

Efficiency of Delayed-Acceptance Random Walk Metropolis Algorithms

Chris Sherlock¹, Alexandre H. Thiery² and Andrew Golightly³

¹c.sherlock@lancaster.ac.uk, ²a.h.thiery@gmail.com and
³andrew.golightly@ncl.ac.uk

Abstract

Delayed-acceptance Metropolis-Hastings and delayed-acceptance pseudo-marginal Metropolis-Hastings algorithms can be applied when it is computationally expensive to calculate the true posterior or an unbiased stochastic approximation thereof, but a computationally cheap deterministic approximation is available. An initial accept-reject stage uses the cheap approximation for computing the Metropolis-Hastings ratio; proposals which are accepted at this stage are subjected to a further accept-reject step which corrects for the error in the approximation. Since the expensive posterior, or the approximation thereof, is only evaluated for proposals which are accepted at the first stage, the cost of the algorithm is reduced and larger scalings may be used.

We focus on the random walk Metropolis (RWM) and consider the delayed-acceptance RWM and the delayed-acceptance pseudo-marginal RWM. We provide a framework for incorporating relatively general deterministic approximations into the theoretical analysis of high-dimensional targets. Justified by diffusion-approximation arguments, we derive expressions for the limiting efficiency and acceptance rates in high dimensional settings. Finally, these theoretical insights are leveraged to formulate practical guidelines for the efficient tuning of the algorithms. The robustness of these guidelines and predicted properties are verified against simulation studies, all of which are strictly outside of the domain of validity of our limit results.

Keywords: Markov Chain Monte Carlo, Delayed-Acceptance, Pseudo-Marginal MCMC, Diffusion limit.

1 Introduction

The Metropolis-Hastings algorithm is widely used to approximately compute expectations with respect to complicated high-dimensional posterior distributions [13, 29]. The algorithm requires that it be possible to evaluate point-wise the posterior density π up to a fixed but arbitrary constant of proportionality. In many cases each such evaluation can be computationally expensive, prompting the use of a surrogate model to accelerate the computations.

The delayed-acceptance Metropolis-Hastings algorithm [4, 15, 16, 33, 38, 54, 55], also called the modified Metropolis algorithm [3, 14], preconditioned MCMC [21] and two-stage MCMC [20], and a special case of the surrogate transition method [36], assumes that the exact posterior π is available up to a constant of integration, but is computationally expensive to evaluate. This

framework is particularly relevant to the Bayesian approach to inverse problems [34, 59] where point estimations of the posterior density typically involve numerically solving sets of partial differential equations. A fast approximation is therefore employed as a first “screening” stage, with proposals which are rejected at the screening stage simply discarded. The correct posterior, π , is only evaluated for proposals which pass the screening stage. A second accept-reject step, which corrects for the error in the fast approximation, is then applied so that the desired true posterior is obtained as the limiting distribution of the Markov chain. The delayed-acceptance Metropolis-Hastings algorithm thus provides a principled method to leverage deterministic approximations to the posterior distribution in inverse problem modeling. In the sequel, we give several examples where a tenfold gain in efficiency is easily obtained by a well-tuned delayed-acceptance strategy.

The pseudo-marginal Metropolis-Hastings algorithm [2, 5] allows Bayesian inference when only an unbiased stochastic estimate of the target density, possibly up to an unknown normalisation constant, is available. The particle marginal Metropolis-Hastings algorithm [1], a special instance of the pseudo-marginal Metropolis-Hastings algorithm when the unbiased estimates are obtained by using a particle filter, is a popular method for estimating parameters in hidden Markov models [31, 35]. The existing literature on tuning the pseudo-marginal Metropolis-Hastings is reviewed in Section 2.2.

The computational expense involved in creating each unbiased stochastic estimate suggests that an initial accept-reject stage using a computationally cheap, deterministic, approximation to the posterior might be beneficial. This motivates the delayed-acceptance pseudo-marginal Metropolis-Hastings algorithm [23, 32, 43, 54, 58, 61]. Although the theoretical understanding of delayed-acceptance methods is still limited, several results are available. [55] compares the ergodicity properties of a delayed-acceptance algorithm with those of the parent MH algorithm, while [28] compares the asymptotic variance of the ergodic average from a delayed-acceptance algorithm with the variance of an importance-sampling estimator which takes as its proposal a sample from an MCMC algorithm targeting a surrogate. Historically, insights into the performance and tuning of MCMC algorithms have been obtained by examining the limiting behaviour of a rescaled version of the Markov chain as the dimension of the statespace increases to infinity [6, 8, 44, 45, 46, 57, 62, 63]. In this article, we focus on random-walk proposals since this class of methods has the advantage of not requiring further information about the target, such as the local gradient or Hessian. Thus, we concentrate on the delayed-acceptance random walk Metropolis (DARWM) and the delayed-acceptance pseudo-marginal random walk Metropolis (DAPsMRWM) algorithms: we obtain tuning and efficiency insights into these important algorithms through a diffusion-approximation.

1.1 Contributions

When an accurate approximate posterior distribution is available, the use of well-tuned DARWM and DAPsMRWM algorithms can lead to large computational savings. Unfortunately, the tuning of these methods is delicate: it involves choosing an appropriate scale for the random walk proposals and, for the DAPsMRWM, a computational budget allocated to the creation of unbiased estimates of the posterior distribution. Tuning these parameters by estimating the Effective Sample Size (ESS) is typically impractical since the ESS is notoriously difficult and computationally expensive to estimate. These tuning difficulties have hindered the adoption of these powerful methods.

We examine the efficiency of the DARWM and DAPsMRWM algorithms when exploring high-dimensional posterior distributions. We express the efficiency of the methods as a function of the scaling of the random walk proposals and, for the DAPsMRWM, of the quality and computational cost of the unbiased estimates of the posterior distribution. One of our main innovations is to

circumvent the difficulty of characterising the infinite variety of problem-specific errors in the cheap approximations to the posterior distribution by assuming that the error is a realisation of a random function – importantly, we empirically demonstrate that, in high-dimensional settings, this framework leads to robust conclusions that can be leveraged to develop efficient tuning guidelines. Under assumptions, we obtain MCMC diffusion limits through homogenization arguments. For the DAPsMRWM algorithm, we focus on a specific standard asymptotic regime which occurs for instance when the unbiased stochastic estimates are obtained through a particle filter or when using a product of importance samplers for panel data.

We imagine that a practitioner has tuned a (pseudo-marginal) RWM algorithm, found it too inefficient, and implemented a delayed-acceptance (pseudo-marginal) RWM algorithm. Our analysis shows that the relative efficiency of the optimally tuned delayed-acceptance algorithm when compared to the optimally tuned parent algorithm, as well as the relative changes in the optimal random-walk scaling and computational budget allocated to the creation of unbiased estimates, can be characterised by two parameters: **(1)** the relative computational cost of the cheap approximation compared to the cost of the posterior distribution **(2)** a measure of the accuracy of the cheap approximation involving the acceptance rate for proposals that have passed the first, screening stage. Crucially, these parameters can be estimated easily from a single additional short MCMC simulation. In practical terms, this means that once the parent algorithm (i.e. RWM or pseudo-marginal RWM) is approximately tuned, a single additional MCMC simulation is sufficient to tune the associated delayed acceptance algorithm.

Simulation studies verify different aspects of the theory, the theoretical predictions, a pivotal result (Lemma 4.1) on the relationship between changes in the posterior and changes in the deterministic approximation, and the tuning advice.

2 Delayed-acceptance Random Walks

Consider a posterior distribution $\pi(d\mathbf{x})$ on a state-space $\mathcal{X} \subseteq \mathbb{R}^d$. We assume throughout this text that π possesses a density $\pi(\mathbf{x})$ with respect to the Lebesgue measure. The Random-Walk Metropolis (RWM) updating scheme provides a general class of algorithms for obtaining approximate samples from the distribution π by constructing a Markov chain that is reversible with respect to π . Given the current value $\mathbf{x} \in \mathcal{X}$ of the Markov chain, a perturbation \mathbf{x}^* distributed as

$$\mathbf{x}^* = \mathbf{x} + \lambda \xi \tag{2.1}$$

is generated, for a standard Gaussian random variable $\xi \sim \mathbf{N}(0, \mathbf{I}_d)$ and a scale parameter $\lambda > 0$. The proposal \mathbf{x}^* is accepted with probability $\alpha(\mathbf{x}; \mathbf{x}^*) = 1 \wedge [\pi(\mathbf{x}^*)/\pi(\mathbf{x})]$. Upon acceptance, the proposal $\mathbf{x}^* \in \mathcal{X}$ becomes the next current value. Otherwise the current value is left unchanged. For a given scale parameter $\lambda > 0$, and $X \sim \pi$, the acceptance rate of the RWM algorithm is defined as

$$\alpha_{\text{rwm}}(\lambda) = \mathbb{E}[\alpha(X, X + \lambda \xi)]. \tag{2.2}$$

This setting is more general than it might appear since for any matrix $V = AA^\top$ with square A , exploration of the posterior of X' using a proposal variance matrix of $\lambda^2 V$ is equivalent to exploring $X = A^{-1}X'$ using (2.1).

2.1 Delayed-Acceptance strategies

As described in the introduction, there are many situations where π is computationally expensive to calculate while a computationally cheap approximation $\pi_a(\mathbf{x})$ to the density $\pi(\mathbf{x})$ is available and can be leveraged within MCMC schemes using the delayed-acceptance algorithm. At the k -th iteration and given the current value $\mathbf{x}_k \in \mathcal{X}$ of the parameter, the DARWM first generates a proposal $\mathbf{x}^* \in \mathcal{X}$ distributed as (2.1) and proceeds as follows.

1. **Stage-One:** compute the approximation $\pi_a(\mathbf{x}^*)$ and the screening acceptance probability $\alpha_1(\mathbf{x}_k; \mathbf{x}^*) = 1 \wedge \frac{\pi_a(\mathbf{x}^*)}{\pi_a(\mathbf{x}_k)}$. With probability $\alpha_1(\mathbf{x}_k; \mathbf{x}^*)$ proceed to Stage-Two. Otherwise set $\mathbf{x}_{k+1} = \mathbf{x}_k$ and iterate.
2. **Stage-Two:** compute the posterior distribution $\pi(\mathbf{x}^*)$ and the second stage probability $\alpha_2(\mathbf{x}_k; \mathbf{x}^*) = 1 \wedge \frac{\pi(\mathbf{x}^*) \pi_a(\mathbf{x}_k)}{\pi(\mathbf{x}_k) \pi_a(\mathbf{x}^*)}$. With probability $\alpha_2(\mathbf{x}_k; \mathbf{x}^*)$, set $\mathbf{x}_{k+1} = \mathbf{x}^*$. Otherwise, set $\mathbf{x}_{k+1} = \mathbf{x}_k$.

This defines a Markov chain that is reversible with respect to the posterior distribution π . The Stage-One acceptance rate is defined as

$$\alpha_1(\lambda) = \mathbb{E}[\alpha_1(X, X^*)] \quad (2.3)$$

where $X \sim \pi$ and $X^* = X + \lambda \xi$. Clearly, the more accurate the approximation π_a , the higher the Stage-Two acceptance probability. The overall acceptance probability is $\alpha_{12}(\mathbf{x}_k; \mathbf{x}^*) = \alpha_1(\mathbf{x}_k; \mathbf{x}^*) \times \alpha_2(\mathbf{x}_k; \mathbf{x}^*)$, whilst the overall acceptance rate is $\alpha_{12}(\lambda) = \mathbb{E}[\alpha_{12}(X, X^*)]$. Our tuning guidelines are based of the conditional Stage-Two acceptance rate $\alpha_{2|1}(\lambda)$ defined as

$$\alpha_{2|1}(\lambda) = \frac{\alpha_{12}(\lambda)}{\alpha_1(\lambda)}. \quad (2.4)$$

Pseudo-marginal Metropolis-Hastings algorithms [2, 5] presume that it is computationally infeasible to evaluate the posterior density $\pi(\mathbf{x})$, even up to a multiplicative constant, but that it is possible to generate a positive and unbiased estimate of it: $\hat{\pi}(\mathbf{x}; \mathbf{u})$. The quantity $\mathbf{u} \in \mathcal{U}$ represents a sample from a source of randomness necessary to produce the stochastic estimate, and $\hat{\pi} : \mathcal{X} \times \mathcal{U} \rightarrow [0, \infty)$ is a deterministic function that, given $\mathbf{x} \in \mathcal{X}$ and a random sample $\mathbf{u} \in \mathcal{U}$, produces the estimate $\hat{\pi}(\mathbf{x}, \mathbf{u})$. Without loss of generality, one can assume that the auxiliary variable $\mathbf{u} \in \mathcal{U}$ is sampled from a fixed and known distribution with density $\rho(\mathbf{u})$. For any $\mathbf{x} \in \mathcal{X}$ we have that $\int \hat{\pi}(\mathbf{x}, \mathbf{u}) \rho(\mathbf{u}) d\mathbf{u} = \pi(\mathbf{x})$. The DAPsMRWM defines a Markov chain on the extended space $\mathcal{X} \times \mathcal{U}$ that can be described as follows. At the k -th iteration, given the current value $(\mathbf{x}_k, \mathbf{u}_k) \in \mathcal{X} \times \mathcal{U}$, the DAPsMRWM generates a Gaussian perturbation $\mathbf{x}^* \in \mathcal{X}$ distributed (2.1). The Stage-One screening procedure is identical to that of the DARWM. If this screening procedure is successful, a new proposal \mathbf{u}^* is generated from $\rho(\mathbf{u})$, leading to an estimate $\hat{\pi}(\mathbf{x}^*; \mathbf{u}^*)$ to the posterior distribution $\pi(\mathbf{x}^*)$. The modified Stage-Two acceptance probability reads

$$\alpha_2(\mathbf{x}_k, \mathbf{u}_k, \mathbf{x}^*, \mathbf{u}^*) = 1 \wedge \frac{\hat{\pi}(\mathbf{x}^*, \mathbf{u}^*) \pi_a(\mathbf{x}_k)}{\hat{\pi}(\mathbf{x}_k, \mathbf{u}_k) \pi_a(\mathbf{x}^*)}.$$

With probability $\alpha_2(\mathbf{x}_k, \mathbf{u}_k, \mathbf{x}^*, \mathbf{u}^*)$ one sets $(\mathbf{x}_{k+1}, \mathbf{u}_{k+1}) = (\mathbf{x}^*, \mathbf{u}^*)$. Otherwise, one sets $(\mathbf{x}_{k+1}, \mathbf{u}_{k+1}) = (\mathbf{x}_k, \mathbf{u}_k)$. Standard arguments show that the DAPsMRWM is reversible with respect to the extended density $\hat{\pi}(\mathbf{x}, \mathbf{u}) \rho(\mathbf{u})$ on $\mathcal{X} \times \mathcal{U}$. Indeed, this extended density has $\pi(\mathbf{x})$ as marginal density. Particle marginal MCMC [1] is a special case of pseudo-marginal MCMC where the unbiased estimate of the posterior is obtained using a particle filter. It has become one of the key

generic methodologies for Bayesian inference of hidden Markov models [17, 27, 31]. The conditional Stage-Two acceptance rate is again defined through (2.4) with $\alpha_1(\lambda)$ as in (2.3), but $\alpha_{12}(\lambda) = \mathbb{E}[\alpha_1(X, X^*) \alpha_2(X, U, X^*, U^*)]$ with (X, U) distributed according to the stationary distribution of the DAPsMRWM Markov chain, $X^* = X + \lambda \xi$ and $U^* \sim \rho(du)$.

2.2 Tuning of RWM algorithms

The efficiency of a given RWM algorithm varies enormously with the scale $\lambda > 0$ of the Gaussian perturbations $\mathbf{x}^* = \mathbf{x} + \lambda \xi$. Small proposed jumps lead to high acceptance rates but little movement across the state-space, whereas large proposed jumps lead to low acceptance rates and again to inefficient exploration of the state-space. Optimisation of the scale of the proposal has been tackled for various shapes of target [6, 10, 44, 46, 50, 51, 52] and has led to the following rule of thumb: choose the scale so that the acceptance rate is approximately $\hat{\alpha}_{\text{rwm}} \approx 23\%$. Although nearly all of the theoretical results are based upon limiting arguments in high dimension, the rule of thumb appears to be applicable even in relatively low dimensions [51].

In discussing the literature on optimising pseudo-marginal algorithms it is helpful to define a standard asymptotic regime which is made precise in Assumptions 1–3 in Section 3.3, where its justification and wide applicability is discussed further.

A relatively tractable lower bound on the efficiency of a pseudo-marginal Metropolis-Hastings algorithm is provided in [19]. Under the standard asymptotic regime, it is shown that the inefficiency of the bounding chain, taking into account the computational cost, is minimised when the variance of the noise in the estimated log-posterior is between 0.92^2 and 1.68^2 . In [57] the pseudo-marginal random walk Metropolis algorithm is examined under various regimes for the noise in the estimate of the posterior. Mixing efficiency is considered in terms of both limiting expected squared jump distance and the speed of a limiting diffusion, and an overall efficiency (ESJD/time) is defined, which takes into account the total computational time. Under the standard asymptotic regime, joint optimisation of this efficiency with respect to the variance of the noise in the log-target and the RWM scale parameter is considered. It is shown that the optimal scaling occurs when the acceptance rate is approximately $\hat{\alpha}_{\text{pm}} \approx 7.0\%$ and the variance of the noise in the estimate of the log-posterior is approximately $\hat{\sigma}_{\text{pm}}^2 \approx 1.82^2$. It is also noted in [57] that for the two different noise distributions considered in the article, the optimal scaling appears to be insensitive to the noise variance, and even to the distribution. This phenomenon is shown to hold across a large class of noise distributions in [53].

This article extends [57] to the corresponding delayed-acceptance algorithm, of which the DARWM is a special case. Results on limiting acceptance rates and mixing efficiency are proved, as is a diffusion limit. For the DARWM, and for the DAPsMRWM under the standard asymptotic regime, efficiency is then considered in detail, leading to the robust, practical tuning advice that we describe and demonstrate next.

2.3 Tuning the DARWM

In the interest of brevity, we focus here on tuning guidelines for the scaling parameter $\lambda_{\text{da}} > 0$ of the DARWM, leading to an estimate, $\hat{\lambda}_{\text{da}}$. The rationale for these is provided in Section 5. Analogous guidelines for tuning both the scaling and the number of particles in a DAPsMRWM algorithm are presented in the *Supplementary Material*.

Assume that a RWM algorithm targeting π has been constructed, with associated approximately optimal scaling $\hat{\lambda}_{\text{rwm}} > 0$ and corresponding empirical acceptance rate $\hat{\alpha}_{\text{rwm}}(\hat{\lambda}_{\text{rwm}})$, perhaps

found using the acceptance-rate heuristic described in Section 2.2, or perhaps by directly maximising an empirical measure of efficiency. In order to accelerate inference, the DARWM algorithm makes use of a computationally cheap approximation π_a . Our tuning guidelines for implementing the DARWM rely on two diagnostics: **(1)** the empirical conditional Stage-Two acceptance rate $\hat{\alpha}_{2|1}(\hat{\lambda}_{\text{rwm}})$, and **(2)** the empirical relative computational cost $\eta > 0$ of evaluating $\pi_a(x)$ compared to $\pi(x)$. Standard timing functions give the latter, whilst the former is the ratio of the number of proposals that were accepted at both stages to the number accepted at Stage-One. The quantity $\hat{\alpha}_{2|1}(\hat{\lambda}_{\text{rwm}})$ is a measure of the accuracy of π_a : if the approximation were perfect, this would equal one.

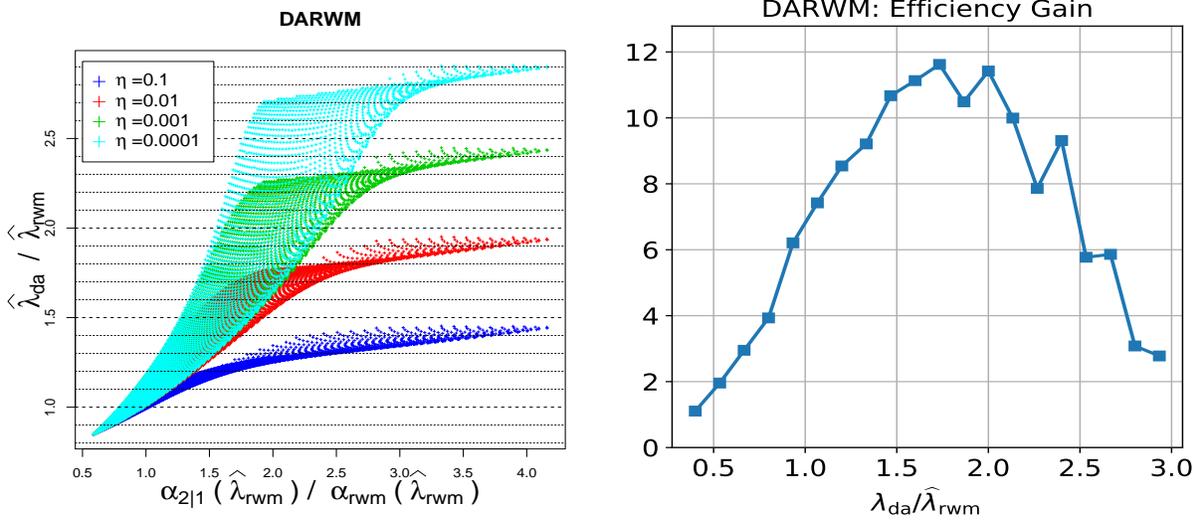


Figure 1: **Left:** Range for $\hat{\lambda}_{\text{da}}/\hat{\lambda}_{\text{rwm}}$ as a function of $\alpha_{2|1}(\hat{\lambda}_{\text{rwm}})/\alpha_{\text{rwm}}(\hat{\lambda}_{\text{rwm}})$ for different values of η . **Right:** Efficiency gains for the ODE model (2.5), defined as the ratio between the minESS/sec when running the DARWM algorithm with parameter λ_{da} , and the minESS/sec when running the optimally tuned RWM with scale parameter $\hat{\lambda}_{\text{rwm}}$.

Figure 1, the creation of which is explained in Section 5.1 and a larger version of which is given in the *Supplementary Material*, then provides a direct look up of the range for $\hat{\lambda}_{\text{da}}/\hat{\lambda}_{\text{rwm}}$ given η and $\hat{\alpha}_{2|1}(\hat{\lambda}_{\text{rwm}})/\hat{\alpha}_{\text{rwm}}(\hat{\lambda}_{\text{rwm}})$; provided that the approximation is reasonably accurate, which is when delayed acceptance is most helpful, this range is narrow. Given a tuned RWM algorithm with optimal scaling $\hat{\lambda}_{\text{rwm}}$ and acceptance rate $\hat{\alpha}_{\text{rwm}}(\hat{\lambda}_{\text{rwm}})$, the tuning of the DARWM proceeds as follows:

1. Run the DARWM with scaling $\hat{\lambda}_{\text{rwm}}$ and estimate $\alpha_{2|1}(\hat{\lambda}_{\text{rwm}})$.
2. Determine η by timing evaluations of π and π_a .
3. Set $\hat{\lambda}_{\text{da}}$ according to Figure 1 (left).

If the envelope of possible values on the y-axis is wide we suggest using the upper bound of the envelope (see the *Supplementary Material*). Moreover, in examples here and in the *Supplementary Material* the predicted optimal scaling is slightly below the true optimum. The practitioner might, therefore, wish to consider a slight upwards shift of the predicted optimal scaling.

We now demonstrate the method and check its efficacy on two examples; further examples are given in the *Supplementary Material*. Throughout, we measure empirical efficiency in terms of the minimum (over all components) number of effective samples per second, minESS/sec. Details of parameter settings, priors *etc* are provided in the *Supplementary Material*.

Example 2.1. A four-state Markov modulated Poisson Process [25] is initialised in state 1, and only transitions $1 \rightarrow 2 \rightarrow 3 \rightarrow 4 \rightarrow 1$ are allowed. The model has 8 unknown parameters. Proposed jumps are $N(0, \lambda^2 I_8)$, and the approximation, π_a is the product of the prior density and a Student- t_5 density centred at the maximum likelihood estimate (MLE) and with parameter Σ^{-1} set to the negative Hessian at the MLE.

The RWM is approximately optimised at $\hat{\lambda}_{\text{rwm}} \approx 0.11$ which gives $\hat{\alpha}_{\text{rwm}}(\hat{\lambda}_{\text{rwm}}) \approx 0.19$ and minESS/sec of 1.28. The single run of the DARWM gives $\hat{\alpha}_{2|1}(\hat{\lambda}_{\text{rwm}}) \approx 0.75$ and $1/\eta \approx 2 \times 10^4$. Figure 1 with this η and with $\hat{\alpha}_{2|1}(\hat{\lambda}_{\text{rwm}})/\hat{\alpha}_{\text{rwm}}(\hat{\lambda}_{\text{rwm}}) \approx 3.93$ suggests $\hat{\lambda}_{\text{da}}/\hat{\lambda}_{\text{rwm}} \approx 2.9$; i.e. $\hat{\lambda}_{\text{da}} \approx 0.32$.

The minESS/sec at $\hat{\lambda}_{\text{da}}$ was ≈ 17.4 , a thirteenfold improvement. A thorough grid search using long MCMC runs suggested the true optimal scaling was ≈ 0.4 , although the minESS/sec was only slightly improved, at 18.7.

