

Markov chain Monte Carlo for exact inference for diffusions

Sermaidis, G.* Papaspiliopoulos, O.† Roberts, G.O.‡ Beskos, A.§ and Fearnhead, P.*

Abstract

We develop exact Markov chain Monte Carlo methods for discretely-sampled, directly and indirectly observed diffusions. The qualification "exact" refers to the fact that the invariant and limiting distribution of the Markov chains is the posterior distribution of the parameters free of any discretisation error. The class of processes to which our methods directly apply are those which can be simulated using the most general to date exact simulation algorithm. The article introduces various methods to boost the performance of the basic scheme, including reparametrisations and auxiliary Poisson sampling. We contrast both theoretically and empirically how this new approach compares to irreducible high frequency imputation, which is the state-of-the-art alternative for the class of processes we consider, and we uncover intriguing connections. All methods discussed in the article are tested on typical examples.

Keywords: Exact inference; Exact simulation; Markov chain Monte Carlo; Stochastic differential equation; Transition density

1 Introduction

Diffusion processes provide a flexible framework for modelling phenomena which evolve randomly and continuously in time and are extensively used throughout Science, e.g. in finance (Ait-Sahalia and Kimmel, 2007), biology (Golightly and Wilkinson, 2006), molecular kinetics (Horenko and Schütte, 2008), pharmacokinetics/pharmacodynamics (Picchini *et al.*, 2010) and spatio-temporal modelling (Brown *et al.*, 2000).

A time-homogeneous diffusion process $V \in \mathbb{R}^d$ is a Markov process defined as the solution to a stochastic differential equation (SDE):

$$dV_s = \beta(V_s; \theta_1)ds + \sigma(V_s; \theta_2)dW_s, \quad V_0 = v, s \geq 0, \quad (1)$$

where W is a d -dimensional standard Brownian motion. The functions $\beta : \mathbb{R}^d \times \Theta_1 \rightarrow \mathbb{R}^d$ and $\sigma : \mathbb{R}^d \times \Theta_2 \rightarrow \mathbb{R}^{d \times d}$ are known as the *drift* and *diffusion coefficient* respectively, and are allowed to depend on an unknown parameter $\theta = (\theta_1, \theta_2) \in \Theta \subset \mathbb{R}^p$. The assumption of distinct parameters for each functional is by no means restrictive and is adopted here for ease of presentation. We assume that σ is invertible and make the usual set of assumptions on β and σ to ensure that (1) has a unique weak non-explosive solution, see for example Theorem 5.2.1 of Øksendal (2003); see also Section 2.

Even though the process is defined in continuous time, the available data consist of observations recorded at a set of discrete time points,

$$Y := \{V_{t_0}, V_{t_1}, \dots, V_{t_n}\}, \quad 0 \leq t_0 < t_1 < \dots < t_n.$$

Statistical inference is pursued in a Bayesian framework where prior beliefs about the parameters, encoded via a prior density $\pi(\theta)$, are updated on the basis of the available data through the discrete-time likelihood to yield the posterior beliefs, encoded via the posterior density:

$$\pi(\theta | Y) \propto \pi(\theta) \prod_{i=1}^n p_{\Delta t_i}(V_{t_{i-1}}, V_{t_i}; \theta), \quad (2)$$

*Lancaster University

†corresponding author, Universitat Pompeu Fabra, omiros.papaspiliopoulos@upf.edu

‡Warwick University

§UCL

where $\Delta t_i = t_i - t_{i-1}$ is the time increment between consecutive observations and

$$p_t(v, w; \theta) = P(V_t \in dw \mid V_0 = v)/dw, \quad t > 0, v, w \in \mathbb{R}^d,$$

is the transition density of the process.

Bayesian and generally likelihood-based inference in this context is hindered by the unavailability of the transition density and sufficiently accurate approximations to the density exist only when t is sufficiently small. One strand of the literature approaches the inference problem by resorting to Monte Carlo data augmentation (DA), according to the following principle. First, a DA scheme is constructed by identifying auxiliary variables such that the joint density of those and the observations, known as complete likelihood, is analytically available. Subsequently, inference is performed by employing a Markov chain Monte Carlo (MCMC) algorithm which targets the posterior density of parameters and auxiliary variables. Early DA schemes were based on imputation of a finite number of points, say M , of the latent diffusion bridges $\{V_s, s \in (t_{i-1}, t_i)\}$, see for example Eraker (2001); Elerian *et al.* (2001). The complete likelihood is still intractable but can be reasonably approximated using an Euler scheme which now operates on smaller time increments. The bias introduced in this approximation is eliminated by increasing M . There are three serious challenges with this approach. First, the simulation of the latent bridges conditionally on the parameters; this simulation is required in the "Imputation" step of a DA algorithm. This problem has been intensively studied, see for example Papaspiliopoulos and Roberts (2012) for a recent account. Second, the choice of M , at least in practice. A good approximation usually requires a large value of M which is typically found by repeated runs of the algorithm until the estimated posterior distributions show no change. This adds a substantial computational burden. Third, Roberts and Stramer (2001) showed that when θ_2 is unknown the mixing time of the MCMC algorithm is $\mathcal{O}(M)$. This is due to the quadratic variation identity, according to which any continuous-time path contains infinite information about θ_2 . Thus, in these early DA schemes reduction in bias comes with unbounded increase in the Monte Carlo variance. Roberts and Stramer (2001) constructed appropriate path-parameter transformations in order to yield a working DA scheme which is valid even in the limit $M \rightarrow \infty$, unlike those of the previous generation. We refer to the limiting case $M \rightarrow \infty$ as path augmentation (PA). An MCMC algorithm based on PA is not implementable in practice, since it involves infinite amount of computation. For finite M , we refer to the DA scheme as high frequency augmentation (HFA) and to the MCMC algorithm which targets it as approximate MCMC (AMCMC), the term reflecting the fact that bias is introduced due to the discretisation of the paths. More details on these schemes are given in Section 3.1.

A new generation of Monte Carlo schemes for diffusions was initiated with the introduction of the exact algorithm (EA) for the exact simulation of non-linear diffusions. The potential of using the EA to build an MCMC algorithm for parameter estimation was sketched in Beskos *et al.* (2006b) for a restrictive class of univariate diffusions.

In this paper we present a novel augmentation scheme, called exact data augmentation (EDA), and develop MCMC algorithms for all diffusions which can be simulated under the broader framework of the so-called EA3 (Beskos *et al.*, 2008). This generation of MCMC algorithms based on EDA is referred to as exact MCMC (EMCMC), and is such that their equilibrium distribution is the exact posterior distribution of the parameters, i.e., free of any discretisation error. We enhance algorithmic performance by designing noncentred reparametrisations and extend our methods to the case of indirect observations where interest lies in estimating both parameters and the latent diffusion process. A further contribution of this work is a theoretical investigation of the connection between EDA and PA. First, it is shown that EDA augments more information than PA. This is rather surprising since the former appears to augment only a small finite-dimensional distribution of the missing paths whereas the latter in principle augments continuous paths and in practice high-frequency approximations thereof. The key is that the extra augmentation in EDA creates conditional independence relationships which are exploited to apply an algorithm which targets an infinite-dimensional state using finite computation. This result also suggests that AMCMC for the same amount of computation is expected to mix faster than EMCMC; this is effectively another instance of the bias-variance tradeoff. This connection motivates a further observation which links the two approaches and suggests a way to improve the convergence rate of EMCMC using auxiliary Poisson sampling.

A comment on the applicability of the methods proposed here is due. The methods rely on a variance-stabilising transformation after which the diffusion has constant diffusion matrix and drift which is of gradient form, see Section 2.1 for details. The transformation poses little limitations for univariate processes, but might not even exist for general multivariate SDEs with coupling in the diffusion. On the other hand, multivariate processes with gradient drift structure and no coupling in the diffusion are rather standard in the framework of physical systems. Note that the variance-stabilising transformation is necessary for the HFA approach as well. In summary, the technology we develop here is not directly applicable to general stochastic volatility models, say, although

exploiting particular structures can push considerably these limitations, see for example Kalogeropoulos *et al.* (2010). Additionally, advances in the exact simulation of diffusions, as for example in Étoré and Martinez (2011); Gonçalves and Roberts (2012) would *eo ipso* lead to exact MCMC methods following the framework of this article. Irreducible DA schemes avoiding this transformation are also currently under investigation, see for example Golightly and Wilkinson (2008).

The article is structured as follows. Section 2 contains the background on assumptions, notations and recalls the EA. Section 3 presents formally the EDA and contrasts it to PA. Section 4 describes noncentred reparametrisations of the EDA and auxiliary Poisson sampling from improving algorithmic performance. Section 5 discusses extensions to indirect observations. Section 6 carries out a careful and extensive numerical comparison of several schemes. Section 7 closes with a discussion and the Appendix contains the proofs of main results.

