

A EXPERIMENTAL SETUPS

For Fig. 1, Fig. 2, Fig. 7(b), we use the FNN with size 49-40-10 for the MNIST classification task, where the input data is resized to 7×7 to reduce the amount of calculation of the root of the parameter covariance matrix. For Fig. 1, Fig. 2, we train the network using GD with the first 1000 images as the training set. We add a dropout layer behind the hidden layer. The dropout rate and learning rate are specified and unchanged in each experiment. For Fig. 7(b), we train the network using SGD with the whole training set, the batch size is 32. We add a dropout layer behind the hidden layer. The learning rate is 0.1 and $p = 0.8$. For the accuracy of SDE simulation, all parameters are used to calculate the Hessian matrix and covariance matrix.

For Fig. 3(a), Fig. 4, Fig. 6, Fig. 8(a), we use the FNN with size 784-50-50-10 for the MNIST classification task. We add a dropout layer behind the second layer. The dropout rate and learning rate are specified and unchanged in each experiment. We only consider the parameter matrix corresponding to the weight and the bias of the fully-connected layer between two hidden layers. Therefore, for experiments in Fig. 3(a), $D = 2500$. For Fig. 3(a), Fig. 4, Fig. 6, we train the network using GD with the first 10000 images as the training set. For Fig. 8(a), we train the network using SGD with the whole training set with batch size 32.

For Fig. 3(b), Fig. 8(b), we use the VGG16 for the CIFAR10 classification task. We train the network using SGD with batch size 128 for Fig. 3(b), and 32 for Fig. 8(b). We add the dropout layer after the first two convolutional layers of each block and the first fully-connected layer. The dropout rate and learning rate are specified and unchanged in each experiment. We only consider the parameter matrix corresponding to the weight and the bias of the first convolutional layer. Therefore, $D = 1728$ for the two experiments. Due to the computational cost of the Hessian matrix and the covariance matrix, we calculate cosine similarity every five epochs.

For Fig. 5(a, c, e, g), we add dropout layers after the convolutional layers, and for each dropout layer, $p = 0.8$. We only consider the parameter matrix corresponding to the weight of the first convolutional layer of the first block of the ResNet-20. Models are trained using full-batch GD on the CIFAR100 classification task for 1200 epochs. The learning rate is initialized at 0.01. Since the Hessian calculation of ResNet takes much time, we only perform it at a specific dropout rate and learning rate.

For Fig. 5(b, d, f, h), we use transformer Vaswani et al. (2017) with $d_{\text{model}} = 50$, $d_k = d_v = 20$, $d_{\text{ff}} = 256$, $h = 4$, $N = 3$, the meaning of the parameters is consistent with the original paper. We only consider the parameter matrix corresponding to the weight of the fully-connected layer whose output is queried in the Multi-Head Attention layer of the first block of the decoder. We apply dropout to the output of each sub-layer before it is added to the sub-layer input and normalized. In addition, we apply dropout to the sums of the embeddings and the positional encodings in both the encoder and decoder stacks. For each dropout layer, $p = 0.9$. For the English-German translation problem, we use the cross-entropy loss with label smoothing trained by full-batch Adam based on the Multi30k dataset. The learning rate strategy is the same as that in Vaswani et al. (2017). The warm-up step is 4000 epochs, the training step is 10000 epochs. We only use the first 2048 examples for training to compromise with the computational burden.

For Fig. 7(a), we use the FNN with size 49-40-40-40-10 for the MNIST classification task, where the input data is resized to 7×7 to reduce the amount of calculation of the root of the parameter covariance matrix. We train the network using GD with the first 1000 images as the training set. We add a dropout layer behind each hidden layer. The learning rate is 0.1 and $p = 0.5$. For the accuracy of SDE simulation, all parameters are used to calculate the Hessian matrix and covariance matrix. The drift term is simulated by sampling the random variable 2000 times.

For Fig. 9(a), we use the ReLU FNN with the width of 1000 to fit the target function as follows,

$$f(x) = \frac{1}{2}\sigma(-x - \frac{1}{3}) + \frac{1}{2}\sigma(x - \frac{1}{3}),$$

where $\sigma(x) = \text{ReLU}(x)$. The learning rate is 1×10^{-3} . For SGD, the batch size is 1. For dropout, the dropout layer is added after the hidden layer with $p = 0.8$. For parameter noise injection, we use the layer noise with the noise standard deviation $\sigma = 0.001$. We initialize the parameters in the linear regime, $\theta \sim N(0, \frac{1}{m^{0.2}})$, where $m = 1000$ is the width of the hidden layer.

For Fig. 9(b), Fig. 9(c), we use the ReLU FNN with the size of 784-500-500-500-10 to classify the MNIST dataset, and the learning rate is 0.01. The batch size is 128. For dropout, the dropout layer is added after the hidden layers with $p = 0.8$. For parameter noise injection, we use the layer noise with the noise standard deviation $\sigma = 0.03$.

B EXTENDED EXPERIMENTS ON NOISE STRUCTURES.

B.1 RANDOM DATA COLLECTION METHODS

We first introduce two types of dynamical datasets collected during dropout training to study the noise structure of dropout. These datasets are different from the training sample \mathcal{S} .

Random trajectory data. The training process of NNs usually consists of two phases: the fast convergence phase and the exploration phase (Shwartz-Ziv and Tishby, 2017). In the exploration phase, the network is often considered to be near a minimum, and the movement of parameters is largely affected by the noise structure. Based on the previous work (Feng and Tu, 2021), we collect parameter sets $\mathcal{D}_{\text{para}} := \{\theta_i\}_{i=1}^N$ from N consecutive training steps in the exploration phase, where θ_i is the network parameter set at i -th sample step. This sampling method requires a large number of training steps, so model parameters often have large fluctuations during the sampling process. To improve the sampling accuracy, we propose another type of random data to characterize the noise structure of dropout as follows.

Random gradient data. We train the network until the loss is near zero and then we freeze the training process, then we sample N realizations of the dropout variable to get the random gradient dataset, i.e., $\mathcal{D}_{\text{grad}} := \{g_i\}_{i=1}^N$. The i -th sample point g_i is obtained as follows: i) Firstly, we generate a realization of the dropout variable δ_i under a given dropout rate; ii) Then, we compute the gradient of the loss function with respect to the parameters, denoted by $g_i(\cdot) := \nabla R_{\mathcal{S}}^{\text{drop}}(\cdot; \delta_i)$. Each element in $\mathcal{D}_{\text{grad}}$ represents an evolution direction of network parameters, determined by the dropout variable. Therefore, studying the structure of $\mathcal{D}_{\text{grad}}$ can help us understand how the dropout noise exerts an impact throughout the training process.

B.2 INVERSE VARIANCE-FLATNESS RELATION

The alignment relation studied above also implies the inverse variance-flatness relation, i.e., the noise variance is large along the sharp direction of the loss landscape, and small along the flat direction. In this subsection, we verify this relation by two sets of experiments. Firstly, we present two different approaches to characterize the flatness of loss landscape and the covariance of noise from the random trajectory data $\mathcal{D}_{\text{para}}$ and random gradient data $\mathcal{D}_{\text{grad}}$, then we numerically demonstrate the inverse variance-flatness relation. For convenience, \mathcal{D} refers to either the dataset $\mathcal{D}_{\text{para}}$ or the dataset $\mathcal{D}_{\text{grad}}$ depending on its context, so is the case for their corresponding covariance Σ and Hessian H . We then proceed to the definitions of **noise variance** and **interval flatness**.

Definition 2 (noise variance). For dataset \mathcal{D} and its covariance Σ , we denote $\lambda_i(\Sigma)$ as the i th eigenvalue of Σ and its corresponding eigen direction as $v_i(\Sigma)$. Then we term $\lambda_i(\Sigma)$ the noise variance of \mathcal{D} at the eigen direction $v_i(\Sigma)$.

The interval flatness below characterizes the flatness of the landscape around a local minimum.

Definition 3 (interval flatness³). For a local minimum θ_0^* , the loss function profile R_v along direction v reads:

$$R_v(\gamma) \equiv R_{\mathcal{S}}(\theta_0^* + \gamma v),$$

where γ represents the distance moved in the v direction. The interval flatness F_v is then defined as the width of the region within which $R_v(\gamma) \leq 2R_v(0)$. We determine F_v by finding two closest points $\theta_v^l < 0$ and $\theta_v^r > 0$ on each side of the minimum that satisfy $R_v(\theta_v^l) = R_v(\theta_v^r) = 2R_v(0)$. The interval flatness is defined as:

$$F_v \equiv \theta_v^r - \theta_v^l. \quad (19)$$

Remark. The experiments show that the result is not sensitive to the selection of the pre-factor 2. A larger value of F_v means a flatter landscape in the direction v .

We use PCA to study the weight variations when the training accuracy is nearly 100%. The networks are trained with full-batch GD for different learning rates and dropout rates under the same random seed. When the loss is small enough, we sample the parameters or gradients of parameters N times ($N = 3000$ for this experiment) and study the relationship between $\{\lambda_i(\Sigma)\}_{i=1}^N$ and $\{F_{v_i(\Sigma)}\}_{i=1}^N$ for both weight dataset $\mathcal{D}_{\text{para}}$ and gradient dataset $\mathcal{D}_{\text{grad}}$.

³This definition is also used in Feng and Tu (2021)




Figure 4: (a, b) The inverse relation between the variance $\{\lambda_i(\Sigma)\}_{i=1}^N$ and the interval flatness $\{F_{v_i}(\Sigma)\}_{i=1}^N$ for different choices of p and learning rate lr with different network structures. The PCA is done for different datasets \mathcal{D} sampled from parameters for the top line and sampled from gradients of parameters for the bottom line. The dashed lines give the approximate slope of the scatter. (c, d) The relation between the variance $\{\text{Var}(\text{Proj}_{v_i}(\mathcal{D}))\}_{i=1}^N$ and the eigenvalue $\{\lambda_i(\mathbf{H})\}_{i=1}^N$ for different choices of p and learning rate lr with different network structures. The projection is done for different datasets \mathcal{D} sampled from parameters for the top line and sampled from gradients of parameters for the bottom line. The dashed lines give the approximate slope of the scatter.

For different learning rates and dropout rates, Fig. 4(a, b) reveal an inverse relationship between the interval flatness of the loss landscape denoted as $\{F_{v_i}(\Sigma)\}_{i=1}^N$, and the noise variance represented by the PCA spectrum $\{\lambda_i(\Sigma)\}_{i=1}^N$. Notably, a power-law relationship can be established between $\{F_{v_i}(\Sigma)\}_{i=1}^N$ and $\{\lambda_i(\Sigma)\}_{i=1}^N$. Specifically, in the low flatness region, the dropout-induced noise exhibits a large variance. As the loss landscape transitions into the high flatness regime, the linear relationship between variance and flatness becomes more evident. Overall, These findings consistently demonstrate the inverse relation between variance and flatness, as exemplified in Fig. 4(a, b). Subsequently, we delve into the definitions of **Projected variance** and **Hessian flatness**.

Definition 4 (projected variance). For a given direction $\mathbf{v} \in \mathbb{R}^D$ and dataset $\mathcal{D} = \{\theta_i\}_{i=1}^N$, where $\theta_i \in \mathbb{R}^D$, the inner product of \mathbf{v} and θ_i is denoted by $\text{Proj}_{\mathbf{v}}(\theta_i) := \langle \theta_i, \mathbf{v} \rangle$, then we can define the projected variance for \mathcal{D} at the direction \mathbf{v} as follows,

$$\text{Var}(\text{Proj}_{\mathbf{v}}(\mathcal{D})) = \frac{\sum_{i=1}^N (\text{Proj}_{\mathbf{v}}(\theta_i) - \mu)^2}{N},$$

where μ is the mean value of $\{\text{Proj}_{\mathbf{v}}(\theta_i)\}_{i=1}^N$.

Definition 5 (Hessian flatness). For Hessian \mathbf{H} , as we denote $\lambda_i(\mathbf{H})$ by the i -th eigenvalue of \mathbf{H} corresponding to the eigenvector $\mathbf{v}_i(\mathbf{H})$, we term $\lambda_i(\mathbf{H})$ the Hessian flatness along direction $\mathbf{v}_i(\mathbf{H})$.

The eigenvalues of the Hessian evaluated at a local minimum often serve as indicators of the flatness of the loss landscape, and larger eigenvalues correspond to sharper directions. In our investigation, we analyze the interplay between the eigenvalues of Hessian \mathbf{H} at the final stage of the training process and the projected variance of dropout at each of the corresponding eigen directions, i.e., $\lambda_i(\mathbf{H})$ v.s. $\{\text{Var}(\text{Proj}_{\mathbf{v}_i(\mathbf{H})}(\mathcal{D}))\}_{i=1}^N$. Specifically, we sample the parameters or gradients of parameters N times ($N = 1000$ for this experiment), and examine the relationship between $\{\lambda_i(\mathbf{H})\}_{i=1}^N$ and $\{\text{Var}(\text{Proj}_{\mathbf{v}_i(\mathbf{H})}(\mathcal{D}))\}_{i=1}^N$ for both the parameter dataset $\mathcal{D}_{\text{para}}$ and the gradient dataset $\mathcal{D}_{\text{grad}}$.

Under various dropout rates and learning rates, Fig. 4(c, d) presents establishes a consistent power-law relationship between $\{\lambda_i(\mathbf{H})\}_{i=1}^N$ and $\{\text{Var}(\text{Proj}_{\mathbf{v}_i(\mathbf{H})}(\mathcal{D}))\}_{i=1}^N$, and this relationship remains robust irrespective of the choice between parameter dataset $\mathcal{D}_{\text{para}}$ or the gradient dataset $\mathcal{D}_{\text{grad}}$. The positive correlation observed between the Hessian flatness and the projection variance provides insights into the structural characteristics of the dropout-induced noise. Specifically, these characteristics have the potential to facilitate the escape from sharp minima and enhance the generalization capabilities of NNs. Additionally, Fig. 4 highlights the distinct linear structure exhibited by gradient sampling in comparison to parameter sampling, which corroborates the discussions outlined in Section B.1.

Furthermore, we verify the inverse relation between the covariance matrix and the Hessian matrix of dropout through different data collection methods and projection methods on larger network




Figure 5: (a, b, c, d) The inverse relation between the variance $\{\lambda_i(\Sigma)\}_{i=1}^N$ and the interval flatness $\{F_{v_i(\Sigma)}\}_{i=1}^N$ for different choices of p and learning rate lr with different network structures. The PCA is done for different datasets D sampled from parameters for the top line and sampled from gradients of parameters for the bottom line. The dashed lines give the approximate slope of the scatter. (e, f, g, h) The relation between the variance $\{\text{Var}(\text{Proj}_{v_i(H)}(D))\}_{i=1}^N$ and the eigenvalue $\{\lambda_i(H)\}_{i=1}^N$ for different choices of p and learning rate lr with different network structures. The projection is done for different datasets D sampled from parameters for the top line and sampled from the scatter.

structures, such as ResNet-20 and transformer, and more complex datasets, such as CIFAR-100 and Multi30k, as shown in Fig. 5.

B.3 THE SIMILARITY OF EIGENVECTORS BETWEEN THE HESSIAN MATRIX AND THE COVARIANCE MATRIX




Figure 6: (a) The cosine similarity between the first 15 eigenvectors of the Hessian matrix and the covariance matrix. The color represents the value of cosine similarity, and the lines on the top and left represent the energy ratio of the first k eigenvalues of the covariance matrix and the Hessian matrix, respectively. (b) The projection ratio of the first 9 eigenvectors of the covariance matrix in different eigendirections of the Hessian matrix.

We study the similarity of eigenvectors between the Hessian matrix and the covariance matrix. The model parameters are derived from the final model represented by the blue line in Fig. 3(a). In Fig. 6(a), we study the cosine similarity between the first 15 eigenvectors of the Hessian matrix and

the covariance matrix. The color represents the value of cosine similarity, and the lines on the top and left represent the energy ratio $E(k)$ of the first k eigenvalues of the covariance matrix and the Hessian matrix respectively, that is, $E(k) = \sum_{i=1}^k (\lambda_i)^2 / \sum_{i=1}^N (\lambda_i)^2$, where λ_i is the i -th eigenvalue, and N is the number of the eigenvalues for the two matrices. It is easy to find that the energy ratio of the first 9 eigenvectors of the two matrices far exceeds the energy ratio of other eigenvectors. Meanwhile, the first 9 eigenvectors of the two matrices are highly similar, which implies the alignment property between the Hessian matrix and the covariance matrices in the eigenspace.

In order to further characterize the similarity between the first few eigenvectors of the two matrices, we study the projection of the first 9 eigenvectors of the covariance matrix in different eigendirections of the Hessian matrix. As shown in Fig. 6(b), we calculate the value of $\sum_{i=1}^k \alpha(\mathbf{v}_i(H), \mathbf{v}_j(\Sigma))^2$ for the index k of the Hessian eigenvector and the index j of the covariance eigenvector, where $\alpha(\mathbf{v}_i, \mathbf{v}_j)$ is the cosine similarity between the two vectors $\mathbf{v}_i, \mathbf{v}_j$. The energy proportions of the first nine eigenvectors of the covariance matrix in the first nine directions of the Hessian matrix are all greater than 0.95, which means that the similarity of the eigenvectors corresponding to the first few eigenvalues of the Hessian matrix and the covariance matrix is much higher than the similarity between them and other feature directions, which further confirms the alignment properties between the Hessian matrix and the covariance matrix in the eigenspace.

C EXTENDED ANALYSIS OF DROPOUT UNDER COMPLEX NETWORK STRUCTURES AND SGD SETTING

In the main text, the theoretical part and experimental verification mainly focus on the settings of the two-layer network and GD. This is mainly due to the derivation of the specific forms of the drift term and covariance matrix. In fact, the SME framework in this work can be applied to more complex network structures and SGD settings without being limited to the explicit calculation of drift terms and covariance matrices. In the following, we conduct a detailed analysis of the feasibility of SME under complex network structures and SGD settings, and verify it through SME numerical simulation and the alignment between the Hessian matrix and covariance matrix under broader settings.

C.1 FEASIBILITY ANALYSIS FOR NETWORK STRUCTURES AND SGD SETTING

The formulation of our stochastic modified equations relies on two critical components. Firstly, the algorithm updates parameters iteratively in a recursive manner, i.e., $\theta_N = F(\theta_{N-1}, \delta_N)$, thereby forming iterations as a time-homogeneous Markov chain. This characteristic is foundational for applying the stochastic modified equation approach to weakly approximate the dynamics of dropout. The second crucial component involves the utilization of Taylor’s theorem with the Lagrange form of the remainder. In light of these two foundational components, we assert that the derivation and analysis of stochastic modified equations extend beyond the confines of two-layer neural networks and dropout. This approach is applicable to a broader spectrum of stochastic algorithms, including but not limited to SGD and ADAM, and extends naturally to deeper or more complex neural network architectures.

One might wonder why we only demonstrate our results limited to two-layer neural networks. For one thing, two-layer neural networks are indeed a meaningful step towards a thorough understanding of dropout. For another, the specific structure of two-layer neural networks, wherein parameters are ‘decoupled’, i.e., $\theta = \mathbf{q}_r = (a_r, \mathbf{w}_r)$, facilitates the demonstration of dropout’s effect through the computation of the network output. Finally, for two-layer neural networks, we are able to give out an explicit expression for the modified loss, since we are able to calculate the first and second moment of the random variable $(\delta)_r$. In the case of multi-layer neural networks, if dropout is applied only to the outermost layer, we can still calculate the explicit expression for the modified loss, since its calculation involves solely the first and second moments of the random variable $(\delta)_r$. However, the situation becomes significantly more challenging when dropout is applied to the inner layers of the multi-layer neural network. In this scenario, obtaining a closed-form expression for the expectation $\mathbb{E}[h(\delta)]$, where h is a highly nonlinear function with respect to δ , becomes nearly impossible. Moreover, the computation of the covariance is even more infeasible in multi-layer neural networks compared to two-layer neural networks.

While the structural dynamics of dropout within nonlinear activations for deep neural networks remain uncertain, we assert that the same structure is retained if we apply dropout in a linear manner to the deep neural networks by computing $\mathbf{f}_\theta(\mathbf{x}) = \sum_{r=1}^m a_r(\delta)_r \sigma(\mathbf{w}_r^T \mathbf{x}^{[L]})$, where $\mathbf{x}^{[L]}$ is the output function of a L -layer neural network for $L \geq 2$, the same structure remains, and the modified loss $L_S(\cdot)$ reads:

$$L_S(\theta) := \frac{1}{2n} \sum_{i=1}^n e_i^2 + \frac{1-p}{2np} \sum_{i=1}^n \sum_{r=1}^m a_r^2 \sigma(\mathbf{w}_r^T \mathbf{x}_i^{[L]})^2.$$

As for the SGD setting, given the nature of SGD as an unbiased estimator with respect to the full sample, we conjecture that an order-1 approximation utilizing SME for SGD combined with dropout shall be in the form:

$$d\theta_t = -\nabla_\theta L_S(\theta_t) dt + \tilde{\Sigma}(\theta_t) d\mathbf{W}_t,$$

wherein the drift term remains invariant regardless of GD or SGD, while the diffusion term $\tilde{\Sigma}(\theta_t)$ combines noise from both dropout and SGD.

As for the scenario where dropout is applied within the nonlinear activations, this exact form of the SME remains unexplored. However, we design numerical experiments to verify the approximation ability of deeper network SMEs in the next subsection, where the drift term here is calculated numerically through multiple samplings on dropout variable δ .

C.2 EXTENDED EXPERIMENTS FOR NETWORK STRUCTURES AND SGD SETTING

In this section, we verify the feasibility of the above analysis through numerical experiments. We all use first-order SME for approximation. Please refer to the Appendix A for specific experimental settings.

C.2.1 SME SIMULATION

In this subsection, we conduct an empirical validation of the feasibility analysis of SMEs for dropout. This validation is conducted through an exploration of the resemblance between the numerical simulation of the SME and the real-time training process of dropout. For the numerical simulation of the SME, unless otherwise specified, we employ the Euler-Maruyama method to approximate its dynamic evolution by the order-1 approximation. To mitigate the computational demands associated with the covariance matrix, we resize the MNIST data to 7×7 , thereby reducing the number of network parameters involved in the simulations.




Figure 7: We train fully connected networks on MNIST. The curves and points are derived from the average results of 20 individual trials, each of which utilized the same initialization distribution. The training loss trajectory obtained by SME simulation or dropout training under four cases of different learning rates and dropout rates. The error bands, portrayed with greater transparency, are derived from the maximum and minimum loss values observed across these 20 random trials at each training step. (a) Four-layer network, each layer width is 40. The learning rate is 0.1 and $p = 0.5$. (b) Two-layer network training with SGD, the batch size is 32. The learning rate is 0.1 and $p = 0.8$.

As shown in Fig. 7, we compare the SME simulation results under deep network and SGD settings with dropout. For deep networks, we add a dropout layer after each hidden layer. Since the drift term cannot be calculated explicitly, we approximate the expectation by sampling the noise multiple times. This also limits the number of parameters in our model. For dropout under the SGD setting, we use the same drift term as under the GD setting, and calculate the noise structure from the combined noise of the two random algorithms, dropout and SGD.