Example 2.2. Consider the ordinary differential equation (ODE) in \mathbb{R}^5 :

$$\dot{x}_t = \varphi(x_t(1 - x_t) + \langle x_t, Ax_t \rangle) \quad (2.5)$$

where the product $x_t(1 - x_t)$ is to be understood component-wise, $A \in \mathbb{R}^{5,5}$ is a skew-symmetric matrix, $\langle \cdot, \cdot \rangle$ is the usual dot-product in \mathbb{R}^5 and the real-valued function $\varphi(u) \equiv 20 \cdot \arctan(u/20)$ is applied component-wise. We collect observations at discrete time $t_k = k \Delta t$ for $\Delta t \leq t_k \leq 4$ and $\Delta t = 0.2$ with additive Gaussian noise: $y_k = x(t_k) + \xi_k$ with $\xi_k \sim \mathbf{N}(0, \sigma_{\text{noise}}^2)$, with $\sigma_{\text{noise}} = 0.03$ fixed and known. For a known initial position x_0 and from the set of noisy observations $\{y_k\}_{k=1}^{19}$, we would like to infer the $d = 10$ unknown coefficients of the skew-symmetric matrix $A \in \mathbb{R}^{5,5}$. We approximate the solution of the ODE with a standard Euler discretization with step $\varepsilon = 10^{-3}$. The posterior distribution π is highly non-isotropic. Proposed jumps are $N(0, \lambda^2 \hat{\Sigma})$, where $\hat{\Sigma}$ is an approximation of the covariance matrix of π estimated from a preliminary RWM run. Optimal efficiency is obtained with $\hat{\lambda}_{\text{rwm}} \approx 0.75$, leading to an acceptance rate of $\hat{\alpha}_{\text{rwm}}(\hat{\lambda}_{\text{rwm}}) \approx 0.15$. As the computationally cheap approximate posterior distribution π_a , we simply use a coarser Euler discretization with step $\varepsilon = 10^{-1}$, leading to a speed-up of $1/\eta \approx 10^2$. Running a DARWM at $\hat{\lambda}_{\text{rwm}}$, one obtains that $\hat{\alpha}_{2|1}(\hat{\lambda}_{\text{rwm}}) \approx 0.55$. Figure 1 with this $\eta = 10^{-2}$ and with $\hat{\alpha}_{2|1}(\hat{\lambda}_{\text{rwm}})/\hat{\alpha}_{\text{rwm}}(\hat{\lambda}_{\text{rwm}}) \approx 0.55/0.15 \approx 3.6$ suggests $\hat{\lambda}_{\text{da}}/\hat{\lambda}_{\text{rwm}} \approx 1.9$; i.e. $\hat{\lambda}_{\text{da}} \approx 1.4$. This is in good agreement with the results presented in Figure 1 (right) and leads to an approximately elevenfold efficiency gain (measured in minESS/sec).

In both examples, choosing $\hat{\lambda}_{\text{rwm}}$ so that $\hat{\alpha}_{\text{rwm}}(\hat{\lambda}_{\text{rwm}}) \approx 23\%$ led to a slightly lower $\hat{\lambda}_{\text{da}}$ and a minESS/sec approximately 80% of the optimum.

3 High dimensional regime

In this section we introduce the high-dimensional asymptotic regime to be analysed in Sections 4 and 5. In Section 3.1, the target distributions are described. In Section 3.2 and 3.3 respectively, we introduce the deterministic and stochastic approximation to the target distribution and the associated notations.

3.1 Product form target distributions

We consider target densities that have a simple product form. A research program along these lines was initiated in the pair of papers [44, 45]. Although only simple exchangeable product form targets were considered, a range of subsequent theoretical analyses confirmed that the results obtained in these articles also hold for more complex target distributions, such as products of one-dimensional distributions with different variances and elliptically symmetric distributions [6, 12, 46, 51, 52]. We consider a target distribution $\pi^{(d)}(d\mathbf{x})$ in \mathbb{R}^d with a density $\pi^{(d)}(\mathbf{x})$ with respect to the Lebesgue measure that can be expressed as

$$\pi^{(d)}(\mathbf{x}) = \pi^{(d)}(x_1, \dots, x_d) = \prod_{i=1}^d \pi(x_i) \quad (3.1)$$

for a one-dimensional density $\pi(x) \equiv \pi^{(1)}(x)$ on the real line. Furthermore, we only consider the random walk algorithm with Gaussian perturbations. For a current position $\mathbf{x} \in \mathbb{R}^d$, the proposal \mathbf{X}^* is distributed as

$$\mathbf{X}^* = \mathbf{x} + \lambda^{(d)} \mathbf{Z}^{(d)} \quad \text{with} \quad \lambda^{(d)} = \mu d^{-1/2} / I \quad (3.2)$$

and a standard centred Gaussian random variable $\mathbf{Z}^{(d)}$ and a tuning parameter $\mu > 0$. The target-dependent coefficient $I > 0$ is given by

$$I^2 = \mathbb{E} [\partial_x(\log \pi)(X)^2] = -\mathbb{E} [\partial_{xx}(\log \pi)(X)] \quad (3.3)$$

for a scalar random variable $X \stackrel{\mathcal{D}}{\sim} \pi$. The second equality in Equation (3.3) follows from an integration by parts that is justified, for example, by the regularity Assumptions 4 described in Section 4. The constant $I > 0$ is introduced to simplify the statements of the results to follow. The scaling $d^{-1/2}$ ensures that, in the high-dimensional regime $d \rightarrow \infty$ the mean acceptance probability of a standard Random Walk Metropolis algorithm with proposals (3.2) and target distribution (3.1) stays bounded away from zero and one. Under mild assumptions, this scaling is optimal [6, 10, 37, 44].

3.2 Deterministic approximation

To circumvent the difficulty of characterising the infinite variety of problem-specific errors in the cheap approximations π_a to the posterior distribution π , we model the discrepancy $s(\mathbf{x}) = \log[\pi_a(\mathbf{x})/\pi(\mathbf{x})]$ as the realisation of a random function. In our setting the target distribution is a d -dimensional product of one-dimensional distributions and we imagine that each of the terms in this product is approximated through an independent realisation of a random function. Thus, the deterministic approximation $\pi_a^{(d)}(\mathbf{x}) = \pi^{(d)}(\mathbf{x}) \times \exp(s^{(d)}(\mathbf{x}))$ to the posterior density $\pi^{(d)}(\mathbf{x})$ has a deterministic error, on a logarithmic scale:

$$s^{(d)}(\mathbf{x}) = \sum_{i=1}^d \mathcal{S}(x_i, \gamma_i), \quad (3.4)$$

where $\{\gamma_i\}_{i \geq 1}$ is the realisation of an i.i.d sequence of auxiliary random variables $\{\Gamma_i\}_{i \geq 1}$. Without loss of generality, we can assume that these auxiliary random variables are uniformly distributed on the interval $[0, 1]$. We assume that the deterministic function $\mathcal{S} : \mathbb{R} \times [0, 1] \rightarrow \mathbb{R}$ in Equation (3.4)

satisfies the regularity Assumptions 4 stated below. The following two properties of the function \mathcal{S} directly influence the limiting efficiency of the delayed acceptance algorithm,

$$\beta_1 = \mathbb{E}[\partial_{xx}\mathcal{S}(X, \Gamma)]/I^2 \quad \text{and} \quad \beta_2 = \{\mathbb{E}[\partial_x\mathcal{S}(X, \Gamma)^2]/I^2\}^{1/2}, \quad (3.5)$$

where expectation is taken over two independent random variables $\Gamma \stackrel{\mathcal{D}}{\sim} \text{Uniform}([0, 1])$ and $X \stackrel{\mathcal{D}}{\sim} \pi$. Equation (3.5) and an integration by parts give that $I^2|\beta_1| = |\mathbb{E}[\partial_{xx}\mathcal{S}(X, \Gamma)]| = |\mathbb{E}[\partial_x(\log \pi)(X) \partial_x\mathcal{S}(X, \Gamma)]|$. The Cauchy-Schwarz inequality and the definition (3.3) of the quantity $I > 0$ then give that the coefficients β_1, β_2 satisfy the inequality

$$-\beta_2 \leq \beta_1 \leq \beta_2. \quad (3.6)$$

3.3 Stochastic approximation

We now describe our modeling assumptions on the stochastic approximations to the target distributions π^d . For modeling purposes, it is more natural to express the stochastic approximation on a logarithmic scale and define the new quantity $w \in \mathbb{R}$ as

$$\widehat{\pi}^{(d)}(\mathbf{x}, \mathbf{u}) = \pi^{(d)}(\mathbf{x}) e^w.$$

In other words, $w = \log[\widehat{\pi}^{(d)}(\mathbf{x}, \mathbf{u})/\pi^{(d)}(\mathbf{x})]$. For a given value of $\mathbf{x} \in \mathbb{R}^d$, the distribution of the quantity $\log[\widehat{\pi}^{(d)}(\mathbf{x}, \mathbf{U}^*)/\pi^{(d)}(\mathbf{x})]$, where $\mathbf{U}^* \sim \rho(d\mathbf{u})$, is denoted as $\pi_{W^*|X^*}(w^*|\mathbf{x})$. The Markov-Chain $\{(\mathbf{x}_k, \mathbf{u}_k)\}_{k \geq 0}$ on $\mathbb{R}^d \times \mathcal{U}$ can equivalently be described as a Markov Chain $\{(\mathbf{x}_k, w_k)\}_{k \geq 0}$ on $\mathbb{R}^d \times \mathbb{R}$. At the k -th iteration, the proposal $(\mathbf{x}_k, \mathbf{u}_k) \mapsto (\mathbf{x}^*, \mathbf{u}^*)$ is equivalently expressed as $(\mathbf{x}_k, w_k) \mapsto (\mathbf{x}^*, w^*)$ where, conditionally upon $\mathbf{x}^* \in \mathbb{R}^d$, the proposal $w^* \in \mathbb{R}$ is distributed as $\pi_{W^*|X^*}(w^*|\mathbf{x}^*)$. The property $\int \widehat{\pi}^{(d)}(\mathbf{x}, \mathbf{u}) \rho(\mathbf{u}) d\mathbf{u} = \pi(\mathbf{x})$ means that for any $\mathbf{x} \in \mathbb{R}^d$ we have $\int e^{w^*} \pi_{W^*|X^*}(w^*|\mathbf{x}) dw^* = 1$. Furthermore, since the Markov Chain $\{(\mathbf{x}_k, \mathbf{u}_k)\}_{k \geq 0}$ is reversible with respect to the density $\widehat{\pi}^{(d)}(\mathbf{x}, \mathbf{u}) \rho(\mathbf{u})$, one can check that the Markov Chain $\{(\mathbf{x}_k, w_k)\}_{k \geq 0}$ is reversible with respect to the density

$$\pi^{(d)}(\mathbf{x}) e^w \pi_{W^*|X^*}(w|\mathbf{x}). \quad (3.7)$$

In the remainder of this article, we write $\alpha_2(\mathbf{x}; w; \mathbf{x}^*; w^*)$ for denoting the Stage-Two acceptance probability of the proposal $(\mathbf{x}, w) \rightarrow (\mathbf{x}^*, w^*)$.

Standard Asymptotic Regime: similarly to the articles [19, 41, 57], we adopt the following three assumptions. These modeling assumptions constitute the standard asymptotic regime alluded to in Section 2.2.

Assumptions 1. The distribution of the additive noise W^* in the estimated log-target is independent of the proposal value itself. There exists a density $\pi_{W^*}(w^*)$ such that, for any $\mathbf{x}^* \in \mathbb{R}^d$, $\pi_{W^*|X^*}(w^*|\mathbf{x}^*) = \pi_{W^*}(w^*)$.

An asymptotic argument justifying this assumption for panel data, where the unbiased estimate is obtained from a product of importance-sampling estimates, and hidden-Markov models, where it is obtained from a particle filter, using the posterior concentration as the number of observations increases to infinity is given in [49]. It follows from Equation (3.7) that the Markov Chain

$\{(\mathbf{x}_k, w_k)\}_{k \geq 0}$ is reversible with respect to distribution $\pi^{(d)} \otimes \pi_W$ where the real valued distribution π_W is given by the change of probability

$$\frac{d\pi_W}{d\pi_{W^*}}(w) = \exp(w). \quad (3.8)$$

This is Lemma 1 of [41]. In Section 5, we examine the behaviour of the algorithm under the following Gaussian assumption.

Assumptions 2. In addition to being independent of the proposal, \mathbf{X}^* , the additive noise in the estimated log-target at the proposal, W^* , is Gaussian:

$$W^* \stackrel{\mathcal{D}}{\sim} \mathbf{N}(-\sigma^2/2, \sigma^2). \quad (3.9)$$

In Equation (3.9) the mean is determined by the variance so as to give an unbiased estimate of the posterior, $\mathbb{E}[\exp(W^*)] = 1$. It follows from (3.8) that at stationarity, under Assumptions 2, we have

$$W \stackrel{\mathcal{D}}{\sim} \mathbf{N}(\sigma^2/2, \sigma^2). \quad (3.10)$$

This article focuses on algorithms where the stochastic approximation to the likelihood is computationally expensive. In most scenarios of interest [26, 31, 32, 35] the stochastic approximation is obtained through Monte-Carlo methods (e.g. importance sampling, particle filter) that converge at the standard $N^{-1/2}$ rate where N designates the number of samples/particles used. To take into account the computational costs necessary to produce a stochastic estimate of the target-density, we thus assume the following in the rest of this article.

Assumptions 3. The computational time required to obtain an estimate of the log-target density with variance σ^2 is: (compute-time) $\propto \sigma^{-2}$.

The article [9] shows that for state-space models (and panel data) the unbiased estimate of the likelihood obtained from standard particle methods [18] (or a product of importance sampling estimators) satisfies a log-normal central limit theorem, as the number of observations and particles (or importance samples) goes to infinity, if this number is of the same order as the number of noisy observations. This justifies the Gaussian approximation (3.9) and shows that the log-error is asymptotically inversely proportional to the number of particles used, justifying Assumptions 3. The article [56] studies the tuning of pseudo-marginal MCMC methods when Assumptions 3 is not appropriate.

4 Asymptotic analysis

In this section we investigate the behaviour of the DAPsMRWM, and hence of the DARWM as a special case, in the high-dimensional regime described in Section 3. We make the following regularity assumptions.

Assumptions 4. The density $\pi : \mathbb{R} \rightarrow (0, \infty)$ and the function $\mathcal{S} : \mathbb{R} \times [0, 1] \rightarrow \mathbb{R}$ satisfy the following.

1. The function $x \mapsto \log \pi$ is thrice differentiable, with second and third derivative bounded and $\mathbb{E}[(\partial_x \log \pi)^2(X)]$ is finite, for $X \stackrel{\mathcal{D}}{\sim} \pi$.

2. The first three derivatives with respect to the first argument of the function $(x, \gamma) \mapsto \mathcal{S}(x, \gamma)$ exist and are bounded over $(x, \gamma) \in \mathbb{R} \times [0, 1]$.

Assumptions 4 are used to control the behaviour of second-order Taylor expansions; they could be relaxed in several directions at the costs of increasing technicality in the proofs. When the current position of the algorithm is $(\mathbf{x}^{(d)}, w^{(d)}) \in \mathbb{R}^d \times \mathbb{R}$, a proposal $(\mathbf{x}^{(d),*}, w^{(d),*})$ is generated, distributed as

$$\mathbf{X}^{(d),*} = \mathbf{x}^{(d)} + \mu d^{-1/2} \mathbf{Z}^{(d)} / I \quad \text{and} \quad W^{(d),*} \stackrel{\mathcal{D}}{\sim} \pi_{W^*}, \quad (4.1)$$

where $\mathbf{Z}^{(d)} \sim \mathbf{N}(\mathbf{0}, \mathbf{I}_d)$. Since the following quantities repeatedly appear in the analysis to follow, we set

$$\begin{cases} q_{\Delta}^{(d)}(\mathbf{x}^{(d)}, \mathbf{x}^{(d),*}) = \log [\pi^{(d)}(\mathbf{x}^{(d),*}) / \pi^{(d)}(\mathbf{x}^{(d)})], \\ s_{\Delta}^{(d)}(\mathbf{x}^{(d)}, \mathbf{x}^{(d),*}) = s^{(d)}(\mathbf{x}^{(d),*}) - s^{(d)}(\mathbf{x}^{(d)}). \end{cases} \quad (4.2)$$

The following lemma, proved in Section E.1 of the *Supplementary Material*, is pivotal to our analysis of the DAPsMRWM algorithm. It shows that the quantities defined in (4.2) converge jointly to a Gaussian distribution whose parameters can be expressed in terms of the scaling $\mu > 0$ of the RWM perturbations, as well as the parameters $(\beta_1, \beta_2) \in \mathbb{R} \times \mathbb{R}^+$ that describe the properties of the deterministic approximation to the target distribution.

Lemma 4.1. *Let the regularity Assumptions 4 hold. Let $\{\gamma_i\}_{i \geq 1}$ be a realisation of the sequence of auxiliary random variable used to described the deterministic approximation (3.4). Let $\{x_i\}_{i \geq 1}$ be the realisation of an i.i.d sequence marginally distributed as π . For $d \geq 1$, set $\mathbf{x}^{(d)} = (x_1, \dots, x_d) \in \mathbb{R}^d$ and let $\mathbf{X}^{(d),*}$ and $\mathbf{Z}^{(d)}$ be as defined in (4.1). For almost all realisations $\{x_i\}_{i \geq 1}$ and $\{\gamma_i\}_{i \geq 1}$ and $w \in \mathbb{R}$, the following limit*

$$\lim_{d \rightarrow \infty} \begin{bmatrix} q_{\Delta}^{(d)}(\mathbf{x}^{(d)}, \mathbf{X}^{(d),*}) \\ s_{\Delta}^{(d)}(\mathbf{x}^{(d)}, \mathbf{X}^{(d),*}) \end{bmatrix} = \begin{bmatrix} Q_{\Delta}^{\infty} \\ S_{\Delta}^{\infty} \end{bmatrix} \stackrel{\mathcal{D}}{\sim} \mathbf{N} \left(-\frac{\mu^2}{2} \begin{bmatrix} 1 \\ -\beta_1 \end{bmatrix}, \mu^2 \begin{bmatrix} 1 & -\beta_1 \\ -\beta_1 & \beta_2^2 \end{bmatrix} \right) \quad (4.3)$$

holds in distribution with parameters β_1 and β_2 defined in (3.5).

That the correlation is $-\beta_1/\beta_2 \in [-1, 1]$ is another manifestation of inequality (3.6). In general, Lemma 4.1 shows that if the approximating density has an average excess of (negative) curvature (i.e. $\beta_1 < 0$), the limiting random variables Q_{Δ}^{∞} and S_{Δ}^{∞} are positively correlated.

The product form Assumptions (3.1) and (3.4) from which we derive the bivariate Gaussian distribution in Lemma 4.1 are chosen for convenience. We expect the same conclusions to hold, at least approximately, in much broader settings. Detailed, empirical verification of Lemma 4.1 for the delayed-acceptance ODE Example 2.2 described in Equation (2.5) and for a delayed-accept pseudo-marginal example is provided in the *Supplementary Material*, demonstrating the robustness of the results proved in this article.

4.1 Limiting acceptance probability

Since the acceptance rates of the DAPsMRWM can be very simply expressed in terms of the quantities defined in Equation (4.2), Lemma 4.1 leads to tractable expression for the acceptance rates as $d \rightarrow \infty$. For $\mathbf{x}^{(d)} \in \mathbb{R}^d$, the Stage-One acceptance rate can be expressed as

$$\alpha_1^{(d)}(\mathbf{x}^{(d)}) = \mathbb{E} \left[F \left(q_{\Delta}^{(d)}(\mathbf{x}^{(d)}, \mathbf{X}^{(d),*}) + s_{\Delta}^{(d)}(\mathbf{x}^{(d)}, \mathbf{X}^{(d),*}) \right) \right]$$

where $F(u) = 1 \wedge \exp(u)$ is the Metropolis-Hastings accept-reject function. Similarly, for $(\mathbf{x}^{(d)}, w^{(d)}) \in \mathbb{R}^d \times \mathbb{R}$, the overall acceptance rate $\alpha_{12}^{(d)}(\mathbf{x}^{(d)}, w^{(d)})$ can be expressed as the expectation of the product

$$F\left(q_{\Delta}^{(d)}(\mathbf{x}^{(d)}, \mathbf{X}^{(d),*}) + s_{\Delta}^{(d)}(\mathbf{x}^{(d)}, \mathbf{X}^{(d),*})\right) \times F\left(W^{(d),*} - w^{(d)} - s_{\Delta}^{(d)}(\mathbf{x}^{(d)}, \mathbf{X}^{(d),*})\right)$$

where $W^{(d),*} \sim \pi_{W^*}$ is independent from all the other sources of randomness. The following proposition, whose proof directly follows from Lemma 4.1 and the dominated convergence theorem, gives the limiting values of these acceptance rates. Convergence is to be understood in the L^2 sense: for a sequence of random variables $\{V_k\}_{k \geq 0}$ and a constant $V_{\infty} \in \mathbb{R}$, the notation $\lim_{k \rightarrow \infty} V_k \stackrel{L^2}{=} V_{\infty}$ indicates that $\mathbb{E}[(V_k - V_{\infty})^2] \rightarrow 0$ as $k \rightarrow \infty$.

Proposition 4.1. *Let Assumptions 1 and 4 hold. For almost every realisation $\{\gamma_i\}_{i \geq 1}$ of the sequence of auxiliary random variables used to describe the deterministic approximation (3.4), we have*

$$\lim_{d \rightarrow \infty} \alpha_1^{(d)}(\mathbf{X}^{(d)}) \stackrel{L^2}{=} \alpha_1 \quad \text{and} \quad \lim_{d \rightarrow \infty} \alpha_{12}^{(d)}(\mathbf{X}^{(d)}, W^{(d)}) \stackrel{L^2}{=} \alpha_{12}$$

where the limiting acceptance rates are given by

$$\begin{cases} \alpha_1 = \mathbb{E}[F(Q_{\Delta}^{\infty} + S_{\Delta}^{\infty})] \\ \alpha_{12} = \mathbb{E}[F(Q_{\Delta}^{\infty} + S_{\Delta}^{\infty}) \times F(W_{\Delta} - S_{\Delta}^{\infty})] \end{cases}$$

for $(Q_{\Delta}^{\infty}, S_{\Delta}^{\infty})$ as described in (4.3) and $W_{\Delta} = W^* - W$ for $(W^*, W) \stackrel{\mathcal{D}}{\sim} \pi_{W^*} \otimes \pi_W$. The dependence of α_1 and α_{12} upon $(\mu, \beta_1, \beta_2, \pi_W)$ is implicit.

For the remainder of our discussion of acceptance rates we suppose that Assumption 2 holds: there is additive Gaussian noise in the logarithm of the stochastic approximation. We also make the dependence of the acceptance rate on the approximation parameters, β_1 and β_2 , explicit. Standard computations (e.g. Proposition 2.4 of [44]) yield that, for $\xi \sim \mathbf{N}(\mu, \sigma^2)$, we have that $\mathbb{E}[F(\xi)] = \Phi(\mu/\sigma) + \exp(\mu + \sigma^2/2) \Phi(-\sigma - \mu/\sigma)$, with $\Phi : \mathbb{R} \rightarrow (0, 1)$ the standard Gaussian cumulative distribution function. This permits straightforward evaluation of the Stage-One acceptance rate $\alpha_1(\mu; \beta_1, \beta_2) = \mathbb{E}[F(Q_{\Delta}^{\infty} + S_{\Delta}^{\infty})]$, the overall acceptance rate

$$\alpha_{12}(\mu, \sigma^2; \beta_1, \beta_2) = \mathbb{E}[F(Q_{\Delta}^{\infty} + S_{\Delta}^{\infty}) \times F(W_{\Delta} - S_{\Delta}^{\infty})], \quad (4.4)$$

as well as the ratio $\alpha_{2|1}(\mu, \sigma^2; \beta_1, \beta_2) = \alpha_{12}(\mu, \sigma^2; \beta_1, \beta_2) / \alpha_1(\mu; \beta_1, \beta_2)$, in terms of standard functions and, for α_{12} , a one-dimensional numerical integral, as detailed in the *Supplementary Material*. The limit as $\beta_1 \rightarrow 0$ and $\beta_2 \rightarrow 0$ corresponds to the case when there is no deterministic error and leads to the usual [37, 44] limiting acceptance rate of $2 \times \Phi(-\mu/2)$. For computing the limiting overall acceptance rate, note that under the Gaussian Assumption 2 we have $W_{\Delta} \stackrel{\mathcal{D}}{\sim} \mathbf{N}(-\sigma^2, 2\sigma^2)$.

The following result, proved in the *Supplementary Material*, shows that it is possible to characterise the (unknown) values of μ and σ^2 in terms of the Stage-One and the conditional Stage-Two acceptance rates.

Proposition 4.2. *Let Assumptions 1, 2 and 4 hold.*

1. For any $\beta_2 > 0$ and $\beta_1 < 1$ the Stage-One acceptance rate $\alpha_1(\mu; \beta_1, \beta_2)$ is a continuous decreasing bijection in μ from $[0, \infty)$ to $(0, 1]$.

2. For any fixed $\mu, \beta_2 > 0$ and β_1 , the conditional Stage-Two acceptance rate $\alpha_{2|1}(\mu, \sigma; \beta_1, \beta_2)$ is a decreasing bijection in σ from $[0, \infty)$ to $(0, \alpha_{2|1}(\mu, 0; \beta_1, \beta_2)]$.

For the DARWM algorithm, $W_\Delta \equiv 0$. Therefore, Equation (4.4) yields that $\alpha_{12}(\mu, 0; \beta_1, \beta_2) = \mathbb{E}[F(Q_\Delta^\infty + S_\Delta^\infty) \times F(-S_\Delta^\infty)]$, which can be evaluated via a one-dimensional numerical integral. When $\beta_1 = \beta_2^2$, which necessitates $\beta_2 \leq 1$ by (3.6), $\text{Cov}[Q_\Delta^\infty + S_\Delta^\infty, -S_\Delta^\infty] = 0$ so that the random variables $Q_\Delta^\infty + S_\Delta^\infty$ and $W_\Delta - S_\Delta^\infty$ are independent. In that case, algebra gives that

$$\alpha_{2|1}(\mu, \sigma^2; \beta_2^2, \beta_2) = 2\Phi\left(-\frac{1}{2}\sqrt{\beta_2^2\mu^2 + 2\sigma^2}\right). \quad (4.5)$$

This is the limiting acceptance probability of a pseudo-marginal RWM algorithm with a scaling of $\beta_2\mu$ and a noise variance of σ^2 [see 57]. Substituting $\sigma^2 = 0$ into (4.5), we find that for the DARWM, $\alpha_{2|1}(\mu, 0; \beta_2^2, \beta_2) = 2\Phi(-\beta_2\mu/2)$, the limiting acceptance probability for a RWM algorithm with a scaling of $\beta_2\mu$ [see 44]. In Section 5 the insights arising from this phenomenon help to motivate our approach to understanding the efficiency and tuning of the DARWM and DAPsMRWM algorithms.

