Pretraining Data Scale Reverses the Muon–AdamW Fine-Tuning Gap

Catherine Li* **Jingchu Gai*** **Andrej Risteski** **Jacob Mitchell Springer†** **Aditi Raghunathan†**
Carnegie Mellon University

{catheri4, jgai}@andrew.cmu.edu

arXiv:2611.03633v1 [cs.LG] 2 Nov 2026

Abstract

Muon has recently emerged as a promising optimizer for pretraining large language models, offering a speedup over AdamW in minimizing pretraining loss. However, it remains unclear whether Muon's advantage persists after fine-tuning. In this paper, we study the downstream behavior of Muon and AdamW and uncover a surprising trend: beyond a sufficiently large pretraining data budget, AdamW outperforms Muon after fine-tuning. Moreover, as model size increases, AdamW overtakes Muon at progressively smaller pretraining token–parameter ratios. To explain this reversal, we view fine-tuning as a perturbation of the outputs and introduce the margin-to-spread ratio, which measures how far the correct logit lies above the wrong ones relative to how widely the wrong logits are spread. In language models, Muon has the larger margin-to-spread ratio at small data budgets and AdamW at large ones, mirroring the fine-tuning trend. In a linear associative memory, we prove that after one update the margin-to-spread ratio of gradient descent gains on that of Muon as the data budget grows, and simulations show that this trend persists after extended training.

1 Introduction

The Muon optimizer has recently shown strong empirical performance for pretraining LLMs, challenging AdamW's status as the standard optimizer for minimizing pretraining loss (Jordan et al., 2024). Prior work has shown that the pretraining optimization choices that minimize pretraining loss do not necessarily yield the best model after fine-tuning (Watts et al., 2026). What ultimately matters downstream is the *learning–forgetting tradeoff*: how well a fine-tuned model learns the target task while retaining its pretrained knowledge. Motivated by this, we ask:

Does Muon's pretraining advantage over AdamW extend to the learning–forgetting tradeoff?

To answer this question, we measure the Muon–AdamW *fine-tuning gap*, which compares the two optimizers on this tradeoff; a negative gap indicates that the Muon pretrained model dominates the AdamW pretrained model after fine-tuning, and a positive gap indicates the reverse. We show that the answer depends on the pretraining data budget, and find that **pretraining data scale reverses the Muon–AdamW fine-tuning gap**. At small budgets, Muon achieves lower fine-tuning loss with less forgetting than AdamW and the fine-tuning gap is negative. At larger budgets, this trend reverses: AdamW achieves a better learning–forgetting tradeoff than Muon, and the fine-tuning gap becomes positive.

In Section 3, we pretrain OLMo-style 30M–600M models (Groeneveld et al., 2024; Team OLMo et al., 2025) on 150M to 154B DCLM tokens (Li et al., 2024), comparing AdamW and Muon, and fine-tune each base model on four datasets while sweeping the fine-tuning optimizer and learning rate, for nearly 600 pretraining runs and 29,000 fine-tuning runs in total. We show that when tuning pretraining hyperparameters for the lowest pretraining loss, Muon consistently achieves lower loss than AdamW. However, with sufficient pretraining data, AdamW outperforms Muon after fine-tuning, reversing the Muon–AdamW fine-tuning gap. This reversal also holds when we disregard pretraining loss, instead sweeping pretraining hyperparameters to obtain the best learning–forgetting tradeoff after fine-tuning. We additionally show that **as model size scales, AdamW surpasses Muon at progressively smaller data budgets**.

* Equal contribution. † Equal advising.

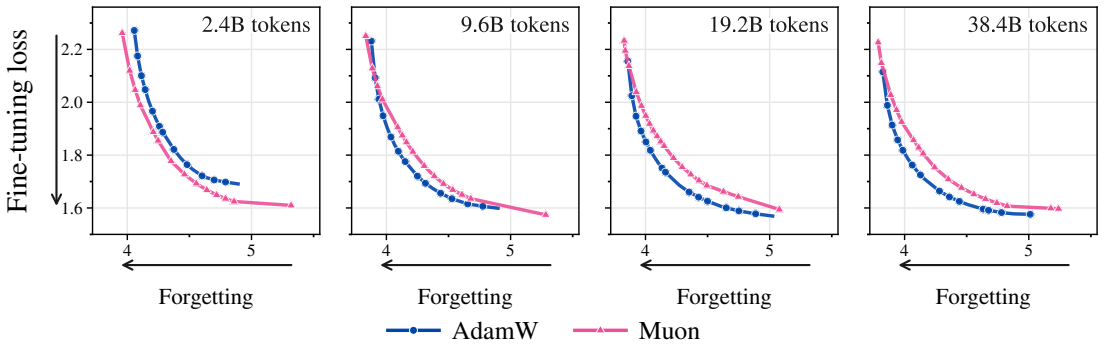


Figure 1: Beyond a sufficiently large data budget, AdamW outperforms Muon. We pretrain OLMo-60M models with Muon and AdamW and fine-tune on StackMathQA. For each pretraining optimizer, we take the pretrained model with the lowest pretraining loss, sweep across all fine-tuning configurations, and plot the resulting learning–forgetting frontier, from 2.4B to 38.4B pretraining tokens. At 2.4B tokens, Muon's frontier dominates AdamW's; at larger budgets the trend reverses and AdamW surpasses Muon. This reversal happens despite Muon achieving lower pretraining loss across all data budgets.

In Section 4, we investigate why the Muon–AdamW fine-tuning gap reverses as the pretraining data budget grows. Fine-tuning changes the output distribution, so we generalize fine-tuning updates with perturbations to the output distribution and study which optimizer produces base models that are more robust to such perturbations. To quantify sensitivity to output perturbations, we introduce the **margin-to-spread ratio (MSR)**, a simple metric derived from the model's output logits that characterizes robustness. For each prediction, the MSR is the difference between the correct logit and the average of the incorrect logits, normalized by the spread of the incorrect logits (Definition 4.1): a prediction is more robust when the MSR is larger.

To understand where this trend comes from, we study a linear associative memory trained with gradient descent (GD) and Muon (Section 4.2). Analyzing a single update from zero initialization, we prove that **the ratio of GD's MSR to Muon's increases with the data budget** (Theorem 4.2). Simulations of many training steps in the linear associative memory setting confirm the same trend: Muon has the larger MSR at small data budgets, while GD has the larger MSR at large data budgets (Section 4.2). Measuring the MSR of our pretrained language models, we find that Muon has the larger MSR at small data budgets, while AdamW has the larger MSR at large data budgets, mirroring the reversal of the fine-tuning gap (Section 4.3).

2 Preliminaries

Our goal is to compare AdamW- and Muon-pretrained models *after fine-tuning*. Since fine-tuning can induce forgetting of the pretraining distribution, performance on the fine-tuning task alone is not sufficient: one base model may learn more but also forget more. We therefore compare base models by the tradeoff between the two, and define metrics to track how this comparison changes with model size and pretraining data budget. Let θ_{PT} denote a pretrained model.

Learning–forgetting tradeoff. Fine-tuning improves downstream performance but may degrade pretrained knowledge (Goodfellow et al., 2014; Kirkpatrick et al., 2017; Springer et al., 2025). When θ_{PT} is fine-tuned to obtain θ_{FT}, we measure learning by the fine-tuning loss $\mathcal{L}_{\mathrm{FT}}(\theta_{\mathrm{FT}})$ and forgetting by the pretraining loss $\mathcal{L}_{\mathrm{PT}}(\theta_{\mathrm{FT}})$. Since both depend on the fine-tuning hyperparameters, a single fine-tuned checkpoint does not capture how well θ_{PT} can be adapted. We therefore summarize each base model by the outcomes reachable by fine-tuning it,

$$\mathcal{T}(\theta_{\mathrm{PT}}) \coloneqq \left\{ \left(\mathcal{L}_{\mathrm{FT}}(\theta_{\mathrm{FT}}), \mathcal{L}_{\mathrm{PT}}(\theta_{\mathrm{FT}})\right) \,\middle|\, \theta_{\mathrm{FT}} \in \Theta_{\mathrm{FT}}(\theta_{\mathrm{PT}}) \right\},$$

where $\Theta_{\mathrm{FT}}(\theta_{\mathrm{PT}})$ is the set of models obtained by fine-tuning θ_{PT}, and compare base models through the lower-left Pareto frontier of this set, which we call the *learning–forgetting frontier*. For a collection of base models $\mathcal{B}$, we compare the corresponding sets

$$\mathcal{T}(\mathcal{B}) \coloneqq \{\mathcal{T}(\theta_{\mathrm{PT}}) \mid \theta_{\mathrm{PT}} \in \mathcal{B}\}.$$

Fine-tuning gap. We compare the two optimizers on this tradeoff with a single metric, *fine-tuning gap* G_{FT}, which we define as the signed mean minimum Euclidean distance between the AdamW and Muon frontier. The fine-tuning gap is positive when the AdamW frontier dominates the Muon frontier, and the reverse else wise. Thus $G_{\mathrm{FT}} > 0$ when AdamW dominates and $G_{\mathrm{FT}} < 0$ when Muon dominates (Section C.5 gives the details).

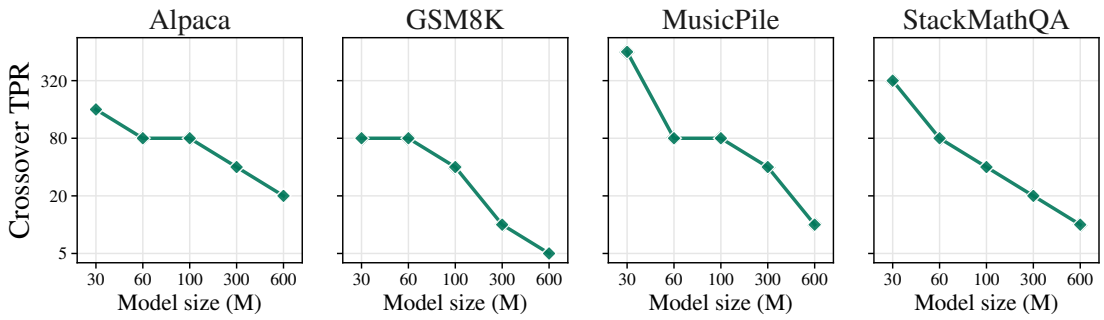


Figure 2: Crossover TPR decreases with model size. In the lowest-pretraining-loss setting, where each pretraining optimizer's lowest-loss pretrained model is swept across all fine-tuning hyperparameters, Muon's advantage reverses at substantially smaller pretraining data budgets as model size increases. The exact crossover TPR varies depending on the fine-tuning dataset, but the high-level trend is consistent.

Crossover token–parameter ratio (TPR). We quantify pretraining data budget using *token–parameter ratio* (TPR), the number of pretraining tokens per model parameter. For a fixed model size and fine-tuning dataset, the *crossover TPR* is the minimum TPR required for the AdamW frontier to overtake the Muon frontier.

3 Experiments

We now compare how the learning–forgetting tradeoffs of Muon and AdamW scale with pretraining data budget and model size. In Section 3.2, we show that as the pretraining data budget increases, Muon's advantage over AdamW shrinks and eventually reverses in favor of AdamW. This reversal happens at progressively smaller data budgets as model size grows.

3.1 Experimental setup

Pretraining. We pretrain OLMo-style 30M–600M models on 150M to 154B DCLM tokens, comparing AdamW and Muon under a warmup–stable–decay (WSD) schedule. We sweep hyperparameters broadly and report the complete pretraining methodology in Section C.

Fine-tuning. For each base model, we evaluate the learning–forgetting tradeoff on four standard datasets: GSM8K (Cobbe et al., 2021), StackMathQA (Zhang, 2024), Alpaca (Taori et al., 2023), and MusicPile (Yuan et al., 2024), sweeping AdamW learning rates from 10^{-4} to 10^{-2} and Muon learning rates from 6×10^{-4} to 2×10^{-2} (Section C.2). We fine-tune AdamW-pretrained models with AdamW and Muon-pretrained models with both AdamW and Muon: AdamW fine-tuning consistently outperforms Muon fine-tuning on AdamW-pretrained models, whereas neither optimizer is consistently better on Muon-pretrained models (Section D.3). We then plot the learning–forgetting frontier over the resulting runs. We choose a broad learning-rate range to ensure we capture the complete learning–forgetting tradeoff spectrum: higher learning rates favor learning at the expense of forgetting, while lower rates preserve pretraining knowledge but learn less.

Evaluating the learning–forgetting tradeoff. For each fine-tuned checkpoint, we compute the fine-tuning loss and pretraining loss on held-out fine-tuning and pretraining (DCLM) data, and use these checkpoints to approximate the learning–forgetting tradeoff. Although this approximation is constructed from a finite set of fine-tuning runs, we conduct a comprehensive sweep over four datasets and across a wide range of fine-tuning learning rates to obtain a close estimate. We evaluate the learning–forgetting tradeoff in two settings: *lowest-pretraining-loss setting*, and the *best-frontier setting*.

1. **Lowest-pretraining-loss setting**: Pretraining hyperparameters are typically selected to minimize pretraining loss. Therefore, for each optimizer, we tune the pretraining hyperparameters to minimize pretraining loss for a given pretraining data budget and model size. We then sweep across all fine-tuning configurations for each of the lowest-loss pretrained models, and evaluate the resulting frontier.
2. **Best-frontier setting**: However, prior studies find that the lowest-loss pretrained model is not always the best model after fine-tuning: models pretrained with other hyperparameters can offer a better learning–forgetting tradeoff (Watts et al., 2026). Therefore, for each pretraining optimizer, we additionally sweep

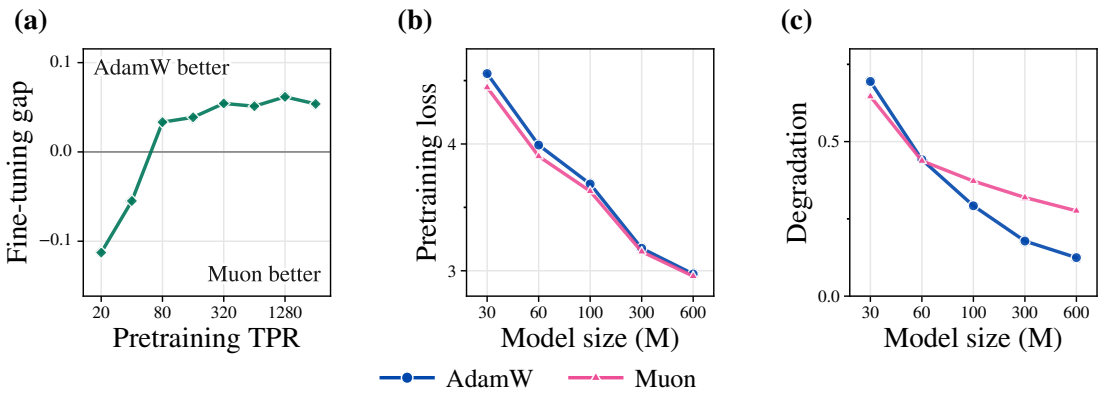


Figure 3: The Muon–AdamW fine-tuning gap widens with increasing pretraining data budget, and Muon's advantage shrinks with model size. (a) The fine-tuning gap after fine-tuning OLMo-60M models on StackMathQA, in the lowest-pretraining-loss setting (each pretraining optimizer's lowest-loss pretrained model swept across all fine-tuning configurations). As the pretraining data budget increases, the fine-tuning gap transitions from negative to positive, with AdamW outperforming Muon beyond the crossover TPR. **(b)** Pretraining loss at 40 TPR. Muon's advantage over AdamW shrinks as model size grows. **(c)** Degradation at 40 TPR, when we accept a StackMathQA loss 10% above the best in exchange for less forgetting. Muon degrades increasingly more than AdamW as model size grows.

across all pretraining and fine-tuning configurations to achieve the best possible learning–forgetting frontier for a given pretraining data budget and model size (Section C.1.3).

3.2 Main result

We now investigate how the Muon–AdamW fine-tuning gap scales with pretraining data budget and model size.

Main result: Pretraining data scale reverses the Muon–AdamW fine-tuning gap. We compare the fine-tuning gap for AdamW and Muon at different pretraining data budgets in the lowest-pretraining-loss setting, using StackMathQA as the fine-tuning dataset (Figure 3a). At small pretraining data budgets (2.4B tokens), Muon Pareto-dominates AdamW (Figure 1) and the fine-tuning gap is negative. At larger pretraining data budgets, this trend reverses: beyond 2.4B tokens, AdamW achieves lower fine-tuning loss with less forgetting than Muon. The trend also holds in the best-frontier setting (Figure 12), and for additional fine-tuning datasets in Section D.

We next investigate how the reversal of the Muon–AdamW fine-tuning gap scales with model size. Specifically, we ask how the amount of pretraining data needed for AdamW to overtake Muon changes as model size grows.

As model size scales, AdamW surpasses Muon at progressively smaller data budgets. We find that the crossover TPR decreases with model size. On StackMathQA, the crossover is 320 TPR for 30M models but only 10 TPR for 600M models (Figure 2). The specific crossover TPR depends on the fine-tuning dataset, but the broader trend persists; as model size grows, the crossover TPR decreases, and Muon's advantage reverses with substantially smaller data budgets.

We hypothesize that in the lowest-pretraining-loss setting, this scaling trend is driven by two factors: Muon's shrinking pretraining loss advantage and increased degradation relative to AdamW (Figure 3b,c). For each optimizer and model size, we compare its base pretraining loss and its degradation $\mathcal{L}_{\mathrm{PT}}(\theta_{\mathrm{FT}}) - \mathcal{L}_{\mathrm{PT}}(\theta_{\mathrm{PT}})$, after fine-tuning on StackMathQA. First, for a fixed TPR, the gap in pretraining loss between Muon and AdamW shrinks as model size increases, reducing Muon's advantage as a lower-loss base model. Additionally, for a fixed TPR and fine-tuning loss, Muon forgets increasingly more than AdamW as model size grows. Consequently, as model size scales, Muon's pretraining advantage over AdamW narrows, while its forgetting relative to AdamW increases, enabling AdamW to overtake Muon at progressively smaller data budgets.

4 Analyzing the Muon–AdamW Forgetting Gap

In Section 3, we show that Muon offers a better learning–forgetting tradeoff at small token budgets, but AdamW overtakes it beyond a crossover budget. This reversal occurs even though Muon achieves lower pretraining loss at

every budget we consider. In this section, we develop a theoretical explanation for this reversal. We compare Muon with exact orthogonalization to full-batch gradient descent (GD), a tractable comparison that captures the qualitative budget trend of Muon and AdamW in language models (Section 4.3).

To formalize the relationship between GD and Muon, we study a linear associative memory model of fact learning. In this setting, GD and Muon learn the facts to a similar degree, but their checkpoints can differ in their sensitivity to perturbations of the output distribution. Since fine-tuning changes the output distribution, we expect checkpoints that preserve factual accuracy under a wider range of perturbations to be more resistant to forgetting during fine-tuning. Our central goal is therefore to characterize this sensitivity across pretraining budgets and use it to explain why the forgetting gap reverses.

To quantify sensitivity to output perturbations, we introduce the **margin-to-spread ratio (MSR)**: the difference between the correct answer's logit and the average incorrect logit, normalized by the spread of the incorrect logits. The MSR captures a checkpoint's robustness to output perturbations and provides an analytically tractable basis for comparing GD and Muon.

Roadmap. We introduce the linear associative memory setting, define the MSR, and establish its connection to sensitivity (Section 4.1). We then prove that, after one update step, the ratio of Muon's MSR to GD's decreases as the data budget grows, indicating that Muon becomes more sensitive to output perturbations relative to GD (Section 4.2). Finally, we show that the MSR tracks the robustness of AdamW and Muon in real language models (Section 4.3).

4.1 Setup and the margin-to-spread ratio

We begin by introducing our linear associative memory setup and defining the margin-to-spread ratio.

Linear associative memory. To capture the behavior of Muon and GD in a tractable setting, we consider a linear associative memory, in which a single weight matrix is trained to associate each of N keys, represented by a vector, with its corresponding value, also represented by a vector. Formally, we represent each fact by a key–value pair $(u_i, v_i) \in \mathbb{R}^d \times \mathbb{R}^d$, with all keys and values drawn independently from $\mathcal{N}(0, I_d)$. The memory is parameterized by a matrix $W \in \mathbb{R}^{d\times d}$. Given a key u_i, it scores each candidate value v_j as $s_{ij}(W) := v_j^\top W u_i$, which we treat as logits, and retrieves the highest-scoring value. Fact i is recovered when $s_{ii}(W) > s_{ij}(W)$ for all $j \neq i$.

We train with the softmax cross-entropy loss. Given a query distribution $q = (q_1, \ldots, q_N)$ over the facts, the predicted probabilities and cross-entropy loss are

$$\hat{p}_W(\cdot \mid i) := \mathrm{softmax}\big(s_{i1}(W), \ldots, s_{iN}(W)\big), \quad \mathcal{L}(W) := -\textstyle\sum_{i=1}^N q_i \log \hat{p}_W(i \mid i). \tag{1}$$

Data budget. We denote the memory dimension by d and the number of facts by N, and grow the two together, keeping their ratio $\gamma = N/d > 1$ fixed, so that $N = \gamma d$ as $d \to \infty$. This mirrors the typical scaling behavior of language models, which store many more facts than they have dimensions. We call γ the **data budget**: it measures facts per dimension and plays the role of token-parameter ratio in Section 2. We compare data budgets by taking this limit at different fixed values of γ. Since $\gamma > 1$, we have $N - 1 \geq d$, and $\log N/d \to 0$, written $\log N = o(d)$. For random variables X and a deterministic $a > 0$, $X = O_p(a)$ means that for every $\varepsilon > 0$ there is a C with $\Pr[|X| > Ca] \leq \varepsilon$ eventually, and $X = o_p(a)$ means that $\Pr[|X| > \varepsilon a] \to 0$ for every $\varepsilon > 0$.

Optimization. For simplicity, we consider the exact orthogonalization variant of Muon. For a matrix A with singular value decomposition $A = P\Sigma Q^\top$, let $\mathrm{polar}(A) := PQ^\top$. Muon updates $W_{t+1} = W_t - \eta\,\mathrm{polar}(\nabla\mathcal{L}(W_t))$. We compare Muon to full-batch gradient descent (GD), whose updates descend along the gradient itself, $W_{t+1} = W_t - \eta\nabla\mathcal{L}(W_t)$.

Margin-to-spread ratio. In a binary classification problem, a prediction is robust if the correct logit is far from the incorrect logit: the larger the margin, the more robust the prediction. In a multiclass classification problem, such as our linear associative memory setting, there are many incorrect logits. The robustness of a correct output to perturbations of the logits thus depends on its separation from all of the incorrect logits. As it turns out, the degree of separation can be usefully summarized by a single statistic, which we call the *margin-to-spread ratio (MSR)*. MSR captures the average separation between the correct logit and all other logits, normalized by the spread of the incorrect logits. A smaller spread, or a larger average separation, would both place the incorrect logits further below the correct logit relative to their spread, which can increase the robustness of the prediction to perturbations. To capture this intuition, we formally define the margin-to-spread ratio as follows:

Definition 4.1 (Margin-to-spread ratio). For fact i, define the pairwise margins $m_{ij} := s_{ii}(W) - s_{ij}(W)$ for $j \neq i$, with mean and variance

$$h_i := \operatorname*{mean}_{j \neq i} m_{ij}, \qquad \kappa_i^2 := \operatorname*{mean}_{j \neq i} (m_{ij} - h_i)^2. \tag{2}$$

The *margin-to-spread ratio* (MSR) for fact i is $\rho_i := h_i / \kappa_i$, for $\kappa_i > 0$. At the model level, we average the mean margins and variances uniformly over facts before taking the ratio:

$$\bar{h} := \operatorname*{mean}_i h_i, \qquad \bar{\kappa}^2 := \operatorname*{mean}_i \kappa_i^2, \qquad \rho := \frac{\bar{h}}{\bar{\kappa}}, \tag{3}$$

where ρ is defined for $\bar{\kappa} > 0$.

Since $s_{ii}(W)$ is the same for every competitor j, κ_i is also the standard deviation of the incorrect logits $s_{ij}(W)$, so the MSR measures the average margin in units of their spread. The definition applies to any model that assigns a logit to each candidate answer; for a language model, $s_{ij}(W)$ are the next-token logits (Section 4.3).

Under isotropic Gaussian noise added to the logits, the MSR gives a lower bound on the fraction of incorrect answers that stay below the correct one: with noise of standard deviation σ on each logit, this fraction is at least $\rho_i^2/(\rho_i^2 + 1 + 2\sigma^2/\kappa_i^2)$ in expectation. For a fixed noise strength relative to the incorrect-logit spread, this bound improves as the MSR increases (Corollary B.5).

4.2 One-step analysis of the margin-to-spread ratio

We have so far introduced our linear associative memory setup and built intuition for how the MSR captures the robustness of a model to perturbations of its logits. Here, we argue that the MSR of GD is greater than that of Muon for sufficiently large data budgets. We precisely track the one-step dynamics of the MSR ratio $\mathcal{R} := \rho^{\mathrm{GD}}/\rho^{\mathrm{Muon}}$ between GD and Muon in the linear associative memory.

We analyze a single update from zero initialization, using uniform query weights $q_i = 1/N$. The weights after many optimization steps depend on the whole training trajectory and are hard to characterize precisely, so theoretical studies of how an optimizer shapes the learned weights often analyze its first step (Ba et al., 2022; Dandi et al., 2024; Moniri et al., 2024). Kim et al. (2026) take the same approach to compare a single step of Muon, SGD, and Newton's method in a linear associative memory.

Theorem 4.2 (Informal). *Consider one GD step and one Muon step from zero in the Gaussian associative memory model, with data budget $\gamma = N/d$. For sufficiently large γ:*

- *GD has a larger MSR than Muon.*
- *Increasing the data budget γ by any fixed percentage increases the ratio $\mathcal{R} = \rho^{\mathrm{GD}}/\rho^{\mathrm{Muon}}$, with high probability as the dimension grows.*
- *The ratio approaches $3\pi/8 \approx 1.18$ from below at a rate of $1/\gamma$:*

$$-\tfrac{C}{\gamma} \leq \mathcal{R} - \tfrac{3\pi}{8} \leq -\tfrac{c}{\gamma},$$

  where $0 < c < C$ are absolute constants.

All claims hold with probability tending to one as $d \to \infty$, for sufficiently large fixed γ.

Proof sketch. Our goal is to compare one update step for Muon and for GD.

1. Since the MSR does not change when W is multiplied by a positive constant, the one-step weights can be written as

$$W^{\mathrm{GD}} = G, \qquad W^{\mathrm{Muon}} = O := \mathrm{polar}(G)$$

   (Appendix, Lemma F.1). We now want to compute $\rho = \bar{h}/\bar{\kappa}$ for these matrices. This requires the mean margins and their standard deviations; we begin with the means.

2. *Compute the mean margins.* Recall that $\bar{h}$ averages the difference between the correct and incorrect logits over facts and incorrect answers. For uniformly weighted facts,

$$G = \textstyle\sum_{k=1}^{N} (v_k - \bar{v}) u_k^\top, \qquad \bar{v} = \frac{1}{N} \sum_{k=1}^{N} v_k.$$

Substituting into this average gives $\bar{h}(W) = \langle W, G\rangle/(N-1)$, and therefore

$$\bar{h}^{\mathrm{GD}} = \tfrac{\|G\|_F^2}{N-1}, \qquad \bar{h}^{\mathrm{Muon}} = \tfrac{\|G\|_*}{N-1},$$

where $\|\cdot\|_F$ is the Frobenius norm and $\|\cdot\|_*$ is the sum of the singular values.

3. *Compute GD's variance.* For each query i, we take the variance across incorrect candidates j, holding the dataset and weights fixed:

$$\kappa_i^2 = \mathrm{Var}_{j\neq i}(v_j^\top G u_i).$$

To calculate its typical size, expand the incorrect logit:

$$v_j^\top G u_i \approx \underbrace{\textstyle\sum_{k\notin\{i,j\}}(v_j^\top v_k)(u_k^\top u_i)}_{\text{other facts}} + \underbrace{\|u_i\|^2(v_j^\top v_i) + \|v_j\|^2(u_j^\top u_i)}_{\text{query and candidate facts}}.$$

Here $\approx$ omits the mean-subtraction term in G, which only affects lower-order terms in the variance. Each other fact contributes a product of two independent, mean-zero inner products, each with variance d, giving variance d^2. In the query and candidate terms, we approximate the squared norms by d, dropping their fluctuations. This gives variance d^3 each, with $O(d^2)$ corrections. Replacing $N-2$ by N also discards only $O(d^2)$ terms.

These contributions are uncorrelated over the Gaussian training data. Including the omitted corrections and concentration of the empirical average, we calculate GD's variance (Appendix, Lemma F.15) to be

$$(\bar{\kappa}^{\mathrm{GD}})^2 = \big(\underbrace{Nd^2}_{\text{other facts}} + \underbrace{2d^3}_{\text{query and candidate}}\big)(1+o_p(1)) = d^3(\gamma+2)(1+o_p(1)).$$

4. *Compute Muon's variance.* Again, the variance is over incorrect candidates for each fixed query. Gaussian isotropy and orthogonality give the leading calculation

$$(\bar{\kappa}^{\mathrm{Muon}})^2 \approx \tfrac{1}{N(N-1)}\textstyle\sum_i\sum_{j\neq i}(v_j^\top O u_i)^2 \approx \tfrac{1}{N}\sum_i\|O u_i\|^2 = \tfrac{1}{N}\sum_i\|u_i\|^2 \approx d.$$

The first approximation drops the squared mean of the incorrect logits for each query. The second neglects departures of the sampled incorrect values from perfect isotropy, and the last drops fluctuations in the average squared key length. We calculate the variance including the omitted corrections and the dependence of O on the training data (Appendix, Lemma F.31) to be

$$(\bar{\kappa}^{\mathrm{Muon}})^2 = d\left[1 + O(\gamma^{-3/2}) + o_p(1)\right].$$

5. *Put the means and standard deviations together.* We calculate the norms in the mean margins using the Gaussian singular-value distribution (Appendix, Lemma F.4). Dividing by the standard deviations gives, to first order in $1/\gamma$,

$$\rho^{\mathrm{GD}} \approx \sqrt{\tfrac{d}{\gamma}}\left(1-\tfrac{1}{\gamma}\right), \qquad \rho^{\mathrm{Muon}} \approx \tfrac{8}{3\pi}\sqrt{\tfrac{d}{\gamma}}\left(1-\tfrac{1}{10\gamma}\right).$$

These approximations omit higher-order budget corrections and relative fluctuations that vanish as $d\to\infty$. Taking their ratio cancels the common factor $\sqrt{d/\gamma}$. Keeping track of the errors gives

$$\mathcal{R} = \tfrac{3\pi}{8}\left(1-\tfrac{9}{10\gamma}\right) + O(\gamma^{-3/2}) + o_p(1).$$

Since $3\pi/8 > 1$, GD has the larger MSR at sufficiently large budgets. The negative $1/\gamma$ term in $\mathcal{R} - 3\pi/8$ dominates the smaller remainder, giving the claimed rate. A fixed percentage increase in γ raises the leading expression by order $1/\gamma$, which also dominates the remainder, establishing the increase with high probability. See Section F for the quantitative bound and full proof, and Figure 4 and Section E.1 for numerical checks.

