SCALING LAWS OF SIGNSGD IN LINEAR REGRESSION: WHEN DOES IT OUTPERFORM SGD?

Anonymous authors

Paper under double-blind review

ABSTRACT

We study scaling laws of signSGD under a power-law random features (PLRF) model that accounts for both feature and target decay. We analyze the expected population risk of a linear model trained with one-pass signSGD on Gaussian-sketched features. We express the risk as a function of model size, training steps, learning rate, and the feature and target decay parameters. Comparing against the SGD risk analyzed by Paquette et al. (2024), we identify a *drift-normalization effect* and a *noise-reshaping effect* unique to signSGD. We then obtain compute-optimal scaling laws under the optimal choice of learning rate. Our analysis shows that the noise-reshaping effect can make the compute-optimal slope of signSGD steeper than that of SGD in regimes where noise is dominant. Finally, we observe that a *stable-decay schedule*—a simplified variant of the widely used warmup-stable-decay (WSD) schedule—further reduces the noise term and sharpens the compute-optimal slope, when feature decay is fast but target decay is slow.

1 Introduction

In large-scale language model training, neural scaling laws are a well-documented empirical regularity: performance tends to improve predictably as data, parameters, and compute increase. Kaplan et al. (2020) observed that the language model cross-entropy loss scales as a power-law of model size M and number of steps N in terms of the risk formula $R(M,N) \approx M^{-\tau_1} + N^{-\tau_2}$. Also, they observe that loss scales as the power of training compute, under optimal allocation of compute between model size and number of steps.

A growing body of theory has sought to explain this phenomenon, most prominently by analyzing the stochastic gradient descent (SGD) optimizer under the power-law random features (PLRF) model (Paquette et al., 2024; Lin et al., 2024; 2025). Yet, in practice, SGD is not the optimizer that powers today's state-of-the-art LLMs. Instead, training is dominated by Adam (Kingma & Ba, 2014) and its variants. While Adam is considerably more difficult to analyze theoretically, it is often approximated in theory by the simpler signSGD (Bernstein et al., 2018a), which captures its coordinate-wise adaptivity. This gap between practice and theory motivates a natural question: how do scaling laws change when we replace SGD with signSGD? Addressing this question can help align theory with optimizer choices used in practice, and clarify how adaptive updates could reshape compute-optimal scaling regimes in the PLRF setting.

1.1 OUR CONTRIBUTION

We study the scaling law of signSGD in the power-law random features (PLRF) model, and our contributions are as follows.

1. We derive a scaling law of signSGD with constant learning rates involving three variables (model size M, training steps N, learning rate γ_0) and two PLRF model parameters (feature decay α , target decay β); see (12). By comparing with the SGD scaling laws of Paquette et al. (2024) and Lin et al. (2024), we observe two effects of signSGD: a *drift-normalization effect* and a *noise-reshaping effect*, inside the scaling law (see Section 4.1).

¹Here \approx denotes equality up to a multiplicative constant, i.e., $f(x) \approx g(x)$ means $c_1 g(x) \leq f(x) \leq c_2 g(x)$ for some constants $c_1, c_2 > 0$.

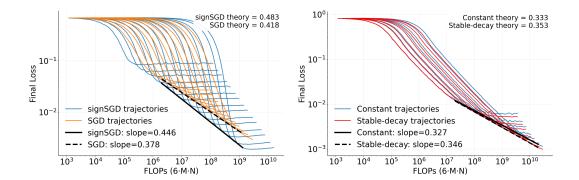


Figure 1: **Left: SGD vs. signSGD; Right: signSGD with constant vs. stable-decay schedules.** Colored lines represent the training trajectories of each algorithm, and black lines denote the compute-optimal curves. The upper right legend shows the theoretical value of the compute-optimal slope. SignSGD achieves a steeper compute-optimal slope than SGD (left panel), and stable-decay scheduling sharpens the compute-optimal slope relative to a constant schedule (right panel), for some parameter configurations. See Appendix C for parameters used in the experiment.

- 2. Under the fixed compute budget, we balance model size M and training steps N, and optimize over learning rate γ_0 . This allows us to characterize the compute-optimal loss decay rate and optimal model size with respect to the compute budget (see Table 1). Comparing against the compute-optimal scaling laws of SGD from Paquette et al. (2024) across regimes of the (α, β) -parameter plane, we find that signSGD can achieve better exponents in the SGD noise bottleneck regimes, due to the noise-reshaping effect (see Figure 1).
- 3. We show that learning rate scheduling can further reduce the stochastic noise of signSGD. We analyze a *stable-decay schedule*, a simplified version of the warmup-stable-decay (WSD) schedule (Wen et al., 2024) widely used in large language model training. By maintaining drift velocity by the stable interval and reducing stochastic noise by the polynomially decaying interval, this schedule increases the compute-optimal slope in the PLRF setting for large α and small β (see Section 4.3 and Figure 1).
- 4. We empirically validate our theory; see Figure 1 and Appendix C for details.

1.2 RELATED WORK

Here we discuss directly relevant results; additional related work is deferred to Appendix B.

Empirical Scaling Laws. Modern empirical work shows that performance improves with scale across data, parameters, and compute, following power laws across many domains (Hestness et al., 2017). In language modeling, Kaplan et al. (2020) document power-law loss trends over multiple orders of magnitude and simple budgeting rules linking model size, data, and compute. Henighan et al. (2020) extend these curves to images, video, and multimodal settings. Building on this, Hoffmann et al. (2022) argue that many LMs were under-trained on tokens and proposed data-optimal scaling that substantially improves accuracy at fixed compute. Tissue et al. (2024) investigate the empirical scaling law with learning rate annealing.

Scaling Law Theory. Our work starts from the SGD scaling law in the PLRF model in Paquette et al. (2024) and Lin et al. (2024). In particular, Paquette et al. (2024) derive a scaling-law formula for one-pass SGD, where M, N, and γ_0 denote the model size, number of training steps, and learning rate, respectively, and α and β are the feature- and target-decay parameters.

$$R(M, N, \gamma_0) \approx \underbrace{M^{-2\alpha + \max(0, 1 - 2\beta)}}_{=:\mathcal{A}(M)} + \underbrace{\left(N\gamma_0\right)^{-\frac{2\alpha + 2\beta - 1}{2\alpha}}}_{=:\mathcal{D}_{al}^{\text{SGD}}(N, \gamma_0)} + \underbrace{M^{-1}(N\gamma_0)^{-\frac{2\alpha - 1}{2\alpha}}}_{=:\mathcal{D}_{dis}^{\text{SGD}}(M, N, \gamma_0)} + \underbrace{\gamma_0(N\gamma_0)^{-\frac{4\alpha - 1}{2\alpha}}}_{=:\mathcal{N}^{\text{SGD}}(N, \gamma_0)}.$$

$$(1)$$

The $\mathcal{A}(M)$ corresponds to the approximation error, i.e., the loss as $N \to \infty$. Paquette et al. (2024) explained that $\mathcal{D}_{\rm al}^{\rm SGD}(N,\gamma_0)$ represents the aligned feature loss, as it coincides with the loss for a diagonal sketch matrix S (see Section 2.2 for formal description). They also explained that $\mathcal{D}_{\rm dis}^{\rm SGD}(M,N,\gamma_0)$ corresponds to the distorted feature loss, arising from projection with a random matrix S, and it decays more slowly than the aligned feature loss. Finally, $\mathcal{N}^{\rm SGD}(N,\gamma_0)$ captures the SGD noise, stemming from the quadratic term in the Taylor expansion of the SGD update.

Several subsequent papers extend this baseline along two axes: (i) optimizer changes and (ii) model/training-protocol changes. On the optimizer side, Ferbach et al. (2025) investigate dimension-adapted Nesterov acceleration in the PLRF model and argued that it gives a better scaling law for $2\alpha > 1$ regime. Kunstner & Bach (2025) compare the gradient descent and sign descent scaling law in the linear bigram model. Comparison with their work is in Appendix B.1. Lin et al. (2025) cover the multi-pass SGD scaling law identifies the effect of data reuse for the scaling law. Discussion on the model side is deferred to Appendix B.

the model side is deferred to A

 SignSGD Dynamics. Bernstein et al. (2018a) give the non-convex convergence rate of signSGD. Xiao et al. (2024) derive the SDE and ODE of signSGD risk. The ODE we derive matches theirs in final form; however, we obtain it in an alternative route that does not require a spectral lower bound on the covariance matrix that they imposed. Detailed comparison with Xiao et al. (2024) is in Appendix B.2. Compagnoni et al. (2024) derive SDEs for adaptive methods, including signSGD.

2 PROBLEM SETUP

2.1 NOTATION

We use bold lowercase letters (e.g., u) to denote vectors and bold uppercase letters (e.g., A) to denote matrices. For vectors u and v, we denote the outer product by $u \otimes v := uv^{\mathsf{T}}$. And $\lambda_i(A)$ denotes the i-th eigenvalue of the matrix A. For positive-valued functions f(x) and g(x), we use $f(x) \lesssim g(x)$ if there exists C > 0 such that $f(x) \leq Cg(x)$ for sufficiently large x, and we use $f(x) \approx g(x)$ if there exist $c \in C$ 0 such that $c \in C$ 1 such that $c \in C$ 2 by $c \in C$ 3 for sufficiently large $c \in C$ 4.

2.2 Model

We consider the power-law random features (PLRF) model, parameterized by $\boldsymbol{\theta} \in \mathbb{R}^M$. Given a feature-label pair $(\boldsymbol{x},y) \in \mathbb{R}^d \times \mathbb{R}$, the parameter $\boldsymbol{\theta}$ plays the role of a linear regression coefficient vector on the sketched features \boldsymbol{Sx} (for some $\boldsymbol{S} \in \mathbb{R}^{M \times d}$), and the population risk function is

$$L(\boldsymbol{\theta}) = \mathbb{E}_{\boldsymbol{x}} [(\langle \boldsymbol{S} \boldsymbol{x}, \boldsymbol{\theta} \rangle - y)^2].$$

The data are generated as follows: the feature vector $\boldsymbol{x} \in \mathbb{R}^d$ is drawn from $\mathcal{N}(0, \boldsymbol{H})$ with $\boldsymbol{H} = \operatorname{diag}(1^{-2\alpha}, 2^{-2\alpha}, \dots, d^{-2\alpha})$, and the label is $y = \langle \boldsymbol{x}, \boldsymbol{w}^* \rangle$ with $\boldsymbol{w}^* = [1^{-\beta}, 2^{-\beta}, \dots, d^{-\beta}]^{\top}$; we call α and β feature-decay and target-decay parameters, respectively. The sketch matrix $\boldsymbol{S} \in \mathbb{R}^{M \times d}$ is a random matrix that has i.i.d. entries $\mathcal{N}(0, 1/M)$, is drawn once and then held fixed throughout training; we refer to M (with $M \leq d$) as the model size. Under these model assumptions,

$$L(\boldsymbol{\theta}) = \|\boldsymbol{H}^{1/2}(\boldsymbol{S}^{\mathsf{T}}\boldsymbol{\theta} - \boldsymbol{w}^*)\|^2.$$

We assume $d \ge rM$ for some r > 1, and let $d/M \to (1, \infty]$ as $d, M \to \infty$ when $2\alpha > 1$, and $d/M \to (1, \infty)$ when $2\alpha < 1$. The projected optimal parameter is

$$\boldsymbol{\theta}^* = (\boldsymbol{S}\boldsymbol{H}\boldsymbol{S}^\mathsf{T})^{-1}\boldsymbol{S}\boldsymbol{H}\boldsymbol{w}^*. \tag{2}$$

Define $w_{\perp} = w^* - S^{\mathsf{T}} \theta^*$ so that $w^* = S^{\mathsf{T}} \theta^* + w_{\perp}$ and $SHw_{\perp} = 0$. The risk decomposes as

$$L(\boldsymbol{\theta}) = \|\boldsymbol{H}^{1/2}\boldsymbol{S}^\mathsf{T}(\boldsymbol{\theta} - \boldsymbol{\theta}^*)\|^2 + \|\boldsymbol{H}^{1/2}\boldsymbol{w}_\perp\|^2,$$

where the second term represents the approximation error.

SignSGD. We estimate the minimizer of the population risk via empirical risk minimization using signSGD. At step k, we draw a fresh sample (x_k, y_k) from the model in Section 2.2 and form the stochastic gradient

$$\mathbf{g}_k = (\langle \mathbf{S} \mathbf{x}_k, \mathbf{\theta}_k \rangle - y_k) \mathbf{S} \mathbf{x}_k. \tag{3}$$

The signSGD update rule is

$$\theta_{k+1} = \theta_k - \gamma_k \operatorname{sign}(g_k) = \theta_k - \gamma_k \operatorname{sign}(\langle Sx_k, \theta_k \rangle - y_k) \operatorname{sign}(Sx_k).$$

2.3 Representation of the Result

Let $R(M,N,\gamma_0)$ denote the $L(\boldsymbol{\theta}_N)$ under learning rate γ_0 and fixed model size M. We define the computational budget in terms of FLOPs as $\mathfrak{f}=MN$, and consider the optimal model size M^\star under fixed \mathfrak{f} , and optimal scaling of learning rate in the form $\gamma_0^\star=M^{-e^\star}$. For SGD, Paquette et al. (2024) derive compute-optimal scaling laws of the following form:

$$M^{\star} \approx \mathfrak{f}^{\xi}, \qquad R\left(M^{\star}, \frac{\mathfrak{f}}{M^{\star}}, \gamma_{0}^{\star}\right) \approx \mathfrak{f}^{-\eta}.$$

Our objective is to derive analogous formulas for signSGD, namely $R(M,N,\gamma_0)$ and $R(M^\star,\frac{f}{M^\star},\gamma_0^\star)$, and to compare them with the corresponding results for SGD.

3 ANALYZING THE SIGNSGD

In this section, we formulate the implicit integral equation for signSGD. We define

$$K = SHS^{\mathsf{T}}, \quad \overline{K} = \operatorname{diag}(K)^{-1/2}K, \quad K_{\sigma} = \arcsin(\operatorname{diag}(K)^{-1/2}K\operatorname{diag}(K)^{-1/2}), \quad (4)$$

where arcsin is applied entry-wise; we use these matrices and notation throughout the paper. We decompose the risk via

$$r_i(N) \coloneqq (\boldsymbol{\theta}_N - \boldsymbol{\theta}^*)^{\mathsf{T}} (\boldsymbol{K} \boldsymbol{u}_i \otimes \boldsymbol{w}_i) (\boldsymbol{\theta}_N - \boldsymbol{\theta}^*),$$

where u_i, w_i are the right/left eigenvectors of \overline{K} corresponding to the *i*th eigenvalue $\lambda_i(\overline{K})$. This modal decomposition matches that of Xiao et al. (2024). For brevity we write $L(N) \equiv L(\theta_N)$.

$$L(N) = \sum_{i=1}^{M} r_i(N) + \|\mathbf{H}^{1/2}\mathbf{w}_{\perp}\|^2.$$
 (5)

In Appendix D.1, we derive the one-step update formula for signSGD on a quadratic objective, using a second-order Taylor expansion and sign–Gaussian identities. Applying this to r_i yields

$$\mathbb{E}[r_i(k+1) - r_i(k) \,|\, \mathcal{F}_k] = -\underbrace{\frac{4\gamma_k}{\pi\sqrt{L(k)}}\,\lambda_i(\overline{\boldsymbol{K}})}_{\text{drift}} \, r_i(k) + \underbrace{\frac{2\gamma_k^2}{\pi}\,\boldsymbol{w}_i^\mathsf{T}\boldsymbol{K}_\sigma\boldsymbol{K}\boldsymbol{u}_i}_{\text{quadratic noise}}.$$
 (6)

- 1. **Drift.** The first term in (6) yields a systematic decrease of mode i: it is proportional to the curvature $\lambda_i(\overline{K})$ and the learning rate γ_k , while the factor $1/\sqrt{L(k)}$ self–normalizes the step. Note that the directions corresponding to larger eigenvalues contract faster.
- 2. Quadratic noise. The second term in (6) is an $O(\gamma_k^2)$ variance injection shaped by curvature and the sign-noise covariance. It is independent of $r_i(k)$ and may set a mode-dependent noise floor, unless γ_k decays.

Overall, one-step progress reflects a balance between drift and quadratic noise: when $r_i(k)$ is large, the drift decreases $r_i(k)$; near the optimum, quadratic noise can dominate and cause $r_i(k)$ to plateau.

Converting the one-step update formula to the continuous-time ODE, we obtain ²

$$\frac{dr_i}{dt} = -\underbrace{\frac{4\gamma_{t/\gamma_0}}{\pi\gamma_0\sqrt{L(t)}}\lambda_i(\overline{K})}_{=:\Phi_i^{\text{prift}}(t)} r_i(t) + \underbrace{\frac{2\gamma_{t/\gamma_0}^2}{\pi\gamma_0} \boldsymbol{w}_i^{\mathsf{T}} \boldsymbol{K}_{\sigma} \boldsymbol{K} \boldsymbol{u}_i}_{=:\Phi_i^{\text{noise}}(t)}.$$
(7)

²We treat L and r_i as their continuous extensions, allowing arbitrary positive real inputs.

Compared to SGD, the drift is self-normalized by $1/\sqrt{L(t)}$ and the quadratic noise term does *not* carry the extra L(t) factor present in SGD. So, for the constant learning rate, the quadratic noise does not decrease over time. The variation-of-constants formula gives the implicit integral representation

$$r_i(N) = r_i(0) \, \exp\left\{-\int_0^N \Phi_i^{\text{drift}}(u) \, du\right\} + \int_0^N \exp\left\{-\int_z^N \Phi_i^{\text{drift}}(u) \, du\right\} \times \Phi_i^{\text{noise}}(z) \, dz. \tag{8}$$

Summing over modes, we define

$$L^{\text{drift}}(N) = \sum_{i=1}^{M} r_i(0) \, \exp\left\{-\int_0^N \Phi_i^{\text{drift}}(u) \, du\right\},\tag{9}$$

$$L^{\text{noise}}(N) = \sum_{i=1}^{M} \int_{0}^{N} \exp\left\{-\int_{z}^{N} \Phi_{i}^{\text{drift}}(u) du\right\} \times \Phi_{i}^{\text{noise}}(z) dz. \tag{10}$$

Exact formulation of $L^{\text{drift}}(N)$ and $L^{\text{noise}}(N)$ can be find in (18) of Appendix D.2. Then by (5) our risk is decomposed as

$$L(N) = L^{\text{drift}}(N) + L^{\text{noise}}(N) + \underbrace{\|\boldsymbol{H}^{1/2}\boldsymbol{w}_{\perp}\|^{2}}_{\text{approx}}.$$
 (11)

4 MAIN RESULTS

4.1 Loss Formula for Constant Learning Rate

We now analyze (11) to get $R(M, N, \gamma_0)$, which is L(N) under learning rate γ_0 and model size M.

• For $L^{\text{drift}}(N)$, we use a deterministic approximation (Appendix D.2.2) similar to Paquette et al. (2024), and obtain the asymptotic self-consistent equation: with $\Gamma_M = M^{\min(\alpha,0.5)}\gamma_0$,

$$L^{\text{drift}}(N) \approx \left(\Gamma_M \int_0^N L^{\text{drift}}(u)^{-1/2} du\right)^{-\frac{2\alpha + 2\beta - 1}{2\alpha}} + M^{-1} \left(\Gamma_M \int_0^N L^{\text{drift}}(u)^{-1/2} du\right)^{-\frac{2\alpha - 1}{2\alpha}}.$$

Solving this yields signSGD counterparts of the aligned- and distorted- feature loss terms in (1), denoted by $\mathcal{D}_{\rm al}^{\rm sign}(M,N,\gamma_0)$ and $\mathcal{D}_{\rm dis}^{\rm sign}(M,N,\gamma_0)$; see (12) below for their precise forms.

ullet For $L^{\mathrm{noise}}(N)$ and approximation term, we calculate the limit loss L_{∞} and get

$$L_{\infty} \approx \max \left\{ \gamma_0^2 M^{2-\min(1,2\alpha)}, \|\boldsymbol{H}^{1/2} \boldsymbol{w}_{\perp}\|^2 \right\}$$

Lastly we use approximation error result from Paquette et al. (2024); Lin et al. (2024),

$$\|\boldsymbol{H}^{1/2}\boldsymbol{w}_{\perp}\|^2 \approx M^{-2\alpha + \max(0, 1 - 2\beta)}.$$

Combining two parts yields a proxy, and we prove that it satisfies the implicit integral equation (11) in Appendix D.3.4 and D.4.4. Finally, we get the following four-term scaling law formula for one-pass signSGD on area $-\alpha + 0.5 < \beta < \alpha + 0.5$:

$$R(M, N, \gamma_{0}) \approx \underbrace{M^{-2\alpha + \max(0, 1 - 2\beta)}}_{=:\mathcal{A}(M)} + \underbrace{\left(M^{\min(\alpha, 0.5)} N \gamma_{0}\right)^{-\frac{2(2\alpha + 2\beta - 1)}{2\alpha - 2\beta + 1}}}_{=:\mathcal{D}_{\mathrm{al}}^{\mathrm{sign}}(M, N, \gamma_{0})} + \underbrace{M^{-\frac{6\alpha - 1}{2\alpha + 1}}(N \gamma_{0})^{-\frac{2(2\alpha - 1)}{2\alpha + 1}}}_{=:\mathcal{D}_{\mathrm{dis}}^{\mathrm{sign}}(M, N, \gamma_{0})} + \underbrace{\gamma_{0}^{2} M^{2 - \min(1, 2\alpha)}}_{=:\mathcal{N}^{\mathrm{sign}}(M, \gamma_{0})}.$$

$$(12)$$

³For the case $\beta > \alpha + 0.5$, $\mathcal{D}_{\rm al}^{\rm sign}(M,N,\gamma_0)$ takes form of $\left(1 - \kappa M^{\min(\alpha,0.5)}N\gamma_0\right)^{-\frac{2(2\alpha+2\beta-1)}{2\alpha-2\beta+1}}$. See Appendix D.5 for more details.

Interpretation. The term $\mathcal{A}(M)$ is the approximation error (irreducible as $N \to \infty$). The terms $\mathcal{D}^{\mathrm{sign}}_{\mathrm{al}}(M,N,\gamma_0)$ and $\mathcal{D}^{\mathrm{sign}}_{\mathrm{dis}}(M,N,\gamma_0)$ arise from the drift's exponential damping $r_i(0) \exp\left\{-\int_0^N \Phi_i^{\mathrm{drift}}(u) \, du\right\}$ and correspond to the aligned and distorted feature losses of SGD scaling law in Paquette et al. (2024). The term $\mathcal{N}^{\mathrm{sign}}(M,\gamma_0)$ captures the quadratic noise from the one-step Taylor expansion, specific to one-pass signSGD.

Comparison. We compare our signSGD scaling law formula with the SGD formula (1) of Paquette et al. (2024). Since the approximation error is optimizer-independent, the term $\mathcal{A}(M)$ remains unchanged. For the N-exponent in $\mathcal{D}_{\rm al}$ and $\mathcal{D}_{\rm dis}$, when the absolute value of the exponent is x for SGD, then it changes to $\frac{2}{2-x}x$ in signSGD, which is strictly larger than x. Therefore, $\mathcal{D}_{\rm al}^{\rm sign}(M,N,\gamma_0)$ and $\mathcal{D}_{\rm dis}^{\rm sign}(M,N,\gamma_0)$ decrease faster in the number of steps N under signSGD. By contrast, the signSGD noise term $\mathcal{N}^{\rm sign}(M,\gamma_0)$ does not decay with N, whereas the SGD noise $\mathcal{N}^{\rm SGD}(N,\gamma_0)$ does.⁴

We discuss underlying mechanism that modifies the drift terms \mathcal{D}_{al} , \mathcal{D}_{dis} , and the noise term \mathcal{N} .

- Drift terms (Drift-normalization effect): In signSGD, the drift in (6) is $\frac{4\gamma_k}{\pi\sqrt{L(k)}}\lambda_i(\overline{\pmb{K}})$, whereas for SGD it is $2\gamma_k\,\lambda_i(\pmb{K})$; see (4) for $\pmb{K},\overline{\pmb{K}}$. The diagonal preconditioning embedded in $\overline{\pmb{K}}$ contributes an extra factor $M^{\min(\alpha,1/2)}$. The normalization by $\sqrt{L(k)}$ replaces the effective flow time $N\gamma_0$ with $\gamma_0\int_0^N L(u)^{-1/2}\,du$, which accelerates progress in training whenever $L(u)\lesssim 1$. Thus, in the aligned/distorted drift terms, $(N\gamma_0)$ is replaced by $M^{\min(\alpha,1/2)}\gamma_0\int_0^N L(u)^{-1/2}\,du$.
- Noise term (Noise-reshaping effect): The signSGD noise in (6) is $\frac{2\gamma_k^2}{\pi} \boldsymbol{w}_i^{\top} \boldsymbol{K}_{\sigma} \boldsymbol{K} \boldsymbol{u}_i$, while for SGD it is $\gamma_k^2 (\boldsymbol{v}_i^{\top} \boldsymbol{K} \boldsymbol{v}_i) L(k)$ with \boldsymbol{v}_i an eigenvector of \boldsymbol{K} . The normalization removes the multiplicative L(k) in signSGD, eliminating the Volterra structure present in Paquette et al. (2024). This difference is crucial: the lack of L(k) in the quadratic term ultimately yields a noise term that does not decay in N. Meanwhile, an additional M-dependence arises from working in the $\overline{\boldsymbol{K}}$ -(rather than \boldsymbol{K} -) eigenbasis due to diagonal preconditioning.

4.2 COMPUTE-OPTIMAL RESULT UNDER OPTIMAL CONSTANT LEARNING RATE

In the constant learning-rate schedule, we allow γ_0 to scale with the model size via $\gamma_0=M^{-e}$. The hyperparameter e directly influences the compute-optimal scaling law. Following Paquette et al. (2024), we distinguish the *maximal* and *optimal* learning rates for SGD. The maximal rate is the largest step that yields a stable (non-exploding) recursion; for signSGD, it leads to a zero compute-optimal slope (see Appendix E.1). We therefore focus on the optimal learning rate γ_0^{\star} , which maximizes the decay exponent η in

$$R(M^{\star}, \mathfrak{f}/M^{\star}, \gamma_0^{\star}) \approx \mathfrak{f}^{-\eta},$$

where M^* denotes the model size minimizing $R(\cdot)$ at fixed compute budget f.

To characterize the compute-optimal scaling, set $\gamma_0 = M^{-e}$, $M = \mathfrak{f}^x$, and $N = \mathfrak{f}^{1-x}$ (with $x \in [0,1]$), and solve

$$(e^*, x^*) \in \operatorname*{arg\,min}_{e, x} R(M, N, \gamma_0) = \operatorname*{arg\,min}_{e, x} R(\mathfrak{f}^x, \mathfrak{f}^{1-x}, \mathfrak{f}^{-ex}). \tag{13}$$

Then $M^\star=\mathfrak{f}^{x^*}$, $N^\star=\mathfrak{f}^{1-x^*}$, and $\gamma_0^\star=(M^\star)^{-e^*}$, and at the optimum

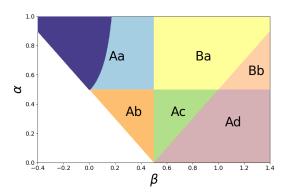
$$R(M^{\star}, \mathfrak{f}/M^{\star}, \gamma_0^{\star}) \approx \mathfrak{f}^{-\eta(\alpha,\beta)},$$

for some $\eta(\alpha, \beta) > 0$, which we refer to as the compute-optimal slope.

In problem (13), each of the four terms in (12) scales as $f^{-\ell_i(e,x)}$, so minimizing R is equivalent to maximizing $\min\{\ell_1,\ell_2,\ell_3,\ell_4\}$. The optimal value (e^*,x^*) is obtained by balancing three active exponents. The resulting formulas and dominant and balancing terms are summarized in Table 1; see Appendix E.2 for details.

⁴As we set γ_0 as M^{-e} later, decay with respect to M depends on the choice of γ_0 .

⁵One may wonder why we do not parameterize by N. Setting $\gamma_0 = M^{-e}$ is without loss of generality, since in the compute-optimal case both M and N are expressed as powers of the total compute \mathfrak{f} .



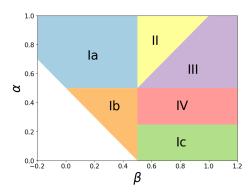


Figure 2: **Left: Phase plane for sign SGD; Right: Phase plane for SGD.** The white region indicates parameter values with no power-law scaling. The dark blue area represents the region where stable-decay scheduling (Section 4.3) yields a better compute-optimal exponent.

Table 1: Dominant and balancing terms, optimal learning rate, compute-optimal model size, and risk across different (α, β) phases. Refer to (12) for the definitions of the terms \mathcal{A} , \mathcal{D}_{al} , \mathcal{D}_{dis} , \mathcal{N} . See Figures 8 to 12 in the Appendix for empirical validation of the theoretical exponents.

	Term structure			Compute-optimal			
Phase	Dominant terms	Balancing terms		γ_0^\star	M^{\star}	$R\left(M^\star,rac{\mathrm{f}}{M^\star},\gamma_0^\star ight)$	
Phase A	$\mathcal{A}, \mathcal{D}_{ m al}, \mathcal{N}$	$\mathcal{A}, \mathcal{D}_{ m al}, \mathcal{N}$			$f^{\frac{1}{2\alpha+1}}$	$\mathfrak{f}^{-\frac{2\alpha+2\beta-1}{2\alpha+1}}$	
			Ab	$M^{-\frac{2\beta+1}{2}}$	$\mathfrak{f}^{rac{1}{2}}$	$\mathfrak{f}^{-\frac{2\alpha+2\beta-1}{2}}$	
			Ac	M^{-1}	$\mathfrak{f}^{\frac{2\alpha+2\beta-1}{2(2\beta-\alpha(2\beta-3)-1)}}$	$\mathfrak{f}^{-\frac{\alpha(2\alpha+2\beta-1)}{2\beta-\alpha(2\beta-3)-1}}$	
			Ad	M^{-1}	$f^{\frac{1}{2-\alpha}}$	$\mathfrak{f}^{-\frac{2\alpha}{2-\alpha}}$	
Phase B	$\mathcal{A}, \mathcal{D}_{ m al}, \mathcal{D}_{ m dis}, \mathcal{N}$	$\mathcal{D}_{ m al},\!\mathcal{D}_{ m dis},\mathcal{N}$	Ba	$M^{-\frac{2\alpha+4\beta-1}{4\beta}}$	$f^{\frac{\beta}{\alpha+\beta}}$	$\mathfrak{f}^{-\frac{2\alpha+2\beta-1}{2\alpha+2\beta}}$	
			Bb	$M^{-\frac{6\alpha+1}{4\alpha+2}}$	$\mathfrak{f}^{\frac{2\alpha+1}{4\alpha+1}}$	$\mathfrak{f}^{-rac{4lpha}{4lpha+1}}$	

We follow Paquette et al. (2024) in defining phases by dominant terms; to avoid confusion with their SGD phases, we label our signSGD phases by uppercase letters. Accordingly, any reference to Phase I–IV hereafter refers exclusively to the SGD phases of Paquette et al. (2024). For signSGD, the phase plane is simpler: when $\alpha>0.5$ and $\beta>0.5$ (Phase B) all four terms are dominant; otherwise (Phase A) the dominant terms are $\mathcal{A}(M)$, $\mathcal{D}_{\rm al}^{\rm sign}(M,N,\gamma_0)$, and $\mathcal{N}^{\rm sign}(M,\gamma_0)$. We declare *subphases* whenever the formula of at least one of $\gamma_0=M^{-e^*}$, M^* , or $R(M^*,\mathfrak{f}/M^*,\gamma_0^*)$ changes. These changes occur across the boundaries $\alpha=0.5$, $\beta=0.5$, and $\beta=\alpha+0.5$, yielding six subphases in total (Phase A split into four, Phase B into two). For context, Paquette et al. (2024) partition the (α,β) -plane into four phases with seven subphases.

Remark 1 (Dominant vs. balancing terms). Dominant terms are those that can lead the risk for some (γ_0, M, N) . Balancing terms are the ones that tie (hence "balancing") at the compute-optimal choice (γ_0^*, M^*, N^*) and therefore determine the slope; they form a subset of the dominant terms.

Comparison of Compute-optimal Results. For the intersection of Phase Aa, Ab, Ac, Ba and Phase I, II, the compute-optimal slope $\eta(\alpha,\beta)$ and optimal model size M^\star are the same for signSGD and SGD. In contrast, for the area of Phase III, IV excluding the case $0.25 < \alpha < 1/3$, $\beta > (1-\alpha)(1-2\alpha)/(2(1-3\alpha))$ (See Figure 4 in the Appendix for the visualization of this area), the compute-optimal slope $\eta(\alpha,\beta)$ for signSGD is *steeper* than that for SGD, and the optimal model size is bigger in signSGD. We refer to this region as the Area III-IV sub. Finally, for the optimal learning rate $\gamma_0 = M^{-e^*}$, the exponent e^* is always bigger in signSGD, which means signSGD always has a smaller optimal learning rate.

4.3 EFFECT OF STABLE-DECAY SCHEDULING

For a stable-decay schedule, we set the learning rate to $\gamma_k = \gamma_0 f(k)$ with

$$f(k) = \begin{cases} 1, & k \le pN, \\ (1 + \tau(k - pN))^{-c}, & k > pN, \end{cases}$$
 (14)

where $p, c \in (0, 1)$ and $\tau > 0$. In other words, the learning rate remains constant for the first pN steps, and then decays polynomially with exponent c for the remaining (1 - p)N steps.

