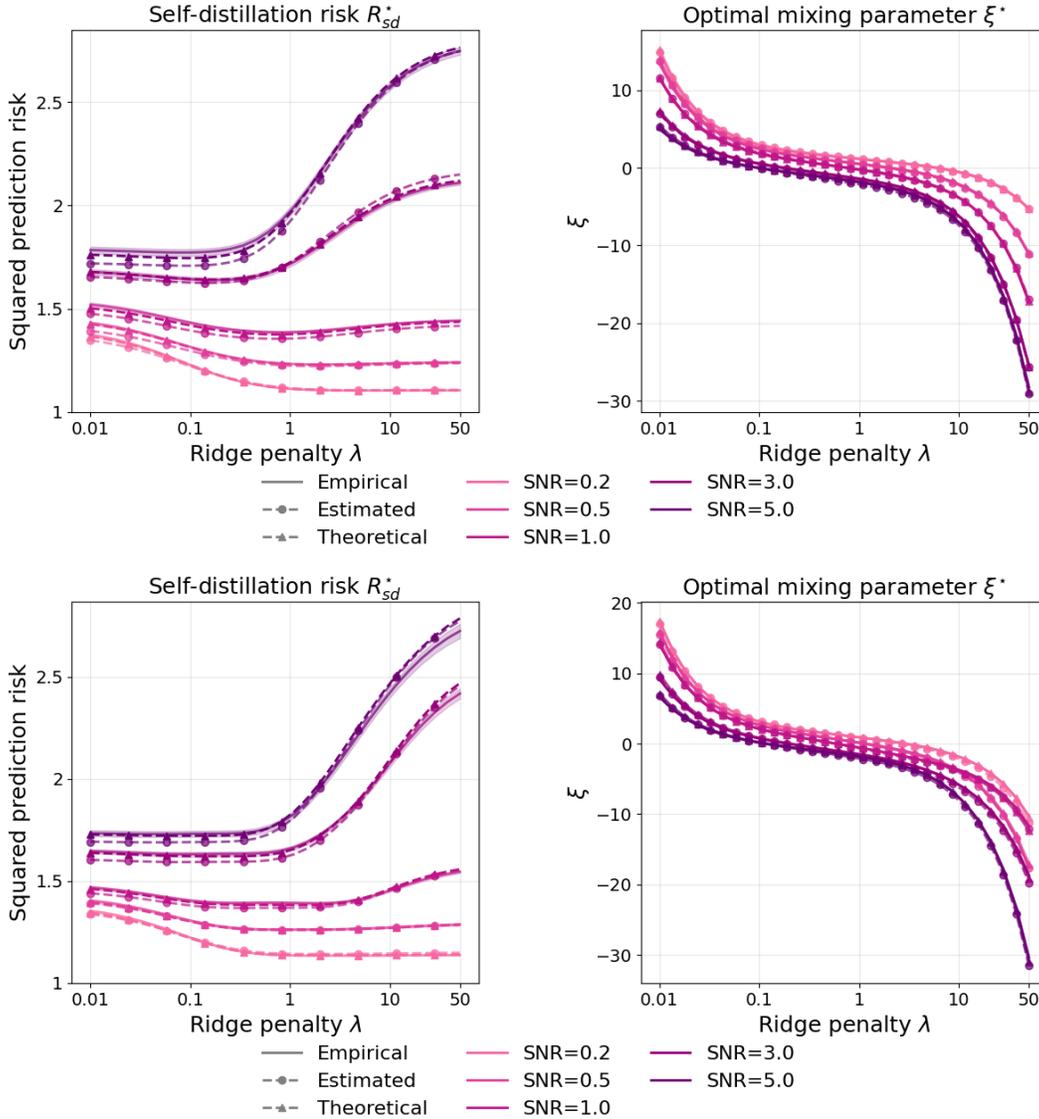


Figure 20: Asymptotic SD risk and optimal mixing parameter over various SNR ratios. Empirical curves are averaged over 30 numerical simulations, and the shaded band represents one standard deviation around the mean. The estimated are obtained using proposed tuning method (Section 4, averaged over 30 runs) and theoretical curves are from Theorem 3.1. $n = 400$, $p = 200$, $\sigma^2 = 1$ and $r^2 = \sigma^2 \text{SNR}$. *Top row*: Data covariance Σ is AR1, deterministic ground-truth signal is aligned with the bottom 10% eigenvalues of Σ , with alignment factor 0.9. *Bottom row*: Data covariance Σ is a spiked covariance matrix, deterministic ground-truth signal is aligned with the top 10% eigenvalues of Σ , with alignment factor 0.9.