4.2 Limiting expected squared jumping distance

A standard measure of efficiency [10, 50, 52] for local algorithms is the Euclidian Expected Squared Jumping Distance (ESJD); see [39, 48] for detailed discussions. Theoretical motivations for our use of the ESJD are given by the diffusion approximation described in Section 4.3. In our d -dimensional setting, it is defined as

$$\text{ESJD}^{(d)} = \mathbb{E}\left[\|\mathbf{X}_{k+1}^{(d)} - \mathbf{X}_k^{(d)}\|^2\right] \quad (4.6)$$

where the Markov chain $\left\{(\mathbf{X}_k^{(d)}, W_k^{(d)})\right\}_{k \geq 0}$ is stationary and $\|\cdot\|$ is the standard Euclidian norm.

Proposition 4.3. *Let Assumptions 1 and 4 hold and let α_{12} be the limit identified in Proposition 4.1. For almost every realisation $\{\gamma_i\}_{i \geq 1}$,*

$$\lim_{d \rightarrow \infty} \text{ESJD}^{(d)} = \alpha_{12} \times \left(\frac{\mu}{I}\right)^2 \equiv J(\mu) \quad (4.7)$$

The dependence of the limiting expected squared jumping distance $J(\mu)$ upon (β, π_W) is implicit.

4.3 Diffusion limit

We wish to prove that the DAPsMRWM algorithm in high dimensions can be well-approximated by an appropriate diffusion limit as this provides theoretical underpinning to our use of the ESJD as measure of efficiency [7, 48]. The connection between ESJD and diffusions arises because the asymptotic jumping distance $\lim_{d \rightarrow \infty} \text{ESJD}^{(d)} = J(\mu)$ is equal to the square of the limiting process's diffusion coefficient and is proportional to the drift coefficient. By a simple time change argument, the asymptotic variance of *any* Monte Carlo estimate of interest is inversely proportional to $J(\mu)$. Consequently, $J(\mu)$ becomes, at least in the high-dimensional limit $d \rightarrow \infty$, unambiguously the right quantity to optimise.

It is important to stress that the existence of the diffusion limit in this argument cannot be circumvented. MCMC algorithms which have non-diffusion limits can behave in very different

ways and ESJD may not be a natural way to compare algorithms. The main result of this section is a diffusion limit for a rescaled version $V^{(d)}$ of the first coordinate process. For time $t \geq 0$ we define the piecewise constant continuous time process

$$V^{(d)}(t) \equiv X_{\lfloor d \times t \rfloor, 1}^{(d)}. \quad (4.8)$$

with the notation $\mathbf{X}_k^{(d)} = (X_{k,1}^{(d)}, \dots, X_{k,d}^{(d)}) \in \mathbb{R}^d$. In general, the process $V^{(d)}$ is not Markovian; the next theorem shows nevertheless that in the limit $d \rightarrow \infty$ the process $V^{(d)}$ can be approximated by a Langevin diffusion.

Theorem 4.1. *Let Assumptions 1 and 4 hold. Let $T > 0$ be a finite time horizon and suppose that for all $d \geq 1$ the DAPsMRWM Markov chain starts at stationarity, $(\mathbf{X}_0^{(d)}, W_0^{(d)}) \stackrel{\mathcal{D}}{\sim} \pi^{(d)} \otimes \pi_W$. Then, as $d \rightarrow \infty$, the sequence of processes $V^{(d)}$ converges weakly to V in the Skorokhod topology on $D([0, T], \mathbb{R})$ where the diffusion process V satisfies the Langevin stochastic differential equation*

$$dV_t = \frac{1}{2} J(\mu) (\log \pi)'(V_t) dt + J^{1/2}(\mu) dB_t \quad (4.9)$$

with initial distribution $V_0 \stackrel{\mathcal{D}}{\sim} \pi$. The process B_t is a standard scalar Brownian motion.

As with Propositions 4.1 and 4.3, the Gaussian Assumption 2 is not necessary for the conclusion of Theorem 4.1 to hold. The proof can be found in Section E.3 of the *Supplementary Material*. It relies on an homogenization argument necessary to average-out the rapidly mixing \mathbf{u} -process. Theorem 4.1 shows that the rescaled first coordinate process converges to a Langevin diffusion V that is a time-change of the diffusion $d\bar{V}_t = \frac{1}{2} (\log \pi)'(\bar{V}_t) dt + dB_t$; indeed, $t \mapsto V_t$ has the same law as $t \mapsto \bar{V}_{J(\mu)t}$. This reveals that when speed of mixing is measured in terms of the number of iterations of the algorithm, the higher $J(\mu)$, the faster the mixing of the Markov chain. See [47] for a detailed discussion and rigorous results. However any measure of overall efficiency should also take into account the computational time required for each iteration of the algorithm, and this is the subject of the next section.

5 Optimising the efficiency

When examining the efficiency of a standard RWM the computational time is usually either not taken into account or is implicitly supposed to be independent of the choice of tuning parameter(s). In any delayed-acceptance scenario, the computational time depends on the number of acceptances at Stage-One; furthermore, in any pseudo-marginal setting the computational time also depends on the variance of the stochastic estimate of $\log \pi$. For this article, we measure the efficiency through a rescaled version of the expected squared jump distance,

$$(\text{Efficiency}) \equiv \frac{(\text{Expected Squared Jump Distance})}{(\text{Averaged one-step computing time})}. \quad (5.1)$$

The quantity $\mathcal{F}(\text{ESJD})/(\text{Averaged one step computing time})$, for any increasing function \mathcal{F} , is a valid measure of efficiency [e.g. 19, 32, 41, 57]; the discussion at the start of Section 4.3 reveals nonetheless, because of the diffusion approximation proved in Theorem 4.1, that (5.1) is the essentially unique measure of efficiency valid in the high-dimensional asymptotic regime considered

in this article. Proposition 4.3 shows that the limiting ESJD equals $\alpha_{12} \times (\mu/I)^2$ where I , defined in Equation (3.3), is a constant irrelevant for the optimisation of the efficiency discussed in this section; the constant also appears in the same form in the limiting ESJD for the equivalent non-delayed acceptance algorithm, and so it may also safely be ignored when calculating relative efficiencies. We examine the efficiency of the DARWM first, then move on to the DAPsMRWM.

5.1 Delayed-acceptance random walk Metropolis

For the DARWM we define an evaluation of π as taking one unit of time and define $\eta > 0$ to be the time for an evaluation of π_a : the one-step cost of a DARWM algorithm is $\eta + \alpha_1$. Following Equation (5.1) and eliminating unnecessary constants, the limiting efficiency of the DARWM can be quantified by the following efficiency functional

$$\text{Eff}_{\text{da}}(\mu) = \frac{\mu^2 \alpha_{12}(\mu, 0)}{\eta + \alpha_1(\mu)}, \quad (5.2)$$

with the dependence upon β_1 and β_2 implicit. In the limit where $\pi_a \equiv \pi$ and $\eta = 0$ the movement of the Markov chain becomes that of a RWM on π , but the efficiency reduces to μ^2 rather than the $\mu^2 \alpha_1$ of a RWM because, in this limit, only acceptances are associated with a computational cost. Using the same timescale, the efficiency of the RWM is $\text{Eff}_{\text{rwm}}(\mu) \equiv 2\mu^2 \Phi(-\mu/2)$ [44], which is optimised at $\mu = \hat{\mu}_{\text{rwm}} \approx 2.38$. We therefore define the relative efficiency of the DARWM algorithm to the optimal efficiency of the RWM:

$$\text{Eff}_{\text{da}}^{\text{rel}}(\mu) \equiv \frac{\text{Eff}_{\text{da}}(\mu)}{\text{Eff}_{\text{rwm}}(\hat{\mu}_{\text{rwm}})}. \quad (5.3)$$

In the special case of $\beta_1 = \beta_2^2$, and as investigated in and around (4.5),

$$\text{Eff}_{\text{da}}^{\text{rel}}(\mu; \beta_2^2, \beta_2) = \frac{\mu^2}{\hat{\mu}_{\text{rwm}}^2} \frac{\alpha_1(\mu; \beta_1, \beta_2) \Phi(-\mu\beta_2/2)}{(\eta + \alpha_1(\mu; \beta_1, \beta_2)) \Phi(-\hat{\mu}_{\text{rwm}}/2)}. \quad (5.4)$$

In the limit as $\eta \downarrow 0$, the efficiency is maximised at $\hat{\mu}_{\text{da}} = \hat{\mu}_{\text{rwm}}/\beta_2$, giving an overall relative efficiency of $\text{Eff}_{\text{da}}^{\text{rel}}(\hat{\mu}_{\text{da}}) = 1/\beta_2^2$. In reality, $\eta > 0$, and if μ is large enough so that $\alpha_1(\mu) \lesssim \eta$ then $\mu^2 \alpha_{12}(\mu; \beta_1, \beta_2)$ decreases rapidly with μ , as does the efficiency. This suggests that $\alpha_{2|1}(\hat{\mu}_{\text{rwm}}; \beta_1, \beta_2)$ might provide insight into the optimal scaling, $\hat{\mu}_{\text{da}}$, provided that $\eta > 0$ is also taken into account. Figure 2 shows $\alpha_{2|1}(\hat{\mu}_{\text{rwm}}; \beta_1, \beta_2)$ and $\hat{\mu}_{\text{da}}/\hat{\mu}_{\text{rwm}}$ as functions of β_1 and β_2 when $\eta = 0.01$. The shapes of the contours are almost identical (as is the shape for $\text{Eff}_{\text{da}}^{\text{rel}}(\hat{\mu}_{\text{da}})$, though not shown), indicating that *whatever the values of β_1 and β_2* , the quantity $\alpha_{2|1}(\hat{\mu}_{\text{rwm}}; \beta_1, \beta_2)$ provides information on the optimal increase in scaling, relative to the optimal scaling for the RWM, as well as the corresponding increase in efficiency. Along the line where $\beta_1 = \beta_2^2$, as predicted, at $\beta_2 = 1$, $\hat{\mu}_{\text{da}} \approx \hat{\mu}_{\text{rwm}}/\beta_2$, but, since $\eta > 0$, as β_2 decreases the optimal scaling increase less quickly than (5.4) suggests.

Figure 1 (left) is in fact a plot of $\hat{\mu}_{\text{da}}/\hat{\mu}_{\text{rwm}}$ vs $\alpha_{2|1}(\hat{\mu}_{\text{rwm}})/\alpha_{\text{rwm}}(\hat{\mu}_{\text{rwm}})$ over the fine grid of values of (β_1, β_2) used to create Figure 2. Since $\mu \propto \lambda$, from (3.2), $\hat{\mu}_{\text{da}}/\hat{\mu}_{\text{rwm}} \equiv \hat{\lambda}_{\text{da}}/\hat{\lambda}_{\text{rwm}}$; the constant of proportionality is unknown, which is why we provide a graph for the ratio. It suggests that $\alpha_{2|1}(\hat{\mu}_{\text{rwm}})$ combined with η does indeed provide information on the relative increase in scaling needed over $\hat{\mu}_{\text{rwm}}$.

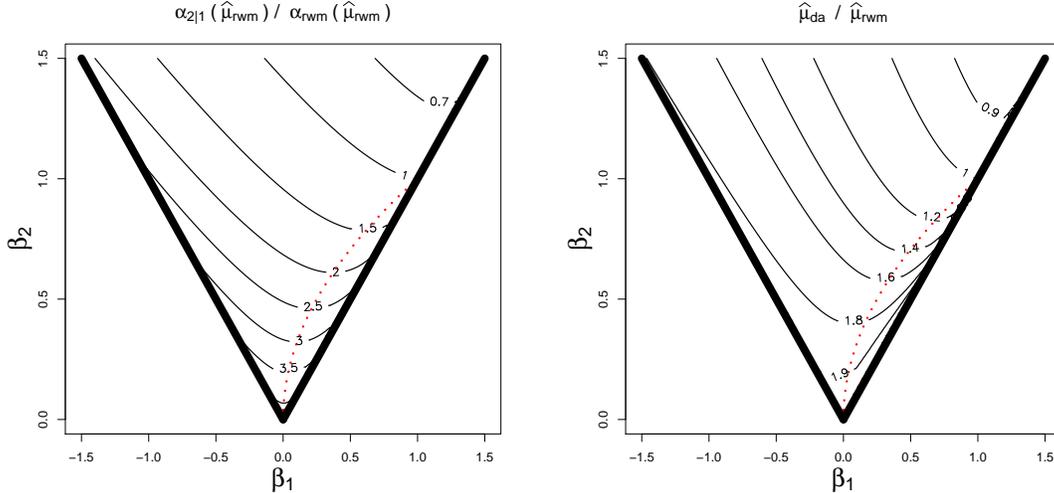


Figure 2: Contour plots of $\alpha_{2|1}(\hat{\mu}_{\text{rwm}}, 0; \beta_1, \beta_2)$ (left) and $\hat{\mu}_{\text{da}}/\hat{\mu}_{\text{rwm}}$ (right), as a function of β_1 and β_2 for $\eta = 0.01$. The red, dotted line satisfies $\beta_1 = \beta_2$.

5.2 Delayed-acceptance pseudo-marginal RWM

For the DAPsMRWM we define an evaluation of $\hat{\pi}$ with $\sigma^2 = 1$ as taking one unit of time, and $\eta > 0$ is defined to be the time for an evaluation of π_a on this scale. Under Assumption 3, the average time needed to compute the stochastic approximation is inversely proportional to the variance, σ^2 , of the estimate of the log-target, which leads to an average computational time for a single iteration of the algorithm of: $\eta + \alpha_1/\sigma^2$. As discussed in Section 3.3, Assumption 3 is reasonable when using particle MCMC to perform inference on the parameters of a hidden-Markov model, or when analysing panel data using a product of importance sampling estimators. Hence, we simplify notation and refer to the resulting efficiency as that of a Delayed-Acceptance Particle-Marginal method. Our efficiency functional is, therefore,

$$\text{Eff}_{\text{dapm}}(\mu, \sigma^2) = \frac{\mu^2 \sigma^2 \alpha_{12}(\mu, \sigma)}{\eta \sigma^2 + \alpha_1(\mu)}. \quad (5.5)$$

Theorem 5.1, which is proved in the *Supplementary Material*, shows that $\text{Eff}_{\text{dapm}}(\mu, \sigma^2)$ possesses intuitive limiting properties.

Theorem 5.1. *Let Assumptions 2, 3 and 4 hold. Then:*

1. For a fixed variance $\sigma^2 > 0$, $\text{Eff}_{\text{dapm}}(\mu, \sigma^2) \rightarrow 0$ as $\mu \rightarrow 0$ or $\mu \rightarrow \infty$.
2. For a fixed jump size $\mu > 0$, $\text{Eff}_{\text{dapm}}(\mu, \sigma^2) \rightarrow 0$ as $\sigma^2 \rightarrow 0$ or $\sigma^2 \rightarrow \infty$.

Using the same time scale as in (5.5), the equivalent efficiency function for the Particle-Marginal RWM is $\text{Eff}_{\text{pm}}(\mu, \sigma^2) \equiv 2\mu^2\sigma^2\Phi\left(-\frac{1}{2}\sqrt{\mu^2 + 2\sigma^2}\right)$, and this is maximised at $\hat{\mu}_{\text{pm}} \approx 2.562$ and $\hat{\sigma}_{\text{pm}}^2 \approx 3.283$ [57]. Thus, we define the efficiency of the DAPsMRWM relative to the maximum achievable efficiency of the Particle-Marginal RWM as:

$$\text{Eff}_{\text{dapm}}^{\text{rel}}(\mu, \sigma^2) \equiv \frac{\text{Eff}_{\text{dapm}}(\mu, \sigma^2)}{\text{Eff}_{\text{pm}}(\hat{\mu}_{\text{pm}}, \hat{\sigma}_{\text{pm}}^2)}. \quad (5.6)$$

An argument analogous to the one used for analyzing the DARWM suggests that $\alpha_{2|1}(\widehat{\mu}_{\text{pm}}\widehat{\sigma}_{\text{pm}}^2)$ and $\eta > 0$ together should be informative on $\widehat{\mu}_{\text{dapm}}$ and $\widehat{\sigma}_{\text{dapm}}^2$. Analogous contour plots to those in Figure 2, provided in the *Supplementary Material*, show the same key property. Scatter plots of $\widehat{\mu}_{\text{dapm}}$ and $\widehat{\sigma}_{\text{dapm}}^2$ against $\alpha_{2|1}(\widehat{\mu}_{\text{pm}}, \widehat{\sigma}_{\text{pm}}^2)$ partitioned by η , analogous to Figure 1 (left), are provided in the *Supplementary Material*. Again the combination of known quantities provides insight on the optimal relative tunings of the DA parameters compared with their non-DA optimal values. An equivalent efficiency plot suggests that it is not worth implementing a DAPsMRWM algorithm if π_a is only ten times faster to evaluate than $\widehat{\pi}$ is with $\sigma^2 = 1$.

As discussed in Section 2.2, an alternative tuning methodology relies on the property of the Particle-Marginal RWM algorithm that the optimal μ for a given σ^2 , $\widehat{\mu}(\sigma^2)$, is almost invariant to σ^2 [53, 57]. This effectively reduces a two-dimensional optimisation problem to two one-dimensional problems. Contour plots in the *Supplementary Material* of $\text{Eff}_{\text{dapm}}^{\text{rel}}$ as a function of μ and σ^2 for specific combinations of $\beta_2 \geq 0$, $|\beta_1| < \beta_2$ and $\eta > 0$. all show a single mode and also show that for a particular variance, the optimal scaling $\widehat{\mu}(\sigma)$ is insensitive to the value of σ , except when $\sigma \lesssim 1$, at which point the optimal scaling increases. Provided the noise variance is not made too small, therefore, μ and σ may also be tuned independently for the DAPsMRWM.

5.3 Tuning guidelines

Theorem 5.1 suggests that our goal of finding the optimal scaling $\widehat{\mu}_{\text{da}} > 0$ or, for the DAPsMRWM, $\widehat{\mu}_{\text{dapm}} > 0$ and $\widehat{\sigma}_{\text{dapm}}^2 > 0$, is sensible. However, β_1 , β_2 and I arise from an idealisation of the form of the target distribution, and the dependence of quantities of interest on these parameters arises from a limiting argument as $d \rightarrow \infty$. In reality, the quantities β_1 and β_2 and I might not exist. Even if they did exist, their values would not be known. This is why the tuning guidelines in Section 2.3 use features that appear to be approximately independent of the specific values of β_1 and β_2 , and for which I is irrelevant. In addition to the two examples in Section 2.3, a wide-ranging simulation study in the *Supplementary Material* provides further evidence for the appropriateness of these guidelines.

The *Supplementary Material* also details tuning guidelines for the both the scaling and variance of the DAPsRWM, and tests these via an extensive simulation study on a discretely observed Lotka-Volterra model [11].

6 Discussion

We have analysed the delayed-acceptance random walk Metropolis (DARWM) and delayed-acceptance pseudo-marginal random walk Metropolis algorithm (DAPsMRWM) in the limit as the dimension of the parameter space tends to infinity. The theory leads to tuning guidelines which we have verified empirically across a wide variety of scenarios.

The theoretical work also supports the intuition that, provided the cheap deterministic approximation is fast and reasonably accurate, the DAPsMRWM and DARWM algorithms should be optimally efficient when μ is much larger than (and the overall acceptance rate is much lower than) that of the equivalent (pseudo-marginal) RWM algorithm.

In some DAPsMRWM scenarios, even the cheap approximation may only be evaluated approximately, *e.g.*, via a particle filter. The expectation of this approximation would then be treated as we have treated π_a . If the noise in the logarithm of the cheap approximation is additive, then an

alternative to Proposition 4.1 can account for this via a term akin to W_Δ . In the *Supplementary Material* we describe this and investigate some consequences.

The surrogate transition method of [36] is a generalisation of delayed acceptance where, from the current position, \mathbf{x} , multiple sub-iterations of a Metropolis-Hastings kernel targeting π_a , are made. The final position of this sub-chain, \mathbf{x}^* , is a proposal that is accepted with a probability analogous to the Stage 2 acceptance probability in delayed-acceptance. For a complex target with a very cheap yet accurate approximation this offers the possibility of even greater efficiency, with a large number of sub-iterations leading to an approximation of an independence sampler proposing directly from a complex π_a . Our theory is based on a limiting diffusion obtained through Taylor expansion about the current point, \mathbf{x} , and so does not apply here, where $\|\mathbf{x}^* - \mathbf{x}\|$ can be large and the limiting process might not even be a diffusion; moreover, our theory also relies on independence between the components of $\mathbf{x}^* - \mathbf{x}$, which does not hold after multiple sub-iterations. When π_a is both cheap and accurate, intuition and our experimentation suggest that the optimal scaling for π is also a sensible scaling for π_a , whilst the number of sub-iterations should be no more than is required for approximate convergence, and fewer if π_a is relatively expensive to evaluate. In Example 2.1 an optimally scaled and iterated surrogate transition kernel was ≈ 1.7 times more efficient than an optimally scaled DARWM because a single jump proposal of the DARWM with its larger optimal scaling was equivalent to approximately 30 sub-iterations. Once π_a is available, however, the surrogate transition method requires little extra coding effort, and, hence, could be worthwhile even for such relatively small gains in efficiency. Just as with PsMRWM, we found that the optimal scaling for the DAPsRWM was insensitive to the choice of the variance of the estimator, turning a two-dimensional optimisation into two one-dimensional optimisations. It is plausible that this insensitivity might hold for the surrogate transition method, offering one possible tuning simplification in the pseudo-marginal setting.

Acknowledgements

AHT acknowledges support from the Singapore Ministry of Education Tier 2 (MOE2016-T2-2-135) and a Young Investigator Award Grant (NUSYIA FY16 P16; R-155-000-180-133).

References

- [1] C. Andrieu, A. Doucet, and R. Holenstein. Particle Markov chain Monte Carlo methods. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 72(3):269–342, 2010.
- [2] C. Andrieu and G. O. Roberts. The pseudo-marginal approach for efficient Monte Carlo computations. *Ann. Statist.*, 37(2):697–725, 2009.
- [3] Siu-Kui Au and James L. Beck. Estimation of small failure probabilities in high dimensions by subset simulation. *Probabilistic Engineering Mechanics*, 16(4):263 – 277, 2001.
- [4] Marco Banterle, Clara Grazian, Anthony Lee, and Christian P. Robert. Accelerating Metropolis-Hastings algorithms by delayed acceptance. *Foundations of Data Science*, 1(2):103–128, 2019.
- [5] M. A. Beaumont. Estimation of population growth or decline in genetically monitored populations. *Genetics*, 164:1139–1160, 2003.

- [6] M. Bédard. Weak convergence of Metropolis algorithms for non-i.i.d. target distributions. *Ann. Appl. Probab.*, 17(4):1222–1244, 2007.
- [7] M. Bédard, R. Douc, and E. Moulines. Scaling analysis of multiple-try MCMC methods. *Stoch. Proc. Appl.*, 122(3):758–786, 2012.
- [8] Mylène Bédard and Jeffrey S. Rosenthal. Optimal scaling of Metropolis algorithms: heading toward general target distributions. *Canad. J. Stat.*, 36:483–503, 2008.
- [9] J. Bérard, P. Del Moral, and A. Doucet. A lognormal central limit theorem for particle approximations of normalizing constants. *arXiv preprint arXiv:1307.0181*, 2013.
- [10] A. Beskos, G. O. Roberts, and A. Stuart. Optimal scalings for local Metropolis-Hastings chains on nonproduct targets in high dimensions. *Ann. Appl. Probab.*, 19(3):863–898, 2009.
- [11] R. J. Boys, D. J. Wilkinson, and T. B. L. Kirkwood. Bayesian inference for a discretely observed stochastic-kinetic model. *Stat. Comput.*, 18:125–135, 2008.
- [12] L. A. Breyer, M. Piccioni, and S. Scarlatti. Optimal scaling of MALA for nonlinear regression. *Ann. Appl. Probab.*, 14(3):1479–1505, 2004.
- [13] S. Brooks, A. Gelman, G. L. Jones, and X.-L. Meng, editors. *Handbook of Markov chain Monte Carlo*. Chapman & Hall/CRC Handbooks of Modern Statistical Methods. CRC Press, Boca Raton, FL, 2011.
- [14] Thomas A. Catanach and James L. Beck. Bayesian Updating and Uncertainty Quantification using Sequential Tempered MCMC with the Rank-One Modified Metropolis Algorithm. *arXiv e-prints*, page arXiv:1804.08738, Apr 2018.
- [15] J. A. Christen and C. Fox. Markov chain Monte Carlo using an approximation. *J. Comp. Graph. Stat.*, 14(4):795–810, 2005.
- [16] T Cui, C Fox, and MJ O’Sullivan. Bayesian calibration of a large-scale geothermal reservoir model by a new adaptive delayed acceptance metropolis hastings algorithm. *Water Resources Research*, 47(10), 2011.
- [17] Johan Dahlin and Thomas B. Schön. Getting Started with Particle Metropolis-Hastings for Inference in Nonlinear Dynamical Models. *arXiv e-prints*, page arXiv:1511.01707, Nov 2019.
- [18] P. Del Moral. *Feynman-Kac formulae*. Probability and its Applications (New York). Springer-Verlag, New York, 2004. Genealogical and interacting particle systems with applications.
- [19] A. Doucet, M. K. Pitt, G. Deligiannidis, and R. Kohn. Efficient implementation of Markov chain Monte Carlo when using an unbiased likelihood estimator. *Biometrika*, 102(2):295–313, 03 2015.
- [20] Y Efendiev, A Datta-Gupta, V Ginting, X Ma, and B Mallick. An efficient two-stage Markov chain Monte Carlo method for dynamic data integration. *Water Resources Research*, 41(12), 2005.
- [21] Y. Efendiev, T. Hou, and W. Luo. Preconditioning Markov chain Monte Carlo simulations using coarse-scale models. *SIAM Journal on Scientific Computing*, 28(2):776–803, 2006.