This theorem demonstrates that in the linear associative memory, when the data budget is sufficiently large, checkpoints learned with GD have a larger MSR than checkpoints learned with Muon after a single step of training. Combined with the robustness bound of Section 4.1, GD's checkpoints therefore carry the stronger guarantee of robustness to perturbations of the output distribution in our theoretical setting.

Remark (Choice of data distribution). Independent standard Gaussian keys and values provide a canonical isotropic model with no preferred direction. Other data distributions can change both the gradient spectrum and the competing-score variances, so the optimizer ordering may differ. The analogous budget trend in the language-model experiments (Section 4.3) suggests that the behavior captured by this simple model is relevant beyond the Gaussian setting.

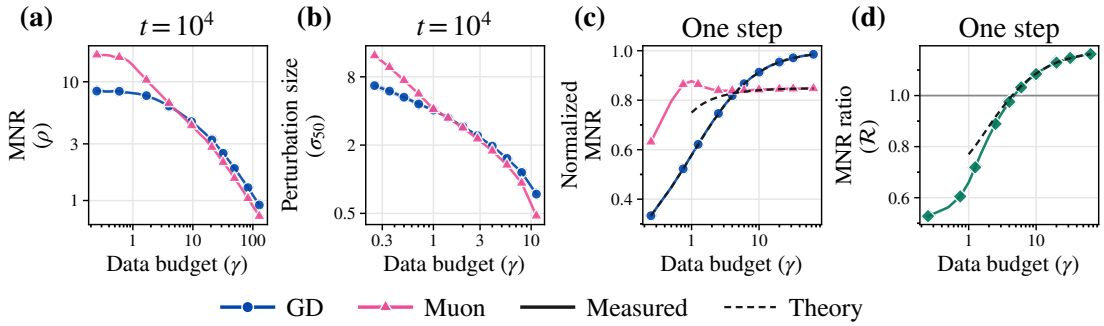


Figure 4: The margin-to-spread ratio of GD gains on Muon's as the data budget grows. (a, b) After $t = 10^4$ steps with model dimension $d = 256$, GD's MSR ρ overtakes Muon's as γ grows, and the ordering of the perturbed robustness σ_{50} turns over with it. **(c, d)** After one step with $d = 512$, GD's normalized MSR $\rho/\sqrt{d/\gamma}$ overtakes Muon's (c), so the MSR ratio $\mathcal{R} = \rho^{\mathrm{GD}}/\rho^{\mathrm{Muon}}$ rises with γ (d), following the theory of Theorem 4.2.

Beyond one step. We next check whether the trend persists after extended training. We train the linear memory for $T = 10^4$ full-batch steps from $W_0 = 0$ at constant learning rates across three seeds. We summarize the saturated MSR by the per-seed median over $t \geq 10^3$ and measure robustness by σ_{50}, the relative perturbation size at which half of the initially recovered facts are lost. We exclude data budgets past the memory's capacity where the two optimizers recover different item sets. The full setup and results for every dimension d appear in Section E.1.

The results in Figure 4 show that $\mathcal{R}$ increases with the data budget. Muon has the larger MSR at small γ, while GD has the larger MSR at large γ. The robustness comparison shifts in the same direction: Muon tolerates stronger perturbations at small γ, and GD does at large γ. Thus, the budget trend persists in the linear setting with GD as the baseline.

4.3 Connection to practice

We test whether the budget trend predicted by Theorem 4.2 also appears in real language models, beyond the assumptions used in our theoretical analysis.

Setup. We evaluate the AdamW and Muon models trained in Section 3, spanning five model sizes, on held-out DCLM text. We compute the MSR at each token position and report the median over correctly predicted positions. At each position, h and κ are the mean and standard deviation of the margins between the correct next-token logit and all other vocabulary logits. We use the median to reduce sensitivity to outliers; Definition 4.1 instead uses the mean margin divided by the square root of the mean variance.

Experimental Result. Figure 5 plots the MSR for the three model sizes with the widest budget ranges, from 30M to 100M. Muon has a larger MSR at small data budgets, while AdamW has a larger MSR at sufficiently large budgets. As the budget grows, the difference between AdamW's and Muon's MSRs increases, and the ratio of AdamW's MSR to Muon's crosses one. This shift toward AdamW is consistent with the theoretical budget trend and parallels the reversal in relative robustness to fine-tuning observed in Section 3.

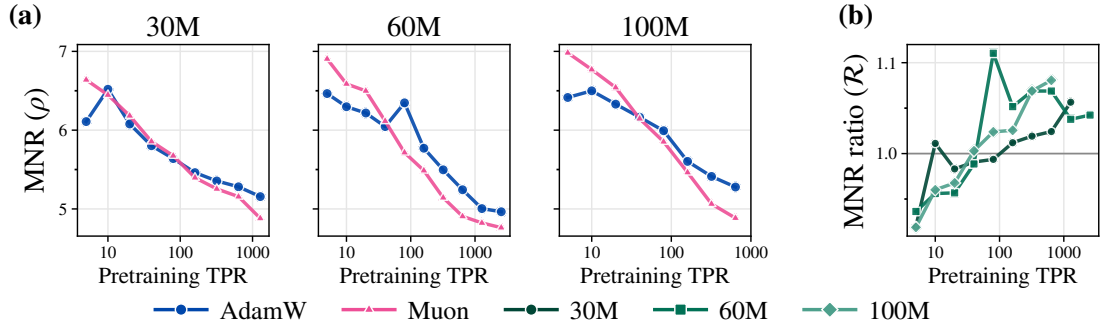


Figure 5: In real language models, the margin-to-spread ratio of AdamW gains on Muon's as the token budget grows. (a) MSR ρ of AdamW and Muon against pretraining TPR for 30M, 60M, and 100M models: Muon has the larger MSR at small budgets and AdamW at large ones. **(b)** The MSR ratio $\rho^{\mathrm{AdamW}}/\rho^{\mathrm{Muon}}$ crosses one as the budget grows, at every model size.

The MSR and fine-tuning crossovers do not necessarily appear at exactly the same budget, since the learning–forgetting tradeoff also depends on pretraining loss. Nevertheless, the MSR captures the direction of the observed shift and connects our theoretical analysis to the behavior of trained language models.

Remark. We summarize the analysis of this section and relate MSR to other metrics used to predict post-training.

- **Takeaway.** In a linear associative memory, we prove that after one update the ratio $\rho^{\mathrm{GD}}/\rho^{\mathrm{Muon}}$ of the two optimizers' MSRs increases with the data budget (Theorem 4.2). Simulations of many training steps show that the same trend persists after extended training (Section 4.2).
- **Extension: Relation of MSR to Other Metrics Predicting Post-Training.** Other work predicts forgetting after post-training with different metrics. For example, Springer et al. (2025) show that the loss increase after adding Gaussian noise to the pretrained weights tracks how much a model degrades under fine-tuning. In Section B, we show that MSR is connected to this robustness to Gaussian perturbations: in the linear memory, a larger MSR provably guarantees that more competitors stay ranked below the correct answer after the perturbation (Theorem B.3), and in both language models and the linear memory, MSR is strongly correlated with perturbation robustness (Section B.2.2).

5 Related Work

Muon. Muon, an optimizer introduced by Jordan et al. (2024), has recently shown strong empirical performance for pretraining LLMs, challenging AdamW's dominance as the status-quo optimizer. Recent work has sought to explain Muon's pretraining advantage through analyses of simplified theoretical models, associative memory, convergence behavior, and loss landscape geometry (Kim et al., 2026; Li et al., 2026; Sato et al., 2025; Su, 2025; Shen et al., 2025; Wang et al., 2026a). Another line of work investigates the training dynamics of Muon, finding that it learns data components more uniformly than AdamW (Vasudeva et al., 2026; Wang et al., 2026b; Dragutinović et al., 2026), while also characterizing the implicit biases of Muon (Fan et al., 2025; Dragutinović et al., 2026). Yet fewer works have examined the downstream consequences of Muon. Ruan et al. (2026) compare Muon and AdamW after fine-tuning; they find that for small data regimes, Muon outperforms AdamW, consistent with our findings. We show that this advantage does not generalize across pretraining data scale; after a sufficiently large pretraining data budget, the trend reverses, and AdamW outperforms Muon.

Muon versus AdamW on pretraining loss. Muon has been reported to reach lower pretraining loss than AdamW for the same compute across a wide range of settings (Liu et al., 2025; Essential AI et al., 2025; Mehta et al., 2025), and frontier models have since been pretrained with Muon or its variants (Kimi Team, 2025; GLM-4.5 Team et al., 2025). The size of this advantage depends on scale: under controlled tuning, Muon's speedup over AdamW shrinks as model size and data grow, and Muon is overtaken by other matrix-based optimizers at large data-to-model ratios (Wen et al., 2026b;a). All of these works measure pretraining loss. Optimizers can also reach qualitatively different solutions rather than faster or slower routes to the same minimum (Pascanu et al., 2025); we ask whether these differences matter after fine-tuning.

Catastrophic forgetting. Catastrophic forgetting, the degradation of previously learned knowledge during subsequent training, remains a central challenge in deep learning (Goodfellow et al., 2014). Numerous approaches have been proposed to mitigate catastrophic forgetting, including regularization-based methods (Kirkpatrick et al., 2017; Zenke et al., 2017; Aljundi et al., 2018), replay techniques (Shin et al., 2017), gradient projection methods (Lopez-Paz & Ranzato, 2017; Chaudhry et al., 2019), and optimization-based approaches (Watts et al., 2026). Our learning–forgetting tradeoff is an instance of the stability–plasticity dilemma, with fine-tuning performance capturing plasticity and knowledge retention capturing stability (Lopez-Paz & Ranzato, 2017). We study forgetting as an intrinsic property of the pretrained checkpoint, rather than as a consequence of the fine-tuning algorithm. Specifically, we investigate how optimization choices during pretraining shape downstream forgetting, showing that Muon and AdamW induce different learning–forgetting tradeoffs.

Pretraining and downstream adaptation. Better pretraining loss does not always yield a better model after fine-tuning. Springer et al. (2025) show that extending pretraining can make language models harder to fine-tune, as the pretrained parameters become increasingly sensitive to modifications, which they measure in part with Gaussian perturbations. Watts et al. (2026) show that sharpness-aware pretraining mitigates forgetting during fine-tuning, and Liu et al. (2026) find that fine-tuning with the same optimizer as pretraining forgets less. We study how the choice of pretraining optimizer interacts with the pretraining data budget, and connect the resulting fine-tuning gap to robustness to weight perturbations.

6 Conclusion

We asked whether the quality of Muon base models persists through fine-tuning, and found that the answer depends on the pretraining budget: beyond a crossover budget, AdamW outperforms Muon after fine-tuning, even though Muon reaches lower pretraining loss. We do not necessarily expect this for every fine-tuning distribution, since fine-tuning on data close to the pretraining data is simply continued pretraining. Still, the trend holds broadly on all datasets we study, including math, instruction following, and music. Furthermore, the crossover budget decreases with model size, which suggests that the problem persists as we scale.

We explain why AdamW outperforms Muon at large budgets by tracking how robust each model is to small isotropic perturbations of its output distribution, which we characterize with the margin-to-spread ratio (MSR). AdamW becomes increasingly more robust than Muon as the data budget grows, and this robustness is a property of the output distribution that the training loss does not capture. Since the MSR measures robustness to perturbations of the output distribution in general, our theory suggests that this robustness is fundamental to the checkpoint learned by each optimizer, rather than specific to fine-tuning, to the perturbations we study, or to a particular architecture or parameterization. It should therefore apply to any update that changes the output distribution.

Although outside the scope of this work, our analysis suggests that pretraining should explicitly optimize for properties beyond the pretraining loss. Optimizer selection requires more than comparing pretraining loss, and we can design objectives that explicitly target adaptability, for which the MSR is a potential candidate.

More broadly, we expect this lesson to extend beyond Muon. Many pretraining choices, such as hyperparameters, data, and architecture, are selected by pretraining loss. Yet, our results show that pretraining loss can be a poor proxy for adaptability, increasingly so as the data budget grows. This suggests that for any of these choices, a lower pretraining loss does not by itself yield a model that can better adapt.

Acknowledgements

We gratefully acknowledge support from Apple, Google, Jane Street, the National Science Foundation and the FLAME cluster at Carnegie Mellon University.

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE2140739. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

The authors thank Kaiyue Wen, Alex Damian, Ishaan Watts, Lawrence Feng, and Gene Yang for helpful discussions.

References

Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 139–154, 2018.

Jimmy Ba, Murat A. Erdogdu, Taiji Suzuki, Zhichao Wang, Denny Wu, and Greg Yang. High-dimensional asymptotics of feature learning: How one gradient step improves the representation. In *Advances in Neural Information Processing Systems*, 2022. URL https://arxiv.org/abs/2205.01445.

Arslan Chaudhry, Marc'Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with A-GEM. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=Hkf2_sC5FX.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems, 2021. URL https://arxiv.org/abs/2110.14168.

Yatin Dandi, Florent Krzakala, Bruno Loureiro, Luca Pesce, and Ludovic Stephan. How two-layer neural networks learn, one (giant) step at a time. *Journal of Machine Learning Research*, 25(349):1–65, 2024. URL https://jmlr.org/papers/v25/23-1543.html.

Sara Dragutinović, Yedi Zhang, and Rajesh Ranganath. To use or not to use Muon: How simplicity bias in optimizers matters, 2026. URL https://arxiv.org/abs/2603.00742.

Essential AI, Ishaan Shah, Anthony M. Polloreno, Karl Stratos, Philip Monk, Adarsh Chaluvaraju, et al. Practical efficiency of Muon for pretraining, 2025. URL https://arxiv.org/abs/2505.02222.

Chen Fan, Mark Schmidt, and Christos Thrampoulidis. Implicit bias of spectral descent and Muon on multiclass separable data. In *Advances in Neural Information Processing Systems*, 2025. URL https://arxiv.org/abs/2502.04664.

GLM-4.5 Team, Aohan Zeng, Xin Lv, Qinkai Zheng, Zhenyu Hou, Bin Chen, et al. GLM-4.5: Agentic, reasoning, and coding (ARC) foundation models, 2025. URL https://arxiv.org/abs/2508.06471.

Ian J. Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. An empirical investigation of catastrophic forgetting in gradient-based neural networks. In *International Conference on Learning Representations*, 2014. URL https://arxiv.org/abs/1312.6211.

Dirk Groeneveld, Iz Beltagy, Pete Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Harsh Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, Shane Arora, David Atkinson, Russell Authur, Khyathi Raghavi Chandu, Arman Cohan, Jennifer Dumas, Yanai Elazar, Yuling Gu, Jack Hessel, Tushar Khot, William Merrill, Jacob Morrison, Niklas Muennighoff, Aakanksha Naik, Crystal Nam, Matthew E. Peters, Valentina Pyatkin, Abhilasha Ravichander, Dustin Schwenk, Saurabh Shah, Will Smith, Emma Strubell, Nishant Subramani, Mitchell Wortsman, Pradeep Dasigi, Nathan Lambert, Kyle Richardson, Luke Zettlemoyer, Jesse Dodge, Kyle Lo, Luca Soldaini, Noah A. Smith, and Hannaneh Hajishirzi. OLMo: Accelerating the science of language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 15789–15809, 2024. URL https://aclanthology.org/2024.acl-long.841.

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. In *Advances in Neural Information Processing Systems*, volume 35, 2022. URL https://arxiv.org/abs/2203.15556.

Keller Jordan, Yuchen Jin, Vlado Boza, Jiacheng You, Franz Cesista, Laker Newhouse, and Jeremy Bernstein. Muon: An optimizer for hidden layers in neural networks. Blog post, 2024. URL https://kellerjordan.github.io/posts/muon/.

Juno Kim, Eshaan Nichani, Denny Wu, Alberto Bietti, and Jason D. Lee. Sharp capacity scaling of spectral optimizers in learning associative memory, 2026. URL https://arxiv.org/abs/2603.26554.

Kimi Team. Kimi K2: Open agentic intelligence, 2025. URL https://arxiv.org/abs/2507.20534.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015. URL https://arxiv.org/abs/1412.6980.

James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13):3521–3526, March 2017. ISSN 1091-6490. doi: 10.1073/pnas.1611835114. URL http://dx.doi.org/10.1073/pnas.1611835114.

Binghui Li, Kaifei Wang, Han Zhong, Pinyan Lu, and Liwei Wang. Muon in associative memory learning: Training dynamics and scaling laws. In *International Conference on Machine Learning*, 2026. URL https://arxiv.org/abs/2602.05725.

Jeffrey Li, Alex Fang, Georgios Smyrnis, Maor Ivgi, Matt Jordan, Samir Yitzhak Gadre, Hritik Bansal, Etash Kumar Guha, Sedrick Keh, Kushal Arora, Saurabh Garg, Rui Xin, Niklas Muennighoff, Reinhard Heckel, Jean Mercat, Mayee F Chen, Suchin Gururangan, Mitchell Wortsman, Alon Albalak, Yonatan Bitton, Marianna Nezhurina, Amro Kamal Mohamed Abbas, Cheng-Yu Hsieh, Dhruba Ghosh, Joshua P Gardner, Maciej Kilian, Hanlin Zhang, Rulin Shao, Sarah M Pratt, Sunny Sanyal, Gabriel Ilharco, Giannis Daras, Kalyani Marathe, Aaron Gokaslan, Jieyu Zhang, Khyathi Chandu, Thao Nguyen, Igor Vasiljevic, Sham M. Kakade, Shuran Song, Sujay Sanghavi, Fartash Faghri, Sewoong Oh, Luke Zettlemoyer, Kyle Lo, Alaaeldin El-Nouby, Hadi Pouransari, Alexander T Toshev, Stephanie Wang, Dirk Groeneveld, Luca Soldaini, Pang Wei Koh, Jenia Jitsev, Thomas Kollar, Alex Dimakis, Yair Carmon, Achal Dave, Ludwig Schmidt, and Vaishaal Shankar. Datacomp-LM: In search of the next generation of training sets for language models. In *The Thirty-eighth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2024. URL https://openreview.net/forum?id=CNWdWn47IE.

Jingyuan Liu, Jianlin Su, Xingcheng Yao, Zhejun Jiang, Guokun Lai, Yulun Du, Yidao Qin, Weixin Xu, Enzhe Lu, Junjie Yan, Yanru Chen, Huabin Zheng, Yibo Liu, Shaowei Liu, Bohong Yin, Weiran He, Han Zhu, Yuzhi Wang, Jianzhou Wang, Mengnan Dong, Zheng Zhang, Yongsheng Kang, Hao Zhang, Xinran Xu, Yutao Zhang, Yuxin Wu, Xinyu Zhou, and Zhilin Yang. Muon is scalable for LLM training, 2025. URL https://arxiv.org/abs/2502.16982.

Yuxing Liu, Jianyu Wang, and Tong Zhang. Optimizer-model consistency: Full finetuning with the same optimizer as pretraining forgets less, 2026. URL https://arxiv.org/abs/2605.06654.

David Lopez-Paz and Marc'Aurelio Ranzato. Gradient episodic memory for continual learning. *Advances in neural information processing systems*, 30, 2017.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations (ICLR)*, 2019.

Sushant Mehta, Raj Dandekar, Rajat Dandekar, and Sreedath Panat. Muon: Training and trade-offs with latent attention and MoE, 2025. URL https://arxiv.org/abs/2509.24406.

Behrad Moniri, Donghwan Lee, Hamed Hassani, and Edgar Dobriban. A theory of non-linear feature learning with one gradient step in two-layer neural networks. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, 2024. URL https://proceedings.mlr.press/v235/moniri24a.html.

Razvan Pascanu, Clare Lyle, Ionut-Vlad Modoranu, Naima Elosegui Borras, Dan Alistarh, Petar Velickovic, Sarath Chandar, Soham De, and James Martens. Optimizers qualitatively alter solutions and we should leverage this, 2025. URL https://arxiv.org/abs/2507.12224.

Tianyu Ruan, Fengzhuo Zhang, Shuche Wang, and Shihua Zhang. Muon learns more robust and transferable features than Adam, 2026. URL https://arxiv.org/abs/2606.09658.

Naoki Sato, Hiroki Naganuma, and Hideaki Iiduka. Convergence bound and critical batch size of Muon optimizer, 2025. URL https://arxiv.org/abs/2507.01598.

Wei Shen, Ruichuan Huang, Minhui Huang, Cong Shen, and Jiawei Zhang. On the convergence analysis of Muon, 2025. URL https://arxiv.org/abs/2505.23737.

Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. In *Advances in Neural Information Processing Systems*, volume 30, 2017. URL https://arxiv.org/abs/1705.08690.

Jacob Mitchell Springer, Sachin Goyal, Kaiyue Wen, Tanishq Kumar, Xiang Yue, Sadhika Malladi, Graham Neubig, and Aditi Raghunathan. Overtrained language models are harder to fine-tune. In *International Conference on Machine Learning*, 2025. URL https://arxiv.org/abs/2503.19206.

Weijie Su. Isotropic curvature model for understanding deep learning optimization: Is gradient orthogonalization optimal?, 2025. URL https://arxiv.org/abs/2511.00674.

Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. Stanford Alpaca: An instruction-following LLaMA model. https://github.com/tatsu-lab/stanford_alpaca, 2023.

Team OLMo, Pete Walsh, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Shane Arora, Akshita Bhagia, Yuling Gu, Shengyi Huang, Matt Jordan, Nathan Lambert, Dustin Schwenk, Oyvind Tafjord, Taira Anderson, David Atkinson, Faeze Brahman, Christopher Clark, Pradeep Dasigi, Nouha Dziri, Allyson Ettinger, Michal Guerquin, David Heineman, Hamish Ivison, Pang Wei Koh, Jiacheng Liu, Saumya Malik, William Merrill, Lester James V. Miranda, Jacob Morrison, Tyler Murray, Crystal Nam, Jake Poznanski, Valentina Pyatkin, Aman Rangapur, Michael Schmitz, Sam Skjonsberg, David Wadden, Christopher Wilhelm, Michael Wilson, Luke Zettlemoyer, Ali Farhadi, Noah A. Smith, and Hannaneh Hajishirzi. 2 OLMo 2 furious. In *Conference on Language Modeling*, 2025. URL https://arxiv.org/abs/2501.00656.

Bhavya Vasudeva, Puneesh Deora, Yize Zhao, Vatsal Sharan, and Christos Thrampoulidis. How Muon's spectral design benefits generalization: A study on imbalanced data. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=YzjS4jcfmS.

Shuche Wang, Fengzhuo Zhang, Jiaxiang Li, Dirk Bergemann, and Zhuoran Yang. Why Muon outperforms Adam: A curvature perspective, 2026a. URL https://arxiv.org/abs/2606.04662.

Shuche Wang, Fengzhuo Zhang, Jiaxiang Li, Cunxiao Du, Chao Du, Tianyu Pang, Zhuoran Yang, Mingyi Hong, and Vincent Y. F. Tan. Muon outperforms Adam in tail-end associative memory learning. In *The Fourteenth International Conference on Learning Representations*, 2026b. URL https://arxiv.org/abs/2509.26030.

Ishaan Watts, Catherine Li, Sachin Goyal, Jacob Mitchell Springer, and Aditi Raghunathan. Sharpness-aware pretraining mitigates catastrophic forgetting. In *Proceedings of the 43rd International Conference on Machine Learning*, 2026. URL https://arxiv.org/abs/2605.02105.

Kaiyue Wen, Xingyu Dang, Kaifeng Lyu, Tengyu Ma, and Percy Liang. Fantastic pretraining optimizers and where to find them II: Hyperball optimization, 2026a. URL https://arxiv.org/abs/2606.16899.

Kaiyue Wen, David Hall, Tengyu Ma, and Percy Liang. Fantastic pretraining optimizers and where to find them. In *The Fourteenth International Conference on Learning Representations*, 2026b. URL https://openreview.net/forum?id=2J51qUZ0iG.

Ruibin Yuan, Hanfeng Lin, Yi Wang, Zeyue Tian, Shangda Wu, Tianhao Shen, Ge Zhang, Yuhang Wu, Cong Liu, Ziya Zhou, Ziyang Ma, Liumeng Xue, Ziyu Wang, Qin Liu, Tianyu Zheng, Yizhi Li, Yinghao Ma, Yiming Liang, Xiaowei Chi, Ruibo Liu, Zili Wang, Pengfei Li, Jingcheng Wu, Chenghua Lin, Qifeng Liu, Tao Jiang, Wenhao Huang, Wenhu Chen, Emmanouil Benetos, Jie Fu, Gus Xia, Roger Dannenberg, Wei Xue, Shiyin Kang, and Yike Guo. ChatMusician: Understanding and generating music intrinsically with LLM. In *Findings of the Association for Computational Linguistics: ACL 2024*, pp. 6252–6271, 2024. URL https://aclanthology.org/2024.findings-acl.373/.

Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *International conference on machine learning*, pp. 3987–3995. PMLR, 2017.

Yifan Zhang. StackMathQA: A Curated Collection of 2 Million Mathematical Questions and Answers Sourced from Stack Exchange. Technical report, Math-AI, 2024.

Appendix Contents

A Definitions **15**

A.1 Optimizers . . . 15

A.2 Learning-rate schedules . . . 15

A.3 Token–parameter ratio . . . 16

B Other Metrics for Analyzing Post-Training **16**

B.1 Gaussian perturbation . . . 16

B.2 Gaussian perturbation and the margin-to-spread ratio . . . 17

C Experimental Details **19**

C.1 Pretraining . . . 19

C.2 Fine-tuning . . . 21

C.3 Fine-tuning regularization . . . 21

C.4 Gaussian perturbations . . . 21

C.5 Evaluation . . . 21

D Additional Results for Section 3 **22**

D.1 Lowest-pretraining-loss setting . . . 22

D.2 Best-frontier setting . . . 23

D.3 Choice of fine-tuning optimizer . . . 25

D.4 Fine-tuning regularization . . . 25

E Additional Results for Section 4 and Appendix B **29**

E.1 Margin-to-spread ratio across data budgets . . . 29

E.2 Perturbation robustness . . . 31

E.3 Margin-to-spread ratio and robustness . . . 32

F Proof of Theorem 4.2 **34**

F.1 The two one-step iterates . . . 34

F.2 Concentration estimates for the data . . . 35

F.3 The GD step . . . 38

F.4 The Muon step . . . 41

F.5 Combining the GD and Muon estimates . . . 47

F.6 Monotonicity in the data budget . . . 48

A Definitions

A.1 Optimizers

Optimizers update a model's parameters during training so as to decrease a loss. Given the gradient of the loss with respect to the parameters, an optimizer determines the direction and the magnitude of each update, and in doing so shapes training stability, convergence speed and the solution that is finally reached. This section recalls the two optimizers compared in the paper.

A.1.1 AdamW

Adam (Kingma & Ba, 2015) combines momentum with a per-coordinate step size. With $g_t = \nabla_\theta \mathcal{L}(\theta_t)$ the gradient at step t, it maintains exponential moving averages of the first and second moments of the gradient,

$$m_t = \beta_1 m_{t-1} + (1-\beta_1) g_t, \qquad v_t = \beta_2 v_{t-1} + (1-\beta_2) g_t^2,$$

with the square taken coordinatewise, corrects their bias towards the zero initialization,

$$\hat{m}_t = \frac{m_t}{1-\beta_1^t}, \qquad \hat{v}_t = \frac{v_t}{1-\beta_2^t},$$

and updates

$$\theta_{t+1} = \theta_t - \alpha \frac{\hat{m}_t}{\sqrt{\hat{v}_t}+\epsilon},$$

where α is the learning rate, $\beta_1, \beta_2 \in [0,1)$ are the decay rates of the two averages, $\epsilon > 0$ guards the division, and the division and the square root are coordinatewise. In effect each coordinate is normalized by the scale of its own recent gradients.

Regularizing Adam with an ℓ_2 penalty $\frac{\lambda}{2}\|\theta\|^2$ adds $\lambda\theta_t$ to g_t, so the penalty is itself divided by $\sqrt{\hat{v}_t}+\epsilon$ and coordinates with large past gradients are regularized less. AdamW (Loshchilov & Hutter, 2019) instead decouples the weight decay from the adaptive step,

$$\theta_{t+1} = \theta_t - \alpha \frac{\hat{m}_t}{\sqrt{\hat{v}_t}+\epsilon} - \alpha\lambda\theta_t,$$

with g_t now the gradient of the unregularized loss and λ the weight decay coefficient, so that every coordinate decays at the same rate. AdamW is the standard optimizer for language-model pretraining and the baseline throughout the paper.

A.1.2 Muon

Muon (Jordan et al., 2024) treats each weight matrix as a whole rather than coordinate by coordinate. For a matrix-valued parameter θ with gradient G_t, it accumulates momentum $M_t = \mu M_{t-1} + G_t$ and updates along the orthogonalized momentum,

$$\theta_{t+1} = \theta_t - \eta \,\mathrm{Orth}(M_t), \tag{4}$$

where η is the learning rate and, for a matrix X with singular value decomposition $X = U\Sigma V^\top$, $\mathrm{Orth}(X) = UV^\top$ is its polar factor: the matrix with the singular vectors of X and all singular values set to one. The update therefore has the same size along every singular direction, and it can be read as the steepest descent direction when the size of a step is measured in the spectral norm (Jordan et al., 2024). An exact singular value decomposition at every step is expensive, so Orth is approximated by a few steps of a Newton–Schulz iteration,

$$X_0 = \frac{M_t}{\|M_t\|_F}, \qquad X_{k+1} = aX_k + b\,(X_k X_k^\top) X_k + c\,(X_k X_k^\top)^2 X_k,$$

whose odd matrix polynomial pushes every singular value of X_0 towards one; the reference implementation uses $(a,b,c) = (3.4445, -4.7750, 2.0315)$ and five steps. Muon is meant for the two-dimensional weight matrices of the hidden layers; embeddings, the output layer and one-dimensional parameters such as normalization gains are updated by AdamW. Large-scale implementations further add decoupled weight decay and rescale the update to the root-mean-square size of AdamW's (Liu et al., 2025). In our experiments the hidden matrices are updated by Muon and all other parameters by AdamW, each with its own learning rate and weight decay (Section 3.1).

