PREVENTING MODEL COLLAPSE UNDER OVER-PARAMETRIZATION: OPTIMAL MIXING RATIOS FOR INTERPOLATION LEARNING AND RIDGE REGRESSION

Anonymous authorsPaper under double-blind review

ABSTRACT

Model collapse occurs when generative models degrade after repeatedly training on their own synthetic outputs. We study this effect in overparameterized linear regression in a setting where each iteration mixes fresh real labels with synthetic labels drawn from the model fitted in the previous iteration. We derive precise generalization error formulae for minimum- ℓ_2 -norm interpolation and ridge regression under this iterative scheme. Our analysis reveals intriguing properties of the optimal mixing weight that minimizes long-term prediction error and provably prevents model collapse. For instance, in the case of min- ℓ_2 -norm interpolation, we establish that the optimal real-data proportion converges to the reciprocal of the golden ratio for fairly general classes of covariate distributions. Previously, this property was known only for ordinary least squares, and additionally in low dimensions. For ridge regression, we further analyze two popular model classes – the random-effects model and the spiked covariance model -demonstrating how spectral geometry governs optimal weighting. In both cases, as well as for isotropic features, we uncover that the optimal mixing ratio should be at least one-half, reflecting the necessity of favoring real-data over synthetic. We validate our theoretical results with extensive simulations.

1 Introduction

Modern AI models are increasingly trained on their own synthetic outputs. However, this practice can lead to *model collapse*, where prediction performance degrades progressively with iterative re-training on AI generated *synthetic data* Shumailov et al. (2024; 2023). The phenomenon has been empirically observed across a wide array of settings Alemohammad et al. (2024); Bertrand et al. (2024); Bohacek & Farid (2023); Briesch et al. (2023); Hataya et al. (2023); Martínez et al. (2023a;b). Motivated by these observations, recent work has rigorously studied model collapse and developed methods to mitigate it in some cases Shumailov et al. (2024); Dohmatob et al. (2024a;b; 2025); Feng et al. (2025); Dey & Donoho (2024); Gerstgrasser et al. (2024); Kazdan et al. (2024). However previous work remains limited to low-dimensional settings or Gaussian features, creating a significant gap in our understanding. This paper breaks this barrier by addressing two fundamental questions: Can mixing fresh real data with synthetic outputs mitigate model collapse in overparametrized problems? What is the optimal mixing ratio that minimizes prediction error in the long run? We provide rigorous answers for overparametrized linear regression, demonstrating how model collapse can be prevented under overparametrization.

Prior work has rigorously studied model collapse across several problem settings, though with important limitations. For low-dimensional Gaussian distribution estimation, Shumailov et al. (2024) shows how repeated use of synthetic data causes the estimated covariance matrix to collapse to zero almost surely, while the sample mean diverges. Similar results hold for linear regression, as established by Dohmatob et al. (2024a), for Gaussian features. Recent work Dohmatob et al. (2024b; 2025); Feng et al. (2025) attributes collapse to a change in scaling laws, with applications to text generation and Gaussian mixture problems. To mitigate model collapse, Gerstgrasser et al. (2024); Dey & Donoho (2024); Kazdan et al. (2024); He et al. (2025) develop a mixing framework, where models are trained on a mixture of real and synthetic data at each iteration. This approach prevents

collapse by ensuring that the test error remains bounded even as the number of iterations increase. Crucially, these papers focus exclusively on low-dimensional problems.

We study two broad classes of estimators—minimum- ℓ_2 -norm interpolators (Section 3.1) and ridge regression (Section 3.2 and Section 3.3). Modern machine learning (ML) algorithms frequently exhibit implicit regularization Zhang & Yu (2005); Soudry et al. (2018); Gunasekar et al. (2018a)—with appropriate initialization and step sizes, algorithms converge to predictors that achieve remarkable generalization in overparameterized regimes. Implicit regularization has become a cornerstone for understanding why overparameterized models generalize well Bartlett et al. (2020). Within this framework, min-norm interpolators have emerged as a fundamental class of predictors that commonly arise as implicitly regularized limits of gradient-based algorithms Bartlett et al. (2020); Deng et al. (2022); Gunasekar et al. (2018a;b); Liang & Sur (2022); Montanari et al. (2019); Muthukumar et al. (2020); Soudry et al. (2018); Zhang & Yu (2005); Wang et al. (2022); Zhou et al. (2022). At the same time, ridge regression represents a fundamental learning paradigm that has historically provided valuable insights into complex algorithms, often illuminating phenomena observed in deep networks Hastie et al. (2022); Patil et al. (2024). We study these popular classes of predictors for understanding model collapse under overparametrization.

Specifically, we adopt the *fresh data augmentation* framework from He et al. (2025): given covariates \mathbf{X} , at iteration t, we generate a new batch of *real* responses \mathbf{y}_t alongside *synthetic* responses $\widetilde{\mathbf{y}}_t$ produced using the estimator from iteration (t-1). The estimator at iteration t is formed using a weighted mixture of these real and synthetic responses. Concretely, for the min- ℓ_2 -norm interpolator and ridge regression with regularization $\lambda > 0$, we define

$$\widehat{\boldsymbol{\beta}}_{t} = (\mathbf{X}^{\top} \mathbf{X})^{\dagger} \mathbf{X}^{\top} (w \mathbf{y}_{t} + (1 - w) \widetilde{\mathbf{y}}_{t}).$$

$$\widehat{\boldsymbol{\beta}}_{t,\lambda} = (\mathbf{X}^{\top} \mathbf{X} + n\lambda I)^{-1} \mathbf{X}^{\top} (w \mathbf{y}_{t} + (1 - w) \widetilde{\mathbf{y}}_{t,\lambda}).$$
(1.1)

Above, \dagger denotes the pseudoinverse. It is well-known that $\widehat{\boldsymbol{\beta}}_t = \lim_{\lambda \to 0^+} \widehat{\boldsymbol{\beta}}_{t,\lambda}$. In the aforementioned setting, our main contributions are as follows:

- (i) Quantifying the generalization error. In an overparametrized regime (stated precisely in Section 2), we characterize the generalization error as $t \to \infty$ for both the min- ℓ_2 -norm interpolator $\widehat{\boldsymbol{\beta}}_t$ (Theorem 3.1) and the ridge estimator $\widehat{\boldsymbol{\beta}}_{t,\lambda}$ (Theorem 3.4). Our results capture the precise dependence of the limiting risk on key problem parameters: the signal strength, feature covariance matrix, regularization level λ , and mixing proportion w. Our work substantially advances the growing literature on interpolation learning and high-dimensional ridge regression Montanari et al. (2019); Deng et al. (2022); Liang & Sur (2022); Wang et al. (2022); Zhou et al. (2022); Bach (2024); Patil et al. (2024); Mallinar et al. (2024); Song et al. (2024), which has previously not examined the impact of synthetic data on these learning problems.
- (ii) Characterizing the optimal mixing ratio. We characterize the optimal weight on real labels that minimizes the long-term prediction error across different settings. For $\min \ell_2$ -norm interpolation, we establish that the asymptotic risk is uniquely minimized at $w^* = 1/\varphi$ (the reciprocal of the golden ratio) for *any* feature covariance matrix with bounded eigenvalues (see equation (2.4)). This phenomenon was previously proved only for ordinary least squares in low-dimensional linear regression with Gaussian features He et al. (2025). For ridge regression, we prove that the risk is log-convex and admits a unique minimum at w^* under several important scenarios: when the covariance is isotropic (Theorem 3.3), or when the covariance follows a spiked model (Section 3.4.2), or when the signal follows a random-effects model (Theorem 3.2). Across all settings, we show $w^* \geq 1/2$, highlighting the necessity of weighting real-data more heavily than synthetic data. Such rigorous analysis of the mixing ratio provides concrete guidance for mitigating model collapse in overparametrized problems, complementing recent theoretical studies that were limited to low-dimensional regression Gerstgrasser et al. (2024); Dey & Donoho (2024).

Paper Structure The rest of the paper is structured as follows. Section 2 formalizes the problem setup and data-augmentation framework. In Section 3, we state our theoretical results – subsection 3.1 considers min- ℓ_2 -norm interpolator and subsections 3.2 and 3.3 consider ridge regression for isotropic and anisotropic covariates, respectively. Subsections 3.4.1 and 3.4.2 provide applications of our main results for random effects model and spike covariance models, respectively. Section 4 provides extensive simulations to complement our theoretical findings. Section 5 presents

a discussion and future research directions. The proofs of all theoretical results are in the supplementary material.

2 PROBLEM SETUP

We consider the fresh data augmentation framework He et al. (2025), but in the context of overparametrized linear regression. Suppose we observe a dataset (y, X) from a linear model, i.e.,

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}, \text{ with } \mathbf{y}, \boldsymbol{\varepsilon} \in \mathbb{R}^n, \boldsymbol{\beta} \in \mathbb{R}^p, \mathbf{X} \in \mathbb{R}^{n \times p}.$$
 (2.1)

We compute an initial ridge estimator $\widehat{\boldsymbol{\beta}}_{0,\lambda}$ using (\mathbf{y},\mathbf{X}) . For $t\geq 1$ we iteratively generate synthetic response vectors $\widetilde{\mathbf{y}}_{t,\lambda}$ using $\widehat{\boldsymbol{\beta}}_{t-1,\lambda}$, then augment these with fresh real responses \mathbf{y}_t . At each step, the next ridge estimator is computed using a mixture of the real and synthetic responses, \mathbf{y}_t and $\widetilde{\mathbf{y}}_{t,\lambda}$ respectively, with a mixing proportion $w\in(0,1)$. The procedure is outlined in Algorithm 1.

Algorithm 1 Iterative ridge with real/synthetic data augmentation

- 1: **Input:** Dataset (y, X); regularization parameter $\lambda > 0$; mixing proportion $w \in (0, 1)$.
- 2: Initialize: $\widehat{\boldsymbol{\beta}}_{0,\lambda} \leftarrow (\mathbf{X}^{\top}\mathbf{X} + n\lambda I)^{-1}\mathbf{X}^{\top}\mathbf{y}$.
- 3: **for** t > 1 **do**
- 4: Generate real responses: $\mathbf{y}_t \leftarrow \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}_t$.
- 5: Generate synthetic responses: $\widetilde{\mathbf{y}}_{t,\lambda} \leftarrow \mathbf{X}\widehat{\boldsymbol{\beta}}_{t-1,\lambda} + \widetilde{\boldsymbol{\varepsilon}}_t$.
- 6: Update estimator:

$$\widehat{\boldsymbol{\beta}}_{t,\lambda} \leftarrow (\mathbf{X}^{\top}\mathbf{X} + n\lambda I)^{-1}\mathbf{X}^{\top}(w\mathbf{y}_{t} + (1-w)\widetilde{\mathbf{y}}_{t,\lambda}).$$

7: end for

To capture an overparametrized regime, we assume that p>n with both diverging at a comparable rate, i.e. $p/n\to\gamma>1$. This means we work with a sequence of problem instances $\{\mathbf{y}(n),\mathbf{X}(n),\boldsymbol{\beta}(n),\varepsilon(n)\}_{n\geq 1}$, with $\mathbf{X}(n)\in\mathbb{R}^{n\times p(n)},\mathbf{y}(n),\varepsilon(n)\in\mathbb{R}^n,\boldsymbol{\beta}(n)\in\mathbb{R}^{p(n)}$, satisfying equation (2.1) and further assume that

$$\lim_{n \to \infty} \|\beta(n)\|^2 = b_* \in (0, \infty).$$
 (2.2)

Below we suppress the dependence on n for conciseness. This regime has seen incredible success in modern ML theory in explaining phenomena observed for deep neural networks and other practical algorithms Hastie et al. (2022); Adlam & Pennington (2020); Montanari et al. (2019); Mei & Montanari (2022); Liang & Sur (2022); Cui et al. (2023); Paquette et al. (2024); Emrullah Ildiz et al. (2025); Lu et al. (2025). The regime has also seen enormous utility in high-dimensional statistics, particularly for the development of new theory and methods in challenging contemporary inference problems Bean et al. (2013); El Karoui (2018); Donoho et al. (2009); Bayati & Montanari (2011); Wang et al. (2017); Sur & Candès (2019); Sur et al. (2019); Fan (2022); Li & Sur (2023); Celentano et al. (2023); Jiang et al. (2025).

In the sequel, we operate under the following assumptions on the covariates and errors that are commonly seen in random matrix theory Bai & Silverstein (2010).

Assumption 2.1. (i) The covariates satisfy $\mathbf{X} = \mathbf{Z} \mathbf{\Sigma}^{1/2}$, where $\mathbf{Z} \in \mathbb{R}^{n \times p}$ are random matrices whose entries Z_{ij} are independent random variables with zero mean and unit variance. We further assume that there exists a constant $\tau > 0$ by which the v-th moment of each entry is bounded for some v > 4:

$$\mathbb{E}\left[\left|Z_{ij}\right|^{v}\right] \le \tau^{-1}.\tag{2.3}$$

We will assume throughout that Σ has bounded eigenvalues s_1, \ldots, s_p :

$$\tau < s_n < \dots < s_1 < \tau^{-1}. \tag{2.4}$$

(ii) The noises ε_t , $\widetilde{\varepsilon}_t$ (defined precisely in Algorithm 1) are assumed to have i.i.d. entries with mean 0, variance σ^2 , and bounded moments up to any order. That is, for any $\phi > 0$, there exists a constant C_{ϕ} such that

$$\mathbb{E}\left[\left|\varepsilon_{t,1}\right|^{\phi}\right] \leq C_{\phi}, \qquad \mathbb{E}\left[\left|\widetilde{\varepsilon}_{t,1}\right|^{\phi}\right] \leq C_{\phi}. \tag{2.5}$$

Together with equation (2.2), Assumption 2.1(ii) implies that the signal-to-noise ratio (SNR) SNR := b_{\star}/σ^2 remains finite as $n, p \to \infty$ ensuring that we work under a non-trivial and interesting regime.