- [22] S. N. Ethier and T. G. Kurtz. *Markov processes: Characterization and convergence*, volume 6. Wiley New York, 1986.
- [23] Richard G. Everitt and Paulina A. Rowińska. Delayed acceptance ABC-SMC. *arXiv e-prints*, page arXiv:1708.02230, Aug 2017.
- [24] P. Fearnhead, V. Giagos, and C. Sherlock. Inference for reaction networks using the Linear Noise Approximation. *Biometrics*, 70:457–466, 2014.
- [25] Paul Fearnhead and Chris Sherlock. An exact Gibbs sampler for the Markov-modulated Poisson process. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 68(5):767–784, 2006.
- [26] Maurizio Filippone and Mark Girolami. Pseudo-marginal Bayesian inference for Gaussian processes. *IEEE Tran. Pattern Anal. Mach. Intell.*, 36(11):2214–2226, 2014.
- [27] T. Flury and N. Shephard. Bayesian inference based only on simulated likelihood: particle filter analysis of dynamic economic models. *Econometric Theory*, 27(05):933–956, 2011.
- [28] Jordan Franks and Matti Vihola. Importance sampling correction versus standard averages of reversible MCMCs in terms of the asymptotic variance. *Stochastic Processes and their Applications*, 2020.
- [29] W. R. Gilks, S. Richardson, and D. J. Spiegelhalter. *Markov Chain Monte Carlo in practice*. Chapman and Hall, London, UK, 1996.
- [30] D. T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.*, 81:2340–2361, 1977.
- [31] A. Golightly and D. J. Wilkinson. Bayesian parameter inference for stochastic biochemical network models using particle Markov chain Monte Carlo. *Interface Focus*, 1(6):807–820, 2011.
- [32] Andrew Golightly, Daniel A. Henderson, and Chris Sherlock. Delayed acceptance particle MCMC for exact inference in stochastic kinetic models. *Statistics and Computing*, 25(5):1039–1055, Sep 2015.
- [33] D. C. Higdon, S. J. Reese, D. Moulton, J. A. Vrugt, and C. Fox. Posterior exploration for computationally intensive forward models. In S. Brooks, A. Gelman, G. L. Jones, and X.-L. Meng, editors, *Handbook of Markov chain Monte Carlo*, chapter 16, pages 401–418. CRC Press, Boca Raton, FL, 2011.
- [34] Jari Kaipio and Erkki Somersalo. *Statistical and computational inverse problems*, volume 160. Springer Science & Business Media, 2006.
- [35] J. Knappe and P. de Valpine. Fitting complex population models by combining particle filters with Markov chain Monte Carlo. *Ecology*, 93(2):256–263, 2012.
- [36] J. S. Liu. *Monte Carlo Strategies In Scientific Computing*. Springer, 2001.
- [37] J. C. Mattingly, N. S. Pillai, and A. M. Stuart. Diffusion limits of the random walk Metropolis algorithm in high dimensions. *Ann. Appl. Probab.*, 22(3):881–930, 2012.

- [38] J. D. Moulton, C. Fox, and D. Svyatskiy. Multilevel approximations in sample-based inversion from the Dirichlet-to-Neumann map. *J. Phys.: Conf. Ser.*, 124(1), 2008.
- [39] C. Pasarica and A. Gelman. Adaptively scaling the Metropolis algorithm using expected squared jumped distance. *Statistica Sinica*, 20(1):343, 2010.
- [40] N. S. Pillai, A. M. Stuart, and A. H. Thiery. Optimal scaling and diffusion limits for the Langevin algorithm in high dimensions. *Ann. Appl. Probab.*, 22(6):2320–2356, 2012.
- [41] M. K. Pitt, R. dos Santos Silva, P. Giordani, and R. Kohn. On some properties of Markov chain Monte Carlo simulation methods based on the particle filter. *Journal of Econometrics*, 171(2):134 – 151, 2012.
- [42] Martyn Plummer, Nicky Best, Kate Cowles, and Karen Vines. Coda: Convergence diagnosis and output analysis for mcmc. *R News*, 6(1):7–11, 2006.
- [43] Matias Quiroz, Minh-Ngoc Tran, Mattias Villani, and Robert Kohn. Speeding up MCMC by delayed acceptance and data subsampling. *Journal of Computational and Graphical Statistics*, 27(1):12–22, 2018.
- [44] G. O. Roberts, A. Gelman, and W. R. Gilks. Weak convergence and optimal scaling of random walk Metropolis algorithms. *Ann. Appl. Probab.*, 7:110–120, 1997.
- [45] G. O. Roberts and J. S. Rosenthal. Optimal scaling of discrete approximations to Langevin diffusions. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 60(1):255–268, 1998.
- [46] G. O. Roberts and J. S. Rosenthal. Optimal scaling for various Metropolis-Hastings algorithms. *Statistical Science*, 16:351–367, 2001.
- [47] G. O. Roberts and J. S. Rosenthal. Complexity bounds for MCMC via diffusion limits. *arXiv preprint arXiv:1411.0712*, 2014.
- [48] G. O. Roberts and J. S. Rosenthal. Minimising MCMC variance via diffusion limits, with an application to simulated tempering. *Ann. Appl. Probab.*, 24(1):131–149, 2014.
- [49] Sebastian M. Schmon, George Deligiannidis, Arnaud Doucet, and Michael K. Pitt. Large Sample Asymptotics of the Pseudo-Marginal Method. *arXiv e-prints*, page arXiv:1806.10060, Jun 2018.
- [50] C. Sherlock. Optimal scaling of the random walk Metropolis: general criteria for the 0.234 acceptance rule. *J. App. Prob.*, 50(1):1–15, 2013.
- [51] C. Sherlock, P. Fearnhead, and G. O. Roberts. The random walk Metropolis: linking theory and practice through a case study. *Statist. Sci.*, 25(2):172–190, 2010.
- [52] C. Sherlock and G. O. Roberts. Optimal scaling of the random walk Metropolis on elliptically symmetric unimodal targets. *Bernoulli*, 15(3):774–798, 2009.
- [53] Chris Sherlock. Optimal scaling for the pseudo-marginal random walk metropolis: Insensitivity to the noise generating mechanism. *Methodology and Computing in Applied Probability*, 18(3):869–884, Sep 2016.

- [54] Chris Sherlock, Andrew Golightly, and Daniel A. Henderson. Adaptive, delayed-acceptance MCMC for targets with expensive likelihoods. *Journal of Computational and Graphical Statistics*, 26(2):434–444, 2017.
- [55] Chris Sherlock and Anthony Lee. Variance bounding of delayed-acceptance kernels. *arXiv e-prints*, page arXiv:1706.02142, Jun 2017.
- [56] Chris Sherlock, Alexandre H Thiery, and Anthony Lee. Pseudo-marginal Metropolis–Hastings sampling using averages of unbiased estimators. *Biometrika*, 104(3):727–734, 2017.
- [57] Chris Sherlock, Alexandre H. Thiery, Gareth O. Roberts, and Jeffrey S. Rosenthal. On the efficiency of pseudo-marginal random walk Metropolis algorithms. *Ann. Statist.*, 43(1):238–275, 02 2015.
- [58] M. E. Smith. Estimating nonlinear economic models using surrogate transitions. Available from https://files.nyu.edu/mes473/public/Smith_Surrogate.pdf, 2011.
- [59] Andrew M Stuart. Inverse problems: a Bayesian perspective. *Acta Numerica*, 19:451–559, 2010.
- [60] N. G. van Kampen. *Stochastic Processes in Physics and Chemistry*. North-Holland, 2001.
- [61] Matti Vihola, Jouni Helske, and Jordan Franks. Importance sampling type estimators based on approximate marginal MCMC. *arXiv e-prints*, page arXiv:1609.02541, Sep 2016.
- [62] Jun Yang, Gareth O Roberts, and Jeffrey S Rosenthal. Optimal scaling of metropolis algorithms on general target distributions. *arXiv preprint arXiv:1904.12157*, 2019.
- [63] Giacomo Zanella, Mylène Bédard, and Wilfrid S Kendall. A Dirichlet form approach to MCMC optimal scaling. *Stochastic Processes and their Applications*, 127(12):4053–4082, 2017.

This document contains the following material. Section A provides further details of the example targets used in Section 2.3. Section B provides explicit expressions for the theoretical acceptance probabilities derived in Section 4.1. Section C provides further plots derived from our theory, including enlargements of the plots required for tuning. Section D provides numerical verification of Lemma 4.1 in a DARWM (D.1) and DAPsMRWM (D.2) setting. Section E provides proofs of important results, as well as of the diffusion limit. Section F gives the proofs of several technical results, whilst Section G describes key results parallel to those in Section 4 but for the case when the error at Stage One is noisy and biased. The method typically recommends a range of tuning parameter values rather than one specific value; Section H contains an investigation of how to choose from the values in the envelope. Section I details a broad simulation study, further validating the tuning advice for the DARWM. Section J describes the tuning strategies for the DAPsRWM and details a simulation study based on noisy observations from the Lotka-Volterra model where the Linear Noise Approximation is used as π_a .

A Example targets

In this section we provide further details on the target distributions used in Section 2.3.

The Markov modulated Poisson process: consider a k -state, continuous-time Markov chain Z_t started from state 1, and a Poisson process N_t whose rate λ_t is a fixed function of Z_t . The doubly-stochastic process is parameterised by the rate matrix for the Markov chain, Q , and a vector of rates for the Poisson process, λ , where λ_i , ($i = 1, \dots, k$) is the rate of N_t when $Z_t = i$.

The event times of N_t are observed over a time window $[0, t_{\text{end}}]$, but the behaviour of Z_t is unknown, and we wish to perform inference on (Q, λ) . Setting $\Lambda = \text{diag}(\lambda)$, the likelihood for the number of events n and the event times t_1, \dots, t_n is [e.g. 25]:

$$L(Q, \lambda; t) = e' \exp[(Q - \Lambda)t_1] \Lambda \exp[(Q - \Lambda)(t_2 - t_1)] \Lambda \dots \Lambda \exp[(Q - \Lambda)(t_{\text{end}} - t_n)] 1,$$

where 1 is the k -vector of ones and $e' = (1, 0, \dots, 0)$. We simulated a dataset using a cyclic four-state Markov chain for a 200-second time window with Q parameters of: $Q_{12} = Q_{23} = Q_{41} = 1.0$, $Q_{34} = 0.25$ and all other off-diagonal rates set to zero. The rate parameters were $\lambda_1 = 20.0$, $\lambda_2 = 5.0$, $\lambda_3 = 1.0$ and $\lambda_4 = 10.0$. We then conducted inference on the natural logarithm of each parameter that was not systematically zero, placing independent $N(0, 2^2)$ priors on each of these.

The ODE model: We set $x_0 = (1, \dots, 1) \in \mathbb{R}^5$ and assume independent centred Gaussian priors with standard deviations of $\sigma_0 = 10$ on the upper-triangular part of $A \in \mathbb{R}^{5,5}$.

B Explicit expressions for the acceptance probabilities

Define $G(a, b) := \mathbb{E}[1 \wedge \exp(\mathbf{N}(a, b^2))] = \Phi(a/b) + \exp(a + b^2/2) \Phi(-b - a/b)$ with $\Phi : \mathbb{R} \rightarrow [0, 1]$ the standard Gaussian cumulative distribution function. Then

$$\alpha_1(\mu, \sigma; \beta_1, \beta_2) = G\left(-\frac{\mu^2}{2}(1 - \beta_1), \mu^2(1 + \beta_2^2 - 2\beta_1)\right). \quad (\text{B.1})$$

Further, we may rewrite

$$Q_{\Delta}^{\infty} = -\frac{1}{2}\mu^2 + \mu\frac{\beta_1}{\beta_2}\xi + \mathbf{N}\left(0, \mu^2 - \mu^2\frac{\beta_1^2}{\beta_2^2}\right) \quad \text{and} \quad S_{\Delta}^{\infty} = \frac{\beta_1}{2}\mu^2 - \mu\beta_2\xi,$$

where $\xi \sim \mathbf{N}(0, 1)$ is independent of any other source of variability. Thus, the quantity $\alpha_{12}(\mu, \sigma^2; \beta_1, \beta_2)$ can also be expressed as

$$\mathbb{E}\left[G\left(-\frac{\mu^2}{2}(1 - \beta_1) + \mu\left(\frac{\beta_1}{\beta_2} - \beta_2\right)\xi, \mu^2 - \mu^2\frac{\beta_1^2}{\beta_2^2}\right) G\left(-\frac{\beta_1}{2}\mu^2 - \sigma^2 + \mu\beta_2\xi, 2\sigma^2\right)\right]. \quad (\text{B.2})$$

C Further plots derived from the theory

Figure 3 shows contour plots of $\alpha_{2|1}(\hat{\mu}_{\text{pm}}, \hat{\sigma}_{\text{pm}}^2)$, $\hat{\mu}_{\text{dapm}}/\hat{\mu}_{\text{pm}}$, $\hat{\sigma}_{\text{dapm}}^2/\hat{\sigma}_{\text{pm}}^2$ and $\text{Eff}_{\text{rel}}(\hat{\mu}_{\text{dapm}}, \hat{\sigma}_{\text{dapm}}^2)$, as a function of β_1 and β_2 for $\eta = 0.01$. The contours all have a very similar shape, which suggests that $\alpha_{2|1}(\hat{\mu}_{\text{pm}}, \hat{\sigma}_{\text{pm}}^2)$, $\hat{\mu}_{\text{dapm}}/\hat{\mu}_{\text{pm}}$ together with η should provide information about the other three quantities.

Figure 4 shows an enlarged scatter plot of $\hat{\lambda}_{\text{da}}/\hat{\lambda}_{\text{rwm}}$ as a function of $\alpha_{2|1}(\hat{\lambda}_{\text{rwm}})/\alpha_{\text{rwm}}(\hat{\lambda}_{\text{rwm}})$. It also shows $\hat{\mu}_{\text{dapm}}/\hat{\mu}_{\text{pm}}$, $\hat{\sigma}_{\text{dapm}}^2/\hat{\sigma}_{\text{pm}}^2$ and $\text{Eff}_{\text{dapm}}^{\text{rel}}(\hat{\mu}_{\text{dapm}}, \hat{\sigma}_{\text{dapm}}^2)$, all vs $\alpha_{2|1}(\hat{\mu}_{\text{pm}}, \hat{\sigma}_{\text{pm}}^2)$, with all plots partitioned by η .

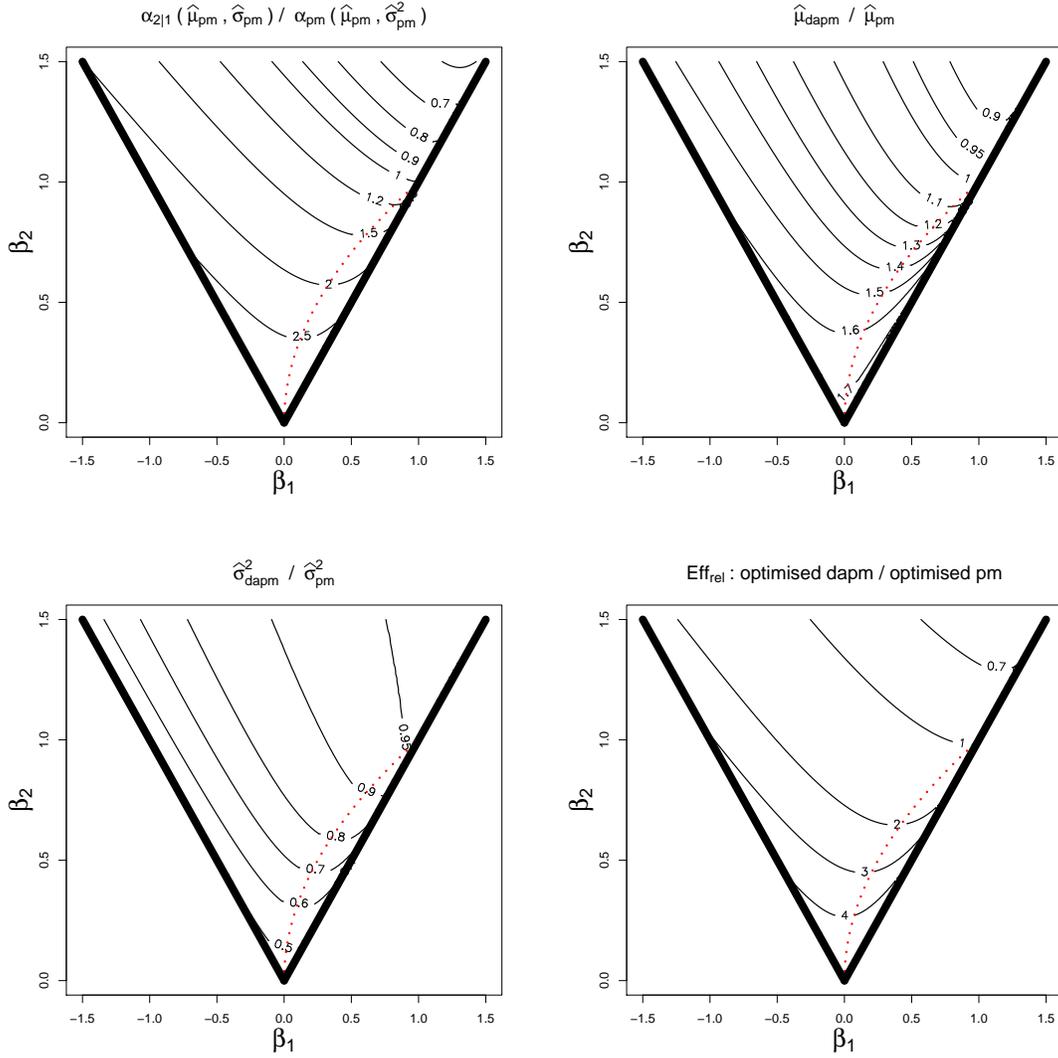


Figure 3: Contour plots of $\alpha_{2|1}(\hat{\mu}_{pm}, \hat{\sigma}_{pm}^2; \beta_1, \beta_2)$ (left), $\hat{\mu}_{dapm}/\hat{\mu}_{pm}$, $\hat{\sigma}_{dapm}^2/\hat{\sigma}_{pm}^2$ and $\text{Eff}_{rel}(\hat{\mu}_{dapm}, \hat{\sigma}_{dapm}^2)$ (right), as a function of β_1 and β_2 for $\eta = 0.01$. The red, dotted line satisfies $\beta_1 = \beta_2$.

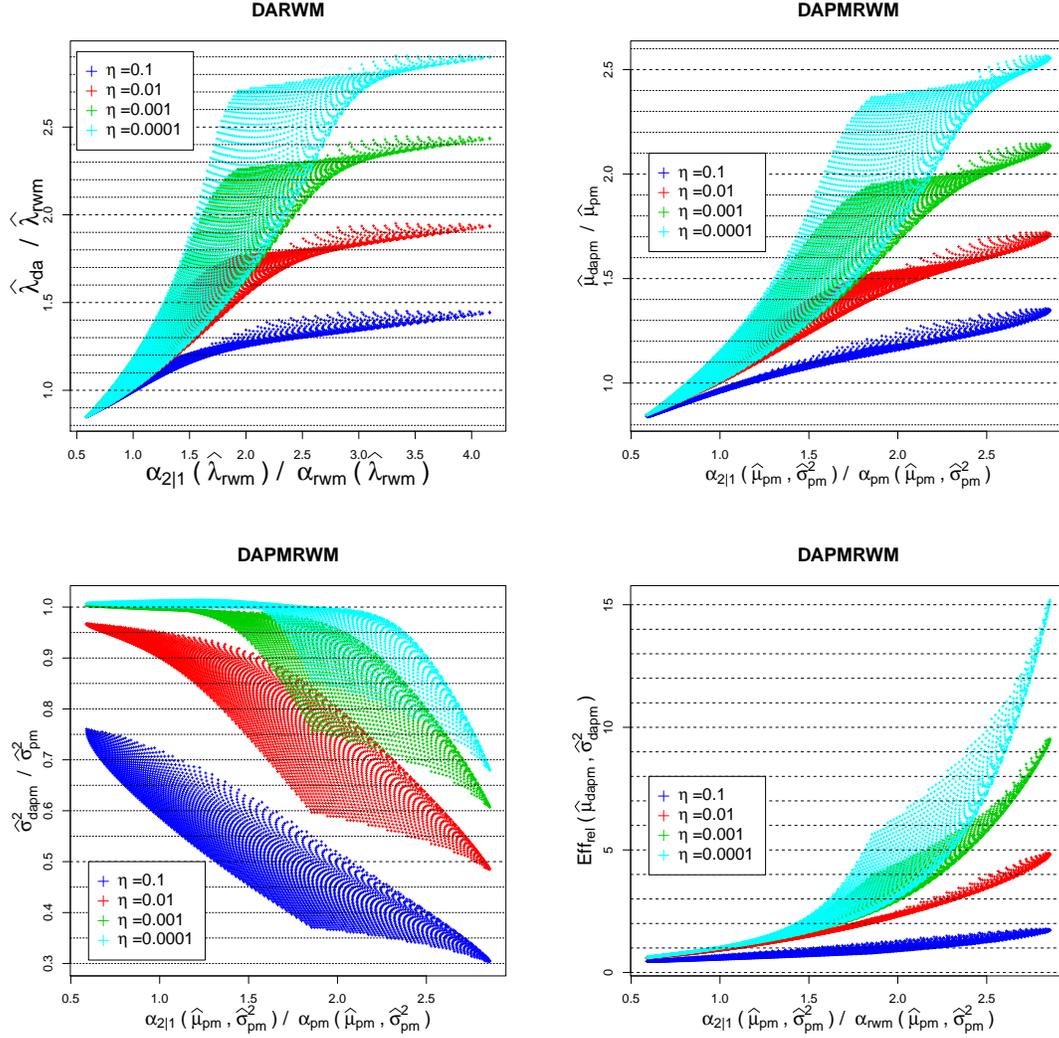


Figure 4: Scatter plots of $\hat{\lambda}_{\text{da}}/\hat{\lambda}_{\text{rwm}}$ as a function of $\alpha_{2|1}(\hat{\lambda}_{\text{rwm}})/\alpha_{\text{rwm}}(\hat{\lambda}_{\text{rwm}})$ for different values of η (top left), then $\hat{\mu}_{\text{dapm}}/\hat{\mu}_{\text{pm}}$ (top right), $\hat{\sigma}_{\text{dapm}}^2/\hat{\sigma}_{\text{pm}}^2$ (bottom left) and $\text{Eff}_{\text{dapm}}^{\text{rel}}(\hat{\mu}_{\text{dapm}}, \hat{\sigma}_{\text{dapm}}^2)$ (bottom right), vs $\alpha_{2|1}(\hat{\mu}_{\text{pm}}, \hat{\sigma}_{\text{pm}}^2)$ (right), partitioned by η .

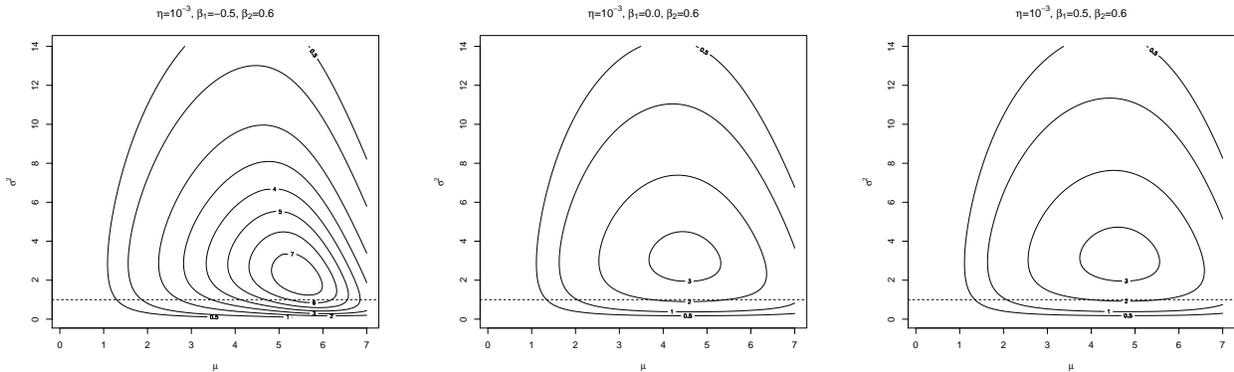


Figure 5: Contour plots of the asymptotic efficiency relative to the optimal efficiency of the equivalent pseudo-marginal RWM algorithm, $\text{Eff}_{\text{dapm}}^{\text{rel}}$, as a function of the scaling, μ , and the variance of the noise in the log-target, σ^2 , for different choices of β_1 , β_2 at $\eta = 10^{-3}$. For comparability, all contours are at 0.5, 1, 2, 3, 4, 5, 6, 7, 8, 10, 12, 15. The horizontal dashed line denotes $\sigma^2 = 1$.

Figure 5, which is typical of many other such figures that we produced, shows contour plots of $\text{Eff}_{\text{dapm}}^{\text{rel}}$ as a function of μ and σ^2 for specific combinations of $\beta_2 \geq 0$, $|\beta_1| < \beta_2$ and $\eta > 0$. Each plot shows a single mode and also shows that for a particular variance, the optimal scaling $\hat{\mu}(\sigma)$ is insensitive to the value of σ , except when $\sigma \lesssim 1$, at which point the optimal scaling increases. Provided the noise variance is not made too small, therefore, μ and σ may also be tuned independently for the DAPsMRWM.

