

# Particle Metropolis-adjusted Langevin algorithms

Christopher Nemeth, Chris Sherlock and Paul Fearnhead

Department of Mathematics and Statistics, Lancaster University, Lancaster  
LA1 4YF U.K.

c.nemeth@lancaster.ac.uk, c.sherlock@lancaster.ac.uk,  
p.fearnhead@lancaster.ac.uk

May 27, 2016

## Abstract

This paper proposes a new sampling scheme based on Langevin dynamics that is applicable within pseudo-marginal and particle Markov chain Monte Carlo algorithms. We investigate this algorithm's theoretical properties under standard asymptotics, which correspond to an increasing dimension of the parameters,  $n$ . Our results show that the behaviour of the algorithm depends crucially on how accurately one can estimate the gradient of the log target density. If the error in the estimate of the gradient is not sufficiently controlled as dimension increases, then asymptotically there will be no advantage over the simpler random-walk algorithm. However, if the error is sufficiently well-behaved, then the optimal scaling of this algorithm will be  $O(n^{-1/6})$  compared to  $O(n^{-1/2})$  for the random walk. Our theory also gives guidelines on how to tune the number of Monte Carlo samples in the likelihood estimate and the proposal step-size.

**Keywords:** Metropolis-adjusted Langevin algorithm; Optimal scaling; Particle Filter; Particle Markov chain Monte Carlo; Pseudo-marginal Markov chain Monte Carlo.

## 1 Introduction

Markov chain Monte Carlo algorithms are a popular and well-studied methodology that can be used to draw samples from posterior distributions. Over the past few years these algorithms have been extended to tackle problems where the model likelihood is intractable (Beaumont, 2003). Andrieu and Roberts (2009) showed that within the Metropolis–Hastings algorithm, if the likelihood is replaced with an unbiased estimate, then the sampler still targets the correct stationary distribution. Andrieu et al. (2010) extended this work further to create a class of

Markov chain algorithms that use sequential Monte Carlo methods, also known as particle filters.

Current implementations of pseudo-marginal and particle Markov chain Monte Carlo use random-walk proposals to update the parameters (e.g., Golightly and Wilkinson, 2011; Knappe and de Valpine, 2012) and shall be referred to herein as particle random-walk Metropolis algorithms. Random walk-based algorithms propose a new value from some symmetric density centred on the current value. This density is not informed by the local properties of the posterior; however, we can often obtain further information about such properties as we obtain our Monte Carlo estimate of the posterior density, and at little or no additional computational overhead. It is therefore natural to consider whether we can use this information to make better proposals for the parameters. In this paper we focus on using Monte Carlo methods to estimate the gradient of the log posterior density, and then use this to guide the proposed parameters towards regions of higher posterior probability. This results in a Monte Carlo version of the Metropolis-adjusted Langevin algorithm (Roberts and Rosenthal, 1998), which we refer to herein as the particle Langevin algorithm.

When the likelihood is tractable, the Metropolis-adjusted Langevin algorithm has better theoretical properties than the random-walk Metropolis algorithm. The mixing properties of these algorithms have been studied in the asymptotic limit as the dimension of the parameters,  $n$ , increases. In this asymptotic regime, the optimal proposal step-size scales as  $n^{-1/2}$  for the random-walk algorithm, but as  $n^{-1/6}$  for the Metropolis-adjusted Langevin algorithm; and the optimal asymptotic acceptance rate is higher; see Roberts et al. (1997), Roberts and Rosenthal (1998) and Roberts and Rosenthal (2001) for more details. It is natural to ask whether these advantages of the Metropolis-adjusted Langevin algorithm over the random-walk algorithm extend to pseudo-marginal and particle Markov chain Monte Carlo algorithms, and, in particular, how they are affected when only noisy estimates of the gradient of the log posterior density are available.

We investigate the asymptotic properties of the particle Langevin algorithm and show that its behaviour depends crucially on the accuracy of the estimate of the gradient of the log posterior density as  $n$  increases. If the error in the estimate of a component of the gradient does not decay with  $n$ , then there will be no benefit over the particle random-walk algorithm. If the error is sufficiently well-behaved, then we find that the particle Langevin algorithm inherits the same asymptotic characteristics as the Metropolis-adjusted Langevin algorithm. The optimal proposal scales as  $n^{-1/6}$ , rather than  $n^{-1/2}$ , and there is a higher optimal acceptance rate. In this well-behaved regime we find that the number of particles should be chosen so that the variance in the estimate of the log posterior density is approximately 3.

Furthermore, we provide explicit guidance for tuning the scaling of the particle Langevin algorithm by aiming for a particular acceptance rate. We show that the optimal acceptance rate depends crucially on how accurately we estimate the log posterior density, a feature that is common to other particle Markov chain Monte Carlo algorithms. As such, tuning the particle

Langevin algorithm using the acceptance rate is only appropriate if we have an estimate of the variance of our estimator of the log posterior density. Additionally, the optimal acceptance rate depends on the accuracy of the gradient estimate. We propose a criterion for choosing an appropriate scaling for the proposal given a fixed but arbitrary number of particles. We provide an acceptance rate to tune to, which is a function of the variance in the log posterior density estimate. This acceptance rate is robust to the unknown accuracy of our estimate of the gradient. Tuning to it will lead to an efficiency of at least 90% of the efficiency of the optimally-scaled particle Langevin algorithm, with the same number of particles and known accuracy of the gradient estimate. Under this criterion, and with sufficient particles so that the variance of the estimate of the log posterior density is approximately 3, we should scale the step-size so that the acceptance rate is 11%.

## 2 Efficient Markov chain Monte Carlo with intractable likelihoods

Let  $p(z | x)$  be a model likelihood, with data  $z \in \mathcal{Z} \subseteq \mathbb{R}^{n_z}$  and model parameters  $x \in \mathcal{X} \subseteq \mathbb{R}^n$ . Using Bayes' theorem, the posterior density over the parameters, up to a constant of proportionality, is  $\pi(x) \propto p(z | x)p(x)$ , where  $p(x)$  is a prior density for  $x$ .

Markov chain Monte Carlo algorithms draw samples,  $(x_1, \dots, x_J)$ , from the posterior distribution. Typically, these samples are generated using the Metropolis–Hastings algorithm, where proposed values  $y$  are sampled from a proposal distribution  $q(\cdot | x)$  and accepted with probability

$$\alpha(y | x) = \min \left\{ 1, \frac{p(z | y)p(y)q(x | y)}{p(z | x)p(x)q(y | x)} \right\}. \quad (1)$$

The Metropolis–Hastings algorithm requires that the likelihood  $p(z | x)$  be tractable, but there are many situations where it can only be evaluated approximately. Andrieu and Roberts (2009) and Andrieu et al. (2010) have shown that the Metropolis–Hastings algorithm is still valid in this setting, provided there is a mechanism for simulating unbiased, non-negative estimates of the likelihood. This technique is known as pseudo-marginal Markov chain Monte Carlo.

The pseudo-marginal approach presupposes that a non-negative unbiased estimator  $\hat{p}(z | x, \mathcal{U}_x)$  of  $p(z | x)$  is available, where  $\mathcal{U}_x \sim p(\cdot | x)$  denotes the random variables used in the sampling mechanism to generate an estimate of the likelihood. We then define a target density on the joint space  $(x, \mathcal{U}_x)$  as,

$$\hat{\pi}(x, \mathcal{U}_x) \propto \hat{p}(z | x, \mathcal{U}_x)p(\mathcal{U}_x | x)p(x). \quad (2)$$

Since the estimate is unbiased, the marginal density of  $x$  is

$$\int \hat{\pi}(x, \mathcal{U}_x) d\mathcal{U}_x \propto \int \hat{p}(z | x, \mathcal{U}_x) p(\mathcal{U}_x | x) p(x) d\mathcal{U}_x = p(z | x) p(x),$$

which is the posterior density of interest.

A valid Markov chain Monte Carlo algorithm targeting (2), with proposal  $q(y | x) p(\mathcal{U}_y | y)$ , has acceptance probability of the form (1), but with the intractable likelihoods,  $p(z | x)$  and  $p(z | y)$  replaced with realizations from their unbiased estimators,  $\hat{p}(z | x, \mathcal{U}_x)$  and  $\hat{p}(z | y, \mathcal{U}_y)$ .

The efficiency of the Metropolis–Hastings algorithm is highly dependent on the choice of proposal distribution  $q(y | x)$ . Ideally, the proposal would use local information about the posterior to sample from areas of higher posterior density. One such proposal is the Metropolis-adjusted Langevin algorithm (Roberts and Rosenthal, 1998) which incorporates the gradient of the log posterior density,  $\nabla \log \pi(x)$ , within the proposal. The asymptotic behaviour of this algorithm, as the number of parameters,  $n$ , increases, gives an optimal step-size of  $O(n^{-1/6})$  (Roberts and Rosenthal, 1998) compared to  $O(n^{-1/2})$  for the random-walk Metropolis algorithm (Roberts et al., 1997). As a result, to maintain a reasonable acceptance rate for large  $n$ , the Metropolis-adjusted Langevin algorithm may propose larger jumps in the posterior than the random-walk Metropolis algorithm, reducing the first order auto-correlation and improving the mixing of the Markov chain.

Using the Metropolis-adjusted Langevin algorithm in the pseudo-marginal setting is challenging because if the likelihood is intractable then typically,  $\nabla \log \pi(x)$  will also be intractable. Therefore, one needs to efficiently estimate both the posterior density  $\hat{\pi}(x)$ , and its log gradient  $\hat{\nabla} \log \pi(x)$ . The resulting algorithm, which we call the particle Langevin algorithm, proposes a new parameter value  $y$  as

$$y = x + \lambda Z + \frac{\lambda^2}{2} \hat{\nabla} \log \pi(x), \quad Z \sim \mathcal{N}(0, \mathbf{I}). \quad (3)$$

It is often possible to generate a Monte Carlo estimate of the gradient of the log posterior density with little additional computational overhead, from the output of the same Monte Carlo method used to estimate the likelihood (Poyiadjis et al., 2011). The efficiency of the particle Langevin algorithm will depend on the choice of scaling parameter  $\lambda$  and the accuracy of the estimator  $\hat{\nabla} \log \pi(x)$ . In the next section we derive asymptotic results which allow us to optimally choose  $\lambda$  and which show how the efficiency of the particle Langevin algorithm depends on the accuracy of the estimator of the gradient.

## 3 Theoretical results

### 3.1 High-dimensional setting

In this section we present two key theoretical results and investigate their practical consequences. These results apply in the general pseudo-marginal setting, but the practical guidance requires specific distributional assumptions and is specific to algorithms where the estimate of the likelihood is obtained using a particle filter. For simplicity, therefore, we continue to use particle Langevin as a general term for both pseudo-marginal and particle Markov chain Monte Carlo algorithms. All proofs are presented in the Supplementary Material.

We consider an infinite sequence of targets  $\pi^n(x^n)$ ,  $n = 1, \dots$ , where  $x^n$  is an  $n$ -dimensional vector. We obtain limiting forms for the acceptance rate and expected squared jump distance,  $J_n$ , for the particle Langevin proposal. The expected squared jumping distance has been used extensively as a measure of mixing of Markov chain Monte Carlo algorithms (e.g., Beskos et al., 2009; Sherlock and Roberts, 2009; Sherlock, 2013), where maximizing it is equivalent to minimizing the first order auto-correlation of the Markov chain. The particle Metropolis-adjusted Langevin kernel itself is not a positive operator; however any kernel with a rejection probability of at least 0.5 for all possible moves is a positive operator and typically we will be tuning our algorithm to give an average acceptance probability between 0.1 and 0.15, so that most moves have a rejection probability in excess of 0.5. Moreover, in the presence of a limiting diffusion, the limiting, scaled expected squared jumping distance is the speed of the diffusion and hence precisely the right measure of efficiency. Sherlock et al. (2015) and Roberts and Rosenthal (1998) show limiting diffusions, respectively, for the particle random-walk Metropolis algorithm and the Metropolis-adjusted Langevin algorithm, suggesting the likely existence of a limiting diffusion for the particle Metropolis-adjusted Langevin kernel.

We start by considering the idealized particle Langevin algorithm where, for any given  $x^n$ , an unbiased stochastic estimate of the target density is used, with an exact gradient of the log target density,  $\nabla \log \pi^n(x^n)$ . This algorithm is unlikely to be usable in practice, but provides a useful reference point for the more general particle Langevin proposal where we assume that we have a noisy and possibly biased estimate of  $\nabla \log \pi^n(x^n)$ . Introducing the possibility of both noise and bias in the estimate allows our results to be applied to a wider range of algorithms that could be used to estimate the gradient of the log target density.

We study a target of the form

$$\pi^n(x^n) = \prod_{i=1}^n f(x_i^n), \quad (4)$$

where  $x_i^n$  denotes the  $i$ th component of an  $n$  dimensional vector  $x^n$ . We set  $g(x) = \log f(x)$

and assume that  $g(x)$  and its derivatives  $g^{(i)}(x)$  satisfy

$$|g(x)|, |g^{(i)}(x)| \leq M_0(x), \quad i = 1, \dots, 8, \quad (5)$$

where  $M_0(x)$  is some polynomial, and

$$\int_{\mathbb{R}} x^k f(x) dx < \infty, \quad k = 1, 2, 3, \dots \quad (6)$$

Our assumptions on the form of the target, (4)–(6), are the same as those in Roberts and Rosenthal (1998). In particular, for tractability, the target is assumed to have a product form. This apparently restrictive assumption is common in much of the literature on high-dimensional limit results, including Roberts et al. (1997), Roberts and Rosenthal (1998), Neal and Roberts (2006), Roberts and Rosenthal (2014) and Sherlock et al. (2015). Some of these results have been extended to more general settings (e.g., Roberts and Rosenthal, 2001; Bédard, 2007; Sherlock and Roberts, 2009; Beskos et al., 2009; Sherlock, 2013), where optimality criteria obtained using a product target have been found to hold for more general statistical applications. The results are also widely used within adaptive Markov chain Monte Carlo algorithms (e.g., Andrieu and Thoms, 2008; Roberts and Rosenthal, 2009; Särkkä et al., 2015).