E.4.2 Varying Aspect Ratios

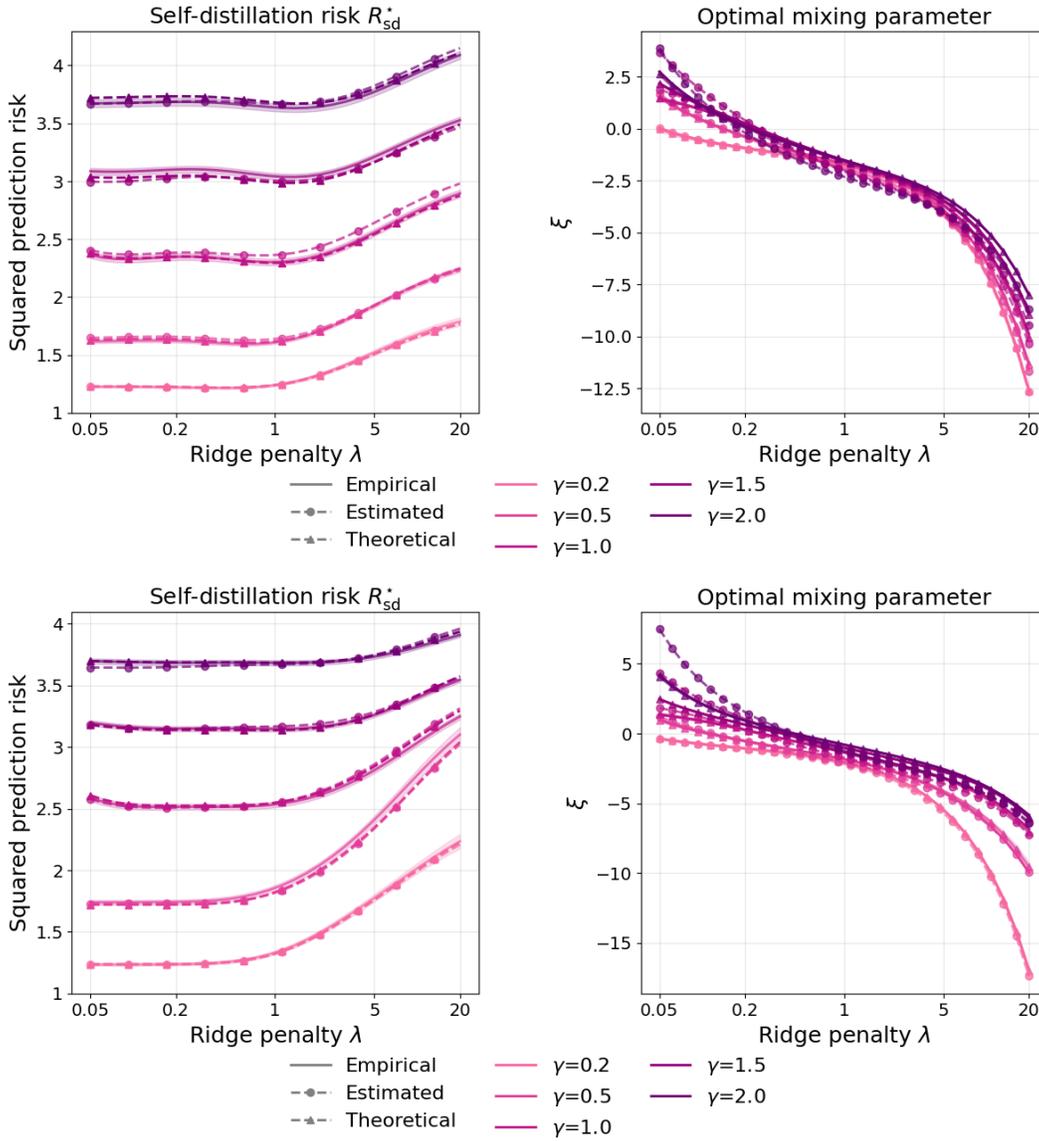


Figure 21: Asymptotic SD risk and optimal mixing parameter over different aspect ratios $\gamma = p/n$. Empirical curves are averaged over 30 numerical simulations, and the shaded band represents one standard deviation around the mean. The estimated are obtained using proposed tuning method (Section 4, averaged over 30 runs) and theoretical curves are from Theorem 3.1. $n = 300$, $p = n\gamma$, $\sigma^2 = 1$ and $r^2 = 5$. *Top row:* Data covariance Σ is AR1, deterministic ground-truth signal is aligned with the top 10% eigenvalues of Σ , with alignment factor 0.9. *Bottom row:* Data covariance Σ is spiked covariance matrix with $\Sigma = I + 5vv^\top$ with v is a random isotropic Gaussian vector, deterministic ground-truth signal is aligned with the top 10% eigenvalues of Σ , with alignment factor 0.9.

