## A Contextualizing our Work

## A. 1 Additional Recent Works

Variations of Adam have been proposed to improve its speed of convergence, generalization, and stability during training. Reddi et al. (2018) observed that Adam does not collect long-term memory of past gradients and therefore the effective learning rate could be increasing in some cases. Hence, they propose AMSGrad that maintains a maximum over the exponential running average of the squared gradients. Zaheer et al. (2018) proposed a more controlled increase in the effective learning rate by switching to additive updates, using a more refined version of AdaGrad (Duchi et al., 2011). Chen et al. (2021) unified the generalization ability of vanilla SGD and the convergence speed of Adam by introducing a new adaptive parameter $p \in\left(0, \frac{1}{2}\right]$ that can be hypertuned for each setting. Other variations include (a) Nadam (Dozat, 2016) that uses Nesterov momentum, (b) AdamW (Loshchilov and Hutter, 2019) that decouples the weight decay from the optimization step, (c) AdaBound (Luo et al., 2019) that maintains a dynamic upper and lower bound on the step size, (d) AdaBelief (Zhuang et al. 2020) uses a decaying average of estimated variance in the gradient in place of the running average of the squared gradients, (e) QHAdam (Ma and Yarats, 2019) that replaces both momentum estimators in Adam with quasi-hyperbolic terms, etc. LAMB (You et al., 2020) used a layerwise adaptive version of Adam to pretrain large language models efficiently.

## A. 2 Broader Impact

Our work is primarily theoretical in nature, but we discuss its broader impacts here. Strubell et al. (2020) highlighted the environmental impact of training large language models. Formal scaling rules remove the need to grid search over hyperparameters: in the case of adaptive algorithms, the grid search is over an even larger space because of the additional adaptivity hyperparameters. Hence, our work reduces the number of times researchers need to train very large models to find ones that work just as well as their smaller, slower counterparts. At the same time, we recognize that the presence of a formal scaling rule may encourage people who were otherwise discouraged by the daunting grid search to train very large models.

## B SDE Approximation Theorem

In this section, we introduce a theorem in aid of our analysis for approximating Stochastic Gradient Algorithm (SGA) with Stochastic Differential Equation (SDE).

We consider Stochastic Gradient Algorithm (SGA) of the following form:

$$
\begin{equation*}
\boldsymbol{x}_{k+1}=\boldsymbol{x}_{k}+\eta_{\mathrm{e}} \boldsymbol{h}_{k}\left(\boldsymbol{x}_{k}, \boldsymbol{\xi}_{k}, \eta_{\mathrm{e}}\right), \tag{4}
\end{equation*}
$$

where $\boldsymbol{x}_{k} \in \mathbb{R}^{D}$ is the parameter vector, $\eta_{\mathrm{e}}$ is the learning rate, $\boldsymbol{h}_{k}$ is a vector-valued function that can depend on the current step $k$, the current parameter $\boldsymbol{x}_{k}$, a random vector $\boldsymbol{\xi}_{k}$, and the learning rate $\eta_{\mathrm{e}}$. The random vector $\boldsymbol{\xi}_{k}$ is sampled from a certain distribution $\Xi\left(\boldsymbol{x}_{k}, \eta_{\mathrm{e}}\right)$ in every step.
We consider Stochastic Differential Equation (SDE) of the following form:

$$
\begin{equation*}
\mathrm{d} \boldsymbol{X}_{t}=\boldsymbol{b}\left(X_{t}, t\right) \mathrm{d} t+\boldsymbol{\sigma}\left(X_{t}, t\right) \mathrm{d} \boldsymbol{W}_{t} \tag{5}
\end{equation*}
$$

where $\boldsymbol{b}: \mathbb{R}^{D} \times \mathbb{R} \rightarrow \mathbb{R}^{D}$ is the drift vector function, $\boldsymbol{\sigma}: \mathbb{R}^{D} \times \mathbb{R} \rightarrow \mathbb{R}^{D \times D}$ is the diffusion matrix function.

Let $\mathcal{P}_{X}(\boldsymbol{x}, s, t)$ be the probability distribution of $\boldsymbol{X}_{t}$ in (5) when the initial condition is $\boldsymbol{X}_{s}=\boldsymbol{x}$. Now we define the following random variables to characterize the one-step changes of SGA and SDE.

$$
\begin{array}{lll}
\boldsymbol{\Delta}(\boldsymbol{x}, k):=\eta_{\mathrm{e}} \boldsymbol{h}_{k}\left(\boldsymbol{x}, \boldsymbol{\xi}, \eta_{\mathrm{e}}\right), & \text { where } \quad \boldsymbol{\xi} \sim \Xi\left(\boldsymbol{x}, \eta_{\mathrm{e}}\right) . \\
\tilde{\boldsymbol{\Delta}}(\boldsymbol{x}, k):=\boldsymbol{X}_{(k+1) \eta_{\mathrm{e}}}-\boldsymbol{x}, & \text { where } \quad \boldsymbol{X}_{(k+1) \eta_{\mathrm{e}}} \sim \mathcal{P}_{X}\left(\boldsymbol{x}, k \eta_{\mathrm{e}},(k+1) \eta_{\mathrm{e}}\right) . \tag{7}
\end{array}
$$

The following regularity condition is needed for our main theorem.
Definition B.1. A function $g: \mathbb{R}^{d} \rightarrow \mathbb{R}$ is said to have polynomial growth if there exist positive integers $\kappa_{1}, \kappa_{2}>0$ such that

$$
|g(\boldsymbol{x})| \leq \kappa_{1}\left(1+\|\boldsymbol{x}\|_{2}^{2 \kappa_{2}}\right),
$$

for all $\boldsymbol{x} \in \mathbb{R}^{d}$. We let $G$ denote the set of all such functions. For each integer $\alpha \geq 1, G^{\alpha}$ denotes the set of $\alpha$-times continuously differentiable functions $\mathbb{R}^{d} \rightarrow \mathbb{R}$ which, together with its partial derivatives up to and including order $\alpha$, belong to $G$. If $g$ depends on additional parameters, we say $g \in G^{\alpha}$ uniformly if the constants $\kappa_{1}, \kappa_{2}$ are independent of those parameters.

Now we are ready to introduce the main theorem for the order of approximation by modeling SGA as SDE. We follow the same proof strategy as Li et al. (2019). First, we show that the moments of the one-step difference for the SGA and for the SDE of our interest are close to each other. Then, under some regularity conditions, we translate the one-step error to an error over a finite interval of time to get the final result. The following is our main theorem to achieve such translation, and it is adapted from Theorem 3 in Li et al. (2019) with slightly different conditions to match our need in the analysis of adaptive gradient methods.
Theorem B. 2 (Adaption of Theorem 3 in Li et al. (2019)). Let $T>0, \eta_{\mathrm{e}} \in(0,1 \wedge T)$ and $N=$ $\left\lfloor T / \eta_{\mathrm{e}}\right\rfloor$. Consider an SDE with drift vector $\boldsymbol{b}\left(\boldsymbol{x}, \eta_{\mathrm{e}}\right)$ and diffusion matrix $\boldsymbol{\sigma}\left(\boldsymbol{x}, \eta_{\mathrm{e}}\right)$, and a stochastic gradient algorithm with initial point $\boldsymbol{x}_{0} \in \mathbb{R}^{D}$ and update rule $\boldsymbol{x}_{k+1}=\boldsymbol{x}_{k}+\eta_{\mathrm{e}} \boldsymbol{h}\left(\boldsymbol{x}_{k}, \boldsymbol{\xi}_{k}, \eta_{\mathrm{e}}\right)$. Let $\mathcal{X}_{k}$ be the support of the random variable $\boldsymbol{x}_{k}$ given $\boldsymbol{x}_{0}$. Assume the following conditions hold:
(1). The drift function $\boldsymbol{b}$ and diffusion function $\boldsymbol{\sigma}$ are Lipschitz and belong to $G^{4}$;
(2). The moments of $\boldsymbol{\Delta} \in \mathbb{R}^{D}$ and $\tilde{\boldsymbol{\Delta}} \in \mathbb{R}^{D}$ satisfy the following. There is a function $K_{1} \in G$ (independent of $\eta_{\mathrm{e}}$ ) such that for all $\boldsymbol{x} \in \mathcal{X}_{k}$ and $1 \leq i, j, l \leq D$,

$$
\begin{aligned}
\left|\mathbb{E}\left[\Delta_{i}(\boldsymbol{x}, k)-\tilde{\Delta}_{i}(\boldsymbol{x}, k)\right]\right| & \leq K_{1}(\boldsymbol{x}) \eta_{\mathrm{e}}^{2} \\
\left|\mathbb{E}\left[\Delta_{i}(\boldsymbol{x}, k) \Delta_{j}(\boldsymbol{x}, k)-\tilde{\Delta}_{i}(\boldsymbol{x}, k) \tilde{\Delta}_{j}(\boldsymbol{x}, k)\right]\right| & \leq K_{1}(\boldsymbol{x}) \eta_{\mathrm{e}}^{2} \\
\left|\mathbb{E}\left[\Delta_{i}(\boldsymbol{x}, k) \Delta_{j}(\boldsymbol{x}, k) \Delta_{l}(\boldsymbol{x}, k)-\tilde{\Delta}_{i}(\boldsymbol{x}, k) \tilde{\Delta}_{j}(\boldsymbol{x}, k) \tilde{\Delta}_{l}(\boldsymbol{x}, k)\right]\right| & \leq K_{1}(\boldsymbol{x}) \eta_{\mathrm{e}}^{2}
\end{aligned}
$$

(3). There exists a subset $P$ of the index set $\{1,2, \ldots, D\}$ such that the following holds. Below we use the notations $\|\boldsymbol{x}\|_{\mathrm{P}}:=\sqrt{\sum_{i \in P} x_{i}^{2}}$ and $\|\boldsymbol{x}\|_{\mathrm{R}}:=\sqrt{\sum_{i \notin P} x_{i}^{2}}$.

- There are constants $C_{1}>0, \omega_{1}>0$ (independent of $\eta_{\mathrm{e}}$ ) so that for all $k \leq N$ and $\boldsymbol{x} \in \mathcal{X}_{k}$,

$$
\begin{aligned}
& \|\mathbb{E} \boldsymbol{\Delta}(\boldsymbol{x}, k)\|_{\mathrm{P}} \leq C_{1} \eta_{\mathrm{e}}\left(1+\|\boldsymbol{x}\|_{\mathrm{P}}\right) \\
& \|\mathbb{E} \boldsymbol{\Delta}(\boldsymbol{x}, k)\|_{\mathrm{R}} \leq C_{1} \eta_{\mathrm{e}}\left(1+\|\boldsymbol{x}\|_{\mathrm{P}}^{\omega_{1}}\right)\left(1+\|\boldsymbol{x}\|_{\mathrm{R}}\right) .
\end{aligned}
$$

- For all $m \geq 1$, there are constants $C_{2 m}, \omega_{2 m}>0$ (independent of $\eta_{\mathrm{e}}$ ) so that for all $k \leq N$ and $\boldsymbol{x} \in \mathcal{X}_{k}$,

$$
\begin{aligned}
& \mathbb{E}\|\boldsymbol{\Delta}(\boldsymbol{x}, k)\|_{\mathrm{P}}^{2 m} \leq C_{2 m} \eta_{\mathrm{e}}^{m}\left(1+\|\boldsymbol{x}\|_{\mathrm{P}}^{2 m}\right), \\
& \mathbb{E}\|\boldsymbol{\Delta}(\boldsymbol{x}, k)\|_{\mathrm{R}}^{2 m} \leq C_{2 m} \eta_{\mathrm{e}}^{m}\left(1+\|\boldsymbol{x}\|_{\mathrm{P}}^{\omega_{2 m}}\right)\left(1+\|\boldsymbol{x}\|_{\mathrm{R}}^{2 m}\right) .
\end{aligned}
$$

Then for each function $g \in G^{4}$, there exists a constant $C>0$ (independent of $\eta_{\mathrm{e}}$ ) such that

$$
\begin{equation*}
\max _{0 \leq k \leq N}\left|\mathbb{E}\left[g\left(\boldsymbol{x}_{k}\right)\right]-\mathbb{E}\left[g\left(\boldsymbol{X}_{k \eta_{\mathrm{e}}}\right)\right]\right| \leq C \eta_{\mathrm{e}} \tag{8}
\end{equation*}
$$

when the SDE starts from the same initial point $x_{0}$ as $S G A$. That is, the $S G A$ is a first-order weak approximation of the $S D E$.

To apply Theorem B. 2 , the major condition to verify is that $\Delta$ and $\tilde{\Delta}$ match in moments. The following lemma computes the moments for $\tilde{\Delta}$ :
Lemma B.3. For drift function $b$ and diffusion function $\sigma$ that belong to $G^{4}$, there is a function $K_{2} \in G$ (independent of $\eta_{\mathrm{e}}$ ) such that for all $\boldsymbol{x} \in \mathbb{R}^{D}$ and $1 \leq i, j, l \leq D$,

$$
\begin{aligned}
\left|\mathbb{E}\left[\tilde{\Delta}_{i}(\boldsymbol{x}, k)\right]-\eta_{\mathrm{e}} b_{i}\left(\boldsymbol{x}, k \eta_{\mathrm{e}}\right)\right| & \leq K_{2}(\boldsymbol{x}) \eta_{\mathrm{e}}^{2} \\
\left|\mathbb{E}\left[\tilde{\Delta}_{i}(\boldsymbol{x}, k) \tilde{\Delta}_{j}(\boldsymbol{x}, k)\right]-\eta_{\mathrm{e}}^{2} \sum_{k=1}^{D} \sigma_{i, k}\left(\boldsymbol{x}, k \eta_{\mathrm{e}}\right) \sigma_{j, k}\left(\boldsymbol{x}, k \eta_{\mathrm{e}}\right)\right| & \leq K_{2}(\boldsymbol{x}) \eta_{\mathrm{e}}^{2} \\
\left|\mathbb{E}\left[\tilde{\Delta}_{i}(\boldsymbol{x}, k) \tilde{\Delta}_{j}(\boldsymbol{x}, k) \tilde{\Delta}_{l}(\boldsymbol{x}, k)\right]\right| & \leq K_{2}(\boldsymbol{x}) \eta_{\mathrm{e}}^{2}
\end{aligned}
$$

## B. 1 Proof for Theorem B. 2

To prove Theorem B.2, we need the following lemma from Li et al. (2019). However, the original version of the lemma does not apply to time-dependent SDEs, i.e., $\boldsymbol{b}$ and $\boldsymbol{\sigma}$ cannot change with time $t$. By carefully scrutinizing the proof, we find that the proof is indeed applicable to time-dependent SDEs.
Lemma B. 4 (Adaption of Proposition 25, Li et al. (2019)). Suppose that the drift function $\boldsymbol{b}$ and diffusion function $\boldsymbol{\sigma}$ are Lipschitz and belong to $G^{\alpha}$ for some $\alpha \geq 1$. Let $s \in[0, T]$ and $g \in G^{\alpha}$. For $t \in[s, T]$, define

$$
u(\boldsymbol{x}, s, t):=\mathbb{E}_{\boldsymbol{X}_{t} \sim \mathcal{P}_{X}(\boldsymbol{x}, s, t)}\left[g\left(\boldsymbol{X}_{t}\right)\right],
$$

Then $u(\cdot, s, t) \in G^{\alpha}$ uniformly in $s, t$.
We need the following lemma to bound the growth of $\boldsymbol{x}_{k}$.
Lemma B.5. Under Condition (3) of Theorem B.2, given the initial point $\boldsymbol{x}_{0}$ and any $m \geq 1$, there exists a constant $C_{2 m}^{\prime}>0$ (depending on $\boldsymbol{x}_{0}$ but independent of $\eta_{e}$ ) such that the parameters $\left\{\boldsymbol{x}_{k}\right\}$ of $S G A$ starting from $\boldsymbol{x}_{0}$ can be uniformly bounded by $\mathbb{E}\left[\left\|\boldsymbol{x}_{k}\right\|_{2}^{2 m}\right] \leq C_{2 m}^{\prime}$ for all $k \leq N:=\left\lfloor T / \eta_{\mathrm{e}}\right\rfloor$.

Proof. It suffices to show that both $\mathbb{E}\left[\left(1+\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{P}}^{2}\right)^{m}\right]$ and $\mathbb{E}\left[\left(1+\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{R}}^{2}\right)^{m}\right]$ are uniformly bounded.
First, we show that for $\mathbb{E}\left[\left(1+\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{P}}^{2}\right)^{m}\right]$. For all $2 \leq j \leq 2 m$, by Jensen's inequality,

$$
\mathbb{E}\left[\left\|\boldsymbol{\Delta}\left(\boldsymbol{x}_{k}, k\right)\right\|_{\mathrm{P}}^{j} \mid \boldsymbol{x}_{k}\right] \leq \mathbb{E}\left[\left\|\boldsymbol{\Delta}\left(\boldsymbol{x}_{k}, k\right)\right\|_{\mathrm{P}}^{2 m} \mid \boldsymbol{x}_{k}\right]^{\frac{j}{2 m}} \leq C_{2 m}^{\frac{j}{2 m}} \eta_{\mathrm{e}}^{j / 2}\left(1+\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{P}}^{2 m}\right)^{\frac{j}{2 m}}
$$

Then there exists a constant $\hat{C}_{1}>0$ so that

$$
\begin{equation*}
\mathbb{E}\left[\left\|\boldsymbol{\Delta}\left(\boldsymbol{x}_{k}, k\right)\right\|_{\mathrm{P}}^{j} \mid \boldsymbol{x}_{k}\right] \leq \hat{C}_{1} \eta_{\mathrm{e}}\left(1+\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{P}}^{j}\right) \tag{9}
\end{equation*}
$$

For any two vectors $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^{D}$, we use the notation $\langle\boldsymbol{x}, \boldsymbol{y}\rangle_{\mathrm{P}}:=\sum_{i \in P} x_{i} y_{i}$ to denote the inner product of $\boldsymbol{x}$ and $\boldsymbol{y}$ restricting on coordinates in $P$. Now we expand $\left(1+\left\|\boldsymbol{x}_{k+1}\right\|_{\mathrm{P}}^{2}\right)^{m}$ using the update rule. Let $\delta_{k}:=2\left\langle\boldsymbol{x}_{k}, \boldsymbol{\Delta}\left(\boldsymbol{x}_{k}, k\right)\right\rangle_{\mathrm{P}}+\left\|\boldsymbol{\Delta}\left(\boldsymbol{x}_{k}, k\right)\right\|_{\mathrm{P}}^{2}$. Then by the binomial theorem,

$$
\begin{aligned}
\mathbb{E}\left[\left(1+\left\|\boldsymbol{x}_{k+1}\right\|_{\mathrm{P}}^{2}\right)^{m} \mid \boldsymbol{x}_{k}\right]= & \mathbb{E}\left[\left(1+\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{P}}^{2}+\delta_{k}\right)^{m} \mid \boldsymbol{x}_{k}\right] \\
= & \left(1+\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{P}}^{2}\right)^{m}+m \mathbb{E}\left[\delta_{k} \mid \boldsymbol{x}_{k}\right]\left(1+\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{P}}^{2}\right)^{m-1} \\
& +\sum_{j=2}^{m}\binom{m}{j} \mathbb{E}\left[\delta_{k}^{j} \mid \boldsymbol{x}_{k}\right]\left(1+\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{P}}^{2}\right)^{m-j}
\end{aligned}
$$

By (9), it can be shown that there is a constant $\hat{C}_{2}$ such that $\mathbb{E}\left[\delta_{k}^{j} \mid \boldsymbol{x}_{k}\right] \leq \hat{C}_{2} \eta_{\mathrm{e}}\left(1+\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{P}}^{2}\right)^{j}$ for $j \geq 2$. Then there exists constant $\hat{C}_{3}, \hat{C}_{4}$ such that

$$
\begin{aligned}
\mathbb{E}\left[\left(1+\left\|\boldsymbol{x}_{k+1}\right\|_{\mathrm{P}}^{2}\right)^{m} \mid \boldsymbol{x}_{k}\right] \leq & \left(1+\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{P}}^{2}\right)^{m}+m \mathbb{E}\left[\delta_{k} \mid \boldsymbol{x}_{k}\right]\left(1+\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{P}}^{2}\right)^{m-1} \\
& +\hat{C}_{3} \eta_{\mathrm{e}}\left(1+\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{P}}^{2}\right)^{m} \\
\leq & \left(1+\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{P}}^{2}\right)^{m}+2 m\left\langle\boldsymbol{x}_{k}, \mathbb{E}\left[\boldsymbol{\Delta}\left(\boldsymbol{x}_{k}, k\right) \mid \boldsymbol{x}_{k}\right]\right\rangle_{\mathrm{P}}\left(1+\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{P}}^{2}\right)^{m-1} \\
& +\hat{C}_{4} \eta_{\mathrm{e}}\left(1+\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{P}}^{2}\right)^{m}
\end{aligned}
$$