We consider the additive noise in the log target density at the current and proposed values:

$$W^n = \log \hat{\pi}^n(x^n, \mathcal{U}_x^n) - \log \pi^n(x^n), \quad V^n = \log \hat{\pi}^n(y^n, \mathcal{U}_y^n) - \log \pi^n(y^n)$$

and their difference

$$B^n = V^n - W^n. \quad (7)$$

As in Pitt et al. (2012), Sherlock et al. (2015) and Doucet et al. (2015), we assume that the distributions of  $V^n$  and  $W^n$  are independent of position. This is unlikely to hold in practice, but simulations in those articles show that it can hold approximately and that guidance obtained from the resulting theory can be robust to variations with position. In the Supplementary Material we investigate and discuss this assumption for the scenarios in Section 4.

For particle Markov chain Monte Carlo, Bérard et al. (2014) examine the particle filter in the limit of a large number,  $N$ , of particles acting on a large number of observations and find that

$$V^n | x^n, y^n, w \sim \mathcal{N}\left(-\frac{1}{2}\sigma^2, \sigma^2\right), \quad (8)$$

for some fixed  $\sigma^2 \propto 1/N$ . From the definition of  $W^n$ , and directly from (2), it follows that

$$W^n \sim \mathcal{N}\left(\frac{1}{2}\sigma^2, \sigma^2\right), \quad B^n \sim \mathcal{N}(-\sigma^2, 2\sigma^2), \quad (9)$$

when the chain is at stationarity (Pitt et al., 2012).

We apply our theoretical results to this common scenario with the assumption (e.g., Pitt et al., 2012; Sherlock et al., 2015; Doucet et al., 2015) that the computational cost is proportional to  $N$  and hence inversely proportional to  $\sigma^2$ . Therefore, our measure of efficiency is, up to a constant of proportionality,

$$\text{Eff}(\ell, \sigma^2) = \sigma^2 J_n(\ell, \sigma^2), \quad (10)$$

where  $\ell$  is related to the scaling of the proposal as in (11) and Theorem 3.3.

We consider a range of levels of control for the bias and variance of the errors in the estimate of each component of the gradient. For a given level of control, we investigate the scaling that is necessary to achieve a non-degenerate limiting acceptance rate, and the behaviour of the efficiency function in that limit. A natural corollary of our analysis is that these scaling requirements, and the resulting general forms for the limiting acceptance rate and expected squared jump distance, would persist even if we were able to use an unbiased estimate of the gradient.

### 3.2 Idealized particle Langevin algorithm

In this section we consider the idealized particle Langevin algorithm, providing general limiting forms for the acceptance rate and expected squared jump distance.

Let the scaling for the proposal on the target  $\pi^n$  be

$$\lambda_n = \ell n^{-1/6}, \quad (11)$$

where  $\ell > 0$  is a tuning parameter. As mentioned earlier, in the idealized particle Langevin algorithm we make the unrealistic assumption that the gradient of the log target density may be evaluated precisely so that the  $i$ th component of the proposal is

$$Y_i^n = x_i^n + \lambda_n Z_i + \frac{1}{2} \lambda_n^2 g'(x_i^n), \quad (12)$$

with  $Z_i \sim \mathcal{N}(0, 1)$  ( $i = 1, \dots, n$ ) independent of all other sources of variation.

Let  $\alpha_n(x, w; y, v)$  be the acceptance probability for the idealized particle Langevin algorithm with current value  $(x, w)$  and proposed value  $(y, v)$ . We are interested in the expected acceptance rate and the expected squared jump distance,

$$\begin{aligned} \bar{\alpha}_n(\ell) &= E \{ \alpha_n(X^n, W^n; Y^n, V^n) \}, \\ J_n(\ell) &= E \{ \|Y^n - X^n\|^2 \alpha_n(X^n, W^n; Y^n, V^n) \}, \end{aligned}$$

where expectation is over  $X^n, W^n, Y^n, V^n$  with distributions as defined in (4), (12), (8) and (9). Our first result is as follows.




Figure 1: Contour plot of the relative efficiency  $\text{Eff}(\ell, \sigma^2)/\text{Eff}(\ell_{\text{opt}}, \sigma_{\text{opt}}^2)$  (left panel), and asymptotic acceptance rate (right panel) plotted against  $\sigma$ , where  $\ell$  is optimized for each  $\sigma$ , for the idealized particle Langevin algorithm and for the particle Langevin algorithm in asymptotic regime (3) of Theorem 3.3.

**Theorem 3.1.** *As  $n \rightarrow \infty$ , the following limits hold in probability:*

$$\bar{\alpha}_n(\ell) \rightarrow \alpha(\ell) = 2E \left\{ \Phi \left( \frac{B}{\ell^3 K} - \frac{\ell^3 K}{2} \right) \right\}, \quad n^{-2/3} J_n(\ell) \rightarrow \ell^2 \alpha(\ell),$$

where in distribution  $B = \lim_{n \rightarrow \infty} B^n$ , and  $B^n$  is defined in (7). Here,

$$K = \left[ \frac{1}{48} E \{ 5g'''(X)^2 - 3g''(X)^3 \} \right]^{1/2} \in \mathbb{R}^+, \quad (13)$$

where expectation in (13) is with respect to the density  $f(x)$ .

The following corollary details the parameters that optimize the efficiency function for the particle filter scenario.

**Corollary 3.2.** *Subject to (8) and (9), the efficiency defined in (10) is maximized when the scaling and noise variance are  $\ell_{\text{opt}} \approx 1.125K^{-1/3}$ , and  $\sigma_{\text{opt}}^2 \approx 3.038$ , at which point, the limiting acceptance rate is  $\alpha_{\text{opt}} \approx 15.47\%$ .*

The optimal variance of the noise in the log target density differs only slightly from that of the particle random-walk Metropolis algorithm, where  $\sigma_{\text{opt}}^2 \approx 3.283$  (Sherlock et al., 2015); however, the optimal asymptotic acceptance rate is increased from 7.00% to 15.47% and the scaling is improved from  $O(n^{-1/2})$  to  $O(n^{-1/6})$ . Therefore, for large  $n$ , the particle Langevin algorithm permits larger jumps leading to a more efficient proposal distribution.

Figure 1 shows the relative efficiency as a function of the scaling and the standard deviation,  $\sigma$ , of the noise, and the optimal acceptance rate as a function of  $\sigma$ . The left panel shows that over a wide range of variances the optimal scaling  $\ell$  is close to  $1.125K^{-1/3}$ , and over a wide range of scalings, the optimal variance  $\sigma^2$  is close to 3.038. This relative insensitivity between the scaling and variance means that the scaling which maximizes the expected squared jump



distance over all possible noise variances will be close to the optimal scaling for any specific noise variance in a large range. The right panel gives the acceptance rate for a range of variances, where  $\ell$  is optimally tuned for each  $\sigma^2$ . The optimal acceptance rate varies considerably over a range of sensible noise variances. This suggests that, given a sensible, but not necessarily optimal noise variance, tuning to achieve an acceptance rate of about 15% may lead to a relatively inefficient algorithm. Instead, one should either choose a scaling which optimizes the effective sample size directly, or estimate the variance in the noise in the log target density, find the acceptance rate that corresponds to the optimal scaling conditional on the estimated variance, and tune to this.

### 3.3 Scaling conditions for the particle Langevin algorithm

In the particle Langevin algorithm we do not have an exact estimate for the gradient of the log target density. In fact, depending on the approach used to estimate the gradient, the estimate may be both biased and noisy. In this section we give conditions on the bias and noise of the gradient estimate that would lead to an efficient proposal distribution.

We start by considering the  $i$ th component of the particle Langevin proposal ( $i = 1, \dots, n$ ):

$$Y_i^n = x_i^n + \lambda_n Z_i + \frac{1}{2} \lambda_n^2 \left[ g'(x_i^n) + \frac{1}{n^\kappa} \{b(x_i^n) + \tau U_{x_i^n}\} \right], \quad (14)$$

where, for all  $i$ ,  $Z_i \sim \mathcal{N}(0, 1)$  ( $i = 1, \dots, n$ ) and  $U_{x_i^n}$  are independent of each other and of all other sources of variation. For any  $x$ ,  $U_x$  is a random variable with a distribution that is independent of  $X$  and  $W$ , with  $E(U_x) = 0$ ,  $\text{var}(U_x) = 1$  and

$$E \left\{ |U_x|^k \right\} < \infty, \quad k > 0. \quad (15)$$

In the Supplementary Material, the assumption that the variance of  $U_{x_i^n}$  is constant, and that  $U_{x_i^n}$  and  $W$  are independent, are checked on the models from Section 4; the variance is shown to change by at most an order of magnitude, and independence is shown to be a good working assumption.

Even though the variance of the noise is fixed, the bias  $b(x_i^n)$  in the estimate of the  $i$ th component of the gradient (at  $x_i^n$ ) can be position specific. Furthermore, we assume that  $b(x)$  and its derivatives  $b^{(i)}(x)$  ( $i = 1, \dots, 7$ ) satisfy

$$|b(x)|, |b^{(i)}(x)| \leq M_0(x), \quad (16)$$

where  $M_0(x)$  is, without loss of generality, the same polynomial as in (5).

The particle Langevin proposal (14) can be viewed as a generalization of the Metropolis-adjusted Langevin proposal, which can be retrieved by setting  $b(x) = \tau = 0$ . The bias and noise components of (14) are scaled by an  $n^{-\kappa}$  term, where  $\kappa \geq 0$ . If  $\kappa = 0$ , as shall be shown

in Part (1) of Theorem 3.3, in order to achieve a non-degenerate limiting acceptance rate, the scaling of the proposal must be chosen so that the particle Langevin proposal has the same limiting behaviour as the particle random-walk algorithm. In addition to the definition of  $K$  in (13), we define

$$K_*^2 = E_f \{b(X)^2\} + \frac{1}{2}\tau^2, \quad (17)$$

$$K_{**} = -\frac{1}{4}E_f \{b'(X)g''(X)\}. \quad (18)$$

where, by assumptions (5), (6) and (16), these expectations are finite.

**Theorem 3.3.** *Define  $\psi(a; B) = \Phi(B/a - a/2)$ , where in distribution  $B = \lim_{n \rightarrow \infty} B^n$ , and  $B^n$  is defined in (7). As  $n \rightarrow \infty$  the following limits hold in probability:*

(1) *If  $\kappa = \frac{1}{3} - \epsilon$ , where  $0 < \epsilon \leq \frac{1}{3}$ , then  $\lambda_n = \ell n^{-1/6-\epsilon}$  for a non-degenerate limiting acceptance rate, whence*

$$\bar{\alpha}_n(\ell) \rightarrow \alpha^{(1)}(\ell) = 2E \{ \psi(\ell K_*; B) \}, \quad n^{-1+\kappa} J_n(\ell) \rightarrow \ell^2 \alpha^{(1)}(\ell).$$

(2) *If  $\kappa = \frac{1}{3}$ , then  $\lambda_n = \ell n^{-1/6}$  for a non-degenerate limiting acceptance rate, whence*

$$\bar{\alpha}_n(\ell) \rightarrow \alpha^{(2)}(\ell) = 2E \left[ \psi \left\{ (\ell^6 K^2 + 2\ell^4 K_{**} + \ell^2 K_*^2)^{1/2}; B \right\} \right], \quad n^{-2/3} J_n(\ell) \rightarrow \ell^2 \alpha^{(2)}(\ell);$$

$$\text{where } \ell^6 K^2 + 2\ell^4 K_{**} + \ell^2 K_*^2 \geq 0.$$

(3) *If  $\kappa > \frac{1}{3}$ , then  $\lambda_n = \ell n^{-1/6}$  for a non-degenerate limiting acceptance rate, whence*

$$\bar{\alpha}_n(\ell) \rightarrow \alpha^{(3)}(\ell) = 2E \{ \psi(\ell^3 K; B) \}, \quad n^{-2/3} J_n(\ell) \rightarrow \ell^2 \alpha^{(3)}(\ell).$$

The theorem highlights the relative contributions of the change in the true posterior and the error in the gradient term appearing in the Metropolis–Hastings acceptance ratio. When  $\kappa < 1/3$ , the contribution from the gradient term must be brought under control by choosing a smaller scaling, but when this smaller scaling is used, the limiting acceptance ratio for the Metropolis-adjusted Langevin algorithm is 1 and so the roughness of the target itself,  $K$ , is irrelevant. It is only at the most extreme end of regime (1), when  $\kappa = 0$ , that the expected squared jump distance is of the same order of magnitude as for the pseudo-marginal random-walk Metropolis algorithm. By contrast, when the scaling is  $\kappa > 1/3$  the effect of the errors in the gradient on the acceptance ratio is negligible; the behaviour is that of the idealized particle Langevin algorithm. The case where  $\kappa = 1/3$  gives a balance between the contributions to the acceptance ratio.

### 3.4 Tuning the particle Langevin algorithm

Theorem 3.3 has two important implications. Firstly, it provides insight into the relative performance of the particle Langevin algorithm compared to the particle random-walk algorithm. Asymptotically, the former algorithm has better mixing properties providing there is some control over the error of each component of the gradient,  $\kappa > 0$ . The greater the control, the better the scaling of the step-size as the number of parameters increases. Under our assumption on the target (4) it would be natural to expect that condition (3) of Theorem 3.3 would hold, and that the optimal scaling would be proportional to  $n^{-1/6}$ . This is because, for the particle Langevin algorithm, we need to control the variance of the estimate of the posterior density as  $n$  increases. This requires the number of particles used to estimate each component of (4) to increase linearly with  $n$  so that the Monte Carlo variance of the estimate of each term in the product (4) is of order  $1/n$ . Under this regime, the Monte Carlo error of the estimate of each component of the gradient would be of order  $n^{-1/2}$ , which corresponds to  $\kappa = 1/2$ . Empirical investigations for two models reported in Section 4, and Section 6.3 of the Supplementary Material, indicate that both fit into case (3).