2 Preliminaries

In this section we collect some necessary background. In terms of notation, $x^{(i)}$ or $x^{(ij)}$ denote the i th or (i, j) th element of a vector or matrix x , $\det[x]$, x^T and x^{-1} denote the determinant, transpose and inverse of a matrix x where appropriate, and I_d denotes the $d \times d$ identity matrix. For two vectors x and y , we define vectors $x \wedge y$ and $x \vee y$ such that $x^{(i)} \wedge y^{(i)} = x^{(i)} \wedge y^{(i)}$ and $x^{(i)} \vee y^{(i)} = x^{(i)} \vee y^{(i)}$. The Euclidean norm is denoted by $\|\cdot\|$. ∇_x and Δ_x denote the Jacobian matrix and Laplacian operators respectively, that is if $x \in \mathbb{R}^d$, then for functions $f_1 : \mathbb{R}^d \rightarrow \mathbb{R}^m$ and $f_2 : \mathbb{R}^d \rightarrow \mathbb{R}$

$$\nabla_x f_1(x) = \left[\frac{f_1^{(j)}(x)}{\partial x^{(i)}} \right]_{i=1, \dots, d; j=1, \dots, m} \quad \Delta_x f_2(x) = \sum_{i=1}^d \frac{\partial^2 f_2(x)}{\partial (x^{(i)})^2}.$$

We define $D : \mathbb{R}^d \times \Theta_2 \rightarrow \mathbb{R}$ as $D = |\det[\sigma]|^{-1}$ and $\gamma : \mathbb{R}^d \times \Theta_2 \rightarrow \mathbb{R}^{d \times d}$ as $\gamma = \sigma \sigma^T$, where σ is the diffusion coefficient. For function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ twice continuously differentiable on its domain, we denote the generator of (1) by

$$A_\theta f(v) = \sum_{i=1}^d \beta^{(i)}(v; \theta_1) \frac{\partial f(v)}{\partial v^{(i)}} + \frac{1}{2} \sum_{i,j=1}^d \gamma^{(ij)}(v; \theta_2) \frac{\partial^2 f(v)}{\partial v^{(i)} \partial v^{(j)}}.$$

Finally, $\mathcal{N}_t(u)$ denotes the density of a Gaussian random variable with mean vector 0 and covariance matrix tI_d evaluated at $u \in \mathbb{R}^d$.

2.1 Reducible diffusions of gradient type

The methods in this paper rely on the existence of a transformation η , known as Lamperti transformation, such that $\eta(V_s; \theta_2)$ solves an SDE with constant diffusion matrix. This transformation can be obtained for univariate diffusions rather trivially. For multivariate processes its existence is a subtle matter. In the elliptic case a sufficient condition is

$$[\text{LAM}]: [\nabla_v \eta(v; \theta_2)]^T = \sigma^{-1}(v; \theta_2),$$

which can be simplified to yield explicit conditions on the elements of σ^{-1} ; see for example Ait-Sahalia (2008). In the rest of the article we will assume the existence of this transformation and denote its inverse by η^{-1} . If $X_s := \eta(V_s; \theta_2)$ is the transformed diffusion, then by Itô's formula X solves

$$dX_s = \alpha(X_s; \theta) ds + dW_s, \quad X_0 = x = \eta(v; \theta_2), \quad s \geq 0, \quad (3)$$

where $\alpha : \mathbb{R}^d \times \Theta \rightarrow \mathbb{R}^d$, with

$$\alpha^{(k)}(u; \theta) = A_\theta \eta^{(k)} \{ \eta^{-1}(u; \theta_2); \theta_2 \}, \quad k = 1, \dots, d.$$

If $\tilde{p}_t(x, z; \theta)$, $x, z \in \mathbb{R}^d$ is the transition density of X , then the transition density of V can be expressed as

$$p_t(v, w; \theta) = D(w; \theta_2) \tilde{p}_t \{ \eta(v; \theta_2), \eta(w; \theta_2); \theta \}. \quad (4)$$

The methodology also requires certain conditions on the drift α of the transformed process. The following set of assumptions should hold for any $\theta \in \Theta$:

[SMOOTH]: $\alpha^{(k)}(\cdot; \theta)$ is continuously differentiable for $k = 1, \dots, d$.

[GRAD]: There exists $H : \mathbb{R}^d \times \Theta \rightarrow \mathbb{R}$ such that $\nabla_x H(x; \theta) = \alpha(x; \theta)$.

[LBOUND]: There exists $l(\theta) > -\infty$, such that $l(\theta) \leq \inf_{u \in \mathbb{R}^d} \frac{1}{2} \{ \|\alpha(u; \theta)\|^2 + \Delta_x H(u; \theta) \}$.

The first is a very weak condition and the third is rather mild too. The second identifies X as a diffusion of gradient-type, where H is called the potential function. When the diffusion is ergodic, its invariant log-density can be expressed directly in terms of H . This condition is trivially satisfied for univariate processes but is more restrictive for multivariate ones. Finally, we require that α is such that the probability law generated by the solution of (3) is absolutely continuous with respect to the Wiener measure. A particularly useful and weak set of conditions are given in Rydberg (1997); in the case of (3) if α is locally bounded the conditions simply require that the SDE be not explosive.

2.2 Exact simulation of diffusions

The EA is a rejection sampling algorithm on the space of diffusion paths, which uses Brownian path proposals and delivers the diffusion path revealed at a finite collection of random points. The path can be filled in later with no further reference to the target process. The main attraction of the algorithm is that the draws are from the exact finite-dimensional distribution. Here we focus on exact *diffusion bridge* simulation, i.e., obtain samples from (1) conditionally on the origin $V_0 = v$ and terminal point $V_t = w$. It turns out that this conditional simulation is really the key to DA methods for parameter estimation.

The target process (1) is transformed into one of unit diffusion matrix as described in Section 2.1. The problem is therefore reduced to the simulation of (3) conditionally on the origin $x = x(\theta_2) := \eta(v; \theta_2)$ and terminal point $y = y(\theta_2) := \eta(w; \theta_2)$. An X -bridge yields a V -bridge by applying the inverse transformation. Let $\mathbb{Q}_\theta^{(t,x,y)}$ denote the law of the X -bridge starting at x and terminating at y at time t , and $\mathbb{W}_\theta^{(t,x,y)}$ the law of a Brownian bridge conditioned on the same endpoints. The following lemma, which is a restatement of Lemma 1 in Beskos *et al.* (2006b), derives the density of the target law with respect to the Brownian bridge law.

Lemma 1. *The law $\mathbb{Q}_\theta^{(t,x,y)}$ is absolutely continuous with respect to $\mathbb{W}_\theta^{(t,x,y)}$ with density*

$$\frac{d\mathbb{Q}_\theta^{(t,x,y)}}{d\mathbb{W}_\theta^{(t,x,y)}}(X) = \frac{\mathcal{N}_t(y-x)}{\tilde{p}_t(x,y;\theta)} \exp \left\{ H(y;\theta) - H(x;\theta) - \frac{1}{2} \int_0^t \{ \|\alpha(X_s;\theta)\|^2 + \Delta_x H(X_s;\theta) \} ds \right\} \quad (5)$$

$$\propto \exp \left\{ - \int_0^t \phi(X_s;\theta) ds \right\} \leq 1, \quad (6)$$

where $\phi : \mathbb{R}^d \times \Theta \rightarrow \mathbb{R}_+$ is defined by

$$\phi(u;\theta) = \frac{1}{2} \{ \|\alpha(u;\theta)\|^2 + \Delta_x H(u;\theta) \} - l(\theta).$$

The EA is based on recognising (6) as the probability of a specific event from an inhomogeneous Poisson process of intensity $\phi(X_s;\theta)$ on $[0, t]$. Such processes can be simulated by constructing an upper bound for the variable intensity and using Poisson thinning. Assume that there exists a finite-dimensional random variable $L := L(X)$ and a positive function r such that

$$r(L;\theta) \geq \sup_{s \in [0,t]} \phi(X_s;\theta),$$

and let $\Phi = \{\Psi, \Upsilon\}$ be a homogeneous Poisson process of intensity $r(L;\theta)$ on $[0, t] \times [0, 1]$, with uniformly distributed points $\Psi = \{\psi_1, \dots, \psi_\kappa\}$ on $[0, t]$ and marks $\Upsilon = \{u_1, \dots, u_\kappa\}$ on $[0, 1]$, where $\kappa \sim \text{Po}[r(L;\theta)t]$. If N is the number of points of Φ below the graph $s \rightarrow \phi(X_s;\theta)/r(L;\theta)$, then

$$P(N = 0 | X) = \exp \left\{ - \int_0^t \phi(X_s;\theta) ds \right\}.$$

This implies a rejection sampler where a proposed path $X \sim \mathbb{W}_\theta^{(t,x,y)}$ is accepted as a path from $\mathbb{Q}_\theta^{(t,x,y)}$ according to the indicator

$$I(L, X, \Phi, v, w, \theta) := \prod_{j=1}^{\kappa} \mathbb{I} \left[\phi(X_{\psi_j}; \theta) / r(L; \theta) < u_j \right]. \quad (7)$$

The EA output is the collection $\{L(X), \Phi, S(X)\}$, where

$$S(X) := \{(0, X_0), (\psi_1, X_{\psi_1}), \dots, (\psi_\kappa, X_{\psi_\kappa}), (t, X_t)\}$$

is a skeleton of the accepted path. The algorithm is presented in Algorithm 1. The technical difficulty that underlies the implementation of the EA is the simulation of $L(X)$, and primarily the conditional simulation of a Brownian bridge given $L(X)$ for the evaluation of (7). This has led to the construction of three EAs that share the rejection sampling principle, but have a different range of applicability.

