

EXACT MEAN SQUARE LINEAR STABILITY ANALYSIS FOR SGD

Anonymous authors

Paper under double-blind review

ABSTRACT

The dynamical stability of optimization methods at the vicinity of minima of the loss has recently attracted significant attention. For gradient descent (GD), stable convergence is possible only to minima that are sufficiently flat w.r.t. the step size, and those have been linked with favorable properties of the trained model. However, while the stability threshold of GD is well-known, to date, no explicit expression has been derived for the exact threshold of stochastic GD (SGD). In this paper, we derive such a closed-form expression. Specifically, we provide an explicit condition on the step size that is both necessary and sufficient for the stability of SGD in the mean square sense. Our analysis sheds light on the precise role of the batch size B . Particularly, we show that the stability threshold is a monotonically non-decreasing function of the batch size, which means that reducing the batch size can only decrease stability. Furthermore, we show that SGD’s stability threshold is equivalent to that of a process which takes in each iteration a full batch gradient step w.p. $1 - p$, and a single sample gradient step w.p. p , where $p \approx 1/B$. This indicates that even with moderate batch sizes, SGD’s stability threshold is very close to that of GD’s. Finally, we prove simple necessary conditions for stability, which depend on the batch size, and are easier to compute than the precise threshold. We demonstrate our theoretical findings through experiments on the MNIST dataset.

1 INTRODUCTION

The dynamical stability of optimization methods has been shown to play a key role in shaping the properties of trained models. For instance, gradient descent (GD) can stably converge only to minima that are sufficiently flat with respect to the step size (Cohen *et al.*, 2021), and in the context of neural networks, such minima were shown to correspond to models with favorable properties. These include smoothness of the predictor function (Ma & Ying, 2021; Nacson *et al.*, 2023; Mulayoff *et al.*, 2021), balancedness of the layers (Mulayoff & Michaeli, 2020), and arguably better generalization Hochreiter & Schmidhuber (1997); Keskar *et al.* (2016); Jastrzębski *et al.* (2017); Wu *et al.* (2017); Ma & Ying (2021). While the stability threshold of GD is well-known, that of stochastic GD (SGD) has yet to be fully understood. Several theoretical works studied SGD’s dynamics using various notions of stability, including mean square (Wu *et al.*, 2018; Granzio *et al.*, 2022; Velikanov *et al.*, 2023), higher moments (Ma & Ying, 2021), and in probability (Ziyin *et al.*, 2023). However, these works either do not provide explicit stability conditions (*e.g.*, presenting the condition as an optimization problem (Wu *et al.*, 2018; Ma & Ying, 2021) or in terms of a moment generating function (Velikanov *et al.*, 2023)), or rely on strong assumptions (*e.g.*, the nature of the Hessian batching noise and infinite network widths (Granzio *et al.*, 2022), momentum parameter close to 1 and “spectrally expressible” dynamics (Velikanov *et al.*, 2023)). Moreover, several empirical works studied SGD’s stability (Cohen *et al.*, 2021; Gilmer *et al.*, 2022; Jastrzębski *et al.*, 2020; 2019), yet its exact stability threshold is still unknown.

In this paper, we analyze the stability of SGD in the mean square sense. We start by considering interpolating minima, which are common in training of overparametrized models. In this case, we provide an explicit threshold on the step size η that is both necessary and sufficient for stability. Our analysis sheds light on the precise role of the batch size B . Particularly, we show that the maximal step size allowing stable convergence is monotonically non-decreasing in the batch size. Namely, decreasing the batch size can only decrease the stability threshold of SGD. Moreover, we show that this threshold is equivalent to that of a process that takes in each iteration a full batch gradient step

w.p. $1 - p$, and a single sample gradient step w.p. p , where $p \approx 1/B$. This suggests that even with moderate batch sizes, SGD’s stability threshold is very close to that of GD’s. Although our result gives an explicit condition on the step size for stability, its computation may still be challenging in practical applications. Thus, we also prove simple necessary criteria for stability, which depend on the batch size and are easier to compute.

Next, we turn to study a broader class of minima which we call *regular*. Specifically, in interpolating minima, the loss of each individual sample has zero gradient and a positive semi-definite (PSD) Hessian. In regular minima, the individual Hessians are still required to be PSD, but the gradients can be arbitrary. Only the mean of the gradients over all samples has to vanish (as in any minimum). In this setting, the dynamics can wander within the null-space of the Hessian, if the gradients have nonzero components in that subspace. However, the interesting question is whether the process is stable within the orthogonal complement of the null space. Here we again provide an explicit condition on the step size that is both necessary and sufficient for stability. We further derive the theoretical limit of the covariance matrix of the dynamics, as well as the limit values of the expected squared distance to the minimum, the expected loss, and the expected squared norm of the gradient, and show how they all decrease when reducing the learning rate. This provides a theoretical explanation to the behavior encountered in common learning rate scheduling strategies.

Finally, we demonstrate our theoretical findings through experiments on the MNIST dataset (LeCun, 1998). These confirm that our theory correctly predicts the stability threshold of SGD, and its dependence on the batch size.

2 BACKGROUND: LINEARIZED DYNAMICS

Let $\ell_i : \mathbb{R}^d \rightarrow \mathbb{R}$ be differentiable almost everywhere for all $i \in [n]$. We consider the minimization of a loss function

$$\mathcal{L}(\boldsymbol{\theta}) = \frac{1}{n} \sum_{i=1}^n \ell_i(\boldsymbol{\theta}) \quad (1)$$

using the SGD iterations

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t - \eta \nabla \hat{\mathcal{L}}_t(\boldsymbol{\theta}_t). \quad (2)$$

Here, η is the step size and $\hat{\mathcal{L}}_t$ is a stochastic approximation of \mathcal{L} obtained as

$$\hat{\mathcal{L}}_t(\boldsymbol{\theta}) = \frac{1}{B} \sum_{i \in \mathfrak{B}_t} \ell_i(\boldsymbol{\theta}), \quad (3)$$

where \mathfrak{B}_t is a batch of size B sampled at iteration t . We assume that the batches $\{\mathfrak{B}_t\}$ are drawn uniformly from the dataset, independently across iterations.

Analyzing the full dynamics of this process is intractable in most cases. Yet, near minima, accurate characterization of the stability of the iterates can be obtained via linearization (Wu *et al.*, 2018; Ma & Ying, 2021; Mulayoff *et al.*, 2021), as is common in stability analysis of nonlinear systems.

Definition 1 (Linearized dynamics). *Let $\boldsymbol{\theta}^*$ be a twice differentiable minimum of \mathcal{L} , and denote*

$$\mathbf{g}_i \triangleq \nabla \ell_i(\boldsymbol{\theta}^*), \quad \mathbf{H}_i \triangleq \nabla^2 \ell_i(\boldsymbol{\theta}^*). \quad (4)$$

Then the linearized dynamics of SGD near $\boldsymbol{\theta}^$ is given by*

$$\boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{H}_i(\boldsymbol{\theta}_t - \boldsymbol{\theta}^*) - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{g}_i. \quad (5)$$

Note that since $\boldsymbol{\theta}^*$ is a minimum point of \mathcal{L} we have that

$$\nabla \mathcal{L}(\boldsymbol{\theta}^*) = \frac{1}{n} \sum_{i=1}^n \mathbf{g}_i = \mathbf{0}. \quad (6)$$

Furthermore, the Hessian of the loss, which we denote by \mathbf{H} , is given by

$$\mathbf{H} \triangleq \nabla^2 \mathcal{L}(\boldsymbol{\theta}^*) = \frac{1}{n} \sum_{i=1}^n \mathbf{H}_i. \quad (7)$$

Thus, the linearized dynamics are in fact SGD iterates on the second-order Taylor expansion of \mathcal{L} at $\boldsymbol{\theta}^*$,

$$\tilde{\mathcal{L}}(\boldsymbol{\theta}) = \mathcal{L}(\boldsymbol{\theta}^*) + \frac{1}{2}(\boldsymbol{\theta} - \boldsymbol{\theta}^*)^\top \mathbf{H}(\boldsymbol{\theta} - \boldsymbol{\theta}^*). \quad (8)$$

3 STABILITY OF FIRST AND SECOND MOMENTS

Our focus is on the stability of SGD’s dynamics. We specifically examine the dynamics within two subspaces, the null space of the Hessian \mathbf{H} at the minimum, and its orthogonal complement. We denote the projection of any vector $\mathbf{v} \in \mathbb{R}^d$ onto the null space of \mathbf{H} by \mathbf{v}^\parallel , and its projection onto the orthogonal complement of the null space by \mathbf{v}^\perp .

Multiple works studied the stability of SGD’s dynamics. Commonly, this was done by analyzing the evolution of the moments of the linearized dynamics (see Sec. 2) over time, with a specific emphasis on the second moment, which is the approach we take here. However, before discussing the evolution of the second moment, let us summarize the behavior of the first moment. Specifically, it is easy to demonstrate that the first moment of SGD’s linearized trajectory $\{\mathbb{E}[\boldsymbol{\theta}_t]\}$ is the same as GD’s. Now, since GD is stable if and only if $\eta \leq 2/\lambda_{\max}(\mathbf{H})$, we have the following (see proof in App. B).

Theorem 1 (Stability of the mean). *Assume that $\boldsymbol{\theta}^*$ is a twice differentiable minimum. Consider the linear dynamics of $\{\boldsymbol{\theta}_t\}$ from Def. 1 and let*

$$\eta_{\text{mean}}^* \triangleq \frac{2}{\lambda_{\max}(\mathbf{H})}. \quad (9)$$

Then

1. $\mathbb{E}[\boldsymbol{\theta}_t^\parallel] = \mathbb{E}[\boldsymbol{\theta}_0^\parallel]$ for all $t \geq 0$;
2. $\limsup_{t \rightarrow \infty} \|\mathbb{E}[\boldsymbol{\theta}_t] - \boldsymbol{\theta}^*\|$ is finite if and only if $\eta \leq \eta_{\text{mean}}^*$;
3. $\lim_{t \rightarrow \infty} \|\mathbb{E}[\boldsymbol{\theta}_t^\perp] - \boldsymbol{\theta}^{*\perp}\| = 0$ if $\eta < \eta_{\text{mean}}^*$.

We next proceed to analyze the dynamics of the second moment, which determine stability in the mean square sense. Note that boundedness of the first moment is a necessary condition for boundedness of the second moment. Therefore, the condition $\eta \leq \eta_{\text{mean}}^*$ is a prerequisite for stability in the mean square sense. However, how much smaller than η_{mean}^* is SGD’s mean square stability threshold, is not currently known in closed form. Here, we determine the precise threshold for the mean square stability of SGD’s linearized dynamics. To achieve this, we leverage the approach taken by Ma & Ying (2021), who investigated the stability of SGD in the context of interpolating minima.

3.1 INTERPOLATING MINIMA

We begin by studying interpolating minima, which are prevalent in the training of overparametrized models. In this case, the model fits the training set perfectly¹, which means that these global minima are also minima for each sample individually. This is expressed mathematically as follows.

Definition 2 (Interpolating minima). *A twice differentiable minimum $\boldsymbol{\theta}^*$ is said to be interpolating if for each sample $i \in [n]$ the gradient $\mathbf{g}_i = \mathbf{0}$ and the Hessian \mathbf{H}_i is PSD.*

In this setting, Ma & Ying (2021) showed that the evolution of any moment of SGD over time is fully tractable. Specifically, for the second moment, they proved the following.

Theorem 2 (Ma & Ying (2021), Thm. 1 + Cor. 3). *Assume that $\boldsymbol{\theta}^*$ is a twice differentiable interpolating minimum. Consider the linear dynamics of $\{\boldsymbol{\theta}_t\}$ from Def. 1, and let*

$$\mathbf{Q}(\eta, B) \triangleq (\mathbf{I} - \eta\mathbf{H}) \otimes (\mathbf{I} - \eta\mathbf{H}) + \frac{n - B}{B(n - 1)} \frac{\eta^2}{n} \sum_{i=1}^n (\mathbf{H}_i \otimes \mathbf{H}_i - \mathbf{H} \otimes \mathbf{H}), \quad (10)$$

where \otimes denotes the Kronecker product. Then $\limsup_{t \rightarrow \infty} \mathbb{E}[\|\boldsymbol{\theta}_t - \boldsymbol{\theta}^*\|^2]$ is finite if and only if

$$\max_{\boldsymbol{\Sigma} \in \mathcal{S}_+(\mathbb{R}^{d \times d})} \frac{\|\mathbf{Q}(\eta, B) \text{vec}(\boldsymbol{\Sigma})\|}{\|\boldsymbol{\Sigma}\|_F} \leq 1, \quad (11)$$

where $\mathcal{S}_+(\mathbb{R}^{d \times d})$ denotes the set of all PSD matrices over $\mathbb{R}^{d \times d}$. Furthermore, if the spectral radius $\rho(\mathbf{Q}(\eta, B)) \leq 1$ then $\limsup_{t \rightarrow \infty} \mathbb{E}[\|\boldsymbol{\theta}_t - \boldsymbol{\theta}^*\|^2]$ is finite.

¹The important minima from a practical standpoint are the ones that benign overfit.

Below, we omit the dependence of \mathbf{Q} on η and B whenever these are not essential for the discussion. In this theorem, Σ represents the second-moment matrix of $\boldsymbol{\theta}_t - \boldsymbol{\theta}^*$. Specifically, the matrix $\Sigma_t = \mathbb{E}[(\boldsymbol{\theta}_t - \boldsymbol{\theta}^*)(\boldsymbol{\theta}_t - \boldsymbol{\theta}^*)^\top]$ evolves over time as $\text{vec}(\Sigma_{t+1}) = \mathbf{Q} \text{vec}(\Sigma_t)$. Therefore, the stability condition of (11) simply states that if the dynamics of the dominant initial state of the system (which is restricted to PSD matrices) is bounded, then Σ_t is bounded and vice versa. However, this characterization leaves us with a complex optimization problem over a high dimension (d^2), which is hard to solve numerically. Therefore, this approach does not reduce the problem into a condition from which we can gain any meaningful theoretical insight into the behavior of SGD.

Our first key result is that the optimization problem (11) can be reduced to an eigenvalue problem. Specifically, we establish (see Sec. 3.3) that when the eigenvectors of the $d^2 \times d^2$ matrix \mathbf{Q} are reshaped into $d \times d$ matrices, they always correspond to either symmetric or skew-symmetric matrices². Moreover, the dominant eigenvalue of \mathbf{Q} is positive and always corresponds to a PSD matrix. Consequently, the maximizer of (11) is the top eigenvector of \mathbf{Q} , which we use, along with some algebraic manipulation, to derive the following result (see proof in App. B).

Theorem 3 (Exact threshold for interpolating minima). *Assume that $\boldsymbol{\theta}^*$ is a twice differentiable interpolating minimum. Consider the linear dynamics of $\{\boldsymbol{\theta}_t\}$ from Def. 1, and let*

$$\mathbf{C} \triangleq \frac{1}{2} \mathbf{H} \oplus \mathbf{H}, \quad \mathbf{D} \triangleq (1-p) \mathbf{H} \otimes \mathbf{H} + p \frac{1}{n} \sum_{i=1}^n \mathbf{H}_i \otimes \mathbf{H}_i, \quad (12)$$

where \oplus denotes the Kronecker sum and $p \triangleq \frac{n-B}{B(n-1)} \in [0, 1]$. Define

$$\eta_{\text{var}}^* \triangleq \frac{2}{\lambda_{\max}(\mathbf{C}^\dagger \mathbf{D})}, \quad (13)$$

where \mathbf{C}^\dagger denotes the Moore-Penrose inverse of \mathbf{C} . Then

1. $\boldsymbol{\theta}_t^\parallel = \boldsymbol{\theta}_0^\parallel$ for all $t \geq 0$;
2. $\limsup_{t \rightarrow \infty} \mathbb{E}[\|\boldsymbol{\theta}_t^\perp - \boldsymbol{\theta}^{*\perp}\|^2]$ is finite if and only if $\eta \leq \eta_{\text{var}}^*$;
3. $\lim_{t \rightarrow \infty} \mathbb{E}[\|\boldsymbol{\theta}_t^\perp - \boldsymbol{\theta}^{*\perp}\|^2] = 0$ if $\eta < \eta_{\text{var}}^*$.

This result provides an explicit characterization of the mean square stability of SGD. First, we observe that the set of step sizes that are stable in the mean square sense, is an interval. This is in contrast to stability in probability, where the stable learning rates can comprise of several disjoint intervals (Ziyin *et al.*, 2023). Moreover, SGD’s threshold, η_{var}^* , has the same form as the threshold for GD, $2/\lambda_{\max}$, but with a different matrix. In App. I we show how Thm. 3 recovers GD’s condition in full batch.

The dependence of η_{var}^* on the batch size B may not be immediate to see from the theorem. However, we can prove the following (see proof in App. D).

Proposition 1. *Assume that $\boldsymbol{\theta}^*$ is a twice differentiable interpolating minimum. Then η_{var}^* is a non-decreasing function of B .*

This result implies that decreasing the batch size can only decrease the stability threshold, which settles with the empirical observations, *e.g.*, in Wu *et al.* (2018). Additionally, since η_{var}^* is non-decreasing with B , and for $B = n$ it equals η_{mean}^* , we have that the gap between $\lambda_{\max}(\mathbf{C}^\dagger \mathbf{D})$ and $\lambda_{\max}(\mathbf{H})$ is non-negative for all $B \in [1, n]$ and non-increasing in B . For stable minima, $\lambda_{\max}(\mathbf{C}^\dagger \mathbf{D})$ is bounded from above by $2/\eta$. This suggests that training with smaller batches leads to lower $\lambda_{\max}(\mathbf{H})$ which results in smoother predictor functions (Mulayoff *et al.*, 2021).

At what rate does η_{var}^* increase with B towards η_{mean}^* ? To understand this, note that \mathbf{D} is a convex combination of two matrices, where p represents the combination weight. The first matrix is $\mathbf{H} \otimes \mathbf{H}$, which is associated with full batch SGD ($B = n$), while the second matrix is $\frac{1}{n} \sum_{i=1}^n \mathbf{H}_i \otimes \mathbf{H}_i$, which is related to single sample SGD ($B = 1$). We can use this fact to explain the effect of the batch size on dynamical stability by presenting an equivalent stochastic process that has the same stability threshold as SGD (see proof in App. E).

²Eigenbases corresponding to eigenvalues of multiplicity greater than one, always have a basis consisting of symmetric and skew-symmetric matrices.

Proposition 2. Let $\text{ALG}(p)$ be a stochastic optimization algorithm in which

$$\boldsymbol{\theta}_{t+1} = \begin{cases} \boldsymbol{\theta}_t - \eta \nabla \ell_{i_t}(\boldsymbol{\theta}_t) & \text{w.p. } p, \\ \boldsymbol{\theta}_t - \eta \nabla \mathcal{L}(\boldsymbol{\theta}_t) & \text{w.p. } 1 - p, \end{cases} \quad (14)$$

where $\{i_t\}$ are i.i.d. random indices distributed uniformly over the training set. Assume that $\boldsymbol{\theta}^*$ is a twice differentiable interpolating minimum. Then when $p = \frac{n-B}{B(n-1)}$, $\text{ALG}(p)$ has the same stability threshold in the vicinity of $\boldsymbol{\theta}^*$ as SGD with batch size B .

In simpler terms, in each iteration the process $\text{ALG}(p)$ takes a gradient step with a batch of one sample ($B = 1$) with probability p and with a full batch ($B = n$) with probability $1 - p$. This result shows that the stability conditions of SGD and of $\text{ALG}(p)$ are the same for $p = \frac{n-B}{B(n-1)}$. When $n \gg B$, we have that $p \approx 1/B$. Therefore, Prop. 2 implies that, in the context of stability, even moderate values of B make mini-batch SGD behave like GD. We note that while propositions 1 and 2 were presented in the context of interpolating minima, they also apply to regular minima (see Sec. 3.2).

It is worthwhile mentioning that if the stability condition is not met, then the linearized dynamics diverge. However, in practice, the full (non-linearized) dynamics can just move to a different point on the loss landscape, where the generalized sharpness $\lambda_{\max}(\mathbf{C}^\dagger \mathbf{D})$ is lower. It was shown that GD possesses such a stabilizing mechanism (Damian *et al.*, 2023). An interesting open question is whether a similar mechanism exists in SGD.

Theorem 3 gives an explicit condition on the step size. However, its computation may still be challenging in practical applications, as it requires inverting, multiplying, and computing the spectral norm of large ($d^2 \times d^2$) matrices. Yet, we can obtain necessary criteria for stability that are simple and easier to verify, and which also depend on the batch size. To do so, we note that the eigenvalues of $\mathbf{C}^\dagger \mathbf{D}$ coincide with those of $(\mathbf{C}^\dagger)^{\frac{1}{2}} \mathbf{D} (\mathbf{C}^\dagger)^{\frac{1}{2}}$. We therefore upper bound η_{var}^* by evaluating $\mathbf{u}^\top ((\mathbf{C}^\dagger)^{\frac{1}{2}} \mathbf{D} (\mathbf{C}^\dagger)^{\frac{1}{2}}) \mathbf{u}$ for interesting directions \mathbf{u} other than the top eigenvector. Specifically, the next result corresponds to $\mathbf{u} = \mathbf{C}^{\frac{1}{2}} \text{vec}(\mathbf{I}) / \|\mathbf{C}^{\frac{1}{2}} \text{vec}(\mathbf{I})\|$ and $\mathbf{u} = \mathbf{C}^{\frac{1}{2}} (\mathbf{v}_{\max} \otimes \mathbf{v}_{\max}) / \|\mathbf{C}^{\frac{1}{2}} (\mathbf{v}_{\max} \otimes \mathbf{v}_{\max})\|$, where \mathbf{v}_{\max} is the top eigenvector of \mathbf{H} (see proof in App. F).

Proposition 3 (Necessary conditions). *The step size η_{var}^* satisfies*

$$\eta_{\text{var}}^* \leq \frac{2\lambda_{\max}(\mathbf{H})}{\lambda_{\max}^2(\mathbf{H}) + \frac{p}{n} \sum_{i=1}^n (\mathbf{v}_{\max}^\top \mathbf{H}_i \mathbf{v}_{\max} - \lambda_{\max}(\mathbf{H}))^2}, \quad (15)$$

as well as

$$\eta_{\text{var}}^* \leq \frac{2\text{Tr}(\mathbf{H})}{(1-p) \|\mathbf{H}\|_{\text{F}}^2 + \frac{p}{n} \sum_{i=1}^n \|\mathbf{H}_i\|_{\text{F}}^2}. \quad (16)$$

From (15), we can deduce a lower bound on the gap between the stability thresholds of GD and SGD. Specifically, when the variance of \mathbf{H}_i over the direction of the top eigenvector of \mathbf{H} is large, η_{var}^* is far from η_{mean}^* for moderate p . In general, this condition is expected to be quite tight when there is a clear dominant direction in \mathbf{H} caused by some \mathbf{H}_i . In contrast, condition (16) is expected to be tight if all $\{\mathbf{H}_i\}$ have roughly the same spectrum but with different bases, *i.e.*, when no sample is dominant and the samples are incoherent.

3.2 NON-INTERPOLATING MINIMA

While for interpolating minima, we saw that $\boldsymbol{\theta}_t^\perp$ can converge to $\boldsymbol{\theta}^{*\perp}$, this is generally not the case for non-interpolating minima. In this section, we explore the dynamics of SGD in the vicinity of a broader class of minima. Particularly, we consider the following definition.

Definition 3 (Regular minima). *A twice differentiable minimum $\boldsymbol{\theta}^*$ is said to be regular if for each sample $i \in [n]$ the Hessian \mathbf{H}_i is PSD.*

This definition encompasses a broader class of minima than Def. 2, as it allows for arbitrary (nonzero) gradients \mathbf{g}_i . Only the mean of the gradients has to vanish (as in any minimum). Intuitively speaking, although a regular minimum does not necessarily fit all the training points, it does not involve a major disagreement among them. This can be understood through the second-order Taylor expansion for each sample, which may be unbounded from below, yet it can only go to minus infinity linearly with the parameters, and not quadratically.

Clearly, having gradients with nonzero components in the null space of the Hessian pushes the dynamics to diverge. Interestingly, for regular minima, the dynamics of SGD in the null space and in its orthogonal complement are separable. Thus, despite having a random walk in the null space, we can give a condition for stability within its orthogonal complement (see proof in App. B).

Theorem 4 (Exact threshold for regular minima). *Assume that θ^* is a twice differentiable regular minimum. Consider the linear dynamics of $\{\theta_t\}$ from Def. 1. Then*

1. $\lim_{t \rightarrow \infty} \mathbb{E}[\|\theta_t^\parallel - \theta^{*\parallel}\|^2] = \infty$ if and only if $\sum_{i=1}^n \|\mathbf{g}_i^\parallel\|^2 > 0$;
2. If $\eta < \eta_{\text{var}}^*$ then $\limsup_{t \rightarrow \infty} \mathbb{E}[\|\theta_t^\perp - \theta^{*\perp}\|^2]$ is finite;
3. If $\limsup_{t \rightarrow \infty} \mathbb{E}[\|\theta_t^\perp - \theta^{*\perp}\|^2]$ is finite then $\eta \leq \eta_{\text{var}}^*$.

We see that η_{var}^* is the stability threshold also for regular minima. Recall that when $\eta < \eta_{\text{var}}^*$, we also have stability of the first moment, and thus $\mathbb{E}[\theta_t^\parallel] = \mathbb{E}[\theta_0^\parallel]$ for any $t \geq 0$. Namely, SGD’s dynamics in the null space is a random walk without drift. Note that moving in the null space does not increase the loss, however it might change the trained model. Furthermore, in the proof, we show that under a mild assumption $\limsup_{t \rightarrow \infty} \mathbb{E}[\|\theta_t^\perp - \theta^{*\perp}\|^2]$ is finite if and only if $0 \leq \eta < \eta_{\text{var}}^*$.

Next, we turn to compute the limit of the second moment of the dynamics (see proof in App. G).

Theorem 5 (Covariance limit). *Assume that θ^* is a twice differentiable regular minimum. Consider the linear dynamics of $\{\theta_t\}$ from Def. 1. If $0 < \eta < \eta_{\text{var}}^*$ then*

$$\lim_{t \rightarrow \infty} \text{vec}(\Sigma_t^\perp) = \eta p (2\mathbf{C} - \eta \mathbf{D})^\dagger \text{vec}(\Sigma_g^\perp), \quad (17)$$

where $\Sigma_g^\perp = \frac{1}{n} \sum_{i=1}^n \mathbf{g}_i^\perp (\mathbf{g}_i^\perp)^\top$.

Using this result we can obtain the mean squared distance to the minimum, the mean of the second-order Taylor expansion of the loss, and the mean of the squared norm of the expansion’s gradient squared at large times (see proof in App. H).