Recall that $\left\|\mathbb{E}\left[\boldsymbol{\Delta}\left(\boldsymbol{x}_{k}, k\right) \mid \boldsymbol{x}_{k}\right]\right\|_{\mathrm{P}} \leq C_{1} \eta_{\mathrm{e}}\left(1+\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{P}}\right)$ by Condition (3). Thus, there exists a constant $\hat{C}_{5}$ (independent of $\eta_{\mathrm{e}},\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{P}}$ ) such that

$$
\mathbb{E}\left[\left(1+\left\|\boldsymbol{x}_{k+1}\right\|_{\mathrm{P}}^{2}\right)^{m} \mid \boldsymbol{x}_{k}\right] \leq\left(1+\hat{C}_{5} \eta_{\mathrm{e}}\right)\left(1+\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{P}}^{2}\right)^{m}
$$

Taking the expectation over $\boldsymbol{x}_{k}$ gives

$$
\mathbb{E}\left[\left(1+\left\|\boldsymbol{x}_{k+1}\right\|_{\mathrm{P}}^{2}\right)^{m}\right] \leq\left(1+\hat{C}_{5} \eta_{\mathrm{e}}\right) \mathbb{E}\left[\left(1+\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{P}}^{2}\right)^{m}\right]
$$

Then taking a telescoping product proves that for all $k \leq N$,

$$
\begin{aligned}
\mathbb{E}\left[\left(1+\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{P}}^{2}\right)^{m}\right] \leq\left(1+\hat{C}_{5} \eta_{\mathrm{e}}\right)^{k}\left(1+\left\|\boldsymbol{x}_{0}\right\|_{\mathrm{P}}^{2}\right)^{m} & \leq \exp \left(\hat{C}_{5} k \eta_{\mathrm{e}}\right)\left(1+\left\|\boldsymbol{x}_{0}\right\|_{\mathrm{P}}^{2}\right)^{m} \\
& \leq \exp \left(\hat{C}_{5} T\right)\left(1+\left\|\boldsymbol{x}_{0}\right\|_{\mathrm{P}}^{2}\right)^{m}
\end{aligned}
$$

So $\mathbb{E}\left[\left(1+\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{P}}^{2}\right)^{m}\right]$ is uniformly bounded (independent of $\eta_{\mathrm{e}}, k$ ).
Now we show that $\mathbb{E}\left[\left(1+\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{R}}^{2}\right)^{m}\right]$ is uniformly bounded. We can repeat the argument above to bound $\mathbb{E}\left[\left(1+\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{R}}^{2}\right)^{m}\right]$, while utilizing the bounds $\|\mathbb{E} \boldsymbol{\Delta}(\boldsymbol{x}, k)\|_{\mathrm{R}} \leq C_{1} \eta_{\mathrm{e}}\left(1+\|\boldsymbol{x}\|_{\mathrm{P}}^{\omega_{1}}\right)\left(1+\|\boldsymbol{x}\|_{\mathrm{R}}\right)$, $\mathbb{E}\|\boldsymbol{\Delta}(\boldsymbol{x}, k)\|_{\mathrm{R}}^{2 m} \leq C_{2 m} \eta_{\mathrm{e}}^{m}\left(1+\|\boldsymbol{x}\|_{\mathrm{P}}^{\omega_{2 m}}\right)\left(1+\|\boldsymbol{x}\|_{\mathrm{R}}^{2 m}\right)$. In the end we can obtain the following for some real constant $\hat{C}_{6}>0$ and some integer constant $\hat{\omega}>0$ :

$$
\mathbb{E}\left[\left(1+\left\|\boldsymbol{x}_{k+1}\right\|_{\mathrm{R}}^{2}\right)^{m}\right] \leq\left(1+\hat{C}_{6} \eta_{\mathrm{e}} \mathbb{E}\left[1+\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{P}}^{2 \hat{\omega}}\right]\right) \mathbb{E}\left[\left(1+\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{R}}^{2}\right)^{m}\right]
$$

As we have shown, $\mathbb{E}\left[1+\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{P}}^{2 \hat{\omega}}\right]$ is uniformly bounded by a constant. Taking a telescoping product proves that $\mathbb{E}\left[\left(1+\left\|\boldsymbol{x}_{k}\right\|_{\mathrm{R}}^{2}\right)^{m}\right]$ is uniformly bounded (independent of $\eta_{\mathrm{e}}, k$ ).

We also need the following lemma adapted from Lemma C.2, Li et al. (2021).
Lemma B. 6 (Adaption of Lemma C.2, Li et al. (2021)). Let $u_{1}, \ldots, u_{N}$ be a set offunctions that belong to $G^{4}$ uniformly. Under Conditions (2), (3) in Theorem B.2. if band $\boldsymbol{\sigma}$ are Lipschitz, then there exists a function $K_{1}^{\prime} \in G$ (independent of $\eta_{e}$ ) such that

$$
\left|\mathbb{E}\left[u_{j}(\boldsymbol{x}+\boldsymbol{\Delta}(\boldsymbol{x}, k))\right]-\mathbb{E}\left[u_{j}(\boldsymbol{x}+\tilde{\boldsymbol{\Delta}}(\boldsymbol{x}, k))\right]\right| \leq K_{1}^{\prime}(\boldsymbol{x}) \eta_{\mathrm{e}}^{2}
$$

for all $1 \leq j \leq N, 1 \leq k \leq N$ and $\boldsymbol{x} \in \mathcal{X}$.
Proof. Since $u_{1}, \ldots, u_{N} \in G^{4}$ uniformly, we can find $K_{0} \in G$ such that $u_{j}(\boldsymbol{x})$ is bounded by $K_{0}(\boldsymbol{x})$ and so are all the partial derivatives of $u_{j}$ up to order 4.
By Taylor's Theorem with Lagrange Remainder, for all $1 \leq j \leq N, 1 \leq k \leq N$ we have

$$
\begin{aligned}
u_{j}(\boldsymbol{x} & +\boldsymbol{\Delta}(\boldsymbol{x}, k))-u_{j}(\boldsymbol{x}+\tilde{\boldsymbol{\Delta}}(\boldsymbol{x}, k)) \\
& =\underbrace{\sum_{s=1}^{3} \frac{1}{s!} \sum_{1 \leq i_{1}, \ldots, i_{s} \leq D} \frac{\partial^{s} u_{j}}{\partial x_{i_{1}} \cdots \partial x_{i_{s}}}(\boldsymbol{x})\left(\prod_{r=1}^{s} \Delta_{i_{r}}(\boldsymbol{x}, k)-\prod_{r=1}^{s} \tilde{\Delta}_{i_{r}}(\boldsymbol{x}, k)\right)}_{=: M_{j}}+R_{j}-\tilde{R}_{j},
\end{aligned}
$$

where the remainders $R_{j}, \tilde{R}_{j}$ are

$$
\begin{aligned}
& R_{j}:=\frac{1}{4!} \sum_{1 \leq i_{1}, \ldots, i_{4} \leq D} \frac{\partial^{4} u_{j}}{\partial x_{i_{1}} \cdots \partial x_{i_{4}}}(\boldsymbol{x}+a \boldsymbol{\Delta}(\boldsymbol{x}, k)) \prod_{r=1}^{4} \Delta_{i_{r}}(\boldsymbol{x}, k) . \\
& \tilde{R}_{j}:=\frac{1}{4!} \sum_{1 \leq i_{1}, \ldots, i_{4} \leq D} \frac{\partial^{4} u_{j}}{\partial x_{i_{1}} \cdots \partial x_{i_{4}}}(\boldsymbol{x}+\tilde{a} \tilde{\boldsymbol{\Delta}}(\boldsymbol{x}, k)) \prod_{r=1}^{4} \tilde{\Delta}_{i_{r}}(\boldsymbol{x}, k) .
\end{aligned}
$$

for some $a, \tilde{a} \in[0,1]$.
By Condition (2), the expectation of $M_{j}$ can be bounded by

$$
\mathbb{E}\left[M_{j}\right] \leq \sum_{s=1}^{3} \frac{1}{s!} \sum_{1 \leq i_{1}, \ldots, i_{s} \leq D}\left|\frac{\partial^{s} u_{j}}{\partial x_{i_{1}} \cdots \partial x_{i_{s}}}(\boldsymbol{x})\right| \cdot K_{1}(\boldsymbol{x}) \eta_{\mathrm{e}}^{2} \leq \sum_{s=1}^{3} \frac{D^{s}}{s!} K_{0}(\boldsymbol{x}) K_{1}(\boldsymbol{x}) \eta_{\mathrm{e}}^{2}
$$

so $\frac{1}{\eta_{e}^{2}} \mathbb{E}\left[M_{j}\right]$ is uniformly bounded by a function in $G$.
Now let $\kappa_{0}, m$ be the constants so that $K_{0}(\boldsymbol{x})^{2} \leq \kappa_{0}^{2}\left(1+\|\boldsymbol{x}\|_{2}^{2 m}\right)$. For $R_{j}$, by Cauchy-Schwarz inequality we have

$$
\begin{aligned}
\mathbb{E}\left[R_{j}\right] & \leq \frac{1}{4!}\left(\sum_{i_{1}, \ldots, i_{4}} \mathbb{E}\left[\left|\frac{\partial^{4} u_{j}}{\partial x_{i_{1}} \cdots \partial x_{i_{4}}}(\boldsymbol{x}+a \boldsymbol{\Delta}(\boldsymbol{x}, k))\right|^{2}\right]\right)^{1 / 2} \cdot\left(\sum_{i_{1}, \ldots, i_{4}} \mathbb{E}\left[\left|\prod_{r=1}^{4} \Delta_{i_{r}}(\boldsymbol{x}, k)\right|^{2}\right]\right)^{1 / 2} \\
& \leq \frac{1}{4!}\left(\sum_{i_{1}, \ldots, i_{4}} \mathbb{E}\left[\left|\frac{\partial^{4} u_{j}}{\partial x_{i_{1} \cdots \partial x_{i_{4}}}}(\boldsymbol{x}+a \boldsymbol{\Delta}(\boldsymbol{x}, k))\right|^{2}\right]\right)^{1 / 2} \cdot \mathbb{E}\left[\|\boldsymbol{\Delta}(\boldsymbol{x}, k)\|_{2}^{8}\right]^{1 / 2} \\
& \leq \frac{1}{4!}\left(D^{2} \cdot K_{0}(\boldsymbol{x}+a \boldsymbol{\Delta}(\boldsymbol{x}, k))\right) \cdot\left(\eta_{\mathrm{e}}^{4} K_{8}(\boldsymbol{x})\right)^{1 / 2}
\end{aligned}
$$

where the last line uses Condition (3) and $K_{8}$ is a function of polynomial growth. For $K_{0}(\boldsymbol{x}+$ $a \boldsymbol{\Delta}(\boldsymbol{x}, k)$ ), we can bound its expectation by

$$
\begin{aligned}
\mathbb{E}\left[K_{0}(\boldsymbol{x}+a \boldsymbol{\Delta}(\boldsymbol{x}, k))\right] & \leq \kappa_{0} \mathbb{E}\left[1+\|\boldsymbol{x}+a \boldsymbol{\Delta}(\boldsymbol{x}, k)\|_{2}^{2 m}\right]^{1 / 2} \\
& \leq \kappa_{0}\left(1+2^{2 m-1} \mathbb{E}\left[\|\boldsymbol{x}\|_{2}^{2 m}+\mathbb{E}\|\boldsymbol{\Delta}(\boldsymbol{x}, k)\|_{2}^{2 m}\right]\right)^{1 / 2} \\
& \leq \kappa_{0}\left(1+2^{2 m-1}\left(\|\boldsymbol{x}\|_{2}^{2 m}+C_{2 m} \eta_{\mathrm{e}}^{2 m}\left(1+\|\boldsymbol{x}\|_{2}^{2 m}\right)\right)\right)^{1 / 2}
\end{aligned}
$$

Combining this with our bound for $\mathbb{E}\left[R_{j}\right]$ proves that $\frac{1}{\eta_{e}^{2}} \mathbb{E}\left[R_{j}\right]$ is uniformly bounded by a function in $G$ :

$$
\mathbb{E}\left[R_{j}\right] \leq \eta_{\mathrm{e}}^{2} \cdot \frac{1}{4!} \cdot D^{4} \cdot \kappa_{0}\left(1+2^{2 m-1}\left(\|\boldsymbol{x}\|_{2}^{2 m}+C_{2 m} \eta_{\mathrm{e}}^{2 m}\left(1+\|\boldsymbol{x}\|_{2}^{2 m}\right)\right)\right)^{1 / 2} \cdot K_{8}^{1 / 2}(\boldsymbol{x})
$$

Then we can repeat the above argument for $R_{j}$ while replacing $\Delta$ as $\tilde{\Delta}$, and conclude that $\frac{1}{\eta_{e}^{2}} \mathbb{E}\left[\tilde{R}_{j}\right]$ is also uniformly bounded by a function in $G$. To do so, we note that $\boldsymbol{b}, \boldsymbol{\sigma}$ are Lipschitz, and apply a similar argument as in Lemma 26 of Li et al. (2019) to show that for all $s \geq 1$ there exists a function $\tilde{K} \in G$ such that for all $1 \leq i_{1}, \ldots, i_{s} \leq D$,

$$
\mathbb{E}\left[\left|\prod_{r=1}^{s} \tilde{\Delta}_{i_{r}}(\boldsymbol{x}, k)\right|\right] \leq \tilde{K}(\boldsymbol{x}) \eta_{\mathrm{e}}^{s}
$$

Finally, we can find a function $K_{1}^{\prime} \in G$ such that $\mathbb{E}\left[M_{j}\right] \leq \frac{1}{3} K_{1}^{\prime}(\boldsymbol{x}) \eta_{\mathrm{e}}^{2}, \mathbb{E}\left[R_{j}\right] \leq \frac{1}{3} K_{1}^{\prime}(\boldsymbol{x}) \eta_{\mathrm{e}}^{2}$, $\mathbb{E}\left[\tilde{R}_{j}\right] \leq \frac{1}{3} K_{1}^{\prime}(\boldsymbol{x}) \eta_{\mathrm{e}}^{2}$. Then $\mathbb{E}\left[u_{j}(\boldsymbol{x}+\boldsymbol{\Delta}(\boldsymbol{x}, k))\right]-\mathbb{E}\left[u_{j}(\boldsymbol{x}+\tilde{\boldsymbol{\Delta}}(\boldsymbol{x}, k))\right] \leq K_{1}^{\prime}(\boldsymbol{x}) \eta_{\mathrm{e}}^{2}$.

Now we are ready to present our proof for TheoremC. 2
Prooffor Theorem B.2 For $0 \leq j \leq k$, let $\hat{\boldsymbol{x}}_{j, k}$ be a random variable that is distributed as the probability distribution $\mathcal{P}_{X}\left(\boldsymbol{x}_{j}, j \eta_{\mathrm{e}}, k \eta_{\mathrm{e}}\right)$ conditioned on $\boldsymbol{x}_{j}$. By definition, $\operatorname{Pr}\left[\hat{\boldsymbol{x}}_{k, k}=\boldsymbol{x}_{k}\right]=1$, $\hat{\boldsymbol{x}}_{0, k} \sim \boldsymbol{X}_{k \eta_{\mathrm{e}}}$. Let $u(\boldsymbol{x}, s, t):=\mathbb{E}_{\boldsymbol{X}_{t} \sim \mathcal{P}_{X}(\boldsymbol{x}, s, t)}\left[g\left(\boldsymbol{X}_{t}\right)\right]$. Then we can do the following decomposition for $\mathbb{E}\left[g\left(\boldsymbol{x}_{k}\right)\right]-\mathbb{E}\left[g\left(\boldsymbol{X}_{k \eta_{\mathrm{e}}}\right)\right]$ :

$$
\begin{aligned}
\left|\mathbb{E}\left[g\left(\boldsymbol{x}_{k}\right)\right]-\mathbb{E}\left[g\left(\boldsymbol{X}_{k \eta_{\mathrm{e}}}\right)\right]\right| & =\sum_{j=0}^{k-1}\left(\mathbb{E}\left[g\left(\hat{\boldsymbol{x}}_{j+1, k}\right)\right]-\mathbb{E}\left[g\left(\hat{\boldsymbol{x}}_{j, k}\right)\right]\right) \\
& =\sum_{j=0}^{k-1}\left(\mathbb{E}\left[u\left(\hat{\boldsymbol{x}}_{j+1, j+1},(j+1) \eta_{\mathrm{e}}, k \eta_{\mathrm{e}}\right)\right]-\mathbb{E}\left[u\left(\hat{\boldsymbol{x}}_{j, j+1},(j+1) \eta_{\mathrm{e}}, k \eta_{\mathrm{e}}\right)\right]\right) .
\end{aligned}
$$

Taking absolute values gives

$$
\left|\mathbb{E}\left[g\left(\boldsymbol{x}_{k}\right)\right]-\mathbb{E}\left[g\left(\boldsymbol{X}_{k \eta_{\mathrm{e}}}\right)\right]\right| \leq \sum_{j=0}^{k-1}\left|\mathbb{E}\left[u\left(\hat{\boldsymbol{x}}_{j+1, j+1},(j+1) \eta_{\mathrm{e}}, k \eta_{\mathrm{e}}\right)\right]-\mathbb{E}\left[u\left(\hat{\boldsymbol{x}}_{j, j+1},(j+1) \eta_{\mathrm{e}}, k \eta_{\mathrm{e}}\right)\right]\right|
$$

Let $u_{j+1}(\boldsymbol{x}):=u\left(\boldsymbol{x},(j+1) \eta_{\mathrm{e}}, k \eta_{\mathrm{e}}\right)$. Note that $\hat{\boldsymbol{x}}_{j+1, j+1} \sim \boldsymbol{x}_{j}+\boldsymbol{\Delta}\left(\boldsymbol{x}_{j}, j\right), \hat{\boldsymbol{x}}_{j, j+1} \sim \boldsymbol{x}_{j}+\tilde{\boldsymbol{\Delta}}\left(\boldsymbol{x}_{j}, j\right)$. We can rewrite the above formula as

$$
\left|\mathbb{E}\left[g\left(\boldsymbol{x}_{k}\right)\right]-\mathbb{E}\left[g\left(\boldsymbol{X}_{k \eta_{\mathrm{e}}}\right)\right]\right| \leq \sum_{j=0}^{k-1}\left|\mathbb{E}\left[u_{j+1}\left(\boldsymbol{x}_{j}+\boldsymbol{\Delta}\left(\boldsymbol{x}_{j}, j\right)\right)\right]-\mathbb{E}\left[u_{j+1}\left(\boldsymbol{x}_{j}+\tilde{\boldsymbol{\Delta}}\left(\boldsymbol{x}_{j}, j\right)\right)\right]\right|
$$

By Lemma B. $4, u \in G^{4}$ uniformly in $s, t$, so $u_{1}, \ldots, u_{N} \in G^{4}$ uniformly. Then by LemmaB.6, we know that there exists a function $K_{1}^{\prime}(\boldsymbol{x})=\kappa_{1}\left(1+\|\boldsymbol{x}\|_{2}^{2 m}\right) \in G$ such that

$$
\left|\mathbb{E}\left[u_{j+1}\left(\boldsymbol{x}_{j}+\boldsymbol{\Delta}\left(\boldsymbol{x}_{j}, j\right)\right)\right]-\mathbb{E}\left[u_{j+1}\left(\boldsymbol{x}_{j}+\tilde{\boldsymbol{\Delta}}\left(\boldsymbol{x}_{j}, j\right)\right)\right]\right| \leq \mathbb{E}\left[K_{1}^{\prime}\left(\boldsymbol{x}_{j}\right) \eta_{\mathrm{e}}^{2}\right],
$$

for all $0 \leq j<N$. Combining this with Lemma B.5. we can bound $\left|\mathbb{E}\left[g\left(\boldsymbol{x}_{k}\right)\right]-\mathbb{E}\left[g\left(\boldsymbol{X}_{k \eta_{\mathrm{e}}}\right)\right]\right|$ by

$$
\begin{aligned}
\left|\mathbb{E}\left[g\left(\boldsymbol{x}_{k}\right)\right]-\mathbb{E}\left[g\left(\boldsymbol{X}_{k \eta_{\mathrm{e}}}\right)\right]\right| \leq \sum_{j=0}^{k-1} \mathbb{E}\left[K_{1}^{\prime}\left(\boldsymbol{x}_{j}\right) \eta_{\mathrm{e}}^{2}\right] & \leq \eta_{\mathrm{e}}^{2} \sum_{j=0}^{k-1} \mathbb{E}\left[\kappa_{1}\left(1+\left\|\boldsymbol{x}_{j}\right\|_{2}^{2 m}\right)\right] \\
& \leq \eta_{\mathrm{e}}^{2} \sum_{j=0}^{k-1} \kappa_{1}\left(1+C_{2 m}^{\prime}\right) \leq \kappa_{1}\left(1+C_{2 m}^{\prime}\right) T \eta_{\mathrm{e}}
\end{aligned}
$$

We can complete the proof by noting that $\kappa_{1}, C_{2 m}^{\prime}, T$ are independent of $\eta_{\mathrm{e}}$.

