

Exploration of the (Non-)Asymptotic Bias and Variance of Stochastic Gradient Langevin Dynamics

Sebastian J. Vollmer[‡], Konstantinos C. Zygalakis[‡] and Yee Whye Teh[‡]

¹Department of Statistics, University of Oxford, 24-29 St Giles' , Oxford, OX1 3LB , UK

²School of Mathematics, University of Edinburgh, Edinburgh, EH9 3JZ, UK

Editor: Lawrence Carin

Abstract

Applying standard Markov chain Monte Carlo (MCMC) algorithms to large data sets is computationally infeasible. The recently proposed stochastic gradient Langevin dynamics (SGLD) method circumvents this problem in three ways: it generates proposed moves using only a subset of the data, it skips the Metropolis-Hastings accept-reject step, and it uses sequences of decreasing step sizes. In Teh et al. (2014), we provided the mathematical foundations for the decreasing step size SGLD, including consistency and a central limit theorem. However, in practice the SGLD is run for a relatively small number of iterations, and its step size is not decreased to zero. The present article investigates the behaviour of the SGLD with fixed step size. In particular we characterise the asymptotic bias explicitly, along with its dependence on the step size and the variance of the stochastic gradient. On that basis a modified SGLD which removes the asymptotic bias due to the variance of the stochastic gradients up to first order in the step size is derived. Moreover, we are able to obtain bounds on the finite-time bias, variance and mean squared error (MSE). The theory is illustrated with a Gaussian toy model for which the bias and the MSE for the estimation of moments can be obtained explicitly. For this toy model we study the gain of the SGLD over the standard Euler method in the limit of large data sets.

Keywords: Markov Chain Monte Carlo, Langevin dynamics, big data, fixed step size.

1. Introduction

A standard approach to estimating expectations under a given target density $\pi(\theta)$ is to construct and simulate from Markov chains whose equilibrium distributions are designed to be π Brooks et al. (2011). A well-studied approach, for example in molecular dynamics Leimkuhler and Matthews (2013); Nawaf and Vanden-Eijnden (2010) and throughout Bayesian statistics Milstein and Tretyakov (2007); Neal (2011), is to use Markov chains constructed as numerical schemes which approximate the time dynamics of stochastic differential equations (SDEs). In this paper we will focus on the case of first order Langevin dynamics, which has the form

$$d\theta(t) = \frac{1}{2} \nabla \log \pi(\theta(t)) dt + dW_t, \quad (1)$$

*. vollmer@stats.ox.ac.uk

†. kzygalak@ed.ac.uk

‡. y.w.teh@stats.ox.ac.uk

where $t \in \mathbb{R}_+$, $\theta \in \mathbb{R}^d$ and W_t is a d -dimensional standard Brownian motion. Under appropriate assumptions on $\pi(\theta)$, it is possible to show that the dynamics generated by Equation (1) are ergodic with respect to $\pi(\theta)$.

The simplest possible numerical scheme for approximating Equation (1) is the Euler-Maruyama method. Let $h > 0$ be a step size. Abusing notation, the diffusion $\theta(k \cdot h)$ at time $k \cdot h$ is approximated by θ_k , which is obtained using the following recursion equation

$$\theta_{k+1} = \theta_k + \frac{h}{2} \nabla \log \pi(\theta_k) + \sqrt{h} \xi_k, \tag{2}$$

where ξ_k is a standard Gaussian random variable on \mathbb{R}^d . One can use the numerical trajectories generated by this scheme for the construction of an empirical measure $\pi_h(\theta)$ either by averaging over one single long trajectory or by averaging over many realisations in order to obtain a finite ensemble average (see for example Milstein and Tretyakov (2007)). However, as discussed in Roberts and Tweedie (1996), one needs to be careful when doing this as it could be the case that the discrete Markov chain generated by Equation (2) is not ergodic. But even if the resulting Markov Chain is ergodic, $\pi_h(\theta)$ will not be equal to $\pi(\theta)$ Mattingly et al. (2010); Abdulle et al. (2014) which thus implies that the resulting sample average is biased. An alternative strategy that avoids this discretization bias and the ergodicity of the numerical procedure, is to use Equation (2) as a proposal for a Metropolis-Hastings (MH) MCMC algorithm Brooks et al. (2011), with an additional accept-reject step which corrects the discretization error.

In this paper we are interested in situations where π arises as the posterior in a Bayesian inference problem with prior density $\pi_0(\theta)$ and a large number $N \gg 1$ of i.i.d. observations X_i with likelihoods $\pi(X_i|\theta)$. In this case, we can write

$$\pi(\theta) \propto \pi_0(\theta) \prod_{i=1}^N \pi(X_i|\theta), \tag{3}$$

and we have the following gradient,

$$\nabla \log \pi(\theta) = \nabla \log \pi_0(\theta) + \sum_{i=1}^N \nabla \log \pi(X_i|\theta). \tag{4}$$

In these situations each update (2) has an impractically high computational cost of $\mathcal{O}(N)$ since it involves computations on all N items in the data set. Likewise, each MH accept-reject step is impractically expensive.

In contrast, the recently proposed stochastic gradient Langevin dynamics (SGLD) algorithm Welling and Teh (2011) circumvents this problem by generating proposals which are only based on a subset of the data, by skipping the accept-reject step and by using a decreasing step-size sequence $(h_k)_{k \geq 0}$. In particular one has

$$\theta_{k+1} = \theta_k + \frac{h_k}{2} \widehat{\nabla \log \pi}(\theta_k) + \sqrt{h_k} \xi_k, \tag{5}$$

$$\widehat{\nabla \log \pi}(\theta_k) = \nabla \log \pi_0(\theta_k) + \frac{N}{n} \sum_{i=1}^n \nabla \log \pi(X_{\tau_{ki}}|\theta_k) \tag{6}$$

where ξ_k are independent standard Gaussian random variables on \mathbb{R}^d , and $\tau_k = (\tau_{k1}, \dots, \tau_{kn})$ is a random subset of $[N] := \{1, \dots, N\}$ of size n , generated, for example, by sampling with or without replacement from $[N]$. The idea behind this algorithm is that, since the stochastic gradient appearing in Equation (5) is an unbiased estimator of the true gradient $\nabla \log \pi(\theta)$, the additional perturbation due to the gradient stochasticity is of order h , smaller than the \sqrt{h} order of the injected noise, and so the limiting dynamics ($k \rightarrow \infty$) of Equation (5) should behave similarly to the case $n = N$. In Teh et al. (2014) it was shown that in this case that the K -step size weighted sample average is consistent and satisfies a CLT with rate depending on the decay of h_k . The optimal rate is limited to $K^{-\frac{1}{3}}$ and achieved by an asymptotic step size decay of $\asymp K^{-\frac{1}{3}}$.

The problem with decaying step sizes is that the efficiency of the algorithm slows the longer it is run for. A common practice for the SGLD and its extensions, the Stochastic Gradient Hamiltonian Monte Carlo Chen et al. (2014) and the Stochastic Gradient Thermostat Monte Carlo algorithm Ding et al. (2014), is to use step sizes that are only decreasing up to a point. The primary aim of this paper is to analyse the behaviour of SGLD with fixed step sizes of $h_k = h$. We provide two complementary analyses in this setting, one asymptotic in nature and one finite time. Let $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$ be a test function whose expectation we are interested in estimating. Using simulations of the dynamics governed by Equation (5), we can estimate the expectation using

$$\mathbb{E}_\pi[\phi(\theta)] \approx \frac{1}{K} \sum_{k=1}^K \phi(\theta_k) \quad (7)$$

for some large number of steps K . Our analyses shed light on the behaviour of this estimator.

In the first analysis, we are interested in the asymptotic bias of the estimator (7) as $K \rightarrow \infty$,

$$\lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \phi(\theta_k) - \mathbb{E}_\pi[\phi(\theta)]. \quad (8)$$

Assuming for the moment that the dynamics governed by Equation (5) is ergodic, with invariant measure $\pi_h(\theta; n)$, the above asymptotic bias simply becomes $\mathbb{E}_{\pi_h(\cdot; n)}[\phi(\theta)] - \mathbb{E}_\pi[\phi(\theta)]$. In the case of Euler-Maruyama, where $n = N$ and the gradient is computed exactly, the asymptotic behaviour of the dynamics is well understood, in particular its asymptotic bias is $\mathcal{O}(h)$ Mattingly et al. (2010). When $n < N$, using the recent generalisations Abdulle et al. (2014); Sato and Nakagawa (2014) of the approach by Talay and Tubaro (1990) reviewed in Section 3, we are able to derive an expansion of the asymptotic bias in powers of the step size h . This allows us to explicitly identify the effect, on the leading order term in the asymptotic bias, of replacing the true gradient (4) with the unbiased estimator (6). In particular, we show in Section 4 that, relative to Euler-Maruyama, the leading term contains an additional factor related to the covariance of the subsampled gradient estimators (6).

Based on this result, in Section 4.2, we propose a modification of the SGLD (referred to simply as mSGLD) which has the same asymptotic bias as the Euler-Maruyama method up to first order in h . The mSGLD is given by

$$\theta_{k+1} = \theta_k + \frac{h}{2} \widehat{\nabla \log \pi}(\theta_k) + \sqrt{h} \left(I - \frac{h}{2} \text{Cov} \left[\widehat{\nabla \log \pi}(\theta_k) \right] \right) \xi_j \quad (9)$$

where $\text{Cov} \left[\widehat{\nabla \log \pi}(\theta_k) \right]$ is the covariance of the gradient estimator. When the covariance is unknown, it can in turn be estimated by subsampling as well. This modification is different from the Stochastic Gradient Fisher Scoring S. Ahn and Welling (2012), a modification of the injected noise in order to better match the Bernstein von Mises posterior. In contrast, the mSGLD is a local modification based on the estimated variance of the stochastic gradient.

The second contribution provides a complementary finite time analysis. In the finite time case both the bias and the variance of the estimator are non-negligible, and our analysis accounts for both by focussing on bounding the mean squared error (MSE) of the estimator (7). Our results, presented in Section 5, show that,

$$\mathbb{E} \left[\left(\frac{1}{K} \sum_{k=0}^{K-1} \phi(\theta_k) - \pi(\phi) \right)^2 \right] \leq C(n) \left(h^2 + \frac{1}{Kh} \right), \quad (10)$$

where the RHS only depends on n through the constant $C(n)$, the h^2 term is a contribution of the (square of the) bias while the $1/Kh$ term is a contribution of the variance. We see that there is a bias-variance trade-off, with bias increasing and variance decreasing monotonically with h . Intuitively, with larger h the Markov chain can converge faster (lower variance) with the same number of steps, but this incurs higher discretization error. Our result is achieved by extending the work of Mattingly et al. (2010) from \mathbb{T}^d to \mathbb{R}^d . The main difficulty in achieving this relates to combining the results of Pardoux and Veretennikov (2001) and Teh et al. (2014), in order to establish the existence of nice, well controlled solutions to the corresponding Poisson equation Mattingly et al. (2010). We can minimise Equation (10) over h , finding that the minimizing h is on the order of $K^{-\frac{1}{3}}$, and yields an MSE of order $K^{-\frac{2}{3}}$. This agrees, surprisingly, with the scaling of $K^{-\frac{1}{3}}$ for the central limit theorem established for the case of decreasing step sizes, for the Euler-Maruyama scheme in Lamberton and Pages (2002) and for SGLD in Teh et al. (2014). This unexpected result, that the decreasing step size and fixed step size discretisations have, up to a constant, the same efficiency seems to be missing from the literature.

Our theoretical findings are confirmed by numerical simulations. More precisely, we start by studying a one dimensional Gaussian toy model both in terms of the asymptotic bias and the MSE of time averages in Section 2. The simplicity of this model allows us to obtain explicit expressions for these quantities and thus illustrate in a clear way the connection with the theory. More precisely, we confirm that the scaling of the step size and the number of steps for a prescribed MSE obtained from the upper bound in Equation (10) matches the scaling derived from the analytic expression for the MSE for this toy model. More importantly, this simplicity allows us to make significant analytic progress in the study of the asymptotic bias and MSE of time averages in the limit of large data sets $N \rightarrow \infty$. In particular, we are able to show that the SGLD reduces the computational complexity by one order of magnitude in N . for the estimation of the second moment in comparison with the Euler method if the error is quantified through the MSE.

In summary, this paper is organised as follows. We present our first explorations of the SGLD applied to a one-dimensional Gaussian toy model in Section 2. For this model we obtain an analytic characterisation of its bias and variance. This serves as intuition and benchmark for the two analyses developed in Sections 3 to 5. In Section 3 we review

some known results about the effect of the numerical discretisation of Equation (1) in terms of the finite time weak error as well as in terms of the invariant measure approximation. In Section 4 we apply these results to analyse the finite and long time properties of the SGLD, as well as to construct the modified SGLD algorithm which, asymptotically in h , behaves exactly as the Euler-Maruyama method ($n = N$) while still sub-sampling the data set at each step. Furthermore, in Section 5 we discuss the properties of the finite time sample averages, include its MSE. In Section 6 we revisit the Gaussian toy model to obtain a more precise understanding of the behaviour of SGLD in a large data and high accuracy regime. This is achieved using analytic expressions of the expectations of the sample average which are obtained using the Mathematica[®] notebook that is available upon request from the first author described in detail in Appendix 10. Finally, in Section 7 we demonstrate the observed performance of SGLD for a Bayesian logistic regression model which matches the theory, while, we conclude this paper in Section 8 with a discussion on some possible extensions of this work.

2. Exploring a one-dimensional Gaussian Toy Model

In this section, we develop results for a simple toy model, which will serve as a benchmark for the theory developed in Sections 3 to 5. In particular, we obtain analytic expressions for the bias and the variance of the sample average of the SGLD, allowing us to characterise its performance in detail.

We consider a one-dimensional linear Gaussian model,

$$\begin{aligned} \theta &\sim \mathcal{N}(0, \sigma_\theta^2), \\ X_i | \theta &\stackrel{i.i.d.}{\sim} \mathcal{N}(\theta, \sigma_x^2) \quad \text{for } i = 1, \dots, N. \end{aligned} \tag{11}$$

The posterior is given by

$$\pi = \mathcal{N}(\mu_p, \sigma_p^2) = \mathcal{N}\left(\frac{\sum_{i=1}^N X_i}{\frac{\sigma_x^2}{\sigma_\theta^2} + N}, \left(\frac{1}{\sigma_\theta^2} + \frac{N}{\sigma_x^2}\right)^{-1}\right). \tag{12}$$

For this choice of π , the Langevin diffusion (1) becomes,

$$d\theta(t) = -\frac{1}{2} \left(\frac{\theta(t) - \mu_p}{\sigma_p^2} \right) dt + dW_t, \tag{13}$$

and its numerical discretisation by the SGLD with step size h reads as follows,

$$\theta_{k+1} = (1 - Ah)\theta_k + B_k h + \sqrt{h}\xi_k, \tag{14}$$

where $\xi_k \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1)$ and

$$\begin{aligned} A &= \frac{1}{2} \left(\frac{1}{\sigma_\theta^2} + \frac{N}{\sigma_x^2} \right), \\ B_k &= \frac{N}{n} \frac{\sum_{i=1}^n X_{\tau_{ki}}}{2\sigma_x^2}, \end{aligned} \tag{15}$$

where $\tau_k = (\tau_{k1}, \dots, \tau_{kn})$ denote a random subset of $[N] = \{1, \dots, N\}$ generated, for example, by sampling with or without replacement from $[N]$, independently for each k . We note that the updates (14) will be stable only if $0 \leq 1 - Ah < 1$, that is, $0 < h < 1/A$ ¹. In the following we will also consider parameterising the step size as $h = r/A$ where $0 < r < 1$.

We denote $B = (B_k)_{k \geq 0}$. At the risk of obfuscating the notation, we will denote by $\text{Var}(B)$ the common variance of B_k for all k . For sampling with replacement, we have

$$\text{Var}(B) = \frac{1}{4\sigma_x^4} \frac{N}{n} \sum_{j=1}^N \left(X_j - \frac{1}{N} \sum_{i=1}^N X_i \right)^2 = \frac{1}{4\sigma_x^4} \frac{N(N-1)}{n} \text{Var}(X),$$

where $\text{Var}(X)$ is the typical unbiased empirical estimate of the variance of $\{X_1, \dots, X_N\}$. For sampling without replacement we have,

$$\text{Var}(B) = \frac{1}{4\sigma_x^4} \frac{N(N-n)}{n(N-1)} \sum_{i=1}^N \left(X_i - \frac{1}{N} \sum_{i=1}^N X_i \right)^2 = \frac{1}{4\sigma_x^4} \frac{N(N-n)}{n} \text{Var}(X). \quad (16)$$

2.1 Analysis of the Asymptotic Bias

We start by inspecting the estimate of the posterior mean. In particular, using Equation (14) and taking expectations with respect to ξ_k , we have

$$\mathbb{E}(\theta_{k+1}|B) = (1 - Ah)\mathbb{E}(\theta_k|B) + B_k h \quad (17)$$

which can be solved in order to obtain

$$\mathbb{E}(\theta_M|B) = (1 - Ah)^M \mathbb{E}(\theta_0) + \sum_{k=0}^{M-1} h(1 - Ah)^k B_{M-k-1}.$$

If we now take the expectation with respect to the random subsets B_k , using the fact that $\mathbb{E}(B_k) = \mathbb{E}(B)$ and take the limit of $M \rightarrow \infty$, we have

$$\mathbb{E}(\theta_\infty) = \sum_{k=0}^{\infty} (1 - Ah)^k h \mathbb{E}(B) = h/(1 - (1 - Ah))\mathbb{E}(B) = \frac{\mathbb{E}(B)}{A} = \frac{\sum_{i=1}^N X_i}{\frac{\sigma_x^2}{\sigma_\theta^2} + N}.$$

We thus see that the SGLD is capturing the correct limiting mean of the posterior independently of the choice of the step size h . In other words, for the test function $\phi(\theta) = \theta$, the asymptotic bias is nil.