The second consequence of Theorem 3.3 is the implementation guidance for the particle Langevin algorithm. In particular, results on optimal acceptance rates are important for tuning the proposal appropriately, and results on the expected squared jump distance aid in the choice of number of particles.

In the particle filter scenario, in an analogous manner to the first part of the proof of Corollary 3.2, the three acceptance rates can be shown to simplify to

$$\begin{aligned}\alpha^{(1)}(\ell, \sigma^2) &= 2\Phi \left\{ -\frac{1}{2}(\ell^2 K_*^2 + 2\sigma^2)^{1/2} \right\}, \\ \alpha^{(2)}(\ell, \sigma^2) &= 2\Phi \left\{ -\frac{1}{2}(\ell^6 K^2 + 2\ell^4 K_{**} + \ell^2 K_*^2 + 2\sigma^2)^{1/2} \right\}, \\ \alpha^{(3)}(\ell, \sigma^2) &= 2\Phi \left\{ -\frac{1}{2}(\ell^6 K^2 + 2\sigma^2)^{1/2} \right\},\end{aligned}$$

where we now make the dependence of the acceptance rates on the distribution of the noise difference,  $B$ , explicit through the parameter  $\sigma^2$ . The  $K_{**}$  (18) term appearing in the acceptance rate for case (2) can be negative, and this can lead to the counter-intuitive situation where increasing the step-size can increase the acceptance rate; see the Supplementary Material.

For regime (3) the optimal variance ( $\sigma^2 \approx 3.0$ ) and acceptance rate ( $\alpha \approx 15.5$ ) are supplied by Corollary 3.2. For regimes (1) and (2) the optimal choices will depend on the relationship between the number of particles,  $K_*$  and  $K_{**}$ , and this relationship is unknown. If  $K_*$  were fixed then the optimal variance for regime (1) would be  $\sigma^2 \approx 3.3$  as for the particle random-walk Metropolis algorithm, because the efficiency has the same general form; however it is reasonable to assume that  $K_*$  and  $K_{**}$  will decrease as  $\sigma^2$  decreases. In this case, we can always slightly increase our efficiency measure (10) by reducing  $\sigma^2$  and increasing  $\ell$  in such

a way that  $\ell^2 K_*^2 + 2\sigma^2$  is fixed. In a real, finite-dimensional, problem our limit theory is not appropriate for very large  $\ell$ . However, the above argument suggests the optimal variance will be less than 3.3. A similar argument holds in case (2), and thus we recommend choosing the number of particles so that the variance of the estimate of the log target density is roughly 3.0. Conditional on a choice of the number of particles or, equivalently, of the variance of the estimator of the log target density, it is possible to provide an acceptance rate that is close to optimal simultaneously across all three regimes. The scaling can therefore be adjusted to obtain this acceptance rate. The idea is to maximize the worst-case performance of the particle Langevin algorithm.

Fix  $\sigma^2$  and assume that the behaviour of the particle Langevin algorithm is described by one of the limiting regimes of Theorem 3.3. Denote the complete details of this regime by  $r = (\kappa, K, K_*, K_{**}) \in \mathcal{R}$ , where  $\mathcal{R}$  denotes the set of possible regimes. Given the counter-intuitive behaviour described above when  $K_{**} < 0$ , we consider only regimes with  $K_{**} \geq 0$ . Denote the asymptotic expected squared jump distance of the particle Langevin algorithm as  $J(\alpha, r)$  for regime  $r$ , where  $\ell$  is chosen to give an average acceptance probability  $\alpha$ . This is well-defined for  $0 < \alpha < 2\Phi(-\sigma/\sqrt{2})$ , as the acceptance rate is continuous and monotonically decreasing with  $\ell$ . Then, for this regime, the relative efficiency of the particle Langevin, with average acceptance rate  $\alpha$ , can be measured as

$$\text{EffR}(\alpha, r) = \frac{J(\alpha, r)}{\max_{\alpha'} J(\alpha', r)},$$

the ratio of the expected squared jump distances for this implementation of the particle Langevin algorithm and for the optimal implementation within regime  $r$ . A robust choice of average acceptance rate to tune to is the value that maximizes the minimum efficiency,

$$\arg \max_{\alpha} \min_r \text{EffR}(\alpha, r).$$

We call this the maximin acceptance rate. Calculating, for any  $\sigma$ , the corresponding maximin acceptance rate is straightforward numerically. In Figure 2 we show the maximin acceptance rate as a function of  $\sigma$  and the corresponding worst-case efficiency. The maximin choice of acceptance rate gives a worst-case relative efficiency of approximately 90% for all values of  $\sigma$ . For  $\sigma^2 \approx 3$  ( $\sigma \approx 1.73$ ) we have a maximin optimal average acceptance rate of  $\approx 11\%$ .

## 4 Inference for state-space models

### 4.1 Particle filtering

In this section we apply the particle Langevin algorithm to two state-space model examples, where, for consistency with notation, we denote  $x \in \mathcal{X}$  as a vector of model parameters and




Figure 2: Plots of the maximin implementation. The left panel shows the maximin relative efficiency (black solid line) against the 90% efficiency level (dashed line). The right panel shows the maximin optimal acceptance rates as a function of  $\sigma$  (black). For comparison we also plot the optimal acceptance rate for regime (3), idealized particle Langevin algorithm, (dotted line); and regime (1) (dashed line).

let  $\{S_t : 1 \leq t \leq T\}$  be a latent Markov process taking values on some measurable space  $\mathcal{S} \subseteq \mathbb{R}^{n_s}$ . The process is fully characterized by its initial density  $p(s_1 | x) = \mu_x(s_1)$  and transition probability density

$$p(s_t | s_{1:t-1}, x) = p(s_t | s_{t-1}, x) = f_x(s_t | s_{t-1}),$$

where we use the notation  $s_{1:t-1}$  in place of  $(s_1, \dots, s_{t-1})$ .

We assume that the process  $S_t$  is not directly observable, but partial observations are received via a second process  $Z_t \subseteq \mathcal{Z}^{n_z}$ . The observations  $z_t$  are conditionally independent given  $S_t$  and are defined by the probability density

$$p(z_t | z_{1:t-1}, s_{1:t}, x) = p(z_t | s_t, x) = g_x(z_t | s_t).$$

The posterior density of the parameters  $\pi(x) \propto p(z_{1:T} | x)p(x)$ , where  $p(x)$  is a prior density for  $x$ , is obtained by integrating out the latent process  $\{S_t\}_{t \geq 1}$  to give the marginal likelihood

$$p(z_{1:T} | x) = p(z_1 | x) \prod_{t=2}^T p(z_t | z_{1:t-1}, x),$$

where

$$p(z_t | z_{1:t-1}, x) = \int g_x(z_t | s_t) \int f_x(s_t | s_{t-1}) p(s_{t-1} | z_{1:t-1}, x) ds_{t-1} ds_t \quad (19)$$

is the predictive likelihood.

In general it is impossible to evaluate the likelihood analytically, but it is often possible to approximate the likelihood using a particle filter (Pitt et al., 2012; Doucet and Johansen, 2011; Fearnhead, 2007), by replacing  $p(s_{t-1} \mid z_{1:t-1}, x)$  in (19) with a particle approximation

$$\hat{p}(ds_{t-1} \mid z_{1:t-1}, x) = \sum_{i=1}^N w_{t-1}^{(i)} \delta_{s_{t-1}^{(i)}}(ds_{t-1}), \quad (20)$$

where  $\delta_s$  is a Dirac mass at  $s$  and  $s_{t-1}^{(i)}$  is the  $i$ th particle at  $t-1$  with normalized weight  $w_{t-1}^{(i)}$ . An approximation to the likelihood (19) is then given by the particle approximation  $\{w_{t-1}^{(i)}, s_{t-1}^{(i)}\}_{i=1}^N$ ,

$$\hat{p}(z_t \mid z_{1:t-1}, x) = \sum_{i=1}^N \frac{\tilde{w}_t^{(i)}}{N},$$

where  $\tilde{w}_t^{(i)}$  is the  $i$ th unnormalized importance weight at  $t$ . Using the Metropolis–Hastings algorithm (1) we can target the exact posterior density  $\pi(x)$  as outlined in Section 2. Full details are given in the Supplementary Material.

The particle Langevin algorithm requires an estimate of the gradient of the log posterior density,  $\nabla \log \pi(x) = \nabla \log p(z_{1:T} \mid x) + \nabla \log p(x)$ . Assuming  $\nabla \log p(x)$  is known, it is then only necessary to approximate the score vector  $\nabla \log p(z_{1:T} \mid x)$  with a particle approximation of Fisher’s identity (Cappé et al., 2005)

$$\nabla \log p(z_{1:T} \mid x) = E \{ \nabla \log p(S_{1:T}, z_{1:T} \mid x) \mid z_{1:T}, x \}, \quad (21)$$

which is the expectation, with respect to  $p(s_{1:T} \mid z_{1:T}, x)$ , of

$$\nabla \log p(s_{1:T}, z_{1:T} \mid x) = \sum_{t=1}^T \nabla \log g_x(z_t \mid s_t) + \nabla \log f_x(s_t \mid s_{t-1})$$

over the path  $s_{1:T}$ , where we have used the notation  $f_x(s_1 \mid s_0) = \mu_x(s_1)$ .

A particle approximation is obtained by running the particle filter for  $t = 1, \dots, T$  and storing the particle path  $s_{1:T}^{(i)}$ . Using the method of Poyiadjis et al. (2011), the score vector is approximated by

$$\nabla \log \hat{p}(z_{1:T} \mid x) = \sum_{i=1}^N w_T^{(i)} \nabla \log p(s_{1:T}^{(i)}, z_{1:T} \mid x),$$

where  $w_T^{(i)}$  is an importance weight.

With this approach the variance of the score estimate increases quadratically with  $T$ . Poyiadjis et al. (2011) suggest an alternative particle filter algorithm, which avoids the quadratically increasing variance, but at the expense of a computational cost that is quadratic in the number of particles. Instead, we use the algorithm of Nemeth et al. (2016), which uses kernel density estimation and Rao–Blackwellization to substantially reduce the Monte Carlo variance, but still

maintains an algorithm whose computational cost is linear in the number of particles; see the Supplementary Material. Importantly, the theory presented in Section 3 is not tied to any particular method for approximating the gradient of the log posterior density, and as such, alternative approaches proposed by Poyiadjis et al. (2011), Ionides et al. (2011), Dahlin et al. (2014) and others, are equally supported by our theoretical results.

## 4.2 Linear Gaussian Model

This section provides simulation results to support the theory outlined in Section 3. We show that, while our theory is based on the limit as the number of parameters tends to infinity, it adequately describes the empirical results for a target with a finite number of parameters.

We start by considering the following linear Gaussian state-space model, where it is possible to estimate the posterior density  $p(x \mid z_{1:T})$ , and its gradient, exactly with the Kalman filter (Durbin and Koopman, 2001),

$$z_t = \alpha + \beta s_t + \tau_\epsilon \nu_t, \quad s_t = \mu + \phi s_{t-1} + \sigma_\epsilon \eta_t, \quad s_0 \sim \mathcal{N}\{\mu/(1-\phi), \sigma_\epsilon^2/(1-\phi^2)\},$$

where  $\nu_t$  and  $\eta_t$  are standard independent Gaussian random variables and the vector of model parameters is  $x = (\alpha, \beta, \tau_\epsilon, \mu, \phi, \sigma_\epsilon)^\top$ . We simulated 500 observations from the model with parameters  $x = (0.2, 1, 1, 0.1, 0.9, 0.15)^\top$ , and defined the following prior distributions:

$$\begin{pmatrix} \alpha \\ \beta \end{pmatrix} \sim \mathcal{N} \left\{ \begin{pmatrix} 0.3 \\ 1.2 \end{pmatrix}, \tau_\epsilon^2 \begin{pmatrix} 0.25 & 0 \\ 0 & 0.5 \end{pmatrix} \right\}, \quad \tau_\epsilon^2 \sim \text{Inverse Gamma}(1, 7/20),$$

$\mu \sim \mathcal{N}(0.15, 0.5)$ ,  $(\phi + 1)/2 \sim \text{Beta}(20, 5)$  and  $\sigma_\epsilon^2 \sim \text{Inverse Gamma}(2, 1/40)$ .

The parameters  $(\phi, \sigma_\epsilon, \tau_\epsilon)$  are constrained as  $|\phi| < 1$ ,  $\sigma_\epsilon > 0$  and  $\tau_\epsilon > 0$ . These parameters are transformed as  $\tanh \phi$ ,  $\log \sigma_\epsilon$  and  $\log \tau_\epsilon$  to implement the particle Langevin and random-walk proposals on the unconstrained space.

For this model it is possible to use the fully adapted particle filter (Pitt and Shephard, 1999) using the optimal proposal for the latent states, which, compared to the simpler bootstrap filter (Gordon et al., 1993), reduces the variance in the posterior estimates. The particle Langevin algorithm was run for 100,000 iterations with  $\lambda^2 = \gamma_i^2 \times 1.125^2/6^{-1/3} \times \hat{V}$ , where  $\gamma = (0.25, 0.5, 0.75, 1, 1.25, 1.5, 1.75, 2)$  and  $\hat{V}$  is the empirical posterior covariance estimated from a pilot run. Estimates of the posterior density and gradient of the log posterior density were calculated using a particle filter with particles  $N = (200, 100, 70, 40, 20, 5, 10, 1)$ ; see the Supplementary Material.

Figure 3 shows the efficiency of the particle Langevin algorithm for various scalings  $\gamma$  and noise  $\sigma^2$ . Dividing the minimum effective sample size, taken over the parameters, by the computational time of the algorithm provides a practical measure of efficiency corresponding to the theoretical measure in (10).