E.4.3 Additional Risk and Gain Curves

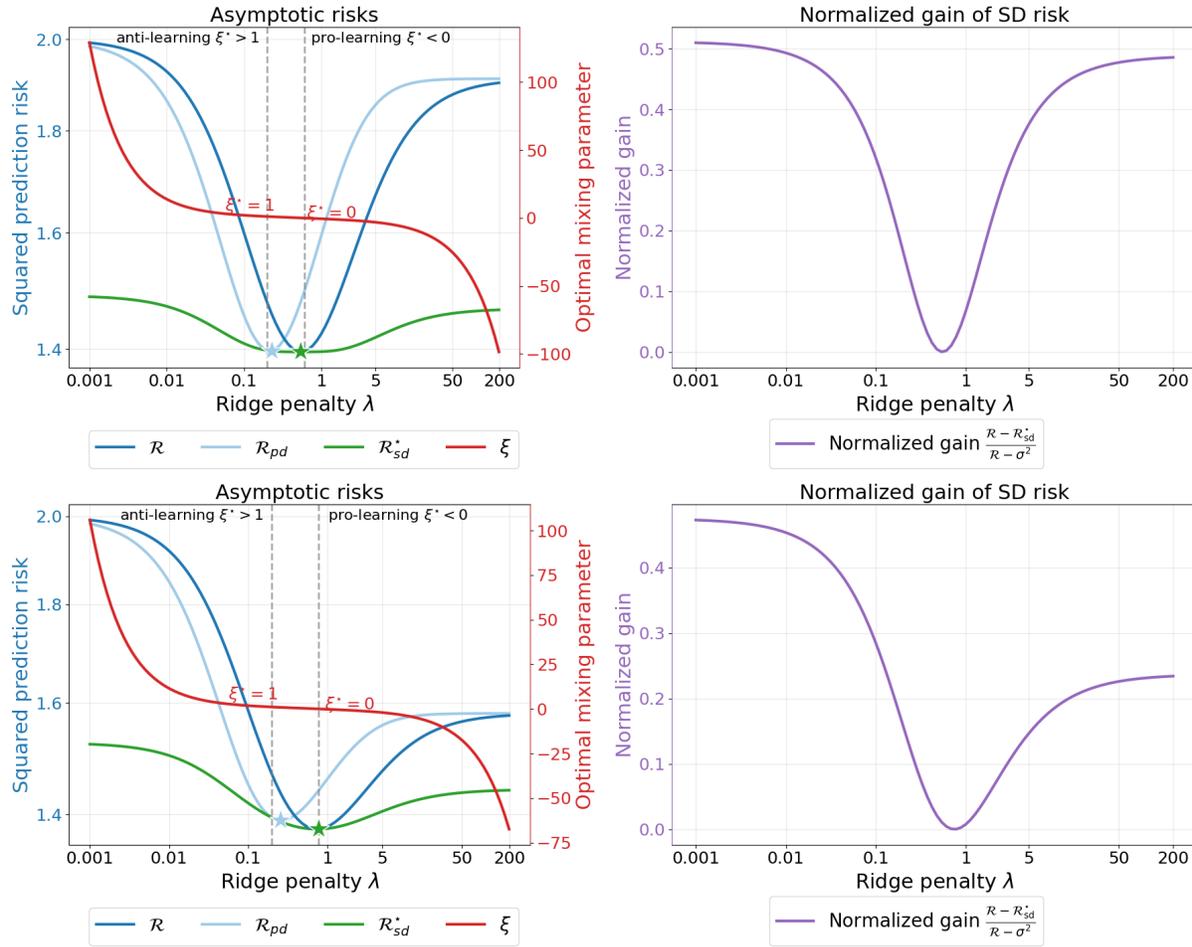


Figure 22: Asymptotic risk curves, asymptotic optimal mixing parameter and gain curves over the penalty λ . *Top row:* Data covariance Σ is isotropic, isotropic signal, $n = 400, p = 200, r^2 = \sigma^2 = 1$. *Bottom row:* Data covariance Σ is AR-1(0.25), deterministic ground-truth signal is aligned with the bottom 10% eigenvalues of Σ , with alignment factor 0.9, $n = 400, p = 200, r^2 = \sigma^2 = 1$.