Corollary 1 (Limit values). *Assume that θ^* is a twice differentiable regular minimum. Consider the linear dynamics of $\{\theta_t\}$ from Def. 1 and the second-order Taylor expansion of the loss, $\tilde{\mathcal{L}}$ of (8). If $\eta < \eta_{\text{var}}^*$ then*

1. $\lim_{t \rightarrow \infty} \mathbb{E}[\|\theta_t^\perp - \theta^{*\perp}\|^2] = \eta p (\text{vec}(\mathbf{I}))^\top (2\mathbf{C} - \eta \mathbf{D})^\dagger \text{vec}(\Sigma_g^\perp)$;
2. $\lim_{t \rightarrow \infty} \mathbb{E}[\tilde{\mathcal{L}}(\theta_t)] - \tilde{\mathcal{L}}(\theta^*) = \frac{1}{2} \eta p (\text{vec}(\mathbf{H}))^\top (2\mathbf{C} - \eta \mathbf{D})^\dagger \text{vec}(\Sigma_g^\perp)$;
3. $\lim_{t \rightarrow \infty} \mathbb{E}[\|\nabla \tilde{\mathcal{L}}(\theta_t)\|^2] = \eta p (\text{vec}(\mathbf{H}^2))^\top (2\mathbf{C} - \eta \mathbf{D})^\dagger \text{vec}(\Sigma_g^\perp)$.

We see that these values depend linearly on the covariance matrix of the gradients. Specifically, if $\Sigma_g = \mathbf{0}$ then we recover the results of interpolating minima. Moreover, note that for $\eta \ll \eta_{\text{var}}^*$, we have that $2\mathbf{C} - \eta \mathbf{D} \approx 2\mathbf{C}$. Therefore, the main dependence on η comes from the factor of η preceding these expressions. We thus get that when decreasing the learning rate, the loss level drops, and the parameters θ_t get closer to the minimum. This explains the empirical behavior observed when decreasing the learning rate in neural network training, which causes the loss level to drop.

3.3 PROOF OUTLINE FOR THEOREM 3

Here we give an outline of the proof of Theorem 3. Ma & Ying (2021) showed that the second moment matrix $\Sigma_t = \mathbb{E}[(\theta_t - \theta^*)(\theta_t - \theta^*)^\top]$ evolves over time as

$$\text{vec}(\Sigma_{t+1}) = \mathbf{Q} \text{vec}(\Sigma_t), \quad (18)$$

where \mathbf{Q} is given in (10). Since Σ_t is PSD by definition, we only care about the effect of \mathbf{Q} on vectorizations of PSD matrices. Hence, $\{\Sigma_t\}$ are bounded if and only if (proof in Ma & Ying (2021))

$$\max_{\Sigma \in \mathcal{S}_+(\mathbb{R}^{d \times d})} \frac{\|\mathbf{Q}(\eta, B) \text{vec}(\Sigma)\|}{\|\Sigma\|_{\mathbb{F}}} \leq 1. \quad (19)$$

To obtain the result of Thm. 3 we first rearrange the terms in \mathbf{Q} as

$$\mathbf{Q}(\eta, B) = (1-p) \times (\mathbf{I} - \eta \mathbf{H}) \otimes (\mathbf{I} - \eta \mathbf{H}) + p \times \frac{1}{n} \sum_{i=1}^n (\mathbf{I} - \eta \mathbf{H}_i) \otimes (\mathbf{I} - \eta \mathbf{H}_i). \quad (20)$$

Then, to relax the optimization problem we use the following theorem (see proof in App. C).

Theorem 6. Assume $\{\mathbf{A}_i\}$ are symmetric matrices over $\mathbb{R}^{d \times d}$. Define

$$\mathbf{Q} = \sum_{i=1}^M \mathbf{A}_i \otimes \mathbf{A}_i, \quad (21)$$

and let \mathbf{z}_{\max} be a top eigenvector of \mathbf{Q} . Then

1. there always exists a set of eigenvectors $\{\mathbf{z}_j\}$ for \mathbf{Q} such that each $\mathbf{Z}_j = \text{vec}^{-1}(\mathbf{z}_j)$ is either a symmetric or a skew-symmetric matrix;
2. the spectral radius $\rho(\mathbf{Q}) = \lambda_{\max}(\mathbf{Q})$, i.e., the dominant eigenvalue is positive;
3. $\text{vec}^{-1}(\mathbf{z}_{\max}) \in \mathcal{S}_+(\mathbb{R}^{d \times d})$, i.e., the top eigenvector corresponds to a PSD matrix.

Applying this theorem we get that the maximizer for the constrained optimization problem in (19) is, in fact, the top eigenvalue of \mathbf{Q} . Hence, the linear system is stable if and only if $\lambda_{\max}(\mathbf{Q}) \leq 1$. Since \mathbf{Q} is symmetric, $\lambda_{\max}(\mathbf{Q}) \leq 1$ is equivalent to $\mathbf{u}^T \mathbf{Q} \mathbf{u} \leq 1$ for all $\mathbf{u} \in \mathbb{S}^{d^2-1}$. It is easy to show that $\mathbf{Q} = \mathbf{I} - 2\eta\mathbf{C} + \eta^2\mathbf{D}$. In App. B we show that

$$\mathbf{u}^T \mathbf{Q} \mathbf{u} = 1 - 2\eta\mathbf{u}^T \mathbf{C} \mathbf{u} + \eta^2\mathbf{u}^T \mathbf{D} \mathbf{u} \leq 1 \quad (22)$$

holds for all $\mathbf{u} \in \mathbb{S}^{d^2-1}$ if and only if

$$\eta \leq \frac{2}{\lambda_{\max}(\mathbf{C}^\dagger \mathbf{D})} = \eta_{\text{var}}^*. \quad (23)$$

The full proof is provided in App. B.

4 EXPERIMENTS

In this section, we experimentally validate our theoretical results. We trained single hidden-layer ReLU networks with varying step sizes and batch sizes on a subset of the MNIST dataset (see App. J for details). Since training with cross-entropy and softmax in overparametrized networks results in infima rather than minima, here we opted to use the quadratic loss. Specifically, each class was labeled with a one-hot vector, and the network was trained to predict the label without softmax. Our primary goal in this experiment is to test the stability threshold of SGD; hence, we initialized the training with large weights to ensure that the minimum closest to the starting point is unstable (large weights imply large Hessians, and are thus more likely to violate the stability criterion). We used the same initial point for all the training runs to eliminate initialization effects. To avoid divergence, we started with a very small learning rate and gradually increased it until it reached its designated value (i.e., LR warm-up). Together, large initialization and warm-up force SGD out of the unstable region until it finds a stable minimum and converges as closely as possible at the stability threshold. Convergence was determined when the loss remained below 10^{-6} for 200 consecutive epochs.

Figure 1(a) visualizes the sharpness of the converged minima versus the learning rate for several values of B . It can be observed that for small batch sizes, $\lambda_{\max}(\mathbf{H})$ is far from $2/\eta$. Yet, for moderate batch sizes and above (e.g., $B \geq 32$), these curves virtually coincide, indicating that, in the context of stability, SGD behaves like GD. Figure 1(b) shows the sharpness versus the batch size for three step sizes. Here the stability threshold of SGD rapidly converges to that of GD as the batch size increases.

Apart for the sharpness $\lambda_{\max}(\mathbf{H})$, we also want to compare the generalized sharpness $\lambda_{\max}(\mathbf{C}^\dagger \mathbf{D})$ to $2/\eta$. Since computing the generalized sharpness is impractical in this task, we underestimate it via a lower bound, which results in a tighter necessary condition than (15). The bound corresponds to restricting the optimization problem in (11) to rank one PSD matrices, and is given by (see App. F.1)

$$\frac{2}{\eta_{\text{var}}^*} = \lambda_{\max}(\mathbf{C}^\dagger \mathbf{D}) \geq \max_{\mathbf{v}: \|\mathbf{v}\|=1} \left\{ \mathbf{v}^T \mathbf{H} \mathbf{v} + p \frac{\frac{1}{n} \sum_{i=1}^n (\mathbf{v}^T \mathbf{H}_i \mathbf{v} - \mathbf{v}^T \mathbf{H} \mathbf{v})^2}{\mathbf{v}^T \mathbf{H} \mathbf{v}} \right\}. \quad (24)$$

We solve this optimization problem numerically, by using GD on the unit sphere with predetermined scheduled geodesic step size. In the following, we present graphs of the sharpness $\lambda_{\max}(\mathbf{H})$ at the minima to which we converged, as well as the bounds (24) and (15) on the generalized sharpness

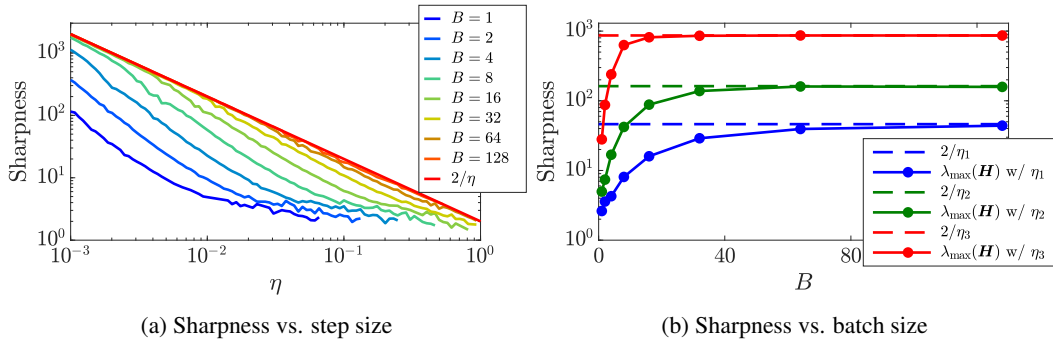


Figure 1: **Sharpness vs. step size and batch size.** We trained single hidden-layer ReLU networks using varying step sizes and batch sizes on a subset of MNIST. Panel (a) visualizes the sharpness of the converged minima versus learning rate for different batch sizes. For small batch sizes, $\lambda_{\max}(\mathbf{H})$ deviates significantly from $2/\eta$. Yet, as the batch size increases to a moderate value, these curves coincide, indicating that in terms of stability, SGD behaves similarly to GD. Panel (b) plots the sharpness against the batch size for three different learning rates $\eta_1 = 0.043$, $\eta_2 = 0.012$, $\eta_3 = 0.002$. Here we see a similar trend where SGD with behaves like GD for $B \geq 32$.

$\lambda_{\max}(\mathbf{C}^\dagger \mathbf{D})$. Using the color coding of Fig. 2, these correspond to

$$\begin{aligned}
 \frac{2}{\eta} &\geq \lambda_{\max}(\mathbf{C}^\dagger \mathbf{D}) \quad \left(= \frac{2}{\eta_{\text{var}}^*} \right) \\
 &\geq \max_{\mathbf{v}: \|\mathbf{v}\|=1} \left\{ \mathbf{v}^\top \mathbf{H} \mathbf{v} + p \frac{\frac{1}{n} \sum_{i=1}^n (\mathbf{v}^\top \mathbf{H}_i \mathbf{v} - \mathbf{v}^\top \mathbf{H} \mathbf{v})^2}{\mathbf{v}^\top \mathbf{H} \mathbf{v}} \right\} \\
 &\geq \lambda_{\max}(\mathbf{H}) + p \frac{\frac{1}{n} \sum_{i=1}^n (\mathbf{v}_{\max}^\top \mathbf{H}_i \mathbf{v}_{\max} - \lambda_{\max}(\mathbf{H}))^2}{\lambda_{\max}(\mathbf{H})} \\
 &\geq \lambda_{\max}(\mathbf{H}),
 \end{aligned} \tag{25}$$

where \mathbf{v}_{\max} denotes the top eigenvector of \mathbf{H} .

Figure 2 depicts the expressions in (25) versus the step size for three batch sizes. We see that for $B = 1$ and $B = 2$, the gap between $2/\eta$ (red) and the optimized bound (24) (purple) upon convergence is small. Particularly, they coincide over a wide range of step sizes η . Since the generalized sharpness $\lambda_{\max}(\mathbf{C}^\dagger \mathbf{D})$ must reside between those two curves, we can deduce two things: (a) Our theory correctly predicts the stability threshold, while SGD converged at the edge of stability (as designed in our experiment); (b) For small batches, the second order moment matrix that maximizes (11) is rank one. As the batch size increases, the two curves draw apart, indicating that the rank of the dominant second-order moment matrix becomes larger. Furthermore, the gap between our simple necessary condition (15) (blue) and the trivial bound of $2/\lambda_{\max}(\mathbf{H})$ (yellow) is large for high learning rates and small for small step sizes. This gap represents the variance of the widths of the minima of the per-sample losses (corresponding to the widths of the quadratic functions $\{(\boldsymbol{\theta} - \boldsymbol{\theta}^*)^\top \mathbf{H}_i (\boldsymbol{\theta} - \boldsymbol{\theta}^*)\}$) in the direction of \mathbf{v}_{\max} , the top eigenvector of \mathbf{H} . Thus we find that for small learning rates, this variance is small and the model is aligned in this direction, and for large learning rates, this variance is high. For more details and experimental results, please see App. J.

5 RELATED WORK

The stability of SGD in the vicinity of minima has been previously studied in multiple works. On the theoretical side, Wu *et al.* (2018) examined stability in the mean square sense and gave an implicit sufficient condition. Granzio *et al.* (2022) used random matrix theory to find the maximal stable learning rate as a function of the batch size. Their work assumes some conditions on the Hessian’s noise caused by batching, and the result holds in the limit of an infinite number of samples and batch size. Velikanov *et al.* (2023) examined SGD with momentum and derived a bound on the maximal learning rate. Their derivation uses “spectrally expressible” approximations and the result is




Figure 2: **(Generalized) Sharpness vs. step size.** We trained single hidden-layer ReLU networks using varying step sizes and batch sizes on MNIST dataset. For each pair of hyper-parameters (η, B) , we measured the sharpness of the minimum (yellow), our necessary condition for stability (blue), and the optimized bound (purple), which their relations are given in (25). We see that for small batch sizes $B = 1$ and $B = 2$, the optimized bound (24) coincides with $2/\eta$, confirming that SGD converged at the edge of stability ($\eta = \eta_{\text{var}}^*$). For additional insights and detail, see discussion in Sec. 4.

given implicitly through a moment-generating function. Ma & Ying (2021) studied the dynamics of higher moments of SGD and gave an implicit necessary and sufficient condition for stability (see Thm. 2 and the discussion following it). Wu *et al.* (2022) gave a necessary condition for stability via alignment property. However, the result assumes and uses a lower bound on a property they coin “alignment” but an analytic bound for this alignment property is lacking for the general case. Ziyin *et al.* (2023) studied the stability of SGD in probability, rather than in mean square. Since convergence in probability is a weaker requirement, theoretically, SGD can converge with high probability to minima which are unstable in the mean square sense. Indeed, their theory predicts that SGD can converge far beyond GD’s threshold. Yet, this did not happen in extensive experiments done in (Cohen *et al.*, 2021, App. G) and (Gilmer *et al.*, 2022). Finally, Mulayoff *et al.* (2021) analyzed the stability in non-differentiable minima, and gave a necessary condition for a minimum to be “strongly stable”, *i.e.*, such that SGD does not escape a ball with a given radius from the minimum.

Liu *et al.* (2021) studied the covariance matrix of the stationary distribution of the iterates in the vicinity of minima. Ziyin *et al.* (2022) improved their results while deriving an implicit equation that relates this covariance to the covariance of the gradient noise. However, these papers do not discuss the conditions under which the dynamics converge to the stationary state. Lee & Jang (2023) studied the stability of SGD along its trajectory and gave an explicit exact condition. Yet, their result does not apply to minima, since the denominator in their condition vanishes at minima.

On the empirical side, Cohen *et al.* (2021) examined the behavior of GD, and showed that it typically converges at the edge of stability. Additionally, for SGD (see their App. G) they found that with large batches, the sharpness behaves similarly to full-batch gradient descent. Moreover, they found that the smaller the batch size, the lower the sharpness at the converged minimum. Gilmer *et al.* (2022) studied how the curvature of the loss affects the training dynamics in multiple settings. They observed that SGD *with momentum* is stable only when the optimization trajectory primarily resides in a region of parameter space where $\eta \lesssim 2/\lambda_{\max}(\mathbf{H})$. Further experimental results in Jastrzȳbski *et al.* (2020; 2019) show that the sharpness along the trajectory of SGD is implicitly regularized.

6 CONCLUSION

We presented an explicit threshold on SGD’s step size, which is both necessary and sufficient for guaranteeing mean-square stability. We showed that this threshold is a monotonically non-decreasing function of the batch size, which implies that decreasing the batch size can only make the process less stable. Additionally, we interpreted the role of the batch size B through an equivalent process that takes in each iteration either a full batch gradient step or a single sample gradient step. Our interpretation highlights that even with moderate batch sizes, SGD’s stability threshold is very close to that of GD. We also proved simpler necessary conditions for stability, which depend on the batch size, and are easier to compute. Finally, we verified our theory through experiments on MNIST.

REFERENCES

- Jeremy Cohen, Simran Kaur, Yuanzhi Li, J Zico Kolter, and Ameet Talwalkar. Gradient descent on neural networks typically occurs at the edge of stability. In *International Conference on Learning Representations*, 2021. 1, 9, 14
- Alex Damian, Eshaan Nichani, and Jason D. Lee. Self-stabilization: The implicit bias of gradient descent at the edge of stability. In *The Eleventh International Conference on Learning Representations*, 2023. 5
- James Allen Fill and Donniell E Fishkind. The moore–penrose generalized inverse for sums of matrices. *SIAM Journal on Matrix Analysis and Applications*, 21(2):629–635, 2000. 34
- Justin Gilmer, Behrooz Ghorbani, Ankush Garg, Sneha Kudugunta, Behnam Neyshabur, David Cardoze, George Edward Dahl, Zachary Nado, and Orhan Firat. A loss curvature perspective on training instabilities of deep learning models. In *International Conference on Learning Representations*, 2022. 1, 9
- Diego Granzio, Stefan Zohren, and Stephen Roberts. Learning rates as a function of batch size: A random matrix theory approach to neural network training. *J. Mach. Learn. Res.*, 23:1–65, 2022. 1, 8
- Sepp Hochreiter and Jürgen Schmidhuber. Flat minima. *Neural Computation*, 9(1):1–42, 1997. 1
- Stanisław Jastrzębski, Zachary Kenton, Devansh Arpit, Nicolas Ballas, Asja Fischer, Yoshua Bengio, and Amos Storkey. Three factors influencing minima in SGD. *arXiv preprint arXiv:1711.04623*, 2017. 1
- Stanisław Jastrzębski, Zachary Kenton, Nicolas Ballas, Asja Fischer, Yoshua Bengio, and Amos Storkey. On the relation between the sharpest directions of DNN loss and the SGD step length. In *International Conference on Learning Representations*, 2019. 1, 9
- Stanisław Jastrzębski, Maciej Szymczak, Stanislav Fort, Devansh Arpit, Jacek Tabor, Kyunghyun Cho*, and Krzysztof Geras*. The break-even point on optimization trajectories of deep neural networks. In *International Conference on Learning Representations*, 2020. 1, 9
- Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv preprint arXiv:1609.04836*, 2016. 1
- Yann LeCun. The MNIST database of handwritten digits. 1998. 2
- Sungyoon Lee and Cheongjae Jang. A new characterization of the edge of stability based on a sharpness measure aware of batch gradient distribution. In *The Eleventh International Conference on Learning Representations*, 2023. 9
- Kangqiao Liu, Liu Ziyin, and Masahito Ueda. Noise and fluctuation of finite learning rate stochastic gradient descent. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 7045–7056. PMLR, 18–24 Jul 2021. 9
- Chao Ma and Lexing Ying. On linear stability of SGD and input-smoothness of neural networks. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021. 1, 2, 3, 6, 9, 19, 31
- Rotem Mulayoff and Tomer Michaeli. Unique properties of flat minima in deep networks. In *International Conference on Machine Learning*, pages 7108–7118. PMLR, 2020. 1
- Rotem Mulayoff, Tomer Michaeli, and Daniel Soudry. The implicit bias of minima stability: A view from function space. *Advances in Neural Information Processing Systems*, 34:17749–17761, 2021. 1, 2, 4, 9, 14
- Mor Shpigel Nacson, Rotem Mulayoff, Greg Ongie, Tomer Michaeli, and Daniel Soudry. The implicit bias of minima stability in multivariate shallow ReLU networks. In *The Eleventh International Conference on Learning Representations*, 2023. 1

- Maksim Velikanov, Denis Kuznedelev, and Dmitry Yarotsky. A view of mini-batch SGD via generating functions: conditions of convergence, phase transitions, benefit from negative momenta. In *The Eleventh International Conference on Learning Representations*, 2023. 1, 8
- Lei Wu, Zhanxing Zhu, *et al.* Towards understanding generalization of deep learning: Perspective of loss landscapes. *arXiv preprint arXiv:1706.10239*, 2017. 1
- Lei Wu, Chao Ma, and E Weinan. How SGD selects the global minima in over-parameterized learning: A dynamical stability perspective. In *Advances in Neural Information Processing Systems*, pages 8279–8288, 2018. 1, 2, 4, 8
- Lei Wu, Mingze Wang, and Weijie Su. The alignment property of sgd noise and how it helps select flat minima: A stability analysis. *Advances in Neural Information Processing Systems*, 35:4680–4693, 2022. 9
- Liu Ziyin, Kangqiao Liu, Takashi Mori, and Masahito Ueda. Strength of minibatch noise in SGD. In *International Conference on Learning Representations*, 2022. 9
- Liu Ziyin, Botao Li, Tomer Galanti, and Masahito Ueda. The probabilistic stability of stochastic gradient descent. *arXiv preprint arXiv:2303.13093*, 2023. 1, 4, 9
- Bruno Zumino. Normal forms of complex matrices. *Journal of Mathematical Physics*, 3(5):1055–1057, 1962. 26

A NOTATIONS AND THE KRONECKER PRODUCT

Throughout our derivations, we use the following notations.

a	Lower case non-bold letters for scalars
\mathbf{a}	Lower case bold letters for vectors
\mathbf{A}	Upper case bold for matrices
\mathbf{A}^\dagger	Moore–Penrose inverse of \mathbf{A}
$\mathbf{P}_\mathcal{V}$	Projection matrix onto the subspace \mathcal{V}
$\mathcal{N}(\mathbf{A})$	Null space of \mathbf{A}
$\mathcal{N}^\perp(\mathbf{A})$	Orthogonal complement of the null space of \mathbf{A}
$\mathcal{R}(\mathbf{A})$	Range of \mathbf{A}
$\mathcal{R}^\perp(\mathbf{A})$	Orthogonal complement of the range of \mathbf{A}
\otimes	Kronecker product
\oplus	Kronecker sum
\odot^k	k 'th Hadamard power
\mathbb{E}	Expectation
\mathbb{P}	Probability
$\ \mathbf{a}\ $	Euclidean norm of \mathbf{a}
$\ \mathbf{A}\ $	Top singular value of \mathbf{A}
$\ \mathbf{A}\ _F$	Frobenius norm of \mathbf{A}
$\rho(\mathbf{A})$	Spectral radius of \mathbf{A}
$\text{vec}(\mathbf{A})$	Vectorization of \mathbf{A} (column stack)
$\text{vec}^{-1}(\mathbf{a})$	Reshaping \mathbf{a} back to $d \times d$ matrix
\mathcal{L}	Loss function
$\boldsymbol{\theta}$	Parameters vector of the loss
$\boldsymbol{\theta}^*$	Minimum point of the loss
d	Dimension of $\boldsymbol{\theta}$
n	Number of training samples
η	Step size
B	Batch size
p	Defined to be $(n - B)/(B(n - 1))$
$\boldsymbol{\Sigma}$	Second moment matrix
\mathbf{H}	Hessian of the full loss at $\boldsymbol{\theta}^*$
\mathbf{H}_i	Hessian of the loss of the sample i at $\boldsymbol{\theta}^*$
\mathbf{g}_i	Gradient of the loss of the sample i at $\boldsymbol{\theta}^*$
\mathbf{a}^\parallel	Projection of \mathbf{a} onto the null space of \mathbf{H}
\mathbf{a}^\perp	Projection of \mathbf{a} onto the orthogonal complement of the null space of \mathbf{H}
$\mathcal{S}_+(\mathbb{R}^{d \times d})$	The set of all positive semi-definite (PSD) matrices over $\mathbb{R}^{d \times d}$
\mathbb{S}^{d-1}	Unit sphere in \mathbb{R}^d

Table 1: Table of notations

Further notations that we use are given below.

$$\begin{aligned}
\boldsymbol{\mu}_t &\triangleq \mathbb{E}[\boldsymbol{\theta}_t - \boldsymbol{\theta}^*], & \boldsymbol{\Sigma}_t &\triangleq \mathbb{E}\left[(\boldsymbol{\theta}_t - \boldsymbol{\theta}^*)(\boldsymbol{\theta}_t - \boldsymbol{\theta}^*)^\top\right], \\
\boldsymbol{\mu}_t^\perp &\triangleq \mathbb{E}[\boldsymbol{\theta}_t^\perp - \boldsymbol{\theta}^{*\perp}], & \boldsymbol{\Sigma}_t^\perp &\triangleq \mathbb{E}\left[(\boldsymbol{\theta}_t^\perp - \boldsymbol{\theta}^{*\perp})(\boldsymbol{\theta}_t^\perp - \boldsymbol{\theta}^{*\perp})^\top\right], \\
\boldsymbol{\mu}_t^\parallel &\triangleq \mathbb{E}[\boldsymbol{\theta}_t^\parallel - \boldsymbol{\theta}^{*\parallel}], & \boldsymbol{\Sigma}_t^\parallel &\triangleq \mathbb{E}\left[(\boldsymbol{\theta}_t^\parallel - \boldsymbol{\theta}^{*\parallel})(\boldsymbol{\theta}_t^\parallel - \boldsymbol{\theta}^{*\parallel})^\top\right].
\end{aligned} \tag{26}$$

Additionally, we make extensive use of the following properties of the Kronecker product. For any matrices M_1, M_2, M_3, M_4 ,

$$\text{vec}(M_1 M_2 M_3) = (M_3^T \otimes M_1) \text{vec}(M_2), \quad (\text{P1})$$

$$(M_1 \otimes M_2)^T = (M_1^T \otimes M_2^T), \quad (\text{P2})$$

$$(M_1 \otimes M_2)(M_3 \otimes M_4) = (M_1 M_3) \otimes (M_2 M_4), \quad (\text{P3})$$

$$[\text{vec}(M_1)]^T (M_2 \otimes M_3) \text{vec}(M_4) = \text{Tr}(M_1^T M_3 M_4 M_2^T). \quad (\text{P4})$$

Finally, we give here the definition of Kronecker sum. If M_1 is $d_1 \times d_1$, M_2 is $d_2 \times d_2$ and I_d denotes the $d \times d$ identity matrix then

$$M_1 \oplus M_2 = M_1 \otimes I_{d_2} + I_{d_1} \otimes M_2. \quad (27)$$

B STABILITY OF THE FIRST AND SECOND MOMENTS

Using our notation (see App. A), for all $v \in \mathbb{R}^d$ we have $v^\perp = P_{\mathcal{N}^\perp(\mathbf{H})}v$ and $v^\parallel = P_{\mathcal{N}(\mathbf{H})}v$. Since \mathbf{H} is symmetric,

$$P_{\mathcal{N}(\mathbf{H})}\mathbf{H} = \mathbf{H}P_{\mathcal{N}(\mathbf{H})} = \mathbf{0}, \quad \text{and} \quad P_{\mathcal{N}^\perp(\mathbf{H})}\mathbf{H} = \mathbf{H}P_{\mathcal{N}^\perp(\mathbf{H})} = \mathbf{H}. \quad (28)$$

If $\mathbf{H}_i \in \mathcal{S}_+(\mathbb{R}^{d \times d})$ for all $i \in [n]$, then the null space of \mathbf{H} is contained in the null space of each \mathbf{H}_i , and therefore we also have that

$$P_{\mathcal{N}(\mathbf{H})}\mathbf{H}_i = \mathbf{H}_i P_{\mathcal{N}(\mathbf{H})} = \mathbf{0}, \quad \text{and} \quad P_{\mathcal{N}^\perp(\mathbf{H})}\mathbf{H}_i = \mathbf{H}_i P_{\mathcal{N}^\perp(\mathbf{H})} = \mathbf{H}_i. \quad (29)$$

B.1 LINEARIZED DYNAMICS

The linearized dynamics near θ^* is

$$\theta_{t+1} = \theta_t - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{H}_i (\theta_t - \theta^*) - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{g}_i. \quad (30)$$

Therefore,

$$\begin{aligned} \theta_{t+1} - \theta^* &= \theta_t - \theta^* - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{H}_i (\theta_t - \theta^*) - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{g}_i \\ &= \left(\mathbf{I} - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{H}_i \right) (\theta_t - \theta^*) - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{g}_i. \end{aligned} \quad (31)$$

Here we assume that the batches are chosen uniformly at random, independently across iterations.