Figure 3: Empirical efficiency measured as the minimum effective sample size per computational second. The left panel gives the efficiency plotted against  $\gamma$  for ( $N = 5$  —,  $N = 10$  - - - -,  $N = 20$  - - - -,  $N = 40$  - · - · -,  $N = 70$  · · · · ·,  $N = 100$  - - - -,  $N = 200$  —). The right panel gives the efficiency plotted against  $\sigma^2$ , estimated at the true parameter, for various scalings ( $\gamma = 0.25$  —,  $\gamma = 0.5$  - - - -,  $\gamma = 0.75$  · · · · ·,  $\gamma = 1$  - · - · -,  $\gamma = 1.25$  - - - -,  $\gamma = 1.5$  - - - -,  $\gamma = 1.75$  —).

The left panel of Fig. 3 shows that, initially, increasing the number of particles leads to a more efficient sampler. However, beyond 20 particles, the increase in computational cost outweighs the further improvement in mixing. Setting  $N = 20$  results in a noisy estimate of the posterior density with  $\sigma^2 \approx 2.6$ , supporting Corollary 3.2; the optimal acceptance rate was 19%, slightly above the theoretical optimum. Also, the insensitivity of the optimal scaling to the noise variance, as shown in Figure 1, is seen here as the efficiency is maximized for  $\gamma$  between 1 and 1.5 regardless of the number of particles; similarly the right panel shows the same insensitivity of the optimal variance to the scaling, with efficiency maximized for  $\sigma^2$  between 1.5 and 3, regardless of the scaling. Both of these insensitivities are predicted by the theory established in Section 3.

### 4.3 Mixture model of autoregressive experts

We now use a real data example from Pitt et al. (2012) to illustrate the improvement of using the particle Langevin algorithm (3) over the particle random-walk algorithm. Moreover, we show that estimating the gradient using the  $\mathcal{O}(N)$  algorithm of Nemeth et al. (2016) is more efficient than the  $\mathcal{O}(N^2)$  algorithm of Poyiadjis et al. (2011).

This example uses a two-component mixture of experts model observed with noise. Each of the experts is represented by a first order autoregressive process, where the mixing of the experts is probabilistic rather than deterministic. The model is defined as

$$\begin{aligned}
 z_t &= s_t + \tau_\epsilon \nu_t, & s_t &= \psi_{J_t} + \phi_{J_t} s_{t-1} + \sigma_{J_t} \eta_t & J_t &= 1, 2 \\
 \text{pr}(J_t = 1 \mid s_{t-1}, s_{t-2}) &= \frac{\exp\{\xi_1 + \xi_2 s_{t-1} + \xi_3 (s_{t-1} - s_{t-2})\}}{1 + \exp\{\xi_1 + \xi_2 s_{t-1} + \xi_3 (s_{t-1} - s_{t-2})\}},
 \end{aligned} \tag{22}$$



Table 1: Empirical effective sample size per computational time

Algorithm	Effective sample size per computation time										
		$\tau_\epsilon$	$\psi_1$	$\psi_2$	$\phi_1$	$\phi_2$	$\sigma_1$	$\sigma_2$	$\xi_1$	$\xi_2$	$\xi_3$
Particle random walk	Min	3.39	2.96	1.65	2.15	1.96	1.38	2.16	2.54	2.05	2.09
	Max	4.65	3.68	3.15	3.68	3.48	2.82	3.56	4.20	3.32	3.71
Particle Langevin $\mathcal{O}(N)$	Min	4.11	3.21	4.77	3.57	4.18	2.60	3.68	4.59	3.32	3.08
	Max	5.12	5.71	6.37	6.12	6.43	5.47	6.22	7.34	7.02	6.10
Poyiadjis $\mathcal{O}(N^2)$	Min	0.76	0.60	1.00	0.96	0.47	0.33	0.90	1.06	0.59	0.59
	Max	1.25	1.19	1.37	1.35	1.35	1.17	1.26	1.82	1.23	0.96

where  $\nu_t$  and  $\eta_t$  are standard independent Gaussian random variables and there are 10 model parameters  $x = (\tau_\epsilon, \psi_1, \psi_2, \phi_1, \phi_2, \sigma_1, \sigma_2, \xi_1, \xi_2, \xi_3)^T$ .

Pitt et al. (2012) used the mixture of autoregressive experts to model the growth of US gross domestic product from the second quarter of 1984 to the third quarter of 2010. This model follows previous observations that economic cycles display nonlinear and non-Gaussian features (Hamilton, 1989). Including measurement noise in the model accounts for adjustments made to the data between the first and final release (Zellner, 1992). We impose the constraint  $\psi_1(1 - \phi_1) < \psi_2(1 - \phi_2)$  to ensure that the mean of expert one is less than that of expert two. This implies that the first expert is identified as a low growth regime.

A particle filter approach to this problem is ideal if we assume measurement error in the data. Standard Markov chain Monte Carlo methods could be applied on this model where the latent states are sampled conditional on the parameters and vice-versa (Pitt et al., 2010). However, this would cause the sampler to mix slowly and would ultimately be less efficient than a particle filter implementation, whereby the latent states are integrated out. We compare the particle Langevin algorithm against the particle random-walk algorithm, as implemented in Pitt et al. (2012). For both methods we implement a fully adapted particle filter, where the number of particles were tuned to give a variance of less than 3 for the log posterior density.

We ran the particle Markov chain Monte Carlo algorithm for 100,000 iterations, discarding the first half as burn-in. We compare the random-walk proposal, scaled as  $\lambda^2 = 2.526^2/10 \times \hat{V}$ , against the particle Langevin proposal, where  $\lambda^2 = 1.125^2/10^{-1/3} \times \hat{V}$  and  $\hat{V}$  is an estimate of the posterior covariance taken from a pilot run. A Gaussian prior density is assumed for  $x$ , where constrained parameters are transformed appropriately, and the hyper-parameters are given in Pitt et al. (2012). Table 1 gives a comparison of the proposals, including a particle Langevin algorithm using the  $\mathcal{O}(N^2)$  gradient estimator of Poyiadjis et al. (2011). The minimum and maximum effective sample size per computational minute, taken over 10 simulations, are reported.

The results from the simulation study are summarized in Table 1. There is a significant improvement in terms of effective sample size when using the particle Langevin proposal compared to the random-walk proposal. The effective sample size of the Poyiadjis et al. (2011)  $\mathcal{O}(N^2)$  algorithm is approximately equal to that of our particle Langevin algorithm, but when

taking into account the computational cost, this proposal performs worse than the random-walk algorithm. Therefore, it is important to estimate the gradient of the log posterior density with the same computational cost used to estimate the log posterior density in order for the gradient information to be beneficial.

## 5 Discussion

Our theory identifies three distinct asymptotic regimes corresponding to three different ranges of control over errors in the estimate of the gradient of the log posterior density. We have shown that, if there is no control of these errors, then the particle Langevin algorithm is asymptotically no more efficient than the particle random-walk Metropolis algorithm. By contrast, if there is sufficient control, the particle Langevin algorithm attains the same asymptotic advantage in efficiency over the particle random-walk algorithm as the Metropolis adjusted Langevin algorithm enjoys over its random-walk Metropolis counterpart.

In the preferred regime, and specifically when the estimate of the log posterior density is generated by a particle filter, we identify an optimal variance for the error in the log posterior density of approximately 3.0 and an optimal acceptance rate of approximately 15%. We also find that the optimal scaling is insensitive to the choice of variance and vice-versa. In general, however, the regime is not known, and so, conditional on a fixed but arbitrary number of particles, we provide a mechanism for tuning the scaling of the proposal by aiming for a maximin acceptance rate that is robust to the regime. This ensures that the resulting algorithm will achieve at least 90% of the efficiency that it would were the regime known and the best scaling for that regime chosen.

Our results are the latest in a number of results concerning at optimal implementations of pseudo-marginal and particle Markov chain Monte Carlo algorithms (Pitt et al., 2012; Sherlock et al., 2015; Doucet et al., 2015). Using similar techniques to those in this article, Sherlock et al. (2015) identified an optimal variance of 3.3 for the target density in the particle random-walk Metropolis algorithm, and found that the optimal variance is insensitive to the scaling and vice-versa. Doucet et al. (2015) analyzed non-asymptotic bounds on the integrated autocorrelation time and suggested that for any Metropolis–Hastings algorithm, the optimal variance should be between 0.85 and 2.82, also suggesting an insensitivity. In high dimensions, one should tune to the variance suggested by the asymptotic theory, but empirical studies on both the particle random-walk and the particle Langevin algorithms have shown that in low dimensions the optimal variance is typically less than 3. Given that all our assumptions hold at best approximately in practice, we would recommend two possible tuning strategies. The first strategy is to evaluate the variance at several points in the main posterior mass and ensure that the largest of these is slightly lower than 3; this is justified both because of the above mentioned empirical findings and because we expect the optimal variance to be smaller in regimes (1) and (2). Then tune the scaling to achieve the acceptance rate given by Figure 2. For the second strategy, start

with a sensible scaling, find the number of particles that optimizes the overall efficiency, for example effective sample size per second, then with this number of particles, find the scaling which optimizes efficiency, for example effective sample size.

## Acknowledgements

The authors are grateful to the editor, associate editor and referees whose comments and advice have greatly improved this paper. We thank Jeff Rosenthal for providing the Mathematica scripts from Roberts and Rosenthal (1998). This research was supported by the EPSRC i-like grant and the STOR-i Centre for Doctoral Training.

## References

- Andrieu, C., Doucet, A., and Holenstein, R. (2010). Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(3):269–342.
- Andrieu, C. and Roberts, G. O. (2009). The pseudo-marginal approach for efficient Monte Carlo computations. *The Annals of Statistics*, 37(2):697–725.
- Andrieu, C. and Thoms, J. (2008). A tutorial on adaptive MCMC. *Statistics and Computing*, 18(4):343–373.
- Beaumont, M. A. (2003). Estimation of population growth or decline in genetically monitored populations. *Genetics*, 164(3):1139–60.
- Bédard, M. (2007). Weak convergence of Metropolis algorithms for non-i.i.d. target distributions. *Ann. Appl. Probab.*, 17(4):1222–1244.
- Bérard, J., Del Moral, P., and Doucet, A. (2014). A lognormal central limit theorem for particle approximations of normalizing constants. *Electronic Journal of Probability*, 19:1–28.
- Beskos, A., Roberts, G., and Stuart, A. (2009). Optimal scalings for local Metropolis–Hastings chains on non-product targets in high dimensions. *Annals of Applied Probability*, 19(3):863–898.
- Cappé, O., Godsill, S., and Moulines, E. (2007). An Overview of Existing Methods and Recent Advances in Sequential Monte Carlo. *Proceedings of the IEEE*, 95(5):899–924.
- Cappé, O., Moulines, E., and Ryden, T. (2005). *Inference in Hidden Markov Models*. Springer, New York.

- Dahlin, J., Lindsten, F., and Schön, T. B. (2014). Particle Metropolis Hastings using gradient and Hessian information. *Statistics and Computing*, 25:81–92.
- Del Moral, P. (2004). *Feynman-Kac Formulae: Genealogical and Interacting Particle Systems with Applications*. Springer, New York.
- Doucet, A., Godsill, S., and Andrieu, C. (2000). On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and Computing*, 10(3):197–208.
- Doucet, A. and Johansen, A. M. (2011). A tutorial on particle filtering and smoothing: Fifteen years later. In Crisan, D. and Rozovskii, B., editors, *The Oxford Handbook of Nonlinear Filtering*, pages 656–704. Oxford University Press, New York.
- Doucet, A., Pitt, M. K., Deligiannidis, G., and Kohn, R. (2015). Efficient implementation of Markov chain Monte Carlo when using an unbiased likelihood estimator. *Biometrika*, 102(2):295–313.
- Durbin, J. and Koopman, S. (2001). *Time Series Analysis by State Space Methods*. Oxford University Press, Oxford.
- Fearnhead, P. (2007). Computational methods for complex stochastic systems: a review of some alternatives to MCMC. *Statistics and Computing*, 18(2):151–171.
- Fearnhead, P., Papaspiliopoulos, O., and Roberts, G. O. (2008). Particle filters for partially observed diffusions. *Journal of the Royal Statistical Society, Series B*, 70:1–28.
- Golightly, A. and Wilkinson, D. J. (2011). Bayesian parameter inference for stochastic biochemical network models using particle Markov chain Monte Carlo. *Interface Focus*, 1(6):807–20.
- Gordon, N., Salmond, D., and Smith, A. F. M. (1993). Novel approach to nonlinear and linear Bayesian state estimation. *IEE Proceedings*, 140(2):107–113.
- Hamilton, J. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2):357–384.
- Ionides, E. L., Bhadra, A., Atchadé, Y., and King, A. (2011). Iterated filtering. *The Annals of Statistics*, 39(3):1776–1802.
- Knape, J. and de Valpine, P. (2012). Fitting complex population models by combining particle filters with Markov chain Monte Carlo. *Ecology*, 93(2):256–263.
- Neal, P. and Roberts, G. O. (2006). Optimal scaling for partially-updating MCMC algorithms. *Ann. Appl. Probab.*, 16(2):475–515.