We now investigate the behaviour of the limiting variance under the SGLD. Starting with the law of total variance,

$$\text{Var}[\theta_{k+1}] = \mathbb{E}(\text{Var}[\theta_{k+1} | B]) + \text{Var}(\mathbb{E}[\theta_{k+1} | B]),$$

a simple calculation now shows that

$$\text{Var}[\theta_{k+1} | B] = (1 - Ah)^2 \text{Var}[\theta_k | B] + h$$

1. Note that the posterior variance is $\asymp 1/A$, so that steps of size $\asymp 1/A$ are $1/\sqrt{A}$ smaller than the width of the posterior. However the injected noise has variance $\asymp 1/A$ which matches the posterior variance.

and

$$\text{Var}(\mathbb{E}[\theta_{k+1} | B]) = (1 - Ah)^2 \text{Var}(\mathbb{E}[\theta_k | B]) + h^2 \text{Var}(B_k).$$

Combining these two results, we see that

$$\text{Var}(\theta_{k+1}) = (1 - Ah)^2 \text{Var}(\theta_k) + h + h^2 \text{Var}(B_k).$$

If we now take the limit of $k \rightarrow \infty$, we have that

$$\text{Var}(\theta_\infty) = \frac{1}{2A - A^2h} + \frac{h \text{Var}(B)}{2A - A^2h}. \quad (18)$$

where $\text{Var}(B)$ is the common value of $\text{Var}(B_k)$ for all $k \geq 0$. We note here that in the case of the Euler-Maruyama method (from here on we will simply refer to this as the Euler method) where $n = N$ and $\text{Var}(B) = 0$, only the first term remains. In other words, the first term is an (over-)estimate of the posterior variance $\sigma_p^2 = 1/2A$ obtained by the Euler-Maruyama discretisation at step size h . Our result here coincides with Zygalakis (2011). On the other hand, the second term is an additional bias term due to the variability of the stochastic gradients. Further, using a Taylor expansion in h of the second summand, we see that the SGLD has an excess bias, relative to the Euler method, with first order term equal to

$$h \frac{\text{Var}(B)}{2A}. \quad (19)$$

Using the fact that $\text{Var}(\theta_\infty) = \mathbb{E}[\theta_\infty^2] - \mathbb{E}[\theta_\infty]^2$, and that the asymptotic bias of estimating $\mathbb{E}[\theta]$ is nil in this simple model, we see that the above gives the asymptotic biases of the Euler method and SGLD in the case of the test function $\phi(\theta) = \theta^2$.

We now consider the modified SGLD given in Equation (9) and to be discussed in Section 4.2. In this case the numerical discretisation of Equation (13) becomes

$$\theta_{k+1} = \theta_k - Ah\theta_k + B_k h + \sqrt{h} \left(1 - \frac{h}{2} \text{Var}(B) \right) \xi_k \quad (20)$$

A similar calculation as for the SGLD shows that

$$\begin{aligned} \mathbb{E}(\theta_\infty) &= \frac{\sum_{i=1}^N X_i}{\frac{\sigma_x^2}{\sigma_\theta^2} + N} \\ \text{Var}(\theta_\infty) &= \frac{1}{2A - A^2h} + \frac{h^2 \text{Var}^2(B)}{4(2A - A^2h)}. \end{aligned} \quad (21)$$

with the last term being the excess asymptotic bias. A Taylor expansion of the excess bias term shows that the term of order h vanishes and the leading term has order h^2 . Hence, for small h , the excess bias is negligible compared to the asymptotic bias of the Euler method, and we can say that, up to first order and in this simple example, the mSGLD has the same asymptotic bias as for the Euler method. In Section 4.2, we will show that these results hold more generally.

It is useful to visualise the above analytic results for the asymptotic biases of the Euler method, SGLD and mSGLD. In Figure 1 we show this for a data set of 1000 points drawn

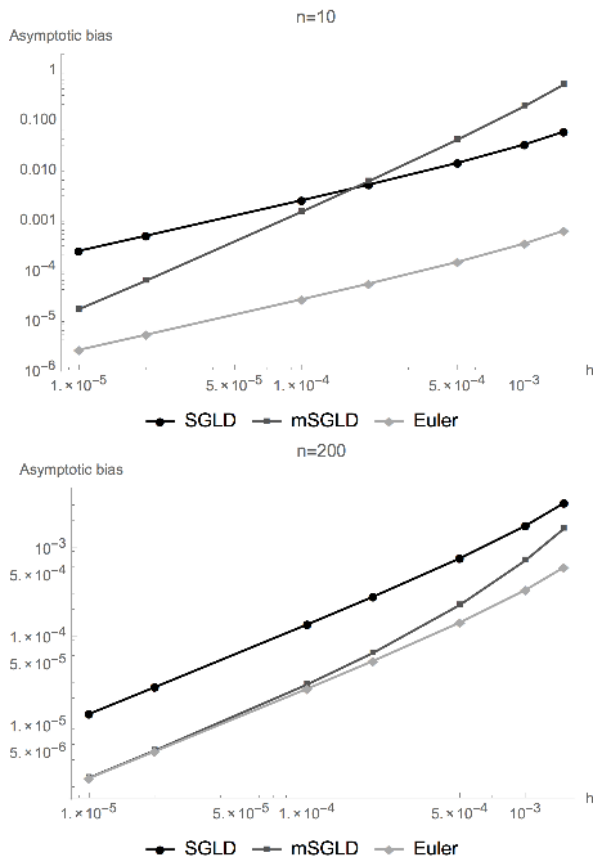


Figure 1: Comparison of the asymptotic biases for the SGLD, the mSGLD and the Euler method for the test function $\phi(\theta) = \theta^2$. For all simulations, we have used $N = 10^3$. We used $n = 10$ on the LHS and $n = 200$ on the RHS.

according to the model. In particular, after choosing the data set, then we calculate analytically $\text{Var}(B)$ for each choice of n using (16), and then use (18), (21) to evaluate the asymptotic bias for each of the methods. The first observation is that the Euler method has lowest asymptotic bias among all three methods (although of course it is also the most computationally expensive; see Section 6). We observe that if we choose $n = 10$ points for each gradient evaluation, for large values of the step size h , the SGLD is superior to the mSGLD. However, as h is reduced, this is no longer the case. Furthermore, if we use a more accurate gradient estimation with $n = 200$ data points, we see that mSGLD outperforms SGLD for all the step sizes used, but more importantly its asymptotic bias is now directly comparable with the Euler method where all the data points are used for evaluating the gradient.

2.2 Finite Time Analysis

In the previous subsection we analysed the behaviours of the three algorithms in terms of their biases in the asymptotic regime. In practice, we can only run our algorithms for a finite number of steps, say K , and it would be interesting to understand the behaviours of the algorithms in this scenario. With a finite number of samples, in addition to bias we also have to account for variance due to the Monte Carlo estimation.

A sensible analysis accounting for both bias and variance is to study the behaviour of the mean squared error (MSE), say in the Monte Carlo estimation for the second moment,

$$\text{MSE}_2 := \mathbb{E} \left(\frac{1}{K} \sum_{k=0}^{K-1} \theta_k^2 - (\mu_p^2 + \sigma_p^2) \right)^2. \quad (22)$$

We can expand the quadratic, and express MSE_2 as a linear combination of terms of the form $\mathbb{E}[\theta_j^p]$ for $p = 1, 2, 3, 4$. Each of terms can be calculated analytically, depending on the data set X , the total number of steps K , the subset size n , as well as the scaled step size parameter $r = hA$. We provide these calculations in Appendix 10.1 and a Mathematica[®] file in the supplementary materials.

In Figure 2 we visualise the behaviour of the resulting MSE_2 for a fixed data set with $N = 1000$ items, and with scaled step size $r = 1/20$. For the same number of steps M , the left figure shows that SGLD and mSGLD behaves similarly, decreasing initially then asymptoting at their asymptotic biases studied in the previous subsection. At $r = 1/20$ mSGLD has lower asymptotic biases than SGLD. Further, both MSE_2 's decrease with increasing subset size n , and are higher than that for the Euler method at $n = 1000$. Since SGLD and mSGLD computational costs per step are linear in n , the right figure instead plots the same MSE_2 's against the (effective) number of passes through the data set, that is, number of steps times n/N . This quantity is now proportional to the computational budget. Now we see that smaller subset sizes produce initial gains, but asymptote at higher biases.

These analytical results for a simple Gaussian model demonstrate the more general theory which forms the core contributions of this paper. Sections 3 and 4 develop a method to study the asymptotic bias as a Taylor expansion in h , while Section 5 provides a finite time analysis in terms of the mean squared error. Both analyses are based on the behaviour of the algorithms for small step sizes, and in this regime we see that mSGLD has better performance than SGLD. In Section 6 we will return to the simple Gaussian model to study the behaviour of the algorithms using different measures of performance and in different regimes. In particular, we will see that for larger step sizes SGLD has better performance than mSGLD.

3. Review of Weak Order Results

In this section we review some existing results regarding the ergodicity and accuracy of numerical approximations of SDEs. We start in Section 3.1 by introducing the framework and notation, the Fokker-Planck and backward Kolmogorov equations, and with some preliminary results on local weak errors of numerical one-step integrators. Section 3.2 presents assumptions necessary for ergodicity, and extends the results to a global error expansion of

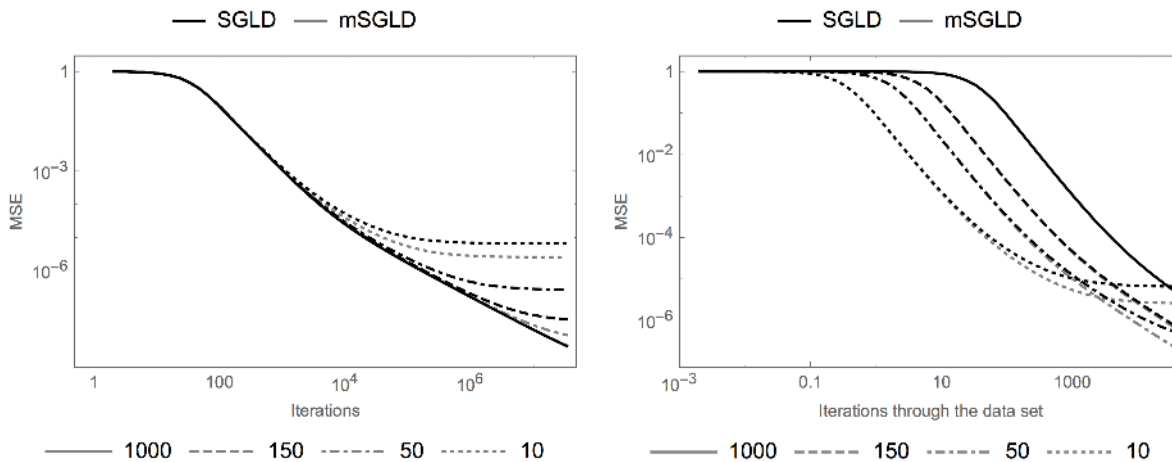


Figure 2: MSE_2 of the sample average for the SGLD and the mSGLD for the second moment of the posterior.

the weak error as well as the error in the approximation of the invariant measure. Finally, in Section 3.3 we apply our results to explicitly calculate the leading order error term of the numerical approximation of an Ornstein-Uhlenbeck solved by the Euler method.

3.1 One-step Numerical Approximations of Langevin Diffusions

Let us denote by $\rho(y, t)$ the probability density of $\theta(t)$ defined by the Langevin diffusion (1) with initial condition $\theta(0) = \theta$ and target density $\pi(y)$. Then $\rho(y, t)$ is the solution of the Fokker-Planck equation,

$$\frac{\partial \rho}{\partial t} = \mathcal{L}^* \rho, \tag{23}$$

with initial condition $\rho(y, 0) = \delta(y - \theta)$, a Dirac mass for the deterministic initial condition, and the operator \mathcal{L}^* given by

$$\mathcal{L}^* \rho = -\frac{1}{2} \nabla_{\theta} \cdot (\nabla \log \pi(\theta) \rho) + \frac{1}{2} \nabla_{\theta} \cdot \nabla_{\theta} \cdot \rho. \tag{24}$$

This operator is the L^2 -adjoint of the generator of the Markov process $(\theta(t))_{t \geq 0}$ given by (1),

$$\mathcal{L} = \frac{1}{2} \nabla_{\theta} \log \pi(\theta) \cdot \nabla_{\theta} + \frac{1}{2} \Delta_{\theta}, \tag{25}$$

Given a test function ϕ , define $u(\theta, t)$ to be the expectation,

$$u(\theta, t) = \mathbb{E}(\phi(\theta(t)) | \theta(0) = \theta), \tag{26}$$

with respect to the diffusion at time t when started with initial condition $\theta(0) = \theta$. We note that $u(\theta, t)$ is the solution of the backward Kolmogorov equation

$$\begin{aligned} \frac{\partial u}{\partial t} &= \mathcal{L}u, \\ u(\theta, 0) &= \phi(\theta). \end{aligned} \tag{27}$$

A formal Taylor series expansion for u in terms of the generator \mathcal{L} was derived in Zygalkakis (2011) and made rigorous by Debussche and Faou (2012) for the case where the state space is $\theta \in \mathbb{T}^d$. The Taylor series is of the following form,

$$u(\theta, h) = \phi(\theta) + \sum_{j=1}^l \frac{h^j}{j!} \mathcal{L}^j \phi(\theta) + h^{l+1} r_l(\theta), \tag{28}$$

for all positive integers l , with the remainder satisfying a bound of the form $|r_l(\theta)| \leq c_l(1 + |\theta|^{\kappa_l})$ for some constants c_l, κ_l depending on π and ϕ .

Remark 1 *Another way to turn $u(\theta, h) = \phi(\theta) + h\mathcal{L}\phi + \frac{h^2}{2}\mathcal{L}^2\phi + \dots$ into a rigorous expansion, see Equation (28), is to follow the approach in (Talay and Tubaro, 1990, Lemma 2) and to assume that $\log \pi$ is C^∞ with bounded derivatives of any order (and this is the approach we follow here). This fact, together with the assumption that*

$$|\phi(\theta)| \leq C(1 + |\theta|^s) \tag{29}$$

for some positive integer s is enough to prove that the solution u of Equation (27) has derivatives of any order that have a polynomial growth of the form of Equation (29), with other constants C, s that are independent of $t \in [0, T]$. In turn, these regularity bounds establish that Equation (28) holds. We mention here that the regularity conditions were relaxed in recent work in Kopec (2014) for the elliptic case and in Kopec (2015) for the hypoelliptic case.

Now assume that one solves Equation (1) numerically with a one step integrator, which we shall denote by,

$$\theta_{n+1} = \Psi(\theta_n, h, \xi_n), \tag{30}$$

where $\theta_0 = \theta(0)$, h denotes the step size, ξ_n are iid $\mathcal{N}(0, 1)$, and θ_n denotes the numerical approximation of $\theta(nh)$ for each $n \in \mathbb{N}$. For example in the case of the Euler method for equation (1) one has

$$\Psi(\theta, h, \xi) = \theta + \frac{h}{2} \nabla \log \pi(\theta) + \sqrt{h} \xi$$

Now, using this formulation we can define

$$U(\theta, h) = \mathbb{E}(\phi(\theta_1) | \theta_0 = \theta), \tag{31}$$

for the expectation of the test function after one step of the numerical integrator starting with the initial condition $\theta_0 = \theta$. We will make the following (easily satisfied) regularity and consistency assumptions about the integrator:

Assumption 2 *We assume that the following hold:*

- $\nabla \log \pi$ is C^∞ with bounded derivatives of all orders.
- For all deterministic initial conditions θ_0 , we have

$$|\mathbb{E}(\theta_1 - \theta_0)| \leq C(1 + |\theta_0|)h, \text{ and } |\theta_1 - \theta_0| \leq M(1 + |\theta_0|)\sqrt{h}, \quad (32)$$

where C is a constant independent of h , for h small enough and M is a random variable that has bounded moments of all orders independent of h and θ_0 .

- Equation (31) has a weak Taylor series expansion of the form

$$U(\theta, h) = \phi(\theta) + hA_0(\pi)\phi(\theta) + h^2A_1(\pi)\phi(\theta) + \dots, \quad (33)$$

where $A_i(\pi)$, $i = 0, 1, 2, \dots$ are linear differential operators with coefficients depending smoothly on the drift function $\nabla \log \pi(\theta)$ and its derivatives (depending on the choice of the integrator).

- $A_0(\pi)$ coincides with the generator \mathcal{L} , in other words, the numerical method has weak order at least one.