## B. 2 Proof for Lemma B. 3

To prove Lemma B.3. we only need to verify the following lemma using Itô-Taylor expansion.
Lemma B.7. Let $\psi: \mathbb{R}^{D} \rightarrow \mathbb{R}$ be a function in $G^{4}$. Define

$$
\mathcal{A}_{t} \psi(\boldsymbol{x}):=\sum_{i \in[D]} b_{i}(\boldsymbol{x}, t) \partial_{i} \psi(\boldsymbol{x})+\frac{1}{2} \sum_{i, j \in[D]}\left(\sum_{l \in[D]} \sigma_{i, l}(\boldsymbol{x}, t) \sigma_{l, j}(\boldsymbol{x}, t)\right) \partial_{i, j}^{2} \psi(\boldsymbol{x})
$$

Then there exists a function $\hat{K} \in G$ such that

$$
\begin{equation*}
\left|\mathbb{E}[\psi(\boldsymbol{x}+\tilde{\boldsymbol{\Delta}}(\boldsymbol{x}, k))]-\psi(\boldsymbol{x})-\eta_{\mathrm{e}} \mathcal{A}_{k \eta_{\mathrm{e}}} \psi(\boldsymbol{x})\right| \leq \hat{K}(\boldsymbol{x}) \eta_{\mathrm{e}}^{2} \tag{10}
\end{equation*}
$$

Proof for Lemma B.3. We can prove the lemma by applying Lemma B. 7 with $\psi(\boldsymbol{x})$ being $\psi(\tilde{\boldsymbol{x}})=$ $\prod_{r=1}^{s}\left(\tilde{x}_{i_{r}}-x_{i_{r}}\right)$ for any tuple $\left(i_{1}, \ldots, i_{s}\right) \in[D]^{s}$ with $s \leq 3$ elements.

Proof for Lemma B.7. WLOG we prove the case of $k=0$, then all the other cases can be proved by shifting the time. Let $\Lambda_{t} \psi(\boldsymbol{x}):=\boldsymbol{\sigma}(\boldsymbol{x}, t)^{\top} \nabla \psi(\boldsymbol{x})$.

$$
\psi\left(\boldsymbol{X}_{\eta}\right)=\psi(\boldsymbol{x})+\int_{0}^{\eta_{\mathrm{e}}} \mathcal{A}_{s} \psi\left(\boldsymbol{X}_{s}\right) \mathrm{d} s+\int_{0}^{\eta_{\mathrm{e}}}\left\langle\Lambda_{s} \psi\left(\boldsymbol{X}_{s}\right), \mathrm{d} \boldsymbol{W}_{s}\right\rangle
$$

Now we further apply the above formula to $\mathcal{A}_{s} \psi\left(\boldsymbol{X}_{s}\right)$. Then we have

$$
\begin{aligned}
\psi\left(\boldsymbol{X}_{\eta}\right)=\psi & \psi(\boldsymbol{x})+\int_{0}^{\eta_{e}}\left(\mathcal{A}_{s} \psi(\boldsymbol{x})+\int_{0}^{s} \mathcal{A}_{r} \mathcal{A}_{s} \psi\left(\boldsymbol{X}_{r}\right) \mathrm{d} r+\int_{0}^{s}\left\langle\Lambda_{r} \mathcal{A}_{s} \psi\left(\boldsymbol{X}_{r}\right), \mathrm{d} \boldsymbol{W}_{r}\right\rangle\right) \mathrm{d} s \\
& +\int_{0}^{\eta_{\mathrm{e}}}\left\langle\Lambda_{s} \psi\left(\boldsymbol{X}_{s}\right), \mathrm{d} \boldsymbol{W}_{s}\right\rangle \\
= & \psi(\boldsymbol{x})+\int_{0}^{\eta_{\mathrm{e}}} \mathcal{A}_{s} \psi(\boldsymbol{x}) \mathrm{d} s+\int_{0}^{\eta_{e}} \int_{0}^{s} \mathcal{A}_{r} \mathcal{A}_{s} \psi\left(\boldsymbol{X}_{r}\right) \mathrm{d} r \mathrm{~d} s \\
& +\int_{0}^{\eta_{\mathrm{e}}} \int_{0}^{s}\left\langle\Lambda_{r} \mathcal{A}_{s} \psi\left(\boldsymbol{X}_{r}\right), \mathrm{d} \boldsymbol{W}_{r}\right\rangle \mathrm{d} s+\int_{0}^{\eta_{e}}\left\langle\Lambda_{s} \psi\left(\boldsymbol{X}_{s}\right), \mathrm{d} \boldsymbol{W}_{s}\right\rangle
\end{aligned}
$$

Taking expectation, the last two integrals vanish. So we have

$$
\mathbb{E} \psi\left(\boldsymbol{X}_{\eta}\right)=\psi(\boldsymbol{x})+\int_{0}^{\eta_{e}} \mathcal{A}_{s} \psi(\boldsymbol{x}) \mathrm{d} s+\int_{0}^{\eta_{e}} \int_{0}^{s} \mathbb{E}\left[\mathcal{A}_{r} \mathcal{A}_{s} \psi\left(\boldsymbol{X}_{r}\right)\right] \mathrm{d} r \mathrm{~d} s
$$

By Lipschitzness of $\boldsymbol{b}$ and $\boldsymbol{\sigma}, \frac{1}{\eta_{\mathrm{e}}}\left(\mathcal{A}_{s} \psi(\boldsymbol{x})-\mathcal{A}_{0} \psi(\boldsymbol{x})\right)$ is bounded by a function of $\boldsymbol{x}$ with polynomial growth. Also, $\mathcal{A}_{r} \mathcal{A}_{s} \psi(\cdot)$ is in $G$ uniformly, then Theorem 19 in (Li et al. 2019) implies that $\mathbb{E}\left[\mathcal{A}_{r} \mathcal{A}_{s} \psi\left(\boldsymbol{X}_{r}\right)\right]$ is also in $G$ uniformly. Then we know that there exists two functions $\hat{K}_{1}, \hat{K}_{2} \in G$ such that

$$
\begin{aligned}
\left|\mathbb{E}\left[\psi\left(\boldsymbol{X}_{\eta_{\mathrm{e}}}\right)\right]-\psi(\boldsymbol{x})-\eta_{\mathrm{e}} \mathcal{A}_{0} \psi(\boldsymbol{x})\right| & \leq \int_{0}^{\eta_{\mathrm{e}}} \eta_{\mathrm{e}} \hat{K}_{1}(\boldsymbol{x}) \mathrm{d} s+\int_{0}^{\eta_{\mathrm{e}}} \int_{0}^{s} \hat{K}_{2}(\boldsymbol{x}) \mathrm{d} r \mathrm{~d} s \\
& \leq \eta_{\mathrm{e}}^{2}\left(\hat{K}_{1}(\boldsymbol{x})+\hat{K}_{2}(\boldsymbol{x})\right)
\end{aligned}
$$

We can conclude the proof by setting $\hat{K}(\boldsymbol{x}):=\hat{K}_{1}(\boldsymbol{x})+\hat{K}_{2}(\boldsymbol{x})$.

## C RMSProp SDE Proof

In this section, we prove the theorem for the SDE approximation of RMSprop.
Definition C. 1 (SDE for RMSprop, matrix form). For constants $\sigma_{0}, \epsilon_{0}$, and $c_{2}$, define the SDE as $\mathrm{d} \boldsymbol{X}_{t}=\boldsymbol{b}\left(\boldsymbol{X}_{t}\right) \mathrm{d} t+\boldsymbol{\sigma}\left(\boldsymbol{X}_{t}\right) \mathrm{d} \boldsymbol{W}_{t}$, where $\boldsymbol{X}_{t} \in \mathbb{R}^{d} \times \mathbb{R}^{d}$, and $\boldsymbol{b}$ and $\boldsymbol{\sigma}$ are defined by

$$
\begin{aligned}
b_{i}(\boldsymbol{\theta}, \boldsymbol{u}) & :=-\frac{1}{\sigma_{0} \sqrt{u_{i}}+\epsilon_{0}} \cdot \partial_{i} f(\boldsymbol{\theta}), & b_{d+i}(\boldsymbol{\theta}, \boldsymbol{u}) & :=c_{2}\left(\Sigma(\boldsymbol{\theta})_{i, i}-u_{i}\right) . \\
\sigma_{i, j}(\boldsymbol{\theta}, \boldsymbol{u}) & :=\frac{1}{\sqrt{u_{i}}+\epsilon_{0} / \sigma_{0}} \cdot\left(\boldsymbol{\Sigma}^{1 / 2}(\boldsymbol{\theta})\right)_{i, j}, & \sigma_{i, d+j}(\boldsymbol{\theta}, \boldsymbol{u}) & :=0, \\
\sigma_{d+i, j}(\boldsymbol{\theta}, \boldsymbol{u}) & :=0, & \sigma_{d+i, d+j}(\boldsymbol{\theta}, \boldsymbol{u}) & :=0,
\end{aligned}
$$

for all $1 \leq i, j \leq d$.

Theorem C.2. Fix constants $\sigma_{0}, c_{2}>0, \epsilon_{0} \geq 0$. Let $T>0, \eta^{2} \in\left(0,1 \wedge T \wedge \frac{1}{2 c_{2}}\right)$ and set $N=\left\lfloor T / \eta^{2}\right\rfloor$. Let $\boldsymbol{u}_{k} \triangleq \boldsymbol{v}_{k} / \sigma^{2}$ and $\boldsymbol{x}_{k} \triangleq\left(\boldsymbol{\theta}_{k}, \boldsymbol{u}_{k}\right) \in \mathbb{R}^{2 d}$. Let $\left\{\boldsymbol{x}_{k}: k \geq 0\right\}$ be the discrete RMSprop iterations defined in Definition 2.1. where $\sigma, \epsilon, \beta$ are set so that $\sigma_{0}=\sigma \eta, \epsilon_{0}=\epsilon \eta$ and $c_{2}=(1-\beta) / \eta^{2}$. For well-behaved NGOS that satisfies the bounded moments and low skewness conditions, the SDE as defined in Definition C.1] is an order-1 weak approximation (Definition 2.4) of discrete RMSprop, if they start with $\boldsymbol{X}_{0}=\boldsymbol{x}_{0}$.

The basic idea is to apply the general theorem (Theorem B.2) with $\eta_{\mathrm{e}}:=\eta^{2}$. However, the SDE above does not satisfy Condition (1) in Theorem B.2 because the denominators such as $\sigma_{0} \sqrt{u_{i}}+\epsilon_{0}$ can be unbounded. To solve this issue, the first step is to reduce Theorem C. 2 to proving the order- 1 weak approximation for the following auxiliary SDE:
Definition C.3. Define $\tau: \mathbb{R} \rightarrow \mathbb{R}$ to be the following smooth transition function:

$$
\tau(z)= \begin{cases}1 & \text { if } z \geq 1  \tag{11}\\ \frac{e^{-1 / z}}{e^{-1 / z}+e^{-1 /(1-z)}} & \text { if } z \in(0,1) \\ 0 & \text { if } z \leq 0\end{cases}
$$

Definition C. 4 (Auxiliary SDE for RMSprop, matrix form). For constants $\sigma_{0}, \epsilon_{0}, c_{2}$ and $u_{\text {min }}$, define $\mu(u)$ as the following function

$$
\begin{equation*}
\mu(u):=\frac{1}{2} u_{\min }+\tau\left(\frac{2 u}{u_{\min }}-1\right) \cdot\left(u-\frac{1}{2} u_{\min }\right) \tag{12}
\end{equation*}
$$

and define the SDE as $\mathrm{d} \boldsymbol{X}_{t}=\boldsymbol{b}\left(\boldsymbol{X}_{t}\right) \mathrm{d} t+\boldsymbol{\sigma}\left(\boldsymbol{X}_{t}\right) \mathrm{d} \boldsymbol{W}_{t}$, where $\boldsymbol{X}_{t} \in \mathbb{R}^{d} \times \mathbb{R}^{d}$, and $\boldsymbol{b}$ and $\boldsymbol{\sigma}$ are defined by

$$
\begin{aligned}
b_{i}(\boldsymbol{\theta}, \boldsymbol{u}) & :=-\frac{1}{\sigma_{0} \sqrt{\mu\left(u_{i}\right)}+\epsilon_{0}} \cdot \partial_{i} f(\boldsymbol{\theta}), & b_{d+i}(\boldsymbol{\theta}, \boldsymbol{u}) & :=c_{2}\left(\Sigma(\boldsymbol{\theta})_{i, i}-u_{i}\right) \\
\sigma_{i, j}(\boldsymbol{\theta}, \boldsymbol{u}) & :=\frac{1}{\sqrt{\mu\left(u_{i}\right)}+\epsilon_{0} / \sigma_{0}} \cdot\left(\boldsymbol{\Sigma}(\boldsymbol{\theta})^{1 / 2}\right)_{i, j}, & \sigma_{i, d+j}(\boldsymbol{\theta}, \boldsymbol{u}) & :=0 \\
\sigma_{d+i, j}(\boldsymbol{\theta}, \boldsymbol{u}) & :=0, & \sigma_{d+i, d+j}(\boldsymbol{\theta}, \boldsymbol{u}) & :=0
\end{aligned}
$$

for all $1 \leq i, j \leq d$.
Theorem C.5. In the setting of Theorem $C .2$, let $u_{\min }=2^{-c_{2} T} \min _{i \in[d]} u_{0, i}$ then the SDE defined by Definition C.4 is an order-1 weak approximation (Definition 2.4) of discrete RMSprop (Definition 2.1), if they start with $\boldsymbol{X}_{0}=\boldsymbol{x}_{0}$.

Proof for Theorem C. 2 Given Theorem C. 5 we only need to show that $\boldsymbol{X}_{t}$ has the same distribution in the original and auxiliary SDEs for all $t \in[0, T]$, when $\boldsymbol{X}_{0}=\left(\boldsymbol{\theta}_{0}, \boldsymbol{u}_{0}\right)$. To see this, we only need to note that in the original SDE

$$
\frac{\mathrm{d} u_{t, i}}{\mathrm{~d} t}=c_{2}\left(\Sigma\left(\boldsymbol{\theta}_{t}\right)_{i, i}-u_{t, i}\right) \geq-c_{2} u_{t, i}
$$

Thus, $u_{t, i} \geq \exp \left(-c_{2} t\right) u_{0, i} \geq u_{\min }$, which means $\operatorname{Pr}\left[u_{t, i}=\mu\left(u_{t, i}\right)\right]=1$ for all $t \in[0, T]$.

It remains to prove TheoremC.5by applying TheoremB.2. In the rest of this section, we verify the three conditions in Theorem B. 2 respectively.
Below we use the notations $\boldsymbol{x}_{k}, \boldsymbol{X}_{t}, \boldsymbol{b}, \boldsymbol{\sigma}$ defined as in Theorem C. 5 Let $D=2 d$. Every $\boldsymbol{x}_{k} \in \mathbb{R}^{D}$ is a concatenation of two $\mathbb{R}^{d}$-vectors $\boldsymbol{\theta}_{k}$ and $\boldsymbol{u}_{k}$. According to the update rule of RMSprop, $\boldsymbol{x}_{k}$ can be seen as SGA $\boldsymbol{x}_{k+1}=\boldsymbol{x}_{k}-\eta_{\mathrm{e}} \boldsymbol{h}_{k}\left(\boldsymbol{x}_{k}, \boldsymbol{z}_{k}, \eta_{\mathrm{e}}\right)$, where $\eta_{\mathrm{e}}=\eta^{2}, \boldsymbol{z}_{k} \sim \mathcal{Z}_{\sigma}\left(\boldsymbol{\theta}_{k}\right)$, and $\boldsymbol{h}_{k}$ is defined below:

$$
\boldsymbol{h}_{k}\left(\boldsymbol{\theta}, \boldsymbol{u}, \boldsymbol{z}, \eta_{\mathrm{e}}\right):=\left[\begin{array}{c}
-(\nabla f(\boldsymbol{\theta})+\sigma \boldsymbol{z}) \odot\left(\sigma_{0} \sqrt{\boldsymbol{u}}+\epsilon_{0}\right)^{-1} \\
c_{2}\left((\nabla f(\boldsymbol{\theta}) / \sigma+\boldsymbol{z})^{2}-\boldsymbol{u}\right)
\end{array}\right] .
$$

We define $\boldsymbol{\Delta}$ and $\tilde{\boldsymbol{\Delta}}$ as in (6) and (7). Fix $\boldsymbol{x}_{0}=\left(\boldsymbol{\theta}_{0}, \boldsymbol{u}_{0}\right)$ with $u_{0, j}>0$ for all $j \in[d]$. Define $u_{\text {min }}$ as in TheoremC.5 Let $\mathcal{X}_{k}$ be the support of the random variable $\boldsymbol{x}_{k}$ given $\boldsymbol{x}_{0}$, then it is easy to show that $\mathcal{X}_{k}$ is a subset of $\left\{(\boldsymbol{\theta}, \boldsymbol{u}): u_{j} \geq u_{\text {min }}\right.$ for all $\left.j \in[d]\right\}$.

## C. 1 Verifying Condition (1)

Lemma C.6. The drift function $\boldsymbol{b}$ and diffusion function $\boldsymbol{\sigma}$ are Lipschitz and belong to $G^{4}$.
Proof. $\boldsymbol{\Sigma}^{1 / 2}(\boldsymbol{\theta})$ is bounded and Lipschitz, so $\boldsymbol{\Sigma}(\boldsymbol{\theta})$ is Lipschitz. Note that the denominators in the fractions in the formulas of $\boldsymbol{b}$ and $\boldsymbol{\sigma}$ are always lower bounded by a constant. Then the Lipschitz property of $\boldsymbol{b}$ and $\boldsymbol{\sigma}$ can be implied by the Lipschitz property of $\nabla f(\boldsymbol{\theta}), \boldsymbol{\Sigma}(\boldsymbol{\theta}), \boldsymbol{\Sigma}^{1 / 2}(\boldsymbol{\theta})$, and $\boldsymbol{b}, \boldsymbol{\sigma} \in G^{4}$ can be implied by $\nabla f, \boldsymbol{\Sigma}^{1 / 2}, \mu \in G^{4}$.

## C. 2 Verifying Condition (2)

To verify Condition (2), we only need to compute the moments of $\boldsymbol{\Delta}$ and $\tilde{\Delta}$ for the discrete RMSprop and the auxiliary SDE, and show that they are close to each other. We compute them by the following two lemmas.
Lemma C.7. For $\boldsymbol{x}=(\boldsymbol{\theta}, \boldsymbol{u}) \in \mathcal{X}_{k}$, if the $N G O S$ is well-behaved and $\mathcal{Z}_{\sigma}$ satisfies the bounded moments and low skewness condition, then the moments of $\boldsymbol{\Delta}:=\boldsymbol{\Delta}(\boldsymbol{x}, k)-\boldsymbol{x}$ can be written as below.