- Nemeth, C., Fearnhead, P., and Mihaylova, L. (2016). Particle approximations of the score and observed information matrix for parameter estimation in state space models with linear computational cost. *Journal of Computational and Graphical Statistics (to appear)*.
- Pitt, M., Giordani, P., and Kohn, R. (2010). Bayesian inference for time series state space models. In Geweke, J., Koop, G., and van Dijk, H., editors, *The Oxford Handbook of Bayesian Econometrics*, pages 61–124. Oxford University Press, New York.
- Pitt, M. K. and Shephard, N. (1999). Filtering via Simulation: Auxiliary Particle Filters. *Journal of the American Statistical Association*, 94(446):590–599.
- Pitt, M. K., Silva, R., Giordani, P., and Kohn, R. (2012). On some properties of Markov chain Monte Carlo simulation methods based on the particle filter. *Journal of Econometrics*, 171(2):134–151.
- Poyiadjis, G., Doucet, A., and Singh, S. S. (2011). Particle approximations of the score and observed information matrix in state space models with application to parameter estimation. *Biometrika*, 98(1):65–80.
- Roberts, G. O., Gelman, A., and Gilks, W. R. (1997). Weak convergence and optimal scaling of random walk Metropolis algorithms. *The Annals of Applied Probability*, 7:110–120.
- Roberts, G. O. and Rosenthal, J. S. (1998). Optimal scaling of discrete approximations to Langevin diffusions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 60(1):255–268.
- Roberts, G. O. and Rosenthal, J. S. (2001). Optimal scaling for various Metropolis–Hastings algorithms. *Statistical Science*, 16(4):351–367.
- Roberts, G. O. and Rosenthal, J. S. (2009). Examples of adaptive MCMC. *Journal of Computational and Graphical Statistics*, 18(2):349–367.
- Roberts, G. O. and Rosenthal, J. S. (2014). Minimising MCMC variance via diffusion limits, with an application to simulated tempering. *Ann. Appl. Probab.*, 24(1):131–149.
- Särkkä, S., Hartikainen, J., Mbalawata, I., and Haario, H. (2015). Posterior inference on parameters of stochastic differential equations via non-linear Gaussian filtering and adaptive MCMC. *Statistics and Computing*, 25(2):427–437.
- Sherlock, C. (2013). Optimal scaling of the random walk Metropolis: general criteria for the 0.234 acceptance rule. *J. App. Prob.*, 50(1):1–15.
- Sherlock, C. (2016). Optimal scaling for the pseudo-marginal random walk Metropolis: insensitivity to the noise generating mechanism. *Methodology and Computing in Applied Probability (to appear)*.

- Sherlock, C. and Roberts, G. (2009). Optimal scaling of the random walk Metropolis on elliptically symmetric unimodal targets. *Bernoulli*, 15(3):774–798.
- Sherlock, C., Thiery, A. H., Roberts, G. O., and Rosenthal, J. R. (2015). On the efficiency of pseudo-marginal random walk Metropolis algorithms. *Annals of Statistics*, 43(1):238–275.
- West, M. (1993). Approximating posterior distributions by mixture. *Journal of the Royal Statistical Society. Series B*, 55(2):409–422.
- Wolfram, S. (2014). *Mathematica 10.0*. Champaign, Illinois.
- Zellner, A. (1992). Commentary. In Belagia, M. and Garfinkel, M., editors, *The Business Cycle: Theories and Evidence: Proceedings of the Sixteenth Annual Economic Policy Conference of the Reserve Bank of St Louis*, pages 148–158. Kluwer Academic Publishers, Boston.

## Supplemental Materials

### 1 Proof of Theorem 3.1

The proposal density for any given component is

$$q(x, y) = (\lambda_n^2 2\pi)^{-1/2} \exp \left\{ -\frac{1}{2\lambda_n^2} \left( y - x - \frac{1}{2}\lambda_n^2 g'(x) \right)^2 \right\}.$$

Define

$$R(x_i^n, Y_i^n) = \log \left\{ \frac{f(Y_i^n)q(Y_i^n, x_i^n)}{f(x_i^n)q(x_i^n, Y_i^n)} \right\}$$

and  $T_n(X^n, Y^n) = \sum_{i=1}^n R(X_i^n, Y_i^n)$ , so that the acceptance probability is

$$\alpha_n(x^n, w^n; Y^n, V^n) = 1 \wedge \exp \{V^n - w^n + T_n(x^n, Y^n)\}, \quad (\text{S1})$$

where  $V^n$  and  $W^n$  are given in (8) and (9). Note, we use the notation  $a \wedge b = \min(a, b)$ .

**Proposition 1.1.**

$$\begin{aligned} R(x_i^n, Y_i^n) &= C_3(x_i^n, Z_i)\lambda_n^3 + C_4(x_i^n, Z_i)\lambda_n^4 + C_5(x_i^n, Z_i)\lambda_n^5 \\ &\quad + C_6(x_i^n, Z_i)\lambda_n^6 + C_7(x_i^n, Z_i, \lambda_n), \end{aligned} \quad (\text{S2})$$

where

$$C_3(x_i^n, Z_i) = -\frac{1}{12}\ell^3 \{3Z_i g'(x_i^n)g''(x_i^n) + Z_i^3 g'''(x_i^n)\},$$

and where  $C_4(x_i^n, Z_i)$ ,  $C_5(x_i^n, Z_i)$  and  $C_6(x_i^n, Z_i)$  are also polynomials in  $Z_i$  and the derivatives of  $g$ . Furthermore, if  $E_Z$  denotes expectation with  $Z \sim \mathcal{N}(0, 1)$  and  $E_X$  denotes expectation with  $X$  having the density  $f(\cdot)$ , then

$$E_X[E_Z \{C_3(X, Z)\}] = E_X[E_Z \{C_4(X, Z)\}] = E_X[E_Z \{C_5(X, Z)\}] = 0, \quad (\text{S3})$$

whereas

$$E_X[E_Z \{C_3(X, Z)^2\}] = \ell^6 K^2 = -2E_X[E_Z \{C_6(X, Z)\}] > 0. \quad (\text{S4})$$

Also

$$\text{var} \{C_4(X, Z)\} < \infty, \quad \text{var} \{C_5(X, Z)\} < \infty, \quad \text{var} \{C_6(X, Z)\} < \infty, \quad (\text{S5})$$

where  $\text{var}$  denotes variance over both  $Z$  and  $X$ . Finally

$$E_Z \{|C_7(x_i^n, Z_i, \lambda_n)|\} \leq n^{-7/6} p(x_i^n), \quad (\text{S6})$$

where  $p(x)$  is a polynomial in  $x$ .

*Proof.* As in Roberts and Rosenthal (1998), equation (S2) follows by Taylor expansion of  $g$  and its derivatives using MATHEMATICA (Wolfram, 2014) and collecting terms in powers of  $n$ . Straightforward inspection shows that  $C_3$  has the claimed form and that  $C_4$ ,  $C_5$  and  $C_6$  are also polynomials in  $Z$  and the derivatives of  $g$ , as claimed. All terms in both  $C_3$  and  $C_5$  contain odd powers of  $Z$  and so their expectations are zero. Equation (S4), and the fact that the expectation of  $C_4$  is zero, follows after integrating by parts where expectations of products of the derivatives of  $g$  are being taken with respect to the density  $e^{g(x)}$ . Thus the equivalent form of  $K$  defined in (13) is real and positive. Equation (S5) follows from the polynomial form for  $C_4$ ,  $C_5$  and  $C_6$  and assumptions (5) and (6).

Using the remainder formula of the Taylor series expansion we may derive the bound

$$|C_7(x_i^n, Z_i, \lambda_n)| \leq n^{-7/6} p_*(x_i^n, w_i),$$

for some polynomial  $p_*$ , with  $|w_i| \leq |Z_i|$ . But for any polynomial  $p_*(x, w) \leq A(1 + x^N)(1 + w^N)$ , with a sufficiently large  $A$  and for a sufficiently large even integer  $N$ , (S6) follows with  $p(x) = AE \{(1 + Z^N)\} (1 + x^N)$ .  $\square$

Proposition 1.1 allows us to find the limiting distribution of one of the key terms in the acceptance probability of the algorithm when the Markov chain on  $(X, W)$  is stationary, (S1).

**Lemma 1.2.**

$$T_n(X^n, Y^n) \Rightarrow T \sim \mathcal{N} \left( -\frac{1}{2} \ell^6 K^2, \ell^6 K^2 \right).$$

*Proof.* First note that, by (S6),

$$E \left\{ \left| \sum_{i=1}^n C_7(X_i^n, Z_i, \lambda_n) \right| \right\} \leq \sum_{i=1}^n E \{|C_7(X_i^n, Z_i, \lambda_n)|\} \leq n^{-1/6} E \{p(X)\}.$$

However  $E \{p(X)\} < \infty$  by (6) so, by Markov's inequality,  $\sum_{i=1}^n C_7(X_i^n, Z_i, \lambda_n) \rightarrow 0$  in probability as  $n \rightarrow \infty$ . By Slutsky's Theorem it is therefore sufficient to show that in probability  $T'_n \rightarrow T \sim \mathcal{N} \left( -\frac{1}{2} \ell^6 K^2, \ell^6 K^2 \right)$ , where we define

$$T'_n = \sum_{i=1}^n \{C_3(X_i^n, Z_i) n^{-1/2} + C_4(X_i^n, Z_i) n^{-2/3} + C_5(X_i^n, Z_i) n^{-5/6} + C_6(X_i^n, Z_i) n^{-1}\}.$$

Combining (S3), (S4) and (S5)

$$\begin{aligned} E(T'_n) &= -\frac{1}{2} \ell^6 K^2, \\ \text{var}(T'_n) &= \text{var} \{C_3(X, Z)\} + O(n^{-1/6}) \rightarrow \ell^6 K^2 \end{aligned}$$

in probability. Moreover  $T'_n$  is the sum of  $n$  independent and identically distributed terms, so the result follows by the central limit theorem.  $\square$



Thus

$$T_n + V^n - W^n \Rightarrow \mathcal{N} \left( B - \frac{1}{2} \ell^6 K^2, \ell^6 K^2 \right). \quad (\text{S7})$$

Now if  $U \sim \mathcal{N}(a, b^2)$  then  $E(1 \wedge e^U) = \Phi(a/b) + e^{a+b^2/2} \Phi(-b - a/b)$  (e.g., Roberts et al. (1997)). Since  $\alpha_n = E(1 \wedge e^{T_n+B^n})$ , we may apply the Bounded Convergence Theorem to see that

$$\lim_{n \rightarrow \infty} \alpha_n = E(1 \wedge e^{T+B}) = 2E \left\{ \Phi \left( \frac{B}{\ell^3 K} - \frac{1}{2} \ell^3 K \right) \right\}.$$

proving the first part of Theorem 3.1.

To prove the second result, we first note the following

**Proposition 1.3.**

$$\begin{aligned} \lim_{n \rightarrow \infty} E \left( n^{-2/3} \|Y^n - X^n\|^2 \right) &= \ell^2 \\ \lim_{n \rightarrow \infty} E \left\{ \left( n^{-2/3} \|Y^n - X^n\|^2 - \ell^2 \right)^2 \right\} &= 0. \end{aligned}$$

*Proof.* To simplify the exposition we suppress the superscripts,  $n$ , in  $X^n$  and  $Y^n$ . Firstly,

$$\begin{aligned} n^{-2/3} E \left( \|Y - X\|^2 \right) &= n^{1/3} E \left\{ (Y_1 - X_1)^2 \right\} \\ &= n^{1/3} E \left[ \left\{ \ell n^{-1/6} Z_1 + \frac{1}{2} \ell^2 n^{-1/3} g'(X_1) \right\}^2 \right] \\ &= \ell^2 + \frac{1}{4} \ell^4 n^{-1/3} E \left\{ g'(X_1)^2 \right\}. \end{aligned}$$

By assumptions (5) and (6),  $E \left\{ g'(X_1)^2 \right\} < \infty$  and the first result follows. Also as  $n \rightarrow \infty$ ,

$$\begin{aligned} \text{var} \left( n^{-2/3} \|Y - X\|^2 - \ell^2 \right) &= n^{-4/3} \text{var} \left\{ \sum_{i=1}^n (Y_1 - X_1)^2 \right\} \\ &= n^{-1/3} \text{var} \left[ \left\{ \ell n^{-1/6} Z_1 + \frac{1}{2} \ell^2 n^{-1/3} g'(X_1) \right\}^2 \right] \rightarrow 0, \end{aligned}$$

by (5) and (6). This, combined with the first part of this proposition proves the second part.  $\square$

To complete the proof, we abbreviate  $\alpha_n(X^n, W^n; Y^n, V^n)$  to  $A_n$ . Now

$$\left| E \left( n^{-2/3} \|Y^n - X^n\|^2 A_n \right) - \ell^2 \bar{\alpha} \right| \leq \left| E \left\{ \left( n^{-2/3} \|Y^n - X^n\|^2 - \ell^2 \right) A_n \right\} \right| + \left| E \left\{ \ell^2 (A_n - \bar{\alpha}) \right\} \right|.$$

The second term on the right hand side converges to zero by the first part of Theorem 3.1. The Cauchy–Schwarz inequality bounds the first term on the right hand side by

$$E \left\{ \left( n^{-2/3} \|Y^n - X^n\|^2 - \ell^2 \right)^2 \right\}^{1/2} E(A_n^2)^{1/2}.$$

The first term converges to zero by Proposition 1.3 and the second term is bounded.

## 2 Proof of Corollary 3.2

First note that for some  $Z \sim N(0, 1)$  that is independent of  $B$ ,

$$\Phi\left(\frac{B}{\ell^3 K} - \frac{\ell^3 K}{2}\right) = \mathbb{P}\left(\ell^3 K Z - B \leq -\frac{1}{2}\ell^6 K^2\right) = \Phi\left\{-\left(\frac{1}{2}\ell^6 K^2 + \sigma^2\right)(\ell^6 K^2 + 2\sigma^2)^{-1/2}\right\}.$$

So

$$\alpha(\ell, \sigma^2) = 2\Phi\left\{-\frac{1}{2}(\ell^6 K^2 + 2\sigma^2)^{-1/2}\right\}.$$

Set  $a^2 = K^2 \ell^6$  and  $b^2 = 2\sigma^2$  then

$$\text{Eff}(\ell, \sigma^2) \propto a^{2/3} b^2 \Phi\left\{-\frac{1}{2}(a^2 + b^2)^{-1/2}\right\}.$$

Given  $a^2 + b^2$ ,  $a^{2/3} b^2$  is maximized when  $b^2 = 3a^2$ , at which point the efficiency is proportional to  $a^{8/3} \Phi(-a)$ . Numerical optimization shows that this function is maximized at  $\hat{a} \approx 1.423$ , and thus the optimal acceptance rate is  $\hat{\alpha} = 2\Phi(-\hat{a}) \approx 15.47\%$ . As a result, the optimal scaling and variance are  $\ell_{\text{opt}} \approx 1.125K^{-1/3}$  and  $\sigma_{\text{opt}}^2 \approx 3.038$ , as given in the statement.