2.1 RISK

The out-of-sample prediction risk of an estimator $\hat{\beta}$ (hereafter simply referred to as risk) at a new data point (y_0, \mathbf{x}_0) is defined as

$$R(\widehat{\boldsymbol{\beta}}; \boldsymbol{\beta}) := \mathbb{E}[(\mathbf{x}_0^{\top} \widehat{\boldsymbol{\beta}} - \mathbf{x}_0^{\top} \boldsymbol{\beta})^2 | \mathbf{X}] = \mathbb{E}[\|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|_{\Sigma}^2 | \mathbf{X}], \tag{2.6}$$

where for a vector \mathbf{v} and matrix Σ , we define $\|\mathbf{v}\|_{\Sigma}^2 = \mathbf{v}^{\top} \Sigma \mathbf{v}$. Note that this risk has a σ^2 difference from the mean-squared prediction error for the new data point, which does not affect the relative performance and is therefore omitted. As defined, the risk involves expectation over both the randomness in the new test point (y_0, \mathbf{x}_0) and that in the noise variables. We define the risk conditional on the feature matrix \mathbf{X} , and our risk characterization results are high probability statements over the randomness in the covariates. Despite this dependence on covariates, we use the notation $R(\hat{\boldsymbol{\beta}}; \boldsymbol{\beta})$ since the context is clear. The risk admits a bias-variance decomposition:

$$R(\widehat{\boldsymbol{\beta}}; \boldsymbol{\beta}) = \underbrace{\|\mathbb{E}(\widehat{\boldsymbol{\beta}}|\mathbf{X}) - \boldsymbol{\beta}\|_{\boldsymbol{\Sigma}}^{2}}_{B(\widehat{\boldsymbol{\beta}}; \boldsymbol{\beta})} + \underbrace{\operatorname{Tr}[\operatorname{Cov}(\widehat{\boldsymbol{\beta}}|\mathbf{X})\boldsymbol{\Sigma}]}_{V(\widehat{\boldsymbol{\beta}}; \boldsymbol{\beta})}.$$
(2.7)

We next state our main results, which involve precise characterization of the risk of the estimator $\hat{\beta}_{t,\lambda}$ from Algorithm 1 and $\hat{\beta}_t$ defined by equation (3.3).

3 MAIN RESULTS

We begin by introducing two measures that feature crucially in the risk of the estimators $\hat{\boldsymbol{\beta}}_{t,\lambda}$ and $\hat{\boldsymbol{\beta}}_t$. Let v_1,\ldots,v_p denote the eigenvectors of $\boldsymbol{\Sigma}$, i.e., $\boldsymbol{\Sigma}=\sum_{k=1}^p s_k v_k v_k^{\top}$. Define the probability measures:

$$\widehat{H}_p(x) = \frac{1}{p} \sum_{k=1}^p \mathbf{1}_{s_k \le x}, \qquad \widehat{G}_p(x) = \frac{1}{\|\boldsymbol{\beta}\|_2^2} \sum_{k=1}^p \langle v_k, \boldsymbol{\beta} \rangle^2 \mathbf{1}_{s_k \le x}.$$
 (3.1)

Throughout we assume \widehat{H}_p and \widehat{G}_p converge weakly to probability measures H and G respectively. For any $z \in \mathbb{C}/\mathbb{R}^+$, define m(z) to be the solution to

$$m(z)^{-1} + z = \gamma \int \frac{x}{1 + m(z)x} dH.$$
 (3.2)

Remark 3.1. Existence and uniqueness of m(z) is well-known (c.f., (Knowles & Yin, 2017, Lemma 2.2)). Further, m(z) is the companion Stieltjes transform of the free convolution of H and MP_{γ} , where MP_{γ} is the Marchenko-Pastur distribution with parameter γ Marčenko & Pastur (1967).

3.1 Min- ℓ_2 -norm interpolator

In this section, we will analyze the behavior of $R(\hat{\beta}_t; \beta)$, where we use

$$\widehat{\boldsymbol{\beta}}_t = (\mathbf{X}^{\top} \mathbf{X})^{\dagger} \mathbf{X}^{\top} (w \mathbf{y}_t + (1 - w) \widetilde{\mathbf{y}}_t), \qquad \widetilde{\mathbf{y}}_t = \mathbf{X} \widehat{\boldsymbol{\beta}}_{t-1} + \widetilde{\boldsymbol{\epsilon}}_t$$
 (3.3)

in place of $\widehat{\beta}_{t,\lambda}$, $\widetilde{\mathbf{y}}_{t,\lambda}$ in Algorithm 1. This estimator is a convex combination of the min- ℓ_2 -norm interpolator computed on $(\mathbf{X}, \mathbf{y}_t)$ and $(\mathbf{X}, \widetilde{\mathbf{y}}_t)$. The generalization error of $\widehat{\boldsymbol{\beta}}_t$ is characterized below.

Theorem 3.1 (Interpolator Risk). In the setting of Section 2, the risk of $\hat{\beta}_t$, defined by equation (3.3), satisifes the following. For any $w \in (0,1)$, we have almost surely over the randomness in the covariates,

$$\lim_{n \to \infty} \lim_{t \to \infty} R(\widehat{\boldsymbol{\beta}}_t; \boldsymbol{\beta}) = \sigma^2 c(w) \mathcal{V} + b_{\star} \mathcal{B}$$
(3.4)

with $c(w) = (w^2 + (1 - w)^2)/w(2 - w)$ and

$$\mathcal{V} = \frac{m'(0)}{m(0)^2} - 1, \qquad \mathcal{B} = \frac{m'(0)}{m(0)^2} \left(\int \frac{x}{(1 + m(0)x)^2} dG \right). \tag{3.5}$$

Moreover, the limiting risk is minimized at $w^* = \varphi^{-1}$, where $\varphi = (1 + \sqrt{5})/2$ is the golden ratio.

Note Theorem 3.1 applies for any Σ obeying our assumptions. For the special case of isotropic features, i.e., $\Sigma = I$, the limiting risk simplifies to

$$\lim_{n\to\infty}\lim_{t\to\infty}R(\widehat{\boldsymbol{\beta}}_t;\boldsymbol{\beta})=\sigma^2c(w)\frac{1}{\gamma-1}+b_\star\Big(1-\frac{1}{\gamma}\Big).$$

Proof of Theorem 3.1 is available in Appendix A.

Effect of number of iteration t and mixing parameter w: In equation (3.4), the first term corresponds to the variance while the second to the bias (recall definitions from equation (2.7)). Note that the quantities $\sigma^2 \mathcal{V}$ and $b_\star \mathcal{B}$ coincides asymptotic variance and bias of min- ℓ_2 -norm interpolators in overparametrized regression (Hastie et al., 2022, Theorem 2). In fact, the bias term $B(\widehat{\boldsymbol{\beta}}_t, \boldsymbol{\beta})$ is independent of both the iteration t and mixing proportion w. The impact of mixing on the generalization error arises through only the variance, captured by the function c(w), which is minimized at $w^\star = \varphi^{-1} \approx 0.618$. This golden-ratio weighting phenomenon was previously observed for ordinary least squares in low-dimensional regression with Gaussian covariates He et al. (2025).

If w=0, then the generalization error $R(\widehat{\boldsymbol{\beta}}_t,\boldsymbol{\beta})\to\infty$ as $t\to\infty$, even for a finite sample size n. This implies that training solely on synthetic data results in model collapse, as seen also by Shumailov et al. (2024); Dohmatob et al. (2024a) for low-dimensional learning problems. Moreover, if the mixing proportion w>1/3, then by equation (A.2), we have the variance $V(\widehat{\boldsymbol{\beta}}_t;\boldsymbol{\beta})$ decreases monotonically with t for any fixed n. Since $w^*=\varphi^{-1}>1/3$, we observe that the generalization error also decreases monotonically for optimal mixing, thereby preventing model collapse.

Dynamic mixing: One might wonder whether the limiting generalization error can be further reduced by selecting the mixing proportion w_t adaptively at each generation to minimize $R(\hat{\boldsymbol{\beta}}_t; \boldsymbol{\beta})$ for any finite sample size n. We show in Section C that the optimal choice w_t^{\star} in this sequential setup satisfies the recursion

$$w_t^{\star} = \frac{1 + w_{t-1}^{\star}}{2 + w_{t-1}^{\star}}, \quad w_0 = 1.$$

It follows immediately that w_t^{\star} is decreasing, so if one is free to adjust the mixing proportion at every generation, the optimal strategy places progressively more weight on the synthetic data. Moreover, using $w_t^{\star} \to w^{\star}$ (as defined in Theorem 3.1), in the long run, the asymptotic risk is the same whether one uses a fixed w^{\star} across all generations or adapts w_t^{\star} dynamically, as proven in Section C.

Next, we study ridge regression. Here, both the variance and bias terms will depend on λ, t, w . For the sake of clarity, we present our results for isotropic covariates in Subsection 3.2. We discuss the case of non-isotropic covariates later in Subsection 3.3.

3.2 RIDGE REGRESSION: ISOTROPIC COVARIANCE ($\Sigma = \alpha I$)

In case of isotropic features, i.e., $\Sigma = \alpha I$, the bias and variance of $\widehat{\beta}_{t,\lambda}$ simplifies since $s_k \equiv \alpha$. Hence, by equation (3.1), $G = H = \delta_{\alpha}$, where δ_x denote the Dirac probability measure at x. This implies m(z), from equation (3.2), simplifies to be the unique solution to

$$m(z)^{-1} + z = \gamma \frac{\alpha}{1 + \alpha m(z)}.$$
 (3.6)

We use the notations $m_1 := m(-\lambda/w), m_2 := m(-\lambda/(2-w)).$

Theorem 3.2 (Isotropic risk). Assume $\Sigma = \alpha I$ and the setting of Section 2. For 0 < w < 1, $\lambda > 0$, we have almost surely over the randomness in the covariates

$$\lim_{n \to \infty} \lim_{t \to \infty} R(\widehat{\boldsymbol{\beta}}_{t,\lambda}; \boldsymbol{\beta}) = \sigma^2 c(w) \, \mathcal{V}_{\lambda} + b_{\star} \, \mathcal{B}_{\lambda},$$

where $c(w) = (w^2 + (1 - w)^2)/w(2 - w)$ and

$$\mathcal{B}_{\lambda} = \frac{\alpha/(1+\alpha m_1)^2}{1-\gamma \alpha^2 m_1^2/(1+\alpha m_1)^2}, \qquad \mathcal{V}_{\lambda} = \frac{w(2-w)}{2(1-w)} \frac{\gamma}{\lambda} \left(\frac{\alpha}{1+\alpha m_1} - \frac{\alpha}{1+\alpha m_2}\right). \tag{3.7}$$

Moreover, the limiting risk is a log-convex function of w and has a unique minimizer.

Theorem 3.2 characterizes the precise asymptotic risk as a function of the regularization parameter λ and the mixing proportion w. Furthermore, the proof of the result shows that the map $w\mapsto c(w)\,\mathcal{V}_\lambda$ is log-convex, and the map $w\mapsto\mathcal{B}_\lambda$ is decreasing and log-convex. Hence, we obtain that the limiting risk is log-convex in w, thereby admitting a unique minimizer. Further, it can be shown that both \mathcal{V}_λ and \mathcal{B}_λ are continuous functions of λ and $\lim_{\lambda\to 0^+}\mathcal{V}_\lambda=\mathcal{V}$ and $\lim_{\lambda\to 0^+}\mathcal{B}_\lambda=\mathcal{B}$, with \mathcal{V},\mathcal{B} defined as in equation (3.5). The proof of Theorem 3.2 shows that the bias at the t-th iterate $B(\widehat{\boldsymbol{\beta}}_{t,\lambda},\boldsymbol{\beta})$ depends on both t and w, unlike the bias of the min- ℓ_2 -norm interpolator. The following result characterizes the behavior of the optimal mixing parameter as a function of λ .

Theorem 3.3 (Isotropic Optimal Mixing). Under the assumptions of Theorem 3.2, let $w^*(\lambda)$ be the unique global minimizer of the limiting risk as defined in Theorem 3.2. Then $w^*(\lambda)$ is a continuous function of λ satisfying

- (i) $w^*(\lambda) \in [0.5, 1]$,
- (ii) $w^*(\lambda) \to \phi^{-1}$ as $\lambda \downarrow 0$,
- (iii) $w^*(\lambda) \to 1$ as $\lambda \uparrow \infty$.

Furthermore, $w^*(\lambda)$ is an increasing function of SNR.

The optimal mixing parameter minimizing the asymptotic risk is always at least one-half, emphasizing the importance of favoring real-data over synthetic data. In Figure 2 (b), we show empirically that w^* can be arbitrarily close to 0.5 for low SNR. Proofs of Theorem 3.2 and Theorem 3.3 are available in Appendix A and Appendix B respectively.

3.3 RIDGE REGRESSION: CORRELATED FEATURES

We now state our most general result for anisotropic covariates, which calculates the limiting generalization error of ridge regression for arbitrary measures \hat{H}_p , \hat{G}_p (recall equation (3.1)).