In Phase Aa, the f-scheduled noise bound can improve over constant LR:

$$L^{\text{noise}}(N) \lesssim \gamma_0^2 M N^{-2c} + \gamma_0^{\frac{1}{2\alpha}} M^{\frac{1}{4\alpha}} N^{-(1-c)(1-\frac{1}{2\alpha})}.$$

Combining this with the drift and approximation terms, and then optimizing over e of $\gamma_0 = M^{-e}$, the decay parameter c, and the model size M, yields the f-scheduled risk bound

$$R_f(M^*, \mathfrak{f}/M^*, (M^*)^{-e^*}) \lesssim \mathfrak{f}^{-\frac{2(4\alpha-1)(2\alpha+2\beta-1)}{16\alpha^2+8\alpha\beta+2\alpha-2\beta-1}}.$$
 (15)

The absolute value of the exponent in (15) exceeds the compute-optimal slope under constant learning rate when $\alpha > 0.5$ and $0.5 - \alpha < \beta < \frac{2\alpha - 1}{2(4\alpha - 1)}$. Thus, stable-decay scheduling yields a strictly larger compute-optimal slope in the upper left region of Phase Aa (marked with dark blue in Figure 2). We will refer to this region as Area Aa* throughout the paper.

Scheduling does not improve the SGD compute-optimal exponent in Phases I–II (see Appendix F.4). Thus, with scheduling, signSGD achieves a larger compute-optimal exponent compared to SGD in Area Aa^{\star} .

5 DISCUSSION: WHERE AND WHY SIGNSGD PROVIDES BENEFITS?

With a constant learning rate $\gamma_0 = M^{-e}$, signSGD yields improvements over SGD in Area III-IV_{sub}. Under stable-decay scheduling, we find signSGD also provides benefits in Area Aa*.

Mechanisms. These gains can be explained by *noise-reshaping*, together with *drift-normalization*. In Paquette et al. (2024), Phases III–IV are the SGD noise-bottleneck regimes. By contrast, noise-reshaping in signSGD can alleviate this bottleneck with a suitable learning-rate choice, yielding improved compute-optimal slopes.

Role of Learning-rate Scaling. The signSGD noise term with constant LR is $\mathcal{N}^{\mathrm{sign}}(M,\gamma_0) = \gamma_0^2 M^{2-\min(1,2\alpha)}$, whereas for SGD it is $\mathcal{N}^{\mathrm{SGD}}(N,\gamma_0) = \gamma_0(N\gamma_0)^{-(4\alpha-1)/(2\alpha)}$. If $\gamma_0 \approx 1$, $\mathcal{N}^{\mathrm{sign}}(M,\gamma_0)$ is much larger than $\mathcal{N}^{\mathrm{SGD}}(N,\gamma_0)$, making the compute-optimal slope asymptotically zero. Hence, we set $\gamma_0 = M^{-e}$ and optimize e to balance terms and obtain a steep compute-optimal curve: decreasing γ_0 lowers $\mathcal{N}^{\mathrm{sign}}(M,\gamma_0)$ while increasing the drift terms $\mathcal{D}^{\mathrm{sign}}_{\mathrm{al}}(M,N,\gamma_0)$ and $\mathcal{D}^{\mathrm{sign}}_{\mathrm{dis}}(M,N,\gamma_0)$, and the optimal e strikes the balance.

Why Gains Arise in Area III-IV_{sub}. For SGD, the shape of $\mathcal{N}^{\mathrm{SGD}}(N,\gamma_0)$ makes it dominate $\mathcal{D}^{\mathrm{SGD}}_{\mathrm{al}}(N,\gamma_0)$ at the compute-optimal point in Phases III–IV. For signSGD, noise-reshaping alters $\mathcal{N}^{\mathrm{sign}}(M,\gamma_0)$ so it can *balance* against $\mathcal{D}^{\mathrm{sign}}_{\mathrm{al}}(M,N,\gamma_0)$, while drift-normalization steepens the decay of $\mathcal{D}^{\mathrm{sign}}_{\mathrm{al}}(M,N,\gamma_0)$. This creates room for a balance in which both terms are smaller than the SGD noise $\mathcal{N}^{\mathrm{SGD}}(N,\gamma_0)$ at optimum, explaining the improvements in Area III-IV_{sub}.

Why Stable-decay Scheduling Helps. For a learning-rate schedule $\gamma_k = \gamma_0 f(k)$ with general f, the drift-only self-consistent solution in Phase Aa takes the form

$$(M^{1/2} \gamma_0 F(N))^{-\frac{2(2\alpha+2\beta-1)}{2\alpha-2\beta+1}}, \quad \text{where} \quad F(N) := \int_0^N f(u) du.$$

⁶Whether scheduling benefits other regions of signSGD or other phases of SGD remains open, since for both methods the scheduled noise upper and lower bounds do not match tightly, even up to constant factors.

This can be viewed as $\mathcal{D}^{\mathrm{sign}}_{\mathrm{al}}(M,N,\gamma_0)$ with N replaced by F(N). This aligns with empirical observations that a loss term can decay polynomially with the area under the learning-rate curve (Tissue et al., 2024).

In contrast, the noise term depends most heavily on the learning rate *near the end* of training, since earlier noise can be damped by later drift; see (8). Stable-decay preserves the total area F(N) asymptotically while shrinking the late-stage learning rate, thereby reducing noise without sacrificing drift. As a result, stable-decay scheduling yields a larger compute-optimal slope in Area Aa* (upper-left Phase Aa; see Section 5.1 for intuition). More broadly, we conjecture that appropriate scheduling can further reduce the signSGD noise term, enabling improvements beyond Area III-IV_{sub}.

5.1 Hypothesis for the Position of the Beneficial Area

Here, we hypothesize why the areas with improved scaling law lie near the left edge (small β) and the right side ($\beta > \alpha$) of the phase plane.

Heuristic Criterion. Let "target decay" denote the decay of the projected optimum θ^* in (2), and "stochastic-gradient decay" the decay of the stochastic gradient in (3). SignSGD is advantageous when the *target decays more slowly* than the stochastic gradient. Under SGD, coordinates with smaller gradients take smaller updates; if the target does not decay much, those coordinates still require learning targets of comparable magnitude, so more iterations are needed—an inefficiency that signSGD mitigates by normalizing per-coordinate updates via the sign operation.

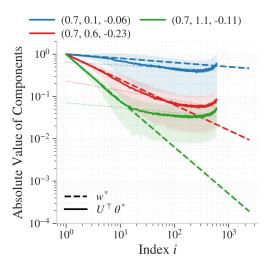


Figure 3: Decay of θ^* in the basis of columns of U compared to w^* . The legend on the top shows $(\alpha, \beta, \text{fitted slope of } U^{\mathsf{T}}\theta^*)$.

When Does This Occur? Observations and Conjecture Writing $SHS^{\top} = U\Lambda U^{\top}$, the expected stochastic-gradient along the U basis decays as $i^{-2\alpha}$. For the projected target $\theta^* = 0$

 $(SHS^{\top})^{-1}SHw^*$, the components $U^{\top}\theta^*$ behave similarly to $w^* = [j^{-\beta}]_{j=1}^d$ on the leading indices. The critical part is the relative decay rate of $U^{\top}\theta^*$ versus $i^{-2\alpha}$. See Appendix G for details of analysis.

Figure 3 illustrates our empirical observations about these patterns. For $(\alpha, \beta) = (0.7, 1.1)$, $U^{\top} \theta^*$ plateaus quickly; for (0.7, 0.6) it decays longer; and for (0.7, 0.1), since w^* hardly decays, the target also shows little decay.

These observations suggest that in the left region (small β) and the right region ($\beta > \alpha$), the targets decay more slowly than the stochastic gradient, whereas in the middle band ($0.5 < \beta < \alpha$) they do not. This could potentially explain why the signSGD-beneficial area appears near the left edge and the right side of the phase plane.

6 CONCLUSION

We derived the scaling law of signSGD under the PLRF model and identified two distinctive effects—drift-normalization and noise-reshaping—relative to SGD. Analyzing compute-optimal tradeoffs, we showed that signSGD achieves steeper slopes than SGD in the noise-bottleneck regimes, and that a stable-decay schedule further improves performance in the Area Aa*. Additionally, in Appendix H, we analyze Adam using the heuristics of Xiao et al. (2024) and observe the same scaling law as signSGD, consistent with our experiments; deriving Adam's scaling law without heuristic assumptions is a compelling direction. We defer limitations and additional future works to Appendix A.

REFERENCES

- Yasaman Bahri, Ethan Dyer, Jared Kaplan, Jaehoon Lee, and Utkarsh Sharma. Explaining neural scaling laws. *Proceedings of the National Academy of Sciences*, 121(27):e2311878121, 2024.
- Lukas Balles, Fabian Pedregosa, and Nicolas Le Roux. The geometry of sign gradient descent. *arXiv* preprint arXiv:2002.08056, 2020.
- Jeremy Bernstein, Yu-Xiang Wang, Kamyar Azizzadenesheli, and Animashree Anandkumar. signsgd: Compressed optimisation for non-convex problems. In *International conference on machine learning*, pp. 560–569. PMLR, 2018a.
- Jeremy Bernstein, Jiawei Zhao, Kamyar Azizzadenesheli, and Anima Anandkumar. signsgd with majority vote is communication efficient and fault tolerant. *arXiv preprint arXiv:1810.05291*, 2018b.
- Blake Bordelon, Alexander Atanasov, and Cengiz Pehlevan. A dynamical model of neural scaling laws. In *International Conference on Machine Learning*, pp. 4345–4382. PMLR, 2024.
- Blake Bordelon, Alexander Atanasov, and Cengiz Pehlevan. How feature learning can improve neural scaling laws. *Journal of Statistical Mechanics: Theory and Experiment*, 2025(8):084002, 2025.
- Enea Monzio Compagnoni, Tianlin Liu, Rustem Islamov, Frank Norbert Proske, Antonio Orvieto, and Aurelien Lucchi. Adaptive methods through the lens of sdes: Theoretical insights on the role of noise. *arXiv preprint arXiv:2411.15958*, 2024.
- Shihong Ding, Haihan Zhang, Hanzhen Zhao, and Cong Fang. Scaling law for stochastic gradient descent in quadratically parameterized linear regression. *arXiv* preprint arXiv:2502.09106, 2025.
- Damien Ferbach, Katie Everett, Gauthier Gidel, Elliot Paquette, and Courtney Paquette. Dimension-adapted momentum outscales sgd. *arXiv preprint arXiv:2505.16098*, 2025.
- Tom Henighan, Jared Kaplan, Mor Katz, Mark Chen, Christopher Hesse, Jacob Jackson, Heewoo Jun, Tom B Brown, Prafulla Dhariwal, Scott Gray, et al. Scaling laws for autoregressive generative modeling. *arXiv preprint arXiv:2010.14701*, 2020.
- Joel Hestness, Sharan Narang, Newsha Ardalani, Gregory Diamos, Heewoo Jun, Hassan Kianinejad, Md Mostofa Ali Patwary, Yang Yang, and Yanqi Zhou. Deep learning scaling is predictable, empirically. *arXiv preprint arXiv:1712.00409*, 2017.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. Training compute-optimal large language models. arXiv preprint arXiv:2203.15556, 2022.
- Marcus Hutter. Learning curve theory. arXiv preprint arXiv:2102.04074, 2021.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.
- Sai Praneeth Karimireddy, Quentin Rebjock, Sebastian Stich, and Martin Jaggi. Error feedback fixes signsgd and other gradient compression schemes. In *International conference on machine learning*, pp. 3252–3261. PMLR, 2019.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint* arXiv:1412.6980, 2014.
- Tanishq Kumar, Zachary Ankner, Benjamin F Spector, Blake Bordelon, Niklas Muennighoff, Mansheej Paul, Cengiz Pehlevan, Christopher Ré, and Aditi Raghunathan. Scaling laws for precision. *arXiv preprint arXiv:2411.04330*, 2024.
 - Frederik Kunstner and Francis Bach. Scaling laws for gradient descent and sign descent for linear bigram models under zipf's law. *arXiv preprint arXiv:2505.19227*, 2025.

- Frederik Kunstner, Jacques Chen, Jonathan Wilder Lavington, and Mark Schmidt. Noise is not the main factor behind the gap between sgd and adam on transformers, but sign descent might be. *arXiv preprint arXiv:2304.13960*, 2023.
 - Licong Lin, Jingfeng Wu, Sham M Kakade, Peter Bartlett, and Jason D Lee. Scaling laws in linear regression: Compute, parameters, and data. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
 - Licong Lin, Jingfeng Wu, and Peter L Bartlett. Improved scaling laws in linear regression via data reuse. *arXiv preprint arXiv:2506.08415*, 2025.
- Bochen Lyu, Di Wang, and Zhanxing Zhu. A solvable attention for neural scaling laws. In *The Thirteenth International Conference on Learning Representations*, 2025.
- Elliot Paquette, Courtney Paquette, Lechao Xiao, and Jeffrey Pennington. 4+ 3 phases of compute-optimal neural scaling laws. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- Tomer Porian, Mitchell Wortsman, Jenia Jitsev, Ludwig Schmidt, and Yair Carmon. Resolving discrepancies in compute-optimal scaling of language models. *Advances in Neural Information Processing Systems*, 37:100535–100570, 2024.
- Utkarsh Sharma and Jared Kaplan. A neural scaling law from the dimension of the data manifold. *arXiv preprint arXiv:2004.10802*, 2020.
- Howe Tissue, Venus Wang, and Lu Wang. Scaling law with learning rate annealing. *arXiv preprint* arXiv:2408.11029, 2024.
- Kaiyue Wen, Zhiyuan Li, Jason Wang, David Hall, Percy Liang, and Tengyu Ma. Understanding warmup-stable-decay learning rates: A river valley loss landscape perspective. arXiv preprint arXiv:2410.05192, 2024.
- Ke Liang Xiao, Noah Marshall, Atish Agarwala, and Elliot Paquette. Exact risk curves of signsgd in high-dimensions: Quantifying preconditioning and noise-compression effects. *arXiv* preprint arXiv:2411.12135, 2024.

SUPPLEMENTARY MATERIALS FOR "SCALING LAWS OF SIGNSGD IN LINEAR REGRESSION: WHEN DOES IT OUTPERFORM SGD?"

USAGE OF LLM

We primarily used LLMs to polish the English writing throughout the paper. They were also employed to help us identify additional related work beyond those we were already familiar with. When preparing well-formatted tables, we relied on LLMs for assistance. We also used LLMs to refine LaTeX code so that complicated formulas appeared clean and readable in the manuscript. Finally, we sought LLM support for debugging code used in our experiments.

OVERVIEW OF APPENDIX

- (1) In Appendix A we discuss limitations and future works.
- (2) In Appendix B we discuss more related works beyond those discussed in Section 1.2, and provide a detailed comparison with closely related works.
- (3) In Appendix C we present experimental results which support our theory.
- (4) In Appendix D we derive the scaling law formula (12) of $R(M, N, \gamma_0)$ under constant learning rate. We first derive a one-step update formula and convert it to an ODE to get an integral equation. We use a deterministic approximation for the integral equation with experimental results. Then we set a proxy of the loss function and verify that it satisfies the integral equation.
- (5) In Appendix E we discuss the maximal learning rate deferred from the main text, and derive the optimal learning rate, compute-optimal loss, and optimal model size in Table 1.
- (6) In Appendix F we derive the result for stable-decay learning rate in Section 4.3.
- (7) In Appendix G we provide analysis for stochastic gradient decay and target decay deferred from Section 5.1.
- (8) In Appendix H we derive scaling law of Adam under heuristic proposed by Xiao et al. (2024), and verify our results with experiment.
- (9) In Appendix I we provide omitted analysis from Appendix D.

A LIMITATION AND FUTURE WORK

Limitation. Our analysis assumes batch size and focuses on the PLRF setting with diagonal covariance \boldsymbol{H} ; extensions to mini-batch and more general covariances are not covered here. For scheduling, we analyze a stable-decay template; exact asymptotics for broader schedules remain open. We also use a deterministic approximation whose accuracy we verify empirically; tightening constants and extending the formal guarantees are left for future work.

Future Work. Combining signSGD with dimension-adapted acceleration (Ferbach et al., 2025) and extending the framework to more complex architectures (e.g., two-layer linear networks or self-attention) are promising avenues.

B ADDITIONAL RELATED WORK

More Related Works on Empirical Scaling Laws. Porian et al. (2024) resolve discrepancy between Kaplan et al. (2020) and Hoffmann et al. (2022). Kumar et al. (2024) investigate precisionaware scaling law.

 More Related Works on Scaling Law Theory. There are lines of work analyzing more complex models compared to the power-law random features (PLRF) model. Bordelon et al. (2025) investigate the scaling law of a two-layer linear neural network with projected gradient descent, and argued the benefit compared to the PLRF model, which is one-layer. Ding et al. (2025) cover the scaling law of quadratically parameterized linear regression with SGD. Lyu et al. (2025) cover the scaling law of linear self-attention under gradient flow.

Sharma & Kaplan (2020) show that test loss scales as a power-law of model size in regression problems. Hutter (2021) investigates binary classification using a tabulation learning algorithm, deriving a power-law scaling with respect to dataset size. Bahri et al. (2024) analyze a linear random features model with SGD, showing a power-law decay in test loss with respect to sample size (or model size, when the other is infinite). Bordelon et al. (2024) derive a power law over model size, dataset size, and time for the linear random features model under gradient flow dynamics.

More Related Works about signSGD and sign descent. Balles et al. (2020) investigate the geometry of sign gradient descent. Kunstner et al. (2023) discover that sign descent could be the key factor making the gap between SGD and Adam on Transformers. Bernstein et al. (2018b) propose signSGD with majority vote, which is communication efficient and fault-tolerant. Karimireddy et al. (2019) prove that error-feedback can make the rate of convergence of signSGD better.

B.1 COMPARISON WITH KUNSTNER & BACH (2025)

First, their work compares the scaling laws of sign descent and gradient descent, whereas our work compares the scaling laws of signSGD and SGD. Second, they analyze for a Linear Bigram Model, while we analyze for the power-law random features (PLRF) model. The advantage of the PLRF model is that it models two parameters each for feature vector decay and target decay, while the Linear Bigram Model has one parameter for data frequency decay. Lastly, they derived a scaling law where the model size goes to infinity; in contrast, our scaling law covers both finite model size and infinite limit by representing the loss as a function of model size, number of steps, and learning rate. This made us possible to analyze the compute-optimal scaling law.

B.2 COMPARISON WITH XIAO ET AL. (2024)

ODE for signSGD in Xiao et al. (2024) is equivalent to the ODE that occurred during our analysis. The reason that we were not able to directly use their ODE is that they derived it under the spectrum lower bound assumption for the covariance matrix. In our case spectrum of the covariance matrix SHS^T is asymptotically the same as $i^{-2\alpha}$, so their assumption does not hold for our setup. So we re-derived the ODE without the spectrum lower bound assumption. Due to the spectrum lower bound assumption, they led to an exponential decay to limit risk, which is completely different from the polynomial neural scaling law derived from our paper. They discussed the noise-reshaping effect on the level of SDE. In contrast, we observed noise reshaping on the level of scaling law and investigated its effect on compute-optimal scaling.

B.3 ADDITIONAL PHASE PLANE PLOTS TO COMPARE WITH PRIOR WORK

Figure 4 indicates the area where signSGD has a steeper compute-optimal slope compared to SGD, by coloring it with Mint green. It lies in Phase Ac, Ad, Ba, Bb, and covers all areas of Phase Bb. In terms of the SGD Phase, it covers all areas of Phase III and most of the areas of Phase IV.

Figure 5 indicates the area where signSGD has a steeper compute-optimal slope compared to DANA-decaying in Ferbach et al. (2025), by coloring it with Lime green. It lies in Phase Ac, Ad, Ba, Bb. It is smaller than the Mint green area, and this is natural, since DANA-decaying in Ferbach et al. (2025) has a steeper slope compared to SGD.

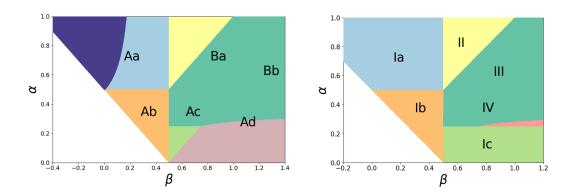


Figure 4: **Phase planes to compare signSGD and SGD.** Mint green area covering all of Phase Bb and III, and some part of Phase Ac, Ad, Ba, IV is the area where signSGD has a steeper compute-optimal slope compared to SGD. The left side is the signSGD phase plane, and the right side is the SGD phase plane. We placed the Mint green area for both of them for clarity. We will call this Mint green area as Area III-IV_{sub}.

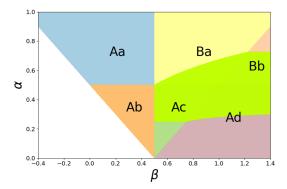


Figure 5: Phase plane to compare signSGD and DANA-decaying in Ferbach et al. (2025). Lime green area covering some part of Phase Ac, Ad, Ba, Bb is the area where signSGD has a steeper compute-optimal slope compared to DANA-decaying in Ferbach et al. (2025).

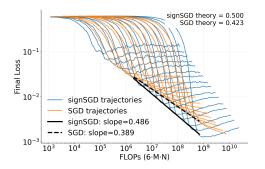
C EXPERIMENTS

C.1 EXPLAINATION FOR FIGURE 1.

Parameters. Left parameters: $(\alpha, \beta) = (0.4, 0.8), \gamma_0 = 0.006, e^* = 1.0 \text{ for signSGD}, e^* = 0.4571 \text{ for SGD}, 20 \text{ runs.}$ Right parameters: $(\alpha, \beta) = (1.0, 0), \gamma_0 = 0.002, e^* = 1.0 \text{ for constant}, e^* = 0.833 \text{ for stable-decay}, <math>c = 0.091, p = 0.1, \tau = 1 \text{ for stable-decay}, 10 \text{ runs}.$

Takeaways. In Figure 1, the left panel demonstrates the steeper compute-optimal slope of signSGD for $(\alpha, \beta) = (0.4, 0.8)$ in the area of Phase Ac. The right panel shows the increase in compute-optimal slope achieved by stable-decay scheduling for $(\alpha, \beta) = (1.0, 0)$. The theoretical and experimental compute-optimal slopes agree within errors of 0.04 (left) and 0.01 (right), which are well within the error margins reported in prior works.

Additionally Figure 6, demonstrates the steeper compute-optimal slope of signSGD for $(\alpha, \beta) = (0.4, 1.0)$ in the Phase Ad and $(\alpha, \beta) = (0.7, 1.1)$ in Phase Ba.



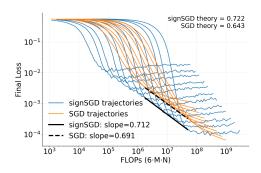
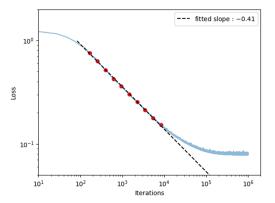


Figure 6: **comparison of SGD and signSGD on Compute-Optimal Scaling.** Colored lines represent the training trajectories of each algorithm, while black lines denote the compute-optimal curves. In both panels, the theoretical compute-optimal predictions closely follow the observed scaling. Both plot shows that signSGD has a steeper compute-optimal slope than SGD. Left parameters: $(\alpha, \beta) = (0.4, 1.0), \gamma_0 = 0.01, e^* = 1.0$ for signSGD, $e^* = 0.533$ for SGD, 5 runs. Right parameters: $(\alpha, \beta) = (0.7, 1.1), \gamma_0 = 0.01, e^* = 1.09$ for signSGD, $e^* = 0$ for SGD, 20 runs.



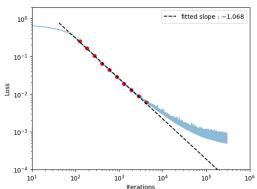


Figure 7: $\mathcal{D}_{\rm al}^{\rm sign}(M,N,\gamma_0)$ term exponent. Blue curves: true signSGD trajectories. Black dotted curves: linear fits over the early-iteration interval in log-log scale. Left: parameters $(\alpha,\beta)=(0.75,0),\ \gamma_0=0.0012,\ f(z)=1,\ M=200,\ d=400.$ The theoretical exponent is $-2(2\alpha+2\beta-1)/(2\alpha-2\beta+1)=-0.4$, which matches the experiment. Right: parameters $(\alpha,\beta)=(1.0,0.2),\ \gamma_0=0.0006,\ f(z)=1,\ M=400,\ d=1600.$ The theoretical exponent is $-2(2\alpha+2\beta-1)/(2\alpha-2\beta+1)=-1.077,$ again consistent with the experiment.

C.2 EXPERIMENT FOR ALIGNED DRIFT

In Figure 7, we examine the exponent of the $\mathcal{D}^{\mathrm{sign}}_{\mathrm{al}}(M,N,\gamma_0)$ term,

$$\left(M^{\min(\alpha,0.5)}\,\gamma_0\,N\right)^{-\frac{2(2\alpha+2\beta-1)}{2\alpha-2\beta+1}},$$

of signSGD. For the Phase Aa, the $\mathcal{D}^{\mathrm{sign}}_{\mathrm{al}}(M,N,\gamma_0)$ term dominates in the early iterations over a sufficient interval, allowing us to evaluate the exponent by line fitting on a log-log plot. The experimental results align well with the theoretical formula $-\frac{2(2\alpha+2\beta-1)}{2\alpha-2\beta+1}$.

C.3 VALIDATION OF THE TABLE 1

In Figure 8 to 12 validates the exponent in Table 1 for various (α, β) . On the left plots, we draw multiple curves with different model size M while setting the learning rate as $\gamma_0 = M^{-e^*}$. Then the lower envelope becomes the compute-optimal curve, and by measuring the slope in a log-log plot, we can validate the compute-optimal loss exponent in the Table 1. On the right plots, we draw

the optimal model size at each flops. Here, the optimal model size is the model size of the curve that meets the lower envelope at that flop. By measuring the slope in a log-log plot, we can validate the optimal model size exponent in the Table 1. Note that we use a similar experimental setting to Paquette et al. (2024). In most cases, the error between the measured exponent and the theoretical exponent was less than 0.04, and the error was less than 0.06 even for the worst case. This error lies within the error margins reported in prior works (Paquette et al., 2024; Ferbach et al., 2025).

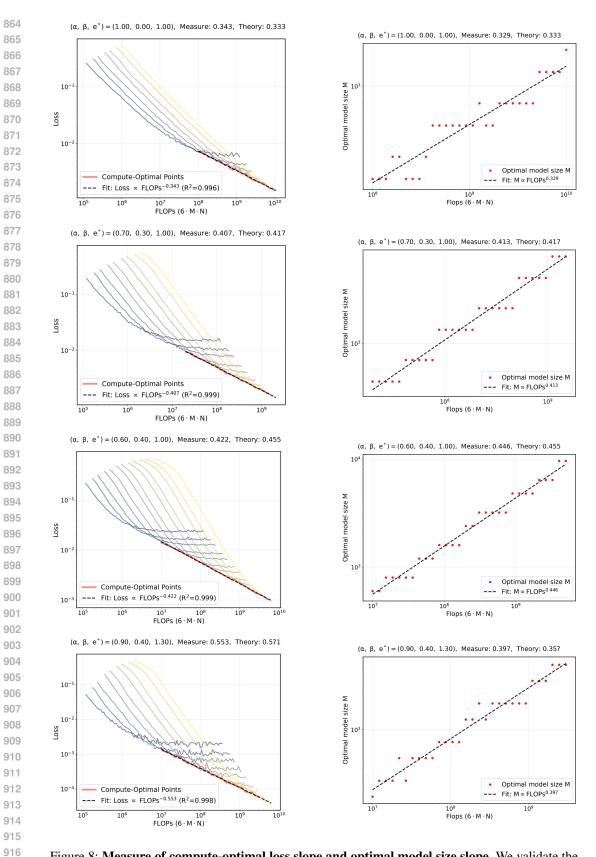


Figure 8: Measure of compute-optimal loss slope and optimal model size slope. We validate the exponent of $R\left(M^\star, \frac{\mathfrak{f}}{M^\star}, \gamma_0^\star\right)$ and M^\star with respect to \mathfrak{f} in the Table 1. The left plot shows the compute-optimal loss with respect to flops 6MN. The right plot shows the optimal model size with respect to flops 6MN. Each plot includes the measured slope and the theoretical slope from the Table 1.

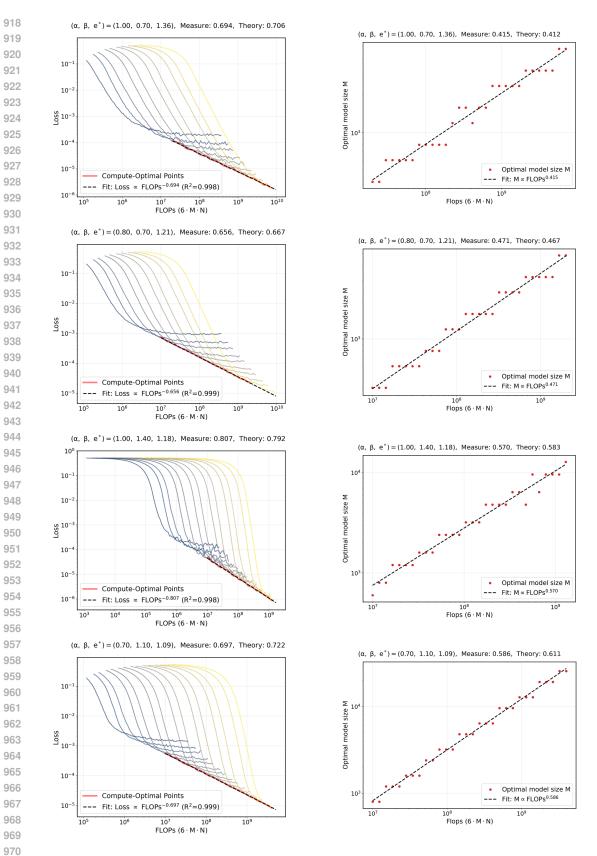


Figure 9: Measure of compute-optimal loss slope and optimal model size slope. We validate the exponent of $R\left(M^\star, \frac{\mathfrak{f}}{M^\star}, \gamma_0^\star\right)$ and M^\star with respect to \mathfrak{f} in the Table 1. The left plot shows the compute-optimal loss with respect to flops 6MN. The right plot shows the optimal model size with respect to flops 6MN.

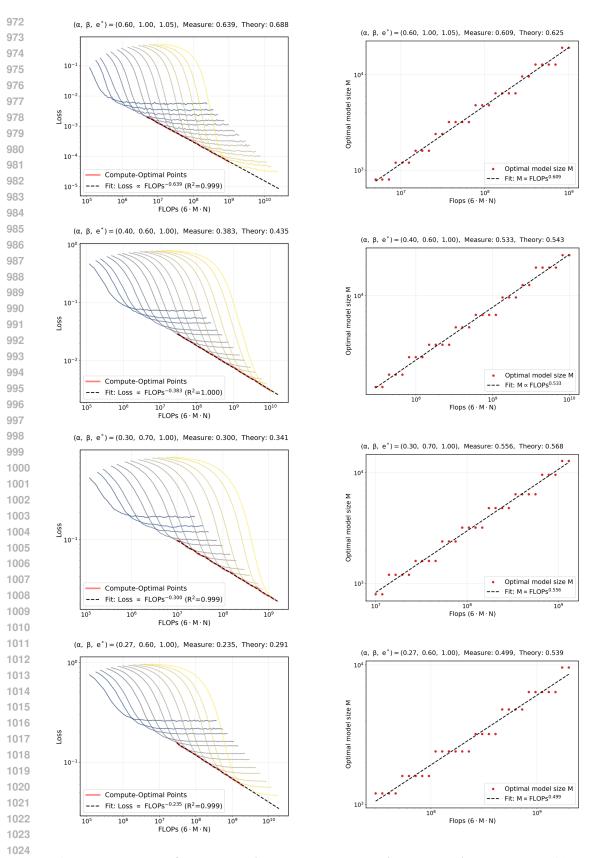


Figure 10: Measure of compute-optimal loss slope and optimal model size slope. We validate the exponent of $R\left(M^\star, \frac{\mathfrak{f}}{M^\star}, \gamma_0^\star\right)$ and M^\star with respect to \mathfrak{f} in the Table 1. The left plot shows the compute-optimal loss with respect to flops 6MN. The right plot shows the optimal model size with respect to flops 6MN.

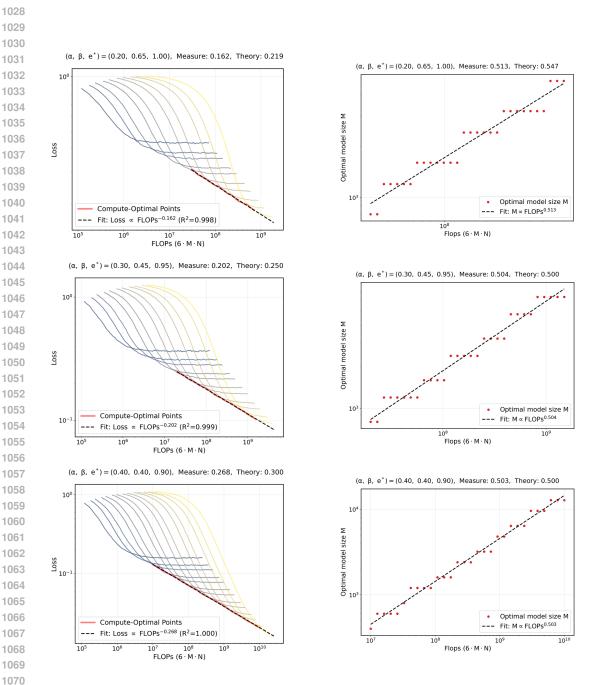


Figure 11: Measure of compute-optimal loss slope and optimal model size slope. We validate the exponent of $R\left(M^\star,\frac{\mathfrak{f}}{M^\star},\gamma_0^\star\right)$ and M^\star with respect to \mathfrak{f} in the Table 1. The left plot shows the compute-optimal loss with respect to flops 6MN. The right plot shows the optimal model size with respect to flops 6MN. Each plot includes the measured slope and the theoretical slope from the Table 1.