## 3 Proof of Theorem 3.3

For the sake of brevity, we shall prove statements (1), (2) and (3) of Theorem 3.3 together rather than separately. Throughout the proof, therefore, the superscript  $*$  will be used to denote a superscript that could be replaced by (1), (2) or (3) according to the case in the statement of Theorem 3.3 that is being considered.

For  $* \in \{(1), (2), (3)\}$  let  $R^*(x_i^n, Y_i^n)$  be the log Metropolis–Hastings ratio where the proposal,  $q^*(x_i^n, Y_i^n)$ , is the particle Langevin proposal given in (14), and let

$$T_n^*(X^n, Y^n) = \sum_{i=1}^n R^*(X_i^n, Y_i^n). \quad (\text{S8})$$

We also define  $U_i = (U_{x_i}, U_{y_i})$ , to be the vector of (zero mean and unit variance) noise terms in the  $i$ th component of the gradient estimate used, respectively, in the particle Langevin proposal from the current value and the proposal for the corresponding reverse move from the proposed value.

The proof commences with an analogous result to Proposition 1.1 from Section 1.

**Proposition 3.1.** *Let  $R(x_i^n, Y_i)$  be the idealized particle Langevin algorithm term from Propo-*

sition 1.1. Then for (\*) in (1), (2) or (3)

$$R^*(x_i^n, Y_i^n) = R(x_i^n, Y_i^n) + C_{1,1}(x_i^n, U_i, Z_i)\lambda_n n^{-\kappa} + C_{2,1}(x_i^n, U_i, Z_i)\lambda_n^2 n^{-\kappa} + C_{3,1}(x_i^n, U_i, Z_i)\lambda_n^3 n^{-\kappa} \\ + C_{4,1}(x_i^n, U_i, Z_i)\lambda_n^4 n^{-\kappa} + C_{2,2}(x_i^n, U_i, Z_i)\lambda_n^2 n^{-2\kappa} + C_r(x_i^n, U_i, Z_i, n), \quad (\text{S9})$$

Here

$$C_{1,1} = -\frac{1}{2}\ell\tau U_{x_i^n} Z_i - \frac{1}{2}\ell\tau U_{y_i^n} Z_i - b(x_i^n)\ell Z_i,$$

and  $C_{2,1}, C_{3,1}, C_{4,1}$  and  $C_{2,2}$  are all polynomials in  $Z$ , in derivatives of  $g(x)$ , and in  $b(x)$  and its derivatives. Furthermore, let  $E_{X,U,Z}$  denote expectation with respect to  $X$  having the density  $f(\cdot)$ ,  $Z \sim \mathcal{N}(0, 1)$ , and with respect to  $U_x$  and  $U_y$  with  $E(U_x) = E(U_y) = 0$  and  $\text{var}(U_x) = \text{var}(U_y) = 1$ . Then

$$E_{X,U,Z} \{C_{1,1}(X, U, Z)\} = E_{X,U,Z} \{C_{2,1}(X, U, Z)\} = E_{X,U,Z} \{C_{3,1}(X, U, Z)\} = 0 \quad (\text{S10})$$

$$E_{X,U,Z} \{C_{1,1}(X, U, Z)^2\} = -2E_{X,U,Z} \{C_{2,2}(X, U, Z)\} = K_*^2 \quad (\text{S11})$$

$$E_{X,U,Z} \{C_3(X, U, Z)C_{1,1}(X, U, Z)\} = -E_{X,U,Z} \{C_{4,1}(X, U, Z)\} = K_{**}. \quad (\text{S12})$$

where  $K_*$  and  $K_{**}$  are defined in (17) and (18). Also

$$\text{var} \{C_{1,1}(X, U, Z)\} < \infty, \text{var} \{C_{2,1}(X, U, Z)\} < \infty, \text{var} \{C_{3,1}(X, U, Z)\} < \infty, \\ \text{var} \{C_{4,1}(X, U, Z)\} < \infty, \text{var} \{C_{2,2}(X, U, Z)\} < \infty, \quad (\text{S13})$$

where  $\text{var}$  denotes variance over  $Z, U$  and  $X$ . Finally

$$E_{U,Z} \{|C_7(x_i^n, U_i, Z_i, \lambda_n)|\} \leq n^{-7/6} p(x_i^n), \quad (\text{S14})$$

where  $p(x)$  is a polynomial in  $x$ .

*Proof.* Writing  $A(x) = b(x) + \tau U_x$  and  $A(y) = b(y) + \tau U_y$ , after some algebra we obtain

$$\log q^*(y_i^n, x_i^n) - \log q^*(x_i^n, y_i^n) = \frac{1}{2}Z_i^2 - \frac{1}{2} \left( Z_i + \frac{\lambda_n}{2} [g'(x_i^n) + g'(y_i^n) + n^{-\kappa}\{A(x_i^n) + A(y_i^n)\}] \right)^2.$$

This, together with a simpler calculation for  $\log \pi(y_i^n) - \log \pi(x_i^n)$ , shows that in a Taylor expansion of  $R^*(x_i^n, y_i^n)$  about  $x_i^n$ , terms in  $n^{-a\kappa}$  ( $a = 1, \dots$ ) must also be multiplied by  $\lambda_n^b$  with  $b \in \{a, a+1, \dots\}$ . Consideration of the maximum possible size of all terms in the Taylor expansion for the three different cases shows that it must be of the form given in (S9) with the largest part of the remainder term being at most  $O(n^{-7/6})$ .

As with Proposition 1.1, the polynomial forms for  $C_{1,1}, C_{2,1}, C_{3,1}, C_{4,1}$  and  $C_{2,2}$  are pro-

Table 1: Coefficients of (S9) terms

	$C_3$	$C_4$	$C_5$	$C_6$	$C_{1,1}$	$C_{2,1}$	$C_{3,1}$	$C_{4,1}$	$C_{2,2}$
(1)	$n^{-\frac{1}{2}-3\epsilon}$	$n^{-\frac{2}{3}-4\epsilon}$	$n^{-\frac{5}{6}-5\epsilon}$	$n^{-1-6\epsilon}$	$n^{-\frac{1}{2}}$	$n^{-\frac{2}{3}-\epsilon}$	$n^{-\frac{5}{6}-2\epsilon}$	$n^{-1-3\epsilon}$	$n^{-1}$
(2)	$n^{-\frac{1}{2}}$	$n^{-\frac{2}{3}}$	$n^{-\frac{5}{6}}$	$n^{-1}$	$n^{-\frac{1}{2}}$	$n^{-\frac{2}{3}}$	$n^{-\frac{5}{6}}$	$n^{-1}$	$n^{-1}$
(3)	$n^{-\frac{1}{2}}$	$n^{-\frac{2}{3}}$	$n^{-\frac{5}{6}}$	$n^{-1}$	$n^{-\frac{1}{2}-\epsilon}$	$n^{-\frac{2}{3}-\epsilon}$	$n^{-\frac{5}{6}-\epsilon}$	$n^{-1-\epsilon}$	$n^{-1-2\epsilon}$

duced using MATHEMATICA (Wolfram, 2014), but this time by also Taylor expanding the term  $b(y)$  in  $y - x$ .

Clearly  $E(C_{1,1}) = 0$  as the terms are multiples of  $U_x$ ,  $U_y$  and odd powers of  $Z$ . The same argument can be used for the expectations of  $C_{3,1}$ ; however for  $C_{2,1}$ ,  $C_{4,1}$  and for the relationships in (S11) and (S12), it must be used in tandem with integration by parts with respect to the target density  $e^{g(x)}$  and using assumptions (5), (6) and (16).

We illustrate this by providing the form for  $C_{2,1}$ :

$$C_{2,1}(X, U, Z) = -\frac{1}{2}\ell^2 Z^2 \{b'(X) + b(X)g'(X) + \tau U_y g'(X)\}.$$

However (5) and (16) imply that  $\int b'(x)e^{g(x)} dx = -\int b(x)g'(x)e^{g(x)}$ .

The final two parts of the proposition follow from analogous arguments to those used in Proposition 1.1 provided (15) holds.  $\square$

Integration by parts, the Cauchy–Schwarz inequality and then further integration by parts gives

$$\begin{aligned} K_{**}^2 &= \frac{1}{16}E[b(X)\{g'(X)g''(X) + g'''(X)\}]^2 \leq \frac{1}{16}E\{b(X)^2\}E[\{g'(X)g''(X) + g'''(X)\}^2] \\ &= \frac{1}{48}E\{b(X)^2\}E[3\{g'''(X)\}^2 - 3\{g''(X)\}^3] \leq K_*^2 K^2, \end{aligned}$$

so that  $\ell^4 K^2 + 2\ell^2 K_{**} + K_*^2 \geq 0$ .

**Lemma 3.2.** For  $* \in \{(1), (2), (3)\}$

$$T_n^*(X^n, Y^n) \Rightarrow T^* \sim \mathcal{N}\left(-\frac{1}{2}a^*, a^*\right),$$

where  $T^*$  is defined in (S8) and

$$a^{(1)} = \ell^2 K_*^2, \quad a^{(2)} = \ell^2 K_*^2 + 2\ell^4 K_{**} + \ell^6 K^2, \quad a^{(3)} = \ell^6 K^2.$$

*Proof.* As proved in Lemma 1.2, by Markov’s inequality,  $\sum_{i=1}^n C_r(X_i^n, U_i, Z_i, n) \rightarrow 0$  in probability and therefore it is sufficient to show that in probability  $T_n^* \rightarrow T \sim \mathcal{N}\left(-\frac{1}{2}a^*, a^*\right)$ , where  $* \in \{(1), (2), (3)\}$  and  $T_n^* = \sum_{i=1}^n \{R^*(x_i^n, Y_i^n, Z_i^n) - C_r(x_i^n, U_i, Z_i, n)\}$ . Table 1 shows the coefficient of each non-remainder term in (S9) in each of the three cases.

Since  $\epsilon > 0$ , as  $n \rightarrow \infty$ , combining (S3), (S4), (S10), (S11) and (S12) gives

$$\begin{aligned} E(T_n^{(1)}) &= -\frac{1}{2}n^{-6\epsilon}\ell^6 K - n^{-3\epsilon}\ell^4 K_{**} - \frac{1}{2}\ell^2 K_* \rightarrow -\frac{1}{2}\ell^2 K_*, \\ E(T_n^{(2)}) &= -\frac{1}{2}\ell^6 K - \ell^4 K_{**} - \frac{1}{2}\ell^2 K_*, \\ E(T_n^{(3)}) &= -\frac{1}{2}\ell^6 K - n^{-\epsilon}\ell^4 K_{**} - \frac{1}{2}n^{-2\epsilon}\ell^2 K_* \rightarrow -\frac{1}{2}\ell^6 K, \end{aligned}$$

in probability. Similarly, using (S5) and (S13),

$$\begin{aligned} \text{var}(T_n^{(1)}) &\rightarrow E(C_{1,1}^2) = \ell^2 K_*^2 \\ \text{var}(T_n^{(2)}) &\rightarrow E\{(C_3 + C_{1,1})^2\} = \ell^2 K_*^2 + 2\ell^4 K_{**} + \ell^6 K^2 \\ \text{var}(T_n^{(3)}) &\rightarrow E(C_3^2) = \ell^6 K^2, \end{aligned}$$

in probability. Moreover,  $T_n^*$  is the sum of  $n$  independent and identically distributed terms, so the result follows by the central limit theorem.  $\square$

The proof for the asymptotic acceptance rate,  $\bar{\alpha}_n(\ell, \sigma^2) \rightarrow \alpha^*(\ell, \sigma^2)$ , is completed using Lemma 3.2 as in the proof of Theorem 3.1 by accounting for the distribution of the noise of the log-target.

Finally, as in the proof of Theorem 3.1, the limit for the squared jump distance,  $J_n(\ell, \sigma^2)$ , follows from Proposition 1.3.

## 4 Implementation details for the particle Langevin algorithm

Particle filters, also known as sequential Monte Carlo algorithms, use importance sampling to sequentially approximate the posterior distribution. In the context of state-space modelling, we are interested in approximating the posterior density  $p(s_t \mid z_{1:t}, x)$  of the filtered latent state  $s_t$ , given a sequence of observations  $z_{1:t}$ . In this section, we shall assume that the model parameters  $x$  are fixed. Approximations of  $p(s_t \mid z_{1:t}, x)$  can be calculated recursively by first approximating  $p(s_1 \mid z_1, x)$ , then  $p(s_2 \mid z_{1:2}, x)$  and so forth for  $t = 1, \dots, T$ . At time  $t$  the posterior of the filtered state is

$$p(s_t \mid z_{1:t}, x) \propto g_x(z_t \mid s_t) \int f_x(s_t \mid s_{t-1}) p(s_{t-1} \mid z_{1:t-1}, x) ds_{t-1} \quad (\text{S15})$$

where  $p(s_{t-1} \mid z_{1:t-1}, x)$  is the posterior density at time  $t - 1$ .

The posterior at time  $t$  can be approximated if we assume that at time  $t - 1$  we have a set of particles  $\{s_{t-1}^{(i)}\}_{i=1}^N$  and corresponding normalized weights  $\{w_{t-1}^{(i)}\}_{i=1}^N$  which produce a discrete

approximation of  $p(s_{t-1} | z_{1:t-1}, x)$ . This induces the following approximation to (S15),

$$\hat{p}(s_t | z_{1:t}, x) \approx cg_x(z_t | s_t) \sum_{i=1}^N w_{t-1}^{(i)} f_x(s_t | s_{t-1}^{(i)}), \quad (\text{S16})$$

where  $c$  is a normalizing constant. The filtered density, as given above, can be updated recursively by propagating and updating the particle set using importance sampling techniques. The resulting algorithms are called particle filters, see Doucet et al. (2000) and Cappé et al. (2007) for a review.