Algorithm 1 EA for diffusion bridges

- 1: simulate $L = L(X)$, where $X \sim \mathbb{W}_\theta^{(L, x, y)}$,
 - 2: simulate $\kappa \sim \text{Po}[r(L; \theta)t]$ and $\Phi = \{(\psi_j, v_j)\}_j$, $1 \leq j \leq \kappa$ uniformly on $[0, t] \times [0, 1]$,
 - 3: conditionally on L sample Brownian bridge X_{ψ_j} ,
 - 4: evaluate $I = I(L, X, \Phi, v, w, \theta)$ as in (7); if $I = 1$ then return $\{L, \Phi, S(X)\}$, otherwise go to 1.
-

2.2.1 The family of EAs

Each EA exploits the specific structure of the drift to construct $L(X)$. The EA1 (Beskos *et al.*, 2006a) is the simplest EA type and its framework is restricted by

Condition 1. $\phi(\cdot; \theta)$ is bounded above.

This condition ensures that $r(L; \theta) \equiv r(\theta)$, implying there is no need for constructing $L(X)$. As a consequence, step 3 of Algorithm 1 merely requires simulation of a Brownian bridge at time instances $\psi_1, \dots, \psi_\kappa$. The EA2 (Beskos *et al.*, 2006a) is applicable only when $d = 1$ and relaxes Condition 1 to a more mild one:

Condition 2. Either $\limsup_{u \rightarrow \infty} \phi(u; \theta) < \infty$ or $\limsup_{u \rightarrow -\infty} \phi(u; \theta) < \infty$.

For simplicity consider only the first case. The algorithm constructs the proposed path by first simulating its minimum, say m , and subsequently the remainder of the path conditioned on m . In this setting $L(X)$ is defined as the two-dimensional random variable $L(X) = \{m, \tau\}$, where τ is the time instance the minimum is attained. Then, the required upper bound is found as $r(L; \theta) = \sup_u \{\phi(u; \theta) : u \geq m\}$. Simulating a Brownian bridge conditionally on its minimum is based on a path transformation of two independent Bessel bridges, see Beskos *et al.* (2006b) for more details.

The EA3 (Beskos *et al.*, 2008) poses no upper boundedness conditions. For the moment assume $d = 1$ and consider a path X with initial point x and terminal point y at time t . The algorithm is based on creating a partition on the path space using a series of lower and upper bounds. For a user-specified constant $\delta > \sqrt{t/3}$ the partition consists of the sets $C(\epsilon, x, y)$, $\epsilon \in \mathbb{N}^*$, defined by

$$\begin{aligned} A(\epsilon, x, y) &= \left\{ \sup_{0 \leq s \leq t} X_s \in [y + (\epsilon - 1)\delta, y + \epsilon\delta] \right\} \cap \left\{ \inf_{0 \leq s \leq t} X_s > x - \epsilon\delta \right\}, \\ B(\epsilon, x, y) &= \left\{ \inf_{0 \leq s \leq t} X_s \in (x - \epsilon\delta, x - (\epsilon - 1)\delta] \right\} \cap \left\{ \sup_{0 \leq s \leq t} X_s < y + \epsilon\delta \right\}, \\ C(\epsilon, x, y) &= A(\epsilon, x, y) \cup B(\epsilon, x, y), \end{aligned}$$

where x, y as defined in the beginning of Section 2.

In the multidimensional case, sets $C_k(\epsilon) := C(\epsilon, x^{(k)}, y^{(k)})$ are constructed for each coordinate $X_s^{(k)}$, $k = 1, \dots, d$, and $L(X)$ is defined as the d -dimensional discrete random variable $L(X) = (L^{(1)}, \dots, L^{(d)})$ where $L^{(k)} = \epsilon_k$, $\epsilon_k \in \mathbb{N}^*$, if $X \in \cap_{k=1}^d C_k(\epsilon_k)$. Hence, $\{L^{(k)} \leq \epsilon_k, k = 1, \dots, d\} \equiv \{x^{(k)} - \epsilon_k\delta < X_s^{(k)} < y^{(k)} + \epsilon_k\delta, 0 \leq s \leq t, k = 1, \dots, d\}$. Figure 1a illustrates the construction for an arbitrary coordinate.

The random variable L is referred to as the *Brownian bridge layer* and a Brownian bridge path conditioned on this layer as the *layered Brownian bridge*. Conditioned on L , and using the continuity of $\phi(\cdot; \theta)$, the Poisson rate can be found as

$$r(L; \theta) = \sup \left\{ \phi(u; \theta) ; u^{(k)} \in (x^{(k)} - L^{(k)}\delta, y^{(k)} + L^{(k)}\delta), k = 1, \dots, d \right\}.$$

The exact mathematical and implementation details of sampling a layered Brownian bridge can be found in the original paper.

2.2.2 Computational considerations

The computational performance of EA depends on its acceptance probability. For any two fixed points x and y , expression (6) implies that the probability of accepting a proposed path is

$$a(x, y, t, \theta) := \mathbb{E}_{\mathbb{W}_\theta^{(t,x,y)}} \left[\exp \left\{ - \int_0^t \phi(X_s; \theta) ds \right\} \right], \quad (8)$$

thus suggesting that the acceptance probability decays exponentially to 0 as d or t increase.

Finally, we note that although EA3 has the widest framework of applicability, EA1 and EA2 should still be preferred whenever possible; simulating a layered Brownian bridge is a non-trivial task and is achieved by means of rejection sampling, thus adding to EA3 an extra level of computational complexity. In particular, an extensive empirical study by Peluchetti and Roberts (2008) suggests as a rule of thumb that EA3 is approximately 10 times slower than EA1.

3 Data augmentation for discretely observed diffusions

In this section we describe irreducible DA approaches for parameter estimation and present formally our EDA scheme. The auxiliary variables involved in EDA are intimately related to the EA output for diffusion bridges and lead to MCMC algorithms which involve no approximation to the statistical model of interest. At first, EDA appears totally different from the PA paradigm of Roberts and Stramer (2001). However, there are close but subtle links and the presentation in this section has been structured to naturally bring those out. Effectively, the PA scheme, which is recalled below, leads to two possibilities. One is its approximation by an HFA with some finite M , which leads to bias. The other is to consider an EA for the simulation of the auxiliary process identified by Roberts and Stramer (2001) and identify finite-dimensional auxiliary variables which then lead to the EDA scheme. Section 3.2 identifies those variables and Section 3.3 gives a theorem which establishes the connection between PA and EDA.

3.1 Irreducible path imputation and finite-dimensional approximations

We first outline the PA approach of Roberts and Stramer (2001). This scheme corresponds to the limiting case $M = \infty$, and leads to an idealised, yet impossible to implement, MCMC algorithm which involves imputing continuous path trajectories. The auxiliary processes are obtained after two path-parameter transformations of the original latent bridges. We then review the HFA scheme which is constructed by approximating the PA scheme with some finite M . The HFA scheme requires only condition [LAM], whereas EDA additionally requires [SMOOTH], [GRAD] and [LBOUND] for being able to employ EA. However, HFA can be considerably improved when [SMOOTH] and [GRAD] are also satisfied.

Consider for the moment only two observations from the diffusion process, $V_0 = v$ and $V_t = w$. The process is first transformed as $V_s \rightarrow X_s = \eta(V_s; \theta_2)$ as in Section 2.1. X -paths over bounded time increments only contain finite information for θ , since all parameters now relate with the drift, see (3). The transformed path starts at $x(\theta_2)$ and terminates at $y(\theta_2)$, which are both deterministic functions of θ_2 and the observations. This suggests that a DA scheme based on X will not work when θ_2 is unknown, since a realisation of X determines θ_2 through its endpoints. An alternative way to see the problem is to note that the collection of dominating measures $\{\mathbb{W}_\theta^{(t,x,y)}, \theta \in \Theta\}$ are mutually singular, and therefore a Gibbs algorithm based on this augmentation would be trapped in the support of one of these measures. This necessitates a further reparametrisation from $X_s \rightarrow \tilde{X}_s$, where

$$\tilde{X}_s := X_s - \left(1 - \frac{s}{t}\right)x(\theta_2) - \frac{s}{t}y(\theta_2), \quad s \in [0, t], \quad (9)$$

which forces the path to start and finish at 0 and essentially transforms the distribution $\mathbb{Q}_\theta^{(t,x,y)}$ so that the dominating measure, now given by $\mathbb{W}^{(t,0,0)}$, is independent of θ .