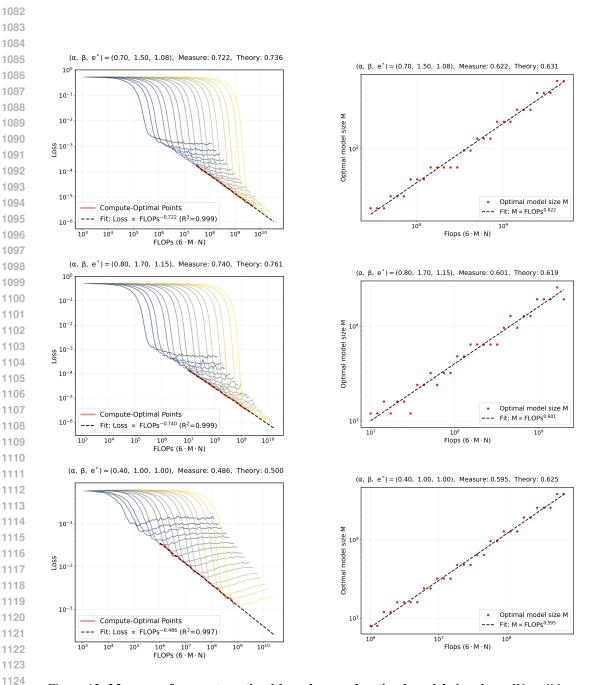


Figure 12: Measure of compute-optimal loss slope and optimal model size slope. We validate the exponent of $R\left(M^\star,\frac{\mathfrak{f}}{M^\star},\gamma_0^\star\right)$ and M^\star with respect to \mathfrak{f} in the Table 1. The left plot shows the compute-optimal loss with respect to flops 6MN. The right plot shows the optimal model size with respect to flops 6MN. Each plot includes the measured slope and the theoretical slope from the Table 1.

D DERIVATION OF THE SCALING LAW FORMULA $R(M, N, \gamma_0)$

In this section, our goal is to derive the scaling law formula (12) of $R(M, N, \gamma_0)$.

On the area $\alpha < 0.5$ or $\beta < 0.5$ with $-\alpha + 0.5 < \beta < \alpha + 0.5$, $\mathcal{D}_{\mathrm{dis}}^{\mathrm{sign}}(M, N, \gamma_0)$ term is smaller than at least one of the other three terms. So it is enough to show

$$R(M,N,\gamma_0) \approx \underbrace{M^{-2\alpha+\max(0,\,1-2\beta)}}_{=:\mathcal{A}(M)} + \underbrace{\left(M^{\min(\alpha,0.5)}N\gamma_0\right)^{-\frac{2(2\alpha+2\beta-1)}{2\alpha-2\beta+1}}}_{=:\mathcal{D}_{\text{al}}^{\text{sign}}(M,N,\gamma_0)} + \underbrace{\gamma_0^2 M^{2-\min(1,2\alpha)}}_{=:\mathcal{N}^{\text{sign}}(M,\gamma_0)}.$$

for that area.

For the area $\alpha > 0.5$ and $\beta > 0.5$ with $-\alpha + 0.5 < \beta < \alpha + 0.5$, as all four terms are dominant, we will prove

$$\begin{split} R(M,N,\gamma_0) \; \approx \; \underbrace{M^{-2\alpha+\max(0,\,1-2\beta)}}_{=:\mathcal{A}(M)} + \underbrace{\left(M^{\min(\alpha,0.5)}N\gamma_0\right)^{-\frac{2(2\alpha+2\beta-1)}{2\alpha-2\beta+1}}}_{=:\mathcal{D}^{\mathrm{sign}}_{\mathrm{all}}(M,N,\gamma_0)} \\ + \underbrace{M^{-\frac{6\alpha-1}{2\alpha+1}}(N\gamma_0)^{-\frac{2(2\alpha-1)}{2\alpha+1}}}_{=:\mathcal{D}^{\mathrm{sign}}_{\mathrm{dis}}(M,N,\gamma_0)} + \underbrace{\gamma_0^2\,M^{2-\min(1,2\alpha)}}_{=:\mathcal{N}^{\mathrm{sign}}(M,\gamma_0)}. \end{split}$$

D.1 ONE-STEP UPDATE FORMULA OF SIGNSGD

Xiao et al. (2024) approximate the signSGD trajectory using SDE and ODE techniques. Their proof relies on a spectral lower bound assumption of the covariance matrix, so their results are not directly applicable to our setting.

For a quadratic function q, by Taylor's theorem, we have

$$\mathbb{E}[q(\boldsymbol{\theta}_{k+1}) - q(\boldsymbol{\theta}_k) \,|\, \mathcal{F}_k] = \mathbb{E}[\langle \nabla q(\boldsymbol{\theta}_k), \boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}_k \rangle \,|\, \mathcal{F}_k] + \frac{1}{2} \,\mathbb{E}\big[\langle \nabla^2 q, (\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}_k)^{\otimes 2} \rangle \,\big|\, \mathcal{F}_k\big],$$
 where $\mathcal{F}_k = \sigma(\boldsymbol{S}, \boldsymbol{\theta}_0, \dots, \boldsymbol{\theta}_k)$. Since

$$\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}_k = -\gamma_k \operatorname{sign}(\langle \boldsymbol{S}\boldsymbol{x}_k, \boldsymbol{\theta}_k \rangle - y_k) \operatorname{sign}(\boldsymbol{S}\boldsymbol{x}_k),$$

We can expand the two terms using sign-Gaussian identities.

Gradient term.

$$\mathbb{E}[\langle \nabla q(\boldsymbol{\theta}_{k}), \boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}_{k} \rangle | \mathcal{F}_{k}]$$

$$= -\gamma_{k} \langle \nabla q(\boldsymbol{\theta}_{k}), \mathbb{E}[\operatorname{sign}(\boldsymbol{S}\boldsymbol{x}_{k}) \operatorname{sign}(\langle \boldsymbol{x}_{k}, \boldsymbol{S}^{\mathsf{T}}\boldsymbol{\theta}_{k} - \boldsymbol{w}^{*} \rangle) | \mathcal{F}_{k}] \rangle$$

$$= -\gamma_{k} \langle \nabla q(\boldsymbol{\theta}_{k}), \frac{2}{\pi} \arcsin\left(\frac{\operatorname{diag}(\boldsymbol{S}\boldsymbol{H}\boldsymbol{S}^{\mathsf{T}})^{-1/2} \boldsymbol{S}\boldsymbol{H} (\boldsymbol{S}^{\mathsf{T}}\boldsymbol{\theta}_{k} - \boldsymbol{w}^{*})}{\sqrt{(\boldsymbol{S}^{\mathsf{T}}\boldsymbol{\theta}_{k} - \boldsymbol{w}^{*})^{\mathsf{T}} \boldsymbol{H} (\boldsymbol{S}^{\mathsf{T}}\boldsymbol{\theta}_{k} - \boldsymbol{w}^{*})}}\right) \rangle$$

$$= -\gamma_{k} \langle \nabla q(\boldsymbol{\theta}_{k}), \frac{2}{\pi} \arcsin\left(\frac{\operatorname{diag}(\boldsymbol{K})^{-1/2} \boldsymbol{K} (\boldsymbol{\theta}_{k} - \boldsymbol{\theta}^{*})}{\|\boldsymbol{H}^{1/2}(\boldsymbol{S}^{\mathsf{T}}\boldsymbol{\theta}_{k} - \boldsymbol{w}^{*})\|}\right) \rangle,$$

where $K = SHS^{\mathsf{T}}$.

Quadratic term.

$$\mathbb{E}\left[\left\langle \nabla^{2}q, \left(\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}_{k}\right)^{\otimes 2}\right\rangle \middle| \mathcal{F}_{k}\right]$$

$$= \gamma_{k}^{2} \left\langle \nabla^{2}q, \mathbb{E}\left[\left(\operatorname{sign}(\boldsymbol{S}\boldsymbol{x}_{k}) \operatorname{sign}(\left\langle \boldsymbol{x}_{k}, \boldsymbol{S}^{\mathsf{T}}\boldsymbol{\theta}_{k} - \boldsymbol{w}^{*}\right\rangle)\right)^{\otimes 2} \middle| \mathcal{F}_{k}\right]\right\rangle$$

$$= \gamma_{k}^{2} \left\langle \nabla^{2}q, \mathbb{E}\left[\left(\operatorname{sign}(\boldsymbol{S}\boldsymbol{x}_{k})\right)^{\otimes 2} \middle| \mathcal{F}_{k}\right]\right\rangle$$

$$= \gamma_{k}^{2} \left\langle \nabla^{2}q, \frac{2}{\pi} \arcsin\left(\operatorname{diag}(\boldsymbol{S}\boldsymbol{H}\boldsymbol{S}^{\mathsf{T}}\right)^{-1/2} \boldsymbol{S}\boldsymbol{H}\boldsymbol{S}^{\mathsf{T}} \operatorname{diag}(\boldsymbol{S}\boldsymbol{H}\boldsymbol{S}^{\mathsf{T}}\right)^{-1/2}\right)\right\rangle$$

$$= \gamma_{k}^{2} \left\langle \nabla^{2}q, \frac{2}{\pi} \arcsin\left(\operatorname{diag}(\boldsymbol{K})^{-1/2} \boldsymbol{K} \operatorname{diag}(\boldsymbol{K})^{-1/2}\right)\right\rangle.$$

One-step update formula. Substituting the gradient and quadratic terms yields the desired one-step update formula for signSGD.

 $\mathbb{E}[q(\boldsymbol{\theta}_{k+1}) - q(\boldsymbol{\theta}_k) \mid \mathcal{F}_k] = -\frac{2\gamma_k}{\pi} \left\langle \nabla q(\boldsymbol{\theta}_k), \arcsin\left(\frac{\overline{K}(\boldsymbol{\theta}_k - \boldsymbol{\theta}^*)}{\sqrt{L(k)}}\right) \right\rangle + \frac{\gamma_k^2}{\pi} \left\langle \nabla^2 q, K_\sigma \right\rangle.$

Let $\lambda_i(\overline{K})$, u_i , and w_i denote the eigenvalue, right eigenvector, and left eigenvector of \overline{K} , respectively. Then $\overline{K} = \sum_{i=1}^M \lambda_i(\overline{K}) u_i \otimes w_i$ and $I = \sum_{i=1}^M u_i \otimes w_i$.

Define

$$r_i(k) = (\boldsymbol{\theta}_k - \boldsymbol{\theta}^*)^\mathsf{T} (\boldsymbol{K} \boldsymbol{u}_i \otimes \boldsymbol{w}_i) (\boldsymbol{\theta}_k - \boldsymbol{\theta}^*).$$

The risk decomposes as

$$L(k) = \left\| \boldsymbol{H}^{1/2} \boldsymbol{S}^{\mathsf{T}} (\boldsymbol{\theta}_k - \boldsymbol{\theta}^*) \right\|^2 + \left\| \boldsymbol{H}^{1/2} \boldsymbol{w}_{\perp} \right\|^2 = (\boldsymbol{\theta}_k - \boldsymbol{\theta}^*)^{\mathsf{T}} \boldsymbol{K} (\boldsymbol{\theta}_k - \boldsymbol{\theta}^*) + \| \boldsymbol{H}^{1/2} \boldsymbol{w}_{\perp} \|^2 = \sum_{i=1}^d r_i(k) + \| \boldsymbol{H}^{1/2} \boldsymbol{w}_{\perp} \|^2.$$

We now apply the one-step update formula to $r_i(k)$. Note that

$$\nabla r_i(k) = K u_i \langle w_i, \theta_k - \theta^* \rangle + w_i \langle K u_i, \theta_k - \theta^* \rangle, \qquad \nabla^2 r_i = K u_i w_i^{\mathsf{T}} + w_i u_i^{\mathsf{T}} K^{\mathsf{T}}.$$

Approximating $\arcsin(x) \approx x$ and using $K^{\mathsf{T}} = K$ together with $K^{\mathsf{T}} \overline{K} = \overline{K}^{\mathsf{T}} K^{\mathsf{T}}$, we obtain

$$\mathbb{E}[r_i(k+1) - r_i(k) | \mathcal{F}_k] \approx -\frac{2\gamma_k}{\pi} \left(\langle \boldsymbol{w}_i, \boldsymbol{\theta}_k - \boldsymbol{\theta}^* \rangle \left\langle \boldsymbol{K} \boldsymbol{u}_i, \frac{\overline{\boldsymbol{K}} (\boldsymbol{\theta}_k - \boldsymbol{\theta}^*)}{\sqrt{L(k)}} \right\rangle + \langle \boldsymbol{K} \boldsymbol{u}_i, \boldsymbol{\theta}_k - \boldsymbol{\theta}^* \rangle \left\langle \boldsymbol{w}_i, \frac{\overline{\boldsymbol{K}} (\boldsymbol{\theta}_k - \boldsymbol{\theta}^*)}{\sqrt{L(k)}} \right\rangle \right)$$

$$+ \frac{2\gamma_k^2}{\pi} \boldsymbol{w}_i^{\mathsf{T}} \boldsymbol{K}_{\sigma} \boldsymbol{K} \boldsymbol{u}_i$$

$$= -\frac{4\gamma_k}{\pi \sqrt{L(k)}} \lambda_i(\overline{\boldsymbol{K}}) r_i(k) + \frac{2\gamma_k^2}{\pi} \boldsymbol{w}_i^{\mathsf{T}} \boldsymbol{K}_{\sigma} \boldsymbol{K} \boldsymbol{u}_i.$$

It is possible to replace the linear approximation $\arcsin(x) \approx x$ by an inequality, and the main results of our paper remain unchanged. We explain it in Appendix I.2. Hence,

$$\mathbb{E}[r_i(k+1) - r_i(k) \,|\, \mathcal{F}_k] \approx -\frac{4\gamma_k}{\pi\sqrt{L(k)}} \,\lambda_i(\overline{\mathbf{K}}) \,r_i(k) + \frac{2\gamma_k^2}{\pi} \,\mathbf{w}_i^\mathsf{T} \mathbf{K}_\sigma \mathbf{K} \mathbf{u}_i.$$

D.2 ODE APPROXIMATION AND IMPLICIT INTEGRAL EQUATION OF SIGNSGD

Let the learning rate be $\gamma_k = \gamma_0 f(k)$. Define $V_i = \boldsymbol{w}_i^\mathsf{T} \boldsymbol{K}_\sigma \boldsymbol{K} \boldsymbol{u}_i$, then our one-step update formula becomes

$$\mathbb{E}[r_i(k+1) - r_i(k) \mid \mathcal{F}_k] = -\frac{4\gamma_k}{\pi\sqrt{L(k)}} \lambda_i(\overline{\mathbf{K}}) r_i(k) + \frac{2\gamma_k^2}{\pi} V_i.$$

Dividing by γ_0 gives

$$\mathbb{E}\left[\frac{r_i(k+1) - r_i(k)}{\gamma_0} \middle| \mathcal{F}_k\right] = -\frac{4}{\pi\sqrt{L(k)}} \lambda_i(\overline{\mathbf{K}}) f(k) r_i(k) + \frac{2f(k)^2 \gamma_0}{\pi} V_i.$$

Interpreting γ_0 as the time step, the discrete index k corresponds to continuous time $t=k\gamma_0$. Let $P(t)=L(t/\gamma_0)$ and $p_i(t)=r_i(t/\gamma_0)$. We then obtain the ODE

$$\frac{dp_i}{dt} = -\frac{4}{\pi\sqrt{P(t)}} \lambda_i(\overline{K}) f(t/\gamma_0) p_i(t) + \frac{2f(t/\gamma_0)^2 \gamma_0}{\pi} V_i.$$
 (16)

From this point onward in the analysis, we treat P, p_i , L, and r_i as their continuous extensions, allowing arbitrary positive real inputs.

Integral formulation. Solving the ODE yields

$$p_i(t) = p_i(0) e^{-\frac{4\lambda_i}{\pi} \int_0^t \frac{f(u/\gamma_0)}{\sqrt{P(u)}} du} + \frac{2\gamma_0}{\pi} V_i \int_0^t e^{-\frac{4\lambda_i}{\pi} \int_s^t \frac{f(u/\gamma_0)}{\sqrt{P(u)}} du} f(s/\gamma_0)^2 ds.$$

Since $P(t) = \sum_{i=1}^{M} p_i(t) + \|\boldsymbol{H}^{1/2}\boldsymbol{w}_{\perp}\|^2$, we obtain

$$P(t) = \|\boldsymbol{H}^{1/2}\boldsymbol{w}_{\perp}\|^{2} + \sum_{i=1}^{M} p_{i}(0) e^{-\frac{4\lambda_{i}}{\pi} \int_{0}^{t} \frac{f(u/\gamma_{0})}{\sqrt{P(u)}} du} + \frac{2\gamma_{0}}{\pi} \sum_{i=1}^{M} V_{i} \int_{0}^{t} e^{-\frac{4\lambda_{i}}{\pi} \int_{s}^{t} \frac{f(u/\gamma_{0})}{\sqrt{P(u)}} du} f(s/\gamma_{0})^{2} ds.$$

Integral equation in discrete form. Note that $L(N) = P(N\gamma_0)$. With a change of variables, we obtain

$$L(N) = \|\boldsymbol{H}^{1/2}\boldsymbol{w}_{\perp}\|^{2} + \sum_{i=1}^{M} r_{i}(0) e^{-\frac{4\lambda_{i}\gamma_{0}}{\pi} \int_{0}^{N} \frac{f(u)}{\sqrt{L(u)}} du} + \frac{2\gamma_{0}^{2}}{\pi} \sum_{i=1}^{M} V_{i} \int_{0}^{N} e^{-\frac{4\lambda_{i}\gamma_{0}}{\pi} \int_{z}^{N} \frac{f(u)}{\sqrt{L(u)}} du} f(z)^{2} dz.$$

$$(17)$$

Drift and noise decomposition. Define

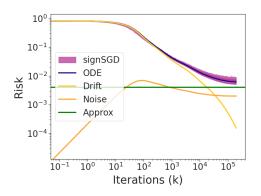
$$L^{\text{drift}}(N) = \sum_{i=1}^{M} r_i(0) \, e^{-\frac{4\lambda_i \gamma_0}{\pi} \int_0^N \frac{f(u)}{\sqrt{L(u)}} \, du}, \quad L^{\text{noise}}(N) = \frac{2\gamma_0^2}{\pi} \sum_{i=1}^{M} V_i \int_0^N e^{-\frac{4\lambda_i \gamma_0}{\pi} \int_z^N \frac{f(u)}{\sqrt{L(u)}} \, du} f(z)^2 \, dz. \tag{18}$$

Then

$$L(N) = \|\mathbf{H}^{1/2}\mathbf{w}_{\perp}\|^{2} + L^{\text{drift}}(N) + L^{\text{noise}}(N), \tag{19}$$

and we will analyze $\|\boldsymbol{H}^{1/2}\boldsymbol{w}_{\perp}\|^2 + L^{\text{drift}}(N)$ and $L^{\text{noise}}(N)$ separately.

Figure 13 show dynamics of three terms $\|\boldsymbol{H}^{1/2}\boldsymbol{w}_{\perp}\|^2$, $L^{\text{drift}}(N)$, $L^{\text{noise}}(N)$ referring each as Approx, Drift, Noise. The right plot in Figure 13 validates the equality in (19).



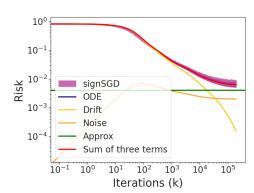


Figure 13: **Dynamics of Drift and Noise.** Left: the purple curve is the 80% confidence interval of the true signSGD trajectory, while the blue curve is the numerical ODE solution. The yellow, orange, and green curves correspond to the approximation, drift, and noise terms in (19). Right: the red curve shows the sum of these three terms, matching both the true trajectory and the ODE solution. Parameters: $\alpha = 1.0$, $\beta = 0$, $\gamma_0 = 0.003$, f(z) = 1, M = 200, d = 800.

D.2.1 TRANSFORMATION OF THE DRIFT TERM AND APPROXIMATION ERROR

Let

$$Q(z) \ = \ \frac{4\gamma_0}{\pi} \int_0^z \frac{f(u)}{\sqrt{L(u)}} \, du, \qquad \overline{\boldsymbol{K}}_1 \ = \ \boldsymbol{H}^{1/2} \boldsymbol{S}^\mathsf{T} \operatorname{diag}(\boldsymbol{S} \boldsymbol{H} \boldsymbol{S}^\mathsf{T})^{-1/2} \boldsymbol{S} \boldsymbol{H}^{1/2}.$$

Then

$$\boldsymbol{K}\,\overline{\boldsymbol{K}}^{p}\ =\ \boldsymbol{S}\boldsymbol{H}^{1/2}\overline{\boldsymbol{K}}_{1}^{p}\boldsymbol{H}^{1/2}\boldsymbol{S}^{\mathsf{T}}$$

1297 holds.

1298 Define

$$oldsymbol{A} = oldsymbol{H}^{1/2} e^{-\overline{oldsymbol{K}}_1 Q(N)} oldsymbol{H}^{1/2}, \qquad oldsymbol{u} = oldsymbol{S}^\mathsf{T} oldsymbol{ heta}_0 - oldsymbol{S}^\mathsf{T} oldsymbol{ heta}^* - oldsymbol{w}_\perp = oldsymbol{S}^\mathsf{T} oldsymbol{ heta}_0 - oldsymbol{w}^*.$$

From $\mathbf{SHw}_{\perp} = 0$ we get

$$\overline{\boldsymbol{K}}_1(\boldsymbol{H}^{1/2}\boldsymbol{w}_\perp) = 0,$$

and this implies

$$e^{-\overline{K}_1 Q(N)}(H^{1/2} w_{\perp}) = e^0(H^{1/2} w_{\perp}) = H^{1/2} w_{\perp}.$$

Thus,

$$oldsymbol{A}oldsymbol{w}_{oldsymbol{\perp}} = oldsymbol{H}oldsymbol{w}_{oldsymbol{\perp}}, \qquad oldsymbol{w}_{oldsymbol{\perp}}^{\mathsf{T}}oldsymbol{A}oldsymbol{w}_{oldsymbol{\perp}} = oldsymbol{w}_{oldsymbol{\perp}}^{\mathsf{T}}oldsymbol{H}oldsymbol{w}_{oldsymbol{\perp}},$$

and

$$oldsymbol{u}^\mathsf{T} A oldsymbol{w}_\perp = oldsymbol{u}^\mathsf{T} H oldsymbol{w}_\perp = oldsymbol{u}^\mathsf{T} H oldsymbol{w}_\perp - oldsymbol{w}_\perp^\mathsf{T} H oldsymbol{w}_\perp = -oldsymbol{w}_\perp^\mathsf{T} H oldsymbol{w}_\perp.$$

Using these identities, we can convert the drift term as follows:

$$\begin{split} L^{\text{drift}}(N) &= \sum_{i=1}^{M} r_i(0) \cdot e^{-\lambda_i(\overline{K})Q(N)} \\ &= \sum_{i=1}^{M} (\boldsymbol{\theta}_0 - \boldsymbol{\theta}^*)^{\mathsf{T}} (K\boldsymbol{u}_i \otimes \boldsymbol{w}_i) (\boldsymbol{\theta}_0 - \boldsymbol{\theta}^*) \cdot e^{-\lambda_i(\overline{K})Q(N)} \\ &= \sum_{i=1}^{M} (\boldsymbol{\theta}_0 - \boldsymbol{\theta}^*)^{\mathsf{T}} \Big((K\boldsymbol{u}_i \otimes \boldsymbol{w}_i) \cdot e^{-\lambda_i(\overline{K})Q(N)} \Big) (\boldsymbol{\theta}_0 - \boldsymbol{\theta}^*) \\ &= (\boldsymbol{\theta}_0 - \boldsymbol{\theta}^*)^{\mathsf{T}} K e^{-\overline{K}Q(N)} (\boldsymbol{\theta}_0 - \boldsymbol{\theta}^*) \\ &= (\boldsymbol{\theta}_0 - \boldsymbol{\theta}^*)^{\mathsf{T}} S H^{1/2} \left(H^{1/2} S e^{-\overline{K}Q(N)} \right) (\boldsymbol{\theta}_0 - \boldsymbol{\theta}^*) \\ &= (\boldsymbol{\theta}_0 - \boldsymbol{\theta}^*)^{\mathsf{T}} S H^{1/2} \left(e^{-\overline{K}_1 Q(N)} H^{1/2} S \right) (\boldsymbol{\theta}_0 - \boldsymbol{\theta}^*) \\ &= (\boldsymbol{u} + \boldsymbol{w}_\perp)^{\mathsf{T}} A (\boldsymbol{u} + \boldsymbol{w}_\perp) \\ &= u^{\mathsf{T}} A \boldsymbol{u} + u^{\mathsf{T}} A \boldsymbol{w}_\perp + w_\perp^{\mathsf{T}} A \boldsymbol{u} + w_\perp^{\mathsf{T}} A \boldsymbol{w}_\perp \\ &= u^{\mathsf{T}} H^{1/2} e^{-\overline{K}_1 Q(N)} H^{1/2} \boldsymbol{u} - w_\perp^{\mathsf{T}} H \boldsymbol{w}_\perp - w_\perp^{\mathsf{T}} H \boldsymbol{w}_\perp + w_\perp^{\mathsf{T}} H \boldsymbol{w}_\perp \\ &= u^{\mathsf{T}} H^{1/2} e^{-\overline{K}_1 Q(N)} H^{1/2} \boldsymbol{u} - \| H^{1/2} \boldsymbol{w}_\perp \|^2. \end{split}$$

Drift term plus approximation error. Adding the approximation error gives

$$\begin{split} L^{\text{drift}}(N) + \|\boldsymbol{H}^{1/2}\boldsymbol{w}_{\perp}\|^2 &= \boldsymbol{u}^{\mathsf{T}}\boldsymbol{H}^{1/2}e^{-\overline{\boldsymbol{K}}_1Q(N)}\boldsymbol{H}^{1/2}\boldsymbol{u} \\ &= \Big\langle e^{-\overline{\boldsymbol{K}}_1Q(N)}, \ \big(\boldsymbol{H}^{1/2}(\boldsymbol{S}^{\mathsf{T}}\boldsymbol{\theta}_0 - \boldsymbol{w}^*)\big)^{\otimes 2} \Big\rangle. \end{split}$$

Also we assume $\theta_0 = 0$, then

$$\langle e^{-\overline{K}_1 Q(N)}, (H^{1/2} (S^\mathsf{T} \boldsymbol{\theta}_0 - \boldsymbol{w}^*))^{\otimes 2} \rangle = \langle e^{-\overline{K}_1 Q(N)}, (H^{1/2} \boldsymbol{w}^*)^{\otimes 2} \rangle.$$

In the next subsection, we will describe how to apply a deterministic approximation, similar to Paquette et al. (2024), to the following term:

$$\mathcal{H} := \langle e^{-\overline{K}_1 Q(N)}, v^{\otimes 2} \rangle,$$

where $oldsymbol{v} := oldsymbol{H}^{1/2} oldsymbol{w}^* \in \mathbb{R}^d$.

D.2.2 DETERMINISTIC APPROXIMATION

Note that we assume $d \geq rM$ for some r > 1, and let $d/M \to (1, \infty]$ as $d, M \to \infty$ when $2\alpha > 1$, and $d/M \to (1, \infty)$ when $2\alpha < 1$. In our setup, $S \in \mathbb{R}^{M \times d}$ have i.i.d. $\mathcal{N}(0, 1/M)$ entries, and we will write the kth column of S^T as $\frac{1}{\sqrt{M}} s_k \in \mathbb{R}^d$; columns are independent.

Define

$$\mathbf{y}_k := \mathbf{H}^{1/2} \mathbf{s}_k \in \mathbb{R}^d, \qquad a_k := \frac{1}{\sqrt{\frac{1}{M} \mathbf{y}_k^\mathsf{T} \mathbf{y}_k}} = \frac{\sqrt{M}}{\sqrt{\mathbf{s}_k^\mathsf{T} \mathbf{H} \mathbf{s}_k}} > 0.$$

The unnormalized baseline and the column–normalized matrices are

$$\widehat{\boldsymbol{K}} := \boldsymbol{H}^{1/2} \boldsymbol{S}^\mathsf{T} \boldsymbol{S} \boldsymbol{H}^{1/2} = \frac{1}{M} \sum_{k=1}^M \boldsymbol{y}_k \boldsymbol{y}_k^\mathsf{T}, \qquad \overline{\boldsymbol{K}}_1 := \boldsymbol{H}^{1/2} \boldsymbol{S}^\mathsf{T} \operatorname{diag} (\boldsymbol{S} \boldsymbol{H} \boldsymbol{S}^\mathsf{T})^{-1/2} \boldsymbol{S} \boldsymbol{H}^{1/2} = \frac{1}{M} \sum_{k=1}^M a_k \, \boldsymbol{y}_k \boldsymbol{y}_k^\mathsf{T}.$$

For $z \in \mathbb{C}^+ := \{z : \Im z > 0\}$, define the resolvents

$$m{L}(z) := (\overline{m{K}}_1 - z m{I})^{-1}, \qquad m{R}^{(k)}(z) := \left(rac{1}{M} \sum_{\ell \neq k} a_\ell m{y}_\ell m{y}_\ell^\mathsf{T} - z m{I}
ight)^{-1}.$$

Note that

$$y_k B y_k \approx \text{Tr}(HB)$$

for matrix B. In particular,

$$\boldsymbol{y}_k^{\mathsf{T}} \boldsymbol{y}_k \approx \operatorname{Tr} \boldsymbol{H}, \qquad a_k \approx \frac{\sqrt{M}}{\sqrt{\operatorname{Tr} \boldsymbol{H}}}.$$

Also note that

$$a_k \, \boldsymbol{y}_k^\mathsf{T} \boldsymbol{R} \, \boldsymbol{R}^{(k)} \, \boldsymbol{y}_k \approx \frac{\sqrt{M}}{\sqrt{\mathrm{Tr} \, \boldsymbol{H}}} \cdot \mathrm{Tr}(\boldsymbol{H} \, \boldsymbol{R} \, \boldsymbol{R}^{(k)}),$$

and

$$a_k \boldsymbol{y}_k^\mathsf{T} \boldsymbol{R}^{(k)} \boldsymbol{y}_k \approx \frac{\sqrt{M}}{\sqrt{\operatorname{Tr} \boldsymbol{H}}} \operatorname{Tr}(\boldsymbol{H} \, \boldsymbol{R}^{(k)}).$$

By the Sherman–Morrison expansion,

$$\boldsymbol{R} = \boldsymbol{R}^{(k)} - \frac{M^{-1}a_k \, \boldsymbol{R}^{(k)} \boldsymbol{y}_k \boldsymbol{y}_k^\mathsf{T} \boldsymbol{R}^{(k)}}{1 + M^{-1}a_k \, \boldsymbol{y}_k^\mathsf{T} \boldsymbol{R}^{(k)} \boldsymbol{y}_k}.$$

Multiplying on the left by R and sandwiching with $y_k^{\mathsf{T}}(\cdot)y_k$, we get

$$a_k \boldsymbol{y}_k^\mathsf{T} \boldsymbol{R} \, \boldsymbol{R} \, \boldsymbol{y}_k = a_k \boldsymbol{y}_k^\mathsf{T} \boldsymbol{R} \, \boldsymbol{R}^{(k)} \boldsymbol{y}_k - \frac{M^{-1} a_k \boldsymbol{y}_k^\mathsf{T} \boldsymbol{R} \, \boldsymbol{R}^{(k)} \boldsymbol{y}_k \cdot a_k \boldsymbol{y}_k^\mathsf{T} \boldsymbol{R}^{(k)} \boldsymbol{y}_k}{1 + M^{-1} a_k \, \boldsymbol{y}_k^\mathsf{T} \boldsymbol{R}^{(k)} \boldsymbol{y}_k}.$$

Now we will replace terms on the right side by

$$a_k \boldsymbol{y}_k^\mathsf{T} \boldsymbol{R} \, \boldsymbol{R}^{(k)} \boldsymbol{y}_k \approx \frac{\sqrt{M}}{\sqrt{\mathrm{Tr} \, \boldsymbol{H}}} \, \mathrm{Tr}(\boldsymbol{H} \, \boldsymbol{R} \, \boldsymbol{R}^{(k)}),$$

and

$$a_k \boldsymbol{y}_k^{\mathsf{T}} \boldsymbol{R}^{(k)} \boldsymbol{y}_k \approx \frac{\sqrt{M}}{\sqrt{\operatorname{Tr} \boldsymbol{H}}} \operatorname{Tr}(\boldsymbol{H} \, \boldsymbol{R}^{(k)}).$$

1401 Thus

$$a_k \boldsymbol{y}_k^\mathsf{T} \boldsymbol{R} \, \boldsymbol{R} \, \boldsymbol{y}_k pprox rac{\sqrt{M}}{\sqrt{\text{Tr} \, \boldsymbol{H}}} \, ext{Tr}(\boldsymbol{H} \, \boldsymbol{R} \, \boldsymbol{R}^{(k)})}{1 + M^{-1} rac{\sqrt{M}}{\sqrt{\text{Tr} \, \boldsymbol{H}}} \, ext{Tr}(\boldsymbol{H} \, \boldsymbol{R}^{(k)})}.$$

Replacing $\mathbf{R}^{(k)}$ by \mathbf{R} and averaging over k, we obtain

$$\frac{1}{M} \sum_{k=1}^{M} a_k \boldsymbol{y}_k^\mathsf{T} \boldsymbol{R} \boldsymbol{R} \boldsymbol{y}_k \approx \frac{p_d \operatorname{Tr}(\boldsymbol{H} \boldsymbol{R} \boldsymbol{R})}{1 + M^{-1} p_d \operatorname{Tr}(\boldsymbol{H} \boldsymbol{R})}, \qquad p_d := \frac{\sqrt{M}}{\sqrt{\operatorname{Tr} \boldsymbol{H}}}.$$

1409 It implies

$$\operatorname{Tr}(\boldsymbol{R}(\boldsymbol{R}^{-1}+z\boldsymbol{I})\boldsymbol{R}) \approx \frac{p_d \operatorname{Tr}(\boldsymbol{R} \boldsymbol{H} \boldsymbol{R})}{1 + M^{-1} p_d \operatorname{Tr}(\boldsymbol{H} \boldsymbol{R})}.$$

1413 This implies

$$L(z)^{-1} + zI \approx \frac{p_d}{1 + M^{-1}p_d \operatorname{Tr}(\boldsymbol{H} L(z))} \boldsymbol{H}.$$

1416 Let

$$m(z/p_d) = \frac{1}{1 + M^{-1}p_d \operatorname{Tr}(\boldsymbol{H} \boldsymbol{L}(z))}.$$

1420 Then

$$L(z) \approx (-z\mathbf{I} + p_d m(z/p_d)\mathbf{H})^{-1}.$$

1422 Thus

$$(\overline{\boldsymbol{K}}_1 - z\boldsymbol{I})^{-1} \approx (-z\boldsymbol{I} + p_d m(z/p_d)\boldsymbol{H})^{-1}.$$

1425 And

$$m(z) = \frac{1}{1 + M^{-1}p_d \operatorname{Tr}(\boldsymbol{H} \, \boldsymbol{R}(p_d z))} \approx \frac{1}{1 + M^{-1} \operatorname{Tr}(\boldsymbol{H}(-z\boldsymbol{I} + m(z)\boldsymbol{H})^{-1})}$$

holds. This fixed-point equation is identical to the one in Paquette et al. (2024).