Theorem 3.4 (Ridge risk under correlated covariates). In the setting of Section 2, suppose $w \in (0,1)$, and $\lambda > 0$. Define $m_1 = m(-\lambda/w)$, $m_2 = m(-\lambda/(2-w))$, where $m(\cdot)$ is the unique solution to equation (3.2). Then almost surely over the randomness in the covariates,

$$\lim_{n \to \infty} \lim_{t \to \infty} R(\widehat{\boldsymbol{\beta}}_{t,\lambda}; \boldsymbol{\beta}) = \sigma^2 c(w) \, \mathcal{V}_{\lambda} + b_{\star} \, \mathcal{B}_{\lambda}, \quad \text{with}$$
 (3.8)

$$\mathcal{B}_{\lambda} = \left(\int \frac{x}{(1+m_1 x)^2} dG \right) \left(1 - \gamma \int \frac{m_1^2 x^2}{(1+m_1 x)^2} dH \right)^{-1}$$
(3.9)

$$\mathcal{V}_{\lambda} = \frac{w(2-w)}{2(1-w)} \frac{\gamma}{\lambda} \left(\int \frac{x}{1+m_1 x} dH - \int \frac{x}{1+m_2 x} dH \right). \tag{3.10}$$

Proof of Theorem 3.4 is available in Appendix A.

Theorem 3.4 is our most general result. It shows for any $w \in (0,1)$ and any $\lambda > 0$, $R(\hat{\boldsymbol{\beta}}_{t,\lambda};\boldsymbol{\beta})$ does not diverge even when t increases. The result highlights the necessity of mixing real-data with synthetic outputs to mitigate model collapse. In its most general form the risk in equation (3.8) is involved to analyze. In the following subsections, we study two popular models where the risk simplifies and the optimal mixing ratio can be studied analytically.

3.4 EXAMPLES

In this section, we study two popular classes of examples: (i) the random effects model (Section 3.4.1) and (ii) the spiked covariance model (Section 3.4.2), where structural assumptions on the signal or the population covariance matrix render the limiting risk more tractable.

3.4.1 RANDOM EFFECTS MODEL

Modern applications, ranging from text classification to economics and genomics, are characterized by dense but weak signals spread across many coordinates (Joachims, 1998; Boyle et al., 2017; Yang et al., 2010; Shen & Xiu, 2025). This setting is well captured by a random-effects model, where each feature contributes a small, independent effect, and it provides a simple yet sophisticated framework for rigorously analyzing interesting high-dimensional predictors (Dobriban & Wager, 2018). Adopting a random-effects framework, we assume that each coordinate β_i of the signal is drawn i.i.d. with $\mathbb{E}\beta_i=0$ and $\mathrm{Var}(\beta_i)=b_\star/p>0$. In this setting, the limiting risk can be characterized as follows.

Proposition 3.1 (Ridge risk under Random-Effects Models). Suppose $\beta_i \stackrel{i.i.d.}{\sim} (0, b_{\star}/p)$. Assume the framework of Section 2 and fix 0 < w < 1, $\lambda > 0$. Let $m(\cdot)$ be the solution to equation (3.2) and define $f(z) = m(-z)^{-1} - z$. Then almost surely over the randomness in the covariates

$$\lim_{n\to\infty} \lim_{t\to\infty} R(\widehat{\boldsymbol{\beta}}_{t,\lambda}; \boldsymbol{\beta}) = \sigma^2 c(w) \, \mathcal{V}_{\lambda} + b_{\star} \, \mathcal{B}_{\lambda}, \quad \text{where}$$

$$\mathcal{B}_{\lambda} = \frac{1}{\gamma} \left(f(\lambda/w) - \frac{\lambda}{w} f'(\lambda/w) \right), \quad \mathcal{V}_{\lambda} = \frac{f\left(\frac{\lambda}{w}\right) - f\left(\frac{\lambda}{2-w}\right)}{\frac{\lambda}{w} - \frac{\lambda}{2-w}}.$$
 (3.11)

Moreover, \mathcal{B}_{λ} is decreasing and log-convex and $c(w)\mathcal{V}_{\lambda}$ is log-convex.

The representation of risk via the risk function f enables to rigorously study properties of the optimal mixing ratio, as presented below.

Proposition 3.2. Under the setup of Theorem 3.1, (i) The generalization error has a unique minimizer w^* , (ii) $w^* \in [0.5, 1]$, with $w^* \to 1$ as $\lambda \uparrow \infty$ and $w^* \to \phi^{-1}$ with $\lambda \downarrow 0$ and (iii) w^* increases with SNR.

We provide some context for the random effects model. Our main Theorem 3.4 makes it clear that in presence of general covariance matrices, the variance in the limiting risk (first term in RHS of equation (3.8)) is determined by the spectrum of Σ , as captured through the limiting spectral distribution H. If the signal were deterministic, without additional assumptions, the bias (second term in the RHS of equation (3.8)) would depend on how β aligns with the eigenbasis of Σ . This is captured through the measure G. But in a random-effects setting, β lies in a generic position relative to Σ , which simplifies both the limiting bias and variance making equation (3.8) tractable to analyze as a function of w. Proofs of Propositions 3.1 and 3.2 are available in Appendix B.

A recent work Dohmatob et al. (2025) studies model collapse in Gaussian random-effects models but they consider a setting where the real and synthetic data are generated from different distributions, subsequently pooling these datasets and studying when model collapse occurs. Crucially, Dohmatob et al. (2025) do not utilize synthetic data generated from a fitted model, unlike in our setting. This leads to fundamental differences in our framework compared to theirs. Additionally, we consider a broad class of random-effects models, without requiring Gaussianity on the signals, and additionally provide guarantees for the optimal mixing ratio.

Remark 3.2. While Propositions 3.1 and 3.2 are stated under the random effects assumption, the conclusions continue to hold for a broader class of parameters (β, Σ) . If the probability measures \widehat{G}_p and \widehat{H}_p from equation (3.1) converge weakly to the same distribution, i.e., G=H, we have the same conclusions. In fact, we prove in Lemma D.1 that the random-effects assumption can be seen as a special case of G=H. Other natural examples where G=H include the isotropic covariance case, $\Sigma=\alpha I$, or where β is drawn uniformly at random from a p-dimensional sphere, or Σ is drawn from an orthogonally invariant ensemble.

3.4.2 SPIKED COVARIANCE MODEL

For our second application, we consider a popular class of covariance matrices – the spiked covariance model Birnbaum et al. (2013); Johnstone (2001); Johnstone & Onatski (2020). In the past two decades, this covariance class has seen exciting applications in population genetics Patterson et al. (2006); Price et al. (2006), finance Knight et al. (2005); Ledoit & Wolf (2022), and signal processing Johnstone & Lu (2009); Wang et al. (2024), among others.

Formally, we assume $\Sigma = I + svv^{\top}$ for some $v \in \mathbb{R}^p$, with $||v||_2 = 1$, and s > 0. The results below extend naturally to spiked models with multiple but finitely many spikes, but for simplicity, we study the problem for the case of a single spike. The limiting risk takes the following form.

Proposition 3.3. In the setting of Section 2, assume that $\Sigma = I + svv^{\top}$ and the signal takes the form $\beta = \theta v + \sqrt{1 - \theta^2}v^{\perp}$, while satisfying equation (2.2), with $v^{\top}v^{\perp} = 0$ and $\|v^{\perp}\|_2 = 1$. If $\theta \equiv \theta(n) \to \theta_{\star}$, then we have almost surely over the randomness of covariates that

$$\lim_{n \to \infty} \lim_{t \to \infty} R(\widehat{\boldsymbol{\beta}}_{t,\lambda}; \boldsymbol{\beta}) = \sigma^2 c(w) \, \mathcal{V}_{\lambda} + b_{\star} \, \mathcal{B}_{\lambda},$$

where V_{λ} is the same as in Theorem 3.2 and

$$\mathcal{B}_{\lambda} = \left(\theta_{\star}^{2} \frac{1+s}{(1+m_{1}(1+s))^{2}} + (1-\theta_{\star}^{2}) \frac{1}{(1+m_{1})^{2}}\right) \left(1-\gamma \frac{m_{1}^{2}}{(1+m_{1})^{2}}\right)^{-1}.$$

Further, the limiting risk is uniquely minimized at a w^* satisfying the conclusions of Theorem 3.3.

Note that, if $\theta_{\star} \neq 0$, then $G \neq H$. Hence, this spiked matrix case differs fundamentally from the settings discussed in Remark 3.2 and the proof of Proposition 3.3 does not follow directly from the proof of Proposition 3.2. Proof of Proposition 3.3 is available in Appendix B.

4 SIMULATIONS

In this section, we conduct numerical experiments to complement our theoretical findings. For all empirical risk plots, we generate an $n \times p$ feature matrix $\mathbf{X} = \mathbf{Z} \mathbf{\Sigma}^{1/2}$, where $Z_{ij} \sim \mathcal{N}(0,1)$ and vary $\mathbf{\Sigma}$ across different plots. We display both the theoretical risk predicted by our formulae (solid lines) and corresponding empirical estimates (× markers). The empirical risks are calculated by averaging over 100 runs.

Risk of min- ℓ_2 -norm Interpolator In Figure 1, we plot the asymptotic generalization error of the min- ℓ_2 -norm interpolator as a function of mixing weight w for two different classes of covariance matrices: i. $\Sigma = I$ (Panel (a)) and ii. $\Sigma_{ij} = 2^{-|i-j|}$ (Panel (b)), corresponding to the correlation matrix of an AR(1) process. We vary $\gamma = 1.5, 2, 3$. The choice of the remaining parameters are as follows: sample size n = 200, number of features $p = \gamma n$, and number of iterations t = 5. To generate β , we first simulate $\widetilde{\beta}$ with $\widetilde{\beta}_i \overset{\text{i.i.d.}}{\sim} \text{Bern}(0.1)$. Set $\beta = \widetilde{\beta}/\|\widetilde{\beta}\|_2$ which yields $b_\star = 1$.

We observe that the empirical risk matches with its theoretical counterpart even for moderate sample size. Further, the generalization error is always minimized at $w^\star = \varphi^{-1} \approx 0.618$ (dashed vertical line), consistent with Theorem 3.1. Panel (c) shows the risk of the min- ℓ_2 -norm interpolator as a function of iteration t. We have computed the risk at optimal mixing weight $w = \varphi^{-1}$ and $\Sigma = I$. We observe that both theoretical and empirical risks stabilize after only a few iterations.

Optimal weights In Figure 2, we consider several properties of generalization error of ridge estimator $\hat{\boldsymbol{\beta}}_{t,\lambda}$. In Figure 2(a), (b), we plot optimal weight w^* as a function of λ . We set $n=200, p=\gamma n$, and t=5, and vary $\gamma=1.2,2,4$. Figure 2 (a) considers isotropic covariance $\Sigma=I$ with high noise variance $\sigma^2=64$. The plot demonstrates that for low SNR, w^* can be arbitrarily close to 0.5. Further, $w^*(\lambda)$ is neither monotone, nor convex as a function of λ with $w^*(\lambda) \to \varphi^{-1}$ as $\lambda \to 0+$.

Figure 2 (b) corresponds to spiked covariance model $\Sigma = I + 5e_1e_1^{\top}$, where e_1, \ldots, e_p denote the canonical basis of \mathbb{R}^p . We set $\beta_1 = 0.5$ and $\boldsymbol{\beta}_{2:p} = \sqrt{1 - 0.5^2} \times \widetilde{\boldsymbol{\beta}} / \|\widetilde{\boldsymbol{\beta}}\|$ where $\widetilde{\beta}_i \overset{\text{i.i.d.}}{\sim} \text{Bern}(0.25)$. This implies that $\|\boldsymbol{\beta}\|^2 = 1$ and $\theta_{\star} = 0.5$, where θ_{\star} defined as in Proposition 3.3. The plot shows that for large regularization parameter λ , we have $w^{\star} = 1$, consistent with Proposition 3.3.

Figure 2 (c) plots the risk of $\widehat{\pmb{\beta}}_{t,\lambda}$ with $\sigma^2=1,$ $w=\varphi^{-1}.$ We set $\pmb{\Sigma}$ to be equicorrelated, i.e.,

$$oldsymbol{\Sigma} = \left(1 - rac{
ho}{\sqrt{p}}
ight)I + rac{
ho}{\sqrt{p}}\, oldsymbol{1}oldsymbol{1}^{ op}, \qquad ext{with }
ho = 1/2,$$

and $\beta_i \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, p^{-1})$. The empirical risk is computed with $\sigma^2 = 1, t = 10, n = 400$. Note that Σ does not satisfy bounded eigenvalue condition in Assumption 2.1, as its largest eigenvalue $\geq C\sqrt{p}$ for some C>0. Nevertheless, the generalization error is accurately predicted by Proposition 3.1 as long as β lies in a generic position relative to Σ . This demonstrates the robustness of our theoretical findings.

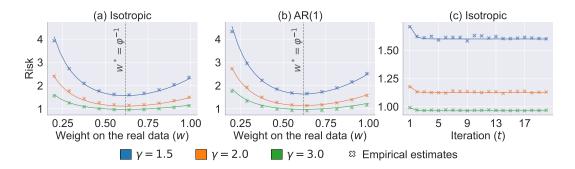


Figure 1: Generalization error of $\min \ell_2$ -norm interpolator as a function of weight w (Panel (a) and (b)) and iterations t (Panel (c)) for different values of γ . Panel (a) considers isotropic covariance $\Sigma = I$ and panel (b) considers anisotropic Σ with $\Sigma_{ij} = 2^{-|i-j|}$, which corresponds to covariance matrix of AR(1) model. In panes (a) and (b), the risk is minimized at $w^* = 1/\varphi$, as proved by Theorem 3.1. Panel (c) shows that both empirical and theoretical risks stabilize in a few iterations.