The PA is based on imputing \tilde{X} . Accounting for all observations, let $x_i(\theta_2) := \eta(V_{t_i}; \theta_2)$ and denote by $\tilde{X}_i = \{\tilde{X}_{i,s}, s \in [0, \Delta t_i]\}$, $i = 1, \dots, n$ the imputed paths. We introduce

$$G_t(X; \theta) = \exp \left(\int_0^t \alpha^T(X_s; \theta) dX_s - \frac{1}{2} \int_0^t \|\alpha(X_s; \theta)\|^2 ds \right), \quad \mu_{i,s}(\theta) = \left(1 - \frac{s}{\Delta t_i}\right)x_{i-1}(\theta_2) + \frac{s}{\Delta t_i}x_i(\theta_2),$$

and note that the inverse transformation of (9), $\tilde{X}_i \rightarrow X_i$ is given by $\tilde{X}_{i,s} + \mu_{i,s}(\theta)$. Then the joint posterior density of θ and imputed paths, $\pi_{PA}(\theta, \{\tilde{X}_i, 1 \leq i \leq n\} | Y)$, is proportional to

$$\pi(\theta) \prod_{i=1}^n D(V_i; \theta_2) \mathcal{N}_{\Delta t_i} \{x_i(\theta_2) - x_{i-1}(\theta_2)\} G_{\Delta t_i} (g_i(\tilde{X}_i; \theta); \theta), \quad (10)$$

with respect to $\text{Leb}^p \otimes_{i=1}^n \mathbb{W}^{(\Delta t_i, 0, 0)}$, where $g_i(\tilde{X}_i; \theta) := \{\tilde{X}_{i,s} + \mu_{i,s}(\theta), s \in [0, \Delta t_i]\}$. For a detailed derivation, the reader is referred to Section 3 of Roberts and Stramer (2001).

The joint posterior density is not be computable since the augmented paths cannot be represented by a finite number of variables, hence the integrals cannot be computed. Instead, the paths are approximated by vectors of size $M + 2$, $\{\tilde{X}_{i,j\Delta t_i/(M+1)}, j = 0, \dots, M + 1\}$, and the integrals are approximated numerically, typically by Riemann sums, to yield $\pi_{HFA,M}(\theta, \{\tilde{X}_i, 1 \leq i \leq n\} | Y)$, where by an abuse of notation we let \tilde{X}_i denote the path and its discretisation. This introduces a bias in the inference for θ . The approximated posterior is targeted by an MCMC algorithm which updates in turns θ and $\tilde{X}_i, i = 1, \dots, n$ according to their conditional densities. Crucially for the efficiency of the algorithm, the auxiliary processes \tilde{X}_i are independent over i conditionally on θ and Y , and thus can be updated sequentially. Each update is typically performed by proposing Brownian bridge skeletons and accepting them according to (10). Algorithm 2 is a typical AMCMC implementation, where updates of θ are obtained from a Metropolis-Hastings step with proposal kernel q .

Algorithm 2 AMCMC

- 1: Choose θ^0 , set $\delta_i = \Delta t_i / (M + 1)$ and $\tilde{X}_i^0 = \{\tilde{X}_{i,j\delta_i}^0\}_j, 0 \leq j \leq M + 1, 1 \leq i \leq n$. Set $t = 0$.
- 2: For $1 \leq i \leq n, 0 \leq j \leq M + 1$
- 3: simulate $X_i^* = \{X_{i,j\delta_i}^*\}_j$ from $\mathbb{W}^{(\Delta t_i, 0, 0)}$, and sample $U \sim \text{Un}(0, 1)$,
- 4: if

$$U < \frac{\pi_{HFA,M}(X_i^* | \theta^t, Y)}{\pi_{HFA,M}(\tilde{X}_i^t | \theta^t, Y)}$$

then set $\tilde{X}_i^{t+1} = X_i^*$, else set $\tilde{X}_i^{t+1} = \tilde{X}_i^t$.

- 5: Sample $\theta^* \sim q(\theta^t, \cdot)$ and $U \sim \text{Un}(0, 1)$,
- 6: if

$$U < \frac{\pi_{HFA,M}(\theta^*, \{\tilde{X}_i^{t+1}, 1 \leq i \leq n\} | Y) q(\theta^*, \theta^t)}{\pi_{HFA,M}(\theta, \{\tilde{X}_i^{t+1}, 1 \leq i \leq n\} | Y) q(\theta, \theta^*)}$$

then set $\theta^{t+1} = \theta^*$, else set $\theta^{t+1} = \theta^t$.

- 7: Set $t = t + 1$ and go to 2.
-

An alternative approximation to π_{PA} exists if conditions [SMOOTH], [GRAD] and [LBOUND] are satisfied. In particular, by using integration by parts we can transform the stochastic integrals in $G_{\Delta t_i}$ into time integrals, and rewrite (10) as

$$\begin{aligned} & \pi(\theta) \exp [H\{x_n(\theta_2); \theta\} - H\{x_0(\theta_2); \theta\} - l(\theta)(t_n - t_0)] \\ & \times \prod_{i=1}^n D(V_i; \theta_2) \mathcal{N}_{\Delta t_i} \{x_i(\theta_2) - x_{i-1}(\theta_2)\} \exp \left\{ - \int_0^{\Delta t_i} \phi(\tilde{X}_{i,s} + \mu_{i,s}(\theta); \theta) ds \right\}. \end{aligned} \quad (11)$$

The finite-dimensional approximation to (11), denoted by $\pi_{HFA,M}$, will typically be less biased than $\pi_{HFA,M}$, as illustrated in Section 6. The relative AMCMC algorithm follows along the lines of Algorithm 2 by replacing $\pi_{HFA,M}$ with $\pi_{HFA,M}$.

3.2 Exact Data Augmentation and MCMC

The main contribution of this section is to demonstrate that an exact rejection sampling algorithm for simulating diffusion bridges implies appropriate auxiliary variables which can be used to design EMCMC algorithms for exact inference for diffusions. We describe the data augmentation and the EMCMC algorithm which corresponds to the EA3 case, and later comment on the simplifications that arise when a more basic EA is applicable to the model of interest.

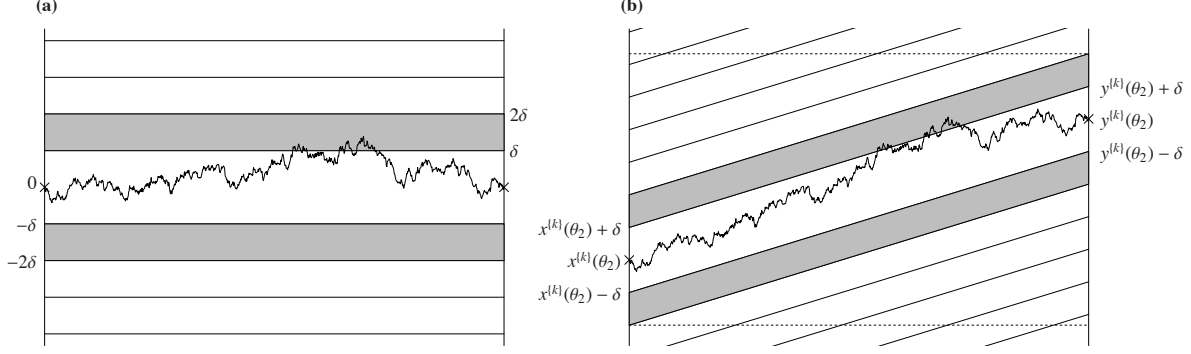


Figure 1: (a) The k th coordinate of a Brownian bridge path \tilde{X} starting and terminating at 0. In this example the event $C_k(2)$ has occurred. (b) The transformed process $\tilde{X}_s^{(k)} + (1 - s/t)x^{(k)}(\theta_2) + sy^{(k)}(\theta_2)/t$ starts at $x^{(k)}(\theta_2)$ and terminates at $y^{(k)}(\theta_2)$, with $x^{(k)}(\theta_2) < y^{(k)}(\theta_2)$. The dashed lines provide a lower and upper bound for the coordinate of the transformed path.

The augmentation is first identified for a pair of consecutive observations, $V_0 = v$ and $V_t = w$, and is then extended to an arbitrary number of observations by using the Markov property. This is achieved using the two main tools developed so far. First, the two path-parameter transformations for irreducible DA; recall that these are $V \rightarrow X$ and $X \rightarrow \tilde{X}$, which can be written in one step as

$$\tilde{X}_s = \eta(V_s; \theta_2) - \left(1 - \frac{s}{t}\right)\eta(v; \theta_2) - \frac{s}{t}\eta(w; \theta_2), \quad s \in [0, t]. \quad (12)$$

Second, the EA for simulating an X -bridge according to $\mathbb{Q}_\theta^{(t,x,y)}$, with $x = x(\theta_2)$ and $y = y(\theta_2)$ as in Section 2.2. Let $\tilde{\mathbb{Q}}_\theta^{(t)}$ denote the measure induced by the linearly transformed bridge \tilde{X} , when X is drawn from $\mathbb{Q}_\theta^{(t,x,y)}$. In passing, recall that when X is drawn from $\mathbb{W}_\theta^{(t,x,y)}$, \tilde{X} is distributed according to $\mathbb{W}^{(t,0,0)}$.