A.2 Learning-rate schedules

A.2.1 Warmup–stable–decay

All pretraining runs use a warmup–stable–decay (WSD) schedule: the learning rate increases linearly from zero to its peak value over the warmup steps, stays at the peak for most of training, and is annealed to zero over the final decay phase.

A.2.2 Cosine

All fine-tuning runs use a cosine schedule: after a linear warmup of 20 steps, the learning rate follows $\eta_t = \frac{1}{2}\eta_{\max}\big(1 + \cos(\pi t/T)\big)$ over the remaining T steps.

A.3 Token–parameter ratio

We measure the pretraining data budget of a model with N parameters trained on D tokens by its token–parameter ratio, $\mathrm{TPR} = D/N$. The compute-optimal budget of Hoffmann et al. (2022) corresponds to $\mathrm{TPR} = 20$; our budgets range from 5 to 2,560 TPR.

B Other Metrics for Analyzing Post-Training

Besides the margin-to-spread ratio, the robustness of a base model to post-training can be probed by perturbing its weights directly. This section reports Gaussian perturbation results in language models and in the linear memory (Section B.1), and then relates them to the MSR (Section B.2).

B.1 Gaussian perturbation

We first check whether Gaussian perturbations of the weights reproduce the budget trend of the learning–forgetting tradeoff, in language models and in the linear memory.

B.1.1 Language models

We test whether the perturbed loss of the pretrained language models shifts from Muon toward AdamW as the token budget grows, as the learning–forgetting tradeoff does.

Setup. Springer et al. (2025) show that adding isotropic Gaussian noise to pretrained weights reproduces key empirical trends in forgetting after fine-tuning. Motivated by this connection, we study Gaussian weight perturbations as a simpler setting in which to investigate why the learning–forgetting gap between Muon and AdamW changes with pretraining budget. At the same checkpoints as in Section 3, we add isotropic Gaussian noise to each weight matrix, scaled to a fraction δ of the matrix's norm, and measure the loss increase $\Delta_{\mathrm{noise}}(\theta_{\mathrm{PT}}; \delta) := \mathcal{L}_{\mathrm{PT}}(\widetilde{\theta}_\delta) - \mathcal{L}_{\mathrm{PT}}(\theta_{\mathrm{PT}})$, where $\widetilde{\theta}_\delta$ denotes the perturbed weights; a smaller increase indicates greater perturbation robustness. See Section E.2 for full experimental details.

Result. Figure 6 plots the gap in loss degradation, $\Delta_{\mathrm{noise}}(\mathrm{Muon}; \delta) - \Delta_{\mathrm{noise}}(\mathrm{AdamW}; \delta)$. This gap is positive across the token budgets tested and grows with the pretraining budget: AdamW incurs a smaller loss increase throughout the sweep, and its advantage grows with budget.

To connect this gap to the learning–forgetting tradeoff, we also account for pretraining loss: loss after perturbation equals pretraining loss plus Δ_{noise}. Muon starts with lower pretraining loss, so it can retain lower loss after perturbation at small budgets despite its larger loss increase. As its pretraining-loss advantage shrinks and the degradation gap grows, AdamW achieves lower loss after perturbation beyond a crossover budget. For sufficiently strong perturbations, this crossover occurs at smaller token budgets as model size increases, paralleling the crossover observed after fine-tuning (Table 2 in Section E.2).

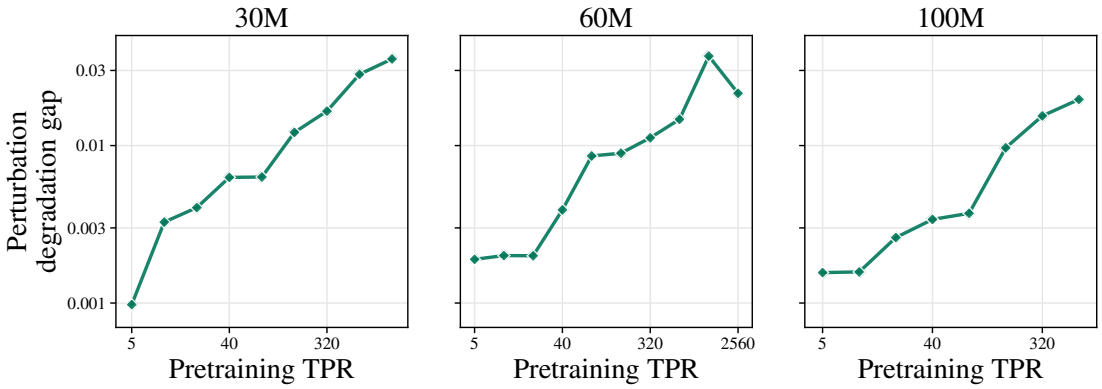


Figure 6: The perturbation degradation gap grows with pretraining data budget. Under a perturbation of magnitude $\delta = 0.02$, Muon degrades increasingly more than AdamW as the pretraining data budget grows. Points show the mean over three noise draws.

B.1.2 Linear associative memory

We repeat the perturbation experiment in the linear memory, where the effect of a perturbation can be measured fact by fact.

Setup. We train the memory of Section 4.1 with GD and Muon to minimize cross-entropy, varying the data budget $\gamma = N/d$ and the dimension d, as in Section 4.2. The trained weights are perturbed to $W' = W + \sigma\|W\|_F\xi$ with $\|\xi\|_F = 1$, over 16 perturbation directions shared by the two optimizers. Since the memory retrieves facts, we measure degradation by the fraction of initially recovered facts that are lost. At a single perturbation magnitude, fact loss can be close to zero or one, making memories with different noise tolerances difficult to distinguish. We therefore compare the perturbation strengths needed to reach a common intermediate level of degradation: σ_{50} is the relative perturbation strength at which half of the initially recovered facts are lost. A larger σ_{50} means that more noise is needed to cause the same fraction of facts to be lost. Data budgets past the memory's capacity, where the two optimizers recover different sets of facts, are excluded. Full details appear in Sections E.1 and E.3.

Result. The robustness ordering of the two optimizers turns over with the data budget: Muon tolerates stronger perturbations at small γ, and GD does at large γ (Figure 4b and the bottom row of Figure 20). The σ_{50} ordering flips at $\gamma \approx 0.5$, 1.3 and 11 at $d = 128$, 256 and 512, and at $d = 64$ GD is already the more robust at the smallest data budget we ran.

B.2 Gaussian perturbation and the margin-to-spread ratio

We next relate robustness to Gaussian perturbations to the MSR, first in theory and then in experiments.

B.2.1 Theoretical analysis

We use the scores $s_{ij}(W) = v_j^\top W u_i$ of Section 4.1 and the margins m_{ij}, h_i, κ_i and ρ_i of Definition 4.1, for arbitrary keys, values and W. The weights are perturbed to $W + \sigma E$, where the entries of E satisfy

$$\mathbb{E}[E_{kl}] = 0, \qquad \mathbb{E}[E_{kl}E_{k'l'}] = \mathbf{1}\{(k,l) = (k',l')\}. \tag{5}$$

Scores are linear in W, so exactly

$$s_{ii}(W + \sigma E) - s_{ij}(W + \sigma E) = m_{ij} + \sigma\,(v_i - v_j)^\top E u_i.$$

We measure robustness by the fraction of competitors still ranked below the correct value,

$$\mathrm{AUC}_i(\sigma) = \operatorname*{mean}_{j \neq i} \Pr_E\left[s_{ii}(W + \sigma E) > s_{ij}(W + \sigma E)\right].$$

Let J be uniform on $\{1, \dots, N\} \setminus \{i\}$, independent of E. Then

$$\mathrm{AUC}_i(\sigma) = \Pr\left[s_{ii}(W + \sigma E) - s_{iJ}(W + \sigma E) > 0\right] = \Pr\left[m_{iJ} + \sigma(v_i - v_J)^\top E u_i > 0\right],$$

so it suffices to bound the probability that one random variable is positive, from its mean and variance.

Lemma B.1 (Cantelli). *If $\mathbb{E}X > 0$ and $\operatorname{Var} X < \infty$, then*

$$\Pr[X \le 0] \;\le\; \frac{\operatorname{Var} X}{\operatorname{Var} X + (\mathbb{E}X)^2}.$$

Proof. For $t \ge 0$, Markov's inequality gives $\Pr[X \le 0] \le \mathbb{E}(\mathbb{E}X - X + t)^2/(\mathbb{E}X + t)^2 = (\operatorname{Var} X + t^2)/(\mathbb{E}X + t)^2$; take $t = \operatorname{Var} X/\mathbb{E}X$. □

Lemma B.2 (Noise adds variance, not bias). *Let E satisfy Equation (5). Then*

$$\mathbb{E}\big[m_{iJ} + \sigma(v_i - v_J)^\top E u_i\big] = h_i,$$
$$\operatorname{Var}\big[m_{iJ} + \sigma(v_i - v_J)^\top E u_i\big] = \kappa_i^2 + \sigma^2\|u_i\|^2 \operatorname*{mean}_{j\neq i}\|v_i - v_j\|^2.$$

Proof. For fixed j, $(v_i - v_j)^\top E u_i = \sum_{k,l}(v_i - v_j)_k E_{kl}(u_i)_l$ has, by Equation (5), mean 0 and second moment $\|v_i - v_j\|^2\|u_i\|^2$. Average over E first, then over J. □

Theorem B.3 (Perturbed AUC and the margin-to-spread ratio). *Let E satisfy Equation (5). For every W, all data, every $\sigma \ge 0$ and every i with $h_i > 0$,*

$$\mathrm{AUC}_i(\sigma) \;\ge\; \frac{h_i^2}{h_i^2 + \kappa_i^2 + \sigma^2\|u_i\|^2 \operatorname{mean}_{j\neq i}\|v_i - v_j\|^2}.$$

Proof. $1 - \mathrm{AUC}_i(\sigma) = \Pr\big[m_{iJ} + \sigma(v_i - v_J)^\top E u_i \le 0\big]$. Apply Lemma B.1 with the mean and variance of Lemma B.2. □

No assumption is made on W, on the shape of the margin distribution, or on d and N.

Corollary B.4 (Gaussian data). *If moreover $u_1, \dots, u_N, v_1, \dots, v_N$ are i.i.d. $\mathcal{N}(0, I_d)$, then*

$$\mathrm{AUC}_i(\sigma) \;\ge\; \frac{h_i^2}{h_i^2 + \kappa_i^2 + 2\sigma^2 d^2\big(1 + O_p(d^{-1/2})\big)} \;=\; \frac{\rho_i^2}{\rho_i^2 + 1 + 2(\sigma d/\kappa_i)^2\big(1 + O_p(d^{-1/2})\big)}.$$

Proof. $\|u_i\|^2 = d\big(1 + O_p(d^{-1/2})\big)$ and $\operatorname{mean}_{j\neq i}\|v_i - v_j\|^2 = 2d\big(1 + O_p(d^{-1/2})\big)$. □

At a fixed noise level in units of κ_i, a larger ρ_i therefore guarantees a larger AUC. For the relative perturbation $W + \delta\|W\|_F E/\|E\|_F$ used in the experiments, with i.i.d. standard Gaussian entries of E, the bounds hold with σ replaced by $\delta\|W\|_F/d$.

Corollary B.5 (Noise on the logits). *Suppose instead that each logit is perturbed directly, to $s_{ij}(W) + \sigma\xi_{ij}$, where the ξ_{ij} have mean 0 and variance 1 and are uncorrelated (for example, i.i.d. $\mathcal{N}(0,1)$). Then for every W, every $\sigma \ge 0$ and every i with $h_i > 0$,*

$$\operatorname*{mean}_{j\neq i} \Pr_{\xi}\big[s_{ii}(W) + \sigma\xi_{ii} > s_{ij}(W) + \sigma\xi_{ij}\big] \;\ge\; \frac{h_i^2}{h_i^2 + \kappa_i^2 + 2\sigma^2} \;=\; \frac{\rho_i^2}{\rho_i^2 + 1 + 2\sigma^2/\kappa_i^2}.$$

Proof. With J uniform on $\{1, \dots, N\} \setminus \{i\}$ and independent of ξ, the variable $m_{iJ} + \sigma(\xi_{ii} - \xi_{iJ})$ has mean h_i and variance $\kappa_i^2 + 2\sigma^2$, since $\xi_{ii} - \xi_{ij}$ has mean 0 and variance 2 for every $j \neq i$. Apply Lemma B.1. □

A fact is mislabeled only if some incorrect answer overtakes the correct one, so by a union bound the probability of mislabeling fact i is at most $N - 1$ times one minus the left-hand side. This bound is informative only when ρ_i^2 is much larger than N.

B.2.2 Empirical validation

We test whether the MSR predicts how much perturbation a checkpoint can tolerate, first in language models and then in the linear memory.

Language models. We measure MSR as in Section 4.3 and degradation by the held-out loss increase of Section B.1.1. The robustness threshold δ^* is the perturbation magnitude at which this increase reaches ΔL nats. Across checkpoints, ρ and δ^* are strongly correlated for both $\Delta L = 0.05$ and 0.1 (Figure 7b): the Spearman correlation over the 72 checkpoints is $+0.83$ at both thresholds, and within each optimizer it is $+0.87$ to $+0.89$ for AdamW and $+0.92$ to $+0.93$ for Muon. Full experimental details appear in Section E.3.

Linear associative memory. With σ_{50} as in Section B.1.2, σ_{50} is strongly correlated with MSR for both optimizers (Figure 7a). Within each optimizer, σ_{50} and ρ move together at every model dimension (Figure 23).

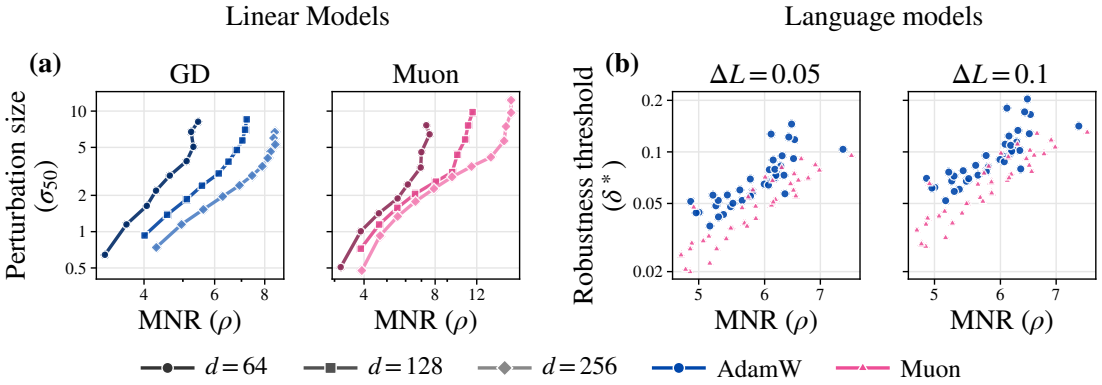


Figure 7: The margin-to-spread ratio tracks robustness to weight perturbations. (a) For the linear models trained under GD and Muon, we plot the perturbation σ_{50} required to lose half of the recovered facts against its MSR, $\bar{h}/\bar{\kappa}$ (Definition 4.1). A higher MSR requires more noise to cause the same loss. **(b)** We measure the MSR of 72 pretrained language models (30M–600M) as the median of h/κ over correctly predicted tokens, and plot the perturbation δ^* at which held-out loss rises by ΔL. The same relationship holds.

C Experimental Details

C.1 Pretraining

Models and data. All language models are OLMo-style decoder-only transformers (Groeneveld et al., 2024; Team OLMo et al., 2025) with 30M, 60M, 100M, 300M and 600M parameters, trained on DCLM (Li et al., 2024) with the OLMo 2 tokenizer. Budgets span 5–1,280 TPR at 30M, 5–2,560 TPR at 60M, 5–640 TPR at 100M, 5–160 TPR at 300M and 5–40 TPR at 600M.

Optimization. Every model is pretrained with sequence length 4096, a global batch of 2^{20} tokens, a warmup–stable–decay (WSD) schedule, weight decay 0.1 and gradient clipping at 1.0. AdamW uses $(\beta_1, \beta_2) = (0.9, 0.98)$. Muon updates the hidden weight matrices; the remaining parameters (embeddings, output layer and norms) are updated by AdamW.

C.1.1 Lowest-pretraining-loss setting

For each model size, budget and optimizer, we sweep the peak learning rate and call the run with the lowest held-out DCLM loss the *loss-optimal* base model. For Muon we sweep the Muon learning rate and fix the learning rate of the AdamW-updated parameters to the tuned AdamW value of the same cell (Section C.1.2). Table 1 lists the resulting learning rates. In the *lowest-pretraining-loss setting*, used throughout the main text, each optimizer is represented by its loss-optimal base model alone: we fine-tune it over the full sweep of Section C.2 and compute the learning–forgetting frontier from its fine-tuning runs.

C.1.2 Learning rate of Muon's AdamW-updated parameters

Muon has a second learning rate for the parameters it leaves to AdamW. Rather than tuning it in every cell, we set it to the tuned AdamW learning rate of the same cell. To check this choice, we pretrain 60M Muon models at 20–160

Table 1: Tuned peak pretraining learning rates per model size and budget. For Muon, the Muon learning rate; the AdamW-parameter learning rate is fixed to the tuned AdamW value.

Size	Optimizer	TPR									
		5	10	20	40	80	160	320	640	1,280	2,560
30M	AdamW	0.08	0.04	0.04	0.014	0.02	0.014	0.02	0.014	0.01	–
	Muon	0.01	0.01	0.01	0.014	0.02	0.01	0.014	0.01	0.014	–
60M	AdamW	0.056	0.04	0.02	0.02	0.01	0.014	0.01	0.01	0.007	0.01
	Muon	0.014	0.01	0.014	0.01	0.014	0.01	0.014	0.007	0.005	0.005
100M	AdamW	0.04	0.02	0.01	0.01	0.01	0.007	0.007	0.007	–	–
	Muon	0.014	0.014	0.014	0.014	0.01	0.014	0.01	0.005	–	–
300M	AdamW	0.001	0.0025	0.0035	0.0025	0.0025	0.0025	–	–	–	–
	Muon	0.02	0.014	0.02	0.014	0.014	0.02	–	–	–	–
600M	AdamW	0.0025	0.0025	0.0025	0.0025	–	–	–	–	–	–
	Muon	0.014	0.014	0.014	0.014	–	–	–	–	–	–

TPR over a grid of Muon learning rates and three to four AdamW learning rates between $0.5\times$ and $2.8\times$ the tuned value. Figure 8 shows the best loss over Muon learning rates for each AdamW learning rate. It varies by at most 0.021 nats, and the tuned value is within 0.016 nats of the best, well below the 0.03–0.20 nats by which Muon beats AdamW at these budgets.

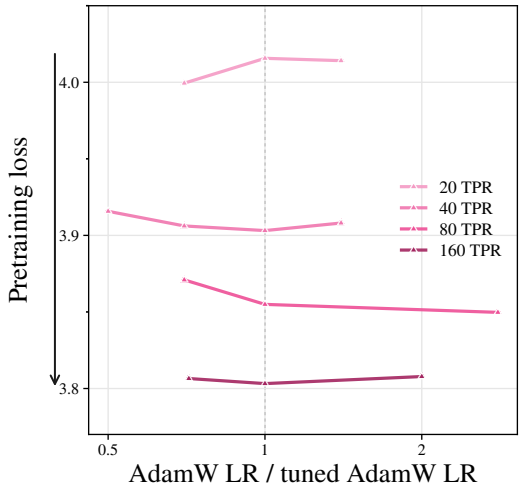


Figure 8: Muon is insensitive to the learning rate of its AdamW-updated parameters. Best pretraining loss over Muon learning rates for each AdamW learning rate, relative to the tuned AdamW value (dashed), for 60M models.

C.1.3 Best-frontier setting

The loss-optimal base model is not always the best starting point for fine-tuning: models pretrained with a larger learning rate can offer a better learning–forgetting tradeoff (Watts et al., 2026). We therefore also compare the two optimizers in the *best-frontier setting*, which sweeps over both pretraining and fine-tuning configurations. For each model size, budget and pretraining optimizer, we fine-tune the loss-optimal base model and the base models pretrained at the next five larger learning rates of the sweep grid $\{0.005, 0.007, 0.01, 0.014, 0.02, 0.028, 0.04, 0.056, 0.08, 0.12, 0.16\}$, pool all of their fine-tuning runs, and compute a single learning–forgetting frontier from the pooled runs. This frontier is at least as good as the one of the lowest-pretraining-loss setting, since it contains the loss-optimal base model's runs. Section D.2 repeats the main results in this setting.

C.2 Fine-tuning

Each base model is fine-tuned for 20M tokens on each of Alpaca (Taori et al., 2023), GSM8K (Cobbe et al., 2021), MusicPile (Yuan et al., 2024) and StackMathQA (Zhang, 2024), with sequence length 1024, a global batch of 65,536 tokens (about 305 steps), a cosine schedule with 20 warmup steps, weight decay 0.1, $(\beta_1, \beta_2) = (0.9, 0.95)$ and gradient clipping at 1.0. AdamW-pretrained models are fine-tuned with AdamW, and Muon-pretrained models with both AdamW and Muon (Section D.3 motivates this choice). We sweep AdamW learning rates from 10^{-4} to 10^{-2}, and Muon learning rates from 6×10^{-4} to 2×10^{-2} with the learning rate of the AdamW-updated parameters set to $0.25\times$ the Muon learning rate. The sweep ranges from runs that barely move the base model to runs that reach the lowest fine-tuning loss at the cost of more than a nat of forgetting.

C.3 Fine-tuning regularization

A natural question is whether the reversal survives when fine-tuning itself is regularized to limit forgetting. We test two standard methods for mitigating catastrophic forgetting, data replay (Shin et al., 2017) and elastic weight consolidation (EWC; Kirkpatrick et al., 2017), on the loss-optimal base models of every size. Both keep the fine-tuning recipe and learning-rate sweep of Section C.2, and each frontier pools all strengths of the regularizer. Section D.4 reports the results.

Replay. We replace a fraction $r \in \{0.1, 0.2, 0.3\}$ of the fine-tuning tokens with DCLM data drawn from a shard that is disjoint from both the pretraining data and the held-out evaluation shard.

EWC. EWC penalizes movement away from the pretrained weights in the directions that matter for pretraining. It minimizes

$$\mathcal{L}_{\mathrm{FT}}(\theta) + \frac{\lambda}{2} \sum_k F_k \big(\theta_k - \theta_{\mathrm{PT},k}\big)^2, \tag{6}$$

where F is the diagonal empirical Fisher information of the pretraining loss at θ_{PT}, estimated from squared gradients over 100 batches of the disjoint DCLM shard used for replay. Following Watts et al. (2026), F is not normalized, so we sweep $\lambda \in \{10^3, 4 \cdot 10^3, 10^4, 4 \cdot 10^4\}$.

C.4 Gaussian perturbations

We perturb each base model by adding independent isotropic Gaussian noise to every parameter tensor W, scaled so that $\|\Delta W\|_F = \delta \|W\|_F$, and measure the perturbation degradation $\Delta_{\mathrm{noise}}(\theta_{\mathrm{PT}}; \delta)$ of Section B.1.1. We average over three noise draws per checkpoint and magnitude, use the loss-optimal checkpoint of each cell, and consider magnitudes from $\delta = 0.002$ to 0.56.

C.5 Evaluation

Losses. Forgetting is the loss on a held-out DCLM shard of about 4.19M tokens that no model is trained on, and learning is the loss on the validation split of the fine-tuning dataset. Both are mean next-token cross-entropy in nats.

Frontiers. Each fine-tuned checkpoint θ_{FT} of a base model gives a point $\big(\mathcal{L}_{\mathrm{PT}}(\theta_{\mathrm{FT}}), \mathcal{L}_{\mathrm{FT}}(\theta_{\mathrm{FT}})\big)$. We compute the learning–forgetting frontier as the lower-left arc of the convex hull of these points, treated as a polyline.

Fine-tuning gap and crossover. We estimate the fine-tuning gap G_{FT} of Section 2 from 100 points sampled uniformly by arc length along the AdamW frontier, each assigned its Euclidean distance to the nearest point of the Muon frontier, negative when that point lies closer to the origin. The crossover TPR is the smallest budget at which $G_{\mathrm{FT}} > 0$. In both settings, the gap and crossover compare the two optimizers' frontiers as defined above: one base model per optimizer in the lowest-pretraining-loss setting, pooled base models in the best-frontier setting.

Degradation. For Figure 3c, the degradation of a base model is the increase in held-out DCLM loss caused by fine-tuning, $\mathcal{L}_{\mathrm{PT}}(\theta_{\mathrm{FT}}) - \mathcal{L}_{\mathrm{PT}}(\theta_{\mathrm{PT}})$. We read $\mathcal{L}_{\mathrm{PT}}(\theta_{\mathrm{FT}})$ off the base model's frontier at a target fine-tuning loss of $1.1\times$ the lowest fine-tuning loss reached by either optimizer at that model size, so that both optimizers are compared at the same amount of learning.

D Additional Results for Section 3

This section extends the fine-tuning results of Section 3. Unless noted otherwise, figures show 60M models, and frontiers are restricted to the region of Section C.5.

D.1 Lowest-pretraining-loss setting

Figure 9 shows that the loss-optimal Muon model reaches a lower pretraining loss than its AdamW counterpart at every budget and model size, so the reversal is not inherited from pretraining loss. Figures 10 and 11 extend the frontiers and gaps of the lowest-pretraining-loss setting (Figures 1 and 3) to all four datasets, at the budgets around each crossover. On every dataset Muon is better than or level with AdamW below the crossover and AdamW is better beyond it.

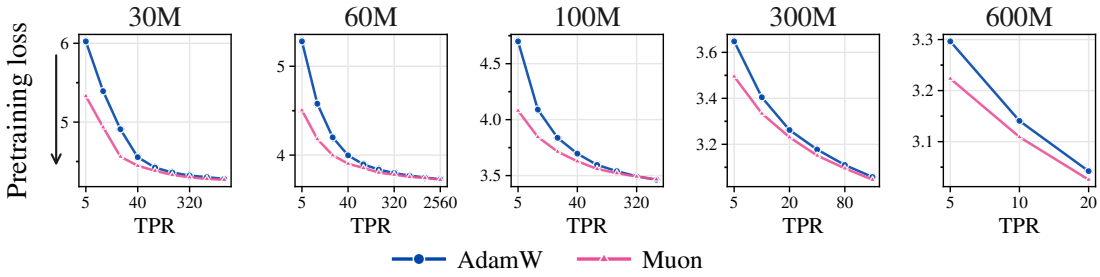


Figure 9: Muon reaches lower pretraining loss than AdamW. DCLM held-out loss of the loss-optimal base models across pretraining data budgets and model sizes.

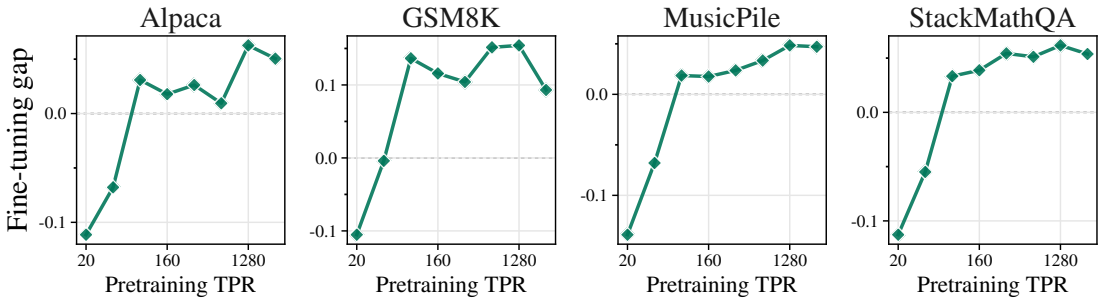


Figure 10: Fine-tuning gap on every fine-tuning dataset, loss-optimal base models (60M). Negative values mean the Muon frontier dominates.

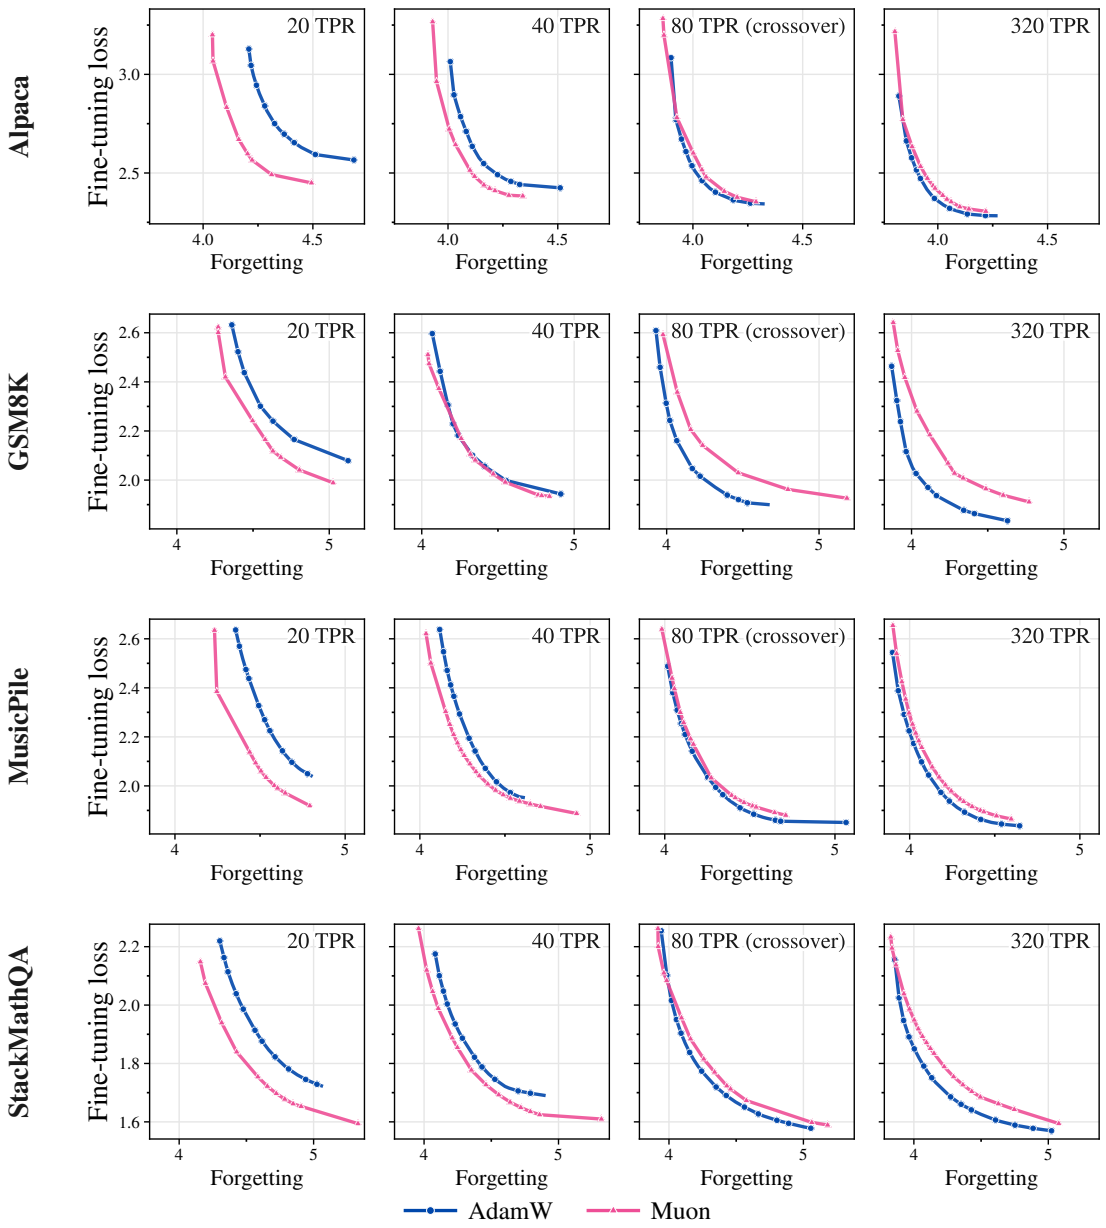


Figure 11: Learning–forgetting frontiers on every fine-tuning dataset, loss-optimal base models (60M). Each row shows the budgets around that dataset's crossover TPR.