1. For $1 \leq i \leq d$, the following holds for the first moments:

$$
\begin{aligned}
\mathbb{E}\left[\Delta_{i}\right]=-\frac{\eta^{2}}{\sigma_{0} \sqrt{u_{i}}+\epsilon_{0}} \partial_{i} f(\boldsymbol{\theta}), \quad \mathbb{E}\left[\Delta_{d+i}\right] & =\eta^{2} c_{2}\left(\Sigma_{i i}(\boldsymbol{\theta})-u_{i}\right)+\frac{\eta^{4} c_{2}}{\sigma_{0}^{2}}\left(\partial_{i} f(\boldsymbol{\theta})\right)^{2} \\
& =\eta^{2} c_{2}\left(\Sigma_{i i}(\boldsymbol{\theta})-u_{i}\right)+\mathcal{O}\left(\eta^{4}\right) .
\end{aligned}
$$

2. For $1 \leq i, j \leq d$, the following holds for the second moments:

$$
\begin{aligned}
\mathbb{E}\left[\Delta_{i} \Delta_{j}\right] & =\frac{\eta^{2} \Sigma_{i j}(\boldsymbol{\theta})}{\left(\sqrt{u_{i}}+\epsilon_{0} / \sigma_{0}\right)\left(\sqrt{u_{j}}+\epsilon_{0} / \sigma_{0}\right)}+\mathcal{O}\left(\eta^{4}\right) & \mathbb{E}\left[\Delta_{i} \Delta_{d+j}\right] & =\mathcal{O}\left(\eta^{4}\right) \\
\mathbb{E}\left[\Delta_{d+i} \Delta_{j}\right] & =\mathcal{O}\left(\eta^{4}\right) & \mathbb{E}\left[\Delta_{d+i} \Delta_{d+j}\right] & =\mathcal{O}\left(\eta^{4}\right)
\end{aligned}
$$

for all $i, j \in[d]$.
3. The third moments are bounded by $\mathbb{E}\left[\boldsymbol{\Delta}^{\otimes 3}\right]=\mathcal{O}\left(\eta^{4}\right)$.

Here the big- $O$ notation $\mathcal{O}(\cdot)$ is used in a way that $\mathcal{O}(1)$ hides constants (independent of $\eta$ and $\boldsymbol{x}$ ) and values that are bounded by a function of $\boldsymbol{x}$ with polynomial growth.

Proof. We note that

$$
\Delta_{i}=-\frac{\eta^{2}}{\sigma_{0} \sqrt{u_{i}}+\epsilon_{0}}\left(\partial_{i} f(\boldsymbol{\theta})+\sigma z_{i}\right), \quad \Delta_{d+i}=(1-\beta)\left(\left(\partial_{i} f(\boldsymbol{\theta}) / \sigma+z_{i}\right)^{2}-u_{i}\right) .
$$

Let $\nu_{i}:=\frac{1}{\sigma_{0} \sqrt{u_{i}}+\epsilon_{0}}$. Since $\boldsymbol{x} \in \mathcal{X}_{k}, \nu_{i} \leq \frac{1}{\sigma_{0} \sqrt{u_{\text {min }}}+\epsilon_{0}}=\mathcal{O}(1)$. Writing $1-\beta$ as $c_{2} \eta^{2}$, we have

$$
\begin{equation*}
\Delta_{i}=-\nu_{i} \eta^{2}\left(\partial_{i} f(\boldsymbol{\theta})+\sigma z_{i}\right), \quad \Delta_{d+i}=c_{2} \eta^{2}\left(\left(\partial_{i} f(\boldsymbol{\theta}) / \sigma+z_{i}\right)^{2}-u_{i}\right) \tag{13}
\end{equation*}
$$

We can now compute the first moments:

$$
\begin{array}{rlrl}
\mathbb{E}\left[\Delta_{i}\right] & =-\nu_{i} \eta^{2} \partial_{i} f(\boldsymbol{\theta}) & \left(\mathbb{E}\left[z_{i}\right]=0\right) \\
\mathbb{E}\left[\Delta_{d+i}\right] & =c_{2} \eta^{2}\left(\mathbb{E}\left[\left(\partial_{i} f(\boldsymbol{\theta}) / \sigma\right)^{2}+\mathbb{E}\left[z_{i}^{2}\right]-\mathbb{E}\left[u_{i}\right]\right)\right. & \\
& =c_{2} \eta^{2}\left(\left(\partial_{i} f(\boldsymbol{\theta}) / \sigma\right)^{2}+\Sigma_{i i}-u_{i}\right) & &
\end{array}
$$

Now we observe that $1 / \sigma=\mathcal{O}(\eta)$, so we can write

$$
\mathbb{E}\left[\Delta_{d+i}\right]=\eta^{2} c_{2}\left(\Sigma_{i i}-u_{i}\right)+\mathcal{O}\left(\eta^{4}\right)
$$

Let $\boldsymbol{\delta}:=\boldsymbol{\Delta}-\mathbb{E}[\boldsymbol{\Delta}]$. That is,

$$
\begin{aligned}
\delta_{i} & =-\nu_{i} \eta^{2} \sigma z_{i} & \delta_{d+i} & =c_{2} \eta^{2}\left(z_{i}^{2}-\Sigma_{i i}+2 z_{i} \partial_{i} f(\boldsymbol{\theta}) / \sigma\right) \\
& =-\nu_{i} \sigma_{0} \eta z_{i} . & & =c_{2} \eta^{2}\left(z_{i}^{2}-\Sigma_{i i}\right)+2 c_{2}\left(\partial_{i} f(\boldsymbol{\theta}) / \sigma_{0}\right) \eta^{3} z_{i} .
\end{aligned}
$$

For convenience we also define $w_{i}=\left(z_{i}^{2}-\Sigma_{i i}\right)+2\left(\partial_{i} f(\boldsymbol{\theta}) / \sigma_{0}\right) \eta z_{i}$ and write $\delta_{d+i}=c_{2} \eta^{2} w_{i}$.
For the second moments we have

$$
\mathbb{E}\left[\Delta_{p} \Delta_{q}\right]=\mathbb{E}\left[\delta_{p} \delta_{q}\right]+\mathbb{E}\left[\Delta_{p}\right] \mathbb{E}\left[\Delta_{q}\right]=\mathbb{E}\left[\delta_{p} \delta_{q}\right]+\mathcal{O}\left(\eta^{4}\right) \quad \text { for all } 1 \leq p, q \leq 2 d
$$

Then it suffices to compute the second moments for $\delta$. For $\mathbb{E}\left[\delta_{i} \delta_{j}\right]$ we have

$$
\mathbb{E}\left[\delta_{i} \delta_{j}\right]=\nu_{i} \nu_{j} \eta^{4} \sigma^{2} \mathbb{E}\left[z_{i} z_{j}\right]=\nu_{i} \nu_{j} \sigma_{0}^{2} \Sigma_{i j}
$$

For $\mathbb{E}\left[\delta_{i} \delta_{d+j}\right]$ we have

$$
\begin{aligned}
\mathbb{E}\left[\delta_{i} \delta_{d+j}\right] & =-c_{2} \nu_{i} \sigma_{0} \eta^{3} \mathbb{E}\left[z_{i} w_{j}\right] \\
& =-c_{2} \nu_{i} \sigma_{0} \eta^{3} \mathbb{E}\left[z_{i} z_{j}^{2}\right]+\mathcal{O}\left(\eta^{4}\right) \\
& =\mathbb{E}\left[z_{i} z_{j}^{2}\right] \cdot \mathcal{O}\left(\eta^{3}\right)+\mathcal{O}\left(\eta^{4}\right) .
\end{aligned}
$$

Similarly, we have $\mathbb{E}\left[\Delta_{d+i} \Delta_{j}\right]=\mathbb{E}\left[z_{i}^{2} z_{j}\right] \cdot \mathcal{O}\left(\eta^{3}\right)+\mathcal{O}\left(\eta^{4}\right)$.
For $\mathbb{E}\left[\delta_{d+i} \delta_{d+j}\right]$, we note that $\boldsymbol{z} \sim \mathcal{Z}_{\sigma}(\boldsymbol{\theta})$ has bounded 4th-order moments, so $\mathbb{E}[g(\boldsymbol{z})]=\mathcal{O}(1)$ for any polynomial $g$ of degree at most 4 , if the coefficients of $g$ are bounded by $\mathcal{O}(1)$. Then we have

$$
\mathbb{E}\left[\delta_{d+i} \delta_{d+j}\right]=c_{2}^{2} \eta^{4} \mathbb{E}\left[w_{i} w_{j}\right]=\mathcal{O}\left(\eta^{4}\right)
$$

Now we can check the third moments.

$$
\begin{aligned}
\mathbb{E}\left[\Delta_{p} \Delta_{q} \Delta_{r}\right] & =\mathbb{E}\left[\delta_{p} \delta_{q} \delta_{r}\right]+\left(\mathbb{E}\left[\delta_{p} \delta_{q}\right] \mathbb{E}\left[\Delta_{r}\right]+\mathbb{E}\left[\delta_{p} \delta_{r}\right] \mathbb{E}\left[\Delta_{q}\right]+\mathbb{E}\left[\delta_{q} \delta_{r}\right] \mathbb{E}\left[\Delta_{p}\right]\right)+\mathbb{E}\left[\Delta_{p}\right] \mathbb{E}\left[\Delta_{q}\right] \mathbb{E}\left[\Delta_{r}\right] \\
& =\mathbb{E}\left[\delta_{p} \delta_{q} \delta_{r}\right]+\mathcal{O}\left(\eta^{2}\right) \cdot \mathcal{O}\left(\eta^{2}\right)+\mathcal{O}\left(\eta^{6}\right) \\
& =\mathbb{E}\left[\delta_{p} \delta_{q} \delta_{r}\right]+\mathcal{O}\left(\eta^{4}\right) .
\end{aligned}
$$

Note that $\delta_{p} \delta_{q} \delta_{r}$ is a polynomial of $\boldsymbol{z}$. For $p=i, q=j, r=k$, by the low skewness condition for $\mathcal{Z}_{\sigma}$ we have

$$
\mathbb{E}\left[\delta_{i} \delta_{j} \delta_{k}\right]=-\nu_{i}^{3} \sigma_{0}^{3} \eta^{3} \mathbb{E}\left[z_{i} z_{j} z_{k}\right]=K_{3}(\boldsymbol{\theta}) / \sigma \cdot \mathcal{O}\left(\eta^{3}\right)=\mathcal{O}\left(\eta^{4}\right)
$$

Except the above case, it can be shown that $\delta_{p} \delta_{q} \delta_{r}$ is a polynomial with coefficients bounded by $\mathcal{O}\left(\eta^{4}\right)$. Combining this with the fact that $\boldsymbol{z} \sim \mathcal{Z}_{\sigma}(\boldsymbol{\theta})$ has bounded moments of any order, we have $\mathbb{E}\left[\delta_{p} \delta_{q} \delta_{r}\right]=\mathcal{O}\left(\eta^{4}\right)$.

Lemma C.8. For $\boldsymbol{x}=(\boldsymbol{\theta}, \boldsymbol{u}) \in \mathcal{X}_{k}$, if the NGOS is well-behaved, then the moments of $\tilde{\boldsymbol{\Delta}}:=\tilde{\boldsymbol{\Delta}}(x, k)$ can be written as below.

1. The first moments are given by

$$
\mathbb{E}\left[\tilde{\Delta}_{i}\right]=-\frac{\eta^{2}}{\sigma_{0} \sqrt{\mu\left(u_{i}\right)}+\epsilon_{0}} \partial_{i} f(\boldsymbol{\theta}), \quad \mathbb{E}\left[\tilde{\Delta}_{d+i}\right]=\eta^{2} c_{2}\left(\Sigma_{i i}(\boldsymbol{\theta})-u_{i}\right)+\mathcal{O}\left(\eta^{4}\right)
$$

for all $i \in[d]$.
2. The second moments are given by

$$
\begin{aligned}
\mathbb{E}\left[\tilde{\Delta}_{i} \tilde{\Delta}_{j}\right] & =\frac{\eta^{2} \Sigma_{i j}(\boldsymbol{\theta})}{\left(\sqrt{\mu\left(u_{i}\right)}+\epsilon_{0} / \sigma_{0}\right)\left(\sqrt{\mu\left(u_{j}\right)}+\epsilon_{0} / \sigma_{0}\right)}+\mathcal{O}\left(\eta^{4}\right) & \mathbb{E}\left[\tilde{\Delta}_{i} \tilde{\Delta}_{d+j}\right]=\mathcal{O}\left(\eta^{4}\right) \\
\mathbb{E}\left[\tilde{\Delta}_{d+i} \tilde{\Delta}_{j}\right] & =\mathcal{O}\left(\eta^{4}\right) & \mathbb{E}\left[\tilde{\Delta}_{d+i} \tilde{\Delta}_{d+j}\right]=\mathcal{O}\left(\eta^{4}\right) .
\end{aligned}
$$

for all $i, j \in[d]$.
3. The third moments are bounded by $\mathbb{E}\left[\tilde{\boldsymbol{\Delta}}^{\otimes 3}\right]=\mathcal{O}\left(\eta^{4}\right)$.

Here the big- $O$ notation $\mathcal{O}(\cdot)$ is used in a way that $\mathcal{O}(1)$ hides constants (independent of $\eta$ and $\boldsymbol{x}$ ) and values that are bounded by a function of $\boldsymbol{x}$ with polynomial growth.

Proof. Applying Lemma B. 3 gives

$$
\mathbb{E}[\tilde{\boldsymbol{\Delta}}]=\eta^{2} \boldsymbol{b}(\boldsymbol{x})+\mathcal{O}\left(\eta^{4}\right), \quad \mathbb{E}\left[\tilde{\boldsymbol{\Delta}} \tilde{\boldsymbol{\Delta}}^{\top}\right]=\eta^{2} \boldsymbol{\sigma}(\boldsymbol{x}) \boldsymbol{\sigma}(\boldsymbol{x})^{\top}+\mathcal{O}\left(\eta^{4}\right), \quad \mathbb{E}\left[\tilde{\boldsymbol{\Delta}}^{\otimes 3}\right]=\mathcal{O}\left(\eta^{4}\right)
$$

Splitting up the formula by indices proves the claim.

## C. 3 Verifying Condition (3)

Lemma C.9. Let $P:=\{1,2, \ldots, d\}$. Then

1. There is a constant $C_{1}>0$ (independent of $\eta_{\mathrm{e}}$ ) so that for all $k \leq N$ and $\boldsymbol{x} \in \mathcal{X}_{k}$,

$$
\begin{aligned}
& \|\mathbb{E} \boldsymbol{\Delta}(\boldsymbol{x}, k)\|_{\mathrm{P}} \leq C_{1} \eta_{\mathrm{e}}\left(1+\|\boldsymbol{x}\|_{\mathrm{P}}\right) \\
& \|\mathbb{E} \boldsymbol{\Delta}(\boldsymbol{x}, k)\|_{\mathrm{R}} \leq C_{1} \eta_{\mathrm{e}}\left(1+\|\boldsymbol{x}\|_{\mathrm{P}}^{2}\right)\left(1+\|\boldsymbol{x}\|_{\mathrm{R}}\right)
\end{aligned}
$$

2. For all $m \geq 1$, there is a constant $C_{2 m}$ (independent of $\eta_{\mathrm{e}}$ ) so that for all $k \leq N$ and $\boldsymbol{x} \in \mathcal{X}_{k}$,

$$
\begin{aligned}
& \mathbb{E}\|\boldsymbol{\Delta}(\boldsymbol{x}, k)\|_{\mathrm{P}}^{2 m} \leq C_{2 m} \eta_{\mathrm{e}}^{m}\left(1+\|\boldsymbol{x}\|_{\mathrm{P}}^{2 m}\right) \\
& \mathbb{E}\|\boldsymbol{\Delta}(\boldsymbol{x}, k)\|_{\mathrm{R}}^{2 m} \leq C_{2 m} \eta_{\mathrm{e}}^{m}\left(1+\|\boldsymbol{x}\|_{\mathrm{P}}^{4 m}\right)\left(1+\|\boldsymbol{x}\|_{\mathrm{R}}^{2 m}\right),
\end{aligned}
$$

Proof. By LemmaC.7, for all $i \in[d]$,

$$
\mathbb{E}\left[\Delta_{i}\right]=-\frac{\eta^{2}}{\sigma_{0} \sqrt{u_{i}}+\epsilon_{0}} \partial_{i} f(\boldsymbol{\theta}), \quad \mathbb{E}\left[\Delta_{d+i}\right]=\eta^{2} c_{2}\left(\Sigma_{i i}(\boldsymbol{\theta})-u_{i}\right)+\frac{\eta^{4} c_{2}}{\sigma_{0}^{2}}\left(\partial_{i} f(\boldsymbol{\theta})\right)^{2}
$$

Combining this with the Lipschitzness of $\nabla f(\boldsymbol{\theta})$ and the boundedness of $\boldsymbol{\Sigma}(\boldsymbol{\theta})$ proves Item 1.
By (13), for all $i \in[d]$,

$$
\begin{aligned}
\left|\Delta_{i}\right| & =\nu_{i} \eta^{2}\left|\partial_{i} f(\boldsymbol{\theta})+\sigma z_{i}\right| \leq \eta^{2} \nu_{i}\left(1+\left|\partial_{i} f(\boldsymbol{\theta})\right|\right)\left(1+\sigma\left|z_{i}\right|\right) \leq \eta \nu_{i}\left(1+\left|\partial_{i} f(\boldsymbol{\theta})\right|\right)\left(1+\sigma_{0}\left|z_{i}\right|\right) \\
\left|\Delta_{d+i}\right| & =c_{2} \eta^{2}\left|\left(\partial_{i} f(\boldsymbol{\theta}) / \sigma+z_{i}\right)^{2}-u_{i}\right| \leq c_{2} \eta^{2}\left(1+\left(\partial_{i} f(\boldsymbol{\theta}) / \sigma+z_{i}\right)^{2}\right)\left(1+u_{i}\right)
\end{aligned}
$$

By the Lipschitzness of $\nabla f(\boldsymbol{\theta})$ and the bounded moments condition for $\mathcal{Z}_{\sigma}$, we can prove Item 2 by taking powers and expectations on both sides of the above inequalities.

## D Adam SDE Proof

In this section, we prove the theorem for the SDE approximation of Adam.
Definition D. 1 (SDE for Adam, matrix form). For constants $\sigma_{0}, \epsilon_{0}, c_{1}$ and $c_{2}$, define $\gamma_{1}(t):=$ $1-e^{-c_{1} t}, \gamma_{2}(t):=1-e^{-c_{2} t}$ and define the SDE as $\mathrm{d} \boldsymbol{X}_{t}=\boldsymbol{b}\left(\boldsymbol{X}_{t}\right) \mathrm{d} t+\boldsymbol{\sigma}\left(\boldsymbol{X}_{t}\right) \mathrm{d} \boldsymbol{W}_{t}$, where $\boldsymbol{X}_{t} \in \mathbb{R}^{d} \times \mathbb{R}^{d} \times \mathbb{R}^{d}, \boldsymbol{b}$ is defined by

$$
\begin{aligned}
b_{i}(\boldsymbol{x}, t) & :=-\frac{\sqrt{\gamma_{2}(t)}}{\gamma_{1}(t)} \cdot \frac{m_{i}}{\sigma_{0} \sqrt{u_{i}}+\epsilon_{0} \sqrt{\gamma_{2}(t)}}, \\
b_{d+i}(\boldsymbol{x}, t) & :=c_{1}\left(\partial_{i} f(\boldsymbol{\theta})-m_{i}\right), \\
b_{2 d+i}(\boldsymbol{x}, t) & :=c_{2}\left(\Sigma(\boldsymbol{\theta})_{i, i}-u_{i}\right),
\end{aligned}
$$

for all $1 \leq i \leq d$, and $\sigma_{d+i, d+j}$ for all $1 \leq i, j \leq d$ is given by

$$
\sigma_{d+i, d+j}(\boldsymbol{x}, t):=\sigma_{0} c_{1}\left(\boldsymbol{\Sigma}^{1 / 2}(\boldsymbol{\theta})\right)_{i, j}
$$

and all the other entries of $\boldsymbol{\sigma}$ are zero.
Theorem D.2. Fix $\sigma_{0}, c_{1}, c_{2}>0, \epsilon_{0} \geq 0$. Let $T>0, \eta^{2} \in\left(0,1 \wedge T \wedge \frac{1}{2 c_{2}}\right)$ and set $N=\left\lfloor T / \eta^{2}\right\rfloor$. Let $\boldsymbol{u}_{k} \triangleq \boldsymbol{v}_{k} / \sigma^{2}$ and $\boldsymbol{x}_{k} \triangleq\left(\boldsymbol{\theta}_{k}, \boldsymbol{m}_{k}, \boldsymbol{u}_{k}\right) \in \mathbb{R}^{3 d}$. Let $\left\{\boldsymbol{x}_{k}: k \geq 0\right\}$ be the discrete Adam iterations defined in Definition 2.2. Set $\sigma, \epsilon, \beta_{1}, \beta_{2}$ so that $\sigma_{0}=\sigma \eta, \epsilon_{0}=\epsilon \eta, c_{1}=\left(1-\beta_{1}\right) / \eta^{2}$ and $c_{2}=\left(1-\beta_{2}\right) / \eta^{2}$. For well-behaved NGOS that satisfies the bounded moments and low skewness conditions, for any constant $t_{0}>0$, the solution $\boldsymbol{X}_{t}\left(t \in\left[t_{0}, T\right]\right)$ of the SDE defined as in Definition D.1 is an order-1 weak approximation (Definition [2.4) of the sequence of Adam iterates $\boldsymbol{x}_{k}$ starting from $k_{0}=\left\lceil t_{0} / \eta^{2}\right\rceil$, if the initial condition of the SDE is set to $\boldsymbol{X}_{t_{0}}=\boldsymbol{x}_{k_{0}}$.