We first sketch an EA for simulating from $\tilde{\mathbb{Q}}_\theta^{(t)}$ using proposals from $\mathbb{W}^{(t,0,0)}$. This is a minor modification of the EA for X , since a proposed path $\tilde{X} \sim \mathbb{W}^{(t,0,0)}$ is accepted as a path from $\tilde{\mathbb{Q}}_\theta^{(t)}$ if and only if $\tilde{X}_s + (1 - s/t)x + sy/t$, $s \in [0, t]$, is accepted as a path from $\mathbb{Q}_\theta^{(t,x,y)}$. Let \tilde{L} denote the layer of a Brownian bridge that starts and terminates at 0, and denote by $\mathbb{M}^{(t)}$ the joint law of (\tilde{L}, \tilde{X}) ; hence if $(\tilde{L}, \tilde{X}) \sim \mathbb{M}^{(t)}$, marginally $\tilde{X} \sim \mathbb{W}^{(t,0,0)}$. The pair \tilde{L}, \tilde{X} imply realisations for the corresponding variables in the X -space. These are easy to obtain, since conditionally on a given \tilde{L} , $\{\tilde{X}_s^{(k)}, s \in [0, t]\}$ moves within $(-\tilde{L}^{(k)}\delta, \tilde{L}^{(k)}\delta)$, and thus

$$x^{(k)}(\theta_2) - \tilde{L}^{(k)}\delta < \tilde{X}_s^{(k)} + \left(1 - \frac{s}{t}\right)x^{(k)}(\theta_2) + \frac{s}{t}y^{(k)}(\theta_2) < y^{(k)}(\theta_2) + \tilde{L}^{(k)}\delta,$$

where we recall that $x^{(k)}(\theta_2) = x^{(k)}(\theta_2) \wedge y^{(k)}(\theta_2)$ and $y^{(k)}(\theta_2) = x^{(k)}(\theta_2) \vee y^{(k)}(\theta_2)$. The above construction and the derivation of the bounds are illustrated in Figure 1.

An EA which samples from $\tilde{\mathbb{Q}}_\theta^{(t)}$ follows easily. If $\Phi = \{\Psi, \Upsilon\}$ is a homogeneous Poisson process of intensity $r(\tilde{L}; \theta)$ on $[0, t] \times [0, 1]$, where Ψ is the projection of the points on $[0, t]$ and Υ the projection on $[0, 1]$, and

$$r(\tilde{L}; \theta) = \sup \left\{ \phi(u; \theta); u^{(k)} \in (x^{(k)}(\theta_2) - \tilde{L}^{(k)}\delta, y^{(k)}(\theta_2) + \tilde{L}^{(k)}\delta), k = 1, \dots, d \right\}, \quad (13)$$

then the algorithm accepts the proposed $(\tilde{L}, \tilde{X}, \Phi)$ according to

$$I(\tilde{L}, \tilde{X}, \Phi, v, w, \theta) := \prod_{j=1}^k \mathbb{I} \left[\frac{1}{r(\tilde{L}; \theta)} \phi \left\{ \tilde{X}_{\psi_j} + \left(1 - \frac{\psi_j}{t}\right)x(\theta_2) + \frac{\psi_j}{t}y(\theta_2); \theta \right\} < u_j \right], \quad (14)$$

which is the familiar indicator function (7), reformulated in terms of $(\tilde{L}, \tilde{X}, \Phi)$.

The EDA scheme for a pair of observations V_0, V_t is now defined by the random variables $(\tilde{L}, \tilde{X}, \Psi)$. Note that we can avoid augmenting Υ and still obtain a tractable density. The density of the auxiliary variables is derived in the following lemma proved in the Appendix.

Lemma 2. *Let $(\tilde{L}, \tilde{X}) \sim \mathbb{M}^{(t)}$ and Ψ be a homogeneous Poisson process of intensity $r(\tilde{L}; \theta)$ on $[0, t]$. If I is the acceptance indicator in (14), then the conditional density of $(\tilde{L}, \tilde{X}, \Psi)$ given $I = 1$, $\pi(\tilde{L}, \tilde{X}, \Psi | u, v, \theta)$, is*

$$\frac{r(\tilde{L}; \theta)^k}{a(x, y, t, \theta)} \exp \left\{ t \left[1 - r(\tilde{L}; \theta) \right] \right\} \prod_{j=1}^k \left[1 - \phi \left\{ \tilde{X}_{\psi_j} + \left(1 - \frac{\psi_j}{t}\right)x(\theta_2) + \frac{\psi_j}{t}y(\theta_2); \theta \right\} / r(\tilde{L}; \theta) \right], \quad (15)$$

with respect to the product measure $\mathbb{M}^{(t)} \times \mathbb{P}^{(t)}$, where $\mathbb{P}^{(t)}$ is the measure of a homogeneous Poisson process on $[0, t]$ with unit intensity, and $a(x, y, t, \theta)$ is the acceptance probability of the EA.

Extending the augmentation scheme to account for all observations is straightforward. Specifically, recall that $x_i(\theta_2) = \eta(V_{t_i}; \theta_2)$ and let $\tilde{L}_i, \Psi_i, \tilde{X}_i = \{\tilde{X}_{i,s}, s \in [0, \Delta t_i]\}$ denote the accepted elements of EA applied to the interval $[t_{i-1}, t_i]$, for $1 \leq i \leq n$. The Markov property of the diffusion process implies that the bridges conditionally on the observations are independent, and thus $\pi(\{\tilde{L}_i, \tilde{X}_i, \Psi_i, 1 \leq i \leq n\} | Y, \theta) = \prod_{i=1}^n \pi(\tilde{L}_i, \tilde{X}_i, \Psi_i | V_{t_{i-1}}, V_{t_i}, \theta)$.

We complete the development of EDA with the following theorem which specifies the joint density of data, auxiliary variables and parameters. This has a simple computable form and it admits the target posterior $\pi(\theta | Y)$ as a marginal with respect to the auxiliary variables and conditional with respect to the data. The key observation is that the joint density is only a function of the finite-dimensional $\{S(\tilde{X}_i), \tilde{L}_i, 1 \leq i \leq n\}$, which are delivered by the EA. Additionally, the intractable normalising constants have been cancelled out. The proof of the theorem is given in the Appendix.

Theorem 1. *The joint density of data Y , the p -dimensional parameters θ and auxiliary variables $\{\tilde{L}_i, \tilde{X}_i, \Psi_i, 1 \leq i \leq n\}$, is given below with respect to the θ -independent dominating measure $\text{Leb}^{n+p} \otimes_{i=1}^n (\mathbb{M}^{(\Delta t_i)} \times \mathbb{P}^{(\Delta t_i)})$:*

$$\begin{aligned} \pi(Y, \theta, \{S(\tilde{X}_i), \tilde{L}_i, 1 \leq i \leq n\}) &= \pi(\theta) \prod_{i=1}^n p_{\Delta t_i}(V_{t_{i-1}}, V_{t_i}; \theta) \prod_{i=1}^n \pi(\tilde{L}_i, \tilde{X}_i, \Psi_i | V_{t_{i-1}}, V_{t_i}, \theta) = \\ \pi(\theta) \exp &\left(H\{x_n(\theta_2); \theta\} - H\{x_0(\theta_2); \theta\} - [l(\theta) - 1](t_n - t_0) - \sum_{i=1}^n r(\tilde{L}_i; \theta) \Delta t_i \right) \\ &\times \prod_{i=1}^n \left\{ D(V_{t_i}; \theta_2) \mathcal{N}_{\Delta t_i}\{x_i(\theta_2) - x_{i-1}(\theta_2)\} r(\tilde{L}_i; \theta)^{\kappa_i} \prod_{j=1}^{\kappa_i} \left[1 - \phi\left\{ \tilde{X}_{i,\psi_{i,j}} + \mu_{i,\psi_{i,j}}(\theta); \theta \right\} / r(\tilde{L}_i; \theta) \right] \right\}, \end{aligned} \quad (16)$$

and it admits (2) as a marginal when $\{\tilde{L}_i, \tilde{X}_i, \Psi_i, 1 \leq i \leq n\}$ is integrated out and Y conditioned upon.

This density can be targeted by MCMC methods; actually at this stage we are only interested in the conditional density given Y , i.e., the joint posterior of parameters and auxiliary variables. We advocate a Gibbs sampler variant since conditionally on Y and θ the auxiliary variables $\{S(\tilde{X}_i), \tilde{L}_i, 1 \leq i \leq n\}$ are independent over i and can be generated using the EA. The conditional density of θ is computable and if it cannot be directly sampled, a Metropolis-Hastings step can be employed. Depending on the EA type used to construct the augmentation scheme, we distinguish between EMCMC1, 2 and 3. A typical implementation of EMCMC3 is given in Algorithm 3.

Algorithm 3 EMCMC3

- 1: Choose θ^0 and set $t = 0$.
- 2: For $1 \leq i \leq n$, set $I_i = 0$ and repeat the following until $I_i = 1$,
- 3: sample layer \tilde{L}_i of Brownian bridge path $\tilde{X}_i \sim \mathbb{W}^{(\Delta t_i, 0, 0)}$,
- 4: sample $\kappa_i \sim \text{Po}[r(\tilde{L}_i; \theta^i) \Delta t_i]$ and $\Phi_i = \{(\psi_{i,j}, u_{i,j})\}_j, 1 \leq j \leq \kappa_i$, uniformly on $[0, \Delta t_i] \times [0, 1]$,
- 5: conditionally on \tilde{L}_i , sample Brownian bridge $\tilde{X}_{i,\psi_{i,j}}$,
- 6: set $I_i = I(\tilde{L}_i, \tilde{X}_i, \Phi_i, V_{t_{i-1}}, V_{t_i}, \theta^i)$ as in (14); if $I_i = 1$, then set $\tilde{L}_i^{t+1} = \tilde{L}_i, \Psi_i^{t+1} = \Psi_i, \tilde{X}_i^{t+1} = \tilde{X}_i$.
- 7: Sample $\theta^* \sim q(\theta^*, \cdot)$ and $U \sim \text{Un}[0, 1]$,
- 8: if

$$U < \frac{\pi(Y, \theta^*, \{S(\tilde{X}_i^{t+1}), \tilde{L}_i^{t+1}, 1 \leq i \leq n\}) q(\theta^*, \theta^i)}{\pi(Y, \theta^i, \{S(\tilde{X}_i^{t+1}), \tilde{L}_i^{t+1}, 1 \leq i \leq n\}) q(\theta^i, \theta^*)}$$

then set $\theta^{t+1} = \theta^*$, else set $\theta^{t+1} = \theta^i$.