Linearized dynamics in the orthogonal complement. Under the assumption that $\mathbf{H}_i \in \mathcal{S}_+(\mathbb{R}^{d \times d})$ for all $i \in [n]$, the linearized dynamics in the orthogonal complement is given by

$$\begin{aligned} \theta_{t+1}^\perp - \theta^{*\perp} &= P_{\mathcal{N}^\perp(\mathbf{H})}(\theta_{t+1} - \theta^*) \\ &= P_{\mathcal{N}^\perp(\mathbf{H})} \left(\mathbf{I} - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{H}_i \right) (\theta_t - \theta^*) - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} P_{\mathcal{N}^\perp(\mathbf{H})} \mathbf{g}_i \\ &= \left(P_{\mathcal{N}^\perp(\mathbf{H})} - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} P_{\mathcal{N}^\perp(\mathbf{H})} \mathbf{H}_i \right) (\theta_t - \theta^*) - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{g}_i^\perp \\ &= \left(P_{\mathcal{N}^\perp(\mathbf{H})} - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{H}_i P_{\mathcal{N}^\perp(\mathbf{H})} \right) (\theta_t - \theta^*) - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{g}_i^\perp \\ &= \left(\mathbf{I} - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{H}_i \right) P_{\mathcal{N}^\perp(\mathbf{H})} (\theta_t - \theta^*) - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{g}_i^\perp \\ &= \left(\mathbf{I} - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{H}_i \right) (\theta_{t+1}^\perp - \theta^{*\perp}) - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{g}_i^\perp. \end{aligned} \quad (32)$$

Linearized dynamics in the null space. Under the assumption that $\mathbf{H}_i \in \mathcal{S}_+(\mathbb{R}^{d \times d})$ for all $i \in [n]$, the linearized dynamics in the null space is given by

$$\begin{aligned}
\boldsymbol{\theta}_{t+1}^{\parallel} - \boldsymbol{\theta}^{*\parallel} &= \mathbf{P}_{\mathcal{N}(\mathbf{H})}(\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}^*) \\
&= \mathbf{P}_{\mathcal{N}(\mathbf{H})} \left(\mathbf{I} - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{H}_i \right) (\boldsymbol{\theta}_t - \boldsymbol{\theta}^*) - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{P}_{\mathcal{N}(\mathbf{H})} \mathbf{g}_i \\
&= \left(\mathbf{P}_{\mathcal{N}(\mathbf{H})} - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{P}_{\mathcal{N}(\mathbf{H})} \mathbf{H}_i \right) (\boldsymbol{\theta}_t - \boldsymbol{\theta}^*) - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{g}_i^{\parallel} \\
&= (\mathbf{P}_{\mathcal{N}(\mathbf{H})} - \mathbf{0}) (\boldsymbol{\theta}_t - \boldsymbol{\theta}^*) - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{g}_i^{\parallel} \\
&= \mathbf{P}_{\mathcal{N}(\mathbf{H})} (\boldsymbol{\theta}_t - \boldsymbol{\theta}^*) - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{g}_i^{\parallel} \\
&= \boldsymbol{\theta}_t^{\parallel} - \boldsymbol{\theta}^{*\parallel} - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{g}_i^{\parallel}. \tag{33}
\end{aligned}$$

Namely,

$$\boldsymbol{\theta}_{t+1}^{\parallel} = \boldsymbol{\theta}_t^{\parallel} - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{g}_i^{\parallel}. \tag{34}$$

Note that if $\mathbf{g}_i^{\parallel} = \mathbf{0}$ for all $i \in [n]$ then

$$\boldsymbol{\theta}_{t+1}^{\parallel} = \boldsymbol{\theta}_t^{\parallel}. \tag{35}$$

B.2 MEAN DYNAMICS (PROOF OF THEOREM 1)

First, we compute the mean of the linearized dynamics.

$$\begin{aligned}
\boldsymbol{\mu}_{t+1} &= \mathbb{E}[\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}^*] = \mathbb{E} \left[\left(\mathbf{I} - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{H}_i \right) (\boldsymbol{\theta}_t - \boldsymbol{\theta}^*) \right] - \mathbb{E} \left[\frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{g}_i \right] \\
&= \mathbb{E} \left[\left(\mathbf{I} - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{H}_i \right) \mathbb{E}[(\boldsymbol{\theta}_t - \boldsymbol{\theta}^*) | \mathfrak{B}_t] \right] \\
&= (\mathbf{I} - \eta \mathbf{H}) \mathbb{E}[(\boldsymbol{\theta}_t - \boldsymbol{\theta}^*)] \\
&= (\mathbf{I} - \eta \mathbf{H}) \boldsymbol{\mu}_t, \tag{36}
\end{aligned}$$

where in the second line we used the law of total expectation, as well as (6). This system is stable if and only if the spectral radius $\rho(\mathbf{I} - \eta \mathbf{H}) \leq 1$. This condition is equivalent to $\lambda_{\max}(\mathbf{H}) \leq 2/\eta$ (see proof in, e.g., Cohen *et al.* (2021); Mulayoff *et al.* (2021)), thus proving point 2 of Theorem 1.

Mean dynamics in the orthogonal complement. In a similar manner, taking the expectation of both sides of (32) and using (6), we get

$$\boldsymbol{\mu}_{t+1}^{\perp} = (\mathbf{I} - \eta \mathbf{H}) \boldsymbol{\mu}_t^{\perp}, \tag{37}$$

Note that for all $t \geq 0$,

$$\boldsymbol{\mu}_{t+1}^{\perp} = \mathbf{P}_{\mathcal{N}^{\perp}(\mathbf{H})} \boldsymbol{\mu}_{t+1}^{\perp} = \mathbf{P}_{\mathcal{N}^{\perp}(\mathbf{H})} (\mathbf{I} - \eta \mathbf{H}) \boldsymbol{\mu}_t^{\perp} = (\mathbf{P}_{\mathcal{N}^{\perp}(\mathbf{H})} - \eta \mathbf{H}) \boldsymbol{\mu}_t^{\perp}. \tag{38}$$

Namely, $\boldsymbol{\mu}_t^{\perp} = (\mathbf{P}_{\mathcal{N}^{\perp}(\mathbf{H})} - \eta \mathbf{H})^t \boldsymbol{\mu}_0^{\perp}$, and thus

$$\|\boldsymbol{\mu}_t^{\perp}\| = \left\| (\mathbf{P}_{\mathcal{N}^{\perp}(\mathbf{H})} - \eta \mathbf{H})^t \boldsymbol{\mu}_0^{\perp} \right\| \leq \|\mathbf{P}_{\mathcal{N}^{\perp}(\mathbf{H})} - \eta \mathbf{H}\|^t \|\boldsymbol{\mu}_0^{\perp}\|. \tag{39}$$

It is easy to show that

$$\|\mathbf{P}_{\mathcal{N}^{\perp}(\mathbf{H})} - \eta \mathbf{H}\| = \max_{\lambda_i(\mathbf{H}) \neq 0} \{|1 - \eta \lambda_i(\mathbf{H})|\}. \tag{40}$$

Therefore, if $0 < \eta < 2/\lambda_{\max}$, we have that $\|\mathbf{P}_{\mathcal{N}^{\perp}(\mathbf{H})} - \eta \mathbf{H}\| < 1$ and thus

$$\lim_{t \rightarrow \infty} \|\boldsymbol{\mu}_t^{\perp}\| \leq \lim_{t \rightarrow \infty} \|\mathbf{P}_{\mathcal{N}^{\perp}(\mathbf{H})} - \eta \mathbf{H}\|^t \|\boldsymbol{\mu}_0^{\perp}\| = 0. \tag{41}$$

This demonstrates point 3 of Theorem 1.

Mean dynamics in the null space. Taking the expectation of both sides of (33) and using (6), we obtain

$$\boldsymbol{\mu}_{t+1}^{\parallel} = \boldsymbol{\mu}_t^{\parallel}. \quad (42)$$

This demonstrates that for all $t \geq 0$,

$$\mathbb{E} [\boldsymbol{\theta}_t^{\parallel} - \boldsymbol{\theta}^{*\parallel}] = \boldsymbol{\mu}_t^{\parallel} = \boldsymbol{\mu}_0^{\parallel} = \mathbb{E} [\boldsymbol{\theta}_0^{\parallel} - \boldsymbol{\theta}^{*\parallel}], \quad (43)$$

so that

$$\mathbb{E} [\boldsymbol{\theta}_t^{\parallel}] = \mathbb{E} [\boldsymbol{\theta}_0^{\parallel}]. \quad (44)$$

This proves Point 1 of Theorem 1.

B.3 COVARIANCE DYNAMICS FOR THE ORTHOGONAL COMPLEMENT

Before providing a complete proof for Theorem 3 (see App. B.7), we next examine the evolution over time of the covariance of the parameter vector. We start by focusing on the orthogonal complement space. Define

$$\mathbf{A}_t = \mathbf{I} - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{H}_i \quad \text{and} \quad \mathbf{v}_t = \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{g}_i, \quad (45)$$

so that (32) can be compactly written as

$$\boldsymbol{\theta}_{t+1}^{\perp} - \boldsymbol{\theta}^{*\perp} = \mathbf{A}_t (\boldsymbol{\theta}_t^{\perp} - \boldsymbol{\theta}^{*\perp}) - \mathbf{v}_t^{\perp}. \quad (46)$$

Recall that this holds under the assumption that $\mathbf{H}_i \in \mathcal{S}_+(\mathbb{R}^{d \times d})$ for all $i \in [n]$. Note that $\{\mathbf{A}_t\}$ are i.i.d. and that $\boldsymbol{\theta}_t^{\perp}$ is constructed from $\mathbf{A}_0, \dots, \mathbf{A}_{t-1}$, so that $\boldsymbol{\theta}_t^{\perp}$ is independent of \mathbf{A}_t . We therefore have

$$\begin{aligned} \boldsymbol{\Sigma}_{t+1}^{\perp} &= \mathbb{E} \left[(\boldsymbol{\theta}_{t+1}^{\perp} - \boldsymbol{\theta}^{*\perp}) (\boldsymbol{\theta}_{t+1}^{\perp} - \boldsymbol{\theta}^{*\perp})^{\top} \right] \\ &= \mathbb{E} \left[(\mathbf{A}_t (\boldsymbol{\theta}_t^{\perp} - \boldsymbol{\theta}^{*\perp}) - \mathbf{v}_t^{\perp}) (\mathbf{A}_t (\boldsymbol{\theta}_t^{\perp} - \boldsymbol{\theta}^{*\perp}) - \mathbf{v}_t^{\perp})^{\top} \right] \\ &= \mathbb{E} \left[\mathbf{A}_t (\boldsymbol{\theta}_t^{\perp} - \boldsymbol{\theta}^{*\perp}) (\boldsymbol{\theta}_t^{\perp} - \boldsymbol{\theta}^{*\perp})^{\top} \mathbf{A}_t^{\top} \right] - \mathbb{E} \left[\mathbf{A}_t (\boldsymbol{\theta}_t^{\perp} - \boldsymbol{\theta}^{*\perp}) (\mathbf{v}_t^{\perp})^{\top} \right] \\ &\quad - \mathbb{E} \left[\mathbf{v}_t^{\perp} (\boldsymbol{\theta}_t^{\perp} - \boldsymbol{\theta}^{*\perp})^{\top} \mathbf{A}_t^{\top} \right] + \mathbb{E} \left[\mathbf{v}_t^{\perp} (\mathbf{v}_t^{\perp})^{\top} \right] \\ &= \mathbb{E} \left[\mathbf{A}_t \mathbb{E} \left[(\boldsymbol{\theta}_t^{\perp} - \boldsymbol{\theta}^{*\perp}) (\boldsymbol{\theta}_t^{\perp} - \boldsymbol{\theta}^{*\perp})^{\top} \right] \mathbf{A}_t^{\top} \right] - \mathbb{E} \left[\mathbf{A}_t \mathbb{E} [\boldsymbol{\theta}_t^{\perp} - \boldsymbol{\theta}^{*\perp}] (\mathbf{v}_t^{\perp})^{\top} \right] \\ &\quad - \mathbb{E} \left[\mathbf{v}_t^{\perp} \mathbb{E} [\boldsymbol{\theta}_t^{\perp} - \boldsymbol{\theta}^{*\perp}]^{\top} \mathbf{A}_t^{\top} \right] + \boldsymbol{\Sigma}_{\mathbf{v}}^{\perp} \\ &= \mathbb{E} \left[\mathbf{A}_t \boldsymbol{\Sigma}_t^{\perp} \mathbf{A}_t^{\top} \right] - \mathbb{E} \left[\mathbf{A}_t \boldsymbol{\mu}_t^{\perp} (\mathbf{v}_t^{\perp})^{\top} \right] - \mathbb{E} \left[\mathbf{v}_t^{\perp} (\boldsymbol{\mu}_t^{\perp})^{\top} \mathbf{A}_t^{\top} \right] + \boldsymbol{\Sigma}_{\mathbf{v}}^{\perp}, \end{aligned} \quad (47)$$

where in the second equality we used the fact that \mathbf{A}_t is independent of $\boldsymbol{\theta}_t^{\perp}$. Using vectorization, the above equation can be written as

$$\begin{aligned} \text{vec} (\boldsymbol{\Sigma}_{t+1}^{\perp}) &= \mathbb{E} \left[\text{vec} \left(\mathbf{A}_t \boldsymbol{\Sigma}_t^{\perp} \mathbf{A}_t^{\top} \right) \right] - \mathbb{E} \left[\text{vec} \left(\mathbf{A}_t \boldsymbol{\mu}_t^{\perp} (\mathbf{v}_t^{\perp})^{\top} \right) \right] - \mathbb{E} \left[\text{vec} \left(\mathbf{v}_t^{\perp} (\boldsymbol{\mu}_t^{\perp})^{\top} \mathbf{A}_t^{\top} \right) \right] + \text{vec} (\boldsymbol{\Sigma}_{\mathbf{v}}^{\perp}) \\ &= \mathbb{E} \left[\mathbf{A}_t \otimes \mathbf{A}_t \right] \text{vec} (\boldsymbol{\Sigma}_t^{\perp}) - \mathbb{E} \left[\mathbf{v}_t^{\perp} \otimes \mathbf{A}_t \right] \boldsymbol{\mu}_t^{\perp} - \mathbb{E} \left[\mathbf{A}_t \otimes \mathbf{v}_t^{\perp} \right] \boldsymbol{\mu}_t^{\perp} + \text{vec} (\boldsymbol{\Sigma}_{\mathbf{v}}^{\perp}) \\ &= \mathbf{Q} \text{vec} (\boldsymbol{\Sigma}_t^{\perp}) - (\mathbb{E} [\mathbf{v}_t^{\perp} \otimes \mathbf{A}_t] + \mathbb{E} [\mathbf{A}_t \otimes \mathbf{v}_t^{\perp}]) \boldsymbol{\mu}_t^{\perp} + \text{vec} (\boldsymbol{\Sigma}_{\mathbf{v}}^{\perp}), \end{aligned} \quad (48)$$

where we denoted

$$\mathbf{Q} \triangleq \mathbb{E} [\mathbf{A}_t \otimes \mathbf{A}_t]. \quad (49)$$

Overall, the joint dynamics of $\boldsymbol{\Sigma}_t^{\perp}$ and $\boldsymbol{\mu}_t^{\perp}$ is given by

$$\begin{pmatrix} \boldsymbol{\mu}_{t+1}^{\perp} \\ \text{vec} (\boldsymbol{\Sigma}_{t+1}^{\perp}) \end{pmatrix} = \begin{pmatrix} \mathbf{I} - \eta \mathbf{H} & \mathbf{0} \\ -\mathbb{E} [\mathbf{v}_t^{\perp} \otimes \mathbf{A}_t] - \mathbb{E} [\mathbf{A}_t \otimes \mathbf{v}_t^{\perp}] & \mathbf{Q} \end{pmatrix} \begin{pmatrix} \boldsymbol{\mu}_t^{\perp} \\ \text{vec} (\boldsymbol{\Sigma}_t^{\perp}) \end{pmatrix} + \begin{pmatrix} \mathbf{0} \\ \text{vec} (\boldsymbol{\Sigma}_{\mathbf{v}}^{\perp}) \end{pmatrix}. \quad (50)$$

In some cases, it is easier to look at a projected version of the transition matrix. In (38) we showed that

$$\boldsymbol{\mu}_{t+1}^{\perp} = (\mathbf{P}_{\mathcal{N}^{\perp}(\mathbf{H})} - \eta \mathbf{H}) \boldsymbol{\mu}_t^{\perp}. \quad (51)$$

Moreover, from (48),

$$\begin{aligned}
\text{vec}(\boldsymbol{\Sigma}_{t+1}^\perp) &= \text{vec}(\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \boldsymbol{\Sigma}_{t+1}^\perp \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})}) \\
&= (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})}) \text{vec}(\boldsymbol{\Sigma}_{t+1}^\perp) \\
&= (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})}) (\mathbf{Q} \text{vec}(\boldsymbol{\Sigma}_t^\perp) - (\mathbb{E}[\mathbf{v}_t^\perp \otimes \mathbf{A}_t] + \mathbb{E}[\mathbf{A}_t \otimes \mathbf{v}_t^\perp]) \boldsymbol{\mu}_t^\perp + \text{vec}(\boldsymbol{\Sigma}_v^\perp)) \\
&= (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})}) \mathbf{Q} \text{vec}(\boldsymbol{\Sigma}_t^\perp) \\
&\quad - (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})}) (\mathbb{E}[\mathbf{v}_t^\perp \otimes \mathbf{A}_t] + \mathbb{E}[\mathbf{A}_t \otimes \mathbf{v}_t^\perp]) \boldsymbol{\mu}_t^\perp + \text{vec}(\boldsymbol{\Sigma}_v^\perp). \tag{52}
\end{aligned}$$

Therefore, the linear system in (50) can be written as

$$\begin{aligned}
\begin{pmatrix} \boldsymbol{\mu}_{t+1}^\perp \\ \text{vec}(\boldsymbol{\Sigma}_{t+1}^\perp) \end{pmatrix} &= \\
&\begin{pmatrix} \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} - \eta \mathbf{H} & \mathbf{0} \\ -(\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})}) (\mathbb{E}[\mathbf{v}_t^\perp \otimes \mathbf{A}_t] + \mathbb{E}[\mathbf{A}_t \otimes \mathbf{v}_t^\perp]) & (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})}) \mathbf{Q} \end{pmatrix} \begin{pmatrix} \boldsymbol{\mu}_t^\perp \\ \text{vec}(\boldsymbol{\Sigma}_t^\perp) \end{pmatrix} \\
&+ \begin{pmatrix} \mathbf{0} \\ \text{vec}(\boldsymbol{\Sigma}_v^\perp) \end{pmatrix}. \tag{53}
\end{aligned}$$

B.4 THE TRANSITION MATRIX OF THE COVARIANCE DYNAMICS

We now proceed to develop an explicit expression for the covariance transition matrix \mathbf{Q} of (49). We have

$$\begin{aligned}
\mathbb{E}[\mathbf{A}_t \otimes \mathbf{A}_t] &= \mathbb{E} \left[\left(\mathbf{I} - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{H}_i \right) \otimes \left(\mathbf{I} - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{H}_i \right) \right] \\
&= \mathbb{E} \left[\mathbf{I} - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} (\mathbf{I} \otimes \mathbf{H}_i + \mathbf{H}_i \otimes \mathbf{I}) + \frac{\eta^2}{B^2} \sum_{i,j \in \mathfrak{B}_t} \mathbf{H}_i \otimes \mathbf{H}_j \right] \\
&= \mathbf{I} - \eta (\mathbf{I} \otimes \mathbf{H} + \mathbf{H} \otimes \mathbf{I}) + \eta^2 \mathbb{E} \left[\frac{1}{B^2} \sum_{i,j \in \mathfrak{B}_t} \mathbf{H}_i \otimes \mathbf{H}_j \right]. \tag{54}
\end{aligned}$$

Note that

$$\begin{aligned}
\mathbb{E} \left[\frac{1}{B^2} \sum_{i,j \in \mathfrak{B}_t} \mathbf{H}_i \otimes \mathbf{H}_j \right] &= \mathbb{E} \left[\frac{1}{B^2} \sum_{i \neq j \in \mathfrak{B}_t} \mathbf{H}_i \otimes \mathbf{H}_j + \frac{1}{B^2} \sum_{i \in \mathfrak{B}_t} \mathbf{H}_i \otimes \mathbf{H}_i \right] \\
&= \frac{1}{B^2} \times B(B-1) \mathbb{E}[\mathbf{H}_i \otimes \mathbf{H}_j | i \neq j \in \mathfrak{B}_t] + \frac{1}{B^2} \mathbb{E} \left[\sum_{i \in \mathfrak{B}_t} \mathbf{H}_i \otimes \mathbf{H}_i \right] \\
&= \frac{B-1}{B} \mathbb{E}[\mathbf{H}_i \otimes \mathbf{H}_j | i \neq j \in \mathfrak{B}_t] + \frac{1}{nB} \sum_{i=1}^n \mathbf{H}_i \otimes \mathbf{H}_i. \tag{55}
\end{aligned}$$

Specifically using symmetry and (7),

$$\begin{aligned}
\mathbb{E}[\mathbf{H}_i \otimes \mathbf{H}_j | i \neq j \in \mathfrak{B}_t] &= \sum_{i \neq j=1}^n \frac{1}{n(n-1)} \mathbf{H}_i \otimes \mathbf{H}_j \\
&= \frac{n}{(n-1)n^2} \sum_{i \neq j=1}^n \mathbf{H}_i \otimes \mathbf{H}_j \\
&= \frac{n}{(n-1)} \left(\mathbf{H} \otimes \mathbf{H} - \frac{1}{n^2} \sum_{i=1}^n \mathbf{H}_i \otimes \mathbf{H}_i \right). \tag{56}
\end{aligned}$$

Hence,

$$\begin{aligned}
\frac{B-1}{B} \mathbb{E} [\mathbf{H}_i \otimes \mathbf{H}_j | i \neq j \in \mathfrak{B}_t] &= \frac{n(B-1)}{B(n-1)} \left(\mathbf{H} \otimes \mathbf{H} - \frac{1}{n^2} \sum_{i=1}^n \mathbf{H}_i \otimes \mathbf{H}_i \right) \\
&= \frac{n(B-1)}{B(n-1)} \mathbf{H} \otimes \mathbf{H} - \frac{B-1}{Bn(n-1)} \sum_{i=1}^n \mathbf{H}_i \otimes \mathbf{H}_i \\
&= \mathbf{H} \otimes \mathbf{H} - \frac{n-B}{B(n-1)} \mathbf{H} \otimes \mathbf{H} - \frac{B-1}{Bn(n-1)} \sum_{i=1}^n \mathbf{H}_i \otimes \mathbf{H}_i.
\end{aligned} \tag{57}$$

Overall,

$$\begin{aligned}
\mathbb{E} \left[\frac{1}{B^2} \sum_{i,j \in \mathfrak{B}_t} \mathbf{H}_i \otimes \mathbf{H}_j \right] &= \mathbf{H} \otimes \mathbf{H} - \frac{n-B}{B(n-1)} \mathbf{H} \otimes \mathbf{H} - \frac{B-1}{Bn(n-1)} \sum_{i=1}^n \mathbf{H}_i \otimes \mathbf{H}_i \\
&\quad + \frac{1}{nB} \sum_{i=1}^n \mathbf{H}_i \otimes \mathbf{H}_i \\
&= \mathbf{H} \otimes \mathbf{H} - \frac{n-B}{B(n-1)} \mathbf{H} \otimes \mathbf{H} + \frac{n-B}{B(n-1)} \times \frac{1}{n} \sum_{i=1}^n \mathbf{H}_i \otimes \mathbf{H}_i \\
&= \mathbf{H} \otimes \mathbf{H} + \frac{n-B}{B(n-1)} \left(\frac{1}{n} \sum_{i=1}^n \mathbf{H}_i \otimes \mathbf{H}_i - \mathbf{H} \otimes \mathbf{H} \right).
\end{aligned} \tag{58}$$

Therefore, we have that \mathbf{Q} can be written as

$$\begin{aligned}
\mathbf{Q}(B, \eta) &= \mathbf{I} - \eta(\mathbf{I} \otimes \mathbf{H} + \mathbf{H} \otimes \mathbf{I}) + \eta^2 \mathbf{H} \otimes \mathbf{H} + \eta^2 \frac{n-B}{B(n-1)} \left(\frac{1}{n} \sum_{i=1}^n \mathbf{H}_i \otimes \mathbf{H}_i - \mathbf{H} \otimes \mathbf{H} \right) \\
&= \mathbf{I} - \eta(\mathbf{I} \otimes \mathbf{H} + \mathbf{H} \otimes \mathbf{I}) + \eta^2 \left[(1-p) \mathbf{H} \otimes \mathbf{H} + p \times \frac{1}{n} \sum_{i=1}^n \mathbf{H}_i \otimes \mathbf{H}_i \right] \\
&= (\mathbf{I} - \eta \mathbf{H}) \otimes (\mathbf{I} - \eta \mathbf{H}) + \eta^2 p \times \left(\frac{1}{n} \sum_{i=1}^n (\mathbf{H}_i - \mathbf{H}) \otimes (\mathbf{H}_i - \mathbf{H}) \right) \\
&= (1-p) \times (\mathbf{I} - \eta \mathbf{H}) \otimes (\mathbf{I} - \eta \mathbf{H}) + p \times \frac{1}{n} \sum_{i=1}^n (\mathbf{I} - \eta \mathbf{H}_i) \otimes (\mathbf{I} - \eta \mathbf{H}_i),
\end{aligned} \tag{59}$$

where

$$p = \frac{n-B}{B(n-1)}. \tag{60}$$

B.5 COVARIANCE MATRIX OF THE GRADIENT NOISE

We now develop an explicit expression for the covariance matrix of the gradient noise \mathbf{v} of (45). We have

$$\begin{aligned}
\boldsymbol{\Sigma}_{\mathbf{v}} = \mathbb{E} [\mathbf{v}_t \mathbf{v}_t^T] &= \left(\frac{\eta}{B} \right)^2 \mathbb{E} \left[\sum_{i,j \in \mathfrak{B}_t} \mathbf{g}_i \mathbf{g}_j^T \right] \\
&= \left(\frac{\eta}{B} \right)^2 \mathbb{E} \left[\sum_{i \neq j \in \mathfrak{B}_t} \mathbf{g}_i \mathbf{g}_j^T + \sum_{i \in \mathfrak{B}_t} \mathbf{g}_i \mathbf{g}_i^T \right] \\
&= \left(\frac{\eta}{B} \right)^2 \left(B(B-1) \mathbb{E} [\mathbf{g}_i \mathbf{g}_j^T | i \neq j \in \mathfrak{B}_t] + \frac{B}{n} \sum_{i=1}^n \mathbf{g}_i \mathbf{g}_i^T \right).
\end{aligned} \tag{61}$$

Observe that

$$\begin{aligned}
\mathbb{E} [\mathbf{g}_i \mathbf{g}_j^\top | i \neq j \in \mathfrak{B}_t] &= \sum_{i \neq j=1}^n \frac{1}{n(n-1)} \mathbf{g}_i \mathbf{g}_j^\top \\
&= \frac{1}{n(n-1)} \left(\sum_{i,j=1}^n \mathbf{g}_i \mathbf{g}_j^\top - \sum_{i=1}^n \mathbf{g}_i \mathbf{g}_i^\top \right) \\
&= \frac{1}{n(n-1)} \left(\left(\sum_{i=1}^n \mathbf{g}_i \right) \left(\sum_{i=1}^n \mathbf{g}_i \right)^\top - \sum_{i=1}^n \mathbf{g}_i \mathbf{g}_i^\top \right) \\
&= -\frac{1}{n(n-1)} \sum_{i=1}^n \mathbf{g}_i \mathbf{g}_i^\top, \tag{62}
\end{aligned}$$

where in the last step we used (6). Thus,

$$\begin{aligned}
\Sigma_v &= \left(\frac{\eta}{B} \right)^2 \left(\frac{B}{n} - \frac{B(B-1)}{n(n-1)} \right) \sum_{i=1}^n \mathbf{g}_i \mathbf{g}_i^\top \\
&= \left(\frac{\eta}{B} \right)^2 \times \frac{B(n-B)}{n(n-1)} \sum_{i=1}^n \mathbf{g}_i \mathbf{g}_i^\top \\
&= \eta^2 \frac{n-B}{B(n-1)} \times \frac{1}{n} \sum_{i=1}^n \mathbf{g}_i \mathbf{g}_i^\top \\
&= \eta^2 p \Sigma_g, \tag{63}
\end{aligned}$$

where we denoted

$$\Sigma_g = \frac{1}{n} \sum_{i=1}^n \mathbf{g}_i \mathbf{g}_i^\top. \tag{64}$$

B.6 THE NULL SPACES OF \mathbf{C} , \mathbf{D} AND \mathbf{E}

Let us now analyze the relation between the null spaces of \mathbf{C} , \mathbf{D} and \mathbf{E} . First, it is easy to see that under the assumption that $\mathbf{H}_i \in \mathcal{S}_+(\mathbb{R}^{d \times d})$ for all $i \in [n]$,

$$\mathcal{N}(\mathbf{H}) = \bigcap_{i=1}^n \mathcal{N}(\mathbf{H}_i), \tag{65}$$

where $\mathcal{N}(\cdot)$ denotes null space of a matrix. Here we show the following.