D Validation of Lemma 4.1

The product form Assumptions (3.1) and (3.4) from which we derive the bivariate Gaussian distribution in Lemma 4.1 are chosen for convenience. We expect the same conclusions to hold, at least approximately, in much broader settings; for example, we believe that extensions of Lemma 4.1 to non i.i.d target distributions similar to those discussed in [6, 8, 10, 12, 40, 52] are possible, at the cost of much less transparent proofs. Here we verify the conclusions of Lemma 4.1 numerically in two examples from the main paper.

D.1 ODE model

First we investigate the ODE Example 2.2 described in Equation (2.5). For values of the jump scaling parameter λ associated with acceptance rates in the range $10\% \leq \alpha_{\text{rwm}}(\lambda) \leq 95\%$, we display the quantities $\text{loc}(S_\Delta)/\lambda^2$, $\text{scale}(S_\Delta)/\lambda$, $\text{loc}(S_\Delta + Q_\Delta)/\lambda^2$ and $\text{scale}(S_\Delta + Q_\Delta)/\lambda$ where $S_\Delta \equiv S(\mathbf{x}_0 + \lambda\xi) - S(\mathbf{x}_0)$ and $S(\mathbf{x}) \equiv \log[\pi(\mathbf{x})/\pi_a(\mathbf{x})]$ and $Q_\Delta \equiv \log[\pi(\mathbf{x}_0 + \lambda\xi)/\pi(\mathbf{x}_0)]$. Here, $\mathbf{x}_0 \in \mathbb{R}^{10}$ is chosen in the bulk of the distribution, the perturbation ξ is a centred Gaussian random variable with covariance approximately matching that of π . For a random variable V , $\text{loc}(V)$ denotes its median, and $\text{scale}(V) \equiv q_{75} - q_{25}$ is the inter-quartile range. We report the location and scale parameters instead of the mean and standard-deviation for increased robustness. Lemma 4.1 predicts that the reported quantities are insensitive to the value of the scale parameter λ , and hence of the acceptance rate $\alpha_{\text{rwm}}(\lambda)$. Figure 6 (right) shows that this property approximately holds true, although departure from the theory is noticeable for low acceptance rate, i.e. large

Gaussian perturbations. Figure 6 (Left) shows that the pair (Q_Δ, S_Δ) is approximately jointly Gaussian, as predicted by Lemma 4.1.

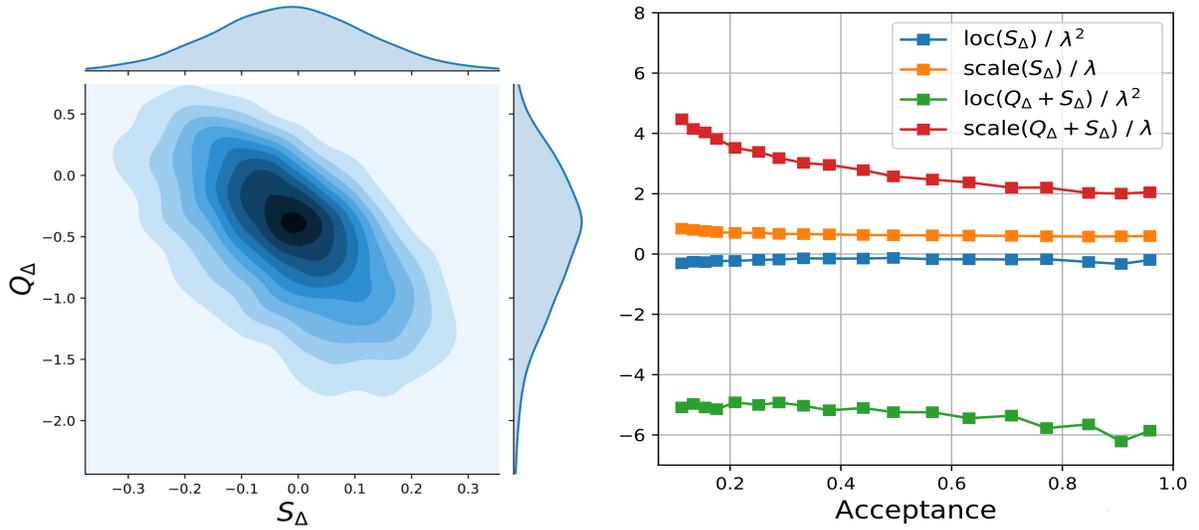


Figure 6: Empirical verification of Lemma 4.1 on the ODE Example 2.2 defined in Equation (2.5).

D.2 Lotka-Volterra model

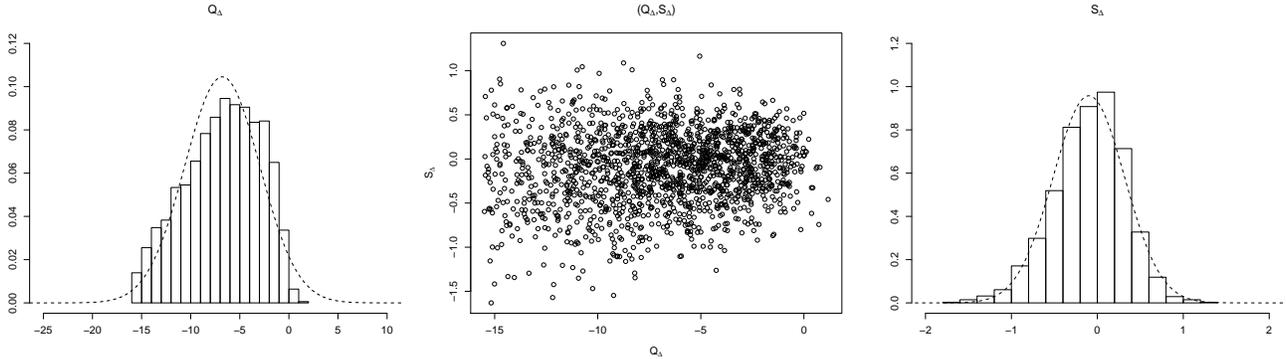


Figure 7: Empirical distribution, for the Lotka-Volterra model with LNA approximation, of $q_\Delta^{(d)}(\mathbf{x}^{(d)}, \mathbf{X}^{(d),*})$ and $s_\Delta^{(d)}(\mathbf{x}^{(d)}, \mathbf{X}^{(d),*})$ evaluated from a fixed current point in the bulk of the target distribution. The dashed lines in the left and right panels show the densities of Gaussian fits to the empirical marginal distributions of $q_\Delta^{(d)}(\mathbf{x}^{(d)}, \mathbf{X}^{(d),*})$ and $s_\Delta^{(d)}(\mathbf{x}^{(d)}, \mathbf{X}^{(d),*})$ respectively.

To validate Lemma 4.1 in the case of the five-dimensional Lotka-Volterra example we picked a point from the bulk of the posterior and used a very large number of particles to obtain an accurate estimate of the posterior, π , then ran the LNA to obtain π_a . From this point we repeatedly proposed jumps, and for each proposed point we also obtained π_a and an accurate estimate of π .

Figure 7 shows histograms of Q_Δ and S_Δ and a scatter plot of the two quantities. Because of the low dimension ($d = 5$), Q_Δ is only approximately Gaussian, although $\mathbb{E}[Q_\Delta] \approx -6.7 \approx -0.5 \times 13.7 \approx -0.5 \times \text{var}[Q_\Delta]$, which fits with the theory. The Gaussian approximation to S_Δ is much more accurate, and the further three quantities of $\mathbb{E}[S_\Delta] \approx -0.096$, $\text{Var}[S_\Delta] \approx 0.17$ and $\text{Cor}[Q_\Delta, S_\Delta] \approx 0.11$ are consistent with the two values $\beta_1 \approx -0.013$ and $\beta_2 \approx 0.11$.

E Proofs

It will be helpful to introduce i.i.d sequences $\{X_i\}_{i \geq 1}$ and $\{\Gamma_i\}_{i \geq 1}$ respectively marginally distributed as π and π_Γ , and corresponding realisations of them, $\{x_i\}_{i \geq 1}$ and $\{\gamma_i\}_{i \geq 1}$. Similarly, we consider an i.i.d sequence $\{Z_{i,k}\}_{i,k \geq 0}$ of standard Gaussian $\mathbf{N}(0, 1)$ random variables, $\{U_k\}_{k \geq 0}$ an i.i.d sequence of random variables uniformly distributed on $[0, 1]$, W a random variable distributed as π_W and $\{W_k^*\}_{k \geq 0}$ an i.i.d sequence distributed as π_{W^*} . For any dimension $d \geq 1$ we set $\mathbf{X}_0^{(d)} = (X_1, \dots, X_d) \in \mathbb{R}^d$ and $W_0^{(d)} = W$ and $X_{k,j}^{(d),*} = X_{k,j}^{(d)} + (\mu/I) d^{-1/2} Z_{k,j}$; we recursively define

$$(\mathbf{X}_{k+1}^{(d)}, W_{k+1}^{(d)}) = \begin{cases} (\mathbf{X}_k^{(d),*}, W_k^*) & \text{if } U_k < \alpha^{(d)}(\mathbf{X}_k^{(d)}, W_k^{(d)}; \mathbf{X}_k^{(d),*}, W_k^*) \\ (\mathbf{X}_k^{(d)}, W_k^{(d)}) & \text{otherwise,} \end{cases}$$

for a proposal $\mathbf{X}_d^{(d),*} = (X_{k,1}^{(d),*}, \dots, X_{k,d}^{(d),*})$. Indeed, the process $(\mathbf{X}_k^{(d)}, W_k^{(d)})$ is a DAPsMRWM Markov chain started at stationarity and targeting $\pi^{(d)} \otimes \pi_W$. We denote by \mathcal{F}_k the σ -algebra generated by the family of random variables $\{\mathbf{X}_t^{(d)}, W_t^{(d)} \mid t \leq k\}$ and use the notation $\mathbb{E}_k[\cdot]$ for designating the conditional expectation $\mathbb{E}[\cdot \mid \mathcal{F}_k]$. Similarly, we use the notation $\mathbb{E}_{\mathbf{x},w}[\cdot]$ instead of $\mathbb{E}[\cdot \mid (\mathbf{X}_0^{(d)}, W_0^{(d)}) = (\mathbf{x}, w)]$. Finally, we set

$$\begin{aligned} q_\Delta^{(d)} &= q_\Delta^{(d)}(\mathbf{x}^{(d)}, \mathbf{x}^{(d),*}), & s_\Delta^{(d)} &= s_\Delta^{(d)}(\mathbf{x}^{(d)}, \mathbf{x}^{(d),*}), & w_\Delta^{(d)} &= w^{(d),*} - w^{(d)}, \\ Q_\Delta^{(d)} &= q_\Delta^{(d)}(\mathbf{x}^{(d)}, \mathbf{X}^{(d),*}), & S_\Delta^{(d)} &= s_\Delta^{(d)}(\mathbf{x}^{(d)}, \mathbf{X}^{(d),*}), & W_\Delta^{(d)} &= W^{(d),*} - w^{(d)}, \\ Q_\Delta^{(d)} &= q_\Delta^{(d)}(\mathbf{X}^{(d)}, \mathbf{X}^{(d),*}), & S_\Delta^{(d)} &= s_\Delta^{(d)}(\mathbf{X}^{(d)}, \mathbf{X}^{(d),*}), & W_\Delta^{(d)} &= W^{(d),*} - W^{(d)}. \end{aligned} \quad (\text{E.1})$$

and use the shorthand notation $\ell(x) \equiv \log \pi(x)$.

E.1 Proof of Lemma 4.1

The Law of Large Numbers and the separability of $L^1(\pi \otimes \pi_\Gamma)$ readily yield that for almost every realisations $\{x_i\}_{i \geq 1}$ and $\{\gamma_i\}_{i \geq 1}$, the following holds,

$$\lim_{n \rightarrow \infty} n^{-1} \sum_{i=1}^n \varphi(x_i, \gamma_i) = \int \varphi(x, \gamma) (\pi \otimes \pi_\Gamma)(dx, d\gamma) \quad \text{for all } \varphi \in L^1(\pi \otimes \pi_\Gamma). \quad (\text{E.2})$$

We can thus safely assume in the remainder of this section that Equation (E.2) holds for the realisation $\{\gamma_i\}_{i \geq 1}$ of the auxiliary random variables used to describe the deterministic approximation (3.4). By the Cramer-Wold device, for proving Lemma 4.1 it suffices to establish that for any coefficient $c_Q, c_S \in \mathbb{R}$ the sequence $c_Q Q_\Delta^{(d)}(\mathbf{x}^{(d)}) + c_S S_\Delta^{(d)}(\mathbf{x}^{(d)})$ converges in law towards $c_Q Q_\Delta^\infty + c_S S_\Delta^\infty$; the boundedness assumption on the derivatives of the functions $x \mapsto \ell(x)$ and

$x \mapsto \mathcal{S}(x, u)$ and a second order Taylor expansion show that this is equivalent to proving that the sum

$$\frac{\mu}{\sqrt{I^2 d}} \sum_{i=1}^d \left\{ c_Q \ell'(x_i) + c_S \partial_x \mathcal{S}(x_i, \gamma_i) \right\} Z_i + \frac{1}{2} \frac{\mu^2}{I^2 d} \sum_{i=1}^d \left\{ c_Q \ell''(x_i) + c_S \partial_{xx} \mathcal{S}(x_i, \gamma_i) \right\} \quad (\text{E.3})$$

converges in law towards $c_Q Q_\Delta^\infty + c_S S_\Delta^\infty$. Definition (3.5) of the coefficient β_1 and β_2 yields that for almost every realisation $\{x_i\}_{i \geq 1}$ and $\{\gamma_i\}_{i \geq 1}$ we have

$$\frac{1}{I^2 d} \sum_{i=1}^d \left(\ell'(x_i)^2, \ell''(x_i), \partial_x \mathcal{S}(x_i, \gamma_i)^2, \partial_{xx} \mathcal{S}(x_i, \gamma_i), \ell'(x_i) \partial_x \mathcal{S}(x_i, \gamma_i) \right) \rightarrow \left(1, -1, \beta_2^2, \beta_1, -\beta_1 \right), \quad (\text{E.4})$$

from which the conclusion directly follows since $c_Q Q_\Delta^\infty + c_S S_\Delta^\infty$ has a Gaussian distribution with mean $\mu^2 (c_S \beta_1 - c_Q)/2$ and variance $\mu^2 (c_Q^2 + c_S^2 \beta_2^2 - 2 c_Q c_S \beta_1)$.

E.2 Proof of Proposition 4.3

The quantity $\text{ESJD}^{(d)}$ can also be expressed as

$$\text{ESJD}^{(d)} = \left(\frac{\mu^2}{I^2 d} \right) \sum_{j=1}^d \mathbb{E} \left[\left(Z_j^{(d)} \right)^2 \times F \left(Q_\Delta^{(d)} + S_\Delta^{(d)} \right) \times F \left(W_\Delta^{(d)} - S_\Delta^{(d)} \right) \right] \quad (\text{E.5})$$

$$= \frac{\mu^2}{I^2} \mathbb{E} \left[\left(Z_1^{(d)} \right)^2 \times F \left(Q_\Delta^{(d)} + S_\Delta^{(d)} \right) \times F \left(W_\Delta^{(d)} - S_\Delta^{(d)} \right) \right] \quad (\text{E.6})$$

for $Q_\Delta^{(d)}$, $S_\Delta^{(d)}$, $W_\Delta^{(d)}$ defined in (E.1); the second equality follows from the exchangeability, at stationarity, of the d coordinates of the Markov chain. One can decompose $Q_\Delta^{(d)}$ and $S_\Delta^{(d)}$ as a sum of a term that is independent of $Z_1^{(d)}$ and a negligible term; we have $Q_\Delta^{(d)} = Q_\Delta^{(d),\perp} + \log [\pi(\mathbf{X}_1^{(d),*})/\pi(\mathbf{X}_1^{(d)})]$ and $S_\Delta^{(d)} = S_\Delta^{(d),\perp} + \mathcal{S}(\mathbf{X}_1^{(d),*}, \gamma_1) - \mathcal{S}(\mathbf{X}_1^{(d)}, \gamma_1)$ with

$$Q_\Delta^{(d),\perp} = \sum_{j=2}^d \log [\pi(\mathbf{X}_j^{(d),*})/\pi(\mathbf{X}_j^{(d)})] \quad \text{and} \quad S_\Delta^{(d),\perp} = \sum_{j=2}^d \mathcal{S}(\mathbf{X}_j^{(d),*}, \gamma_j) - \mathcal{S}(\mathbf{X}_j^{(d)}, \gamma_j). \quad (\text{E.7})$$

Note that $Q_\Delta^{(d),\perp}$ and $S_\Delta^{(d),\perp}$ are independent of $Z_1^{(d)}$. Under Assumption 4, the moments of order two of the differences $Q_\Delta^{(d)} - Q_\Delta^{(d),\perp}$ and $S_\Delta^{(d)} - S_\Delta^{(d),\perp}$ are finite and converges to zero as $d \rightarrow \infty$. The Cauchy-Schwarz inequality and the fact that F is bounded and Lipschitz yield that $\text{ESJD}^{(d)}/(\mu/I)^2$ can also be expressed as

$$\mathbb{E} \left[\left(Z_1^{(d)} \right)^2 \times F \left(Q_\Delta^{(d),\perp} + S_\Delta^{(d),\perp} \right) \times F \left(W_\Delta^{(d)} - S_\Delta^{(d),\perp} \right) \right] \quad (\text{E.8})$$

$$+ \mathbb{E} \left[\left(Z_1^{(d)} \right)^2 \times F \left(Q_\Delta^{(d),\perp} + S_\Delta^{(d),\perp} \right) \times \left\{ F \left(W_\Delta^{(d)} - S_\Delta^{(d)} \right) - F \left(W_\Delta^{(d)} - S_\Delta^{(d),\perp} \right) \right\} \right] \quad (\text{E.9})$$

$$+ \mathbb{E} \left[\left(Z_1^{(d)} \right)^2 \times \left\{ F \left(Q_\Delta^{(d)} + S_\Delta^{(d)} \right) - F \left(Q_\Delta^{(d),\perp} + S_\Delta^{(d),\perp} \right) \right\} \times F \left(W_\Delta^{(d)} - S_\Delta^{(d)} \right) \right] \quad (\text{E.10})$$

$$= \mathbb{E} \left[F \left(Q_\Delta^{(d),\perp} + S_\Delta^{(d),\perp} \right) \times F \left(W_\Delta^{(d)} - S_\Delta^{(d),\perp} \right) \right] + o(1) \quad (\text{E.11})$$

$$= \mathbb{E} [F(Q_\Delta^\infty + S_\Delta^\infty) \times F(W_\Delta - S_\Delta^\infty)] + o(1) = \alpha_{12} + o(1), \quad (\text{E.12})$$

as required. We have used the fact that for almost every realisation of the auxiliary random variable $\{\Gamma_j\}_{j \geq 1}$ the sequence $(Q_\Delta^{(d),\perp}, S_\Delta^{(d),\perp})$ converges in distribution to $(Q_\Delta^\infty, S_\Delta^\infty)$, which readily follows from Lemma 4.1.

E.3 Proof of Theorem 4.1

The proof is a generalisation of the generator approach of [6, 44] coupled with an homogenization argument. We introduce the subsampled processes $\tilde{\mathbf{X}}^{(d)}$ and $\tilde{W}^{(d)}$ defined by

$$\tilde{\mathbf{X}}_k^{(d)} = \mathbf{X}_{k \times T^{(d)}}^{(d)} \quad \text{and} \quad \tilde{W}_k^{(d)} = W_{k \times T^{(d)}}^{(d)} \quad (\text{E.13})$$

for an intermediary time scale defined as $T^{(d)} = \lfloor d^\gamma \rfloor$ where γ is an arbitrary exponent such that $\gamma \in (0, 1/4)$. One step of the process $\tilde{\mathbf{X}}^{(d)}$ (resp. $\tilde{W}^{(d)}$) corresponds to $T^{(d)}$ steps of the process $\mathbf{X}^{(d)}$ (resp. $W^{(d)}$). We then define an accelerated version $\tilde{V}^{(d)}$ of the subsampled process $\tilde{X}^{(d)}$. In order to prove a diffusion limit for the process $X^{(d)}$, one needs to accelerate time by a factor of d ; consequently, in order to prove a diffusion limit for the process $\tilde{X}^{(d)}$, one needs to accelerate time by a factor $d/T^{(d)}$ and thus define $\tilde{V}^{(d)}$ by

$$\tilde{V}^{(d)}(t) := \tilde{X}_{\lfloor td/T^{(d)} \rfloor, 1}^{(d)}.$$

The proof then consists of showing that the sequence $\tilde{V}^{(d)}$ converges weakly in the Skorohod topology towards the limiting diffusion (4.9) and verifying that $\|\tilde{V}^{(d)} - V^{(d)}\|_{\infty, [0, T]}$ converges to zero in probability; this is enough to prove that the sequence $V^{(d)}$ converges weakly in the Skorohod topology towards the limiting diffusion (4.9). We denote by \mathcal{L} the generator of the limiting diffusion (4.9). Similarly, we define $\mathcal{L}^{(d)}$ and $\tilde{\mathcal{L}}^{(d)}$ the approximate generators of the first coordinate processes $X_1^{(d)}$ and $\tilde{X}_1^{(d)}$; for any smooth and compactly supported test function $\varphi : \mathbb{R} \rightarrow \mathbb{R}$, vector $\mathbf{x} = (x_1, \dots, x_d) \in \mathbb{R}^d$ and scalar $x, w \in \mathbb{R}$ we have

$$\begin{cases} \mathcal{L}^{(d)}\varphi(\mathbf{x}, w) &= \mathbb{E}_{\mathbf{x}, w}[\varphi(X_{1,1}^{(d)}) - \varphi(X_{0,1}^{(d)})]/\delta \\ \tilde{\mathcal{L}}^{(d)}\varphi(\mathbf{x}, w) &= \mathbb{E}_{\mathbf{x}, w}[\varphi(\tilde{X}_{1,1}^{(d)}) - \varphi(\tilde{X}_{0,1}^{(d)})]/(T^{(d)} \times \delta) \\ \mathcal{L}\varphi(x) &= \frac{1}{2} J(\mu) \times (\ell'(x)\varphi'(x) + \varphi''(x)). \end{cases} \quad \text{for } \delta \equiv 1/d \quad (\text{E.14})$$

Note that although φ is a scalar function, the functions $\mathcal{L}^{(d)}\varphi$ and $\tilde{\mathcal{L}}^{(d)}\varphi$ are defined on $\mathbb{R}^d \times \mathbb{R}$. The law of iterated conditional expectation yields the important identity between the generators $\mathcal{L}^{(d)}$ and $\tilde{\mathcal{L}}^{(d)}$,

$$\tilde{\mathcal{L}}^{(d)}\varphi(\mathbf{x}, w) = \frac{1}{T^{(d)}} \mathbb{E}_{\mathbf{x}, w} \left[\sum_{k=0}^{T^{(d)}-1} \mathcal{L}^{(d)}\varphi(\mathbf{X}_k^{(d)}, W_k^{(d)}) \right]. \quad (\text{E.15})$$

For clarity, the proof of Theorem 4.1 is divided into several steps.

E.3.1 The finite dimensional marginals of $\tilde{V}^{(d)}$ converge to those of the diffusion (4.9)

Since the limiting process is a scalar diffusion, the set of smooth and compactly supported functions is a core for the generator of the limiting diffusion ([22], Theorem 2.1, Chapter 8); in the sequel, one

can thus work with test functions belonging to this core only. Because the processes are started at stationarity, it suffices to show ([22], Chapter 4, Theorem 8.2, Corollary 8.4) that for any smooth and compactly supported function $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ the following limit holds,

$$\lim_{d \rightarrow \infty} \mathbb{E} \left[\left| \widetilde{\mathcal{L}}^{(d)} \varphi(X_1, \dots, X_d, W) - \mathcal{L} \varphi(X_1) \right|^2 \right] = 0. \quad (\text{E.16})$$

The proof of Equation (E.16) spans the remaining of this section and is based on an asymptotic expansion that we now describe. For every $x, w \in \mathbb{R}$ we define the approximated generator $\mathcal{A}\varphi : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ by

$$\mathcal{A}\varphi(x, w) = \left(\frac{\mu}{I} \right)^2 \left\{ A(w) \ell'(x) + \left(\frac{1}{2} \alpha_{12} + [A(w) - B(w)] \partial_x \mathcal{S}(x, \gamma_1) \right) \varphi''(x) \right\} \quad (\text{E.17})$$

where $A, B : \mathbb{R} \rightarrow (0; \infty)$ are two bounded and continuous functions defined by

$$\begin{cases} A(w) &= \mathbb{E} \left[F'(Q_\Delta^\infty + S_\Delta^\infty) \times F(W^* - w - S_\Delta^\infty) \right] \\ B(w) &= \mathbb{E} \left[F(Q_\Delta^\infty + S_\Delta^\infty) \times F'(W^* - w - S_\Delta^\infty) \right] \end{cases} \quad (\text{E.18})$$

for $W^* \stackrel{\mathcal{D}}{\sim} \pi_{W^*}$ and $F'(u) = e^u \mathbb{I}_{u < 0}$ and $(Q_\Delta^\infty, S_\Delta^\infty)$ as defined in (4.3). The functions $A, B : \mathbb{R} \rightarrow \mathbb{R}_+$ are such that

$$\mathbb{E} [A(W)] = \mathbb{E} [B(W)] = \frac{1}{2} \alpha_{12}. \quad (\text{E.19})$$

The proof of (E.19) can be found in Appendix F.3. It follows from (E.19) that for any fixed $x \in \mathbb{R}$ we have

$$\mathbb{E} [\mathcal{A}\varphi(x, W)] = \mathcal{L}\varphi(x) \quad (\text{E.20})$$

for a random variable $W \stackrel{\mathcal{D}}{\sim} \pi_W$.