C.2.2 HESSIAN-VARIANCE ALIGNMENT

In this section, we mainly study the impact of dropout of complex network structures on Hessian-variance alignment under the SGD setting. In the main text, we have conducted relevant experiments on complex network structures under the GD setting. To investigate the Hessian-Variance alignment relation for dropout with SGD, we study the cosine similarity quantity $\alpha(\theta_t)$ between the covariance matrix $\Sigma_t := \Sigma(\theta_t)$ and the Hessian matrix $\mathbf{H}_t := \mathbf{H}(\theta_t)$ at each time step t under the SGD setting. Σ_t is the covariance matrix of $\mathcal{D}_{\text{grad}}$, a collection of gradients calculated with different dropout variables δ and input data x sampled at the t th step. On the other hand, \mathbf{H}_t is the Hessian of the loss function evaluated at the t th iteration. Then the crucial cosine similarity metric $\alpha(\theta_t)$ is formally expressed as:

$$\alpha(\theta_t) = \frac{\text{Tr}(\mathbf{H}_t \Sigma_t)}{\|\mathbf{H}_t\|_F \|\Sigma_t\|_F} \quad (20)$$

As depicted in Fig. 8, it is evident that throughout the training process, $\alpha(\theta_t)$ consistently attains values surpassing 0.85 in Fig. 8(a) and 0.65 in Fig. 8(b), and these observations hold true across

varying learning rates and dropout rates. It’s worth noting that, based on 100 samples, the average cosine similarity between the two random matrices based on the selected parameters is only 8.6×10^{-4} . The eigenvalues of the two random matrices are derived from the eigenvalues of the Hessian matrix and covariance matrix of the selected parameters, and the corresponding eigenvectors are sampled from a normal distribution and normalized. The model parameters are derived from the final model represented by the blue line in Fig. 8(a). Consequently, the introduced noise is highly anisotropic in that it aligns well with the Hessian matrix across all directions. We acknowledged that due to computational constraints, this experiment limits the trace calculation to a subset of parameters, which can be effectively regarded as the projection of the Hessian and the noise into specific directions.




Figure 8: The cosine similarity $\alpha(\theta_t)$ between the Hessian of the loss function and the covariance of the dropout noise at each training epoch t for different choices of dropout rate and learning rate. (a) The FNN with size 784-50-50-10 is trained on the MNIST dataset with a batch size of 32. The dropout layer is added after the first hidden layer. (b) The VGG16 is trained on the CFIAR-10 dataset with a batch size of 32. The dropout layers are added after the first two convolutional layers of each block and the first fully-connected layer. The calculation of $\alpha(\theta_t)$ is performed every five epochs.

D EFFECT OF NOISE FROM DIFFERENT TRAINING STRATEGIES ON RESULTS.

In this subsection, we compare the effect of noise from different training strategies on model output and generalization. We mainly compared three training strategies: SGD, dropout, and parameter noise injection. Parameter noise injection is well known as the explicit regularizer on the trace of the hessian (Orvieto et al., 2023). We first study the results of ReLU NNs training under three training strategies, as shown in Fig. 9(a). It is easy to see that the model output of GD+dropout has better smoothness and is more in line with expectations. In Fig. 9(b), we compare the generalization capabilities of three randomized algorithms. We follow the settings in Orvieto et al. (2023) and conduct experiments using MLP1 in the corresponding code, with the network size 784-500-500-500-10. We use a batch size of 128 for the three algorithms. The mean (solid line) and standard deviation (shaded range) were calculated from five independent experiments for each setting. For dropout, we add the dropout layer under each hidden layer, and we set $p = 0.8$. For parameter noise injection, we set the noise standard deviation $\sigma = 0.03$. At the same time, we compared the traces of the Hessian matrix of the three training strategies, which is a metric widely used to characterize flatness. As shown in Fig. 9(c), the trace of the Hessian matrix of dropout is smaller than that of the other two training strategies during the entire training process, suggesting a stronger tendency for the flat solution.




Figure 9: (a) Two-layer ReLU NN output under different training strategies. The width of the hidden layer is 1000, and the learning rate is 1×10^{-3} . For SGD, the batch size is 1. For dropout, the dropout layer is added after the hidden layer with $p = 0.8$. For parameter noise injection, we use the layer noise with the noise standard deviation $\sigma = 0.001$. (b, c) The deep FNN trained on MNIST under different training strategies. The network size is 784-500-500-500-10, and the learning rate is 0.01. The batch size is 128. For dropout, the dropout layer is added after the hidden layers with $p = 0.8$. For parameter noise injection, we use the layer noise with the noise standard deviation $\sigma = 0.03$. (b) Test accuracy. (c) Trace of the Hessian matrix.

E EXTENDED EXPERIMENTS OF DROPOUT UNDER MEAN-FIELD LIMIT

In the above, we do not treat the network width as the variable under study, but treat it as a constant value of order one. However, if we take the mean-field limit of the two-layer network, i.e.,

$$f_{\theta}(\mathbf{x}) := \frac{1}{m} \sum_{r=1}^m a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}),$$

where $m \rightarrow \infty$, then the modified equation has the following form:

$$L_S(\theta) := \frac{1}{2n} \sum_{i=1}^n e_i^2 + \frac{1-p}{2n p m^2} \sum_{i=1}^n \sum_{r=1}^m a_r^2 \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i)^2.$$

The additional regularity of the dropout modified equation disappears in the infinite width limit, which is contrary to our common sense. Intuitively speaking, if we take each neuron to be independent and identically distributed, even if we ‘drop’ some neurons, dropout does not seem to exert an impact on the training process. We carefully verified this by experimentally studying different wide networks with and without dropout. As shown in Fig. 10, for narrow networks, dropout will have a certain impact on the training process, and this impact will become insignificant when the network is wide enough.

We design experiments with and without dropout at different hidden layer widths under the mean-field setting. The quantity R_1 we study is defined as $\frac{1-p}{2n m^2 p} \sum_{i=1}^n \sum_{r=1}^m a_r^2 \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i)^2$. As shown in Fig. 10, for the wide network, the model loss trajectory is almost consistent with or without dropout, and the cosine similarity between the two model parameter vectors i.e., flatten all the parameters of the model into one vector, is 0.9999998, but there is no such similarity for the narrow network.




Figure 10: The two-layer network with and without dropout at different hidden layer widths m under the mean-field setting. The learning rate is 0.01 for all experiments, and $p = 0.9$ for experiments with dropout. These settings can make the dropout noise have little effect on the results.

F PRELIMINARIES

F.1 NOTATIONS

We adhere wherever possible to the following notation. Dimensional indices are written as subscripts with a bracket to avoid confusion with other sequential indices (e.g. time, iteration number), which do not have brackets. When more than one indices are present, we separate them with a comma, e.g. $\mathbf{x}_{k,(i)}$ is the i -th coordinate of the vector \mathbf{x}_k , the k^{th} member of a sequence.

We set a special vector $(1, 1, 1, \dots, 1)^\top$ by $\mathbf{1} := (1, 1, 1, \dots, 1)^\top$ whose dimension varies. We set n for the number of input samples, m for the width of the neural network, and $D := m(d + 1)$ hereafter in this paper. We let $[n] = \{1, 2, \dots, n\}$. We set $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ as the normal distribution with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$. We denote \otimes as the Kronecker tensor product, $\langle \cdot, \cdot \rangle$ for standard inner product between two vectors, and $\mathbf{A} : \mathbf{B}$ for the Frobenius inner product between two matrices \mathbf{A} and \mathbf{B} . We denote vector L^2 norm as $\|\cdot\|_2$, vector or function L_∞ norm as $\|\cdot\|_\infty$, function L_1 norm as $\|\cdot\|_1$, matrix infinity norm as $\|\cdot\|_{\infty \rightarrow \infty}$, matrix spectral (operator) norm as $\|\cdot\|_{2 \rightarrow 2}$, and matrix Frobenius norm as $\|\cdot\|_F$. Finally, we denote the set of continuous functions $f(\cdot) : \mathbb{R}^D \rightarrow \mathbb{R}$ possessing continuous derivatives of order up to and including r by $\mathcal{C}^r(\mathbb{R}^D)$, and for a Polish space \mathcal{X} , we denote the space of bounded measurable functions by $\mathcal{B}_b(\mathcal{X})$, and the space of bounded continuous functions by $\mathcal{C}_b(\mathcal{X})$. In the mathematical discipline of general topology, a Polish space is a separable complete metric space.

F.2 PROBLEM SETUP

For the empirical risk minimization problem given by the quadratic loss:

$$\min_{\boldsymbol{\theta}} R_{\mathcal{S}}(\boldsymbol{\theta}) = \frac{1}{2n} \sum_{i=1}^n (f_{\boldsymbol{\theta}}(\mathbf{x}_i) - y_i)^2, \quad (21)$$

where $\mathcal{S} := \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ is the training sample, $f_{\boldsymbol{\theta}}(\mathbf{x})$ is the prediction function, $\boldsymbol{\theta}$ are the parameters to be optimized over, and their dependence is modeled by a two-layer neural network (NN) with m hidden neurons

$$f_{\boldsymbol{\theta}}(\mathbf{x}) := \sum_{r=1}^m a_r \sigma(\mathbf{w}_r^\top \mathbf{x}), \quad (22)$$

where $\mathbf{x} \in \mathbb{R}^d$, $\boldsymbol{\theta} = \text{vec}(\boldsymbol{\theta}_a, \boldsymbol{\theta}_w)$ with $\boldsymbol{\theta}_a = \text{vec}(\{a_r\}_{r=1}^m)$, $\boldsymbol{\theta}_w = \text{vec}(\{\mathbf{w}_r\}_{r=1}^m)$ is the set of parameters, $\sigma(\cdot)$ is the activation function applied coordinate-wisely to its input, and σ is 1-Lipschitz with $\sigma \in \mathcal{C}^\infty(\mathbb{R})$. More precisely, $\boldsymbol{\theta} = \text{vec}(\{\mathbf{q}_r\}_{r=1}^m)$ whereas for each $r \in [m]$, $\mathbf{q}_r := (a_r, \mathbf{w}_r^\top)^\top$. We remark that the bias term b_r can be incorporated by expanding \mathbf{x} and \mathbf{w}_r to $(\mathbf{x}^\top, 1)^\top$ and $(\mathbf{w}_r^\top, b_r)^\top$.

Given fixed learning rate $\eta > 0$, then at the N -th iteration, where

$$t_N := N\eta,$$

and a scaling vector $\boldsymbol{\delta}_N \in \mathbb{R}^m$ is sampled with independent random coordinates: For each $k \in [m]$,

$$(\boldsymbol{\delta}_N)_k = \begin{cases} \frac{1}{p} & \text{with probability } p, \\ 0 & \text{with probability } 1 - p, \end{cases} \quad (23)$$

and we observe that $\{\boldsymbol{\delta}_N\}_{N \geq 1}$ is an i.i.d. Bernulli sequence with $\mathbb{E}\boldsymbol{\delta}_1 = \mathbf{1}$, and naturally, with slight abuse of notations, the σ -fields $\mathcal{F}_N := \{\sigma(\boldsymbol{\delta}_1, \boldsymbol{\delta}_2, \dots, \boldsymbol{\delta}_N)\}$ forms a filtration.

We then apply dropout to two-layer NNs by computing

$$\mathbf{f}_{\boldsymbol{\theta}}(\mathbf{x}; \boldsymbol{\delta}) := \sum_{r=1}^m (\boldsymbol{\delta})_r a_r \sigma(\mathbf{w}_r^\top \mathbf{x}), \quad (24)$$

and we denote the empirical risk associated with dropout by

$$\begin{aligned} R_S^{\text{drop}}(\boldsymbol{\theta}; \boldsymbol{\delta}) &:= \frac{1}{2n} \sum_{i=1}^n (\mathbf{f}_{\boldsymbol{\theta}}(\mathbf{x}_i; \boldsymbol{\delta}) - y_i)^2 \\ &= \frac{1}{2n} \sum_{i=1}^n \left(\sum_{r=1}^m (\boldsymbol{\delta})_{r a_r} \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i) - y_i \right)^2. \end{aligned} \quad (25)$$

We observe that the parameters at the N -th step are updated via back propagation as follows:

$$\boldsymbol{\theta}_N = \boldsymbol{\theta}_{N-1} - \eta \nabla_{\boldsymbol{\theta}} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N), \quad (26)$$

where $\boldsymbol{\theta}_0 := \boldsymbol{\theta}(0)$. Finally, we denote hereafter that for all $i \in [n]$,

$$e_i^N := e_i(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N) := \mathbf{f}_{\boldsymbol{\theta}_{N-1}}(\mathbf{x}_i; \boldsymbol{\delta}_N) - y_i,$$

hence the empirical risk associated with dropout $R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N)$ can be written into

$$R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N) = \frac{1}{2n} \sum_{i=1}^n (e_i^N)^2,$$

thus the dropout iteration (26) reads

$$\boldsymbol{\theta}_N - \boldsymbol{\theta}_{N-1} = -\eta \nabla_{\boldsymbol{\theta}} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N) = -\frac{\eta}{n} \sum_{i=1}^n e_i^N \nabla_{\boldsymbol{\theta}} e_i^N,$$

and we may proceed to the introduction of the SME approximation.

G STOCHASTIC MODIFIED EQUATIONS FOR DROPOUT

G.1 MODIFIED LOSS

Recall that the parameters at the N -th step are updated as follows:

$$\boldsymbol{\theta}_N = \boldsymbol{\theta}_{N-1} - \frac{\eta}{n} \sum_{i=1}^n e_i^N \nabla_{\boldsymbol{\theta}} e_i^N, \quad (27)$$

and since $\{\boldsymbol{\delta}_N\}_{N \geq 1}$ is an i.i.d. sequence, then the dropout iteration (27) updates the parameters in a recursion form of

$$\boldsymbol{\theta}_N = \mathbf{F}(\boldsymbol{\theta}_{N-1}, \boldsymbol{\delta}_N), \quad (28)$$

where $\mathbf{F}(\cdot, \cdot) : \mathbb{R}^D \times \mathbb{R}^m \rightarrow \mathbb{R}^D$ is a smooth (\mathcal{C}^∞) function, and $\{\boldsymbol{\delta}_N\}_{N \geq 1}$ is a disturbance sequence on \mathbb{R}^m , whose marginal distribution possesses a density supported on an open subset of \mathbb{R}^m . Then, based on the results in Meyn and Tweedie (2012), the dropout iterations (27) forms a time-homogeneous Markov chain. Thus, we may misuse $\mathbb{E}[\cdot | \mathcal{F}_N]$, the conditional expectation given \mathcal{F}_N , with $\mathbb{E}_{\boldsymbol{\theta}_{N-1}}[\cdot]$, the conditional expectation given $\boldsymbol{\theta}_{N-1}$. Then, for each $k \in [m]$, the conditional expectation of the increment restricted to \mathbf{q}_k reads

$$\mathbb{E}_{\boldsymbol{\theta}_{N-1}} \left[\sum_{i=1}^n e_i^N \nabla_{\mathbf{q}_k} e_i^N \right] = \mathbb{E}_{\boldsymbol{\theta}_{N-1}} \left[\sum_{i=1}^n e_i^N (\boldsymbol{\delta}_N)_k \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)) \right],$$

and since

$$\begin{aligned} \mathbb{E}_{\boldsymbol{\theta}_{N-1}} [e_i^N (\boldsymbol{\delta}_N)_k] &= \mathbb{E}_{\boldsymbol{\theta}_{N-1}} \left[\sum_{r=1, r \neq k}^m (\boldsymbol{\delta}_N)_r a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i) - y_i \right] \mathbb{E}_{\boldsymbol{\theta}_{N-1}} [(\boldsymbol{\delta}_N)_k] \\ &\quad + \mathbb{E}_{\boldsymbol{\theta}_{N-1}} [(\boldsymbol{\delta}_N)_k^2] a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \\ &= \left(\sum_{r=1, r \neq k}^m a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i) - y_i \right) + \frac{1}{p} a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \\ &= \left(\sum_{r=1}^m a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i) - y_i \right) + \left(\frac{1}{p} - 1 \right) a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i). \end{aligned}$$

For simplicity, given fixed $k \in [m]$, for any $i \in [n]$, we denote hereafter that

$$\begin{aligned} e_i &:= e_i(\boldsymbol{\theta}) := \sum_{r=1}^m a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i) - y_i, \\ e_{i, \setminus k} &:= e_{i, \setminus k}(\boldsymbol{\theta}) := \sum_{r=1, r \neq k}^m a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i) - y_i, \end{aligned}$$

we remark that compared with e_i^N , e_i and $e_{i, \setminus k}$ do not depend on the random variable $\boldsymbol{\delta}_N$. Then $\mathbb{E}_{\boldsymbol{\theta}_{N-1}} (e_i^N (\boldsymbol{\delta}_N)_k)$ can be written in short by

$$\begin{aligned} \mathbb{E}_{\boldsymbol{\theta}_{N-1}} [e_i^N (\boldsymbol{\delta}_N)_k] &= e_{i, \setminus k} + \frac{1}{p} a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \\ &= e_i + \left(\frac{1}{p} - 1 \right) a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i). \end{aligned} \quad (29)$$

Hence for each $k \in [m]$, expectation of the increment restricted to \mathbf{q}_k reads

$$\begin{aligned} &\mathbb{E}_{\boldsymbol{\theta}_{N-1}} \left[\sum_{i=1}^n e_i^N (\boldsymbol{\delta}_N)_k \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)) \right] \\ &= \sum_{i=1}^n e_i \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)) + \sum_{i=1}^n \left(\frac{1}{p} - 1 \right) a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)), \end{aligned}$$

then we define the *modified loss* $L_S(\cdot) : \mathbb{R}^{m(d+1)} \rightarrow \mathbb{R}$ for dropout:

$$L_S(\boldsymbol{\theta}) := \frac{1}{2n} \sum_{i=1}^n e_i^2 + \frac{1-p}{2np} \sum_{i=1}^n \sum_{r=1}^m a_r^2 \sigma(\mathbf{w}_r^\top \mathbf{x}_i)^2, \quad (30)$$

since as $\boldsymbol{\theta}_{N-1}$ is given, then by taking the conditional expectation, increment of the dropout iteration (27) reads

$$\boldsymbol{\theta}_N - \boldsymbol{\theta}_{N-1} = -\eta \mathbb{E}_{\boldsymbol{\theta}_{N-1}} \left[\nabla_{\boldsymbol{\theta}} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N) \right] = -\eta \nabla_{\boldsymbol{\theta}} L_S(\boldsymbol{\theta}) \Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_{N-1}},$$

which implies that in the sense of expectations, $\{\boldsymbol{\theta}_N\}_{N \geq 0}$ follows close to the gradient descent trajectory of $L_S(\boldsymbol{\theta})$ with fixed learning rate η .

G.2 STOCHASTIC MODIFIED EQUATIONS

We then follow the strategy of Li et al. (2017) to derive the stochastic modified equations (SME) for dropout. Firstly, from the results in Section G.1, we observe that given $\boldsymbol{\theta}_{N-1}$,

$$\boldsymbol{\theta}_N - \boldsymbol{\theta}_{N-1} = -\eta \nabla_{\boldsymbol{\theta}} L_S(\boldsymbol{\theta}) \Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_{N-1}} + \sqrt{\eta} \mathbf{V}(\boldsymbol{\theta}_{N-1}), \quad (31)$$

where $L_S(\cdot) : \mathbb{R}^{m(d+1)} \rightarrow \mathbb{R}$ is the modified loss defined in (30), and $\mathbf{V}(\cdot) : \mathbb{R}^{m(d+1)} \rightarrow \mathbb{R}^{m(d+1)}$ is a $m(d+1)$ -dimensional random vector, and when given $\boldsymbol{\theta}_{N-1}$, $\mathbf{V}(\boldsymbol{\theta}_{N-1})$ has mean $\mathbf{0}$ and covariance $\eta \boldsymbol{\Sigma}(\boldsymbol{\theta}_{N-1})$, where $\boldsymbol{\Sigma}(\cdot) : \mathbb{R}^{m(d+1)} \rightarrow \mathbb{R}^{m(d+1) \times m(d+1)}$ is the covariance of $\nabla_{\boldsymbol{\theta}} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N)$. Recall that $\boldsymbol{\theta} = \text{vec}(\{\mathbf{q}_r\}_{r=1}^m) = \text{vec}(\{(a_r, \mathbf{w}_r)\}_{r=1}^m)$, and for any $k, r \in [m]$, we denote that

$$\boldsymbol{\Sigma}_{kr}(\boldsymbol{\theta}_{N-1}) := \text{Cov} \left(\nabla_{\mathbf{q}_k} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N), \nabla_{\mathbf{q}_r} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N) \right),$$

then

$$\boldsymbol{\Sigma} = \begin{bmatrix} \boldsymbol{\Sigma}_{11} & \boldsymbol{\Sigma}_{12} & \cdots & \boldsymbol{\Sigma}_{1m} \\ \boldsymbol{\Sigma}_{21} & \boldsymbol{\Sigma}_{22} & \cdots & \boldsymbol{\Sigma}_{2m} \\ \vdots & \vdots & \vdots & \vdots \\ \boldsymbol{\Sigma}_{m1} & \boldsymbol{\Sigma}_{m2} & \cdots & \boldsymbol{\Sigma}_{mm} \end{bmatrix}.$$

For each $k \in [m]$, we obtain that

$$\begin{aligned} \boldsymbol{\Sigma}_{kk}(\boldsymbol{\theta}_{N-1}) &= \text{Cov} \left(\nabla_{\mathbf{q}_k} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N), \nabla_{\mathbf{q}_k} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N) \right) \\ &= \left(\frac{1}{p} - 1 \right) \left(\frac{1}{n} \sum_{i=1}^n \left(e_{i, \setminus k} + \frac{1}{p} a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \right) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)) \right) \\ &\quad \otimes \left(\frac{1}{n} \sum_{i=1}^n \left(e_{i, \setminus k} + \frac{1}{p} a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \right) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)) \right) \\ &\quad + \left(\frac{1}{p^2} - \frac{1}{p} \right) \sum_{l=1, l \neq k}^m \left(\frac{1}{n} \sum_{i=1}^n a_l \sigma(\mathbf{w}_l^\top \mathbf{x}_i) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)) \right) \\ &\quad \otimes \left(\frac{1}{n} \sum_{i=1}^n a_l \sigma(\mathbf{w}_l^\top \mathbf{x}_i) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)) \right), \end{aligned}$$

and for each $k, r \in [m]$ with $k \neq r$,

$$\begin{aligned} \Sigma_{kr}(\boldsymbol{\theta}_{N-1}) &= \text{Cov} \left(\nabla_{\mathbf{q}_k} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N), \nabla_{\mathbf{q}_r} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N) \right) \\ &= \left(\frac{1}{p} - 1 \right) \left(\frac{1}{n} \sum_{i=1}^n \left(e_{i, \setminus k, \setminus r} + \frac{1}{p} a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) + \frac{1}{p} a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i) \right) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)) \right) \\ &\quad \otimes \left(\frac{1}{n} \sum_{i=1}^n a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i)) \right) \\ &\quad + \left(\frac{1}{p} - 1 \right) \left(\frac{1}{n} \sum_{i=1}^n a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)) \right) \\ &\quad \otimes \left(\frac{1}{n} \sum_{i=1}^n \left(e_{i, \setminus k, \setminus r} + a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) + \frac{1}{p} a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i) \right) \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i)) \right), \end{aligned}$$

where we denote hereafter that

$$e_{i, \setminus k, \setminus r} := e_{i, \setminus k, \setminus r}(\boldsymbol{\theta}) := \sum_{l=1, l \neq k, l \neq r}^m a_l \sigma(\mathbf{w}_l^\top \mathbf{x}_i) - y_i,$$

and compared with e_i^N , $e_{i, \setminus k, \setminus r}$ still does not depend on the random variable $\boldsymbol{\delta}_N$. We remark that the expression above is consistent in that for the extreme case where $p = 1$, dropout ‘degenerates’ to gradient descent (GD), hence the covariance matrix degenerates to a zero matrix, i.e., $\Sigma = \mathbf{0}_{D \times D}$. We remark that details for the derivation of Σ is deferred to Section J.