- 9: Set $t = t + 1$ and go to 2.
-

Special cases: EMCMC1 and EMCMC2

Certain simplifications are feasible when a simpler EA can be applied to the process of interest. The EDA based on EA1 requires less imputation than that of EA3, since exact simulation from $\tilde{\mathbb{Q}}_\theta^{(t)}$ no longer requires the variable \tilde{L} . Thus, the augmentation scheme involves only $\{\tilde{X}_i, \Psi_i, 1 \leq i \leq n\}$. The joint density analogous to (16) can be easily obtained and amounts to simply replacing $r(\tilde{L}; \theta)$ by $r(\theta)$. The conditional of θ trivially follows; originally it was given in Theorem 3 of Beskos *et al.* (2006b).

A DA scheme can be built using the auxiliary variables used in EA2. We do not present the details of this, since it is not a direct modification of the general scheme, as it is the case with EA1, but instead involves a different construction of bridges. Details can be found in Chapter 7 of Sermaidis (2010).

3.3 Interpreting EDA in terms of PA

The following result provides the connection between EDA and PA. It effectively shows that the PA scheme is a *collapsed* version of EDA, i.e., when we integrate out a subset of the latent variables we obtain the distribution which is targeted by PA.

Theorem 2. *Let $\pi(\theta, \{\tilde{X}_i, 1 \leq i \leq n\} | Y)$ be the density obtained from (16) by conditioning on Y , and marginalising with respect to $\{\tilde{L}_i, \Psi_i, 1 \leq i \leq n\}$, and $\pi_{PA}(\theta, \{\tilde{X}_i, 1 \leq i \leq n\} | Y)$ the density targeted by the irreducible path imputation algorithm defined in (11). Then, the two densities are equal a.s.*

The result is insightful towards the comparison of the computational efficiency of EMCMC to that of AMCMC, as it suggests that the mixing time of the former might generally be larger due to its higher degree of augmentation; a price one has to pay in order to eliminate the discretisation error. Nonetheless, Section 4 shows a variety of ways with which one can increase the performance of EMCMC and achieve good mixing rates. A numerical comparison of EMCMC and AMCMC is investigated in Section 6 in concrete examples.

3.4 Qualitative characteristics of EMCMC and AMCMC

We can make some qualitative remarks about the efficiency of the MCMC schemes based on HFA and EDA. These remarks are based on general properties of DA methods, as for example discussed in Papaspiliopoulos *et al.* (2007); Yu and Meng (2011) but also the particular structure of the models at hand. These qualitative statements are backed up by numerical evidence in Section 6, and they motivate the three approaches we propose in the next section to boost the algorithmic efficiency.

To fix terminology, we will identify DA with a Gibbs-type sampler which updates auxiliary variables and parameters according to their conditional distributions. In general, DA works well when the fraction of missing information is not too large relative to the observed one, i.e., when the augmented dataset is not considerably more informative than the observed regarding the parameters.

In that respect, both EMCMC and AMCMC become more efficient when the time increment t between a pair of observations decreases. In the context of AMCMC, this is due to the fact that as t decreases, the latent bridges increasingly look like Brownian bridges, hence they do not carry information about the drift over and above the one contained in the observed endpoints and the augmented information converges to the observed one. The efficiency of EMCMC improves as the Poisson rate, $r(\tilde{L}; \theta)t$, decreases. In the limiting case when the rate is 0, the missing data and parameters are independent, since the skeleton is empty, and EMCMC achieves maximal efficiency.

On the other hand, as t increases, the augmented paths contain increasing amount of information about the parameters relative to the information content in the observed data, therefore any algorithm which iteratively simulates from the conditional distributions of parameters and missing data will degrade in this sparse-data limit, even if the conditional distributions can be simulated exactly and efficiently. EDA has a further weakness over HFA because the number of points in a EA skeleton is informative about $r(\tilde{L}; \theta)$, and thus about θ . However, we can deal with this dependence, which increases with t , by means of a reparametrisation that is described in Section 4.

Additionally, both EMCMC and AMCMC suffer at the step of updating missing data given parameters as t increases. In EMCMC the acceptance rate of EA decays to 0 exponentially with t , implying that the algorithm can spend a large amount of time by simulating proposed paths until acceptance. Similarly in AMCMC the acceptance rate of the independent Metropolis-Hastings step decays to 0 at the same rate (expression (11)), suggesting that the algorithm will be rarely updating the last accepted path, thus leading to a slower exploration of the state space. However, this problem can be improved by resorting to alternative update schemes for the imputation step. One option for AMCMC is to use local algorithms, see for example Cotter *et al.* (2012). An other alternative, which can apply to both algorithms, is to update the paths in smaller time segments by imputing $O(t)$ additional points between pairs of observations. This can turn the complexity from exponential to linear in t .

4 Boosting EMCMC efficiency

4.1 Noncentred reparametrisation

One general approach for improving efficiency of data augmentation in hierarchical models and auxiliary variable models is to adopt a reparametrisation. Indeed, we have already done so in HFA and in EDA by transforming $V \rightarrow \tilde{X}$ as in (12). Following Papaspiliopoulos *et al.* (2007), for a generic random variable E and data Y , a reparametrisation of an augmentation scheme E is defined by any random pair (\tilde{E}, θ) together with a function h such that $E = h(\tilde{E}, \theta, Y)$, where h need not be 1-1. A reparametrisation is called *noncentred* when the distribution of \tilde{E} is independent of θ . Intuitively, in cases where Y is not strongly informative about E , a noncentred scheme can perform well due to the prior independence of \tilde{E} and θ . We will attempt to reduce the dependence between the Poisson process and θ by resorting to a noncentred reparametrisation.

Noncentred reparametrisations for decoupling the dependence between Poisson processes and their intensity were originally proposed in Roberts *et al.* (2004). Applying their idea in this context, for two observations $V_0 = v$, $V_t = w$, if Ψ is a Poisson process of rate $r(\tilde{L}; \theta)$ on $[0, t]$ and $\tilde{\Psi}$ is a Poisson process of unit intensity on $[0, t] \times [0, \infty)$ with point-coordinates $\{(\tilde{\psi}_j, \tilde{\xi}_j)\}$ then

$$\Psi = h(\tilde{\Psi}, \tilde{L}, \theta) = \{\tilde{\psi}_j; \tilde{\xi}_j < r(\tilde{L}; \theta)\}. \quad (17)$$

Although $\tilde{\Psi}$ includes an infinite number of points, Ψ only depends on those for which the second coordinate is below $r(\tilde{L}; \theta)$. Accounting for all the observations, the noncentred reparametrisation is $(\theta, \{\tilde{L}_i, \tilde{X}_i, \Psi_i, 1 \leq i \leq n\}) \rightarrow (\theta, \{\tilde{L}_i, \tilde{X}_i, \tilde{\Psi}_i, 1 \leq i \leq n\})$, where $\Psi_i = h(\tilde{\Psi}_i, \tilde{L}_i, \theta)$. The theorem below derives the conditional density of θ given the latent variables and observations, and is proved in the Appendix.

Theorem 3. *The conditional density of θ given the auxiliary variables $\{\tilde{L}_i, \tilde{X}_i, \tilde{\Psi}_i, 1 \leq i \leq n\}$, $\pi_{nc}(\theta \mid \{\tilde{L}_i, \tilde{X}_i, \tilde{\Psi}_i, 1 \leq i \leq n\}, Y)$ is proportional to*

$$\begin{aligned} & \pi(\theta) \exp[H\{x_n(\theta); \theta\} - H\{x_0(\theta); \theta\} - l(\theta)(t_n - t_0)] \times \prod_{i=1}^n D(V_i; \theta_2) \mathcal{N}_{\Delta t_i} \{x_i(\theta_2) - x_{i-1}(\theta_2)\} \\ & \times \prod_{i=1}^n \prod_{j=1}^{\infty} [1 - \mathbb{I}[\tilde{\xi}_{i,j} < r(\tilde{L}_i; \theta)]] \phi\{\tilde{X}_{i,\tilde{\psi}_{i,j}} + \mu_{i,\tilde{\psi}_{i,j}}(\theta); \theta\} / r(\tilde{L}_i; \theta). \end{aligned} \quad (18)$$

Notice that evaluation of (18) for any value of θ requires only finite computation and therefore discretisations are avoided.