E.5 Additional Illustrations on Extreme Regularized Risks

E.5.1 Isotropic Covariance, Isotropic Signal

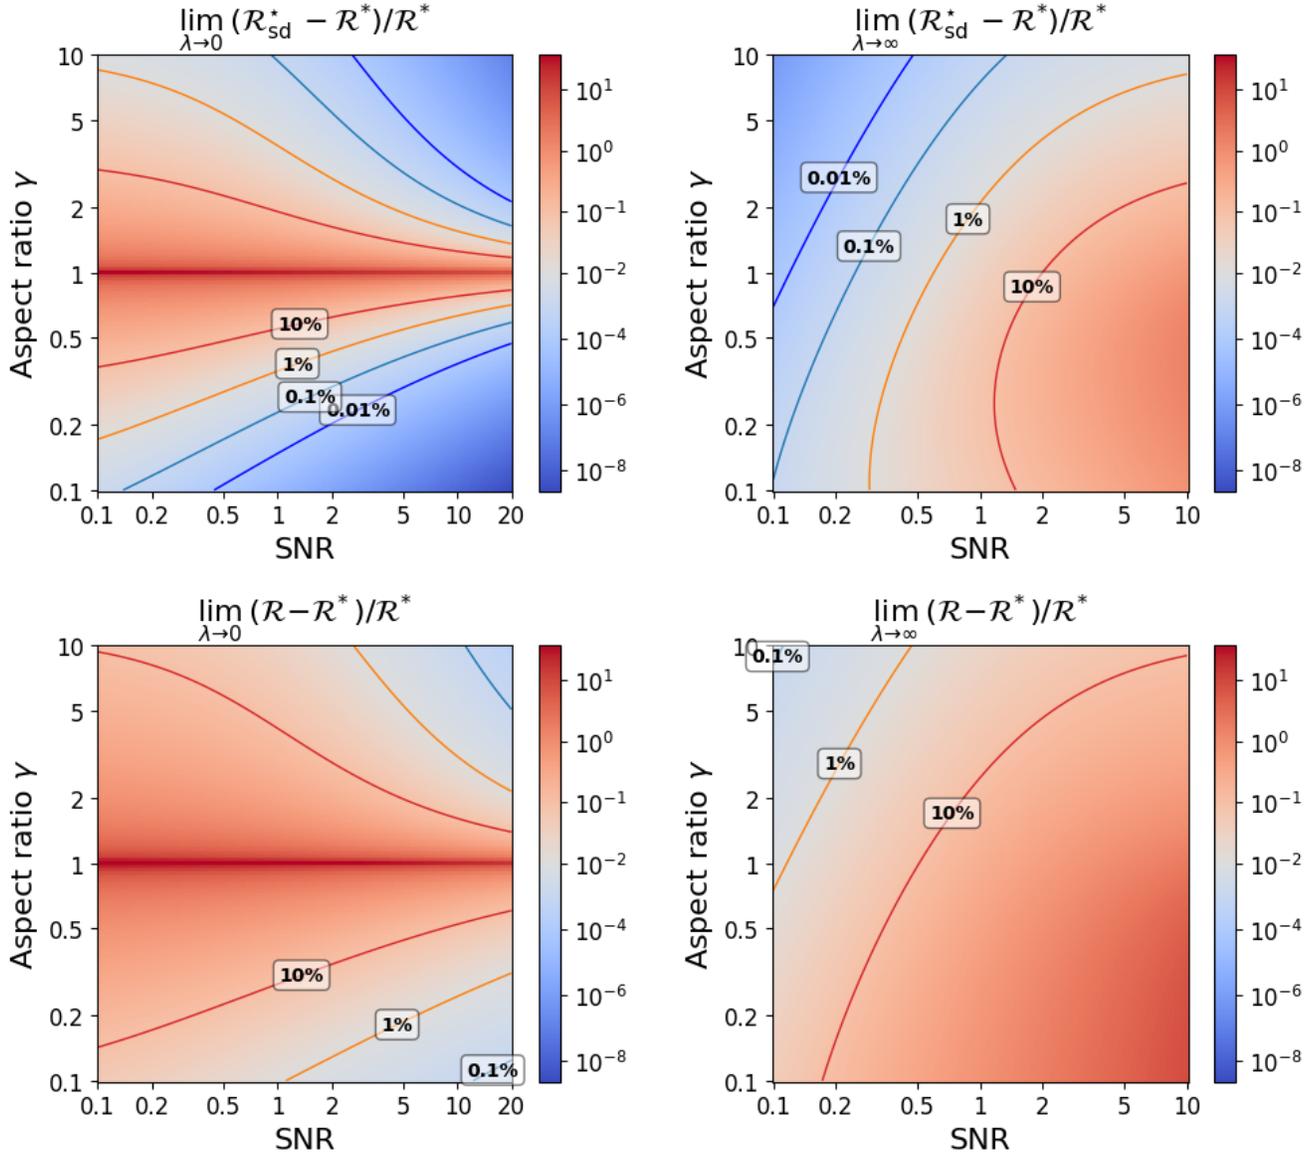


Figure 23: Percentage difference between $\mathcal{R}(\text{SNR}, \gamma)$ and $\mathcal{R}_{sd}^*(\text{SNR}, \gamma)$ to the optimal ridge $\mathcal{R}^*(\text{SNR}, \gamma)$ (also the best predictor in this setting), and with isotropic design $\Sigma = I_p$, random signal follow an isotropic Gaussian distribution. The values plotted are calculated from Proposition 3.3.