The proof strategy is essentially the same as that for RMSprop. Similar to what we have done for RMSprop, we turn to prove the approximation order for the following auxiliary SDE.

Definition D. 3 (Auxiliary SDE for Adam, matrix form). For constants $\sigma_{0}, \epsilon_{0}, c_{1}, c_{2}$ and $u_{\text {min }}$, define $\gamma_{1}(t):=1-e^{-c_{1} t}, \gamma_{2}(t):=1-e^{-c_{2} t}, \mu(u):=\frac{1}{2} u_{\min }+\tau\left(\frac{2 u}{u_{\min }}-1\right) \cdot\left(u-\frac{1}{2} u_{\min }\right)$, and define the SDE as $\mathrm{d} \boldsymbol{X}_{t}=\boldsymbol{b}\left(\boldsymbol{X}_{t}\right) \mathrm{d} t+\boldsymbol{\sigma}\left(\boldsymbol{X}_{t}\right) \mathrm{d} \boldsymbol{W}_{t}$, where $\boldsymbol{X}_{t} \in \mathbb{R}^{d} \times \mathbb{R}^{d} \times \mathbb{R}^{d}, \boldsymbol{b}$ is defined by

$$
\begin{aligned}
b_{i}(\boldsymbol{x}, t) & :=-\frac{\sqrt{\gamma_{2}(t)}}{\gamma_{1}(t)} \cdot \frac{m_{i}}{\sigma_{0} \sqrt{\mu\left(u_{i}\right)}+\epsilon_{0} \sqrt{\gamma_{2}(t)}}, \\
b_{d+i}(\boldsymbol{x}, t) & :=c_{1}\left(\partial_{i} f(\boldsymbol{\theta})-m_{i}\right) \\
b_{2 d+i}(\boldsymbol{x}, t) & :=c_{2}\left(\Sigma(\boldsymbol{\theta})_{i, i}-u_{i}\right),
\end{aligned}
$$

for all $1 \leq i \leq d$, and $\sigma_{d+i, d+j}$ for all $1 \leq i, j \leq d$ is given by

$$
\sigma_{d+i, d+j}(\boldsymbol{x}, t):=\sigma_{0} c_{1}\left(\boldsymbol{\Sigma}^{1 / 2}(\boldsymbol{\theta})\right)_{i, j}
$$

and all the other entries of $\boldsymbol{\sigma}$ are zero.
Theorem D.4. In the setting of Theorem D.2. let $u_{\min }=2^{-\left(T-t_{0}\right)} \min _{i \in[d]} u_{k_{0}, i}$, For any constant $t_{0}>0$, the solution $\boldsymbol{X}_{t}\left(t \in\left[t_{0}, T\right]\right)$ of the $S D E$ defined as in Definition D. 3 is an order- 1 weak approximation (Definition 2.4) of the sequence of Adam iterates $\boldsymbol{x}_{k}$ starting from $k_{0}=\left\lceil t_{0} / \eta^{2}\right\rceil$, if the initial condition of the $\overline{S D E}$ is set to $\boldsymbol{X}_{t_{0}}=\boldsymbol{x}_{k_{0}}$.

Proof for Theorem D.2. Given Theorem D.4 we only need to show that $\boldsymbol{X}_{t}$ has the same distribution in the original and auxiliary SDEs for all $t \in[0, T]$, when $\boldsymbol{X}_{0}=\left(\boldsymbol{\theta}_{0}, \boldsymbol{m}_{0}, \boldsymbol{u}_{0}\right)$. To see this, we only need to note that in the original SDE

$$
\frac{\mathrm{d} u_{t, i}}{\mathrm{~d} t}=c_{2}\left(\Sigma\left(\boldsymbol{\theta}_{t}\right)_{i, i}-u_{t, i}\right) \geq-c_{2} u_{t, i} .
$$

Thus, $u_{t, i} \geq \exp \left(-c_{2} t\right) u_{0, i} \geq u_{\text {min }}$, which means $\operatorname{Pr}\left[u_{t, i}=\mu\left(u_{t, i}\right)\right]=1$ for all $t \in[0, T]$.
It remains to prove Theorem D. 4 by applying Theorem B. 2 In the rest of this section, we verify the three conditions in Theorem B. 2 respectively.
Below we use the notations $\boldsymbol{x}_{k}, \boldsymbol{X}_{t}, \boldsymbol{b}, \boldsymbol{\sigma}$ defined as in Theorem D.4. Let $D=3 d$. Every $\boldsymbol{x}_{k} \in \mathbb{R}^{D}$ is a concatenation of three $\mathbb{R}^{d}$-vectors $\boldsymbol{\theta}_{k}, \boldsymbol{m}_{k}$ and $\boldsymbol{u}_{k}$. According to the update rule of Adam, $\boldsymbol{x}_{k}$ can be seen as SGA $\boldsymbol{x}_{k+1}=\boldsymbol{x}_{k}-\eta_{\mathrm{e}} \boldsymbol{h}_{k}\left(\boldsymbol{x}_{k}, \boldsymbol{z}_{k}, \eta_{\mathrm{e}}\right)$, where $\eta_{\mathrm{e}}=\eta^{2}, \boldsymbol{z}_{k} \sim \mathcal{Z}_{\sigma}\left(\boldsymbol{\theta}_{k}\right)$, and $\boldsymbol{h}_{k}$ is defined below:

$$
\boldsymbol{h}_{k}\left(\boldsymbol{\theta}, \boldsymbol{u}, \boldsymbol{z}, \eta_{\mathrm{e}}\right):=\left[\begin{array}{c}
-\frac{\sqrt{1-\beta_{2}^{k}}}{1-\beta_{1}^{k+1}} \boldsymbol{m} \odot\left(\sigma_{0} \sqrt{\boldsymbol{u}}+\epsilon_{0} \sqrt{1-\beta_{2}^{k}}\right)^{-1} \\
c_{1}(\nabla f(\boldsymbol{\theta})+\sigma \boldsymbol{z}-\boldsymbol{m}) \\
c_{2}\left((\nabla f(\boldsymbol{\theta}) / \sigma+\boldsymbol{z})^{2}-\boldsymbol{u}\right)
\end{array}\right] .
$$

We define $\Delta$ and $\tilde{\Delta}$ as in (6) and (7).
Fix $\boldsymbol{x}_{0}=\left(\boldsymbol{\theta}_{0}, \boldsymbol{m}_{0}, \boldsymbol{u}_{0}\right)$ with $u_{0, j}>0$ for all $j \in[d]$. Define $u_{\text {min }}$ as in TheoremD.4 Let $\mathcal{X}_{k}$ be the support of the random variable $\boldsymbol{x}_{k}$ given $\boldsymbol{x}_{0}$, then it is easy to show that $\mathcal{X}_{k}$ is a subset of $\left\{(\boldsymbol{\theta}, \boldsymbol{m}, \boldsymbol{u}): u_{j} \geq u_{\text {min }}\right.$ for all $\left.j \in[d]\right\}$.

## D. 1 Verifying Condition (1)

Lemma D.5. The drift function $\boldsymbol{b}$ and diffusion function $\boldsymbol{\sigma}$ are Lipschitz and belong to $G^{4}$.
Proof. Same argument as for LemmaC. 6.

## D. 2 Verifying Condition (2)

Let $\boldsymbol{x}=(\boldsymbol{\theta}, \boldsymbol{m}, \boldsymbol{u}) \in \mathbb{R}^{3 d}$. For ease of notation, let $\hat{\gamma}_{1}=1-\beta_{1}^{k+1}$ and $\hat{\gamma}_{2}=1-\beta_{2}^{k}$. These are not constants across time steps like the other constants, but they are deterministic and upper bounded by constants for $k \geq t_{0} / \eta^{2}$.
To verify Condition 2 , we only need to compute the moments of $\boldsymbol{\Delta}$ and $\tilde{\Delta}$ for the discrete Adam and the auxiliary SDE, and show that they are close to each other. We compute them by the following two lemmas.

Lemma D.6. For $\boldsymbol{x}=(\boldsymbol{\theta}, \boldsymbol{u}) \in \mathcal{X}_{k}$, if the $N G O S$ is well-behaved and $\mathcal{Z}_{\sigma}$ satisfies the bounded moments and low skewness condition, then the moments of $\boldsymbol{\Delta}:=\boldsymbol{\Delta}(\boldsymbol{x}, k)$ can be written as below.

1. The first moments are given by

$$
\begin{aligned}
\mathbb{E}\left[\Delta_{i}\right] & =-\frac{\sqrt{\hat{\gamma}_{2}}}{\hat{\gamma}_{1}} \cdot \frac{\eta^{2}}{\sigma_{0} \sqrt{u_{i}}+\epsilon_{0} \sqrt{\hat{\gamma}_{2}}}\left(m_{i}+c_{1} \eta^{2}\left(\partial_{i} f(\boldsymbol{\theta})-m_{i}\right)\right) \\
& =-\frac{\sqrt{\hat{\gamma}_{2}}}{\hat{\gamma}_{1}} \cdot \frac{\eta^{2} m_{i}}{\sigma_{0} \sqrt{u_{i}}+\epsilon_{0} \sqrt{\hat{\gamma}_{2}}}+\mathcal{O}\left(\eta^{4}\right) \\
\mathbb{E}\left[\Delta_{d+i}\right] & =c_{1} \eta^{2}\left(\partial_{i} f(\boldsymbol{\theta})-m_{i}\right) \\
\mathbb{E}\left[\Delta_{2 d+i}\right] & =c_{2} \eta^{2}\left(\left(\partial_{i} f(\boldsymbol{\theta}) / \sigma\right)^{2}+\Sigma_{i i}-u_{i}\right) \\
& =c_{2} \eta^{2}\left(\Sigma_{i i}-u_{i}\right)+\mathcal{O}\left(\eta^{4}\right) .
\end{aligned}
$$

for all $i \in[d]$.
2. The second moments are given by

$$
\mathbb{E}\left[\Delta_{p} \Delta_{q}\right]= \begin{cases}c_{1}^{2} \sigma_{0}^{2} \eta^{2} \Sigma_{i j}+\mathcal{O}\left(\eta^{4}\right) & \text { if } p=d+i, q=d+j \quad \text { for some } \quad i, j \in[d] \\ \mathcal{O}\left(\eta^{4}\right) & \text { otherwise. }\end{cases}
$$

for all $p, q \in[3 d]$.
3. The third moments are bounded by $\mathbb{E}\left[\boldsymbol{\Delta}^{\otimes 3}\right]=\mathcal{O}\left(\eta^{4}\right)$.

Here the big- $O$ notation $\mathcal{O}(\cdot)$ is used in a way that $\mathcal{O}(1)$ hides constants (independent of $\eta$ and $\boldsymbol{x}$ ) and values that are bounded by a function of $\boldsymbol{x}$ with polynomial growth.

Proof. For notational convenience, we write $\Delta, \Sigma$ instead of $\Delta(\boldsymbol{x}), \Sigma(\boldsymbol{\theta})$. We note that

$$
\begin{aligned}
\Delta_{i} & =-\frac{\sqrt{\hat{\gamma}_{2}}}{\hat{\gamma}_{1}} \cdot \frac{\eta^{2} m_{i}+\eta^{2}\left(1-\beta_{1}\right)\left(\partial_{i} f(\boldsymbol{\theta})+\sigma z_{i}-m_{i}\right)}{\sigma_{0} \sqrt{u_{i}}+\epsilon_{0} \sqrt{\hat{\gamma}_{2}}} \\
\Delta_{d+i} & =\left(1-\beta_{1}\right)\left(\partial_{i} f(\boldsymbol{\theta})+\sigma z_{i}-m_{i}\right) \\
\Delta_{2 d+i} & =\left(1-\beta_{2}\right)\left(\left(\partial_{i} f(\boldsymbol{\theta}) / \sigma+z_{i}\right)^{2}-u_{i}\right)
\end{aligned}
$$

Let $\nu_{i}:=\frac{\sqrt{\hat{\gamma}_{2}}}{\hat{\hat{\gamma}_{1}}} \cdot \frac{1}{\sigma_{0} \sqrt{u_{i}}+\epsilon_{0} \sqrt{\hat{\gamma}_{2}}}$. And we write $1-\beta=c_{2} \eta^{2}$. Then we have

$$
\begin{align*}
\Delta_{i} & =-\nu_{i} \eta^{2}\left(m_{i}+c_{1} \eta^{2}\left(\partial_{i} f(\boldsymbol{\theta})+\sigma z_{i}-m_{i}\right)\right) \\
\Delta_{d+i} & =c_{1} \eta^{2}\left(\partial_{i} f(\boldsymbol{\theta})+\sigma z_{i}-m_{i}\right)  \tag{14}\\
\Delta_{2 d+i} & =c_{2} \eta^{2}\left(\left(\partial_{i} f(\boldsymbol{\theta}) / \sigma+z_{i}\right)^{2}-u_{i}\right)
\end{align*}
$$

The first moments follow directly:

$$
\begin{aligned}
\mathbb{E}\left[\Delta_{i}\right] & =-\nu_{i} \eta^{2}\left(m_{i}+c_{1} \eta^{2}\left(\partial_{i} f(\boldsymbol{\theta})-m_{i}\right)\right)=-\nu_{i} \eta^{2} m_{i}+\mathcal{O}\left(\eta^{4}\right) \\
\mathbb{E}\left[\Delta_{d+i}\right] & =c_{1} \eta^{2}\left(\partial_{i} f(\boldsymbol{\theta})-m_{i}\right) \\
\mathbb{E}\left[\Delta_{2 d+i}\right] & =c_{2} \eta^{2}\left(\left(\partial_{i} f(\boldsymbol{\theta}) / \sigma\right)^{2}+\Sigma_{i i}-u_{i}\right)=c_{2} \eta^{2}\left(\Sigma_{i i}-u_{i}\right)+\mathcal{O}\left(\eta^{4}\right) .
\end{aligned}
$$

Let $\boldsymbol{\delta}:=\boldsymbol{\Delta}-\mathbb{E}[\boldsymbol{\Delta}]$. That is,

$$
\begin{aligned}
\delta_{i} & =-\nu_{i} c_{1} \eta^{4} \sigma z_{i}=-\nu_{i} c_{1} \sigma_{0} \eta^{3} z_{i} \\
\delta_{d+i} & =c_{1} \eta^{2} \sigma z_{i}=c_{1} \sigma_{0} \eta z_{i} \\
\delta_{2 d+i} & =c_{2} \eta^{2}\left(z_{i}^{2}-\Sigma_{i i}+2 z_{i} \partial_{i} f(\boldsymbol{\theta}) / \sigma\right) \\
& =c_{2} \eta^{2}\left(z_{i}^{2}-\Sigma_{i i}\right)+2 c_{2}\left(\partial_{i} f(\boldsymbol{\theta}) / \sigma_{0}\right) \eta^{3} z_{i}
\end{aligned}
$$

For convenience we also define $w_{i}=\left(z_{i}^{2}-\Sigma_{i i}\right)+2\left(\partial_{i} f(\boldsymbol{\theta}) / \sigma_{0}\right) \eta z_{i}$ and write $\delta_{d+i}=c_{2} \eta^{2} w_{i}$.
Similar as the proof for LemmaC.7, for the second moments we have

$$
\mathbb{E}\left[\Delta_{p} \Delta_{q}\right]=\mathbb{E}\left[\delta_{p} \delta_{q}\right]+\mathbb{E}\left[\Delta_{p}\right] \mathbb{E}\left[\Delta_{q}\right]=\mathbb{E}\left[\delta_{p} \delta_{q}\right]+\mathcal{O}\left(\eta^{4}\right) \quad \text { for all } 1 \leq p, q \leq 2 d
$$

Then it suffices to compute the second order moments for $\delta$. A direct computation gives the following:

$$
\begin{aligned}
\mathbb{E}\left[\delta_{i} \delta_{j}\right] & =\nu_{i} \nu_{j} c_{1}^{2} \sigma_{0}^{2} \eta^{6} \mathbb{E}\left[z_{i} z_{j}\right]=\mathcal{O}\left(\eta^{6}\right) . \\
\mathbb{E}\left[\delta_{i} \delta_{d+j}\right] & =-\nu_{i} c_{1}^{2} \sigma_{0}^{2} \eta^{4} \mathbb{E}\left[z_{i} z_{j}\right]=\mathcal{O}\left(\eta^{4}\right) . \\
\mathbb{E}\left[\delta_{i} \delta_{2 d+j}\right] & =-\nu_{i} c_{1} c_{2} \sigma_{0} \eta^{5} \mathbb{E}\left[z_{i} w_{j}\right]=\mathcal{O}\left(\eta^{5}\right) . \\
\mathbb{E}\left[\delta_{d+i} \delta_{d+j}\right] & =c_{1}^{2} \sigma_{0}^{2} \eta^{2} \mathbb{E}\left[z_{i} z_{j}\right]=c_{1}^{2} \sigma_{0}^{2} \eta^{2} \Sigma_{i j} . \\
\mathbb{E}\left[\delta_{d+i} \delta_{2 d+j}\right] & =c_{1} c_{2} \sigma_{0} \eta^{3} \mathbb{E}\left[z_{i} w_{j}\right]=c_{1} c_{2} \sigma_{0} \eta^{3} \mathbb{E}\left[z_{i} z_{j}^{2}\right]+\mathcal{O}\left(\eta^{4}\right) . \\
\mathbb{E}\left[\delta_{2 d+i} \delta_{2 d+j}\right] & =c_{2}^{2} \eta^{4} \mathbb{E}\left[w_{i} w_{j}\right]=\mathcal{O}\left(\eta^{4}\right) .
\end{aligned}
$$

Now we check the third moments. Similar as the proof for LemmaC.7 $\mathbb{E}\left[\Delta_{p} \Delta_{q} \Delta_{r}\right]=\mathbb{E}\left[\delta_{p} \delta_{q} \delta_{r}\right]+$ $\mathcal{O}\left(\eta^{4}\right)$. Note that $\delta_{p} \delta_{q} \delta_{r}$ is a polynomial of $\boldsymbol{z}$. For $p=d+i, q=d+j, r=d+k$, by the low skewness condition for $\mathcal{Z}_{\sigma}$ we have

$$
\mathbb{E}\left[\delta_{d+i} \delta_{d+j} \delta_{d+k}\right]=c_{1}^{3} \sigma_{0}^{3} \eta^{3} \mathbb{E}\left[z_{i} z_{j} z_{k}\right]=K_{3}(\boldsymbol{\theta}) / \sigma \cdot \mathcal{O}\left(\eta^{3}\right)=\mathcal{O}\left(\eta^{4}\right)
$$

Except the above case, it can be shown that $\delta_{p} \delta_{q} \delta_{r}$ is a polynomial with coefficients bounded by $\mathcal{O}\left(\eta^{4}\right)$. Combining this with the fact that $\boldsymbol{z} \sim \mathcal{Z}_{\sigma}(\boldsymbol{\theta})$ has bounded moments of any order, we have $\mathbb{E}\left[\delta_{p} \delta_{q} \delta_{r}\right]=\mathcal{O}\left(\eta^{4}\right)$.