D.2 Best-frontier setting

We repeat the analyses of Section 3 in the best-frontier setting (Section C.1.3), where each pretraining optimizer is swept across all pretraining and fine-tuning configurations. Figures 12 and 13 show the learning–forgetting frontiers and fine-tuning gaps of 60M models on all four fine-tuning datasets, at the budgets around each crossover. On every dataset, the Muon frontier is better below the crossover; at the crossover AdamW has caught up with or overtaken Muon, and beyond it the AdamW frontier dominates. The fine-tuning gap turns positive at the crossover and stays positive, broadly growing with the data budget. Figure 14 shows that, as in the lowest-pretraining-loss setting (Figure 2), the crossover TPR decreases with model size.

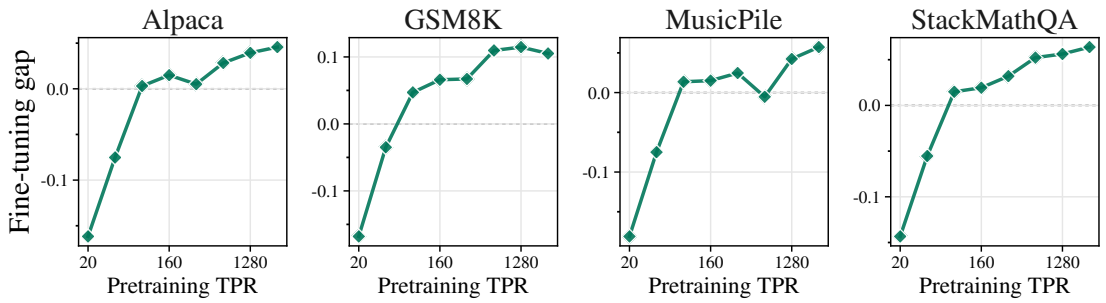


Figure 12: Fine-tuning gap on every fine-tuning dataset, best-frontier setting (60M). Negative values mean the Muon frontier dominates.

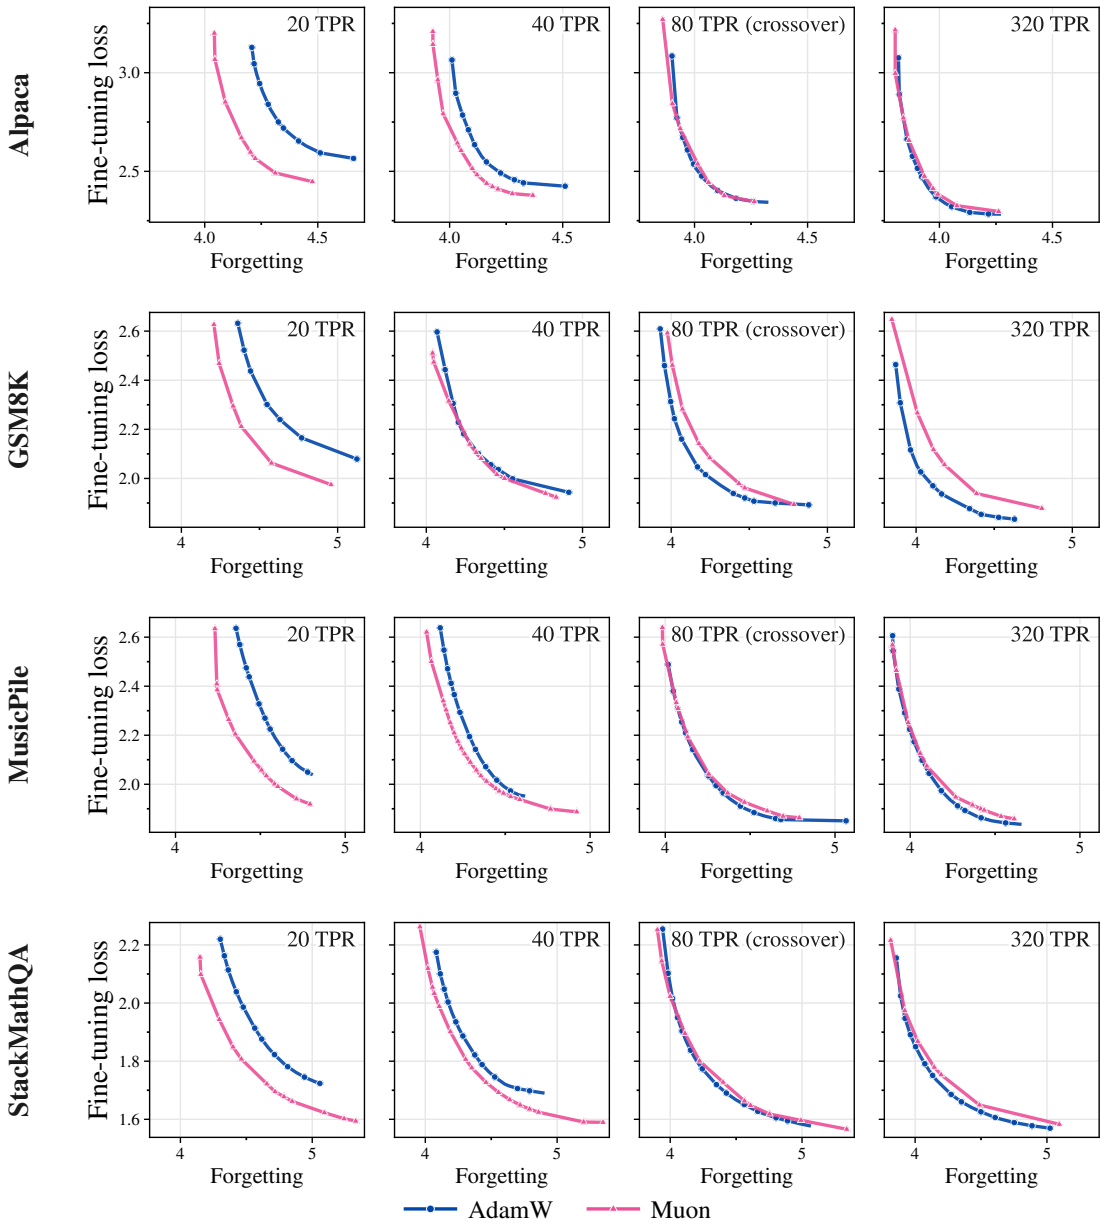


Figure 13: Learning–forgetting frontiers on every fine-tuning dataset, best-frontier setting (60M). Each row shows the budgets around that dataset's crossover TPR.

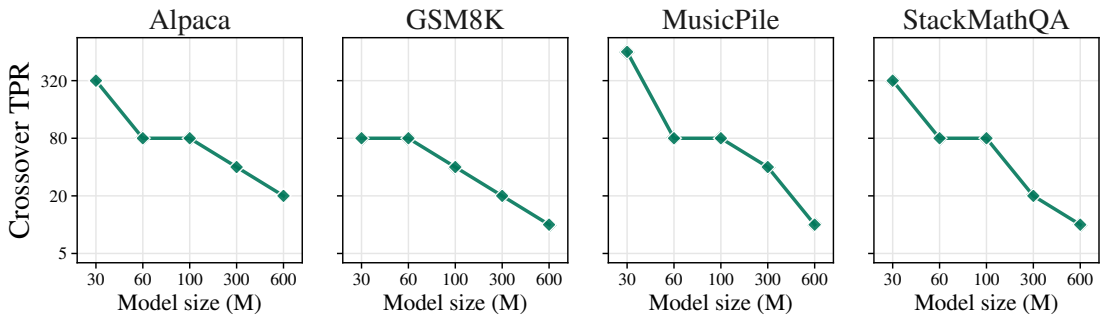


Figure 14: Crossover TPR decreases with model size, in the best-frontier setting. As model size increases, Muon's advantage reverses with substantially smaller pretraining data budgets.

D.3 Choice of fine-tuning optimizer

Our frontiers fine-tune AdamW-pretrained models with AdamW and Muon-pretrained models with both AdamW and Muon. Figure 15 motivates this choice, and shows that Muon's greater forgetting at large budgets is a property of the pretrained model rather than of the fine-tuning recipe. For each pretraining optimizer and budget, we take the best 60M base model among those fine-tuned with both optimizers, and plot its AdamW- and Muon-fine-tuned frontiers separately. We use GSM8K, the dataset on which AdamW-pretrained models were fine-tuned with both optimizers across the full range of budgets.

Two observations follow. First, the fine-tuning optimizer matters for AdamW-pretrained models but not for Muon-pretrained ones: on AdamW-pretrained models, AdamW fine-tuning dominates Muon fine-tuning at every budget, whereas the two frontiers of a Muon-pretrained model lie close together with no consistent winner. Second, the reversal does not depend on how the Muon-pretrained model is fine-tuned: at 1.2B tokens both of its frontiers lie below those of the AdamW-pretrained model, and from 19.2B tokens onward the AdamW-pretrained, AdamW-fine-tuned frontier dominates both.

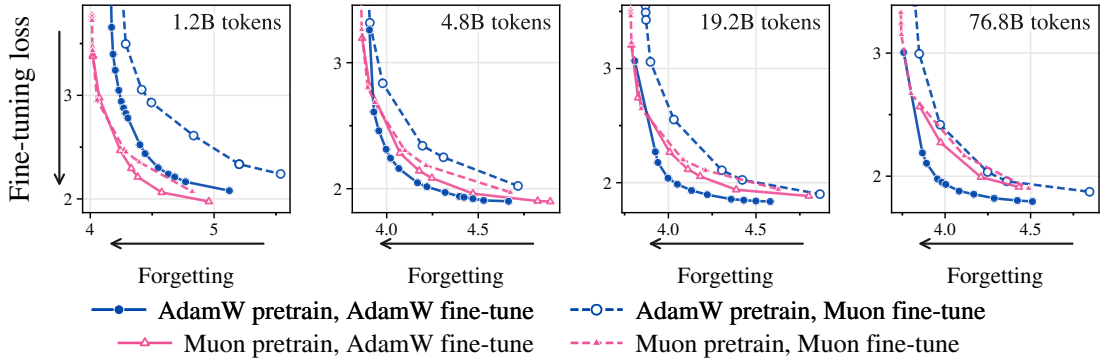


Figure 15: Muon's greater forgetting comes from the pretrained model, not the fine-tuning optimizer. Learning–forgetting frontiers of the best 60M base models on GSM8K, split by fine-tuning optimizer (solid: AdamW; dashed: Muon). AdamW fine-tuning is always better for AdamW-pretrained models, while the two fine-tuning optimizers give nearly the same frontier for Muon-pretrained models. At large budgets, the AdamW-pretrained model dominates under either choice.

D.4 Fine-tuning regularization

We repeat the comparison when fine-tuning is regularized with replay or EWC (setup in Section C.3), using the loss-optimal 60M base models. Figures 16 and 17 show the frontiers on all four datasets; each row shows the budgets around that dataset's crossover under the regularizer, a quarter and half of the crossover TPR, the crossover itself, and four times it. The crossover is the first budget at which the fine-tuning gap turns positive (Section C.5). Both regularizers reduce forgetting and can move the crossover to larger budgets (from 80 to 320 TPR on Alpaca and StackMathQA), but AdamW's frontier is better at large budgets on every dataset under both (Figure 18), with one exception: under EWC on StackMathQA the gap turns negative again at the largest budget, 2,560 TPR. With replay,

Muon's frontier is better than or level with AdamW's below the crossover on every dataset. With EWC the picture below the crossover is mixed: Muon is clearly better at small budgets on GSM8K, while on the other datasets the two frontiers are close.

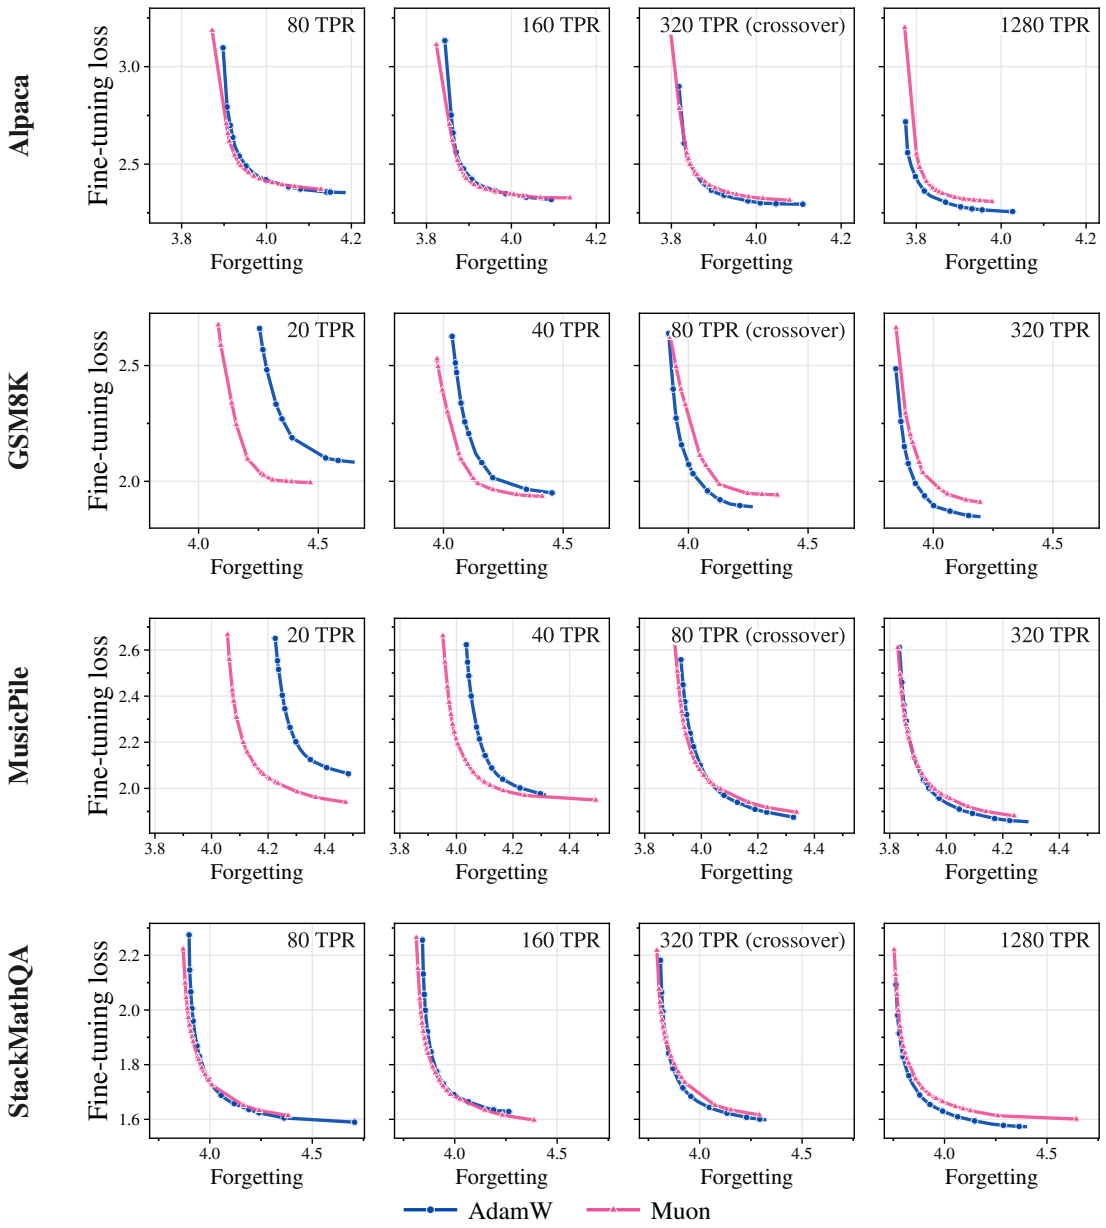


Figure 16: Learning–forgetting frontiers with replay of 10–30% DCLM data, loss-optimal 60M base models. Each row shows the budgets around that dataset's crossover under replay.

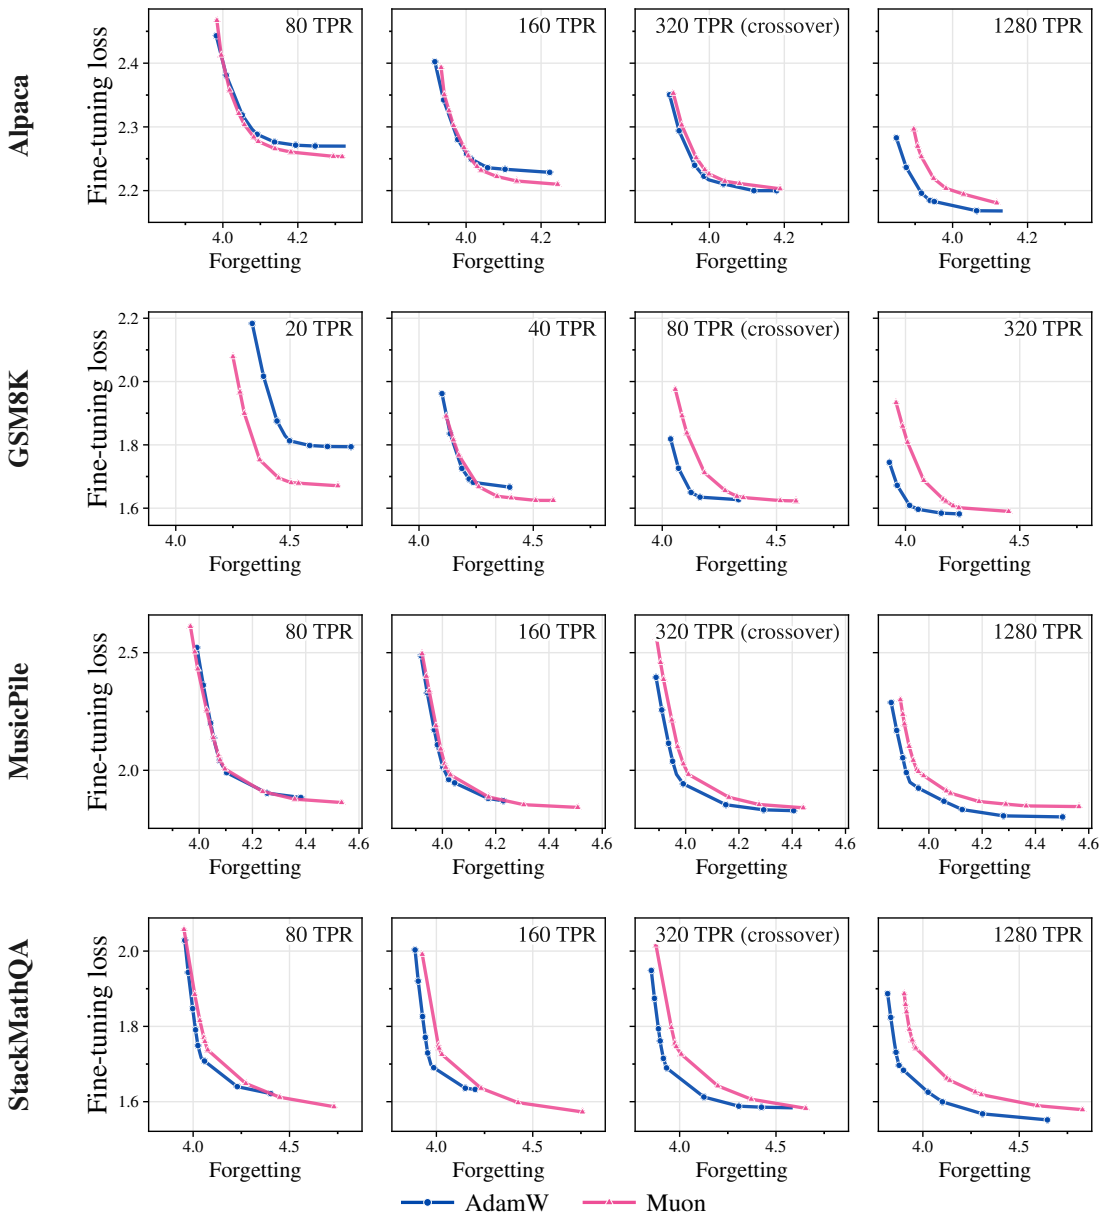


Figure 17: Learning–forgetting frontiers with EWC, $\lambda \in \{10^3, 4 \cdot 10^3, 10^4, 4 \cdot 10^4\}$, loss-optimal 60M base models. Each row shows the budgets around that dataset's crossover under EWC.

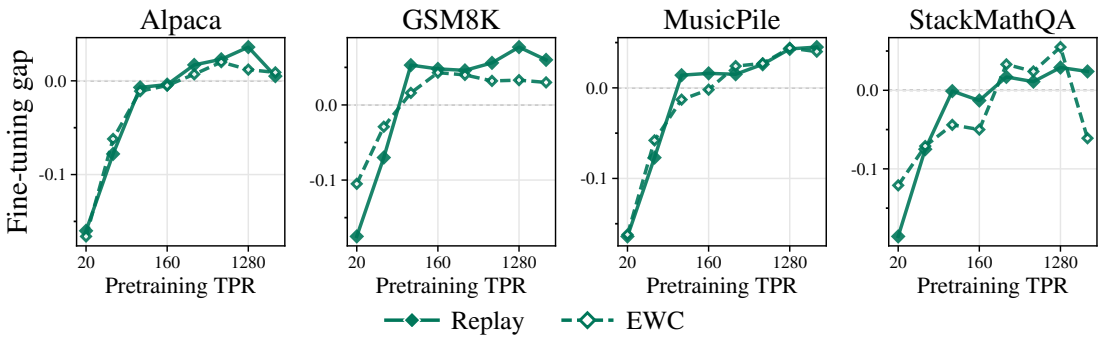


Figure 18: Fine-tuning gap under fine-tuning regularization, with replay (solid) and EWC (dashed), loss-optimal 60M base models. Negative values mean the Muon frontier dominates.

E Additional Results for Section 4 and Appendix B

This section collects details and additional results for the analysis of Section 4 and for the perturbation experiments of Section B: the margin-to-spread ratio across data budgets in the linear memory (Section E.1, for Section 4.2), perturbation robustness at every magnitude (Section E.2, for Section B.1), and the margin-to-spread ratio and robustness in the linear memory and in language models (Section E.3, for Sections B.2 and 4.3).

E.1 Margin-to-spread ratio across data budgets

E.1.1 One step

This is the regime Theorem 4.2 is about: the two solutions are the exact first step from $W_0 = 0$, so every quantity below can be compared against a closed form.

Setup Data. Keys and values are drawn as in Section 4.1 at $d = 512$, with $N = \gamma d$ and $\gamma \in \{\frac{1}{4}, \frac{1}{2}, \frac{3}{4}, 1, \frac{5}{4}, \frac{3}{2}, 2, \frac{5}{2}, 3, 4, 5, 6, 8, 10, 12, 16, 20, 24, 32, 48, 64\}$, five random seeds per data budget. Section 4.1 requires $N - 1 \geq d$, so the points with $\gamma \leq 1$ lie outside the range in which Theorem 4.2 is stated and are drawn only for context.

One step. The gradient of the loss of Section 4.1 at $W_0 = 0$ is computed exactly, with equal item weights. GD takes it as W^{GD}; Muon takes its polar factor $W^{\mathrm{Muon}} = \mathrm{Orth}(W^{\mathrm{GD}})$, computed exactly by singular value decomposition rather than by the Newton–Schulz iteration of Section A.1.2. As a baseline we also record the first AdamW step, which with zero initialization and the usual ϵ is the entrywise $\mathrm{sign}(W^{\mathrm{GD}})$; Theorem 4.2 says nothing about it.

Observables. Every h_i and κ_i is evaluated in closed form for all i as in Definition 4.1, giving $\bar{h}$, $\bar{\kappa}$, ρ and $\mathcal{R}$.

Plotting. Solid lines are means over the five seeds and bands are $\pm$ one standard deviation; dashed lines are the closed forms behind Theorem 4.2, with $\Phi(\gamma)$ replaced by its limit $\frac{8}{3\pi}$. In the $\bar{h}$ and $\bar{\kappa}$ panels each solution is normalized to $\|W\|_F = 1$, which is needed because the two steps have different scales while ρ does not see them; the $\bar{h}$ and ρ panels are divided by $\sqrt{d/\gamma}$, the unit in which Theorem 4.2 is stated.

Result Figure 19 (top) shows $\mathcal{R}$ rising monotonically from 0.53 at $\gamma = \frac{1}{4}$ to 1.16 at $\gamma = 64$ and crossing 1 near $\gamma = 5$, with the closed form tracking it once $\gamma \gtrsim 2$. The remaining panels separate the two factors. Normalized, GD's margin is constant in γ, which is the identity $\bar{h}^{\mathrm{GD}} = \|W^{\mathrm{GD}}\|_F^2/(N-1)$ read at $\|W\|_F = 1$, while Muon's climbs to about $\frac{8}{3\pi}$; both variances approach the same limit, GD's from above and Muon's from below. The limiting ratio $\frac{3\pi}{8}$ is therefore carried by the margin. Its approach to that limit comes from both factors: Muon's margin contributes the $+\frac{1}{10\gamma}$ correction, while the $-\frac{1}{\gamma}$ part of the $\frac{9}{10\gamma}$ term comes from GD's variance, the factor $\gamma + 2$ in $(\bar{\kappa}^{\mathrm{GD}})^2$ (Section 4.2). The bottom row runs the same comparison against the AdamW surrogate: its $\mathcal{R}$ has the same shape but reaches only 0.93 at $\gamma = 64$, so over the data budgets we ran the one-step reversal is a statement about GD, not about AdamW.

E.1.2 Many steps

The theory covers one step. This experiment runs the same memory to convergence and asks whether the ordering it predicts survives training.

Setup Data. Keys and values are drawn as in Section 4.1, with item i queried with probability $q_i \propto 1/i$ (Zipf). We use $d \in \{64, 128, 256, 512\}$, γ log-spaced in $[\frac{1}{4}, 128]$, and three seeds per (d, γ) shared by the two optimizers.

Training. From $W_0 = 0$ we take $T = 10^4$ full-batch steps of $W_{t+1} = W_t + 0.02\, G_t$ for GD and $W_{t+1} = W_t + 0.005\, \mathrm{Orth}(G_t)$ for Muon, where $G_t = \sum_i q_i \big(1 - \hat{p}_{W_t}(i \mid i)\big) u_i v_i^\top$ is the Hebbian part of the gradient of the loss of Section 4.1, at constant learning rate, without momentum or weight decay, and with the polar factor again taken by singular value decomposition.

Observables. h_i and κ_i are as in Definition 4.1, with equal item weights rather than the q_i used in training. The saturated value of an observable is the per-seed median over $t \geq 10^3$; $\mathcal{R}$ is formed per seed and only then averaged. Curves are means over seeds and bands are $\pm$ one standard deviation.

Perturbed robustness. We perturb the trained weights as in Section B.1.2, $W' = W_T + \sigma \|W_T\|_F \xi$ with $\|\xi\|_F = 1$, over 16 perturbation directions shared by the two optimizers, and read off σ_{50}, the σ at which half of the items correct

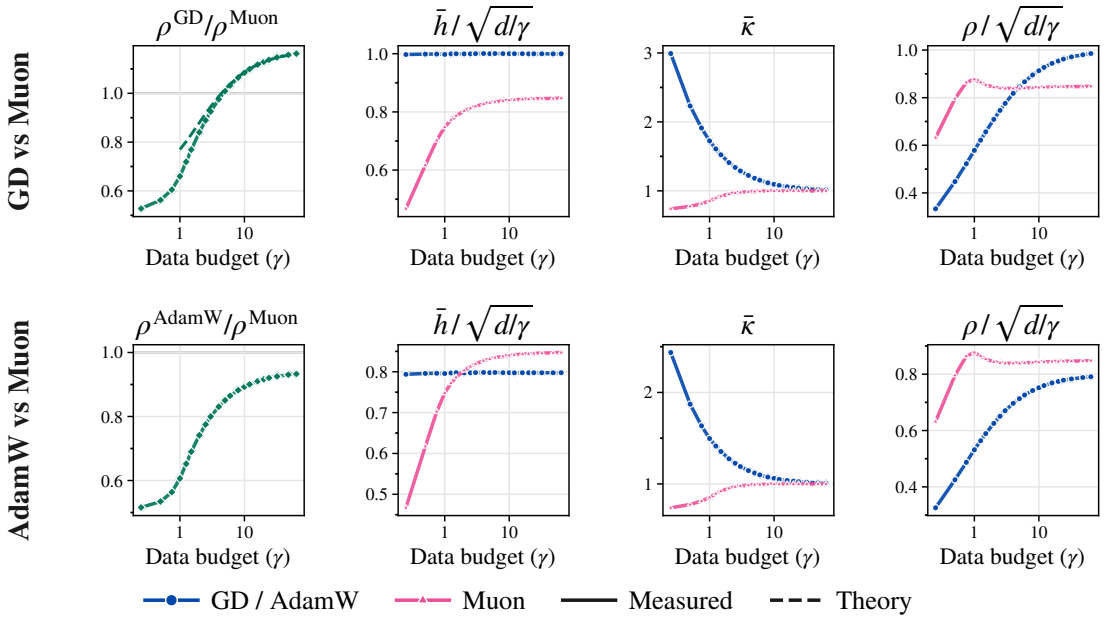


Figure 19: One step at $d = 512$, GD against Muon (top) and the AdamW surrogate against Muon (bottom); solid is experiment, dashed the closed form of Theorem 4.2.

at $\sigma = 0$ are lost. For $N < d$ the Muon solution is rescaled to its component in the span of the keys, since outside that span it is an arbitrary orthogonal completion that carries norm but no information. Data budgets at which either optimizer is already past the capacity of the memory, meaning accuracy below 0.99 at $\sigma = 0$, are dropped from the σ_{50} panels: there the two optimizers recover different sets of items and the two σ_{50} are not comparable.