Remark 3 *Equation (33) holds for almost any Taylor based method applied to (1) but also to general SDEs with multiplicative noise. This is discussed further in Abdulle et al. (2012), which also contains examples numerical methods, other than the Euler-Maryuama method, for which (33) holds.*

Assumptions 2 immediately imply the existence of a rigorous expansion

$$U(\theta, h) = \phi(\theta) + \sum_{i=0}^l h^{i+1} A_i(\pi)\phi(\theta) + h^{l+2} R_l(\theta) \quad (34)$$

for all positive integers l , with a remainder satisfying $|R_l(\theta)| \leq C_l(1 + |\theta|^{K_l})$ for some constants C_l, K_l . We say that the numerical solution has local weak order p if the first p terms in the expansion (33) of the numerical approximation agrees with that (28) for the exact diffusion. In this case, it is easy to see that the following local error formula holds,

$$\mathbb{E}(\phi(\theta(h))|\theta(0) = \theta) - \mathbb{E}(\phi(\theta_1)|\theta_0 = \theta) = h^{p+1} \left(\frac{\mathcal{L}^{p+1}}{(p+1)!} - A_p \right) \phi(\theta) + \mathcal{O}(h^{p+2}). \quad (35)$$

3.2 Global Weak Error Expansion

In this subsection, we will extend the local weak error expansion to a global one. Specifically, after M steps of the numerical integrator with step size h , we are interested in the difference between θ_M and the exact diffusion $\theta(T)$ where $T = Mh$, as evaluated by the difference between the corresponding expectations of ϕ ,

$$E(\phi, h, T) = \mathbb{E}(\phi(\theta(T))|\theta(0) = \theta) - \mathbb{E}(\phi(\theta_M)|\theta_0 = \theta), \quad (36)$$

In order for this study to make sense (when considering the limit $T \rightarrow \infty$), we will require that the SDE and its numerical approximation are both ergodic. We make standard assumptions in order for the Langevin diffusion $(\theta(t))_{t \geq 0}$ as given by (1) to be ergodic (see Hasminskii (1980)):

Assumption 4 *We assume that the following hold for the Langevin diffusion $(\theta(t))_{t \geq 0}$:*

- $\nabla \log \pi$ is C^∞ with bounded derivatives of all orders.
- there exists $\beta > 0$ and a compact set $K \subset \mathbb{R}^d$ such that $\forall \theta \in \mathbb{R}^d \setminus K$,

$$\langle \theta, \nabla \log \pi(\theta) \rangle \leq -\beta \|\theta\|_2^2.$$

The question of the ergodicity of the numerical approximation (θ_n) is considerably more intricate in general. There exist cases where the underlying Langevin diffusion is ergodic, but its numerical approximation is not ergodic, or does not converge exponentially fast Roberts and Tweedie (1996). This relates mainly to the properties of the drift coefficient and its behaviour at infinity. For the Euler-Maruyama and the Milstein scheme this has been investigated in Talay and Tubaro (1990). In what follows we will simply assume that the Markov chain (θ_n) defined by the numerical approximation is indeed ergodic. Under this assumption the following theorem, which combines results derived by Talay and Tubaro (1990) and Milstein (1986), can be shown (see Abdulle et al. (2014) for a proof):

Theorem 5 *Suppose that the state space is \mathbb{R}^d , that Assumptions 2 and 4 hold, and that the Markov chain $(\theta_n)_{n \geq 0}$ defined by the one step integrator (30) is ergodic. If the numerical approximation has local weak order p , that is, Equation (35) holds, then we have the following expansion of the global error (36), for all $\phi \in C_P^{2p+4}(\mathbb{R}^d, \mathbb{R})$,*

$$E(\phi, h, T) = h^p \int_0^T \mathbb{E}(e(\theta(s), s)) ds + \mathcal{O}(h^{p+1}), \quad (37)$$

where $e(\theta, t)$ is given by

$$e(\theta, t) = \left(\frac{1}{(p+1)!} \mathcal{L}^{p+1} - A_p \right) v(\theta, t), \quad (38)$$

with $v(\theta, t) = \mathbb{E}(\phi(\theta(T)) | \theta(t) = \theta)$ satisfying

$$\begin{aligned} \frac{\partial v}{\partial t} &= -\mathcal{L}v, \\ v(\theta, T) &= \phi(\theta). \end{aligned} \quad (39)$$

The expression (37) was proved by Talay and Tubaro (1990) for specific methods (e.g. the Euler-Maruyama or the Milstein methods), while the general procedure to infer the global weak order from the local weak order is due to Milstein (1986) (see also (Milstein and Tretyakov, 2004, Chapter 2.2)). However, the formulation of the error function (38) here is in terms of the generator \mathcal{L} and the operators A_i in Assumption 2, and does not contain any time derivatives as in Talay and Tubaro (1990); Milstein (1986). This formulation will be particularly useful for obtaining our main results.

Using Theorem 5, one can obtain a similar expansion to that in Equation (37) for the difference between the true and the numerical ergodic averages:

Theorem 6 *Suppose that Assumption 4 holds, that our numerical method with deterministic initial condition is ergodic and of weak order p , and that $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$ is a smooth function satisfying Equation (29). Then,*

$$\lim_{K \rightarrow \infty} \frac{1}{K} \sum_{n=0}^{K-1} \phi(\theta_n) - \int_{\mathbb{R}^d} \phi(y) \pi(y) dy = -\lambda_p h^p + \mathcal{O}(h^{p+1}) \quad (40)$$

where λ_p is defined as

$$\lambda_p = \int_{\mathbb{R}^d} \int_0^\infty \left(\frac{1}{(p+1)!} \mathcal{L}^{p+1} - A_p \right) u(y, t) \pi(y) dt dy, \quad (41)$$

and $u(y, t)$ satisfies Equation (27).

The proof is given in Abdulle et al. (2014), and is similar to that in (Talay and Tubaro, 1990, Theorem 4) with the main difference being that Equation (37) is used as the starting point, instead of the specific formula for the Euler-Maruyama method used in Talay and Tubaro (1990).

Theorem 6 provides an explicit expression for the leading order term of the asymptotic bias of the numerical method. It will thus be the key result in our analysis of the asymptotic behaviour of SGLD later. Intuitively Equation (41) says that if we want to calculate the error between the numerical and the true ergodic averages, we need to take into account the long time ($t \rightarrow \infty$) discrepancy between the true and the numerical solution given by

$$\left(\frac{1}{(p+1)!} \mathcal{L}^{p+1} - A_p \right) u(y, t),$$

and then average over all possible initial conditions y with respect to invariant measure $\pi(y)$.

3.3 An Illustrative Example

We illustrate the weak order results above in the case of the Euler-Maruyama scheme applied to the Ornstein-Uhlenbeck process. For ease of notation, let $f(x) = \frac{1}{2} \nabla \log \pi(\theta)$. The Euler-Maruyama update steps are,

$$\theta_{n+1} = \theta_n + hf(\theta_n) + \sqrt{h} \xi_n. \quad (42)$$

A straightforward calculation Zygalakis (2011) yields that the differential operator A_1 in (33) is given by

$$A_1 \psi = \frac{1}{2} f^T \nabla^2 \psi f + \frac{1}{2} \sum_{i=1}^d \psi'''(e_i, e_i, f) + \frac{1}{8} \sum_{i,j=1}^d \psi^{(4)}(e_i, e_i, e_j, e_j) \quad (43)$$

where e_1, \dots, e_d denotes the canonical basis of \mathbb{R}^d and $\psi'''(\cdot, \cdot, \cdot)$ and $\psi^{(4)}(\cdot, \cdot, \cdot, \cdot)$, are the derivatives of ψ , which are trilinear and quadrilinear forms, respectively. In dimension $d = 1$, it reduces to

$$A_1 \psi = \frac{1}{2} f^2 \psi'' + \frac{1}{2} \psi''' + \frac{1}{8} \psi^{(4)}$$

In the case where $\theta \in \mathbb{R}$ and $\pi(\theta) = e^{-(\theta-\mu)^2/2\sigma^2}$, the Langevin diffusion (1) corresponds to the one dimensional Ornstein-Uhlenbeck process,

$$d\theta(t) = -\frac{1}{2} \left(\frac{\theta(t) - \mu}{\sigma^2} \right) dt + dW_t \quad (44)$$

For the test function $\phi(\theta) = \theta^2$, a simple calculation reveals that the solution of Equation (27) is

$$u(\theta, t) = \sigma^2 \left(1 - e^{-t/\sigma^2} \right) + \theta^2 e^{-t/\sigma^2} + \frac{\mu\theta}{\sigma^2} (1 - e^{-t/2\sigma^2}) e^{-t/2\sigma^2} + \frac{\mu^2}{4\sigma^4} (1 - e^{-t/2\sigma^2})^2. \quad (45)$$

The Euler-Maruyama scheme has weak order $p = 1$, and (see Zygalkakis (2011) for more details),

$$\frac{1}{2} \mathcal{L}^2 - A_1 = \frac{1}{8\sigma^4} (\theta - \mu) \frac{d}{d\theta} - \frac{1}{4\sigma^2} \frac{d^2}{d^2\theta}.$$

Using this together with Equation (45) for $u(\theta, t)$, we find

$$\left(\frac{1}{2} \mathcal{L}^2 - A_1 \right) u(\theta, t) = \frac{e^{-t/\sigma^2}}{2\sigma^2} - \frac{e^{-t/2\sigma^2} (\mu - \theta) \left((1 - e^{-t/2\sigma^2}) \mu + e^{-t/2\sigma^2} \theta \right)}{4\sigma^4}$$

Formula (41) now gives

$$\begin{aligned} \lambda_1 &= - \int_0^\infty \int_{-\infty}^{+\infty} \left(\frac{e^{-t/\sigma^2}}{2\sigma^2} - \frac{e^{-t/2\sigma^2} (\mu - \theta) \left((1 - e^{-t/2\sigma^2}) \mu + e^{-t/2\sigma^2} \theta \right)}{4\sigma^4} \right) \frac{e^{-(\theta-\mu)^2/2\sigma^2}}{\sqrt{2\pi}\sigma} d\theta dt \\ &= \frac{1}{4}. \end{aligned} \quad (46)$$

This is in agreement with known results on the stationary distribution of the Euler-Maruyama approximation to the Ornstein-Uhlenbeck process, see Zygalkakis (2011):

$$\pi_h \sim N(\mu, \sigma_h^2) \quad \text{where} \quad \sigma_h^2 = \frac{\sigma^2}{1 - \frac{h}{4}\sigma^{-2}} = \sigma^2 + \frac{1}{4}h + \mathcal{O}(h^2).$$

4. Weak Convergence Analysis

We study the weak convergence properties of the SGLD method in the light of Theorems 5 and 6. The analysis in Section 4.1 implies that at leading order there is a cost associated with not calculating the likelihood over all points. Thus, we introduce in Section 4.2 a modification of the original algorithm which has an error that is, asymptotically in h , identical to the error of the Euler method, when all data points are taken into account in the calculation of each likelihood gradient.

4.1 Stochastic Gradient Langevin Dynamics

Theorems 5 and 6 imply that in order to characterise the leading order error term both for the weak convergence and the invariant measure, we need to calculate the corresponding

differential operators A_0, A_1, \dots in Equation (33). To simplify the presentation and to illustrate the main ideas, we present the calculations only in the case where $\theta(t)$ is one dimensional. We start our calculations by rewriting the SGLD method in the following form

$$\theta_{j+1} = \theta_j + h\hat{f}_j(\theta_j) + \sqrt{h}\xi_j, \quad (47)$$

where

$$\hat{f}_j(\theta) = \frac{1}{2} \left(\nabla \log \pi_0(\theta) + \frac{N}{n} \sum_{i=1}^n \nabla \log \pi(X_{\tau_{ji}}|\theta) \right),$$

τ_j is the subset (possibly with repetition) chosen at step j and,

$$\mathbb{E}_{\tau_j} \hat{f}_j(\theta) = f(\theta) := \frac{1}{2} \nabla \log \pi(\theta), \quad \forall n \leq N. \quad (48)$$

Expanding $\phi(\theta_{j+1})$ in powers of h and then taking expectations with respect to the injected random noise ξ_j ,

$$\begin{aligned} \mathbb{E}_{\xi_j}(\phi(\theta_{j+1})|\theta_j) &= \phi(\theta_j) + h \left(\hat{f}_j(\theta_j) \phi'(\theta_j) + \frac{1}{2} \phi''(\theta_j) \right) \\ &\quad + \frac{h^2}{2} \left(\hat{f}_j^2(\theta_j) \phi''(\theta_j) + \hat{f}_j(\theta_j) \phi^{(3)}(\theta_j) + \frac{1}{4} \phi^{(4)}(\theta_j) \right) + \mathcal{O}(h^3). \end{aligned}$$

If we now take expectations with respect to τ_j ,

$$\mathbb{E}(\phi(\theta_{j+1})|\theta_j) = \phi(\theta_j) + h\mathcal{L}\phi(\theta_j) + \frac{h^2}{2} \left(\mathbb{E}_{\tau_j}(\hat{f}_j^2(\theta_j)) \phi''(\theta_j) + f(\theta_j) \phi^{(3)}(\theta_j) + \frac{1}{4} \phi^{(4)}(\theta_j) \right) + \mathcal{O}(h^3), \quad (49)$$

where \mathcal{L} is the generator (25) of Equation (1). We thus see that the SGLD method is a first order weak method and, dropping the indexing by j for notational convenience from now on,

$$A_1(\pi)\phi = \frac{1}{2} \left(\mathbb{E}_{\tau}(\hat{f}^2(\theta)) \phi'' + f(\theta) \phi^{(3)} + \frac{1}{4} \phi^{(4)} \right).$$

The asymptotic bias in Equation (41) has an expansion based on the differential operator,

$$\begin{aligned} \frac{1}{2} \mathcal{L}^2 - A_1 &= \frac{1}{2} \left(f(\theta) f'(\theta) + \frac{1}{2} f''(\theta) \right) \frac{d}{d\theta} + \frac{1}{2} \left(f'(\theta) + f^2(\theta) - \mathbb{E}_{\tau}(\hat{f}^2(\theta)) \right) \frac{d^2}{d\theta^2} \\ &= \frac{1}{2} \left(f(\theta) f'(\theta) + \frac{1}{2} f''(\theta) \right) \frac{d}{d\theta} + \frac{1}{2} \left(f'(\theta) - \text{Var}(\hat{f}(\theta)) \right) \frac{d^2}{d\theta^2} \end{aligned} \quad (50)$$

We thus see that in the case of SGLD the leading order error term contains an extra factor of $-\frac{1}{2} \text{Var}(\hat{f}(\theta)) \nabla_{\theta}^2$ when compared to the Euler method ($n = N$), in which case $\text{Var}(\hat{f}(\theta)) = 0$. This can be understood as the penalty associated with not using all the available points for calculating the likelihood at every time step. It results in an extra term in the corresponding error expressions given in Theorems 5 and 6 when compared with the Euler method. More precisely, for $n \ll N$ the term $-\frac{1}{2} \text{Var}(\hat{f}(\theta))$ is of size $\mathcal{O}(N^2)$ thus making the leading order error term $\mathcal{O}(hN^2)$ in Equation (37).

Example 1 We illustrate the above findings on the toy model discussed in Section 2. In particular, using the expression for $u(\theta, t)$ from Section 3.3 (replacing μ and σ by μ_p and σ_p respectively), and that $\text{Var}(\hat{f}(\theta)) = \text{Var}(B)$ for this simple model, the extra term in Equation (50) when compared with the Euler method is now given by

$$-\frac{1}{2} \text{Var}(B) \partial_{\theta}^2 u(\theta, t) = -\text{Var}(B) e^{-t/\sigma_p^2}.$$

A simple integration of this term according to the formula (41) gives that the overall contribution of the extra term, which is,

$$\sigma_p \text{Var}(B) = \frac{\text{Var}(B)}{2A},$$

and thus agreeing with Equation (19) derived in Section 2.1.

4.2 Modified SGLD

As we have seen in the previous section, the SGLD method introduces an extra term $-\frac{1}{2} \text{Var}(\hat{f}(\theta)) \nabla_{\theta}^2$ in the leading order error term related to the weak error (Theorem 5) and to the ergodic averages (Theorem 6). When $n \ll N$, this term is of order $\mathcal{O}(hN^2)$. In this section we will explore a modification of SGLD (mSGLD) for which this term is removed, so that the leading order term is exactly the same as for the Euler-Maruyama scheme. Specifically, the mSGLD updates are,

$$\theta_{j+1} = \theta_j + h\hat{f}(\theta_j) + \sqrt{h} \left(1 - \frac{h}{2} \text{Var}\hat{f}(\theta_j) \right) \xi_j. \quad (51)$$

We can again derive the weak order expansion as in the previous subsection. Our first step is to expand $\phi(\theta_{j+1})$ in powers of h and then take expectations with respect to the random variable ξ_j . In particular, we obtain

$$\begin{aligned} \mathbb{E}_{\xi_j}(\phi(\theta_{j+1})) &= \phi(\theta_j) + h \left(\hat{f}_j(\theta_j) \phi'(\theta_j) + \frac{1}{2} \phi''(\theta_j) \right) \\ &\quad + \frac{h^2}{2} \left(\left[\hat{f}_j^2(\theta_j) - \text{Var}\hat{f}(\theta_j) \right] \phi''(\theta_j) + \hat{f}_j(\theta_j) \phi^{(3)}(\theta_j) + \frac{1}{4} \phi^{(4)}(\theta_j) \right) + \mathcal{O}(h^3). \end{aligned}$$

Taking expectations with respect to the random sampling and using Equation (48), we obtain

$$\begin{aligned} \mathbb{E}(\phi(\theta_{j+1})) &= \phi(\theta_j) + h\mathcal{L}\phi(\theta_j) \\ &\quad + \frac{h^2}{2} \left(\left[\mathbb{E}_{\tau_j}(\hat{f}_j^2(\theta_j)) - \text{Var}\hat{f}(\theta_j) \right] \phi''(\theta_j) + f(\theta_j) \phi^{(3)}(\theta_j) + \frac{1}{4} \phi^{(4)}(\theta_j) \right) + \mathcal{O}(h^3), \end{aligned} \quad (52)$$

where \mathcal{L} is the generator of Equation (1). We thus see that the mSGLD is a first order weak method and

$$A_1(\pi)\phi = \frac{1}{2} \left(\left[\mathbb{E}_{\tau}(\hat{f}^2(\theta)) - \text{Var}\hat{f}(\theta) \right] \phi'' + f(\theta) \phi^{(3)} + \frac{1}{4} \phi^{(4)} \right)$$

Using the expression for \mathcal{L}^2 as in the case of SGLD, we have that,

$$\begin{aligned} \frac{1}{2}\mathcal{L}^2 - A_1 &= \frac{1}{2} \left(f(\theta)f'(\theta) + \frac{1}{2}f''(\theta) \right) \frac{d}{d\theta} + \frac{1}{2} \left(f'(\theta) + f^2(\theta) - \mathbb{E}_\tau(\hat{f}^2(\theta)) + \text{Var}\hat{f}(\theta) \right) \frac{d^2}{d\theta^2} \\ &= \frac{1}{2} \left(f(\theta)f'(\theta) + \frac{1}{2}f''(\theta) \right) \frac{d}{d\theta} + \frac{1}{2}f'(\theta)\frac{d^2}{d\theta^2} \end{aligned} \quad (53)$$

We see that the leading order term in the weak error and the error for the ergodic averages is the same as for the Euler method, which uses all data at every step. In higher dimensions, a similar calculation gives the mSGLD updates,

$$\theta_{j+1} = \theta_j + h\hat{f}_j(\theta_j) + \sqrt{h} \left(I - \frac{h}{2}\text{Cov}\hat{f}(\theta_j) \right) \xi_j \quad (54)$$

where

$$\text{Cov}\hat{f}(\theta) = \mathbb{E} \left[\left(\hat{f}(\theta) - \mathbb{E}(\hat{f}(\theta)) \right) \left(\hat{f}(\theta) - \mathbb{E}(\hat{f}(\theta)) \right)^\top \right]$$

and ξ_j is a d -dimensional standard normal random variable.