Lemma 1. Assume that $\mathbf{H}_i \in \mathcal{S}_+(\mathbb{R}^{d \times d})$ for all $i \in [n]$ and let

$$\begin{aligned}
\mathbf{C} &= \frac{1}{2} \mathbf{H} \oplus \mathbf{H}, \\
\mathbf{D} &= (1-p) \mathbf{H} \otimes \mathbf{H} + p \frac{1}{n} \sum_{i=1}^n \mathbf{H}_i \otimes \mathbf{H}_i, \\
\mathbf{E} &= \frac{1}{n} \sum_{i=1}^n (\mathbf{H}_i - \mathbf{H}) \otimes (\mathbf{H}_i - \mathbf{H}). \tag{66}
\end{aligned}$$

Then $\mathcal{N}(\mathbf{C}) \subseteq \mathcal{N}(\mathbf{D})$ and $\mathcal{N}(\mathbf{C}) \subseteq \mathcal{N}(\mathbf{E})$.

Proof. Let $\mathbf{u} \in \mathcal{N}(\mathbf{C})$ and denote $\mathbf{U} = \text{vec}^{-1}(\mathbf{u})$, then

$$\mathbf{0} = 2\mathbf{C}\mathbf{u} = \mathbf{H} \oplus \mathbf{H}\mathbf{u} = (\mathbf{H} \otimes \mathbf{I} + \mathbf{I} \otimes \mathbf{H})\mathbf{u}. \tag{67}$$

In matrix form we get

$$\mathbf{U}\mathbf{H} + \mathbf{H}\mathbf{U} = \mathbf{0}. \tag{68}$$

Let us take the Frobenius norm, then

$$\|\mathbf{UH} + \mathbf{HU}\|_{\mathbb{F}}^2 = \|\mathbf{UH}\|_{\mathbb{F}}^2 + \|\mathbf{HU}\|_{\mathbb{F}}^2 + 2\text{Tr}((\mathbf{UH})^{\text{T}}\mathbf{HU}) = 0, \quad (69)$$

where

$$\text{Tr}((\mathbf{UH})^{\text{T}}\mathbf{HU}) = \text{Tr}(\mathbf{HU}^{\text{T}}\mathbf{HU}) = \text{Tr}\left(\mathbf{H}^{\frac{1}{2}}\mathbf{U}^{\text{T}}\mathbf{H}\mathbf{U}\mathbf{H}^{\frac{1}{2}}\right) \geq 0 \quad (70)$$

because $\mathbf{H}^{\frac{1}{2}}\mathbf{U}^{\text{T}}\mathbf{H}\mathbf{U}\mathbf{H}^{\frac{1}{2}} = (\mathbf{H}^{\frac{1}{2}}\mathbf{U}\mathbf{H}^{\frac{1}{2}})^{\text{T}}(\mathbf{H}^{\frac{1}{2}}\mathbf{U}\mathbf{H}^{\frac{1}{2}})$ is PSD. This implies that

$$\|\mathbf{UH}\|_{\mathbb{F}}^2 = \|\mathbf{HU}\|_{\mathbb{F}}^2 = 0. \quad (71)$$

Thus, $\mathbf{u} \in \mathcal{N}(\mathbf{C})$ if and only if $\mathbf{UH} = \mathbf{0}$ and $\mathbf{HU} = \mathbf{0}$. Since the null space of \mathbf{H} is the intersection of $\{\mathbf{H}_i\}$ we have that \mathbf{U} also satisfies $\mathbf{H}_i\mathbf{U} = \mathbf{U}\mathbf{H}_i = \mathbf{0}$ for all $i \in [n]$. Now,

$$\mathbf{D}\mathbf{u} = (1-p)\mathbf{H} \otimes \mathbf{H}\mathbf{u} + p \frac{1}{n} \sum_{i=1}^n \mathbf{H}_i \otimes \mathbf{H}_i\mathbf{u}, \quad (72)$$

and in matrix form,

$$\text{vec}^{-1}(\mathbf{D}\mathbf{u}) = (1-p)\mathbf{H}\mathbf{U}\mathbf{H} + p \frac{1}{n} \sum_{i=1}^n \mathbf{H}_i\mathbf{U}\mathbf{H}_i = \mathbf{0}. \quad (73)$$

Namely, $\mathbf{u} \in \mathcal{N}(\mathbf{D})$. Similarly,

$$\mathbf{E}\mathbf{u} = \frac{1}{n} \sum_{i=1}^n (\mathbf{H}_i - \mathbf{H}) \otimes (\mathbf{H}_i - \mathbf{H})\mathbf{u}, \quad (74)$$

and in matrix form,

$$\text{vec}^{-1}(\mathbf{E}\mathbf{u}) = \frac{1}{n} \sum_{i=1}^n (\mathbf{H}_i - \mathbf{H})\mathbf{U}(\mathbf{H}_i - \mathbf{H}) = \mathbf{0}. \quad (75)$$

Namely, $\mathbf{u} \in \mathcal{N}(\mathbf{E})$. ■

B.7 PROOF OF THEOREM 3

We are now ready to prove Theorem 3.

First statement. In (35) we showed that for interpolating minima $\boldsymbol{\theta}_{t+1}^{\parallel} = \boldsymbol{\theta}_t^{\parallel}$, which completes the proof for the first statement of the theorem.

Second statement. Ma & Ying (2021) showed that the second moment $\boldsymbol{\Sigma}_t = [(\boldsymbol{\theta}_t - \boldsymbol{\theta}^*)(\boldsymbol{\theta}_t - \boldsymbol{\theta}^*)^{\text{T}}]$ for interpolating minima evolves over time as

$$\text{vec}(\boldsymbol{\Sigma}_{t+1}) = \mathbf{Q} \text{vec}(\boldsymbol{\Sigma}_t), \quad (76)$$

where \mathbf{Q} is given in (10). Since $\boldsymbol{\Sigma}_t$ is PSD by definition, we only care about the effect of \mathbf{Q} on vectorizations of PSD matrices. Therefore, we have that $\{\boldsymbol{\Sigma}_t\}$ are bounded if and only if (see proof in (Ma & Ying, 2021))

$$\max_{\boldsymbol{\Sigma} \in \mathcal{S}_+(\mathbb{R}^{d \times d})} \frac{\|\mathbf{Q}(\eta, B) \text{vec}(\boldsymbol{\Sigma})\|}{\|\boldsymbol{\Sigma}\|_{\mathbb{F}}} \leq 1. \quad (77)$$

To obtain the second result of Thm. 3 we first rearrange the terms in \mathbf{Q} as

$$\mathbf{Q}(\eta, B) = (1-p) \times (\mathbf{I} - \eta\mathbf{H}) \otimes (\mathbf{I} - \eta\mathbf{H}) + p \times \frac{1}{n} \sum_{i=1}^n (\mathbf{I} - \eta\mathbf{H}_i) \otimes (\mathbf{I} - \eta\mathbf{H}_i). \quad (78)$$

By applying Theorem 6, we have that the spectral radius of \mathbf{Q} equals its top eigenvalue, and the corresponding eigenvector is a vectorization of a PSD matrix (see proof of Thm. 6 in App. C). Thus, we have that the maximizer for the constrained optimization problem in (77) is, in fact, the top

eigenvalue of \mathbf{Q} . Hence, the linear system is stable if and only if $\lambda_{\max}(\mathbf{Q}) \leq 1$. Regathering the terms of \mathbf{Q} gives

$$\begin{aligned}\mathbf{Q} &= \mathbf{I} - \eta(\mathbf{H} \oplus \mathbf{H}) + \eta^2 \left[(1-p)\mathbf{H} \otimes \mathbf{H} + p \times \frac{1}{n} \sum_{i=1}^n \mathbf{H}_i \otimes \mathbf{H}_i \right] \\ &= \mathbf{I} - 2\eta\mathbf{C} + \eta^2\mathbf{D}.\end{aligned}\quad (79)$$

Because \mathbf{Q} is symmetric, the condition $\lambda_{\max}(\mathbf{Q}) \leq 1$ is equivalent to the requirement that $\mathbf{u}^T \mathbf{Q} \mathbf{u} \leq 1$ for all $\mathbf{u} \in \mathbb{S}^{d^2-1}$. Recall from App. B.6 that $\mathcal{N}(\mathbf{C}) \subseteq \mathcal{N}(\mathbf{D})$. Therefore, if $\mathbf{u} \in \mathcal{N}(\mathbf{C})$ then also $\mathbf{u} \in \mathcal{N}(\mathbf{D})$ and we get

$$1 \geq \mathbf{u}^T \mathbf{Q} \mathbf{u} = 1 - 2\eta \mathbf{u}^T \mathbf{C} \mathbf{u} + \eta^2 \mathbf{u}^T \mathbf{D} \mathbf{u} = 1. \quad (80)$$

Namely, directions in the null space of \mathbf{C} do not impose any constraint on the learning rate, and thus can be ignored. Additionally, if $\mathbf{u} \in \mathcal{N}(\mathbf{D})$ but $\mathbf{u} \notin \mathcal{N}(\mathbf{C})$, then

$$1 \geq \mathbf{u}^T \mathbf{Q} \mathbf{u} = 1 - 2\eta \mathbf{u}^T \mathbf{C} \mathbf{u} + \eta^2 \mathbf{u}^T \mathbf{D} \mathbf{u} = 1 - 2\eta \mathbf{u}^T \mathbf{C} \mathbf{u}, \quad (81)$$

which is true for all $\eta \geq 0$, because \mathbf{C} is PSD. Now,

$$\mathbf{u}^T \mathbf{Q} \mathbf{u} = 1 - 2\eta \mathbf{u}^T \mathbf{C} \mathbf{u} + \eta^2 \mathbf{u}^T \mathbf{D} \mathbf{u} \leq 1 \quad (82)$$

happens for all $\mathbf{u} \notin \mathcal{N}(\mathbf{D})$ (which also results in $\mathbf{u} \notin \mathcal{N}(\mathbf{C})$), if and only if

$$0 \leq \eta \leq 2 \inf_{\mathbf{u} \in \mathbb{S}^{d^2-1}: \mathbf{u} \notin \mathcal{N}(\mathbf{D})} \left\{ \frac{\mathbf{u}^T \mathbf{C} \mathbf{u}}{\mathbf{u}^T \mathbf{D} \mathbf{u}} \right\}. \quad (83)$$

Therefore, the stability threshold η_{var}^* is given by

$$\begin{aligned}\eta_{\text{var}}^* &= 2 \inf_{\mathbf{u} \in \mathbb{S}^{d^2-1}: \mathbf{u} \notin \mathcal{N}(\mathbf{D})} \left\{ \frac{\mathbf{u}^T \mathbf{C} \mathbf{u}}{\mathbf{u}^T \mathbf{D} \mathbf{u}} \right\} \\ &= 2 \left(\sup_{\mathbf{u} \in \mathbb{S}^{d^2-1}: \mathbf{u} \notin \mathcal{N}(\mathbf{D})} \left\{ \frac{\mathbf{u}^T \mathbf{D} \mathbf{u}}{\mathbf{u}^T \mathbf{C} \mathbf{u}} \right\} \right)^{-1}.\end{aligned}\quad (84)$$

Note that \mathbf{u} 's norm has no effect, and therefore we can relax the constraint $\mathbf{u} \in \mathbb{S}^{d^2-1}$. Additionally, we can also reduce the constraint $\mathbf{u} \notin \mathcal{N}(\mathbf{D})$ to $\mathbf{u} \notin \mathcal{N}(\mathbf{C})$, because the supremum in (84) is over a non-negative function, and will not be affected by adding to the domain points at which the function vanishes. Since $\mathcal{N}(\mathbf{C}) \subseteq \mathcal{N}(\mathbf{D})$ we have that $\mathcal{N}^\perp(\mathbf{D}) \subseteq \mathcal{N}^\perp(\mathbf{C})$ and therefore

$$\mathbf{P}_{\mathcal{N}^\perp(\mathbf{C})} \mathbf{D} = \mathbf{D} \mathbf{P}_{\mathcal{N}^\perp(\mathbf{C})} = \mathbf{D}, \quad (85)$$

where $\mathbf{P}_{\mathcal{N}^\perp(\mathbf{C})}$ is the projection matrix onto the orthogonal complement of the null space of \mathbf{C} .

Additionally, \mathbf{C} is PSD, and therefore $\mathbf{C}^{\frac{1}{2}}$ exists and is also PSD, so that

$$\mathbf{P}_{\mathcal{N}^\perp(\mathbf{C})} = \left(\mathbf{C}^{\frac{1}{2}} \right)^\dagger \mathbf{C}^{\frac{1}{2}} = \mathbf{C}^{\frac{1}{2}} \left(\mathbf{C}^{\frac{1}{2}} \right)^\dagger. \quad (86)$$

Therefore,

$$\begin{aligned}\sup_{\mathbf{u} \notin \mathcal{N}(\mathbf{C})} \left\{ \frac{\mathbf{u}^T \mathbf{D} \mathbf{u}}{\mathbf{u}^T \mathbf{C} \mathbf{u}} \right\} &= \sup_{\mathbf{u} \notin \mathcal{N}(\mathbf{C})} \left\{ \frac{\mathbf{u}^T \mathbf{P}_{\mathcal{N}^\perp(\mathbf{C})} \mathbf{D} \mathbf{P}_{\mathcal{N}^\perp(\mathbf{C})} \mathbf{u}}{\mathbf{u}^T \mathbf{C} \mathbf{u}} \right\} \\ &= \sup_{\mathbf{u} \notin \mathcal{N}(\mathbf{C})} \left\{ \frac{\mathbf{u}^T \mathbf{C}^{\frac{1}{2}} \left(\mathbf{C}^{\frac{1}{2}} \right)^\dagger \mathbf{D} \left(\mathbf{C}^{\frac{1}{2}} \right)^\dagger \mathbf{C}^{\frac{1}{2}} \mathbf{u}}{\mathbf{u}^T \mathbf{C}^{\frac{1}{2}} \mathbf{C}^{\frac{1}{2}} \mathbf{u}} \right\} \\ &= \sup_{\mathbf{u} \notin \mathcal{N}(\mathbf{C})} \left\{ \frac{\left(\mathbf{C}^{\frac{1}{2}} \mathbf{u} \right)^T \left(\mathbf{C}^{\frac{1}{2}} \right)^\dagger \mathbf{D} \left(\mathbf{C}^{\frac{1}{2}} \right)^\dagger \left(\mathbf{C}^{\frac{1}{2}} \mathbf{u} \right)}{\left(\mathbf{C}^{\frac{1}{2}} \mathbf{u} \right)^T \left(\mathbf{C}^{\frac{1}{2}} \mathbf{u} \right)} \right\},\end{aligned}\quad (87)$$

where in the first step we used (85), and in the second step we used (86). By a simple change of variables $\mathbf{y} = \mathbf{C}^{\frac{1}{2}} \mathbf{u} \in \mathcal{N}^\perp(\mathbf{C})$ we get

$$\begin{aligned} \max_{\mathbf{y} \in \mathcal{N}^\perp(\mathbf{C})} \left\{ \frac{\mathbf{y}^\top (\mathbf{C}^{\frac{1}{2}})^\dagger \mathbf{D} (\mathbf{C}^{\frac{1}{2}})^\dagger \mathbf{y}}{\mathbf{y}^\top \mathbf{y}} \right\} &= \max_{\mathbf{y} \in \mathbb{R}^{d^2}} \left\{ \frac{\mathbf{y}^\top (\mathbf{C}^{\frac{1}{2}})^\dagger \mathbf{D} (\mathbf{C}^{\frac{1}{2}})^\dagger \mathbf{y}}{\mathbf{y}^\top \mathbf{y}} \right\} \\ &= \lambda_{\max} \left((\mathbf{C}^{\frac{1}{2}})^\dagger \mathbf{D} (\mathbf{C}^{\frac{1}{2}})^\dagger \right), \end{aligned} \quad (88)$$

where in the first step we used the fact that adding to \mathbf{y} a component in $\mathcal{N}(\mathbf{C})$ will increase the denominator by $\|\mathbf{P}_{\mathcal{N}(\mathbf{C})} \mathbf{y}\|^2$ but will not affect the numerator. Namely, the optimum cannot be attained by $\mathbf{y} \notin \mathcal{N}^\perp(\mathbf{C})$. Now, let $(\lambda_i, \mathbf{y}_i)$ be an eigenpair of $(\mathbf{C}^{\frac{1}{2}})^\dagger \mathbf{D} (\mathbf{C}^{\frac{1}{2}})^\dagger$, then we have

$$\lambda_i \mathbf{y}_i = (\mathbf{C}^{\frac{1}{2}})^\dagger \mathbf{D} (\mathbf{C}^{\frac{1}{2}})^\dagger \mathbf{y}_i. \quad (89)$$

Since we only care about nonzero eigenvalues, we can assume that $\lambda_i \neq 0$, i.e., $\mathbf{y}_i \notin \mathcal{N}(\mathbf{C})$. Multiplying by $(\mathbf{C}^{\frac{1}{2}})^\dagger$ from the left we get

$$\begin{aligned} \lambda_i (\mathbf{C}^{\frac{1}{2}})^\dagger \mathbf{y}_i &= (\mathbf{C}^{\frac{1}{2}})^\dagger (\mathbf{C}^{\frac{1}{2}})^\dagger \mathbf{D} (\mathbf{C}^{\frac{1}{2}})^\dagger \mathbf{y}_i \\ &= \mathbf{C}^\dagger \mathbf{D} (\mathbf{C}^{\frac{1}{2}})^\dagger \mathbf{y}_i. \end{aligned} \quad (90)$$

Namely, $(\mathbf{C}^{\frac{1}{2}})^\dagger \mathbf{y}_i \neq \mathbf{0}$ is an eigenvector of $\mathbf{C}^\dagger \mathbf{D}$ with eigenvalue λ_i , which means that $(\mathbf{C}^{\frac{1}{2}})^\dagger \mathbf{D} (\mathbf{C}^{\frac{1}{2}})^\dagger$ and $\mathbf{C}^\dagger \mathbf{D}$ have the same eigenvalues³, and in particular have the same top eigenvalue. Therefore,

$$\lambda_{\max} \left((\mathbf{C}^{\frac{1}{2}})^\dagger \mathbf{D} (\mathbf{C}^{\frac{1}{2}})^\dagger \right) = \lambda_{\max} (\mathbf{C}^\dagger \mathbf{D}). \quad (91)$$

Overall, we showed that the condition in (11) is equivalent to

$$\eta \leq \frac{2}{\lambda_{\max} (\mathbf{C}^\dagger \mathbf{D})}. \quad (92)$$

This completes the proof for the second statement of the theorem.

Third statement. For the third statement of the theorem, note that for all $t \geq 0$

$$\begin{aligned} \text{vec} (\boldsymbol{\Sigma}_{t+1}^\perp) &= \text{vec} (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \boldsymbol{\Sigma}_{t+1}^\perp \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})}) \\ &= (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})}) \text{vec} (\boldsymbol{\Sigma}_{t+1}^\perp) \\ &= (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})}) \mathbf{Q} \text{vec} (\boldsymbol{\Sigma}_t^\perp). \end{aligned} \quad (93)$$

Namely, $\text{vec} (\boldsymbol{\Sigma}_t^\perp) = [(\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})}) \mathbf{Q}]^t \text{vec} (\boldsymbol{\Sigma}_0^\perp)$, and thus

$$\|\text{vec} (\boldsymbol{\Sigma}_t^\perp)\| = \|[(\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})}) \mathbf{Q}]^t \text{vec} (\boldsymbol{\Sigma}_0^\perp)\| \leq \|(\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})}) \mathbf{Q}\|^t \|\text{vec} (\boldsymbol{\Sigma}_0^\perp)\|. \quad (94)$$

Here

$$\begin{aligned} (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})}) \mathbf{Q} &= (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})}) \\ &\quad \left[(1-p)(\mathbf{I} - \eta \mathbf{H}) \otimes (\mathbf{I} - \eta \mathbf{H}) + p \times \frac{1}{n} \sum_{i=1}^n (\mathbf{I} - \eta \mathbf{H}_i) \otimes (\mathbf{I} - \eta \mathbf{H}_i) \right] \\ &= (1-p)(\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} - \eta \mathbf{H}) \otimes (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} - \eta \mathbf{H}) \\ &\quad + p \times \frac{1}{n} \sum_{i=1}^n (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} - \eta \mathbf{H}_i) \otimes (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} - \eta \mathbf{H}_i). \end{aligned} \quad (95)$$

³Since from Lemma 1 $\mathcal{N}(\mathbf{C}) \subseteq \mathcal{N}(\mathbf{D})$ and thus $\dim(\mathcal{N}((\mathbf{C}^{\frac{1}{2}})^\dagger \mathbf{D} (\mathbf{C}^{\frac{1}{2}})^\dagger)) = \dim(\mathcal{N}(\mathbf{C}^\dagger \mathbf{D}))$.