Lemma D.7. For $\boldsymbol{x}=(\boldsymbol{\theta}, \boldsymbol{u}) \in \mathcal{X}_{k}$, if the NGOS is well-behaved, then the moments of $\tilde{\boldsymbol{\Delta}}:=\tilde{\boldsymbol{\Delta}}(\boldsymbol{x}, k)$ can be written as below.

## 1. The first moments are given by

$$
\begin{aligned}
\mathbb{E}\left[\Delta_{i}\right] & =-\frac{\sqrt{\hat{\gamma}_{2}}}{\hat{\gamma}_{1}} \cdot \frac{\eta^{2} m_{i}}{\sigma_{0} \sqrt{u_{i}}+\epsilon_{0} \sqrt{\hat{\gamma}_{2}}}+\mathcal{O}\left(\eta^{4}\right) \\
\mathbb{E}\left[\Delta_{d+i}\right] & =c_{1} \eta^{2}\left(\partial_{i} f(\boldsymbol{\theta})-m_{i}\right)+\mathcal{O}\left(\eta^{4}\right) \\
\mathbb{E}\left[\Delta_{2 d+i}\right] & =c_{2} \eta^{2}\left(\Sigma_{i i}-u_{i}\right)+\mathcal{O}\left(\eta^{4}\right)
\end{aligned}
$$

for all $i \in[d]$.
2. The second moments are given by

$$
\mathbb{E}\left[\Delta_{p} \Delta_{q}\right]= \begin{cases}c_{1}^{2} \sigma_{0}^{2} \eta^{2} \Sigma_{i j}+\mathcal{O}\left(\eta^{4}\right) & \text { if } p=d+i, q=d+j \quad \text { for some } \quad i, j \in[d] \\ \mathcal{O}\left(\eta^{4}\right) & \text { otherwise. }\end{cases}
$$

for all $p, q \in[3 d]$.
3. The third moments are bounded by $\mathbb{E}\left[\boldsymbol{\Delta}^{\otimes 3}\right]=\mathcal{O}\left(\eta^{4}\right)$.

Here the big- $O$ notation $\mathcal{O}(\cdot)$ is used in a way that $\mathcal{O}(1)$ hides constants (independent of $\eta$ and $\boldsymbol{x}$ ) and values that are bounded by a function of $\boldsymbol{x}$ with polynomial growth.

Proof. Applying Lemma B. 3 gives

$$
\mathbb{E}[\tilde{\boldsymbol{\Delta}}]=\eta^{2} \boldsymbol{b}(\boldsymbol{x})+\mathcal{O}\left(\eta^{4}\right), \quad \mathbb{E}\left[\tilde{\boldsymbol{\Delta}} \tilde{\boldsymbol{\Delta}}^{\top}\right]=\eta^{2} \boldsymbol{\sigma}(\boldsymbol{x}) \boldsymbol{\sigma}(\boldsymbol{x})^{\top}+\mathcal{O}\left(\eta^{4}\right), \quad \mathbb{E}\left[\tilde{\boldsymbol{\Delta}}^{\otimes 3}\right]=\mathcal{O}\left(\eta^{4}\right)
$$

Noting that $\gamma_{1}\left(k \eta_{\mathrm{e}}\right)=\hat{\gamma}_{1}+\mathcal{O}\left(\eta^{2}\right)$ and $\gamma_{2}\left(k \eta_{\mathrm{e}}\right)=\hat{\gamma}_{2}+\mathcal{O}\left(\eta^{2}\right)$, we can prove the claim.

## D. 3 Verifying Condition (3)

Lemma D.8. Let $P:=\{1,2, \ldots, 2 d\}$. Then

1. There is a constant $C_{1}>0$ (independent of $\eta_{\mathrm{e}}$ ) so that for all $k \leq N$ and $\boldsymbol{x} \in \mathcal{X}_{k}$,

$$
\begin{aligned}
\|\mathbb{E} \boldsymbol{\Delta}(\boldsymbol{x}, k)\|_{\mathrm{P}} & \leq C_{1} \eta_{\mathrm{e}}\left(1+\|\boldsymbol{x}\|_{\mathrm{P}}\right) \\
\|\mathbb{E} \boldsymbol{\Delta}(\boldsymbol{x}, k)\|_{\mathrm{R}} & \leq C_{1} \eta_{\mathrm{e}}\left(1+\|\boldsymbol{x}\|_{\mathrm{P}}^{2}\right)\left(1+\|\boldsymbol{x}\|_{\mathrm{R}}\right)
\end{aligned}
$$

2. For all $m \geq 1$, there is a constant $C_{2 m}$ (independent of $\eta_{\mathrm{e}}$ ) so that for all $k \leq N$ and $\boldsymbol{x} \in \mathcal{X}_{k}$,

$$
\begin{aligned}
& \mathbb{E}\|\boldsymbol{\Delta}(\boldsymbol{x}, k)\|_{\mathrm{P}}^{2 m} \leq C_{2 m} \eta_{\mathrm{e}}^{m}\left(1+\|\boldsymbol{x}\|_{\mathrm{P}}^{2 m}\right) \\
& \mathbb{E}\|\boldsymbol{\Delta}(\boldsymbol{x}, k)\|_{\mathrm{R}}^{2 m} \leq C_{2 m} \eta_{\mathrm{e}}^{m}\left(1+\|\boldsymbol{x}\|_{\mathrm{P}}^{4 m}\right)\left(1+\|\boldsymbol{x}\|_{\mathrm{R}}^{2 m}\right) .
\end{aligned}
$$

Proof. By Lemma D.6 for all $i \in[d]$,

$$
\begin{aligned}
\mathbb{E}\left[\Delta_{i}\right] & =-\nu_{i} \eta^{2}\left(m_{i}+c_{1} \eta^{2}\left(\partial_{i} f(\boldsymbol{\theta})-m_{i}\right)\right) \\
\mathbb{E}\left[\Delta_{d+i}\right] & =c_{1} \eta^{2}\left(\partial_{i} f(\boldsymbol{\theta})-m_{i}\right) \\
\mathbb{E}\left[\Delta_{2 d+i}\right] & =c_{2} \eta^{2}\left(\left(\partial_{i} f(\boldsymbol{\theta}) / \sigma\right)^{2}+\Sigma_{i i}-u_{i}\right)
\end{aligned}
$$

Combining this with the Lipschitzness of $\nabla f(\boldsymbol{\theta})$ and the boundedness of $\boldsymbol{\Sigma}(\boldsymbol{\theta})$ proves Item 1.
By (14), for all $i \in[d]$, one can show that there exists a constant $\hat{C}$ such that

$$
\begin{aligned}
\left|\Delta_{i}\right| & \leq \eta^{2} \hat{C}\left(1+\left|m_{i}\right|+\left|\partial_{i} f(\boldsymbol{\theta})\right|\right)\left(1+\left|z_{i}\right|\right) \\
\left|\Delta_{d+i}\right| & \leq \eta^{2} \hat{C}\left(1+\left|m_{i}\right|+\left|\partial_{i} f(\boldsymbol{\theta})\right|\right)\left(1+\left|z_{i}\right|\right) \\
\left|\Delta_{2 d+i}\right| & \leq \eta^{2} \hat{C}\left(1+\left|\partial_{i} f(\boldsymbol{\theta})\right|^{2}+z_{i}^{2}\right)\left(1+u_{i}\right)
\end{aligned}
$$

By the Lipschitzness of $\nabla f(\boldsymbol{\theta})$ and the bounded moments condition for $\mathcal{Z}_{\sigma}$, we can prove Item 2 by taking powers and expectations on both sides of the above inequalities.

## E Analysis of SVAG Operator

Lemma E.1. Let $\mathcal{G}_{\sigma}=\left(f, \boldsymbol{\Sigma}, \mathcal{Z}_{\sigma}\right)$ be a NGOS and $\widehat{\mathcal{G}}_{\ell \sigma}=\left(f, \boldsymbol{\Sigma}, \widehat{\mathcal{Z}}_{\ell \sigma}\right)$ be the NGOS after applying the SVAG operator with hyperparameter $\ell>0$. Then $\widehat{\mathcal{G}}_{\ell \sigma}$ is indeed an NGOS. That is, $\widehat{\mathcal{Z}}_{\ell \sigma}(\boldsymbol{\theta})$ is well-defined and has mean zero, covariance $\boldsymbol{\Sigma}(\boldsymbol{\theta})$.

Proof. Let $\widehat{\mathcal{Z}}_{\ell \sigma}(\boldsymbol{\theta})$ be the distribution of $\hat{\boldsymbol{z}}:=\frac{1}{\ell}\left(r_{1}(\ell) \boldsymbol{z}_{1}+r_{2}(\ell) \boldsymbol{z}_{2}\right)$ when $\boldsymbol{z}_{1}, \boldsymbol{z}_{2} \sim \mathcal{Z}_{\sigma}(\boldsymbol{\theta})$. Then it is easy to check that $\hat{\boldsymbol{g}}$ has the same distribution as $\nabla f(\boldsymbol{\theta})+\ell \sigma \hat{\boldsymbol{z}}$, since

$$
\begin{aligned}
\hat{\boldsymbol{g}}=r_{1}(\ell) \boldsymbol{g}_{1}+r_{2}(\ell) \boldsymbol{g}_{2} & \sim r_{1}(\ell)\left(\nabla f(\boldsymbol{\theta})+\sigma \boldsymbol{z}_{1}\right)+r_{2}(\ell)\left(\nabla f(\boldsymbol{\theta})+\sigma \boldsymbol{z}_{2}\right) \\
& =\left(r_{1}(\ell)+r_{2}(\ell)\right) \nabla f(\boldsymbol{\theta})+\sigma\left(r_{1}(\ell) \boldsymbol{z}_{1}+r_{2}(\ell) \boldsymbol{z}_{2}\right) \\
& =\nabla f(\boldsymbol{\theta})+\ell \sigma \hat{\boldsymbol{z}}
\end{aligned}
$$

where the last equality uses the fact that $r_{1}(\ell)+r_{2}(\ell)=1$. Hence $\widehat{\mathcal{Z}}_{\ell \sigma}(\boldsymbol{\theta})$ is well-defined.
Now we check the mean and covariance of $\hat{\boldsymbol{z}} \sim \widehat{\mathcal{Z}}_{\ell \sigma}(\boldsymbol{\theta})$. By linearity of expectation and linearity of variance (for independent variables), we have

$$
\begin{aligned}
\mathbb{E}[\hat{\boldsymbol{z}}] & =\frac{1}{\ell}\left(r_{1}(\ell) \mathbb{E}\left[\boldsymbol{z}_{1}\right]+r_{2}(\ell) \mathbb{E}\left[\boldsymbol{z}_{2}\right]\right)=\frac{1}{\ell}(\mathbf{0}+\mathbf{0})=\mathbf{0}, \\
\operatorname{Cov}(\hat{\boldsymbol{z}}) & =\frac{1}{\ell^{2}}\left(r_{1}^{2}(\ell) \operatorname{Cov}\left(\boldsymbol{z}_{1}\right)+r_{2}^{2}(\ell) \operatorname{Cov}\left(\boldsymbol{z}_{2}\right)\right)=\frac{1}{\ell^{2}}\left(r_{1}^{2}(\ell)+r_{2}^{2}(\ell)\right) \boldsymbol{\Sigma}(\boldsymbol{\theta})=\boldsymbol{\Sigma}(\boldsymbol{\theta}),
\end{aligned}
$$

where the last equality is due to $r_{1}^{2}(\ell)+r_{2}^{2}(\ell)=\left(\frac{1+\sqrt{2 \ell^{2}-1}}{2}\right)^{2}+\left(\frac{1-\sqrt{2 \ell^{2}-1}}{2}\right)^{2}=\frac{1+2 \ell^{2}-1}{2}=\ell^{2}$.
Lemma E.2. Let $\mathcal{G}_{\sigma}=\left(f, \boldsymbol{\Sigma}, \mathcal{Z}_{\sigma}\right)$ be a NGOS with scale $\sigma$. Applying the SVAG operator with hyperparameter $\ell \geq 1$, we obtain $\widehat{\mathcal{G}}_{\hat{\sigma}}=\left(f, \boldsymbol{\Sigma}, \widehat{\mathcal{Z}}_{\hat{\sigma}}\right)$ with scale $\hat{\sigma}=\ell \sigma$. Fixing $\sigma$ and changing $\ell$ produces $\widehat{\mathcal{G}}_{\hat{\sigma}}$ for all scales $\hat{\sigma} \geq \sigma$. If $\mathcal{G}_{\sigma}$ is well-behaved and satisfies the bounded moments condition, then $\widehat{\mathcal{G}}_{\hat{\sigma}}$ is also well-behaved and satisfies the bounded moments condition. Furthermore, $\widehat{\mathcal{G}}_{\hat{\sigma}}$ satisfies the low-skewness condition.

Proof. The loss function $f$ and covariance function $\Sigma$ are not changed after applying the SVAG operator, so $\widehat{\mathcal{G}}_{\hat{\sigma}}$ is well-behaved.
Now we verify the bounded moments condition. Let $\hat{\boldsymbol{z}}=\frac{1}{\ell}\left(r_{1}(\ell) \boldsymbol{z}_{1}+r_{2}(\ell) \boldsymbol{z}_{2}\right)$, where $\boldsymbol{z}_{1}, \boldsymbol{z}_{2} \sim$ $\mathcal{Z}_{\sigma}(\boldsymbol{\theta})$. Then we have

$$
\begin{aligned}
\mathbb{E}\left[\|\hat{\boldsymbol{z}}\|_{2}^{2 m}\right]^{\frac{1}{2 m}} & \leq \frac{1}{\ell}\left(\left|r_{1}(\ell)\right| \mathbb{E}\left[\left\|\boldsymbol{z}_{1}\right\|_{2}^{2 m}\right]^{\frac{1}{2 m}}+\left|r_{2}(\ell)\right| \mathbb{E}\left[\left\|\boldsymbol{z}_{2}\right\|_{2}^{2 m}\right]^{\frac{1}{2 m}}\right) \\
& \leq \frac{1}{\ell}\left(\left|r_{1}(\ell)\right|+\left|r_{2}(\ell)\right|\right) \cdot \mathbb{E}\left[\left\|\boldsymbol{z}_{1}\right\|_{2}^{2 m}\right]^{\frac{1}{2 m}} \\
& \leq(1+\sqrt{2}) \mathbb{E}\left[\left\|\boldsymbol{z}_{1}\right\|_{2}^{2 m}\right]^{\frac{1}{2 m}}
\end{aligned}
$$

By the bounded moments condition for $\mathcal{G}_{\sigma}$, there exists a constant $C_{2 m}$ such that $\mathbb{E}_{\boldsymbol{z} \sim \mathcal{Z}_{\sigma}(\boldsymbol{\theta})}\left[\|\boldsymbol{z}\|_{2}^{2 m}\right]^{\frac{1}{2 m}} \leq C_{2 m}\left(1+\|\boldsymbol{\theta}\|_{2}\right)$ for all $\boldsymbol{\theta} \in \mathbb{R}^{d}$. So $\mathbb{E}\left[\|\hat{\boldsymbol{z}}\|_{2}^{2 m}\right]^{\frac{1}{2 m}} \leq(1+\sqrt{2}) C_{2 m}\left(1+\|\boldsymbol{\theta}\|_{2}\right)$ for all $\boldsymbol{\theta} \in \mathbb{R}^{d}$.
We now verify the low skewness condition by showing that third moment of $\hat{\boldsymbol{z}}$ is $\mathcal{O}(1 / \ell)$. By the bounded moments condition with $m=2$ and Jensen's inequality,

$$
\left|\mathbb{E}\left[\boldsymbol{z}_{1}^{\otimes 3}\right]\right| \leq \mathbb{E}\left[\left\|\boldsymbol{z}_{1}\right\|_{2}^{3}\right] \leq \mathbb{E}\left[\left\|\boldsymbol{z}_{1}\right\|_{2}^{4}\right]^{3 / 4} \leq\left(C_{4}\left(1+\|\boldsymbol{\theta}\|_{2}\right)\right)^{3 / 4}
$$

So $\left|\mathbb{E}\left[\boldsymbol{z}_{1}^{\otimes 3}\right]\right|$ is bounded by $\tilde{K}_{3}(\boldsymbol{\theta}):=\left(C_{4}\left(1+\|\boldsymbol{\theta}\|_{2}\right)\right)^{3 / 4}$ of polynomial growth.
Let $r=\sqrt{2 \ell^{2}-1}$. Then $r_{1}(\ell)=\frac{1}{2}(1+r), r_{2}(\ell)=\frac{1}{2}(1-r)$. Since the third moments of two independent random vectors are additive,

$$
\begin{aligned}
\mathbb{E}\left[\hat{\boldsymbol{z}}^{\otimes 3}\right] & =\frac{1}{8 \ell^{3}} \mathbb{E}\left[\left((1+r) \boldsymbol{z}_{1}+(1-r) \boldsymbol{z}_{2}\right)^{\otimes 3}\right] \\
& =\frac{1}{8 \ell^{3}}\left(\mathbb{E}\left[(1+r)^{3} \boldsymbol{z}_{1}^{\otimes 3}\right]+\left[(1-r)^{3} \boldsymbol{z}_{2}^{\otimes 3}\right]\right) \\
& =\frac{1}{8 \ell^{3}}\left((1+r)^{3}+(1-r)^{3}\right) \mathbb{E}\left[\boldsymbol{z}_{1}^{\otimes 3}\right] \\
& =\frac{1}{8 \ell^{3}}\left(2+6 r^{2}\right) \mathbb{E}\left[\boldsymbol{z}_{1}^{\otimes 3}\right] \\
& =\frac{1}{8 \ell^{3}}\left(12 \ell^{2}-4\right) \mathbb{E}\left[\boldsymbol{z}_{1}^{\otimes 3}\right] \leq \frac{1}{\ell} \tilde{K}_{3}(\boldsymbol{\theta})
\end{aligned}
$$

where the 4th equality is due to $(1+r)^{3}+(1-r)^{3}=\left(1+3 r+3 r^{2}+r^{3}\right)+\left(1-3 r+3 r^{2}-r^{3}\right)=2+6 r^{2}$. Therefore, the low skewness condition is verified.

## F Miscellaneous Theoretical Arguments

## F. 1 How does the noise scale change with batch size?

We discuss how the noise scale $\sigma$ in the NGOS (Definition 2.3) changes when the batch size changes: in particular, $\sigma \sim 1 / \sqrt{B}$. The argument follows from the linearity of covariance and is an already well-known result - we reproduce it here for clarity. We first fix a parameter vector $\boldsymbol{\theta}$. Let $\boldsymbol{g}^{(1)}, \ldots, \boldsymbol{g}^{(B)}$ be the gradients evaluated at data points from a batch of size $B$, where $\boldsymbol{g}^{(b)}=\nabla f(\boldsymbol{\theta})+\boldsymbol{z}^{(b)}$, and every $\boldsymbol{z}^{(b)}$ is a gradient noise vector drawn i.i.d. from a distribution with mean $\mathbf{0}$ and covariance $\boldsymbol{\Sigma}(\boldsymbol{\theta})$. For sampling with replacement on a finite dataset of size $n$, where $f_{1}(\boldsymbol{\theta}), \ldots, f_{n}(\boldsymbol{\theta})$ are the loss functions for the $n$ data points (and the average of these $n$ functions is $f(\boldsymbol{\theta})$ ), this covariance matrix can be explicitly written as:

$$
\boldsymbol{\Sigma}(\boldsymbol{\theta})=\frac{1}{n} \sum_{i=1}^{n}\left(\nabla f_{i}(\boldsymbol{\theta})-\nabla f(\boldsymbol{\theta})\right)\left(\nabla f_{i}(\boldsymbol{\theta})-\nabla f(\boldsymbol{\theta})\right)^{\top}
$$

The average gradient over the batch is $\boldsymbol{g}:=\frac{1}{B} \sum_{b=1}^{B} \boldsymbol{g}^{(b)}=\nabla f(\boldsymbol{\theta})+\frac{1}{B} \sum_{b=1}^{B} \boldsymbol{z}^{(b)}$. As $\boldsymbol{z}^{(1)}, \ldots, \boldsymbol{z}^{(B)}$ are sampled i.i.d., their average $\frac{1}{B} \sum_{b=1}^{B} \boldsymbol{z}^{(b)}$ has mean $\mathbf{0}$ and covariance $\boldsymbol{\Sigma}(\boldsymbol{\theta})$ by linearity of expectation and covariance. We can set $\mathcal{Z}_{\sigma}(\boldsymbol{\theta})$ to be the distribution of the random variable $\frac{1}{\sqrt{B}} \sum_{b=1}^{B} \boldsymbol{z}_{b}$, where $\sigma=\frac{1}{\sqrt{B}}$, then $\boldsymbol{g}$ has the same distribution as the stochastic gradient produced by the NGOS $\mathcal{G}_{\sigma}=\left(f, \boldsymbol{\Sigma}, \mathcal{Z}_{\sigma}\right)$.