Remark 7 *Except for special cases, $\text{Var}f(\theta_j)$ does not have a closed form. The simplest possible way to proceed without it is to replace it by an unbiased estimator, for example in case of sampling without replacement,*

$$\widehat{\text{Var}}\hat{f}_j(\theta) := \frac{N(N-n)}{n(n-1)} \sum_{i=1}^n \left(\nabla \log \pi(x_{\tau_{ji}} | \theta) - \frac{\hat{f}_j(\theta)}{N} \right)^2.$$

This replacement does not change Equation (53) because the smallest order contribution to Equation (49) is of the form

$$-h^2 \mathbb{E} \left[\widehat{\text{Var}}f(\theta_j) \xi_j^2 \right] = -h^2 \text{Var}f(\theta_j).$$

However, estimating the variance of the stochastic gradient will affect higher order terms in h . For fixed h these terms may have larger contribution to the overall error depending on the choice of n and N . In fact, this is true even if we use the exact variance for the toy model in Section 2.1. More precisely, we compare the bias of the mSGLD and the SGLD in Equation (70) notice that h^2 term might be larger depending on the choice of n and N .

5. Finite Time Sample Averages

Having focused on the SGLD in the asymptotic regime, we will now provide non-asymptotic analysis of the mean squared error (MSE) of the finite time sample averages of the SGLD. In particular, we will decompose the MSE into bias and variance. The main result of this section will be of the form

$$\begin{aligned} \text{Bias:} & \quad \left| \mathbb{E} \frac{1}{K} \sum_{i=0}^{K-1} \phi(\theta_i) - \int \phi(x)\pi(x)dx \right| = O \left(h + \frac{1}{Kh} \right) \\ \text{MSE:} & \quad \mathbb{E} \left(\frac{1}{K} \sum_{i=0}^{K-1} \phi(\theta_i) - \int \phi(x)\pi(x)dx \right)^2 = O \left(h^2 + \frac{1}{Kh} \right) \end{aligned} \quad (55)$$

Remark 8 In Teh et al. (2014), a central limit theorem was provided for the decreasing step size SGLD which shows a convergence rate of $O(K^{-\frac{1}{3}})$. At first sight, the bound in Equation (55) seems better because of the $\frac{1}{Kh}$ term in the upper bound. However, due to the bias, an additional term of order $O(h^2)$ appears. In order to compare (55) with the previous result of Teh et al. (2014), we optimise the sum of both terms over the step size h . This results in a bound on the MSE of the SGLD of order $O(K^{-\frac{2}{3}})$ and agrees with the rate achieved by the decreasing step size SGLD. This agreement between decreasing step size discretisation and fixed step size discretisation is, to our knowledge, not a widely-known observation in the literature. In contrast, for standard MCMC algorithms the MSE is bounded by order $O(K^{-1})$ due to the Metropolis-Hastings correction that removes the bias. Nevertheless, experimental results in the literature demonstrate that the SGLD might be advantageous in the initial transient phase of learning, see e.g. Patterson and Teh (2013); Chen et al. (2014)

In Section 5.2 we will focus on establishing the bound in Equation (55) which is an extension of the work by Mattingly et al. (2010). The authors obtained similar results for finite time sample averages of discretisations of diffusions of the form

$$d\theta_t = f(\theta_t) + g(\theta_t)dW_t \quad (56)$$

on the torus which we review subsequently in Section 5.1.

5.1 Preliminaries on the Poisson Equation and Time Averages

In the following a connection between time averages of the diffusion and the corresponding Poisson equation will be presented. For a more elaborate description of this technique we point the reader to Section 4.2 of Mattingly et al. (2010) and references therein.

The Poisson equation is an elliptic PDE on the basis of the generator associated with Equation (56). The generator of Equation (56) is

$$\mathcal{L}\psi = \nabla\psi \cdot \nabla f + \frac{1}{2}g(\theta)^\top \nabla^2 \psi g(\theta),$$

while the Poisson equation is given by

$$\mathcal{L}\psi = \phi - \bar{\phi} \text{ on } \mathbb{R}^d \quad (57)$$

where ϕ is a test function and $\bar{\phi} := \int \phi(x)\pi(dx)$ with π being the invariant distribution of (56). For applications in Bayesian statistics π represents the posterior and the quantity $\bar{\phi}$ the posterior expectation of interest. The posterior expectation $\bar{\phi}$ is estimated by the time average $\frac{1}{t} \int_0^t \phi(\theta(s))ds$ of the Langevin dynamics. The difference between the two can be expressed explicitly by using Itô's formula on the solution ψ of the Poisson equation

$$\begin{aligned} \psi(\theta(t)) - \psi(\theta(0)) &= \int_0^t \phi(\theta(s)) - \bar{\phi} ds + \int_0^t \nabla\psi(\theta(s)) \cdot g(\theta(s))dW_s, \\ \frac{1}{t} \int_0^t \phi(\theta(s))ds - \bar{\phi} &= \frac{1}{t} (\psi(\theta(t)) - \psi(\theta(0))) - \frac{1}{t} \int_0^t \nabla\psi(\theta(s)) \cdot g(\theta(s))dW_s. \end{aligned}$$

If the first term and the variance of the second term (the martingale term) on the right hand side can be bounded, an error bound for the time average is obtained.

In this article, we are interested in the time average of the Euler discretisation and the SGLD. We can build on the ideas of Section 5 in Mattingly et al. (2010) which considers time discretisations of Equation (56) of the following form

$$\theta_{k+1} = \theta_k + hf(\theta_k, h) + \sqrt{h}g(\theta_k, h)\eta_k, \quad \eta_k \sim \mathcal{N}(0, I).$$

In Mattingly et al. (2010) a Taylor expansion is used to express

$$\Delta\psi(\theta_{k+1}) := \psi(\theta_{k+1}) - \psi(\theta_k) = h(A_0\psi)(\theta_k) + R_k$$

where R_k is the remainder term. The term A_0 was introduced in Equation (33) in Section 3.1.

Using that $\mathcal{L}\psi = \phi - \bar{\phi}$, summing over k and dividing by hK yields

$$\hat{\phi}_K := \frac{1}{K} \sum_{k=0}^{K-1} \phi(\theta_k) = \bar{\phi} + \frac{1}{Kh} (\psi(\theta_K) - \psi(\theta_0)) - \frac{1}{hK} \sum_{k=0}^{K-1} h(A_0 - \mathcal{L})\psi(\theta_k) - \frac{1}{Kh} \sum_{k=0}^{K-1} R_k.$$

Controlling $A_0 - \mathcal{L}$ and the remainder gives rise to Theorem 5.1 and 5.2 in Mattingly et al. (2010) stating that

$$\begin{aligned} \left| \mathbb{E} \hat{\phi}_K - \bar{\phi} \right| &\leq C \left(h + \frac{1}{h \cdot K} \right) \\ \mathbb{E} \left(\hat{\phi}_K - \bar{\phi} \right)^2 &\leq C \left(h^2 + \frac{1}{h \cdot K} \right). \end{aligned} \tag{58}$$

In particular, these results were derived for discretisations of SDEs on the torus. This simplifies the presentation because the derivatives of ψ are bounded on a compact set. However, the same arguments hold if the following assumption is imposed instead

$$\sup_k \mathbb{E} \left\| \psi^{(i)}(\theta_k) \right\| < \infty \text{ for } i = 1, \dots, 4. \tag{59}$$

verifying this condition will allow us to work on \mathbb{R}^d .

5.2 The Bias and the MSE of Finite Time SGLD Averages

We consider the SDE

$$d\theta_t = f(\theta_t)dt + g(\theta_t) dW_t. \tag{60}$$

with $g = I$ being the identity matrix but we keep g in order to make the presentation clearer. Based on this setup the recursion of the corresponding SGLD reads as follows

$$\Delta_{k+1} = \theta_{k+1} - \theta_k = \hat{f}_k h + h^{\frac{1}{2}} g_k \xi_{k+1}$$

where \hat{f}_k is an unbiased estimate of f . The focus of this section is to establish results similar to Equation (58) for the SGLD. They will be formulated in Theorem 9.

For the readability of the subsequent calculation we use the following notations

$$\Delta_{k+1} = \theta_{k+1} - \theta_k, \phi_k = \phi(\theta_k),$$

$\hat{f}_k = \hat{f}(\theta_k, \tau_k, h)$ for the estimate of the drift, $g_k = g(\theta_k, h) = I$, $\psi_k = \psi(\theta_k)$, $V_k = V(\theta_k)$ and $D^k \psi_k = D^k \psi(\theta_k)$. The term A_0 , as introduced in Equation (33) in Section 3.1, satisfies $A_0 = \mathcal{L}$ but we keep A_0 for clarity. Thus, we have

$$A_0 \psi_k = \nabla \psi \cdot \mathbb{E}_\tau \hat{f}(\theta_k, \tau, h) + \frac{1}{2} \text{trace} (S(\theta_k)^t \nabla^2 \psi(\theta_k))$$

where $S(\theta) = g(\theta)g(\theta)^T = I$.

We use the following third order Taylor expansion on $\psi(\theta_{k+1}) - \psi(\theta_k)$ in order to obtain a bound on $\frac{1}{K} \sum_{k=0}^{K-1} (\phi_k - \bar{\phi})$

$$\begin{aligned} \psi_{k+1} &= \psi_k + \nabla \psi_k \cdot \Delta_{k+1} + \frac{1}{2} \Delta_{k+1}^T \nabla^2 \psi_k \Delta_{k+1} + \frac{1}{6} \psi_{k+1}^{(3)}(\Delta_{k+1}, \Delta_{k+1}, \Delta_{k+1}) + R_{k+1} \\ R_{k+1} &= \frac{1}{6} \int_0^1 s^3 \psi^{(4)}(s\theta_k + (1-s)\theta_{k+1})(\Delta_{k+1}, \Delta_{k+1}, \Delta_{k+1}, \Delta_{k+1}) ds. \end{aligned}$$

Here $\psi^{(3)}$ and $\psi^{(4)}$ are the third and fourth order derivative in the form of a trilinear and a quadrilinear form, respectively. In this setting, a third order expansion is required in order to obtain the h^2 term in the $h^2 + \frac{1}{T}$ bound in the MSE (see Equation (58) or Theorem 9). More precisely, the remainder of this expansion is forth order which together with the term $\sqrt{h} \xi_m$ in Equation (5) contributes to the h^2 error term. In order to make the connection to the Poisson equation, we write the expansion above in terms of A_0 . This yields

$$\begin{aligned} \psi_{k+1} &= \psi_k + h A_0 \psi_k + h^{\frac{1}{2}} \nabla \psi_k \cdot (g_k \xi_{k+1}) + h \nabla \psi_k \cdot \underbrace{\left(\hat{f}_k - \mathbb{E}_\tau \hat{f}(\theta_k, \tau, h) \right)}_{H_k} + h^{\frac{3}{2}} (g_k \xi_{k+1})^T \nabla^2 \psi_k \hat{f}_k \\ &\quad + \frac{1}{2} \hat{f}_k^T h^2 \nabla^2 \psi_k \hat{f}_k + \frac{1}{6} \psi_k^{(3)}(\Delta_{k+1}, \Delta_{k+1}, \Delta_{k+1}) + r_{k+1} + R_{k+1} \end{aligned}$$

where $r_{k+1} = h^{\frac{1}{2}} \left((g_k \xi_{k+1})^T \nabla^2 \psi_k (g_k \xi_{k+1}) - S(x, h) \right)$.

Notice that $\frac{1}{hK} \sum_{k=0}^{K-1} h A_0 \psi_k = \frac{1}{K} \sum_{k=0}^{K-1} (\phi_k - \bar{\phi})$ is the error of interest. In order to control this error, we sum the expression for ψ_{k+1} for $k = 0, \dots, K-1$ and divide by

$T = hK$. Grouping the terms for subsequent inspection gives

$$\begin{aligned}
 \frac{\psi_K - \psi_0}{Kh} &= \frac{1}{K} \sum_{k=0}^{K-1} (\phi_k - \bar{\phi}) + \underbrace{\sum_{k=0}^{K-1} (\mathcal{L} - A_0) \psi_k}_0 \\
 &+ \underbrace{\frac{1}{T} \sum_{k=0}^{K-1} r_{k+1}}_{M_{1,K}} + \underbrace{\frac{1}{T} h^{\frac{1}{2}} \sum_{k=0}^{K-1} \nabla \psi_k (g_k \xi_{k+1})}_{M_{2,K}} + \underbrace{\frac{1}{T} h^{\frac{3}{2}} \sum_{k=0}^{K-1} \hat{f}_k^T \nabla^2 \psi_k (g_k \xi_{k+1})}_{M_{3,K}} \\
 &+ \underbrace{\frac{1}{T} h \sum_{k=0}^{K-1} \nabla \psi_k \cdot (\hat{f}_k - \mathbb{E}_\tau f(\theta_k, \tau, h))}_{M_{4,K}} + \underbrace{\frac{1}{T} \frac{1}{2} \sum_{k=0}^{K-1} h^2 \hat{f}_k^T \nabla^2 \psi_k \hat{f}_k}_{S_{1,K}} \\
 &+ \underbrace{\frac{1}{T} \sum_{k=0}^{K-1} R_{k+1}}_{S_{2,K}} + \underbrace{\frac{1}{T} \frac{1}{6} \sum_{k=0}^{K-1} \psi_k^{(3)} (\Delta_{k+1}, \Delta_{k+1}, \Delta_{k+1})}_{S_{3,K}}. \tag{61}
 \end{aligned}$$

where the $M_{i,k}$ indicate the martingale terms and the $S_{i,k}$ other remainder terms. We split

$$S_{3,K} = M_{0,K} + \tilde{M}_{0,K} + S_{0,K} + \tilde{S}_{0,K}$$

in terms of

$$\begin{aligned}
 M_{0,K} &= \frac{1}{6} h^{\frac{3}{2}} \sum_{k=0}^{K-1} \left(\psi_k^{(3)} ((g_k \eta_{k+1}), (g_k \eta_{k+1}), (g_k \eta_{k+1})) \right) \\
 \tilde{M}_{0,K} &= \frac{1}{2} \sum_{k=0}^{K-1} h^{\frac{5}{2}} \psi_k^{(3)} (\hat{f}_k, \hat{f}_k, g_k \eta_{k+1}) \\
 S_{0,K} &= \frac{1}{6} \sum_{k=0}^{K-1} 3h^2 \mathfrak{z} \psi_k^{(3)} (g_k \eta_{k+1}, g_k \eta_{k+1}, \hat{f}_k) \\
 \tilde{S}_{0,K} &= \frac{1}{6} \sum_{k=0}^{K-1} h^3 \psi_k^{(3)} (\hat{f}_k, \hat{f}_k, \hat{f}_k).
 \end{aligned}$$

Rearranging Equation (61) for $\frac{1}{K} \sum_{k=0}^{K-1} (\phi_k - \bar{\phi})$ and controlling the resulting right hand side of Equation (61) gives rise to the following theorem.