Now, as we consider the stochastic differential equation (SDE),

$$d\boldsymbol{\Theta}_t = \mathbf{b}(\boldsymbol{\Theta}_t) dt + \boldsymbol{\sigma}(\boldsymbol{\Theta}_t) d\mathbf{W}_t, \quad \boldsymbol{\Theta}_0 = \boldsymbol{\Theta}(0), \quad (32)$$

where \mathbf{W}_t is a standard $m(d+1)$ -dimensional standard Wiener process, whose Euler–Maruyama discretization with step size $\eta > 0$ at the N -th step reads

$$\boldsymbol{\Theta}_{\eta N} = \boldsymbol{\Theta}_{\eta(N-1)} + \eta \mathbf{b}(\boldsymbol{\Theta}_{\eta(N-1)}) + \sqrt{\eta} \boldsymbol{\sigma}(\boldsymbol{\Theta}_{\eta(N-1)}) \mathbf{Z}_N,$$

where $\mathbf{Z}_N \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{m(d+1)})$ and $\boldsymbol{\Theta}_0 = \boldsymbol{\Theta}(0)$. Thus, if we set

$$\begin{aligned} \mathbf{b}(\boldsymbol{\Theta}) &:= -\nabla_{\boldsymbol{\Theta}} L_S(\boldsymbol{\Theta}), \\ \boldsymbol{\sigma}(\boldsymbol{\Theta}) &:= \sqrt{\eta} (\Sigma(\boldsymbol{\Theta}))^{\frac{1}{2}}, \\ \boldsymbol{\Theta}_0 &:= \boldsymbol{\theta}_0, \end{aligned} \quad (33)$$

then we would expect (32) to be a ‘good’ approximation of (31) with the time identification $t = \eta N$. Based on the earlier work of Li et al. (2017), since the path of dropout and the counterpart of SDE are driven by noises sampled in different spaces. Firstly, notice that the stochastic process $\{\boldsymbol{\theta}_N\}_{N \geq 0}$ induces a probability measure on the product space $\mathbb{R}^D \times \mathbb{R}^D \times \dots \times \mathbb{R}^D \times \dots$, whereas $\{\boldsymbol{\Theta}_t\}_{t \geq 0}$ induces a probability measure on $\mathcal{C}([0, \infty), \mathbb{R}^D)$. To compare them, one can form a piece-wise linear interpolation of the former. Alternatively, as we do in this work, we sample a discrete number of points from the latter. Secondly, the process $\{\boldsymbol{\theta}_N\}_{N \geq 0}$ is adapted to the filtration generated by \mathcal{F}_N whereas the process $\{\boldsymbol{\Theta}_t\}_{t \geq 0}$ is adapted to an independent Wiener filtration \mathcal{F}_t . Hence, it is not appropriate to compare individual sample paths. Rather, we define below a sense of *weak* approximations (Kloeden and Platen, 2011, Section 9.7) by comparing the distributions of the two processes.

To compare different discrete time approximations, we need to take the rate of weak convergence into consideration, and we also need to choose an appropriate class of functions as the space of test functions. We introduce the following set of smooth functions:

$$\mathcal{C}_b^M(\mathbb{R}^{m(d+1)}) = \left\{ f \in \mathcal{C}^M(\mathbb{R}^{m(d+1)}) \mid \|f\|_{\mathcal{C}^M} := \sum_{|\beta| \leq M} \|D^\beta f\|_\infty < \infty \right\},$$

where D is the usual differential operator. We remark that $\mathcal{C}_b^M(\mathbb{R}^D)$ is a subset of $\mathcal{G}(\mathbb{R}^D)$, the class of functions with polynomial growth, which is chosen to be the space of test functions in previous works (Li et al., 2017; Kloeden and Platen, 2011; Malladi et al., 2022).

Before we proceed to the definition of weak approximation, to ensure the rigor and validity of our analysis, we shall assert an assumption regarding the existence and uniqueness of solutions to the SDE (32).

Assumption 2. *There exists $T^* > 0$, such that for any time $t \in [0, T^*]$, there exists a unique t -continuous solution Θ_t of the initial value problem:*

$$d\Theta_t = \mathbf{b}(\Theta_t) dt + \sigma(\Theta_t) d\mathbf{W}_t, \quad \Theta_0 = \Theta(0),$$

with the property that Θ_t is adapted to the filtration \mathcal{F}_t generated by \mathbf{W}_s for all time $s \leq t$. Furthermore, for any $t \in [0, T^*]$,

$$\mathbb{E} \int_0^t \|\Theta_s(\cdot)\|_2^2 ds < \infty.$$

Moreover, we assume that the second, fourth and sixth moments of the solution to SDE (32) are uniformly bounded with respect to time t , i.e., for each $l \in [3]$, there exists $C(T^*, \Theta_0) > 0$, such that

$$\sup_{0 \leq s \leq T^*} \mathbb{E} \|\Theta_s(\cdot)\|_2^{2l} \leq C(T^*, \Theta_0). \quad (34)$$

As for the dropout iterations (27), we assume further that the second, fourth and sixth moments of the dropout iterations (27) are uniformly bounded with respect to the number of iterations N , i.e., let $0 < \eta < 1$, $T > 0$ and set $N_{T,\eta} := \lfloor \frac{T}{\eta} \rfloor$, then for each $l \in [3]$, there exists $T^* > 0$ and $\eta_0 > 0$, such that for any given learning rate $\eta \leq \eta_0$ and all $N \in [0 : N_{T^*,\eta}]$, there exists $C(T^*, \theta_0, \eta_0) > 0$, such that

$$\sup_{0 \leq N \leq [N_{T^*,\eta}]} \mathbb{E} \|\theta_N\|_2^{2l} \leq C(T^*, \theta_0, \eta_0). \quad (35)$$

We remark that if $\mathcal{G}(\mathbb{R}^D)$ is chosen to be the test functions in Li et al. (2019), then similar relations to (34) and (35) shall be imposed, except that in our cases, we only require the second, fourth and sixth moments to be uniformly bounded, while in their cases, all $2l$ -moments are required for $l \geq 1$. Establishments of the validity of Assumption 2 regarding the existence and uniqueness of the SDE will be exhibited in Section I.

The definition of weak approximation is stated out as follows.

Definition 6. *The SDE (32) is an order α weak approximation to the dropout (27), if for every $g \in \mathcal{C}_b^M(\mathbb{R}^{m(d+1)})$, there exists $C > 0$ and $\eta_0 > 0$, such that given any $\eta \leq \eta_0$ and $T \leq T^*$, then for all $N \in [N_{T,\eta}]$,*

$$|\mathbb{E}g(\Theta_{\eta N}) - \mathbb{E}g(\theta_N)| \leq C(T, g, \eta_0)\eta^\alpha. \quad (36)$$

H SEMIGROUP AND PROOF DETAILS FOR THE MAIN THEOREM

In this section, we use a semigroup approach (Feng et al., 2018) to study the time-homogeneous Markov chains (processes) formed by dropout.

H.1 DISCRETE AND CONTINUOUS SEMIGROUP

Definition 7. A Markov operator over a Polish space \mathcal{X} is a bounded linear operator $\mathcal{P} : \mathcal{B}_b(\mathcal{X}) \rightarrow \mathcal{B}_b(\mathcal{X})$ satisfying

- $\mathcal{P}\mathbf{1} = \mathbf{1}$;
- $\mathcal{P}\varphi$ is positive whenever φ is positive;
- If a sequence $\{\varphi_n\} \subset \mathcal{B}_b(\mathcal{X})$ converges pointwise to an element $\varphi \in \mathcal{B}_b(\mathcal{X})$, then $\mathcal{P}\varphi_n$ converges pointwise to $\mathcal{P}\varphi$;

To demonstrate further inequalities that Markov operators satisfy, we offer the following proposition

Proposition 1. A Markov operator $\mathcal{P} : \mathcal{B}_b(\mathcal{X}) \rightarrow \mathcal{B}_b(\mathcal{X})$ over a Polish space \mathcal{X} satisfies

- $(\mathcal{P}f(\mathbf{x}))^+ \leq \mathcal{P}f^+(\mathbf{x})$;
- $(\mathcal{P}f(\mathbf{x}))^- \leq \mathcal{P}f^-(\mathbf{x})$;
- $|\mathcal{P}f(\mathbf{x})| \leq \mathcal{P}|f(\mathbf{x})|$.

Moreover, if the Polish space \mathcal{X} is equipped with a measure μ , a function $f : \mathcal{X} \rightarrow \mathbb{R}$ is said to be an element of $\mathcal{L}^1(\mathcal{X})$ if

$$\int_{\mathcal{X}} |f| d\mu < \infty.$$

Then for every $f \in \mathcal{L}^1(\mathcal{X})$, the following holds

- $\|\mathcal{P}f\|_1 \leq \|f\|_1$.

In mathematics, the positive part of a real function is defined by the formula

$$f^+(\mathbf{x}) = \max(f(\mathbf{x}), 0) = \begin{cases} f(\mathbf{x}) & \text{if } f(\mathbf{x}) > 0, \\ 0 & \text{otherwise.} \end{cases}$$

Similarly, the negative part of f is defined as

$$f^-(\mathbf{x}) = \max(-f(\mathbf{x}), 0) = -\min(f(\mathbf{x}), 0) = \begin{cases} -f(\mathbf{x}) & \text{if } f(\mathbf{x}) < 0, \\ 0 & \text{otherwise.} \end{cases}$$

We proceed to the proof for Proposition 1

Proof. From the definition of f^+ and f^- , it follows that

$$\begin{aligned} (\mathcal{P}f)^+ &= (\mathcal{P}f^+ - \mathcal{P}f^-)^+ = \max(0, \mathcal{P}f^+ - \mathcal{P}f^-) \\ &\leq \max(0, \mathcal{P}f^+) = \mathcal{P}f^+. \end{aligned}$$

Similarly, we obtain that

$$\begin{aligned} (\mathcal{P}f)^- &= (\mathcal{P}f^+ - \mathcal{P}f^-)^- = \max(0, \mathcal{P}f^- - \mathcal{P}f^+) \\ &\leq \max(0, \mathcal{P}f^-) = \mathcal{P}f^-. \end{aligned}$$

Hence for the last inequality

$$\begin{aligned} |\mathcal{P}f| &= (\mathcal{P}f)^+ + (\mathcal{P}f)^- \\ &\leq \mathcal{P}f^+ + \mathcal{P}f^- \\ &= \mathcal{P}(f^+ + f^-) = \mathcal{P}|f|. \end{aligned}$$

Finally, by integrating the above relation over \mathcal{X} , we obtain that

$$\begin{aligned}\|\mathcal{P}f\|_1 &= \int_{\mathcal{X}} |\mathcal{P}f| d\mu \\ &\leq \int_{\mathcal{X}} \mathcal{P}|f| d\mu = \int_{\mathcal{X}} |f| d\mu = \|f\|_1.\end{aligned}\tag{37}$$

□

Inequality (37) is extremely important, and any operator \mathcal{P} that satisfies it is called a contraction. This relation is known as the contractive property of \mathcal{P} . To illustrate its power, note that for any $f \in \mathcal{L}^1(\mathcal{X})$, we have

$$\|\mathcal{P}^n f\|_1 = \|\mathcal{P} \circ \mathcal{P}^{n-1} f\|_1 \leq \|\mathcal{P}^{n-1} f\|_1.$$

As we consider Markov processes with continuous time, it is natural to consider a family of Markov operators indexed by time. We call such a family a Markov semigroup (Hairer, 2008), provided that it satisfies the relation

$$\mathcal{P}_{t+s} = \mathcal{P}_t \circ \mathcal{P}_s, \quad \text{for any time } s, t > 0.\tag{38}$$

And if given $A \in \mathcal{B}(\mathcal{X})$, where $\mathcal{B}(\mathcal{X})$ is the Borel σ -algebra on \mathcal{X} , and given any two times $s < t$, if the following holds almost surely

$$\mathbb{P}(\mathbf{X}_t \in A \mid \mathbf{X}_s) = (\mathcal{P}_{t-s} \mathbf{1}_A)(\mathbf{X}_s),$$

then we call \mathbf{X}_t a time-homogeneous Markov process with semigroup $\{\mathcal{P}_t\}_{t \geq 0}$.

In our case for dropout, we set the Polish space $\mathcal{X} = \mathbb{R}^D$, and since $\mathcal{C}_b^M(\mathbb{R}^D) \subset \mathcal{B}_b(\mathbb{R}^D)$, then WLOG we fix $g \in \mathcal{C}_b^M(\mathbb{R}^D)$ and define

$$\mathcal{P}_\eta g(\tilde{\theta}) := \mathbb{E} \left[g \left(\tilde{\theta} - \eta \nabla_{\theta} R_S^{\text{drop}}(\theta; \delta) \mid_{\theta=\tilde{\theta}} \right) \right].\tag{39}$$

We conclude that the dropout iterations (27) forms a time-homogeneous Markov chain with discrete Markov semigroup $\{\mathcal{P}_\eta^n\}_{n \geq 0}$.

As for the SDE (32), based on Assumption 2 and combined with the results in (Hairer, 2008, Example 2.11), the Markov semigroup $\{\mathcal{P}_t\}_{t \geq 0}$ associated to the solutions of the SDE reads: For any $g \in \mathcal{B}_b(\mathbb{R}^D)$,

$$\partial_t \mathcal{P}_t g = \mathcal{L} \mathcal{P}_t g,$$

where \mathcal{L} is termed the *generator* of the diffusion process (32), which reads

$$\mathcal{L}g := \langle \mathbf{b}, \nabla_{\Theta} g \rangle + \frac{1}{2} \boldsymbol{\sigma} \boldsymbol{\sigma}^\top : \nabla_{\Theta}^2 g.\tag{40}$$

Moreover, for a fixed test function $g \in \mathcal{C}_b^M(\mathbb{R}^D)$, then for any two times $s, t \geq 0$,

$$\mathcal{P}_t g(\Theta_s) := \exp(t\mathcal{L})g(\Theta_s) := \mathbb{E}_{\Theta_s} [g(\Theta_{t+s})],\tag{41}$$

and $\{\mathcal{P}_t\}_{t \geq 0}$ forms a continuous Markov semigroup for the SDE (32).

H.2 SEMIGROUP EXPANSION WITH ACCURACY OF ORDER ONE

Our results are essentially based on Itô-Taylor expansions (Kloeden and Platen, 2011) or Taylor's theorem with the Lagrange form of the remainder (Li et al., 2019, Lemma 27).

Theorem 2 (Order-1 accuracy). *Fix time $T \leq T^*$, if we choose*

$$\begin{aligned}\mathbf{b}(\Theta) &:= -\nabla_{\Theta} L_S(\Theta), \\ \boldsymbol{\sigma}(\Theta) &:= \sqrt{\eta} (\boldsymbol{\Sigma}(\Theta))^{\frac{1}{2}},\end{aligned}$$

then for all $t \in [0, T]$, the stochastic processes Θ_t satisfying

$$d\Theta_t = \mathbf{b}(\Theta_t) dt + \boldsymbol{\sigma}(\Theta_t) d\mathbf{W}_t, \quad \Theta_0 = \Theta(0),\tag{42}$$

is an order-1 approximation of dropout (27), i.e., given any test function $g \in \mathcal{C}_b^4(\mathbb{R}^D)$, there exists $\eta_0 > 0$ and $C(T, \|g\|_{\mathcal{C}^4}, \eta_0) > 0$, such that for any $\eta \leq \eta_0$ and $T \leq T^*$, and for all $N \in [N_T, \eta]$, the following holds:

$$|\mathbb{E}g(\theta_N) - \mathbb{E}g(\Theta_{\eta N})| \leq C(T, \|g\|_{\mathcal{C}^4}, \theta_0, \eta_0) \delta,\tag{43}$$

where $\theta_0 = \Theta_0$.

Proof. By application of Taylor's theorem with the Lagrange form of the remainder, we have that for some $\alpha \geq 1$,

$$\begin{aligned} g(\boldsymbol{\vartheta}) - g(\tilde{\boldsymbol{\vartheta}}) &= \sum_{s=1}^{\alpha} \frac{1}{s!} \sum_{i_1, \dots, i_j=1}^D \prod_{j=1}^s [\boldsymbol{\vartheta}_{(i_j)} - \tilde{\boldsymbol{\vartheta}}_{(i_j)}] \frac{\partial^s g}{\partial \boldsymbol{\vartheta}_{(i_1)} \dots \partial \boldsymbol{\vartheta}_{(i_j)}}(\tilde{\boldsymbol{\vartheta}}) \\ &\quad + \frac{1}{(\alpha+1)!} \sum_{i_1, \dots, i_j=1}^D \prod_{j=1}^{\alpha+1} [\boldsymbol{\vartheta}_{(i_j)} - \tilde{\boldsymbol{\vartheta}}_{(i_j)}] \frac{\partial^{\alpha+1} g}{\partial \boldsymbol{\vartheta}_{(i_1)} \dots \partial \boldsymbol{\vartheta}_{(i_j)}}(\gamma \boldsymbol{\vartheta} + (1-\gamma)\tilde{\boldsymbol{\vartheta}}), \end{aligned}$$

for some $\gamma \in (0, 1)$. We adopt the Einstein's summation convention, where repeated (spatial) indices are summed, i.e.,

$$\mathbf{x}_{(i)} \mathbf{x}_{(i)} := \sum_{i=1}^D \mathbf{x}_{(i)} \mathbf{x}_{(i)}.$$

As we choose $\boldsymbol{\vartheta} := \boldsymbol{\theta}_1$, $\tilde{\boldsymbol{\vartheta}} := \boldsymbol{\theta}_0$ and $\alpha = 1$, then we obtain that

$$\begin{aligned} g(\boldsymbol{\theta}_1) - g(\boldsymbol{\theta}_0) &= \langle \nabla_{\boldsymbol{\theta}} g(\boldsymbol{\theta}_0), \boldsymbol{\theta}_1 - \boldsymbol{\theta}_0 \rangle \\ &\quad + \frac{1}{2} \nabla_{\boldsymbol{\theta}}^2 g(\gamma \boldsymbol{\theta}_1 + (1-\gamma)\boldsymbol{\theta}_0) : (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_0) \otimes (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_0) \\ &= \langle \nabla_{\boldsymbol{\theta}} g(\boldsymbol{\theta}_0), \boldsymbol{\theta}_1 - \boldsymbol{\theta}_0 \rangle + \frac{1}{2} \nabla_{\tilde{\boldsymbol{\theta}}}^2 g(\tilde{\boldsymbol{\theta}}_0) : (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_0) \otimes (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_0), \end{aligned}$$

where $\tilde{\boldsymbol{\theta}}_0 := \gamma \boldsymbol{\theta}_1 + (1-\gamma)\boldsymbol{\theta}_0$, and we observe that since

$$\boldsymbol{\theta}_1 - \boldsymbol{\theta}_0 = -\eta \nabla_{\boldsymbol{\theta}} L_{\mathcal{S}}(\boldsymbol{\theta}) \Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0} + \sqrt{\eta} \mathbf{V}(\boldsymbol{\theta}_0),$$

then

$$\begin{aligned} \mathbb{E}g(\boldsymbol{\theta}_1) - \mathbb{E}g(\boldsymbol{\theta}_0) &= \langle \nabla_{\boldsymbol{\theta}} g(\boldsymbol{\theta}_0), \mathbb{E}\boldsymbol{\theta}_1 - \mathbb{E}\boldsymbol{\theta}_0 \rangle + \frac{1}{2} \mathbb{E} \left[\nabla_{\tilde{\boldsymbol{\theta}}}^2 g(\tilde{\boldsymbol{\theta}}_0) : (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_0) \otimes (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_0) \right] \\ &= -\eta \left\langle \nabla_{\boldsymbol{\theta}} g(\boldsymbol{\theta}_0), \nabla_{\boldsymbol{\theta}} L_{\mathcal{S}}(\boldsymbol{\theta}) \Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0} \right\rangle + E_{\eta}^1(\boldsymbol{\theta}_0), \end{aligned}$$

where the remainder term $E_{\eta}^1(\cdot) : \mathbb{R}^D \rightarrow \mathbb{R}$, whose expression reads

$$E_{\eta}^1(\boldsymbol{\theta}_0) := \frac{1}{2} \mathbb{E} \left[\nabla_{\tilde{\boldsymbol{\theta}}}^2 g(\tilde{\boldsymbol{\theta}}_0) : (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_0) \otimes (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_0) \right], \quad (44)$$

and we remark that $\tilde{\boldsymbol{\theta}}_0$ and $\boldsymbol{\theta}_1$ are implicitly defined by $\boldsymbol{\theta}_0$. Then, directly from Assumption 2, we obtain that

$$\begin{aligned} E_{\eta}^1(\boldsymbol{\theta}_0) &= \frac{1}{2} \mathbb{E} \left[\nabla_{\tilde{\boldsymbol{\theta}}}^2 g(\tilde{\boldsymbol{\theta}}_0) : (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_0) \otimes (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_0) \right] \\ &\leq \frac{1}{2} \|g\|_{C^4} \mathbb{E} \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_0\|_2^2 = \eta^2 \|g\|_{C^4} \mathbb{E} \left[\left\| \nabla_{\boldsymbol{\theta}} R_{\mathcal{S}}^{\text{drop}}(\boldsymbol{\theta}_0; \boldsymbol{\delta}_1) \right\|_2^2 \right] \\ &\leq \eta^2 \|g\|_{C^4} C(T^*, \boldsymbol{\theta}_0, \eta), \end{aligned}$$

since $\nabla_{\boldsymbol{\theta}} L_{\mathcal{S}}(\boldsymbol{\theta})$ and $\boldsymbol{\Sigma}(\boldsymbol{\theta})$ can be bounded above by the second and fourth moments of the dropout iteration (27).