In this paper the particle approximations of the latent process are created with the auxiliary particle filter of Pitt and Shephard (1999). This filter can be viewed as a general filter from which simpler filters are given as special cases (Fearnhead et al., 2008). The aim is to view the target (S16) as defining a joint distribution on the particle at time  $t - 1$  and the value of a new particle at time  $t$ . The probability of sampling particle  $s_{t-1}^{(i)}$  and  $s_t$  is

$$cw_{t-1}^{(i)}g_x(z_t | s_t)f_x(s_t | s_{t-1}^{(i)}).$$

We approximate this with  $\xi_t^{(i)}q(s_t | s_{t-1}^{(i)}, z_t, x)$ , where  $q(s_t | s_{t-1}^{(i)}, z_t, x)$  is a density function that can be sampled from and  $\{\xi_t^{(i)}\}_{i=1}^N$  are a set of probabilities. This defines a proposal which we can simulate from by first choosing particle  $s_{t-1}^{(i)}$  with probability  $\xi_t^{(i)}$ , and then, conditional on this, a new particle value,  $s_t$ , is sampled from  $q(s_t | s_{t-1}^{(i)}, z_t, x)$ . The weight assigned to our new particle is then

$$\tilde{w}_t = \frac{w_{t-1}^{(i)}g_x(z_t | s_t)f_x(s_t | s_{t-1}^{(i)})}{\xi_t^{(i)}q(s_t | s_{t-1}^{(i)}, z_t, x)}.$$

Details are summarized in Algorithm 1.

The optimal proposal density, in terms of minimizing the variance of the weights (Doucet et al., 2000), is available when  $q(s_t | s_{t-1}^{(i)}, z_t, x) = p(s_t | s_{t-1}^{(i)}, z_t, x)$  and  $\xi_t^{(i)} \propto w_{t-1}^{(i)}p(z_t | s_{t-1}^{(i)})$ . This filter is said to be fully adapted as all the weights  $w_t^{(i)}$  will equal  $1/N$ . Generally, it is not possible to sample from the optimal proposal, but alternative proposals can be used which approximate the fully adapted filter.

One of the benefits of using the particle filter is that an estimate for the likelihood  $p(z_{1:T} | x)$  is given for free from the particle filter output. We can estimate  $p(z_t | z_{1:t-1}, x)$  by

$$\hat{p}(z_t | z_{1:t-1}, x) = \sum_{i=1}^N \frac{\tilde{w}_t^{(i)}}{N}, \quad (\text{S17})$$

where  $\tilde{w}_t^{(i)}$  are unnormalized weights. An unbiased estimate of the likelihood (Del Moral, 2004) is then

$$\hat{p}(z_{1:T} | x) = \hat{p}(z_1 | x) \prod_{t=2}^T \hat{p}(z_t | z_{1:t-1}, x).$$

---

**Algorithm 1** Auxiliary Particle Filter
 

---

*Step 1:* Iteration  $t = 1$ .

- (a) For  $i = 1, \dots, N$ , sample particles  $\{s_1^{(i)}\}$  from the prior  $p(s_1 | x)$  and set  $\tilde{w}_1^{(i)} = p(z_1 | s_1^{(i)})$ .  
 (b) Calculate  $C_1 = \sum_{i=1}^N \tilde{w}_1^{(i)}$ ; set  $\hat{p}(z_1) = C_1/N$ ; and calculate normalized weights  $w_1^{(i)} = \tilde{w}_1^{(i)}/C_1$  for  $i = 1, \dots, N$ .

*Step 2:* Iteration  $t = 2, \dots, T$ . Assume a user-defined set of proposal weights  $\{\xi_t^{(i)}\}_{i=1}^N$  and family of proposal distributions  $q(s_t | s_{t-1}^{(i)}, z_t, x)$ .

- (a) Sample indices  $\{k_1, k_2, \dots, k_N\}$  from  $\{1, \dots, N\}$  with probabilities  $\xi_t^{(i)}$ .  
 (b) Propagate particles  $s_t^{(i)} \sim q(\cdot | s_{t-1}^{(k_i)}, z_t, x)$ .  
 (c) Weight particles  $\tilde{w}_t^{(i)} = \frac{w_{t-1}^{(k_i)} g_x(z_t | s_t^{(i)}) f_x(s_t^{(i)} | s_{t-1}^{(k_i)})}{\xi_t^{(k_i)} q(s_t^{(i)} | s_{t-1}^{(k_i)}, z_t, x)}$  and calculate  $C_t = \sum_{i=1}^N \tilde{w}_t^{(i)}$ .  
 (d) Obtain an estimate of the predictive likelihood,  $\hat{p}(z_t | z_{1:t-1}, x) = C_t/N$ , and calculate normalized weights  $w_t^{(i)} = \tilde{w}_t^{(i)}/C_t$  for  $i = 1, \dots, N$ .
- 

Implementing the particle Langevin algorithm requires an approximation of the gradient of the log posterior density  $\nabla \log \pi(x)$ , where  $\nabla \log \pi(x) = \nabla \log p(z_{1:T} | x) + \nabla \log p(x)$ . As outlined in the Section 4 we can use the Poyiadjis et al. (2011) algorithm to approximate the gradient, however, the variance of this approximation increases quadratically with  $t$ . An alternative method proposed by Nemeth et al. (2016) has been shown to produce estimates of the gradient with only linearly increasing variance. We shall use this method to create the particle Langevin proposal, details of which are as follows.

For each particle at a time  $t - 1$ , there is an associated path, defined by tracing the ancestry of each particle back in time. With slight abuse of notation denote this path by  $s_{1:t-1}^{(i)}$ . We can thus associate with particle  $i$  at time  $t - 1$  a value  $\alpha_{t-1}^{(i)} = \nabla \log p(s_{1:t-1}^{(i)}, z_{1:t-1} | x)$ . These values can be updated recursively. Remember that in step 2(b) of Algorithm 1 we sample  $k_i$ , which is the index of the particle at time  $t - 1$  that is propagated to produce the  $i$ th particle at time  $t$ . Thus we have

$$\alpha_t^{(i)} = \alpha_{t-1}^{(k_i)} + \nabla \log g_x(z_t | s_t^{(i)}) + \nabla \log f_x(s_t^{(i)} | s_{t-1}^{(k_i)}). \quad (\text{S18})$$

The main idea behind the Nemeth et al. (2016) approach is to use kernel density estimation to replace each discrete  $\alpha_{t-1}^{(i)}$  value by a Gaussian distribution:

$$\alpha_{t-1}^{(i)} \sim \mathcal{N}(m_{t-1}^{(i)}, V_{t-1}). \quad (\text{S19})$$

The mean of this distribution is obtained by shrinking  $\alpha_{t-1}^{(i)}$  towards the mean of  $\alpha_{t-1}$ ,

$$m_{t-1}^{(i)} = \zeta \alpha_{t-1}^{(i)} + (1 - \zeta) \sum_{i=1}^N w_{t-1}^{(i)} \alpha_{t-1}^{(i)}.$$

Here  $0 < \zeta < 1$  is a user-defined shrinkage parameter. The idea of this shrinkage is that it corrects for the increase in variability introduced through the kernel density estimation of West

(1993). For a definition of  $V_{t-1}$  see Nemeth et al. (2016), however, its actual value does not affect the following details.

The resulting model for the  $\alpha_t$ 's, including their updates (S18), is linear Gaussian. Hence we can use Rao–Blackwellization to avoid sampling  $\alpha_t^{(i)}$ , and instead calculate the parameters of the kernel (S19) directly. This gives the following recursion for the means,

$$m_t^{(i)} = \zeta m_{t-1}^{(k_i)} + (1 - \zeta) \sum_{i=1}^N w_{t-1}^{(i)} m_{t-1}^{(i)} + \nabla \log g_x(z_t | s_t^{(i)}) + \nabla \log f_x(s_t^{(i)} | s_{t-1}^{(k_i)}).$$

The final score estimate depends only on these means, and is

$$\nabla \log \hat{p}(z_{1:t} | x) = \sum_{i=1}^N w_t^{(i)} m_t^{(i)}.$$

See Algorithm 2 for a summary.

---

**Algorithm 2** Rao-Blackwellized Kernel Density Estimate of the Score Vector

---

Add the following steps to Algorithm 1.

*Step 1:*

(c) Set  $\nabla \log \hat{p}(z_1 | x) = \nabla \log g_x(z_1 | s_1^{(i)}) + \nabla \log \mu_x(s_1^{(i)})$ .

*Step 2:*

(e) For  $i = 1, \dots, N$ , calculate

$$m_t^{(i)} = \zeta m_{t-1}^{(k_i)} + (1 - \zeta) \sum_{i=1}^N w_{t-1}^{(i)} m_{t-1}^{(i)} + \nabla \log g_x(z_t | s_t^{(i)}) + \nabla \log f_x(s_t^{(i)} | s_{t-1}^{(k_i)}).$$

(f) Update and store the score vector

$$\nabla \log \hat{p}(z_{1:t} | x) = \sum_{i=1}^N w_t^{(i)} m_t^{(i)}.$$


---

When  $\zeta = 1$  the recursion simplifies to the method given by Poyiadjis et al. (2011), where the variance of the score estimate will increase quadratically with  $t$ . The use of a shrinkage parameter  $\zeta < 1$  alleviates the degeneracy problems that affect the estimation of the score and significantly reduces the estimate's variance. As a rule of thumb, setting  $\zeta = 0.95$  produces reliable estimates and we shall use this tuning for all examples in the Section 4. Decreasing  $\zeta$  leads to a decrease in variance, but at the cost of increasing the bias in the estimate of the gradient. Nemeth et al. (2016) have shown that reliable results can be obtained for a wide range of  $\zeta$  and that values in the range  $0.5 < \zeta < 0.99$  work particularly well.



## 5 Negative $K_{**}$

The  $K_{**}$  term, that appears in the acceptance rate for regime (2) can be negative. This term depends on the interaction between the bias in our estimate of the gradient and the curvature of the posterior. A negative value corresponds to a case where the bias in our estimate of the gradient is beneficial and actually improves the mixing of the algorithm. Intuitively, for these cases, the bias is correcting for the error in the Euler discretization of the Langevin diffusion that is used to obtain the Metropolis-adjusted Langevin proposal. A negative  $K_{**}$  value can lead to the counter-intuitive situation where increasing the step-size can sometimes increase the acceptance rate.

To see how this happens, we present a simple example. We assume that the target distribution has independent and identically distributed standard Gaussian components, and that the bias of the estimate of the gradient for the component of interest is  $b(x)/n^{1/3} = -x/n^{1/3}$ . So as to emphasize the effect that the bias is having, we will consider the case where  $\sigma^2 = \tau^2 = 0$ , so the likelihood is estimated without error, and the only error in the gradient is due to the bias. We are considering regime (2), where  $\kappa = 1/3$ .

Simple calculations give  $K = 1/4$ ,  $K_*^2 = 1$  and  $K_{**} = -1/4$ . The limiting acceptance rate is thus

$$\alpha^{(2)}(\ell, 0) = 2\Phi \left\{ -\frac{1}{2} (\ell^6/16 - \ell^4/2 + \ell^2) \right\}.$$

This limiting acceptance rate is equal to 1 either when  $\ell = 0$ , or when  $\ell = 2$ .

The Langevin dynamics for the component of interest are defined by the stochastic differential equation

$$dX_t = -\frac{1}{2}X_t dt + dB_t.$$

The standard Langevin algorithm will propose, using an Euler approximation,

$$Y = x \left( 1 - \lambda^2/2 \right) + \lambda Z, \tag{S20}$$

where  $x$  is the current value of the chain, and  $Z$  is an independent standard Gaussian random variable. The particle Langevin algorithm will have proposal

$$Y = x \left( 1 - \lambda^2/2 - \frac{\lambda^2}{2n^{1/3}} \right) + \lambda Z, \tag{S21}$$

where the difference is due to the bias in the estimate of the gradient.

Now it is straightforward to show that if  $\lambda < 1$  a proposal of the form

$$Y = x \left( 1 - \lambda^2 \right)^{1/2} + \lambda Z \tag{S22}$$

will have acceptance rate of 1, as this is the true transition density of the Langevin dynamics over a time-step of size  $-\log(1 - \lambda^2)$ .

For our asymptotic regime (2) we have  $\lambda = \ell n^{-1/6}$  and we let  $n \rightarrow \infty$ . We can expand the coefficient of  $x$  in (S22) to give

$$(1 - \lambda^2)^{1/2} = 1 - \frac{1}{2}\lambda^2 - \frac{1}{8}\lambda^4 + O(\lambda^6).$$

The standard Langevin proposal (S20) is the same as the ideal proposal (S22) up to order  $\lambda^2$ . By comparison there is a local maximum of our limiting acceptance rate at  $\ell = 2$ , and, with this scaling, the particle Langevin proposal (S21) is better as it is the same as the ideal proposal (S22) up to order  $\lambda^4$ .

## 6 Empirical analysis of assumptions pertaining to theoretical results

Our theoretical results are posited on a number of simplifying assumptions. Some, such as the shape of the target and the independence between position and the distribution of the noise in the log-target are discussed at the start of Section 3. Others, such as the asymptotic distribution of the particle filter estimates, are based on previous theory (Bérard et al., 2014) and have been investigated previously (e.g., Sherlock et al., 2015; Doucet et al., 2015). Others pertain to the estimates of the gradient of the log-target and are entirely new. In this section, we verify that many of these assumptions hold approximately for the two examples in our simulation study.

Our theoretical results also show three possible regimes, with the final regime, where the effect of the error in the gradient is negligible, being the most desirable. We describe diagnostics that relate to the regime and we use these to show that both of our simulation studies are in the desirable regime (3).