Lemma E.1. *Let Assumptions 4 hold. We have*

$$\lim_{d \rightarrow \infty} \mathbb{E} \left[\left| \mathcal{L}^{(d)} \varphi(X_1, \dots, X_d, W) - \mathcal{A}\varphi(X_1, W) \right|^2 \right] = 0. \quad (\text{E.21})$$

The proof of Lemma (E.1) consists in second order Taylor expansion and an averaging argument; details are in Section F.4. For proving Equation (E.16), note that identity (E.15) and Jensen's inequality yield the quantity inside the limit described in Equation (E.16) is less than two times the expectation of

$$\left\{ \frac{\sum_{k=0}^{T^{(d)}} \mathcal{L}^{(d)} \varphi \left(\mathbf{X}_k^{(d)}, W_k^{(d)} \right) - \mathcal{A}\varphi \left(\mathbf{X}_{k,1}^{(d)}, W_k^{(d)} \right)}{T^{(d)}} \right\}^2 + \left\{ \frac{\sum_{k=0}^{T^{(d)}} \mathcal{A}\varphi \left(\mathbf{X}_{k,1}^{(d)}, W_k^{(d)} \right) - \mathcal{L}\varphi \left(\mathbf{X}_{0,1}^{(d)} \right)}{T^{(d)}} \right\}^2. \quad (\text{E.22})$$

The expectation of the first term is less than $\mathbb{E} \left[\left| \mathcal{L}^{(d)} \varphi(X_1, \dots, X_d, W) - \mathcal{A}\varphi(X_1, W) \right|^2 \right]$ and Lemma (E.1) shows that this quantity goes to zero as $d \rightarrow \infty$. To finish the proof it thus remains to verify that the expectation of the second term also converges to zero; to prove so, note

that the second term is less than two times

$$\frac{\sum_{k=0}^{T^{(d)}} \left| \mathcal{A}\varphi(X_{k,1}^{(d)}, W_k^{(d)}) - \mathcal{A}\varphi(X_{0,1}^{(d)}, W_k^{(d)}) \right|^2}{T^{(d)}} + \left\{ \sum_{k=0}^{T^{(d)}} \mathcal{A}\varphi(X_{0,1}^{(d)}, W_k^{(d)}) - \mathcal{L}\varphi(X_{0,1}^{(d)}) T^{(d)} \right\}^2. \quad (\text{E.23})$$

Under the assumptions of Theorem 4.1, it is straightforward to verify that the function $\mathcal{A}\varphi$ is globally Lipschitz in the sense that there exists a constant $\|\mathcal{A}\varphi\|_{\text{Lip}}$ such that for every $x_1, x_2, w \in \mathbb{R}$ we have $|\mathcal{A}\varphi(x_1, w) - \mathcal{A}\varphi(x_2, w)| \leq \|\mathcal{A}\varphi\|_{\text{Lip}} \times |x_1 - x_2|$; since $\mathbb{E}[(X_{k,1}^{(d)} - X_{0,1}^{(d)})^2] \lesssim k^2/d$, it follows that

$$\mathbb{E} \left(\left| \mathcal{A}\varphi(X_{k,1}^{(d)}, W_k^{(d)}) - \mathcal{A}\varphi(X_{0,1}^{(d)}, W_k^{(d)}) \right|^2 \right) \lesssim \frac{k^2}{d}.$$

Consequently, the expectation of the first term in (E.23) converges to zero. For proving that the second term also converges to zero, we make use of the following ergodic averaging Lemma whose proof can be found in Section F.5.

Lemma E.2. *Let $h : \mathbb{R} \rightarrow \mathbb{R}$ be a bounded and measurable test function. We have*

$$\lim_{d \rightarrow \infty} \mathbb{E} \left[\left| \frac{\sum_{k=0}^{T^{(d)}-1} h(W_k^{(d)})}{T^{(d)}} - \mathbb{E}[h(W)] \right|^2 \right] = 0,$$

for a random variable $W \stackrel{\mathcal{D}}{\sim} \pi_W$ independent from any other sources of randomness.

Identity (E.20), a standard conditioning argument and Lemma E.2 yield that the expectation of the second term in Equation (E.23) also converges to zero; this finishes the proof of the convergence of the finite dimensional marginals of \tilde{V}^d to those of the limiting diffusion (4.9).

E.3.2 The sequence \tilde{V}^d converges weakly towards the diffusion (4.9)

The finite dimensional marginals of the sequence process \tilde{V}^d converges to those of the diffusion (4.9). To prove that the sequence \tilde{V}^d actually converges to the diffusion (4.9), it thus suffices to verify that the sequence \tilde{V}^d is relatively weak compact in the Skorohod topology: since the process $\tilde{V}^{(d)}$ is started at stationarity and the space of smooth functions with compact support is an algebra that strongly separates points, ([22], Chapter 4, Corollary 8.6) states that it suffices to show that for any smooth and compactly supported test function φ the sequence $d \mapsto \mathbb{E} \left| \tilde{\mathcal{L}}^{(d)} \varphi(X_1, \dots, X_d, W) \right|^2$ is bounded. Equation (E.16) shows that it suffices to verify that $\mathbb{E} \left| \mathcal{L}\varphi(X) \right|^2 < \infty$ for $X \stackrel{\mathcal{D}}{\sim} \pi$, which is obvious since φ is assumed to be smooth with compact support.

E.3.3 The sequence V^d converges weakly towards the diffusion (4.9)

Because the sequence \tilde{V}^d converges weakly to the diffusion (4.9), it suffices to prove that the difference $\|V^d - \tilde{V}^d\|_{\infty, [0, T]}$ goes to zero in probability. To this end, it suffices to prove that the supremum

$$\sup \left\{ \left| X_{kT^{(d)}+i,1}^{(d)} - X_{kT^{(d)},1}^{(d)} \right| : k \times T^{(d)} \leq d \times T, i \leq T^{(d)} \right\} \quad (\text{E.24})$$

converges to zero in probability. The quantity $\left| X_{kT^{(d)+i,1}^{(d)}} - X_{kT^{(d),1}^{(d)}} \right|$ is less than a constant times

$$\frac{1}{d^{1/2}} \left\{ |Z_{kT^{(d),1}^{(d)}}| + \dots + |Z_{(k+1)T^{(d)-1,1}^{(d)}}| \right\}.$$

Therefore, for any $\varepsilon > 0$ and integer $p \geq 1$ and an i.i.d sequence of standard Gaussian random variables $\{\xi_k\}_{k \geq 0}$, the union bound and Markov's inequality show the probability that the supremum in Equation (E.24) is larger than ε is less than a constant multiple of

$$\frac{d}{T^{(d)}} \times \mathbb{P}[|\xi_1| + \dots + |\xi_{T^{(d)}}| > \varepsilon d^{1/2}] \leq \frac{d}{T^{(d)}} \times \frac{\mathbb{E}[(|\xi_1| + \dots + |\xi_{T^{(d)}}|)^p]}{\varepsilon^p d^{p/2}}$$

Since $T^{(d)} = d^\gamma$ and for every integer $p \geq 1$ we have that $\mathbb{E}[(|\xi_1| + \dots + |\xi_n|)^p] \leq C(p) n^p$ for a constant $C(p)$ that only depends on p , it follows that

$$\frac{d}{T^{(d)}} \times \mathbb{P}[|\xi_1| + \dots + |\xi_{T^{(d)}}| > \varepsilon d^{1/2}] \lesssim d^{1-\gamma+p[\gamma-1/2]}.$$

Since $\gamma < 1/2$, one can choose p large enough such that $1 - \gamma + p[\gamma - 1/2] < 0$. This concludes the proof of Theorem 4.1.

F Proof of technical results

In this section we denote by $\Phi(x) = \int_{-\infty}^x \varphi(u) du$ the cumulative Gaussian function with $\varphi(u) = e^{-u^2/2}/\sqrt{2\pi}$. The bound $1 - \Phi(x) < \varphi(x)/x$ for $x > 0$ is used in several places.

F.1 Proof of Proposition 4.2

The only not entirely trivial parts of this proposition involve establishing that α_1 and $\alpha_{2|1}$ are decreasing in μ and σ respectively. For proving that $\alpha_1 = G\left(-\frac{\mu^2}{2}(1 - \beta_1), \mu^2(1 + \beta_2^2 - 2\beta_1)\right)$ is decreasing as a function of μ when $\beta_1 < 1$, note that since $|\beta_1| < \beta_2$, $1 + \beta_2^2 - 2\beta_1 \geq (1 - \beta_1)^2$; hence it suffices to show that for any positive constant $c > 0$ the function $h : \mu \mapsto G(-\mu^2, c^2\mu^2)$ is decreasing. Since $h(\mu) = \mathbb{E}[F(-\mu^2 + c\mu\xi)]$ for a random variable $\xi \stackrel{\mathcal{D}}{\sim} \mathbf{N}(0, 1)$ and $F'(x) = e^x \mathbb{I}(x < 0)$ it follows that

$$h'(\mu) = \int_{z \in \mathbb{R}} F'(-\mu^2 + c\mu z) (-2\mu + cz) \varphi(z) dz = \int_{z < \mu/c} F'(-\mu^2 + c\mu z) (-2\mu + cz) \varphi(z) dz.$$

This quantity is negative since $-2\mu + cz < 0$ on the event $\{z : z < \mu/c\}$. Proving that $\alpha_{2|1}$ is decreasing as a function of σ readily follows from the fact that for any fixed $a \in \mathbb{R}$ the derivative of the function $\sigma \mapsto G(-\sigma^2 + a, 2\sigma^2) < -\sqrt{2}\varphi(-\sigma/\sqrt{2} + a/(\sigma\sqrt{2})) < 0$ and differentiation under the integral sign.

F.2 Proof of Theorem 5.1

Since $\text{Eff}(\mu, \sigma^2) = \frac{\mu^2 \alpha_{12}(\mu, \sigma)}{\eta + \alpha_1(\mu)/\sigma^2}$, for a fixed value of scaling $\mu > 0$ the efficiency functional goes to zero as $\sigma \rightarrow 0$ and $\sigma \rightarrow \infty$. Similarly, the fact that the efficiency goes to zero as $\mu \rightarrow 0$ for any fixed value of $\sigma > 0$ is straightforward; it remains to verify that the efficiency also converge

to zero as $\mu \rightarrow \infty$. It suffices to show that $\mu^2 \alpha_{12}(\mu, \sigma) \rightarrow 0$; since for any $x, y \in \mathbb{R}$ we have $\min(1, e^x) \min(1, e^y) \leq \min(1, e^{x+y})$,

$$\alpha_{12} \leq \mathbb{E}[F(Q_\Delta^\infty + W_\Delta)] = 2\Phi \left\{ -\frac{(\mu^2 + 2\sigma^2)^{1/2}}{2} \right\}$$

and the conclusion readily follows.

F.3 Proof of Equation (E.19)

Equation 3.8 yields that $R \equiv W^* - W$ for $(W^*, W) \sim \pi_{W^*} \otimes \pi_W$ has a density π_R such that the function $r \mapsto e^{r/2} \pi_R(r)$ is symmetric i.e. $e^{r/2} \pi_R(r) = e^{-r/2} \pi_R(-r)$. Similarly, algebra reveals that the joint Gaussian density $\pi_{Q,S}(q, s)$ of the pair $(Q_\Delta^\infty, S_\Delta^\infty)$ described in Lemma 4.1 is such that

$$e^{q/2} \pi_{Q,S}(q, s) = e^{-q/2} \pi_{Q,S}(-q, -s).$$

That is because $-\log \pi_{Q,S}(q, s) = aq^2 + bs^2 + cqs - q/2 + (\text{constant})$ for some coefficients $a, b, c \in \mathbb{R}$. Consequently, since the accept reject function F is such that $e^{-u}F(u) = F(-u)$ for any $u \in \mathbb{R}$, the function

$$\begin{aligned} g(q, r, s) &= e^{(q+s)/2} F(r-s) \pi_{Q,S}(q, s) \pi_R(r) \\ &= e^{-(r-s)/2} F(r-s) (e^{q/2} \pi_{Q,S}(q, s)) (e^{r/2} \pi_R(r)) \end{aligned}$$

is such that $g(q, r, s) = g(-q, -r, -s)$. It follows that

$$\begin{aligned} \mathbb{E}[A(W)] &= \mathbb{E}[F'(Q_\Delta^\infty + S_\Delta^\infty) \times F(R - S_\Delta^\infty)] = \iiint_{\mathbb{R}^3} F'(q+s) F(r-s) \pi_{Q,S}(q, s) \pi_R(r) dq dr ds \\ &= \iiint_{\mathbb{R}^3} e^{-(q+s)/2} F'(q+s) g(q, r, s) dq dr ds = \iiint_{\mathbb{R}^3} e^{(q+s)/2} F'(-[q+s]) g(q, r, s) dq dr ds \\ &= \mathbb{E}[e^{Q_\Delta^\infty + S_\Delta^\infty} F'(-[Q_\Delta^\infty + S_\Delta^\infty]) F(R - S_\Delta^\infty)]. \end{aligned}$$

Consequently, since $F'(u) + e^u F'(-u) = F(u)$ for $u \in \mathbb{R}$, it follows that

$$\begin{aligned} 2 \times \mathbb{E}[A(W)] &= \mathbb{E}[F'(Q_\Delta^\infty + S_\Delta^\infty) \times F(R - S_\Delta^\infty) + e^{Q_\Delta^\infty + S_\Delta^\infty} F'(-[Q_\Delta^\infty + S_\Delta^\infty]) F(R - S_\Delta^\infty)] \\ &= \mathbb{E}[F(Q_\Delta^\infty + S_\Delta^\infty)] \equiv \alpha_{12}. \end{aligned}$$

The proof that $\mathbb{E}[B(W)] = \alpha_{12}/2$ is similar and thus omitted.

F.4 Proof of Lemma E.1

In this section we need to consider asymptotic expansions of the type $\mathbb{E}_{\mathbf{x}, w}[\dots] = \Psi(\mathbf{x}, w) + (\text{error term})$, where $(\mathbf{x}, w) \in \mathbb{R}^d \times \mathbb{R}$ and $(\text{error term}) = \varepsilon_d(\mathbf{x}, w)$ for a function $\varepsilon_d : \mathbb{R}^d \times \mathbb{R} \rightarrow \mathbb{R}$. We use the notation $(\text{error term}) = o_{L^2}(1)$ to indicate that, under the equilibrium distribution, the moment of order two of the error term is asymptotically negligible, $\mathbb{E}[\varepsilon_d(\mathbf{X}^{(d)}, W)^2] \rightarrow 0$ as $d \rightarrow \infty$ for $(\mathbf{X}^{(d)}, W) \stackrel{\mathcal{D}}{\sim} \pi^{(d)} \otimes \pi_W$. Since φ is smooth with compact support, a second order Taylor expansion reveals that

$$\mathcal{L}^{(d)}\varphi(\mathbf{x}, w) = (\text{drift term}) \varphi'(\mathbf{x}) + (1/2) (\text{volatility term}) \varphi''(\mathbf{x}) + o_{L^2}(1) \quad (\text{F.1})$$

where the drift and volatility terms are given by the following conditional expectations,

$$\begin{cases} \text{(drift term)} & = (1/\delta) \times \mathbb{E}_{\mathbf{x},w} \left[\left(X_{1,1}^{(d),*} - x_1 \right) \alpha_{12}^{(d)}(\mathbf{x}, w, \mathbf{X}^{(d),*}, W^{(d),*}) \right] \\ \text{(volatility term)} & = (1/\delta) \times \mathbb{E}_{\mathbf{x},w} \left[\left(X_{1,1}^{(d),*} - x_1 \right)^2 \alpha_{12}^{(d)}(\mathbf{x}, w, \mathbf{X}^{(d),*}, W^{(d),*}) \right] \end{cases} \quad (\text{F.2})$$

with $\mathbf{X}_1^{(d),*} = \mathbf{x} + (\mu/I) \delta^{1/2} \mathbf{Z}^{(d)}$ and standard centred Gaussian random variable $\mathbf{Z}^{(d)} = (Z_1, \dots, Z_d)$

- It readily follows from Lemma 4.1 that for π -almost every \mathbf{x} we have

$$\text{(volatility term)} = \alpha_{12} \times (\mu/I)^2 = J(\mu) + o_{L^2}(1). \quad (\text{F.3})$$

- For the drift term, we make use of the following integration-by-part formula, also known as Stein's identity,

$$\mathbb{E}[Z \times g(Z)] = \mathbb{E}[g'(Z)] \quad \text{for} \quad Z \stackrel{\mathcal{D}}{\sim} \mathbf{N}(0, 1), \quad (\text{F.4})$$

which holds for any continuous and piecewise continuously differentiable function $g : \mathbb{R} \rightarrow \mathbb{R}$ such that $x \mapsto \max(g(x), g'(x))$ is polynomially bounded. In what follows, $F'(u) = e^u \mathbb{I}_{u < 0}$. The expression for $\alpha_{12}^{(d)}(\mathbf{x}, w, \mathbf{X}^{(d),*}, W^{(d),*})$, identity (F.4) and standard algebraic manipulations yield that

$$\begin{aligned} \text{(drift term)} &= \delta^{-1/2} (\mu/I) \mathbb{E}_{\mathbf{x},w} [Z_1 \alpha_{12}^{(d)}(\mathbf{x}, w, \mathbf{X}^{(d),*}, W^{(d),*})] \\ &= (\mu/I)^2 \mathbb{E}_{\mathbf{x},w} [F'(\mathbf{Q}_\Delta^{(d)} + \mathbf{S}_\Delta^{(d)}) F(\mathbf{W}_\Delta^{(d)} - \mathbf{S}_\Delta^{(d)}) \{ \ell'(X_{1,1}^{(d),*}) + \partial_x \mathcal{S}(X_{1,1}^{(d),*}, \gamma_1) \}] \\ &\quad - (\mu/I)^2 \mathbb{E}_{\mathbf{x},w} [F(\mathbf{Q}_\Delta^{(d)} + \mathbf{S}_\Delta^{(d)}) F'(\mathbf{W}_\Delta^{(d)} - \mathbf{S}_\Delta^{(d)}) \partial_x \mathcal{S}(X_{1,1}^{(d),*}, \gamma_1)] \\ &= (\mu/I)^2 A(w) \ell'(x_1) + (\mu/I)^2 [A(w) - B(w)] \partial_x \mathcal{S}(x_1, \gamma_1) + o_{L^2}(1), \end{aligned} \quad (\text{F.5})$$

where the functions $A, B : \mathbb{R} \rightarrow \mathbb{R}^+$ are defined in Equation (E.18) and the quantities $\mathbf{Q}_\Delta^{(d)}, \mathbf{S}_\Delta^{(d)}$ and $\mathbf{W}_\Delta^{(d)}$ in Equation (E.1).

Plugging (F.5) and (F.3) into (F.1) shows that the limit

$$\lim_{d \rightarrow \infty} \mathbb{E} \left[\left| \mathcal{L}^{(d)} \varphi(\mathbf{X}^{(d)}, W) - \mathcal{A} \varphi(X_1^{(d)}, W) \right|^2 \right] = 0$$

holds for $(\mathbf{X}^{(d)}, W) \sim \pi^{(d)} \otimes \pi_W$, as required.

F.5 Proof of Lemma E.2

The strategy of the proof is as follows. We define three stochastic processes $\{W_{\clubsuit, k}^{(d)}\}_{k \geq 0}$, $\{W_{\spadesuit, k}^{(d)}\}_{k \geq 0}$, $\{W_{\blacksquare, k}\}_{k \geq 0}$ such that

$$\begin{cases} \lim_{d \rightarrow \infty} \mathbb{P} \left(W_{\clubsuit, k}^{(d)} = W_k^{(d)} : 0 \leq k \leq T^{(d)} \right) = 1, \\ \left(W_{\clubsuit, k}^{(d)} = W_k^{(d)} : 0 \leq k \leq T^{(d)} \right) \stackrel{\text{law}}{=} \left(W_{\spadesuit, k}^{(d)} = W_k^{(d)} : 0 \leq k \leq T^{(d)} \right), \\ \lim_{d \rightarrow \infty} \mathbb{P} \left(W_{\spadesuit, k}^{(d)} = W_{\blacksquare, k} : 0 \leq k \leq T^{(d)} \right) = 1, \\ \{W_{\blacksquare, k}\}_{k \geq 0} \text{ is a Markov chain that is ergodic with respect to } \pi_W. \end{cases} \quad (\text{F.6})$$

Once (F.6) is proved, Lemma E.2 immediately follows. Let us now define these three processes and verify that Equation (F.6) holds. To do so, let us consider i.i.d sequences $\{X_i\}_{i \geq 1}$ and $\{W_i^*\}_{i \geq 1}$ and $\{Z_{i,k}\}_{i,k \geq 1}$ and $\{U_k\}_{k \geq 0}$ respectively marginally distributed as π and π_{W^*} and $\mathbf{N}(0, 1)$ and $\text{Uniform}([0, 1])$. We consider $\{x_i\}_{i \geq 1}$ a realisation of $\{X_i\}_{i \geq 1}$ and for any index $d \geq 1$ we set $\mathbf{X}_0^{(d)} = (x_1, \dots, x_d)$ and $W_0^{(d)} \stackrel{\mathcal{D}}{\sim} \pi_W$ and recursively define $(\mathbf{X}_{k+1}^{(d)}, W_{k+1}^{(d)}) = (\mathbf{X}_k^{(d)*}, W_k^*)$, with $\mathbf{X}_k^{(d)*} = \mathbf{X}_k^{(d)} + (\mu/I) \delta^{1/2} \mathbf{Z}_k^{(d)}$ and $\mathbf{Z}_k^{(d)} = (Z_{1,k}, \dots, Z_{d,k})$, if

$$U_k \leq F \left(Q_{\Delta,k}^{(d)} + S_{\Delta,k}^{(d)} \right) \times F \left(W_k^* - W_k^{(d)} - S_{\Delta,k}^{(d)} \right) \quad (\text{F.7})$$

and $(\mathbf{X}_{k+1}^{(d)}, W_{k+1}^{(d)}) = (\mathbf{X}_k^{(d)}, W_k^{(d)})$ otherwise. In the above

$$\begin{cases} Q_{\Delta,k}^{(d)} &= \sum_{i=1}^d \ell \left(X_{k,i}^{(d)*} \right) - \ell \left(X_{k,i}^{(d)} \right) \\ S_{\Delta,k}^{(d)} &= \sum_{i=1}^d \mathcal{S} \left(X_{k,i}^{(d)*}, \gamma_i \right) - \mathcal{S} \left(X_{k,i}^{(d)}, \gamma_i \right). \end{cases}$$

Indeed, for any index $d \geq 1$ the process $\left\{ \left(\mathbf{X}_k^{(d)}, W_k^{(d)} \right) \right\}_{k \geq 0}$ is a DAPsMRWM Markov chain that targets $\pi^{(d)} \otimes \pi_W$. Let us now define the processes $W_{\clubsuit}, W_{\spadesuit}, W_{\blacksquare}$.

- We set $W_{\clubsuit,0}^{(d)} = W_0^{(d)}$ and recursively define $W_{\clubsuit,k+1}^{(d)} = W_k^*$ if

$$U_k \leq F \left(Q_{\clubsuit,\Delta,k}^{(d)} + S_{\clubsuit,\Delta,k}^{(d)} \right) \times F \left(W_k^* - W_{\clubsuit,k}^{(d)} - S_{\clubsuit,\Delta,k}^{(d)} \right) \quad (\text{F.8})$$

and $W_{\clubsuit,k+1}^{(d)} = W_{\clubsuit,k}^{(d)}$ otherwise; we have used the notations

$$\begin{cases} Q_{\clubsuit,\Delta,k}^{(d)} &= (\mu\delta/I) \sum_{i=1}^d \ell'(x_i) Z_{i,k} + (\mu^2\delta^2/2I^2) \sum_{i=1}^d \ell''(x_i) \\ S_{\clubsuit,\Delta,k}^{(d)} &= (\mu\delta/I) \sum_{i=1}^d \mathcal{S}'(x_i, \gamma_i) Z_{i,k} + (\mu^2\delta^2/2I^2) \sum_{i=1}^d \mathcal{S}''(x_i, \gamma_i). \end{cases}$$

- Similarly, we set $W_{\spadesuit,0}^{(d)} = W_0^{(d)}$ and recursively define $W_{\spadesuit,k+1}^{(d)} = W_k^*$ if

$$U_k \leq F \left(Q_{\spadesuit,\Delta,k}^{(d)} + S_{\spadesuit,\Delta,k}^{(d)} \right) \times F \left(W_k^* - W_{\spadesuit,k}^{(d)} - S_{\spadesuit,\Delta,k}^{(d)} \right) \quad (\text{F.9})$$

and $W_{\spadesuit,k+1}^{(d)} = W_{\spadesuit,k}^{(d)}$ otherwise; we have used the notations $(Q_{\spadesuit,\Delta,k}^{(d)}, S_{\spadesuit,\Delta,k}^{(d)})$ to designate a Gaussian random variable in \mathbb{R}^2 , independent from any other source of randomness, with same law as $(Q_{\clubsuit,\Delta,k}^{(d)}, S_{\clubsuit,\Delta,k}^{(d)})$.

- Finally, we set $W_{\blacksquare,0}^{(d)} = W_0^{(d)}$ and recursively define $W_{\blacksquare,k+1}^{(d)} = W_k^*$ if

$$U_k \leq F \left(Q_{\Delta,k}^{(\infty)} + S_{\Delta,k}^{(\infty)} \right) \times F \left(W_k^* - W_{\blacksquare,k}^{(d)} - S_{\Delta,k}^{(\infty)} \right) \quad (\text{F.10})$$

and $W_{\blacksquare,k+1}^{(d)} = W_{\blacksquare,k}^{(d)}$ otherwise; in the above $\left\{ \left(Q_{\Delta,k}^{(\infty)}, S_{\Delta,k}^{(\infty)} \right) \right\}_{k \geq 0}$ is an i.i.d sequence marginally distributed as $(Q_{\Delta}^{(\infty)}, S_{\Delta}^{(\infty)})$; see Lemma 4.1.