E.5.2 AR1 Covariance, Isotropic Signal

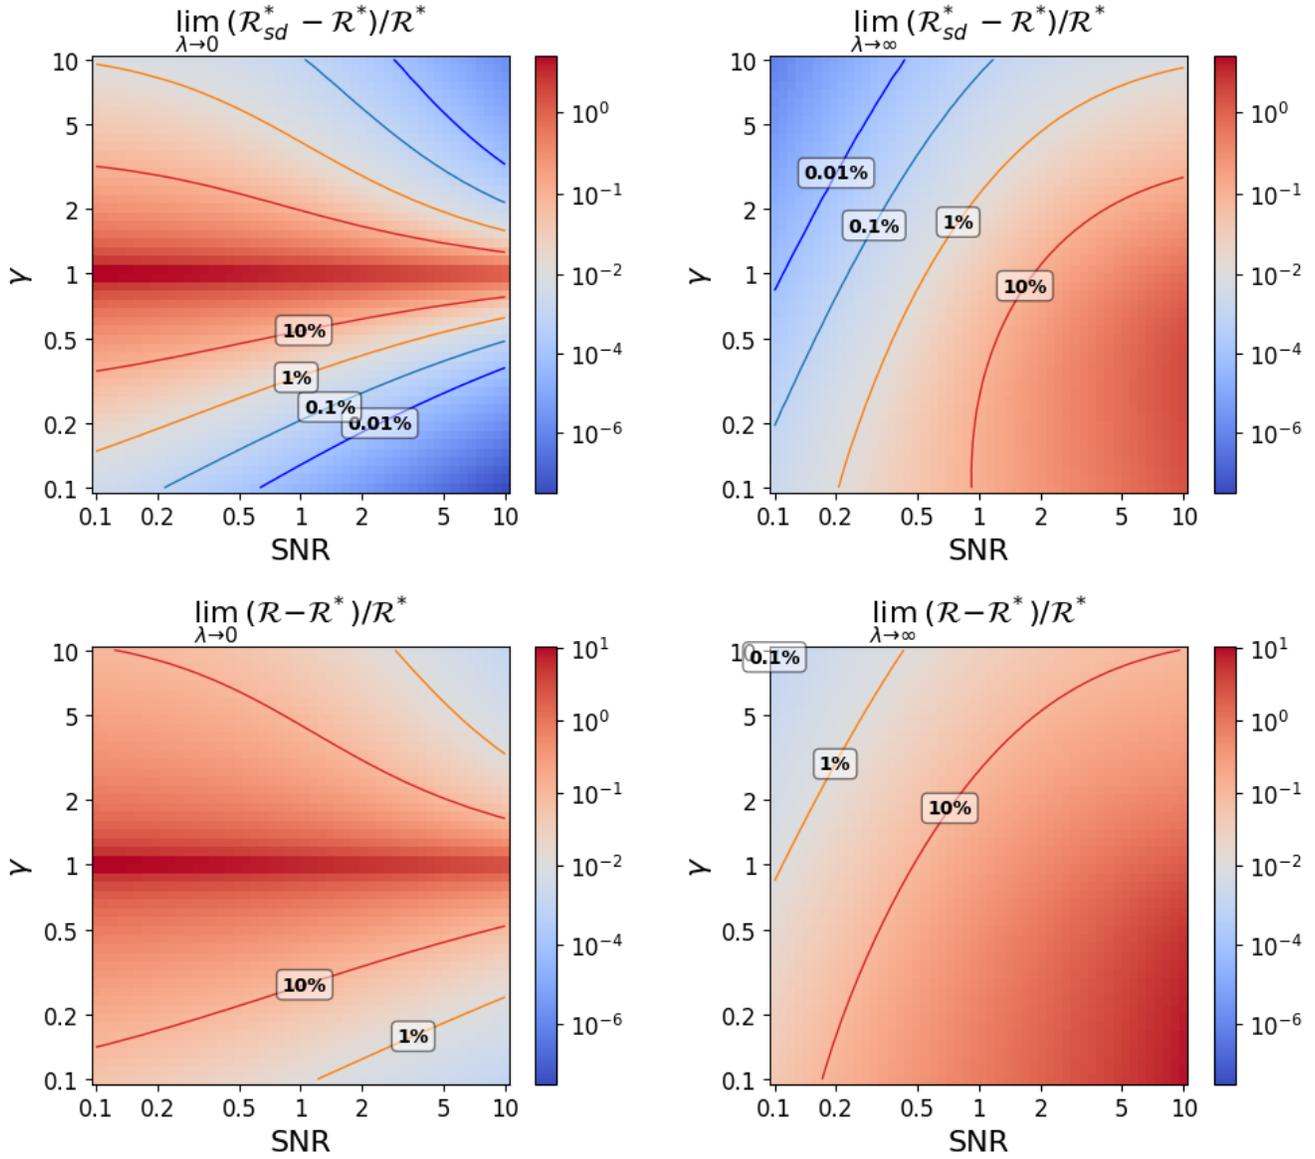


Figure 24: Percentage difference between $\mathcal{R}(\text{SNR}, \gamma)$ and $\mathcal{R}_{sd}^*(\text{SNR}, \gamma)$ to the optimal ridge $\mathcal{R}^*(\text{SNR}, \gamma)$ (also the best predictor in this setting), with AR1 covariance design and random signal follows an isotropic Gaussian distribution. The ratios plotted are calculated using risk formulations at Theorem 3.1 at $\lambda = 10^{-3}$ and $\lambda = 10^6$.