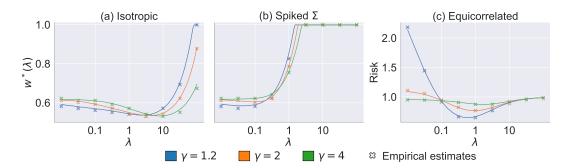


Figure 2: In panels (a) and (b), we plot the optimal mixing weight w^\star as a function of λ for different values of γ and two classes of covariance matrices:. Panel (a) considers isotropic covariance $\Sigma = I$ with high noise $\sigma^2 = 64$, demonstrating w^\star can be close to 0.5 for low SNR. Panel (b) plots w^\star for the spiked covariance matrix showing $w^\star = 1$ for large λ . Panel (c) plots the generalization error as a function of λ for Σ equicorrelated matrix. Here, empirical risks align with theoretical predictions given by Proposition 3.1, even though Σ violates Assumption 2.1, illustrating the robustness of our results.

5 Discussion

We provide a rigorous analysis of model collapse under overparametrization for linear models. As an overarching theme, we demonstrate how mixing real-data with synthetic outputs mitigates model collapse, and identify optimal mixing ratios that minimize prediction error in this context. As a promising next direction, understanding how model collapse affects interpolators in other ℓ_p geometries would be crucial. Such interpolators arise as implicit regularized limits of popular algorithms Gunasekar et al. (2018b) and typically require techniques beyond random matrix theory—a critical tool employed in our analysis. Additionally, approaches that extend linear arguments to non-linear high-dimensional problems (c.f., Sur & Candès (2019); Hu & Lu (2022)) should enable our qualitative conclusions to generalize to structured non-linear models. This includes generalized linear models, single-index models, and even non-parametric models through parametric-to-non-parametric equivalence techniques introduced in Lahiry & Sur (2023). Finally, extending our results to more complex architectures remains an important challenge. One promising approach involves studying multi-index models and their sequential variants, which capture classes of neural networks and transformer architectures Troiani et al. (2024; 2025); Cui (2025). Investigating model collapse and mitigation strategies in these contexts presents an exciting avenue for future research.

REFERENCES

- Ben Adlam and Jeffrey Pennington. Understanding double descent requires a fine-grained biasvariance decomposition. *Advances in neural information processing systems*, 33:11022–11032, 2020.
- Sina Alemohammad, Josue Casco-Rodriguez, Lorenzo Luzi, Ahmed Imtiaz Humayun, Hossein Babaei, Daniel LeJeune, Ali Siahkoohi, and Richard G Baraniuk. Self-consuming generative models go mad. International Conference on Learning Representations (ICLR), 2024.
- Francis Bach. High-dimensional analysis of double descent for linear regression with random projections. SIAM Journal on Mathematics of Data Science, 6(1):26–50, 2024.
- Zhidong Bai and Jack W Silverstein. *Spectral analysis of large dimensional random matrices*, volume 20. 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.
- Mohsen Bayati and Andrea Montanari. The lasso risk for gaussian matrices. *IEEE Transactions on Information Theory*, 58(4):1997–2017, 2011.
 - Derek Bean, Peter J Bickel, Noureddine El Karoui, and Bin Yu. Optimal m-estimation in high-dimensional regression. *Proceedings of the National Academy of Sciences*, 110(36):14563–14568, 2013.
- Quentin Bertrand, Avishek Joey Bose, Alexandre Duplessis, Marco Jiralerspong, and Gauthier Gidel. On the stability of iterative retraining of generative models on their own data. In *ICLR*, 2024.
- Aharon Birnbaum, Iain M Johnstone, Boaz Nadler, and Debashis Paul. Minimax bounds for sparse pca with noisy high-dimensional data. *Annals of statistics*, 41(3):1055, 2013.
- Matyas Bohacek and Hany Farid. Nepotistically trained generative-ai models collapse. *arXiv e-prints*, pp. arXiv–2311, 2023.
- Evan A Boyle, Yang I Li, and Jonathan K Pritchard. An expanded view of complex traits: from polygenic to omnigenic. *Cell*, 169(7):1177–1186, 2017.
- Martin Briesch, Dominik Sobania, and Franz Rothlauf. Large language models suffer from their own output: An analysis of the self-consuming training loop. *CoRR*, 2023.
- Michael Celentano, Andrea Montanari, and Yuting Wei. The lasso with general gaussian designs with applications to hypothesis testing. *The Annals of Statistics*, 51(5):2194–2220, 2023.
- Hugo Cui. High-dimensional learning of narrow neural networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2025(2):023402, 2025.
- Hugo Cui, Florent Krzakala, and Lenka Zdeborová. Bayes-optimal learning of deep random networks of extensive-width. In *International Conference on Machine Learning*, pp. 6468–6521. PMLR, 2023.
 - Zeyu Deng, Abla Kammoun, and Christos Thrampoulidis. A model of double descent for high-dimensional binary linear classification. *Information and Inference: A Journal of the IMA*, 11(2): 435–495, 2022.
 - Apratim Dey and David Donoho. Universality of the $\pi^2/6$ pathway in avoiding model collapse. arXiv preprint arXiv:2410.22812, 2024.
 - Edgar Dobriban and Stefan Wager. High-dimensional asymptotics of prediction: Ridge regression and classification. *The Annals of Statistics*, 46(1):247–279, 2018.
 - Elvis Dohmatob, Yunzhen Feng, and Julia Kempe. Model collapse demystified: The case of regression. *Advances in Neural Information Processing Systems*, 37:46979–47013, 2024a.

- Elvis Dohmatob, Yunzhen Feng, Pu Yang, Francois Charton, and Julia Kempe. A tale of tails: model collapse as a change of scaling laws. In *Proceedings of the 41st International Conference on Machine Learning*, ICML'24. JMLR.org, 2024b.
 - Elvis Dohmatob, Yunzhen Feng, Arjun Subramonian, and Julia Kempe. Strong model collapse. In Y. Yue, A. Garg, N. Peng, F. Sha, and R. Yu (eds.), *International Conference on Representation Learning*, volume 2025, pp. 15656–15691, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/file/284afdc2309f9667d2d4fb9290235b0c-Paper-Conference.pdf.
 - David L Donoho, Arian Maleki, and Andrea Montanari. Message-passing algorithms for compressed sensing. *Proceedings of the National Academy of Sciences*, 106(45):18914–18919, 2009.
 - Noureddine El Karoui. On the impact of predictor geometry on the performance on high-dimensional ridge-regularized generalized robust regression estimators. *Probability Theory and Related Fields*, 170(1):95–175, 2018.
 - 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 *13th International Conference on Learning Representations*, 2025.
 - Zhou Fan. Approximate message passing algorithms for rotationally invariant matrices. *The Annals of Statistics*, 50(1):197–224, 2022.
 - Yunzhen Feng, Elvis Dohmatob, Pu Yang, Francois Charton, and Julia Kempe. Beyond model collapse: Scaling up with synthesized data requires verification. In *The Thirteenth International Conference on Learning Representations*, 2025.
 - 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.
 - Suriya Gunasekar, Jason Lee, Daniel Soudry, and Nathan Srebro. Characterizing implicit bias in terms of optimization geometry. In *International Conference on Machine Learning*, pp. 1832–1841. PMLR, 2018a.
 - Suriya Gunasekar, Jason D Lee, Daniel Soudry, and Nati Srebro. Implicit bias of gradient descent on linear convolutional networks. *Advances in neural information processing systems*, 31, 2018b.
 - 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.
 - Ryuichiro Hataya, Han Bao, and Hiromi Arai. Will large-scale generative models corrupt future datasets? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 20555–20565, 2023.
 - Hengzhi He, Shirong Xu, and Guang Cheng. Golden ratio weighting prevents model collapse. *arXiv* preprint arXiv:2502.18049, 2025.
 - Hong Hu and Yue M Lu. Universality laws for high-dimensional learning with random features. *IEEE Transactions on Information Theory*, 69(3):1932–1964, 2022.
 - Kuanhao Jiang, Rajarshi Mukherjee, Subhabrata Sen, and Pragya Sur. A new central limit theorem for the augmented ipw estimator: Variance inflation, cross-fit covariance and beyond. *The Annals of Statistics*, 53(2):647–675, 2025.
 - Thorsten Joachims. Text categorization with support vector machines: Learning with many relevant features. In *European conference on machine learning*, pp. 137–142. Springer, 1998.
 - Iain M Johnstone. On the distribution of the largest eigenvalue in principal components analysis. *The Annals of statistics*, 29(2):295–327, 2001.

- Iain M Johnstone and Arthur Yu Lu. On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association*, 104(486):682–693, 2009.
 - Iain M Johnstone and Alexei Onatski. Testing in high-dimensional spiked models. *The Annals of Statistics*, 48(3), 2020.
 - Joshua Kazdan, Rylan Schaeffer, Apratim Dey, Matthias Gerstgrasser, Rafael Rafailov, David L Donoho, and Sanmi Koyejo. Collapse or thrive: Perils and promises of synthetic data in a self-generating world. In *Forty-second International Conference on Machine Learning*, 2024.
 - John L Knight, Stephen Satchell, and Chris Adcock. *Linear factor models in finance*. Elsevier, 2005.
 - Antti Knowles and Jun Yin. Anisotropic local laws for random matrices. *Probability Theory and Related Fields*, 169(1):257–352, 2017.
 - Samriddha Lahiry and Pragya Sur. Universality in block dependent linear models with applications to nonparametric regression. *arXiv preprint arXiv:2401.00344*, 2023.
 - Olivier Ledoit and Michael Wolf. The power of (non-) linear shrinking: A review and guide to covariance matrix estimation. *Journal of Financial Econometrics*, 20(1):187–218, 2022.
 - Yufan Li and Pragya Sur. Spectrum-aware debiasing: A modern inference framework with applications to principal components regression. *arXiv* preprint arXiv:2309.07810, 2023.
 - Tengyuan Liang and Pragya Sur. A precise high-dimensional asymptotic theory for boosting and minimum- ℓ_1 -norm interpolated classifiers. *The Annals of Statistics*, 50(3), 2022.
 - Yue M Lu, Mary Letey, Jacob A Zavatone-Veth, Anindita Maiti, and Cengiz Pehlevan. Asymptotic theory of in-context learning by linear attention. *Proceedings of the National Academy of Sciences*, 122(28):e2502599122, 2025.
 - Neil Mallinar, Austin Zane, Spencer Frei, and Bin Yu. Minimum-norm interpolation under covariate shift. *arXiv preprint arXiv:2404.00522*, 2024.
 - Vladimir A Marčenko and Leonid Andreevich Pastur. Distribution of eigenvalues for some sets of random matrices. *Mathematics of the USSR-Sbornik*, 1(4):457, 1967.
 - Gonzalo Martínez, Lauren Watson, Pedro Reviriego, José Alberto Hernández, Marc Juarez, and Rik Sarkar. Combining generative artificial intelligence (ai) and the internet: Heading towards evolution or degradation? *arXiv preprint arXiv:2303.01255*, 2023a.
 - Gonzalo Martínez, Lauren Watson, Pedro Reviriego, José Alberto Hernández, Marc Juarez, and Rik Sarkar. Towards understanding the interplay of generative artificial intelligence and the internet. In *International Workshop on Epistemic Uncertainty in Artificial Intelligence*, pp. 59–73. Springer, 2023b.
 - Song Mei and Andrea Montanari. The generalization error of random features regression: Precise asymptotics and the double descent curve. *Communications on Pure and Applied Mathematics*, 75(4):667–766, 2022.
 - Andrea Montanari, Feng Ruan, Youngtak Sohn, and Jun Yan. The generalization error of max-margin linear classifiers: Benign overfitting and high dimensional asymptotics in the over-parametrized regime. *The Annals of Statistics (to appear)*, 2019.
 - Vidya Muthukumar, Kailas Vodrahalli, Vignesh Subramanian, and Anant Sahai. Harmless interpolation of noisy data in regression. *IEEE Journal on Selected Areas in Information Theory*, 1(1): 67–83, 2020.
 - Elliot Paquette, Courtney Paquette, Lechao Xiao, and Jeffrey Pennington. 4+ 3 phases of compute-optimal neural scaling laws. *Advances in Neural Information Processing Systems*, 37:16459–16537, 2024.