We see that $(\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}$ is a sum of Kronecker products, where each product is a symmetric matrix multiplied by itself. This means that Theorem 6 applies to $(\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}$, and thus we have that $\|(\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}\| = \lambda_{\max}((\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q})$. Moreover, it is easy to show that $\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} = \mathbf{P}_{\mathcal{N}^\perp(\mathbf{D})}$, and $\mathbf{P}_{\mathcal{N}^\perp(\mathbf{D})}\mathbf{C} = \mathbf{C}\mathbf{P}_{\mathcal{N}^\perp(\mathbf{D})}$. Combining this with (79), we have

$$(\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q} = \mathbf{P}_{\mathcal{N}^\perp(\mathbf{D})} - 2\eta\mathbf{C}\mathbf{P}_{\mathcal{N}^\perp(\mathbf{D})} + \eta^2\mathbf{D}. \quad (96)$$

Thus, for all $\mathbf{u} \in \mathcal{N}(\mathbf{D})$ we have

$$(\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}\mathbf{u} = \mathbf{P}_{\mathcal{N}^\perp(\mathbf{D})}\mathbf{u} - 2\eta\mathbf{C}\mathbf{P}_{\mathcal{N}^\perp(\mathbf{D})}\mathbf{u} + \eta^2\mathbf{D}\mathbf{u} = \mathbf{0}. \quad (97)$$

Since the eigenvectors of symmetric matrices are orthogonal, and $\mathcal{N}(\mathbf{D})$ is an eigenspace, we get that the top eigenvector of $(\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}$ should be in $\mathcal{N}^\perp(\mathbf{D})$. Now, for $\mathbf{u} \in \mathcal{N}^\perp(\mathbf{D}) \cap \mathbb{S}^{d^2-1}$

$$\begin{aligned} \mathbf{u}^\top (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}\mathbf{u} &= \mathbf{u}^\top \mathbf{P}_{\mathcal{N}^\perp(\mathbf{D})}\mathbf{u} - 2\eta\mathbf{u}^\top \mathbf{C}\mathbf{P}_{\mathcal{N}^\perp(\mathbf{D})}\mathbf{u} + \eta^2\mathbf{u}^\top \mathbf{D}\mathbf{u} \\ &= 1 - 2\eta\mathbf{u}^\top \mathbf{C}\mathbf{u} + \eta^2\mathbf{u}^\top \mathbf{D}\mathbf{u}, \end{aligned} \quad (98)$$

where in the second step we used the fact that $\mathbf{u} \in \mathcal{N}^\perp(\mathbf{D})$, and therefore $\mathbf{P}_{\mathcal{N}^\perp(\mathbf{D})}\mathbf{u} = \mathbf{u}$. Additionally, note that

$$\inf_{\mathbf{u} \in \mathbb{S}^{d^2-1}, \mathbf{u} \notin \mathcal{N}(\mathbf{D})} \left\{ \frac{\mathbf{u}^\top \mathbf{C}\mathbf{u}}{\mathbf{u}^\top \mathbf{D}\mathbf{u}} \right\} = \inf_{\mathbf{u} \in \mathbb{S}^{d^2-1} \cap \mathcal{N}^\perp(\mathbf{D})} \left\{ \frac{\mathbf{u}^\top \mathbf{C}\mathbf{u}}{\mathbf{u}^\top \mathbf{D}\mathbf{u}} \right\}. \quad (99)$$

Namely, having a component of \mathbf{u} in $\mathcal{N}(\mathbf{D})$ can only be non-optimal, since the denominator is invariant to vectors in $\mathcal{N}(\mathbf{D})$, while the numerator can only increase (\mathbf{C} is PSD). Now, assuming $\eta > 0$ we have from the derivation of η_{var}^* in the second statement (see (84))

$$\begin{aligned} \eta &< \eta_{\text{var}}^* \\ \Leftrightarrow \eta &< 2 \inf_{\mathbf{u} \in \mathbb{S}^{d^2-1}, \mathbf{u} \notin \mathcal{N}(\mathbf{D})} \left\{ \frac{\mathbf{u}^\top \mathbf{C}\mathbf{u}}{\mathbf{u}^\top \mathbf{D}\mathbf{u}} \right\} \\ \Leftrightarrow \eta &< 2 \inf_{\mathbf{u} \in \mathbb{S}^{d^2-1} \cap \mathcal{N}^\perp(\mathbf{D})} \left\{ \frac{\mathbf{u}^\top \mathbf{C}\mathbf{u}}{\mathbf{u}^\top \mathbf{D}\mathbf{u}} \right\} \\ \Leftrightarrow \eta &< 2 \frac{\mathbf{u}^\top \mathbf{C}\mathbf{u}}{\mathbf{u}^\top \mathbf{D}\mathbf{u}} \quad \forall \mathbf{u} \in \mathbb{S}^{d^2-1} \cap \mathcal{N}^\perp(\mathbf{D}) \\ \Leftrightarrow \eta &< 2 \frac{\mathbf{u}^\top \mathbf{C}\mathbf{u}}{\mathbf{u}^\top \mathbf{D}\mathbf{u}} \quad \forall \mathbf{u} \in \mathbb{S}^{d^2-1} \cap \mathcal{N}^\perp(\mathbf{D}) \\ \Leftrightarrow \eta \mathbf{u}^\top \mathbf{D}\mathbf{u} &< 2\mathbf{u}^\top \mathbf{C}\mathbf{u} \quad \forall \mathbf{u} \in \mathbb{S}^{d^2-1} \cap \mathcal{N}^\perp(\mathbf{D}) \\ \Leftrightarrow \eta^2 \mathbf{u}^\top \mathbf{D}\mathbf{u} &< 2\eta \mathbf{u}^\top \mathbf{C}\mathbf{u} \quad \forall \mathbf{u} \in \mathbb{S}^{d^2-1} \cap \mathcal{N}^\perp(\mathbf{D}) \quad (\eta > 0) \\ \Leftrightarrow -2\eta \mathbf{u}^\top \mathbf{C}\mathbf{u} + \eta^2 \mathbf{u}^\top \mathbf{D}\mathbf{u} &< 0 \quad \forall \mathbf{u} \in \mathbb{S}^{d^2-1} \cap \mathcal{N}^\perp(\mathbf{D}) \\ \Leftrightarrow 1 - 2\eta \mathbf{u}^\top \mathbf{C}\mathbf{u} + \eta^2 \mathbf{u}^\top \mathbf{D}\mathbf{u} &< 1 \quad \forall \mathbf{u} \in \mathbb{S}^{d^2-1} \cap \mathcal{N}^\perp(\mathbf{D}) \\ \Leftrightarrow \mathbf{u}^\top (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}\mathbf{u} &< 1 \quad \forall \mathbf{u} \in \mathbb{S}^{d^2-1} \cap \mathcal{N}^\perp(\mathbf{D}) \\ \Leftrightarrow \lambda_{\max}((\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}) &< 1 \end{aligned} \quad (100)$$

where in the penultimate step we used (98). Overall we have that $0 < \eta < \eta_{\text{var}}^*$ if and only if $\lambda_{\max}((\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}) < 1$ (we will use this fact later on). Therefore, $\lambda_{\max}((\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}) = \|(\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}\| < 1$. Hence, from (94) we get

$$\lim_{t \rightarrow \infty} \|\text{vec}(\boldsymbol{\Sigma}_t^+) \| \leq \|(\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}\|^t \|\text{vec}(\boldsymbol{\Sigma}_0^+)\| = 0, \quad (101)$$

which proves the statement.

B.8 PROOF OF THEOREM 4

First statement. Let us start by proving the first statement. In (34) we showed that if the minimum is regular then

$$\boldsymbol{\theta}_{t+1}^\parallel - \boldsymbol{\theta}^{*\parallel} = \boldsymbol{\theta}_t^\parallel - \boldsymbol{\theta}^{*\parallel} - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{g}_i^\parallel. \quad (102)$$

Let us compute the expected squared norm. We have

$$\begin{aligned}
\mathbb{E} \left[\|\boldsymbol{\theta}_{t+1}^{\parallel} - \boldsymbol{\theta}^{*\parallel}\|^2 \right] &= \mathbb{E} \left[\left\| \boldsymbol{\theta}_t^{\parallel} - \boldsymbol{\theta}^{*\parallel} - \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{g}_i^{\parallel} \right\|^2 \right] \\
&= \mathbb{E} \left[\|\boldsymbol{\theta}_t^{\parallel} - \boldsymbol{\theta}^{*\parallel}\|^2 \right] + \mathbb{E} \left[\left\| \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{g}_i^{\parallel} \right\|^2 \right] - 2\mathbb{E} \left[(\boldsymbol{\theta}_t^{\parallel} - \boldsymbol{\theta}^{*\parallel})^{\top} \left(\frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{g}_i^{\parallel} \right) \right] \\
&= \mathbb{E} \left[\|\boldsymbol{\theta}_t^{\parallel} - \boldsymbol{\theta}^{*\parallel}\|^2 \right] + \mathbb{E} \left[\left\| \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{g}_i^{\parallel} \right\|^2 \right] - 2\mathbb{E} \left[\boldsymbol{\theta}_t^{\parallel} - \boldsymbol{\theta}^{*\parallel} \right]^{\top} \mathbb{E} \left[\frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{g}_i^{\parallel} \right] \\
&= \mathbb{E} \left[\|\boldsymbol{\theta}_t^{\parallel} - \boldsymbol{\theta}^{*\parallel}\|^2 \right] + \mathbb{E} \left[\left\| \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{g}_i^{\parallel} \right\|^2 \right], \tag{103}
\end{aligned}$$

where in the second line we used the fact that $\boldsymbol{\theta}_t^{\parallel}$ is independent of \mathfrak{B}_t and in the last line we used the fact that

$$\mathbb{E} \left[\frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{g}_i^{\parallel} \right] = \frac{\eta}{n} \sum_{i=1}^n \mathbf{g}_i^{\parallel} = \mathbf{P}_{\mathcal{N}(\mathbf{H})} \frac{\eta}{n} \sum_{i=1}^n \mathbf{g}_i = \mathbf{0}. \tag{104}$$

Using (63) we have that (see definition in (45))

$$\begin{aligned}
\mathbb{E} \left[\left\| \frac{\eta}{B} \sum_{i \in \mathfrak{B}_t} \mathbf{g}_i^{\parallel} \right\|^2 \right] &= \mathbb{E} \left[\|\mathbf{P}_{\mathcal{N}(\mathbf{H})} \mathbf{v}_t\|^2 \right] \\
&= \text{Tr} \left(\mathbf{P}_{\mathcal{N}(\mathbf{H})} \mathbb{E} \left[\mathbf{v}_t \mathbf{v}_t^{\top} \right] \mathbf{P}_{\mathcal{N}(\mathbf{H})} \right) \\
&= \eta^2 \frac{n-B}{B(n-1)} \frac{1}{n} \sum_{i=1}^n \text{Tr} \left(\mathbf{P}_{\mathcal{N}(\mathbf{H})} \mathbf{g}_i \mathbf{g}_i^{\top} \mathbf{P}_{\mathcal{N}(\mathbf{H})} \right) \\
&= \eta^2 p \frac{1}{n} \sum_{i=1}^n \|\mathbf{P}_{\mathcal{N}(\mathbf{H})} \mathbf{g}_i\|^2 \\
&= \eta^2 p \frac{1}{n} \sum_{i=1}^n \|\mathbf{g}_i^{\parallel}\|^2. \tag{105}
\end{aligned}$$

Unrolling (103) we have that

$$\mathbb{E} \left[\|\boldsymbol{\theta}_t^{\parallel} - \boldsymbol{\theta}^{*\parallel}\|^2 \right] = \mathbb{E} \left[\|\boldsymbol{\theta}_0^{\parallel} - \boldsymbol{\theta}^{*\parallel}\|^2 \right] + t \times \eta^2 p \frac{1}{n} \sum_{i=1}^n \|\mathbf{g}_i^{\parallel}\|^2. \tag{106}$$

Thus, $\lim_{t \rightarrow \infty} \mathbb{E}[\|\boldsymbol{\theta}_t^{\parallel} - \boldsymbol{\theta}^{*\parallel}\|^2] = \infty$ if and only if $\sum_{i=1}^n \|\mathbf{g}_i^{\parallel}\|^2 > 0$.

Second and third statements. Next, we turn to prove the second and third statements of the theorem. In App. B.9 we show the following.

Lemma 2. *Assume that $\boldsymbol{\theta}^*$ is a twice differentiable regular minimum. Consider the linear dynamics of $\{\boldsymbol{\theta}_t\}$ from Def. 1.*

1. If $\lambda_{\max}((\mathbf{P}_{\mathcal{N}^{\perp}(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^{\perp}(\mathbf{H})})\mathbf{Q}) < 1$ then $\limsup_{t \rightarrow \infty} \mathbb{E}[\|\boldsymbol{\theta}_t^{\perp} - \boldsymbol{\theta}^{*\perp}\|^2]$ is finite.
2. If $\limsup_{t \rightarrow \infty} \mathbb{E}[\|\boldsymbol{\theta}_t^{\perp} - \boldsymbol{\theta}^{*\perp}\|^2]$ is finite then $\lambda_{\max}((\mathbf{P}_{\mathcal{N}^{\perp}(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^{\perp}(\mathbf{H})})\mathbf{Q}) \leq 1$.
3. Let \mathbf{z}_{\max} denote the top eigenvector of $(\mathbf{P}_{\mathcal{N}^{\perp}(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^{\perp}(\mathbf{H})})\mathbf{Q}$, and assume that $\mathbf{z}_{\max}^{\top} \text{vec}(\boldsymbol{\Sigma}_{\mathbf{g}}^{\perp}) \neq 0$. If $\limsup_{t \rightarrow \infty} \mathbb{E}[\|\boldsymbol{\theta}_t^{\perp} - \boldsymbol{\theta}^{*\perp}\|^2]$ is finite then $\lambda_{\max}((\mathbf{P}_{\mathcal{N}^{\perp}(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^{\perp}(\mathbf{H})})\mathbf{Q}) < 1$.

Moreover, in (100) we showed that $\lambda_{\max}((\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}) < 1$ if and only if $0 < \eta < \eta_{\text{var}}^*$, which proves the second and third statements. Note that under the mild assumption that $\mathbf{z}_{\max}^\top \text{vec}(\boldsymbol{\Sigma}_g^\perp) \neq 0$ we get that $\limsup_{t \rightarrow \infty} \mathbb{E}[\|\boldsymbol{\theta}_t^\perp - \boldsymbol{\theta}^{*\perp}\|^2]$ is finite if and only if $0 \leq \eta < \eta_{\text{var}}^*$.

B.9 PROOF OF LEMMA 2

First statement. Here we assume $\lambda_{\max}((\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}) < 1$, and show this implies that $\limsup_{t \rightarrow \infty} \mathbb{E}[\|\boldsymbol{\theta}_t^\perp - \boldsymbol{\theta}^{*\perp}\|^2]$ is finite. The (projected) transition matrix that governs the dynamics of $\boldsymbol{\Sigma}_t^\perp$ and $\boldsymbol{\mu}_t^\perp$ in (53) is given by

$$\Xi = \begin{pmatrix} \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} - \eta\mathbf{H} & \mathbf{0} \\ -(\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})(\mathbb{E}[\mathbf{v}_t^\perp \otimes \mathbf{A}_t] + \mathbb{E}[\mathbf{A}_t \otimes \mathbf{v}_t^\perp]) & (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q} \end{pmatrix}. \quad (107)$$

Since this matrix is a block lower triangular matrix, its eigenvalues are

$$\{\lambda_j(\Xi)\} = \{\lambda_j(\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} - \eta\mathbf{H})\} \cup \{\lambda_j((\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q})\}. \quad (108)$$

Note that if $\rho((\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}) < 1$ then $\rho(\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} - \eta\mathbf{H}) < 1$ (see App. B.9.1). Therefore, all the eigenvalues of Ξ are less than 1 in absolute value. Therefore, $\|\text{vec}(\boldsymbol{\Sigma}_t^\perp)\|_2 = \|\boldsymbol{\Sigma}_t^\perp\|_F$ is bounded. Since $\boldsymbol{\Sigma}_t^\perp$ is PSD we have

$$\|\boldsymbol{\Sigma}_t^\perp\|_F = \sqrt{\sum_{j=1}^d \lambda_j^2(\boldsymbol{\Sigma}_t^\perp)} \geq \frac{1}{\sqrt{d}} \sum_{j=1}^d \lambda_j(\boldsymbol{\Sigma}_t^\perp) = \frac{1}{\sqrt{d}} \text{Tr}(\boldsymbol{\Sigma}_t^\perp) = \frac{1}{\sqrt{d}} \mathbb{E}[\|\boldsymbol{\theta}_t^\perp - \boldsymbol{\theta}^{*\perp}\|^2]. \quad (109)$$

Therefore, $\mathbb{E}[\|\boldsymbol{\theta}_t^\perp - \boldsymbol{\theta}^{*\perp}\|^2]$ is bounded.

Second statement. Here we assume that $\limsup_{t \rightarrow \infty} \mathbb{E}[\|\boldsymbol{\theta}_t^\perp - \boldsymbol{\theta}^{*\perp}\|^2]$ is finite, then we show $\lambda_{\max}((\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}) \leq 1$. The matrix $(\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}$ can be written as

$$\begin{aligned} (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q} &= (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} - \eta\mathbf{H}) \otimes (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} - \eta\mathbf{H}) \\ &\quad + \eta^2 p \left(\frac{1}{n} \sum_{i=1}^n (\mathbf{H}_i - \mathbf{H}) \otimes (\mathbf{H}_i - \mathbf{H}) \right). \end{aligned} \quad (110)$$

This expression is a sum of Kronecker products, where each product is a symmetric matrix with itself. Therefore, according to Theorem 6, we get $\lambda_{\max}((\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}) = \rho((\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q})$ and $\mathbf{Z}_{\max} = \text{vec}^{-1}(\mathbf{z}_{\max})$ is a PSD matrix, where \mathbf{z}_{\max} is a normalized top eigenvector of $(\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}$. Now, set⁴ $\boldsymbol{\Sigma}_0^\perp = \mathbf{Z}_{\max}$ and $\boldsymbol{\mu}_0 = \mathbf{0}$, then in this case $\boldsymbol{\mu}_t^\perp = (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} - \eta\mathbf{H})^t \boldsymbol{\mu}_0^\perp = \mathbf{0}$ for all $t > 0$. Therefore, from (53)

$$\text{vec}(\boldsymbol{\Sigma}_{t+1}^\perp) = (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q} \text{vec}(\boldsymbol{\Sigma}_t^\perp) + \text{vec}(\boldsymbol{\Sigma}_v^\perp). \quad (111)$$

Thus,

$$\begin{aligned} \mathbf{z}_{\max}^\top \text{vec}(\boldsymbol{\Sigma}_{t+1}^\perp) &= \mathbf{z}_{\max}^\top (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q} \text{vec}(\boldsymbol{\Sigma}_t^\perp) + \mathbf{z}_{\max}^\top \text{vec}(\boldsymbol{\Sigma}_v^\perp) \\ &= \lambda_{\max}((\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}) \mathbf{z}_{\max}^\top \text{vec}(\boldsymbol{\Sigma}_t^\perp) + \text{Tr}(\mathbf{Z}_{\max} \boldsymbol{\Sigma}_v^\perp) \\ &\geq \lambda_{\max}((\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}) \mathbf{z}_{\max}^\top \text{vec}(\boldsymbol{\Sigma}_t^\perp), \end{aligned} \quad (112)$$

where we used the fact that $\mathbf{Z}_{\max}^{\frac{1}{2}} \boldsymbol{\Sigma}_v^\perp \mathbf{Z}_{\max}^{\frac{1}{2}}$ is a PSD matrix, and thus

$$\text{Tr}(\mathbf{Z}_{\max} \boldsymbol{\Sigma}_v^\perp) = \text{Tr}\left(\mathbf{Z}_{\max}^{\frac{1}{2}} \boldsymbol{\Sigma}_v^\perp \mathbf{Z}_{\max}^{\frac{1}{2}}\right) \geq 0. \quad (113)$$

⁴ $\mathbf{Z}_{\max} \in \mathcal{N}^\perp(\mathbf{D})$, i.e., $\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \mathbf{Z}_{\max} \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} = \mathbf{Z}_{\max}$, and therefore this initialization is possible (see discussion below (97)).

Therefore,

$$\begin{aligned} \mathbf{z}_{\max}^T \text{vec}(\boldsymbol{\Sigma}_t^\perp) &\geq \lambda_{\max}^t((\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q})\mathbf{z}_{\max}^T \text{vec}(\boldsymbol{\Sigma}_0) \\ &= \lambda_{\max}^t((\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}), \end{aligned} \quad (114)$$

where in the last step we used $\boldsymbol{\Sigma}_0 = \mathbf{Z}_{\max}$ and $\|\mathbf{Z}_{\max}\|_{\text{F}} = 1$. Additionally, for all $t > 0$

$$\begin{aligned} \mathbf{z}_{\max}^T \text{vec}(\boldsymbol{\Sigma}_t^\perp) &\leq \|\mathbf{z}_{\max}\|_2 \|\text{vec}(\boldsymbol{\Sigma}_t^\perp)\|_2 \\ &= \|\boldsymbol{\Sigma}_t^\perp\|_{\text{F}} \\ &= \sqrt{\sum_{j=1}^d \lambda_j^2(\boldsymbol{\Sigma}_t^\perp)} \\ &\leq \sum_{j=1}^d \lambda_j(\boldsymbol{\Sigma}_t^\perp) \\ &= \text{Tr}(\boldsymbol{\Sigma}_t^\perp) \\ &= \mathbb{E}[\|\boldsymbol{\theta}_t^\perp - \boldsymbol{\theta}^{*\perp}\|^2]. \end{aligned} \quad (115)$$

Overall, combining (114) and (115) results with

$$\lambda_{\max}^t((\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}) \leq \mathbb{E}[\|\boldsymbol{\theta}_t^\perp - \boldsymbol{\theta}^{*\perp}\|^2]. \quad (116)$$

Since $\mathbb{E}[\|\boldsymbol{\theta}_t^\perp - \boldsymbol{\theta}^{*\perp}\|^2]$ is bounded then $\lambda_{\max}((\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}) \leq 1$.

Third statement. Furthermore, if $\mathbf{z}_{\max}^T \text{vec}(\boldsymbol{\Sigma}_v^\perp) \neq 0$ we get from (113)

$$\mathbf{z}_{\max}^T \text{vec}(\boldsymbol{\Sigma}_v^\perp) > 0. \quad (117)$$

Assume by contradiction that $\lambda_{\max}((\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}) = 1$, then (112) gives

$$\mathbf{z}_{\max}^T \text{vec}(\boldsymbol{\Sigma}_{t+1}^\perp) = \mathbf{z}_{\max}^T \text{vec}(\boldsymbol{\Sigma}_t^\perp) + \mathbf{z}_{\max}^T \text{vec}(\boldsymbol{\Sigma}_v^\perp). \quad (118)$$

Unrolling this equation gives $\mathbf{z}_{\max}^T \text{vec}(\boldsymbol{\Sigma}_t^\perp) = t\mathbf{z}_{\max}^T \text{vec}(\boldsymbol{\Sigma}_v^\perp)$. Then, by (115) we get

$$\mathbb{E}[\|\boldsymbol{\theta}_t - \boldsymbol{\theta}^*\|^2] \geq \mathbf{z}_{\max}^T \text{vec}(\boldsymbol{\Sigma}_t^\perp) = t\mathbf{z}_{\max}^T \text{vec}(\boldsymbol{\Sigma}_v^\perp). \quad (119)$$

Since $\mathbf{z}_{\max}^T \text{vec}(\boldsymbol{\Sigma}_v^\perp) > 0$, then $\mathbb{E}[\|\boldsymbol{\theta}_t - \boldsymbol{\theta}^*\|^2] \rightarrow \infty$ and we have a contradiction. Therefore $\lambda_{\max}((\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}) < 1$.

B.9.1 $\rho(\mathbf{Q}) < 1$ IMPLIES $\rho(\mathbf{I} - \eta\mathbf{H}) < 1$

The matrix $(\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}$ can be written as

$$\begin{aligned} (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q} &= (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} - \eta\mathbf{H}) \otimes (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} - \eta\mathbf{H}) \\ &\quad + \eta^2 p \left(\frac{1}{n} \sum_{i=1}^n (\mathbf{H}_i - \mathbf{H}) \otimes (\mathbf{H}_i - \mathbf{H}) \right). \end{aligned} \quad (120)$$

Note that $\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} - \eta\mathbf{H}$ is a symmetric matrix, and thus its principal eigenvector $\tilde{\mathbf{v}} \in \mathbb{S}^{d-1}$ satisfies $\rho(\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} - \eta\mathbf{H}) = |\tilde{\mathbf{v}}^T(\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} - \eta\mathbf{H})\tilde{\mathbf{v}}|$. Additionally, $\|\tilde{\mathbf{v}} \otimes \tilde{\mathbf{v}}\| = \|\tilde{\mathbf{v}}\tilde{\mathbf{v}}^T\|_{\text{F}} = 1$, i.e., $\tilde{\mathbf{v}} \otimes \tilde{\mathbf{v}} \in \mathbb{S}^{d^2-1}$. Now, assume that the spectral radius $\rho((\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}) < 1$ then

$$\begin{aligned} 1 &> \rho((\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q}) \\ &\geq |[\tilde{\mathbf{v}} \otimes \tilde{\mathbf{v}}]^T (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})})\mathbf{Q} [\tilde{\mathbf{v}} \otimes \tilde{\mathbf{v}}]| \\ &= (\tilde{\mathbf{v}}^T (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} - \eta\mathbf{H}) \tilde{\mathbf{v}})^2 + \eta^2 p \frac{1}{n} \sum_{i=1}^n (\tilde{\mathbf{v}}^T (\mathbf{H}_i - \mathbf{H}) \tilde{\mathbf{v}})^2 \\ &= \rho^2(\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} - \eta\mathbf{H}) + \eta^2 p \frac{1}{n} \sum_{i=1}^n (\tilde{\mathbf{v}}^T (\mathbf{H}_i - \mathbf{H}) \tilde{\mathbf{v}})^2 \\ &\geq \rho^2(\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} - \eta\mathbf{H}). \end{aligned} \quad (121)$$

C PROOF OF THEOREM 6

First statement. Let $\mathbf{z} \in \mathbb{R}^{d^2}$ be an eigenvector of \mathbf{Q} , then we can look at its matrix form $\mathbf{Z} = \text{vec}^{-1}(\mathbf{z})$, where $\mathbf{Z} \in \mathbb{R}^{d \times d}$. First, we show that \mathbf{Q} always has a set of eigenvectors that correspond to either symmetric or skew-symmetric matrices $\{\mathbf{Z}_j\}$ (see proof in App. C.1).

Proposition 4. Let $\{\mathbf{A}_i\}$ be symmetric matrices over $\mathbb{R}^{d \times d}$, and define

$$\mathbf{Q} = \sum_{i=1}^M \mathbf{A}_i \otimes \mathbf{A}_i. \quad (122)$$

Then there exists a set of eigenvectors $\{\mathbf{z}_j\}$ such that each $\mathbf{Z}_j = \text{vec}^{-1}(\mathbf{z}_j)$ is either a symmetric or a skew-symmetric matrix.