Contour representation. Let $oldsymbol{v}:=oldsymbol{H}^{1/2}oldsymbol{w}^*\in\mathbb{R}^d$ and consider

$$\mathcal{H} := \left\langle e^{-\overline{K}_1 Q(N)}, \, v^{\otimes 2} \right\rangle.$$

For any analytic g on a contour Γ_2 enclosing $\operatorname{Spec}(\overline{K}_1)$,

$$g(\overline{K}_1) = -\frac{1}{2\pi i} \oint_{\Gamma_2} g(z) (\overline{K}_1 - zI)^{-1} dz.$$

We prove

$$c_1 M^{\min(0.5,\alpha)} I \leq \operatorname{diag}(\mathbf{SHS}^{\mathsf{T}})^{-1/2} \leq c_2 M^{\min(0.5,\alpha)} I$$

in Section I.4. It leads to

$$c_1 M^{\min(0.5,\alpha)} \widehat{\mathbf{K}} \preceq \overline{\mathbf{K}} \preceq c_2 M^{\min(0.5,\alpha)} \widehat{\mathbf{K}}.$$

 \overline{K}_1 has eigenvalues scaled by $M^{\min(0.5,\alpha)}$ compared to \widehat{K} excluding constant. Note that $p_d \approx M^{\min(0.5,\alpha)}$. So, there exists a contour Γ_2 enclosing the spectrum of \overline{K}_1 , and its $1/p_d$ -scaled version

 Γ encloses the spectrum of K.

Taking $q(z) = e^{-Q(N)z}$.

$$\begin{split} \mathcal{H} &= -\frac{1}{2\pi i} \oint_{\Gamma_2} e^{-Q(N)z} \left\langle (\overline{\boldsymbol{K}}_1 - z\boldsymbol{I})^{-1}, \, \boldsymbol{v}^{\otimes 2} \right\rangle \, dz \\ &\approx -\frac{1}{2\pi i} \oint_{\Gamma_2} e^{-Q(N)z} \left\langle (-z\boldsymbol{I} + p_d m(z/p_d)\boldsymbol{H})^{-1}, \, \boldsymbol{v}^{\otimes 2} \right\rangle \, dz \\ &= -\frac{1}{2\pi i} \oint_{\Gamma} e^{-p_d Q(N)z} \left\langle (-z\boldsymbol{I} + m(z)\boldsymbol{H})^{-1}, \, \boldsymbol{v}^{\otimes 2} \right\rangle \, dz. \end{split}$$

Let $\mathcal{R}(z) = (-z\mathbf{I} + m(z)\mathbf{H})^{-1}$, then our objective converts to

$$\mathcal{H} \approx -\frac{1}{2\pi i} \oint_{\Gamma} e^{-p_d Q(N)z} \left\langle \mathcal{R}(z), \, \boldsymbol{v}^{\otimes 2} \right\rangle \, dz.$$

D.2.3 FINAL TRANSFORMATION RESULT

Paquette et al. (2024) evaluate the contour integrals with $\mathcal{R}(z)$. When $\alpha < 0.5$ or $\beta < 0.5$, they show

$$-\frac{1}{2\pi i} \oint_{\Gamma} \left(1 - 2\gamma Bz + \gamma^2 B(B+1)z^2 \right)^r \left\langle \mathcal{L}(z), v^{\otimes 2} \right\rangle dz \approx M^{-2\alpha + \max(0, 1-2\beta)} + (2\gamma Br)^{-\frac{2\alpha + 2\beta - 1}{2\alpha}}. \tag{20}$$

When $\alpha > 0.5$ and $\beta > 0.5$, they obtained

$$-\frac{1}{2\pi i} \oint_{\Gamma} \left(1 - 2\gamma B z + \gamma^2 B (B+1) z^2 \right)^r \left\langle \mathcal{L}(z), v^{\otimes 2} \right\rangle dz \approx M^{-2\alpha + \max(0, 1 - 2\beta)}$$

$$+ (2\gamma B r)^{-\frac{2\alpha + 2\beta - 1}{2\alpha}}$$

$$+ M^{-1} (2\gamma B r)^{-2 + \frac{1}{2\alpha}}. \tag{21}$$

For the case $\alpha < 0.5$ or $\beta < 0.5$, applying a similar method to our objective yields

$$-\frac{1}{2\pi i} \oint_{\Gamma} e^{-p_d Q(N)z} \left\langle \mathcal{L}(z), v^{\otimes 2} \right\rangle dz \approx M^{-2\alpha + \max(0, 1 - 2\beta)} + \left(M^{\min(\alpha, 0.5)} Q(N) \right)^{-\frac{2\alpha + 2\beta - 1}{2\alpha}}, \tag{22}$$

with details provided in Appendix I.1. Hence,

$$\left\langle e^{-\overline{K}_1 Q(N)}, (\boldsymbol{H}^{1/2} \boldsymbol{w}^*)^{\otimes 2} \right\rangle \approx M^{-2\alpha + \max(0, 1 - 2\beta)} + \left(M^{\min(\alpha, 0.5)} Q(N) \right)^{-\frac{2\alpha + 2\beta - 1}{2\alpha}}.$$
 (23)

For the case $\alpha > 0.5$ and $\beta > 0.5$, a similar argument gives

$$-\frac{1}{2\pi i} \oint_{\Gamma} e^{-p_d Q(N)z} \left\langle \mathcal{L}(z), v^{\otimes 2} \right\rangle dz \approx M^{-2\alpha + \max(0, 1 - 2\beta)}$$

$$+ \left(M^{\min(\alpha, 0.5)} Q(N) \right)^{-\frac{2\alpha + 2\beta - 1}{2\alpha}}$$

$$+ M^{-1} \left(M^{\min(\alpha, 0.5)} Q(N) \right)^{-1 + \frac{1}{2\alpha}}, \qquad (24)$$

with details in Appendix I.1. Consequently,

$$\left\langle e^{-\overline{K}_1 Q(N)}, (\boldsymbol{H}^{1/2} \boldsymbol{w}^*)^{\otimes 2} \right\rangle \approx M^{-2\alpha + \max(0, 1 - 2\beta)}$$

$$+ \left(M^{\min(\alpha, 0.5)} Q(N) \right)^{-\frac{2\alpha + 2\beta - 1}{2\alpha}}$$

$$+ M^{-1} \left(M^{\min(\alpha, 0.5)} Q(N) \right)^{-1 + \frac{1}{2\alpha}}.$$
 (25)

In summary, we obtain

$$L^{\text{drift}}(N) + \|\boldsymbol{H}^{1/2}\boldsymbol{w}_{\perp}\|^{2} \approx M^{-2\alpha + \max(0, 1 - 2\beta)} + \left(M^{\min(\alpha, 0.5)}Q(N)\right)^{-\frac{2\alpha + 2\beta - 1}{2\alpha}}, \quad (26)$$

for $\alpha < 0.5$ or $\beta < 0.5$, and

$$L^{\text{drift}}(N) + \|\boldsymbol{H}^{1/2}\boldsymbol{w}_{\perp}\|^{2} \approx M^{-2\alpha + \max(0, 1 - 2\beta)} + \left(M^{\min(\alpha, 0.5)} Q(N)\right)^{-\frac{2\alpha + 2\beta - 1}{2\alpha}} + M^{-1} \left(M^{\min(\alpha, 0.5)} Q(N)\right)^{-1 + \frac{1}{2\alpha}}, \tag{27}$$

for $\alpha > 0.5$ and $\beta > 0.5$. Figure 14 shows our transformed result in (26) and (27) based on deterministic approximation matches the true signSGD trajectory up to a constant factor. When interpreting the figure, note that our analysis is asymptotic; hence, discrepancies may appear in the very early iterations.

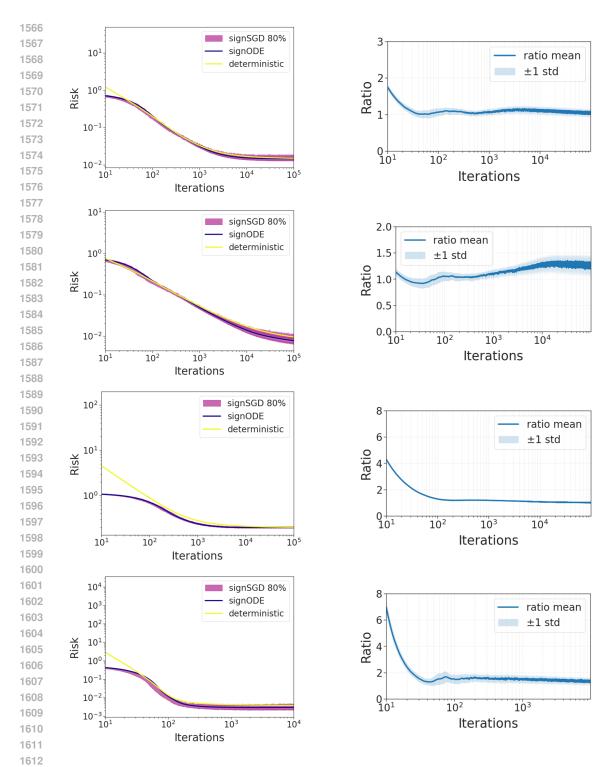


Figure 14: Verification of the deterministic approximation and drift-term transformation. Left: the purple curve denotes the 80% confidence interval of the true signSGD trajectory, the blue curve represents the numerical ODE solution, and the yellow curve corresponds to the deterministic approximation after drift-term transformation in (26) and (27). Deterministic approximation matches the true trajectory up to a constant factor. It should be noted that our analysis is asymptotic, and thus, discrepancies may occur in the very early iterations. Right: the ratio between the approximation and the true trajectory remains bounded by a constant factor, confirming the validity of our approach. Parameters: $(\alpha, \beta) = (0.7, 0.3), (1.0, 0), (0.4, 0.4), (0.7, 1.1)$ from top to bottom, $\gamma_0 = 0.003, f(z) = 1, M = 200, d = 800, 100$ runs.

D.3 Constant Learning Rate: Proxy and Verification for the Case $\alpha < 0.5$ or $\beta < 0.5$ (Phase A)

Throughout this section, we set $f(z) \equiv 1$; hence

$$Q(N) = \frac{4\gamma_0}{\pi} \int_0^N \frac{du}{\sqrt{L(u)}}.$$

Applying the drift/approximation-term transformation to the ODE solution yields the implicit relation

$$L(N) \approx \underbrace{M^{-2\alpha + \max(0, 1 - 2\beta)}}_{\text{approx}} + \underbrace{\left(M^{\min(\alpha, 0.5)} Q(N)\right)^{-\frac{2\alpha + 2\beta - 1}{2\alpha}}}_{\text{drift}}$$
(28)

$$+ \underbrace{\frac{2\gamma_0^2}{\pi} \sum_{i=1}^M V_i \int_0^N \exp\left(-\frac{4\gamma_0}{\pi} \lambda_i(\overline{K}) \int_z^N \frac{du}{\sqrt{L(u)}}\right) dz}_{\text{2.2}}.$$
 (29)

D.3.1 EARLY STAGE (DOMINANCE OF THE DRIFT TERM)

At N=0, the noise integral is 0, the approximation term is independent of N, and the drift term is large and decreases with N. Thus, in the early phase,

$$L(N) \approx \left(M^{\min(\alpha, 0.5)} Q(N)\right)^{-\frac{2\alpha+2\beta-1}{2\alpha}}.$$
 (30)

Since $Q(N) = \frac{4\gamma_0}{\pi} \int_0^N L(u)^{-1/2} du$, (30) is equivalent (up to absolute constants) to

$$L(N)^{-\frac{2\alpha}{2\alpha+2\beta-1}} \approx M^{\min(\alpha, 0.5)} \gamma_0 \int_0^N \frac{du}{\sqrt{L(u)}}.$$
 (31)

To obtain a proxy profile, we replace \approx by equality in (31) and differentiate both sides:

$$-\frac{2\alpha}{2\alpha + 2\beta - 1} L(t)^{-\frac{2\alpha}{2\alpha + 2\beta - 1} - 1} L'(t) = M^{\min(\alpha, 0.5)} \gamma_0 \frac{1}{\sqrt{L(t)}}.$$
 (32)

Solving (32) for L'(t) and separating variables gives the separable ODE

$$\frac{dL}{dt} = -\kappa L^{\zeta}, \qquad \zeta = \frac{2\alpha}{2\alpha + 2\beta - 1} + \frac{1}{2}, \qquad \kappa = \frac{2\alpha + 2\beta - 1}{2\alpha} M^{\min(\alpha, 0.5)} \gamma_0.$$

Assuming $\zeta > 1$ (i.e. $2\alpha + 2\beta < 4\alpha + 1$), we integrate to obtain

$$-\frac{L(t)^{-(\zeta-1)}}{\zeta-1} = -\kappa t + \text{constant} \implies L(t) = \left[(\zeta - 1) \kappa t \right]^{-1/(\zeta-1)}. \tag{33}$$

Substituting $\zeta = \frac{2\alpha}{2\alpha + 2\beta - 1} + \frac{1}{2}$ and $\kappa = \frac{2\alpha + 2\beta - 1}{2\alpha} M^{\min(\alpha, 0.5)} \gamma_0$ yields the early-phase proxy

$$L(N) \approx (\gamma_0 M^{\min(\alpha, 0.5)} N)^{-p}, \qquad p := \frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta}.$$
 (34)

By construction, (34) satisfies (31) (hence (30)) up to absolute constants.

D.3.2 LIMIT STAGE (STATIONARY ANALYSIS AND FLOOR)

With $f \equiv 1$, the mode-wise ODE is

$$\frac{dp_i}{dt} = -\frac{4}{\pi\sqrt{P(t)}}\,\lambda_i(\overline{K})\,p_i(t) + \frac{2\gamma_0}{\pi}\,V_i.$$

At stationarity, $p_i(t) \to s_i$ and $P(t) \to L_{\infty}$, we must have

$$-\frac{4}{\pi\sqrt{L_{\infty}}}\,\lambda_{i}(\overline{K})\,s_{i} + \frac{2\gamma_{0}}{\pi}\,V_{i} = 0 \quad \Longrightarrow \quad s_{i} = \frac{\gamma_{0}\sqrt{L_{\infty}}}{2\,\lambda_{i}(\overline{K})}\,V_{i} = \frac{\gamma_{0}\sqrt{L_{\infty}}}{2\,\lambda_{i}(\overline{K})}\,(\boldsymbol{w}_{i}^{\mathsf{T}}\boldsymbol{K}_{\sigma}\boldsymbol{K}\boldsymbol{u}_{i}).$$

Using the risk decomposition $P(t) = \sum_{i=1}^M p_i(t) + \|\boldsymbol{H}^{1/2}\boldsymbol{w}_{\perp}\|^2$, we obtain

$$L_{\infty} = \sum_{i=1}^{M} s_i + \|\boldsymbol{H}^{1/2}\boldsymbol{w}_{\perp}\|^2 = \frac{\gamma_0}{2} \left(\sum_{i=1}^{M} \frac{\boldsymbol{w}_i^{\mathsf{T}} \boldsymbol{K}_{\sigma} \boldsymbol{K} \boldsymbol{u}_i}{\lambda_i(\overline{K})} \right) \sqrt{L_{\infty}} + \|\boldsymbol{H}^{1/2} \boldsymbol{w}_{\perp}\|^2$$
$$= \frac{\gamma_0}{2} \operatorname{Tr} \left(\operatorname{diag}(\boldsymbol{K})^{1/2} \boldsymbol{K}_{\sigma} \right) \sqrt{L_{\infty}} + \|\boldsymbol{H}^{1/2} \boldsymbol{w}_{\perp}\|^2 = \frac{\gamma_0 \pi}{4} \operatorname{Tr} \left(\operatorname{diag}(\boldsymbol{K})^{1/2} \right) \sqrt{L_{\infty}} + \|\boldsymbol{H}^{1/2} \boldsymbol{w}_{\perp}\|^2.$$

Solving the quadratic in $\sqrt{L_{\infty}}$ gives

$$L_{\infty} = \left(\frac{\frac{\gamma_0 \pi}{4} \operatorname{Tr}(\operatorname{diag}(\boldsymbol{K})^{1/2}) + \sqrt{\left(\frac{\gamma_0 \pi}{4} \operatorname{Tr}(\operatorname{diag}(\boldsymbol{K})^{1/2})\right)^2 + 4\|\boldsymbol{H}^{1/2}\boldsymbol{w}_{\perp}\|^2}}{2}\right)^2$$
(35)

$$\approx \max \left\{ \left(\gamma_0 \operatorname{Tr}(\operatorname{diag}(\boldsymbol{K})^{1/2}) \right)^2, \ \|\boldsymbol{H}^{1/2} \boldsymbol{w}_{\perp}\|^2 \right\}. \tag{36}$$

Under our setup,

$$\operatorname{Tr}(\operatorname{diag}(\boldsymbol{K})^{1/2}) = \sum_{i=1}^{M} \sqrt{(\boldsymbol{S}\boldsymbol{H}\boldsymbol{S}^{\mathsf{T}})_{ii}} \approx M \cdot \sqrt{\frac{1}{M} M^{\max(1-2\alpha,0)}} \approx M^{1-\min(\alpha,\,0.5)}.$$

By the results from Paquette et al. (2024); Lin et al. (2024), and note in Appendix I.3,

$$\|\boldsymbol{H}^{1/2}\boldsymbol{w}_{\perp}\|^2 \approx M^{-2\alpha + \max(0, 1-2\beta)}$$

Hence

$$L_{\infty} \approx \max \left\{ \gamma_0^2 M^{2-2\min(\alpha, 0.5)}, M^{-2\alpha + \max(0, 1-2\beta)} \right\}.$$
 (37)

D.3.3 Proxy

Combining the early-phase decay (34) with the floor (37), we adopt

$$L_{\rm px}(N) := \left(\gamma_0 M^{\min(\alpha, 0.5)} N\right)^{-p} + \underbrace{\gamma_0^2 M^{2-2\min(\alpha, 0.5)} + M^{-2\alpha + \max(0, 1-2\beta)}}_{=:C}, \qquad p = \frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta}.$$
(38)

D.3.4 VERIFICATION OF THE PROXY

We show that $L_{\rm px}$ satisfies (28) up to absolute constants. Equivalently, writing $Q_{L_{\rm px}}(N):=\frac{4\gamma_0}{\pi}\int_0^N\frac{du}{\sqrt{L_{\rm px}(u)}}$, we establish

$$\underbrace{\left(M^{\min(\alpha,\,0.5)}\,Q_{L_{\mathrm{px}}}(N)\right)^{-\frac{2\alpha+2\beta-1}{2\alpha}}}_{\mathbf{drift}} + \underbrace{M^{-2\alpha+\max(0,\,1-2\beta)}}_{\mathbf{approx}} + \underbrace{\frac{2\gamma_0^2}{\pi}\sum_{i=1}^M V_i \int_0^N \exp\Bigl(-\frac{4\gamma_0}{\pi}\lambda_i(\overline{K})\int_z^N \frac{du}{\sqrt{L_{\mathrm{px}}(u)}}\Bigr)\,dz}_{\mathbf{noise}}$$

$$= \underbrace{\left(\gamma_0\,M^{\min(\alpha,\,0.5)}\,N\right)^{-p} + C}_{L_{\mathrm{px}}(N)}.$$

$$(40)$$

Lower Bound We prove

$$\mathbf{drift} + \mathbf{approx} + \mathbf{noise} \gtrsim \left(\gamma_0 M^{\min(\alpha, 0.5)} N\right)^{-p} + C. \tag{41}$$

Since $L_{px}(u) \ge (\gamma_0 M^{\min(\alpha, 0.5)} u)^{-p}$,

$$\begin{aligned} \textbf{drift} &= \left(M^{\min(\alpha,\,0.5)}\,Q_{L_{\mathrm{px}}}(N)\right)^{-\frac{2\alpha+2\beta-1}{2\alpha}} \,\gtrsim\, \left(M^{\min(\alpha,\,0.5)}\cdot\gamma_0\int_0^N \!\!\left(\gamma_0\,M^{\min(\alpha,\,0.5)}\,u\right)^{p/2}du\right)^{-\frac{2\alpha+2\beta-1}{2\alpha}} \\ &\approx\, \left(\gamma_0\,M^{\min(\alpha,\,0.5)}\,N\right)^{-p}. \end{aligned}$$

Since $L_{\mathrm{px}}(u) \geq C$ for all u,

$$\int_{z}^{N} \frac{du}{\sqrt{L_{\text{px}}(u)}} \le \frac{N-z}{\sqrt{C}}.$$

Hence

$$\mathbf{noise} \ge \frac{2\gamma_0^2}{\pi} \sum_{i=1}^M V_i \int_0^N \exp\left(-\frac{4\gamma_0}{\pi} \lambda_i(\overline{K}) \frac{N-z}{\sqrt{C}}\right) dz \tag{42}$$

$$=\frac{2\gamma_0^2}{\pi}\sum_{i=1}^M V_i \frac{\sqrt{C}}{\frac{4\gamma_0}{\pi}\lambda_i(\overline{K})} \left(1 - e^{-\frac{4\gamma_0}{\pi}\lambda_i(\overline{K})\frac{N}{\sqrt{C}}}\right) \tag{43}$$

$$\gtrsim \gamma_0 \sqrt{C} \sum_{i=1}^M \frac{V_i}{\lambda_i(\overline{K})} = \frac{\gamma_0}{2} \operatorname{Tr} \left(\operatorname{diag}(K)^{1/2}\right) \sqrt{C} \approx \gamma_0 M^{1-\min(\alpha, 0.5)} \sqrt{C} \gtrsim \gamma_0^2 M^{2-2\min(\alpha, 0.5)}.$$
(44)

Adding the approximation term $M^{-2\alpha+\max(0,\,1-2\beta)}$ gives **noise** + **approx** $\gtrsim C$. Combining with the drift contribution yields (41).

Upper Bound We establish

$$\mathbf{drift} + \mathbf{approx} + \mathbf{noise} \lesssim \left(\gamma_0 M^{\min(\alpha, 0.5)} N\right)^{-p} + C. \tag{45}$$

Let

$$A(N) := \max \{ (\gamma_0 M^{\min(\alpha, 0.5)} N)^{-p}, C \}, \qquad p = \frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta}.$$

Then $L_{\rm px}(N) \approx A(N)$. Define N_0 by $(\gamma_0 M^{\min(\alpha, 0.5)} N_0)^{-p} = C$, i.e.

$$A(N) = \begin{cases} (\gamma_0 M^{\min(\alpha, 0.5)} N)^{-p}, & N \le N_0, \\ C, & N > N_0. \end{cases}$$

There exists a constant $B \ge 1$ such that

$$L_{\rm px}(N) \le B A(N) \qquad (\forall N \ge 0).$$
 (46)

Upper bound for the drift term. Since $L \leq BA$ by (64) and Q is decreasing in its denominator,

$$\mathbf{drift} = \left(M^{\min(\alpha,\,0.5)}\,Q_L(N)\right)^{-\frac{2\alpha+2\beta-1}{2\alpha}} \, \lesssim \, \left(M^{\min(\alpha,\,0.5)}\,Q_{BA}(N)\right)^{-\frac{2\alpha+2\beta-1}{2\alpha}}.$$

We evaluate the right-hand side by cases.

Case $N \leq N_0$. Then $A(u) = (\gamma_0 M^{\min(\alpha, 0.5)} u)^{-p}$ for $u \leq N$, so

$$Q_{BA}(N) = \frac{4\gamma_0}{\pi} \int_0^N \frac{du}{\sqrt{BA(u)}} = \frac{c}{\sqrt{B}} \gamma_0 \int_0^N (\gamma_0 M^{\min(\alpha, 0.5)} u)^{p/2} du$$

for an absolute constant c > 0, which implies

$$\mathbf{drift} \, \lesssim \, \left(\gamma_0 \, M^{\min(\alpha, \, 0.5)} \, N \right)^{-\frac{2\alpha + 2\beta - 1}{2\alpha} (1 + p/2)} = \left(\gamma_0 \, M^{\min(\alpha, \, 0.5)} \, N \right)^{-p}.$$

Case $N > N_0$. Split the integral at N_0 :

$$M^{\min(\alpha, 0.5)}Q_{BA}(N) = \frac{c}{\sqrt{B}} \gamma_0 M^{\min(\alpha, 0.5)} \left[\int_0^{N_0} (\gamma_0 M^{\min(\alpha, 0.5)} u)^{p/2} du + \int_{N_0}^N \frac{du}{\sqrt{BC}} \right]$$
$$= \frac{c}{\sqrt{B}} \left[(\gamma_0 M^{\min(\alpha, 0.5)} N_0)^{1+p/2} + \gamma_0 M^{\min(\alpha, 0.5)} \frac{N - N_0}{\sqrt{BC}} \right].$$

Raising to the power $-\frac{2\alpha+2\beta-1}{2\alpha}$ and using $(\gamma_0\,M^{\min(\alpha,\,0.5)}\,N_0)^{-p}=C$,

$$\mathbf{drift} \lesssim \left[C^{-\frac{1+p/2}{p}} + \gamma_0 \, \frac{N-N_0}{\sqrt{BC}}\right]^{-\frac{2\alpha+2\beta-1}{2\alpha}} \, \leq \, \left(C^{-\frac{1}{(2\alpha+2\beta-1)/(2\alpha)}}\right)^{-\frac{2\alpha+2\beta-1}{2\alpha}} = C.$$

Combining the two cases,

$$\mathbf{drift} \lesssim \left(\gamma_0 \, M^{\min(\alpha, \, 0.5)} \, N\right)^{-p} + C. \tag{47}$$

Upper bound for the noise integral. By the monotonicity of $r \mapsto r^{-1/2}$,

$$\int_z^N \frac{du}{\sqrt{L(u)}} \, \geq \, \frac{1}{\sqrt{B}} \int_z^N \frac{du}{\sqrt{A(u)}}.$$

Therefore,

$$\mathbf{noise} \leq \frac{2\gamma_0^2}{\pi} \sum_{i=1}^M V_i \int_0^N \exp\left(-\frac{4\gamma_0}{\pi\sqrt{B}} \lambda_i(\overline{K}) \int_z^N \frac{du}{\sqrt{A(u)}}\right) dz. \tag{48}$$

We again split into two cases.

Case $N \leq N_0$. Then $A(u) = (\gamma_0 M^{\min(\alpha, 0.5)} u)^{-p}$ on [0, N], hence

$$\int_{z}^{N} \frac{du}{\sqrt{A(u)}} = \left(\gamma_0 M^{\min(\alpha, 0.5)}\right)^{p/2} \int_{z}^{N} u^{p/2} du = \left(\gamma_0 M^{\min(\alpha, 0.5)}\right)^{p/2} \frac{N^{1+p/2} - z^{1+p/2}}{1 + p/2}.$$

Plugging this into (48) and factoring,

$$\begin{split} \mathbf{noise} &= \frac{2\gamma_0^2}{\pi} \sum_{i=1}^{M} V_i \int_0^N \exp\Bigl(-\frac{4\gamma_0}{\pi\sqrt{B}} \, \lambda_i(\overline{K}) \, \bigl(\gamma_0 \, M^{\min(\alpha,\,0.5)}\bigr)^{p/2} \, \frac{N^{1+p/2} - z^{1+p/2}}{1+p/2} \Bigr) \, dz \\ &= \frac{2\gamma_0^2}{\pi} \sum_{i=1}^{M} V_i \, \exp\Bigl(-\frac{4\gamma_0}{\pi\sqrt{B}} \, \lambda_i(\overline{K}) \, \bigl(\gamma_0 \, M^{\min(\alpha,\,0.5)}\bigr)^{p/2} \, \frac{N^{1+p/2}}{1+p/2} \Bigr) \\ &\qquad \times \int_0^N \exp\Bigl(\frac{4\gamma_0}{\pi\sqrt{B}} \, \lambda_i(\overline{K}) \, \bigl(\gamma_0 \, M^{\min(\alpha,\,0.5)}\bigr)^{p/2} \, \frac{z^{1+p/2}}{1+p/2} \Bigr) \, dz. \end{split}$$

Make the change of variables $y=z^{1+p/2}$ so that $dz=\frac{1}{1+p/2}y^{\frac{1}{1+p/2}-1}dy$ and the upper limit becomes $N^{1+p/2}$:

$$\begin{split} \mathbf{noise} &= \frac{2\gamma_0^2}{\pi} \sum_{i=1}^M V_i \, e^{-\alpha_i N^{1+p/2}} \int_0^{N^{1+p/2}} e^{\alpha_i y} \, \frac{1}{1+p/2} \, y^{\frac{1}{1+p/2}-1} \, dy, \\ &\alpha_i := \frac{4\gamma_0}{\pi \sqrt{B}} \, \lambda_i(\overline{K}) \, \frac{\left(\gamma_0 \, M^{\min(\alpha, \, 0.5)}\right)^{p/2}}{1+p/2}. \end{split}$$

Let $X := N^{1+p/2}$ and

$$g(y) := \frac{1}{1 + p/2} y^{\frac{1}{1 + p/2} - 1} = \frac{1}{1 + p/2} y^{-\frac{p}{2 + p}}.$$

Since $e^{\alpha_i y}$ is increasing and g(y) is decreasing on (0, X], Chebyshev's integral inequality (oppositely monotone) yields

$$\frac{1}{X} \int_0^X e^{\alpha_i y} g(y) \, dy \, \leq \, \Big(\frac{1}{X} \int_0^X e^{\alpha_i y} \, dy \Big) \Big(\frac{1}{X} \int_0^X g(y) \, dy \Big).$$

Hence

$$\begin{split} e^{-\alpha_i X} \int_0^X e^{\alpha_i y} g(y) \, dy & \leq e^{-\alpha_i X} \, \frac{e^{\alpha_i X} - 1}{\alpha_i} \, \frac{1}{X} \int_0^X g(y) \, dy \\ & = \frac{1 - e^{-\alpha_i X}}{\alpha_i} \, \frac{1}{1 + p/2} \cdot \frac{1}{1 - \frac{p}{2 + p}} \, X^{-\frac{p}{2 + p}} \\ & = \frac{1 - e^{-\alpha_i X}}{\alpha_i} \, X^{-\frac{p}{2 + p}} \quad \left(\text{since } (1 - \frac{p}{2 + p})(1 + \frac{p}{2}) = 1 \right) \\ & \leq \frac{1}{\alpha_i} \, X^{-\frac{p}{2 + p}} = \frac{1}{\alpha_i} \, N^{-p/2}. \end{split}$$

Therefore

$$\label{eq:noise} \mbox{noise} \ \leq \ \frac{2\gamma_0^2}{\pi} \sum_{i=1}^M V_i \, \frac{1}{\alpha_i} \, N^{-p/2},$$

and with $\alpha_i = \frac{4\gamma_0}{\pi\sqrt{B}} \, \lambda_i(\overline{K}) \, \frac{\left(\gamma_0 \, M^{\min(\alpha, \, 0.5)}\right)^{p/2}}{1+p/2}$ this becomes

$$\mathbf{noise} \ \leq \ \frac{\gamma_0 \sqrt{B}}{2} \left(1 + p/2\right) \sum_{i=1}^M \frac{V_i}{\lambda_i(\overline{K})} \left(\gamma_0 \, M^{\min(\alpha, \, 0.5)} \, N\right)^{-p/2}.$$

Using
$$\sum_i \frac{V_i}{\lambda_i(\overline{K})} = \text{Tr}(\text{diag}(K)^{1/2}) \approx M^{1-\min(\alpha, 0.5)}$$
, we get

$$\begin{split} & \mathbf{noise} \lesssim \gamma_0 \, M^{1-\min(\alpha,\,0.5)} \, (\gamma_0 \, M^{\min(\alpha,\,0.5)} \, N)^{-p/2} \\ &= \gamma_0 \, M^{1-\min(\alpha,\,0.5)} \, (\gamma_0 \, M^{\min(\alpha,\,0.5)} \, N)^{p/2} \, (\gamma_0 \, M^{\min(\alpha,\,0.5)} \, N)^{-p} \\ &\leq \gamma_0 \, M^{1-\min(\alpha,\,0.5)} \, \frac{1}{\sqrt{C}} \, (\gamma_0 \, M^{\min(\alpha,\,0.5)} \, N)^{-p} \, \lesssim \, (\gamma_0 \, M^{\min(\alpha,\,0.5)} \, N)^{-p}, \end{split}$$

where we used $(\gamma_0 M^{\min(\alpha, 0.5)} N)^{p/2} \le (\gamma_0 M^{\min(\alpha, 0.5)} N_0)^{p/2} = C^{-1/2}$.