- Pratik Patil, Jin-Hong Du, and Ryan J Tibshirani. Optimal ridge regularization for out-of-distribution prediction. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 39908–39954, 2024.
- Nick Patterson, Alkes L Price, and David Reich. Population structure and eigenanalysis. *PLoS genetics*, 2(12):e190, 2006.
 - Alkes L Price, Nick J Patterson, Robert M Plenge, Michael E Weinblatt, Nancy A Shadick, and David Reich. Principal components analysis corrects for stratification in genome-wide association studies. *Nature genetics*, 38(8):904–909, 2006.
 - René L Schilling, Renming Song, and Zoran Vondraček. *Bernstein functions: theory and applications*. Walter de Gruyter, 2009.
 - Zhouyu Shen and Dacheng Xiu. Can machines learn weak signals? Technical report, National Bureau of Economic Research, 2025.
 - Ilia Shumailov, Zakhar Shumaylov, Yiren Zhao, Yarin Gal, Nicolas Papernot, and Ross Anderson. The curse of recursion: Training on generated data makes models forget. *arXiv preprint arXiv:2305.17493*, 2023.
 - Ilia 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.
 - Yanke Song, Sohom Bhattacharya, and Pragya Sur. Generalization error of min-norm interpolators in transfer learning. *arXiv preprint arXiv:2406.13944*, 2024.
 - Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro. The implicit bias of gradient descent on separable data. *Journal of Machine Learning Research*, 19(70): 1–57, 2018.
 - Pragya Sur and Emmanuel J Candès. A modern maximum-likelihood theory for high-dimensional logistic regression. *Proceedings of the National Academy of Sciences*, 116(29):14516–14525, 2019.
 - Pragya Sur, Yuxin Chen, and Emmanuel J Candès. The likelihood ratio test in high-dimensional logistic regression is asymptotically a rescaled chi-square. *Probability theory and related fields*, 175(1):487–558, 2019.
 - Emanuele Troiani, Yatin Dandi, Leonardo Defilippis, Lenka Zdeborová, Bruno Loureiro, and Florent Krzakala. Fundamental computational limits of weak learnability in high-dimensional multi-index models. *arXiv preprint arXiv:2405.15480*, 2024.
 - Emanuele Troiani, Hugo Cui, Yatin Dandi, Florent Krzakala, and Lenka Zdeborová. Fundamental limits of learning in sequence multi-index models and deep attention networks: High-dimensional asymptotics and sharp thresholds. *arXiv* preprint arXiv:2502.00901, 2025.
 - Guillaume Wang, Konstantin Donhauser, and Fanny Yang. Tight bounds for minimum ℓ_1 -norm interpolation of noisy data. In *International Conference on Artificial Intelligence and Statistics*, pp. 10572–10602. PMLR, 2022.
 - Shuaiwen Wang, Haolei Weng, and Arian Maleki. Which bridge estimator is optimal for variable selection? *arXiv preprint arXiv:1705.08617*, 2017.
 - Zhichao Wang, Denny Wu, and Zhou Fan. Nonlinear spiked covariance matrices and signal propagation in deep neural networks. In *The Thirty Seventh Annual Conference on Learning Theory*, pp. 4891–4957. PMLR, 2024.
 - Jian Yang, Beben Benyamin, Brian P McEvoy, Scott Gordon, Anjali K Henders, Dale R Nyholt, Pamela A Madden, Andrew C Heath, Nicholas G Martin, Grant W Montgomery, et al. Common snps explain a large proportion of the heritability for human height. *Nature genetics*, 42(7):565–569, 2010.

Tong Zhang and Bin Yu. Boosting with early stopping: Convergence and consistency. *The Annals of Statistics*, 33(4):1538, 2005.

Lijia Zhou, Frederic Koehler, Pragya Sur, Danica J Sutherland, and Nati Srebro. A non-asymptotic moreau envelope theory for high-dimensional generalized linear models. *Advances in Neural Information Processing Systems*, 35:21286–21299, 2022.

Supplementary Material

A Proof of Theorem 3.1, Theorem 3.2, and Theorem 3.4

A.1 NOTATIONS

 Our objective is to derive the asymptotic generalization error of $\hat{\beta}_{t,\lambda}$ given by Algorithm 1 and $\hat{\beta}_t$ defined by equation (3.3). We define the scaled sample covariance matrix and its resolvent as

$$M := \mathbf{X}^{\top} \mathbf{X}, \qquad A_{\lambda} = (M + n\lambda I)^{-1}.$$

Note that A_{λ} and M are simultaneously diagonalizable and thus commute with each other. If $\lambda=0$, we use the notation $A_0=(\mathbf{X}^{\top}\mathbf{X})^{\dagger}$. The mixing proportion of the real-data is denoted by w and we define $\widetilde{w}:=1-w$. For two sequence of random variables u_n,v_n , we use the notation $u_n\sim v_n$ if $u_n/v_n\stackrel{P}{\longrightarrow} 1$ as $n\to\infty$.

A.2 VARIANCE

We use the notations of Subsection A.1 throughout the proof. First, we consider the case $\lambda > 0$.

Lemma A.1 (Ridge $\hat{\boldsymbol{\beta}}_{t,\lambda}$ Covariance). For $\lambda > 0$, the covariance matrix $Cov[\hat{\boldsymbol{\beta}}_{t,\lambda}|\mathbf{X}]$ is given by

$$\sigma^{2} \left[(w^{2} + \widetilde{w}^{2}) \sum_{k=1}^{t} \widetilde{w}^{2(k-1)} M^{2k-1} A_{\lambda}^{2k} + \widetilde{w}^{2t} M^{2t+1} A_{\lambda}^{2t+2} \right]$$
(A.1)

Proof of Lemma A.1. We will prove equation (A.1) by induction. For the base case (t = 0), we have $\widehat{\boldsymbol{\beta}}_{0,\lambda} = A_{\lambda} \mathbf{X}^{\top} \mathbf{y}$. Hence,

$$\begin{aligned} \mathsf{Cov}[\widehat{\boldsymbol{\beta}}_{0,\lambda}|\mathbf{X}] &= A_{\lambda}\mathbf{X}^{\top}\mathsf{Cov}[\mathbf{y}|\mathbf{X}]\mathbf{X}A_{\lambda} \\ &= A_{\lambda}\mathbf{X}^{\top}(\sigma^{2}I)\mathbf{X}A_{\lambda} = \sigma^{2}A_{\lambda}^{2}M \end{aligned}$$

This proves the base case. Now let's assume equation (A.1) holds for some t. Then, for iteration (t+1), we have

$$\begin{split} \mathsf{Cov}[\widehat{\pmb{\beta}}_{t+1,\lambda}|\mathbf{X}] &= A_{\lambda}\mathbf{X}^{\top}\mathsf{Cov}[w\mathbf{y}_{t} + \widetilde{w}\mathbf{y}_{t}|\mathbf{X}]\mathbf{X}A_{\lambda} \\ &= A_{\lambda}\mathbf{X}^{\top}(w^{2}\sigma^{2}I + \widetilde{w}^{2}\mathbf{X}\mathsf{Cov}[\widehat{\pmb{\beta}}_{t,\lambda}|\mathbf{X}]\mathbf{X}^{\top} + \sigma^{2}\widetilde{w}^{2}I)\mathbf{X}A_{\lambda} \\ &= \widetilde{w}^{2}A_{\lambda}M\mathsf{Cov}[\widehat{\pmb{\beta}}_{t,\lambda}|\mathbf{X}]MA_{\lambda} + \sigma^{2}(w^{2} + \widetilde{w}^{2})MA_{\lambda}^{2} \\ &= \widetilde{w}^{2}A_{\lambda}M\left(\sigma^{2}\left[(w^{2} + \widetilde{w}^{2})\sum_{k=1}^{t}\widetilde{w}^{2(k-1)}M^{2k-1}A_{\lambda}^{2k} + \widetilde{w}^{2t}M^{2t+1}A_{\lambda}^{2t+2}\right]\right)MA_{\lambda} \\ &+ \sigma^{2}(w^{2} + \widetilde{w}^{2})MA_{\lambda}^{2} \\ &= \sigma^{2}\left[(w^{2} + \widetilde{w}^{2})\sum_{k=1}^{t}\widetilde{w}^{2k}M^{2k+1}A_{\lambda}^{2k+2} + \widetilde{w}^{2t+2}M^{2t+3}A_{\lambda}^{2t+4}\right] \\ &+ \sigma^{2}(w^{2} + \widetilde{w}^{2})MA_{\lambda}^{2} \\ &= \sigma^{2}\left[(w^{2} + \widetilde{w}^{2})\sum_{k=1}^{t+1}\widetilde{w}^{2(k-1)}M^{2k-1}A_{\lambda}^{2k} + \widetilde{w}^{2t+2}M^{2t+3}A_{\lambda}^{2t+4}\right] \end{split}$$

This completes the proof of equation (A.1).

Corollary A.1. Using equation (2.7), $V(\widehat{\beta}_{t,\lambda}; \beta)$ is given by

$$V(\boldsymbol{\beta}_{t,\lambda};\boldsymbol{\beta}) = \sigma^2 \left[(w^2 + \widetilde{w}^2) \sum_{k=1}^t \widetilde{w}^{2(k-1)} \operatorname{Tr}[B_{\lambda}^k] + \widetilde{w}^{2t} \operatorname{Tr}[B_{\lambda}^{t+1}] \right]$$
(A.2)

where $B_{\lambda}^{k}=M^{2k-1}A_{\lambda}^{2k}\mathbf{\Sigma}$

Next, note that $\sigma_{\max}(MA_{\lambda}) < 1$ and therefore $\mathsf{Cov}[\widehat{\boldsymbol{\beta}}_{t,\lambda}|\mathbf{X}]$ converges in matrix norm to the following limit:

$$\begin{split} \lim_{t \to \infty} \mathsf{Cov}[\widehat{\boldsymbol{\beta}}_{t,\lambda} | \mathbf{X}] &= \sigma^2(w^2 + (1-w)^2) \sum_{k=0}^{\infty} (1-w)^{2k} (\mathbf{X}^{\top} \mathbf{X})^{2k+1} A_{\lambda}^{2k+2} \\ &= \sigma^2(w^2 + (1-w)^2) \mathbf{X}^{\top} \mathbf{X} A_{\lambda}^2 (I - (1-w)^2 (\mathbf{X}^{\top} \mathbf{X})^2 A_{\lambda}^2)^{-1} \\ &= \sigma^2(w^2 + \widetilde{w}^2) \mathbf{X}^{\top} \mathbf{X} (n\lambda + w \mathbf{X}^{\top} \mathbf{X})^{-1} (n\lambda + (2-w) \mathbf{X}^{\top} \mathbf{X})^{-1} \\ &= \frac{\sigma^2(w^2 + (1-w)^2)}{2(1-w)} \left[(w \mathbf{X}^{\top} \mathbf{X} + n\lambda I)^{-1} - ((2-w) \mathbf{X}^{\top} \mathbf{X} + n\lambda I)^{-1} \right] \\ \Longrightarrow \lim_{t \to \infty} V(\widehat{\boldsymbol{\beta}}_{t,\lambda}; \boldsymbol{\beta}) &= \frac{\sigma^2(w^2 + (1-w)^2)}{2(1-w)} \left(\frac{1}{w} \operatorname{Tr}[A_{\lambda/w} \boldsymbol{\Sigma}] - \frac{1}{2-w} \operatorname{Tr}[A_{\lambda/(2-w)} \boldsymbol{\Sigma}] \right) \end{split}$$

where the second to last equality follows using the following equality

$$\frac{x}{(n\lambda+wx)(n\lambda+(2-w)x)} = \frac{1}{2(1-w)} \left(\frac{1}{n\lambda+wx} - \frac{1}{n\lambda+(2-w)x} \right),$$

combined with diagonalization of $\mathbf{X}^{\top}\mathbf{X}$. Thus, to compute the limiting variance, it is enough to calculate

$$\lim_{n \to \infty} \text{Tr}[A_{-z} \mathbf{\Sigma}] = \lim_{n \to \infty} \frac{1}{n} \text{Tr}\left[(n^{-1} \mathbf{X}^{\top} \mathbf{X} - zI_p)^{-1} \mathbf{\Sigma} \right],$$

for $z \in \mathbb{C}/\mathbb{R}^+$. Using averaged local law Knowles & Yin (2017), we obtain

$$\frac{1}{n}\operatorname{Tr}\left[(n^{-1}\mathbf{X}^{\top}\mathbf{X} - zI_p)^{-1}\mathbf{\Sigma}\right] \xrightarrow{\text{a.s.}} \frac{-1}{z}\gamma\int \frac{x}{mx+1}dH(x),$$

where m(z) is as defined in equation (3.2). Combining the above display with Corollary A.1, we obtain that

$$\lim_{n \to \infty} \lim_{t \to \infty} V(\widehat{\boldsymbol{\beta}}_t; \boldsymbol{\beta}) = \sigma^2 \frac{\gamma}{\lambda} \frac{(w^2 + (1 - w)^2)}{2(1 - w)} \left[\int \frac{x}{1 + m_1 x} dH - \int \frac{x}{1 + m_2 x} dH \right], \quad (A.3)$$

where m_1 and m_2 are as defined in equation (3.10). This completes the proof of the expression of variance for $\lambda > 0$.

Next, we turn our attention to the case $\lambda = 0$. Here, we have the following expression of covariance.

Lemma A.2. For $t \ge 1$, the covariance matrix $Cov[\widehat{\boldsymbol{\beta}}_t | \mathbf{X}]$ is given by

$$\mathsf{Cov}[\widehat{\boldsymbol{\beta}}_t | \mathbf{X}] = \sigma^2 (\mathbf{X}^\top \mathbf{X})^\dagger \left[(w^2 + (1 - w)^2) \sum_{k=1}^t (1 - w)^{2(k-1)} + (1 - w)^{2t} \right], \tag{A.4}$$

where $\hat{\beta}_t$ is given by equation (3.3).

The proof follows the same steps as proof of Lemma A.1 once we note that $(\mathbf{X}^{\top}\mathbf{X})^{2k-1}((\mathbf{X}^{\top}\mathbf{X})^{\dagger})^{2k} = (\mathbf{X}^{\top}\mathbf{X})^{\dagger}$ for any $k \geq 1$.