Theorem 9 *Suppose that there exists a function V such that the following three assumptions hold:*

1. *There are $p_{\psi,1}, \dots, p_{\psi,4} \in (0, \infty)$ such that the derivatives of the solution ψ to the Poisson equation satisfy the following bound*

$$\|\psi^{(k)}\| \lesssim V^{p_{\psi,k}}, \quad \text{for } k = 0, \dots, 4. \tag{62}$$

2. The drift f and the error from the estimate $H := \hat{f}(\theta, \tau) - f(\theta)$ satisfy

$$\begin{aligned} \mathbb{E}_\tau H(\theta, \tau)^{2p} &\lesssim V(\theta)^p \quad \forall p \leq p^* \\ \|f\|^2 &\lesssim V. \end{aligned} \quad (63)$$

for $p^* = \max\{2p_{\psi,2} + 2, 2p_{\psi,4} + 4, 2p_{\psi,3} + 1, 2p_{\psi,3} + 3\}$. Moreover, we suppose that the $\mathbb{E}V^p(\theta_k)$ is bounded from above and that this bound is independent of k , that is

$$\sup_k \mathbb{E}V^p(\theta_k) < \infty, \quad \forall p \leq p^*. \quad (64)$$

3. V satisfies

$$\sup_s V(s\theta_1 + (1-s)\theta_2)^p \lesssim V(\theta_1)^p + V(\theta_2)^p, \quad \text{for all } \theta_1, \theta_2, p \leq p^*. \quad (65)$$

Under these assumptions there exists $h_0 > 0$ and constant C such that for all $h < h_0$

$$\text{Bias}(\hat{\phi}_K) = \left| \mathbb{E}\hat{\phi}_K - \bar{\phi} \right| \leq C \left(h + \frac{1}{Kh} \right) \quad (66)$$

$$\mathbb{E}(\hat{\phi}_K - \bar{\phi})^2 \leq C \left(h^2 + \frac{1}{Kh} \right) \quad (67)$$

where

$$\hat{\phi}_K = \frac{1}{K} \sum_{k=0}^{K-1} \phi(\theta_k) \quad \text{and} \quad \bar{\phi} = \mathbb{E}_\pi \phi.$$

Proof For each term we bound the term inside the sum by a power of V^p and then obtain an overall bound using $\sup_i \mathbb{E}V_i^p < \infty$. For example, $\frac{1}{T}\mathbb{E}S_{1,K}$ can be bounded as follows

$$\begin{aligned} \frac{1}{T}\mathbb{E}S_{1,K} &\lesssim \frac{1}{T}\mathbb{E} \sum_{k=0}^{K-1} h^2 V_k^{p_{\psi,2}} \mathbb{E}_{\tau_k} \left\| \hat{f}_k \right\|^2 \\ &\lesssim \frac{1}{T} \sum_{k=0}^{K-1} h^2 \sup_i \mathbb{E}V_i^{p_{\psi,2}+1} \lesssim \frac{1}{T} h^2 K \lesssim h. \end{aligned}$$

The details of this computation are contained in Appendix 9. ■

Theorem 9 and the results for the decreasing step size SGLD Teh et al. (2014) hold under assumptions formulated in terms of the solution ψ of the Poisson equation. More precisely, the crucial step is to establish a bound of the form

$$\sup_k \mathbb{E} \left\| \psi_{k+1}^{(i)}(\theta_k) \right\| < \infty \quad \text{for } k = 1, \dots, 4.$$

This bound is established using Equations (62) and (64)

$$\sup_k \mathbb{E} \left\| D^{(i)}\psi(\theta_k) \right\| \lesssim \sup_k \mathbb{E} \|V\|^{p_{\psi,i}} < \infty \quad \text{for } i = 1, \dots, 4.$$

Thus, we are left with finding an appropriate Lyapunov function V such that Equations (62) and (64) hold. In Appendix 9.1, we formulate strong sufficient conditions on π that ensure that these assumptions are satisfied and that Theorem 9 is applicable.

6. An Analytic Investigation of the Toy Model

We now extend our analysis of the one-dimensional Gaussian toy model introduced in Section 2 beyond the general results of the previous two sections. More precisely, in Section 6.1, we compare the Euler method, the SGLD and the mSGLD by comparing the computational cost for fixed level accuracy specified in terms of the mean square error in estimating the second moment (MSE_2), optimising over the step size h , the subsample size n and the number of steps M . A numerical solution to the resulting optimisation problem demonstrates that the SGLD is advantageous in the lower accuracy regime while it degenerates to $n = N$ in the high accuracy regime. On the other hand the mSGLD does not degenerate and seems to maintain a constant speed up compared to the Euler method. In Section 6.2 we then consider the MSE_2 and use an analytic expression to study the behaviour of these algorithms for growing N . This allows us to extend the analysis of Sections 2 and 4 (in which we only consider the case limit $h \rightarrow 0$) and study the asymptotic bias of the SGLD and the mSGLD by scaling both n and h in N .

In Section 6.3 we finally adopt a different viewpoint by considering a fixed value of our parameter θ , denoted by θ^\dagger , while we take expectations with respect to the realisation of the data $\{X_i\}$. This enables us to study how $\mathbb{E}_X(\text{MSE}_2)$ behaves in the limit of $N \rightarrow \infty$. In particular, we find that for the case of the SGLD, the computational cost in order for $\mathbb{E}_X(\text{MSE}_2) \rightarrow 0$ is reduced by a factor of N when compared to the Euler method. A similar analysis for the expected relative error in estimating the posterior variance (ERE)

$$\text{ERE} = \mathbb{E}_{\{\theta_i\}_i} \frac{\frac{1}{K} \sum_{i=0}^{K-1} \theta_i^2 - \left(\frac{1}{K} \sum_{i=0}^{K-1} \theta_i \right)^2}{\sigma_p^2} - 1. \quad (68)$$

reveals that under the constraint $\mathbb{E}_X(\text{ERE}) \rightarrow 0$ the Euler method and the SGLD have the same computational cost on the algebraic scale in N .


6.1 Minimising Computational Effort for Constrained MSE_2

In Section 2.2 we compared the Euler method, the SGLD and the mSGLD for the same choice of $r = \frac{h}{A}$. In the following we numerically minimise the computational effort with respect to the condition $\text{MSE}_2 \leq \epsilon^2$. We assume that the computational cost is proportional to $M \cdot n$ which leads to the problem of solving


$$\begin{aligned} \min_F \quad & M \cdot n & (69) \\ \text{subject to} \quad & \text{MSE}_2(r, M, n) \leq \epsilon^2 \\ \text{w.r.t.} \quad & r < 1, M, n. \end{aligned}$$

Even though we have analytic expressions for the MSE_2 , the solution to the optimisation problem does not have a closed form. To conclude our analysis, we illustrate the numerical solution to this problem for $N = 1000$ for the Euler method, the SGLD and the mSGLD. The results, depicted in Figure 3, can be summarised as follows:

1. as ϵ becomes smaller, the gain of the SGLD over the Euler method in terms of computational effort decreases (due to the fact that n increases);



(a) Minimised computational cost for the Euler, the SGLD and the mSGLD algorithm (b) Subset size n for minimised computational cost



(c) Step size $r = \frac{h}{A}$ for minimised computational cost

Figure 3: Minimisation of computational cost $\propto M \cdot n$ subject to $MSE \leq \epsilon^2$

2. as ϵ becomes smaller, the mSGLD gains efficiency over the SGLD (the reason being that n seems to asymptote as ϵ decreases).

The upper bound obtained in Equation (67) suggests a scaling of $M \sim \epsilon^{-3}$ and $r \sim \epsilon$ to obtain an MSE of order ϵ^2 with minimal computational effort. The numerical minimisation of M with respect to r and M subject to the condition $MSE(r, M) \leq \epsilon^2$ confirms this scaling empirically, see Figure 4.

6.2 The MSE_2 for fixed and increasing N

We now consider the behaviour for growing data size N where a new data set X is generated in each instance. Figure 5 depicts the MSE_2 for $N = 10^i$ with $i = 1, \dots, 4$ for the subset choices $n = N^{0.1}, N^{0.5}$ and $N^{0.9}$; each compared to the Euler method corresponding to $n = N$. In this plot we notice that the SGLD outperforms the mSGLD for $n = N^{0.1}$ and $n = N^{0.5}$. The behaviour in Figure 5 suggests that the mSGLD has a larger bias than the SGLD which seems to contradict the findings of Section 2 and 4. Previously, we have just considered the asymptotic of $h \rightarrow 0$. In contrast, we scale both $h = r \frac{1}{A(N)}$ and $n = N^p$ in




Figure 4: Scaling of $r(\epsilon)$ and $M(\epsilon)$ for minimal computational cost subject to the $\text{MSE}_2(r(\epsilon), M(\epsilon)) \leq \epsilon^2$

terms of N in Figure 5. In the following we investigate this relationship further by using the explicit formula for the asymptotic bias which has been made available in Section 2 .

For simplicity we consider sampling without replacement, using the expression in Equation (16). Similar conclusions hold for sampling with replacement. First we consider the mSGLD. From Equation (21) and using the parameterisation $h = r/A$ of the step size, the excess asymptotic bias becomes

$$h^2 \frac{\text{Var}(B)^2}{4(2A - A^2h)} = \frac{(r/A)^2 \left(\frac{N-n}{n}\right)^2 N^2 \text{Var}(X)^2}{4(2-r)A}.$$

Because $A \sim N$, we see that the excess bias stays bounded for large N if and only if $n \gtrsim N^{\frac{1}{2}}$. In contrast, the same consideration for the SGLD shows that the excess bias due to subsampling vanishes so long as $n \rightarrow \infty$ when $N \rightarrow \infty$.

We can also identify the regime in which the mSGLD has a smaller asymptotic bias than the SGLD. From Equations (18) and (21) we see that this is the case when

$$\frac{h^2 \text{Var}^2(B)}{4(2A - A^2h)} \leq \frac{h \text{Var}(B)}{2A - A^2h}. \tag{70}$$

Using Equation (16), the above can be rearranged to

$$1 \geq \frac{h}{4} \frac{1}{16\sigma_x^2} \frac{N(N-n)}{N} \text{Var}(X).$$

Let $c = \frac{n}{N}$ be the relative size of the subsampling. Using $h = r/A$ where A is given in Equation (15), we get

$$c \geq \frac{2rN\text{Var} X}{16\sigma_x^2 \left(\frac{1}{\sigma_\theta^2} + \frac{N}{\sigma_x^2}\right) + 2rN\text{Var} X}$$

which in the limit of large data sets $N \rightarrow \infty$ yields

$$c \geq \frac{2r\text{Var} X}{16 + 2r\text{Var} X} > 0.$$

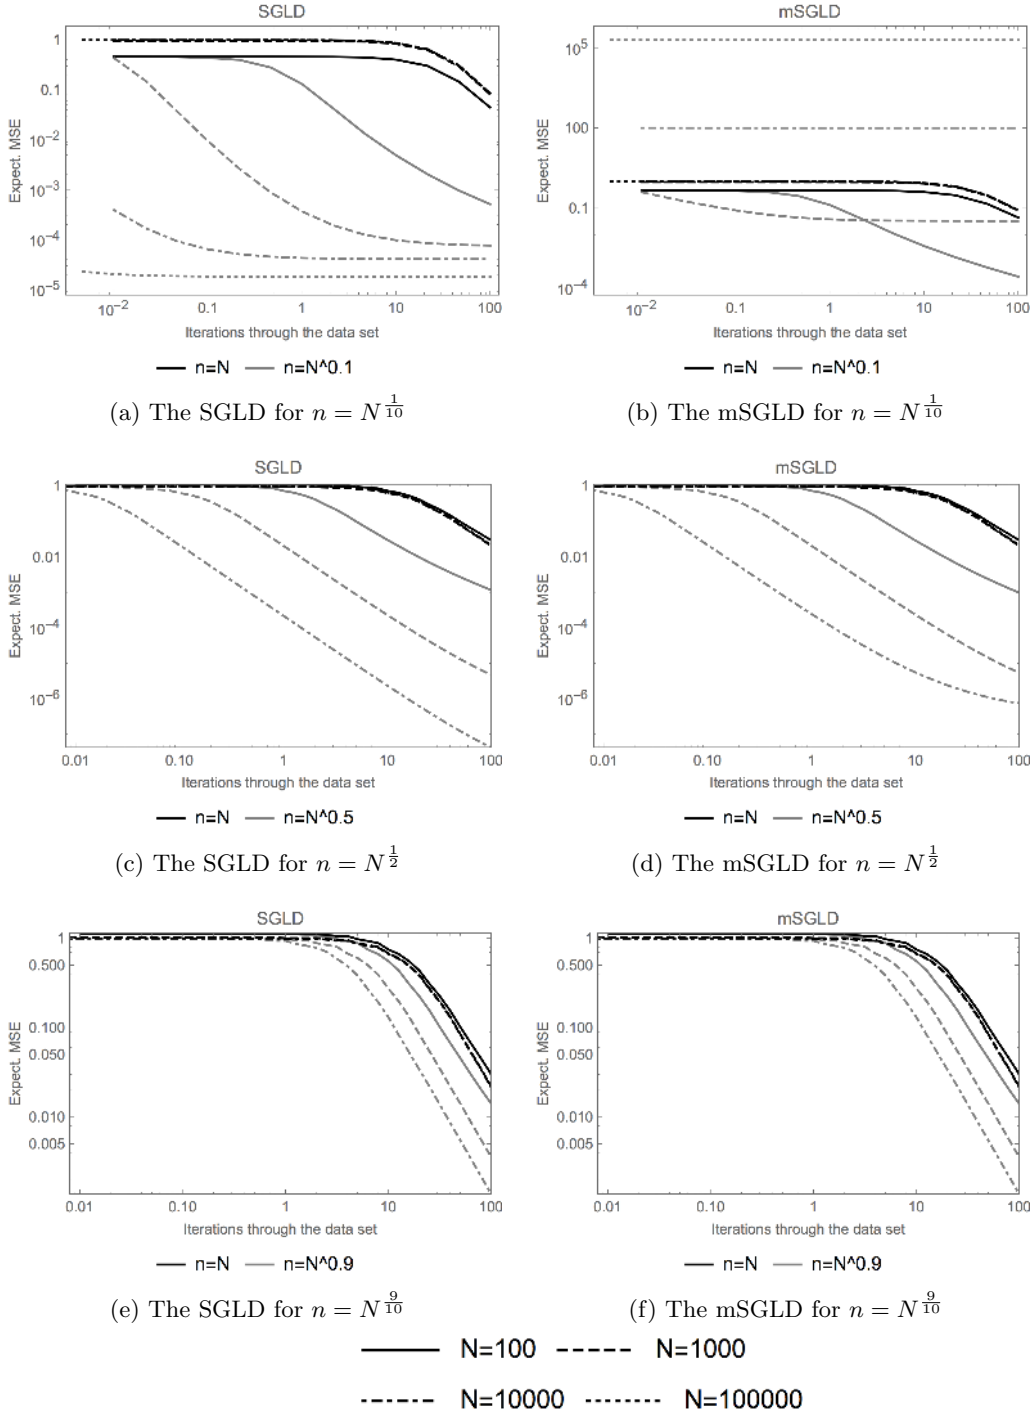


Figure 5: MSE_2 of the time average for the SGLD and the mSGLD for the second moment of the posterior as $N \rightarrow \infty$. Notice that Figures (c) and (d) and (e) and (f) have the same scaling respectively. Moreover, figures (a) and (b) have separate scaling because of the instability of the mSGLD.

In conclusion, for a fixed step size given by r/A , the mSGLD has a smaller bias than the SGLD if the above holds. In other words, the subsampling size has to be linear in the size of the data set for a fixed choice of r .

6.3 Limit of the MSE and ERE for well-specified Data as $N \rightarrow \infty$

In order to investigate the limit of $N \rightarrow \infty$, we need to specify the behaviour of the data as well. We study this in the well-specified case, in other words, we assume that the data is generated by the model for $\theta^\dagger = 1$. Previously, we have obtained an analytic expression for the expectation of the MSE_2 with respect to the realisation of the noise driving the algorithm. In contrast, we take expectations with respect to the realisation of the data X and the noise driving the algorithm in the following results. These results are formulated in analytic expressions² for the MSE_2 and the ERE depending only on M, n, r and N . We then choose M, n and r as functions of N and study the limit $N \rightarrow \infty$ and how this affects the computational cost and the behaviour of the ERE and the MSE_2 as $N \rightarrow \infty$ for the different algorithms.

For the Euler method ($n = N$) we need to take $h < \frac{1}{A} \asymp \frac{1}{N}$ in order to make Equation (17) stable. Moreover, we need the number of steps M to be of order N to approximate the diffusion to a time of order $\mathcal{O}(1)$. Because we evaluate N data points per step, this heuristic argument suggests that the complexity is of order $\mathcal{O}(N^2)$. Furthermore, we verify (using Mathematica[®]), that for the Euler method ($n = N$)

$$\lim_{N \rightarrow \infty} \mathbb{E}_X \text{MSE} = 0, \quad \lim_{N \rightarrow \infty} \mathbb{E}_X \text{ERE} = 0 \tag{71}$$

for the choices $M = N^{1+2\epsilon}$ and $r = N^{-\epsilon}$ for any $\epsilon > 0$. The computational cost for fixed N is $M \cdot n = N^{2+2\epsilon}$. Thus, this confirms the heuristics we used for the Euler method in Section 6.2.

A natural next question to ask in terms of the SGLD is if one can have Equation (71) to hold but for smaller computational complexity than the Euler method. Using Mathematica[®], we obtain the following theorem for the MSE

Theorem 10 *For any $\epsilon > 0$ and the choices $h = N^{-1-\epsilon}$, $M = N^{1+2\epsilon}$ and $n = 1$, the SGLD satisfies*

$$\lim_{N \rightarrow \infty} \mathbb{E}_{\theta_{ii}, X} \left(\frac{1}{M} \sum_{k=0}^{M-1} \theta_j^2 - (\mu_p^2 + \sigma_p^2) \right)^2 = 0.$$

This constitutes a substantial gain compared to the Euler method because it reduces the computational complexity in the data size N from being almost quadratic to almost linear.