We observe that

$$\boldsymbol{\Theta}_{\eta} - \boldsymbol{\Theta}_0 = \int_0^{\eta} \mathbf{b}(\boldsymbol{\Theta}_s) ds + \int_0^{\eta} \boldsymbol{\sigma}(\boldsymbol{\Theta}_s) d\mathbf{W}_s.$$

As we choose $\boldsymbol{\vartheta} := \boldsymbol{\Theta}_{\eta}$, $\tilde{\boldsymbol{\vartheta}} := \boldsymbol{\Theta}_0$ and $\alpha = 1$, then we obtain that

$$\begin{aligned} g(\boldsymbol{\Theta}_{\eta}) - g(\boldsymbol{\Theta}_0) &= \langle \nabla_{\boldsymbol{\Theta}} g(\boldsymbol{\Theta}_0), \boldsymbol{\Theta}_{\eta} - \boldsymbol{\Theta}_0 \rangle \\ &\quad + \frac{1}{2} \nabla_{\tilde{\boldsymbol{\Theta}}}^2 g(\tilde{\boldsymbol{\Theta}}_0) : (\boldsymbol{\Theta}_{\eta} - \boldsymbol{\Theta}_0) \otimes (\boldsymbol{\Theta}_{\eta} - \boldsymbol{\Theta}_0), \end{aligned}$$

where

$$\tilde{\boldsymbol{\Theta}}_0 := \gamma \boldsymbol{\Theta}_{\eta} + (1-\gamma)\boldsymbol{\Theta}_0,$$

for some $\gamma \in (0, 1)$. Then

$$\begin{aligned} & \mathbb{E}g(\Theta_\eta) - \mathbb{E}g(\Theta_0) \\ &= \langle \nabla_{\Theta} g(\Theta_0), \mathbb{E}\Theta_\eta - \mathbb{E}\Theta_0 \rangle + \frac{1}{2} \mathbb{E} \left[\nabla_{\Theta}^2 g(\tilde{\Theta}_0) : (\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0) \right] \\ &= \left\langle \nabla_{\Theta} g(\Theta_0), \int_0^\eta \mathbb{E}[\mathbf{b}(\Theta_s)] ds \right\rangle + \frac{1}{2} \mathbb{E} \left[\nabla_{\Theta}^2 g(\tilde{\Theta}_0) : (\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0) \right], \end{aligned}$$

and since

$$\langle \nabla_{\Theta} g(\Theta_0), \mathbb{E}[\mathbf{b}(\Theta_s)] \rangle = \langle \nabla_{\Theta} g(\Theta_0), \mathbb{E}[\mathbf{b}(\Theta_0)] \rangle + \int_0^s \mathcal{L} \langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle (\Theta_v) dv,$$

then we obtain that

$$\begin{aligned} \mathbb{E}g(\Theta_\eta) - \mathbb{E}g(\Theta_0) &= \eta \langle \nabla_{\Theta} g(\Theta_0), \mathbb{E}[\mathbf{b}(\Theta_0)] \rangle + \int_0^\eta \int_0^s \mathcal{L} \langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle (\Theta_v) dv ds \\ &\quad + \frac{1}{2} \mathbb{E} \left[\nabla_{\Theta}^2 g(\tilde{\Theta}_0) : (\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0) \right] \\ &= \eta \langle \nabla_{\Theta} g(\Theta_0), \mathbf{b}(\Theta_0) \rangle + \eta^2 \bar{E}_\eta^1(\Theta_0), \end{aligned}$$

where the remainder term $\bar{E}_\eta^1(\cdot) : \mathbb{R}^D \rightarrow \mathbb{R}$, whose expression reads

$$\begin{aligned} \bar{E}_\eta^1(\Theta_0) &:= \int_0^\eta \int_0^s \mathcal{L} \langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle (\Theta_v) dv ds \\ &\quad + \frac{1}{2} \mathbb{E} \left[\nabla_{\Theta}^2 g(\tilde{\Theta}_0) : (\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0) \right], \end{aligned} \tag{45}$$

and we remark that $\tilde{\Theta}_0$ and Θ_η are implicitly defined by Θ_0 . As we choose

$$\begin{aligned} \mathbf{b}(\Theta) &= -\nabla_{\Theta} L_S(\Theta), \\ \sigma(\Theta) &= \sqrt{\eta} (\Sigma(\Theta))^{\frac{1}{2}}, \end{aligned}$$

then we carry out the computation for $\mathcal{L} \langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle (\Theta_v)$,

$$\begin{aligned} \mathcal{L} \langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle (\Theta_v) &= \langle \nabla_{\Theta} L_S(\Theta_v), \nabla_{\Theta} \langle \nabla_{\Theta} g(\Theta_0), \nabla_{\Theta} L_S(\Theta) \rangle |_{\Theta=\Theta_v} \rangle \\ &\quad + \frac{\eta}{2} \Sigma(\Theta_v) : \nabla_{\Theta}^2 (\langle \nabla_{\Theta} g(\Theta_0), \nabla_{\Theta} L_S(\Theta) \rangle) |_{\Theta=\Theta_v}, \end{aligned}$$

since $\nabla_{\Theta} L_S(\Theta)$, $\nabla_{\Theta}^2 L_S(\Theta)$, $\nabla_{\Theta}^3 L_S(\Theta)$ and $\Sigma(\Theta)$ can be bounded above by the second, fourth and sixth moments of the solution to SDE (32), hence we may apply the mean value theorem to (45) and obtain that

$$\begin{aligned} |\bar{E}_\eta^1(\Theta_0)| &= \left| \int_0^\eta s \mathcal{L} \langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle (\tilde{\Theta}_s) ds + \frac{1}{2} \mathbb{E} \left[\nabla_{\Theta}^2 g(\tilde{\Theta}_0) : (\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0) \right] \right| \\ &\leq \int_0^\eta s \|g\|_{C^4} C(T^*, \Theta_0) ds + \frac{1}{2} \|g\|_{C^4} \mathbb{E} \|\Theta_\eta - \Theta_0\|_2^2 \\ &\leq \frac{\eta^2}{2} \|g\|_{C^4} C(T^*, \Theta_0) + \|g\|_{C^4} \mathbb{E} \left\| \int_0^\eta \mathbf{b}(\Theta_s) ds + \int_0^\eta \sigma(\Theta_s) d\mathbf{W}_s \right\|_2^2 \\ &\leq \frac{\eta^2}{2} \|g\|_{C^4} C(T^*, \Theta_0) + 2 \|g\|_{C^4} \mathbb{E} \left\| \int_0^\eta \mathbf{b}(\Theta_s) ds \right\|_2^2 \\ &\quad + 2 \|g\|_{C^4} \mathbb{E} \left\| \int_0^\eta \sigma(\Theta_s) d\mathbf{W}_s \right\|_2^2 \\ &\leq \frac{\eta^2}{2} \|g\|_{C^4} C(T^*, \Theta_0) + 2 \|g\|_{C^4} \eta^2 \mathbb{E} \left\| \nabla_{\Theta} L_S(\tilde{\Theta}_0) \right\|_2^2 \\ &\quad + 2 \|g\|_{C^4} \mathbb{E} \int_0^\eta \|\sigma(\Theta_s)\|_{\mathbb{F}}^2 ds \\ &\leq \frac{\eta^2}{2} \|g\|_{C^4} C(T^*, \Theta_0) + 2 \|g\|_{C^4} \eta^2 \mathbb{E} \left\| \nabla_{\Theta} L_S(\tilde{\Theta}_0) \right\|_2^2 \\ &\quad + 2 \|g\|_{C^4} \eta \mathbb{E} \left[\eta \left\| \Sigma(\tilde{\Theta}_0) \right\|_{\mathbb{F}} \right] \leq \eta^2 \|g\|_{C^4} C(T^*, \Theta_0). \end{aligned}$$

To sum up for now,

$$\begin{aligned} |\mathbb{E}g(\boldsymbol{\theta}_1) - \mathbb{E}g(\boldsymbol{\Theta}_\eta)| &= \left| \mathbb{E}g(\boldsymbol{\theta}_0) - \eta \left\langle \nabla_{\boldsymbol{\theta}}g(\boldsymbol{\theta}_0), \nabla_{\boldsymbol{\theta}}L_S(\boldsymbol{\theta}) \Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0} \right\rangle + E_\eta^1(\boldsymbol{\theta}_0) \right. \\ &\quad \left. - \mathbb{E}g(\boldsymbol{\Theta}_0) - \eta \langle \nabla_{\boldsymbol{\Theta}}g(\boldsymbol{\Theta}_0), \mathbf{b}(\boldsymbol{\Theta}_0) \rangle + \bar{E}_\eta^1(\boldsymbol{\Theta}_0) \right|, \end{aligned}$$

since $\boldsymbol{\theta}_0 = \boldsymbol{\Theta}_0$ and $\mathbf{b}(\boldsymbol{\Theta}_0) = -\nabla_{\boldsymbol{\Theta}}L_S(\boldsymbol{\Theta}) \Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0}$, thus

$$\begin{aligned} |\mathcal{P}_\eta^1 g(\boldsymbol{\theta}_0) - \mathcal{P}_\eta g(\boldsymbol{\Theta}_0)| &= |\mathbb{E}g(\boldsymbol{\theta}_1) - \mathbb{E}g(\boldsymbol{\Theta}_\eta)| \\ &\leq |E_\eta^1(\boldsymbol{\theta}_0)| + |\bar{E}_\eta^1(\boldsymbol{\Theta}_0)| \\ &\leq \eta^2 \|g\|_{C^4} C(T^*, \boldsymbol{\theta}_0, \eta_0) + \eta^2 \|g\|_{C^4} C(T^*, \boldsymbol{\Theta}_0) \\ &= \mathcal{O}(\eta^2). \end{aligned} \tag{46}$$

For the N -th step iteration, since

$$|\mathbb{E}g(\boldsymbol{\theta}_N) - \mathbb{E}g(\boldsymbol{\Theta}_{\eta N})| = |\mathcal{P}_\eta^N g(\boldsymbol{\theta}_0) - \mathcal{P}_{\eta N} g(\boldsymbol{\Theta}_0)|,$$

and the RHS of the above equation can be written into a telescoping sum as

$$\mathcal{P}_\eta^N g(\boldsymbol{\theta}_0) - \mathcal{P}_{\eta N} g(\boldsymbol{\Theta}_0) = \sum_{l=1}^N (\mathcal{P}_\eta^{N-l+1} \circ \mathcal{P}_{(l-1)\eta} g(\boldsymbol{\theta}_0) - \mathcal{P}_\eta^{N-l} \circ \mathcal{P}_{l\eta} g(\boldsymbol{\Theta}_0)),$$

hence by application of Proposition 1, we obtain that

$$\begin{aligned} |\mathbb{E}g(\boldsymbol{\theta}_N) - \mathbb{E}g(\boldsymbol{\Theta}_{\eta N})| &\leq \sum_{l=1}^N |\mathcal{P}_\eta^{N-l+1} \circ \mathcal{P}_{(l-1)\eta} g(\boldsymbol{\theta}_0) - \mathcal{P}_\eta^{N-l} \circ \mathcal{P}_{l\eta} g(\boldsymbol{\Theta}_0)| \\ &\leq \sum_{l=1}^N |\mathcal{P}_\eta^{N-l} \circ (\mathcal{P}_\eta^1 \circ \mathcal{P}_{(l-1)\eta} - \mathcal{P}_\eta \circ \mathcal{P}_{(l-1)\eta}) g(\boldsymbol{\Theta}_0)|, \end{aligned}$$

since $(\mathcal{P}_\eta^1 \circ \mathcal{P}_{(l-1)\eta} - \mathcal{P}_\eta \circ \mathcal{P}_{(l-1)\eta}) g(\boldsymbol{\Theta}_0)$ can be regarded as $\mathcal{L}^1(\mathbb{R}^D)$ if we choose measure μ to be the delta measure concentrated on $\boldsymbol{\Theta}_0$. i.e.,

$$\mu := \delta_{\boldsymbol{\Theta}_0},$$

hence by the contraction property of Markov operators, we obtain further that

$$\begin{aligned} |\mathbb{E}g(\boldsymbol{\theta}_N) - \mathbb{E}g(\boldsymbol{\Theta}_{\eta N})| &\leq \sum_{l=1}^N |(\mathcal{P}_\eta^1 \circ \mathcal{P}_{(l-1)\eta} - \mathcal{P}_\eta \circ \mathcal{P}_{(l-1)\eta}) g(\boldsymbol{\Theta}_0)| \\ &\leq \sum_{l=1}^N |\mathcal{P}_\eta^1 g(\boldsymbol{\Theta}_{(l-1)\eta}) - \mathcal{P}_\eta g(\boldsymbol{\Theta}_{(l-1)\eta})|. \end{aligned}$$

By taking expectation conditioned on $\boldsymbol{\Theta}_{(l-1)\eta}$, then similar to the relation (46), the following holds

$$\begin{aligned} |\mathcal{P}_\eta^1 g(\boldsymbol{\Theta}_{(l-1)\eta}) - \mathcal{P}_\eta g(\boldsymbol{\Theta}_{(l-1)\eta})| &= \mathbb{E} \left[\left| \mathbb{E}g(\boldsymbol{\theta}_l) - \mathbb{E}g(\boldsymbol{\Theta}_\eta l) \right| \Big| \boldsymbol{\Theta}_{(l-1)\eta} \right] \\ &\leq \mathbb{E} |E_\eta^1(\boldsymbol{\Theta}_{(l-1)\eta})| + \mathbb{E} |\bar{E}_\eta^1(\boldsymbol{\Theta}_{(l-1)\eta})| \\ &\leq \eta^2 \|g\|_{C^4} C(T^*, \boldsymbol{\theta}_0, \eta_0) + \eta^2 \|g\|_{C^4} C(T^*, \boldsymbol{\Theta}_0) \\ &= \mathcal{O}(\eta^2). \end{aligned}$$

We remark that the last line of the above relation is essentially based on Assumption 2, since $\mathbb{E} |E_\eta^1(\boldsymbol{\Theta}_{(l-1)\eta})|$ and $\mathbb{E} |\bar{E}_\eta^1(\boldsymbol{\Theta}_{(l-1)\eta})|$ can be bounded above by the second, fourth and sixth moments of the solution to SDE (32), hence we may apply dominated convergence theorem to obtain the last line of the above relation.

To sum up, as

$$|\mathcal{P}_\eta^N g(\boldsymbol{\theta}_0) - \mathcal{P}_{\eta N} g(\boldsymbol{\Theta}_0)| \leq \sum_{l=1}^N |\mathcal{P}_\eta^{N-l+1} \circ \mathcal{P}_{(l-1)\eta} g(\boldsymbol{\theta}_0) - \mathcal{P}_\eta^{N-l} \circ \mathcal{P}_{l\eta} g(\boldsymbol{\Theta}_0)| = N\mathcal{O}(\eta^2),$$

hence for $N = N_{T,\eta}$,

$$|\mathcal{P}_\eta^N g(\boldsymbol{\theta}_0) - \mathcal{P}_{\eta N} g(\boldsymbol{\Theta}_0)| = N\mathcal{O}(\eta^2) = N\eta\mathcal{O}(\eta) \leq T\mathcal{O}(\eta) = \mathcal{O}(\eta).$$

□

H.3 SEMIGROUP EXPANSION WITH ACCURACY OF ORDER TWO

Theorem 3 (Order-2 accuracy). *Fix time $T \leq T^*$, if we choose*

$$\begin{aligned} \mathbf{b}(\Theta) &= -\nabla_{\Theta} \left(L_S(\Theta) + \frac{\eta}{4} \|\nabla_{\Theta} L_S(\Theta)\|_2^2 \right), \\ \sigma(\Theta) &= \sqrt{\eta} (\Sigma(\Theta))^{\frac{1}{2}}, \end{aligned}$$

then for all $t \in [0, T]$, the stochastic processes Θ_t satisfying

$$d\Theta_t = \mathbf{b}(\Theta_t) dt + \sigma(\Theta_t) d\mathbf{W}_t, \quad \Theta_0 = \Theta(0), \quad (47)$$

is an order-2 approximation of dropout (27), i.e., given any test function $g \in \mathcal{C}_b^6(\mathbb{R}^D)$, there exists $\eta_0 > 0$ and $C(T, \|g\|_{C^6}, \eta_0) > 0$, such that for any $\eta \leq \eta_0$ and $T \leq T^*$, and for all $N \in [N_{T,\eta}]$, the following holds:

$$|\mathbb{E}g(\theta_N) - \mathbb{E}g(\Theta_{\eta N})| \leq C(T, \|g\|_{C^6}, \theta_0, \eta_0) \delta, \quad (48)$$

where $\theta_0 = \Theta_0$.

Proof. By application of Taylor's theorem with the Lagrange form of the remainder, we have that for some $\alpha \geq 1$,

$$\begin{aligned} g(\vartheta) - g(\tilde{\vartheta}) &= \sum_{s=1}^{\alpha} \frac{1}{s!} \sum_{i_1, \dots, i_j=1}^D \prod_{j=1}^s [\vartheta_{(i_j)} - \tilde{\vartheta}_{(i_j)}] \frac{\partial^s g}{\partial \vartheta_{(i_1)} \dots \partial \vartheta_{(i_j)}}(\tilde{\vartheta}) \\ &\quad + \frac{1}{(\alpha+1)!} \sum_{i_1, \dots, i_j=1}^D \prod_{j=1}^{\alpha+1} [\vartheta_{(i_j)} - \tilde{\vartheta}_{(i_j)}] \frac{\partial^{\alpha+1} g}{\partial \vartheta_{(i_1)} \dots \partial \vartheta_{(i_j)}}(\gamma \vartheta + (1-\gamma) \tilde{\vartheta}), \end{aligned}$$

for some $\gamma \in (0, 1)$.

As we choose $\vartheta := \theta_1$, $\tilde{\vartheta} := \theta_0$ and $\alpha = 2$, with slight misuse of the Frobenius inner product notation, we obtain that

$$\begin{aligned} g(\theta_1) - g(\theta_0) &= \langle \nabla_{\theta} g(\theta_0), \theta_1 - \theta_0 \rangle + \frac{1}{2} \nabla_{\theta}^2 g(\theta_0) : (\theta_1 - \theta_0) \otimes (\theta_1 - \theta_0) \\ &\quad + \frac{1}{6} \nabla_{\theta}^3 g(\gamma \theta_1 + (1-\gamma) \theta_0) : (\theta_1 - \theta_0) \otimes (\theta_1 - \theta_0) \otimes (\theta_1 - \theta_0) \\ &= \langle \nabla_{\theta} g(\theta_0), \theta_1 - \theta_0 \rangle + \frac{1}{2} \nabla_{\theta}^2 g(\theta_0) : (\theta_1 - \theta_0) \otimes (\theta_1 - \theta_0) \\ &\quad + \frac{1}{6} \nabla_{\theta}^3 g(\tilde{\theta}_0) : (\theta_1 - \theta_0) \otimes (\theta_1 - \theta_0) \otimes (\theta_1 - \theta_0), \end{aligned}$$

where $\tilde{\theta}_0 := \gamma \theta_1 + (1-\gamma) \theta_0$, and we observe that since

$$\theta_1 - \theta_0 = -\eta \nabla_{\theta} L_S(\theta) \Big|_{\theta=\theta_0} + \sqrt{\eta} \mathbf{V}(\theta_0),$$

then

$$\begin{aligned} \mathbb{E}g(\theta_1) - \mathbb{E}g(\theta_0) &= \langle \nabla_{\theta} g(\theta_0), \mathbb{E}\theta_1 - \mathbb{E}\theta_0 \rangle + \frac{1}{2} \nabla_{\theta}^2 g(\theta_0) : \mathbb{E}[(\theta_1 - \theta_0) \otimes (\theta_1 - \theta_0)] \\ &\quad + \frac{1}{6} \mathbb{E} \left[\nabla_{\theta}^3 g(\tilde{\theta}_0) : (\theta_1 - \theta_0) \otimes (\theta_1 - \theta_0) \otimes (\theta_1 - \theta_0) \right] \\ &= -\eta \left\langle \nabla_{\theta} g(\theta_0), \nabla_{\theta} L_S(\theta) \Big|_{\theta=\theta_0} \right\rangle \\ &\quad + \frac{\eta^2}{2} \nabla_{\theta}^2 g(\theta_0) : \left(\nabla_{\theta} L_S(\theta) \Big|_{\theta=\theta_0} \otimes \nabla_{\theta} L_S(\theta) \Big|_{\theta=\theta_0} + \Sigma(\theta_0) \right) \\ &\quad + E_{\eta}^2(\theta_0), \end{aligned}$$

where the remainder term $E_{\eta}^2(\cdot) : \mathbb{R}^D \rightarrow \mathbb{R}$, whose expression reads

$$E_{\eta}^2(\theta_0) := \frac{1}{6} \mathbb{E} \left[\nabla_{\theta}^3 g(\tilde{\theta}_0) : (\theta_1 - \theta_0) \otimes (\theta_1 - \theta_0) \otimes (\theta_1 - \theta_0) \right], \quad (49)$$

and we remark that $\tilde{\boldsymbol{\theta}}_0$ and $\boldsymbol{\theta}_1$ are implicitly defined by $\boldsymbol{\theta}_0$. Then, directly from Assumption 2, we obtain that

$$\begin{aligned} E_\eta^2(\boldsymbol{\theta}_0) &\leq \frac{1}{6} \|g\|_{C^6} \mathbb{E} \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_0\|_2^3 = \eta^3 \|g\|_{C^6} \mathbb{E} \left[\left\| \nabla_{\boldsymbol{\theta}} R_S^{\text{drop}}(\boldsymbol{\theta}_0; \boldsymbol{\delta}_1) \right\|_2^3 \right] \\ &\leq \eta^3 \|g\|_{C^6} C(T^*, \boldsymbol{\theta}_0, \eta_0), \end{aligned}$$

since $\nabla_{\boldsymbol{\theta}} L_S(\boldsymbol{\theta})$ and $\Sigma(\boldsymbol{\theta})$ can be bounded above by the second and fourth moments of the dropout iteration (27).