As a consequence, we have

$$\lim_{t \to \infty} V(\widehat{\boldsymbol{\beta}}_t; \boldsymbol{\beta}) = \sigma^2 \frac{w^2 + (1 - w)^2}{w(2 - w)} \operatorname{Tr}[(\mathbf{X}^\top \mathbf{X})^\dagger \boldsymbol{\Sigma}]$$
(A.5)

The limit of the above quantity as $n \to \infty$ is given by (Hastie et al., 2022, Theorem 2) and is equal to the variance given by equation (3.5).

A.3 BIAS

We use the notation of Subsection A.1 throughout this subsection.

 Lemma A.3 (Ridge $\hat{\beta}_{t,\lambda}$ Expectation). For $\lambda > 0$, we have

$$\mathbb{E}[\widehat{\boldsymbol{\beta}}_{t,\lambda}|\mathbf{X}] = (\widetilde{w}A_{\lambda}M)^t A_{\lambda}M\boldsymbol{\beta} + w \sum_{i=0}^{t-1} (\widetilde{w}A_{\lambda}M)^i A_{\lambda}M\boldsymbol{\beta}$$
(A.6)

Proof. We prove equation (A.6) by induction. Let's start with the base case (t = 0).

$$\mathbb{E}[\widehat{\boldsymbol{\beta}}_{0,\lambda}|\mathbf{X}] = (\mathbf{X}^{\top}\mathbf{X} + n\lambda)^{-1}\mathbf{X}^{\top}\mathbb{E}[\mathbf{y}|\mathbf{X}]$$
$$= A_{\lambda}M\boldsymbol{\beta}$$

This proves the base case. Now assume equation (A.6) holds for t = k. Let's prove it for t = k + 1.

$$\begin{split} \mathbb{E}[\widehat{\boldsymbol{\beta}}_{k+1,\lambda}|\mathbf{X}] &= A_{\lambda}M(w\boldsymbol{\beta} + (1-w)\mathbb{E}[\widehat{\boldsymbol{\beta}}_{k,\lambda}|\mathbf{X}]) \\ &= wA_{\lambda}M\boldsymbol{\beta} + \widetilde{w}A_{\lambda}M\mathbb{E}[\widehat{\boldsymbol{\beta}}_{k,\lambda}|\mathbf{X}]) \\ &= wA_{\lambda}M\boldsymbol{\beta} + \widetilde{w}A_{\lambda}M\left[(\widetilde{w}A_{\lambda}M)^{k}A_{\lambda}M\boldsymbol{\beta} + w\sum_{i=0}^{k-1}(\widetilde{w}A_{\lambda}M)^{i}A_{\lambda}M\boldsymbol{\beta}\right] \\ &= (\widetilde{w}A_{\lambda}M)^{k+1}A_{\lambda}M\boldsymbol{\beta} + w\sum_{i=0}^{k}(\widetilde{w}A_{\lambda}M)^{i}A_{\lambda}M\boldsymbol{\beta} \end{split}$$

This completes the proof of Lemma A.3.

Corollary A.2. Bias for the Ridge Estimator is given by

$$B(\widehat{\boldsymbol{\beta}}_{t,\lambda};\boldsymbol{\beta}) = \left(\frac{n\lambda}{w}\right)^2 \|A_{\lambda/w}(I - (\widetilde{w}A_{\lambda}M)^{t+1})\boldsymbol{\beta}\|_{\boldsymbol{\Sigma}}^2$$
(A.7)

$$\lim_{t \to \infty} B(\widehat{\boldsymbol{\beta}}_{t,\lambda}; \boldsymbol{\beta}) = \left(\frac{n\lambda}{w}\right)^2 \|A_{\lambda/w}\boldsymbol{\beta}\|_{\boldsymbol{\Sigma}}^2 = \left(\frac{n\lambda}{w}\right)^2 \operatorname{Tr}[\boldsymbol{\beta}\boldsymbol{\beta}^\top A_{\lambda/w}\boldsymbol{\Sigma} A_{\lambda/w}] \tag{A.8}$$

Proof. From equation (A.6), we have

$$\boldsymbol{\beta} - \mathbb{E}[\widehat{\boldsymbol{\beta}}_{t,\lambda}|\mathbf{X}] = \boldsymbol{\beta} - (\widetilde{w}A_{\lambda}M)^{t}A_{\lambda}M\boldsymbol{\beta} - w(I - \widetilde{w}A_{\lambda}M)^{-1}(I - (\widetilde{w}A_{\lambda}M)^{t})A_{\lambda}M\boldsymbol{\beta}$$
(A.9)

Diagonalizing $A_{\lambda}M$ and and simplifying using the following identity

$$1 - \widetilde{w}^t x^{t+1} - w \frac{1 - \widetilde{w}^t x^t}{1 - \widetilde{w} x} x = \frac{(1 - x)(1 - \widetilde{w}^{1+t} x^{1+t})}{1 - \widetilde{w} x},$$

tells us that

$$\boldsymbol{\beta} - \mathbb{E}[\widehat{\boldsymbol{\beta}}_t | \mathbf{X}] = (I - A_{\lambda} M) (I - \widetilde{w} A_{\lambda} M)^{-1} (I - (\widetilde{w} A_{\lambda} M)^{1+t}) \boldsymbol{\beta}$$

Next, we diagonalize M use the following identity on product of first two matrices.

$$\left(1 - \frac{x}{x + n\lambda}\right) \left(1 - \frac{\widetilde{w}x}{x + n\lambda}\right)^{-1} = \frac{n\lambda}{wx + n\lambda} = \frac{n\lambda}{w} \left(x + n\frac{\lambda}{w}\right)^{-1}$$

to conclude $(I-A_{\lambda}M)(I-\widetilde{w}A_{\lambda}M)^{-1}=\frac{n\lambda}{w}A_{\lambda/w}$. This completes the proof of equation (A.7). equation (A.8) follows since $\sigma_{\max}(\widetilde{w}A_{\lambda}M)<\widetilde{w}<1$.

Lemma A.4. Suppose Assumption 2.1 holds. For any deterministic sequence of symmetric matrices $C \in \mathbb{R}^{p \times p}$ with bounded operator norm and $z \in \mathbb{C} \setminus \mathbb{R}^+$, we have

$$z^2 \operatorname{Tr}[CQ_n(z)\Sigma Q_n(z)] \sim \operatorname{Tr}[C(I+m(z)\Sigma)^{-2}\Sigma] \frac{1}{1-\frac{1}{2}df_2(1/m(z))}$$
 (A.10)

where $df_2(\kappa) = \text{Tr}[\mathbf{\Sigma}^2(\kappa I + \mathbf{\Sigma})^{-2}]$ and $Q_n(z) = (n^{-1}M - zI)^{-1}$.

 Proof. By plugging in A=C and $B=\Sigma$ in (Bach, 2024, Eq (3.9)) and noticing that $(I+m(z)\Sigma)$ and Σ are simultaneously diagonalizable (and therefore commute), we get

$$z^{2}\operatorname{Tr}[CQ_{n}(z)\Sigma Q_{n}(z)] \sim \operatorname{Tr}[C(I+m(z)\Sigma)^{-2}\Sigma]$$

$$+\operatorname{Tr}[C(I+m(z)\Sigma)^{-2}\Sigma] \cdot \frac{\operatorname{Tr}[\Sigma^{2}\left(m(z)^{-1}\Sigma+I\right)^{-2}]}{n-df_{2}(m(z)^{-1})}$$

$$\sim \operatorname{Tr}[C(I+m(z)\Sigma)^{-2}\Sigma] \cdot \left(1+\frac{df_{2}(m(z)^{-1})}{n-df_{2}(m(z)^{-1})}\right)$$

$$\sim \operatorname{Tr}[C(I+m(z)\Sigma)^{-2}\Sigma] \cdot \frac{1}{1-n^{-1}df_{2}(m(z)^{-1})}$$

This completes the proof of equation (A.10).

We now show that Equation (3.9) holds. Start by noticing that $nA_z = Q_n(-z)$. Then, equation (A.7) combined with Lemma A.4 tells us that

$$\lim_{t \to \infty} B(\widehat{\boldsymbol{\beta}}_{t,\lambda}; \boldsymbol{\beta}) = \left(\frac{n\lambda}{w}\right)^2 \text{Tr}[\boldsymbol{\beta}\boldsymbol{\beta}^{\top} A_{\lambda/w} \boldsymbol{\Sigma} A_{\lambda/w}]$$
$$\sim \text{Tr}[\boldsymbol{\beta}\boldsymbol{\beta}^{\top} (I + m_1 \boldsymbol{\Sigma})^{-2} \boldsymbol{\Sigma}] \cdot \frac{1}{1 - n^{-1} df_2(m_1^{-1})}$$

where $m_1 = m(-\lambda/w)$. All that is left to get equation (3.9) is realizing that

$$\lim_{n \to \infty} \text{Tr}[\boldsymbol{\beta} \boldsymbol{\beta}^{\top} (I + m_1 \boldsymbol{\Sigma})^{-2} \boldsymbol{\Sigma}] = b_{\star} \int \frac{x}{(1 + m_1 x)^2} dG$$
and
$$\lim_{n \to \infty} n^{-1} df_2(m_1^{-1}) = \gamma \int \frac{m_1^2 x^2}{(1 + m_1 x)^2} dH$$
(A.11)

This completes the proof of limiting value of bias.

PROOF OF THEOREM 3.4

The asymptotic variance of $\hat{\beta}_{t,\lambda}$ is given by equation (A.3) for general Σ . The asymptotic bias is given by equation (A.11). This completes the proof of Theorem 3.4.

PROOF OF THEOREM 3.2

Plugging $H = G = \delta_{\alpha}$ in equation (3.10) and equation (3.9) yields the asymptotic variance and bias for $\widehat{\beta}_{t,\lambda}$ when $\Sigma = I$. The proof of log-convexity of risk follows from the proof of Theorem 3.3

PROOF OF THEOREM 3.1

The asymptotic variance of $\hat{\boldsymbol{\beta}}_t$ is obtained by equation (A.5). To obtain the bias, recall that $\hat{\boldsymbol{\beta}}_t = (\mathbf{X}^\top \mathbf{X})^\dagger \mathbf{X}^\top (w \mathbf{y}_t + (1-w) \tilde{\mathbf{y}}_t)$. Since $\hat{\boldsymbol{\beta}}_0 = (\mathbf{X}^\top \mathbf{X})^\dagger \mathbf{X}^\top \mathbf{y}$, we have $\mathbb{E}(\hat{\boldsymbol{\beta}}_0 | \mathbf{X}) = (\mathbf{X}^\top \mathbf{X})^\dagger \mathbf{X}^\top \mathbf{X} \boldsymbol{\beta}$. Now, we want to prove by induction that

$$\mathbb{E}(\widehat{\boldsymbol{\beta}}_t|\mathbf{X}) = (\mathbf{X}^{\top}\mathbf{X})^{\dagger}\mathbf{X}^{\top}\mathbf{X}\boldsymbol{\beta}$$
(A.12)

for all t. To this end note that,

$$\mathbb{E}(\widehat{\boldsymbol{\beta}}_t|\mathbf{X}) = (\mathbf{X}^{\top}\mathbf{X})^{\dagger}\mathbf{X}^{\top} \Big(w\mathbb{E}(\mathbf{y}_t|\mathbf{X}) + (1-w)\mathbb{E}(\widetilde{\mathbf{y}}_t|\mathbf{X}) \Big)$$
(A.13)

Since $\mathbf{y}_t = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}_t$, we have $\mathbb{E}(\mathbf{y}_t|\mathbf{X}) = \mathbf{X}\boldsymbol{\beta}$. For the synthetic data, $\widetilde{\mathbf{y}}_t = \mathbf{X}\widehat{\boldsymbol{\beta}}_{t-1} + \widehat{\boldsymbol{\varepsilon}}_t$. Therefore,

$$\mathbb{E}(\widetilde{\mathbf{y}}_t|\mathbf{X}) = \mathbf{X}\mathbb{E}(\widehat{\boldsymbol{\beta}}_{t-1}|\mathbf{X}) = \mathbf{X}(\mathbf{X}^{\top}\mathbf{X})^{\dagger}\mathbf{X}^{\top}\mathbf{X}\boldsymbol{\beta}$$

by induction hypothesis. Using equation (A.13) we have

$$\mathbb{E}(\widehat{\boldsymbol{\beta}}_t|\mathbf{X}) = (\mathbf{X}^{\top}\mathbf{X})^{\dagger}\mathbf{X}^{\top} \Big(w\mathbf{X}\boldsymbol{\beta} + (1-w)\mathbf{X}(\mathbf{X}^{\top}\mathbf{X})^{\dagger}\mathbf{X}^{\top}\mathbf{X}\boldsymbol{\beta} \Big) = (\mathbf{X}^{\top}\mathbf{X})^{\dagger}\mathbf{X}^{\top}\mathbf{X}\boldsymbol{\beta}.$$

This proves equation (A.12) by induction. This in turn implies using equation (2.7) that for any n, t,

$$B(\widehat{\boldsymbol{\beta}}_t; \boldsymbol{\beta}) = \boldsymbol{\beta}^{\top} P_{\mathbf{X}} \boldsymbol{\Sigma} P_{\mathbf{X}} \boldsymbol{\beta}$$

where $P_{\mathbf{X}} = I - (\mathbf{X}^{\top} \mathbf{X})^{\dagger} (\mathbf{X}^{\top} \mathbf{X})$. Using (Hastie et al., 2022, Theorem 2), we obtain the analytic value of the bias.

Finally, note that the asymptotic risk depends on w only via c(w). Since c(w) is minimized at $1/\varphi$, this completes the proof of Theorem 3.1.