Case $N > N_0$. Split the z-integral at N_0 :

$$\mathbf{noise} \leq \frac{2\gamma_0^2}{\pi} \sum_{i=1}^M V_i \left[\int_0^{N_0} \exp\Bigl(-\frac{4\gamma_0}{\pi\sqrt{B}} \, \lambda_i(\overline{K}) \int_z^{N_0} \frac{du}{\sqrt{A(u)}} \Bigr) \, dz + \int_{N_0}^N \exp\Bigl(-\frac{4\gamma_0}{\pi\sqrt{B}} \, \lambda_i(\overline{K}) \int_z^N \frac{du}{\sqrt{A(u)}} \Bigr) \, dz \right].$$

The first integral is the $N=N_0$ case just handled, hence

$$\int_0^{N_0} \cdots dz \lesssim (\gamma_0 M^{\min(\alpha, 0.5)} N_0)^{-p} = C.$$

For the second integral, we use that $A \equiv C$ on $[N_0, N]$:

$$\int_{N_0}^{N} \exp\left(-\frac{4\gamma_0}{\pi\sqrt{B}} \lambda_i(\overline{K}) \int_{z}^{N} \frac{du}{\sqrt{A(u)}}\right) dz = \int_{N_0}^{N} \exp\left(-\frac{4\gamma_0}{\pi\sqrt{B}} \lambda_i(\overline{K}) \frac{N-z}{\sqrt{C}}\right) dz$$

$$= \frac{\sqrt{C}}{\frac{4\gamma_0}{\pi\sqrt{B}} \lambda_i(\overline{K})} \left(1 - e^{-\frac{4\gamma_0}{\pi\sqrt{B}} \lambda_i(\overline{K}) \frac{N-N_0}{\sqrt{C}}}\right)$$

$$\leq \frac{\pi\sqrt{B}}{4\gamma_0} \frac{\sqrt{C}}{\lambda_i(\overline{K})}.$$

Therefore,

$$\begin{split} & \textbf{noise} \lesssim (\gamma_0 \, M^{\min(\alpha, \, 0.5)} \, N_0)^{-p} + \frac{2\gamma_0^2}{\pi} \sum_{i=1}^M V_i \cdot \frac{\pi \sqrt{B}}{4\gamma_0} \, \frac{\sqrt{C}}{\lambda_i(\overline{K})} \\ &= C + \frac{\gamma_0 \sqrt{B}}{2} \, \sqrt{C} \sum_{i=1}^M \frac{V_i}{\lambda_i(\overline{K})} = C + \frac{\gamma_0 \sqrt{B}}{2} \, \sqrt{C} \, \operatorname{Tr}(\operatorname{diag}(\boldsymbol{K})^{1/2}) \\ &\lesssim C + \gamma_0 \, M^{\min(\alpha, \, 0.5)} \, \sqrt{C} \, \lesssim \, C + \sqrt{C} \cdot \sqrt{C} \, \lesssim \, C. \end{split}$$

Combining both cases,

$$\mathbf{noise} \lesssim \left(\gamma_0 \, M^{\min(\alpha, \, 0.5)} \, N\right)^{-p} \, + \, C. \tag{49}$$

Conclusion of the upper bound. From (47), (49), and approx $= M^{-2\alpha + \max(0, 1-2\beta)} \le C$, we obtain (45).

Finally, combining the lower bound (41) and the upper bound (45) proves (40). Therefore, the proxy (58) satisfies the implicit relation (28) up to absolute constants, with the three contributions labeled as **approx**, **drift**, and **noise**.

D.4 Constant Learning Rate: Proxy and Verification for the Case $\alpha>0.5$ and $\beta>0.5$ (Phase B)

We now handle the case $\alpha>0.5$ and $\beta>0.5$. Since $\alpha>0.5$, we have $\min(\alpha,0.5)=0.5$, and because $\beta>0.5$, we have $\min(2\alpha,2\alpha+2\beta-1)=2\alpha$. Applying the drift/approximation-term transformation to the ODE solution yields

$$L(N) \approx \underbrace{M^{-2\alpha}}_{\text{approx}} + \underbrace{\left(M^{1/2}Q(N)\right)^{-\frac{2\alpha+2\beta-1}{2\alpha}}}_{\text{drift}_1} + \underbrace{M^{-1}\left(M^{1/2}Q(N)\right)^{-1+\frac{1}{2\alpha}}}_{\text{drift}_2} \tag{50}$$

$$+\underbrace{\frac{2\gamma_0^2}{\pi}\sum_{i=1}^M V_i \int_0^N \exp\left(-\frac{4\gamma_0}{\pi}\lambda_i(\overline{K})\int_z^N \frac{du}{\sqrt{L(u)}}\right) dz}_{\text{poise}},$$
 (51)

where

$$Q(N) = \frac{4\gamma_0}{\pi} \int_0^N \frac{du}{\sqrt{L(u)}}.$$

D.4.1 EARLY STAGE PROXIES (DRIFT₁ AND DRIFT₂)

We extract proxies from the two drift terms in (50) by the same differentiate-and-separate trick as before.

drift₁: $(M^{1/2}Q(N))^{-(2\alpha+2\beta-1)/(2\alpha)}$. Assuming this term dominates and replacing π by equality,

$$L(N)^{-\frac{2\alpha}{2\alpha+2\beta-1}} = M^{1/2}\gamma_0 \int_0^N \frac{du}{\sqrt{L(u)}}.$$

Differentiation gives the separable ODE $L'(t) = -\kappa_1 L(t)^{\beta_1}$ with

$$\beta_1 = \frac{2\alpha}{2\alpha + 2\beta - 1} + \frac{1}{2}, \qquad \kappa_1 = \frac{2\alpha + 2\beta - 1}{2\alpha} M^{1/2} \gamma_0.$$

For $\beta_1 > 1$ (equivalently $2\alpha + 2\beta < 4\alpha + 1$) we obtain

$$L_1(N) \approx (\gamma_0 M^{1/2} N)^{-p_1}, \qquad p_1 = \frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta}.$$
 (52)

drift₂: $M^{-1}(M^{1/2}Q(N))^{-1+\frac{1}{2\alpha}}$. Assume $\alpha > \frac{1}{2}$ and, in the early phase, the second drift term dominates:

$$L(N) \approx M^{-1} (M^{1/2} Q(N))^{-\frac{2\alpha - 1}{2\alpha}}, \qquad Q(N) \approx \gamma_0 \int_0^N \frac{du}{\sqrt{L(u)}}.$$

Expanding the M-exponent,

$$(M^{1/2}Q)^{-\frac{2\alpha-1}{2\alpha}} = M^{\frac{-(2\alpha-1)}{4\alpha}} Q^{-\frac{2\alpha-1}{2\alpha}},$$

hence

$$L(N) \approx M^{-\frac{6\alpha-1}{4\alpha}} \left(\gamma_0 I(N)\right)^{-\frac{2\alpha-1}{2\alpha}}, \qquad I(N) := \int_0^N \frac{du}{\sqrt{L(u)}}.$$
 (53)

Raise both sides of (53) to the power $-\frac{2\alpha}{2\alpha-1}$ so that the integral becomes linear:

$$L(N)^{-\frac{2\alpha}{2\alpha-1}} = M^{\frac{6\alpha-1}{4\alpha-2}} \gamma_0 I(N) \approx M^{\frac{6\alpha-1}{4\alpha-2}} \gamma_0 \int_0^N \frac{du}{\sqrt{L(u)}}.$$
 (54)

Differentiating (54) with respect to t yields

$$-\frac{2\alpha}{2\alpha - 1} L(t)^{-\frac{2\alpha}{2\alpha - 1} - 1} L'(t) = M^{\frac{6\alpha - 1}{4\alpha - 2}} \gamma_0 \frac{1}{\sqrt{L(t)}}.$$

Rearranging gives a separable ODE of the usual power form

$$L'(t) = -\kappa_2 L(t)^{\beta_2}, \qquad \beta_2 = \frac{2\alpha}{2\alpha - 1} + \frac{1}{2} = \frac{6\alpha - 1}{4\alpha - 2} > 1,$$
 (55)

with

$$\kappa_2 = \frac{2\alpha - 1}{2\alpha} \gamma_0 M^{\frac{6\alpha - 1}{4\alpha - 2}} > 0.$$
 (56)

Since $\beta_2 > 1$, solving (55) gives

$$L(t)^{-(\beta_2-1)} = (\beta_2-1) \kappa_2 t + \text{const.}$$

Absorbing harmless absolute constants into $\overline{\sim}$ and setting t=N,

$$L_2(N) \approx \left(\gamma_0 M^{\frac{6\alpha - 1}{4\alpha - 2}} N\right)^{-p_2}, \qquad p_2 = \frac{1}{\beta_2 - 1} = \left| \frac{2(2\alpha - 1)}{2\alpha + 1} \right|.$$
 (57)

Crossover scale. Equating (52) and (57) gives

$$N_1 \approx \gamma_0^{-1} M^{\eta}, \qquad \eta = \frac{2\alpha + 1 - 4\beta}{4\beta},$$

so R_1 dominates for $N \lesssim N_1$ and L_2 for $N \gtrsim N_1$ (when $\alpha > 0.5$ and $0.5 < \beta < \alpha + 0.5$).

D.4.2 LIMIT STAGE (APPROX AND NOISE FLOORS)

As in the case $\alpha < 0.5$ or $\beta < 0.5$, the stationary analysis with $f \equiv 1$ yields

$$L_{\infty} \approx \max \{ \gamma_0^2 \operatorname{Tr}(\operatorname{diag}(\boldsymbol{K})^{1/2})^2, \|\boldsymbol{H}^{1/2}\boldsymbol{w}_{\perp}\|^2 \}.$$

Under our standing model $\operatorname{Tr}(\operatorname{diag}(\boldsymbol{K})^{1/2}) \approx M^{0.5}$ and by the results from Paquette et al. (2024); Lin et al. (2024), and note in Appendix I.3, $\|\boldsymbol{H}^{1/2}\boldsymbol{w}_{\perp}\|^2 \approx M^{-2\alpha}$, hence the floor

$$C := \gamma_0^2 M + M^{-2\alpha}.$$

D.4.3 COMBINED PROXY

$$L_{px}(N) := L_1(N) + L_2(N) + C$$

= $(\gamma_0 M^{0.5} N)^{-p_1} + (\gamma_0 M^{\frac{6\alpha - 1}{4\alpha - 2}} N)^{-p_2} + C,$ (58)

where

$$p_1 = \frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta}, \qquad p_2 = \frac{4\alpha - 2}{2\alpha + 1}.$$

D.4.4 VERIFICATION OF THE PROXY

We show that $L_{\rm DX}$ satisfies (50) up to absolute constants.

Lower bound. We claim

$$\underbrace{\left(\underline{M^{0.5}Q_{L_{\mathrm{px}}}(N)\right)^{-\frac{2\alpha+2\beta-1}{2\alpha}}}_{\text{drift}_{1}} + \underbrace{M^{-1}\left(M^{0.5}Q_{L_{\mathrm{px}}}(N)\right)^{-1+\frac{1}{2\alpha}}}_{\text{drift}_{2}} + \underbrace{M^{-2\alpha}}_{\text{approx}}$$
(59)

$$+\underbrace{\frac{2\gamma_0^2}{\pi}\sum_{i=1}^M V_i \int_0^N \exp\left(-\frac{4\gamma_0}{\pi}\lambda_i(\overline{K}) \int_z^N \frac{du}{\sqrt{L_{\text{px}}(u)}}\right) dz}_{\text{poise}} \gtrsim L_{\text{px}}(N). \tag{60}$$

Drift part. Using $L_{\text{px}} \geq R_1$ inside Q,

$$\left(M^{0.5} Q_{L_{\text{px}}}(N) \right)^{-\frac{2\alpha + 2\beta - 1}{2\alpha}} \gtrsim \left(M^{0.5} \gamma_0 \int_0^N \frac{du}{\sqrt{L_1(u)}} \right)^{-\frac{2\alpha + 2\beta - 1}{2\alpha}} \approx (\gamma_0 M^{0.5} N)^{-p_1} \approx L_1(N).$$

Similarly, using $L_{\text{px}} \geq L_2$ inside Q,

$$M^{-1} \left(M^{0.5} Q_{L_{\text{px}}}(N) \right)^{-1 + \frac{1}{2\alpha}} \gtrsim M^{-1} \left(M^{0.5} \gamma_0 \int_0^N \frac{du}{\sqrt{L_2(u)}} \right)^{-1 + \frac{1}{2\alpha}} \approx \left(\gamma_0 M^{\frac{6\alpha - 1}{4\alpha - 2}} N \right)^{-p_2} \approx L_2(N).$$

Therefore.

$$drift_1 + drift_2 \gtrsim L_1(N) + L_2(N). \tag{61}$$

Noise + approx. Since $L_{px} \geq C$,

$$\int_{z}^{N} \frac{du}{\sqrt{L_{\mathrm{px}}(u)}} \leq \frac{N-z}{\sqrt{C}}.$$

As in the Equation 44,

$$\text{noise} \ \gtrsim \ \gamma_0 \sqrt{C} \sum_{i=1}^M \frac{V_i}{\lambda_i(\overline{K})} \ = \ \frac{\gamma_0}{2} \ \text{Tr} \big(\text{diag}(K)^{1/2} \big) \, \sqrt{C} \ \eqsim \ \gamma_0 \, M^{0.5} \sqrt{C} \ \gtrsim \ \gamma_0^2 \, M.$$

Thus noise + approx $\gtrsim C$. Together with (61), this proves (60).

Upper bound. We will prove

$$\underbrace{\left(M^{0.5}Q_{L_{\mathrm{px}}}(N)\right)^{-\frac{2\alpha+2\beta-1}{2\alpha}}}_{\text{drift}_{1}} + \underbrace{M^{-1}\left(M^{0.5}Q_{L_{\mathrm{px}}}(N)\right)^{-1+\frac{1}{2\alpha}}}_{\text{drift}_{2}} + \underbrace{M^{-2\alpha}}_{\text{approx}}$$
(62)

$$+\underbrace{\frac{2\gamma_0^2}{\pi}\sum_{i=1}^M V_i \int_0^N \exp\left(-\frac{4\gamma_0}{\pi}\lambda_i(\overline{K}) \int_z^N \frac{du}{\sqrt{L_{\text{px}}(u)}}\right) dz}_{} \lesssim L_{\text{px}}(N). \tag{63}$$

Let

$$A(N) = \begin{cases} (\gamma_0 M^{0.5} N)^{-p_1}, & N \leq N_1, \\ (\gamma_0 M^{\frac{6\alpha - 1}{4\alpha - 2}} N)^{-p_2}, & N_1 \leq N \leq N_2, \\ C, & N > N_2, \end{cases}$$

where N_1 and N_2 are the crossover points between the three terms. There exists a constant $B \ge 1$ such that

$$L_{\rm px}(N) \le B A(N) \qquad (\forall N \ge 0).$$
 (64)

It suffices to show

$$\underbrace{\left(M^{0.5}Q_{B\cdot A}(N)\right)^{-\frac{2\alpha+2\beta-1}{2\alpha}}}_{\text{drift}_1} + \underbrace{M^{-1}\left(M^{0.5}Q_{B\cdot A}(N)\right)^{-1+\frac{1}{2\alpha}}}_{\text{drift}_2} + \underbrace{M^{-2\alpha}}_{\text{approx}}$$
(65)

$$+ \underbrace{\frac{2\gamma_0^2}{\pi} \sum_{i=1}^M V_i \int_0^N \exp\left(-\frac{4\gamma_0}{\pi} \lambda_i(\overline{K}) \int_z^N \frac{du}{\sqrt{B \cdot A(u)}}\right) dz}_{i=1} \lesssim L_{\text{px}}(N). \tag{66}$$

Case $N \leq N_1$. It is enough to prove

$$\underbrace{\left(M^{0.5}Q_{B\cdot A}(N)\right)^{-\frac{2\alpha+2\beta-1}{2\alpha}}}_{\text{drift}_1} + \underbrace{M^{-1}\left(M^{0.5}Q_{B\cdot A}(N)\right)^{-1+\frac{1}{2\alpha}}}_{\text{drift}_2} + \underbrace{M^{-2\alpha}}_{\text{approx}}$$
(67)

$$+\underbrace{\frac{2\gamma_0^2}{\pi}\sum_{i=1}^M V_i \int_0^N \exp\left(-\frac{4\gamma_0}{\pi}\lambda_i(\overline{K})\int_z^N \frac{du}{\sqrt{B\cdot A(u)}}\right) dz}_{\text{poiss}} \lesssim (\gamma_0 M^{0.5} N)^{-p_1}. \tag{68}$$

We have $M^{-2\alpha} \lesssim (\gamma_0 M^{0.5} N)^{-p_1}$ directly. Also, the following holds with straightforward integration.

$$(M^{0.5}Q_{B\cdot A}(N))^{-\frac{2\alpha+2\beta-1}{2\alpha}} \approx (\gamma_0 M^{0.5}N)^{-p_1}.$$

Since $N \leq N_1 \approx \gamma_0^{-1} \, M^{\eta}$ with $\eta = \frac{2\alpha + 1 - 4\beta}{4\beta}$, following holds by integration and calculation.

$$M^{-1}(M^{0.5}Q_{B\cdot A}(N))^{-1+\frac{1}{2\alpha}} \lesssim (\gamma_0 M^{0.5}N)^{-p_1}.$$

Finally, arguing as in the $N \leq N_0$ case of Section D.3.4,

$$\frac{2\gamma_0^2}{\pi} \sum_{i=1}^M V_i \int_0^N \exp \left(-\frac{4\gamma_0}{\pi} \lambda_i(\overline{K}) \int_z^N \frac{du}{\sqrt{B \cdot A(u)}} \right) dz \, \lesssim \, (\gamma_0 \, M^{0.5} N)^{-p_1}.$$

Hence, the claim holds for $N \leq N_1$.

Case $N_1 \leq N \leq N_2$. We will show

$$\underbrace{\left(M^{0.5}Q_{BA}(N)\right)^{-\frac{2\alpha+2\beta-1}{2\alpha}}}_{\text{drift}_1} + \underbrace{M^{-1}\left(M^{0.5}Q_{BA}(N)\right)^{-1+\frac{1}{2\alpha}}}_{\text{drift}_2} + \underbrace{M^{-2\alpha}}_{\text{approx}}$$
(69)

$$+\underbrace{\frac{2\gamma_0^2}{\pi}\sum_{i=1}^M V_i \int_0^N \exp\left(-\frac{4\gamma_0}{\pi}\lambda_i(\overline{K})\int_z^N \frac{du}{\sqrt{BA(u)}}\right) dz}_{\text{poise}} \lesssim \left(\gamma_0 M^{\frac{6\alpha-1}{4\alpha-2}} N\right)^{-p_2},\tag{70}$$

where

$$p_1 = \frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta}, \qquad p_2 = \frac{2(2\alpha - 1)}{2\alpha + 1}, \qquad A(u) = \begin{cases} (\gamma_0 M^{0.5} u)^{-p_1}, & u \le N_1, \\ (\gamma_0 M^{\frac{6\alpha - 1}{4\alpha - 2}} u)^{-p_2}, & N_1 < u \le N, \end{cases}$$

and
$$Q_{BA}(N) = \frac{4\gamma_0}{\pi} \int_0^N \frac{du}{\sqrt{BA(u)}}$$
.

Approx term. Since $N \leq N_2$,

$$M^{-2\alpha} \lesssim \left(\gamma_0 M^{\frac{6\alpha-1}{4\alpha-2}} N\right)^{-p_2}$$

Drift term. If $N_1 \leq N \leq 2N_1$, using the case $N \leq N_1$, we get an inequality for two drift terms.

$$\underbrace{\left(M^{0.5}Q_{BA}(N)\right)^{-\frac{2\alpha+2\beta-1}{2\alpha}}}_{\text{drift}} + \underbrace{M^{-1}\left(M^{0.5}Q_{BA}(N)\right)^{-1+\frac{1}{2\alpha}}}_{\text{drift}} \tag{71}$$

$$\leq \underbrace{\left(M^{0.5}Q_{BA}(N_{1})\right)^{-\frac{2\alpha+2\beta-1}{2\alpha}}}_{\text{drift}_{1}} + \underbrace{M^{-1}\left(M^{0.5}Q_{BA}(N_{1})\right)^{-1+\frac{1}{2\alpha}}}_{\text{drift}_{2}} \tag{72}$$

$$\lesssim (\gamma_0 M^{0.5} N_1)^{-p_1} \lesssim (\gamma_0 M^{\frac{6\alpha - 1}{4\alpha - 2}} N)^{-p_2}.$$
 (73)

So while covering the drift term, we will temporarily assume $2N_1 \leq N$.

Lower bound on $Q_{BA}(N)$. Split the integral at N_1 :

$$Q_{BA}(N) \approx \gamma_0 \int_0^{N_1} \frac{du}{\sqrt{A(u)}} + \gamma_0 \int_{N_1}^{N} \frac{du}{\sqrt{A(u)}} =: \gamma_0 (I_1 + I_2). \tag{74}$$

For the first part, using $A(u) = (\gamma_0 M^{0.5} u)^{-p_1}$ on $[0, N_1]$,

$$I_1 = (\gamma_0 M^{0.5})^{p_1/2} \int_0^{N_1} u^{p_1/2} du = \frac{(\gamma_0 M^{0.5})^{p_1/2}}{1 + p_1/2} N_1^{1 + p_1/2}.$$
 (75)

2110 2111 For the second part, using $A(u)=(\gamma_0 M^{\frac{6\alpha-1}{4\alpha-2}}u)^{-p_2}$ on $[N_1,N]$

$$I_2 = \left(\gamma_0 M^{\frac{6\alpha - 1}{4\alpha - 2}}\right)^{p_2/2} \int_{N_1}^N u^{p_2/2} \, du = \frac{\left(\gamma_0 M^{\frac{6\alpha - 1}{4\alpha - 2}}\right)^{p_2/2}}{1 + p_2/2} \left(N^{1 + p_2/2} - N_1^{1 + p_2/2}\right). \tag{76}$$

Since we temporarily assumed $N \geq 2N_1$, we have

$$I_2 \gtrsim (\gamma_0 M^{\frac{6\alpha-1}{4\alpha-2}})^{p_2/2} N^{1+p_2/2}$$

Hence, from (74),

$$Q_{BA}(N) \gtrsim \gamma_0 \left(\gamma_0 M^{\frac{6\alpha - 1}{4\alpha - 2}}\right)^{p_2/2} N^{1 + p_2/2}.$$
 (77)

drift₁ vs. **drift**₂. From $N \ge N_1$ and (77), we have $Q_{BA}(N) \ge Q_{BA}(N_1)$. It follows that

$$\mathrm{drift}_1 = \left(M^{0.5}Q_{BA}(N)\right)^{-\frac{2\alpha+2\beta-1}{2\alpha}} \ \leq \ M^{-1} \left(M^{0.5}Q_{BA}(N)\right)^{-1+\frac{1}{2\alpha}} = \mathrm{drift}_2,$$

so it suffices to control drift₂.

drift₂ **bound.** Using (77),

$$\operatorname{drift}_{2} = M^{-1} \left(M^{0.5} Q_{BA}(N) \right)^{-1 + \frac{1}{2\alpha}} \\
 \lesssim M^{-1} \left(M^{0.5} \cdot \gamma_{0}^{1 + p_{2}/2} M^{\frac{6\alpha - 1}{4\alpha - 2} \cdot \frac{p_{2}}{2}} N^{1 + p_{2}/2} \right)^{-1 + \frac{1}{2\alpha}}.$$
(78)

Now compute the exponents of N, γ_0 , and M separately.

(i) N-exponent:

$$(1 + \frac{p_2}{2}) \left(-1 + \frac{1}{2\alpha} \right) = \left(1 + \frac{2\alpha - 1}{2\alpha + 1} \right) \left(\frac{1}{2\alpha} - 1 \right) = \frac{4\alpha}{2\alpha + 1} \cdot \left(-\frac{2\alpha - 1}{2\alpha} \right) = -\frac{2(2\alpha - 1)}{2\alpha + 1} = -p_2.$$

- 2137 (ii) γ_0 -exponent: the same calculation as in (i) gives $-p_2$.
- (iii) M-exponent: the total exponent equals

$$-1 + \left(-1 + \frac{1}{2\alpha}\right) \left(0.5 + \frac{6\alpha - 1}{4\alpha - 2} \cdot \frac{p_2}{2}\right).$$

A direct simplification shows this equals $-\frac{6\alpha-1}{4\alpha-2}p_2$. Therefore, from (78),

$$drift_2 \lesssim \left(\gamma_0 M^{\frac{6\alpha - 1}{4\alpha - 2}} N\right)^{-p_2}. \tag{79}$$

Since drift₁ \leq drift₂, we also have drift₁ $\lesssim (\gamma_0 M^{\frac{6\alpha-1}{4\alpha-2}} N)^{-p_2}$.

Noise bound. It suffices to show

$$\frac{2\gamma_0^2}{\pi} \sum_{i=1}^M V_i \int_0^{N_1} \exp\left(-\frac{4\gamma_0}{\pi} \lambda_i(\overline{K}) \int_z^N \frac{du}{\sqrt{B \cdot A(u)}}\right) dz \tag{80}$$

$$+ \frac{2\gamma_0^2}{\pi} \sum_{i=1}^M V_i \int_{N_1}^N \exp\left(-\frac{4\gamma_0}{\pi} \lambda_i(\overline{K}) \int_z^N \frac{du}{\sqrt{B \cdot A(u)}}\right) dz \lesssim \left(\gamma_0 M^{\frac{6\alpha - 1}{4\alpha - 2}} N\right)^{-p_2}. \tag{81}$$

Integral over $[N_1, N]$. As in the case $N \leq N_0$ of Section D.3.4, with $A(u) = (\gamma_0 M^{\frac{6\alpha - 1}{4\alpha - 2}} u)^{-p_2}$ on $[N_1, N]$,

$$\frac{2\gamma_0^2}{\pi} \sum_{i=1}^M V_i \int_{N_1}^N \exp\left(-\frac{4\gamma_0}{\pi} \lambda_i(\overline{K}) \int_z^N \frac{du}{\sqrt{B \cdot A(u)}}\right) dz \lesssim \left(\gamma_0 M^{\frac{6\alpha - 1}{4\alpha - 2}} N\right)^{-p_2}.$$

Integral over $[0, N_1]$. First,

$$\frac{2162}{2163} \qquad \frac{2\gamma_0^2}{\pi} \sum_{i=1}^{M} V_i \int_0^{N_1} \exp \left(-\frac{4\gamma_0}{\pi} \lambda_i(\overline{K}) \int_z^N \frac{du}{\sqrt{B \cdot A(u)}} \right) dz \\ \leq \frac{2\gamma_0^2}{\pi} \sum_{i=1}^{M} V_i \int_0^{N_1} \exp \left(-\frac{4\gamma_0}{\pi} \lambda_i(\overline{K}) \int_z^{N_1} \frac{du}{\sqrt{B \cdot A(u)}} \right) dz.$$

As in the case $N \leq N_0$ of Section D.3.4,

$$\frac{2\gamma_0^2}{\pi} \sum_{i=1}^{M} V_i \int_0^{N_1} \exp\left(-\frac{4\gamma_0}{\pi} \lambda_i(\overline{K}) \int_z^{N_1} \frac{du}{\sqrt{B \cdot A(u)}}\right) dz \lesssim \sqrt{C} \left(\gamma_0 M^{0.5} N\right)^{-p_1/2} \lesssim \left(\gamma_0 M^{\frac{6\alpha - 1}{4\alpha - 2}} N_1\right)^{-p_2}.$$

If $N \leq 2N_1$, this already implies

$$\frac{2\gamma_0^2}{\pi} \sum_{i=1}^M V_i \int_0^{N_1} \exp\Bigl(-\tfrac{4\gamma_0}{\pi} \lambda_i(\overline{K}) \int_z^N \tfrac{du}{\sqrt{B \cdot A(u)}}\Bigr) \, dz \; \lesssim \; \bigl(\gamma_0 \, M^{\tfrac{6\alpha-1}{4\alpha-2}} \, N\bigr)^{-p_2}.$$

If $N > 2N_1$, then

$$\int_0^{N_1} \exp\Bigl(-\tfrac{4\gamma_0}{\pi}\lambda_i(\overline{K}) \int_z^N \tfrac{du}{\sqrt{B\cdot A(u)}}\Bigr) \, dz \; \leq \; N_1 \, \exp\Bigl(-\tfrac{4\gamma_0}{\pi}\lambda_i(\overline{K}) \int_{N_1}^N \tfrac{du}{\sqrt{B\cdot A(u)}}\Bigr),$$

and, using $e^{-x} \leq 1/x$ together with the lower bound $\int_{N_1}^{N} \frac{du}{\sqrt{B \cdot A(u)}} \gtrsim N^{1+p_2/2} \left(\gamma_0 M^{\frac{6\alpha-1}{4\alpha-2}}\right)^{p_2/2}$, we get

$$\begin{split} \frac{2\gamma_0^2}{\pi} \sum_{i=1}^M V_i \int_0^{N_1} \cdots \, dz &\lesssim \frac{2\gamma_0^2}{\pi} \sum_{i=1}^M V_i \, \frac{N_1}{\frac{4\gamma_0}{\pi} \lambda_i(\overline{K}) \int_{N_1}^N \frac{du}{\sqrt{B \cdot A(u)}}} \\ &\lesssim \gamma_0 \sum_{i=1}^M \frac{V_i}{\lambda_i(\overline{K})} \left(\gamma_0 M^{\frac{6\alpha - 1}{4\alpha - 2}} N \right)^{-p_2/2} \\ &\approx \gamma_0 \, M^{0.5} \left(\gamma_0 M^{\frac{6\alpha - 1}{4\alpha - 2}} N \right)^{-p_2/2} \lesssim \left(\gamma_0 \, M^{\frac{6\alpha - 1}{4\alpha - 2}} N \right)^{-p_2}, \end{split}$$

where the last step uses $\gamma_0 M^{0.5} \leq (\gamma_0 M^{\frac{6\alpha-1}{4\alpha-2}} N)^{-p_2/2}$ which holds from $N \leq N_2$.

Combining the $[N_1, N]$ and $[0, N_1]$ bounds yields

noise
$$\lesssim (\gamma_0 M^{\frac{6\alpha-1}{4\alpha-2}} N)^{-p_2}$$
,

as required for the case $N_1 \leq N \leq N_2$.

Case $N \ge N_2$. We have $M^{-2\alpha} \lesssim C$ directly. As in the above case,

$$(M^{0.5}Q_{B\cdot A}(N))^{-\frac{2\alpha+2\beta-1}{2\alpha}} \le M^{-1}(M^{0.5}Q_{B\cdot A}(N))^{-1+\frac{1}{2\alpha}}$$

Using the estimate from the previous case,

$$M^{-1} \left(M^{0.5} Q_{B \cdot A}(N) \right)^{-1 + \frac{1}{2\alpha}} \lesssim M^{-1} \left(M^{0.5} Q_{B \cdot A}(N_2) \right)^{-1 + \frac{1}{2\alpha}} \lesssim \left(\gamma_0 M^{\frac{6\alpha - 1}{4\alpha - 2}} N_2 \right)^{-p_2} \lesssim C.$$

Finally, as in the $N > N_0$ case of Section D.3.4,

$$\frac{2\gamma_0^2}{\pi} \sum_{i=1}^M V_i \int_0^N \exp\left(-\frac{4\gamma_0}{\pi} \lambda_i(\overline{K}) \int_z^N \frac{du}{\sqrt{B \cdot A(u)}}\right) dz \lesssim C.$$

Therefore, the bound holds for $N \geq N_2$ as well.

D.5 Note on the Regime $\beta > \alpha + 0.5$

When $\beta>\alpha+0.5$, the assumption $\zeta>1$ used in step 33 no longer holds. In this case, the first drift term takes a different form:

$$L_{\text{drift}_1}(N) \approx \left(1 - \kappa \gamma_0 M^{\min(\alpha, 0.5)} N\right)^{\frac{2(2\alpha + 2\beta - 1)}{2\beta - 2\alpha - 1}},$$

for a finite horizon and some constant κ . Inserting the max function, we can represent it as a global function.

 $L_{\text{drift}_1}(N) \approx \left(\max \left(1 - \kappa \gamma_0 \, M^{\min(\alpha, 0.5)} \, N, 0 \right) \right)^{\frac{2(2\alpha + 2\beta - 1)}{2\beta - 2\alpha - 1}}.$

Now we explain the behavior of the term. When N is asymptotically smaller than $(\gamma_0 M^{\min(\alpha,0.5)})^{-1}$, the term is asymptotically constant. On $N = (\gamma_0 M^{\min(\alpha,0.5)})^{-1}$, the term suddenly drops form constant scale to 0.

For the case $\alpha < 0.5$ or $\beta < 0.5$ the valid proxy is

$$2223 \\ 2224 \qquad L_{\mathrm{px}}(N) := \left(\max \left(1 - \kappa \, \gamma_0 \, M^{\min(\alpha, 0.5)} \, N, 0 \right) \right)^{\frac{2(2\alpha + 2\beta - 1)}{2\beta - 2\alpha - 1}} + \gamma_0^2 \, M^{\, 2 - 2 \min(\alpha, \, 0.5)} + M^{-2\alpha + \max(0, \, 1 - 2\beta)},$$

2225 and for the case $\alpha > 0.5$ and $\beta > 0.5$ the valid proxy is

$$2227 \\ 2228 \qquad L_{\mathrm{px}}(N) := \left(\max \left(1 - \kappa \, \gamma_0 \, M^{0.5} \, N, 0 \right) \right)^{\frac{2(2\alpha + 2\beta - 1)}{2\beta - 2\alpha - 1}} + M^{-\frac{6\alpha - 1}{2\alpha + 1}} (N\gamma_0)^{-\frac{2(2\alpha - 1)}{2\alpha + 1}} + \gamma_0^2 \, M + M^{-2\alpha}.$$

These satisfy the implicit integral equation, same as Sections D.3.4 and D.4.4.