E.5.3 Spiked Covariance, Isotropic Signal

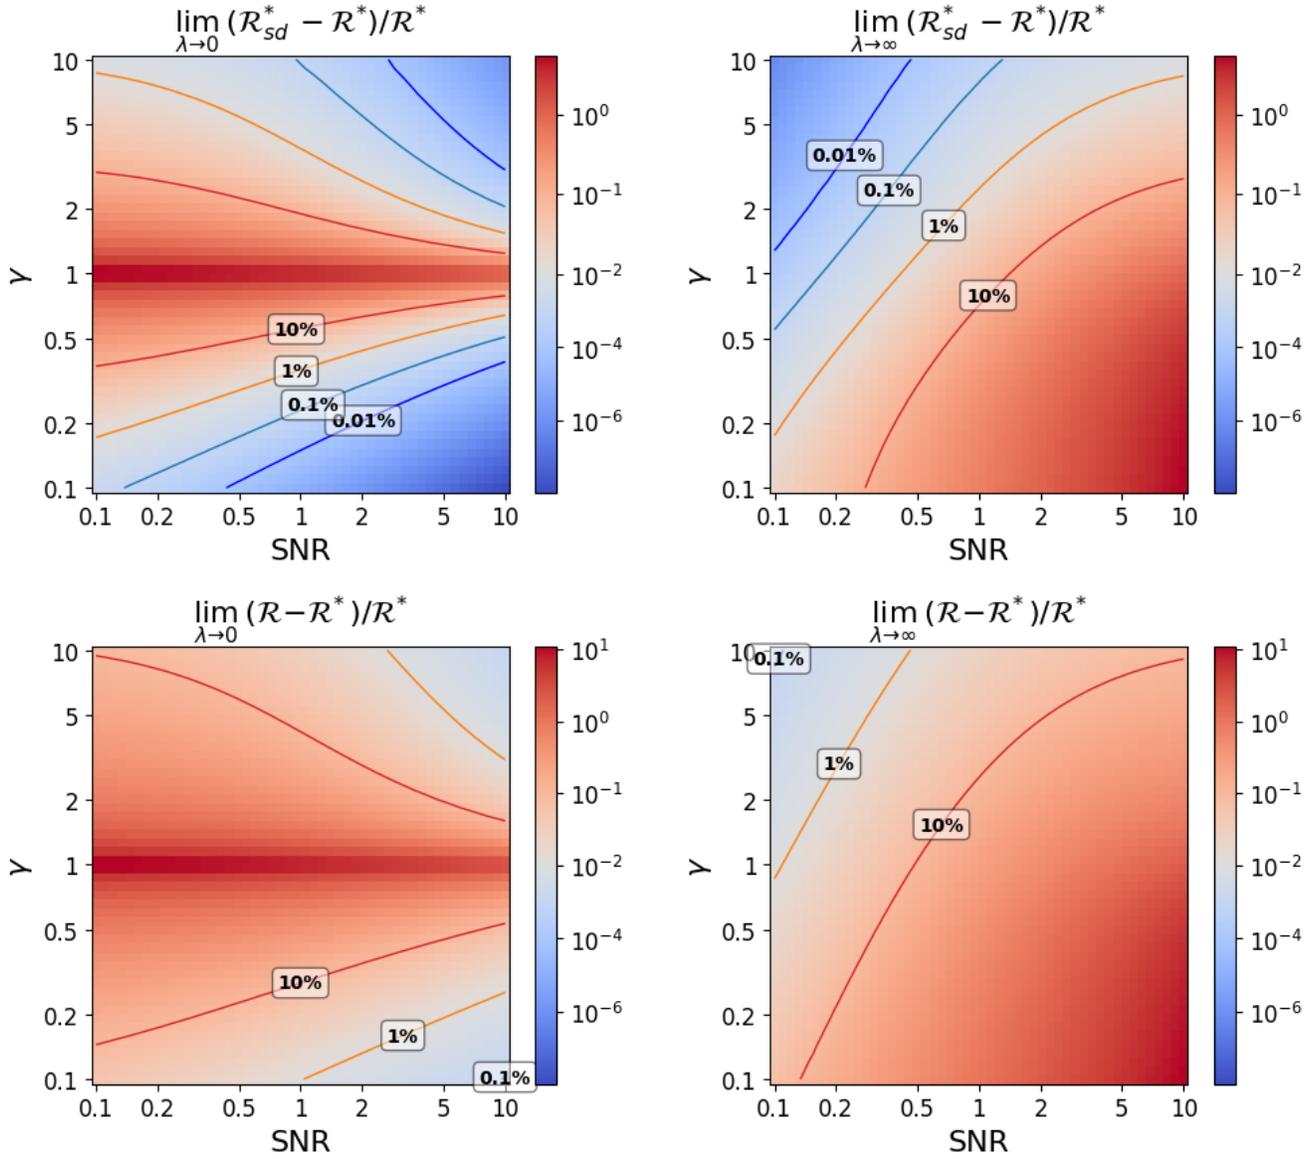


Figure 25: Percentage difference between $\mathcal{R}(\text{SNR}, \gamma)$ and $\mathcal{R}_{sd}^*(\text{SNR}, \gamma)$ to the optimal ridge $\mathcal{R}^*(\text{SNR}, \gamma)$ (also the best predictor in this setting), with spiked covariance design and random signal follow an isotropic Gaussian distribution. The ratios plotted are calculated using risk formulations at Theorem 3.1 at $\lambda = 10^{-3}$ and $\lambda = 10^6$.

F Experiment Details

F.1 Real-World Regression Tasks: Datasets Description, Splitting and Pre-processing

For the experiments on real-world regression task, including UCI Blog Feedback (Buza, 2014), UCI Communities and Crime (Redmond, 2002), UCI Air Quality (Vito, 2008) datasets, we split the data into the training and test set (detailed below). Then, we process the data by removing records with missing values and we centered and standardized all the variables using the training set’s mean and standard deviation. All the variables in the test set are also centered and standardized using the information from training set only.

All the risk curves shown at Figures 2 and 3 are the squared risk measured on the test set. The optimal mixing parameter is computed using the formula from Equation (8) where R , R_{pd} and C are calculated using the test data.