We now draw our attention to the expected relative error in estimating the posterior variance, abbreviated by

$$\text{ERE} := \mathbb{E}_{\{\theta_i\}_i} \frac{\frac{1}{K} \sum_{i=0}^{K-1} \theta_i^2 - \left(\frac{1}{K} \sum_{i=0}^{K-1} \theta_i \right)^2}{\sigma_p^2} - 1. \tag{72}$$

2. see Appendix 10.2 for a sketch of the derivation for the MSE_2 (the derivation for ERE is similar)

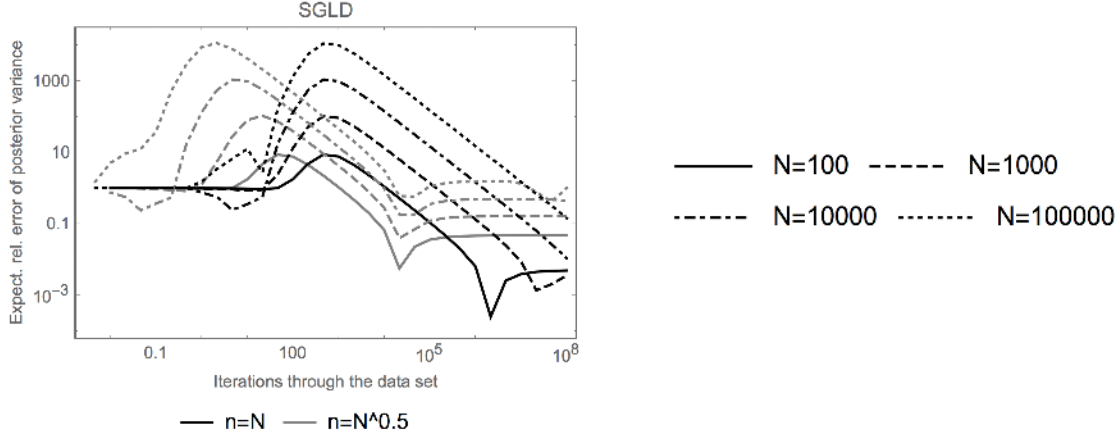


Figure 6: Expected relative error of the estimate of the variance of the posterior based on the SGLD.

Because the posterior variance goes to zero as $N \Rightarrow \infty$, it is conceivable that it requires more computational effort to ensure that $\lim_{N \rightarrow \infty} \mathbb{E}_X \text{ERE} = 0$. In order to illustrate the behaviour of the ERE, we first consider the behaviour for a fixed data set and repeat the experiment of Figure 5 in Figure 6. The latter demonstrates that the asymptotic bias for the choice $n = N^{\frac{1}{2}}$ (the asymptotes for the grey lines) have an increasing value in N . We used $h = \frac{1}{20A} \sim \frac{1}{N}$ which decreases with N . However, we show below that this requirement cancels exactly the gain from $n \ll N$ at least on the algebraic scale in N .

In particular, we now choose $r = N^{-\alpha}$, $n = N^\beta$ and $M = N^\gamma$. Hence the computational cost is $N^{\beta+\gamma}$. The step size $h = \frac{r}{A}$ satisfies $h \sim N^{-1-\alpha}$ because $A \sim N$. Since the algorithm performs $M = N^\gamma$ steps, we expect it to approximate the diffusion on the time interval $h \cdot M = N^{-1-\alpha+\gamma}$. Therefore it is reasonable to require that $\gamma > 1 + \alpha$. Under this assumption and with the help of Mathematica[®], we reduced the limit above to

$$\begin{aligned} \lim_{N \rightarrow \infty} \mathbb{E}_X \text{ERE} &= \lim_{N \rightarrow \infty} \left(\frac{2 \cdot N^{-\alpha-\beta+3}}{(N+1)^2 (N^{-\alpha} - 2)^2} - \frac{N^{-2\alpha-\beta+3}}{(N+1)^2 (N^{-\alpha} - 2)^2} \right) \\ &= \begin{cases} 0 & \text{if } \alpha + \beta > 1 \\ \infty & \text{if } \alpha + \beta < 1. \end{cases} \end{aligned}$$

Thus for $\lim_{N \rightarrow \infty} \mathbb{E}_X \text{ERE} = 0$ it is necessary that $\alpha + \beta \geq 1$. This condition in turn implies that the computational complexity satisfies

$$\underbrace{N^\gamma}_{\text{steps}} \times \underbrace{N^\beta}_{\text{cost per step}} = N^{1+\alpha} N^\beta = N^{1+\alpha+\beta} \gtrsim N^2.$$

Thus, there is no computational gain for the ERE in the limit $N \rightarrow \infty$ on the algebraic scale in N . We note that picking $\theta^* = \frac{1}{N}$ instead of $\theta^* = 1$ does not change the results. Thus, a closer initialisation does not change the result for the ERE.

7. Logistic Regression

In the following we present numerical simulations for a Bayesian logistic regression model. The data items are given by covariates $x_i \in \mathbb{R}^d$ that are labeled by $y_i \in \{-1, 1\}$. We assume the data $y_i \in \{-1, 1\}$ is modelled by

$$p(y_i|x_i, \beta) = \sigma(y_i \beta^t x_i) \tag{73}$$

where $\sigma(z) = \frac{1}{1+\exp(-z)} \in [0, 1]$. The model posses the assumption that y_i depends on x_i through the linear relationship $\beta^t x_i$. Nevertheless, logistic regression is commonly used after a preprocessing has taken place and is therefore used here for numerical illustration.

We put a Gaussian prior $\mathcal{N}(0, C_0)$ on β , for simplicity we use $C_0 = I$ subsequently. By Bayes' rule the posterior π satisfies

$$\pi(\beta) \propto \exp\left(-\frac{1}{2} \|\beta\|_{C_0}^2\right) \prod_{i=1}^N \sigma(y_i \beta^T x_i).$$

We consider $d = 3$ and $N = 1000$ data points and choose the covariate to be

$$x = \begin{pmatrix} x_{1,1} & x_{1,2} & 1 \\ x_{2,1} & x_{2,2} & 1 \\ \vdots & \vdots & \vdots \\ x_{1000,1} & x_{1000,2} & 1 \end{pmatrix}$$

for a fixed sample of $x_{i,j} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$ for $i = 1, \dots, 1000$ and $j = 1, 2$. We use a long run of the Random-Walk-Metropolis algorithm to estimate the posterior mean.

On that basis we estimate the MSE of the SGLD based mean estimate using 100 runs of the algorithm with step size $h = 0.002$ for various subset sizes. Figure 7 depicts the MSE as function of the iterations and effective iterations through the data set. Notice that for this example the variance of the stochastic gradient $\hat{f}(\theta)$ depends on both θ and all the data items. For this reason we replace $\text{Var } \hat{f}(\theta)$ by an estimate $\widehat{\text{Var}} \hat{f}(\theta)$, see also Remark 7. We note that the mSGLD is superior for $n = 150$, inferior for $n = 50$ and for $n = 10$ the MSE of the mSGLD does not drop below 1.

8. Conclusion

This article presents the mathematical foundations that are necessary for posterior sampling for stochastic gradient methods with fixed step sizes. We derived an error expansion of the asymptotic bias in terms of powers of the step size and identified how the constant in the leading order term depends on the unbiased estimator of the gradient. We construct a modified SGLD to match the Euler method in this asymptotic expansion. These asymptotic results are complemented by upper bounds on the bias and the MSE over a finite time horizon. Minimising the MSE with respect to the step size yields a decay of the error at the same rate as the decreasing step size SGLD, see Remark 8. These theoretical findings are completed with extensive analytic investigations of a one dimensional toy model that allows the derivation of analytic expressions for the sample average and its moments. Finally, this




Figure 7: Expected MSE of time average for the SGLD and the mSGLD for the mean of the posterior

yields an exact quantification of expected errors. The results of this investigation can be summarised as follows:

- In the high accuracy regime the SGLD deteriorates to the Euler method while the mSGLD prevails.
- For small data batches the bias of the mSGLD is larger than for the SGLD.
- In the limit as the number of data items goes to infinity the SGLD reduces computational complexity of estimating the second moment with vanishing MSE by one power of the number of data items.

This recommends the construction of new and a study of existing modifications of the SGLD such as the Stochastic Gradient Hamiltonian Monte Carlo Chen et al. (2014) and the Stochastic Gradient Thermostat Monte Carlo Ding et al. (2014) algorithms.

Acknowledgement

SJV and YWT acknowledge EPSRC for research funding through grant EP/K009850/1, EP/N000188/1 and EP/K009362/1. The authors thank Rémi Bardenet for fruitful discussions.

9. Appendix A: Finite time Ergodic Average based Poisson Equation

Proof [Proof of Theorem 9] Rearranging Equation (61) for $\frac{1}{K} \sum_{k=0}^{K-1} (\phi_k - \bar{\phi})$, the bias and the MSE can be controlled as follows:

$$\mathbb{E} \frac{1}{Kh} (\psi_K - \psi_0) \lesssim \frac{1}{T} \sup_i \mathbb{E} V_i^{p_{\psi,0}} \lesssim \frac{1}{T}.$$

Because $A_0 = \mathcal{L}$ for the SGLD it is left to bound to bound $\mathbb{E} \frac{1}{T} S_{i,K}$ for $i = 0 \dots 4$. First we consider $i = 1$ and use Equation (A.6)

$$\begin{aligned}
 \frac{1}{T}\mathbb{E}S_{1,K} &\lesssim \frac{1}{T}\mathbb{E}\sum_{k=0}^{K-1}h^2V_k^{p\psi,2}\mathbb{E}_{\tau_k}\|\hat{f}_k\|^2 \\
 &\lesssim \frac{1}{T}\sum_{k=0}^{K-1}h^2\sup_i\mathbb{E}V_i^{p\psi,2+1}\lesssim\frac{1}{T}h^2K\lesssim h.
 \end{aligned}$$

This procedure will be used over and over again. It can be summarised as follows:

1. bounding the terms in the sum by a power of V^p , using Equation (A.6) and the assumption on derivatives of ψ ;
2. then derive the bound using $\sup_i\mathbb{E}V_i^p < \infty$.

For $i = 2$ we additionally use that Equation (65) implies

$$\int_0^1 s^3\psi^{(4)}(s\theta_k+(1-s)\theta_{k+1})(\Delta_{k+1},\Delta_{k+1},\Delta_{k+1},\Delta_{k+1})ds\lesssim(V_k^{p\psi,4}+V_{k+1}^{p\psi,4})\|\Delta_k\|^4.$$

which allows us to follow the general procedure

$$\begin{aligned}
 \frac{1}{T}\mathbb{E}S_{2,K} &\lesssim \frac{1}{T}\mathbb{E}\sum_{k=0}^{K-1}R_{k+1} \\
 &\lesssim \frac{1}{T}\mathbb{E}\sum_{k=0}^{K-1}h^2(V_k^{p\psi,4}+V_{k+1}^{p\psi,4})\mathbb{E}_{\tau}\|\Delta_{k+1}\|^4 \\
 &\lesssim \frac{1}{T}\mathbb{E}\sum_{k=0}^{K-1}h^2\sup_i\mathbb{E}V^{p\psi,4+2}(\theta_i)\lesssim\frac{1}{T}h^2\lesssim h.
 \end{aligned}$$

We apply the general procedure to $\mathbb{E}S_{0,K}$

$$\begin{aligned}
 \frac{1}{T}\mathbb{E}S_{0,K} &\lesssim \frac{1}{T}\mathbb{E}\sum_{k=0}^{K-1}h^2V_k^{p\psi,3}\mathbb{E}_{\tau}\|g_k\eta_{k+1}\|^2\|\hat{f}_k\| \\
 &\lesssim \frac{1}{T}\mathbb{E}\sum_{k=0}^{K-1}h^2\sup_i\mathbb{E}V^{p\psi,3+\frac{1}{2}}(\theta_i) \\
 &\lesssim \frac{1}{T}h^2\lesssim h. \\
 \frac{1}{T}\mathbb{E}\tilde{S}_{0,K} &\lesssim \frac{1}{T}\mathbb{E}\sum_{k=0}^{K-1}h^3V_k^{p\psi,3}\mathbb{E}_{\tau}\|\hat{f}_k\|^3 \\
 &\lesssim \frac{1}{T}\sum_{k=0}^{K-1}h^3\sup_i\mathbb{E}V^{p\psi,3+\frac{3}{2}}(\theta_i)\lesssim\frac{1}{T}Kh^3\lesssim h^2.
 \end{aligned}$$

Thus, we have established the bound on the bias given by Equation (A.3).

In order to establish the bound on the MSE in Equation (67) we note that Equation (61) yields

$$\begin{aligned} \mathbb{E} \left(\frac{1}{K} \sum_{k=0}^{K-1} (\phi - \bar{\phi}) \right)^2 &\lesssim \mathbb{E} \frac{(\psi_K - \psi)^2}{T^2} \\ &\quad + \frac{1}{T^2} \sum_{i=0}^2 \mathbb{E} S_{i,K}^2 + \frac{1}{T^2} \sum_{i=0}^2 \mathbb{E} M_{i,K}^2 \end{aligned}$$

First we note that

$$\mathbb{E} \frac{(\psi_K - \psi)^2}{T^2} \lesssim \frac{1}{T^2} \sup_i \mathbb{E} V_i^{2p_\psi, 0} \lesssim \frac{1}{T^2}.$$

The $S_{i,K}^2$ terms can be bound in similar way as above with the additional use of Cauchy-Schwartz inequality to break the correlation between V_i^p and V_j^p .

$$\begin{aligned} \frac{1}{T^2} \mathbb{E} S_{1,K}^2 &\lesssim \frac{1}{T^2} \mathbb{E} \sum_{i,j=0}^{K-1} h^4 \nabla^2 \psi_i [\hat{f}_i, \hat{f}_i] \nabla^2 \psi_j [\hat{f}_j, \hat{f}_j] \\ &\lesssim \frac{1}{T^2} \mathbb{E} \sum_{i,j=0}^{K-1} h^4 \|\nabla^2 \psi_i\| \mathbb{E}_\tau \|\hat{f}_i\|^2 \|\nabla^2 \psi_j\| \mathbb{E}_\tau \|\hat{f}_j\|^2 \\ &\lesssim \frac{1}{T^2} \sum_{i,j=0}^{K-1} h^4 \mathbb{E} V_i^{p_\psi, 2+1} V_j^{p_\psi, 2+1} \\ &\lesssim \frac{1}{T^2} \sum_{i,j=0}^{K-1} h^4 \left(\mathbb{E} V_i^{2p_\psi, 2+2} \right)^{\frac{1}{2}} \left(\mathbb{E} V_j^{2p_\psi, 2+2} \right)^{\frac{1}{2}} \\ &\lesssim \frac{1}{T^2} \sum_{i,j=0}^{K-1} h^4 \left(\sup_i \mathbb{E} V_i^{2p_\psi, 2+2} \right) \lesssim \frac{K^2 h^4}{T^2} \lesssim h^2 \end{aligned}$$

Similarly, we bound

$$\begin{aligned}
 \frac{1}{T^2} \mathbb{E} S_{2,K}^2 &\lesssim \frac{1}{T^2} \mathbb{E} \sum_{i,j=0}^{K-1} R_{i+1} R_{j+1} \\
 &\lesssim \frac{1}{T^2} \mathbb{E} \sum_{i,j=0}^{K-1} (V_i^{p_\psi,4} + V_{i+1}^{p_\psi,4}) (V_j^{p_\psi,4} + V_{j+1}^{p_\psi,4}) \mathbb{E}_\tau \|\Delta_{i+1}\|^4 \mathbb{E}_\tau \|\Delta_{j+1}\|^4 \\
 &\lesssim \frac{h^4}{T^2} \mathbb{E} \sum_{i,j=0}^{K-1} (V_i^{p_\psi,4} + V_{i+1}^{p_\psi,4}) (V_j^{p_\psi,4} + V_{j+1}^{p_\psi,4}) V_i^2 V_j^2 \\
 &\lesssim \frac{h^4}{T^2} \mathbb{E} \sum_{i,j=0}^{K-1} (V_i^{p_\psi,4} + V_{i+1}^{p_\psi,4}) (V_j^{p_\psi,4} + V_{j+1}^{p_\psi,4}) V_i^2 V_j^2 \\
 &\lesssim \frac{h^4}{T^2} \sum_{i,j=0}^{K-1} \left(\sup_i \mathbb{E} V_i^{2p_\psi,4+4} \right) \lesssim \frac{K^2 h^4}{T^2} \lesssim h^2.
 \end{aligned}$$

Similar bounds can also be obtained for $\frac{1}{T^2} \mathbb{E} S_{0,K}^2$ and $\frac{1}{T^2} \mathbb{E} \tilde{S}_{0,K}^2$

$$\begin{aligned}
 \frac{1}{T^2} \mathbb{E} S_{0,K}^2 &\lesssim \frac{h^4}{T^2} \sum_{i,j=0}^{K-1} \mathbb{E} V_i^{p_\psi,3} \|\hat{f}_i\| V_j^{p_\psi,3} \|\hat{f}_j\| \\
 &\lesssim \frac{h^4 K^2}{T^2} \sup_i \mathbb{E} V_i^{2p_\psi,3+1} \lesssim h^2 \\
 \frac{1}{T^2} \mathbb{E} \tilde{S}_{0,K}^2 &\lesssim \frac{h^6}{T^2} \sum_{i,j=0}^{K-1} \mathbb{E} V_i^{p_\psi,3} \|\hat{f}_i\|^3 V_j^{p_\psi,3} \|\hat{f}_j\|^3 \\
 &\lesssim \frac{h^6 K^2}{T^2} \sup_i \mathbb{E} V_i^{2p_\psi,3+2} \lesssim h^4
 \end{aligned}$$