## F. 2 What happens when the noise does not dominate the gradient?

We discuss the linear warm-up setting described in Section 4.1. Recall that when ignoring the effect of $\epsilon$, the RMSprop update can be written as

$$
\boldsymbol{\theta}_{k+1} \approx \boldsymbol{\theta}_{k}-\eta \boldsymbol{g}_{k} \odot\left(\overline{\boldsymbol{g}}^{2}+\sigma^{2} \mathbf{1}\right)^{-1 / 2}
$$

From the above equation, it is clear that the dynamics of $\boldsymbol{\theta}$ depends on the relationship between the noise scale $\sigma$ and the gradient $\|\overline{\boldsymbol{g}}\|$. In Section 4.1, we discuss the case where $\sigma \gg\|\overline{\boldsymbol{g}}\|$, which is the regime where the SDE approximation can exist.


Figure 4: We compare the norm of the average gradient with the noise scale for different batch sizes during training of ResNet-50 and VGG-16 model with RMSprop on the CIFAR-10 dataset. Here, $(\eta, \beta)=\left(10^{-3}, 0.999\right)$ for batch size 256 and scaled with our proposed square root scaling rule (Definition 5.1) for the other batch sizes. We show the results for $\epsilon$ at both small (of order $10^{-30}$ ) and large scale (of order $10^{-8}$ ). We observe that for small batches, the noise in the gradient dominates the signal in the gradient, supporting our hypothesis. For larger batches, the hypothesis seems to hold true towards the end of training.

Here, we argue that when $\sigma \ll\|\overline{\boldsymbol{g}}\|$, no SDE approximation can exist for the discrete trajectory. In this case, the RMSprop update would instead be $\boldsymbol{\theta}_{k+1} \approx \boldsymbol{\theta}_{k}-\eta \boldsymbol{U}^{-1} \boldsymbol{g}_{k}$, where $\boldsymbol{U}=\operatorname{diag}\left(\sqrt{\overline{\boldsymbol{g}}^{2}}\right)$. Combining this with $\boldsymbol{g}_{k} \sim \mathcal{N}(\overline{\boldsymbol{g}}, \sigma \boldsymbol{I})$ yields that $\boldsymbol{\theta}_{k+1}-\boldsymbol{\theta}_{k} \sim \mathcal{N}\left(\eta \boldsymbol{U}^{-1} \overline{\boldsymbol{g}}, \eta^{2} \sigma^{2} \boldsymbol{U}^{-2}\right)$ approximately. We can again take a telescoping sum to obtain the marginal distribution of $\boldsymbol{\theta}_{k}$ : $\boldsymbol{\theta}_{k} \sim \mathcal{N}\left(k \eta \boldsymbol{U}^{-1} \overline{\boldsymbol{g}}, k \eta^{2} \sigma^{2} \boldsymbol{U}^{-2}\right)$ approximately.

However, it is impossible to make the above distribution fixed even as $\sigma$ changes, so no SDE approximation exists. In particular, we need to make both $k \eta$ and $k \eta^{2} \sigma^{2}$ fixed, so $\eta \sigma$ must be a constant. If $\sigma \ll\|\overline{\boldsymbol{g}}\|$, then $\frac{1}{\eta} \ll\|\overline{\boldsymbol{g}}\|$ too. This requirement on $\eta$ implies that no SDE approximation exists when $\sigma \ll\|\overline{\boldsymbol{g}}\|$, and hence motivates us to study the case of $\sigma \gg\|\overline{\boldsymbol{g}}\|$.

## G Experimental Verification of Assumptions

In this section, we take measurements and perform experiments to verify that the various assumptions made in our theory do not harm the applicability of our findings to realistic settings.

## G. 1 Noise Dominates the Gradient

Our analysis in Section 4.1 suggests that an SDE approximation cannot exist when the the gradient $\overline{\boldsymbol{g}}$ dominates the noise scale $\sigma$. Note that Section 4.1 performs a rough analysis under the assumption of a linear loss function (i.e., fixed gradient throughout training), which is far from practice. In the more general setting, we require that for every step $k, \mathbb{E}\left\|\boldsymbol{z}_{k}\right\|^{2}$ (i.e., the gradient variance) dominates $\left\|\mathbb{E} \boldsymbol{g}_{k}\right\|$ (i.e., the norm of the average gradient), where the expectations are taken over sampling seeds, in order for the SDE approximation to exist. Figure 4 shows that our assumption holds for small batches and for large batches near the end of training.


Figure 5: We compare the norm of the average gradient with the noise scale for different batch sizes during training of ResNet-50 model with Adam on the CIFAR-10 dataset. Here, $\left(\eta, \beta_{1}, \beta_{2}\right)=\left(10^{-3}, 0.999,0.999\right)$ for batch size 256 and scaled with our proposed square root scaling rule (Definition 5.1) for the other batch sizes. We show the results for $\epsilon$ at both small (of order $10^{-30}$ ) and large scale (of order $10^{-8}$ ). We observe that for small batches, the noise in the gradient dominates the signal in the gradient, supporting our hypothesis. For larger batches, the hypothesis seems to hold true towards the end of training.

## G. 2 Using $\boldsymbol{v}_{k}$ instead of $\boldsymbol{v}_{k+1}$ in the update rule

In Definitions 2.1 and 2.2, we slightly modify the standard implementation of RMSprop and Adam by using $\boldsymbol{v}_{k}$ in the update rule instead of $\boldsymbol{v}_{k+1}$. Here, we verify for Adam that this modification of the optimization algorithms does not significantly harm performance. Figure 6shows the behavior of ResNet-50 and VGG-16 trained with Adam on CIFAR-10 with the above modification of the optimization algorithm. We observe a small drop ( $\approx 1 \%$ ) in test accuracies. However, the behavior of the trajectories across different batch sizes for the proposed scaling rule stays the same, i.e. we observe a maximum of $3 \%$ test accuracy gap between training batch size 256 and 8192 . Moreover, the behavior of the test functions match across the trajectories of different batch sizes.

## H SVAG Experiments

Recall that the SVAG algorithm (Definition 6.2) is a computationally efficient simulation of the SDEs corresponding to RMSprop and Adam. The SVAG algorithm requires a hyperparameter $\ell$, and the resulting parameters after $k \ell^{2}$ steps should match the parameters on the corresponding discrete optimization trajectory after $k$ steps. In particular, Theorem 6.3 shows that the SVAG algorithm is an order-1 weak approximation (Definition 2.4) of the SDE, and the approximation error scales as $1 / \ell$. One may be initially concerned that realistic deep learning settings require $1 / \ell$ to be very small, which would make $\ell$ large and hence computationally intractable. Li et al. (2021) showed that the SVAG trajectories appear to converge to the SDE trajectory for computationally tractable small values of $\ell$. We similarly find that our proposed SVAG-like algorithms in Definition 6.2 appear to converge for small $\ell$ in various settings.

CIFAR-10. Figures 7 and 8 show that SVAG converges and closely tracks RMSprop at smaller $\epsilon\left(=10^{-30}\right)$, and at larger $\epsilon\left(=10^{-8}\right)$ respectively. Figures 3 and 9 show that SVAG converges and closely tracks Adam at smaller $\epsilon\left(=10^{-30}\right)$, and at larger $\epsilon\left(=10^{-8}\right)$ respectively. All experiments follow the setting in Appendix J. 2 for batch size 256.


Figure 6: We repeat the Square Root Scaling experiments with Adam for ResNet-50 and VGG-16 on CIFAR-10 dataset, where we slightly modify the standard implementation of Adam by using $\boldsymbol{v}_{k}$ in the update rule instead of $\boldsymbol{v}_{k+1}$. For batch size $256,\left(\eta, \epsilon, \beta_{1}, \beta_{2}\right)=\left(10^{-3}, 10^{-8}, 0.999,0.999\right)$ and the hyperparameters are scaled according to the square root scaling rule for other batch sizes. We observe a small drop ( $\approx 1 \%$ ) in test accuracies. However, the behavior of the trajectories across different batch sizes stays the same.


Figure 7: SVAG experiments on ResNet-50 trained on CIFAR-10 with RMSprop. We use batch size 256 and the hyperparameters $\eta=10^{-3}, \beta=0.999, \epsilon=10^{-30}$ and a weight decay of $10^{-4}$. Since SVAG takes $\ell$ smaller steps to simulate the continuous dynamics in $\eta$ time, we plot accuracy against continuous time defined as $k \times \eta^{2} / \ell^{2}$.

Wikipedia + Books (Academic BERT). Figure 10 shows that SVAG converges and closely tracks Adam. We use the experimental setting for batch size 1024 in Appendix J. 4 except the hyperparameters $\beta_{1}$ and $\beta_{2}$ are fixed at 0.9 and 0.98 respectively.

WikiText-103 (GPT). Figure 23 shows that SVAG converges and closely tracks Adam. We use the experimental setting for batch size 1024 in Appendix J. 4 except the hyperparameters $\beta_{1}$ and $\beta_{2}$ are fixed at 0.9 and 0.98 respectively. Additionally, for computational reasons, we pretrain on sequences of length 64 .

## I Square Root Scaling Experiments

We experimentally evaluate the scaling rules proposed in Definitions 5.1 and 5.2 by training models with different batch sizes and modifying the optimization hyperparameters accordingly. The number of gradient steps were modified to keep the total amount of continuous time same across all the


Figure 8: SVAG experiments on ResNet-50 trained on CIFAR-10 with RMSprop. We use batch size 256 and the hyperparameters $\eta=10^{-3}, \beta=0.999, \epsilon=10^{-8}$, and a weight decay of $10^{-4}$. Since SVAG takes $\ell$ smaller steps to simulate the continuous dynamics in $\eta$ time, we plot accuracy against continuous time defined as $k \times \eta^{2} / \ell^{2}$.


Figure 9: SVAG experiments on ResNet-50 trained on CIFAR-10 with Adam. We use batch size 256 and the hyperparameters $\eta=10^{-3}, \beta_{1}=0.9, \beta_{2}=0.999, \epsilon=10^{-30}$, and a weight decay of $10^{-4}$. Since SVAG takes $\ell$ smaller steps to simulate the continuous dynamics in $\eta$ time, we plot accuracy against continuous time defined as $k \times \eta^{2} / \ell^{2}$.


Figure 10: SVAG experiments on RoBERTa pretrained on Bookcorpus + Wikipedia dataset with Adam. We use batch size 1024 and the hyperparameters $\eta=10^{-3}, \beta_{1}=0.9, \beta_{2}=0.98, \epsilon=$ $2 \times 10^{-6}$. Since SVAG takes $\ell$ smaller steps to simulate the continuous dynamics in $\eta$ time, we plot accuracy against continuous time defined as $k \times \eta^{2} / \ell^{2}$.


Figure 11: SVAG experiments on GPT pretrained on WikiText-103 with Adam. We use batch size 1024 and the hyperparameters $\eta=10^{-3}, \beta_{1}=0.9, \beta_{2}=0.98, \epsilon=2 \times 10^{-8}$. Since SVAG takes $\ell$ smaller steps to simulate the continuous dynamics in $\eta$ time, we plot accuracy against continuous time defined as $k \times \eta^{2} / \ell^{2}$. For computational constraints, we train on sequences of length 64 .
batches. The warmup schedule and the learning rate decay schedule were kept the same. See Appendix JJor details about the baseline runs for each dataset.
We perform two types of ablation studies. In Appendix I.1, we compare our proposed scaling rule to variants that only scale some subset of the optimization hyperparameters. We find that our proposed scaling rule is the best at preserving the validation accuracy and the other test functions across different batch sizes. In Appendix I.2, we compare the proposed square root scaling rule to a linear one and find that they perform comparably on CIFAR-10, though the square root scaling significantly outperforms linear scaling on the ImageNet dataset. We hypothesize that for simpler datasets, like CIFAR-10, the differences between the two scaling rules is not reflected clearly in the validation accuracies because the task is too easy to learn.

CIFAR-10. Figures 12 and 14 show the performance of VGG-16 when trained with different batch sizes with RMSprop at smaller $\epsilon\left(=10^{-30}\right)$, and at larger $\epsilon\left(=10^{-8}\right)$ respectively. Figures 13 and 15 show the performance of VGG-16 when trained with different batch sizes with Adam at smaller $\epsilon\left(=10^{-30}\right)$, and at larger $\epsilon\left(=10^{-8}\right)$ respectively. For the corresponding experiments on ResNet-50, please refer to Figures 16 to 19 For the details on the experimental setting, please refer to Appendix J. 2 .

We observe that in all settings the test accuracies vary by at most $3 \%$ across batch sizes when using the proposed square root scaling rule. Moreover, the test functions stay close across multiple batch sizes, signifying that the trajectories stay close across the batch sizes using the Square root scaling rule.

ImageNet. Figures 20 and 21 show the performance of ResNet- 50 when trained with different batch sizes with RMSprop at smaller $\epsilon\left(=10^{-30}\right)$, and at larger $\epsilon\left(=10^{-8}\right)$ respectively. For the details on the experimental setting, please refer to Appendix J. 3
We observe that on the validation set, the loss and accuracy behavior for the model is very similar, when trained with different batch sizes. Moreover, the difference between the validation accuracies is atmost $3 \%$ between the batch sizes $256,1024,4096$, and 16384 for smaller $\epsilon$. The difference between the validation accuracies is atmost $1.5 \%$ between the batch sizes $256,1024,4096,16384$, and 32768 for larger $\epsilon$.

Wikipedia + Books (Academic BERT). Figure 22 shows the performance of a RoBERTa model when pretrained with different batch sizes with ADAM. The scaling rule is applied to modify the peak values of the optimization hyperparameters. We also evaluate the pretrained models on several downstream tasks, and show the results in Table 1. For the details on the setting for both pretraining and downstream experiments, please refer to Appendix J. 4

| Pretrain batch size $B$ | CoLA | SST-2 | MRPC | STS-B | QQP | MNLI | QNLI | RTE |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 1024 | 0.585 | 0.92 | 0.73 | 0.866 | 0.873 | 0.836 | 0.906 | 0.682 |
| 2048 | 0.563 | 0.928 | 0.803 | 0.869 | 0.875 | 0.826 | 0.897 | 0.653 |
| 4096 | 0.581 | 0.921 | 0.778 | 0.869 | 0.875 | 0.839 | 0.892 | 0.675 |
| 8192 | 0.626 | 0.929 | 0.778 | 0.884 | 0.877 | - | 0.9 | 0.675 |

Table 1: Performance of the pretrained RoBERTa models when finetuned on different downstream tasks in GLUE Wang et al. (2019). F1 scores are reported for QQP and MRPC, Spearman correlations are reported for STS-B, Matthews correlations for CoLA, and accuracy scores are reported for the other tasks. Here, $B$ denotes the batch size used for pretraining. We run an extensive grid search (Table 4 ) to find the best performance of each pretrained model on each of the downstream tasks.

We observe that the log perplexity on the training and validation datasets matches across different batch sizes throughout pretraining. Moreover, we also observe that the gradient norms across different batch sizes remain close throughout the pretraining. Furthermore, the models pretrained across different batch sizes can achieve very similar performance in the downstream tasks.

WikiText-103 (Academic GPT). Figure 23 shows the log perplexity behavior of GPT on training and validation datasets, across different batch sizes of pretraining. We observe that the log perplexity matches across different batch sizes. Moreover, we also observe that the gradient norms across different batch sizes remain close throughout the training.


Figure 12: VGG-16 trained on CIFAR-10 using RMSprop are close for different batch sizes when the optimization hyperparameters are varied according to the proposed scaling rule for RMSprop (Definition 5.2). We use a baseline setting of $\eta=10^{-3}, \epsilon=10^{-8}$, and $\beta=0.999$ for batch size 256 . We use a weight decay factor of $10^{-4}$. We observe a gap of at most $3 \%$ among the different batch sizes under consideration.


Figure 13: VGG-16 trained on CIFAR-10 using Adam are close for different batch sizes when the optimization hyperparameters are varied according to the proposed scaling rule for Adam (Definition 5.2. We use a baseline setting of $\eta=10^{-3}, \epsilon=10^{-8}$, and $\left(\beta_{1}, \beta_{2}\right)=(0.999,0.999)$ for batch size 256 . We use a weight decay factor of $10^{-4}$. We observe a gap of at most $3 \%$ among the different batch sizes under consideration.


Figure 14: VGG-16 trained on CIFAR-10 using RMSprop are close for different batch sizes when the optimization hyperparameters are varied according to the proposed scaling rule for RMSprop (Definition 5.2). We use a baseline setting of $\eta=10^{-3}, \epsilon=10^{-30}$, and $\beta=0.999$ for batch size 256. We use a weight decay factor of $10^{-4}$. We observe a gap of at most $3 \%$ among the different batch sizes under consideration.


Figure 15: VGG-16 trained on CIFAR-10 using Adam are close for different batch sizes when the optimization hyperparameters are varied according to the proposed scaling rule for Adam (Definition 5.2. We use a baseline setting of $\eta=10^{-3}, \epsilon=10^{-30}$, and $\left(\beta_{1}, \beta_{2}\right)=(0.999,0.999)$ for batch size 256 . We use a weight decay factor of $10^{-4}$. We observe a gap of at most $3 \%$ among the different batch sizes under consideration.


Figure 16: ResNet-50 trained on CIFAR-10 using RMSprop are close for different batch sizes when the optimization hyperparameters are varied according to the proposed scaling rule for RMSprop (Definition 5.2. We use a baseline setting of $\eta=10^{-3}, \epsilon=10^{-8}$, and $\beta=0.999$ for batch size 256. $\epsilon=10^{-30} \approx 0$ for all experiments. We use a weight decay factor of $10^{-4}$. We observe a gap of at most $3 \%$ among the different batch sizes under consideration.


Figure 17: ResNet-50 trained on CIFAR-10 using Adam are close for different batch sizes when the optimization hyperparameters are varied according to the proposed scaling rule for Adam (Definition 5.2). We use a baseline setting of $\eta=10^{-3}$ and $\left(\beta_{1}, \beta_{2}\right)=(0.999,0.999)$ for batch size 256. $\epsilon=10^{-30} \approx 0$ for all experiments. We use a weight decay factor of $10^{-4}$. We observe a gap of at most $3 \%$ among the different batch sizes under consideration.


Figure 18: ResNet-50 trained on CIFAR-10 using RMSprop are close for different batch sizes when the optimization hyperparameters are varied according to the proposed scaling rule for RMSprop (Definition 5.2). We use a baseline setting of $\eta=10^{-3}, \epsilon=10^{-8}$, and $\beta=0.999$ for batch size 256 . We use a weight decay factor of $10^{-4}$. We observe a gap of at most $3 \%$ among the different batch sizes under consideration.


Figure 19: ResNet-50 trained on CIFAR-10 using Adam are close for different batch sizes when the optimization hyperparameters are varied according to the proposed scaling rule for Adam (Definition 5.2. We use a baseline setting of $\eta=10^{-3}, \epsilon=10^{-8}$, and $\left(\beta_{1}, \beta_{2}\right)=(0.999,0.999)$ for batch size 256 . We use a weight decay factor of $10^{-4}$. We observe a gap of at most $3 \%$ among the different batch sizes under consideration.