It is obvious that $\{W_{\clubsuit,k}^{(d)}\}_{k \geq 0}$ and $\{W_{\spadesuit,k}^{(d)}\}_{k \geq 0}$ have the same law. The fact that $\{W_{\blacksquare,k}^{(d)}\}_{k \geq 0}$ is a Markov chain ergodic with respect to π_W readily follows from the fact that it is reversible with respect to π_W ; it is a standard Gaussian computation. The proof of the first and third equation in (F.6) is based on the following basic remark. For convenience, let us denote by $\mathcal{E}_k^{(d)}, \mathcal{E}_{k,\clubsuit}^{(d)}, \mathcal{E}_{k,\spadesuit}^{(d)}, \mathcal{E}_{k,\infty}^{(d)}$ the Bernoulli random variables indicating whether or not the respective events (F.7),(F.8),(F.9), (F.10) are realised or not. We have

$$1 - \mathbb{P}\left(W_{\clubsuit,k}^{(d)} = W_k^{(d)} : 0 \leq k \leq T^{(d)}\right) \leq \sum_{k=0}^{T^{(d)}-1} \mathbb{P}\left(\mathcal{E}_k^{(d)} \neq \mathcal{E}_{k,\clubsuit}^{(d)} \mid W_{\clubsuit,k}^{(d)} = W_k^{(d)}\right) \quad (\text{F.11})$$

and the conditional probability $\mathbb{P}\left(\mathcal{E}_k^{(d)} \neq \mathcal{E}_{k,\clubsuit}^{(d)} \mid W_{\clubsuit,k}^{(d)} = W_k^{(d)}\right)$ is less than the expectation, conditioned upon the event $\{W_{\clubsuit,k}^{(d)} = W_k^{(d)}\}$, of the absolute difference

$$\left|F\left(Q_{\Delta,k}^{(d)} + S_{\Delta,k}^{(d)}\right) F\left(W_k^* - W_k^{(d)} - S_{\Delta,k}^{(d)}\right) - F\left(Q_{\clubsuit,\Delta,k}^{(d)} + S_{\clubsuit,\Delta,k}^{(d)}\right) F\left(W_k^* - W_{\clubsuit,k}^{(d)} - S_{\clubsuit,\Delta,k}^{(d)}\right)\right|. \quad (\text{F.12})$$

Because the $[0, 1]$ -valued function F is assumed to be Lipschitz, if $W_{\clubsuit,k}^{(d)} = W_k^{(d)}$ the absolute difference in (F.12) is less than $2 \times \|F\|_{\text{Lip}} \times \left\{\left|Q_{\Delta,k}^{(d)} - Q_{\clubsuit,\Delta,k}^{(d)}\right| + \left|S_{\Delta,k}^{(d)} - S_{\clubsuit,\Delta,k}^{(d)}\right|\right\}$. Because the second and third derivatives of the log-likelihood function ℓ are globally bounded, a third order Taylor expansion yield that

$$\begin{aligned} \mathbb{E}\left|Q_{\Delta,k}^{(d)} - Q_{\clubsuit,\Delta,k}^{(d)}\right| &\lesssim d^{-1/2} \mathbb{E}\left|\sum_{i=1}^d \left(\ell'(X_{k,i}^{(d)}) - \ell'(x_i)\right) Z_{i,k}\right| + d^{-1} \mathbb{E}\left|\sum_{i=1}^d \left(\ell''(X_{k,i}^{(d)}) - \ell''(x_i)\right) Z_{i,k}^2\right| + \mathcal{O}(d^{-1/2}) \\ &\lesssim d^{-1/2} \left\{\sum_{i=1}^d \mathbb{E}\left[\left(\ell'(X_{k,i}^{(d)}) - \ell'(x_i)\right)^2\right]\right\}^{1/2} + d^{-1} \left\{\sum_{i=1}^d \mathbb{E}\left[\left(\ell''(X_{k,i}^{(d)}) - \ell''(x_i)\right)^2\right]\right\}^{1/2} \\ &= \mathcal{O}(k d^{-1/2}). \end{aligned}$$

We have used the fact that for any exponent $p \geq 1$ we have $\mathbb{E}\left[\left|X_{k,i}^{(d)} - x_i\right|^p\right]^{1/p} \lesssim k d^{-1/2}$, which readily follows from the triangular inequality. Similarly, we have that $\mathbb{E}\left|S_{\Delta,k}^{(d)} - S_{\clubsuit,\Delta,k}^{(d)}\right| \lesssim k d^{-1/2}$. Plugging these estimates in (F.11) shows that

$$1 - \mathbb{P}\left(W_{\clubsuit,k}^{(d)} = W_k^{(d)} : 0 \leq k \leq T^{(d)}\right) \lesssim d^{-1/2} \sum_{k=0}^{T^{(d)}-1} k \rightarrow 0$$

since $T^{(d)} = d^\gamma$ for some exponent $\gamma \in (0, 1/4)$; we have thus proved that $\mathbb{P}\left(W_{\clubsuit,k}^{(d)} = W_k^{(d)} : 0 \leq k \leq T^{(d)}\right)$ converges to one as $d \rightarrow \infty$. The proof of the estimate $\mathbb{P}\left(W_{\spadesuit,k}^{(d)} = W_k^{(d)} : 0 \leq k \leq T^{(d)}\right) \rightarrow 1$ uses the same ingredients and is thus omitted.

G Pseudo-marginal estimate at Stage 1

Throughout the main article we have assumed that $\pi_a(x)$ is a computationally cheap, *deterministic* approximation to $\pi(x)$. However, in the DAPsMRWM setting, where only a noisy but unbiased

non-negative estimator, $\widehat{\pi}(x, U)$, of $\pi(x)$ is available one might have a cheaper, noisy, non-negative unbiased estimator, $\widehat{\pi}_a(x, U_a)$ to use in Stage One. Indeed, $\widehat{\pi}_a(x, U_a)$ could even be biased, and the delayed-acceptance algorithm would still be valid [e.g., see 32]. Here, we provide an equivalent to Lemma 4.1 and Proposition 4.1 for this general scenario, and then investigate the use of a noisy, unbiased estimator at Stage One; this suggests an upper bound on the efficiency achievable in the case of a noisy, biased estimator.

Analogously to Section 3.3, we may write

$$\widehat{\pi}_a^{(d)}(x, u_a) = \pi_a^{(d)}(x)e^{w_a}.$$

Assumptions 1 and 2 lead to distributions $W_a^* \sim N(-\sigma_a^2/2, \sigma_a^2)$ for the noise in the density at the proposed value, and $W_a \sim N(\sigma_a^2/2, \sigma_a^2)$ for the noise in the density at the current value. Similarly, Assumption 3 gives the compute time for Stage One as being proportional to σ_a^{-2} .

Section G.1 sketches the proof, by steps analogous to those in Section 4, that the following limits exist: $\alpha_1(\mu, \sigma_a^2; \beta_1, \beta_2) = \lim_{d \rightarrow \infty} \alpha_1^{(d)}(X^{(d)}, W_a^{(d)})$ and $\alpha_{12}(\mu, \sigma^2, \sigma_a^2; \beta_1, \beta_2) = \lim_{d \rightarrow \infty} \alpha_{12}(X^{(d)}, W_a^{(d)}, W^{(d)})$. We do not pursue a complete analysis of the joint tuning of three parameters over the range of possible values of β_1 and β_2 ; instead we investigate the limit as $\beta_1 \rightarrow 0$ and $\beta_2 \downarrow 0$. This describes the behaviour as the bias in the cheap approximation approaches 0, and might provide an approximate upper bound on the possible improvements achievable when the Stage One approximation is, in fact, biased. In this limit may obtain the following values, analogous to the expressions in Appendix B (where G is defined):

$$\alpha_1(\mu, \sigma_a^2; 0, 0) = G\left(-\frac{1}{2}\mu^2 - \sigma_a^2, \mu^2 + 2\sigma_a^2\right) \quad (\text{G.1})$$

$$\alpha_{12}(\mu, \sigma_a^2, \sigma^2; 0, 0) = \mathbb{E}\left[G\left(-\frac{1}{2}\mu^2 - \sigma_a^2 + \sqrt{2}\sigma_a\xi, \mu^2\right) G\left(-\sigma^2 + \sigma_a^2 - \sqrt{2}\sigma_a\xi, 2\sigma^2\right)\right], \quad (\text{G.2})$$

where $\xi \sim \mathcal{N}(0, 1)$. An analogous proof to that of Proposition 4.3 then gives $\lim_{d \rightarrow \infty} ESJD^{(d)} \propto \alpha_{12}\mu^2$.

For the DAPsRWM, η was the relative computational cost of the deterministic approximation to the pseudo marginal approximation when $\sigma^2 = 1$. When the approximation at Stage One is also random we define η to be the relative computational cost of $\widehat{\pi}_a$ when $\sigma_a^2 = 1$ to the cost of $\widehat{\pi}$ when $\sigma^2 = 1$. Thus, following Section 5.2, the total computational cost is proportional to $\eta/\sigma_a^2 + \alpha_1/\sigma^2$, and we obtain a limiting algorithm efficiency of

$$\text{Eff}_{\text{dapmpm}}(\mu, \sigma^2, \sigma^2) = \mu^2 \frac{\alpha_{12}(\mu, \sigma^2, \sigma_a^2)}{\eta/\sigma_a^2 + \alpha_1(\mu, \sigma_a^2)/\sigma^2}. \quad (\text{G.3})$$

With $\eta = 0.1$ the optimal efficiency in (G.3) is only a factor of ≈ 2.1 better than standard pseudo-marginal MCMC and we conclude that, as with the DAPsRWM algorithm (see Section 5.2), when only an order of magnitude cheaper, $\widehat{\pi}_a$ is of dubious utility. Finally, at $\eta = 0.01$ the optimal efficiency factor is ≈ 6.8 and the algorithm could be worth implementing provided that the bias is small. In [32], a biased, noisy estimator at Stage One (with $\eta \approx 0.1$) is found to lead to an efficiency of at most double that of the PsMRWM, agreeing with the above. In contrast, for the same target, a deterministic approximation at Stage One (with $\eta \approx 1/300$) leads to an efficiency increase of over an order of magnitude relative to the PsMRWM.

In [32] it is also suggested (though not implemented) that the noise in $\widehat{\pi}_a$ could be made positively correlated with the noise in $\widehat{\pi}$ since this should increase the Stage Two acceptance

probability on average. For example, in particle MCMC-based inference on the parameters of an SDE, $\hat{\pi}_a$ could use a coarser Euler-Maruyama time step than $\hat{\pi}$ and both could use the same driving Brownian motion. Use of the same Brownian motion would produce the highest correlation possible between the two estimators; however, the particle-filter resampling steps would ensure this was below 1.

The article [9] considers the output from a single particle filter, rather than the joint output from two correlated particle filters, and it is beyond the scope of this work to rigorously generalise this result. If, however, the joint noise vector $(W_{a,\Delta}^\infty, W_\Delta^\infty)$ in Proposition G.1 has a bivariate Gaussian distribution with a correlation of ρ then it is straightforward to generalise (G.2) and, for a given ρ , to optimise the efficiency with respect to $(\mu, \sigma_a^2, \sigma^2)$. Figure 8 plots the optimal $(\mu, \sigma_a^2, \sigma^2)$ and the efficiency relative to optimally tuned PsMRWM, all against the correlation, ρ , and under the assumption of a bivariate Gaussian density.

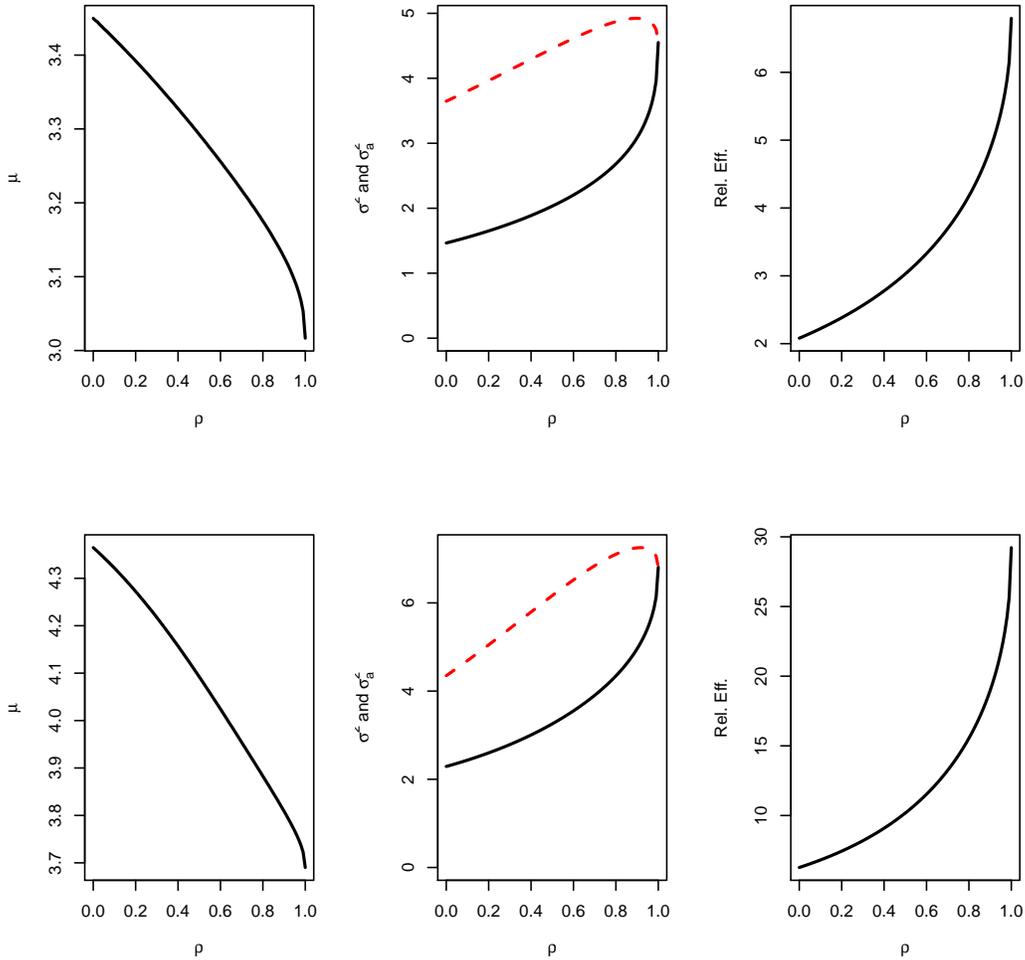


Figure 8: Optimal values for μ (left) and σ_a^2 (red dashed line) and σ^2 (solid black line) (centre), as well as the efficiency relative to optimised pseudo-marginal RWM (right), when $\eta = 0.1$ (top) and $\eta = 0.01$ (bottom). Plots created under the assumption of a bivariate Gaussian density for $(W_{a,\Delta}^\infty, W_\Delta^\infty)$.

As anticipated, the optimal efficiency increases with ρ . Also unsurprisingly, $\sigma_a^2 \geq \sigma^2$; further, the optimal variance tends to increase with ρ because a large positive correlation reduces the difference $W_\Delta^\infty - W_{a,\Delta}^\infty$. However, when $\rho = 1$, $\widehat{\sigma}_a^2 = \widehat{\sigma}^2$ precisely so that (in this limit where $\beta_1 = \beta_2 = 0$) $W_\Delta^\infty - W_{a,\Delta}^\infty = 0$ and acceptance is guaranteed at Stage Two. In this case the algorithm becomes PsRWM on the unbiased $\widehat{\pi}_a$, but the cost is $(\eta + \alpha_1)/\sigma^2$ instead of $1/\sigma^2$, potentially leading to a very large improvement in efficiency. In reality, when $\widehat{\pi}_a$ is biased, but the bias is small, strong positive correlation in the noises of the two approximations should, therefore, lead to substantial improvements in efficiency.

G.1 Steps in the derivation of (G.1) and (G.2)

The derivation of (G.1) and (G.2) follows analogous steps to those in Section 4 of the main text. We describe the intermediate results; the proofs are either special cases of or slight variations of the proofs of the equivalent steps for the results in the main text and are omitted.

Lemma 4.1 is still applicable, since $(q_\Delta^{(d)}(x^{(d)}, X^{(d),*}), s_\Delta^{(d)}(x^{(d)}, X^{(d),*}))$ and $W_a^{(d),*} - W_a^{(d)}$ are independent; we include the joint limit with $W_a^{(d),*} - W_a^{(d)}$ and $W^{(d),*} - W^{(d)}$ to aid with the exposition on correlated noise.

Proposition G.1. *Let Assumptions 1 and 2 hold for both the noise in the Stage One approximation and the noise in the Stage Two approximation, with variances of σ_a^2 and σ^2 , respectively. Further, let Part 1 of Assumptions 4 hold. Let $\{x_i\}_{i \geq 1}$ be the realisation of an i.i.d sequence marginally distributed as π . For $d \geq 1$, set $\mathbf{x}^{(d)} = (x_1, \dots, x_d) \in \mathbb{R}^d$ and let $\mathbf{X}^{(d),*}$ and $\mathbf{Z}^{(d)}$ be as defined in (4.1). For almost all realisations $\{x_i\}_{i \geq 1}$, the following limit*

$$\lim_{d \rightarrow \infty} \begin{bmatrix} q_\Delta^{(d)}(\mathbf{x}^{(d)}, \mathbf{X}^{(d),*}) \\ s_\Delta^{(d)}(\mathbf{x}^{(d)}, \mathbf{X}^{(d),*}) \\ W_a^{(d),*} - W_a^{(d)} \\ W^{(d),*} - W^{(d)} \end{bmatrix} = \begin{bmatrix} Q_\Delta^\infty \\ S_\Delta^\infty \\ W_{a,\Delta}^\infty \\ W_\Delta^\infty \end{bmatrix} \stackrel{\mathcal{D}}{\sim} \mathbf{N} \left(\begin{bmatrix} -\frac{1}{2}\mu^2 \\ \frac{1}{2}\mu^2\beta_1 \\ -\sigma_a^2 \\ -\sigma^2 \end{bmatrix}, \begin{bmatrix} \mu^2 & -\mu^2\beta_1 & 0 & 0 \\ -\mu^2\beta_1 & \mu^2\beta_2^2 & 0 & 0 \\ 0 & 0 & 2\sigma_a^2 & 0 \\ 0 & 0 & 0 & 2\sigma^2 \end{bmatrix} \right) \quad (\text{G.4})$$

holds in distribution.

Analogously to Proposition 4.1 (but, for simplicity of presentation, taking Assumption 2 as well) we then obtain the limiting acceptance probabilities as follows.

Proposition G.2. *Let Assumptions 1 and 2 hold both for the Stage One approximation and the Stage Two approximation, with variances of σ_a^2 and σ^2 respectively. Let Part 1 of Assumptions 4 hold. Then*

$$\lim_{d \rightarrow \infty} \alpha_1^{(d)}(\mathbf{X}^{(d)}, W_a^{(d)}) \stackrel{L^2}{=} \alpha_1 \quad \text{and} \quad \lim_{d \rightarrow \infty} \alpha_{12}^{(d)}(\mathbf{X}^{(d)}, W_a^{(d)}, W^{(d)}) \stackrel{L^2}{=} \alpha_{12}$$

where the limiting acceptance rates are given by

$$\begin{cases} \alpha_1 = \mathbb{E} [F(Q_\Delta^\infty + S_\Delta^\infty + W_{a,\Delta}^\infty)] \\ \alpha_{12} = \mathbb{E} [F(Q_\Delta^\infty + S_\Delta^\infty + W_{a,\Delta}^\infty) \times F(W_\Delta^\infty - S_\Delta^\infty - W_{a,\Delta}^\infty)] \end{cases} \quad (\text{G.5})$$

for $(Q_\Delta^\infty, S_\Delta^\infty, W_{a,\Delta}^\infty, W_\Delta^\infty)$ as described in (G.4). The dependence of α_1 and α_{12} upon $(\mu, \sigma_a^2, \sigma^2, \beta_1, \beta_2)$ is implicit.

H Envelope width

Figures 1 (left) and ?? in the main text provide a look-up, given η and $\alpha_{2|1}$ at the optimal parameter setting, of the ratio of the optimal tuning parameter (scaling for DARWM, scaling and variance for DAPMRWM) for the DA algorithm to the optimal value for the parent algorithm. However, in places the envelope of possible values is relatively wide. Here, we investigate this envelope.

Figure 9 shows the plots in Figures 1 (left) and ??, but specifically for $\eta = 0.01$, and with points coloured according to β_2 . The exception is the top-right plot, which repeats Figure 1 (left) but coloured according to theoretical efficiency.

Firstly we examine the choice of scaling. The theoretical efficiency curve shows clearly that for the DARWM with $\eta = 0.01$, moderate to large efficiencies are only obtained when $\alpha_{2|1}$ is large. This pattern is repeated, though not shown here, across different η values and for PMRWM (where for $\eta = 0.01$, the ‘yellow’ region, corresponding to between 1/3 and 2/3 of the maximum value, is at an x axis value of 2.0 rather than 3.0). When $\alpha_{2|1}$ is large the envelope for $\hat{\mu}$ is narrow. Thus for the more successful DA strategies there is little uncertainty in the recommended choice of scaling.

For large $\alpha_{2|1}$, the fine spray of points above the main line correspond to scenarios where $\beta_1 \approx \beta_2$. Recall from Section 3.2 that $|\beta_1| \leq \beta_2$, and notice that $\beta_1 = \beta_2$ corresponds to the case where, in deriving the bound, the Cauchy-Schwarz inequality is exact and so the gradient of the random function, S that is the error in the approximation, *is proportional to* the gradient of $\log \pi$, with a positive coefficient of proportionality; the random function from which our approximation derives is, in fact, deterministic. This is counter to the generality of our set up of using a realisation from a *random* function, and, moreover, corresponds to an approximation which matches the mode(s) of the target and has steeper gradients everywhere else, which would not be sensible. We, therefore, recommend ignoring this fine spray of points (indeed we have already removed points with $|\beta_1|/\beta_2 > 0.9$) and using the more solid part of the envelope.

The top-left and bottom-left panels show that amongst the less efficient DA strategies, the more efficient of these occur nearer the top of the envelope. With a chicken-and-egg assumption that one would not be using a DA approximation unless it was at least reasonably efficient, we would suggest choose a value towards the top of the envelope.

In contrast to the behaviour for the scaling parameter, for relatively large $\alpha_{2|1}$, the range of possible values for σ^2 is large, and all except the largest values correspond to the lower β_2 values. Interestingly, in the example in Section J.2 the curve for efficiency as a function of σ^2 has a flat peak which covers roughly a doubling of the variance. Thus we conjecture that choosing any value within the envelope will lead to a close-to-optimal algorithm.

I DARWM: Gaussian target with logistic approximation

We consider a scenario where the true target is a product of standard Gaussians and the deterministic approximation is a product of logistic densities with a mode at φ_1 and inverse-scale parameter φ_2 ,

$$\pi(x) \propto \exp \left\{ -\frac{1}{2} \sum_{i=1}^d x_i^2 \right\} \quad \text{and} \quad \hat{\pi}_a(x) \propto \prod_{i=1}^d \frac{e^{\varphi_2(x_i - \varphi_1)}}{(1 + e^{\varphi_2(x_i - \varphi_1)})^2}. \quad (\text{I.1})$$

We consider fourteen scenarios: ten different combinations of values for (φ_1, φ_2) , three approximations where the values of φ_1 or φ_2 vary from component to component, and the ‘perfect approximation’, $\pi_a = \pi$; see Table 1 for further details.

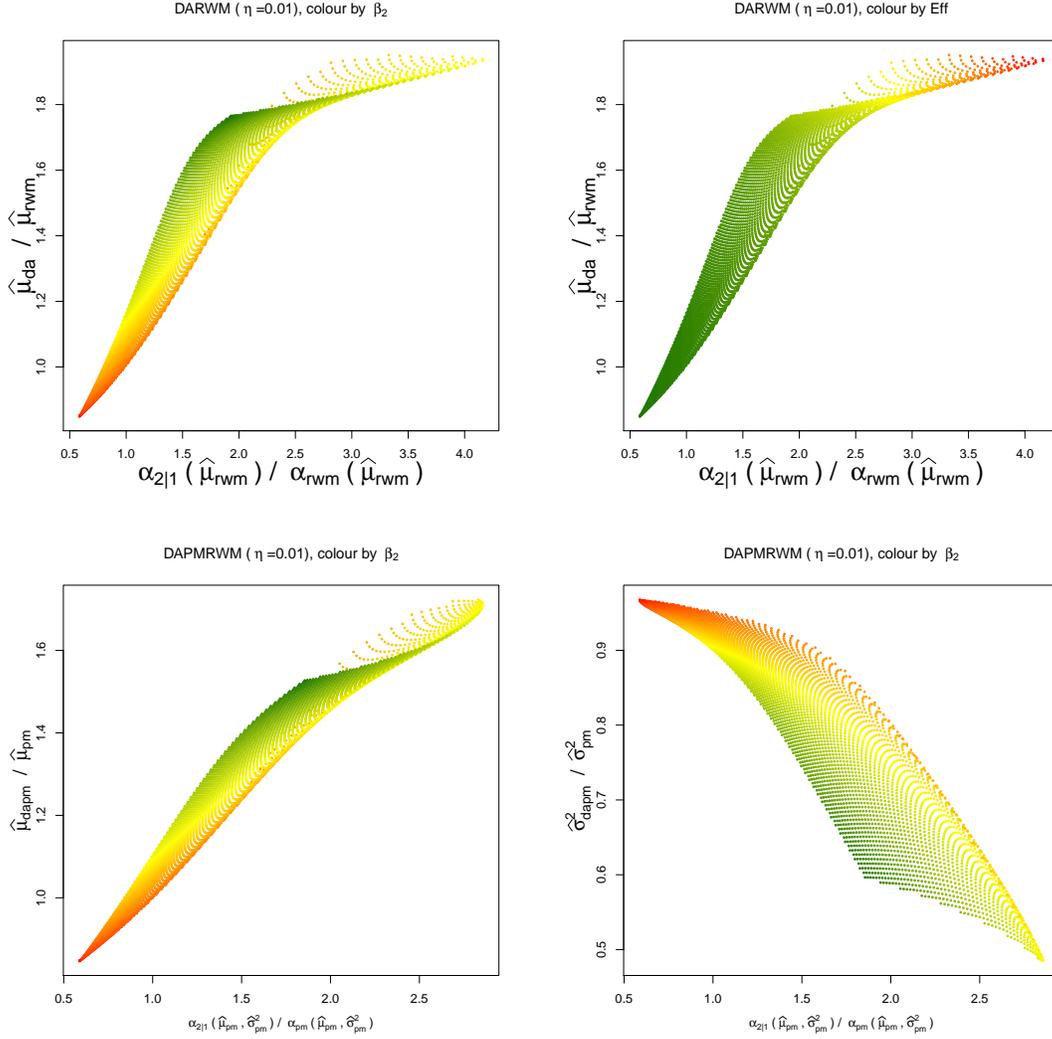


Figure 9: Scatter plots of (top) $\hat{\mu}_{\text{darwm}}/\hat{\mu}_{\text{rwm}}$, (bottom left) $\hat{\mu}_{\text{dapm}}/\hat{\mu}_{\text{pm}}$ and (bottom right) $\hat{\sigma}_{\text{dapm}}^2/\hat{\sigma}_{\text{pm}}^2$ (right) all against $\alpha_{2|1}(\hat{\mu}_{\text{pm}}, \hat{\sigma}_{\text{pm}}^2)$. Colour is according to β_2 except for the top right plot where colour is by theoretical efficiency; lowest values appear in dark green and highest values appear in red. All plots are for $\eta = 0.01$.