Therefore, for the case $\alpha < 0.5$, $\beta > \alpha + 0.5$,

$$R(M, N, \gamma_0) = \left(\max \left(1 - \kappa \gamma_0 M^{\alpha} N, 0 \right) \right)^{\frac{2(2\alpha + 2\beta - 1)}{2\beta - 2\alpha - 1}} + \gamma_0^2 M^{2 - 2\alpha} + M^{-2\alpha}, \tag{82}$$

and for the case $\alpha > 0.5$ and $\beta > 0.5$,

$$R(M, N, \gamma_0) = \left(\max\left(1 - \kappa \gamma_0 M^{0.5} N, 0\right)\right)^{\frac{2(2\alpha + 2\beta - 1)}{2\beta - 2\alpha - 1}} + M^{-\frac{6\alpha - 1}{2\alpha + 1}} (N\gamma_0)^{-\frac{2(2\alpha - 1)}{2\alpha + 1}} + \gamma_0^2 M + M^{-2\alpha}.$$
(83)

E DERIVATION OF THE COMPUTE-OPTIMAL RESULT

E.1 COMPUTE-OPTIMAL RESULT FOR MAXIMAL LEARNING RATE

We now discuss the maximal learning rate case deferred from the main text. Note that Paquette et al. (2024) showed that the maximal learning rate for SGD is $\gamma_0 \approx 1$ when $\alpha > \frac{1}{2}$, and $\gamma_0 \approx M^{-(1-2\alpha)}$ when $\alpha < \frac{1}{2}$.

Now, we discuss the maximal learning rate for signSGD. Because the noise term is $\gamma_0^2 M^{2-\min(1,2\alpha)}$, stability requires

$$\gamma_0^2 M^{2-\min(1,2\alpha)} \lesssim 1.$$

Otherwise, the signSGD noise term explodes as M grows. This condition is satisfied by choosing

$$\gamma_0 = M^{-1 + \min(\alpha, 0.5)},$$

which ensures $\gamma_0^2 M^{2-\min(1,2\alpha)} \approx 1$ while the other terms still decay appropriately.

For $\alpha < 0.5$, the term

$$\left(M^{\min(\alpha, \, 0.5)} \, N \, \gamma_0 \right)^{-\frac{2(2\alpha + 2\beta - 1)}{2\alpha - 2\beta + 1}} \; = \; \left(M^{-(1 - 2\alpha)} \, N \right)^{-\frac{2(2\alpha + 2\beta - 1)}{2\alpha - 2\beta + 1}}$$

decreases with N but increases with M. However, under a fixed compute budget $\mathfrak{f}=MN$, one can allocate resources so that this term does not cause an exploding loss; hence we do not classify it as unstable.

Thus, the maximal learning rate for signSGD is

$$\gamma_0 = M^{-1 + \min(\alpha, 0.5)}$$

In this case, however, we obtain $R(M, N, \gamma_0) \approx 1$, so the slope of the compute-optimal curve is always zero.

E.2 DERIVATION OF COMPUTE-OPTIMAL RESULT FOR OPTIMAL LEARNING RATE

We assume $\alpha + \beta > 0.5$ throughout, even for the case where it is not specified.

2272 E.2.1 $\alpha > 0.5$, $\beta < 0.5$ (Phase Aa)

We start from

$$R(M, N, \gamma_0) \approx (M^{1/2}N\gamma_0)^{-\frac{2(2\alpha+2\beta-1)}{2\alpha+1-2\beta}} + M^{-(2\alpha+2\beta-1)} + \gamma_0^2 M.$$

Substitute

$$\gamma_0 = M^{-e}, \qquad N = \frac{\mathfrak{f}}{M}, \qquad M = \mathfrak{f}^x,$$

so that, up to constant factors,

$$R \approx f^{\max\{\ell_1(x), \ell_2(x), \ell_3(x)\}}$$

where

$$\ell_1(x) = -(2\alpha + 2\beta - 1) x,$$

$$\ell_2(x) = \frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta} \left(e + \frac{1}{2} \right) x - \frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta},$$

$$\ell_3(x) = (1 - 2e) x.$$

We minimize the convex, piecewise–linear function $f(x,e) = \max_i \ell_i(x,e)$ over $x \in (0,1)$ and $e \in \mathbb{R}$. By convexity, any interior minimizer must occur at a kink where at least two lines are active. In our regime $\alpha + \beta > 0.5$ and $\beta < \alpha + 0.5$, the only admissible triple intersection is $\{\ell_1, \ell_2, \ell_3\}$. Solving $\ell_1 = \ell_3$ and $\ell_2 = \ell_3$ yields

$$e^* = \alpha + \beta,$$
 $x^* = \frac{1}{2\alpha + 1},$ $h^* = \ell_1(x^*) = \ell_2(x^*) = \ell_3(x^*) = -\frac{2\alpha + 2\beta - 1}{2\alpha + 1}.$

To verify that this kink is the global minimizer, note first that $x^* \in (0,1)$ when $\alpha > 0.5$, hence it is interior. Next, the subgradient optimality condition for convex max-of-lines problems requires $(0,0) \in \partial f(x^*,e^*)$. At (x^*,e^*) the active lines have slopes that straddle zero in both coordinates:

$$\partial_x \ell_1 = -(2\alpha + 2\beta - 1) < 0, \quad \partial_x \ell_2 = \frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta} \left(e^* + \frac{1}{2} \right) > 0, \quad \partial_x \ell_3 = 1 - 2e^* = 1 - 2(\alpha + \beta) < 0,$$

and

$$\partial_e \ell_1 = 0, \qquad \partial_e \ell_2 = \frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta} x^* > 0, \qquad \partial_e \ell_3 = -2x^* < 0.$$

Since 0 lies in the convex hull of the active slopes in both x and e, we have $(0,0) \in \partial f(x^*,e^*)$, so the interior triple intersection is the global minimizer; no boundary check is needed.

$$\gamma_0 = M^{-(\alpha+\beta)}, \quad M^* \approx \mathfrak{f}^{1/(2\alpha+1)}, \quad R\left(M^*, \frac{\mathfrak{f}}{M^*}\right) \approx \mathfrak{f}^{-\frac{2\alpha+2\beta-1}{2\alpha+1}}.$$

E.2.2 $\alpha < 0.5, \beta < 0.5$ (Phase Ab)

We start from

$$R(M, N, \gamma_0) = (M^{\alpha} N \gamma_0)^{-\frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta}} + M^{-(2\alpha + 2\beta - 1)} + \gamma_0^2 M^{2 - 2\alpha}.$$

2317 Substitute

$$\gamma_0 = M^{-e}, \qquad N = \frac{\mathfrak{f}}{M}, \qquad M = \mathfrak{f}^x,$$

so that, up to constant factors,

$$R \approx \mathfrak{f}^{\max\{\ell_1(x),\ell_2(x),\ell_3(x)\}},$$

where

$$\begin{split} \ell_1(x) &= -\left(2\alpha + 2\beta - 1\right)x, \\ \ell_2(x) &= -\frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta}\left(\alpha - e - 1\right)x - \frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta}, \\ \ell_3(x) &= \left(2 - 2\alpha - 2e\right)x. \end{split}$$

We minimize the convex, piecewise–linear function $f(x,e) = \max_i \ell_i(x,e)$ over $x \in (0,1)$ and $e \in \mathbb{R}$. Under our standing assumptions $\alpha + \beta > 0.5$ and $\beta < \alpha + 0.5$, the only admissible triple intersection is $\{\ell_1, \ell_2, \ell_3\}$. Solving $\ell_1 = \ell_3$ and $\ell_2 = \ell_3$ gives

$$e^* = \beta + \frac{1}{2},$$
 $x^* = \frac{1}{2},$ $h^* = \ell_1(x^*) = \ell_2(x^*) = \ell_3(x^*) = -\frac{2\alpha + 2\beta - 1}{2}.$

To certify optimality, note that $x^* \in (0,1)$ (since $x^* = \frac{1}{2}$) and check the subgradient condition $(0,0) \in \partial f(x^*,e^*)$. At (x^*,e^*) the active lines have slopes straddling zero in both coordinates:

$$\partial_x \ell_1 = -(2\alpha + 2\beta - 1) < 0, \quad \partial_x \ell_2 = \frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta} (e^* + 1 - \alpha) > 0, \quad \partial_x \ell_3 = 2 - 2\alpha - 2e^* = 1 - 2(\alpha + \beta) < 0,$$

and

$$\partial_e \ell_1 = 0, \qquad \partial_e \ell_2 = \frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta} x^* > 0, \qquad \partial_e \ell_3 = -2x^* < 0.$$

Hence 0 lies in the convex hull of the active slopes in both variables, so the interior kink (x^*, e^*) is the global minimizer; no boundary check is required.

$$\gamma_0 = M^{-(\beta+0.5)}, \quad M^{\star} \approx \mathfrak{f}^{1/2}, \quad R\left(M^{\star}, \frac{\mathfrak{f}}{M^{\star}}\right) \approx \mathfrak{f}^{-\frac{2\alpha+2\beta-1}{2}}.$$

E.2.3
$$\alpha < 0.5, 0.5 < \beta < \alpha + 0.5$$
 (Phase Ac)

We start from

$$R(M, N, \gamma_0) = (M^{\alpha} N \gamma_0)^{-\frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta}} + M^{-2\alpha} + \gamma_0^2 M^{2 - 2\alpha}.$$

Substitute

$$\gamma_0 = M^{-e}, \qquad N = \frac{\mathfrak{f}}{M}, \qquad M = \mathfrak{f}^x,$$

so that, up to constant factors,

$$R \, \eqsim \, \mathfrak{f}^{\max\{\ell_1(x),\,\ell_2(x),\,\ell_3(x)\}}$$

where

$$\ell_1(x) = -2\alpha x,$$

$$\ell_2(x) = -\frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta} \left(\alpha - e - 1\right) x - \frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta},$$

$$\ell_3(x) = (2 - 2\alpha - 2e) x.$$

We minimize the convex, piecewise–linear objective $f(x,e) = \max_i \ell_i(x,e)$ over $x \in (0,1)$ and $e \in \mathbb{R}$. In the regime $\alpha + \beta > 0.5$ and $\beta < \alpha + 0.5$ (with $\alpha < 0.5 < \beta$), the only admissible triple intersection is $\{\ell_1, \ell_2, \ell_3\}$. Solving $\ell_1 = \ell_3$ and $\ell_2 = \ell_3$ yields

$$e^* = 1,$$
 $x^* = \frac{2\alpha + 2\beta - 1}{-4\alpha\beta + 6\alpha + 4\beta - 2},$ $h^* = \ell_1(x^*) = \ell_2(x^*) = \ell_3(x^*) = -\frac{2\alpha(2\alpha + 2\beta - 1)}{-4\alpha\beta + 6\alpha + 4\beta - 2}.$

One checks that the denominator is positive in this regime and exceeds the positive numerator $2\alpha + 2\beta - 1$, hence $x^* \in (0,1)$.

Interior optimality. At (x^*, e^*) the active lines' slopes straddle zero in both coordinates:

$$\partial_x \ell_1 = -2\alpha < 0, \quad \partial_x \ell_2 = \frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta} (e^* + 1 - \alpha) > 0, \quad \partial_x \ell_3 = 2 - 2\alpha - 2e^* = -2\alpha < 0,$$

2376 and

$$\partial_e \ell_1 = 0, \qquad \partial_e \ell_2 = \frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta} x^* > 0, \qquad \partial_e \ell_3 = -2x^* < 0.$$

Thus $(0,0) \in \partial f(x^*,e^*)$ and, with $x^* \in (0,1)$, the interior kink is the global minimizer; no boundary check is required.

$$\gamma_0 = M^{-1}, \quad M^{\star} \approx \mathfrak{f}^{\frac{2\alpha + 2\beta - 1}{-4\alpha\beta + 6\alpha + 4\beta - 2}}, \quad R(M^{\star}, \frac{\mathfrak{f}}{M^{\star}}) \approx \mathfrak{f}^{-\frac{2\alpha(2\alpha + 2\beta - 1)}{-4\alpha\beta + 6\alpha + 4\beta - 2}}.$$

E.2.4 $\alpha > 0.5, 0.5 < \beta < \alpha + 0.5$ (Phase Ba)

We start from

$$R(M, N, \gamma_0) = \left(M^{1/2} N \gamma_0\right)^{-\frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta}} + \left(M^{\frac{6\alpha - 1}{4\alpha - 2}} N \gamma_0\right)^{-\frac{2(2\alpha - 1)}{2\alpha + 1}} + M^{-2\alpha} + \gamma_0^2 M.$$

Substitute $\gamma_0 = M^{-e}$, $N = \mathfrak{f}/M$, $M = \mathfrak{f}^x$. Then, up to \mathfrak{f} -independent factors,

$$R \approx \mathfrak{f}^{\max_{i=1,\dots,4}\ell_i(x,e)}$$

where

$$\begin{split} \ell_1(x) &= -2\alpha \, x, \\ \ell_2(x) &= \frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta} \left(e + \frac{1}{2} \right) x \, - \, \frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta} \\ \ell_3(x) &= \left(\frac{2(2\alpha - 1)}{2\alpha + 1} \, e - 1 \right) x \, - \, \frac{2(2\alpha - 1)}{2\alpha + 1}, \\ \ell_4(x) &= (1 - 2e) \, x. \end{split}$$

We minimize the convex, piecewise–linear function $f(x,e) = \max_i \ell_i(x,e)$ over $x \in (0,1), e \in \mathbb{R}$. In the regime $\alpha > 0.5, \ \beta > 0.5$, the only admissible interior kink with three active lines is $\{\ell_2,\ell_3,\ell_4\}$. Solving $\ell_2 = \ell_4$ and $\ell_3 = \ell_4$ yields

$$e^* = \frac{2\alpha + 4\beta - 1}{4\beta}, \qquad x^* = \frac{\beta}{\alpha + \beta}, \qquad h^* = \ell_2(x^*, e^*) = \ell_3(x^*, e^*) = \ell_4(x^*, e^*) = -\frac{2\alpha + 2\beta - 1}{2\alpha + 2\beta}.$$

Interior optimality. First, $x^* \in (0,1)$ since $\alpha, \beta > 0.5$. Second, the subgradient condition $(0,0) \in \partial f(x^*, e^*)$ holds because the active slopes straddle zero in both variables:

$$\partial_x \ell_2 = \frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta} \left(e^* + \frac{1}{2} \right) > 0, \quad \partial_x \ell_3 = \frac{2(2\alpha - 1)}{2\alpha + 1} e^* - 1 < 0, \quad \partial_x \ell_4 = 1 - 2e^* = \frac{1 - 2\alpha - 2\beta}{2\beta} < 0,$$

and

$$\partial_e \ell_2 = \frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta} x^* > 0, \qquad \partial_e \ell_3 = \frac{2(2\alpha - 1)}{2\alpha + 1} x^* > 0, \qquad \partial_e \ell_4 = -2x^* < 0.$$

Hence (x^*, e^*) is the global minimizer among interior points. It remains to exclude ℓ_1 at (x^*, e^*) :

$$\ell_1(x^*) = -2\alpha \frac{\beta}{\alpha + \beta} \le -\frac{2\alpha + 2\beta - 1}{2(\alpha + \beta)} = h^*,$$

since $4\alpha\beta - 2\alpha - 2\beta + 1 = 4(\alpha - \frac{1}{2})(\beta - \frac{1}{2}) \ge 0$ for $\alpha, \beta > 0.5$. Therefore the triple intersection $\{\ell_2, \ell_3, \ell_4\}$ is the global optimum.

$$\boxed{\gamma_0 = \boldsymbol{M}^{-\frac{2\alpha + 4\beta - 1}{4\beta}}, \quad \boldsymbol{M}^\star \ \eqsim \ \mathfrak{f}^{\frac{\beta}{\alpha + \beta}}, \quad \boldsymbol{R}\!\left(\boldsymbol{M}^\star, \frac{\mathfrak{f}}{\boldsymbol{M}^\star}\right) \ \eqsim \ \mathfrak{f}^{-\frac{2\alpha + 2\beta - 1}{2\alpha + 2\beta}}.}$$

E.2.5 $\alpha < 0.5, \ \beta > \alpha + 0.5 \ (Phase Ad)$

Recall the loss formula (82)

$$R(M, N, \gamma_0) = \left(\max \left(1 - \kappa \gamma_0 M^{\alpha} N, 0 \right) \right)^{\frac{2(2\alpha + 2\beta - 1)}{2\beta - 2\alpha - 1}} + \gamma_0^2 M^{2 - 2\alpha} + M^{-2\alpha}.$$

- Note that the drift term vanishes at $N = (\gamma_0 M^{\alpha})^{-1}$.
- Let $\gamma_0 = M^{-e}$. Note that because of the approximation error $M^{-2\alpha}$, there is no gain from setting e bigger than 1. So we will only consider the case $e \le 1$. In that case, loss is a constant scale before $N = M^{e-\alpha}$, and it drops to the scale of $M^{-2e-2\alpha+2}$.
- Since a constant scale loss cannot be compute-optimal, the loss $M^{-2e-2\alpha+2}$ at $N = M^{e-\alpha}$ will be a candidate for the compute-optimal point. In that case $\mathfrak{f}=MN=M^{1+e-\alpha}$ holds and it leads to
- $M = \int_{1+e-\alpha}^{1} dx$. So the loss $M^{-2e-2\alpha+2}$ has the size $\int_{1+e-\alpha}^{-2e-2\alpha+2} dx$.
- Since e=1 minimizes $\frac{-2e-2\alpha+2}{1+e-\alpha}$, $\gamma_0=M^{-1}$ is the optimal learning rate. This leads to the follow-ing result.

$$\gamma_0 = M^{-1}, \quad M^* \approx \mathfrak{f}^{\frac{1}{2-\alpha}}, \quad R\left(M^*, \frac{\mathfrak{f}}{M^*}\right) \approx \mathfrak{f}^{-\frac{2\alpha}{2-\alpha}}.$$

- E.2.6 $\alpha > 0.5, \ \beta > \alpha + 0.5 \ (Phase Bb)$
- Recall the loss formula (83)

$$R(M, N, \gamma_0) = \left(\max\left(1 - \kappa \gamma_0 M^{0.5} N, 0\right)\right)^{\frac{2(2\alpha + 2\beta - 1)}{2\beta - 2\alpha - 1}} + M^{-\frac{6\alpha - 1}{2\alpha + 1}} (N\gamma_0)^{-\frac{2(2\alpha - 1)}{2\alpha + 1}} + \gamma_0^2 M + M^{-2\alpha}.$$

- Note that the first term vanishes at $N = (\gamma_0 M^{\alpha})^{-1}$. At that point second term becomes
- $M^{-\frac{6\alpha-1}{2\alpha+1}}(N\gamma_0)^{-\frac{2(2\alpha-1)}{2\alpha+1}} \approx M^{-\frac{4\alpha}{2\alpha+1}}.$
- As we optimize over three parameters N, M, γ_0 , and one constraint $\mathfrak{f} = MN$, we have two degrees of freedom. So this means three terms may balance together at the compute-optimal point.
- The first possible case is the balance of the first three terms, and in this case, $\gamma_0^2 M = M^{-\frac{4\alpha}{2\alpha+1}}$ and
- $N = (\gamma_0 M^{\alpha})^{-1}$ must hold. Here, the loss is $M^{-\frac{4\alpha}{2\alpha+1}}$ and $\mathfrak{f} = MN = M^{\frac{2\alpha+1}{4\alpha+1}}$ holds, so the loss is
- The second possible case is the balance of the last three terms, and after solving the equations, the loss is $f^{-\frac{2\alpha}{2\alpha+1}}$.
 - The first case has a steeper decay, so it is the compute-optimal. This leads to the following result.

$$\boxed{\gamma_0 = M^{-\frac{6\alpha+1}{4\alpha+2}}, \quad M^\star \ \eqsim \ \mathfrak{f}^{\frac{2\alpha+1}{4\alpha+1}}, \quad R\left(M^\star, \frac{\mathfrak{f}}{M^\star}\right) \ \eqsim \ \mathfrak{f}^{-\frac{4\alpha}{4\alpha+1}}.}$$

DERIVATION FOR THE STABLE-DECAY SCHEDULING

We set the learning rate as $\gamma_k = \gamma_0 f(k)$. Previously, we considered the constant-learning-rate case $(f \equiv 1)$. In this section, we start with a general decaying learning rate by taking f to be a decreasing function, and then substitute the stable-decay scheduling. Throughout, for simplicity, we assume $\alpha > 0.5$ and $\beta < 0.5$ (Phase Aa).

Recall the implicit integral equation (17):

$$L(N) = \|\boldsymbol{H}^{1/2}\boldsymbol{w}_{\perp}\|^{2} + \sum_{i=1}^{M} r_{i}(0) \exp\left(-\frac{4\lambda_{i}\gamma_{0}}{\pi} \int_{0}^{N} \frac{f(u)}{\sqrt{L(u)}} du\right)$$
(84)

$$+\frac{2\gamma_0^2}{\pi} \sum_{i=1}^{M} V_i \int_0^N \exp\left(-\frac{4\lambda_i \gamma_0}{\pi} \int_z^N \frac{f(u)}{\sqrt{L(u)}} du\right) f(z)^2 dz.$$
 (85)

Also recall Equation 18 and 19.

$$L^{\text{drift}}(N) = \sum_{i=1}^{M} r_i(0) e^{-\frac{4\lambda_i \gamma_0}{\pi} \int_0^N \frac{f(u)}{\sqrt{L(u)}} du}, \quad L^{\text{noise}}(N) = \frac{2\gamma_0^2}{\pi} \sum_{i=1}^M V_i \int_0^N e^{-\frac{4\lambda_i \gamma_0}{\pi} \int_z^N \frac{f(u)}{\sqrt{L(u)}} du} f(z)^2 dz.$$
(86)

2484
$$L(N) = \|\boldsymbol{H}^{1/2}\boldsymbol{w}_{\perp}\|^2 + L^{\text{drift}}(N) + L^{\text{noise}}(N).$$
 (87)

Recall also the drift/approximation transformation (26):

$$L^{\text{drift}}(N) + \|\boldsymbol{H}^{1/2}\boldsymbol{w}_{\perp}\|^{2} \approx M^{-(2\alpha+2\beta-1)} + (M^{0.5}Q_{L}(N))^{-\frac{2\alpha+2\beta-1}{2\alpha}},$$
$$Q_{L}(z) := \frac{4\gamma_{0}}{\pi} \int_{0}^{z} \frac{f(u)}{\sqrt{L(u)}} du.$$

Hence,

$$L(N) \approx M^{-(2\alpha+2\beta-1)} + \left(M^{0.5}Q_L(N)\right)^{-\frac{2\alpha+2\beta-1}{2\alpha}}$$

$$+ \frac{2\gamma_0^2}{\pi} \sum_{i=1}^{M} V_i \int_0^N \exp\left(-\frac{4\lambda_i\gamma_0}{\pi} \int_{-\infty}^N \frac{f(u)}{\sqrt{L(u)}} du\right) f(z)^2 dz.$$
(89)

Remark 2 (Early-iteration proxy). In early iterations the drift term $\left(M^{0.5}Q_L(N)\right)^{-\frac{2\alpha+2\beta-1}{2\alpha}}$ dominates. Solving $L(N) \approx \left(M^{0.5}Q_L(N)\right)^{-\frac{2\alpha+2\beta-1}{2\alpha}}$ yields

$$L(N) \ \approx \ \left(M^{0.5} \gamma_0 F(N) \right)^{-\frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta}}, \qquad F(N) := \int_0^N f(u) \, du.$$

Now we move on to stable-decay scheduling.

Stable-decay schedule. We henceforth consider the following stable-decay learning-rate schedule:

$$f(k) = \begin{cases} 1, & k \le pN, \\ (1 + \tau(k - pN))^{-c}, & k > pN, \end{cases} \quad p, c \in (0, 1), \ \tau > 0.$$

That is, the learning rate is constant for the first pN steps and then decays polynomially with exponent c over the remaining (1-p)N steps.

Remark 3. Note that f depends on the total training steps N. To be precise, we have to represent it as f_N , but for simplicity, we write it as f throughout the analysis.

First, we will make an upper bound on the noise term under stable-decay scheduling.

F.1 UPPER BOUND OF THE NOISE TERM

Fix p < q < 1 close to 1 and split $L^{\text{noise}}(N)$ as

$$\begin{split} L^{\text{noise}}(N) &= \frac{2\gamma_0^2}{\pi} \sum_{i=1}^M V_i \int_0^{qN} \exp\Bigl(-\frac{4\lambda_i \gamma_0}{\pi} \int_z^N \frac{f(u)}{\sqrt{L(u)}} \, du\Bigr) f(z)^2 \, dz \\ &+ \frac{2\gamma_0^2}{\pi} \sum_{i=1}^M V_i \int_{qN}^N \exp\Bigl(-\frac{4\lambda_i \gamma_0}{\pi} \int_z^N \frac{f(u)}{\sqrt{L(u)}} \, du\Bigr) f(z)^2 \, dz =: T_{\leq qN} + T_{>qN}. \end{split}$$

Bounding $T_{>qN}$. Note that f(N) = f(z) holds for qN < z < N. So

$$\int_{qN}^{N} \exp\left(-\frac{4\lambda_i \gamma_0}{\pi} \int_{z}^{N} \frac{f(u)}{\sqrt{L(u)}} du\right) f(z)^2 dz \approx f(N)^2 \int_{qN}^{N} \exp\left(-\frac{4\lambda_i \gamma_0}{\pi} \int_{z}^{N} \frac{f(u)}{\sqrt{L(u)}} du\right) dz.$$

For q sufficiently close to 1, there exist constants $c_0, c_1 > 0$ such that for qN < z < N

$$c_0 \frac{(N-z)f(N)}{\sqrt{L(N)}} \leq \int_z^N \frac{f(u)}{\sqrt{L(u)}} du \leq c_1 \frac{(N-z)f(N)}{\sqrt{L(N)}}.$$

Therefore,

$$T_{>qN} \le \frac{2\gamma_0^2}{\pi} f(N)^2 \sum_{i=1}^M V_i \int_{qN}^N \exp\left(-\frac{4\lambda_i \gamma_0}{\pi} c_0 \frac{(N-z)f(N)}{\sqrt{L(N)}}\right) dz$$

$$= \frac{2\gamma_0^2}{\pi} f(N)^2 \sum_{i=1}^M V_i \frac{\pi \sqrt{L(N)}}{4\lambda_i \gamma_0 c_0 f(N)} = \gamma_0 f(N) \sqrt{L(N)} \sum_{i=1}^M \frac{V_i}{\lambda_i}$$

$$= \gamma_0 f(N) \sqrt{L(N)} \operatorname{Tr}(\operatorname{diag}(K)^{1/2}) = \gamma_0 f(N) \sqrt{L(N)} M^{0.5}.$$

To summarize, we have

$$T_{>qN} \lesssim \gamma_0 f(N) \sqrt{L(N)} M^{0.5} \approx \gamma_0 M^{1/2} N^{-c} \sqrt{L(N)}$$

Bounding
$$T_{\leq qN}$$
. Let $Q(z,N) = \frac{4\gamma_0}{\pi} \int_z^N \frac{f(u)}{\sqrt{L(u)}} du$. Then

$$T_{\leq qN} = \frac{2\gamma_0^2}{\pi} \sum_{i=1}^{M} (w_i^\mathsf{T} \mathbf{K}_\sigma K u_i) \int_0^{qN} e^{-\frac{4\gamma_0}{\pi} \lambda_i(\overline{\mathbf{K}}) \int_z^N \frac{f(u)}{\sqrt{L(u)}} du} f(z)^2 dz$$

$$= \frac{2\gamma_0^2}{\pi} \sum_{i=1}^{M} \mathrm{Tr}(\mathbf{K}_\sigma K u_i w_i^\mathsf{T}) \int_0^{qN} e^{-\lambda_i(\overline{\mathbf{K}})Q(z,N)} f(z)^2 dz$$

$$= \frac{2\gamma_0^2}{\pi} \int_0^{qN} \sum_{i=1}^{M} \mathrm{Tr}(\mathbf{K}_\sigma K u_i w_i^\mathsf{T}) e^{-\lambda_i(\overline{\mathbf{K}})Q(z,N)} f(z)^2 dz$$

$$= \frac{2\gamma_0^2}{\pi} \int_0^{qN} \mathrm{Tr}(\mathbf{K}_\sigma K \sum_{i=1}^{M} e^{-\lambda_i(\overline{\mathbf{K}})Q(z,N)} u_i w_i^\mathsf{T}) f(z)^2 dz$$

$$= \frac{2\gamma_0^2}{\pi} \int_0^{qN} \mathrm{Tr}(\mathbf{K}_\sigma K e^{-\overline{\mathbf{K}}Q(z,N)}) f(z)^2 dz.$$

Using $\arcsin x \approx x$ approximation on $\mathbf{K}_{\sigma} = \arcsin(\operatorname{diag}(K)^{-1/2} \cdot K \cdot \operatorname{diag}(K)^{-1/2})$, we get $\operatorname{Tr}(\mathbf{K}_{\sigma}Ke^{-\overline{\mathbf{K}}Q(z,N)}) = \operatorname{Tr}(\mathbf{K}_{\sigma}SH^{1/2}e^{-\overline{\mathbf{K}}_{1}Q(z,N)}H^{1/2}S^{\mathsf{T}})$

$$\operatorname{Tr}(\boldsymbol{K}_{\sigma}\boldsymbol{K}e^{-\boldsymbol{K}_{\sigma}\boldsymbol{S}}H^{T}e^{-\boldsymbol{K}_{\sigma}\boldsymbol{S}}H^{T}e^{-\boldsymbol{K}_{1}Q(z,N)}) \approx \operatorname{Tr}(\boldsymbol{K}_{1}^{2}e^{-\boldsymbol{K}_{1}Q(z,N)}).$$

Using same contour representation method and deterministic approximation with Section D.2.2 we get

$$T_{\leq qN} \approx \frac{2\gamma_0^2}{\pi} \int_0^{qN} \operatorname{Tr}(\overline{\boldsymbol{K}}_1^2 e^{-\overline{\boldsymbol{K}}_1 Q(z,N)}) f(z)^2 dz$$

$$= \frac{2\gamma_0^2}{\pi} \int_0^{qN} \operatorname{Tr}\left(\frac{-1}{2\pi i} \oint_{\Gamma_2} z_1^2 e^{-Q(z,N)z_1} (\overline{\boldsymbol{K}}_1 - z_1 \boldsymbol{I})^{-1} dz_1\right) f(z)^2 dz$$

$$\approx \frac{2\gamma_0^2}{\pi} \int_0^{qN} \operatorname{Tr}\left(\frac{-1}{2\pi i} \oint_{\Gamma} p_d^2 z_1^2 e^{-p_d Q(z,N)z_1} \mathcal{R}(z_1) dz_1\right) f(z)^2 dz$$

$$\approx \frac{2\gamma_0^2}{\pi} M \int_0^{qN} \operatorname{Tr}\left(\frac{-1}{2\pi i} \oint_{\Gamma} z_1^2 e^{-p_d Q(z,N)z_1} \mathcal{R}(z_1) dz_1\right) f(z)^2 dz$$

Adopting the method in Paquette et al. (2024) same as Section I.1, we get

$$\operatorname{Tr}\left(\frac{-1}{2\pi i} \oint_{\Gamma} z_1^2 e^{-p_d Q(z,N)z_1} \mathcal{R}(z_1) dz_1\right) \approx (p_d Q(z,N))^{-2+1/(2\alpha)} \approx (M^{1/2} Q(z,N))^{-2+1/(2\alpha)}.$$

It leads to

$$T_{\leq qN} \approx \frac{2\gamma_0^2}{\pi} M \int_0^{qN} (M^{1/2}Q(z,N))^{-2+1/(2\alpha)} f(z)^2 dz$$
$$\approx \gamma_0^2 M^{1/(4\alpha)} \int_0^{qN} (Q(z,N))^{-2+1/(2\alpha)} f(z)^2 dz$$

Finally,
$$\gamma_0^2 M^{1/(4\alpha)} \int_0^{qN} (Q(z,N))^{-2+1/(2\alpha)} f(z)^2 \, dz$$

$$> \gamma_0^{1/(2\alpha)} M^{1/(4\alpha)} \int_0^{pN} \frac{f(z)^2}{\left(\int_z^N \frac{f(u)}{\sqrt{L(u)}} du\right)^{2-1/(2\alpha)}} \, dz + \gamma_0^{1/(2\alpha)} M^{1/(4\alpha)} \int_{pN}^{qN} \frac{f(z)^2}{\left(\int_z^N \frac{f(u)}{\sqrt{L(u)}} du\right)^{2-1/(2\alpha)}} \, dz$$

$$> \gamma_0^{1/(2\alpha)} M^{1/(4\alpha)} \int_0^{pN} \frac{1}{\left(\frac{p^{N-z}}{\sqrt{L(u)}} + \frac{1}{\sqrt{L(pN)}} \int_{pN}^N f(u) du\right)^{2-1/(2\alpha)}} \, dz$$

$$> \gamma_0^{1/(2\alpha)} M^{1/(4\alpha)} \int_0^{qN} \frac{1}{\left(\frac{p^{N-z}}{\sqrt{L(pN)}} + \frac{1}{\sqrt{N}} \int_{pN}^N f(u) du\right)^{2-1/(2\alpha)}} \, dz$$

$$> \gamma_0^{1/(2\alpha)} M^{1/(4\alpha)} \int_p^{qN} \frac{f(z)^2}{\left(\frac{1}{\sqrt{L(pN)}} \int_{qN}^N f(u) du\right)^{2-1/(2\alpha)}} \, dz$$

$$> \gamma_0^{1/(2\alpha)} M^{1/(4\alpha)} \int_0^{pN} \frac{1}{\left(\frac{p^{N-z}}{\sqrt{L(pN)}} + \frac{N^{1-c}}{\sqrt{L(pN)}}\right)^{2-1/(2\alpha)}} \, dz + \gamma_0^{1/(2\alpha)} M^{1/(4\alpha)} \int_{pN}^{qN} \frac{f(z)^2}{\left(\frac{N^{1-c}}{\sqrt{L(pN)}}\right)^{2-1/(2\alpha)}} \, dz$$

$$> \gamma_0^{1/(2\alpha)} M^{1/(4\alpha)} \sqrt{L(0)} \left(\frac{N^{1-c}}{\sqrt{L(pN)}} \right)^{1/(2\alpha)-1} - \left(pN + \frac{N^{1-c}}{\sqrt{L(pN)}} \right)^{1/(2\alpha)-1} \right)$$

$$+ \gamma_0^{1/(2\alpha)} M^{1/(4\alpha)} N^{\max(1-2c,0)} \left(\frac{N^{1-c}}{\sqrt{L(pN)}} \right)^{1/(2\alpha)-2}$$

$$\ge \gamma_0^{1/(2\alpha)} M^{1/(4\alpha)} N^{-(1-c)(1-1/(2\alpha))} L(pN)^{(1/2-1/(4\alpha))} \lesssim \gamma_0^{1/(2\alpha)} M^{1/(4\alpha)} N^{-(1-c)(1-1/(2\alpha))}$$

So we have

$$T_{\leq qN} \lesssim \gamma_0^{1/(2\alpha)} M^{1/(4\alpha)} N^{-(1-c)(1-1/(2\alpha))}$$
.