Second and third statement. Our next step is to bring the matrices $\{\mathbf{Z}_j\}$ to a normal (canonical) form. Here we assume without loss of generality that the eigenvectors are normalized, *i.e.*, $\|\mathbf{Z}_j\|_F = 1$. For symmetric matrix \mathbf{Z} , we have the spectral decomposition theorem, and thus $\mathbf{Z} = \mathbf{V}\mathbf{S}\mathbf{V}^T$, where \mathbf{V} is an orthogonal matrix and \mathbf{S} is diagonal. We can also bring a skew-symmetric matrix to a similar form, but with \mathbf{S} a block diagonal matrix, with $\lfloor d/2 \rfloor$ blocks of size 2×2 . Specifically, these blocks are in the form of (Zumino, 1962)

$$\begin{bmatrix} 0 & s_i \\ -s_i & 0 \end{bmatrix}. \quad (123)$$

If the dimension d is odd, then \mathbf{S} has a row and column at the end filled with zeros. For numerical purposes, this normal (canonical) form can be computed by using the real Schur decomposition. Now, for symmetric matrices, we define the vector $\mathbf{s}_{\text{sym}} \in \mathbb{R}^d$ to be the diagonal of \mathbf{S} , and for skew-symmetric matrices we define $\mathbf{s}_{\text{skew}} \in \mathbb{R}^{\lfloor d/2 \rfloor}$ to be $[s_1, s_2, \dots, s_{\lfloor d/2 \rfloor}]^T$. In App. C.2 we show that for symmetric matrix \mathbf{Z} , the corresponding vector form \mathbf{z} satisfies

$$\mathbf{z}^T \mathbf{Q} \mathbf{z} = \mathbf{s}_{\text{sym}}^T \sum_{i=1}^M \mathbf{M}_i^{\odot 2} \mathbf{s}_{\text{sym}}, \quad (124)$$

where $\mathbf{M}_i = \mathbf{V}^T \mathbf{A}_i \mathbf{V}$, the superscript \odot^k denotes the Hadamard power and $\|\mathbf{s}_{\text{skew}}\| = 1$. For skew-symmetric matrices, we define for $1 \leq \ell, p \leq \lfloor d/2 \rfloor$

$$\mathbf{T}_{i[\ell, p]} = \mathbf{M}_{i[2\ell-1, 2p-1]} \mathbf{M}_{i[2\ell, 2p]} - \mathbf{M}_{i[2\ell-1, 2p]} \mathbf{M}_{i[2\ell-1, 2p]}. \quad (125)$$

Namely, $\mathbf{T}_i \in \mathbb{R}^{\lfloor d/2 \rfloor \times \lfloor d/2 \rfloor}$ is the determinant of each 2×2 block of \mathbf{M}_i without overlap. We show in App. C.2 that for skew-symmetric matrix \mathbf{Z} , the corresponding vector form \mathbf{z} satisfies

$$\mathbf{z}^T \mathbf{Q} \mathbf{z} = 2 \mathbf{s}_{\text{skew}}^T \sum_{i=1}^M \mathbf{T}_i \mathbf{s}_{\text{skew}}, \quad (126)$$

where $\|\mathbf{s}_{\text{skew}}\| = 1/\sqrt{2}$. Let us define the projection matrix $\mathbf{P} \in \mathbb{R}^{\lfloor d/2 \rfloor \times d}$ as

$$\mathbf{P} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 & 0 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & \cdots & 0 & 0 \\ \vdots & & & \ddots & & & & \vdots \\ 0 & 0 & 0 & 0 & 0 & \cdots & 1 & 1 \end{bmatrix}. \quad (127)$$

This matrix is semi-orthogonal, *i.e.*, it satisfies $\mathbf{P}\mathbf{P}^T = \mathbf{I}$. Note that

$$\left[\mathbf{P} \mathbf{M}_i^{\odot 2} \mathbf{P}^T \right]_{[\ell, p]} = \frac{1}{2} \left(\mathbf{M}_{i[2\ell-1, 2p-1]}^2 + \mathbf{M}_{i[2\ell-1, 2p]}^2 + \mathbf{M}_{i[2\ell-1, 2p]}^2 + \mathbf{M}_{i[2\ell, 2p]}^2 \right). \quad (128)$$

Therefore, for all ℓ, p and i we have

$$|\mathbf{T}_{i[\ell, p]}| \leq \left[\mathbf{P} \mathbf{M}_i^{\odot 2} \mathbf{P}^T \right]_{[\ell, p]}. \quad (129)$$

Now, given orthogonal matrix \mathbf{V} , for all $\mathbf{s}_{\text{skew}} \in \mathbb{R}^{\lfloor d/2 \rfloor}$ such that $\|\mathbf{s}_{\text{skew}}\| = 1/\sqrt{2}$, we set $\mathbf{u} = \sqrt{2}\mathbf{s}_{\text{skew}}$ and then

$$\begin{aligned}
\left| 2\mathbf{s}_{\text{skew}}^{\text{T}} \sum_{i=1}^M \mathbf{T}_i \mathbf{s}_{\text{skew}} \right| &= \left| \mathbf{u}^{\text{T}} \sum_{i=1}^M \mathbf{T}_i \mathbf{u} \right| \\
&\leq \sum_{\ell=1}^d \sum_{p=1}^d \sum_{i=1}^M |\mathbf{u}_{\ell}| |\mathbf{T}_i[\ell,p]| |\mathbf{u}_p| \\
&\leq \sum_{\ell=1}^d \sum_{p=1}^d \sum_{i=1}^M |\mathbf{u}_{\ell}| \left[\mathbf{P} \mathbf{M}_i^{\odot 2} \mathbf{P}^{\text{T}} \right]_{[\ell,p]} |\mathbf{u}_p| \\
&\leq \lambda_{\max} \left(\sum_{i=1}^M \mathbf{P} \mathbf{M}_i^{\odot 2} \mathbf{P}^{\text{T}} \right) \\
&= \lambda_{\max} \left(\mathbf{P} \left(\sum_{i=1}^M \mathbf{M}_i^{\odot 2} \right) \mathbf{P}^{\text{T}} \right) \\
&\leq \lambda_{\max} \left(\sum_{i=1}^M \mathbf{M}_i^{\odot 2} \right), \tag{130}
\end{aligned}$$

where in the last step we used the Cauchy interlacing theorem (a.k.a. Poincaré separation theorem). Thus, for every $\mathbf{s}_{\text{skew}} \in \mathbb{R}^{\lfloor d/2 \rfloor}$ such that $\|\mathbf{s}_{\text{skew}}\| = 1/\sqrt{2}$, which corresponds to a skew-symmetric matrix \mathbf{z}_{skew} , we can always find $\mathbf{z}_{\text{sym}} \in \mathbb{S}^{d-1}$ (the top eigenvector of $\sum_{i=1}^M \mathbf{M}_i^{\odot 2}$) which corresponds to a symmetric matrix \mathbf{z}_{sym} with the same basis \mathbf{V} such that

$$|\mathbf{z}_{\text{skew}}^{\text{T}} \mathbf{Q} \mathbf{z}_{\text{skew}}| \leq \mathbf{z}_{\text{sym}}^{\text{T}} \mathbf{Q} \mathbf{z}_{\text{sym}}. \tag{131}$$

Therefore, the principal eigenvector of \mathbf{Q} corresponds to a symmetric matrix rather than a skew-symmetric one. Denoted the principal eigenvector of \mathbf{Q} by $\tilde{\mathbf{z}}$ and let $\tilde{\mathbf{Z}} = \tilde{\mathbf{V}} \tilde{\mathbf{S}} \tilde{\mathbf{V}}^{\text{T}}$ be its spectral decomposition. Set

$$\mathbf{\Psi} = \sum_{i=1}^M \tilde{\mathbf{M}}_i^{\odot 2}, \quad \text{s.t.} \quad \mathbf{M}_i = \tilde{\mathbf{V}}^{\text{T}} \mathbf{A}_i \tilde{\mathbf{V}}. \tag{132}$$

Since \mathbf{Q} is symmetric, then by the spectral theorem we have that all its eigenvectors and eigenvalues are real, and they are given by the quadratic form using the corresponding eigenvectors. Thus,

$$\begin{aligned}
\rho(\mathbf{Q}) &= |\tilde{\mathbf{z}}^{\text{T}} \mathbf{Q} \tilde{\mathbf{z}}| \\
&= \left| \tilde{\mathbf{s}}_{\text{sym}}^{\text{T}} \sum_{i=1}^M \tilde{\mathbf{M}}_i^{\odot 2} \tilde{\mathbf{s}}_{\text{sym}} \right| \\
&= \left| \tilde{\mathbf{s}}_{\text{sym}}^{\text{T}} \mathbf{\Psi} \tilde{\mathbf{s}}_{\text{sym}} \right| \\
&= \left| \sum_{i=1}^d \sum_{j=1}^d \tilde{\mathbf{s}}_{\text{sym},i} \mathbf{\Psi}_{[i,j]} \tilde{\mathbf{s}}_{\text{sym},j} \right| \\
&\leq \sum_{i=1}^d \sum_{j=1}^d |\tilde{\mathbf{s}}_{\text{sym},i}| |\mathbf{\Psi}_{[i,j]}| |\tilde{\mathbf{s}}_{\text{sym},j}| \\
&= |\tilde{\mathbf{s}}_{\text{sym}}|^{\text{T}} \mathbf{\Psi} |\tilde{\mathbf{s}}_{\text{sym}}| \\
&= \hat{\mathbf{z}}^{\text{T}} \mathbf{Q} \hat{\mathbf{z}}, \tag{133}
\end{aligned}$$

where $\hat{\mathbf{z}} = \text{vec}(\tilde{\mathbf{V}} |\tilde{\mathbf{S}}| \tilde{\mathbf{V}}^{\text{T}})$. Namely, the vector $\hat{\mathbf{z}}$ that corresponds to the matrix built by the element-wise absolute value of the spectrum of $\tilde{\mathbf{Z}}$ yields a greater or equal result than the spectral radius of \mathbf{Q} , while still having a unit Euclidean norm. Thus, if $\tilde{\mathbf{s}}_{\text{sym}} \neq |\tilde{\mathbf{s}}_{\text{sym}}|$ then both $\tilde{\mathbf{z}}$ and $\hat{\mathbf{z}}$ are principal eigenvectors (or else we get a contradiction). Note that $\text{vec}^{-1}(\hat{\mathbf{z}})$ is in fact a PSD matrix. Therefore,

there is always a principal eigenvector for \mathbf{Q} which corresponds to a PSD matrix. Additionally, since $\max_j |\lambda_j(\mathbf{Q})| = \rho(\mathbf{Q}) = \hat{\mathbf{z}}^\top \mathbf{Q} \hat{\mathbf{z}}$, then $\hat{\mathbf{z}}$ is also a top eigenvector which corresponds to $\lambda_{\max}(\mathbf{Q})$, *i.e.*, $\rho(\mathbf{Q}) = \lambda_{\max}(\mathbf{Q})$.

C.1 PROOF OF PROPOSITION 4

Let (λ, \mathbf{z}) be an eigenpair of \mathbf{Q} , *i.e.*, $\mathbf{Q}\mathbf{z} = \lambda\mathbf{z}$. Set $\mathbf{Z} = \text{vec}^{-1}(\mathbf{z})$, then \mathbf{Z} satisfies

$$\begin{aligned}
\lambda\mathbf{Z} &= \text{vec}^{-1}(\lambda\mathbf{z}) \\
&= \text{vec}^{-1}(\mathbf{Q}\mathbf{z}) \\
&= \text{vec}^{-1}\left(\sum_{i=1}^M \mathbf{A}_i \otimes \mathbf{A}_i \mathbf{z}\right) \\
&= \sum_{i=1}^M \text{vec}^{-1}(\mathbf{A}_i \otimes \mathbf{A}_i \mathbf{z}) \\
&= \sum_{i=1}^M \mathbf{A}_i \mathbf{Z} \mathbf{A}_i^\top \\
&= \sum_{i=1}^M \mathbf{A}_i \mathbf{Z} \mathbf{A}_i.
\end{aligned} \tag{134}$$

Now, by taking a transpose on both ends of this equation we have

$$\begin{aligned}
\lambda\mathbf{Z}^\top &= \sum_{i=1}^M \mathbf{A}_i^\top \mathbf{Z}^\top \mathbf{A}_i \\
&= \sum_{i=1}^M \mathbf{A}_i \mathbf{Z}^\top \mathbf{A}_i,
\end{aligned} \tag{135}$$

where in the last step we used the fact that $\{\mathbf{A}_i\}$ are symmetric. Thus, we have that $\text{vec}(\mathbf{Z}^\top)$ is also an eigenvector of \mathbf{Q} . If λ has multiplicity one, then it must be that $\mathbf{Z}^\top = \pm\mathbf{Z}$, *i.e.*, symmetric or skew-symmetric matrix. If the multiplicity is greater than one and $\mathbf{Z}^\top \neq \pm\mathbf{Z}$, then any linear combination between \mathbf{Z} and \mathbf{Z}^\top will also be an eigenvector corresponding to λ . Particularly,

$$\hat{\mathbf{Z}}_1 = \frac{1}{2}(\mathbf{Z} + \mathbf{Z}^\top) \quad \hat{\mathbf{Z}}_2 = \frac{1}{2}(\mathbf{Z} - \mathbf{Z}^\top). \tag{136}$$

Now, $\hat{\mathbf{Z}}_1$ and $\hat{\mathbf{Z}}_2$ are symmetric and skew-symmetric eigenvectors, which correspond to λ . This procedure can be repeated while projecting the next eigenvectors of λ onto the orthogonal complement of the already found vectors, until we find all the eigenvectors of λ . In this way, we can always find a set of eigenvectors that is comprised solely of symmetric and skew-symmetric vectors.

C.2 QUADRATIC FORM CALCULATION FOR SYMMETRIC AND SKEW-SYMMETRIC MATRICES

Let $\mathbf{z} = \text{vec}(\mathbf{Z})$, and assume that $\mathbf{Z} = \mathbf{V}\mathbf{S}\mathbf{V}^\top$ where \mathbf{V} is orthogonal matrix. Then

$$\begin{aligned}
\mathbf{z}^\top \mathbf{Q} \mathbf{z} &= \left[\text{vec}(\mathbf{V}\mathbf{S}\mathbf{V}^\top) \right]^\top \mathbf{Q} \text{vec}(\mathbf{V}\mathbf{S}\mathbf{V}^\top) \\
&= [(\mathbf{V} \otimes \mathbf{V}) \text{vec}(\mathbf{S})]^\top \mathbf{Q} (\mathbf{V} \otimes \mathbf{V}) \text{vec}(\mathbf{S}) \\
&= [\text{vec}(\mathbf{S})]^\top (\mathbf{V}^\top \otimes \mathbf{V}^\top) \sum_{i=1}^M \mathbf{A}_i \otimes \mathbf{A}_i (\mathbf{V} \otimes \mathbf{V}) \text{vec}(\mathbf{S}) \\
&= [\text{vec}(\mathbf{S})]^\top \sum_{i=1}^M (\mathbf{V}^\top \mathbf{A}_i \mathbf{V}) \otimes (\mathbf{V}^\top \mathbf{A}_i \mathbf{V}) \text{vec}(\mathbf{S}) \\
&= \sum_{i=1}^M [\text{vec}(\mathbf{S})]^\top \mathbf{M}_i \otimes \mathbf{M}_i \text{vec}(\mathbf{S}).
\end{aligned} \tag{137}$$

Now, for each $i \in [n]$ we have

$$[\text{vec}(\mathbf{S})]^\top \mathbf{M}_i \otimes \mathbf{M}_i \text{vec}(\mathbf{S}) = \sum_{m=1}^{d^2} \sum_{k=1}^{d^2} [\mathbf{M}_i \otimes \mathbf{M}_i]_{[m,k]} \text{vec}(\mathbf{S})_m \text{vec}(\mathbf{S})_k. \quad (138)$$

Set $m = d(m_2 - 1) + m_1$ and $k = d(k_2 - 1) + k_1$ where $m_1, m_2, k_1, k_2 \in [d]$, then

$$[\text{vec}(\mathbf{S})]^\top \mathbf{M}_i \otimes \mathbf{M}_i \text{vec}(\mathbf{S}) = \sum_{m_2=1}^d \sum_{m_1=1}^d \sum_{k_2=1}^d \sum_{k_1=1}^d \mathbf{M}_i_{[m_2,k_2]} \mathbf{M}_i_{[m_1,k_1]} \mathbf{S}_{[m_2,m_1]} \mathbf{S}_{[k_2,k_1]}. \quad (139)$$

C.2.1 SYMMETRIC EIGENVECTORS

Assume that \mathbf{Z} is symmetric, then \mathbf{S} is a diagonal matrix. Therefore, we only need to consider the terms in the series above for which $m_1 = m_2 = p$ and $k_1 = k_2 = \ell$.

$$\sum_{p=1}^d \sum_{\ell=1}^d \mathbf{M}_i_{[p,\ell]} \mathbf{M}_i_{[p,\ell]} \mathbf{S}_{[p,p]} \mathbf{S}_{[\ell,\ell]} = \sum_{p=1}^d \sum_{\ell=1}^d \mathbf{M}_i^2_{[p,\ell]} \mathbf{s}_{\text{sym},p} \mathbf{s}_{\text{sym},\ell} = \mathbf{s}_{\text{sym}}^\top \mathbf{M}_i^{\odot 2} \mathbf{s}_{\text{sym}}. \quad (140)$$

Overall,

$$\mathbf{z}^\top \mathbf{Q} \mathbf{z} = \mathbf{s}_{\text{sym}}^\top \sum_{i=1}^M \mathbf{M}_i^{\odot 2} \mathbf{s}_{\text{sym}}. \quad (141)$$

C.2.2 SKEW-SYMMETRIC EIGENVECTORS

Assume that \mathbf{Z} is skew-symmetric, then \mathbf{S} is a block diagonal matrix, where each block is 2×2 in the form of

$$\begin{bmatrix} 0 & s_j \\ -s_j & 0 \end{bmatrix}. \quad (142)$$

If the dimension d is odd, then \mathbf{S} has a row and column at the end filled with zeros.

In this scenario, we have four different cases to consider.

Case I: $m_1 = 2p - 1, m_2 = 2p, k_1 = 2\ell - 1, k_2 = 2\ell$.

$$\sum_{p=1}^{\lfloor d/2 \rfloor} \sum_{\ell=1}^{\lfloor d/2 \rfloor} \mathbf{M}_i_{[2p,2\ell]} \mathbf{M}_i_{[2p-1,2\ell-1]} \mathbf{S}_{[2p,2p-1]} \mathbf{S}_{[2\ell,2\ell-1]} = \sum_{p=1}^{\lfloor d/2 \rfloor} \sum_{\ell=1}^{\lfloor d/2 \rfloor} \mathbf{M}_i_{[2p,2\ell]} \mathbf{M}_i_{[2p-1,2\ell-1]} \mathbf{s}_{\text{skew},p} \mathbf{s}_{\text{skew},\ell}. \quad (143)$$

Case II: $m_1 = 2p, m_2 = 2p - 1, k_1 = 2\ell - 1, k_2 = 2\ell$.

$$\sum_{p=1}^{\lfloor d/2 \rfloor} \sum_{\ell=1}^{\lfloor d/2 \rfloor} \mathbf{M}_i_{[2p-1,2\ell]} \mathbf{M}_i_{[2p,2\ell-1]} \mathbf{S}_{[2p-1,2p]} \mathbf{S}_{[2\ell,2\ell-1]} = - \sum_{p=1}^{\lfloor d/2 \rfloor} \sum_{\ell=1}^{\lfloor d/2 \rfloor} \mathbf{M}_i_{[2p-1,2\ell]} \mathbf{M}_i_{[2p,2\ell-1]} \mathbf{s}_{\text{skew},p} \mathbf{s}_{\text{skew},\ell}. \quad (144)$$

Case III: $m_1 = 2p - 1, m_2 = 2p, k_1 = 2\ell, k_2 = 2\ell - 1$.

$$\sum_{p=1}^{\lfloor d/2 \rfloor} \sum_{\ell=1}^{\lfloor d/2 \rfloor} \mathbf{M}_i_{[2p,2\ell-1]} \mathbf{M}_i_{[2p-1,2\ell]} \mathbf{S}_{[2p,2p-1]} \mathbf{S}_{[2\ell-1,2\ell]} = - \sum_{p=1}^{\lfloor d/2 \rfloor} \sum_{\ell=1}^{\lfloor d/2 \rfloor} \mathbf{M}_i_{[2p,2\ell-1]} \mathbf{M}_i_{[2p-1,2\ell]} \mathbf{s}_{\text{skew},p} \mathbf{s}_{\text{skew},\ell}. \quad (145)$$

Case IV: $m_1 = 2p, m_2 = 2p - 1, k_1 = 2\ell, k_2 = 2\ell - 1$.

$$\sum_{p=1}^{\lfloor d/2 \rfloor} \sum_{\ell=1}^{\lfloor d/2 \rfloor} \mathbf{M}_i_{[2p-1,2\ell-1]} \mathbf{M}_i_{[2p,2\ell]} \mathbf{S}_{[2p-1,2p]} \mathbf{S}_{[2\ell-1,2\ell]} = \sum_{p=1}^{\lfloor d/2 \rfloor} \sum_{\ell=1}^{\lfloor d/2 \rfloor} \mathbf{M}_i_{[2p-1,2\ell-1]} \mathbf{M}_i_{[2p,2\ell]} \mathbf{s}_{\text{skew},p} \mathbf{s}_{\text{skew},\ell}. \quad (146)$$

Summing over all these cases we get

$$\begin{aligned}
\mathbf{z}^\top \mathbf{Q} \mathbf{z} &= \sum_{i=1}^M \sum_{p=1}^{\lfloor d/2 \rfloor} \sum_{\ell=1}^{\lfloor d/2 \rfloor} 2 \left(\mathbf{M}_i^{[2\ell-1, 2p-1]} \mathbf{M}_i^{[2\ell, 2p]} - \mathbf{M}_i^{[2\ell-1, 2p]} \mathbf{M}_i^{[2\ell-1, 2p]} \right) \mathbf{s}_{\text{skew}, p} \mathbf{s}_{\text{skew}, \ell} \\
&= 2 \sum_{i=1}^M \sum_{p=1}^{\lfloor d/2 \rfloor} \sum_{\ell=1}^{\lfloor d/2 \rfloor} \mathbf{T}_i^{[p, \ell]} \mathbf{s}_{\text{skew}, p} \mathbf{s}_{\text{skew}, \ell} \\
&= 2 \sum_{i=1}^M \mathbf{s}_{\text{skew}}^\top \mathbf{T}_i \mathbf{s}_{\text{skew}}.
\end{aligned} \tag{147}$$

D PROOF OF PROPOSITION 1

Here we focus on interpolating minima for simplicity. A similar proof can be derived for regular minima. To begin with, note that (see (91))

$$\lambda_{\max}(\mathbf{C}^\dagger \mathbf{D}) = \lambda_{\max} \left(\left(\mathbf{C}^{\frac{1}{2}} \right)^\dagger \mathbf{D} \left(\mathbf{C}^{\frac{1}{2}} \right)^\dagger \right). \tag{148}$$

Additionally,

$$\begin{aligned}
\mathbf{D} &= (1-p) \mathbf{H} \otimes \mathbf{H} + p \frac{1}{n} \sum_{i=1}^n \mathbf{H}_i \otimes \mathbf{H}_i \\
&= \mathbf{H} \otimes \mathbf{H} + p \frac{1}{n} \sum_{i=1}^n (\mathbf{H}_i - \mathbf{H}) \otimes (\mathbf{H}_i - \mathbf{H}) \\
&= \mathbf{H} \otimes \mathbf{H} + p \mathbf{E},
\end{aligned} \tag{149}$$

where $\mathbf{E} \triangleq \frac{1}{n} \sum_{i=1}^n (\mathbf{H}_i - \mathbf{H}) \otimes (\mathbf{H}_i - \mathbf{H})$. Let $\mathbf{y} \in \mathbb{S}^{d^2-1}$ be the top eigenvector of $(\mathbf{C}^{\frac{1}{2}})^\dagger \mathbf{D} (\mathbf{C}^{\frac{1}{2}})^\dagger$, then since $(\mathbf{C}^{\frac{1}{2}})^\dagger \mathbf{D} (\mathbf{C}^{\frac{1}{2}})^\dagger$ is symmetric we have

$$\begin{aligned}
\frac{\partial}{\partial p} \lambda_{\max} \left(\left(\mathbf{C}^{\frac{1}{2}} \right)^\dagger \mathbf{D} \left(\mathbf{C}^{\frac{1}{2}} \right)^\dagger \right) &= \mathbf{y}^\top \left[\left(\mathbf{C}^{\frac{1}{2}} \right)^\dagger \left(\frac{\partial}{\partial p} \mathbf{D} \right) \left(\mathbf{C}^{\frac{1}{2}} \right)^\dagger \right] \mathbf{y} \\
&= \mathbf{y}^\top \left[\left(\mathbf{C}^{\frac{1}{2}} \right)^\dagger \mathbf{E} \left(\mathbf{C}^{\frac{1}{2}} \right)^\dagger \right] \mathbf{y}.
\end{aligned} \tag{150}$$

In App. D.1 we show that \mathbf{y} has the form of $\mathbf{y} = \mathbf{C}^{\frac{1}{2}} \mathbf{u}$ such that $\text{vec}^{-1}(\mathbf{u}) \in \mathcal{S}_+(\mathbb{R}^{d \times d})$. Plugging this into the equation above we get

$$\begin{aligned}
\mathbf{y}^\top \left[\left(\mathbf{C}^{\frac{1}{2}} \right)^\dagger \mathbf{E} \left(\mathbf{C}^{\frac{1}{2}} \right)^\dagger \right] \mathbf{y} &= \mathbf{u}^\top \mathbf{C}^{\frac{1}{2}} \left(\mathbf{C}^{\frac{1}{2}} \right)^\dagger \mathbf{E} \left(\mathbf{C}^{\frac{1}{2}} \right)^\dagger \mathbf{C}^{\frac{1}{2}} \mathbf{u} \\
&= \mathbf{u}^\top \mathbf{P}_{\mathcal{N}^\perp(\mathbf{C})} \mathbf{E} \mathbf{P}_{\mathcal{N}^\perp(\mathbf{C})} \mathbf{u} \\
&= \mathbf{u}^\top \mathbf{E} \mathbf{u},
\end{aligned} \tag{151}$$

where in the first step we used the fact that \mathbf{C} is symmetric. Additionally, in the second step, we used the fact that $\mathbf{C}^{\frac{1}{2}} (\mathbf{C}^{\frac{1}{2}})^\dagger$ and $(\mathbf{C}^{\frac{1}{2}})^\dagger \mathbf{C}^{\frac{1}{2}}$ are projection matrices onto the column space of \mathbf{C} . Since the null space of \mathbf{E} contains the null space of \mathbf{C} , we have that these projections can be removed (see App. B.6). Note that $\text{vec}^{-1}(\mathbf{u})$ is PSD, and let $\mathbf{V} \mathbf{S} \mathbf{V}^\top$ be its spectral decomposition, then in App. C.2 we show that in this case

$$\mathbf{u}^\top \mathbf{E} \mathbf{u} = \mathbf{s}^\top \sum_{i=1}^n \mathbf{M}_i^{\odot 2} \mathbf{s}, \tag{152}$$

where $\mathbf{M}_i = \mathbf{V}^\top (\mathbf{H}_i - \mathbf{H}) \mathbf{V}$ and \mathbf{s} is a vector containing the eigenvalues of $\text{vec}^{-1}(\mathbf{u})$. Since $\text{vec}^{-1}(\mathbf{u})$ is PSD, we have the right-hand side of (152) is a sum over non-negative terms. Namely,

$$\frac{\partial}{\partial p} \lambda_{\max} \left(\left(\mathbf{C}^{\frac{1}{2}} \right)^\dagger \mathbf{D} \left(\mathbf{C}^{\frac{1}{2}} \right)^\dagger \right) = \mathbf{u}^\top \mathbf{E} \mathbf{u} = \mathbf{s}^\top \sum_{i=1}^n \mathbf{M}_i^{\odot 2} \mathbf{s} \geq 0. \tag{153}$$

Therefore, $\lambda_{\max}(\mathbf{C}^\dagger \mathbf{D})$ is monotonically non-decreasing in p , which means that $\eta_{\text{var}}^* = 2/\lambda_{\max}(\mathbf{C}^\dagger \mathbf{D})$ is monotonically non-decreasing with B .