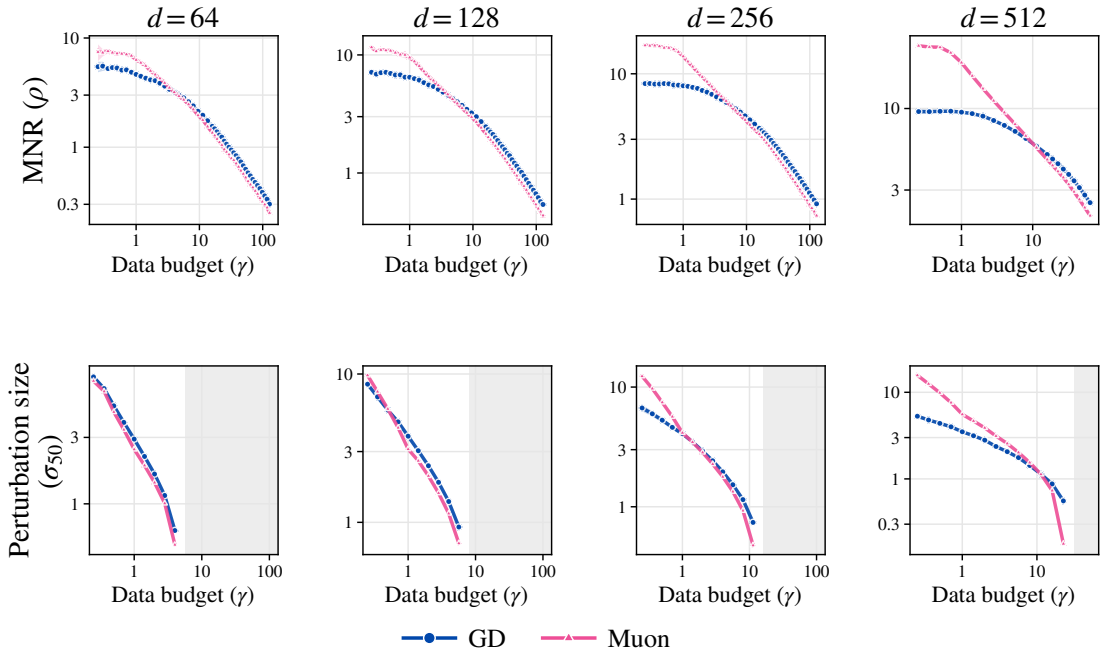


Figure 20: Top: saturated ρ against the data budget, the two curves behind Figure 21. **Bottom: perturbed robustness σ_{50} against the data budget**, below capacity only. One column per d.

Result Training does not undo the one-step picture. Figure 21 shows $\mathcal{R}$ increasing with γ at every d, over a range from about 0.4 to about 1.2, and the crossing $\mathcal{R} = 1$ moving to larger data budgets as d grows: $\gamma^* \approx 4.0$, 4.4, 5.4 and 9.9 at $d = 64$, 128, 256 and 512. Figure 20 shows the two ρ separately: both fall with the data budget, Muon's from

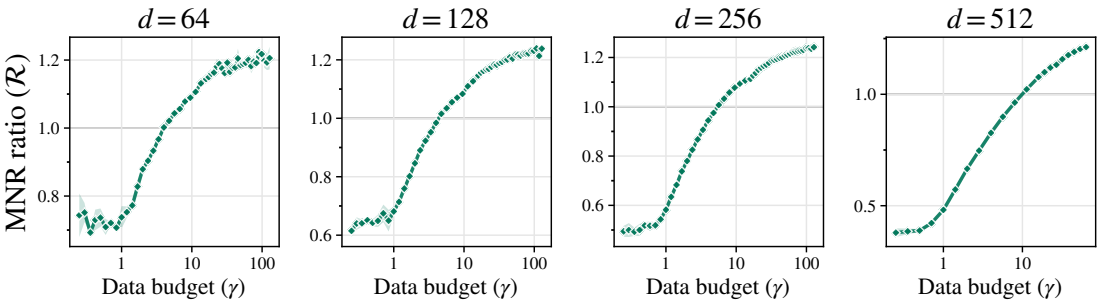


Figure 21: Saturated $\mathcal{R}$ against the data budget, one panel per d.

roughly twice GD's at $\gamma = \frac{1}{4}$ to below it past γ^*. The bottom row shows that the robustness ordering turns over as well, GD's σ_{50} gaining on Muon's at every d. The two turnovers need not coincide, and they do not: the σ_{50} ordering flips at $\gamma \approx 0.5$, 1.3 and 11 at $d = 128$, 256 and 512, and at $d = 64$ GD is already the more robust at the smallest data budget we ran.

E.2 Perturbation robustness

E.2.1 Setup

For every (model size, data budget) cell and each optimizer we take the checkpoint attaining the lowest held-out DCLM loss over the full pretraining learning-rate grid. A Gaussian perturbation is added independently to each parameter tensor of that checkpoint, scaled so that $\|\Delta W\|_F \approx \delta \|W\|_F$, and we report the increase in held-out DCLM loss over the checkpoint without perturbation, averaged over three perturbation draws. The grid spans three model sizes, 30M to 100M, token–parameter ratios from 5 to 2,560 TPR, and perturbation magnitudes δ from 0.005 to 0.28. Magnitudes below $\delta \approx 0.02$ are omitted from the figures because the increase in loss is then at the level of the seed-to-seed variation for both optimizers.

Table 2 and the robustness counts below use the same protocol on the loss-optimal checkpoints of all five model sizes, 30M to 600M: 36 (model size, data budget) cells and perturbation magnitudes δ from 0.002 to 0.56.

E.2.2 Result

Figure 6 fixes the perturbation magnitude; Figure 22 reports the same difference of degradations at every magnitude we sweep. The left three panels fix δ and vary the model size, the right panel fixes the model size at 100M and varies δ. The difference is positive at every budget, model size and magnitude we tested, and grows with the budget at every δ. Magnitudes are not comparable across panels, since the degradation of both optimizers grows steeply with δ; what the panels share is the shape of the curve.

Over the 36 cells, the perturbation magnitude δ^* at which the loss has risen by 0.05, 0.1 or 0.2 nats is larger for AdamW in every cell and at every threshold, by a median factor of 1.4, and Muon's increase in loss exceeds AdamW's in every cell at every δ from 0.005 to 0.2. The ordering does not depend on which checkpoint of a cell is used: giving each optimizer its most perturbation-robust checkpoint within 0.02 nats of the cell's best base loss still leaves Muon degrading 20–400% more.

Loss under perturbation. Table 2 compares the loss under perturbation itself, rather than its increase, and places the result next to the fine-tuning crossover of Figure 2. At every model size, Muon's loss under perturbation is the lower one at small budgets and AdamW's at large ones, apart from the three marked entries. The crossover moves to smaller budgets as δ grows, so its value is not comparable to the fine-tuning crossover; its trend with model size is. From $\delta = 0.1$ on, the crossover budget does not increase with model size, as for every fine-tuning dataset; at $\delta = 0.05$ it is not monotone.

Table 2: Crossover budgets, in TPR (tokens per parameter). Left: the smallest budget at which AdamW's loss under a perturbation of magnitude δ is below Muon's (loss-optimal checkpoints, mean over three perturbation draws); > 20: Muon's is lower up to the largest 600M budget, 20 TPR; ≤ 5: AdamW's is lower already at the smallest budget; †: Muon's is lower again at 160 TPR, by 0.003 nats. Right: the fine-tuning crossover of Figure 2.

	Perturbation, δ				Fine-tuning			
Size	**0.05**	**0.1**	**0.13**	**0.2**	**Alpaca**	**GSM8K**	**MusicPile**	**StackMathQA**
30M	160	40	40	20	160	80	640	320
60M	80	40	40	20	80	80	80	80
100M	160	40	20	10	80	40	80	40
300M	$80^{\dagger}$	40	20	10	40	10	40	20
600M	> 20	20	10	≤ 5	20	5	10	10

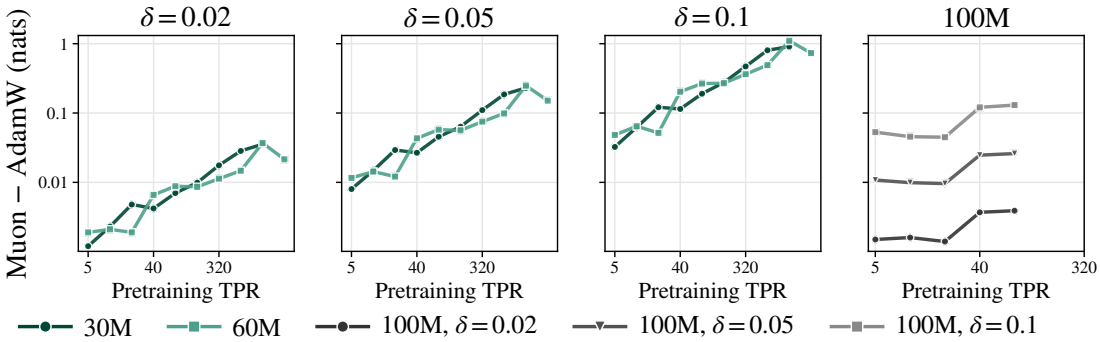


Figure 22: The perturbation degradation gap is positive at every perturbation magnitude. Muon minus AdamW increase in held-out DCLM loss. Left three panels: one δ each; right: 100M across δ.

E.3 Margin-to-spread ratio and robustness

E.3.1 Setup

Linear model. Each cell is a data budget $\gamma = N/d$ and a model dimension d, run with three seeds. We train W from $W_0 = 0$ for $T = 10^4$ full-batch steps on Zipf-distributed item frequencies, with GD at a learning rate of 0.02 and Muon at 0.005, and read ρ off the trained W. The same W is then perturbed, $W' = W + \sigma\|W\|_F\xi$ with $\|\xi\|_F = 1$, over 16 perturbation directions shared by the two optimizers, and σ_{50} is the perturbation level at which half of the items recovered at $\sigma = 0$ are lost, interpolated between grid points on a log scale.

Real language models. The analogue of a stored fact is a next-token prediction, and the analogue of the competitors of Definition 4.1 is the rest of the vocabulary. We evaluate the checkpoint on held-out DCLM text (128 sequences of 4096 tokens, never trained on). At a position with gold token t and logit vector z, the margins are

$$m_j \;=\; z_t - z_j, \qquad j \neq t, \tag{7}$$

over the whole vocabulary, and h and κ^2 are their mean and variance over those $|V| - 1$ competitors, exactly as in Definition 4.1. Two choices make this well defined. First, only the positions at which the gold token is already the argmax, $\min_{j\neq t} m_j > 0$, are used: elsewhere the prediction is wrong before any perturbation is added and the margin is not a robustness margin. Second, the position-level values h/κ are summarised by their median over positions rather than their mean, since the distribution over positions is heavy-tailed. The reduction over the vocabulary axis is carried out in double precision and chunked over positions, so the margin vector is never materialised. Both quantities are measured on the same checkpoints: for every (model size, token budget) cell, the checkpoint with the lowest held-out DCLM loss over the pretraining learning-rate grid, which gives 72 checkpoints over five model sizes. For robustness we take δ^*, the relative weight-perturbation level at which the held-out loss rises by ΔL nats, interpolated on a log–log scale and averaged over three perturbation draws.

Learning-rate ablation. At two data budgets, each optimizer is run at its learning rate and at that rate scaled by $\frac{1}{4}$ and by 4, all other settings unchanged, at $d = 128$ and $T = 10^4$ over several seeds.

γ	optimizer	lr	accuracy	final ρ	saturated ρ	$\mathcal{R}$
16	GD	0.005 ($\times\frac{1}{4}$)	0.2057 ± 0.0031	2.4610 ± 0.0142	2.4550 ± 0.0136	1.1354
	Muon	0.00125 ($\times\frac{1}{4}$)	0.1580 ± 0.0031	2.1470 ± 0.0125	2.1624 ± 0.0133	
	GD	0.02	0.1978 ± 0.0014	2.4609 ± 0.0145	2.4610 ± 0.0144	1.1468
	Muon	0.005	0.1566 ± 0.0024	2.1418 ± 0.0121	2.1461 ± 0.0123	
	GD	0.08 ($\times 4$)	0.1969 ± 0.0017	2.4607 ± 0.0146	2.4608 ± 0.0145	1.1445
	Muon	0.02 ($\times 4$)	0.1577 ± 0.0042	2.1492 ± 0.0111	2.1502 ± 0.0112	
64	GD	0.005 ($\times\frac{1}{4}$)	0.0356 ± 0.0013	0.9413 ± 0.0129	0.9384 ± 0.0128	1.1992
	Muon	0.00125 ($\times\frac{1}{4}$)	0.0272 ± 0.0013	0.7755 ± 0.0070	0.7825 ± 0.0073	
	GD	0.02	0.0344 ± 0.0012	0.9411 ± 0.0128	0.9412 ± 0.0129	1.2145
	Muon	0.005	0.0269 ± 0.0015	0.7731 ± 0.0070	0.7750 ± 0.0070	
	GD	0.08 ($\times 4$)	0.0341 ± 0.0011	0.9410 ± 0.0128	0.9410 ± 0.0128	1.2138
	Muon	0.02 ($\times 4$)	0.0271 ± 0.0014	0.7748 ± 0.0074	0.7753 ± 0.0073	

Table 3. *Learning-rate ablation in the linear model at $d = 128$, $T = 10^4$, mean $\pm$ sd over seeds; "final" is ρ at the last step and "saturated" its median over $t \geq 10^3$. Each $\mathcal{R}$ pairs the two rows beside it. Accuracy and ρ move by less than 1% over the 16$\times$ range, and so does $\mathcal{R}$.*

E.3.2 Result

Within each optimizer σ_{50} and ρ move together at every model dimension, over more than an order of magnitude in σ_{50}, and the association is monotone in ρ rather than in d or γ separately (Figure 23). In the language models, the Spearman correlation between ρ and δ^* over the 72 checkpoints is $+0.83$ at $\Delta L = 0.05$ and 0.1 nats and $+0.84$ at 0.2 nats; within each optimizer it is $+0.87$ to $+0.89$ for AdamW and $+0.92$ to $+0.93$ for Muon (Figure 23). Across the 16$\times$ learning-rate range, accuracy moves by less than one percentage point and ρ by less than 1%, far below the gap between the two optimizers, and the ratio $\mathcal{R}$ between them by at most 1.3% (Table 3); the comparison in Section 4.2 is therefore not an artefact of the learning rates.

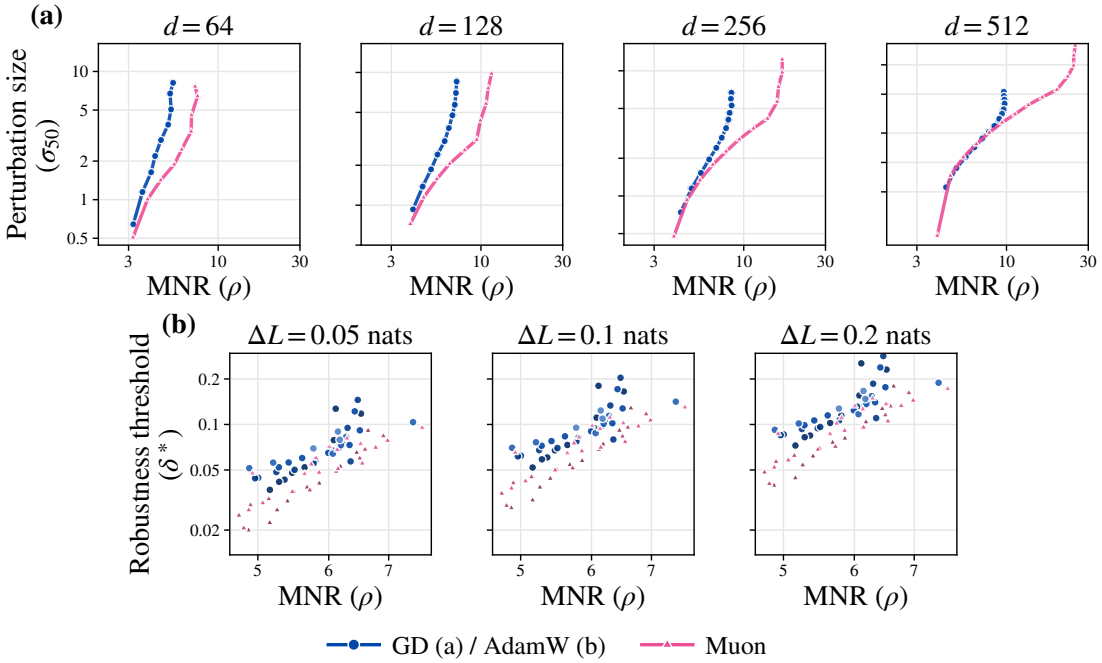


Figure 23: Top: σ_{50} against ρ in the linear model, one panel per model dimension, one point per γ cell, below capacity only. **Bottom: real language models, δ^* against ρ at every threshold**; marker shape is the model size.

F Proof of Theorem 4.2

This section proves the quantitative bound underlying the informal Theorem 4.2. For every fixed data budget $\gamma = N/d \geq 16$, we establish

$$\left|\mathcal{R} - \frac{3\pi}{8}\left(1 - \frac{9}{10\gamma}\right)\right| \leq C_1\gamma^{-3/2} + o_p(1) \tag{8}$$

as $d \to \infty$, where C_1 is an absolute constant. The fixed-percentage increase is established in Lemma F.33; the other conclusions of Theorem 4.2 follow from the negative $1/\gamma$ term in $\mathcal{R} - 3\pi/8$, as detailed below.

We first compute the one-step GD and Muon iterates and collect concentration estimates for the Gaussian data. We then evaluate the mean margin $\bar{h}$ and the margin spread $\bar{\kappa}$ of Definition 4.1 separately for the two iterates, and finally take the ratio of the two margin-to-spread ratios. The GD estimates hold for every $\gamma > 1$; the Muon estimates, and hence the theorem, require $\gamma \geq 16$.

Setup. We work in the setting of Section 4.1 with *uniform query weights*, $q_i = 1/N$ for every fact i; this is the setting of Theorem 4.2 and is assumed throughout. The keys and values $u_1, \dots, u_N, v_1, \dots, v_N$ are independent $\mathcal{N}(0, I_d)$ vectors, $N = \gamma d$ with the data budget $\gamma > 1$ held fixed, $N - 1 \geq d$, and $d \to \infty$; O_p and o_p are as in Section 4.1.

Notation. We collect the data in $U = [u_1 \cdots u_N] \in \mathbb{R}^{d\times N}$ and $V = [v_1 \cdots v_N] \in \mathbb{R}^{d\times N}$, and write $\mathbf{1} = (1, \dots, 1)^\top$, e_i for the i-th standard basis vector and $\langle X, Y\rangle = \operatorname{tr}(X^\top Y)$. The operator, Frobenius and nuclear norms $\|\cdot\|_{\mathrm{op}}$, $\|\cdot\|_F$ and $\|\cdot\|_*$ satisfy

$$\begin{aligned} &\|X\|_{\mathrm{op}} \leq \|X\|_F \leq \|X\|_* \leq \operatorname{rank}(X)\|X\|_{\mathrm{op}}, \qquad \|X\|_F \leq \sqrt{\operatorname{rank}(X)}\,\|X\|_{\mathrm{op}}, \\ &|\langle X, Y\rangle| \leq \|X\|_*\|Y\|_{\mathrm{op}}. \end{aligned} \tag{9}$$

For $X \in \mathbb{R}^{d\times d}$ with compact SVD $X = P\Sigma Q^\top$ we write $\operatorname{Orth}(X) = PQ^\top$ for its polar factor, so that $\operatorname{Orth}(cX) = \operatorname{Orth}(X)$ for $c > 0$ and $\|\operatorname{Orth}(X)\|_{\mathrm{op}} \leq 1$.

Scores and observables. By Definition 4.1, for any W,

$$\begin{aligned} \bar{h} &= \operatorname*{mean}_i u_i^\top W v_i - \operatorname*{mean}_i \operatorname*{mean}_{j\neq i} u_j^\top W v_i, \\ \bar{\kappa}^2 &= \frac{1}{N(N-1)}\sum_{i=1}^N \sum_{j\neq i} \left(u_j^\top W v_i\right)^2 - \operatorname*{mean}_i \left(\operatorname*{mean}_{j\neq i} u_j^\top W v_i\right)^2, \end{aligned} \tag{10}$$

where the scores are written as $u_j^\top W v_i$ rather than $v_j^\top W u_i$ as in Section 4.1. The two conventions are exchanged by swapping the roles of U and V: reading Section 4.1 with (U, V) replaced by (V, U) turns every score $v_j^\top W u_i$ into $u_j^\top W v_i$ and its one-step iterate into Equation (11), leaving the loss and every h_i and κ_i unchanged. Since U and V are i.i.d., the law of ρ is the same under either convention.

F.1 The two one-step iterates

Lemma F.1 (One step). *Up to a positive multiple, one step of GD and one step of Muon from $W_0 = 0$ give*

$$W^{\mathrm{GD}} = UV^\top - \frac{1}{N}(U\mathbf{1})(V\mathbf{1})^\top = U\left(I_N - \frac{\mathbf{1}\mathbf{1}^\top}{N}\right)V^\top, \qquad W^{\mathrm{Muon}} = \operatorname{Orth}\left(W^{\mathrm{GD}}\right). \tag{11}$$

The positive multiple is immaterial for ρ.

Proof. Gradient. With $\partial_W(u_j^\top W v_i) = u_j v_i^\top$, the gradient of the loss of Section 4.1 with uniform weights $q_i = 1/N$ is

$$-\nabla_W \mathcal{L}(W) = \frac{1}{N}\sum_{i=1}^N \left(u_i - \sum_{j=1}^N \hat{p}_W(j \mid i) u_j\right) v_i^\top. \tag{12}$$

At $W = 0$ all scores vanish, so $\hat{p}_0(j \mid i) = 1/N$; moreover $\sum_j u_j = U\mathbf{1}$, $\sum_i u_i v_i^\top = UV^\top$ and $\sum_i v_i^\top = (V\mathbf{1})^\top$, whence

$$-\nabla_W \mathcal{L}(0) = \frac{1}{N}\left(UV^\top - \frac{1}{N}(U\mathbf{1})(V\mathbf{1})^\top\right) = \frac{1}{N}U\left(I_N - \frac{\mathbf{1}\mathbf{1}^\top}{N}\right)V^\top. \tag{13}$$

One step. With step size $\eta > 0$, GD gives $-\eta \nabla_W \mathcal{L}(0) = \frac{\eta}{N} W^{\mathrm{GD}}$ and Muon gives $\eta \,\mathrm{Orth}(-\nabla_W \mathcal{L}(0)) = \eta \,\mathrm{Orth}(W^{\mathrm{GD}})$, using $\mathrm{Orth}(cX) = \mathrm{Orth}(X)$ for $c > 0$. *Normalization.* Replacing W by cW with $c > 0$ multiplies every score $u_j^\top W v_i$ by c, hence $h_i \mapsto c h_i$ and $\kappa_i \mapsto c\kappa_i$, so $\bar{h} \mapsto c\bar{h}$, $\bar{\kappa} \mapsto c\bar{\kappa}$ and ρ is unchanged. □

Lemma F.2 (Full rank and orthogonality). *Almost surely* $\mathrm{rank}\, W^{\mathrm{GD}} = d$. *Consequently* W^{Muon} *is an orthogonal matrix:* $\|W^{\mathrm{Muon}} x\| = \|x\|$, $\|W^{\mathrm{Muon}}\|_{\mathrm{op}} = 1$ *and* $\|W^{\mathrm{Muon}}\|_F^2 = d$.

Proof. $\det W^{\mathrm{GD}}$ is a polynomial in the entries of (U, V). Taking $U = [\, I_d \;\; 0\,]$ and $V = [\, I_d \;\; -\sum_{k \le d} e_k \;\; 0\,]$, which is possible because $N \ge d + 1$, gives $V\mathbf{1} = 0$ and $W^{\mathrm{GD}} = UV^\top = I_d$, so the polynomial is not identically zero; since (U, V) has a density, $\Pr[\det W^{\mathrm{GD}} = 0] = 0$. At full rank the factors P, Q of the compact SVD are $d \times d$ orthogonal matrices, so $W^{\mathrm{Muon}} = PQ^\top$ is orthogonal and $\|W^{\mathrm{Muon}}\|_F^2 = \mathrm{rank}\, W^{\mathrm{GD}} = d$. □

F.2 Concentration estimates for the data

Lemma F.3 (Concentration for U, V). *As $d \to \infty$ along the sequence above:*

(i) $\mathrm{mean}_k \|u_k\|^2 = d\big(1 + O_p((Nd)^{-1/2})\big)$, *and likewise for* V.

(ii) $\mathrm{mean}_k \left(\|u_k\|\|v_k\| - d\right)^2 = O_p(d)$.

(iii) $\|U\mathbf{1}\|^2/(Nd) \to 1$ *in probability; in particular* $\|U\mathbf{1}\|^2 = O_p(Nd)$. *Likewise for* V.

(iv) $\|U\|_{\mathrm{op}}/\sqrt{N} \to 1 + \gamma^{-1/2}$ *almost surely, hence* $\|U\|_{\mathrm{op}}^2 \le 4N(1 + o_p(1))$; *and for every* $m \ge 1$ *there is* C_m *with* $\mathbb{E}\|U\|_{\mathrm{op}}^{2m} \le C_m N^m$. *Likewise for* V.

(v) $\|UU^\top - NI_d\|_{\mathrm{op}} = O_p(\sqrt{Nd})$. *If* $\gamma \ge 16$, *then with probability tending to one both* $S = \frac{1}{N} VV^\top$ *and* $S = \frac{1}{N}(VV^\top - \frac{1}{N}(V\mathbf{1})(V\mathbf{1})^\top)$ *satisfy* $\|S - I_d\|_{\mathrm{op}} \le 3\gamma^{-1/2}$ *and* $\lambda_{\min}(S) \ge \frac{1}{8}$.

(vi) $\|VV^\top\|_F^2 = O_p(N^2 d)$, *and* $\mathbb{E}\|VV^\top\|_F^{2m} \le C_{2m}(N^2 d)^m$. *Likewise for* $UU^\top$.

(vii) $\|UV^\top\|_F^2 = Nd^2(1 + O_p(d^{-1}))$ *and* $\|UV^\top\|_F^2 - d\sum_k \|v_k\|^2 = O_p(\sqrt{N} d^{3/2} + Nd)$.

(viii) $UV^\top \overset{d}{=} (UU^\top)^{1/2}\Xi$ *with* $\Xi \in \mathbb{R}^{d \times d}$ *having i.i.d.* $\mathcal{N}(0,1)$ *entries independent of* U; *with probability tending to one,* $\|UV^\top\|_{\mathrm{op}} \le 5\sqrt{Nd}$.

Proof. (i) $\sum_k \|u_k\|^2 \sim \chi^2_{Nd}$ has mean Nd and variance $2Nd$, so Chebyshev gives $\sum_k \|u_k\|^2 = Nd + O_p(\sqrt{Nd})$; divide by N.

(ii) For $x \ge 0$ one has $\sqrt{x} \ge 1 + \frac{x-1}{2} - \frac{(x-1)^2}{2}$: the difference g of the two sides satisfies $g(0) = g(1) = g'(1) = 0$ and $g''(x) = 1 - \frac{1}{4}x^{-3/2}$ changes sign once, at $x_0 = 4^{-2/3}$, so g is concave on $[0, x_0]$ and convex on $[x_0, \infty)$, and the two boundary values force $g \ge 0$. With $u \sim \mathcal{N}(0, I_d)$, $\mathbb{E}\|u\|^2/d = 1$ and $\mathrm{var}(\|u\|^2/d) = 2/d$, so $\mathbb{E}\|u\| \ge \sqrt{d}(1 - \frac{1}{d})$ and $(\mathbb{E}\|u\|)^2 \ge d - 2$. Hence for independent u, v,

$$\mathbb{E}\big(\|u\|\|v\| - d\big)^2 = \mathbb{E}\|u\|^2\, \mathbb{E}\|v\|^2 - 2d\, \mathbb{E}\|u\|\, \mathbb{E}\|v\| + d^2 \le 2d^2 - 2d(d-2) = 4d, \tag{14}$$

and Markov's inequality gives the claim.

(iii) $U\mathbf{1} = \sum_k u_k \sim \mathcal{N}(0, NI_d)$, so $\|U\mathbf{1}\|^2/N \sim \chi^2_d$ has mean d and variance $2d$; Chebyshev gives $\|U\mathbf{1}\|^2/(Nd) = 1 + O_p(d^{-1/2})$.

(iv) $\|U\|_{\mathrm{op}}^2 = \lambda_{\max}(UU^\top)$, and by the Bai–Yin law $\lambda_{\max}(UU^\top/N) \to (1 + \gamma^{-1/2})^2 \le 4$ almost surely, since $\gamma \ge 1$. For the moments, $U \mapsto \|U\|_{\mathrm{op}}$ is 1-Lipschitz for $\|\cdot\|_F$ and $\mathbb{E}\|U\|_{\mathrm{op}} \le \sqrt{N} + \sqrt{d} \le 2\sqrt{N}$ by Gordon's theorem, so Gaussian concentration gives $\Pr[\|U\|_{\mathrm{op}} \ge 2\sqrt{N} + t] \le e^{-t^2/2}$; integrating the tail bounds all moments.