F.2 COMBINING TERMS

Combining the bounds,

$$L(N) \lesssim M^{-(2\alpha+2\beta-1)} + \left(M^{0.5}\gamma_0 N\right)^{-\frac{2(2\alpha+2\beta-1)}{2\alpha+1-2\beta}} + \gamma_0 M^{0.5} N^{-c} \sqrt{L(N)} + \gamma_0^{\frac{1}{2\alpha}} M^{\frac{1}{4\alpha}} N^{-(1-c)(1-\frac{1}{2\alpha})}.$$

We replaced the drift part with $(M^{0.5}\gamma_0 N)^{-\frac{2(2\alpha+2\beta-1)}{2\alpha+1-2\beta}}$ temporarily based on Remark 2, and justify this on our selected parameters in Remark 4. Solving the inequality asymptotically yields

$$L(N) \lesssim M^{-(2\alpha+2\beta-1)} + \left(M^{0.5}\gamma_0 N\right)^{-\frac{2(2\alpha+2\beta-1)}{2\alpha+1-2\beta}} + \gamma_0^2 M N^{-2c} + \gamma_0^{\frac{1}{2\alpha}} M^{\frac{1}{4\alpha}} N^{-(1-c)(1-\frac{1}{2\alpha})}.$$

Finally, substituting $\gamma_0 = M^{-e}$ and $N = \mathfrak{f}/M$ yields

$$R(M,\mathfrak{f}) \lesssim M^{-(2\alpha+2\beta-1)} + \left(M^{-e-0.5}\mathfrak{f}\right)^{-\frac{2(2\alpha+2\beta-1)}{2\alpha+1-2\beta}} + M^{1+2c-2e}\mathfrak{f}^{-2c} + M^{\frac{1}{4\alpha}-\frac{e}{2\alpha}+(1-c)(1-\frac{1}{2\alpha})}\mathfrak{f}^{-(1-c)(1-\frac{1}{2\alpha})}.$$

Optimizing over M gives a bound of the form $R(M^*,\mathfrak{f}) \leq \mathfrak{f}^{-h(\alpha,\beta,c,e)}$, and we then optimize over c,e to maximize $h(\alpha,\beta,c,e)$.

F.3 OPTIMIZING OVER c, e TO MAXIMIZE $h(\alpha, \beta, c, e)$

Assume throughout $\alpha > 0.5$, $\beta < 0.5$, and $2\alpha + 2\beta > 1$. Consider the upper bound

$$R_U(M,\mathfrak{f}) = M^{-(2\alpha+2\beta-1)} + \left(M^{-e-0.5}\,\mathfrak{f}\right)^{-\frac{2(2\alpha+2\beta-1)}{2\alpha+1-2\beta}} + M^{1+2c-2e}\,\mathfrak{f}^{-2c} + M^{\frac{1}{4\alpha}-\frac{e}{2\alpha}+(1-c)\left(1-\frac{1}{2\alpha}\right)}\,\mathfrak{f}^{-(1-c)\left(1-\frac{1}{2\alpha}\right)}.$$

For large f, define

$$R_{\min}(\mathfrak{f}) := \min_{M>0} R_U(M,\mathfrak{f}).$$

We show $R_{\min}(\mathfrak{f}) \approx \mathfrak{f}^{h^{\star}(\alpha,\beta)}$ with $h^{\star}(\alpha,\beta) < 0$, and identify $c^{\star}(\alpha,\beta)$, $e^{\star}(\alpha,\beta)$, and $M = \mathfrak{f}^{m^{\star}(\alpha,\beta)}$.

Logarithmic reduction to exponent balancing

Let $M = \mathfrak{f}^m$ with $m \in \mathbb{R}$. Writing each term as \mathfrak{f}^{L_i} gives

$$L_1(m) = -(2\alpha + 2\beta - 1) m, (90)$$

$$L_2(m,e) = -\frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta} + \frac{2(2\alpha + 2\beta - 1)}{2\alpha + 1 - 2\beta} m(e + 0.5), \tag{91}$$

$$L_3(m, c, e) = m(1 + 2c - 2e) - 2c, (92)$$

$$L_4(m, c, e) = m\left(\frac{1}{4\alpha} - \frac{e}{2\alpha} + (1 - c)\left(1 - \frac{1}{2\alpha}\right)\right) - (1 - c)\left(1 - \frac{1}{2\alpha}\right). \tag{93}$$

Thus minimizing R_U is equivalent to

$$\min_{m,e \in \mathbb{R}, \ 0 < c < 1} \max\{L_1, L_2, L_3, L_4\}. \tag{94}$$

Introduce $h \in \mathbb{R}$ and rewrite as

$$\min_{m,c,e,h} h \quad \text{s.t.} \quad L_i(m,c,e) \le h \ (i = 1, 2, 3, 4), \ \ 0 < c < 1. \tag{95}$$

At an interior optimum (0 < c < 1), constraints equalize:

$$L_1 = L_2 = L_3 = L_4 = h. (96)$$

Solving the equality yields

$$c^* = \frac{-8\alpha\beta + 2\alpha + 2\beta - 1}{16\alpha^2 + 8\alpha\beta - 6\alpha - 2\beta + 1},\tag{97}$$

$$e^* = \frac{8\alpha^2 + 16\alpha\beta - 4\alpha - 4\beta + 1}{2(4\alpha - 1)},\tag{98}$$

$$m^* = \frac{2(4\alpha - 1)}{16\alpha^2 + 8\alpha\beta + 2\alpha - 2\beta - 1},\tag{99}$$

$$h^* = -\frac{2(4\alpha - 1)(2\alpha + 2\beta - 1)}{16\alpha^2 + 8\alpha\beta + 2\alpha - 2\beta - 1}.$$
 (100)

Feasibility. Since $\alpha > 0.5$, denominators are positive. The condition $c^* > 0$ is equivalent to

$$-8\alpha\beta + 2\alpha + 2\beta - 1 > 0 \iff \beta < \frac{2\alpha - 1}{2(4\alpha - 1)} := B^{\star}(\alpha),$$

which is stricter than $\beta < 0.5$. Moreover, $c^* < 1$ holds automatically for $\beta > 0$. Hence, the interior solution is feasible whenever

$$0.5 - \alpha < \beta < B^*(\alpha) \qquad \text{with} \quad B^*(\alpha) = \frac{2\alpha - 1}{2(4\alpha - 1)}. \tag{101}$$

In this band,

$$M = \mathfrak{f}^{m^*}, \qquad R_{\min}(\mathfrak{f}) \approx \mathfrak{f}^{h^*}$$

with m^*, h^* as in (99)–(100). Note $m^* > 0$ and $h^* < 0$.

Result As $\mathfrak{f} \to \infty$, the choice $M = \mathfrak{f}^{m^*}$ with

$$m^{\star} = \frac{2(4\alpha - 1)}{16\alpha^2 + 8\alpha\beta + 2\alpha - 2\beta - 1},$$

$$c^{\star} = \frac{-8\alpha\beta + 2\alpha + 2\beta - 1}{16\alpha^2 + 8\alpha\beta - 6\alpha - 2\beta + 1},$$

$$e^{\star} = \frac{8\alpha^2 + 16\alpha\beta - 4\alpha - 4\beta + 1}{2(4\alpha - 1)},$$

$$h^{\star} = -\frac{2(4\alpha - 1)(2\alpha + 2\beta - 1)}{16\alpha^2 + 8\alpha\beta + 2\alpha - 2\beta - 1}$$

is optimal for $\alpha > 0.5, 0.5 - \alpha < \beta < B^*(\alpha)$, where

$$B^{\star}(\alpha) = \frac{2\alpha - 1}{2(4\alpha - 1)}.$$

This choice minimizes $\max\{L_1, L_2, L_3, L_4\}$ in (94). Consequently,

$$R_{\min}(\mathfrak{f}) \approx \mathfrak{f}^{h^{\star}(\alpha,\beta)}$$
 with $h^{\star}(\alpha,\beta) < 0$.

Remark 4 (Justification on drift term conversion). Note that $M = \mathfrak{f}^{M^*}$ and $N = \mathfrak{f}^{1-M^*}$ holds for the selected parameters.

For pN iterations the stable-decay scheduling behaves same as the constant learning rate. Let N_0 be the crossover point in constant learning rate. Note that $N \gtrsim N_0$ holds, and N is asymptotically strictly bigger than N_0 . So $L(u) \lesssim \gamma_0^2 M + M^{-2\alpha - 2\beta + 1}$ holds for $u \geq N_0$.

Also for selected $\gamma_0 = M^{-e^*}$, $\gamma_0^2 M \gtrsim M^{-2\alpha - 2\beta + 1}$ holds.

So we have $L(u) \lesssim \gamma_0^2 M$ for $u \geq N_0$.

$$(M^{0.5}Q_L(N))^{-\frac{2\alpha+2\beta-1}{2\alpha}} = \left(M^{0.5}\frac{4\gamma_0}{\pi} \int_0^N \frac{f(u)}{\sqrt{L(u)}} du\right)^{-\frac{2\alpha+2\beta-1}{2\alpha}}$$

$$\lesssim \left(M^{0.5}\frac{4\gamma_0}{\pi} \int_{N_0}^{pN} \frac{f(u)}{\sqrt{L(u)}} du\right)^{-\frac{2\alpha+2\beta-1}{2\alpha}}$$

$$= \left(\gamma_0 M^{0.5} \int_{N_0}^{pN} \frac{1}{\sqrt{L(u)}} du\right)^{-\frac{2\alpha+2\beta-1}{2\alpha}}$$

$$= \left(\gamma_0 M^{0.5} \int_{N_0}^{pN} \frac{1}{\sqrt{\gamma_0^2 M}} du\right)^{-\frac{2\alpha+2\beta-1}{2\alpha}}$$

$$= (pN - N_0)^{-\frac{2\alpha+2\beta-1}{2\alpha}} = N^{-\frac{2\alpha+2\beta-1}{2\alpha}}$$

And for selected parameters $M=\mathfrak{f}^{M^\star},\,N=\mathfrak{f}^{1-M^\star},\,c^*,$ and $\gamma_0=M^{-e^*}$ following holds. $N^{-\frac{2\alpha+2\beta-1}{2\alpha}}\lesssim \gamma_0^2MN^{-2c^*}.$

As $(M^{0.5}Q_L(N))^{-\frac{2\alpha+2\beta-1}{2\alpha}} \lesssim \gamma_0^2 M N^{-2c^*}$, replacing the drift term with proxy does not makes problem.

F.4 SCHEDULING ON SGD

In this subsection, we explain that the scheduling does not lift the compute-optimal exponent of SGD in the Phase I and Phase II. Assume a bounded scheduling function f, and define $F(k) = \int_0^k f(z) \, dz$.

Ferbach et al. (2025) proved

$$R_f(M, N, \gamma_0) \gtrsim M^{-2\alpha + \max(0, 1 - 2\beta)} + (\gamma_0 F(N))^{-(2\alpha + 2\beta - 1)/(2\alpha)} + M^{-1}(\gamma_0 F(N))^{-1 + 1/(2\alpha)}$$

for the risk $R_f(M, N, \gamma_0)$ with general bounded scheduling function f.

Since f is bounded, we have $F(N) \lesssim N$. Therefore,

$$R_f(M, N, \gamma_0) \gtrsim M^{-2\alpha + \max(0, 1 - 2\beta)} + (\gamma_0 F(N))^{-(2\alpha + 2\beta - 1)/(2\alpha)} + M^{-1}(\gamma_0 F(N))^{-1 + 1/(2\alpha)}$$
$$\gtrsim M^{-2\alpha + \max(0, 1 - 2\beta)} + (\gamma_0 N)^{-(2\alpha + 2\beta - 1)/(2\alpha)} + M^{-1}(\gamma_0 N)^{-1 + 1/(2\alpha)}$$
$$\gtrsim R_1(M, N, \gamma_0),$$

where $R_1(M, N, \gamma_0)$ is the loss under a constant schedule $f \equiv 1$.

Thus, scheduling does not improve the compute-optimal exponent of SGD in Phase I and Phase II.

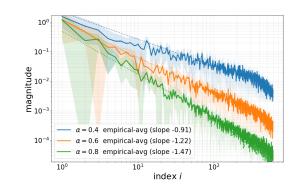


Figure 15: **Decay of gradient under the basis of** U. Colored solid lines show the average of gradients under the basis of U for the parameter $(\alpha, \beta) = (0.4, 0.5), (0.6, 0.5), (0.8, 0.5)$. On the legend, we only noted the α . The dotted line is fitted for the average of gradients, and we noted its slope in the legend. Slope is similar to 2α within error 0.13.

G ANALYSIS ABOUT HYPOTHESIS FOR THE POSITION OF THE BENEFICIAL AREA

In this section, we cover the analysis of stochastic gradient decay and target decay, which were deferred from Section 5.1.

We first examine the decaying structure of the stochastic gradient. Assume a feature vector \boldsymbol{x} is drawn from the distribution $\mathcal{N}(0, \boldsymbol{H})$, and its label is $y = \langle \boldsymbol{x}, \boldsymbol{w}^* \rangle$. Then the stochastic gradient for that feature vector is

$$g = (\langle Sx_t, \theta_{t-1} \rangle - y) Sx_t$$

= $Sxx^TS^T(\theta - \theta^*) - Sxx^Tw_{\perp}$.

Taking the expectation of the stochastic gradient and using $SHw_{\perp}=0$, we obtain

$$egin{aligned} \mathbb{E}[g] &= SHS^{\mathsf{T}}(heta - heta^*) - SHw_{\perp} \ &= SHS^{\mathsf{T}}(heta - heta^*). \end{aligned}$$

Lin et al. (2024) proved that the eigenvalues λ_i of SHS^T satisfy $\lambda_i \approx i^{-2\alpha}$. Let the eigenvalue decomposition of SHS^T be $SHS^T = U \operatorname{diag}(\lambda_i)U^T$. Then

$$\boldsymbol{U}^{\mathsf{T}}\mathbb{E}[\boldsymbol{g}] = \operatorname{diag}(\lambda_i) \, \boldsymbol{U}^{\mathsf{T}}(\boldsymbol{\theta} - \boldsymbol{\theta}^*),$$

which provides the intuition that $\mathbb{E}[g]$, expressed in the basis of the columns of U, decays as $i^{-2\alpha}$. Figure 15 shows that the expected gradient decays similarly to $i^{-2\alpha}$. Also, it is important to note that a larger α leads to a steeper gradient decay.

Next, we examine how the target decays in the basis of the columns of U. Using the heuristic $S^{\mathsf{T}}S \approx I$,

$$\boldsymbol{\theta}^* \ = \ (\boldsymbol{S}\boldsymbol{H}\boldsymbol{S}^\mathsf{T})^{-1}\boldsymbol{S}\boldsymbol{H}\boldsymbol{w}^* \ \approx \ (\boldsymbol{S}\boldsymbol{H}\boldsymbol{S}^\mathsf{T})^{-1}\boldsymbol{S}\boldsymbol{H}\boldsymbol{S}^\mathsf{T}\boldsymbol{S}\boldsymbol{w}^* \ = \ \boldsymbol{S}\boldsymbol{w}^*.$$

Since $SHS^T = U \operatorname{diag}(\lambda_i)U^T$, by the similarity between U and S in the early columns, $U^T\theta^* = U^TSw^*$ may exhibit a decay pattern similar to w^* on the early indices. The important part is how long $U^T\theta^*$ decays in conjunction with w^* . We observe it in Figure 3 in the main text, and discussed it in Section 5.1.

H SCALING LAW OF ADAM WITH HEURISTIC

First, we recall the Adam (Kingma & Ba, 2014) update and notation. For the stochastic gradient

$$g_k = (\langle Sx_k, \theta_k \rangle - y_k) Sx_k.$$

Adam maintains first and second moment estimates

$$m_k = \beta_1 m_{k-1} + (1 - \beta_1) g_k,$$

 $v_k = \beta_2 v_{k-1} + (1 - \beta_2) g_k^{\circ 2},$

with bias corrections $\hat{m}_k = m_k/(1-\beta_1^k)$, $\hat{v}_k = v_k/(1-\beta_2^k)$. The update is

$$oldsymbol{ heta}_{k+1} = oldsymbol{ heta}_k - \gamma_k \, \hat{oldsymbol{m}}_k \odot \left(\epsilon \mathbf{1} + \hat{oldsymbol{v}}_k\right)^{-1/2},$$

where \odot denotes elementwise multiplication and the (-1/2) power is taken elementwise; $\epsilon > 0$ is the usual damping (we will set $\epsilon = 0$ in the asymptotic analysis).

Xiao et al. (2024) proposed a heuristic for Adam: take β_2 sufficiently close to 1 so that the second moment can be treated as an expectation, and assume (in high dimensions) approximate independence between the coordinatewise square $(Sx_k)^{\odot 2}$ and the squared residual $(\langle Sx_k, \theta_{k-1} \rangle - y_k)^2$.

We present results under a same heuristic. In addition, Ferbach et al. (2025) prove that SGD with momentum obeys the same scaling law as SGD; motivated by this, we set $\beta_1 = 0$ and omit the first-moment term for simplicity.

Second-moment proxy and normalized update. Under the heuristic of Xiao et al. (2024),

$$egin{aligned} \hat{m{v}}_k &pprox \mathbb{E}\!\!\left[(m{S}m{x}_k)^{\odot 2} \left(\langle m{S}m{x}_k, m{ heta}_k
angle - y_k
ight)^2 \, \middle| \, \mathcal{F}_k
ight] \ &pprox \mathbb{E}\!\!\left[(m{S}m{x}_k)^{\odot 2}
ight] \, \mathbb{E}\!\!\left[\left(\langle m{S}m{x}_k, m{ heta}_k
angle - y_k
ight)^2 \, \middle| \, \mathcal{F}_k
ight] \ &= \mathrm{diag}(m{S}m{H}m{S}^\mathsf{T}) \cdot L(m{ heta}_k). \end{aligned}$$

Hence, the (elementwise) normalized update is

$$m{ heta}_{k+1} - m{ heta}_k ~pprox ~ rac{\left(\left\langle m{S}m{x}_k, m{ heta}_k
ight
angle - y_k
ight) m{S}m{x}_k}{\sqrt{\mathrm{diag}(m{S}m{H}m{S}^{\mathsf{T}}) \cdot L(m{ heta}_k)}}$$

One-step update formula. Recalling the Taylor expansion used for signSGD,

$$\mathbb{E}[q(\boldsymbol{\theta}_{k+1}) - q(\boldsymbol{\theta}_k) \mid \mathcal{F}_k] = \mathbb{E}[\langle \nabla q(\boldsymbol{\theta}_k), \boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}_k \rangle \mid \mathcal{F}_k] + \frac{1}{2} \mathbb{E}[\langle \nabla^2 q, (\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}_k)^{\otimes 2} \rangle \mid \mathcal{F}_k].$$

Gradient term:

$$\mathbb{E}\!\!\left[\left\langle \nabla q(\boldsymbol{\theta}_k),\,\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}_k \right\rangle \,\middle|\, \mathcal{F}_k \right] \;\approx\; \left\langle \nabla q(\boldsymbol{\theta}_k),\, \frac{\boldsymbol{S}\boldsymbol{H}\boldsymbol{S}^\mathsf{T}\,\boldsymbol{\theta}_k - \boldsymbol{S}\boldsymbol{H}\boldsymbol{w}^*}{\sqrt{\operatorname{diag}(\boldsymbol{S}\boldsymbol{H}\boldsymbol{S}^\mathsf{T})\cdot L(\boldsymbol{\theta}_k)}} \right\rangle.$$

Quadratic term:

$$\mathbb{E}\left[\left\langle \nabla^{2}q, (\boldsymbol{\theta}_{k+1} - \boldsymbol{\theta}_{k})^{\otimes 2} \right\rangle \middle| \mathcal{F}_{k} \right]$$

$$= \mathbb{E}\left[\left\langle \nabla^{2}q, \operatorname{diag}(\boldsymbol{S}\boldsymbol{H}\boldsymbol{S}^{\mathsf{T}})^{-1/2} \boldsymbol{S}\boldsymbol{x}_{k} \boldsymbol{x}_{k}^{\mathsf{T}} \boldsymbol{S}^{\mathsf{T}} \operatorname{diag}(\boldsymbol{S}\boldsymbol{H}\boldsymbol{S}^{\mathsf{T}})^{-1/2} \right\rangle \frac{\left(\left\langle \boldsymbol{S}\boldsymbol{x}_{k}, \boldsymbol{\theta}_{k} \right\rangle - y_{k}\right)^{2}}{L(\boldsymbol{\theta}_{k})} \middle| \mathcal{F}_{k} \right]$$

$$= \frac{1}{L(\boldsymbol{\theta}_{k})} \left(\left\langle \operatorname{diag}(\boldsymbol{S}\boldsymbol{H}\boldsymbol{S}^{\mathsf{T}})^{-1/2} \nabla^{2}q \operatorname{diag}(\boldsymbol{S}\boldsymbol{H}\boldsymbol{S}^{\mathsf{T}})^{-1/2}, \boldsymbol{S}\boldsymbol{H}\boldsymbol{S}^{\mathsf{T}} \right\rangle L(\boldsymbol{\theta}_{k})$$

$$+ 2 \left\langle \boldsymbol{S}\boldsymbol{H}\boldsymbol{S}^{\mathsf{T}} \operatorname{diag}(\boldsymbol{S}\boldsymbol{H}\boldsymbol{S}^{\mathsf{T}})^{-1/2} \nabla^{2}q \operatorname{diag}(\boldsymbol{S}\boldsymbol{H}\boldsymbol{S}^{\mathsf{T}})^{-1/2} \boldsymbol{S}\boldsymbol{H}\boldsymbol{S}^{\mathsf{T}}, (\boldsymbol{\theta}_{k} - \boldsymbol{\theta}^{*})^{\otimes 2} \right\rangle \right).$$

Combining the two contributions and inserting a stepsize γ_k ,

$$\mathbb{E}[q(\boldsymbol{\theta}_{k+1}) - q(\boldsymbol{\theta}_{k}) \mid \mathcal{F}_{k}] = -\frac{\gamma_{k}}{\sqrt{L(\boldsymbol{\theta}_{k})}} \left\langle \nabla q(\boldsymbol{\theta}_{k}), \overline{\boldsymbol{K}}(\boldsymbol{\theta}_{k} - \boldsymbol{\theta}^{*}) \right\rangle + \frac{\gamma_{k}^{2}}{2} \left\langle \nabla^{2} q, \boldsymbol{K}_{\tau} \right\rangle \\ + \frac{\gamma_{k}^{2}}{L(\boldsymbol{\theta}_{k})} \left\langle \boldsymbol{S} \boldsymbol{H} \boldsymbol{S}^{\mathsf{T}} \operatorname{diag}(\boldsymbol{S} \boldsymbol{H} \boldsymbol{S}^{\mathsf{T}})^{-1/2} \nabla^{2} q \operatorname{diag}(\boldsymbol{S} \boldsymbol{H} \boldsymbol{S}^{\mathsf{T}})^{-1/2} \boldsymbol{S} \boldsymbol{H} \boldsymbol{S}^{\mathsf{T}}, (\boldsymbol{\theta}_{k} - \boldsymbol{\theta}^{*})^{\otimes 2} \right\rangle,$$

where $K_{\tau} := \operatorname{diag}(\mathbf{SHS}^{\mathsf{T}})^{-1/2} \mathbf{SHS}^{\mathsf{T}} \operatorname{diag}(\mathbf{SHS}^{\mathsf{T}})^{-1/2}$.

Mode-wise recursion. For $r_i(k) := (\boldsymbol{\theta}_k - \boldsymbol{\theta}^*)^\mathsf{T} (\boldsymbol{K} \boldsymbol{u}_i \otimes \boldsymbol{w}_i) (\boldsymbol{\theta}_k - \boldsymbol{\theta}^*)$ (cf. Appendix D.1),

$$\mathbb{E}[r_i(k+1) - r_i(k) \mid \mathcal{F}_k] = -\frac{2\gamma_k}{\sqrt{L(\boldsymbol{\theta}_k)}} \lambda_i(\overline{\boldsymbol{K}}) r_i(k) + \gamma_k^2 (\boldsymbol{w}_i^{\mathsf{T}} \boldsymbol{K}_{\tau} \boldsymbol{K} \boldsymbol{u}_i) + \frac{2\gamma_k^2}{L(\boldsymbol{\theta}_k)} \lambda_i(\overline{\boldsymbol{K}}) r_i(k)$$

$$= -\left(\frac{2\gamma_k}{\sqrt{L(\boldsymbol{\theta}_k)}} - \frac{2\gamma_k^2}{L(\boldsymbol{\theta}_k)}\right) \lambda_i(\overline{\boldsymbol{K}}) r_i(k) + \gamma_k^2 (\boldsymbol{w}_i^{\mathsf{T}} \boldsymbol{K}_{\tau} \boldsymbol{K} \boldsymbol{u}_i).$$

We now assume $f \equiv 1$, and $\gamma_k = \gamma_0$ for simplicity. Passing to the ODE limit as in Section D.2 we get following ODE for $P(t) = L(t/\gamma_0)$ and $p_i(t) = r_i(t/\gamma_0)$.

$$\frac{dp_i}{dt} = -2\left(\frac{1}{\sqrt{P(t)}} - \frac{\gamma_0}{P(t)}\right) \lambda_i(\overline{K}) \, p_i(t) + \gamma_0 \, V_i'. \tag{102}$$

Interpreting the solution of the ODE as an implicit integral equation and summing over i, similar to Section D.2, and writing

$$Q_2(N):=2\gamma_0\int_0^N \Bigl(\frac{1}{\sqrt{L(u)}}-\gamma_0\frac{1}{L(u)}\Bigr)\,du,$$

we obtain the implicit integral equation

$$L(N) = \|\boldsymbol{H}^{1/2}\boldsymbol{w}_{\perp}\|^{2} + \sum_{i=1}^{M} r_{i}(0) \exp\left(-\lambda_{i}Q_{2}(N)\right)$$
$$+ \gamma_{0}^{2} \sum_{i=1}^{M} V_{i}' \int_{0}^{N} \exp\left(-2\lambda_{i}\gamma_{0} \int_{z}^{N} \left(\frac{1}{\sqrt{L(u)}} - \gamma_{0} \frac{1}{L(u)}\right) du\right) dz,$$

where $V_i' := \boldsymbol{w}_i^\mathsf{T} \boldsymbol{K}_{ au} \boldsymbol{K} \boldsymbol{u}_i$

Drift transformation and limit phase. By the same drift/approximation transformation as in equation 26,

$$\begin{split} L(N) &= M^{-2\alpha + \max(0,\, 1-2\beta)} + \left(M^{\min(\alpha,0.5)}\,Q_2(N)\right)^{-\frac{2\alpha + 2\beta - 1}{2\alpha}} \\ &\quad + \, \mathbf{1}_{\{\alpha > 0.5,\,\, \beta > 0.5\}}\,\,M^{-1}\,\left(M^{\min(\alpha,0.5)}Q_2(N)\right)^{-1 + \frac{1}{2\alpha}} + \gamma_0^2 \sum_{i=1}^M V_i' \int_0^N e^{-2\lambda_i \gamma_0 \int_z^N \left(\frac{f(u)}{\sqrt{L(u)}} - \gamma_0 \frac{f(u)^2}{L(u)}\right) \, du} \, f(z)^2 \, dz. \end{split}$$

We will first handle the limit phase, similar to Section D.3.2. At stationarity, let $p_i(t) \to s_i$ and $P(t) \to L_{\infty}$, we must have

$$-2\left(\frac{1}{\sqrt{L_{\infty}}} - \frac{\gamma_0}{L_{\infty}}\right)\lambda_i(\overline{\boldsymbol{K}})\,s_i + \gamma_0\,V_i = 0 \quad \Longrightarrow \quad s_i = \frac{\gamma_0\sqrt{L_{\infty}}}{2\,\lambda_i(\overline{K})}\,V_i'\frac{1}{1 - \frac{\gamma_0}{\sqrt{L_{\infty}}}} = \frac{\gamma_0\sqrt{L_{\infty}}}{2\,\lambda_i(\overline{K})}\,(\boldsymbol{w}_i^{\mathsf{T}}\boldsymbol{K}_{\tau}\boldsymbol{K}\boldsymbol{u}_i)\frac{1}{1 - \frac{\gamma_0}{\sqrt{L_{\infty}}}}.$$

Using the risk decomposition $P(t) = \sum_{i=1}^{M} p_i(t) + \|\boldsymbol{H}^{1/2}\boldsymbol{w}_{\perp}\|^2$, we obtain

$$\begin{split} L_{\infty} &= \sum_{i=1}^{M} s_{i} + \|\boldsymbol{H}^{1/2}\boldsymbol{w}_{\perp}\|^{2} = \frac{\gamma_{0}}{2} \Big(\sum_{i=1}^{M} \frac{\boldsymbol{w}_{i}^{\mathsf{T}}\boldsymbol{K}_{\tau}\boldsymbol{K}u_{i}}{\lambda_{i}(\overline{K})} \Big) \sqrt{L_{\infty}} \frac{1}{1 - \frac{\gamma_{0}}{\sqrt{L_{\infty}}}} + \|\boldsymbol{H}^{1/2}\boldsymbol{w}_{\perp}\|^{2} \\ &= \frac{\gamma_{0}}{2} \operatorname{Tr} \Big(\operatorname{diag}(K)^{1/2}\boldsymbol{K}_{\tau} \Big) \sqrt{L_{\infty}} \frac{1}{1 - \frac{\gamma_{0}}{\sqrt{L_{\infty}}}} + \|\boldsymbol{H}^{1/2}\boldsymbol{w}_{\perp}\|^{2} \\ &= \frac{\gamma_{0}}{2} \operatorname{Tr} \Big(\operatorname{diag}(K)^{1/2} \Big) \sqrt{L_{\infty}} \frac{1}{1 - \frac{\gamma_{0}}{\sqrt{L_{\infty}}}} + \|\boldsymbol{H}^{1/2}\boldsymbol{w}_{\perp}\|^{2}. \end{split}$$

And

$$L_{\infty} \; \approx \; \max \Bigl\{ \gamma_0^2 \; \text{Tr} \bigl(\text{diag}(K)^{1/2} \bigr)^2, \; \| \boldsymbol{H}^{1/2} \boldsymbol{w}_{\perp} \|^2 \Bigr\} \; \approx \; \max \Bigl\{ \gamma_0^2 \, M^{\, 2 - 2 \min(\alpha, 0.5)}, \; M^{-2\alpha + \max(0, \, 1 - 2\beta)} \Bigr\}.$$

asymptotically satisfies the equation. So we have the same floor as for signSGD.

Since f is bounded and $L(N) \ge \gamma_0^2 M^{2-2\min(\alpha,0.5)}$,

$$\frac{\frac{f(u)}{\sqrt{L(u)}}}{\gamma_0\frac{f(u)^2}{L(u)}} = \frac{\sqrt{L(u)}}{\gamma_0f(u)} \ \gtrsim \ M^{\,1-\min(\alpha,0.5)},$$

so the subtraction inside Q_2 is asymptotically negligible and $Q_2(N) \approx Q(N)$. Hence, the drift contribution coincides with that of signSGD.

Scaling law (constant learning rate). For $f \equiv 1$, Adam (under this heuristic) follows the same scaling law as signSGD:

$$\begin{split} R(M,N,\gamma_0) \; \approx \; M^{-2\alpha+\max(0,\,1-2\beta)} + \left(M^{\min(\alpha,0.5)}\,N\,\gamma_0\right)^{-\frac{2(2\alpha+2\beta-1)}{2\alpha+1-2\beta}} \\ + \left(M^{\frac{6\alpha-1}{4\alpha-2}}\,N\,\gamma_0\right)^{-\frac{2(2\alpha-1)}{2\alpha+1}} + \gamma_0^2\,M^{\,2-2\,\min(\alpha,0.5)}. \end{split}$$

Since the loss formula $R(M,N,\gamma_0)$ is the same as signSGD, the compute-optimal scaling law will also be the same as signSGD. So we expect that Adam has the compute-optimal scaling law in Table 1. Figure 16 shows that exponents in the Table 1 and measured compute-optimal loss slope and optimal model size slope (in log-log plot) for Adam match well.