Figure 20: ResNet-50 trained on ImageNet using Adam are close for different batch sizes when the optimization hyperparameters are varied according to the proposed scaling rule for Adam (Definition 5.2. We use a baseline setting of $\eta=3 \times 10^{-4}, \epsilon=10^{-8}$, and $\left(\beta_{1}, \beta_{2}\right)=(0.999,0.999)$ for batch size 256 . We use a weight decay factor of $10^{-4}$. We achieve around $74 \%$ validation accuracy with batch size 256 and the accuracy drops by at most $1.5 \%$ at batch size 32768 .


Figure 21: ResNet-50 trained on ImageNet using Adam are close for different batch sizes when the optimization hyperparameters are varied according to the proposed scaling rule for Adam (Definition 5.2. We use a baseline setting of $\eta=3 \times 10^{-4}$ and $\left(\beta_{1}, \beta_{2}\right)=(0.999,0.999)$ for batch size 256. $\epsilon=10^{-30} \approx 0$ for all experiments. We use a weight decay factor of $10^{-4}$. We achieve around $74 \%$ validation accuracy with batch size 1024 and the accuracy drops by at most $3 \%$ at batch size 16384.

## I. 1 Ablation study on the proposed scaling rule

CIFAR-10. We conduct an ablation study on whether all the parameters $\eta, \epsilon, \beta_{1}, \beta_{2}$ need to be scaled in our proposed scaling rule. To do so, we compare the performance of a ResNet-50 model trained with batch size 256 and hypeparameters $\left(\epsilon=10^{-8}, \beta_{1}=0.999, \beta_{2}=0.999\right.$ ) with the performance at a larger batch size, across 5 runs representing 5 different scaling rules: (a) Scale $\eta$, keeping others fixed, (b) Scale $\eta, \epsilon$, keeping others fixed, (c) Scale $\eta, \epsilon, \beta_{1}$, keeping others fixed, (d) Scale $\eta, \epsilon, \beta_{2}$, keeping others fixed, and (e) Scale $\eta, \epsilon, \beta_{1}, \beta_{2}$. Please check the behaviors of the different scaling rules at batch size 2048, 4096, 8192 and 16384 in Appendix I.2. We found that (e) consistently beats others in terms of the test functions and the validation accuracies at all batch sizes. The closest scaling rule (c) involved scaling only $\eta, \epsilon$, and $\beta_{1}$ while keeping $\beta_{2}$ fixed.

## I. 2 Ablation against linear scaling rules

CIFAR-10. We compared the proposed scaling rule against possible linear scaling rules, that scale the hyperparameters linearly with the increase in training batch size. We focused on ResNet-50 training with Adam. The linear scaling rules that we tried were: (a) Scale $\eta$ linearly, keeping $\beta_{1}, \beta_{2}, \epsilon$


Figure 22: The train and validation log perplexities of RoBERTa-large trained on the Wiki+Books corpus using Adam (for 48 hours) are close for moderate batch sizes using the Square Root Scaling Rule on Adam. $\eta=10^{-3}$ and $\left(\beta_{1}, \beta_{2}\right)=(0.99375,0.996)$ for batch size $1024, \epsilon=2 \times 10^{-6}$ for batch size 1024 and scaled likewise for other batch sizes. We achieve a validation log perplexity of $2.1 \pm 0.1$ for batch size $1024,2048,4096$ and 8192 . Training with batch size 256 is computationally inefficient, but follows the same behavior during its 48-hour trajectory.


Figure 23: The train and validation perplexities of 12 layer GPT trained on the Wikitext corpus using Adam (for 48 hours) are close for moderate batch sizes Square Root Scaling Rule on Adam on RoBERTa-large. $\eta=10^{-3}$ and $\left(\beta_{1}, \beta_{2}\right)=(0.9875,0.996)$ for batch size $1024, \epsilon=2 \times 10^{-8}$ for batch size 1024 and scaled likewise for other batch sizes. We achieve a validation log perplexity of $5 \pm 0.1$ for all the batch sizes under consideration. Moreover, we observe an alignment in the behavior of the trajectories across the different batch sizes (except 8192 in the first half of the training).
fixed, (b) Scale $\eta, 1-\beta_{1}$ linearly, keeping $\beta_{2}, \epsilon$ fixed, (c) Scale $\eta, 1-\beta_{2}$ linearly, keeping $\beta_{1}, \epsilon$ fixed, and (d) Scale $\eta, 1-\beta_{1}, 1-\beta_{2}$ linearly, keeping $\epsilon$ fixed. Fig. 25 shows the behavior of these scaling rules at batch size 8192 and 16384, at different values of $\epsilon$.
The linear scaling rule in (d) seems to perform as well as the proposed Square Root scaling rule (Definition 5.2 ) in terms of validation accuracies. However, with a closer look on the train accuracy plots, we observe that the Square Root scaling rule tracks the smaller batch training trajectory better than the linear scaling rule. The linear scaling rules seem to catch up, only after the learning rate is decayed. Our hypothesis is that the CIFAR-10 dataset is simple enough for different scaling rules to work well.

ImageNet. We conduct ablation experiments on ResNet-50, trained with Adam on ImageNet, where we follow a linear scaling rule to scale the hyperparameters across batch sizes. Due to computational issues, we didn't conduct extensive experiments, as was done for CIFAR-10. The linear scaling rule for Adam is as follows: the hyperparameters $\eta, 1-\beta_{1}$ and $1-\beta_{2}$ are scaled by $\kappa$, when the batch size is scaled by $\kappa$, and $\epsilon$ isn't scaled. As was noted earlier, the definition of continuous time will change to $\eta \times \#$ gradient steps $\}$. We keep the number of training epochs equal to 90 as before, follow the
same learning rate schedule, and show the performance of the models in Figure 26. We observe that scaling the hyperparameters to larger batch training with the proposed LSR results in a big drop in validation accuracies.

## J Experiment Configuration Details

## J. 1 A note on learning rate schedule and warm-up

We have used a learning rate schedule and a warm-up phase in all our experiments. We have to admit that our current theorems do not directly apply to time-varying learning rates or batch sizes. But our experiments demonstrate that our scaling rules continue to hold for learning rate schedulers with a special warm-up, even if they go beyond the scope of our theoretical setting. Technically, the extensions of our theorems to time-varying learning rates or batch sizes are interesting, and we believe they can indeed be shown following the same proof strategy. The corresponding SDE approximations should have hyperparameters changing with time.

## J. 2 CIFAR-10

There are 50000 images in the training set and 10000 images in the validation set of CIFAR-10 Krizhevsky et al.

Architecture. We used the architecture of ResNet-56 from He et al. (2015a) without modification. We used the same architecture of VGG-16 with batch normalization from Simonyan and Zisserman (2014). However, we kept the final layer of the architecture fixed throughout training, to make the model 1-homogenous and avoid the optimization difficulties of 2-homogenous networks (Li and Arora (2020).

RMSprop. To fix a baseline to compare against, we first trained the models with batches of size 256 , sampled with replacement, at peak learning rate $\eta=10^{-3}$ and $\beta=0.999$. The model was trained for $500 \times\lfloor(50000 / 256)\rfloor=97500$ gradient steps (or 500 epochs). We followed an initial warmup for the first $2 \%$ of the total gradient steps. The learning rate schedule during the warmup phase is given by $\eta \times 10^{-3} \times\left(10^{3}\right)^{\# \text { epochs } / 10}$. We also followed a learning rate decay schedule, the learning rate was decayed by 0.1 when the model reaches $80 \%$ ( 400 epoch) and $90 \%$ ( 450 epoch) of the total continuous time respectively. We did experiments at two values of $\epsilon$, small $\left(=10^{-30}\right)$ and large $\left(=10^{-8}\right)$.

We then made multiple runs of the same model with batches of varying sizes in $\{1024,4096,16384\}$, with the hyperparameters $\eta, \epsilon$ and $\beta$ modified appropriately according to the scaling rule. The number of gradient steps were modified to keep the total amount of continuous time same across all the batches (which amounted to 500 epochs by the equivalence between continuous time and the number of training epochs). The warmup schedule and the learning rate decay schedule were kept the same.

Adam. We first trained the models with batches of size 256, sampled with replacement, at peak learning rate $\eta=10^{-3}$ and $\left(\beta_{1}, \beta_{2}\right)=(0.999,0.999)$. The total continuous time of training (or number of epochs), the amount of continuous time in the warmup phase, and the learning rate schedule in the warmup phase were same as RMSprop. The only difference was the learning rate schedule after the warmup phase, the learning rate was decayed by 0.1 when the model reaches $60 \%$ of the total continuous time (or 300 epochs).

We then made multiple runs of the same model with batches of varying sizes in $\{1024,4096,16384\}$, with the hyperparameters $\eta, \epsilon$ and $\beta_{1}, \beta_{2}$ modified appropriately according to the scaling rule. The number of gradient steps were modified to keep the total amount of continuous time same across all the batches. The warmup schedule and the learning rate decay schedule were kept the same.

## J. 3 ImageNet

There are 1281167 images in the training set and 50000 images in the validation set of ImageNet (Deng et al., 2009).


Figure 24: Ablation study for the square root scaling rule on Resnet-50 trained with Adam on CIFAR-10. We compare the performance of a model trained with batch size 256 and hypeparameters $\left(\epsilon=10^{-8}, \beta_{1}=0.999, \beta_{2}=0.999\right)$ with the performance at a larger batch size, across 5 runs representing 5 variations of the square root scaling rule: (a) Scale $\eta$, keeping others fixed, (b) Scale $\eta, \epsilon$, keeping others fixed, (c) Scale $\eta, \epsilon, \beta_{1}$, keeping others fixed, (d) Scale $\eta, \epsilon, \beta_{2}$, keeping others fixed, and (e) Scale $\eta, \epsilon, \beta_{1}, \beta_{2}$. We use a weight decay of $10^{-4}$ in all the experiments. We observe that scaling all the hyperparameters consistently gives better performance at higher batch size. Scaling rule (c) is close second.


Figure 25: Ablation study against (possible) linear scaling rules on Resnet-50 trained with Adam on CIFAR-10. We compare the performance of a model trained with batch size 256 and hypeparameters $\left(\epsilon=10^{-8} / 10^{-30}, \beta_{1}=0.999, \beta_{2}=0.999\right)$ with the performance at a larger batch size, across 5 runs representing 5 possible linear scaling rules: (a) Scale $\eta$ linearly, keeping $\beta_{1}, \beta_{2}, \epsilon$ fixed, (b) Scale $\eta, 1-\beta_{1}$ linearly, keeping $\beta_{2}, \epsilon$ fixed, (c) Scale $\eta, 1-\beta_{2}$ linearly, keeping $\beta_{1}, \epsilon$ fixed, and (d) Scale $\eta, 1-\beta_{1}, 1-\beta_{2}$ linearly, keeping $\epsilon$ fixed. We use a weight decay of $10^{-4}$ in all the experiments. We also compare the behavior of the linear scaling rules against the square root scaling rule. Since the continuous time definition varies across the scaling rules, we plot against the number of epochs trained. A closer look at the training accuracy plots shows that the square root scaling rule tracks the smaller batch training trajectory better than the linear scaling rule.


Figure 26: Ablation study against a (possible) linear scaling rule on Resnet-50 trained with Adam on Imagenet. We compare the performance of a model trained with batch size 256 and hypeparameters ( $\eta=3 \times 10^{-4}, \epsilon=10^{-8}$, and $\left(\beta_{1}, \beta_{2}\right)=(0.999,0.999)$ ) with the performance at a larger batch size, when the hyperparameters are scaled as follows: scale $\eta, 1-\beta_{1}, 1-\beta_{2}$ by $\kappa$, if the batch size is scaled by $\kappa$, keeping $\epsilon$ fixed. We use a weight decay factor of $10^{-4}$. We clearly observe a decrease in performance at larger batch size, in comparison to a model trained with square root scaling rule (see Figure 20).

Architecture. We trained a ResNet-50 (He et al. 2015a) model without modification.

Adam. To fix a baseline to compare against, we first trained the models with batches of size 256, sampled with replacement, at learning rate $\eta=3 \times 10^{-4}$ and $\left(\beta_{1}, \beta_{2}\right)=(0.999,0.999)$. The model was trained for a total of $90 \times\lfloor(1281167 / 256)\rfloor=450360$ gradient steps (or 90 epochs). We followed an initial warmup for the first $\frac{1}{18}$ fraction of the total continuous time (or the first 5 epochs). The learning rate schedule during the warmup phase is increased linearly with epoch, i.e. the learning rate is given by $\eta \times 10^{-3} \times\left(10^{3}\right)^{\# \text { epochs } / 5}$. We then followed a learning rate decay schedule, where the learning rate was decayed by 0.1 when the model reaches $\frac{5}{18}$ fraction (or 50 epoch ) and $\frac{8}{9}$ fraction (or 80 epoch) of the total continuous time respectively. We use two different values for $\epsilon$, small $\epsilon\left(=10^{-30} \approx 0\right)$ and a larger $\epsilon\left(=10^{-8}\right)$.

## J. 4 Books and Wikipedia (Academic BERT)

We use a combination of Bookcorpus (Zhu et al. 2015) plus English Wikipedia, which totals 16 GB of uncompressed text. We split the data uniformly with a ratio $9: 1$, to create training and validation datasets for pretraining.

Architecture. We pretrain a 24-layer RoBERTa Liu et al. (2019) model. We pretrain on sequences of length 128 .

Pre-training with Adam. We use the code from Wettig et al. (2022). We follow the optimization recipe from Izsak et al. (2021) for efficient pre-training. To fix a baseline, we first trained our model with batch size 1024, with the optimization parameters given in table 2 . In the warmup phase, the learning rate is increased linearly over the interval, i.e. the learning rate at step $k$ in the warmup phase is given by $\frac{k}{k_{\text {warmup }}} \eta$, where $k_{\text {warmup }}$ denotes the total number of warmup steps and $\eta$ denotes the peak learning rate. Moreover, after the warmup phase, the learning rate is decayed linearly to 0 , i.e. the learning rate at step $k$ after the warmup phase is given by $\frac{k-k_{\text {warmup }}}{k_{\text {max }}-k_{\text {warmup }}} \eta$, where $\eta$ denotes the peak learning rate and $k_{\text {max }}$ denotes the maximum number of gradient steps intended for pretraining.
Fine-tuning with Adam: We also validate the performance of the pretrained models from the previous section on the GLUE(Wang et al. 2019) datasets. For each downstream task, we run an extensive grid search on the hyperparameters for finetuning each pretrained model. We focused on the following hyperparameters for grid search: batch size, the peak learning rate and the total number of training epochs. Please see Table 4 for the hyperparameter grid. The rest of the hyperparameters

| Hyperparameter | Value |
| :---: | :---: |
| Dropout | 0.1 |
| Attention Dropout | 0.1 |
| Warmup Steps | 5520 |
| Peak Learning Rate | $10^{-3}$ |
| Batch Size | 1024 |
| Weight Decay | $10^{-4}$ |
| Max Steps | 92000 |
| Learning Rate Decay | Linear |
| Adam $\beta_{1}$ | 0.99375 |
| Adam $\beta_{2}$ | 0.996 |
| Adam $\epsilon$ | $2 \times 10^{-6}$ |
| Gradient Clipping | 0.0 |
| Position embeddings | 128 |
|  |  |

Table 2: Optimization hyperparameters of baseline RoBERTa model during pretraining.

| Hyperparameter | Value |
| :---: | :---: |
| Dropout | 0.1 |
| Attention Dropout | 0.1 |
| Warmup Steps | $6 \%$ of total |
| Weight Decay | 0.1 |
| Learning Rate Decay | Linear |
| Adam $\beta_{1}$ | 0.9 |
| Adam $\beta_{2}$ | 0.98 |
| Adam $\epsilon$ | $10^{-6}$ |
| Gradient Clipping | 0.0 |
|  |  |

Table 3: Optimization hyperparameters of all pretrained RoBERTa models during finetuning.
are fixed for all the runs and are given in Table 3 We follow a similar learning rate schedule during the warmup phase as was used for pretraining: the learning rate at step $k$ in the warmup phase is given by $\frac{k}{k_{\text {warmup }}} \eta$, where $k_{\text {warmup }}$ denotes the total number of warmup steps and $\eta$ denotes the peak learning rate. Moreover, after the warmup phase, the learning rate is decayed linearly to 0 , i.e. the learning rate at step $k$ after the warmup phase is given by $\frac{k-k_{\text {warmup }}}{k_{\text {max }}-k_{\text {warmup }}} \eta$, where $\eta$ denotes the peak learning rate and $k_{\text {max }}$ denotes the maximum number of gradient steps intended for finetuning.

| Dataset | Finetune batch size | Peak Learning rate | Total training epochs |
| :---: | :---: | :---: | :---: |
| CoLA | $\{16,32\}$ | $\left\{10^{-5}, 3 \times 10^{-5}, 5 \times 10^{-5}, 8 \times 10^{-5}\right\}$ | $\{3,5,10\}$ |
| SST-2 | $\{16,32\}$ | $\left\{10^{-5}, 3 \times 10^{-5}, 5 \times 10^{-5}, 8 \times 10^{-5}\right\}$ | $\{3,5,10\}$ |
| MRPC | $\{16,32\}$ | $\left\{10^{-5}, 3 \times 10^{-5}, 5 \times 10^{-5}, 8 \times 10^{-5}\right\}$ | $\{3,5,10\}$ |
| STS-B | $\{16,32\}$ | $\left\{10^{-5}, 3 \times 10^{-5}, 5 \times 10^{-5}, 8 \times 10^{-5}\right\}$ | $\{3,5,10\}$ |
| RTE | $\{16,32\}$ | $\left\{10^{-5}, 3 \times 10^{-5}, 5 \times 10^{-5}, 8 \times 10^{-5}\right\}$ | $\{3,5,10\}$ |
| QQP | $\{32\}$ | $\left\{5 \times 10^{-5}, 8 \times 10^{-5}\right\}$ | $\{3,5\}$ |
| MNLI | $\{32\}$ | $\left\{5 \times 10^{-5}, 8 \times 10^{-5}\right\}$ | $\{3,5\}$ |
| QNLI | $\{32\}$ | $\left\{5 \times 10^{-5}, 8 \times 10^{-5}\right\}$ | $\{3,5\}$ |

Table 4: Hyperparameter grid for pretrained RoBERTa $_{\text {LARGE }}$ on the downstream tasks.

| Hyperparameter | Value |
| :---: | :---: |
| Dropout | 0.1 |
| Attention Dropout | 0.1 |
| Warmup Steps | 1000 |
| Peak Learning Rate | $10^{-3}$ |
| Batch Size | 1024 |
| Weight Decay | $10^{-4}$ |
| Max Steps | 12500 |
| Learning Rate Decay | Linear |
| Adam $\beta_{1}$ | 0.9875 |
| Adam $\beta_{2}$ | 0.996 |
| Adam $\epsilon$ | $2 \times 10^{-8}$ |
| Gradient Clipping | 0.0 |
| Position embeddings | 128 |
|  |  |

Table 5: Optimization hyperparameters of baseline GPT model during pretraining.

## J. 5 WikiText-103 (GPT)

WikiText-103 (Merity et al. 2017) is a dataset with 103 million tokens extracted from Wikipedia. We split the data uniformly with a ratio $9: 1$, to create training and validation datasets for pretraining. We use Adam (Kingma and Ba, 2015) optimization algorithm.

Architecture. We pretrain a 12-layer GPT (Brown et al., 2020) model without modification. The model has 12 layers with hidden dimension 768 , feedforward network dimension 3072, 12 attention heads in each attention layer, and attention head size 64 . We pretrain on sequences of length 128 (unless stated otherwise).

Adam. We use the code from Wettig et al. (2022). To fix a baseline, we first trained our model with batch size 1024, with the optimization parameters given in Table 5. In the warmup phase, the learning rate is increased linearly over the interval, i.e. the learning rate at step $k$ in the warmup phase is given by $\frac{k}{k_{\text {warmup }}} \eta$, where $k_{\text {warmup }}$ denotes the total number of warmup steps and $\eta$ denotes the peak learning rate. Moreover, after the warmup phase, the learning rate is decayed linearly to 0 , i.e. the learning rate at step $k$ after the warmup phase is given by $\frac{k-k_{\text {warmup }}}{k_{\text {max }}-k_{\text {warmup }}} \eta$, where $\eta$ denotes the peak learning rate and $k_{\text {max }}$ denotes the maximum number of gradient steps intended for pretraining.