(v) By Bai–Yin, $\lambda_{\max}(UU^\top/N) \to (1 + \gamma^{-1/2})^2$ and $\lambda_{\min}(UU^\top/N) \to (1 - \gamma^{-1/2})^2$ almost surely, so $\|UU^\top/N - I_d\|_{\mathrm{op}} \to \max\{(1+\gamma^{-1/2})^2 - 1, 1 - (1-\gamma^{-1/2})^2\} = 2\gamma^{-1/2} + \gamma^{-1}$, which is at most $3\gamma^{-1/2} < 4\gamma^{-1/2}$; multiplying by N gives the first claim. For $\gamma \ge 16$ we have $\gamma^{-1/2} \le \frac{1}{4}$, so in probability $\|VV^\top/N - I_d\|_{\mathrm{op}} \to 2\gamma^{-1/2} + \gamma^{-1} \le \frac{9}{4}\gamma^{-1/2} < 3\gamma^{-1/2}$ and $\lambda_{\min}(VV^\top/N) \to (1 - \gamma^{-1/2})^2 \ge \frac{9}{16} > \frac{1}{8}$, and the inequalities being strict they hold with

probability tending to one. For the centred version, write $\frac{1}{N}\big(VV^\top - \frac{1}{N}(V\mathbf{1})(V\mathbf{1})^\top\big) = \frac{VV^\top}{N} - \frac{1}{N^2}(V\mathbf{1})(V\mathbf{1})^\top$, where $\frac{1}{N^2}(V\mathbf{1})(V\mathbf{1})^\top \succeq 0$ has operator norm $\|V\mathbf{1}\|^2/N^2 = \frac{\|V\mathbf{1}\|^2}{Nd}\cdot\frac{d}{N} \to \gamma^{-1}$ in probability by (iii). By Weyl's inequality,

$$\begin{aligned}\lambda_{\min}\Big(\frac{1}{N}\big(VV^\top - \frac{1}{N}(V\mathbf{1})(V\mathbf{1})^\top\big)\Big) &\geq \lambda_{\min}\Big(\frac{VV^\top}{N}\Big) - \frac{\|V\mathbf{1}\|^2}{N^2},\\ \Big\|\frac{1}{N}\big(VV^\top - \frac{1}{N}(V\mathbf{1})(V\mathbf{1})^\top\big) - I_d\Big\|_{\mathrm{op}} &\leq \Big\|\frac{VV^\top}{N} - I_d\Big\|_{\mathrm{op}} + \frac{\|V\mathbf{1}\|^2}{N^2},\end{aligned} \tag{15}$$

and in probability the two right-hand sides converge to

$$(1-\gamma^{-1/2})^2 - \gamma^{-1} \geq \frac{9}{16} - \frac{1}{16} = \frac{1}{2} > \frac{1}{8}, \qquad 2\gamma^{-1/2} + 2\gamma^{-1} \leq \frac{5}{2}\gamma^{-1/2} < 3\gamma^{-1/2}, \tag{16}$$

using $\gamma^{-1/2} \leq \frac{1}{4}$; the inequalities being strict, they hold with probability tending to one.

(vi) $VV^\top$ has rank at most d, so $\|VV^\top\|_F \leq \sqrt{d}\|VV^\top\|_{\mathrm{op}} = \sqrt{d}\|V\|_{\mathrm{op}}^2$ by Equation (9); now use (iv).

(vii) Expanding, $\|UV^\top\|_F^2 = \sum_{k,l}(u_k^\top u_l)(v_k^\top v_l) = \sum_k \|u_k\|^2\|v_k\|^2 + \sum_{k\neq l}(u_k^\top u_l)(v_k^\top v_l)$.

Mean. $\mathbb{E}\|u_k\|^2\|v_k\|^2 = d^2$ and, for $k \neq l$, $\mathbb{E}(u_k^\top u_l)(v_k^\top v_l) = 0$, so $\mathbb{E}\|UV^\top\|_F^2 = Nd^2$.

Variance. The variables $\|u_k\|^2\|v_k\|^2$, $k \in [N]$, are independent with $\mathrm{var}(\|u_k\|^2\|v_k\|^2) = (d^2+2d)^2 - d^4 \leq 8d^3$. Next, $\sum_{k\neq l}(u_k^\top u_l)(v_k^\top v_l) = 2\sum_{k<l}(u_k^\top u_l)(v_k^\top v_l)$; for unordered pairs $\{k,l\} \neq \{k',l'\}$ one of the indices k, l, k', l' occurs exactly once, and integrating over the u with that index gives $\mathbb{E}\big[(u_k^\top u_l)(v_k^\top v_l)(u_{k'}^\top u_{l'})(v_{k'}^\top v_{l'})\big] = 0$; since $\mathbb{E}(u_k^\top u_l)^2(v_k^\top v_l)^2 = d^2$, this gives $\mathrm{var}\sum_{k\neq l}(u_k^\top u_l)(v_k^\top v_l) = 4\cdot\frac{N(N-1)}{2}\cdot d^2 \leq 2N^2d^2$. For $k' \in [N]$ and $k \neq l$, one of k, l differs from k', and integrating over the u with that index gives $\mathbb{E}\big[(\|u_{k'}\|^2\|v_{k'}\|^2 - d^2)(u_k^\top u_l)(v_k^\top v_l)\big] = 0$, so the two sums are uncorrelated. Therefore

$$\mathrm{var}\|UV^\top\|_F^2 \leq 8Nd^3 + 2N^2d^2, \qquad \frac{\big(\mathrm{var}\|UV^\top\|_F^2\big)^{1/2}}{\mathbb{E}\|UV^\top\|_F^2} \leq \sqrt{\frac{8}{Nd} + \frac{2}{d^2}} \leq \frac{4}{d}, \tag{17}$$

and Chebyshev gives the first claim.

Centred version. $\|UV^\top\|_F^2 - d\sum_k\|v_k\|^2 = \sum_k(\|u_k\|^2 - d)\|v_k\|^2 + \sum_{k\neq l}(u_k^\top u_l)(v_k^\top v_l)$ has mean zero. The variables $(\|u_k\|^2 - d)\|v_k\|^2$ are independent with variance $\mathbb{E}(\|u\|^2 - d)^2\,\mathbb{E}\|v\|^4 = 2d(d^2+2d) \leq 6d^3$; the second sum has variance at most $2N^2d^2$ as above; and for $k' \in [N]$ and $k \neq l$, integrating over the u whose index among k, l differs from k' gives $\mathbb{E}\big[(\|u_{k'}\|^2 - d)\|v_{k'}\|^2(u_k^\top u_l)(v_k^\top v_l)\big] = 0$. Hence the variance is at most $6Nd^3 + 2N^2d^2$, its square root is at most $\sqrt{6}\sqrt{N}d^{3/2} + \sqrt{2}Nd$, and Chebyshev gives the second claim.

(viii) Conditionally on U, the columns of $UV^\top$ are i.i.d. $\mathcal{N}(0, UU^\top)$, i.e. of the form $(UU^\top)^{1/2}g$ with $g \sim \mathcal{N}(0, I_d)$; stacking the d vectors g as the columns of Ξ gives the stated equality in distribution, and the conditional law of Ξ does not depend on U, so Ξ is independent of U. For the norm, $\|(UU^\top)^{1/2}\Xi\|_{\mathrm{op}} \leq \|U\|_{\mathrm{op}}\|\Xi\|_{\mathrm{op}}$, and almost surely $\|U\|_{\mathrm{op}}/\sqrt{N} \to 1 + \gamma^{-1/2} \leq 2$ by (iv) while $\|\Xi\|_{\mathrm{op}}/\sqrt{d} \to 2$ by Bai–Yin with aspect ratio 1. □

Lemma F.4 (Norms of W^{GD}).

(i) $\|W^{\mathrm{GD}}\|_F^2 = (N-1)d^2(1+O_p(d^{-1}))$, *and* $\|W^{\mathrm{GD}}\|_F^2 - d\,\mathrm{tr}\,\big(VV^\top - \frac{1}{N}(V\mathbf{1})(V\mathbf{1})^\top\big) = O_p(\sqrt{N}d^{3/2} + Nd)$.

(ii) $\|W^{\mathrm{GD}}\|_{\mathrm{op}} \leq \|U\|_{\mathrm{op}}\|V\|_{\mathrm{op}}$, *and with probability tending to one* $\|W^{\mathrm{GD}}\|_{\mathrm{op}} \leq 7\sqrt{Nd}$.

(iii) With probability tending to one, $\frac{1}{8}d\sqrt{Nd} \leq \|W^{\mathrm{GD}}\|_* \leq 7d\sqrt{Nd}$.

(iv) In probability, $\|W^{\mathrm{GD}}\|_*/(Nd) \to \Phi(\gamma)/\sqrt{\gamma}$, *where*

$$\Phi(\gamma) := \frac{1}{\pi}\int_0^1 \left[\sqrt{1-\frac{p}{\gamma}}\,\frac{1+p}{\sqrt{p}} + \frac{\sqrt{p}\,(1-p)}{\gamma\sqrt{1-p/\gamma}}\right]\mathrm{d}p, \tag{18}$$

and for $\gamma \geq 16$

$$\Phi(\gamma) = \frac{8}{3\pi}\Big(1 - \frac{1}{10\gamma} + r(\gamma)\Big), \qquad |r(\gamma)| \leq \gamma^{-2}. \tag{19}$$

Part (iv) is where the constant $3\pi/8$ enters: Muon replaces W^{GD} by its polar factor, which turns the diagonal score into $\|W^{\mathrm{GD}}\|_*$, and $8/(3\pi)$ is the limiting Marčenko–Pastur integral at $\gamma = \infty$.

Proof. (i) By Equation (11) and $\|X-Y\|_F^2 = \|X\|_F^2 - 2\langle X, Y\rangle + \|Y\|_F^2$,

$$\|W^{\mathrm{GD}}\|_F^2 = \|UV^\top\|_F^2 - \frac{2}{N}(U\mathbf{1})^\top UV^\top(V\mathbf{1}) + \frac{1}{N^2}\|U\mathbf{1}\|^2\|V\mathbf{1}\|^2, \tag{20}$$

where the cross term is $O_p(N^2 d)/N = O_p(Nd)$ and the last term is $O_p(d^2)$ by Theorem F.3(iii),(iv),(viii). Substituting $\|UV^\top\|_F^2 = Nd^2(1+O_p(d^{-1}))$ from Theorem F.3(vii) and using $d^2 \le Nd$ gives the first claim. For the second, $\mathrm{tr}(VV^\top) = \sum_k \|v_k\|^2$ and $\mathrm{tr}((V\mathbf{1})(V\mathbf{1})^\top) = \|V\mathbf{1}\|^2$, so the same expansion gives

$$\begin{aligned}\|W^{\mathrm{GD}}\|_F^2 - d\,\mathrm{tr}\left(VV^\top - \frac{1}{N}(V\mathbf{1})(V\mathbf{1})^\top\right) &= \left(\|UV^\top\|_F^2 - d\sum_k \|v_k\|^2\right) + \frac{d\|V\mathbf{1}\|^2}{N} \\ &\quad - \frac{2}{N}(U\mathbf{1})^\top UV^\top(V\mathbf{1}) + \frac{\|U\mathbf{1}\|^2\|V\mathbf{1}\|^2}{N^2},\end{aligned} \tag{21}$$

whose four terms are $O_p(\sqrt{N}d^{3/2} + Nd)$ by the second half of Theorem F.3(vii), $O_p(d^2)$, $O_p(Nd)$ and $O_p(d^2)$ respectively; since $d^2 \le Nd$, the claim follows.

(ii) $\|I_N - \frac{\mathbf{1}\mathbf{1}^\top}{N}\|_{\mathrm{op}} = 1$ gives the first inequality. For the second, the triangle inequality and the operator norm of a rank-one matrix give $\|W^{\mathrm{GD}}\|_{\mathrm{op}} \le \|UV^\top\|_{\mathrm{op}} + \|U\mathbf{1}\|\|V\mathbf{1}\|/N \le 5\sqrt{Nd} + \frac{3}{2}d \le 7\sqrt{Nd}$ with probability tending to one, by Theorem F.3(iii),(viii) and $d \le \sqrt{Nd}$.

(iii) For any X, $\|X\|_F^2 = \sum_a s_a^2 \le s_1 \sum_a s_a = \|X\|_{\mathrm{op}}\|X\|_*$, and $\|X\|_* \le \mathrm{rank}(X)\|X\|_{\mathrm{op}}$. Combining with (i) and (ii),

$$\begin{aligned}\|W^{\mathrm{GD}}\|_* &\ge \frac{\|W^{\mathrm{GD}}\|_F^2}{\|W^{\mathrm{GD}}\|_{\mathrm{op}}} \ge \frac{(N-1)d^2(1-o_p(1))}{7\sqrt{Nd}} \ge \frac{1}{8}d\sqrt{Nd}, \\ \|W^{\mathrm{GD}}\|_* &\le d\|W^{\mathrm{GD}}\|_{\mathrm{op}} \le 7d\sqrt{Nd}.\end{aligned} \tag{22}$$

(iv) *Removing the centring.* By the triangle inequality for the nuclear norm and the nuclear norm of a rank-one matrix, $|\|UV^\top\|_* - \|W^{\mathrm{GD}}\|_*| \le \|U\mathbf{1}\|\|V\mathbf{1}\|/N = O_p(d)$, which is $o(Nd)$; so it suffices to prove the limit for $\|UV^\top\|_*$.

Limit. By Theorem F.3(viii), $\{s_a(UV^\top)^2/(Nd)\}_{a\le d}$ has the same law as the eigenvalues of $\frac{1}{d}(\frac{UU^\top}{N})^{1/2}\Xi\Xi^\top(\frac{UU^\top}{N})^{1/2}$ with Ξ independent of U: conditionally on U this is a sample covariance matrix with population covariance $\frac{1}{N}UU^\top$ and aspect ratio 1, and the spectral distribution of $\frac{1}{N}UU^\top$ converges almost surely to the Marčenko–Pastur law H with ratio $1/\gamma$ and mean 1. By the Marčenko–Pastur–Silverstein theorem the empirical distribution of $\{s_a(UV^\top)^2/(Nd)\}$ converges weakly in probability to a law ν whose Stieltjes transform m satisfies

$$z = -\frac{1}{m} + \int \frac{t}{1+tm}\,\mathrm{d}H(t). \tag{23}$$

By Theorem F.3(viii) all $s_a^2/(Nd) \le 25$ with probability tending to one, so both the empirical law and ν are supported in $[0,25]$; taking the limit for the bounded continuous function $\lambda \mapsto \min(\sqrt{\lambda}, 5)$ gives $\frac{1}{d}\sum_a s_a/\sqrt{Nd} \to \int \sqrt{\lambda}\,\mathrm{d}\nu$ in probability, that is

$$\frac{\|UV^\top\|_*}{Nd} = \sqrt{\frac{d}{N}}\cdot\frac{1}{d}\sum_a \frac{s_a}{\sqrt{Nd}} \longrightarrow \frac{1}{\sqrt{\gamma}}\int \sqrt{\lambda}\,\mathrm{d}\nu(\lambda). \tag{24}$$

Closed form. Let $\psi(x) := \int \frac{\lambda x}{1-\lambda x}\mathrm{d}\nu$. From $\frac{\lambda x}{1-\lambda x} = -1 + \frac{1/x}{1/x-\lambda}$ we get $\psi(x) = -1 - \frac{1}{x}m(1/x)$, i.e. $m(1/x) = -x(1+\psi(x))$. Taking $z = 1/x$ in Equation (23) and multiplying by $x(1+\psi(x))$,

$$\psi(x) = \int \frac{x(1+\psi(x))t}{1-x(1+\psi(x))t}\,\mathrm{d}H(t). \tag{25}$$

On the other hand the Stieltjes transform of the Marčenko–Pastur law with ratio $1/\gamma$ and mean 1 satisfies $\frac{1}{\gamma}zm_H(z)^2 - (1-\frac{1}{\gamma}-z)m_H(z)+1 = 0$, and $\frac{ut}{1-ut} = -1 + \frac{1/u}{1/u-t}$ gives $m_H(1/u) = -u\big(1+\int \frac{ut}{1-ut}\mathrm{d}H(t)\big)$; substituting $z = 1/u$ into the quadratic yields, for all u,

$$\frac{u}{\gamma}\Big(1+\int\frac{ut}{1-ut}\mathrm{d}H(t)\Big)^2 + \Big(1-\frac{1}{\gamma}\Big)u\Big(1+\int\frac{ut}{1-ut}\mathrm{d}H(t)\Big) - \int\frac{ut}{1-ut}\mathrm{d}H(t) = 0. \tag{26}$$

Take $u = x(1+\psi(x))$. By Equation (25) the integral then equals $\psi(x)$, so

$$\frac{1}{\gamma}x(1+\psi)^3 + \Big(1-\frac{1}{\gamma}\Big)x(1+\psi)^2 - \psi = 0, \qquad \text{i.e.} \qquad x = \frac{\psi(x)}{(1+\psi(x))^2\,(1+\psi(x)/\gamma)}. \tag{27}$$

Using $\pi\sqrt{\lambda} = \int_0^\infty \frac{\lambda}{\lambda+w}\frac{\mathrm{d}w}{\sqrt{w}}$ and Tonelli's theorem, and then substituting $p = -\psi(-1/w) \in (0,1)$,

$$\begin{aligned} \int \sqrt{\lambda}\,\mathrm{d}\nu &= \frac{1}{\pi}\int_0^\infty \frac{-\psi(-1/w)}{\sqrt{w}}\,\mathrm{d}w, \\ w &= \frac{(1-p)^2(1-p/\gamma)}{p}, \qquad -\frac{\mathrm{d}w}{\mathrm{d}p} = w\Big[\frac{2}{1-p} + \frac{1/\gamma}{1-p/\gamma} + \frac{1}{p}\Big], \\ \int \sqrt{\lambda}\,\mathrm{d}\nu &= \frac{1}{\pi}\int_0^1 \sqrt{1-\frac{p}{\gamma}}\Big[2\sqrt{p} + \frac{\sqrt{p}(1-p)/\gamma}{1-p/\gamma} + \frac{1-p}{\sqrt{p}}\Big]\mathrm{d}p; \end{aligned} \tag{28}$$

then $2\sqrt{p} + \frac{1-p}{\sqrt{p}} = \frac{1+p}{\sqrt{p}}$ and $\sqrt{1-p/\gamma}\cdot\frac{\sqrt{p}(1-p)/\gamma}{1-p/\gamma} = \frac{\sqrt{p}(1-p)}{\gamma\sqrt{1-p/\gamma}}$ give Equation (18).

Expansion. For $0 \le s \le \frac{1}{16}$ one has

$$\Big|\sqrt{1-s} - 1 + \frac{s}{2}\Big| \le \frac{s^2}{5}, \qquad 0 \le (1-s)^{-1/2} - 1 \le \frac{2s}{3}. \tag{29}$$

For the first, $g(s) := 1 - \frac{s}{2} - \sqrt{1-s}$ satisfies $g(0) = 0$, $g'(s) = \frac{1}{2}((1-s)^{-1/2} - 1) \ge 0$ and $g''(s) = \frac{1}{4}(1-s)^{-3/2} \le \frac{1}{4}(1-\frac{1}{16})^{-3/2} < \frac{2}{5}$ on $[0, \frac{1}{16}]$, so $g(s) \le \frac{s^2}{2}\max g'' < \frac{s^2}{5}$. For the second, $f(s) := (1-s)^{-1/2} - 1 - \frac{2s}{3}$ satisfies $f(0) = 0$ and $f'(s) = \frac{1}{2}(1-s)^{-3/2} - \frac{2}{3} < 0$ for $s < 1 - (\frac{3}{4})^{2/3} \approx 0.175$, hence on $[0, \frac{1}{16}]$, so $f \le 0$ there. Take $s = p/\gamma$ with $\gamma \ge 16$. From

$$\frac{1}{\pi}\int_0^1 \frac{1+p}{\sqrt{p}}\,\mathrm{d}p = \frac{8}{3\pi}, \qquad -\frac{1}{2\gamma\pi}\int_0^1 \sqrt{p}(1+p)\,\mathrm{d}p + \frac{1}{\gamma\pi}\int_0^1 \sqrt{p}(1-p)\,\mathrm{d}p = \frac{8}{3\pi}\Big(-\frac{1}{10\gamma}\Big), \tag{30}$$

and Equation (18), the remainder is bounded by

$$\begin{aligned} \Big|\Phi(\gamma) - \frac{8}{3\pi}\Big(1-\frac{1}{10\gamma}\Big)\Big| &\le \frac{1}{\pi\gamma^2}\Big[\frac{1}{5}\int_0^1 p^{3/2}(1+p)\,\mathrm{d}p + \frac{2}{3}\int_0^1 p^{3/2}(1-p)\,\mathrm{d}p\Big] \\ &= \frac{1}{\pi\gamma^2}\Big[\frac{24}{175} + \frac{8}{105}\Big] < \frac{8}{3\pi}\gamma^{-2}, \end{aligned} \tag{31}$$

which is Equation (19). □

F.3 The GD step

By Equation (10) with $W = W^{\mathrm{GD}}$,

$$\bar{h}^{\mathrm{GD}} = \underbrace{\operatorname*{mean}_i u_i^\top W^{\mathrm{GD}} v_i}_{\text{diagonal}} - \underbrace{\operatorname*{mean}_i \operatorname*{mean}_{j\neq i} u_j^\top W^{\mathrm{GD}} v_i}_{\text{off-diagonal}}. \tag{32}$$

Lemma F.5 (Diagonal-sum identity). *The diagonal scores of W^{GD} sum to*

$$\sum_{i=1}^N u_i^\top W^{\mathrm{GD}} v_i = \|W^{\mathrm{GD}}\|_F^2 + \frac{1}{N}(U\mathbf{1})^\top W^{\mathrm{GD}}(V\mathbf{1}).$$

Proof. A scalar equals its own trace and $\sum_i u_i v_i^\top = UV^\top$, so $\sum_i u_i^\top W^{\mathrm{GD}} v_i = \operatorname{tr}\big((W^{\mathrm{GD}})^\top \sum_i u_i v_i^\top\big) = \langle W^{\mathrm{GD}}, UV^\top\rangle$. Rearranging Equation (11) gives $UV^\top = W^{\mathrm{GD}} + \frac{1}{N}(U\mathbf{1})(V\mathbf{1})^\top$, and bilinearity together with $\langle Z, ab^\top\rangle = a^\top Z b$ gives the claim. □

Lemma F.6 (Rank-one cross term). $\frac{1}{N^2}|(U\mathbf{1})^\top W^{\mathrm{GD}}(V\mathbf{1})| = O_p(d)$.

Proof. By Lemma F.4(ii) and Theorem F.3(iii),(iv),

$$\frac{|(U\mathbf{1})^\top W^{\mathrm{GD}}(V\mathbf{1})|}{N^2} \le \frac{\|U\mathbf{1}\|\,\|W^{\mathrm{GD}}\|_{\mathrm{op}}\,\|V\mathbf{1}\|}{N^2} = \frac{O_p(\sqrt{Nd})\cdot O_p(N)\cdot O_p(\sqrt{Nd})}{N^2} = O_p(d). \tag{33}$$

□

Lemma F.7. $\mathrm{mean}_i\, u_i^\top W^{\mathrm{GD}} v_i = d^2(1+O_p(d^{-1}))$.

Proof. Divide Theorem F.5 by N. By Lemma F.4(i), $\|W^{\mathrm{GD}}\|_F^2/N = \frac{N-1}{N} d^2(1+O_p(d^{-1})) = d^2(1+O_p(d^{-1}))$, using $\frac{1}{N} \le \frac{1}{d}$; by Theorem F.6 the second term is $O_p(d) = d^2 \cdot O_p(d^{-1})$. □

Lemma F.8 (Column-sum identity). *For every $W \in \mathbb{R}^{d\times d}$ and every i, $(N-1)\,\mathrm{mean}_{j\neq i}\, u_j^\top W v_i = (U\mathbf{1})^\top W v_i - u_i^\top W v_i$.*

Proof. $(N-1)\,\mathrm{mean}_{j\neq i}\, u_j^\top W v_i = \sum_j u_j^\top W v_i - u_i^\top W v_i$ and $\sum_j u_j = U\mathbf{1}$. □

Lemma F.9 (Two sums of squares). $\sum_i \big((U\mathbf{1})^\top W^{\mathrm{GD}} v_i\big)^2 = O_p(N^4 d)$ *and* $\sum_i \big(u_i^\top W^{\mathrm{GD}} v_i\big)^2 = O_p(N^4 d)$.

Proof. The numbers $(U\mathbf{1})^\top W^{\mathrm{GD}} v_i$ are the entries of $V^\top (W^{\mathrm{GD}})^\top U\mathbf{1}$, so by Lemma F.4(ii) and Theorem F.3(iii),(iv),

$$\big\|V^\top (W^{\mathrm{GD}})^\top U\mathbf{1}\big\|^2 \le \|V\|_{\mathrm{op}}^2 \|W^{\mathrm{GD}}\|_{\mathrm{op}}^2 \|U\mathbf{1}\|^2 = O_p(N\cdot N^2\cdot Nd) = O_p(N^4 d). \tag{34}$$

For the second sum, $u_i^\top W^{\mathrm{GD}} v_i = [U^\top W^{\mathrm{GD}} V]_{ii}$ and $\sum_i X_{ii}^2 \le \|X\|_F^2$; by Equation (11), $U^\top W^{\mathrm{GD}} V = (U^\top U)(I_N - \frac{\mathbf{1}\mathbf{1}^\top}{N})(V^\top V)$, and since the nonzero eigenvalues of $V^\top V$ and $VV^\top$ coincide, $\|V^\top V\|_F = \|VV^\top\|_F$, so by Theorem F.3(iv),(vi),

$$\sum_i \big(u_i^\top W^{\mathrm{GD}} v_i\big)^2 \le \|U^\top W^{\mathrm{GD}} V\|_F^2 \le \|U\|_{\mathrm{op}}^4 \|VV^\top\|_F^2 = O_p(N^2\cdot N^2 d) = O_p(N^4 d). \tag{35}$$

□

Lemma F.10. $\mathrm{mean}_i \big(\mathrm{mean}_{j\neq i}\, u_j^\top W^{\mathrm{GD}} v_i\big)^2 = O_p(\gamma^2 d^2)$, *hence* $|\mathrm{mean}_i\, \mathrm{mean}_{j\neq i}\, u_j^\top W^{\mathrm{GD}} v_i| = O_p(\gamma d)$.

Proof. Square Theorem F.8 with $W = W^{\mathrm{GD}}$, use $(x-y)^2 \le 2x^2 + 2y^2$, sum over i and apply Theorem F.9:

$$\sum_i \Big((N-1)\operatorname*{mean}_{j\neq i} u_j^\top W^{\mathrm{GD}} v_i\Big)^2 = O_p(N^4 d) \le O_p(N^5), \tag{36}$$

using $d \le N$. Dividing by $N(N-1)^2 \ge \frac{1}{4}N^3$ and substituting $N = \gamma d$ gives the first claim; the second follows from Cauchy–Schwarz, $|\mathrm{mean}_i\, x_i| \le (\mathrm{mean}_i\, x_i^2)^{1/2}$. □

Lemma F.11 ($\bar{h}$ for GD). $\bar{h}^{\mathrm{GD}} = d^2\big(1+O_p(\gamma/d)\big)$.

Proof. By Equation (32), Theorem F.7 and Theorem F.10, $\bar{h}^{\mathrm{GD}} = d^2(1+O_p(d^{-1})) - O_p(\gamma d) = d^2(1+O_p(d^{-1}) + O_p(\gamma/d))$, and $\gamma \ge 1$ absorbs the first error term into the second. □

We turn to $\bar{\kappa}^{\mathrm{GD}}$. By Equation (10),

$$(\bar{\kappa}^{\mathrm{GD}})^2 = \underbrace{\frac{1}{N(N-1)}\sum_{i=1}^N \sum_{j\neq i} \big(u_j^\top W^{\mathrm{GD}} v_i\big)^2}_{\text{second moment}} - \underbrace{\operatorname*{mean}_i \Big(\operatorname*{mean}_{j\neq i} u_j^\top W^{\mathrm{GD}} v_i\Big)^2}_{\text{squared mean}}. \tag{37}$$

Lemma F.12 (Expectation of the second moment)**.**

$$\begin{aligned}\mathbb{E}\sum_{i=1}^N \sum_{j\neq i} \big(u_j^\top W^{\mathrm{GD}} v_i\big)^2 &= N(N-1)d^2(N+2d+2) \\ &\quad - d^2\Big[\Big(1+\frac{1}{N}\Big)(N+d+1)^2 - 2(d+2)^2 - 2(N-1)\Big] \\ &= N(N-1)d^2(N+2d+2)\big(1+O(N^{-1})\big).\end{aligned} \tag{38}$$

Proof. Write $u_j^\top W^{\mathrm{GD}} v_i = [U^\top W^{\mathrm{GD}} V]_{ji}$; by Equation (11), $U^\top W^{\mathrm{GD}} V = (U^\top U)(I_N - \frac{\mathbf{1}\mathbf{1}^\top}{N})(V^\top V)$, i.e. $u_j^\top W^{\mathrm{GD}} v_i = \sum_{k,l}(u_j^\top u_k)(\delta_{kl} - \frac{1}{N})(v_l^\top v_i)$, with U and V independent. Two second moments, identical for u and v, are used:

$$\mathbb{E}\big[(u_k^\top u_l)(u_l^\top u_m)\big] = \begin{cases} d^2+2d, & k=l=m \\ d, & k=m\neq l \\ 0, & \text{otherwise,} \end{cases} \qquad \mathbb{E}\big[(u_i^\top u_k)(u_i^\top u_{k'})\big] = \delta_{kk'} d\big[1+(d+1)\delta_{ki}\big]. \tag{39}$$

The first of these gives $\big[\mathbb{E}(U^\top U)^2\big]_{km} = \sum_l \mathbb{E}\big[(u_k^\top u_l)(u_l^\top u_m)\big] = \delta_{km}\big[(d^2+2d)+(N-1)d\big] = \delta_{km} d(d+N+1)$, i.e. $\mathbb{E}(U^\top U)^2 = d(d+N+1)I_N$, and likewise for V. *All pairs* (i,j). Since $\|X\|_F^2 = \mathrm{tr}(X^\top X)$ and by cyclicity,

$$\sum_{i,j}\big(u_j^\top W^{\mathrm{GD}} v_i\big)^2 = \|U^\top W^{\mathrm{GD}} V\|_F^2 = \mathrm{tr}\Big[(U^\top U)^2\Big(I_N - \frac{\mathbf{1}\mathbf{1}^\top}{N}\Big)(V^\top V)^2\Big(I_N - \frac{\mathbf{1}\mathbf{1}^\top}{N}\Big)\Big], \tag{40}$$

and by independence and $\mathrm{tr}[(I_N - \frac{\mathbf{1}\mathbf{1}^\top}{N})^2] = N-1$,

$$\mathbb{E}\sum_{i,j}\big(u_j^\top W^{\mathrm{GD}} v_i\big)^2 = \big[d(d+N+1)\big]^2(N-1). \tag{41}$$

Diagonal $j=i$. Using the independence of U and V, the second display above for both factors, and then $(\delta_{kl}-\frac{1}{N})^2 = \delta_{kl}(1-\frac{2}{N}) + \frac{1}{N^2}$ together with $\sum_k[1+(d+1)\delta_{ki}] = N+d+1$ and $\sum_k[1+(d+1)\delta_{ki}]^2 = N-1+(d+2)^2$,

$$\begin{aligned} \mathbb{E}\big(u_i^\top W^{\mathrm{GD}} v_i\big)^2 &= \sum_{k,l,k',l'} \mathbb{E}\big[(u_i^\top u_k)(u_i^\top u_{k'})\big]\,\mathbb{E}\big[(v_l^\top v_i)(v_{l'}^\top v_i)\big]\Big(\delta_{kl}-\frac{1}{N}\Big)\Big(\delta_{k'l'}-\frac{1}{N}\Big) \\ &= d^2\sum_{k,l}\big[1+(d+1)\delta_{ki}\big]\big[1+(d+1)\delta_{li}\big]\Big(\delta_{kl}-\frac{1}{N}\Big)^2 \\ &= d^2\Big[\Big(1-\frac{2}{N}\Big)\big(N-1+(d+2)^2\big) + \frac{(N+d+1)^2}{N^2}\Big]. \end{aligned} \tag{42}$$

Subtracting. $\mathbb{E}\sum_i\sum_{j\neq i} = \mathbb{E}\sum_{i,j} - N\mathbb{E}(u_i^\top W^{\mathrm{GD}} v_i)^2$; rearranging with $(N+d+1)^2 - (d+2)^2 - (N-1) = (N-1)(N+2d+2)$ gives the first equality. Since $d \le N$, the discarded terms are bounded by $38N^2d^2 \le \frac{38}{N-1}\cdot N(N-1)d^2(N+2d+2)$, which is the stated relative error. □

Lemma F.13 (Variance of the second moment). *There is an absolute constant C with* $\mathrm{var}\big[\sum_i\sum_{j\neq i}(u_j^\top W^{\mathrm{GD}} v_i)^2\big] \le CN^7 d$.