B PROOFS OF REMAINING RESULTS

In this Section, we will show that the expression of generalization error $R(\widehat{\boldsymbol{\beta}}_{t,\lambda})$ can be simplified further if the probability measures \widehat{G}_p and \widehat{H}_p defined by equation (3.1) weakly converges to the same probability distribution, i.e., G=H. In this special case, the generalization error has unique minima w.r.to w.

Recall the definitions of V_{λ} and \mathcal{B}_{λ} given by equation (3.10) and equation (3.9) respectively. We will first rewriting V_{λ} and \mathcal{B}_{λ} in a different form. To this end, differentiate both sides of equation (3.2) w.r.to z to obtain

$$-\frac{m'(z)}{m^2(z)} + 1 = \gamma \int \frac{-x^2 m'(z)}{(1 + xm(z))^2} dH$$
 (B.1)

The above equality will be helpful in writing the integrals concisely. Also define $f(z) = m(-z)^{-1} - z$. We know that m(z) is a Stieltjes transform of a non-negative random variable by Remark 3.1. We will need the following technical Lemma whose proof we defer.

Lemma B.1. There exists some measure μ on \mathbb{R}^+ with $|f(1)| < \infty$ such that

$$f(z) = a + \int \frac{z}{z+t} \mu(dt). \tag{B.2}$$

Invoking Lemma B.1, equation (3.10) tells us that

$$\frac{2(1-w)}{w(2-w)} \mathcal{V}_{\lambda} = \frac{\gamma}{\lambda} \left(\int \frac{x}{1+m_1 x} dH - \int \frac{x}{1+m_2 x} dH \right)$$

$$= \frac{1}{\lambda} \left(\frac{1}{m_1} - \frac{\lambda}{w} - \frac{1}{m_2} + \frac{\lambda}{2-w} \right) \qquad \text{(By equation (3.2))} \qquad \text{(B.3)}$$

$$= \frac{1}{\lambda} \left(f(\lambda/w) - f(\lambda/2 - w) \right)$$

$$= \frac{1}{\lambda} \left(\int \frac{\lambda}{\lambda + wt} \mu(dt) - \int \frac{\lambda}{\lambda + (2-w)t} \mu(dt) \right)$$

$$= \int \frac{2(1-w)t}{(\lambda + wt)(\lambda + (2-w)t)} \mu(dt).$$

Thus, we obtain that for some measure μ on \mathbb{R}^+ , we have

$$\implies \mathcal{V}_{\lambda} = \int t \frac{w(2-w)}{(\lambda + wt)(\lambda + (2-w)t)} \mu(dt) \tag{B.5}$$

Similarly, we can also simplify the Bias. If H = G, we have from equation (3.9),

$$\mathcal{B}_{\lambda} = \left(\int \frac{x}{(1+m_1 x)^2} dH \right) \left(1 - \gamma \int \frac{m_1^2 x^2}{(1+m_1 x)^2} dH \right)^{-1}$$

We first multiply both sides of equation (B.1) by $-m_1^2/m'(-\lambda/w)$ to obtain

$$\gamma \int \frac{m_1^2 x^2}{(1+m_1 x)^2} dH = 1 - \frac{m_1^2}{m'(-\lambda/w)}$$

Next, we decompose the first term of the bias as

$$\frac{x}{(1+m_1x)^2} = \frac{x}{1+m_1x} - \frac{m_1x^2}{(1+m_1x)^2}$$

$$\implies \gamma \int \frac{x}{(1+m_1x)^2} dH = \gamma \int \frac{x}{1+m_1x} dH - \gamma \int \frac{m_1x^2}{(1+m_1x)^2} dH$$

$$= \frac{1}{m_1} - \frac{\lambda}{w} - \left(\frac{1}{m_1} - \frac{m_1}{m'(\lambda/w)}\right)$$

$$= \frac{m_1}{m'(\lambda/w)} - \frac{\lambda}{w}$$

Combing the two equalities above, we get

$$\mathcal{B}_{\lambda} = \gamma^{-1} \left(\frac{m_1}{m'(-\lambda/w)} - \frac{\lambda}{w} \right) \frac{m'(-\lambda/w)}{m_1^2}$$

$$= \gamma^{-1} \left(\frac{1}{m(\frac{-\lambda}{w})} - \frac{\lambda}{w} \frac{m'(\frac{-\lambda}{w})}{m(\frac{-\lambda}{w})^2} \right)$$

$$= \gamma^{-1} \left(f(\lambda/w) - \frac{\lambda}{w} f'(\lambda/w) \right)$$
(B.6)

Let $z = \lambda/w$, we use equation (B.2) to get

$$\gamma \mathcal{B}_{\lambda} = a + \int \frac{z}{z+t} - \frac{zt}{(z+t)^2} \mu(dt)$$

$$= a + \int \frac{z^2}{(z+t)^2} \mu(dt)$$

$$= a + \int \frac{\lambda^2}{(\lambda + wt)^2} \mu(dt)$$
(B.7)

Since the sum of log-convex functions are log-convex, this implies that \mathcal{B}_{λ} is decreasing and log-convex. Next we show that $c(w)\mathcal{V}_{\lambda}$ is log-convex. By equation (B.5), we have

$$c(w)\mathcal{V}_{\lambda} = \int t\gamma \frac{w^2 + (1-w)^2}{(\lambda + wt)(\lambda + (2-w)t)} \mu(dt)$$

The log-convexity of the above expression simply follows from the fact that $w^2 + (1-w)^2$, $(\lambda + wt)^{-1}$ and $(\lambda + (2-w)t)^{-1}$ are all log-convex in $w \in [0,1]$ for all $t, \lambda \geq 0$ and the fact that sums and products of log-convex functions are log convex.

As long as $\mu \neq \delta_0$, we further have that $c(w)\mathcal{V}_{\lambda}$ is strictly log-convex, a fact crucial for uniqueness of the minimizer of the risk. It can be readily verified from the definition of f(z) that $\mu = \delta_0 \implies m(z) = (a-z)^{-1}$, that is m must be a Stieltjes transform of a degenerate random variable. However, recall that by Remark 3.1, m is the Steiltjes transform of a free convolution between MP_{γ} and H. Since MP_{γ} is non-degenerate, we must have that $m(z) \neq (a-z)^{-1}$ and hence the variance is strictly convex.

As long as $\mu(\mathbb{R}^+)>0$, we further have that $c(w)\mathcal{V}_\lambda$ is strictly log-convex, a fact crucial for uniqueness of the minimizer of the risk. It can be readily verified from the definition of f(z) that $\mu\equiv 0 \implies m(z)=(a-z)^{-1}$, that is m must be a Stieltjes transform of a degenerate random variable. However, recall that by Remark 3.1, m is the Steiltjes transform of a free convolution between MP_γ and H. Since MP_γ is non-degenerate, we must have that $m(z)\neq (a-z)^{-1}$ and hence the variance is strictly convex.

Next, we show that w^* is a continuous function of λ . Define the quantity

$$\mathcal{R}(w,\lambda) = \lim_{n \to \infty} \lim_{t \to \infty} R(\widehat{\boldsymbol{\beta}}_{t,\lambda}, \widehat{\boldsymbol{\beta}}).$$

Note that, $\mathcal{R}(w,\lambda)$ is continuous in both w and λ . Define any convergent sequence $\lambda_n \to \widetilde{\lambda}$. Let w_k^{\star} to be the unique minimizer of $\mathcal{R}(w,\lambda_k)$ (unique minimizer because \mathcal{R} is strictly log-convex is

all three results of this section). Since w_n^\star is a sequence in the compact set [0,1], there exists a convergent subsequence $w_{n_k}^\star$ converging to some limit \widetilde{w}^\star . For any $\forall w \in [0,1]$,

$$\mathcal{R}(w_{n_k}, \lambda_{n_k}) \leq \mathcal{R}(w, \lambda_{n_k}) \implies \mathcal{R}(\widetilde{w}, \widetilde{\lambda}) \leq \mathcal{R}(w, \widetilde{\lambda}).$$

That is, \widetilde{w}^{\star} is the minimizer of $\mathcal{R}(\cdot,\widetilde{\lambda})$. Since the convergent subsequence n_k we picked was arbitrary, we have shown that every convergent subsequence of w_k converges to \widetilde{w}^{\star} and thus we must have that $w_n^{\star} \to \widetilde{w}^{\star}$. This proves that w^{\star} is a continuous function of λ .

Next, we propose to show that under G=H, the optimal mixing proportion w^\star is in [1/2,1], with $w^\star(\lambda) \to \phi^{-1}$ as $\lambda \downarrow 0$ and $w^\star(\lambda) \to 1$ as $\lambda \to \infty$. Recall that both $c(w)\mathcal{V}_\lambda$ and \mathcal{B}_λ are continuous and log-convex, and \mathcal{B}_λ is also decreasing. This immediately tells us that the limiting generalization error is log-convex and hence it has a unique minimizer.

$$\lim_{\lambda \to 0} c(w) \mathcal{V}_{\lambda} = \frac{w^2 + (1-w)^2}{w(2-w)} \gamma \int \frac{1}{t} \mu(dt) \quad \text{ and } \quad \lim_{\lambda \to 0} \mathcal{B}_{\lambda} = \frac{a}{\gamma} \tag{B.8}$$

The minimizer of the of the risk at λ is clearly only dependent on c(w), which is minimized at ϕ^{-1} . Next, as $\lambda \to \infty$, we have $c(w)\mathcal{V}_{\lambda} \to 0$. However, since \mathcal{B}_{λ} is a decreasing function of w and $\liminf_{\lambda \to \infty} \mathcal{B}_{\lambda} > 0$, we must have that $w^{\star} \to 1$. Finally, we need to show that $w^{\star} \geq 1/2$. To this end, write the variance as

$$c(w)\mathcal{V}_{\lambda} = (w^2 + (1-w)^2) \int \frac{t\gamma}{(\lambda + wt)(\lambda + (2-w)t)} \mu(dt)$$

The expression outside of the integral is minimized at w=1/2, while the factor inside the integral is a decreasing function of w for $w \in [0,1]$. This can be seen by calculating the derivative of the integrand

$$\frac{d}{dw}\frac{1}{(\lambda+wt)(\lambda+(2-w)t)} = \frac{-2t^2(1-w)}{(\lambda+wt)^2(\lambda+(2-w)t)^2}$$
 (B.9)

Since \mathcal{B}_{λ} is also a decreasing function of w, we must have that risk at w < 1/2 must be larger than the risk at 1/2. This completes the proof that $w^* \ge 1/2$.

PROOF OF PROPOSITION 3.1

We obtain the asymptotic bias and variance for the random effects model from equation (B.6) and equation (B.4) respectively. The log-convexity of variance follows from equation (B.5). \mathcal{B}_{λ} is log-convex and decreasing using equation (B.7). This completes the proof of theorem 3.1.

PROOF OF PROPOSITION 3.2

By Corollary D.1, we obtain that G=H. Further, under the assumption $\beta_i^{\text{i.i.d.}}(0,b_\star/p)$, we obtain $\|\beta\|^2 \to b_\star$ almost surely. Then, by the argument above, we have $w^\star \in [1/2,1]$, with $w^\star(\lambda) \to \phi^{-1}$ as $\lambda \downarrow 0$ and $w^\star(\lambda) \to 1$ as $\lambda \to \infty$. Finally, to see that w^\star increases with SNR $=b_\star/\sigma^2$, recall that the risk is $\mathcal{R}(w,b_\star)=\sigma^2c(w)\mathcal{V}_\lambda+b_\star\mathcal{B}_\lambda$. Note that dividing the risk by σ^2 does not change does not change w^\star . Thus, it is enough to show that w^\star increases with b_\star . By the implicit function theorem,

$$\partial_{b_{\star}} w^{\star} = -\frac{\partial_{wb_{\star}} \mathcal{R}}{\partial_{ww} \mathcal{R}} = -\frac{\partial_{w} \mathcal{B}_{\lambda}}{\partial_{ww} \mathcal{R}}$$

Since Bias is a decreasing function of w and risk is a strictly convex function of w, must have that $\partial_{b_{\star}} w^{\star} \geq 0$, thus proving w^{\star} is a non decreasing function of b_{\star} . 3.2.

PROOF OF THEOREM 3.3

If $\Sigma = \alpha I$, we have G = H, and we obtain the desired conclusion by the argument above.

PROOF OF PROPOSITION 3.3

For spike covariance matrix Σ , we obtain $\widehat{H}_p = p^{-1}\delta_{1+s} + (1-p^{-1})\delta_1$, where δ_x is the Dirac measure at the point x. Therefore, $\widehat{H}_p \Rightarrow H = \delta_1$. To compute \widehat{G}_p , we write the signal $\boldsymbol{\beta}$ as $\boldsymbol{\beta} = \theta v + \sqrt{1-\theta^2}v^{\perp}$, where $v^{\top}v^{\perp} = 0$ and $\|v^{\perp}\|_2 = 1$. This implies $\|\boldsymbol{\beta}\|_2 = 1$ and hence $\widehat{G}_p = \theta^2 \delta_{1+s} + (1-\theta^2)\delta_1$. If $\theta = \theta(n) \to 0$ as $n \to \infty$, we obtain $\widehat{G}_p \Rightarrow H = \delta_1$ and the conclusion of Theorem 3.2 holds. Hence we will restrict ourselves to the case $\theta(n) \to \theta_{\star} \neq 0$. Here we have $G = \theta_{\star}^2 \delta_{1+s} + (1-\theta_{\star}^2)\delta_1$. Since $H = \delta_1$, the limiting expression of variance is still $\sigma^2 c(w) \mathcal{V}_{\lambda}$, where c(w) and \mathcal{V}_{λ} are the same as Theorem 3.2.