UCI Blogfeedback. The UCI Blogfeedback dataset (Buza, 2014) originates from blog posts, where the raw HTML-documents of the blog posts were crawled and processed. We use only the training set of this dataset, which contains 52,397 samples and perform a random 5-95 split (5% for training and 95% for test set). The dataset consists of 280 features that capture many aspects of blog content and metadata, such as post length, number of links, number of comments in the first 24 hours after the publication of the blog post. The target variable is the number of comments in the next 24 hours (relative to base time). Thus, the training set in our experiment contains $p = 280$ covariates and $n = 2,619$ samples. The additional results of different split ratios are shown at Figure 17.

UCI Communities and Crime. The UCI Communities and Crime dataset (Redmond, 2002) are authentic data that combines socio-economic data from the 1990 US Census, law enforcement data from the 1990 US LEMAS survey, and crime data from the 1995 FBI UCR. The dataset consists of 1994 samples with 127 covariates. The goal is to predict the value of “ViolentCrimesPerPop”, which is the rate of violent crimes per 100,000 population. For this dataset, we removed 5 nonpredictive identifiers and 23 covariates from the LEMAS survey that contain 1,675 missing values out of 1,994 instances. The remaining 100 covariates have no missing values and used in our experiments. We perform a random 20-80 split (20% for training and 80% for test set). Thus, the training set in our experiment contains $p = 99$ covariates and $n = 398$ samples. The additional results of different split ratios are shown at Figure 18.