We observe that

$$\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0 = \int_0^\eta \mathbf{b}(\boldsymbol{\Theta}_s) ds + \int_0^\eta \boldsymbol{\sigma}(\boldsymbol{\Theta}_s) d\mathbf{W}_s.$$

As we choose $\boldsymbol{\vartheta} := \boldsymbol{\Theta}_\eta$, $\tilde{\boldsymbol{\vartheta}} := \boldsymbol{\Theta}_0$ and $\alpha = 3$, then we obtain that

$$\begin{aligned} g(\boldsymbol{\Theta}_\eta) - g(\boldsymbol{\Theta}_0) &= \langle \nabla_{\boldsymbol{\Theta}} g(\boldsymbol{\Theta}_0), \boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0 \rangle \\ &\quad + \frac{1}{2} \nabla_{\boldsymbol{\Theta}}^2 g(\boldsymbol{\Theta}_0) : (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \otimes (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \\ &\quad + \frac{1}{6} \nabla_{\boldsymbol{\Theta}}^3 g(\boldsymbol{\Theta}_0) : (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \otimes (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \otimes (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \\ &\quad + \frac{1}{24} \nabla_{\boldsymbol{\Theta}}^4 g(\tilde{\boldsymbol{\Theta}}_0) : (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \otimes (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \otimes (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \otimes (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0), \end{aligned}$$

where

$$\tilde{\boldsymbol{\Theta}}_0 := \gamma \boldsymbol{\Theta}_\eta + (1 - \gamma) \boldsymbol{\Theta}_0,$$

for some $\gamma \in (0, 1)$. Then

$$\begin{aligned} &\mathbb{E}g(\boldsymbol{\Theta}_\eta) - \mathbb{E}g(\boldsymbol{\Theta}_0) \\ &= \langle \nabla_{\boldsymbol{\Theta}} g(\boldsymbol{\Theta}_0), \mathbb{E}\boldsymbol{\Theta}_\eta - \mathbb{E}\boldsymbol{\Theta}_0 \rangle + \frac{1}{2} \nabla_{\boldsymbol{\Theta}}^2 g(\boldsymbol{\Theta}_0) : \mathbb{E}[(\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \otimes (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0)] \\ &\quad + \frac{1}{6} \nabla_{\boldsymbol{\Theta}}^3 g(\boldsymbol{\Theta}_0) : \mathbb{E}[(\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \otimes (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \otimes (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0)] \\ &\quad + \frac{1}{24} \mathbb{E} \left[\nabla_{\boldsymbol{\Theta}}^4 g(\tilde{\boldsymbol{\Theta}}_0) : (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \otimes (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \otimes (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \otimes (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \right] \\ &= \left\langle \nabla_{\boldsymbol{\Theta}} g(\boldsymbol{\Theta}_0), \int_0^\eta \mathbb{E}[\mathbf{b}(\boldsymbol{\Theta}_s)] ds \right\rangle + \frac{1}{2} \nabla_{\boldsymbol{\Theta}}^2 g(\boldsymbol{\Theta}_0) : \mathbb{E}[(\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \otimes (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0)] \\ &\quad + \frac{1}{6} \nabla_{\boldsymbol{\Theta}}^3 g(\boldsymbol{\Theta}_0) : \mathbb{E}[(\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \otimes (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \otimes (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0)] \\ &\quad + \frac{1}{24} \mathbb{E} \left[\nabla_{\boldsymbol{\Theta}}^4 g(\tilde{\boldsymbol{\Theta}}_0) : (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \otimes (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \otimes (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \otimes (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \right], \end{aligned}$$

and since

$$\langle \nabla_{\boldsymbol{\Theta}} g(\boldsymbol{\Theta}_0), \mathbb{E}[\mathbf{b}(\boldsymbol{\Theta}_s)] \rangle = \langle \nabla_{\boldsymbol{\Theta}} g(\boldsymbol{\Theta}_0), \mathbb{E}[\mathbf{b}(\boldsymbol{\Theta}_0)] \rangle + \int_0^s \mathcal{L} \langle \nabla_{\boldsymbol{\Theta}} g(\boldsymbol{\Theta}_0), \mathbf{b} \rangle(\boldsymbol{\Theta}_v) dv,$$

then we obtain that

$$\begin{aligned} \mathbb{E}g(\boldsymbol{\Theta}_\eta) - \mathbb{E}g(\boldsymbol{\Theta}_0) &= \eta \langle \nabla_{\boldsymbol{\Theta}} g(\boldsymbol{\Theta}_0), \mathbb{E}[\mathbf{b}(\boldsymbol{\Theta}_0)] \rangle + \int_0^\eta \int_0^s \mathcal{L} \langle \nabla_{\boldsymbol{\Theta}} g(\boldsymbol{\Theta}_0), \mathbf{b} \rangle(\boldsymbol{\Theta}_v) dv ds \\ &\quad + \frac{1}{2} \nabla_{\boldsymbol{\Theta}}^2 g(\boldsymbol{\Theta}_0) : \mathbb{E}[(\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \otimes (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0)] \\ &\quad + \frac{1}{6} \nabla_{\boldsymbol{\Theta}}^3 g(\boldsymbol{\Theta}_0) : \mathbb{E}[(\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \otimes (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \otimes (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0)] \\ &\quad + \frac{1}{24} \mathbb{E} \left[\nabla_{\boldsymbol{\Theta}}^4 g(\tilde{\boldsymbol{\Theta}}_0) : (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \otimes (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \otimes (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \otimes (\boldsymbol{\Theta}_\eta - \boldsymbol{\Theta}_0) \right], \end{aligned}$$

and once again since

$$\mathcal{L} \langle \nabla_{\boldsymbol{\Theta}} g(\boldsymbol{\Theta}_0), \mathbf{b} \rangle(\boldsymbol{\Theta}_v) = \mathcal{L} \langle \nabla_{\boldsymbol{\Theta}} g(\boldsymbol{\Theta}_0), \mathbf{b} \rangle(\boldsymbol{\Theta}_0) + \int_0^v \mathcal{L}(\mathcal{L} \langle \nabla_{\boldsymbol{\Theta}} g(\boldsymbol{\Theta}_0), \mathbf{b} \rangle)(\boldsymbol{\Theta}_u) du,$$

then we obtain that

$$\begin{aligned}
\mathbb{E}g(\Theta_\eta) - \mathbb{E}g(\Theta_0) &= \eta \langle \nabla_{\Theta} g(\Theta_0), \mathbb{E}[\mathbf{b}(\Theta_0)] \rangle + \int_0^\eta \int_0^s \mathcal{L} \langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle (\Theta_v) dv ds \\
&\quad + \frac{1}{2} \nabla_{\Theta}^2 g(\Theta_0) : \mathbb{E}[(\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0)] \\
&\quad + \frac{1}{6} \nabla_{\Theta}^3 g(\Theta_0) : \mathbb{E}[(\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0)] \\
&\quad + \frac{1}{24} \mathbb{E} \left[\nabla_{\Theta}^4 g(\tilde{\Theta}_0) : (\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0) \right] \\
&= \eta \langle \nabla_{\Theta} g(\Theta_0), \mathbb{E}[\mathbf{b}(\Theta_0)] \rangle + \int_0^\eta \int_0^s \mathcal{L} \langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle (\Theta_0) dv ds \\
&\quad + \int_0^\eta \int_0^s \int_0^v \mathcal{L} (\mathcal{L} \langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle) (\Theta_u) dudv ds \\
&\quad + \frac{1}{2} \nabla_{\Theta}^2 g(\Theta_0) : \mathbb{E}[(\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0)] \\
&\quad + \frac{1}{6} \nabla_{\Theta}^3 g(\Theta_0) : \mathbb{E}[(\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0)] \\
&\quad + \frac{1}{24} \mathbb{E} \left[\nabla_{\Theta}^4 g(\tilde{\Theta}_0) : (\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0) \right] \\
&= \eta \langle \nabla_{\Theta} g(\Theta_0), \mathbb{E}[\mathbf{b}(\Theta_0)] \rangle + \frac{\eta^2}{2} \mathcal{L} \langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle (\Theta_0) \\
&\quad + \frac{1}{2} \nabla_{\Theta}^2 g(\Theta_0) : \mathbb{E}[(\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0)] + \bar{E}_\eta^2(\Theta_0),
\end{aligned}$$

where the remainder term $\bar{E}_\eta^2(\cdot) : \mathbb{R}^D \rightarrow \mathbb{R}$, whose expression reads

$$\begin{aligned}
\bar{E}_\eta^2(\Theta_0) &:= \int_0^\eta \int_0^s \int_0^v \mathcal{L} (\mathcal{L} \langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle) (\Theta_u) dudv ds \\
&\quad + \frac{1}{6} \nabla_{\Theta}^3 g(\Theta_0) : \mathbb{E}[(\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0)] \\
&\quad + \frac{1}{24} \mathbb{E} \left[\nabla_{\Theta}^4 g(\tilde{\Theta}_0) : (\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0) \right],
\end{aligned} \tag{50}$$

and we remark that $\tilde{\Theta}_0$ and Θ_η are implicitly defined by Θ_0 . As we choose

$$\begin{aligned}
\mathbf{b}(\Theta) &= -\nabla_{\Theta} \left(L_S(\Theta) + \frac{\eta}{4} \|\nabla_{\Theta} L_S(\Theta)\|_2^2 \right), \\
\sigma(\Theta) &= \sqrt{\eta} (\Sigma(\Theta))^{\frac{1}{2}},
\end{aligned}$$

then we carry out the computation for $\mathcal{L} (\mathcal{L} \langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle) (\Theta_u)$,

$$\begin{aligned}
&\mathcal{L} (\mathcal{L} \langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle) (\Theta_u) \\
&= \mathcal{L} (\langle \mathbf{b}, \nabla_{\Theta} (\langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle) \rangle) (\Theta_u) + \mathcal{L} \left(\frac{\eta}{2} \Sigma : \nabla_{\Theta}^2 (\langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle) \right) (\Theta_u) \\
&= \langle \mathbf{b}, \nabla_{\Theta} (\langle \mathbf{b}, \nabla_{\Theta} (\langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle) \rangle) \rangle + \frac{\eta}{2} \Sigma : \nabla_{\Theta} (\langle \mathbf{b}, \nabla_{\Theta}^2 (\langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle) \rangle) \\
&\quad + \frac{\eta}{2} \langle \mathbf{b}, \nabla_{\Theta} (\Sigma : \nabla_{\Theta}^2 (\langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle)) \rangle + \frac{\eta^2}{4} \Sigma : \nabla_{\Theta}^2 (\Sigma : \nabla_{\Theta}^2 (\langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle)) \\
&= \mathbf{b}^\top \nabla_{\Theta} (\mathbf{b}^\top \nabla_{\Theta} \mathbf{b} \nabla_{\Theta} g(\Theta_0)) (\Theta_u) + \eta R_\eta(\Theta_u) \\
&= \left\langle \nabla_{\Theta} L_S(\Theta_u), \nabla_{\Theta} \left(\left\langle \frac{1}{2} \nabla_{\Theta} (\|\nabla_{\Theta} L_S(\Theta_u)\|_2^2), \nabla_{\Theta} g(\Theta_0) \right\rangle \right) \right\rangle + \eta R'_\eta(\Theta_u),
\end{aligned}$$

since $\nabla_{\Theta} L_{\mathcal{S}}(\Theta)$, $\nabla_{\Theta}^2 L_{\mathcal{S}}(\Theta)$, $\nabla_{\Theta}^3 L_{\mathcal{S}}(\Theta)$, $\Sigma(\Theta)$, $R_{\eta}(\Theta_u)$ and $R'_{\eta}(\Theta_u)$ can be bounded above by the second, fourth and sixth moments of the solution to SDE (32). Moreover, we observe that

$$\begin{aligned} & \mathbb{E} [(\Theta_{\eta} - \Theta_0) \otimes (\Theta_{\eta} - \Theta_0) \otimes (\Theta_{\eta} - \Theta_0)] \\ = & \mathbb{E} \left[\left(\int_0^{\eta} \mathbf{b}(\Theta_s) ds + \int_0^{\eta} \sigma(\Theta_s) d\mathbf{W}_s \right) \otimes \left(\int_0^{\eta} \mathbf{b}(\Theta_s) ds + \int_0^{\eta} \sigma(\Theta_s) d\mathbf{W}_s \right) \right. \\ & \left. \otimes \left(\int_0^{\eta} \mathbf{b}(\Theta_s) ds + \int_0^{\eta} \sigma(\Theta_s) d\mathbf{W}_s \right) \right], \end{aligned}$$

and its entry can be categorized into four types. The first one is the pure drift part, i.e.,

$$\int_0^{\eta} \mathbf{b}(\Theta_s) ds \otimes \int_0^{\eta} \mathbf{b}(\Theta_s) ds \otimes \int_0^{\eta} \mathbf{b}(\Theta_s) ds,$$

then by application of the mean value theorem and the fact that $\nabla_{\Theta} L_{\mathcal{S}}(\Theta)$, $\nabla_{\Theta}^2 L_{\mathcal{S}}(\Theta)$, $\nabla_{\Theta}^3 L_{\mathcal{S}}(\Theta)$, and $\Sigma(\Theta)$ can be bounded above by the second, fourth and sixth moments of the solution to SDE (32), we obtain that

$$\begin{aligned} & \mathbb{E} \int_0^{\eta} \mathbf{b}(\Theta_s) ds \otimes \int_0^{\eta} \mathbf{b}(\Theta_s) ds \otimes \int_0^{\eta} \mathbf{b}(\Theta_s) ds \\ = & \eta^3 \mathbb{E} \mathbf{b}(\tilde{\Theta}_s) \otimes \mathbf{b}(\tilde{\Theta}_s) \otimes \mathbf{b}(\tilde{\Theta}_s) = \mathcal{O}(\eta^3). \end{aligned}$$

The second one is the pure noise part, i.e.,

$$\left(\int_0^{\eta} \sigma(\Theta_s) d\mathbf{W}_s \right) \otimes \left(\int_0^{\eta} \sigma(\Theta_s) d\mathbf{W}_s \right) \otimes \left(\int_0^{\eta} \sigma(\Theta_s) d\mathbf{W}_s \right),$$

and as the odd moments of zero mean Gaussian variables are zero, hence we have

$$\mathbb{E} \left[\left(\int_0^{\eta} \sigma(\Theta_s) d\mathbf{W}_s \right) \otimes \left(\int_0^{\eta} \sigma(\Theta_s) d\mathbf{W}_s \right) \otimes \left(\int_0^{\eta} \sigma(\Theta_s) d\mathbf{W}_s \right) \right] = \mathbf{0},$$

the third and fourth one are both of the mixed part, for the third one

$$\int_0^{\eta} \mathbf{b}(\Theta_s) ds \otimes \int_0^{\eta} \mathbf{b}(\Theta_s) ds \otimes \left(\int_0^{\eta} \sigma(\Theta_s) d\mathbf{W}_s \right),$$

whose expectation is of course zero since the drift part and the noise part is independent, and the fact the odd moments of zero mean Gaussian variables are zero, and for the fourth one

$$\int_0^{\eta} \mathbf{b}(\Theta_s) ds \otimes \left(\int_0^{\eta} \sigma(\Theta_s) d\mathbf{W}_s \right) \otimes \left(\int_0^{\eta} \sigma(\Theta_s) d\mathbf{W}_s \right),$$

we obtain that

$$\begin{aligned} & \mathbb{E} \left[\int_0^{\eta} \mathbf{b}(\Theta_s) ds \otimes \left(\int_0^{\eta} \sigma(\Theta_s) d\mathbf{W}_s \right) \otimes \left(\int_0^{\eta} \sigma(\Theta_s) d\mathbf{W}_s \right) \right] \\ = & \eta \mathbb{E} \mathbf{b}(\tilde{\Theta}_s) \otimes \mathbb{E} \left[\left(\int_0^{\eta} \sigma(\Theta_s) d\mathbf{W}_s \right) \otimes \left(\int_0^{\eta} \sigma(\Theta_s) d\mathbf{W}_s \right) \right] = \mathcal{O}(\eta^3). \end{aligned}$$

As we denote

$$\bar{R}^3(\Theta_0) := \mathbb{E} [(\Theta_{\eta} - \Theta_0) \otimes (\Theta_{\eta} - \Theta_0) \otimes (\Theta_{\eta} - \Theta_0)],$$

then we obtain that

$$\|\text{vec}(\bar{R}^3(\Theta_0))\|_2 \leq \eta^3 C(T^*, \Theta_0).$$

Hence we may apply the mean value theorem to (50) and obtain that

$$\begin{aligned}
|\bar{E}_\eta^2(\Theta_0)| &= \left| \int_0^\eta \int_0^s v \mathcal{L}(\mathcal{L} \langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle)(\tilde{\Theta}_u) dv ds \right. \\
&\quad + \frac{1}{6} \nabla_{\Theta}^3 g(\Theta_0) : \mathbb{E}[(\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0)] \\
&\quad \left. + \frac{1}{24} \mathbb{E} \left[\nabla_{\Theta}^4 g(\tilde{\Theta}_0) : (\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0) \right] \right| \\
&\leq \int_0^\eta \int_0^s v \|g\|_{C^6} C(T^*, \Theta_0) dv ds + \frac{1}{6} \|g\|_{C^6} \eta^3 C(T^*, \Theta_0) \\
&\quad + \frac{1}{24} \|g\|_{C^6} \|\Theta_\eta - \Theta_0\|_2^4 \\
&= \frac{\eta^3}{6} \|g\|_{C^6} C(T^*, \Theta_0) + \frac{1}{6} \|g\|_{C^6} \eta^3 C(T^*, \Theta_0) \\
&\quad + \frac{1}{24} \|g\|_{C^6} \mathbb{E} \left\| \int_0^\eta \mathbf{b}(\Theta_s) ds + \int_0^\eta \boldsymbol{\sigma}(\Theta_s) d\mathbf{W}_s \right\|_2^4 \\
&\leq \eta^3 \|g\|_{C^6} C(T^*, \Theta_0) + \frac{1}{6} \|g\|_{C^6} \eta^3 C(T^*, \Theta_0) \\
&\quad + \frac{4}{24} \|g\|_{C^6} \eta^3 \mathbb{E} \left\| \nabla_{\Theta} L_S(\tilde{\Theta}_0) \right\|_2^2 + \frac{4}{24} \|g\|_{C^6} \mathbb{E} \left\| \int_0^\eta \boldsymbol{\sigma}(\Theta_s) d\mathbf{W}_s \right\|_2^4 \\
&\leq \eta^3 \|g\|_{C^6} C(T^*, \Theta_0) + \frac{1}{6} \|g\|_{C^6} \eta^3 C(T^*, \Theta_0) \\
&\quad + \frac{4}{24} \|g\|_{C^6} \eta^3 \mathbb{E} \left\| \nabla_{\Theta} L_S(\tilde{\Theta}_0) \right\|_2^2 + \frac{C}{24} \|g\|_{C^6} \mathbb{E} \int_0^\eta \|\boldsymbol{\sigma}(\Theta_s)\|_F^4 ds \\
&\leq \eta^3 \|g\|_{C^6} C(T^*, \Theta_0) + \frac{1}{6} \|g\|_{C^6} \eta^3 C(T^*, \Theta_0) \\
&\quad + \eta^3 \|g\|_{C^6} \mathbb{E} \left\| \nabla_{\Theta} L_S(\tilde{\Theta}_0) \right\|_2^2 + C \|g\|_{C^6} \eta \mathbb{E} \left[\eta^2 \left\| \boldsymbol{\Sigma}(\tilde{\Theta}_0) \right\|_F^2 \right] \\
&\leq \eta^3 \|g\|_{C^6} C(T^*, \Theta_0).
\end{aligned}$$

We remark that for the last but third line we apply the Burkholder-Davis-Gundy inequality.

To sum up for now,

$$\begin{aligned}
\mathbb{E}g(\theta_1) - \mathbb{E}g(\theta_0) &= -\eta \left\langle \nabla_{\theta} g(\theta_0), \nabla_{\theta} L_S(\theta) \Big|_{\theta=\theta_0} \right\rangle \\
&\quad + \frac{\eta^2}{2} \nabla_{\theta}^2 g(\theta_0) : \left(\nabla_{\theta} L_S(\theta) \Big|_{\theta=\theta_0} \otimes \nabla_{\theta} L_S(\theta) \Big|_{\theta=\theta_0} + \boldsymbol{\Sigma}(\theta_0) \right) + E_\eta^2(\theta_0),
\end{aligned}$$

and

$$\begin{aligned}
\mathbb{E}g(\Theta_\eta) - \mathbb{E}g(\Theta_0) &= \eta \langle \nabla_{\Theta} g(\Theta_0), \mathbb{E}[\mathbf{b}(\Theta_0)] \rangle + \frac{\eta^2}{2} \mathcal{L} \langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle (\Theta_0) \\
&\quad + \frac{1}{2} \nabla_{\Theta}^2 g(\Theta_0) : \mathbb{E} [(\Theta_\eta - \Theta_0) \otimes (\Theta_\eta - \Theta_0)] + \bar{E}_\eta^2(\Theta_0) \\
&= \eta \langle \nabla_{\Theta} g(\Theta_0), \mathbb{E}[\mathbf{b}(\Theta_0)] \rangle + \frac{\eta^2}{2} \mathcal{L} \langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle (\Theta_0) \\
&\quad + \frac{1}{2} \nabla_{\Theta}^2 g(\Theta_0) : \mathbb{E} \left[\left(\int_0^\eta \mathbf{b}(\Theta_s) ds + \int_0^\eta \boldsymbol{\sigma}(\Theta_s) d\mathbf{W}_s \right) \right. \\
&\quad \quad \left. \otimes \left(\int_0^\eta \mathbf{b}(\Theta_s) ds + \int_0^\eta \boldsymbol{\sigma}(\Theta_s) d\mathbf{W}_s \right) \right] + \bar{E}_\eta^2(\Theta_0) \\
&= \eta \langle \nabla_{\Theta} g(\Theta_0), \mathbb{E}[\mathbf{b}(\Theta_0)] \rangle + \frac{\eta^2}{2} \mathcal{L} \langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle (\Theta_0) \\
&\quad + \frac{1}{2} \nabla_{\Theta}^2 g(\Theta_0) : \mathbb{E} \left[\int_0^\eta \mathbf{b}(\Theta_s) ds \otimes \int_0^\eta \mathbf{b}(\Theta_s) ds \right] \\
&\quad + \frac{1}{2} \nabla_{\Theta}^2 g(\Theta_0) : \mathbb{E} \left[\int_0^\eta \boldsymbol{\sigma}(\Theta_s) d\mathbf{W}_s \otimes \int_0^\eta \boldsymbol{\sigma}(\Theta_s) d\mathbf{W}_s \right] + \bar{E}_\eta^2(\Theta_0) \\
&= \eta \langle \nabla_{\Theta} g(\Theta_0), \mathbb{E}[\mathbf{b}(\Theta_0)] \rangle + \frac{\eta^2}{2} \mathcal{L} \langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle (\Theta_0) \\
&\quad + \frac{1}{2} \nabla_{\Theta}^2 g(\Theta_0) : \mathbb{E} \left[\int_0^\eta \int_0^\eta \mathbf{b}(\Theta_s) \otimes \mathbf{b}(\Theta_u) ds du \right] \\
&\quad + \frac{1}{2} \nabla_{\Theta}^2 g(\Theta_0) : \mathbb{E} \left[\int_0^\eta \boldsymbol{\sigma}(\Theta_s) d\mathbf{W}_s \otimes \int_0^\eta \boldsymbol{\sigma}(\Theta_s) d\mathbf{W}_s \right] + \bar{E}_\eta^2(\Theta_0),
\end{aligned}$$

we observe that

$$\begin{aligned}
&\frac{1}{2} \nabla_{\Theta}^2 g(\Theta_0) : \mathbb{E} \left[\int_0^\eta \boldsymbol{\sigma}(\Theta_s) d\mathbf{W}_s \otimes \int_0^\eta \boldsymbol{\sigma}(\Theta_s) d\mathbf{W}_s \right] \\
&= \mathbb{E} \left[\int_0^\eta \frac{1}{2} \nabla_{\Theta}^2 g(\Theta_0) : \boldsymbol{\sigma} \boldsymbol{\sigma}^\top(\Theta_s) ds \right] \\
&= \frac{\eta}{2} \mathbb{E} \left[\int_0^\eta \nabla_{\Theta}^2 g(\Theta_0) : \boldsymbol{\Sigma}(\Theta_s) ds \right],
\end{aligned}$$

thus

$$\begin{aligned}
\mathbb{E}g(\Theta_\eta) - \mathbb{E}g(\Theta_0) &= \eta \langle \nabla_{\Theta} g(\Theta_0), \mathbb{E}[\mathbf{b}(\Theta_0)] \rangle + \frac{\eta^2}{2} \mathcal{L} \langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle (\Theta_0) \\
&\quad + \frac{1}{2} \nabla_{\Theta}^2 g(\Theta_0) : \mathbb{E} \left[\int_0^\eta \int_0^\eta \mathbf{b}(\Theta_s) \otimes \mathbf{b}(\Theta_u) ds du \right] \\
&\quad + \frac{\eta}{2} \mathbb{E} \left[\int_0^\eta \nabla_{\Theta}^2 g(\Theta_0) : \boldsymbol{\Sigma}(\Theta_s) ds \right] + \bar{E}_\eta^2(\Theta_0).
\end{aligned}$$