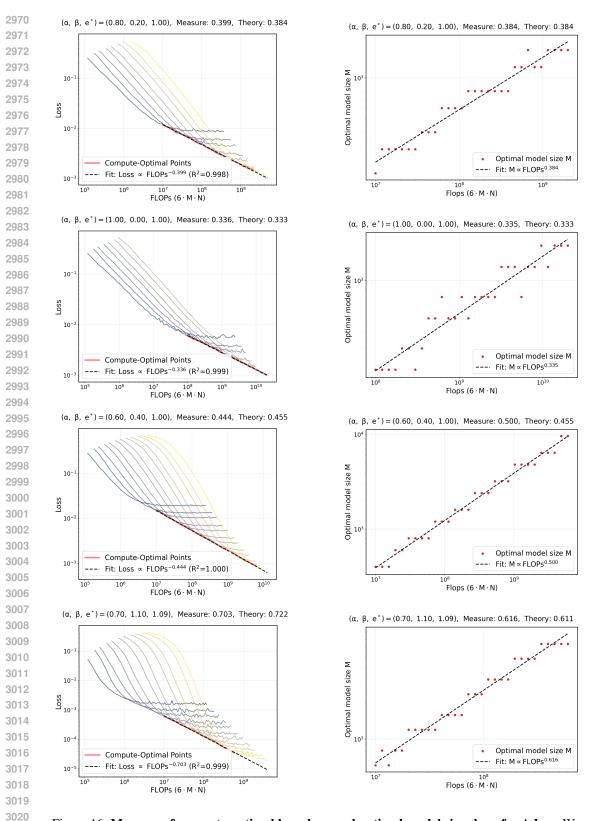


Figure 16: Measure of compute-optimal loss slope and optimal model size slope for Adam. We validate the exponent of $R\left(M^\star,\frac{f}{M^\star},\gamma_0^\star\right)$ and M^\star with respect to f in the Table 1. The left plot shows the compute-optimal loss with respect to flops 6MN. The right plot shows the optimal model size with respect to flops 6MN. Each plot includes the measured slope and the theoretical slope from the Table 1. Parameters: $\beta_1=0.9,\,\beta_2=0.999,\,\epsilon=10^{-8},\,\gamma_0=0.002$.

I OMITTED ANALYSIS FROM SECTION D

I.1 OMITTED PROOF OF (22) AND (24)

In this section, we cover omitted proof of (22) and (24). Note that the proof is almost similar to Paquette et al. (2024), but we cover it briefly for completeness. Refer to Appendix F, G, H of Paquette et al. (2024) for more details.

It is enough to prove

$$-\frac{1}{2\pi i} \oint_{\Gamma} e^{-p_d Q(N)z} \left\langle \mathcal{L}(z), v^{\otimes 2} \right\rangle dz \approx M^{-2\alpha + \max(0, 1 - 2\beta)}$$

$$+ \left(M^{\min(\alpha, 0.5)} Q(N) \right)^{-\frac{2\alpha + 2\beta - 1}{2\alpha}}$$

$$+ 1_{\{\alpha > 0.5, \beta > 0.5\}} M^{-1} \left(M^{\min(\alpha, 0.5)} Q(N) \right)^{-1 + \frac{1}{2\alpha}}.$$

From now on, we will use similar notation to Paquette et al. (2024), except in the inevitable case, to facilitate easy comparison for the reader. Note that we use M and d for model size and initial dimension before projection, while Paquette et al. (2024) uses d and v.

We use Γ for the contour containing the spectrum of K, while Paquette et al. (2024) used $\Gamma \cup \Gamma_0$ for that, where Γ_0 is a small circle containing the origin.

Let

$$\mathcal{F}(N) := -\frac{1}{2\pi i} \oint_{\Gamma} \langle \mathcal{R}(z), (H^{1/2} w^*)^{\otimes 2} \rangle e^{-p_d Q(N) z} dz.$$
 (103)

The exponential kernel $e^{-p_dQ(N)z}$ replaces all polynomial weights in the analysis of Paquette et al. (2024). The resulting leading orders remain the same while constants and exponents are altered in a transparent way; precise statements follow.

We can split the $\mathcal{F}(N)$ by splitting the keyhole contour Γ . We let

$$\mathcal{F}(N) = \mathcal{F}_0(N) + \mathcal{F}_{\text{caps}}(N) + \mathcal{F}_C(N) + \text{(lower-order)}, \tag{104}$$

where \mathcal{F}_0 collects the small circle around the origin, $\mathcal{F}_{\rm caps}$ collects the right/left caps adjacent to the positive real axis, and \mathcal{F}_C collects the central arc close to [0,1]. Refer to Appendix F of Paquette et al. (2024) for more details about the picture of contour and decomposition of contour.

In the following proposition the function $(x)_+ := \max(x, 0)$ is used.

Proposition I.1. $\mathcal{F}_0(N)$ is independent of N and obeys

$$\left| \mathcal{F}_0(0) - \sum_{j=1}^d \frac{j^{-2\alpha - 2\beta}}{1 + j^{-2\alpha} M^{2\alpha} \kappa(d/M)} \right| \le C M^{-2\alpha + (2\beta - 1)_+ - 1}.$$

Sketch. Putting z=0 to the exponential leads to 1, so we can reduce to the analysis of Paquette et al. (2024). So the error bound is identical.

After this $\mathcal{F}_0(N) \approx M^{-2\alpha+\max(0,\,1-2\beta)}$ holds by identical procedure calculating $\sum_{j=1}^d \frac{j^{-2\alpha-2\beta}}{1+j^{-2\alpha}M^{2\alpha}\kappa(d/M)}$.

Proposition I.2. There exist functions f, q > 0 with

$$f(N) \le C \exp(-p_d Q(N) M^{-2\alpha}), \qquad g(N) \le C \exp(-p_d Q(N)),$$

so that

$$|\mathcal{F}_{\text{caps}}(N)| \le C f(N) M^{-2\alpha + (1-2\beta)_+} + C g(N).$$

Sketch. Use $|m(z)-1| \lesssim M^{-\min\{2\alpha,1\}}$ (as in Paquette et al. (2024)) on a cap pushed $\mathcal{O}(1)$ -close to [0,1] to replace $\langle \mathcal{R}(z), (H^{1/2}w^*)^{\otimes 2} \rangle$ by a simple partial fraction, and control the remainder by the real part of z.

The main contribution arises from the arc parameterized by $z(u) = u + i\eta(u)$ with $u \in [M^{-2\alpha}, 1]$ and $|\eta(u)| \ll u$. Along this arc we have the uniform approximation

$$\left| m(z(u)) - \left(1 - \frac{\pi}{2\alpha}(c(u) + i) u^{-1/(2\alpha)} M^{-1} \right) \right| \le \varepsilon u^{-1/(2\alpha)} M^{-1}$$
 (105)

for some bounded real c(u). Inserting (105) in $\mathcal{R}(z) = (-zI + m(z)H)^{-1}$ and extracting the imaginary part produces two canonical integrals,

$$\mathcal{F}_{pp}(N) := \frac{1}{2\alpha} \int_0^1 u^{(2\beta - 1)/(2\alpha)} e^{-p_d Q(N)u} du, \qquad \mathcal{F}_{ac}(N) := \frac{c_\beta}{2\alpha} \int_{M^{-2\alpha}}^1 u^{-1/(2\alpha)} M^{-1} e^{-p_d Q(N)u} du, \tag{106}$$

with $c_{\beta} = \sum_{j>1} j^{-2\beta}$ if $2\beta > 1$ and $c_{\beta} = 0$ otherwise.

Proposition I.3. There exists C>0 such that for all $N\geq 0$, $|\mathcal{F}_C(N)|\leq C\big(\mathcal{F}_{pp}(N)+\mathcal{F}_{ac}(N)\big)$. Moreover, there are A>0 and a bounded function C(N)>0 with $C(N)\leq 1+\varepsilon$ whenever $p_dQ(N)\in [A,M^{2\alpha}/A]$, and

$$\frac{1}{C(N)} \left(\mathcal{F}_{pp}(N) + \mathcal{F}_{ac}(N) \right) \leq \mathcal{F}_{C}(N) \leq C(N) \left(\mathcal{F}_{pp}(N) + \mathcal{F}_{ac}(N) \right).$$

Sketch. Parameterize Γ_C by u and use (105) to separate real/imaginary parts. The imaginary terms integrate exactly to (106), while the real part is smaller by a factor $\mathcal{O}(\varepsilon)$ since $|\eta(u)| \ll u$.

Proposition I.4 (Asymptotics of \mathcal{F}_{pp}). Assume $2\alpha + 2\beta > 1$ and set $X := p_d Q(N)$. For any $\varepsilon > 0$ there exists A > 0 such that for $X \ge A$,

$$\left| \mathcal{F}_{pp}(N) - g_{pp}(N) \right| \le \varepsilon g_{pp}(N),$$

where

$$g_{pp}(N) := (2\alpha)^{-1} X^{-(1+\beta/\alpha)+1/(2\alpha)} \Gamma\left(\frac{\beta}{\alpha} - \frac{1}{2\alpha} + 1\right).$$

Moreover, if $X \leq \widetilde{A}$ then $c \leq \mathcal{F}_{pp}(N) \leq C$ for constants c, C > 0, and if $X \geq \widetilde{A} M^{2\alpha}$ then $\mathcal{F}_{pp}(N) \leq \widetilde{C} \mathcal{F}_0(N)$ for some $\widetilde{C} > 0$ independent of M.

Sketch. With the change of variables w = Xu, we get

$$\mathcal{F}_{pp}(N) = (2\alpha)^{-1} X^{-(1+\beta/\alpha)+1/(2\alpha)} \int_{0}^{X} w^{(2\beta-1)/(2\alpha)} e^{-w} dw.$$

Comparing to the complete gamma integral yields the relative error bound in terms of the upper incomplete gamma tail, which can be made $\leq \varepsilon$ by choosing A large. The remaining bounds follow by monotonicity and elementary estimates.

Proposition I.5 (Asymptotics of \mathcal{F}_{ac}). Let $X := p_d Q(N)$. There exists $C(\alpha, \beta) > 0$ such that

$$\mathcal{F}_{ac}(N) \le \begin{cases} C \mathcal{F}_0(N), & 2\beta > 1, \ 2\alpha < 1, \\ 0, & 2\beta < 1. \end{cases}$$

If in addition $2\alpha > 1$ and $2\beta > 1$, then for any $\varepsilon > 0$ there is A > 0 such that whenever $X \in [A, M^{2\alpha}/A]$,

$$|\mathcal{F}_{ac}(N) - g_{ac}(N)| \le \varepsilon \, g_{ac}(N), \qquad g_{ac}(N) := \left(\sum_{j=1}^{\nu} j^{-2\beta}\right) (2\alpha)^{-1} \, \Gamma\left(1 - \frac{1}{2\alpha}\right) X^{-1 + 1/(2\alpha)} \, M^{-1}.$$

Furthermore, for any $\widetilde{A} > 0$ there exist constants C, c > 0 (independent of M) such that

$$\mathcal{F}_{ac}(N) \le \begin{cases} C M^{-1}, & X \le \widetilde{A}, \\ c \mathcal{F}_0(N), & X \ge \widetilde{A} M^{2\alpha}. \end{cases}$$

Sketch. Compare the truncated integral in (106) with its extension to $[0,\infty)$ and control the two tails $[0,M^{-2\alpha}]$ and $[1,\infty)$ separately. The first is at most $\widetilde{c}M^{-2\alpha}$; the second is bounded by $M^{-1}X^{-1}e^{-X}$. Normalizing by $g_{ac}(N)$ shows both are relatively small for $X\in [A,M^{2\alpha}/A]$ with A large. The endpoint bounds follow from dropping the exponential and from a crude $\int e^{-Xu}du \leq X^{-1}e^{-XM^{-2\alpha}}$ estimate when $X\gtrsim M^{2\alpha}$.

Finally we get

3133
$$-\frac{1}{2\pi i} \oint_{\Gamma} e^{-p_d Q(N)z} \left\langle \mathcal{L}(z), v^{\otimes 2} \right\rangle dz \approx \mathcal{F}_0(N) + \mathcal{F}_{\text{caps}}(N) + \mathcal{F}_C(N)$$
3136
$$\approx M^{-2\alpha + \max(0, 1 - 2\beta)} + \left(p_d Q(N)\right)^{-\frac{2\alpha + 2\beta - 1}{2\alpha}}$$
3138
$$+ 1_{\{\alpha > 0.5, \beta > 0.5\}} M^{-1} \left(p_d Q(N)\right)^{-1 + \frac{1}{2\alpha}} M^{-2\alpha + \max(0, 1 - 2\beta)}$$
3140
$$\approx M^{-2\alpha + \max(0, 1 - 2\beta)} + \left(M^{\min(\alpha, 0.5)} Q(N)\right)^{-\frac{2\alpha + 2\beta - 1}{2\alpha}}$$
3143
$$+ 1_{\{\alpha > 0.5, \beta > 0.5\}} M^{-1} \left(M^{\min(\alpha, 0.5)} Q(N)\right)^{-1 + \frac{1}{2\alpha}}$$
3145

I.2 Note on the $\arcsin x \approx x$ Approximation

We explain that it is possible to replace the linear approximation $\arcsin x \approx x$ by an inequality, and the main results of our paper remain unchanged.

Replacing the \arcsin -linearization by a uniform sandwich. Fix $0 < \rho \le 1$ and define

$$c_1(\rho) \ := \ \inf_{|t| \le \rho} \frac{\arcsin t}{t} = 1, \qquad c_2(\rho) \ := \ \sup_{|t| \le \rho} \frac{\arcsin t}{t} = \frac{\arcsin \rho}{\rho} \ \le \ \frac{\pi}{2}.$$

For $x \in \mathbb{R}^d$ with $||x||_{\infty} \leq \rho$, the entrywise odd and monotone map $t \mapsto \arcsin t$ satisfies the componentwise bounds

$$c_1(\rho) x \le \arcsin(x) \le c_2(\rho) x.$$

In our update, put $v_k \coloneqq \boldsymbol{\theta}_k - \boldsymbol{\theta}^*$ and

$$x_k := \frac{\overline{K} v_k}{\sqrt{L(\boldsymbol{\theta}_k)}}, \quad \text{so that} \quad \arcsin(x_k) = D_k x_k,$$

for some diagonal $D_k = \operatorname{diag}(\kappa_{k,1}, \dots, \kappa_{k,d})$ with $c_1(\rho) \leq \kappa_{k,j} \leq c_2(\rho)$. Using $K^\mathsf{T} = K$ and $K^\mathsf{T}\overline{K} = \overline{K}^\mathsf{T}K^\mathsf{T}$, the one–step drift can be written as

$$\mathbb{E}[r_i(k+1) - r_i(k) \mid \mathcal{F}_k] = -\frac{2\gamma_k}{\pi\sqrt{L(\boldsymbol{\theta}_k)}} v_k^\mathsf{T} \left(K\boldsymbol{u}_i\boldsymbol{w}_i^\mathsf{T} + \boldsymbol{w}_i\boldsymbol{u}_i^\mathsf{T}K\right) D_k \overline{K} v_k + \frac{2\gamma_k^2}{\pi} (\boldsymbol{w}_i^\mathsf{T}K_\sigma K \boldsymbol{u}_i).$$

Since D_k is diagonal with $c_1(\rho)I \leq D_k \leq c_2(\rho)I$, the quadratic form is sandwiched between the same expression with D_k replaced by $c_1(\rho)I$ and $c_2(\rho)I$. Recalling the identity used earlier,

$$v_k^{\mathsf{T}} \left(K \boldsymbol{u}_i \boldsymbol{w}_i^{\mathsf{T}} + \boldsymbol{w}_i \boldsymbol{u}_i^{\mathsf{T}} K \right) \overline{K} \, v_k = 2 \, \lambda_i(\overline{K}) \, r_i(k),$$

we obtain the two-sided one-step bound

$$-\frac{4 c_2(\rho) \gamma_k}{\pi \sqrt{L(\boldsymbol{\theta}_k)}} \lambda_i(\overline{K}) r_i(k) + \frac{2 \gamma_k^2}{\pi} (\boldsymbol{w}_i^\mathsf{T} K_\sigma K \boldsymbol{u}_i) \leq \mathbb{E}[r_i(k+1) - r_i(k) \mid \mathcal{F}_k]$$

$$\leq -\frac{4 c_1(\rho) \gamma_k}{\pi \sqrt{L(\boldsymbol{\theta}_k)}} \lambda_i(\overline{K}) r_i(k) + \frac{2 \gamma_k^2}{\pi} (\boldsymbol{w}_i^\mathsf{T} K_\sigma K \boldsymbol{u}_i).$$

Consequences for the ODE limit and the implicit integral equation. Let $\gamma_k = \gamma_0 f(k)$, $t = k\gamma_0$, $p_i(t) := r_i(k)$, and $P(t) := L(\theta_k)$, as in Appendix D.2. Then we obtain the differential inequalities

$$-\frac{4\,c_2(\rho)}{\pi\sqrt{P(t)}}\,\lambda_i(\overline{K})\,f(t/\gamma_0)\,p_i(t) + \frac{2\gamma_0}{\pi}f(t/\gamma_0)^2V_i \,\leq\, \dot{p}_i(t) \,\leq\, -\frac{4\,c_1(\rho)}{\pi\sqrt{P(t)}}\,\lambda_i(\overline{K})\,f(t/\gamma_0)\,p_i(t) + \frac{2\gamma_0}{\pi}f(t/\gamma_0)^2V_i,$$

with $V_i := \boldsymbol{w}_i^\mathsf{T} K_\sigma K \boldsymbol{u}_i$. Solving these linear comparison inequalities yields the bounds

$$p_i^{(c_2)}(t) \, \leq \, p_i(t) \, \leq \, p_i^{(c_1)}(t), \qquad P^{(c_2)}(t) \, \leq \, P(t) \, \leq \, P^{(c_1)}(t),$$

where $p_i^{(c)}(\cdot)$ and $P^{(c)}(\cdot)$ denote the solutions of the ODE/integral equations from Appendix D.2 with the factor $\frac{4}{\pi}$ replaced by $\frac{4c}{\pi}$. Equivalently, defining

$$Q_c(N) := \frac{4c\gamma_0}{\pi} \int_0^N \frac{f(u)}{\sqrt{P(u)}} du,$$

the drift/noise expressions remain valid with Q(N) replaced by $Q_c(N)$, and all proofs carry through verbatim.

Only multiplicative constants change; scaling exponents and phases do not. Every appearance of Q(N) in the final formulas enters either through an exponential $e^{-\lambda Q(N)}$ or through a polynomial factor $(M^{\mu}Q(N))^{-\nu}$. Replacing Q by $Q_c=c\,Q$ only multiplies these terms by constants: $e^{-\lambda cQ}$ converts to $(M^{\mu}cQ)^{-\nu}=c^{-\nu}(M^{\mu}Q)^{-\nu}$. Hence the rates, exponents, and phase boundaries of the scaling laws are unchanged; only the prefactors are rescaled by fixed constants depending on $c_1(\rho), c_2(\rho) \in [1, \pi/2]$. In particular, all " \approx " statements (equalities up to absolute constants) remain valid with the same exponents.

I.3 NOTE ON APPROXIMATION ERROR

Though proof of Paquette et al. (2024) implicitly implies

$$\left\| \boldsymbol{H}^{1/2} \boldsymbol{w}_{\perp} \right\|^2 \approx M^{-2\alpha + \max(0, 1 - 2\beta)}.$$

It was not explicitly specified. So we clarify it here.

First,

$$-\frac{1}{2\pi i} \oint_{|z|=\varepsilon} \left\langle (\widehat{\boldsymbol{K}} - z\boldsymbol{I})^{-1}, \, (\boldsymbol{H}^{1/2}\boldsymbol{w}^*)^{\otimes 2} \right\rangle dz \approx M^{-2\alpha + \max(0, \, 1 - 2\beta)},$$

is directly implied from Proposition H.3 of Paquette et al. (2024). So it is enough to prove the following claim.

Claim. Let

$$\widehat{m{K}} = m{H}^{1/2} m{S}^\mathsf{T} m{S} \, m{H}^{1/2}, \qquad m{w}^* = m{S}^\mathsf{T} m{ heta}^* + m{w}_\perp, \qquad m{S} \, m{H} \, m{w}_\perp = m{0}.$$

For a sufficiently small circle $|z|=\varepsilon$ enclosing only the eigenvalue 0 of \widehat{K} ,

$$-\frac{1}{2\pi i} \oint_{|z|=\varepsilon} \left\langle (\widehat{\boldsymbol{K}} - z\boldsymbol{I})^{-1}, (\boldsymbol{H}^{1/2}\boldsymbol{w}^*)^{\otimes 2} \right\rangle dz = \|\boldsymbol{H}^{1/2}\boldsymbol{w}_{\perp}\|^2.$$

Proof. By the Riesz projection theorem (Dunford–Riesz functional calculus), for a small circle $|z| = \varepsilon$ enclosing only the eigenvalue 0 of \widehat{K} ,

$$\Pi_0 := -\frac{1}{2\pi i} \oint_{|z|=\varepsilon} (\widehat{K} - z\mathbf{I})^{-1} dz$$

is the spectral Riesz projector onto the 0-eigenspace; since \widehat{K} is Hermitian, Π_0 is the *orthogonal* projector onto $\ker(\widehat{K})$.

And we have

$$-\frac{1}{2\pi i} \oint_{|z|=\varepsilon} \left\langle (\widehat{\boldsymbol{K}} - z\boldsymbol{I})^{-1}, (\boldsymbol{H}^{1/2}\boldsymbol{w}^*)^{\otimes 2} \right\rangle dz = \left\langle \Pi_0, (\boldsymbol{H}^{1/2}\boldsymbol{w}^*)^{\otimes 2} \right\rangle = \left\| \Pi_0 \boldsymbol{H}^{1/2}\boldsymbol{w}^* \right\|_2^2$$

Since $\ker(\widehat{\boldsymbol{K}}) = \{ \boldsymbol{x} : \boldsymbol{S} \boldsymbol{H}^{1/2} \boldsymbol{x} = \boldsymbol{0} \} = \left(\operatorname{Im}(\boldsymbol{H}^{1/2} \boldsymbol{S}^{\mathsf{T}}) \right)^{\perp}$, We have the orthogonal decomposition

$$m{H}^{1/2}m{w}^* = \underbrace{m{H}^{1/2}m{S}^{\mathsf{T}}m{ heta}^*}_{\in \operatorname{Im}(m{H}^{1/2}m{S}^{\mathsf{T}})} + \underbrace{m{H}^{1/2}m{w}_{\perp}}_{\in (\operatorname{Im}(m{H}^{1/2}m{S}^{\mathsf{T}}))^{\perp}},$$

where the second membership uses $m{S}\,m{H}\,m{w}_\perp = m{0}$. Hence $\Pi_0\,m{H}^{1/2}m{w}^* = m{H}^{1/2}m{w}_\perp$, and therefore

$$\left\langle \Pi_0, \left(\boldsymbol{H}^{1/2} \boldsymbol{w}^* \right)^{\otimes 2} \right\rangle = \left\| \boldsymbol{H}^{1/2} \boldsymbol{w}_{\perp} \right\|^2. \quad \Box$$

I.4 Proof of Matrix Inequality for $\operatorname{diag}(\boldsymbol{SHS}^{\top})^{-1/2}$

We will prove the inequality in the following form in this section.

$$c_1 M^{\min(0.5,\alpha)} I \leq \operatorname{diag}(\mathbf{SHS}^{\mathsf{T}})^{-1/2} \leq c_2 M^{\min(0.5,\alpha)} I$$

Setup. Let $S \in \mathbb{R}^{M \times d}$ have i.i.d. entries $S_{ij} \sim \mathcal{N}(0, 1/M)$, and let

$$H = diag(1^{-2\alpha}, 2^{-2\alpha}, \dots, d^{-2\alpha}), \qquad \alpha > 0.$$

Then, for each $i \in \{1, \dots, M\}$,

$$\left[\operatorname{diag}(\boldsymbol{SHS}^{\mathsf{T}})\right]_{ii} = \sum_{j=1}^{d} H_{jj} S_{ij}^2 = \frac{1}{M} \sum_{j=1}^{d} j^{-2\alpha} \chi_j^2,$$

where $\chi_1^2, \dots, \chi_d^2$ are i.i.d. $\chi^2(1)$.

Remark 5 (Rough intuition for what we will prove).

$$\left[\mathrm{diag}(\boldsymbol{SHS}^\mathsf{T})\right]_{ii} = \frac{1}{M} \sum_{j=1}^d j^{-2\alpha} \chi_j^2 \ \approx \ \begin{cases} M^{-1}, & \alpha > \frac{1}{2}, \\ M^{-1} \, d^{\, 1-2\alpha} \approx M^{-2\alpha}, & \alpha \leq \frac{1}{2} \text{ with } d \approx M, \end{cases}$$

So, we want to obtain $\operatorname{diag}(\mathbf{SHS}^{\mathsf{T}})^{-1/2} \approx M^{\min(0.5,\alpha)}I$.

Define

$$S_d(\alpha) := \sum_{i=1}^d j^{-2\alpha} \chi_j^2 \implies \left[\operatorname{diag}(\mathbf{SHS}^\mathsf{T}) \right]_{ii} = \frac{1}{M} S_d(\alpha).$$

Hence, any high–probability upper/lower bounds on $S_d(\alpha)$ translate into corresponding bounds on $\operatorname{diag}(\boldsymbol{SHS}^\mathsf{T})^{-1/2}$ via

$$\frac{1}{M} S_d(\alpha) \leq U \implies \left[\operatorname{diag}(\boldsymbol{SHS}^{\mathsf{T}}) \right]^{-1/2} \succeq \sqrt{\frac{M}{U}} I,$$

$$\frac{1}{M} S_d(\alpha) \geq L \implies \left[\operatorname{diag}(\boldsymbol{SHS}^{\mathsf{T}}) \right]^{-1/2} \preceq \sqrt{\frac{M}{L}} I.$$

We consider two regimes and then unify them through $M^{\min(0.5,\alpha)}$.

REGIME I: $\alpha > \frac{1}{2}$ (SUMMABLE WEIGHTS)

In this regime, $\sum_{j=1}^{\infty} j^{-2\alpha} = \zeta(2\alpha) < \infty$. Write $X_j := j^{-2\alpha}(\chi_j^2 - 1)$, so that

$$S_d(\alpha) = \mathbb{E}[S_d(\alpha)] + \sum_{j=1}^d X_j, \qquad \mathbb{E}[S_d(\alpha)] = \sum_{j=1}^d j^{-2\alpha} \le \zeta(2\alpha).$$

Moreover, $Var(S_d(\alpha)) = 2 \sum_{j=1}^d j^{-4\alpha} \le 2\zeta(4\alpha)$.

Upper tail (to lower–bound diag^{-1/2}). For $\lambda = \frac{1}{2}$,

$$\mathbb{E}[e^{\lambda X_j}] = e^{-\lambda j^{-2\alpha}} (1 - 2\lambda j^{-2\alpha})^{-1/2} \le \exp(\frac{1}{2}j^{-4\alpha}),$$

hence

$$\mathbb{E}\left[e^{\frac{1}{2}\left(S_d(\alpha) - \mathbb{E}S_d(\alpha)\right)}\right] \le \exp\left(\frac{1}{2}\sum_{i=1}^d j^{-4\alpha}\right) \le \exp\left(\frac{1}{2}\zeta(4\alpha)\right).$$

By Markov and a union bound over the M diagonal entries, setting the per–entry failure probability to $\delta_0 := \delta_{\text{total}}/M$,

$$\Pr\left(S_d(\alpha) \le \zeta(2\alpha) + \zeta(4\alpha) + 2\log\frac{M}{\delta_{\text{total}}}\right) \ge 1 - \delta_{\text{total}}.$$

Therefore, with probability at least $1 - \delta_{\text{total}}$,

$$\operatorname{diag}(\boldsymbol{SHS}^{\mathsf{T}})^{-1/2} \;\succeq\; \frac{\sqrt{M}}{\left(\zeta(2\alpha) + \zeta(4\alpha) + 2\log\frac{M}{\delta_{\operatorname{logal}}}\right)^{1/2}}\,I.$$

Lower tail (to upper-bound diag^{-1/2}). A Chernoff bound on the lower tail of $S_d(\alpha)$ (via the mgf of $e^{-t j^{-2\alpha} \chi^2}$) gives, for any $\delta \in (0,1)$, the existence of a constant

$$c_{\downarrow}(\alpha) = \left(\frac{2\alpha-1}{2}\right)^{2\alpha-1}/2^{2\alpha-1}$$

such that

$$\Pr\Big(S_d(\alpha) \geq c_{\downarrow}(\alpha) \left(\log(1/\delta)\right)^{-(2\alpha-1)}\Big) \geq 1 - \delta.$$

With $\delta = \delta_0 = \delta_{\text{total}}/M$ and a union bound over the M rows, with probability at least $1 - \delta_{\text{total}}$,

$$\mathrm{diag}(\boldsymbol{SHS}^\mathsf{T})^{-1/2} \; \preceq \; \frac{\sqrt{M}}{\left(c_{\downarrow}(\alpha)\right)^{1/2}} \left(\log \frac{M}{\delta_{\mathsf{total}}}\right)^{\frac{2\alpha-1}{2}} I.$$

Conclusion for $\alpha > \frac{1}{2}$. Combining the two displays,

$$\frac{\sqrt{M}}{\left(\zeta(2\alpha) + \zeta(4\alpha) + 2\log\frac{M}{\delta_{\text{total}}}\right)^{1/2}}I \leq \operatorname{diag}(\boldsymbol{SHS}^{\mathsf{T}})^{-1/2} \leq \frac{\sqrt{M}}{\left(c_{\downarrow}(\alpha)\right)^{1/2}} \left(\log\frac{M}{\delta_{\text{total}}}\right)^{\frac{2\alpha - 1}{2}}I \qquad (\alpha > \frac{1}{2}).$$

Regime II: $\alpha \leq \frac{1}{2}$ (diverging weights)

Assume $d \ge r M$ for some fixed r > 1 (as in our setup). Then

$$\mathbb{E}[S_d(\alpha)] = \sum_{i=1}^d j^{-2\alpha} \quad \text{satisfies} \quad \frac{(d+1)^{1-2\alpha}-1}{1-2\alpha} \ \leq \ \mathbb{E}[S_d(\alpha)] \ \leq \ 1 + \frac{d^{1-2\alpha}-1}{1-2\alpha}.$$

Hence $\mathbb{E}[S_d(\alpha)] = d^{1-2\alpha}$. Moreover,

$$\operatorname{Var}(S_d(\alpha)) = 2 \sum_{i=1}^d j^{-4\alpha} \begin{cases} = O(1), & \alpha > \frac{1}{4}, \\ = \Theta(d^{1-4\alpha}), & \alpha < \frac{1}{4}, \end{cases}$$

so in all cases $\sqrt{\operatorname{Var}(S_d(\alpha))} = o(\mathbb{E}[S_d(\alpha)])$ as $d \to \infty$. Thus, by Bernstein and a union bound over the M rows, for all sufficiently large M we get, with probability at least $1 - \delta_{\text{total}}$,

$$\frac{1}{2} \mathbb{E}[S_d(\alpha)] \leq S_d(\alpha) \leq \frac{3}{2} \mathbb{E}[S_d(\alpha)].$$

Using $d \geq rM$ and the integral bounds for $\mathbb{E}[S_d(\alpha)]$,

$$\frac{(rM)^{1-2\alpha}-1}{2(1-2\alpha)} \le S_d(\alpha) \le \frac{3}{2} \left(1 + \frac{(rM)^{1-2\alpha}-1}{1-2\alpha}\right).$$

Dividing by M and inverting the square–root yields constants

$$C_L(\alpha, r) := \left(\frac{3}{1 - 2\alpha} r^{1 - 2\alpha}\right)^{-1/2}, \qquad C_U(\alpha, r) := \left(\frac{1}{2(1 - 2\alpha)} r^{1 - 2\alpha}\right)^{-1/2},$$

such that, with probability at least $1 - \delta_{\text{total}}$,

$$\boxed{C_L(\alpha, r) M^{\alpha} I \leq \operatorname{diag}(\mathbf{SHS}^{\mathsf{T}})^{-1/2} \leq C_U(\alpha, r) M^{\alpha} I} \quad (\alpha \leq \frac{1}{2}).$$

Unified statement Combining Regimes I and II, there exist positive constants $c_1(\alpha, r, \delta_{total})$ and $c_2(\alpha, r, \delta_{total})$ such that, with probability at least $1 - \delta_{\text{total}}$, $c_1(\alpha, r, \delta_{\text{total}}) M^{\min(0.5, \alpha)} I \leq \operatorname{diag}(\boldsymbol{SHS}^{\mathsf{T}})^{-1/2} \leq c_2(\alpha, r, \delta_{\text{total}}) M^{\min(0.5, \alpha)} I$ with the following explicit choices: • If $\alpha > \frac{1}{2}$: $c_1(\alpha, \cdot, \delta_{\text{total}}) = \left(\zeta(2\alpha) + \zeta(4\alpha) + 2\log\frac{M}{\delta_{\text{total}}}\right)^{-1/2}, \qquad c_2(\alpha, \cdot, \delta_{\text{total}}) = \left(c_{\downarrow}(\alpha)\right)^{-1/2} \left(\log\frac{M}{\delta_{\text{total}}}\right)^{\frac{2\alpha - 1}{2}},$ where one admissible choice is $c_{\downarrow}(\alpha)=\left(\frac{2\alpha-1}{2}\right)^{2\alpha-1}\!/\,2^{2\alpha-1}.$ • If $\alpha \leq \frac{1}{2}$ and $d \geq rM$: $c_1(\alpha, r, \cdot) = C_L(\alpha, r), \qquad c_2(\alpha, r, \cdot) = C_U(\alpha, r),$ with C_L, C_U as defined above.