### 6.1 Noise in the log posterior density

Theorems 3.1 and 3.3 both assume that the distribution of the noise in the log posterior density is independent of the position in the parameter value,  $x$ . Corollary 3.2, and our maximin procedure, specify further that the noise is Gaussian (8) with a variance that is inversely proportional to the number of particles. These three assumptions have been made before (Doucet et al., 2015; Sherlock et al., 2015; Pitt et al., 2012); the first of them, in particular, is unlikely to hold in practice but has been found to hold approximately. The second and third are suggested by particle filter theory (Del Moral, 2004; Bérard et al., 2014). We now check these assumptions for our simulation study examples.

Figure S1 shows a histogram of the variance of the noise in the log posterior density evaluated at 100 points sampled at random from the posterior. It can be seen that the variance fluctuates by about half an order of magnitude either side of a central value. Sherlock (2016) shows that for random walk-based algorithms a key quantity of interest, the optimal scaling, is




Figure S1: Base 10 logarithm of the variance of the log posterior density at random points in the posterior for the linear Gaussian model (left panel) and the mixture of experts model (right panel).

robust to changes in the global distributions of  $V^n$  and  $W^n$ ; Figure 1 suggests a similar robustness for the particle Langevin algorithm. In moderate to high dimensions the particle Langevin algorithm can require many iterations to traverse the posterior. Provided the variance in the noise changes sufficiently slowly, the variance will appear to be approximately constant for many consecutive iterations; thus, tuning to the optimal scaling that would apply to the current variance if it were global should be close to optimal locally. Since the optimal scaling is robust to the variance it seems plausible, that, as suggested by our empirical findings, guidance from our theory may be robust to (sufficiently slow) local variations in the distributions of  $V^n$  and  $W^n$ .

Figure S2 shows, for each of our two examples, kernel density estimates of the log posterior density based on 500 point estimates at each of two points sampled from the posterior. The noise in the log posterior density is, at least approximately, Gaussian. This is an important check as the theory that predicts a Gaussian distribution is based upon the use of a large number of particles, but for the linear Gaussian and mixture of experts models we needed respectively only 20 and 100 particles.

For each example, Figure S3 plots an estimate of the logged-variance (obtained using 500 repeated estimations of the log posterior density for each number of particles) evaluated at the same random point in the posterior against the logged number of particles. The straight line has gradient  $-1$  and shows that the variance is indeed inversely proportional to the number of particles.

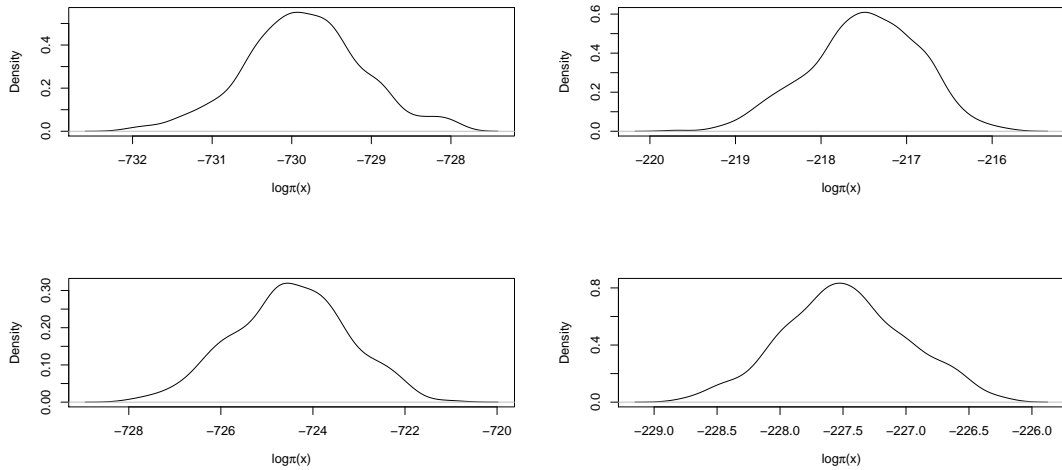


Figure S2: Empirical log posterior density, taken at two random points in the posterior, for the linear Gaussian model (left panels) and the mixture of experts model (right panels).




Figure S3: Number of particles against the empirical variance, taken at a random point in the posterior, for the linear Gaussian model (left panel) and the mixture of experts model (right panel). The diagonal lines have slope  $-1$ .

## 6.2 Noise in the estimate of the gradient

Theorem 3.3 allows for an error in the estimate of a given component of the gradient in the log posterior density. The variance of this error is assumed to be independent of position and the error is assumed to be independent of the error in the estimate of the log posterior density. This independence is by no means certain since both estimates are created from the same run of a particle filter.

To test these assumptions, in each of our two scenarios in Section 4, the linear Gaussian and mixture of experts examples, we sampled 100 points independently from the posterior.




Figure S4: Estimates of the log posterior density against the first and second components of the gradient of the log posterior density for the linear Gaussian model (left panel) and mixture of experts model (right panel).

For each of these points we ran the particle filter 500 times, creating 500 estimates of the log posterior density and 500 estimates of the gradient of the log posterior density.

Figure S4 plots, for one of these points in the posterior, the estimate of the log posterior density against the first and second components of the estimate of  $\nabla \log \pi$ . This lack of any visible pattern was repeated over the remaining 4 and 8 components of the linear Gaussian and mixture of experts models, respectively, and also over other points in the posterior.

Figure S5 presents a histogram of the variance of these estimates in the gradient for each of the 6 parameters in the linear Gaussian model. It shows that the variation in this variance across the posterior is typically of an order of magnitude or less.

Figure S6 presents kernel density estimates of the distribution of the noise in the estimates of the first two components of the gradient in the log posterior density. The shapes suggest that this density has light tails, in line with our assumption of finite moments (15). Additionally, although we did not require this, it is interesting that the noise in the gradient appears, at least approximately, to be Gaussian.

### 6.3 Regime diagnostics

Suppose for simplicity that we know precisely the log posterior density at the current value  $x$ . We then estimate the log posterior density,  $\log \hat{\pi}(x')$ , at a proposed value,  $x'$ . The change in the log posterior density,  $\Delta := \log \hat{\pi}(x') - \log \pi(x)$  can be split in to three separate contributions:

- $\Delta_A$  The change in the log posterior density that would have resulted if we had proposed a new value using the true gradient,  $\nabla \log \pi(x)$ .

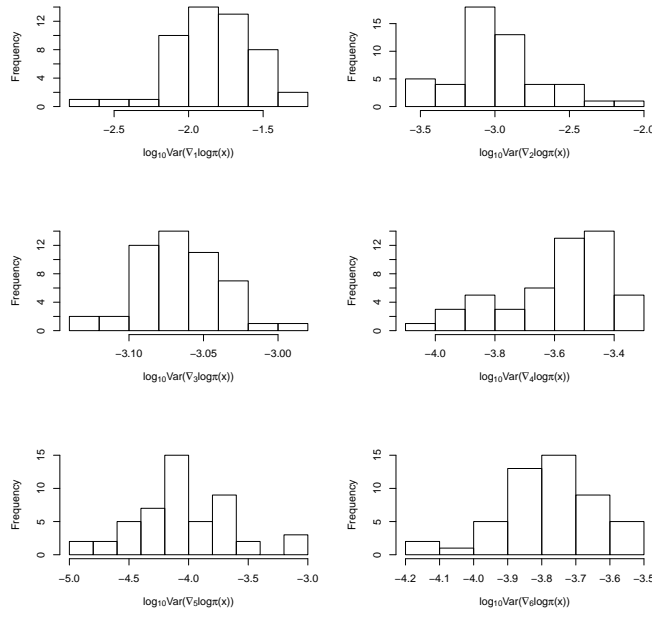


Figure S5: Histogram of log base 10 variances for each component of the gradient of the log posterior density for the linear Gaussian model.

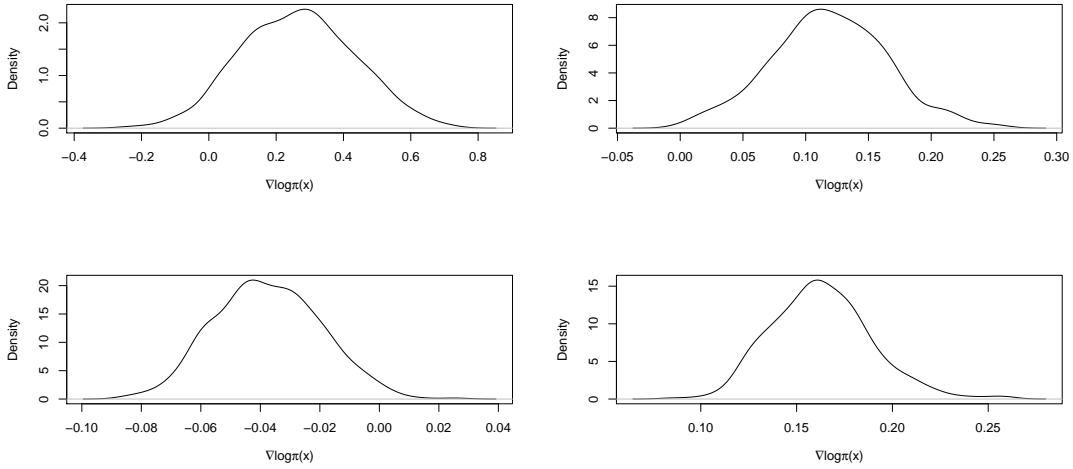


Figure S6: Empirical density of the first and fourth components of the gradient of the log posterior density, taken at a random point in the posterior, for the linear Gaussian model (left panel) and the mixture of experts model (right panel).

$\Delta_B$  The additional change in the log posterior density because we actually used an approximate gradient,  $\hat{\nabla} \log \pi(x)$ .

$\Delta_C$  The error in the log posterior density at the proposed new value.

Throughout Theorem 3.3,  $\Delta_C$  is assumed to have a variance of  $\sigma^2$  which we expect to be  $O(1)$ . In regime (1), however  $|\Delta_C| \sim |\Delta_B| \gg |\Delta_A|$ , whereas in Regime (3)  $|\Delta_C| \sim |\Delta_A| \gg |\Delta_B|$ . In Regime (2) the terms are all of similar magnitudes.

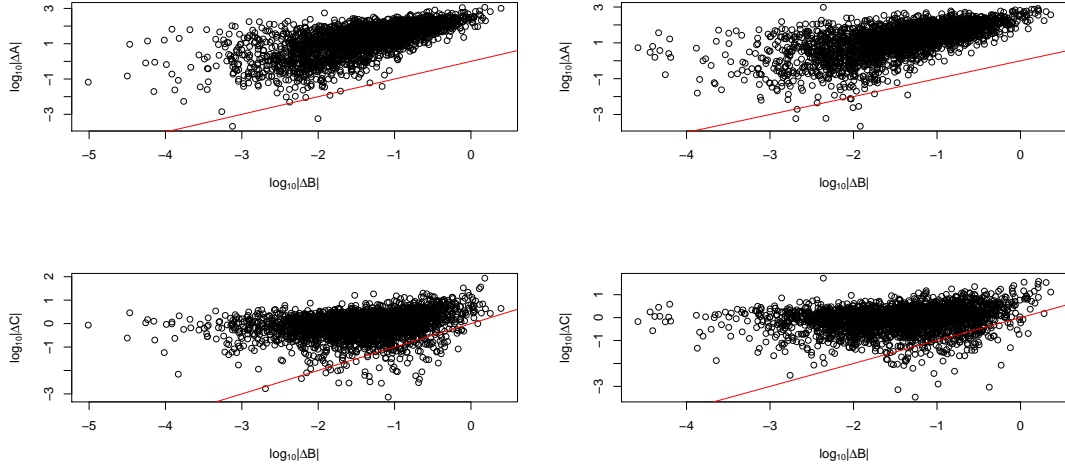


Figure S7: Regime diagnostics for the linear Gaussian model (left panel) and the mixture of experts model (right panel) in log base 10. The red line in each plot represents equality.

To be specific, define

$$\begin{aligned}
 x^* &= x + \lambda Z + \frac{\lambda^2}{2} \nabla \log \pi(x) \\
 x' &= x + \lambda Z + \frac{\lambda^2}{2} \hat{\nabla} \log \pi(x),
 \end{aligned}$$

where  $Z \sim \mathcal{N}(0, I)$ . The first proposal is the standard Metropolis-adjusted Langevin proposal where the gradient is known exactly and the second is the particle Langevin proposal. Then

$$\begin{aligned}
 \Delta_A &= \log \pi(x^*) - \log \pi(x), \\
 \Delta_B &= \log \pi(x') - \log \pi(x^*), \\
 \Delta_C &= \log \hat{\pi}(x') - \log \pi(x').
 \end{aligned}$$

For each example in Section 4, and for each of 50 points in the posterior (each representing a value of  $x$ ), we performed the following. We ran the particle filter 50 times to obtain 50 estimates,  $\hat{\nabla} \log \pi$ , and then, for the mixture of experts model, one further time with a very large number of particles to get a very accurate estimate of  $\nabla \log \pi(x)$  (for the linear Gaussian model this was calculated exactly using a Kalman filter). For each of the 50 estimates of  $\nabla \log \pi$  we also simulated a vector of Gaussian random variables  $Z$ . This lead to 50 pairs of  $(x^*, x')$  values. For each  $x^*$  and  $x'$ , for the mixture of experts model we ran the particle filter with a very large number of particles to obtain a very good estimate of the true log posterior density (for the linear Gaussian model this was obtained from the Kalman filter), we also ran the particle filter with  $N$  particles (where  $N$  is the same as in Section 4) to obtain an estimate of the log posterior density at  $x'$ . Thus for each of the 50 points we obtained 50 estimates of  $\Delta_A, \Delta_B$  and  $\Delta_C$ .

Figure S7 plots  $\Delta A$  and  $\Delta_C$  against  $\Delta_B$  for each of the 2,500 points. It can be seen from the plot that for both examples,  $|\Delta_C| \sim |\Delta_A| \gg |\Delta_B|$ , confirming empirically that we expect to be in regime (3) of Theorem 3.3.