UCI Air Quality. The UCI Air Quality dataset (Vito, 2008) contains records of hourly averaged responses from an array of 5 metal oxide chemical sensors. Data were recorded at an Italy city from March 2004 to February 2005. Similar as Pareek et al. (2024), we use $p = 8$ covariates for this task, which include 5 metal oxide chemical readings $PT08.S1(CO)$, $PT08.S2(NMHC)$, $PT08.S3(NO_x)$, $PT08.S4(NO_2)$, $PT08.S5(O_3)$ and 3 other covariates including Temperature T , Relative Humidity RH , and Absolute Humidity AH . The goal is to predict the Nitrogen Dioxide $NO_2(GT)$. After removing missing records for these variables, we are left with $n = 7,393$ samples. We perform a *sequential* 70-30 split since the data is heavily time-dependent. Thus, the training set in our experiment contains $p = 8$ covariates and $n = 5,175$ samples. The additional results of different split ratios are shown at Figure 19.

For the multi-round distillation experiment shown in Figure 8, they are performed on a sequential 70-30 split of Air Quality dataset and a random 20-80 split of Communities and Crime dataset.

For the kernel ridge regression experiment shown in Figure 11, we use a Gaussian kernel with

bandwidth estimated as the median of the ℓ_2 distances between covariates in the training set.

F.2 CIFAR10 and CIFAR100 Experiments

For the experiment on CIFAR10 and CIFAR100 datasets, we extract the pretrained ResNet-18 and ResNet-34 features (He et al., 2016), that trained on ImageNet dataset (available in Pytorch). For CIFAR10, we randomly sample 2,000 samples for training and 2,000 samples for the test set. For CIFAR100, we randomly sample 20,000 samples for training and 10,000 samples for the test set. Let K is the number of classes. We then perform ridge regression on the last-layer features of the pretrained models. The predictor is now defined as the vector-valued function $f : \mathbb{R}^{512} \rightarrow \mathbb{R}^K$. To aggregate the risk, for an one-hot label vector $y \in \mathbb{R}^K$ where , we simply sum the mean squared error over the K input dimensions,

$$R(f) = \mathbb{E}_{(x,y)} \left[\sum_{k=1}^K (y_k - f(x)_k)^2 \right], \quad (105)$$

where $f_k : \mathbb{R}^{512} \rightarrow \mathbb{R}$ is a ridge predictor that predict the probability that the input belong to class k . The optimal mixing parameter is calculated from Equation (8).

F.3 Synthetic Asymptotic Experiments

We give more details here about the data covariance Σ and signal β (defined in Assumption A) used in the proportional asymptotic experiments in Sections 3, E.4 and E.5.

- Isotropic covariance: $\Sigma = I_p$
- AR1 covariance: ρ -autoregressive covariance with $\rho = 0.25$, $\Sigma_{ij} = \rho^{|i-j|}$ for all i, j .
- Spiked covariance: $\Sigma = I_p + 5vv^\top$ where $v \in \mathbb{R}^p$ is a random isotropic Gaussian vector.
- Isotropic signal: $\beta \sim \mathcal{N}(0, (r^2/p)I_p)$.
- Top-aligned signal with alignment ratio $m\%$ and alignment factor of a : let $k = \frac{m}{100} \cdot p$, then $\beta \sim \mathcal{N}(0, \Sigma_\beta)$ where $\Sigma_\beta = pV \text{diag}(\frac{a}{k}, \frac{a}{k}, \dots, \frac{a}{k}, \frac{1-a}{p-k}, \dots, \frac{1-a}{p-k})V^\top$ where $V = [v_1, \dots, v_p]$ contain the eigenvectors of Σ with v_k corresponds to k -th largest eigenvalue.
- Bottom-aligned signal with alignment ratio $m\%$ and alignment factor of a : let $k = \frac{m}{100} \cdot p$, then $\beta \sim \mathcal{N}(0, \Sigma_\beta)$ where $\Sigma_\beta = pV \text{diag}(\frac{1-a}{p-k}, \dots, \frac{1-a}{p-k}, \frac{a}{k}, \dots, \frac{a}{k})V^\top$ where $V = [v_1, \dots, v_p]$ contain the eigenvectors of Σ with v_k corresponds to k -th largest eigenvalue.