Algorithm	φ_1	φ_2	β_1	β_2	α_1	$\alpha_{2 1}$
RWM					0.2616	
DA	0.0	0.6	0.834	0.834	0.261	0.128
DA	0.0	1.2	0.441	0.449	0.069	0.533
DA	0.0	1.8	-0.042	0.262	0.041	0.738
DA	0.0	2.3	-0.467	0.649	0.034	0.595
DA	0.0	2.7	-0.810	1.025	0.032	0.492
DA	0.5	1.2	0.466	0.552	0.370	0.547
DA	1.0	1.2	0.535	0.763	0.140	0.151
DA	1.5	1.2	0.630	0.979	0.482	0.276
DA	0.6	1.8	0.056	0.681	0.0650	0.279
DA	0.5	2.3	-0.351	0.941	0.049	0.289
DA	0.0	1.5–2.0			0.248	0.772
DA	0.0	1.2–2.7			0.238	0.609
DA	0.0–1.0	1.2			0.377	0.517

Table 1: Values of φ_1 and φ_2 used in (I.1), and the corresponding values of β_1 and β_2 (where calculable), $\alpha_1(\widehat{\lambda}_{\text{rwm}})$ and $\alpha_{2|1}(\widehat{\lambda}_{\text{rwm}})$.

Empirical effective sample sizes (ESSs) for each of the d components are calculated using the coda package in R [42]; the overall ESS is taken to be the average of the ESSs over the d individual components. All algorithms were run for 10^6 iterations.

We first obtained the optimal scaling, $\widehat{\lambda}_{\text{rwm}}$, for a RWM targeting π by optimising the empirical ESS, and evaluated $\alpha_{\text{rwm}}(\widehat{\lambda}_{\text{rwm}})$ as well as the empirical ESS at this tuning. Then we ran the DA algorithm with this scaling to find $\alpha_{2|1}(\widehat{\lambda}_{\text{rwm}})$. Next, we artificially induced three different values of η : 0.1, 0.01, 0.001 and evaluated the efficiency, (empirical ESS-100) / CPU time) over a grid of possible scalings, λ , to find the optimal scaling. The regularisation penalty is needed because for very poorly mixing chains the empirical ESS tends to overestimate the true efficiency.

Figure 10 reproduces Figure 1 for $\eta \in \{0.1, 0.01, 0.001\}$, but in three shades of grey, then plots $\widehat{\lambda}_{\text{da}}/\widehat{\lambda}_{\text{rwm}}$ ($d = 10$ and $d = 25$); a similar plot for the predicted and realised relative efficiency ($d = 25$) against $\alpha_{2|1}(\widehat{\lambda}_{\text{da}})/\alpha_{\text{rwm}}(\widehat{\lambda}_{\text{rwm}})$ is also provided. At $d = 10$ the theory sometimes slightly overestimates the increase in scaling that is required, although (not shown) the predicted range of gains in efficiency is accurate except when η is small and $\alpha_{2|1}(\widehat{\mu}_{\text{rwm}})/\alpha_{\text{rwm}}(\widehat{\mu}_{\text{rwm}})$ is large, but by $d = 25$ the theoretical prediction of the ratio is quite accurate, as is the predicted efficiency gain. Essentially, with a larger scaling and a smaller dimension the diffusion approximation is less accurate.

J DAPsRWM: tuning advice and Lotka Volterra simulation study

J.1 DAPsRWM tuning advice

Our theory predicts behaviour in terms of the theoretical scaling, μ , and the variance σ^2 , whereas the quantities the user wishes to tune are the actual scaling, λ , and the number of particles, m .

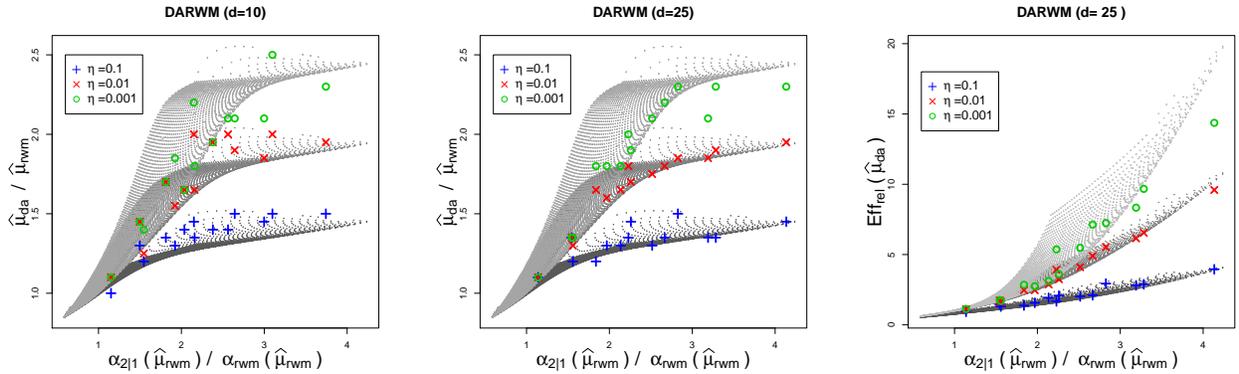


Figure 10: Scatter plots of $\widehat{\mu}_{\text{da}}/\widehat{\mu}_{\text{rwm}}$ ($d = 10$, left and $d = 25$, centre) and $\text{Eff}_{\text{da}}^{\text{rel}}(\widehat{\mu}_{\text{da}})$ ($d = 25$, right), vs $\alpha_{2|1}(\widehat{\mu}_{\text{rwm}})$, partitioned by η .

Since $\lambda \propto \mu$ (3.2), and $\sigma^2 \propto 1/m$ (see the discussion following Assumptions 3), but with unknown proportionality constants, we use our theory to predict $\widehat{\mu}_{\text{da}}/\widehat{\mu}_{\text{rwm}} = \widehat{\lambda}_{\text{da}}/\widehat{\lambda}_{\text{rwm}}$ and $\widehat{\sigma}_{\text{dapm}}^2/\widehat{\sigma}_{\text{pm}}^2 = \widehat{m}_{\text{pm}}/\widehat{m}_{\text{dapm}}$. Importantly, our tuning guidelines are based on quantities that can straightforwardly and robustly be estimated from a short MCMC trajectory.

As described in Section 5.2 there are two possible tuning strategies for the DAPsMRWM.

- Analogously to a strategy for the pseudo-marginal RWM (see Section 2.2), since the effect of altering the number of particles is approximately orthogonal to the effect of altering the scaling provided $\sigma^2 \geq 1$, choose a number of particles that gives $\sigma^2 \geq 1$, conditional on this tune the scaling to optimise efficiency and then, with this scaling, choose the number of particles to optimise efficiency.
- Alternatively, given an approximately optimally tuned pseudo-marginal RWM algorithm, together with the parameters \widehat{m}_{pm} and $\widehat{\lambda}_{\text{pm}}$, a single run of the DAPsMRWM with these parameters provides $\alpha_{2|1}(\widehat{\lambda}_{\text{pm}}, \widehat{\sigma}_{\text{pm}}^2)$ and the value of $\eta > 0$. One may then obtain the ratios $\widehat{\lambda}_{\text{dapm}}/\widehat{\lambda}_{\text{pm}}$ and $\widehat{m}_{\text{pm}}/\widehat{m}_{\text{dapm}}$ from the tuning scatter plots in Figure 4.

J.2 Lotka-Volterra simulation study

To illustrate the advice for the DAPsMRWM, and provide a check on its validity, we consider a Lotka-Volterra predator-prey model [11]. The model describes the continuous time evolution of $\mathbf{U}_t = (U_{1,t}, U_{2,t})$ where $U_{1,t}$ (prey) and $U_{2,t}$ (predator) are non-negative integer-values processes. Starting from an initial value, which is assumed known for simplicity, \mathbf{U}_t evolves according to a Markov jump process (MJP) parameterised by rate constants $\mathbf{c} = (c_1, c_2, c_3)$; details of the state transitions are provided in Section J.3. The process is easily simulated via the Gillespie algorithm [30] and the pseudo-marginal RWM scheme is straightforward to apply [31]. We assume that the MJP is observed with Gaussian error every time unit for n time units, $t = 1, \dots, n$: $Y_{1,t} \sim \mathbf{N}(u_{1,t}, s_1^2)$ and $Y_{2,t} \sim \mathbf{N}(u_{2,t}, s_2^2)$, independently. As all of the parameters of interest must be strictly positive, we consider inference for

$$\mathbf{x} = (\log(c_1), \log(c_2), \log(c_3), \log(s_1), \log(s_2)).$$

Parameter values and prior distributions are provided in Section J.3. The DAPsMRWM scheme requires a computationally cheap approximation of the MJP. We follow [32] by constructing a linear noise approximation (see e.g. [60]), detailed in Section J.3.

For a pseudo-marginal RWM scheme [57] suggests that for a Gaussian target (where, for each principal component, I is known) proposals with a variance of $\mathbf{V}_{\text{Gauss}} = (2.56^2/d) \times \text{Var}(\mathbf{X})$ would be optimal. We propose Gaussian jumps with a variance of $\mathbf{V}_{\text{prop}} = \gamma^2 \widehat{\mathbf{V}}_{\text{Gauss}}$, where $\text{Var}(\mathbf{X})$, has been replaced with an approximation, $\widehat{\text{Var}}(\mathbf{X})$, created from an initial run. If the target were in fact a high-dimensional Gaussian, and the variance approximation were exact, then γ would correspond exactly to the theoretical scaling, $\mu/\widehat{\mu}_{\text{pm}}$, and $\widehat{\gamma}_{\text{pm}}$ would be 1.0. We found that the pseudo-marginal RWM was optimised at $\widehat{\gamma}_{\text{pm}} \approx 1.2$. [57] suggests that the optimal number of particles should lead to a variance in $\log \widehat{\pi}$ of approximately 3.3. We found an optimum of $m = 180$, which occurred when the $\text{Var}[\log \widehat{\pi}(x_*)]$ (with x_* an initial estimate of the componentwise posterior median) was approximately 2.9. The mean acceptance probability at this optimal tuning was $\alpha_{\text{pm}} \approx 8.0\%$ and the empirical efficiency (minimum, over each parameter component, effective sample size per second) was 0.067.

The DAPsMRWM with $\gamma = 1.2$ and $m = 180$ gave $\alpha_{2|1} \approx 20.7\%$, so that $\alpha_{2|1}/\alpha_{\text{pm}} \approx 2.6$; timing diagnostics gave $\eta = 0.0014$. For this combination, the tuning scatter plots suggest increasing the scaling by a factor of around 2.0, decreasing the variance by a factor of between 0.7 and 0.8, and that this should lead to an increase in efficiency of a factor of between 6 and 7. The tuning suggestions translate to $\gamma \approx 2.4$ and $m \approx 225 - 255$. Alternatively, Figure 5 suggests that provided $\sigma^2 > 1$, m and γ may be tuned independently.

To confirm that the practical advice is reasonable and to test some of the other predictions of our theory, the number of particles m was varied between 80 and 2000 and, for each m , the scaling γ was varied between 1 and 4.5. For each (m, γ) pair, a long MCMC run (of at least 4×10^5 iterations) was performed. Figure 11 shows empirical efficiency as a function of the scaling γ (with a varying number of particles m) and as a function of m (for various γ) and provides empirical evidence of the insensitivity of the optimal choice of scaling, γ , to the value of σ^2 , for values of $\sigma^2 \geq 0.89$; furthermore, for variances below 0.89 the optimal scaling increases, as predicted by our theory. Empirical efficiencies for each run, as well as empirical Stage-One and Conditional Stage-Two acceptance rates are given in Table 2, and back up the heuristic from the figures, that $\gamma = 2.5$ gives close to the optimal efficiency, with $\widehat{\gamma} \approx 3.1$, and $\widehat{m} \approx 220 - 250$ as predicted. The empirical efficiency gain from using the DAPsMRWM algorithm compared to the pseudo-marginal RWM algorithm was $0.441/0.067 \approx 6.6$, which is in the centre of the range predicted by the theory.

Proposition 4.2 proves that, subject to assumptions, the Stage 2 acceptance probability decreases as the variance in the log-posterior (σ^2) increases and the Stage 1 acceptance probability decreases as the scaling increases; Table 2 shows that these patterns are observed in our experiments.

As with standard MCMC, the samples from tuning runs could be combined with the sample from the run at the optimal parameter values to decrease the variance of any estimator still further.

J.3 Lotka-Volterra details

The Lotka-Volterra MJP is characterised by transitions over $(t, t + dt]$ of the form

$$\begin{aligned} \mathbb{P}(U_{1,t+dt} = u_{1,t} + 1, U_{2,t+dt} = u_{2,t} | u_{1,t}, u_{2,t}) &= c_1 u_{1,t} dt + o(dt), \\ \mathbb{P}(U_{1,t+dt} = u_{1,t} - 1, U_{2,t+dt} = u_{2,t} + 1 | u_{1,t}, u_{2,t}) &= c_2 u_{1,t} u_{2,t} dt + o(dt), \\ \mathbb{P}(U_{1,t+dt} = u_{1,t}, U_{2,t+dt} = u_{2,t} - 1 | u_{1,t}, u_{2,t}) &= c_3 u_{2,t} dt + o(dt). \end{aligned}$$

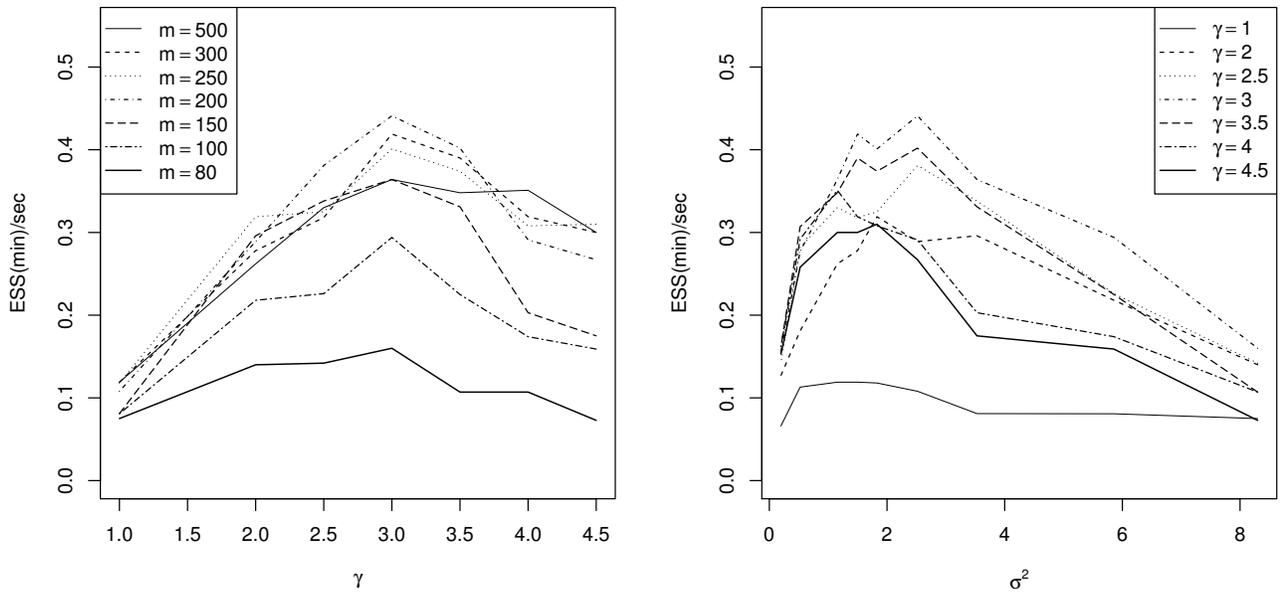


Figure 11: Empirical efficiency measured as the effective sample size per CPU second. The left-hand panel gives the efficiency plotted against γ for various numbers of particles. The right-hand panel gives the efficiency plotted against σ (estimated at the posterior median), for various scalings.

γ	m	80	100	150	200	250	300	500	800	2000
	σ^2	8.30	5.86	3.53	2.52	1.83	1.50	0.89	0.52	0.20
1	mESS/s	0.0750	0.0808	0.0810	0.108	0.118	0.119	0.119	0.113	0.0661
	$\hat{\alpha}_1$	0.256	0.255	0.257	0.255	0.257	0.254	0.254	0.255	0.258
	$\hat{\alpha}_{2 1}$	0.0651	0.0883	0.170	0.237	0.289	0.341	0.447	0.547	0.692
2	mESS/s	0.140	0.218	0.296	0.289	0.319	0.278	0.262	0.181	0.127
	$\hat{\alpha}_1$	0.0556	0.0514	0.0489	0.0503	0.0517	0.0520	0.0513	0.0517	0.0505
	$\hat{\alpha}_{2 1}$	0.0619	0.0895	0.163	0.213	0.286	0.313	0.438	0.522	0.674
2.5	mESS/s	0.142	0.226	0.338	0.381	0.325	0.318	0.330	0.282	0.142
	$\hat{\alpha}_1$	0.0244	0.0237	0.0234	0.0259	0.0264	0.0234	0.0241	0.0230	0.0250
	$\hat{\alpha}_{2 1}$	0.0600	0.0815	0.159	0.218	0.252	0.312	0.434	0.523	0.675
3	mESS/s	0.160	0.294	0.364	0.441	0.401	0.419	0.364	0.277	0.156
	$\hat{\alpha}_1$	0.0143	0.0123	0.0119	0.0114	0.0131	0.0120	0.0114	0.0124	0.0121
	$\hat{\alpha}_{2 1}$	0.0416	0.101	0.152	0.233	0.274	0.320	0.426	0.516	0.673
3.5	mESS/s	0.107	0.225	0.331	0.402	0.374	0.390	0.348	0.307	0.162
	$\hat{\alpha}_1$	0.00629	0.00789	0.00763	0.00684	0.00669	0.00663	0.00725	0.00634	0.00694
	$\hat{\alpha}_{2 1}$	0.0550	0.0869	0.170	0.237	0.273	0.312	0.424	0.534	0.673
4	mESS/s	0.107	0.174	0.176	0.291	0.308	0.319	0.351	0.292	0.162
	$\hat{\alpha}_1$	0.00343	0.00318	0.00401	0.00388	0.00372	0.00357	0.00377	0.00402	0.00418
	$\hat{\alpha}_{2 1}$	0.0680	0.105	0.151	0.215	0.287	0.310	0.407	0.500	0.681
4.5	mESS/s	0.0728	0.159	0.150	0.267	0.310	0.300	0.300	0.258	0.153
	$\hat{\alpha}_1$	0.00220	0.00183	0.00207	0.00247	0.00230	0.00256	0.00224	0.00249	0.00226
	$\hat{\alpha}_{2 1}$	0.0527	0.111	0.143	0.213	0.265	0.280	0.424	0.491	0.658

Table 2: Minimum effective sample size (mESS) per second, stage 1 acceptance probability $\hat{\alpha}_1$ and stage 2 acceptance probability $\hat{\alpha}_{2|1}$ as functions of the number of particles m and scaling γ . The variance (σ^2) of the estimated log-posterior at the median is also shown for each choice of m .

Data were simulated using an initial value $\mathbf{u}_0 = (71, 79)$ for $n = 50$ time units with $\mathbf{c} = (1.0, 0.005, 0.6)$ and $s_1 = s_2 = 8$. The parameters $\mathbf{x} = (\log c_1, \log c_2, \log c_3, \log s_1, \log s_2)$ were assumed to be independent *a priori* with proper Uniform densities on the interval $[-8, 8]$ ascribed to X_i , ($i = 1, \dots, 5$).

Under the linear noise approximation (LNA) we have that $\mathbf{U}_t \stackrel{\mathcal{D}}{\sim} \mathbf{N}(\mathbf{z}_t + \mathbf{m}_t, \mathbf{V}_t)$ where \mathbf{z}_t , \mathbf{m}_t and \mathbf{V}_t satisfy a coupled ODE system

$$\begin{cases} \dot{\mathbf{z}}_t &= \mathbf{S}\mathbf{h}(\mathbf{z}_t, \mathbf{c}) \\ \dot{\mathbf{m}}_t &= \mathbf{F}_t\mathbf{m}_t \\ \dot{\mathbf{V}}_t &= \mathbf{V}_t\mathbf{F}_t^T + \mathbf{S}\text{diag}\{\mathbf{h}(\mathbf{z}_t, \mathbf{c})\}\mathbf{S}^T + \mathbf{F}_t\mathbf{V}_t \end{cases} \quad (\text{J.1})$$

For the Lotka-Volterra model, the rate vector $\mathbf{h}(\mathbf{z}_t, \mathbf{c})$, stoichiometry matrix \mathbf{S} and Jacobian matrix \mathbf{F}_t are given by $\mathbf{h}(\mathbf{z}_t, \mathbf{c}) = (c_1 z_{1,t}, c_2 z_{1,t} z_{2,t}, c_3 z_{2,t})$ as well as

$$\mathbf{S} = \begin{pmatrix} 1 & -1 & 0 \\ 0 & 1 & -1 \end{pmatrix}, \quad \text{and} \quad \mathbf{F}_t = \begin{pmatrix} c_1 - c_2 z_{2,t} & -c_2 z_{1,t} \\ c_2 z_{2,t} & c_2 z_{1,t} - c_3 \end{pmatrix}.$$

We now describe an algorithm for evaluating the posterior (up to proportionality) under the LNA. For further details regarding the LNA and its use as an approximation to a MJP, we refer the reader to [24] and [32]. For simplicity of exposition we assume an observation regime of the form $\mathbf{Y}_t = \mathbf{U}_t + \boldsymbol{\varepsilon}_t$ with $\boldsymbol{\varepsilon}_t \sim \mathbf{N}(0, \boldsymbol{\Sigma})$ where $\boldsymbol{\varepsilon}_t$ is a length- d_x Gaussian random vector. Suppose that \mathbf{U}_1 is fixed at some value \mathbf{u}_1 . The marginal likelihood (and hence the posterior up to proportionality) under the LNA, $\pi_a(\mathbf{y}_{1:n}|\mathbf{x})$ can be obtained as follows.

1. Initialisation. Compute $\pi_a(\mathbf{y}_1|\mathbf{x}) = \varphi(\mathbf{y}_1; \mathbf{u}_1, \boldsymbol{\Sigma})$ where $\varphi(\mathbf{y}_1; \mathbf{u}_1, \boldsymbol{\Sigma})$ denotes the Gaussian density with mean vector \mathbf{u}_1 and variance matrix $\boldsymbol{\Sigma}$. Set $\mathbf{a}_1 = \mathbf{u}_1$ and \mathbf{C} to be the $d_x \times d_x$ matrix of zeros.
2. For times $t = 1, 2, \dots, n - 1$,
 - (a) Prior at $t + 1$. Initialise the LNA with $\mathbf{z}_t = \mathbf{a}_t$, $\mathbf{m}_t = \mathbf{0}$ and $\mathbf{V}_t = \mathbf{C}_t$. Note that $\mathbf{m}_s = \mathbf{0}$ for all $s > t$. Integrate the ODE system (J.1) forward to $t + 1$ to obtain \mathbf{z}_{t+1} and \mathbf{V}_{t+1} . Hence $\mathbf{X}_{t+1}|\mathbf{y}_{1:t} \sim \mathbf{N}(\mathbf{z}_{t+1}, \mathbf{V}_{t+1})$.
 - (b) One-step forecast. Using the observation equation, we have that $\mathbf{Y}_{t+1}|\mathbf{y}_{1:t} \sim \mathbf{N}(\mathbf{z}_{t+1}, \mathbf{V}_{t+1} + \boldsymbol{\Sigma})$. Compute $\pi_a(\mathbf{y}_{1:t+1}|\mathbf{x}) = \pi_a(\mathbf{y}_{1:t}|\mathbf{x}) \varphi(\mathbf{y}_{t+1}; \mathbf{z}_{t+1}, \mathbf{V}_{t+1} + \boldsymbol{\Sigma})$.
 - (c) Posterior at $t + 1$. Combining the distributions in (a) and (b) gives $\mathbf{U}_{t+1}|\mathbf{y}_{1:t+1} \sim \mathbf{N}(\mathbf{a}_{t+1}, \mathbf{C}_{t+1})$ where $\mathbf{a}_{t+1} = \mathbf{z}_{t+1} + \mathbf{V}_{t+1}(\mathbf{V}_{t+1} + \boldsymbol{\Sigma})^{-1}(\mathbf{y}_{t+1} - \mathbf{z}_{t+1})$ and $\mathbf{C}_{t+1} = \mathbf{V}_{t+1} - \mathbf{V}_{t+1}(\mathbf{V}_{t+1} + \boldsymbol{\Sigma})^{-1}\mathbf{V}_{t+1}$.