Turning to the characterization of asymptotic bias, we again use the function m(z) satisfying $m(z)^{-1} + z = \gamma \frac{1}{1+m(z)}$. Defining $m_1 = m(-\lambda/w)$, we have

$$\mathcal{B}_{\lambda} = \left(\theta_{\star}^{2} \frac{1+s}{(1+m_{1}(1+s))^{2}} + (1-\theta_{\star}^{2}) \frac{1}{(1+m_{1})^{2}}\right) \left(1 - \gamma \frac{m_{1}^{2}}{(1+m_{1})^{2}}\right)^{-1}$$

We know from Theorem 3.2 that $\phi_1(m_1) = \frac{1}{(1+m_1)^2 - \gamma m_1^2}$ is a decreasing and log-convex function in w. From the above display, we obtain that

$$\mathcal{B}_{\lambda} = \theta_{\star}^{2} (1+s) \underbrace{\frac{(1+m_{1})^{2}}{(1+m_{1}(1+s))^{2}}}_{\phi_{2}(m_{1})} \phi_{1}(m_{1}) + (1-\theta_{\star}^{2})\phi_{1}(m_{1}).$$

Since $\phi_2(m_1)$ is also decreasing and log-convex as function of w, this implies that $\lim_{t\to\infty}\lim_{n\to\infty}B(\widehat{\boldsymbol{\beta}}_{t,\lambda};\boldsymbol{\beta})$ is decreasing and log-convex in w. Therefore, asymptotic generalization error has a unique minima w^* . To see the properties of w^* , note that by equation (B.9), we have $w^*\geq 1/2$. Using equation (B.8), we obtain that $w^*\to 1/\varphi$ as $\lambda\to 0+$. Finally similar to the proof of Theorem 3.2, we have $\mathcal{V}_\lambda\to 0$ as $\lambda\to\infty$ and $\liminf_{\lambda\to\infty}\mathcal{B}_\lambda>0$. Since \mathcal{B}_λ is decreasing, we again have $\lim_{\lambda\to\infty}w^*(\lambda)=1$. This completes the proof of Proposition 3.3.

We conclude the section with with the proof of Lemma B.1.

PROOF OF LEMMA B.1

We need the following definition.

Definition B.1 (Stieltjes Function (SF)). A function $f: \mathbb{R}^+ \to \mathbb{R}^+$ is called a Stieltjes function if it can be written as

$$f(x) = \frac{a}{x} + b + \int_{\mathbb{R}^+} \frac{1}{x+t} \mu(dt),$$

where $a, b \geq 0$ and μ is a positive measure on \mathbb{R}^+ with $\int_{\mathbb{R}^+} \frac{1}{1+t} \, \mu(dt) < \infty$.

Recall, the definition $m(z)=\mathbb{E}(A-z)^{-1}$ for some non-negative random variable A. Using the definition above, the map $z\to m(-z)$ is a Stieltjes function. Further $\psi(z)=1/z$ is a Stieltjes function by Definition B.1. Using (Schilling et al., 2009, Theorem 6.2(ii) and Corollary 7.9), we obtain that 1/(zm(-z)) is a Stieltjes function. Note that 1/(zm(-z))=(f(z)/z)-1. Hence, by Definition B.1, we have

$$f(z) + z = a + bz + \int \frac{z}{z+t} \mu(dt)$$

with $|f(1)| < \infty$. It remains to show that b = 1. This follows from

$$b = \lim_{z \to \infty} \frac{f(z)}{z} = \lim_{z \to \infty} \mathbb{E}[z(A+z)^{-1}]^{-1} = 1$$

where the last equality follows from dominated convergence theorem.

C DYNAMIC MIXING

In this Section we prove the following claim for min- ℓ_2 -norm interpolator: Suppose we select the mixing proportion w_t adaptively at each generation to minimize $R(\hat{\beta}_t; \beta)$ for any finite sample size n. Then w_t^* satisfies

$$w_t^{\star} = \frac{1 + w_{t-1}^{\star}}{2 + w_{t-1}^{\star}}, \quad w_0^{\star} = 1.$$
 (C.1)

Further, under this setup

$$\lim_{n \to \infty} \lim_{t \to \infty} R(\widehat{\boldsymbol{\beta}}_t; \boldsymbol{\beta}) = \sigma^2 c(w^*) \mathcal{V} + b_* \mathcal{B}, \tag{C.2}$$

where V, \mathcal{B} as in equation (3.5) and $w^* = 1/\varphi$. To see the claim, note that the bias of $\widehat{\beta}_t$ is independent of mixing proportion, hence it converges to \mathcal{B} . Regarding the variance, we will prove the following by induction

$$w_t^{\star} = \operatorname*{arg\,min}_{t} V(\boldsymbol{\beta}_t; \boldsymbol{\beta}), \quad \mathsf{Cov}[\boldsymbol{\beta}_t | \mathbf{X}] = \sigma^2 w_t^{\star} (\mathbf{X}^{\top} \mathbf{X})^{\dagger} \boldsymbol{\Sigma}.$$
 (C.3)

Suppose t = 1. For any $0 < w_1 < 1$, we have

$$\begin{split} \mathsf{Cov}[\widehat{\boldsymbol{\beta}}_t|\mathbf{X}] &= (\mathbf{X}^{\top}\mathbf{X})^{\dagger}\mathbf{X}^{\top}(w_1^2\sigma^2I + 2(1-w_1)^2\sigma^2\mathbf{X}(\mathbf{X}^{\top}\mathbf{X})^{\dagger}\mathbf{X}^{\top})\mathbf{X}(\mathbf{X}^{\top}\mathbf{X})^{\dagger} \\ &= \sigma^2(w_1^2 + 2(1-w_1)^2)(\mathbf{X}^{\top}\mathbf{X})^{\dagger}. \end{split}$$

Hence the variance is minimized at $w_1^* = 2/3 = (1+w_0^*)/(2+w_0^*)$. Further, $w_1^{*2} + (1-w_1^*)^2 = w_1^*$ proving equation (C.3) for t=1. Now, for any t>1 and any mixing proportion w_t , we have by induction hypothesis

$$\begin{split} & \mathsf{Cov}[\widehat{\boldsymbol{\beta}}_t|\mathbf{X}] = (\mathbf{X}^{\top}\mathbf{X})^{\dagger}\mathbf{X}^{\top}(w_t^2\sigma^2I + (1-w_t)^2\mathsf{Cov}(\widetilde{y}_{t-1}|\mathbf{X}))\mathbf{X}(\mathbf{X}^{\top}\mathbf{X})^{\dagger} \\ &= w_t^2\sigma^2(\mathbf{X}^{\top}\mathbf{X})^{\dagger} + (1-w_t)^2(\mathbf{X}^{\top}\mathbf{X})^{\dagger}\mathbf{X}^{\top} \left[\mathbf{X}\mathsf{Cov}(\widehat{\boldsymbol{\beta}}_{t-1}|\mathbf{X})\mathbf{X}^{\top} + \sigma^2I\right]\mathbf{X}(\mathbf{X}^{\top}\mathbf{X})^{\dagger} \\ &= (w_t^2 + (1-w_t)^2(1+w_{t-1}^{\star}))\sigma^2(\mathbf{X}^{\top}\mathbf{X})^{\dagger}, \end{split}$$

which is minimized at $w_t^\star = \frac{1+w_{t-1}^\star}{2+w_{t-1}^\star}$. Further we have $w_t^{\star 2} + (1-w_t^\star)^2(1+w_{t-1}^\star) = w_t^\star$, proving equation (C.3) by induction principle. Since by equation (C.1), we have $w_t^\star \to w^\star$, we immediately have equation (C.2).

D RANDOM EFFECTS

Lemma D.1 (Equal weak limits). Let $\Sigma = \sum_{k=1}^p s_k v_k v_k^{\top}$ with eigenvalues s_1, \ldots, s_p and orthonormal eigenvectors v_1, \ldots, v_p Define

$$\widehat{H}_p(x) = \frac{1}{p} \sum_{k=1}^p \mathbf{1}_{s_k \le x}, \qquad \widehat{G}_p(x) = \frac{1}{\|\boldsymbol{\beta}\|_2^2} \sum_{k=1}^p \langle v_k, \boldsymbol{\beta} \rangle^2 \mathbf{1}_{s_k \le x}.$$

Assume $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)$ has i.i.d. entries with $\mathbb{E}\beta_i = 0$, $\mathbb{E}\beta_i^2 = \tau^2 \in (0, \infty)$ and $\mathbb{E}\beta_i^4 < \infty$, independent of $\boldsymbol{\Sigma}$.

Then, conditional on Σ *, for every bounded Lipschitz test function* ψ *,*

$$\int \psi \, d\widehat{G}_p - \int \psi \, d\widehat{H}_p \stackrel{\mathbb{P}}{\to} 0$$

as $p \to \infty$. Consequently, if $\widehat{H}_p \Rightarrow H$ weakly, then also $\widehat{G}_p \Rightarrow H$ weakly.

Proof. Fix any bounded Lipschitz continuous function $\psi : \mathbb{R} \to \mathbb{R}$ such that $\|\psi\|_{\infty} \leq M$. Set $a_k := \psi(s_k)$ (which is deterministic given Σ). Let $\xi_k := \langle v_k, \beta \rangle$. Since β_i 's are i.i.d. and independent of

 Σ , orthonormality gives that, conditional on Σ , the variables ξ_k identically distributed and mutually uncorrelated with $\mathbb{E}\xi_k = 0$, $\mathbb{E}\xi_k^2 = \tau^2$, $\mathbb{E}\xi_k^4 < \infty$. Then

$$\int \psi \, d\widehat{G}_p = \frac{\sum_{k=1}^p a_k \, \xi_k^2}{\sum_{j=1}^p \xi_j^2}, \qquad \int \psi \, d\widehat{H}_p = \frac{1}{p} \sum_{k=1}^p a_k =: A_p.$$

Define centered averages

$$B_p := \frac{1}{p} \sum_{k=1}^p a_k (\xi_k^2 - \tau^2), \qquad C_p := \frac{1}{p} \sum_{j=1}^p (\xi_j^2 - \tau^2).$$

Then

$$\frac{\sum_k a_k \xi_k^2}{\sum_j \xi_j^2} = \frac{p(\tau^2 A_p + B_p)}{p(\tau^2 + C_p)} = A_p + \frac{B_p - A_p C_p}{\tau^2 + C_p}.$$

Since ψ is bounded, $|a_k| \leq M$. Using $\mathbb{E}(\xi_1^2) = \tau^2$ and $\mathbb{E}(\xi_1^4) < \infty$,

$$\operatorname{Var}(B_p) = \frac{1}{p^2} \sum_{k=1}^p a_k^2 \operatorname{Var}(\xi_k^2) \le \frac{M^2}{p} \operatorname{Var}(\xi_1^2) = O\left(\frac{1}{p}\right), \quad \operatorname{Var}(C_p) = \frac{1}{p} \operatorname{Var}(\xi_1^2) = O\left(\frac{1}{p}\right).$$

Hence $B_p = O_{\mathbb{P}}(p^{-1/2})$ and $C_p = O_{\mathbb{P}}(p^{-1/2})$ conditional on Σ . Also $|A_p| \leq M$ and $\tau^2 + C_p \xrightarrow{\mathbb{P}} \tau^2 > 0$. Therefore

$$\left| \int \psi \, d\widehat{G}_p - \int \psi \, d\widehat{H}_p \right| = \left| \frac{B_p - A_p C_p}{\tau^2 + C_p} \right| = O_p \left(\frac{1}{\sqrt{p}} \right) \xrightarrow{\mathbb{P}} 0,$$

again conditional on Σ . By characterizations of weak convergence, we have that $\widehat{G}_p - \widehat{H}_p \Rightarrow 0$ in probability, and thus any weak limit of \widehat{H}_p is also a weak limit of \widehat{G}_p .

Corollary D.1. Let $\widehat{H}_p, H, \widehat{G}_p, G$ and Σ be as defined in Lemma D.1. Assume $\beta = \eta_p \omega$ where $\eta_p \neq 0$ is some arbitrary sequence of real numbers and $\omega = (\omega_1, \ldots, \omega_p)$ has i.i.d. entries with $\mathbb{E}\omega_i = 0$, $\mathbb{E}\omega_i^2 = \tau^2 \in (0, \infty)$ and $\mathbb{E}\omega_i^4 < \infty$, independent of Σ .

Then, conditional on Σ *, for every bounded Lipschitz test function* ψ *,*

$$\int \psi \, d\widehat{G}_p - \int \psi \, d\widehat{H}_p \stackrel{\mathbb{P}}{\to} 0$$

as $p \to \infty$. Consequently, if $\hat{H}_p \Rightarrow H$ weakly, then also $\hat{G}_p \Rightarrow H$ weakly.

Proof. Follows from the fact that

$$\widehat{G}_p(x) = \frac{1}{\|\boldsymbol{\beta}\|_2^2} \sum_{k=1}^p \langle v_k, \boldsymbol{\beta} \rangle^2 \, \mathbf{1}_{s_k \le x} = \frac{1}{\|\boldsymbol{\omega}\|_2^2} \sum_{k=1}^p \langle v_k, \boldsymbol{\omega} \rangle^2 \, \mathbf{1}_{s_k \le x} =: \widetilde{G}_p(x),$$

and applying Lemma D.1 on \widetilde{G}_p .