D.1 TOP EIGENVECTOR OF $(\mathbf{C}^{\frac{1}{2}})^\dagger \mathbf{D} (\mathbf{C}^{\frac{1}{2}})^\dagger$

Using the stability condition of Ma & Ying (2021), we have that $\{\mathbb{E}[\|\boldsymbol{\theta}_t - \boldsymbol{\theta}\|^2]\}$ is bounded if and only if (see proof in (Ma & Ying, 2021))

$$\max_{\boldsymbol{\Sigma} \in \mathcal{S}_+(\mathbb{R}^{d \times d})} \frac{\|\mathbf{Q}(\eta, B) \text{vec}(\boldsymbol{\Sigma})\|}{\|\boldsymbol{\Sigma}\|_{\text{F}}} \leq 1. \quad (154)$$

Let us repeat the same steps from Sec. 3.3 but *without* relaxing the constraint of PSD matrices. Applying Thm. 6 we get that the maximizer for (154) is the top eigenvalue of \mathbf{Q} which corresponds to a PSD matrix. Then,

$$\mathbf{u}^\top \mathbf{Q} \mathbf{u} = 1 - 2\eta \mathbf{u}^\top \mathbf{C} \mathbf{u} + \eta^2 \mathbf{u}^\top \mathbf{D} \mathbf{u} \leq 1 \quad (155)$$

holds for any $\mathbf{u} \in \mathbb{S}^{d^2-1}$ such that $\text{vec}^{-1}(\mathbf{u}) \in \mathcal{S}_+(\mathbb{R}^{d \times d})$ and $\mathbf{u} \notin \mathcal{N}(\mathbf{D})$, if and only if

$$\eta \leq \frac{2}{\lambda^*} = \eta_{\text{var}}^*, \quad (156)$$

where

$$\lambda^* = \sup_{\mathbf{u} \in \mathbb{S}^{d^2-1}} \frac{\mathbf{u}^\top \mathbf{D} \mathbf{u}}{\mathbf{u}^\top \mathbf{C} \mathbf{u}} \quad \text{s.t.} \quad \text{vec}^{-1}(\mathbf{u}) \in \mathcal{S}_+(\mathbb{R}^{d \times d}) \text{ and } \mathbf{u} \notin \mathcal{N}(\mathbf{D}). \quad (157)$$

(Note that $\mathbf{u} \in \mathcal{N}(\mathbf{D})$ do not contribute any conditions on the learning rate, and therefore can be ignored - see App. B.7). With some algebra (see (87) and (88))

$$\lambda^* = \max_{\mathbf{y} \in \mathbb{S}^{d^2-1}} \mathbf{y}^\top \left(\mathbf{C}^{\frac{1}{2}}\right)^\dagger \mathbf{D} \left(\mathbf{C}^{\frac{1}{2}}\right)^\dagger \mathbf{y} \quad \text{s.t.} \quad \mathbf{y} = \left(\mathbf{C}^{\frac{1}{2}}\right) \mathbf{u} \text{ and } \text{vec}^{-1}(\mathbf{u}) \in \mathcal{S}_+(\mathbb{R}^{d \times d}). \quad (158)$$

Since the alternative form of η_{var}^* in (156) has to be equal to the definition in (13) (or else we will get a contradiction), we get

$$\lambda_{\max} \left(\left(\mathbf{C}^{\frac{1}{2}}\right)^\dagger \mathbf{D} \left(\mathbf{C}^{\frac{1}{2}}\right)^\dagger \right) = \lambda^*. \quad (159)$$

Namely, the top eigenvector \mathbf{y} of $(\mathbf{C}^{\frac{1}{2}})^\dagger \mathbf{D} (\mathbf{C}^{\frac{1}{2}})^\dagger$ has the form of $\mathbf{y} = \mathbf{C}^{\frac{1}{2}} \mathbf{u}$ such that $\text{vec}^{-1}(\mathbf{u}) \in \mathcal{S}_+(\mathbb{R}^{d \times d})$.

E PROOF OF PROPOSITION 2

Here we focus on interpolating minima for simplicity. A similar proof can be derived for regular minima. Let $\{\beta_t\}$ and $\{\kappa_t\}$ be i.i.d. random variables such that $\beta_t \sim \text{Bernoulli}(p)$ and $\kappa_t \sim \mathcal{U}(\{1, \dots, n\})$, then

$$\mathfrak{B}_t = \begin{cases} \kappa_t & \text{if } \beta_t = 1, \\ \{1, \dots, n\} & \text{otherwise.} \end{cases} \quad (160)$$

Let us consider the following stochastic loss function

$$\hat{\mathcal{L}}_t(\boldsymbol{\theta}) = \frac{1}{|\mathfrak{B}_t|} \sum_{j \in \mathfrak{B}_t} \ell_j(\boldsymbol{\theta}), \quad (161)$$

and define the following notation

$$\mathbf{A}_t = \mathbf{I} - \frac{\eta}{|\mathfrak{B}_t|} \sum_{i \in \mathfrak{B}_t} \mathbf{H}_i. \quad (162)$$

First, for interpolating minima we have

$$\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}^* = \left(\mathbf{I} - \frac{\eta}{|\mathfrak{B}_t|} \sum_{i \in \mathfrak{B}_t} \mathbf{H}_i \right) (\boldsymbol{\theta}_t - \boldsymbol{\theta}^*) = \mathbf{A}_t (\boldsymbol{\theta}_t - \boldsymbol{\theta}^*). \quad (163)$$

Thus,

$$\begin{aligned}
\boldsymbol{\Sigma}_{t+1} &= \mathbb{E} \left[(\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}^*) (\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}^*)^\top \right] \\
&= \mathbb{E} \left[\mathbf{A}_t (\boldsymbol{\theta}_t - \boldsymbol{\theta}^*) (\boldsymbol{\theta}_t - \boldsymbol{\theta}^*)^\top \mathbf{A}_t \right] \\
&= \mathbb{E} \left[\mathbf{A}_t \mathbb{E} \left[(\boldsymbol{\theta}_t - \boldsymbol{\theta}^*) (\boldsymbol{\theta}_t - \boldsymbol{\theta}^*)^\top \middle| \mathfrak{B}_t \right] \mathbf{A}_t \right] \\
&= \mathbb{E} \left[\mathbf{A}_t \mathbb{E} \left[(\boldsymbol{\theta}_t - \boldsymbol{\theta}^*) (\boldsymbol{\theta}_t - \boldsymbol{\theta}^*)^\top \right] \mathbf{A}_t \right] \\
&= \mathbb{E} \left[\mathbf{A}_t \boldsymbol{\Sigma}_t \mathbf{A}_t \right].
\end{aligned} \tag{164}$$

Using vectorization we get

$$\text{vec}(\boldsymbol{\Sigma}_{t+1}) = \mathbb{E} [\mathbf{A}_t \otimes \mathbf{A}_t] \text{vec}(\boldsymbol{\Sigma}_t). \tag{165}$$

Let us compute this term.

$$\begin{aligned}
\mathbb{E} [\mathbf{A}_t \otimes \mathbf{A}_t] &= \mathbb{P}(\beta_t = 0) \mathbb{E} [\mathbf{A}_t \otimes \mathbf{A}_t | \beta_t = 0] + \mathbb{P}(\beta_t = 1) \mathbb{E} [\mathbf{A}_t \otimes \mathbf{A}_t | \beta_t = 1] \\
&= (1-p) \mathbf{Q}(\eta, B=n) + p \mathbf{Q}(\eta, B=1) \\
&= (1-p) \times (\mathbf{I} - \eta \mathbf{H}) \otimes (\mathbf{I} - \eta \mathbf{H}) + p \times \frac{1}{n} \sum_{i=1}^n (\mathbf{I} - \eta \mathbf{H}_i) \otimes (\mathbf{I} - \eta \mathbf{H}_i).
\end{aligned} \tag{166}$$

This is the same matrix that we had for mini-batch SGD with batch size B such that $p = \frac{n-B}{B(n-1)}$. Therefore, the stability threshold is the same.

F PROOF OF PROPOSITION 3

The stability threshold given by Theorem 3 and Theorem 4 is

$$\eta_{\text{var}}^* = \frac{2}{\lambda_{\max}(\mathbf{C}^\dagger \mathbf{D})} \tag{167}$$

where

$$\mathbf{C} = \frac{1}{2} \mathbf{H} \oplus \mathbf{H}, \quad \mathbf{D} = (1-p) \mathbf{H} \otimes \mathbf{H} + p \frac{1}{n} \sum_{i=1}^n \mathbf{H}_i \otimes \mathbf{H}_i. \tag{168}$$

This threshold corresponds to a necessary and sufficient condition for stability. Here we derive simplified necessary conditions for stability. In App. B.7 we show that

$$\frac{2}{\lambda_{\max}(\mathbf{C}^\dagger \mathbf{D})} = 2 \inf_{\mathbf{u} \in \mathbb{S}^{d^2-1}, \mathbf{u} \notin \mathcal{N}(\mathbf{D})} \left\{ \frac{\mathbf{u}^\top \mathbf{C} \mathbf{u}}{\mathbf{u}^\top \mathbf{D} \mathbf{u}} \right\}. \tag{169}$$

We shall upper bound the stability threshold by considering non-optimal yet interesting vectors \mathbf{u} . Specifically, in the following we look at $\mathbf{u} = \mathbf{v}_{\max} \otimes \mathbf{v}_{\max}$, where \mathbf{v}_{\max} is the top eigenvector of \mathbf{H} , and $\mathbf{u} = \text{vec}(\mathbf{I})$ to obtain the results of Proposition 3.

F.1 SETTING $\mathbf{u} = \mathbf{v}_{\max} \otimes \mathbf{v}_{\max}$

Let $\mathbf{u} = \mathbf{v} \otimes \mathbf{v} \notin \mathcal{N}(\mathbf{D})$ where $\|\mathbf{v}\| = 1$, then

$$\begin{aligned}
\mathbf{u}^\top \mathbf{C} \mathbf{u} &= \frac{1}{2} \mathbf{u}^\top (\mathbf{H} \otimes \mathbf{I} + \mathbf{I} \otimes \mathbf{H}) \mathbf{u} \\
&= \frac{1}{2} [(\mathbf{v}^\top \otimes \mathbf{v}^\top) (\mathbf{H} \otimes \mathbf{I}) (\mathbf{v} \otimes \mathbf{v}) + (\mathbf{v}^\top \otimes \mathbf{v}^\top) (\mathbf{I} \otimes \mathbf{H}) (\mathbf{v} \otimes \mathbf{v})] \\
&= \frac{1}{2} [(\mathbf{v}^\top \mathbf{H} \mathbf{v}) \otimes (\mathbf{v}^\top \mathbf{v}) + (\mathbf{v}^\top \mathbf{v}) \otimes (\mathbf{v}^\top \mathbf{H} \mathbf{v})] \\
&= \frac{1}{2} [(\mathbf{v}^\top \mathbf{H} \mathbf{v}) \otimes 1 + 1 \otimes (\mathbf{v}^\top \mathbf{H} \mathbf{v})] \\
&= \mathbf{v}^\top \mathbf{H} \mathbf{v}.
\end{aligned} \tag{170}$$

Similarly,

$$\begin{aligned}
\mathbf{u}^\top (\mathbf{H} \otimes \mathbf{H}) \mathbf{u} &= (\mathbf{v}^\top \otimes \mathbf{v}^\top) (\mathbf{H} \otimes \mathbf{H}) (\mathbf{v} \otimes \mathbf{v}) \\
&= (\mathbf{v}^\top \mathbf{H} \mathbf{v}) \otimes (\mathbf{v}^\top \mathbf{H} \mathbf{v}) \\
&= (\mathbf{v}^\top \mathbf{H} \mathbf{v}) \otimes (\mathbf{v}^\top \mathbf{H} \mathbf{v}) \\
&= (\mathbf{v}^\top \mathbf{H} \mathbf{v})^2.
\end{aligned} \tag{171}$$

And again

$$\begin{aligned}
\mathbf{u}^\top \left(\frac{1}{n} \sum_{i=1}^n \mathbf{H}_i \otimes \mathbf{H}_i \right) \mathbf{u} &= \frac{1}{n} \sum_{i=1}^n (\mathbf{v}^\top \otimes \mathbf{v}^\top) (\mathbf{H}_i \otimes \mathbf{H}_i) (\mathbf{v} \otimes \mathbf{v}) \\
&= \frac{1}{n} \sum_{i=1}^n (\mathbf{v}^\top \mathbf{H}_i \mathbf{v}) \otimes (\mathbf{v}^\top \mathbf{H}_i \mathbf{v}) \\
&= \frac{1}{n} \sum_{i=1}^n (\mathbf{v}^\top \mathbf{H}_i \mathbf{v})^2.
\end{aligned} \tag{172}$$

Thus,

$$\begin{aligned}
\mathbf{u}^\top \mathbf{D} \mathbf{u} &= (1-p) \mathbf{u}^\top (\mathbf{H} \otimes \mathbf{H}) \mathbf{u} + p \mathbf{u}^\top \left(\frac{1}{n} \sum_{i=1}^n \mathbf{H}_i \otimes \mathbf{H}_i \right) \mathbf{u} \\
&= (1-p) (\mathbf{v}^\top \mathbf{H} \mathbf{v})^2 + p \frac{1}{n} \sum_{i=1}^n (\mathbf{v}^\top \mathbf{H}_i \mathbf{v})^2 \\
&= (\mathbf{v}^\top \mathbf{H} \mathbf{v})^2 + p \left[\frac{1}{n} \sum_{i=1}^n (\mathbf{v}^\top \mathbf{H}_i \mathbf{v})^2 - (\mathbf{v}^\top \mathbf{H} \mathbf{v})^2 \right] \\
&= (\mathbf{v}^\top \mathbf{H} \mathbf{v})^2 + p \left[\frac{1}{n} \sum_{i=1}^n (\mathbf{v}^\top \mathbf{H}_i \mathbf{v})^2 - 2 (\mathbf{v}^\top \mathbf{H} \mathbf{v}) (\mathbf{v}^\top \mathbf{H} \mathbf{v}) + (\mathbf{v}^\top \mathbf{H} \mathbf{v})^2 \right] \\
&= (\mathbf{v}^\top \mathbf{H} \mathbf{v})^2 + p \frac{1}{n} \sum_{i=1}^n \left[(\mathbf{v}^\top \mathbf{H}_i \mathbf{v})^2 - 2 (\mathbf{v}^\top \mathbf{H}_i \mathbf{v}) (\mathbf{v}^\top \mathbf{H} \mathbf{v}) + (\mathbf{v}^\top \mathbf{H} \mathbf{v})^2 \right] \\
&= (\mathbf{v}^\top \mathbf{H} \mathbf{v})^2 + p \frac{1}{n} \sum_{i=1}^n (\mathbf{v}^\top \mathbf{H}_i \mathbf{v} - \mathbf{v}^\top \mathbf{H} \mathbf{v})^2.
\end{aligned} \tag{173}$$

Therefore, for general $\mathbf{u} = \mathbf{v} \otimes \mathbf{v}$ we get

$$\eta_{\text{var}}^* \leq 2 \frac{\mathbf{u}^\top \mathbf{C} \mathbf{u}}{\mathbf{u}^\top \mathbf{D} \mathbf{u}} = \frac{2 \mathbf{v}^\top \mathbf{H} \mathbf{v}}{(\mathbf{v}^\top \mathbf{H} \mathbf{v})^2 + \frac{p}{n} \sum_{i=1}^n (\mathbf{v}^\top \mathbf{H}_i \mathbf{v} - \mathbf{v}^\top \mathbf{H} \mathbf{v})^2}. \tag{174}$$

Specifically, for $\mathbf{u} = \mathbf{v}_{\max} \otimes \mathbf{v}_{\max}$ we get

$$\eta_{\text{var}}^* \leq \frac{2 \lambda_{\max}(\mathbf{H})}{\lambda_{\max}^2(\mathbf{H}) + \frac{p}{n} \sum_{i=1}^n (\mathbf{v}_{\max}^\top \mathbf{H}_i \mathbf{v}_{\max} - \lambda_{\max}(\mathbf{H}))^2}. \tag{175}$$

Finally, from (174) we get the following result which we used in Sec.4.

$$\lambda_{\max}(\mathbf{C}^\dagger \mathbf{D}) = \frac{2}{\eta_{\text{var}}^*} \geq \mathbf{v}^\top \mathbf{H} \mathbf{v} + p \frac{\frac{1}{n} \sum_{i=1}^n (\mathbf{v}^\top \mathbf{H}_i \mathbf{v} - \mathbf{v}^\top \mathbf{H} \mathbf{v})^2}{\mathbf{v}^\top \mathbf{H} \mathbf{v}}. \tag{176}$$

Since this inequality holds for every $\mathbf{v} \notin \mathcal{N}(\mathbf{H})$, we can take the maximum to obtain (24).

F.2 SETTING $\mathbf{u} = \text{vec}(\mathbf{I})$

Let $\mathbf{u} = \text{vec}(\mathbf{I}) \notin \mathcal{N}(\mathbf{D})$, then

$$\begin{aligned} \mathbf{u}^\top \mathbf{C} \mathbf{u} &= \frac{1}{2} \mathbf{u}^\top (\mathbf{H} \otimes \mathbf{I} + \mathbf{I} \otimes \mathbf{H}) \mathbf{u} \\ &= \frac{1}{2} \left([\text{vec}(\mathbf{I})]^\top (\mathbf{H} \otimes \mathbf{I}) \text{vec}(\mathbf{I}) + [\text{vec}(\mathbf{I})]^\top (\mathbf{I} \otimes \mathbf{H}) \text{vec}(\mathbf{I}) \right) \\ &= \frac{1}{2} \left(\text{Tr}(\mathbf{H}^\top) + \text{Tr}(\mathbf{H}) \right) \\ &= \text{Tr}(\mathbf{H}), \end{aligned} \quad (177)$$

where we used (P4). Moreover,

$$\mathbf{u}^\top (\mathbf{H} \otimes \mathbf{H}) \mathbf{u} = [\text{vec}(\mathbf{I})]^\top (\mathbf{H} \otimes \mathbf{H}) \text{vec}(\mathbf{I}) = \text{Tr}(\mathbf{H}\mathbf{H}^\top) = \|\mathbf{H}\|_{\text{F}}^2.$$

Similarly,

$$\mathbf{u}^\top (\mathbf{H}_i \otimes \mathbf{H}_i) \mathbf{u} = [\text{vec}(\mathbf{I})]^\top (\mathbf{H}_i \otimes \mathbf{H}_i) \text{vec}(\mathbf{I}) = \text{Tr}(\mathbf{H}_i \mathbf{H}_i^\top) = \|\mathbf{H}_i\|_{\text{F}}^2.$$

$$\begin{aligned} \mathbf{u}^\top \mathbf{D} \mathbf{u} &= (1-p) \mathbf{u}^\top (\mathbf{H} \otimes \mathbf{H}) \mathbf{u} + p \frac{1}{n} \sum_{i=1}^n \mathbf{u}^\top \mathbf{H}_i \otimes \mathbf{H}_i \mathbf{u} \\ &= (1-p) \|\mathbf{H}\|_{\text{F}}^2 + p \frac{1}{n} \sum_{i=1}^n \|\mathbf{H}_i\|_{\text{F}}^2. \end{aligned} \quad (178)$$

Therefore,

$$\eta_{\text{var}}^* \leq 2 \frac{\mathbf{u}^\top \mathbf{C} \mathbf{u}}{\mathbf{u}^\top \mathbf{D} \mathbf{u}} = \frac{2\text{Tr}(\mathbf{H})}{(1-p) \|\mathbf{H}\|_{\text{F}}^2 + p \frac{1}{n} \sum_{i=1}^n \|\mathbf{H}_i\|_{\text{F}}^2}. \quad (179)$$

G PROOF OF THEOREM 5

In this section, we use the following result on the Moore–Penrose inverse of a sum of two matrices.

Theorem 7 (Fill & Fishkind (2000), Thm. 3). *Let $\mathbf{X}, \mathbf{Y} \in \mathbb{R}^{p \times p}$ with $\text{rank}(\mathbf{X} + \mathbf{Y}) = \text{rank}(\mathbf{X}) + \text{rank}(\mathbf{Y})$. Then*

$$(\mathbf{X} + \mathbf{Y})^\dagger = (\mathbf{I} - \mathbf{L})\mathbf{X}^\dagger(\mathbf{I} - \mathbf{O}) + \mathbf{L}\mathbf{Y}^\dagger\mathbf{O}, \quad (180)$$

where

$$\mathbf{L} = (\mathbf{P}_{\mathcal{R}(\mathbf{Y}^\top)} \mathbf{P}_{\mathcal{R}^\perp(\mathbf{X}^\top)})^\dagger \quad \text{and} \quad \mathbf{O} = (\mathbf{P}_{\mathcal{R}^\perp(\mathbf{X})} \mathbf{P}_{\mathcal{R}(\mathbf{Y})})^\dagger. \quad (181)$$

Moreover, we use the following relations.

$$\begin{aligned} \mathcal{R}(\mathbf{D}) &= \mathcal{R}(\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})}), \\ \mathcal{R}(\mathbf{C}) &= \mathcal{R}(\mathbf{I} - \mathbf{P}_{\mathcal{N}(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}(\mathbf{H})}), \\ \mathcal{R}(\mathbf{P}_{\mathcal{N}(\mathbf{D})} \mathbf{C}) &= \mathcal{R}(\mathbf{P}_{\mathcal{N}(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} + \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}(\mathbf{H})}). \end{aligned} \quad (182)$$

The dynamics of $\boldsymbol{\mu}_t^\dagger$ and $\boldsymbol{\Sigma}_t^\dagger$ are given by (see (53))

$$\begin{pmatrix} \boldsymbol{\mu}_{t+1}^\dagger \\ \text{vec}(\boldsymbol{\Sigma}_{t+1}^\dagger) \end{pmatrix} = \Xi \begin{pmatrix} \boldsymbol{\mu}_t^\dagger \\ \text{vec}(\boldsymbol{\Sigma}_t^\dagger) \end{pmatrix} + \begin{pmatrix} \mathbf{0} \\ \text{vec}(\boldsymbol{\Sigma}_v^\dagger) \end{pmatrix}. \quad (183)$$

where

$$\begin{aligned} \Xi &= \begin{pmatrix} \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} - \eta \mathbf{H} & \mathbf{0} \\ -(\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})}) (\mathbb{E}[\mathbf{v}_t^\dagger \otimes \mathbf{A}_t] + \mathbb{E}[\mathbf{A}_t \otimes \mathbf{v}_t^\dagger]) & (\mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})}) \mathbf{Q} \end{pmatrix} \\ &\triangleq \begin{pmatrix} \Xi_{1,1} & \Xi_{1,2} \\ \Xi_{2,1} & \Xi_{2,2} \end{pmatrix}. \end{aligned} \quad (184)$$

In App. B.8 we show that if $0 < \eta < \eta_{\text{var}}^*$ then the spectral radius of Ξ is less than one. Therefore, the dynamical system is stable, and the asymptotic values of μ_t^\pm and Σ_t^\pm as $t \rightarrow \infty$ are given by

$$\lim_{t \rightarrow \infty} \begin{pmatrix} \mu_t^\pm \\ \text{vec}(\Sigma_t^\pm) \end{pmatrix} = (\mathbf{I} - \Xi)^{-1} \begin{pmatrix} \mathbf{0} \\ \text{vec}(\Sigma_v^\pm) \end{pmatrix}. \quad (185)$$

Using the inversion formula for block matrix and the fact that $\Xi_{1,2} = \mathbf{0}$ we have that

$$\begin{aligned} (\mathbf{I} - \Xi)^{-1} &= \begin{pmatrix} \mathbf{I} - \Xi_{1,1} & -\Xi_{1,2} \\ -\Xi_{2,1} & \mathbf{I} - \Xi_{2,2} \end{pmatrix}^{-1} \\ &= \begin{pmatrix} (\mathbf{I} - \Xi_{1,1} - \Xi_{1,2}(\mathbf{I} - \Xi_{2,2})^{-1}\Xi_{2,1})^{-1} & \mathbf{0} \\ \mathbf{0} & (\mathbf{I} - \Xi_{2,2} - \Xi_{2,1}(\mathbf{I} - \Xi_{1,1})^{-1}\Xi_{1,2})^{-1} \end{pmatrix} \\ &\quad \begin{pmatrix} \mathbf{I} & \Xi_{1,2}(\mathbf{I} - \Xi_{2,2})^{-1} \\ \Xi_{2,1}(\mathbf{I} - \Xi_{1,1})^{-1} & \mathbf{I} \end{pmatrix} \\ &= \begin{pmatrix} (\mathbf{I} - \Xi_{1,1})^{-1} & \mathbf{0} \\ \mathbf{0} & (\mathbf{I} - \Xi_{2,2})^{-1} \end{pmatrix} \begin{pmatrix} \mathbf{I} & \mathbf{0} \\ \Xi_{2,1}(\mathbf{I} - \Xi_{1,1})^{-1} & \mathbf{I} \end{pmatrix}. \end{aligned} \quad (186)$$

Therefore,

$$\begin{aligned} \lim_{t \rightarrow \infty} \begin{pmatrix} \mu_t^\pm \\ \text{vec}(\Sigma_t^\pm) \end{pmatrix} &= \begin{pmatrix} (\mathbf{I} - \Xi_{1,1})^{-1} & \mathbf{0} \\ \mathbf{0} & (\mathbf{I} - \Xi_{2,2})^{-1} \end{pmatrix} \begin{pmatrix} \mathbf{I} & \mathbf{0} \\ \Xi_{2,1}(\mathbf{I} - \Xi_{1,1})^{-1} & \mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{0} \\ \text{vec}(\Sigma_v^\pm) \end{pmatrix} \\ &= \begin{pmatrix} (\mathbf{I} - \Xi_{1,1})^{-1} & \mathbf{0} \\ \mathbf{0} & (\mathbf{I} - \Xi_{2,2})^{-1} \end{pmatrix} \begin{pmatrix} \mathbf{0} \\ \text{vec}(\Sigma_v^\pm) \end{pmatrix}. \end{aligned} \quad (187)$$

Namely,

$$\lim_{t \rightarrow \infty} \mu_t^\pm = \mathbf{0} \quad \text{and} \quad \lim_{t \rightarrow \infty} \text{vec}(\Sigma_t^\pm) = (\mathbf{I} - \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}\mathbf{Q})^{-1} \text{vec}(\Sigma_v^\pm). \quad (188)$$

Now,

$$\begin{aligned} (\mathbf{I} - \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}\mathbf{Q})^{-1} &= (\mathbf{I} - \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}\mathbf{Q})^\dagger \\ &= (\mathbf{P}_{\mathcal{N}(\mathcal{D})} + \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})} - \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}\mathbf{Q})^\dagger \\ &= (\mathbf{P}_{\mathcal{N}(\mathcal{D})} + \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}(\mathbf{I} - \mathbf{Q}))^\dagger. \end{aligned} \quad (189)$$

Let us apply Thm. 7 on $(\mathbf{P}_{\mathcal{D}} + \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}(\mathbf{I} - \mathbf{Q}))^\dagger$. Here, $\mathbf{X}_1 = \mathbf{P}_{\mathcal{N}(\mathcal{D})}$ and $\mathbf{Y}_1 = \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}(\mathbf{I} - \mathbf{Q})$. Note that $\mathcal{R}(\mathbf{X}_1) = \mathcal{R}^\perp(\mathcal{D})$ and $\mathcal{R}(\mathbf{Y}_1) = \mathcal{R}(\mathcal{D})$ and therefore $\text{rank}(\mathbf{X}_1 + \mathbf{Y}_1) = \text{rank}(\mathbf{X}_1) + \text{rank}(\mathbf{Y}_1)$. Additionally,

$$\mathbf{P}_{\mathcal{R}(\mathbf{Y}_1^\top)} = \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}, \quad \mathbf{P}_{\mathcal{R}^\perp(\mathbf{X}_1^\top)} = \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}, \quad \mathbf{P}_{\mathcal{R}^\perp(\mathbf{X}_1)} = \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}, \quad \mathbf{P}_{\mathcal{R}(\mathbf{Y}_1)} = \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}. \quad (190)$$

Hence,

$$\begin{aligned} \mathbf{L}_1 &= (\mathbf{P}_{\mathcal{R}(\mathbf{Y}_1^\top)}\mathbf{P}_{\mathcal{R}^\perp(\mathbf{X}_1^\top)})^\dagger = (\mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}\mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})})^\dagger = \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}, \\ \mathbf{O}_1 &= (\mathbf{P}_{\mathcal{R}^\perp(\mathbf{X}_1)}\mathbf{P}_{\mathcal{R}(\mathbf{Y}_1)})^\dagger = (\mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}\mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})})^\dagger = \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}. \end{aligned} \quad (191)$$