Since

$$\begin{aligned}
&\nabla_{\Theta}^2 g(\Theta_0) : \mathbb{E} [\mathbf{b}(\Theta_s) \otimes \mathbf{b}(\Theta_u)] \\
&= \nabla_{\Theta}^2 g(\Theta_0) : \mathbb{E} [\mathbf{b}(\Theta_s) \otimes \mathbf{b}(\Theta_0)] + \int_0^u \mathcal{L} (\nabla_{\Theta}^2 g(\Theta_0) : \mathbf{b}(\Theta_s) \otimes \mathbf{b}(\Theta_v)) dv \\
&= \nabla_{\Theta}^2 g(\Theta_0) : \mathbb{E} [\mathbf{b}(\Theta_0) \otimes \mathbf{b}(\Theta_0)] + \int_0^s \mathcal{L} (\nabla_{\Theta}^2 g(\Theta_0) : \mathbf{b}(\Theta_w) \otimes \mathbf{b}(\Theta_0)) dw \\
&\quad + \int_0^u \mathcal{L} (\nabla_{\Theta}^2 g(\Theta_0) : \mathbf{b}(\Theta_s) \otimes \mathbf{b}(\Theta_v)) dv,
\end{aligned}$$

and since

$$\begin{aligned} & \nabla_{\Theta}^2 g(\Theta_0) : \mathbb{E}[\Sigma(\Theta_s)] \\ &= \nabla_{\Theta}^2 g(\Theta_0) : \mathbb{E}[\Sigma(\Theta_0)] + \int_0^s \mathcal{L}(\nabla_{\Theta}^2 g(\Theta_0) : \Sigma(\Theta_s)) dv, \end{aligned}$$

we are one step away to finish our proof,

$$\begin{aligned} \mathbb{E}g(\Theta_\eta) - \mathbb{E}g(\Theta_0) &= \eta \langle \nabla_{\Theta} g(\Theta_0), \mathbb{E}[\mathbf{b}(\Theta_0)] \rangle + \frac{\eta^2}{2} \mathcal{L} \langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle (\Theta_0) \\ &\quad + \frac{1}{2} \nabla_{\Theta}^2 g(\Theta_0) : \mathbb{E} \left[\int_0^\eta \int_0^\eta \mathbf{b}(\Theta_0) \otimes \mathbf{b}(\Theta_0) ds du \right] \\ &\quad + \frac{\eta}{2} \mathbb{E} \left[\int_0^\eta \nabla_{\Theta}^2 g(\Theta_0) : \Sigma(\Theta_0) ds \right] + \bar{E}_\eta^2(\Theta_0), \end{aligned}$$

where we misuse our notations for $\bar{E}_\eta^2(\Theta_0)$, and the term

$$\begin{aligned} & \int_0^\eta \int_0^\eta \int_0^s \mathcal{L}(\nabla_{\Theta}^2 g(\Theta_0) : \mathbf{b}(\Theta_w) \otimes \mathbf{b}(\Theta_0)) dw ds du \\ &+ \int_0^\eta \int_0^\eta \int_0^u \mathcal{L}(\nabla_{\Theta}^2 g(\Theta_0) : \mathbf{b}(\Theta_s) \otimes \mathbf{b}(\Theta_v)) dv ds du \\ &+ \int_0^\eta \int_0^s \mathcal{L}(\nabla_{\Theta}^2 g(\Theta_0) : \Sigma(\Theta_s)) dv ds, \end{aligned}$$

is included, and $\bar{E}_\eta^2(\Theta_0)$ is still of order $\mathcal{O}(\eta^3)$ by similar reasoning and we omit its demonstration. Thus

$$\begin{aligned} \mathbb{E}g(\Theta_\eta) - \mathbb{E}g(\Theta_0) &= \eta \langle \nabla_{\Theta} g(\Theta_0), \mathbb{E}[\mathbf{b}(\Theta_0)] \rangle + \frac{\eta^2}{2} \langle \mathbf{b}(\Theta_0), \nabla_{\Theta} \langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle (\Theta_0) \rangle \\ &\quad + \frac{\eta^3}{2} \Sigma(\Theta_0) : \nabla_{\Theta}^2 \langle \nabla_{\Theta} g(\Theta_0), \mathbf{b} \rangle (\Theta_0) \\ &\quad + \frac{\eta^2}{2} \nabla_{\Theta}^2 g(\Theta_0) : \mathbb{E}[\mathbf{b}(\Theta_0) \otimes \mathbf{b}(\Theta_0)] \\ &\quad + \frac{\eta^2}{2} \mathbb{E}[\nabla_{\Theta}^2 g(\Theta_0) : \Sigma(\Theta_0)] + \bar{E}_\eta^2(\Theta_0), \end{aligned}$$

and recall that since we choose

$$\begin{aligned} \mathbf{b}(\Theta) &= -\nabla_{\Theta} \left(L_S(\Theta) + \frac{\eta}{4} \|\nabla_{\Theta} L_S(\Theta)\|_2^2 \right), \\ \sigma(\Theta) &= \sqrt{\eta} (\Sigma(\Theta))^{\frac{1}{2}}, \end{aligned}$$

then

$$\begin{aligned} \mathbb{E}g(\Theta_\eta) - \mathbb{E}g(\Theta_0) &= -\eta \langle \nabla_{\Theta} g(\Theta_0), \nabla_{\Theta} (L_S(\Theta)) |_{\Theta=\Theta_0} \rangle \\ &\quad - \frac{\eta^2}{4} \left\langle \nabla_{\Theta} g(\Theta_0), \nabla_{\Theta} \left(\|\nabla_{\Theta} L_S(\Theta)\|_2^2 \right) |_{\Theta=\Theta_0} \right\rangle \\ &\quad + \frac{\eta^2}{2} \langle \nabla_{\Theta} (L_S(\Theta)) |_{\Theta=\Theta_0}, \nabla_{\Theta} \langle \nabla_{\Theta} g(\Theta_0), \nabla_{\Theta} (L_S(\Theta)) \rangle |_{\Theta=\Theta_0} \rangle \\ &\quad + \frac{\eta^2}{2} \nabla_{\Theta}^2 g(\Theta_0) : (\nabla_{\Theta} (L_S(\Theta)) |_{\Theta=\Theta_0} \otimes \nabla_{\Theta} (L_S(\Theta)) |_{\Theta=\Theta_0}) \\ &\quad + \frac{\eta^2}{2} \nabla_{\Theta}^2 g(\Theta_0) : \Sigma(\Theta_0) + \bar{E}_\eta^2(\Theta_0) \\ &= -\eta \langle \nabla_{\Theta} g(\Theta_0), \nabla_{\Theta} (L_S(\Theta)) |_{\Theta=\Theta_0} \rangle \\ &\quad + \frac{\eta^2}{2} \nabla_{\Theta}^2 g(\Theta_0) : (\nabla_{\Theta} (L_S(\Theta)) |_{\Theta=\Theta_0} \otimes \nabla_{\Theta} (L_S(\Theta)) |_{\Theta=\Theta_0}) \\ &\quad + \frac{\eta^2}{2} \nabla_{\Theta}^2 g(\Theta_0) : \Sigma(\Theta_0) + \bar{E}_\eta^2(\Theta_0), \end{aligned}$$

thus, we have

$$\begin{aligned}
|\mathbb{E}g(\boldsymbol{\theta}_1) - \mathbb{E}g(\boldsymbol{\Theta}_\eta)| &= \left| \mathbb{E}g(\boldsymbol{\theta}_0) - \eta \left\langle \nabla_{\boldsymbol{\theta}} g(\boldsymbol{\theta}_0), \nabla_{\boldsymbol{\theta}} L_S(\boldsymbol{\theta}) \Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0} \right\rangle \right. \\
&\quad + \frac{\eta^2}{2} \nabla_{\boldsymbol{\theta}}^2 g(\boldsymbol{\theta}_0) : \left(\nabla_{\boldsymbol{\theta}} L_S(\boldsymbol{\theta}) \Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0} \otimes \nabla_{\boldsymbol{\theta}} L_S(\boldsymbol{\theta}) \Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0} + \Sigma(\boldsymbol{\theta}_0) \right) \\
&\quad + E_\eta^2(\boldsymbol{\theta}_0) \\
&\quad - \mathbb{E}g(\boldsymbol{\Theta}_0) + \eta \left\langle \nabla_{\boldsymbol{\Theta}} g(\boldsymbol{\Theta}_0), \nabla_{\boldsymbol{\Theta}} (L_S(\boldsymbol{\Theta})) \Big|_{\boldsymbol{\Theta}=\boldsymbol{\Theta}_0} \right\rangle \\
&\quad - \frac{\eta^2}{2} \nabla_{\boldsymbol{\Theta}}^2 g(\boldsymbol{\Theta}_0) : \left(\nabla_{\boldsymbol{\Theta}} (L_S(\boldsymbol{\Theta})) \Big|_{\boldsymbol{\Theta}=\boldsymbol{\Theta}_0} \otimes \nabla_{\boldsymbol{\Theta}} (L_S(\boldsymbol{\Theta})) \Big|_{\boldsymbol{\Theta}=\boldsymbol{\Theta}_0} \right) \\
&\quad - \frac{\eta^2}{2} \nabla_{\boldsymbol{\Theta}}^2 g(\boldsymbol{\Theta}_0) : \Sigma(\boldsymbol{\Theta}_0) + \bar{E}_\eta^2(\boldsymbol{\Theta}_0) \Big| \\
&\leq |E_\eta^2(\boldsymbol{\theta}_0)| + |\bar{E}_\eta^2(\boldsymbol{\Theta}_0)| \\
&\leq \eta^3 \|g\|_{C^6} C(T^*, \boldsymbol{\theta}_0, \eta_0) + \eta^3 \|g\|_{C^6} C(T^*, \boldsymbol{\Theta}_0) \\
&= \mathcal{O}(\eta^3).
\end{aligned}$$

For the N -th step iteration, since

$$|\mathbb{E}g(\boldsymbol{\theta}_N) - \mathbb{E}g(\boldsymbol{\Theta}_{\eta N})| = |\mathcal{P}_\eta^N g(\boldsymbol{\theta}_0) - \mathcal{P}_{\eta N} g(\boldsymbol{\Theta}_0)|,$$

and the RHS of the above equation can be written into a telescoping sum as

$$\mathcal{P}_\eta^N g(\boldsymbol{\theta}_0) - \mathcal{P}_{\eta N} g(\boldsymbol{\Theta}_0) = \sum_{l=1}^N (\mathcal{P}_\eta^{N-l+1} \circ \mathcal{P}_{(l-1)\eta} g(\boldsymbol{\theta}_0) - \mathcal{P}_\eta^{N-l} \circ \mathcal{P}_{l\eta} g(\boldsymbol{\Theta}_0)),$$

hence by application of Proposition 1, we obtain that

$$\begin{aligned}
|\mathbb{E}g(\boldsymbol{\theta}_N) - \mathbb{E}g(\boldsymbol{\Theta}_{\eta N})| &\leq \sum_{l=1}^N |\mathcal{P}_\eta^{N-l+1} \circ \mathcal{P}_{(l-1)\eta} g(\boldsymbol{\theta}_0) - \mathcal{P}_\eta^{N-l} \circ \mathcal{P}_{l\eta} g(\boldsymbol{\Theta}_0)| \\
&\leq \sum_{l=1}^N |\mathcal{P}_\eta^{N-l} \circ (\mathcal{P}_\eta^1 \circ \mathcal{P}_{(l-1)\eta} - \mathcal{P}_\eta \circ \mathcal{P}_{(l-1)\eta}) g(\boldsymbol{\Theta}_0)|,
\end{aligned}$$

since $(\mathcal{P}_\eta^1 \circ \mathcal{P}_{(l-1)\eta} - \mathcal{P}_\eta \circ \mathcal{P}_{(l-1)\eta}) g(\boldsymbol{\Theta}_0)$ can be regarded as $\mathcal{L}^1(\mathbb{R}^D)$ if we choose measure μ to be the delta measure concentrated on $\boldsymbol{\Theta}_0$. i.e.,

$$\mu := \delta_{\boldsymbol{\Theta}_0},$$

hence by the contraction property of Markov operators, we obtain further that

$$\begin{aligned}
|\mathbb{E}g(\boldsymbol{\theta}_N) - \mathbb{E}g(\boldsymbol{\Theta}_{\eta N})| &\leq \sum_{l=1}^N |(\mathcal{P}_\eta^1 \circ \mathcal{P}_{(l-1)\eta} - \mathcal{P}_\eta \circ \mathcal{P}_{(l-1)\eta}) g(\boldsymbol{\Theta}_0)| \\
&\leq \sum_{l=1}^N |\mathcal{P}_\eta^1 g(\boldsymbol{\Theta}_{(l-1)\eta}) - \mathcal{P}_\eta g(\boldsymbol{\Theta}_{(l-1)\eta})|.
\end{aligned}$$

By taking expectation conditioned on $\boldsymbol{\Theta}_{(l-1)\eta}$, then similar to the relation (46), the following holds

$$\begin{aligned}
|\mathcal{P}_\eta^1 g(\boldsymbol{\Theta}_{(l-1)\eta}) - \mathcal{P}_\eta g(\boldsymbol{\Theta}_{(l-1)\eta})| &= \mathbb{E} \left[\left| \mathbb{E}g(\boldsymbol{\theta}_l) - \mathbb{E}g(\boldsymbol{\Theta}_{\eta l}) \Big|_{\boldsymbol{\Theta}_{(l-1)\eta}} \right| \right] \\
&\leq \mathbb{E} |E_\eta^2(\boldsymbol{\Theta}_{(l-1)\eta})| + \mathbb{E} |\bar{E}_\eta^2(\boldsymbol{\Theta}_{(l-1)\eta})| \\
&\leq \eta^3 \|g\|_{C^6} C(T^*, \boldsymbol{\theta}_0, \eta_0) + \eta^3 \|g\|_{C^6} C(T^*, \boldsymbol{\Theta}_0) \\
&= \mathcal{O}(\eta^3).
\end{aligned}$$

We remark that the last line of the above relation is essentially based on Assumption 2, since $\mathbb{E} |E_\eta^2(\boldsymbol{\Theta}_{(l-1)\eta})|$ and $\mathbb{E} |\bar{E}_\eta^2(\boldsymbol{\Theta}_{(l-1)\eta})|$ can be bounded above by the second, fourth and sixth moments of the solution to SDE (32), hence we may apply dominated convergence theorem to obtain the last line of the above relation.

To sum up, as

$$|\mathcal{P}_\eta^N g(\boldsymbol{\theta}_0) - \mathcal{P}_{\eta N} g(\boldsymbol{\Theta}_0)| \leq \sum_{l=1}^N |\mathcal{P}_\eta^{N-l+1} \circ \mathcal{P}_{(l-1)\eta} g(\boldsymbol{\theta}_0) - \mathcal{P}_\eta^{N-l} \circ \mathcal{P}_{l\eta} g(\boldsymbol{\Theta}_0)| = N\mathcal{O}(\eta^3),$$

hence for $N = N_{T,\eta}$,

$$|\mathcal{P}_\eta^N g(\boldsymbol{\theta}_0) - \mathcal{P}_{\eta N} g(\boldsymbol{\Theta}_0)| = N\mathcal{O}(\eta^3) = N\eta\mathcal{O}(\eta) \leq T\mathcal{O}(\eta^2) = \mathcal{O}(\eta^2).$$

□

I VALIDATION FOR ASSUMPTION 1

In this section, we endeavor to demonstrate the validity of Assumption 1. We begin this section by making some estimates on the modified loss L_S and covariance Σ .

I.1 ESTIMATES ON MODIFIED LOSS AND COVARIANCE

For the modified loss, recall that $\theta = \text{vec}(\{\mathbf{q}_r\}_{r=1}^m) = \text{vec}(\{(a_r, \mathbf{w}_r)\}_{r=1}^m)$, as we have

$$\nabla_{\mathbf{q}_k} L_S(\Theta) = \frac{1}{n} \sum_{i=1}^n e_i \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)) + \frac{1-p}{np} \sum_{i=1}^n a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)),$$

and under the usual convention that for all $i \in [n]$,

$$\frac{1}{c} \leq \|\mathbf{x}_i\|_2, \quad |y_i| \leq c,$$

where c is some universal constant, and that $\sigma(0) = 0$, we obtain that

$$\begin{aligned} |e_i| &= \left| \sum_{r=1}^m a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i) - y_i \right| \\ &\leq 1 + \sum_{r=1}^m |a_r| \|\mathbf{w}_r\|_2 \\ &\leq 1 + \frac{1}{2} \sum_{r=1}^m (|a_r|^2 + \|\mathbf{w}_r\|_2^2) \\ &\leq 1 + \|\Theta\|_2^2, \end{aligned}$$

hence

$$\|\nabla_{\mathbf{q}_k} L_S(\Theta)\|_2 \leq (1 + \|\Theta\|_2^2) \|\mathbf{q}_k\|_2 + \frac{1-p}{p} \|\mathbf{q}_k\|_2^3,$$

thus we have

$$\begin{aligned} \|\nabla_{\Theta} L_S(\Theta)\|_2 &\leq (1 + \|\Theta\|_2^2) \|\Theta\|_2 + \frac{1-p}{p} \|\Theta\|_2^3 \\ &\leq C_p (1 + \|\Theta\|_2^3). \end{aligned}$$

Moreover, since

$$\begin{aligned} \nabla_{\Theta}^2 L_S(\Theta) &= \frac{1}{n} \sum_{i=1}^n (\nabla_{\Theta} e_i \otimes \nabla_{\Theta} e_i + e_i \nabla_{\Theta}^2 e_i) \\ &\quad + \frac{1-p}{np} \sum_{i=1}^n \text{diag} \{ \nabla_{\mathbf{q}_k}^2 (a_k^2 \sigma(\mathbf{w}_k^\top \mathbf{x}_i)^2) \}, \end{aligned}$$

as we denote only for now \times as matrix multiplication,

$$\begin{aligned} &\nabla_{\Theta}^2 L_S(\Theta) \nabla_{\Theta} L_S(\Theta) \\ &= \left(\frac{1}{n} \sum_{i=1}^n (\nabla_{\Theta} e_i \otimes \nabla_{\Theta} e_i + e_i \nabla_{\Theta}^2 e_i) + \frac{1-p}{np} \sum_{i=1}^n \text{diag} \{ \nabla_{\mathbf{q}_k}^2 (a_k^2 \sigma(\mathbf{w}_k^\top \mathbf{x}_i)^2) \} \right) \\ &\quad \times \left(\frac{1}{n} \sum_{i=1}^n e_i \nabla_{\Theta} e_i + \frac{1-p}{np} \sum_{i=1}^n \nabla_{\Theta} (a_k^2 \sigma(\mathbf{w}_k^\top \mathbf{x}_i)^2) \right), \end{aligned}$$

then the components in $\nabla_{\Theta}^2 L_S(\Theta) \nabla_{\Theta} L_S(\Theta)$ can be categorized into six different types: Firstly,

$$\begin{aligned} &\|(\nabla_{\Theta} e_i \otimes \nabla_{\Theta} e_i) e_j \nabla_{\Theta} e_j\|_2 \\ &\leq |e_j| \|\nabla_{\Theta} e_i\|_2^2 \|\nabla_{\Theta} e_j\|_2 \\ &\leq (1 + \|\Theta\|_2^2) \|\Theta\|_2^3 \\ &\leq (1 + \|\Theta\|_2^5). \end{aligned}$$

Secondly,

$$\begin{aligned}
& \| (e_i \nabla_{\Theta}^2 e_i) e_j \nabla_{\Theta} e_j \|_2 \\
& \leq \left(1 + \|\Theta\|_2^2\right)^2 \|\nabla_{\Theta}^2 e_i\|_{2 \rightarrow 2} \|\nabla_{\Theta} e_j\|_2 \\
& \leq \left(1 + \|\Theta\|_2^4\right) \|\Theta\|_2^2 \\
& \leq \left(1 + \|\Theta\|_2^6\right).
\end{aligned}$$

Thirdly,

$$\begin{aligned}
& \| (\text{diag} \{ \nabla_{q_k}^2 (a_k^2 \sigma(\mathbf{w}_k^T \mathbf{x}_i)^2) \}) e_j \nabla_{\Theta} e_j \|_2 \\
& \leq \left(1 + \|\Theta\|_2^2\right) \|\text{diag} \{ \nabla_{q_k}^2 (a_k^2 \sigma(\mathbf{w}_k^T \mathbf{x}_i)^2) \}\|_{2 \rightarrow 2} \|\Theta\|_2 \\
& \leq \left(1 + \|\Theta\|_2^2\right) \left(1 + \|\Theta\|_2^3\right) \|\Theta\|_2 \\
& \leq \left(1 + \|\Theta\|_2^6\right).
\end{aligned}$$

Fourthly,

$$\begin{aligned}
& \| (\nabla_{\Theta} e_i \otimes \nabla_{\Theta} e_i) \nabla_{\Theta} (a_k^2 \sigma(\mathbf{w}_k^T \mathbf{x}_j)^2) \|_2 \\
& \leq \|\nabla_{\Theta} e_i\|_2^2 \|\Theta\|_2^3 \\
& \leq \left(1 + \|\Theta\|_2^5\right).
\end{aligned}$$

Fifthly,

$$\begin{aligned}
& \| (e_i \nabla_{\Theta}^2 e_i) \nabla_{\Theta} (a_k^2 \sigma(\mathbf{w}_k^T \mathbf{x}_j)^2) \|_2 \\
& \leq \left(1 + \|\Theta\|_2^2\right) \|\nabla_{\Theta}^2 e_i\|_{2 \rightarrow 2} \|\Theta\|_2^3 \\
& \leq \left(1 + \|\Theta\|_2^2\right) \|\Theta\|_2^4 \\
& \leq \left(1 + \|\Theta\|_2^6\right).
\end{aligned}$$

Finally,

$$\begin{aligned}
& \| (\text{diag} \{ \nabla_{q_k}^2 (a_k^2 \sigma(\mathbf{w}_k^T \mathbf{x}_i)^2) \}) \nabla_{\Theta} (a_k^2 \sigma(\mathbf{w}_k^T \mathbf{x}_j)^2) \|_2 \\
& \leq \|\text{diag} \{ \nabla_{q_k}^2 (a_k^2 \sigma(\mathbf{w}_k^T \mathbf{x}_i)^2) \}\|_{2 \rightarrow 2} \|\Theta\|_2^3 \\
& \leq \left(1 + \|\Theta\|_2^3\right) \|\Theta\|_2^3 \\
& \leq \left(1 + \|\Theta\|_2^6\right).
\end{aligned}$$

To sum up, for the drift term $\mathbf{b}(\Theta)$, regardless of the choice of first order or second order accuracy, we obtain that

$$\|\mathbf{b}(\Theta)\|_2 \leq 1 + \|\Theta\|_2^6.$$

As for the covariance Σ , recall that $\theta = \text{vec}(\{q_r\}_{r=1}^m) = \text{vec}(\{(a_r, \mathbf{w}_r)\}_{r=1}^m)$, then we obtain that the covariance Σ reads

$$\Sigma = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} & \cdots & \Sigma_{1m} \\ \Sigma_{21} & \Sigma_{22} & \cdots & \Sigma_{2m} \\ \vdots & \vdots & \vdots & \vdots \\ \Sigma_{m1} & \Sigma_{m2} & \cdots & \Sigma_{mm} \end{bmatrix}.$$