The MCMC algorithm based on this reparametrisation is practically a small modification of that based on the original scheme (Algorithm 3). First, we wish to draw from the distribution of $\{\tilde{L}_i, \tilde{X}_i, \tilde{\Psi}_i, 1 \leq i \leq n\}$ given Y and the current parameter value, say θ' . It is clear from (17) that $\tilde{\Psi}_i$ need only be revealed on $[0, \Delta t_i] \times [0, r(\tilde{L}_i; \theta')]$, which involves a finite number of points $\tilde{\kappa}_i \sim \text{Po}[r(\tilde{L}_i; \theta') \Delta t_i]$ and is sufficient for the implementation of the EA. The i th output consists of $\tilde{\Psi}_i$ partially observed at $\{(\tilde{\psi}_{i,j}, \tilde{\xi}_{i,j}), 1 \leq j \leq \tilde{\kappa}_i\}$ and the pair $\{\tilde{L}_i, \tilde{X}_i\}$ discretely observed at times $\tilde{\psi}_{i,j}$. However, sampling from the distribution of θ given the latent variables and Y is more tricky. Specifically, if the proposed value, say θ^* , is such that $r(\tilde{L}_i; \theta^*) > r(\tilde{L}_i; \theta')$, then evaluating (18) at θ^* requires revealing \tilde{X}_i at additional time points $\{\tilde{\psi}_{i,j}; r(\tilde{L}_i; \theta') < \tilde{\xi}_{i,j} < r(\tilde{L}_i; \theta^*)\}$, which have not been revealed in the EA output. Notice that this does not occur when $r(\tilde{L}_i; \theta^*) < r(\tilde{L}_i; \theta')$.

We propose two ways to overcome this. The first is based on prospectively revealing $\tilde{\Psi}_i$ at any additional time instances by simply simulating extra $\tilde{\kappa}_i^*$ uniform random variates on $[0, \Delta t_i] \times [r(\tilde{L}_i; \theta'), r(\tilde{L}_i; \theta^*)]$, where $\tilde{\kappa}_i^* \sim \text{Po}\{[r(\tilde{L}_i; \theta^*) - r(\tilde{L}_i; \theta')] \Delta t_i\}$. The path \tilde{X}_i is then filled in at the additional points by Brownian bridge interpolations. The second is closest in spirit to the retrospective nature of the EA. In particular, θ^* can be simulated prior to the application of the EA and therefore $r(\tilde{L}_i; \theta')$ and $r(\tilde{L}_i; \theta^*)$ are known before the simulation of the latent path. Consequently, we can simulate $\tilde{\Psi}_i$ directly on $[0, \Delta t_i] \times [0, r(\tilde{L}_i; \theta^*)]$ and reveal \tilde{X}_i at all required time instances during the implementation of the EA.

In this paper, we adopt the retrospective approach because it can be applied in a similar fashion to all three EAs. The prospective approach is simple in the EA1 case due to simple Brownian bridge interpolations, but becomes more involved in the EA2 and EA3 cases.

Again, a simplification can be achieved when EA1 is applicable, where the transformation in that case becomes $\{\tilde{X}_i, \Psi_i, 1 \leq i \leq n\} \rightarrow \{\tilde{X}_i, \tilde{\Psi}_i, 1 \leq i \leq n\}$, where $\Psi_i = h(\tilde{\Psi}_i, \theta) = \{\tilde{\psi}_{i,j}; \tilde{\xi}_{i,j} < r(\theta)\}$. The full conditional density of θ is essentially given by expression (18) replacing $r(\tilde{L}_i; \theta)$ with $r(\theta)$. Noncentred reparametrisations for EMCMC2 also exist; details can be found in Section 7.4.3 of Sermaidis (2010).

Algorithm 4 Noncentred EMCMC3

- 1: Choose θ^0 and set $t = 0$.
 - 2: Sample $\theta^* \sim q(\theta^*, \cdot)$ and $U \sim \text{Un}[0, 1]$.
 - 3: For $1 \leq i \leq n$, set $I_i = 0$ and repeat the following until $I_i = 1$,
 - 4: sample layer \bar{L}_i of Brownian bridge path $\bar{X}_i \sim \mathbb{W}^{(\Delta t_i, 0, 0)}$, and set $r_{\max} = r(\bar{L}_i; \theta^*) \vee r(\bar{L}_i; \theta^t)$,
 - 5: sample $\bar{\kappa}_i \sim \text{Po}[r_{\max} \Delta t_i]$ and $\{(\bar{\psi}_{i,j}, \bar{u}_{i,j}, \bar{\xi}_{i,j})\}_j$, $1 \leq j \leq \bar{\kappa}_i$ uniformly on $[0, \Delta t_i] \times [0, 1] \times [0, r_{\max}]$,
 set $\bar{\Psi}_i = \{(\bar{\psi}_{i,j}, \bar{\xi}_{i,j})\}_j$, $\bar{\Upsilon}_i = \{(\bar{u}_{i,j}, \bar{\xi}_{i,j})\}_j$,
 - 6: set $\Phi_i = \{\Psi_i, \Upsilon_i\}$, where $\Psi_i = h(\bar{\Psi}_i, \bar{L}_i, \theta^t)$ and $\Upsilon_i = h(\bar{\Upsilon}_i, \bar{L}_i, \theta^t)$ as in (17),
 - 7: conditionally on \bar{L}_i , sample Brownian bridge $\tilde{X}_{i, \bar{\psi}_{i,j}}$,
 - 8: set $I_i = I(\bar{L}_i, \tilde{X}_i, \Phi_i, V_{i-1}, V_i, \theta^t)$ as in (14); if $I_i = 1$, then set $\bar{L}_i^{t+1} = \bar{L}_i$, $\bar{\Psi}_i^{t+1} = \bar{\Psi}_i$ and $\bar{X}_i^{t+1} = \tilde{X}_i$.
 - 9: IF

$$U < \frac{\pi_{nc}(\theta^* \mid \{\bar{L}_i^{t+1}, \bar{X}_i^{t+1}, \bar{\Psi}_i^{t+1}, 1 \leq i \leq n\}, Y) q(\theta^*, \theta^t)}{\pi_{nc}(\theta^t \mid \{\bar{L}_i^{t+1}, \bar{X}_i^{t+1}, \bar{\Psi}_i^{t+1}, 1 \leq i \leq n\}, Y) q(\theta^t, \theta^*)}$$
 then set $\theta^{t+1} = \theta^*$, else set $\theta^{t+1} = \theta^t$.
 - 10: Set $t = t + 1$ and go to 2.
-

An interweaving strategy

When a noncentred transformation is available, it is not necessary to choose between that and the original parametrisation. Yu and Meng (2011) propose instead to *interweave* the two, by creating a single algorithm which mixes steps of both algorithms. This requires practically no extra coding work, but as it is demonstrated in the article, in certain cases the interweaved algorithm outperforms its parent algorithms even when taking the added computational cost into account.

In the EMCMC context, if θ^t and $\{\bar{L}_i^t, \bar{X}_i^t, \bar{\Psi}_i^t, 1 \leq i \leq n\}$ denote the current state of the chain, then the algorithm can effectively be described in four steps. The first two are identical to sampling from the noncentred algorithm, i.e., the latent variables are updated to $\{\bar{L}_i^{t+1}, \bar{X}_i^{t+1}, \bar{\Psi}_i^{t+1}, 1 \leq i \leq n\}$ using the EA and then the parameter is updated by drawing $\theta^{t+\frac{1}{2}}$ conditionally on these latent data and observations. Subsequently, the latent data are transformed to their original parametrisation using $\Psi_i = h(\bar{\Psi}_i^{t+1}, \theta^{t+\frac{1}{2}})$; notice that this step does not impose any computational difficulties, since it merely involves a deterministic transformation. Finally, the parameter is re-drawn under the original parametrisation, $\theta^{t+1} \sim \pi(\cdot \mid \{S(\bar{X}_i^{t+1}), \bar{L}_i^{t+1}, 1 \leq i \leq n\}, Y)$.

4.2 Auxiliary Poisson sampling

An alternative way to improve the mixing time by increasing the computational cost is to exploit the connection between EDA and PA. For simplicity we present the result for EMCMC1, and then discuss extensions to EMCMC3.

Note that if $r(\theta)$ is the Poisson rate, then it is valid to apply EA1 with any Poisson sampling rate $R(\theta) > r(\theta)$; the acceptance probability is invariant to that choice. Increasing the value of $R(\theta)$ leads to an increase in the computational cost since the number of points at which the path is evaluated gets larger. Thus, in terms of computing time it is optimal to implement the algorithm with the smallest possible $R(\theta)$. This is not the case though for EMCMC1 that iterates between imputation and estimation. As it turns out, the dependence between missing data and parameters decreases with $R(\theta)$ and optimal implementation in terms of execution time and Monte Carlo error can be achieved for $R(\theta) > r(\theta)$. Thus, we consider data augmentation where the auxiliary variables are chosen according to the output of EA1, as described in Section 3.2, but where the Poisson rate is $R(\theta)$.

A theoretical result (not included here) goes along the following lines. Let $R(\theta) = r(\theta) + \lambda$, where $\lambda \geq 0$ is a user-specified constant independent of θ . Then, the joint law of $(\theta, \{\bar{X}_i, \Psi_i, 1 \leq i \leq n\})$ after a step of the EMCMC1 algorithm with parameter λ , converges to the law of one step of PA when $\lambda \rightarrow \infty$. An argument for proving this is based on properties of series expansions for exponential functionals, as discussed for example in Papaspiliopoulos (2011).