Therefore,

$$\begin{aligned} (\mathbf{I} - \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}\mathbf{Q})^{-1} &= (\mathbf{P}_{\mathcal{N}(\mathcal{D})} + \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}(\mathbf{I} - \mathbf{Q}))^\dagger \\ &= (\mathbf{X}_1 + \mathbf{Y}_1)^\dagger \\ &= (\mathbf{I} - \mathbf{L}_1)\mathbf{X}_1^\dagger(\mathbf{I} - \mathbf{O}_1) + \mathbf{L}_1\mathbf{Y}_1^\dagger\mathbf{O}_1 \\ &= (\mathbf{I} - \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})})(\mathbf{P}_{\mathcal{N}(\mathcal{D})})^\dagger(\mathbf{I} - \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}) \\ &\quad + \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}(\mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}(\mathbf{I} - \mathbf{Q}))^\dagger\mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})} \\ &= (\mathbf{I} - \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})})\mathbf{P}_{\mathcal{N}(\mathcal{D})}(\mathbf{I} - \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}) \\ &\quad + \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}(\mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}(\mathbf{I} - \mathbf{Q}))^\dagger\mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})} \\ &= \mathbf{P}_{\mathcal{N}(\mathcal{D})} + \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}(\mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}(\mathbf{I} - \mathbf{Q}))^\dagger\mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}, \end{aligned} \quad (192)$$

where in the third step we used Thm. 7. Thus we get the following intermediate result

$$\begin{aligned} \lim_{t \rightarrow \infty} \text{vec}(\Sigma_t^\perp) &= (\mathbf{I} - \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}\mathbf{Q})^{-1} \text{vec}(\Sigma_v^\perp) \\ &= \left(\mathbf{P}_{\mathcal{N}(\mathcal{D})} + \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})} (\mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}(\mathbf{I} - \mathbf{Q}))^\dagger \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})} \right) \text{vec}(\Sigma_v^\perp) \\ &= \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})} (\mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}(\mathbf{I} - \mathbf{Q}))^\dagger \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})} \text{vec}(\Sigma_v^\perp), \end{aligned} \quad (193)$$

where in the final step we used $\mathbf{P}_{\mathcal{N}(\mathcal{D})} \text{vec}(\Sigma_v^\perp) = \mathbf{0}$. Now, note that $\mathcal{R}(\mathbf{P}_{\mathcal{N}(\mathcal{D})}\mathbf{C}) = \mathcal{R}(\mathbf{P}_{\mathcal{N}(\mathcal{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})} + \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})} \otimes \mathbf{P}_{\mathcal{N}(\mathcal{H})})$, whereas $\text{vec}(\Sigma_v^\perp) \in \mathcal{R}(\mathcal{D}) = \mathcal{R}(\mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})})$ and therefore $(\mathbf{P}_{\mathcal{N}(\mathcal{D})}\mathbf{C})^\dagger \text{vec}(\Sigma_v^\perp) = \mathbf{0}$. Hence,

$$\begin{aligned} \lim_{t \rightarrow \infty} \text{vec}(\Sigma_t^\perp) &= \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})} (\mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}(\mathbf{I} - \mathbf{Q}))^\dagger \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})} \text{vec}(\Sigma_v^\perp) \\ &= \left((2\eta \mathbf{P}_{\mathcal{N}(\mathcal{D})}\mathbf{C})^\dagger + \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})} (\mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}(\mathbf{I} - \mathbf{Q}))^\dagger \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})} \right) \text{vec}(\Sigma_v^\perp). \end{aligned} \quad (194)$$

Let us apply again Thm. 7 but in the other direction. This time, $\mathbf{X}_2 = 2\eta \mathbf{P}_{\mathcal{N}(\mathcal{D})}\mathbf{C}$ and $\mathbf{Y}_2 = \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}(\mathbf{I} - \mathbf{Q})$. Note that $\mathcal{R}(\mathbf{X}_2) = \mathcal{R}(\mathbf{P}_{\mathcal{N}(\mathcal{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})} + \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})} \otimes \mathbf{P}_{\mathcal{N}(\mathcal{H})})$ and $\mathcal{R}(\mathbf{Y}_2) = \mathcal{R}(\mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})})$ and therefore $\text{rank}(\mathbf{X}_2 + \mathbf{Y}_2) = \text{rank}(\mathbf{X}_2) + \text{rank}(\mathbf{Y}_2)$. Additionally,

$$\begin{aligned} \mathbf{P}_{\mathcal{R}(\mathbf{Y}_2^\top)} &= \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})}, \\ \mathbf{P}_{\mathcal{R}(\mathbf{Y}_2)} &= \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})}, \\ \mathbf{P}_{\mathcal{R}^\perp(\mathbf{X}_2^\top)} &= \mathbf{P}_{\mathcal{N}(\mathcal{H})} \otimes \mathbf{P}_{\mathcal{N}(\mathcal{H})} + \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})}, \\ \mathbf{P}_{\mathcal{R}^\perp(\mathbf{X}_2)} &= \mathbf{P}_{\mathcal{N}(\mathcal{H})} \otimes \mathbf{P}_{\mathcal{N}(\mathcal{H})} + \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})}. \end{aligned} \quad (195)$$

Hence,

$$\begin{aligned} \mathbf{L}_2 &= \left(\mathbf{P}_{\mathcal{R}(\mathbf{Y}_2^\top)} \mathbf{P}_{\mathcal{R}^\perp(\mathbf{X}_2^\top)} \right)^\dagger \\ &= \left(\mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})} (\mathbf{P}_{\mathcal{N}(\mathcal{H})} \otimes \mathbf{P}_{\mathcal{N}(\mathcal{H})} + \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})}) \right)^\dagger \\ &= \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})} \\ &= \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}, \\ \mathbf{O}_2 &= \left(\mathbf{P}_{\mathcal{R}^\perp(\mathbf{X}_2)} \mathbf{P}_{\mathcal{R}(\mathbf{Y}_2)} \right)^\dagger \\ &= \left((\mathbf{P}_{\mathcal{N}(\mathcal{H})} \otimes \mathbf{P}_{\mathcal{N}(\mathcal{H})} + \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})}) \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})} \right)^\dagger \\ &= \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})} = \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}. \end{aligned} \quad (196)$$

Moreover, since $\mathcal{R}(\mathbf{X}_2) = \mathcal{R}(\mathbf{P}_{\mathcal{N}(\mathcal{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})} + \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})} \otimes \mathbf{P}_{\mathcal{N}(\mathcal{H})})$ and $\mathcal{R}^\perp(\mathbf{Y}_2) = \mathcal{R}(\mathbf{P}_{\mathcal{N}(\mathcal{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})} + \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})} \otimes \mathbf{P}_{\mathcal{N}(\mathcal{H})} + \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})} \otimes \mathbf{P}_{\mathcal{N}^\perp(\mathcal{H})})$ we have that $\mathcal{R}(\mathbf{X}_2) \subseteq \mathcal{R}^\perp(\mathbf{Y}_2) = \mathcal{N}(\mathcal{D})$. Therefore,

$$(\mathbf{I} - \mathbf{L}_2) \mathbf{X}_2^\dagger (\mathbf{I} - \mathbf{O}_2) = (\mathbf{I} - \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}) \mathbf{X}_2^\dagger (\mathbf{I} - \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}) = \mathbf{P}_{\mathcal{N}(\mathcal{D})} \mathbf{X}_2^\dagger \mathbf{P}_{\mathcal{N}(\mathcal{D})} = \mathbf{X}_2^\dagger. \quad (197)$$

Therefore, applying Thm. 7 we get

$$\begin{aligned} (2\eta \mathbf{P}_{\mathcal{N}(\mathcal{D})}\mathbf{C})^\dagger + \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})} (\mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}(\mathbf{I} - \mathbf{Q}))^\dagger \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})} & \\ &= \mathbf{X}_2^\dagger + \mathbf{L}_2 \mathbf{Y}_2^\dagger \mathbf{O}_2 \\ &= (\mathbf{I} - \mathbf{L}_2) \mathbf{X}_2^\dagger (\mathbf{I} - \mathbf{O}_2) + \mathbf{L}_2 \mathbf{Y}_2^\dagger \mathbf{O}_2 \\ &= (\mathbf{X}_2 + \mathbf{Y}_2)^\dagger \\ &= (2\eta \mathbf{P}_{\mathcal{N}(\mathcal{D})}\mathbf{C} + \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}(\mathbf{I} - \mathbf{Q}))^\dagger \\ &= (2\eta \mathbf{P}_{\mathcal{N}(\mathcal{D})}\mathbf{C} + 2\eta \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}\mathbf{C} - \eta^2 \mathbf{P}_{\mathcal{N}^\perp(\mathcal{D})}\mathbf{D})^\dagger \\ &= (2\eta \mathbf{C} - \eta^2 \mathbf{D})^\dagger. \end{aligned} \quad (198)$$

where in the second step we used (197), and in the third step we used Thm. 7. Overall, together with (63) we get

$$\begin{aligned} \lim_{t \rightarrow \infty} \text{vec}(\boldsymbol{\Sigma}_t^\perp) &= (2\eta\mathbf{C} - \eta^2\mathbf{D})^\dagger \text{vec}(\boldsymbol{\Sigma}_v^\perp) \\ &= \left(\frac{1}{\eta}(2\mathbf{C} - \eta\mathbf{D})^\dagger\right) (\eta^2 p \text{vec}(\boldsymbol{\Sigma}_g^\perp)) \\ &= \eta p (2\mathbf{C} - \eta\mathbf{D})^\dagger \text{vec}(\boldsymbol{\Sigma}_g^\perp). \end{aligned} \quad (199)$$

H PROOF OF COROLLARY 1

From Thm. 5 we have that if $0 < \eta < \eta_{\text{var}}^*$ then

$$\lim_{t \rightarrow \infty} \text{vec}(\boldsymbol{\Sigma}_t^\perp) = \eta p (2\mathbf{C} - \eta\mathbf{D})^\dagger \text{vec}(\boldsymbol{\Sigma}_g^\perp), \quad (200)$$

Using this result, we prove Corollary 1.

First statement. If $0 < \eta < \eta_{\text{var}}^*$ then by Prop. 5

$$\begin{aligned} \lim_{t \rightarrow \infty} \mathbb{E}[\|\boldsymbol{\theta}_t^\perp - \boldsymbol{\theta}^{*\perp}\|^2] &= (\text{vec}(\mathbf{I}))^\top \lim_{t \rightarrow \infty} \text{vec}(\boldsymbol{\Sigma}_t^\perp) \\ &= (\text{vec}(\mathbf{I}))^\top \left(\eta p (2\mathbf{C} - \eta\mathbf{D})^\dagger \text{vec}(\boldsymbol{\Sigma}_g^\perp) \right) \\ &= \eta p (\text{vec}(\mathbf{I}))^\top (2\mathbf{C} - \eta\mathbf{D})^\dagger \text{vec}(\boldsymbol{\Sigma}_g^\perp). \end{aligned} \quad (201)$$

Second statement. Similarly, let us compute the limit of the expected value of the loss function to obtain point 2.

$$\begin{aligned} \lim_{t \rightarrow \infty} \mathbb{E}[\tilde{\mathcal{L}}(\boldsymbol{\theta}_t)] - \mathcal{L}(\boldsymbol{\theta}^*) &= \frac{1}{2} \lim_{t \rightarrow \infty} \mathbb{E}[(\boldsymbol{\theta}_t - \boldsymbol{\theta}^*)^\top \mathbf{H}(\boldsymbol{\theta}_t - \boldsymbol{\theta}^*)] \\ &= \frac{1}{2} \lim_{t \rightarrow \infty} \mathbb{E}[(\boldsymbol{\theta}_t - \boldsymbol{\theta}^*)^\top \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \mathbf{H} \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} (\boldsymbol{\theta}_t - \boldsymbol{\theta}^*)] \\ &= \frac{1}{2} \lim_{t \rightarrow \infty} \mathbb{E}[(\boldsymbol{\theta}_t^\perp - \boldsymbol{\theta}^{*\perp})^\top \mathbf{H}(\boldsymbol{\theta}_t^\perp - \boldsymbol{\theta}^{*\perp})] \\ &= \frac{1}{2} \text{Tr} \left(\mathbf{H} \lim_{t \rightarrow \infty} \mathbb{E}[(\boldsymbol{\theta}_t^\perp - \boldsymbol{\theta}^{*\perp})(\boldsymbol{\theta}_t^\perp - \boldsymbol{\theta}^{*\perp})^\top] \right) \\ &= \frac{1}{2} \text{Tr} \left(\mathbf{H} \lim_{t \rightarrow \infty} \boldsymbol{\Sigma}_t^\perp \right) \\ &= \frac{1}{2} (\text{vec}(\mathbf{H}))^\top \lim_{t \rightarrow \infty} \text{vec}(\boldsymbol{\Sigma}_t^\perp) \\ &= \frac{1}{2} (\text{vec}(\mathbf{H}))^\top \left(\eta p (2\mathbf{C} - \eta\mathbf{D})^\dagger \text{vec}(\boldsymbol{\Sigma}_g^\perp) \right) \\ &= \frac{1}{2} \eta p (\text{vec}(\mathbf{H}))^\top (2\mathbf{C} - \eta\mathbf{D})^\dagger \text{vec}(\boldsymbol{\Sigma}_g^\perp). \end{aligned} \quad (202)$$

Third statement. Finally, we prove point 3. The gradient of the second-order Taylor expansion of the loss is given by

$$\nabla \tilde{\mathcal{L}}(\boldsymbol{\theta}) = \mathbf{H}(\boldsymbol{\theta} - \boldsymbol{\theta}^*). \quad (203)$$

Therefore

$$\begin{aligned}
\lim_{t \rightarrow \infty} \mathbb{E} \left[\left\| \nabla \tilde{\mathcal{L}}(\boldsymbol{\theta}_t) \right\|^2 \right] &= \lim_{t \rightarrow \infty} \mathbb{E} [(\boldsymbol{\theta}_t - \boldsymbol{\theta}^*)^\top \mathbf{H}^2 (\boldsymbol{\theta}_t - \boldsymbol{\theta}^*)] \\
&= \lim_{t \rightarrow \infty} \mathbb{E} [(\boldsymbol{\theta}_t - \boldsymbol{\theta}^*)^\top \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} \mathbf{H}^2 \mathbf{P}_{\mathcal{N}^\perp(\mathbf{H})} (\boldsymbol{\theta}_t - \boldsymbol{\theta}^*)] \\
&= \lim_{t \rightarrow \infty} \mathbb{E} [(\boldsymbol{\theta}_t^\perp - \boldsymbol{\theta}^{*\perp})^\top \mathbf{H}^2 (\boldsymbol{\theta}_t^\perp - \boldsymbol{\theta}^{*\perp})] \\
&= \text{Tr} \left(\mathbf{H}^2 \lim_{t \rightarrow \infty} \mathbb{E} [(\boldsymbol{\theta}_t^\perp - \boldsymbol{\theta}^{*\perp})(\boldsymbol{\theta}_t^\perp - \boldsymbol{\theta}^{*\perp})^\top] \right) \\
&= \text{Tr} \left(\mathbf{H}^2 \lim_{t \rightarrow \infty} \boldsymbol{\Sigma}_t^\perp \right) \\
&= (\text{vec}(\mathbf{H}^2))^\top \lim_{t \rightarrow \infty} \text{vec}(\boldsymbol{\Sigma}_t^\perp) \\
&= (\text{vec}(\mathbf{H}^2))^\top \left(\eta p (2\mathbf{C} - \eta \mathbf{D})^\dagger \text{vec}(\boldsymbol{\Sigma}_g^\perp) \right) \\
&= \eta p (\text{vec}(\mathbf{H}^2))^\top (2\mathbf{C} - \eta \mathbf{D})^\dagger \text{vec}(\boldsymbol{\Sigma}_g^\perp). \tag{204}
\end{aligned}$$

I RECOVERING GD'S STABILITY CONDITION

In this section we show how our stability condition for SGD reduces to GD's when $B = n$. In this case $p = 0$ and thus

$$\mathbf{C} = \frac{1}{2} \mathbf{H} \oplus \mathbf{H}, \quad \mathbf{D} = \mathbf{H} \otimes \mathbf{H}. \tag{205}$$

For simplicity, here we assume that \mathbf{H} has full rank, where the case that \mathbf{H} has null space requires minor adjustments. Let $\mathbf{H} = \mathbf{V} \boldsymbol{\Lambda} \mathbf{V}^\top$ be the eigenvalue decomposition of \mathbf{H} , where $\mathbf{V} \mathbf{V}^\top = \mathbf{V}^\top \mathbf{V} = \mathbf{I}$, then

$$\begin{aligned}
\mathbf{C} &= \frac{1}{2} \mathbf{H} \oplus \mathbf{H} \\
&= \frac{1}{2} (\mathbf{H} \otimes \mathbf{I} + \mathbf{H} \otimes \mathbf{I}) \\
&= \frac{1}{2} \left((\mathbf{V} \boldsymbol{\Lambda} \mathbf{V}^\top) \otimes (\mathbf{V} \mathbf{V}^\top) + (\mathbf{V} \mathbf{V}^\top) \otimes (\mathbf{V} \boldsymbol{\Lambda} \mathbf{V}^\top) \right) \\
&= \frac{1}{2} \left((\mathbf{V} \otimes \mathbf{V}) (\boldsymbol{\Lambda} \otimes \mathbf{I}) (\mathbf{V}^\top \otimes \mathbf{V}^\top) + (\mathbf{V} \otimes \mathbf{V}) (\mathbf{I} \otimes \boldsymbol{\Lambda}) (\mathbf{V}^\top \otimes \mathbf{V}^\top) \right) \\
&= (\mathbf{V} \otimes \mathbf{V}) \left(\frac{1}{2} \boldsymbol{\Lambda} \otimes \mathbf{I} + \frac{1}{2} \mathbf{I} \otimes \boldsymbol{\Lambda} \right) (\mathbf{V} \otimes \mathbf{V})^\top. \tag{206}
\end{aligned}$$

Note that

$$(\mathbf{V} \otimes \mathbf{V})^\top (\mathbf{V} \otimes \mathbf{V}) = (\mathbf{V}^\top \otimes \mathbf{V}^\top) (\mathbf{V} \otimes \mathbf{V}) = (\mathbf{V}^\top \mathbf{V}) \otimes (\mathbf{V}^\top \mathbf{V}) = \mathbf{I} \otimes \mathbf{I} = \mathbf{I}, \tag{207}$$

i.e., $(\mathbf{V} \otimes \mathbf{V})$ is an orthogonal matrix. Since $\frac{1}{2}(\boldsymbol{\Lambda} \otimes \mathbf{I} + \mathbf{I} \otimes \boldsymbol{\Lambda})$ is diagonal, then the last result in (206) is an eigenvalue decomposition of \mathbf{C} . Similarly,

$$\begin{aligned}
\mathbf{D} &= \mathbf{H} \otimes \mathbf{H} \\
&= (\mathbf{V} \boldsymbol{\Lambda} \mathbf{V}^\top) \otimes (\mathbf{V} \boldsymbol{\Lambda} \mathbf{V}^\top) \\
&= (\mathbf{V} \otimes \mathbf{V}) (\boldsymbol{\Lambda} \otimes \boldsymbol{\Lambda}) (\mathbf{V}^\top \otimes \mathbf{V}^\top) \\
&= (\mathbf{V} \otimes \mathbf{V}) (\boldsymbol{\Lambda} \otimes \boldsymbol{\Lambda}) (\mathbf{V} \otimes \mathbf{V})^\top. \tag{208}
\end{aligned}$$

We have that \mathbf{C} and \mathbf{D} have the same eigenvectors. Thus, set $\lambda_\ell = \boldsymbol{\Lambda}_{[\ell, \ell]} = \lambda_\ell(\mathbf{H})$, and define

$$f(x) = \begin{cases} \frac{1}{x}, & x \neq 0, \\ 0, & x = 0. \end{cases} \tag{209}$$

Then

$$\lambda_{\max}(\mathbf{C}^\dagger \mathbf{D}) = \max_{\ell, p \in [d]} \left\{ \lambda_\ell \lambda_p f\left(\frac{1}{2}(\lambda_\ell + \lambda_p)\right) \right\}. \quad (210)$$

Note that the objective vanishes whenever $\lambda_\ell = 0$ or $\lambda_p = 0$. Restricting to only positive eigenvalues gives

$$\lambda_{\max}(\mathbf{C}^\dagger \mathbf{D}) = \max_{\lambda_\ell, \lambda_p > 0} \left\{ \frac{\lambda_\ell \lambda_p}{\frac{1}{2}(\lambda_\ell + \lambda_p)} \right\}. \quad (211)$$

Additionally $\sqrt{\lambda_\ell \lambda_p} \leq \frac{1}{2}(\lambda_\ell + \lambda_p)$ holds for all $\lambda_\ell, \lambda_p > 0$, therefore

$$\begin{aligned} \frac{\lambda_\ell \lambda_p}{\frac{1}{2}(\lambda_\ell + \lambda_p)} &= \frac{\sqrt{\lambda_\ell \lambda_p} \sqrt{\lambda_\ell \lambda_p}}{\frac{1}{2}(\lambda_\ell + \lambda_p)} \\ &\leq \sqrt{\lambda_\ell \lambda_p} \\ &\leq \lambda_{\max}. \end{aligned} \quad (212)$$

Yet, for $\lambda_\ell = \lambda_p = \lambda_{\max}$ we have that

$$\frac{\lambda_\ell \lambda_p}{\frac{1}{2}(\lambda_\ell + \lambda_p)} = \lambda_{\max}. \quad (213)$$

Hence we have

$$\lambda_{\max}(\mathbf{C}^\dagger \mathbf{D}) = \max_{\lambda_\ell, \lambda_p > 0} \left\{ \frac{\lambda_\ell \lambda_p}{\frac{1}{2}(\lambda_\ell + \lambda_p)} \right\} = \lambda_{\max}(\mathbf{H}). \quad (214)$$

J ADDITIONAL EXPERIMENTAL RESULTS AND DETAIL

In this section, we complete the technical detail of the experiment shown in Sec. 4. For the experiment, we used a single-hidden layer ReLU network with fully connected layers (with bias vectors). The number of neurons is 1024, and the total number of parameters is 807,940. We used four classes from MNIST, 256 samples from each class, with a total of 1024 samples. To get large initialization, we used standard torch initialization and multiplied the initial weights by a factor of 15. The maximal number of epochs was set to 4×10^4 . If SGD did not converge within this number of epochs, then we removed this run from the plots.

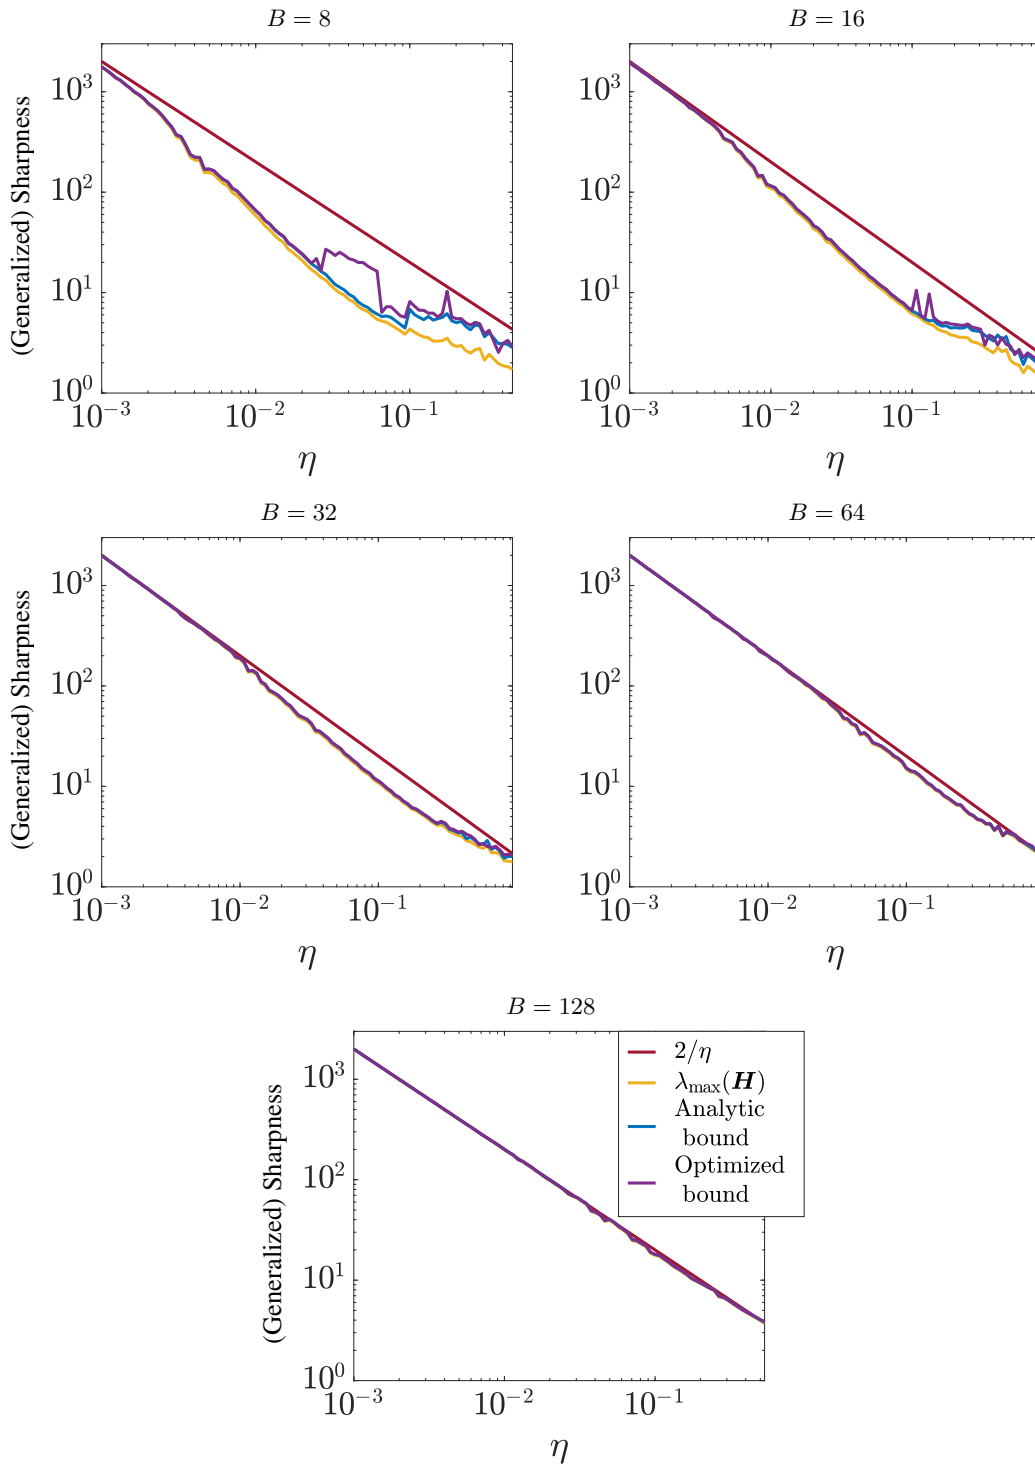


Figure 3: **Sharpness vs. learning rate.** Additional results for the experiment in Sec. 4. These five figures complete the results of Fig. 1. Here we see that SGD with big batch sizes behaves like GD.