For Martingale terms the cross terms vanish which allows us to obtain the following bounds

$$\begin{aligned}
 \frac{1}{T^2} \mathbb{E} M_{1,K}^2 &= \frac{1}{T^2} \sum_{i=0}^{K-1} \left(\mathbb{E} D_{lm}^2 \psi_i \left(g_i^{l,a} \eta_{i+1}^a g_i^{m,b} \eta_{i+1}^b - g_i^{k,l} g_i^{k,m} \right) \right)^2 \\
 &\lesssim \frac{1}{T^2} \sum_{i=0}^{K-1} \sup_i \mathbb{E} V_i^{2p_\psi,2}.
 \end{aligned}$$

The following term is the crucial Martingale term as it yields the $\mathcal{O}\left(\frac{1}{T}\right)$ contribution

$$\begin{aligned}
 \frac{1}{T^2} \mathbb{E} M_{2,K}^2 &\lesssim \frac{1}{T^2} h \sum_{k=0}^{K-1} \mathbb{E} \|\nabla \psi_k\|^2 \|g_k \xi_{k+1}\|^2 \\
 &\lesssim \frac{1}{T^2} h \sum_{k=0}^{K-1} \mathbb{E} V_k^{2p_\psi,1} \lesssim \frac{1}{T}.
 \end{aligned}$$

Similarly, we estimate

$$\begin{aligned} \frac{1}{T^2} \mathbb{E} M_{3,K}^2 &\lesssim \frac{1}{T^2} h^3 \sum_{k=0}^{K-1} \mathbb{E} \|\nabla^2 \psi_k\|^2 \|\hat{f}_k\|^2 \|g_k \xi_{k+1}\|^2 \\ &\lesssim \frac{1}{T^2} h^3 \sum_{k=0}^{K-1} \mathbb{E} V_k^{2p_{\psi,2}+1} \lesssim \frac{h^2}{T}. \end{aligned}$$

The terms $\frac{1}{T^2} \mathbb{E} M_{0,K}^2$ and $\frac{1}{T^2} \mathbb{E} \tilde{M}_{0,K}^2$ can be bounded in the same way

$$\begin{aligned} \frac{1}{T^2} \mathbb{E} M_{0,K}^2 &\lesssim \frac{h^3}{T^2} \sum_{k=0}^{K-1} \mathbb{E} V_k^{2p_{\psi,3}} \lesssim \frac{h^3 K}{T^2} \leq \frac{h^2}{T} \\ \frac{1}{T^2} \mathbb{E} \tilde{M}_{0,K}^2 &\lesssim \frac{h^5}{T^2} \sum_{k=0}^{K-1} \mathbb{E} V_k^{2p_{\psi,3}+2} \lesssim \frac{h^5 K}{T^2} \leq \frac{h^4}{T} \end{aligned}$$

The additional part for the SGLD is the term corresponding to the Martingale $M_{4,K}$

$$\begin{aligned} \frac{1}{T^2} M_{4,K}^2 &\lesssim \frac{1}{T^2} \mathbb{E} h^2 \sum_{k=0}^{K-1} (\nabla \psi_k(H_k))^2 \\ &\lesssim \frac{1}{T^2} \mathbb{E} h^2 \sum_{k=0}^{K-1} V^{2p_{\psi,1}} \mathbb{E}_\tau \|H_k\|^2 \\ &\lesssim \frac{1}{T^2} h^2 \sum_{k=0}^{K-1} \mathbb{E} V^{2p_{\psi,1}+1} \lesssim \frac{h}{T} \end{aligned}$$

For all these calculations to go through need $\sup_i \mathbb{E} V_i^{p^*}$ to be bounded. Collecting the orders present, we see that

$$p^* = \max \{2p_{\psi,2} + 2, 2p_{\psi,4} + 4, 2p_{\psi,3} + 1, 2p_{\psi,3} + 3\}$$

is sufficient. ■

9.1 Sufficient Conditions on π Ensuring Finite Time Bounds on Bias and MSE

We formulate a sufficient condition on π that ensure that Theorem 9 is applicable. This hinges on deriving a sufficient condition for Equations (62), (64) and (65). The aim of this section is to establish and motivate the sufficient condition formulated in the following theorem.

Theorem 11 *Suppose the following condition holds*

$$\langle \theta, \nabla \log \pi_0(\theta) \rangle \leq -\alpha \|\theta\|^2 + \beta \tag{A.1}$$

$$\langle \theta, \nabla \log \pi(X_i | \theta) \rangle \leq -\alpha \|\theta\|^2 + \beta \text{ for } i = 1, \dots, N. \tag{A.2}$$

then Theorem 9 is applicable for polynomially bounded and continuous ϕ , that is

$$\begin{aligned} \text{Bias}(\hat{\phi}_K) &= \left| \mathbb{E} \hat{\phi}_K - \bar{\phi} \right| \leq C \left(h + \frac{1}{Kh} \right) \\ \mathbb{E} \left(\hat{\phi}_K - \bar{\phi} \right)^2 &\leq C \left(h^2 + \frac{1}{Kh} \right) \end{aligned}$$

Proof This result follows from Theorem 5.2 and results quoted below. ■

First we appeal to a sufficient condition for Equation(64) before summarising a regularity results of Pardoux and Veretennikov (2001) which allows us to establish Equation (62). We believe that the sufficient conditions above can be weakened, but this requires improving the results of Pardoux and Veretennikov (2001) which is out of the scope of this article.

The condition $\sup_k \mathbb{E} V^p(\theta_k) < \infty$ (that is Equation (64)) is established for all $p \leq p^*$ by Lemma 5 in Teh et al. (2014) if p^* satisfies the following assumption.

Assumption 12 *The drift term $\theta \mapsto \frac{1}{2} \nabla \log \pi(\theta)$ is continuous. The function $V : \mathbb{R}^d \rightarrow [1, \infty)$ tends to infinity as $\|\theta\| \rightarrow \infty$, is twice differentiable with bounded second derivatives and satisfies the following conditions:*

1. *V is a Lyapunov function for the Langevin dynamics, i.e. there are constants $\alpha, \beta > 0$ such that for every $\theta \in \mathbb{R}^d$ we have*

$$\left\langle \nabla V(\theta), \frac{1}{2} \nabla \log \pi(\theta) \right\rangle \leq -\alpha V(\theta) + \beta. \tag{A.3}$$

2. *The following bounds hold*

- *There exists an exponent $p_H \geq 2$ such that*

$$\mathbb{E}[\|H(\theta, \mathcal{U})\|^{2p_H}] \lesssim V^{p_H}(\theta). \tag{A.4}$$

Moreover, this implies that $\mathbb{E}[\|H(\theta, \mathcal{U})\|^{2p}] \lesssim V^p(\theta)$ for any exponent $0 \leq p \leq p_H$.

- *For every $\theta \in \mathbb{R}^d$ we have*

$$\|\nabla V(\theta)\|^2 + \|\nabla \log \pi(\theta)\|^2 \lesssim V(\theta). \tag{A.5}$$

Notice that for \hat{f} based on subsampling we obtain that $p_H = \infty$ if

$$\|\nabla \log \pi(x|\theta)^2\| \leq C(x)V(\theta).$$

Notice that Assumption 12 also implies

$$\begin{aligned} \mathbb{E} \|\theta_{k+1} - \theta_k\|^{2p} &\leq V_k^p \\ \mathbb{E} \|\hat{f}_k\|^{2p} &\leq V_k^p \end{aligned} \tag{A.6}$$

if for $p \leq p_H$. Equation (65) could now simply be formulated as an additional assumption, however currently we need even stronger assumptions to verify Equation (62). Subsequently, we show how the results of Pardoux and Veretennikov (2001) can be used to establish Equation (62) if Equations (A.3) and (A.5) hold for $V = \|\theta\|^2 + 1$.

Theorem 1 and 2 of Pardoux and Veretennikov (2001) characterise the smoothness and growth of the solution to the Poisson equation associated with Equation (60). This is important for our results because the key ingredient for the proof of Theorem 9 is

$$\sup_k \mathbb{E} \left\| \psi^{((i))}(\theta_k) \right\| < \infty \text{ for } k = 1, \dots, 4.$$

Because we have already established in Section 4 that

$$\sup_i \mathbb{E} V^p(\theta_i) < \infty \text{ it is left to verify}$$

$$\left\| \psi^{(k)} \right\| \lesssim V^{p_{\psi,k}}, \quad \text{for } k = 0, \dots, 4 \quad (62)$$

The assumptions needed to apply the Theorem 9 results are

$$\left\langle f(\theta), \frac{\theta}{\|\theta\|} \right\rangle \leq -r \|\theta\|, \quad \|\theta\| \geq M_0 \quad (A.7)$$

$$0 < \lambda_- \leq \left\langle g(\theta) g^*(\theta) \frac{\theta}{\|\theta\|}, \frac{\theta}{\|\theta\|} \right\rangle \leq \lambda_+ < \infty. \quad (A.8)$$

This holds if Assumption 12 is satisfied with $V(\theta) = \|\theta\|^2 + 1$.

Theorem 13 *Pardoux and Veretennikov (2001) Let $\bar{f} = 0$ and Equations (A.7) and (A.8) are satisfied. Then there exists a solution $\psi \in W_{loc}^2$ to the Poisson equation*

$$\mathcal{L}\psi = \phi - \bar{\phi}$$

1. *If there is a C such that*

$$|\phi(\theta)| \leq C(1 + \|\theta\|)^\beta$$

for some $\beta < 0$, then ψ is bounded. Moreover,

$$\sup_{\theta} |\psi(\theta)| \leq C \sup_{\theta} |f| (1 + \|\theta\|)^{-\beta}$$

and

$$\|\nabla\psi\| \leq C$$

2. *if there exists a constant C and some $\beta > 0$ such that*

$$|\phi(\theta)| \leq C(1 + \|\theta\|)^\beta$$

then there exist a constant such that

$$|\psi(\theta)| \leq C'(1 + \|\theta\|)^\beta.$$

Finally there exists C such that

$$\|\nabla\psi\| \leq C(1 + \|\theta\|^\beta).$$

Remark 14 *We believe that assumption Theorem 13 can be weakened to be of the form of Equation (A.3) but it is out of the scope of this article to explore this direction.*

In order to iterate Theorem 13 we note that the derivatives ψ can be expressed as solution to Poisson equations with different RHSs.

$$\mathcal{L}\psi = \phi - \bar{\phi} \tag{A.9}$$

$$A\partial_i\psi = \partial_i\phi - \frac{1}{2}\nabla\psi \cdot \partial_i f \tag{A.10}$$

$$A\partial_{ij}\psi = \partial_{ij}\phi - \frac{1}{2}\nabla\partial_j\psi \cdot \partial_i f - \frac{1}{2}\nabla\psi \cdot \partial_{ij}f - \frac{1}{2}\nabla\psi \cdot \partial_j f \tag{A.11}$$

$$\begin{aligned} A\partial_{ijk}\psi &= \partial_{ijk}\phi - \frac{1}{2}\nabla\partial_{jk}\psi \cdot \partial_i f - \frac{1}{2}\nabla\partial_j\psi \cdot \partial_{ik}f - \frac{1}{2}\nabla\partial_k\psi \cdot \partial_{ij}f - \frac{1}{2}\nabla\psi \cdot \partial_{ijk}f \\ &\quad - \frac{1}{2}\nabla\partial_k\psi \cdot \partial_{jk}f - \frac{1}{2}\nabla\psi \cdot \partial_{jk}f - \frac{1}{2}\nabla\psi \cdot \partial_k f \end{aligned} \tag{A.12}$$

We will denote by $\beta_{\psi,i}$ numbers that satisfy

$$\sup_{|\alpha|=i} \|\partial^\alpha\psi\| \lesssim \left(1 + |\theta|^{\beta_{\psi,i}}\right) \tag{A.13}$$

where we used multi-index notation for derivatives. We use a similar notation for the derivatives of f , that is $\beta_{f,i}$ and assume that these bounds are a priori given.

Using Theorem 13 we can obtain $p_{\psi,i}$ to satisfy Equation (62) in terms of the β 's which we formulate as the following lemma.

Lemma 15 *Suppose that ϕ and its derivatives are bounded and Assumption 12 and Equations Equations (A.7) and (A.8) hold. Then the choice*

$$\begin{aligned} p_{\psi,0} &= 0 \\ p_{\psi,1} &= 0 \\ p_{\psi,2} &= \frac{\beta_{f,1}}{2} \\ p_{\psi,3} &= \beta_{f,1} \vee \frac{\beta_{f,2}}{2} \\ p_{\psi,4} &= \frac{1}{2} (3\beta_{f,1} \vee (\beta_{f,1} + \beta_{f,2}) \vee \beta_{f,3}) \end{aligned}$$

satisfies Equation (62).

Proof Assumption 12 yields that $\beta_{f,0} = 1$ is a valid choice. Applying Theorem 13 to Equation (A.9) implies that $\beta_{\psi,0} := \beta_{\psi,1} := 0$ satisfies Equation (A.13). Applying Theorem 13 to Equation (A.10) yields that

$$\beta_{\psi,2} \leq \beta_{f,1}.$$

Applying Theorem 13 to Equation (A.11) yields that

$$\beta_{\psi,3} \leq 2\beta_{f,1} \vee \beta_{f,2}.$$

Applying Theorem 13 to Equation (A.12) yields that

$$\begin{aligned}\beta_{\psi,4} &\leq (\beta_{f,1} + (2\beta_{f,1} \vee \beta_{f,2\cdot})) \vee (\beta_{f,1} + \beta_{f,2\cdot}) \vee \beta_{f,3} \\ &\leq 3\beta_{f,1} \vee (\beta_{f,1} + \beta_{f,2\cdot}) \vee \beta_{f,3}.\end{aligned}$$

Thus we have established Equation (62). \blacksquare

10. Analytic expressions for the Gaussian Toy Model

We sketch the derivation of the analytic expression of the MSE are used for the plots in Section 6.

10.1 Expected MSE for fixed Data Sets

We outline how an analytic expression for the MSE of the sample average can be derived. The following method generalises to any polynomial test function but we concentrate on the sample average for the second moment of the posterior given by $S_2 = \frac{1}{M} \sum_{j=0}^{M-1} \theta_j^2$. Its MSE can be expressed using Equation (12)

$$MSE = \mathbb{E} \left(\frac{1}{M} \sum_{k=0}^{M-1} \theta_j^2 - (\mu_p^2 + \sigma_p^2) \right)^2 = \mathbb{E} S_2^2 - 2\mathbb{E} S_2 (\mu_p^2 + \sigma_p^2) + (\mu_p^2 + \sigma_p^2)^2. \quad (\text{C.1})$$

In order to express $\mathbb{E} S_2^2$ in Equation (C.1) we derive the recurrence equations for $\mathbb{E} \theta_j^i$ for $i = 1, \dots, 4$ by taking the expectations of

$$\begin{aligned}\theta_{j+1} &= (1 - Ah)\theta_j + hB_j + \sqrt{h}\eta_j \\ \theta_{j+1}^2 &= (1 - Ah)^2\theta_j^2 + h^2B_j^2 + h\eta_j^2 + 2(1 - Ah)\theta_j hB_j + 2(1 - Ah)\theta_j \sqrt{h}\eta_j + 2hB_j \sqrt{h}\eta_j\end{aligned} \quad (\text{C.2})$$

$$\begin{aligned}\theta_{j+1}^3 &= \left((1 - Ah)\theta_j + hB_j + \sqrt{h}\eta_j \right)^3 \\ &= (1 - Ah)^3\theta_j^3 + 3(1 - Ah)^2\theta_j^2 hB_j + 3(1 - Ah)\theta_j h^2B_j^2 + h^3B_j^3 \\ &\quad + 3\sqrt{h}\eta(\dots) + 3\eta^2 h((1 - Ah)\theta_j + hB_j) + \eta^3 h^{\frac{3}{2}}\end{aligned}$$

$$\theta_{j+1}^4 = (1 - Ah)^4\theta_j^4 + 4(1 - Ah)^3\theta_j^3 hB_j + 6(1 - Ah)^2\theta_j^2 h^2B_j^2 \quad (\text{C.3})$$

$$+ 4(1 - Ah)\theta_j h^3B_j^3 + h^4B_j^4 + 4\eta(\dots) + 4\eta^3(\dots) \quad (\text{C.4})$$

$$+ 6h\eta^2((1 - Ah)^2\theta_j^2 + 2(1 - Ah)\theta_j hB_j + h^2B_j^2) + \eta^4 h^2. \quad (\text{C.5})$$

The recurrent equation for $\mathbb{E} \theta_j$ is linear and first order and can therefore be solved explicitly. Plugging the result into the equation for $\mathbb{E} \theta_j^2$ turns it into a first order linear equation as well. Repeating this processes yields explicit expressions for $\mathbb{E} \theta_j^i$ $i = 1, \dots, 4$. The sums can be carried out explicitly because the terms are of the form of a geometric sum or a geometric term with a polynomial factor. This allows us to obtain an analytic expression for $\mathbb{E} S_2^2$ by reducing it to $\mathbb{E} \theta_j^i$ as follows

$$\mathbb{E} S_2^2 = \frac{1}{M^2} \mathbb{E} \left(\sum_{i=0}^{M-1} \theta_i^4 + 2 \sum_{i=0}^{M-1} \theta_i^2 \sum_{j=i+1}^{M-1} \theta_j^2 \right). \quad (\text{C.6})$$