Proof. By Equation (11), $\sum_i\sum_{j\neq i}(u_j^\top W^{\mathrm{GD}} v_i)^2 = \|U^\top W^{\mathrm{GD}} V\|_F^2 - \sum_i [U^\top W^{\mathrm{GD}} V]_{ii}^2$, a smooth function of the $2Nd$ independent standard Gaussian entries of (U,V). By the Gaussian Poincaré inequality the variance is at most $\mathbb{E}\|\nabla_U \cdot\|_F^2 + \mathbb{E}\|\nabla_V \cdot\|_F^2$.

The gradient of $M \mapsto \|M\|_F^2 - \sum_i M_{ii}^2$ is $2M - 2\,\mathrm{diag}(M)$, of Frobenius norm at most $2\|M\|_F$; the directional derivative of $U \mapsto (U^\top U)(I_N - \frac{\mathbf{1}\mathbf{1}^\top}{N})(V^\top V)$ along H has Frobenius norm at most $2\|H\|_F\|U\|_{\mathrm{op}}\|V\|_{\mathrm{op}}^2$. With $\|U^\top W^{\mathrm{GD}} V\|_F \le \|U\|_{\mathrm{op}}^2\|VV^\top\|_F$, the chain rule gives

$$\Big\|\nabla_U \sum_i\sum_{j\neq i}(u_j^\top W^{\mathrm{GD}} v_i)^2\Big\|_F \le 4\|U\|_{\mathrm{op}}^3\|V\|_{\mathrm{op}}^2\|VV^\top\|_F. \tag{43}$$

Squaring, taking expectations and applying Hölder together with Theorem F.3(iv),(vi) gives $\mathbb{E}\|\nabla_U\cdot\|_F^2 \le 16(\mathbb{E}\|U\|_{\mathrm{op}}^{18})^{1/3}(\mathbb{E}\|V\|_{\mathrm{op}}^{12})^{1/3}(\mathbb{E}\|VV^\top\|_F^6)^{1/3} \le 16(C_9C_6^2)^{1/3}N^7 d$. The expression is invariant under transposition combined with swapping U and V, so the same bound holds for ∇_V; adding the two gives the claim with $C = 32(C_9C_6^2)^{1/3}$. □

Lemma F.14 (Second moment). *The second moment of the off-diagonal GD scores satisfies*

$$\frac{1}{N(N-1)}\sum_{i=1}^N\sum_{j\neq i}\big(u_j^\top W^{\mathrm{GD}} v_i\big)^2 = d^3(2+\gamma)\big(1+O_p(\sqrt{\gamma}/d)\big).$$

Proof. By Equation (38) and $N(N-1)d^2(N+2d+2) \geq N^2(N-1)d^2 \geq \frac{1}{2}N^3d^2$ for $N \geq 2$, the expectation is at least $\frac{1}{4}N^3d^2$ for N large. Combined with Theorem F.13, the relative standard deviation is at most $\sqrt{CN^7d}/(\frac{1}{4}N^3d^2) = 4\sqrt{C}\sqrt{N/d^3} = 4\sqrt{C}\sqrt{\gamma}/d$, so Chebyshev gives

$$\sum_i \sum_{j\neq i} \left(u_j^\top W^{\mathrm{GD}} v_i\right)^2 = \Big(\mathbb{E}\sum_i \sum_{j\neq i} \left(u_j^\top W^{\mathrm{GD}} v_i\right)^2\Big)\big(1 + O_p(\sqrt{\gamma}/d)\big). \tag{44}$$

Since $N^{-1} = \gamma^{-1}d^{-1} \leq \sqrt{\gamma}/d$, the relative error $O(N^{-1})$ of Theorem F.12 is absorbed; dividing by $N(N-1)$ leaves $d^2(N+2d+2) = d^3(\gamma + 2 + \frac{2}{d}) = d^3(2+\gamma)(1 + \frac{2}{d(2+\gamma)})$ and $\frac{2}{d(2+\gamma)} \leq \frac{1}{d} \leq \sqrt{\gamma}/d$. □

Lemma F.15 ($\bar{\kappa}$ for GD). $(\bar{\kappa}^{\mathrm{GD}})^2 = d^3(2+\gamma)\big(1 + O_p(\gamma/d)\big)$.

Proof. By Theorem F.10 the squared mean in Equation (37) is $O_p(\gamma^2 d^2) = d^3(2+\gamma)\cdot O_p\big(\frac{\gamma^2}{d(2+\gamma)}\big) = d^3(2+\gamma)\cdot O_p(\gamma/d)$. Subtracting it from Theorem F.14 and using $\sqrt{\gamma} \leq \gamma$ gives the claim. □

Lemma F.16 (ρ for GD). *The margin-to-spread ratio of the GD iterate is*

$$\rho^{\mathrm{GD}} = \sqrt{\frac{d}{2+\gamma}}\big(1 + O_p(\gamma/d)\big) = \sqrt{\frac{d}{\gamma}}\Big(1 - \frac{1}{\gamma} + O(\gamma^{-2}) + o_p(1)\Big).$$

Both $\bar{h}^{\mathrm{GD}}$ and $\bar{\kappa}^{\mathrm{GD}}$ are a power of d times a function of γ alone, so ρ^{GD} depends on d only through the prefactor $\sqrt{d}$.

Proof. By Theorem F.11, Theorem F.15 and $\sqrt{1 + O_p(\gamma/d)} = 1 + O_p(\gamma/d)$,

$$\rho^{\mathrm{GD}} = \frac{\bar{h}^{\mathrm{GD}}}{\bar{\kappa}^{\mathrm{GD}}} = \frac{d^2\big(1 + O_p(\gamma/d)\big)}{d^{3/2}\sqrt{2+\gamma}\big(1 + O_p(\gamma/d)\big)} = \sqrt{\frac{d}{2+\gamma}}\big(1 + O_p(\gamma/d)\big). \tag{45}$$

With γ fixed and $d \to \infty$, $O_p(\gamma/d) = o_p(1)$; moreover $\sqrt{\gamma/(2+\gamma)} = (1 + 2/\gamma)^{-1/2} = 1 - \gamma^{-1} + O(\gamma^{-2})$. □

F.4 The Muon step

Throughout this subsection $\gamma \geq 16$. By Equation (10) with $W = W^{\mathrm{Muon}}$ the same two decompositions hold as in Equation (32) and Equation (37), with W^{GD} replaced by W^{Muon}. We use throughout that W^{Muon} is orthogonal (Theorem F.2).

Lemma F.17 (Diagonal-sum identity). *The diagonal scores of W^{Muon} sum to*

$$\sum_{i=1}^N u_i^\top W^{\mathrm{Muon}} v_i = \|W^{\mathrm{GD}}\|_* + \frac{1}{N}(U\mathbf{1})^\top W^{\mathrm{Muon}}(V\mathbf{1}).$$

Proof. As in Theorem F.5, $\sum_i u_i^\top W^{\mathrm{Muon}} v_i = \langle W^{\mathrm{Muon}}, UV^\top\rangle = \langle W^{\mathrm{Muon}}, W^{\mathrm{GD}}\rangle + \frac{1}{N}(U\mathbf{1})^\top W^{\mathrm{Muon}}(V\mathbf{1})$. Writing the compact SVD $W^{\mathrm{GD}} = P\Sigma Q^\top$, so that $W^{\mathrm{Muon}} = PQ^\top$, and using $P^\top P = Q^\top Q = I$,

$$\langle W^{\mathrm{Muon}}, W^{\mathrm{GD}}\rangle = \operatorname{tr}\left(QP^\top P\Sigma Q^\top\right) = \operatorname{tr}\left(\Sigma Q^\top Q\right) = \operatorname{tr}\Sigma = \|W^{\mathrm{GD}}\|_*. \tag{46}$$

□

Lemma F.18 (Rank-one cross term). $\frac{1}{N^2}|(U\mathbf{1})^\top W^{\mathrm{Muon}}(V\mathbf{1})| = O_p(d/N)$.

Proof. By $\|W^{\mathrm{Muon}}\|_{\mathrm{op}} \leq 1$ and Theorem F.3(iii), $|(U\mathbf{1})^\top W^{\mathrm{Muon}}(V\mathbf{1})|/N^2 \leq \|U\mathbf{1}\|\|V\mathbf{1}\|/N^2 = O_p(Nd)/N^2 = O_p(d/N)$. □

Lemma F.19. $\operatorname{mean}_i u_i^\top W^{\mathrm{Muon}} v_i = \frac{\|W^{\mathrm{GD}}\|_*}{N}\big(1 + O_p((Nd)^{-1/2})\big)$.

Proof. Divide Theorem F.17 by N and use Theorem F.18. By Lemma F.4(iii), with probability tending to one $\|W^{\mathrm{GD}}\|_*/N \geq d\sqrt{Nd}/(8N) = d^{3/2}/(8\sqrt{N})$, so the error term satisfies $O_p(d/N) = \frac{\|W^{\mathrm{GD}}\|_*}{N} O_p\big(\frac{d}{N}\frac{8\sqrt{N}}{d^{3/2}}\big) = \frac{\|W^{\mathrm{GD}}\|_*}{N} O_p((Nd)^{-1/2})$. □

Lemma F.20 (Two sums of squares). $\sum_i \big((U\mathbf{1})^\top W^{\mathrm{Muon}} v_i\big)^2 = O_p(N^2 d)$ *and* $\sum_i \big(u_i^\top W^{\mathrm{Muon}} v_i\big)^2 = O_p(Nd^2)$.

Proof. For the first, W^{Muon} is orthogonal, so $\|(W^{\mathrm{Muon}})^\top U\mathbf{1}\| = \|U\mathbf{1}\|$, and $\sum_i v_i v_i^\top = VV^\top$, whence by Theorem F.3(iii),(iv)

$$\sum_i \big((U\mathbf{1})^\top W^{\mathrm{Muon}} v_i\big)^2 = (U\mathbf{1})^\top W^{\mathrm{Muon}} VV^\top (W^{\mathrm{Muon}})^\top (U\mathbf{1}) \le \|V\|_{\mathrm{op}}^2 \|U\mathbf{1}\|^2 = O_p(N \cdot Nd). \tag{47}$$

For the second, Cauchy–Schwarz and $\|W^{\mathrm{Muon}} v_i\| = \|v_i\|$ give $(u_i^\top W^{\mathrm{Muon}} v_i)^2 \le \|u_i\|^2 \|v_i\|^2$, and $x^2 \le 2(x-d)^2 + 2d^2$ with Theorem F.3(ii) gives $\sum_i \|u_i\|^2\|v_i\|^2 \le 2\sum_i (\|u_i\|\|v_i\| - d)^2 + 2Nd^2 = N \cdot O_p(d) + 2Nd^2 = O_p(Nd^2)$. □

Lemma F.21. $\mathrm{mean}_i \big(\mathrm{mean}_{j\neq i}\, u_j^\top W^{\mathrm{Muon}} v_i\big)^2 = O_p(d/N)$, *hence* $|\mathrm{mean}_i\, \mathrm{mean}_{j\neq i}\, u_j^\top W^{\mathrm{Muon}} v_i| = O_p(\sqrt{d/N})$.

Proof. Square Theorem F.8 with $W = W^{\mathrm{Muon}}$, use $(x-y)^2 \le 2x^2 + 2y^2$, sum over i and apply Theorem F.20:

$$\sum_i \Big((N-1) \operatorname*{mean}_{j\neq i} u_j^\top W^{\mathrm{Muon}} v_i\Big)^2 \le 2O_p(N^2 d) + 2O_p(Nd^2) = O_p(N^2 d), \tag{48}$$

using $d \le N$. Dividing by $N(N-1)^2 \ge \frac{1}{4}N^3$ gives the first claim, and Cauchy–Schwarz the second. □

Lemma F.22 ($\bar{h}$ for Muon). *The mean margin of the Muon iterate is*

$$\bar{h}^{\mathrm{Muon}} = \frac{\|W^{\mathrm{GD}}\|_*}{N}\big(1 + O_p(d^{-1})\big) = \frac{d\,\Phi(\gamma)}{\sqrt{\gamma}}\big(1 + o_p(1)\big).$$

Proof. By Theorem F.19 and Theorem F.21, $\bar{h}^{\mathrm{Muon}} = \frac{\|W^{\mathrm{GD}}\|_*}{N}(1 + O_p((Nd)^{-1/2})) - O_p(\sqrt{d/N})$. By Lemma F.4(iii), with probability tending to one $\frac{\|W^{\mathrm{GD}}\|_*}{N} \ge \frac{d^{3/2}}{8\sqrt{N}}$, so

$$O_p\big(\sqrt{d/N}\big) = \frac{\|W^{\mathrm{GD}}\|_*}{N} \cdot O_p\Big(\sqrt{\frac{d}{N}} \cdot \frac{8\sqrt{N}}{d^{3/2}}\Big) = \frac{\|W^{\mathrm{GD}}\|_*}{N} \cdot O_p(d^{-1}), \tag{49}$$

and $(Nd)^{-1/2} \le d^{-1}$, which gives the first equality. For the second, Lemma F.4(iv) gives $\frac{\|W^{\mathrm{GD}}\|_*}{Nd} = \frac{\Phi(\gamma)}{\sqrt{\gamma}} + o_p(1)$, and $\frac{\Phi(\gamma)}{\sqrt{\gamma}} \ge \frac{1}{8\sqrt{\gamma}} > 0$ by Lemma F.4(iii), so the additive $o_p(1)$ may be written multiplicatively. □

We now turn to $\bar{\kappa}^{\mathrm{Muon}}$. The orthogonality of W^{Muon} turns the second moment into three terms.

Lemma F.23 (Three-term identity). *Almost surely*

$$\begin{aligned}\frac{1}{N(N-1)} \sum_{i=1}^N \sum_{j\neq i} \big(u_j^\top W^{\mathrm{Muon}} v_i\big)^2 &= \frac{N}{N-1} \operatorname*{mean}_i \|v_i\|^2 - \frac{\mathrm{mean}_i \big(u_i^\top W^{\mathrm{Muon}} v_i\big)^2}{N-1} \\ &\quad + \frac{\mathrm{tr}\big[(UU^\top - NI_d) W^{\mathrm{Muon}} VV^\top (W^{\mathrm{Muon}})^\top\big]}{N(N-1)}.\end{aligned} \tag{50}$$

Proof. Since $(U^\top x)_j = u_j^\top x$ and $\|W^{\mathrm{Muon}} v_i\| = \|v_i\|$,

$$\begin{aligned}\sum_{j\neq i} \big(u_j^\top W^{\mathrm{Muon}} v_i\big)^2 &= \|U^\top W^{\mathrm{Muon}} v_i\|^2 - \big(u_i^\top W^{\mathrm{Muon}} v_i\big)^2, \\ \|U^\top W^{\mathrm{Muon}} v_i\|^2 &= N\|v_i\|^2 + \big(W^{\mathrm{Muon}} v_i\big)^\top \big(UU^\top - NI_d\big)\big(W^{\mathrm{Muon}} v_i\big).\end{aligned} \tag{51}$$

Summing over i and using $x^\top Z x = \mathrm{tr}(Zxx^\top)$ with $\sum_i v_i v_i^\top = VV^\top$ turns the middle term into the stated trace; dividing by $N(N-1)$ and using $\sum_i \|v_i\|^2 = N\,\mathrm{mean}_i \|v_i\|^2$ gives the claim. □

Lemma F.24 (Norm term). $\frac{N}{N-1} \mathrm{mean}_i \|v_i\|^2 = d\big(1 + O(N^{-1}) + O_p((Nd)^{-1/2})\big)$.

Proof. Theorem F.3(i) gives $\mathrm{mean}_i \|v_i\|^2 = d(1 + O_p((Nd)^{-1/2}))$, and $\frac{N}{N-1} = 1 + \frac{1}{N-1} = 1 + O(N^{-1})$. □

The trace term requires two auxiliary facts.

Lemma F.25 (Decomposition of $UU^\top$). *Almost surely $VV^\top - \frac{1}{N}(V\mathbf{1})(V\mathbf{1})^\top$ is invertible and there is $\Gamma \in \mathbb{R}^{(N-1-d)\times d}$ with*

$$UU^\top - \frac{1}{N}(U\mathbf{1})(U\mathbf{1})^\top = W^{\mathrm{GD}}\Big(VV^\top - \frac{1}{N}(V\mathbf{1})(V\mathbf{1})^\top\Big)^{-1}(W^{\mathrm{GD}})^\top + \Gamma^\top\Gamma, \tag{52}$$

where, conditionally on V, the entries of Γ are i.i.d. $\mathcal{N}(0,1)$ and independent of W^{GD}. Consequently, for every symmetric $Y \in \mathbb{R}^{d\times d}$ that is a function of (V, W^{GD}),

$$\begin{aligned}\mathbb{E}\big[\operatorname{tr}(\Gamma^\top\Gamma Y) \mid V, W^{\mathrm{GD}}\big] &= (N-1-d)\operatorname{tr} Y,\\ \operatorname{var}\big[\operatorname{tr}(\Gamma^\top\Gamma Y) \mid V, W^{\mathrm{GD}}\big] &= 2(N-1-d)\|Y\|_F^2.\end{aligned} \tag{53}$$

Proof. Invertibility. $VV^\top - \frac{1}{N}(V\mathbf{1})(V\mathbf{1})^\top = V(I_N - \frac{\mathbf{1}\mathbf{1}^\top}{N})V^\top \succeq 0$ and its determinant is a polynomial in V which is nonzero at $V = [\,I_d \;\; -\sum_{k\le d} e_k \;\; 0\,]$, where it equals $\det(I_d + (\sum_k e_k)(\sum_k e_k)^\top) = 1 + d \neq 0$; so it is nonzero almost surely.

Basis. If $c_0\mathbf{1}^\top + c^\top V = 0$ for some $c_0 \in \mathbb{R}$ and $c \in \mathbb{R}^d$, then multiplying on the right by $(I_N - \frac{\mathbf{1}\mathbf{1}^\top}{N})V^\top c$, and using $\mathbf{1}^\top(I_N - \frac{\mathbf{1}\mathbf{1}^\top}{N}) = 0$, gives $c^\top(VV^\top - \frac{1}{N}(V\mathbf{1})(V\mathbf{1})^\top)c = 0$, so $c = 0$ and then $c_0 = 0$. Hence $\{x \in \mathbb{R}^N : \mathbf{1}^\top x = 0, Vx = 0\}$ has dimension $N-1-d$ almost surely; stack an orthonormal basis of it as the columns of $R \in \mathbb{R}^{N\times(N-1-d)}$, so that $R^\top R = I$, $\mathbf{1}^\top R = 0$ and $VR = 0$, and set $\Gamma := R^\top U^\top$.

Projection identity. Because $\mathbb{R}^N = \operatorname{span}\mathbf{1} \oplus \operatorname{col}\big((I_N - \frac{\mathbf{1}\mathbf{1}^\top}{N})V^\top\big) \oplus \operatorname{col} R$ with the three subspaces pairwise orthogonal and of dimensions $1 + d + (N-1-d) = N$,

$$I_N - \frac{\mathbf{1}\mathbf{1}^\top}{N} = \Big(I_N - \frac{\mathbf{1}\mathbf{1}^\top}{N}\Big)V^\top\Big(VV^\top - \frac{1}{N}(V\mathbf{1})(V\mathbf{1})^\top\Big)^{-1}V\Big(I_N - \frac{\mathbf{1}\mathbf{1}^\top}{N}\Big) + RR^\top. \tag{54}$$

Both sides are symmetric, so it suffices to check them on the three subspaces. Write $P := I_N - \frac{\mathbf{1}\mathbf{1}^\top}{N}$ and $G := VV^\top - \frac{1}{N}(V\mathbf{1})(V\mathbf{1})^\top$, and use $P^2 = P, VPV^\top = G$ and $PR = R - \frac{1}{N}\mathbf{1}(\mathbf{1}^\top R) = R$. For $x = \mathbf{1}$: $Px = 0, VPx = 0$ and $R^\top x = 0$, so both sides vanish. For $x = PV^\top y$: $Px = x$, while $VPx = VPV^\top y = Gy$, so the first term of the right-hand side is $PV^\top G^{-1}Gy = x$, and $R^\top x = (VR)^\top y = 0$, so the second vanishes. For $x = Rz$: $Px = Rz = x$, while $VPx = VRz = 0$ kills the first term, and $RR^\top Rz = Rz = x$. In all three cases both sides agree. Multiplying Equation (54) on the left by U and on the right by $U^\top$, and using $U(I_N - \frac{\mathbf{1}\mathbf{1}^\top}{N})U^\top = UU^\top - \frac{1}{N}(U\mathbf{1})(U\mathbf{1})^\top$ and $U(I_N - \frac{\mathbf{1}\mathbf{1}^\top}{N})V^\top = W^{\mathrm{GD}}$, gives the decomposition with $URR^\top U^\top = \Gamma^\top\Gamma$.

Distribution. Conditionally on V, the rows of U are i.i.d. $\mathcal{N}(0, I_N)$; the a-th column of Γ is $R^\top U^\top e_a \sim \mathcal{N}(0, R^\top R) = \mathcal{N}(0, I_{N-1-d})$ with independent columns, so the entries of Γ are i.i.d. standard Gaussian. The a-th row of W^{GD} is $\big(V(I_N - \frac{\mathbf{1}\mathbf{1}^\top}{N})U^\top e_a\big)^\top$ and $\operatorname{Cov}\big(V(I_N - \frac{\mathbf{1}\mathbf{1}^\top}{N})U^\top e_a, R^\top U^\top e_a\big) = V(I_N - \frac{\mathbf{1}\mathbf{1}^\top}{N})R = 0$; being jointly Gaussian conditionally on V, Γ and W^{GD} are conditionally independent.

Moments. The rows of Γ are i.i.d. $\mathcal{N}(0, I_d)$ and $\operatorname{tr}(\Gamma^\top\Gamma Y) = \sum_m (\Gamma^\top e_m)^\top Y(\Gamma^\top e_m)$; for $g \sim \mathcal{N}(0, I_d)$ and fixed symmetric Y, $\mathbb{E}[g^\top Y g] = \operatorname{tr} Y$ and $\operatorname{var}[g^\top Y g] = 2\|Y\|_F^2$; summing over the $N-1-d$ rows gives the claim. □

Lemma F.26 (Trace of a commutator). *Almost surely, with $S := \frac{1}{N}\big(VV^\top - \frac{1}{N}(V\mathbf{1})(V\mathbf{1})^\top\big)$,*

$$\Big|\operatorname{tr}\Big[W^{\mathrm{GD}}(NS)^{-1}(W^{\mathrm{GD}})^\top W^{\mathrm{Muon}}(NS)(W^{\mathrm{Muon}})^\top\Big] - \|W^{\mathrm{GD}}\|_F^2\Big| \le \frac{2\|S - I_d\|_{\mathrm{op}}^2}{\lambda_{\min}(S)}\|W^{\mathrm{GD}}\|_F^2. \tag{55}$$

Proof. A matrix inequality. For symmetric $S \succ 0$ and symmetric T,

$$\big|\operatorname{tr}[S^{-1}TST] - \|T\|_F^2\big| \le \frac{2\|S - I_d\|_{\mathrm{op}}^2\|T\|_F^2}{\lambda_{\min}(S)}. \tag{56}$$

Indeed, cyclicity gives $\operatorname{tr}[S^{-1}STT] = \|T\|_F^2$ and $\operatorname{tr}[(TS - ST)T] = 0$, so with $S^{-1} = I_d - S^{-1}(S - I_d)$,

$$\operatorname{tr}[S^{-1}TST] - \|T\|_F^2 = \operatorname{tr}\big[S^{-1}(TS - ST)T\big] = -\operatorname{tr}\big[S^{-1}(S - I_d)(TS - ST)T\big], \tag{57}$$

and $|\operatorname{tr}(ABC)| \le \|A\|_{\mathrm{op}}\|B\|_F\|C\|_F$, $\|S^{-1}(S - I_d)\|_{\mathrm{op}} \le \|S - I_d\|_{\mathrm{op}}/\lambda_{\min}(S)$ and $\|TS - ST\|_F \le 2\|S - I_d\|_{\mathrm{op}}\|T\|_F$ give Equation (56).

Polar decomposition. With $W^{\mathrm{GD}} = P\Sigma Q^\top$ and $W^{\mathrm{Muon}} = PQ^\top$ we have $(W^{\mathrm{GD}})^\top W^{\mathrm{Muon}} = Q\Sigma Q^\top = (W^{\mathrm{Muon}})^\top W^{\mathrm{GD}}$ and $\|Q\Sigma Q^\top\|_F = \|\Sigma\|_F = \|W^{\mathrm{GD}}\|_F$.

Substitution. Moving the leading W^{GD} to the end by cyclicity,

$$\operatorname{tr}\Big[W^{\mathrm{GD}}(NS)^{-1}(W^{\mathrm{GD}})^\top W^{\mathrm{Muon}}(NS)(W^{\mathrm{Muon}})^\top\Big] = \operatorname{tr}\Big[(NS)^{-1}\big(Q\Sigma Q^\top\big)(NS)\big(Q\Sigma Q^\top\big)\Big], \tag{58}$$

and applying Equation (56) with $T = Q\Sigma Q^\top$ (the factors N cancel) gives the claim. □

Lemma F.27 (Trace term). *If $\gamma \geq 16$, then with probability tending to one*

$$\frac{1}{N(N-1)d}\Big|\operatorname{tr}\Big[\big(UU^\top - NI_d\big)W^{\mathrm{Muon}}VV^\top(W^{\mathrm{Muon}})^\top\Big]\Big| \leq 145\gamma^{-2}. \tag{59}$$

The assumption $\gamma \geq 16$ is tailored to this estimate: through Theorem F.3(v) it controls how far $\frac{1}{N}VV^\top$ is from the identity, so that the fluctuation of $UU^\top$ around NI_d, seen through the Muon iterate, contributes only $O(\gamma^{-2})$ to the second moment.

Proof. Write $B := W^{\mathrm{Muon}}\big(VV^\top - \frac{1}{N}(V\mathbf{1})(V\mathbf{1})^\top\big)(W^{\mathrm{Muon}})^\top$. By Theorem F.2 orthogonal conjugation preserves $\|\cdot\|_F$, $\|\cdot\|_{\mathrm{op}}$ and the trace; moreover $VV^\top - \frac{1}{N}(V\mathbf{1})(V\mathbf{1})^\top = V(I_N - \frac{\mathbf{1}\mathbf{1}^\top}{N})V^\top$ satisfies $0 \preceq VV^\top - \frac{1}{N}(V\mathbf{1})(V\mathbf{1})^\top \preceq VV^\top$. Hence, by Theorem F.3(iii),(iv),(vi)

$$\begin{aligned}
\|B\|_F &\leq \|VV^\top\|_F + \frac{\|V\mathbf{1}\|^2}{N} = O_p(N\sqrt{d}),\\
\|B\|_{\mathrm{op}} &\leq \|V\|_{\mathrm{op}}^2 = O_p(N),\\
\operatorname{tr} B &= \sum_k \|v_k\|^2 - \frac{\|V\mathbf{1}\|^2}{N} = O_p(Nd).
\end{aligned} \tag{60}$$

Rank-one replacement. $VV^\top - (VV^\top - \frac{1}{N}(V\mathbf{1})(V\mathbf{1})^\top) = \frac{1}{N}(V\mathbf{1})(V\mathbf{1})^\top$, so by Theorem F.3(iii),(v),

$$\begin{aligned}
&\Big|\operatorname{tr}\big[(UU^\top - NI_d)W^{\mathrm{Muon}}VV^\top(W^{\mathrm{Muon}})^\top\big] - \operatorname{tr}\big[(UU^\top - NI_d)B\big]\Big|\\
&\qquad \leq \frac{1}{N}\|UU^\top - NI_d\|_{\mathrm{op}}\|V\mathbf{1}\|^2 = N(N-1)d\cdot O_p\big(\sqrt{d}N^{-3/2}\big).
\end{aligned} \tag{61}$$

Splitting. By Theorem F.25 and $d + (N-1-d) + 1 = N$,

$$UU^\top - NI_d = \Big[W^{\mathrm{GD}}\big(NS\big)^{-1}(W^{\mathrm{GD}})^\top - dI_d\Big] + \Big[\Gamma^\top\Gamma - (N-1-d)I_d\Big] + \Big[\frac{1}{N}(U\mathbf{1})(U\mathbf{1})^\top - I_d\Big], \tag{62}$$

and we bound the trace of each bracket against B separately.