For each $k \in [m]$, we obtain that

$$\begin{aligned} \Sigma_{kk}(\Theta) &= \left(\frac{1}{p} - 1\right) \left(\frac{1}{n} \sum_{i=1}^n \left(e_{i,\setminus k} + \frac{1}{p} a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \right) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)) \right) \\ &\quad \otimes \left(\frac{1}{n} \sum_{i=1}^n \left(e_{i,\setminus k} + \frac{1}{p} a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \right) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)) \right) \\ &\quad + \left(\frac{1}{p^2} - \frac{1}{p}\right) \sum_{l=1, l \neq k}^m \left(\frac{1}{n} \sum_{i=1}^n a_l \sigma(\mathbf{w}_l^\top \mathbf{x}_i) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)) \right) \\ &\quad \otimes \left(\frac{1}{n} \sum_{i=1}^n a_l \sigma(\mathbf{w}_l^\top \mathbf{x}_i) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)) \right), \end{aligned}$$

and for each $k, r \in [m]$ with $k \neq r$,

$$\begin{aligned} \Sigma_{kr}(\Theta) &= \left(\frac{1}{p} - 1\right) \left(\frac{1}{n} \sum_{i=1}^n \left(e_{i,\setminus k, \setminus r} + \frac{1}{p} a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) + \frac{1}{p} a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i) \right) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)) \right) \\ &\quad \otimes \left(\frac{1}{n} \sum_{i=1}^n a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i)) \right) \\ &\quad + \left(\frac{1}{p} - 1\right) \left(\frac{1}{n} \sum_{i=1}^n a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)) \right) \\ &\quad \otimes \left(\frac{1}{n} \sum_{i=1}^n \left(e_{i,\setminus k, \setminus r} + a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) + \frac{1}{p} a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i) \right) \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i)) \right), \end{aligned}$$

hence we obtain that

$$\begin{aligned} \|\Sigma_{kk}(\Theta)\|_{\mathbb{F}}^2 &\leq C_p \left(\left| e_{i,\setminus k} + \frac{1}{p} a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \right|^2 + \sum_{l=1, l \neq k}^m a_l^2 \sigma(\mathbf{w}_l^\top \mathbf{x}_i)^2 \right) \|\nabla_{\Theta} e_i\|_2^2 \\ &\leq C_p (1 + \|\Theta\|_2^2) \|\Theta\|_2^2 \\ &\leq (1 + \|\Theta\|_2^6), \end{aligned}$$

and by similar reasoning

$$\|\Sigma_{kr}(\Theta)\|_{\mathbb{F}}^2 \leq (1 + \|\Theta\|_2^6).$$

1.2 EXISTENCE, UNIQUENESS AND MOMENT ESTIMATES OF THE SOLUTION TO SDE

Existence of the solution to SDE (32) is proved by a truncation procedure: For each $M \geq 1$, define the truncation function

$$\mathbf{b}_M(\Theta) := \begin{cases} \mathbf{b}(\Theta) & \text{if } \|\Theta\|_2 \leq M, \\ \mathbf{b}(M \frac{\Theta}{\|\Theta\|_2}) & \text{if } \|\Theta\|_2 > M. \end{cases}$$

We also perform similar truncation to $\sigma(\Theta)$ and obtain its truncation $\sigma_M(\Theta)$. Then \mathbf{b}_M and σ_M satisfy the Lipschitz condition and the linear growth condition, hence by application of the classical results (Oksendal, 2013, Theorem 5.2.1) in SDE, there exists a unique solution $\Theta_M(\cdot)$ to the truncated SDE

$$d\Theta_t = \mathbf{b}_M(\Theta_t) dt + \sigma_M(\Theta_t) d\mathbf{W}_t, \quad \Theta_0 = \Theta(0). \quad (51)$$

We may choose M large enough, such that

$$\|\Theta_0\|_2 < M,$$

and the solution to SDE (32) coincides with the solution to SDE (51) at least for a period of time $T^* > 0$ since $\|\Theta_0\|_2 < M$. We remark that T^* is the desired time in Assumption 2. We also remark that not only for any time $t \in [0, T^*]$, the second, fourth and sixth moments of the solution to SDE

(32) are uniformly bounded with respect to time t , but also that for any time $t \in [0, T^*]$, all moments of the solution to SDE (32) are uniformly bounded with respect to time t .

At this point, it is important to discuss that we prove is that for fixed time T , we can take the learning rate $\eta > 0$ small enough so that the SME is a good approximation of the distribution of the dropout iterates. What we did not prove is that for fixed η , the approximations hold for arbitrary time T . In particular, it is not hard to construct systems where for fixed η , both the SME and the asymptotic expansion fails when time T is large enough.

I.3 MOMENT ESTIMATES OF THE DROPOUT ITERATION

Recall that the dropout iteration reads

$$\boldsymbol{\theta}_N = \boldsymbol{\theta}_{N-1} - \eta \nabla_{\boldsymbol{\theta}} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N),$$

then we obtain that

$$\mathbb{E} \|\boldsymbol{\theta}_N\|_2^{2l} = \mathbb{E} \|\boldsymbol{\theta}_{N-1}\|_2^{2l} - 2l\eta \mathbb{E} \left[\|\boldsymbol{\theta}_{N-1}\|_2^{2l-2} \left\langle \boldsymbol{\theta}_{N-1}, \nabla_{\boldsymbol{\theta}} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N) \right\rangle \right] + \mathcal{O}(\eta^2),$$

then for learning rate η small enough, we observe that $\{\mathbb{E} \|\boldsymbol{\theta}_N\|_2^{2l}\}_{N \geq 0}$ follows close to the trajectory of an ordinary differential equation (ODE). Moreover, from the estimates obtained in Section I.1,

$$\begin{aligned} & \|\boldsymbol{\theta}_{N-1}\|_2^{2l-2} \left\langle \boldsymbol{\theta}_{N-1}, \nabla_{\boldsymbol{\theta}} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N) \right\rangle \\ & \leq \|\boldsymbol{\theta}_{N-1}\|_2^{2l-1} \left\| \nabla_{\boldsymbol{\theta}} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N) \right\|_2 \\ & = \|\boldsymbol{\theta}_{N-1}\|_2^{2l-1} \|e_i^N\| \|\nabla_{\boldsymbol{\theta}} e_i^N\|_2 \\ & \leq \|\boldsymbol{\theta}_{N-1}\|_2^{2l-1} C_p (1 + \|\boldsymbol{\theta}_{N-1}\|_2^2) \|\boldsymbol{\theta}_{N-1}\|_2 \\ & \leq C_p (1 + \|\boldsymbol{\theta}_{N-1}\|_2^{2l+2}), \end{aligned}$$

we remark that as the above estimates hold almost surely, then for learning rate η small enough, we may apply Gronwall inequality to $\{\mathbb{E} \|\boldsymbol{\theta}_N\|_2^{2l}\}_{N \geq 0}$ and shows that for some N^* , all moments of the dropout iterations are uniformly bounded with respect to N , since for the ODE

$$\frac{du}{dt} = 1 + u^{1+\lambda}, \quad u_0 := u(0), \quad (52)$$

with $\lambda > 0$. There exists time $T^* > 0$, such that for any time $t \in [0, T^*]$, its solution $\{u_t\}_{t \geq 0}$ is uniformly bounded with respect to time t . And since for small enough learning rate, all moments of the dropout iterations $\{\mathbb{E} \|\boldsymbol{\theta}_N\|_2^{2l}\}_{N \geq 0}$ follows close to the trajectory of ODEs of (52) type, hence all these moments are also uniformly bounded with respect to N .

J EXPLICIT FORM AND DERIVATION ON THE COVARIANCE

In this section, we present the expression for Σ . As $\boldsymbol{\theta} = \text{vec}(\{\mathbf{q}_r\}_{r=1}^m) = \text{vec}(\{(a_r, \mathbf{w}_r)\}_{r=1}^m)$, then covariance of $\nabla_{\boldsymbol{\theta}} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N)$ equals to $\Sigma(\boldsymbol{\theta}_{N-1})$. We denote

$$\Sigma_{kr}(\boldsymbol{\theta}_{N-1}) := \text{Cov}\left(\nabla_{\mathbf{q}_k} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N), \nabla_{\mathbf{q}_r} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N)\right),$$

then

$$\Sigma = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} & \cdots & \Sigma_{1m} \\ \Sigma_{21} & \Sigma_{22} & \cdots & \Sigma_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ \Sigma_{m1} & \Sigma_{m2} & \cdots & \Sigma_{mm} \end{bmatrix}.$$

For each $k \in [m]$, we obtain that

$$\begin{aligned} \Sigma_{kk}(\boldsymbol{\theta}_{N-1}) &= \text{Cov}\left(\nabla_{\mathbf{q}_k} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N), \nabla_{\mathbf{q}_k} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N)\right) \\ &= \left(\frac{1}{p} - 1\right) \left(\frac{1}{n} \sum_{i=1}^n \left(e_{i,\setminus k} + \frac{1}{p} a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)\right) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i))\right) \\ &\quad \otimes \left(\frac{1}{n} \sum_{i=1}^n \left(e_{i,\setminus k} + \frac{1}{p} a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)\right) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i))\right) \\ &\quad + \left(\frac{1}{p^2} - \frac{1}{p}\right) \sum_{k'=1, k' \neq k}^m \left(\frac{1}{n} \sum_{i=1}^n a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i))\right) \\ &\quad \otimes \left(\frac{1}{n} \sum_{i=1}^n a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i))\right), \end{aligned}$$

where $e_{i,\setminus k} := e_{i,\setminus k}(\boldsymbol{\theta}) := \sum_{l=1, l \neq k}^m a_l \sigma(\mathbf{w}_l^\top \mathbf{x}_i) - y_i$, and for each $k, r \in [m]$ with $k \neq r$,

$$\begin{aligned} \Sigma_{kr}(\boldsymbol{\theta}_{N-1}) &= \text{Cov}\left(\nabla_{\mathbf{q}_k} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N), \nabla_{\mathbf{q}_r} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N)\right) \\ &= \left(\frac{1}{p} - 1\right) \sum_{k'=1, k' \neq k, k' \neq r}^m \left(\frac{1}{n} \sum_{i=1}^n a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i))\right) \\ &\quad \otimes \left(\frac{1}{n} \sum_{i=1}^n a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i))\right) \\ &\quad + \left(\frac{1}{p} - 1\right) \left(\frac{1}{n} \sum_{i=1}^n \left(e_{i,\setminus k, \setminus r} + \frac{1}{p} a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) + \frac{1}{p} a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i)\right) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i))\right) \\ &\quad \otimes \left(\frac{1}{n} \sum_{i=1}^n a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i))\right) \\ &\quad + \left(\frac{1}{p} - 1\right) \left(\frac{1}{n} \sum_{i=1}^n a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i))\right) \\ &\quad \otimes \left(\frac{1}{n} \sum_{i=1}^n \left(e_{i,\setminus k, \setminus r} + a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) + \frac{1}{p} a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i)\right) \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i))\right), \end{aligned}$$

where $e_{i,\setminus k, \setminus r} := e_{i,\setminus k, \setminus r}(\boldsymbol{\theta}) := \sum_{l=1, l \neq k, l \neq r}^m a_l \sigma(\mathbf{w}_l^\top \mathbf{x}_i) - y_i$. We remark that such expression is consistent in that for the extreme case where $p = 1$, dropout ‘degenerates’ to GD, hence the covariance matrix degenerates to a zero matrix, i.e., $\Sigma = \mathbf{0}_{D \times D}$. The following part is the specific derivation process of the covariance matrix.

J.1 ELEMENTS ON THE DIAGONAL

In this part, we compute Σ_{kk} for all $k \in [m]$.

$$\begin{aligned}\Sigma_{kk}(\boldsymbol{\theta}_{N-1}) &= \text{Cov} \left(\nabla_{\mathbf{q}_k} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N), \nabla_{\mathbf{q}_k} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N) \right) \\ &= \frac{1}{n^2} \sum_{i,j=1}^n \text{Cov} \left(e_i^N(\boldsymbol{\delta}_N)_k, e_j^N(\boldsymbol{\delta}_N)_k \right) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)) \otimes \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_j)),\end{aligned}$$

in order to compute $\text{Cov} \left(e_i^N(\boldsymbol{\delta}_N)_k, e_j^N(\boldsymbol{\delta}_N)_k \right)$, we need to compute firstly $\mathbb{E} \left[e_i^N e_j^N (\boldsymbol{\delta}_N)_k^2 \right]$, and since $\mathbb{E} \left[e_i^N e_j^N (\boldsymbol{\delta}_N)_k^2 \right]$ consists of four parts, one of which is

$$\begin{aligned}& \mathbb{E} \left[\left(\sum_{k'=1, k' \neq k}^m (\boldsymbol{\delta}_N)_{k'} a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) - y_i \right) \left(\sum_{l=1, l \neq k}^m (\boldsymbol{\delta}_N)_l a_l \sigma(\mathbf{w}_l^\top \mathbf{x}_j) - y_j \right) (\boldsymbol{\delta}_N)_k^2 \right] \\ &= \mathbb{E} \left[\left(\sum_{k'=1, k' \neq k}^m (\boldsymbol{\delta}_N)_{k'} a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) - y_i \right) \left(\sum_{l=1, l \neq k}^m (\boldsymbol{\delta}_N)_l a_l \sigma(\mathbf{w}_l^\top \mathbf{x}_j) - y_j \right) \right] \mathbb{E} \left[(\boldsymbol{\delta}_N)_k^2 \right] \\ &= \frac{1}{p} \left(\mathbb{E} \left[\sum_{k'=1, k' \neq k}^m (\boldsymbol{\delta}_N)_{k'}^2 a_{k'}^2 \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_j) \right] + \mathbb{E} \left[\sum_{k' \neq l, k', l \neq k} (\boldsymbol{\delta}_N)_{k'} (\boldsymbol{\delta}_N)_l a_{k'} a_l \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) \sigma(\mathbf{w}_l^\top \mathbf{x}_j) \right] \right) \\ &\quad - y_i \mathbb{E} \left[\sum_{k'=1, k' \neq k}^m (\boldsymbol{\delta}_N)_{k'} a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_j) \right] - y_j \mathbb{E} \left[\sum_{k'=1, k' \neq k}^m (\boldsymbol{\delta}_N)_{k'} a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) \right] + y_i y_j \\ &= \frac{1}{p^2} \sum_{k'=1, k' \neq k}^m a_{k'}^2 \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_j) + \frac{1}{p} \sum_{k' \neq l, k', l \neq k} a_{k'} a_l \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) \sigma(\mathbf{w}_l^\top \mathbf{x}_j) \\ &\quad - \frac{y_i}{p} \sum_{k'=1, k' \neq k}^m a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_j) - \frac{y_j}{p} \sum_{k'=1, k' \neq k}^m a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) + \frac{y_i y_j}{p} \\ &= \frac{1}{p} \left[\sum_{k'=1, k' \neq k}^m a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) - y_i \right] \left[\sum_{k'=1, k' \neq k}^m a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_j) - y_j \right] \\ &\quad + \left(\frac{1}{p^2} - \frac{1}{p} \right) \left(\sum_{k'=1, k' \neq k}^m a_{k'}^2 \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_j) \right),\end{aligned}$$

and the second part reads

$$\begin{aligned}& \mathbb{E} \left[(\boldsymbol{\delta}_N)_k a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \left(\sum_{l=1, l \neq k}^m (\boldsymbol{\delta}_N)_l a_l \sigma(\mathbf{w}_l^\top \mathbf{x}_j) - y_j \right) (\boldsymbol{\delta}_N)_k^2 \right] \\ &= \frac{a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)}{p^2} \left(\sum_{k'=1, k' \neq k}^m a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_j) - y_j \right),\end{aligned}$$

and by symmetry, the third part reads

$$\begin{aligned}& \mathbb{E} \left[(\boldsymbol{\delta}_N)_k a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_j) \left(\sum_{l=1, l \neq k}^m (\boldsymbol{\delta}_N)_l a_l \sigma(\mathbf{w}_l^\top \mathbf{x}_i) - y_i \right) (\boldsymbol{\delta}_N)_k^2 \right] \\ &= \frac{a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_j)}{p^2} \left(\sum_{k'=1, k' \neq k}^m a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) - y_i \right),\end{aligned}$$

and finally, the fourth part reads

$$\mathbb{E} [(\boldsymbol{\delta}_N)_k a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) (\boldsymbol{\delta}_N)_k a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_j) (\boldsymbol{\delta}_N)_k^2] = \frac{1}{p^3} a_k^2 \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \sigma(\mathbf{w}_k^\top \mathbf{x}_j).$$

To sum up,

$$\begin{aligned} \mathbb{E} [e_i^N e_j^N (\boldsymbol{\delta}_N)_k^2] &= \left(\frac{1}{p^2} - \frac{1}{p} \right) \left(\sum_{k'=1, k' \neq k}^m a_{k'}^2 \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_j) \right) \\ &\quad + \frac{1}{p} e_{i, \setminus k} e_{j, \setminus k} + \frac{a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_j)}{p^2} e_{i, \setminus k} + \frac{a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)}{p^2} e_{j, \setminus k} \\ &\quad + \frac{1}{p^3} a_k^2 \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \sigma(\mathbf{w}_k^\top \mathbf{x}_j), \end{aligned}$$

and

$$\begin{aligned} &\mathbb{E} [e_i^N (\boldsymbol{\delta}_N)_k] \mathbb{E} [e_j^N (\boldsymbol{\delta}_N)_k] \\ &= \left(e_{i, \setminus k} + \frac{1}{p} a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \right) \left(e_{j, \setminus k} + \frac{1}{p} a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_j) \right) \\ &= e_{i, \setminus k} e_{j, \setminus k} + \frac{a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_j)}{p} e_{i, \setminus k} + \frac{a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)}{p} e_{j, \setminus k} + \frac{1}{p^2} a_k^2 \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \sigma(\mathbf{w}_k^\top \mathbf{x}_j), \end{aligned}$$

hence

$$\begin{aligned} &\text{Cov} (e_i^N (\boldsymbol{\delta}_N)_k, e_j^N (\boldsymbol{\delta}_N)_k) \\ &= \mathbb{E} [e_i^N e_j^N (\boldsymbol{\delta}_N)_k^2] - \mathbb{E} [e_i^N (\boldsymbol{\delta}_N)_k] \mathbb{E} [e_j^N (\boldsymbol{\delta}_N)_k] \\ &= \left(\frac{1}{p^2} - \frac{1}{p} \right) \left(\sum_{k'=1, k' \neq k}^m a_{k'}^2 \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_j) \right) \\ &\quad + \left(\frac{1}{p} - 1 \right) e_{i, \setminus k} e_{j, \setminus k} + \left(\frac{1}{p^2} - \frac{1}{p} \right) a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) e_{j, \setminus k} \\ &\quad + \left(\frac{1}{p^2} - \frac{1}{p} \right) a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_j) e_{i, \setminus k} + \left(\frac{1}{p^3} - \frac{1}{p^2} \right) a_k^2 \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \sigma(\mathbf{w}_k^\top \mathbf{x}_j) \\ &= \left(\frac{1}{p} - 1 \right) \mathbb{E} (e_i^N (\boldsymbol{\delta}_N)_k) \mathbb{E} (e_j^N (\boldsymbol{\delta}_N)_k) + \left(\frac{1}{p^2} - \frac{1}{p} \right) \left(\sum_{k'=1, k' \neq k}^m a_{k'}^2 \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_j) \right), \end{aligned}$$

by summation over the indices i and j , for each $k \in [m]$, the covariance matrix reads:

$$\begin{aligned} \boldsymbol{\Sigma}_{kk}(\boldsymbol{\theta}_{N-1}) &= \text{Cov} \left(\nabla_{\mathbf{q}_k} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N), \nabla_{\mathbf{q}_k} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N) \right) \\ &= \left(\frac{1}{p} - 1 \right) \left(\frac{1}{n} \sum_{i=1}^n \left(e_{i, \setminus k} + \frac{1}{p} a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \right) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)) \right) \\ &\quad \otimes \left(\frac{1}{n} \sum_{i=1}^n \left(e_{i, \setminus k} + \frac{1}{p} a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \right) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)) \right) \\ &\quad + \left(\frac{1}{p^2} - \frac{1}{p} \right) \sum_{l=1, l \neq k}^m \left(\frac{1}{n} \sum_{i=1}^n a_l \sigma(\mathbf{w}_l^\top \mathbf{x}_i) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)) \right) \\ &\quad \otimes \left(\frac{1}{n} \sum_{i=1}^n a_l \sigma(\mathbf{w}_l^\top \mathbf{x}_i) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)) \right). \end{aligned}$$

J.2 ELEMENTS OFF THE DIAGONAL

In this part, we compute Σ_{kr} for all $k, r \in [m]$, where $k \neq r$.