The cross terms can be removed using Equation (C.2) so that

$$\begin{aligned}
 \theta_j^2 &= (1 - Ah)^{2(j-i)} \theta_i^2 + \sum_{k=0}^{j-1-i} (1 - Ah)^{2k} \\
 &\quad \left[h^2 B_{j-1-k}^2 + h^2 \eta_{j-1-k}^2 + 2(1 - Ah) B_{j-1-k} h \theta_{j-1-k} + \right. \\
 &\quad \left. + 2(1 - Ah) \eta_{j-1-k} \theta_{j-1-k} \sqrt{h} + 2h^{\frac{3}{2}} B_{j-1-k} \eta_{j-1-k} \right].
 \end{aligned} \tag{C.7}$$

Plugging this into Equation (C.6) yields

$$\begin{aligned}
 \mathbb{E} S_2^2 &= \mathbb{E} \frac{1}{M^2} \sum_{i=0}^{M-1} \theta_i^4 \left(1 + 2 \sum_{j=i+1}^{M-1} (1 - Ah)^{2(j-i)} \right) \\
 &\quad + \mathbb{E} \frac{1}{M^2} \sum_{i=0}^{M-1} \theta_i^2 \sum_{j=i+1}^{M-1} \sum_{k=0}^{j-1-i} (1 - Ah)^{2k} \\
 &\quad \left[h^2 B_{j-1-k}^2 + h \eta_{j-1-k}^2 + 2(1 - Ah) B_{j-1-k} h \theta_{j-1-k} \right. \\
 &\quad \left. + 2(1 - Ah) \eta_{j-1-k} \sqrt{h} + 2h^{\frac{3}{2}} B_{j-1-k} \eta_{j-1-k} \right].
 \end{aligned}$$

Using the recurrence Equation to express θ_{j-1-k} in terms of θ_i we conclude that $\mathbb{E} S_2^2$ is equal to

$$\begin{aligned}
 &\frac{1}{M^2} \sum_{i=0}^{M-1} \mathbb{E} \theta_i^4 \left(1 + 2 \sum_{j=i+1}^{M-1} (1 - Ah)^{j-i} \right) \\
 &+ \frac{1}{M^2} \sum_{i=0}^{M-1} \mathbb{E} \theta_i^2 \sum_{j=i+1}^{M-1} \sum_{k=0}^{j-1-i} (1 - Ah)^{2k} [h^2 \mathbb{E} B^2 + h]. \\
 &+ \mathbb{E} \frac{1}{M^2} \sum_{i=0}^{M-1} \theta_i^2 \sum_{j=i+1}^{M-1} \sum_{k=0}^{j-1-i} (1 - Ah)^{2k} 2(1 - Ah) B_{j-1-k} h \\
 &\quad \left((1 - Ah)^{(j-1-k)-i} \theta_i + \sum_{l=0}^{(j-1-k)-i-1} (1 - Ah)^l \left(h B_{j-1-k-l-1} + \sqrt{h} \eta_{j-1-k-l-1} \right) \right)
 \end{aligned}$$

Taking the expectations into the sum yields

$$\begin{aligned}
 \mathbb{E}S_2^2 &= \frac{1}{M^2} \sum_{i=0}^{M-1} \mathbb{E}\theta_i^4 \left(1 + 2 \sum_{j=i+1}^{M-1} (1 - Ah)^{j-i} \right) \\
 &+ \frac{1}{M^2} \sum_{i=0}^{M-1} \mathbb{E}\theta_i^2 \sum_{j=i+1}^{M-1} \sum_{k=0}^{j-1-i} (1 - Ah)^{2k} [h^2 \mathbb{E}B^2 + h]. \\
 &+ \mathbb{E} \frac{1}{M^2} \sum_{i=0}^{M-1} \mathbb{E}\theta_i^3 \sum_{j=i+1}^{M-1} \sum_{k=0}^{j-1-i} (1 - Ah)^{2k} 2(1 - Ah) \mathbb{E}Bh(1 - Ah)^{(j-1-k)-i} \\
 &+ \frac{1}{M^2} \sum_{i=0}^{M-1} \mathbb{E}\theta_i^2 \sum_{j=i+1}^{M-1} \sum_{k=0}^{j-1-i} (1 - Ah)^{2k} 2(1 - Ah) \mathbb{E}Bh \sum_{l=0}^{(j-1-k)-i-1} (1 - Ah)^l h \mathbb{E}B.
 \end{aligned}$$

We have an expressions for $\mathbb{E}B$ but in the following we derive the expressions for $\mathbb{E}B^2$ required to express $\mathbb{E}S_2^2$. The terms $\mathbb{E}B^3$ and $\mathbb{E}B^4$ are needed for the derivation of $\mathbb{E}\theta_j^4$. In order to derive expressions for $\mathbb{E}B^p$, we introduce the power sums

$$p_k = \sum_{i=1}^N X_i^k$$

and the elementary symmetric polynomials

$$e_0 = 1, e_1 = \sum_{i=1}^N X_i, e_2 = \sum_{1 \leq i < j \leq N} X_i X_j, \dots, e_N = \prod_{i=1}^N X_i. \quad (\text{C.8})$$

Computing e_i naively has complexity of order $\mathcal{O}(N^i)$ for $i \ll N$. Using Newton's identities

$$e_1 = p_1, e_2 = \frac{1}{2}(e_1 p_1 - p_2), e_3 = \frac{1}{3}(-e_1 p_2 + e_2 p_1 + p_3), e_4 = \frac{1}{4}(e_1 p_3 - e_2 p_2 + e_3 p_1 - p_4)$$

e_i can be expressed in terms of p_k , $k \leq i$ each of which can be computed with complexity of order $\mathcal{O}(N)$.

We consider the term $B = \frac{N}{n} \frac{\sum_{i=1}^n X_{\tau_i}}{2\sigma_x^2}$ where τ_i are sampled with replacement from a fixed data set $\{1, \dots, N\}$. The second moment can be calculated as follows

$$\begin{aligned}
 \mathbb{E}B^2 &= \left(\frac{N}{n2\sigma_x^2} \right)^2 \sum_{i,j} \mathbb{E}X_{\tau_i} X_{\tau_j} \\
 &= \left(\frac{N}{n2\sigma_x^2} \right)^2 \left(n(n-1) \underbrace{\mathbb{E}X_{\tau_1} X_{\tau_2}}_{\text{Mom}_{2,1}} + n \underbrace{\mathbb{E}X_{\tau_1} X_{\tau_1}}_{\text{Mom}_{2,2}} \right).
 \end{aligned}$$

We use Newton's identities to express $\text{Mom}_{2,1}$ and $\text{Mom}_{2,2}$

$$\text{Mom}_{2,1} = \frac{2e_2}{N(N-1)} \quad \text{Mom}_{2,2} = \frac{p_2}{N}.$$

Similarly, we obtain

$$\begin{aligned}
 \mathbb{E}B^3 &= \left(\frac{N}{n2\sigma_x^2}\right)^3 \sum_{i,j,k} \mathbb{E}X_{\tau_i}X_{\tau_j}X_{\tau_k} \\
 &= \left(\frac{N}{n2\sigma_x^2}\right)^3 \left(n(n-1)(n-2) \underbrace{\mathbb{E}X_{\tau_1}X_{\tau_2}X_{\tau_3}}_{\text{Mom}_{3,1}} + 3n(n-1) \underbrace{\mathbb{E}X_{\tau_1}X_{\tau_2}^2}_{\text{Mom}_{3,2}} + n \underbrace{\mathbb{E}X_{\tau_1}^3}_{\text{Mom}_{3,3}} \right) \\
 \text{Mom}_{3,1} &= \frac{\sum_{i \neq j \neq l} X_i X_j X_l}{N(N-1)(N-2)} = \frac{6e_3}{N(N-1)(N-2)} \\
 \text{Mom}_{3,2} &= \frac{p_1 p_2 - p_3}{N(N-1)}, \quad \text{Mom}_{3,3} = \frac{p_3}{N}.
 \end{aligned}$$

A similar calculation yields a representation of $\mathbb{E}B^4$ in terms of p_1, \dots, p_4 .

10.2 Expected MSE for Random Data

We sketch the derivation of the MSE

$$\mathbb{E}_{\theta, X} \left(\frac{1}{M} \sum_{k=0}^{M-1} \theta_j^2 - (\mu_p^2 + \sigma_p^2) \right)^2$$

where we take expectation with respect to $X_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\theta^\dagger, \sigma_X^2)$ for $i = 1, \dots, N$ and the randomness in the recursion for θ_j . We obtain an analytic expression for the MSE by deriving expressions for $\mathbb{E}_{X, \theta} \theta_j^p$. We illustrate the computation for $p = 2$, noting that we assume $\theta_0 = 0$ a.s.. We know that

$$\begin{aligned}
 \theta_j^2 &= \sum_{k=0}^{j-1} (1-Ah)^{2k} \left[h^2 B_{j-1-k}^2 + h \eta_{j-1-k}^2 + 2(1-Ah) \eta_{j-1-k} \theta_{j-1-k} \sqrt{h} \right. \\
 &\quad \left. + 2h^{\frac{3}{2}} B_{j-1-k} \eta_{j-1-k} + 2(1-Ah) B_{j-1-k} h \theta_{j-1-k} \right] \tag{C.9} \\
 &= \sum_{k=0}^{j-1} (1-Ah)^{2k} \left[h^2 B_{j-1-k}^2 + 2(1-Ah) \eta_{j-1-k} \theta_{j-1-k} \sqrt{h} + 2h^{\frac{3}{2}} B_{j-1-k} \eta_{j-1-k} \right. \\
 &\quad \left. h \eta_{j-1-k}^2 + 2(1-Ah) B_{j-1-k} h \sum_{l=0}^{j-1-k-1} (1-Ah)^l \left(h B_{j-k-2-l} + \sqrt{h} \eta_{j-k-2-l} \right) \right].
 \end{aligned}$$

The expectation $\mathbb{E}_{X, \theta} \theta_j$ therefore boils down to calculating $\mathbb{E}B^2$, $\mathbb{E}B B'$ and $\mathbb{E}B$ where B and B' are independent samples of Equation (15). We start by calculating

$$\begin{aligned}
\mathbb{E}BB' &= \frac{N^2}{n^2 4\sigma_x^4} \sum_{i=1}^n \sum_{j=1}^n \mathbb{E}X_{\tau_i} X_{\tilde{\tau}_j} \\
&= \frac{N^2}{n^2 4\sigma_x^4} \sum_{i=1}^n \sum_{j=1}^n \left(\frac{1}{N} \mathbb{E}X^2 + \frac{N-1}{N} \mathbb{E}X\tilde{X} \right) \\
&= \frac{N^2}{n^2 4\sigma_x^4} \left(\sum_{i=1}^n \sum_{j=1}^n \frac{1}{N} (\theta^{\dagger 2} + \sigma_X^2) + \frac{N-1}{N} \theta^{\dagger 2} \right).
\end{aligned}$$

Similarly, we obtain

$$\mathbb{E}B^2 = \frac{N^2}{n^2 4\sigma_x^4} \frac{n(\theta^{\dagger 2} + \sigma_\theta^2) + n(n-1)\theta^{\dagger 2}}{1}.$$

Deriving $\mathbb{E}\theta_j^p$ for $p = 1, 2, 3, 4$ requires the calculation of $\mathbb{E}B_1^{\alpha_1} B_2^{\alpha_2} B_3^{\alpha_3} B_4^{\alpha_4}$ where B_i are i.i.d. following the distribution of Equation (15) for $\alpha_i \geq 0$ and $\sum_i \alpha_i \leq 4$. The arguments so far allow us to derive T_1 and T_3 in

$$\begin{aligned}
MSE &= \mathbb{E} \left(S_2^2 - 2S_2(\mu_p^2 + \sigma_p^2) + (\mu_p^2 + \sigma_p^2)^2 \right) \\
&= \underbrace{\mathbb{E}S_2^2}_{T_1} - 2 \underbrace{\mathbb{E}S_2\mu_p^2}_{T_2} - 2 \underbrace{\mathbb{E}S_2\sigma_p^2}_{T_3} + \underbrace{\mathbb{E}(\mu_p^4 + 2\mu_p^2\sigma_p^2 + \sigma_p^4)}_{T_4}.
\end{aligned}$$

Recall that the posterior for this toy model is given by

$$\mathcal{N}(\mu_p, \sigma_p^2) = \mathcal{N} \left(\frac{\sum_{i=1}^N X_i}{\frac{\sigma_x^2}{\sigma_\theta^2} + N}, \left(\frac{1}{\sigma_\theta^2} + \frac{N}{\sigma_x^2} \right)^{-1} \right)$$

and hence T_4 can be computed explicitly. The summands of T_2 can be derived similarly to Equation (C.9) in terms of the quantities $\mathbb{E}B\mu_p^2$, $\mathbb{E}B^2\mu_p^2$ and $\mathbb{E}BB'\mu_p^2$. The explicit expression can be obtained from the supplemented Mathematica[®] file.

References

- A. Abdulle, D. Cohen, G. Vilmart, and K. C. Zygalakis. High order weak methods for stochastic differential equations based on modified equations. *SIAM J. Sci. Comput.*, 34(3):1800–1823, 2012.
- A. Abdulle, G. Vilmart, and K. Zygalakis. High order numerical approximation of the invariant measure of ergodic sdes. *SIAM Journal on Numerical Analysis*, 52(4):1600–1622, 2014.
- Steve Brooks, Andrew Gelman, Galin Jones, and Xiao-Li Meng. *Handbook of Markov Chain Monte Carlo*. Chapman and Hall/CRC, 2011.
- T. Chen, E. Fox, and C. Guestrin. Stochastic gradient hamiltonian monte carlo. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 1683–1691, 2014.

- A. Debussche and E. Faou. Weak backward error analysis for SDEs. *SIAM J. Numer. Anal.*, 50(3):1735–1752, 2012. ISSN 0036-1429.
- N. Ding, Y. Fang, R. Babbush, C. Chen, R. Skeel, and H. Neven. Bayesian sampling using stochastic gradient thermostats. *Advances in Neural Information Processing Systems (NIPS), Montreal, Canada*, 2014.
- Rafail Hasminskii. *Stochastic stability of differential equations*. Sijthoff and Noordhoff, 1980.
- M. Kopec. Weak backward error analysis for overdamped langevin processes. *IMA Journal of Numerical Analysis*, 2014.
- M. Kopec. Weak backward error analysis for langevin process. *BIT Numerical Mathematics*, pages 1–47, 2015.
- D. Lamberton and G. Pages. Recursive computation of the invariant distribution of a diffusion. *Bernoulli*, 8(3):367–405, 2002.
- B. Leimkuhler and C. Matthews. Rational construction of stochastic numerical methods for molecular sampling. *Applied Mathematics Research eXpress*, 2013(1):34–56, 2013.
- J. C. Mattingly, A. M. Stuart, and M. V. Tretyakov. Convergence of Numerical Time-Averaging and Stationary Measures via Poisson Equations. *SIAM J. Numer. Anal.*, 48(2):552–577, 2010. ISSN 0036-1429.
- G. N. Milstein and M. V. Tretyakov. Computing ergodic limits for Langevin equations. *Phys. D*, 229(1):81–95, 2007. ISSN 0167-2789.
- G.N. Milstein. Weak approximation of solutions of systems of stochastic differential equations. *Theory Probab. Appl.*, 30(4):750–766, 1986. ISSN 0040585X.
- G..N Milstein and M.V. Tretyakov. *Stochastic numerics for mathematical physics*. Scientific Computing. Springer-Verlag, Berlin and New York, 2004.
- B-R Nawaf and E. Vanden-Eijnden. Pathwise accuracy and ergodicity of metropolized integrators for SDEs. *Communications on Pure and Applied Mathematics*, 63(5):655–696, 2010.
- R. M. Neal. Markov chain monte carlo using hamiltonian dynamics. In *Handbook of Markov Chain Monte Carlo*. Chapman & Hall/CRC, 2011.
- E. Pardoux and A. Yu. Veretennikov. On the Poisson equation and diffusion approximation. I. *The Annals of Probability*, 29(3):1061–1085, 2001.
- S. Patterson and Y. W. Teh. Stochastic gradient Riemannian Langevin dynamics on the probability simplex. In *Advances in Neural Information Processing Systems*, 2013.
- G. O. Roberts and R. L. Tweedie. Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli*, 2(4):pp. 341–363, 1996.

- A. Koratticara S. Ahn and M. Welling. Bayesian posterior sampling via stochastic gradient fisher scoring. In *ICML*, 2012.
- I. Sato and H. Nakagawa. Approximation analysis of stochastic gradient langevin dynamics by using fokker-planck equation and ito process. In Tony Jebara and Eric P. Xing, editors, *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 982–990. JMLR Workshop and Conference Proceedings, 2014.
- D. Talay and L. Tubaro. Expansion of the global error for numerical schemes solving stochastic differential equations. *Stochastic Anal. Appl.*, 8(4):483–509 (1991), 1990. ISSN 0736-2994.
- Y. W. Teh, A. H. Thiéry, and S. J. Vollmer. Consistency and fluctuations for stochastic gradient Langevin dynamics. *ArXiv e-prints*, 2014. Accepted by J. Mach. Learn. Res.
- M. Welling and Y. W. Teh. Bayesian learning via Stochastic Gradient Langevin Dynamics. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 681–688, 2011.
- K. C. Zygalakis. On the existence and applications of modified equations for stochastic differential equations. *SIAM J. Sci. Comput.*, 33:102–130, 2011.