First bracket. Since $\operatorname{tr} B$ appears, write

$$\begin{aligned}
\operatorname{tr}\big[(W^{\mathrm{GD}}(NS)^{-1}(W^{\mathrm{GD}})^\top - dI_d)B\big] = \big(&\operatorname{tr}[W^{\mathrm{GD}}(NS)^{-1}(W^{\mathrm{GD}})^\top B] - \|W^{\mathrm{GD}}\|_F^2\big)\\
&+ \big(\|W^{\mathrm{GD}}\|_F^2 - d\operatorname{tr} B\big).
\end{aligned} \tag{63}$$

By Theorem F.26, Theorem F.3(v) ($\gamma \geq 16$) and Lemma F.4(i), with probability tending to one the first difference is at most

$$\frac{2(3\gamma^{-1/2})^2}{1/8}\|W^{\mathrm{GD}}\|_F^2 = 144\gamma^{-1}(N-1)d^2\big(1 + O_p(d^{-1})\big) = N(N-1)d\cdot 144\gamma^{-2}\big(1 + O_p(d^{-1})\big), \tag{64}$$

while the second is $O_p(\sqrt{N}d^{3/2} + Nd) = N(N-1)d\cdot O_p(\sqrt{d}N^{-3/2} + N^{-1})$ by Lemma F.4(i).

Second bracket. B is symmetric and a function of (V, W^{GD}), so Theorem F.25 gives $\mathbb{E}\big[\operatorname{tr}[(\Gamma^\top\Gamma - (N-1-d)I_d)B] \mid V, W^{\mathrm{GD}}\big] = 0$ and, by Equation (60), conditional variance $2(N-1-d)\|B\|_F^2 = O(N)\cdot O_p(N^2d) = O_p(N^3d)$. The conditional Chebyshev inequality then gives $\operatorname{tr}[(\Gamma^\top\Gamma - (N-1-d)I_d)B] = O_p(N^{3/2}d^{1/2}) = N(N-1)d\cdot O_p((Nd)^{-1/2})$.

Third bracket. By $\operatorname{tr}(aa^\top Y) = a^\top Ya \le \|a\|^2\|Y\|_{\mathrm{op}}$, Equation (60) and Theorem F.3(iii),

$$\begin{aligned}\left|\operatorname{tr}\left[\left(\frac{1}{N}(U\mathbf{1})(U\mathbf{1})^\top - I_d\right)B\right]\right| &\le \frac{\|U\mathbf{1}\|^2}{N}\|B\|_{\mathrm{op}} + |\operatorname{tr} B| \\ &= O_p(d)\cdot O_p(N) + O_p(Nd) = N(N-1)d\cdot O_p(N^{-1}).\end{aligned} \tag{65}$$

Total. Adding the four contributions and dividing by $N(N-1)d$, with probability tending to one the left-hand side is at most $144\gamma^{-2}(1+O_p(d^{-1})) + o_p(1) = 144\gamma^{-2} + o_p(1)$, which is at most $145\gamma^{-2}$ since γ is fixed. □

The remaining term of Theorem F.23, the diagonal one, needs the following sandwich.

Lemma F.28 (Nuclear-norm sandwich). *For all $a, b \in \mathbb{R}^d$,*

$$a^\top \operatorname{Orth}\left(W^{\mathrm{GD}} - ab^\top\right)b \;\le\; \|W^{\mathrm{GD}}\|_* - \|W^{\mathrm{GD}} - ab^\top\|_* \;\le\; a^\top W^{\mathrm{Muon}} b. \tag{66}$$

Proof. For all X, Y, $\langle X, \operatorname{Orth}(Y)\rangle \le \|X\|_*\|\operatorname{Orth}(Y)\|_{\mathrm{op}} \le \|X\|_*$ by Equation (9), and writing the compact SVD $Y = P\Sigma Q^\top$ gives $\langle Y, \operatorname{Orth}(Y)\rangle = \operatorname{tr}(Q\Sigma P^\top PQ^\top) = \|Y\|_*$; also $\langle ab^\top, Z\rangle = a^\top Zb$. Hence

$$\begin{aligned}a^\top \operatorname{Orth}\left(W^{\mathrm{GD}} - ab^\top\right)b &= \left\langle W^{\mathrm{GD}}, \operatorname{Orth}(W^{\mathrm{GD}} - ab^\top)\right\rangle - \left\langle W^{\mathrm{GD}} - ab^\top, \operatorname{Orth}(W^{\mathrm{GD}} - ab^\top)\right\rangle \\ &\le \|W^{\mathrm{GD}}\|_* - \|W^{\mathrm{GD}} - ab^\top\|_*,\end{aligned}$$

and symmetrically

$$\begin{aligned}\|W^{\mathrm{GD}}\|_* - \|W^{\mathrm{GD}} - ab^\top\|_* &= \langle W^{\mathrm{GD}}, W^{\mathrm{Muon}}\rangle - \|W^{\mathrm{GD}} - ab^\top\|_* \\ &\le \langle W^{\mathrm{GD}}, W^{\mathrm{Muon}}\rangle - \langle W^{\mathrm{GD}} - ab^\top, W^{\mathrm{Muon}}\rangle = a^\top W^{\mathrm{Muon}} b.\end{aligned}$$

□

Lemma F.29 (Self term). *If $\gamma \ge 16$, then with probability tending to one*

$$0 \le \frac{1}{(N-1)d}\operatorname*{mean}_i \left(u_i^\top W^{\mathrm{Muon}} v_i\right)^2 \le 15\gamma^{-3/2}. \tag{67}$$

This bounds the diagonal term of Theorem F.23; it is what forces the $\gamma^{-3/2}$ error in $\bar\kappa^{\mathrm{Muon}}$, and the only place where Theorem F.28 is used.

Proof. Nonnegativity is clear.

Step 1 (reduction to a first absolute moment). By $\|W^{\mathrm{Muon}}\|_{\mathrm{op}} \le 1$, $|u_i^\top W^{\mathrm{Muon}} v_i| \le \|u_i\|\|v_i\|$, so

$$\begin{aligned}\operatorname*{mean}_i \left(u_i^\top W^{\mathrm{Muon}} v_i\right)^2 \le d \operatorname*{mean}_i \left|u_i^\top W^{\mathrm{Muon}} v_i\right| \\ + \left(\operatorname*{mean}_i \left(u_i^\top W^{\mathrm{Muon}} v_i\right)^2\right)^{1/2}\left(\operatorname*{mean}_i \left(\|u_i\|\|v_i\| - d\right)^2\right)^{1/2}.\end{aligned} \tag{68}$$

For $x, c_0, c_1 \ge 0$, $x^2 \le c_0 + c_1 x$ forces $x \le \frac{c_1 + \sqrt{c_1^2 + 4c_0}}{2} \le c_1 + \sqrt{c_0}$ and hence $x^2 \le 2c_1^2 + 2c_0$; taking $x = (\operatorname{mean}_i (u_i^\top W^{\mathrm{Muon}} v_i)^2)^{1/2}$, $c_0 = d \operatorname{mean}_i |u_i^\top W^{\mathrm{Muon}} v_i|$ and $c_1 = (\operatorname{mean}_i (\|u_i\|\|v_i\| - d)^2)^{1/2}$, and using Theorem F.3(ii),

$$\operatorname*{mean}_i \left(u_i^\top W^{\mathrm{Muon}} v_i\right)^2 \le O_p(d) + 2d \operatorname*{mean}_i \left|u_i^\top W^{\mathrm{Muon}} v_i\right|. \tag{69}$$

Step 2 (a signed mean). By Theorem F.17 and Theorem F.18, $\operatorname{mean}_i u_i^\top W^{\mathrm{Muon}} v_i = \frac{\|W^{\mathrm{GD}}\|_*}{N} + O_p(d/N)$. For real $x \ge \ell$ one has

$$|x| \le x + 2|\ell| : \tag{70}$$

this is clear for $x \ge 0$, and for $x < 0$ it follows from $|x| = -x \le -\ell \le |\ell| \le x + 2|\ell|$, the last step because $x + 2|\ell| \ge \ell + 2|\ell| \ge |\ell|$.

Step 3 (conditioning). Fix i and set

$$a := \frac{u_i - \frac{1}{N}U\mathbf{1}}{\sqrt{1 - 1/N}}, \qquad b := \frac{v_i - \frac{1}{N}V\mathbf{1}}{\sqrt{1 - 1/N}}.$$

Since $\mathbf{1}^\top(e_i - \frac{1}{N}\mathbf{1}) = 0$ and $\|e_i - \frac{1}{N}\mathbf{1}\|^2 = 1 - \frac{1}{N}$, Equation (11) gives

$$W^{\mathrm{GD}} - ab^\top = U\Big[\Big(I_N - \frac{\mathbf{1}\mathbf{1}^\top}{N}\Big) - \frac{(e_i - \frac{1}{N}\mathbf{1})(e_i - \frac{1}{N}\mathbf{1})^\top}{1 - 1/N}\Big]V^\top, \tag{71}$$

where the bracket, call it M, is the orthogonal projection onto $\{x \in \mathbb{R}^N : \mathbf{1}^\top x = 0,\ (e_i - \frac{1}{N}\mathbf{1})^\top x = 0\}$. Indeed $M = P - \Pi$ with $P := I_N - \frac{\mathbf{1}\mathbf{1}^\top}{N}$ and $\Pi := (e_i - \frac{1}{N}\mathbf{1})(e_i - \frac{1}{N}\mathbf{1})^\top / \|e_i - \frac{1}{N}\mathbf{1}\|^2$ the orthogonal projections onto $\mathbf{1}^\perp$ and onto $\mathrm{span}(e_i - \frac{1}{N}\mathbf{1})$; since $\mathbf{1}^\top(e_i - \frac{1}{N}\mathbf{1}) = 0$ the range of Π lies in that of P, so $P\Pi = \Pi P = \Pi$ and hence $M^\top = M$, $M^2 = M$ and $M(e_i - \frac{1}{N}\mathbf{1}) = 0$. Consequently a and b are standard Gaussian and independent of $W^{\mathrm{GD}} - ab^\top$. For a: its k-th coordinate is $(e_i - \frac{1}{N}\mathbf{1})^\top U^\top e_k / \sqrt{1 - 1/N}$ and the k-th row of UM is $MU^\top e_k$; conditionally on V the vectors $U^\top e_k$ are i.i.d. $\mathcal{N}(0, I_N)$, and

$$\mathrm{Cov}\Big((e_i - \frac{1}{N}\mathbf{1})^\top U^\top e_k,\ MU^\top e_k\Big) = M(e_i - \frac{1}{N}\mathbf{1}) = 0, \tag{72}$$

so, being jointly Gaussian, a is independent of UM, and $\|(e_i - \frac{1}{N}\mathbf{1})/\sqrt{1 - 1/N}\| = 1$ gives $a \sim \mathcal{N}(0, I_d)$. The same argument applied to the rows of V gives $b \sim \mathcal{N}(0, I_d)$ independent of VM, and $U \perp V$ makes a, b independent of each other; since $W^{\mathrm{GD}} - ab^\top = (UM)(VM)^\top$ by idempotence of M, the pair (a, b) is independent of $W^{\mathrm{GD}} - ab^\top$.

Step 4 (two mean-square bounds). For independent $a, b \sim \mathcal{N}(0, I_d)$ and Z independent of them, $\mathbb{E}[(a^\top Z b)^2 \mid Z] = \|Z\|_F^2$; taking $Z = \mathrm{Orth}(W^{\mathrm{GD}} - ab^\top)$, whose squared Frobenius norm is its rank, at most d, and averaging over i,

$$\Big(\operatorname*{mean}_i \big(a^\top \mathrm{Orth}(W^{\mathrm{GD}} - ab^\top)b\big)^2\Big)^{1/2} = O_p(\sqrt{d}). \tag{73}$$

Substituting $u_i = \sqrt{1 - \frac{1}{N}}\, a + \frac{1}{N}U\mathbf{1}$ and $v_i = \sqrt{1 - \frac{1}{N}}\, b + \frac{1}{N}V\mathbf{1}$,

$$\begin{aligned} u_i^\top W^{\mathrm{Muon}} v_i - \Big(1 - \frac{1}{N}\Big) a^\top W^{\mathrm{Muon}} b = \sqrt{1 - \frac{1}{N}}\, a^\top W^{\mathrm{Muon}} \frac{V\mathbf{1}}{N} + \sqrt{1 - \frac{1}{N}}\, \frac{(U\mathbf{1})^\top}{N} W^{\mathrm{Muon}} b \\ + \frac{(U\mathbf{1})^\top}{N} W^{\mathrm{Muon}} \frac{V\mathbf{1}}{N}, \end{aligned} \tag{74}$$

so by $(x + y + z)^2 \le 3(x^2 + y^2 + z^2)$, $\|W^{\mathrm{Muon}}\|_{\mathrm{op}} \le 1$, $\mathbb{E}\|a\|^2 = \mathbb{E}\|b\|^2 = d$ (so that $\mathrm{mean}_i \|a\|^2 = O_p(d)$ by Markov) and Theorem F.3(iii),

$$\begin{aligned} &\operatorname*{mean}_i \Big(u_i^\top W^{\mathrm{Muon}} v_i - \big(1 - \frac{1}{N}\big) a^\top W^{\mathrm{Muon}} b\Big)^2 \\ &\qquad \le 3\Big[\operatorname*{mean}_i \|a\|^2 \frac{\|V\mathbf{1}\|^2}{N^2} + \frac{\|U\mathbf{1}\|^2}{N^2} \operatorname*{mean}_i \|b\|^2 + \frac{\|U\mathbf{1}\|^2 \|V\mathbf{1}\|^2}{N^4}\Big] = O_p\Big(\frac{d^2}{N}\Big), \end{aligned} \tag{75}$$

that is

$$\Big(\operatorname*{mean}_i \Big(u_i^\top W^{\mathrm{Muon}} v_i - \big(1 - \frac{1}{N}\big) a^\top W^{\mathrm{Muon}} b\Big)^2\Big)^{1/2} = O_p\big(d/\sqrt{N}\big). \tag{76}$$

Step 5 (combining). By Theorem F.28 and $1 - \frac{1}{N} > 0$, $u_i^\top W^{\mathrm{Muon}} v_i \ge (1 - \frac{1}{N}) a^\top \mathrm{Orth}(W^{\mathrm{GD}} - ab^\top) b - |u_i^\top W^{\mathrm{Muon}} v_i - (1 - \frac{1}{N}) a^\top W^{\mathrm{Muon}} b|$. Using Equation (70) of Step 2 with this lower bound, averaging over i, and applying Equation (73), Equation (76) and $\|W^{\mathrm{GD}}\|_* \le 7d\sqrt{Nd}$ from Lemma F.4(iii),

$$\operatorname*{mean}_i \big|u_i^\top W^{\mathrm{Muon}} v_i\big| \le \frac{\|W^{\mathrm{GD}}\|_*}{N} + O_p(d/N) + O_p(\sqrt{d}) + O_p(d/\sqrt{N}) \le \frac{7d^{3/2}}{\sqrt{N}} + O_p(\sqrt{d}). \tag{77}$$

Substituting into Equation (69) and dividing by $(N - 1)d$,

$$\begin{aligned} \frac{1}{(N-1)d} \operatorname*{mean}_i \big(u_i^\top W^{\mathrm{Muon}} v_i\big)^2 &\le \frac{O_p(d)}{(N-1)d} + \frac{2}{N-1}\Big(\frac{7d^{3/2}}{\sqrt{N}} + O_p(\sqrt{d})\Big) \\ &= 14\gamma^{-3/2}\Big(1 + \frac{1}{N-1}\Big) + o_p(1), \end{aligned} \tag{78}$$

which is at most $15\gamma^{-3/2}$ with probability tending to one. □

Lemma F.30 (Second moment for Muon). *If $\gamma \geq 16$, then $\frac{1}{N(N-1)d}\sum_{i=1}^{N}\sum_{j\neq i}\left(u_j^\top W^{\mathrm{Muon}} v_i\right)^2 = 1 + O(\gamma^{-3/2}) + o_p(1)$, where with probability tending to one the $O(\gamma^{-3/2})$ term has absolute value at most $52\gamma^{-3/2}$.*

Proof. Divide Theorem F.23 by d and subtract 1. By Theorem F.24 the norm term contributes $O(N^{-1}) + O_p((Nd)^{-1/2}) = o_p(1)$; by Lemma F.27 the trace term is at most $145\gamma^{-2} \leq \frac{145}{4}\gamma^{-3/2}$ in absolute value for $\gamma \geq 16$; by Lemma F.29 the diagonal term lies in $[-15\gamma^{-3/2}, 0]$. Adding, $\frac{145}{4} + 15 \leq 52$. □

Lemma F.31 ($\bar{\kappa}$ for Muon). *If $\gamma \geq 16$, then $(\bar{\kappa}^{\mathrm{Muon}})^2 = d\big(1 + O(\gamma^{-3/2}) + o_p(1)\big)$ with the same constant as in Theorem F.30.*

Proof. By Equation (37) with $W = W^{\mathrm{Muon}}$, Theorem F.30 and Theorem F.21, $(\bar{\kappa}^{\mathrm{Muon}})^2 = d(1+O(\gamma^{-3/2})+o_p(1)) - O_p(d/N) = d(1 + O(\gamma^{-3/2}) + o_p(1))$. □

Lemma F.32 (ρ for Muon). *If $\gamma \geq 16$, then*

$$\rho^{\mathrm{Muon}} = \sqrt{\frac{d}{\gamma}}\Big[\frac{8}{3\pi}\Big(1 - \frac{1}{10\gamma}\Big) + O(\gamma^{-3/2}) + o_p(1)\Big]. \tag{79}$$

Unlike for GD, the data budget enters ρ^{Muon} a second time, through $\Phi(\gamma)$.

Proof. By Theorem F.22, $\bar{h}^{\mathrm{Muon}}/\sqrt{d} = \sqrt{d/\gamma}\,\Phi(\gamma)(1 + o_p(1))$. Set $y := (\bar{\kappa}^{\mathrm{Muon}})^2/d$. By Theorem F.31 and $\gamma \geq 16$, with probability tending to one

$$|y - 1| \leq 52\gamma^{-3/2} + o_p(1) \leq \frac{52}{64} + o_p(1) \leq \frac{9}{10},$$

so $y \geq \frac{1}{10}$. For $y \geq \frac{1}{10}$,

$$|y^{-1/2} - 1| = \frac{|1-y|}{\sqrt{y}\,(1+\sqrt{y})} \leq \sqrt{10}\,|y - 1|,$$

which gives $\sqrt{d}/\bar{\kappa}^{\mathrm{Muon}} = 1 + O(\gamma^{-3/2}) + o_p(1)$. Multiplying the two,

$$\rho^{\mathrm{Muon}} = \frac{\bar{h}^{\mathrm{Muon}}}{\sqrt{d}} \cdot \frac{\sqrt{d}}{\bar{\kappa}^{\mathrm{Muon}}} = \sqrt{\frac{d}{\gamma}}\,\Phi(\gamma)\big(1 + O(\gamma^{-3/2}) + o_p(1)\big), \tag{80}$$

and Lemma F.4(iv) gives $\Phi(\gamma) = \frac{8}{3\pi}(1 - \frac{1}{10\gamma} + r(\gamma))$ with $|r(\gamma)| \leq \gamma^{-2} \leq \frac{1}{4}\gamma^{-3/2}$, so $0 < \Phi(\gamma) \leq \frac{8}{3\pi}(1+\gamma^{-2}) < 1$ and $\Phi(\gamma)(1 + O(\gamma^{-3/2}) + o_p(1)) = \frac{8}{3\pi}(1 - \frac{1}{10\gamma}) + O(\gamma^{-3/2}) + o_p(1)$.

Two by-products of this computation are used in the proof of Theorem 4.2. First, the constant implied by the $O(\gamma^{-3/2})$ in the displayed multiplicative form is absolute, so there is an absolute $\gamma_1 \geq 16$ such that for $\gamma \geq \gamma_1$ that factor lies in $[\frac{1}{2}, \frac{3}{2}]$ with probability tending to one. Second, for every $\gamma \geq 16$ the same display bounds ρ^{Muon} away from zero without any expansion: $y \leq \frac{19}{10}$ gives $\sqrt{d}/\bar{\kappa}^{\mathrm{Muon}} \geq (19/10)^{-1/2} > \frac{7}{10}$, and $\Phi(\gamma) \geq \frac{1}{8}$ by Lemma F.4(iii),(iv), so with probability tending to one

$$\rho^{\mathrm{Muon}} \;\geq\; \frac{7}{80}\sqrt{d/\gamma}\,(1 - o_p(1)) \;\geq\; \frac{1}{16}\sqrt{d/\gamma}. \tag{81}$$

□

F.5 Combining the GD and Muon estimates

Proof of Equation (8). *Reduction to large γ.* By Lemma F.16 and $2 + \gamma \geq \gamma$, $\rho^{\mathrm{GD}} \leq \sqrt{d/\gamma}\,(1 + o_p(1))$, so Equation (81) gives $\mathcal{R} \leq 17$ with probability tending to one, for every $\gamma \geq 16$; together with $0 \leq \frac{3\pi}{8}(1 - \frac{9}{10\gamma}) \leq \frac{3\pi}{8}$ this makes the left-hand side of Equation (8) at most 19 there. On a bounded range $16 \leq \gamma \leq \gamma_1$ we have $19 \leq 19\gamma_1^{3/2}\gamma^{-3/2}$, so Equation (8) holds on that range as soon as $C_1 \geq 19\gamma_1^{3/2}$. It therefore suffices to prove it for $\gamma \geq \gamma_1$, with γ_1 the absolute constant of Lemma F.32; on that range the error factor of Lemma F.32 lies in $[\frac{1}{2}, \frac{3}{2}]$ with probability tending to one, so its reciprocal is again $1 + O(\gamma^{-3/2}) + o_p(1)$.

Expansion. By Lemma F.16 and Lemma F.32, and since $\Phi(\gamma) > 0$,

$$\begin{aligned}\mathcal{R} = \frac{\rho^{\mathrm{GD}}}{\rho^{\mathrm{Muon}}} &= \sqrt{\frac{d}{2+\gamma}}\big(1+o_p(1)\big)\cdot\sqrt{\frac{\gamma}{d}}\,\frac{1}{\Phi(\gamma)\big(1+O(\gamma^{-3/2})+o_p(1)\big)} \\ &= \frac{\sqrt{\gamma/(2+\gamma)}}{\Phi(\gamma)}\big(1+O(\gamma^{-3/2})+o_p(1)\big).\end{aligned} \tag{82}$$

The deterministic prefactor is expanded in $1/\gamma$: $\sqrt{\gamma/(2+\gamma)} = (1+2/\gamma)^{-1/2} = 1-\gamma^{-1}+O(\gamma^{-2})$, while Lemma F.4(iv) gives $\Phi(\gamma)^{-1} = \frac{3\pi}{8}\big(1+\frac{1}{10\gamma}+O(\gamma^{-2})\big)$, so

$$\frac{\sqrt{\gamma/(2+\gamma)}}{\Phi(\gamma)} = \frac{3\pi}{8}\Big(1-\frac{1}{\gamma}+O(\gamma^{-2})\Big)\Big(1+\frac{1}{10\gamma}+O(\gamma^{-2})\Big) = \frac{3\pi}{8}\Big(1-\frac{9}{10\gamma}\Big)+O(\gamma^{-2}). \tag{83}$$

For $\gamma \geq 16$ the prefactor is bounded, $0 < \sqrt{\gamma/(2+\gamma)}/\Phi(\gamma) \leq \frac{3\pi}{8}(1-\frac{1}{160}-\frac{1}{256})^{-1}$, so multiplying by $1+O(\gamma^{-3/2})+o_p(1)$ leaves the error of the same order; and $\gamma^{-2} \leq \frac{1}{4}\gamma^{-3/2}$ in that range absorbs the $O(\gamma^{-2})$ term. Hence

$$\mathcal{R} = \frac{3\pi}{8}\Big(1-\frac{9}{10\gamma}\Big)+O(\gamma^{-3/2})+o_p(1), \tag{84}$$

where, with probability tending to one, the $O(\gamma^{-3/2})$ term has absolute value at most $C_1\gamma^{-3/2}$ with C_1 independent of γ, d, N. This is Equation (8). □

F.6 Monotonicity in the data budget

The expansion Equation (8) has two consequences for how $\mathcal{R}$ changes with γ. The first compares a fixed pair of data budgets; the second concerns the limit $d \to \infty$.

Lemma F.33 (Fixed ratio). *Let $\lambda > 0$ and*

$$\gamma > \gamma_0(\lambda) := \max\Big\{16, \Big(\frac{160C_1(1+\lambda)}{27\pi\lambda}\Big)^2\Big\}.$$

For the same d, let $N_1 = \gamma d$ and $N_2 = (1+\lambda)\gamma d$, and let $\mathcal{R}_{N_1,d}$ and $\mathcal{R}_{N_2,d}$ be the corresponding values of $\mathcal{R}$. Then, with probability tending to one, $\mathcal{R}_{N_2,d} > \mathcal{R}_{N_1,d}$.

Proof. Apply Equation (8) once at γ and once at $(1+\lambda)\gamma$, both of which are at least 16: for every $\varepsilon > 0$, with probability tending to one,

$$\begin{aligned}\mathcal{R}_{N_2,d} - \mathcal{R}_{N_1,d} &\geq \frac{3\pi}{8}\Big(1-\frac{9}{10(1+\lambda)\gamma}\Big) - \frac{3\pi}{8}\Big(1-\frac{9}{10\gamma}\Big) \\ &\quad - C_1\big((1+\lambda)\gamma\big)^{-3/2} - C_1\gamma^{-3/2} - 2\varepsilon \\ &\geq \frac{27\pi}{80}\cdot\frac{\lambda}{(1+\lambda)\gamma} - 2C_1\gamma^{-3/2} - 2\varepsilon.\end{aligned} \tag{85}$$

Moreover $\frac{27\pi}{80}\cdot\frac{\lambda}{(1+\lambda)\gamma} > 2C_1\gamma^{-3/2}$ is equivalent to $\gamma^{1/2} > \frac{160C_1(1+\lambda)}{27\pi\lambda}$, which holds by assumption; taking $\varepsilon = \frac{1}{4}\big(\frac{27\pi}{80}\cdot\frac{\lambda}{(1+\lambda)\gamma} - 2C_1\gamma^{-3/2}\big)$ gives the claim. □

Deduction of the informal theorem. Write $A = 3\pi/8$ and $B = 27\pi/80$. By Equation (8), $\mathcal{R} - A = -B/\gamma + O(\gamma^{-3/2}) + o_p(1)$. For sufficiently large fixed γ, the $O(\gamma^{-3/2})$ term has absolute value at most $B/(4\gamma)$ with probability tending to one, and the $o_p(1)$ term has absolute value at most $B/(4\gamma)$ with probability tending to one as $d \to \infty$. Consequently,

$$-\frac{3B}{2\gamma} \leq \mathcal{R} - A \leq -\frac{B}{2\gamma}$$

with probability tending to one. This gives the claimed two-sided $1/\gamma$ rate, and $\mathcal{R} > 1$ once γ is sufficiently large because $A > 1$. Together with Lemma F.33, this proves all three conclusions of Theorem 4.2.

Lemma F.34 (Pointwise monotonicity). *Suppose that for every sufficiently large γ the limit in probability $\mathcal{R}_\infty(\gamma) := \lim_{d\to\infty}\mathcal{R}$ exists, and that $\mathcal{R}_\infty$ is analytic as a function of $1/\gamma$ in a neighbourhood of 0, i.e. $\mathcal{R}_\infty(\gamma) = \sum_{k\geq 0} c_k\gamma^{-k}$ for all sufficiently large γ. Then there is γ_2 such that $\mathcal{R}'_\infty(\gamma) > 0$ for $\gamma > \gamma_2$.*

Proof. By Equation (8) and convergence in probability, $|\mathcal{R}_\infty(\gamma) - \frac{3\pi}{8}(1 - \frac{9}{10\gamma})| \le C_1\gamma^{-3/2}$, that is

$$\Big|\sum_{k\ge 0} c_k\gamma^{-k} - \frac{3\pi}{8} + \frac{27\pi}{80}\gamma^{-1}\Big| \le C_1\gamma^{-3/2}. \tag{86}$$

Letting $\gamma \to \infty$ gives $c_0 = \frac{3\pi}{8}$; dividing both sides by γ^{-1} and letting $\gamma \to \infty$ gives $c_1 = -\frac{27\pi}{80}$. Hence

$$\mathcal{R}'_\infty(\gamma) = -\sum_{k\ge 1} k\, c_k\gamma^{-k-1} = \gamma^{-2}\Big(\frac{27\pi}{80} - \sum_{k\ge 2} k\, c_k\gamma^{-(k-1)}\Big), \tag{87}$$

and the series in the parentheses tends to 0 as $\gamma \to \infty$, so the parenthesized quantity is positive for all sufficiently large γ. □

Remark F.35. Neither hypothesis of Theorem F.34 is proved here. The estimates above pin $(\bar\kappa^{\mathrm{Muon}})^2/d$ down only to within $O(\gamma^{-3/2})$ (Theorem F.31), so they do not by themselves give the limit in probability; $\rho^{\mathrm{GD}}/\sqrt{d}$ and $\bar h^{\mathrm{Muon}}/d$, by contrast, do converge, by Lemma F.16 and Theorem F.22. Analyticity does not follow from Equation (8) alone either: with $x = 1/\gamma$, the function $1 - x + x^{3/2}\sin(1/x)$ admits an expansion of the same form and is not monotone near $x = 0$.

Numerical consequences. Since $\frac{3\pi}{8} \approx 1.178 > 1$, Equation (8) puts $\mathcal{R}$ above one for all large γ: past some data budget the GD iterate has the larger margin-to-spread ratio. The leading term $\frac{3\pi}{8}(1 - \frac{9}{10\gamma})$ is increasing in γ; Theorem F.33 turns this into a statement about $\mathcal{R}$ itself, and Theorem F.34 into one about its limit, under the hypotheses discussed in Theorem F.35. The leading term crosses one at $\gamma \approx 5.95$, below the range $\gamma \ge 16$ in which Lemma F.32 is proved, so the location of the crossing is established empirically rather than by Theorem 4.2.