$$\begin{aligned}\Sigma_{kr}(\boldsymbol{\theta}_{N-1}) &= \text{Cov} \left(\nabla_{\mathbf{q}_k} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N), \nabla_{\mathbf{q}_r} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N) \right) \\ &= \frac{1}{n^2} \sum_{i,j=1}^n \text{Cov} \left(e_i^N(\boldsymbol{\delta}_N)_k, e_j^N(\boldsymbol{\delta}_N)_r \right) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)) \otimes \nabla_{\mathbf{q}_r} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_j)),\end{aligned}$$

in order to compute $\text{Cov} \left(e_i^N(\boldsymbol{\delta}_N)_k, e_j^N(\boldsymbol{\delta}_N)_r \right)$, we need to compute firstly $\mathbb{E} \left[e_i^N e_j^N(\boldsymbol{\delta}_N)_k(\boldsymbol{\delta}_N)_r \right]$, and since $\mathbb{E} \left[e_i^N e_j^N(\boldsymbol{\delta}_N)_k(\boldsymbol{\delta}_N)_r \right]$ consists of nine parts, one of which is

$$\begin{aligned}& \mathbb{E} \left[\left(\sum_{k'=1, k' \neq k, k' \neq r}^m (\boldsymbol{\delta}_N)_{k'} a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) - y_i \right) \left(\sum_{l=1, l \neq k, l \neq r}^m (\boldsymbol{\delta}_N)_l a_l \sigma(\mathbf{w}_l^\top \mathbf{x}_j) - y_j \right) (\boldsymbol{\delta}_N)_k (\boldsymbol{\delta}_N)_r \right] \\ &= \mathbb{E} \left[\left(\sum_{k'=1, k' \neq k, k' \neq r}^m (\boldsymbol{\delta}_N)_{k'} a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) - y_i \right) \left(\sum_{l=1, l \neq k, l \neq r}^m (\boldsymbol{\delta}_N)_l a_l \sigma(\mathbf{w}_l^\top \mathbf{x}_j) - y_j \right) \right] \mathbb{E} [(\boldsymbol{\delta}_N)_k (\boldsymbol{\delta}_N)_r] \\ &= \frac{1}{p} \sum_{k'=1, k' \neq k, k' \neq r}^m a_{k'}^2 \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_j) + \sum_{k' \neq l \text{ and } k', l \neq k, r} a_{k'} a_l \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) \sigma(\mathbf{w}_l^\top \mathbf{x}_j) \\ &\quad - y_i \sum_{k'=1, k' \neq k, k' \neq r}^m a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_j) - y_j \sum_{k'=1, k' \neq k, k' \neq r}^m a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) + y_i y_j \\ &= \left[\sum_{k'=1, k' \neq k, k' \neq r}^m a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) - y_i \right] \left[\sum_{k'=1, k' \neq k, k' \neq r}^m a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_j) - y_j \right] \\ &\quad + \left(\frac{1}{p} - 1 \right) \left(\sum_{k'=1, k' \neq k, k' \neq r}^m a_{k'}^2 \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_j) \right) \\ &= e_{i, \setminus k, \setminus r} e_{j, \setminus k, \setminus r} + \left(\frac{1}{p} - 1 \right) \left(\sum_{k'=1, k' \neq k, k' \neq r}^m a_{k'}^2 \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_j) \right),\end{aligned}$$

and the second part reads

$$\begin{aligned}& \mathbb{E} \left[\left(\sum_{k'=1, k' \neq k, k' \neq r}^m (\boldsymbol{\delta}_N)_{k'} a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) - y_i \right) (\boldsymbol{\delta}_N)_k a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_j) (\boldsymbol{\delta}_N)_k (\boldsymbol{\delta}_N)_r \right] \\ &= \mathbb{E} \left[\sum_{k'=1, k' \neq k, k' \neq r}^m (\boldsymbol{\delta}_N)_{k'} a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) - y_i \right] a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_j) \mathbb{E} [(\boldsymbol{\delta}_N)_k^2 (\boldsymbol{\delta}_N)_r] = \frac{a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_j)}{p} e_{i, \setminus k, \setminus r},\end{aligned}$$

by similar reasoning and symmetry, the third part reads

$$\begin{aligned}& \mathbb{E} \left[\left(\sum_{k'=1, k' \neq k, k' \neq r}^m (\boldsymbol{\delta}_N)_{k'} a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) - y_i \right) (\boldsymbol{\delta}_N)_r a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_j) (\boldsymbol{\delta}_N)_k (\boldsymbol{\delta}_N)_r \right] \\ &= \mathbb{E} \left[\sum_{k'=1, k' \neq k, k' \neq r}^m (\boldsymbol{\delta}_N)_{k'} a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) - y_i \right] a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_j) \mathbb{E} [(\boldsymbol{\delta}_N)_k (\boldsymbol{\delta}_N)_r^2] = \frac{a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_j)}{p} e_{i, \setminus k, \setminus r},\end{aligned}$$

also by similar reasoning and symmetry, the fourth part reads

$$\begin{aligned} & \mathbb{E} \left[\left(\sum_{k'=1, k' \neq k, k' \neq r}^m (\delta_N)_{k'} a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_j) - y_j \right) (\delta_N)_k a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) (\delta_N)_k (\delta_N)_r \right] \\ &= \mathbb{E} \left[\sum_{k'=1, k' \neq k, k' \neq r}^m (\delta_N)_{k'} a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_j) - y_j \right] a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \mathbb{E} [(\delta_N)_k^2 (\delta_N)_r] = \frac{a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)}{p} e_{j, \setminus k, \setminus r}, \end{aligned}$$

and the fifth part reads

$$\begin{aligned} \mathbb{E} [(\delta_N)_k a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) (\delta_N)_k a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_j) (\delta_N)_k (\delta_N)_r] &= \mathbb{E} [(\delta_N)_k^3 (\delta_N)_r a_k^2 \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \sigma(\mathbf{w}_k^\top \mathbf{x}_j)] \\ &= \frac{1}{p^2} a_k^2 \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \sigma(\mathbf{w}_k^\top \mathbf{x}_j), \end{aligned}$$

and the sixth part reads

$$\begin{aligned} & \mathbb{E} [(\delta_N)_k a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) (\delta_N)_r a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_j) (\delta_N)_k (\delta_N)_r] \\ &= \mathbb{E} [(\delta_N)_k^2 (\delta_N)_r^2 a_k a_r \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \sigma(\mathbf{w}_r^\top \mathbf{x}_j)] = \frac{1}{p^2} a_k a_r \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \sigma(\mathbf{w}_r^\top \mathbf{x}_j), \end{aligned}$$

also by similar reasoning and symmetry, the seventh part reads

$$\begin{aligned} & \mathbb{E} \left[\left(\sum_{k'=1, k' \neq k, k' \neq r}^m (\delta_N)_{k'} a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_j) - y_j \right) (\delta_N)_r a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i) (\delta_N)_k (\delta_N)_r \right] \\ &= \mathbb{E} \left[\sum_{k'=1, k' \neq k, k' \neq r}^m (\delta_N)_{k'} a_{k'} \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_j) - y_j \right] a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i) \mathbb{E} [(\delta_N)_k (\delta_N)_r^2] = \frac{a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i)}{p} e_{j, \setminus k, \setminus r}, \end{aligned}$$

and the eighth part reads

$$\begin{aligned} \mathbb{E} [(\delta_N)_r a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i) (\delta_N)_k a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_j) (\delta_N)_k (\delta_N)_r] &= \mathbb{E} [(\delta_N)_k^2 (\delta_N)_r^2 a_k a_r \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \sigma(\mathbf{w}_r^\top \mathbf{x}_j)] \\ &= \frac{1}{p^2} a_k a_r \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \sigma(\mathbf{w}_r^\top \mathbf{x}_j), \end{aligned}$$

and the ninth part reads

$$\begin{aligned} & \mathbb{E} [(\delta_N)_r a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i) (\delta_N)_r a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_j) (\delta_N)_k (\delta_N)_r] \\ &= \mathbb{E} [(\delta_N)_k (\delta_N)_r^3 a_r^2 \sigma(\mathbf{w}_r^\top \mathbf{x}_i) \sigma(\mathbf{w}_r^\top \mathbf{x}_j)] = \frac{1}{p^2} a_r^2 \sigma(\mathbf{w}_r^\top \mathbf{x}_i) \sigma(\mathbf{w}_r^\top \mathbf{x}_j). \end{aligned}$$

To sum up,

$$\begin{aligned} & \mathbb{E} [e_i^N e_j^N (\delta_N)_k (\delta_N)_r] \\ &= e_{i, \setminus k, \setminus r} e_{j, \setminus k, \setminus r} + \left(\frac{1}{p} - 1 \right) \left(\sum_{k'=1, k' \neq k, k' \neq r}^m a_{k'}^2 \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_j) \right) + \frac{a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_j)}{p} e_{i, \setminus k, \setminus r} \\ & \quad + \frac{a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_j)}{p} e_{i, \setminus k, \setminus r} + \frac{a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)}{p} e_{j, \setminus k, \setminus r} + \frac{1}{p^2} a_k^2 \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \sigma(\mathbf{w}_k^\top \mathbf{x}_j) \\ & \quad + \frac{1}{p^2} a_k a_r \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \sigma(\mathbf{w}_r^\top \mathbf{x}_j) + \frac{a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i)}{p} e_{j, \setminus k, \setminus r} \\ & \quad + \frac{1}{p^2} a_k a_r \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \sigma(\mathbf{w}_r^\top \mathbf{x}_j) + \frac{1}{p^2} a_r^2 \sigma(\mathbf{w}_r^\top \mathbf{x}_i) \sigma(\mathbf{w}_r^\top \mathbf{x}_j), \end{aligned}$$

and

$$\begin{aligned}
& \mathbb{E} [e_i^N(\boldsymbol{\delta}_N)_k] \mathbb{E} [e_j^N(\boldsymbol{\delta}_N)_r] \\
&= \left(e_{i,\setminus k,\setminus r} + a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i) + \frac{1}{p} a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \right) \left(e_{j,\setminus k,\setminus r} + a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_j) + \frac{1}{p} a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_j) \right) \\
&= e_{i,\setminus k,\setminus r} e_{j,\setminus k,\setminus r} + e_{i,\setminus k,\setminus r} a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_j) + \frac{1}{p} e_{i,\setminus k,\setminus r} a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_j) + a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i) e_{j,\setminus k,\setminus r} \\
&\quad + a_r a_k \sigma(\mathbf{w}_r^\top \mathbf{x}_i) \sigma(\mathbf{w}_k^\top \mathbf{x}_j) + \frac{1}{p} a_r^2 \sigma(\mathbf{w}_r^\top \mathbf{x}_i) \sigma(\mathbf{w}_r^\top \mathbf{x}_j) + \frac{1}{p} a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) e_{j,\setminus k,\setminus r} \\
&\quad + \frac{1}{p} a_k^2 \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \sigma(\mathbf{w}_k^\top \mathbf{x}_j) + \frac{1}{p^2} a_r a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \sigma(\mathbf{w}_r^\top \mathbf{x}_j),
\end{aligned}$$

hence

$$\begin{aligned}
& \text{Cov} (e_i^N(\boldsymbol{\delta}_N)_k, e_j^N(\boldsymbol{\delta}_N)_r) \\
&= \mathbb{E} [e_i^N e_j^N(\boldsymbol{\delta}_N)_k(\boldsymbol{\delta}_N)_r] - \mathbb{E} [e_i^N(\boldsymbol{\delta}_N)_k] \mathbb{E} [e_j^N(\boldsymbol{\delta}_N)_r] \\
&= \left(\frac{1}{p} - 1 \right) \left(\sum_{k'=1, k' \neq k, k' \neq r}^m a_{k'}^2 \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_i) \sigma(\mathbf{w}_{k'}^\top \mathbf{x}_j) \right) + \left(\frac{1}{p} - 1 \right) a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_j) e_{i,\setminus k,\setminus r} \\
&\quad + \left(\frac{1}{p} - 1 \right) a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i) e_{j,\setminus k,\setminus r} + \left(\frac{1}{p^2} - \frac{1}{p} \right) a_r^2 \sigma(\mathbf{w}_r^\top \mathbf{x}_i) \sigma(\mathbf{w}_r^\top \mathbf{x}_j) \\
&\quad + \left(\frac{1}{p^2} - \frac{1}{p} \right) a_k^2 \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \sigma(\mathbf{w}_k^\top \mathbf{x}_j) + \left(\frac{1}{p^2} - 1 \right) a_r a_k \sigma(\mathbf{w}_r^\top \mathbf{x}_i) \sigma(\mathbf{w}_k^\top \mathbf{x}_j),
\end{aligned}$$

by summation over the indices i and j , the covariance matrix reads

$$\begin{aligned}
& \boldsymbol{\Sigma}_{kr}(\boldsymbol{\theta}_{N-1}) = \text{Cov} \left(\nabla_{\mathbf{q}_k} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N), \nabla_{\mathbf{q}_r} R_S^{\text{drop}}(\boldsymbol{\theta}_{N-1}; \boldsymbol{\delta}_N) \right) \\
&= \left(\frac{1}{p} - 1 \right) \left(\frac{1}{n} \sum_{i=1}^n \left(e_{i,\setminus k,\setminus r} + \frac{1}{p} a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) + \frac{1}{p} a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i) \right) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)) \right) \\
&\quad \otimes \left(\frac{1}{n} \sum_{i=1}^n a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i)) \right) \\
&\quad + \left(\frac{1}{p} - 1 \right) \left(\frac{1}{n} \sum_{i=1}^n a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i) \nabla_{\mathbf{q}_k} (a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i)) \right) \\
&\quad \otimes \left(\frac{1}{n} \sum_{i=1}^n \left(e_{i,\setminus k,\setminus r} + a_k \sigma(\mathbf{w}_k^\top \mathbf{x}_i) + \frac{1}{p} a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i) \right) \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i)) \right),
\end{aligned}$$

K THE STRUCTURAL SIMILARITY BETWEEN HESSIAN AND COVARIANCE

We can derive the Hessian of the loss function in the expectation sense with respect to the dropout noise δ and the covariance matrix of dropout noise under intuitive approximations. We first show our assumptions as follows:

Assumption 1. *The NN piece-wise linear activation.*

Assumption 2. *The parameters of NN's output layer are fixed during training.*

Assumption 3. *We study the loss landscape after training reaches a stable stage, i.e., the loss function in the sense of expectation is small enough,*

$$\mathbb{E}_{\delta} \nabla_{\theta} R_S^{\text{drop}}(\theta; \delta) \approx \mathbf{0}.$$

Hessian matrix with dropout regularization Based on the Assumption 1, 2, the Hessian matrix of the loss function with respect to $f_{\theta, \delta}^{\text{drop}}(\mathbf{x})$ can be written in the mean sense as:

$$\mathbf{H}(\theta) \approx \frac{1}{n} \sum_{i=1}^n \left[\nabla_{\theta} f_{\theta}(\mathbf{x}_i) \otimes \nabla_{\theta} f_{\theta}(\mathbf{x}_i) + \frac{1-p}{p} \sum_{r=1}^m \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i)) \otimes \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i)) \right],$$

where $\mathbf{H}(\theta) := \nabla_{\theta}^2 L_S(\theta)$.

Proof. We first compute the Hessian matrix after taking expectations with respect to the dropout variable,

$$\nabla_{\theta}^2 L_S(\theta) = \nabla_{\theta}^2 R_S(\theta) + \frac{1-p}{2np} \sum_{i=1}^n \sum_{r=1}^m \nabla_{\mathbf{q}_r}^2 (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i))^2. \quad (53)$$

The first and second terms on the RHS of the Eq. (53) are as follows,

$$\begin{aligned} \nabla_{\theta}^2 R_S(\theta) &= \frac{1}{n} \sum_{i=1}^n (\nabla_{\theta} f_{\theta}(\mathbf{x}_i) \otimes \nabla_{\theta} f_{\theta}(\mathbf{x}_i) + (f_{\theta}(\mathbf{x}_i) - y_i) \cdot \nabla_{\theta}^2 f_{\theta}(\mathbf{x}_i)) \\ &= \frac{1-p}{2np} \sum_{i=1}^n \sum_{r=1}^m \nabla_{\mathbf{q}_r}^2 (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i))^2 \\ &= \frac{1-p}{np} \sum_{i=1}^n \sum_{r=1}^m \left(\nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i)) \otimes \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i)) + (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i)) \cdot \nabla_{\mathbf{q}_r}^2 (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i))^2 \right). \end{aligned}$$

Note that for linear activate function, $\nabla_{\theta}^2 f_{\theta}(\mathbf{x}_i) = \nabla_{\mathbf{q}_r}^2 (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i))^2 = \mathbf{0}$, a.e. $\forall i \in [n], \forall r \in [m]$, we have

$$\begin{aligned} \nabla_{\theta}^2 R_S(\theta) &= \frac{1}{n} \sum_{i=1}^n \nabla_{\theta} f_{\theta}(\mathbf{x}_i) \otimes \nabla_{\theta} f_{\theta}(\mathbf{x}_i) \\ \frac{1-p}{2np} \sum_{i=1}^n \sum_{r=1}^m \nabla_{\mathbf{q}_r}^2 (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i))^2 &= \frac{1-p}{np} \sum_{i=1}^n \sum_{r=1}^m \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i)) \otimes \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i)). \end{aligned}$$

Thus the Eq. (53) can be rewritten as

$$\mathbf{H}(\theta) = \frac{1}{n} \sum_{i=1}^n \left(\nabla_{\theta} f_{\theta}(\mathbf{x}_i) \otimes \nabla_{\theta} f_{\theta}(\mathbf{x}_i) + \frac{1-p}{p} \sum_{r=1}^m \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i)) \otimes \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i)) \right). \quad \square$$

Covariance matrix with dropout regularization Based on the Assumption 3, the covariance matrix of the loss function under the randomness of dropout variable δ and data \mathbf{x} can be written as:

$$\Sigma(\theta) \approx \frac{1}{n} \sum_{i=1}^n \left[l_{i,1} \nabla_{\theta} f_{\theta}(\mathbf{x}_i) \otimes \nabla_{\theta} f_{\theta}(\mathbf{x}_i) + l_{i,2} \frac{1-p}{p} \sum_{r=1}^m \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i)) \otimes \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i)) \right],$$

where $l_{i,1} := (e_i)^2 + \frac{1-p}{p} \sum_{r=1}^m a_r^2 \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i)^2$, $l_{i,2} := (e_i)^2$.

Proof. For simplicity, we approximate the loss function through Taylor expansion, which is also used in Wei et al. (2020),

$$\ell(f_{\theta}(\mathbf{x}_i; \boldsymbol{\delta}), y_i) \approx \ell(f_{\theta}(\mathbf{x}_i), y_i) + (f_{\theta}(\mathbf{x}_i) - y_i) \sum_{r=1}^m a_r (\boldsymbol{\delta} - \mathbf{1})_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i),$$

where $\ell(f_{\theta}(\mathbf{x}_i; \boldsymbol{\delta}), y_i) = \frac{1}{2} (f_{\theta}(\mathbf{x}_i; \boldsymbol{\delta}) - y_i)^2$, $\ell(f_{\theta}(\mathbf{x}_i), y_i) = \frac{1}{2} (f_{\theta}(\mathbf{x}_i) - y_i)^2$. The covariance matrix under dropout regularization is

$$\begin{aligned} \boldsymbol{\Sigma}(\boldsymbol{\theta}) &\approx \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\boldsymbol{\delta}} (\nabla_{\boldsymbol{\theta}} \ell(f_{\theta}(\mathbf{x}_i; \boldsymbol{\delta}), y_i) \otimes \nabla_{\boldsymbol{\theta}} \ell(f_{\theta}(\mathbf{x}_i; \boldsymbol{\delta}), y_i)) - \nabla_{\boldsymbol{\theta}} \mathbb{E}_{\boldsymbol{\delta}} R_{\mathcal{S}}^{\text{drop}}(\boldsymbol{\theta}; \boldsymbol{\delta}) \otimes \nabla_{\boldsymbol{\theta}} \mathbb{E}_{\boldsymbol{\delta}} R_{\mathcal{S}}^{\text{drop}}(\boldsymbol{\theta}; \boldsymbol{\delta}) \\ &\approx \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\boldsymbol{\delta}} (\nabla_{\boldsymbol{\theta}} \ell(f_{\theta}(\mathbf{x}_i; \boldsymbol{\delta}), y_i) \otimes \nabla_{\boldsymbol{\theta}} \ell(f_{\theta}(\mathbf{x}_i; \boldsymbol{\delta}), y_i)). \end{aligned}$$

Combining the properties of the dropout variable $\boldsymbol{\delta}$, we have,

$$\begin{aligned} \boldsymbol{\Sigma}(\boldsymbol{\theta}) &\approx \frac{1}{n} \sum_{i=1}^n \nabla_{\boldsymbol{\theta}} \ell(f_{\theta}(\mathbf{x}_i), y_i) \otimes \nabla_{\boldsymbol{\theta}} \ell(f_{\theta}(\mathbf{x}_i), y_i) \\ &\quad + \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\boldsymbol{\delta}} \left(\sum_{r=1}^m (\boldsymbol{\delta} - \mathbf{1})_r \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i) e_i) \otimes \sum_{r=1}^m (\boldsymbol{\delta} - \mathbf{1})_r \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i) e_i) \right) \\ &= \frac{1}{n} \sum_{i=1}^n \left(\nabla_{\boldsymbol{\theta}} \ell(f_{\theta}(\mathbf{x}_i), y_i) \otimes \nabla_{\boldsymbol{\theta}} \ell(f_{\theta}(\mathbf{x}_i), y_i) + \frac{1-p}{p} \sum_{r=1}^m \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i) e_i) \otimes \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i) e_i) \right) \\ &:= \frac{1}{n} \sum_{i=1}^n \left(\boldsymbol{\Sigma}_1(\mathbf{x}_i, y_i) + \frac{1-p}{p} \boldsymbol{\Sigma}_2(\mathbf{x}_i, y_i) \right). \end{aligned} \tag{54}$$

We calculate the two terms on the RHS of the Eq. (54) separately:

$$\boldsymbol{\Sigma}_1(\mathbf{x}_i, y_i) = (e_i)^2 \cdot \nabla_{\boldsymbol{\theta}} f_{\theta}(\mathbf{x}_i) \otimes \nabla_{\boldsymbol{\theta}} f_{\theta}(\mathbf{x}_i),$$

$$\begin{aligned} \boldsymbol{\Sigma}_2(\mathbf{x}_i, y_i) &= (e_i)^2 \sum_{r=1}^m \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i)) \otimes \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i)) + \nabla_{\boldsymbol{\theta}} f_{\theta}(\mathbf{x}_i) \otimes \nabla_{\boldsymbol{\theta}} f_{\theta}(\mathbf{x}_i) \sum_{r=1}^m (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i))^2 \\ &\quad + 2 \sum_{r=1}^m e_i a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i) \cdot \nabla_{\boldsymbol{\theta}} e_i \otimes \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i)) \\ &= (e_i)^2 \sum_{r=1}^m \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i)) \otimes \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i)) + \nabla_{\boldsymbol{\theta}} f_{\theta}(\mathbf{x}_i) \otimes \nabla_{\boldsymbol{\theta}} f_{\theta}(\mathbf{x}_i) \sum_{r=1}^m (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i))^2 \\ &\quad + \frac{1}{2} \sum_{r=1}^m \nabla_{\boldsymbol{\theta}} (e_i)^2 \otimes \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i))^2. \end{aligned}$$

Under the assumption that $\nabla_{\boldsymbol{\theta}} (e_i)^2 = 2 \cdot \nabla_{\boldsymbol{\theta}} \ell(f_{\theta}(\mathbf{x}_i), y_i) = \mathbf{0}$, $\forall i \in [n]$, we have

$$\boldsymbol{\Sigma}_2(\mathbf{x}_i, y_i) = (e_i)^2 \sum_{r=1}^m \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i)) \otimes \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i)) + \nabla_{\boldsymbol{\theta}} f_{\theta}(\mathbf{x}_i) \otimes \nabla_{\boldsymbol{\theta}} f_{\theta}(\mathbf{x}_i) \sum_{r=1}^m (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i))^2.$$

Thus the Eq. (54) can be rewritten as

$$\begin{aligned} \boldsymbol{\Sigma}(\boldsymbol{\theta}) &= \frac{1}{n} \sum_{i=1}^n \nabla_{\boldsymbol{\theta}} f_{\theta}(\mathbf{x}_i) \otimes \nabla_{\boldsymbol{\theta}} f_{\theta}(\mathbf{x}_i) \left((e_i)^2 + \frac{1-p}{p} \sum_{r=1}^m (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i))^2 \right) \\ &\quad + \frac{1-p}{np} \sum_{i=1}^n \sum_{r=1}^m (e_i)^2 \cdot \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i)) \otimes \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^{\top} \mathbf{x}_i)). \end{aligned}$$

Note that

$$(e_i)^2 + \frac{1-p}{p} \sum_{r=1}^m (a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i))^2 = \mathbb{E}_{\boldsymbol{\delta}} 2\ell(f_{\boldsymbol{\theta}}(\mathbf{x}_i; \boldsymbol{\delta}), y_i),$$

we have

$$\begin{aligned} \boldsymbol{\Sigma}(\boldsymbol{\theta}) &= \frac{2}{n} \sum_{i=1}^n \mathbb{E}_{\boldsymbol{\delta}} \ell(f_{\boldsymbol{\theta}}(\mathbf{x}_i; \boldsymbol{\delta}), y_i) \cdot \nabla_{\boldsymbol{\theta}} f_{\boldsymbol{\theta}}(\mathbf{x}_i) \otimes \nabla_{\boldsymbol{\theta}} f_{\boldsymbol{\theta}}(\mathbf{x}_i) \\ &\quad + \frac{2(1-p)}{np} \sum_{i=1}^n \sum_{r=1}^m (\ell(f_{\boldsymbol{\theta}}(\mathbf{x}_i), y_i)) \cdot \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i)) \otimes \nabla_{\mathbf{q}_r} (a_r \sigma(\mathbf{w}_r^\top \mathbf{x}_i)). \end{aligned}$$

□