Hence, in a sense the auxiliary variables Ψ_i are effectively integrated out by increasing computation and an improvement in the convergence of EMCMC is expected as λ increases. Similar arguments are valid for the noncentred algorithm since it merely involves a reparametrisation of Ψ_i . A similar property is enjoyed by a

generic EMCMC3. In fact, the limiting algorithm as λ increases is a PA which also imputes the layer, although the latter is immaterial in the limit since it does not contain additional information about the parameters.

5 Diffusion observed with error

The methodology described so far can be easily extended to account for cases where the diffusion is not directly observed. We assume that the observations inform only indirectly about the value of the process (1) at discrete times $t_i, i = 0, 1, \dots, n$, according to the following observation equation

$$Z_i \sim q(\cdot | V_i, \tau),$$

where $Z := \{Z_{t_0}, Z_{t_1}, \dots, Z_{t_n}\}$ are conditionally independent given $Y = \{V_{t_0}, V_{t_1}, \dots, V_{t_n}\}$, and q is a known density function parametrised by an unknown parameter τ .

The EDA described in Section 3.2 is not appropriate anymore since the end points V_{t_i} are not directly observed. On the other hand, Theorem 1 can be used to design an augmentation scheme where apart from the auxiliary variables involved in EDA, the latent points Y are imputed as well. A direct application of Bayes' theorem yields that the joint posterior $\pi(Y, \theta, \tau, \{S(\tilde{X}_i), \tilde{L}_i, 1 \leq i \leq n\} | Z)$ is proportional to

$$\pi(\theta, \tau) \pi(Y, \{S(\tilde{X}_i), \tilde{L}_i, 1 \leq i \leq n\} | \theta) \pi(V_{t_0} | \theta) \prod_{i=0}^n q(Y_{t_i} | V_{t_i}, \tau), \quad (19)$$

where the second term is obtained directly from Theorem 1, $\pi(\cdot | \theta)$ is a prior density for the initial point of the diffusion process, and $\pi(\theta, \tau)$ is a prior density of the parameters. As before, a direct simplification is available when the EA1 can be applied to simulate from \tilde{X}_i .

A simple scheme for simulating from (19) is by a component-wise updating algorithm. $\{S(\tilde{X}_i), \tilde{L}_i, 1 \leq i \leq n\}$ and θ are simulated conditionally on Y and τ , using any of the EMCMC schemes we have proposed, and subsequently Y and τ conditionally on $\{S(\tilde{X}_i), \tilde{L}_i, 1 \leq i \leq n\}$ and θ according to the conditional derived from (19). When $\pi(\theta, \tau) = \pi(\theta)\pi(\tau)$, τ is conditionally independent from $\{S(\tilde{X}_i), \tilde{L}_i, 1 \leq i \leq n\}$ and θ given Y , and may have a conditional density which can be easily simulated. Simulation of the latent points Y can be done with various ways. The simplest is to update them one-at-a-time according to their conditional density, an approach often called single-site updating. Such schemes for time series are known to be in general problematic, especially when the latent process exhibits high persistence and the observations are not very informative about the latent points, see for instance Pitt and Shephard (1999); Papaspiliopoulos *et al.* (2007). In the example we consider in the next section we adopt this simplistic approach since it works quite well. In applications where the process Y exhibits very high persistence a joint update of the endpoints can be done using a Metropolis-adjusted Langevin algorithm, or more general version of such algorithms, as discussed for example in Girolami and Calderhead (2011). Other possibility is to resort to an overlapping block scheme, as for example in Pitt and Shephard (1999); Golightly and Wilkinson (2008).

6 Numerical investigation of MCMC algorithms

We investigate the numerical performance of the several algorithms we have presented on some standard examples. One is a diffusion which belongs in the Pearson family, see for example Forman and Sørensen (2008), and it is an example of a process that can be simulated using the EA1. The second is a univariate double well potential model, typical of models that are used to describe processes with metastable behaviour, see for example Metzner *et al.* (2006). This is an example of a process that can only be simulated using the EA3. The third is a double well potential model in two dimensions.

We compare several algorithms. The plain-vanilla EMCMC together with its elaborations: using noncentred reparametrisation, the interweaved strategy and auxiliary Poisson sampling. Additionally, we compare against both versions of AMCMC described in Section 3.1. The first is the plain one as introduced in Roberts and Stramer (2001) and is based on (10); we refer to this basic version simply as AMCMC. The second eliminates the stochastic integrals by integration by parts as in (11) and will be denoted in the text by ‘‘AMCMC (int-by-parts)’’. This is expected to enhance algorithmic performance, hence we evaluate the effect of this approach. For both versions, we add the suffix $-M$ to indicate the number of imputed points. In general, the quality of each MCMC output is assessed with an adjusted effective sample size, defined as the ratio of the effective sample size (ESS) to the

computational (CPU) time to run each algorithm. The adjusted ESS is essentially the number of independent draws per second generated by the Markov chain. The ESS is calculated with the R (R Development Core Team, 2010) package coda (Plummer *et al.*, 2010).

The existence of EMCMC allows us to have a realistic evaluation of the performance of the biased approaches. We carry out bootstrap Kolmogorov-Smirnov tests (Sekhon, 2007) for checking whether the marginal posterior distributions on the parameters obtained by different levels of imputation are significantly different from the exact samples. These evaluations are based on thinning the original Markov chain output so that to obtain practically independent draws from the corresponding distributions.

6.1 A Pearson diffusion

Consider the univariate diffusion process specified by

$$dV_s = -\rho(V_s - \mu)ds + \sigma \sqrt{1 + V_s^2}dW_s,$$

where $\sigma > 0$, $\rho > 0$ is a mean reverting parameter and $\mu \in \mathbb{R}$ is the stationary mean. The parameter vectors are identified as $\theta_1 = (\rho, \mu)^T$ and $\theta_2 = \sigma$. This model belongs to a rich class of diffusion processes, known as the Pearson diffusions, and admits a stationary distribution with skewness and heavy tails that decay at the same rate as those of a t -distribution.

The unit volatility process is obtained as $X_s := \sinh(V_s)/\sigma$, with drift given by

$$a(x; \theta) = -\left(\frac{\rho}{\sigma} + \frac{\sigma}{2}\right) \tanh(\sigma x) + \frac{\rho\mu}{\sigma \cosh(\sigma x)}.$$

This is a process which belongs in the EA1 class, and exact inference can be performed with

$$l(\theta) = -\frac{1}{2} \left(\rho + \frac{\sigma^2}{2} + \frac{\rho\mu}{2} \right), \quad r(\theta) = \frac{1}{8} \left\{ \rho(6\mu + 8) + 3\sigma^2 + \frac{4\rho^2}{\sigma^2}(\mu^2 + \mu + 1) \right\}. \quad (20)$$

We test the methods on a simulated data set from this process, based on $n = 1000$ (excluding the initial point) equidistant points with $\Delta t_i = 1$, $V_0 = 1$ and parameter values $(\rho, \mu, \sigma) = (1/2, 1, 1/2)$ (Figure 2a). We have used improper prior densities for the parameters, $\pi(\rho) \propto 1$, $\pi(\mu) \propto 1$, and $\pi(\sigma) \propto 1/\sigma$. For all algorithms, sampling from the conditional density of the parameters was performed using a block Metropolis-Hastings step. The chains were run for 10^5 iterations.

Figure 3 shows the autocorrelation plots along with posterior density estimates derived from the Markov chains. Starting from EMCMC1 under the original parametrisation, notice that for $\lambda = 0$ the chain exhibits strong serial dependence even at large lags, particularly for ρ . This is due to strong a priori dependence between the parameter and the number of Poisson points, as shown in (20), which remains significant in the posterior distribution. In particular, the posterior correlation between $\sum_i^j \kappa_i$ and ρ was estimated equal to 0.93, thus suggesting that non-centring the Poisson process can result in better mixing rates, as Figure 3d confirms. To improve the performance of the exact methods, we consider various values for $\lambda = \{2, 5, 10\}$, thus revealing the path at additionally 2, 5 and 10 points between consecutive observations respectively. The increase in performance is reflected in the autocorrelation function, which now decays to 0 more quickly. As expected, the chains of the AMCMC algorithms mix more rapidly than the exact ones due to the less amount of augmentation. The posterior distributions estimated from the HFA algorithms provide evidence of bias even for $M = 30$ (Figures 3g to 3i).

Table 1 presents posterior summary statistics. Notice that even for $M = 30$, the AMCMC algorithm fails to pass the Kolmogorov-Smirnov tests at a 5% significance level, whereas less amount of imputation ($M = 10$) combined with the integration by parts yields less biased approximations, clearly illustrating the importance of eliminating the stochastic integrals as in (11). In terms of computational performance, the interweaved algorithm with $\lambda = 2$ outperforms the rest and exhibits adjusted ESS for μ and σ which are respectively 36% and 23% larger than that of the sufficiently accurate AMCMC methods.

The algorithms were also run using proper priors, an exponential distribution for ρ , a Gaussian for μ and an inverse Gamma for σ^2 , yielding no significant differences from the results presented above.

6.2 A double well potential model

We consider the solution process to

$$dV_s = -\rho V_s (V_s^2 - \mu) ds + \sigma dW_s,$$




Figure 2: Simulated datasets from (a) the Pearson model, (b) the double well potential model, (c) the double well observed with error, and (d),(e) the two-dimensional double well model.

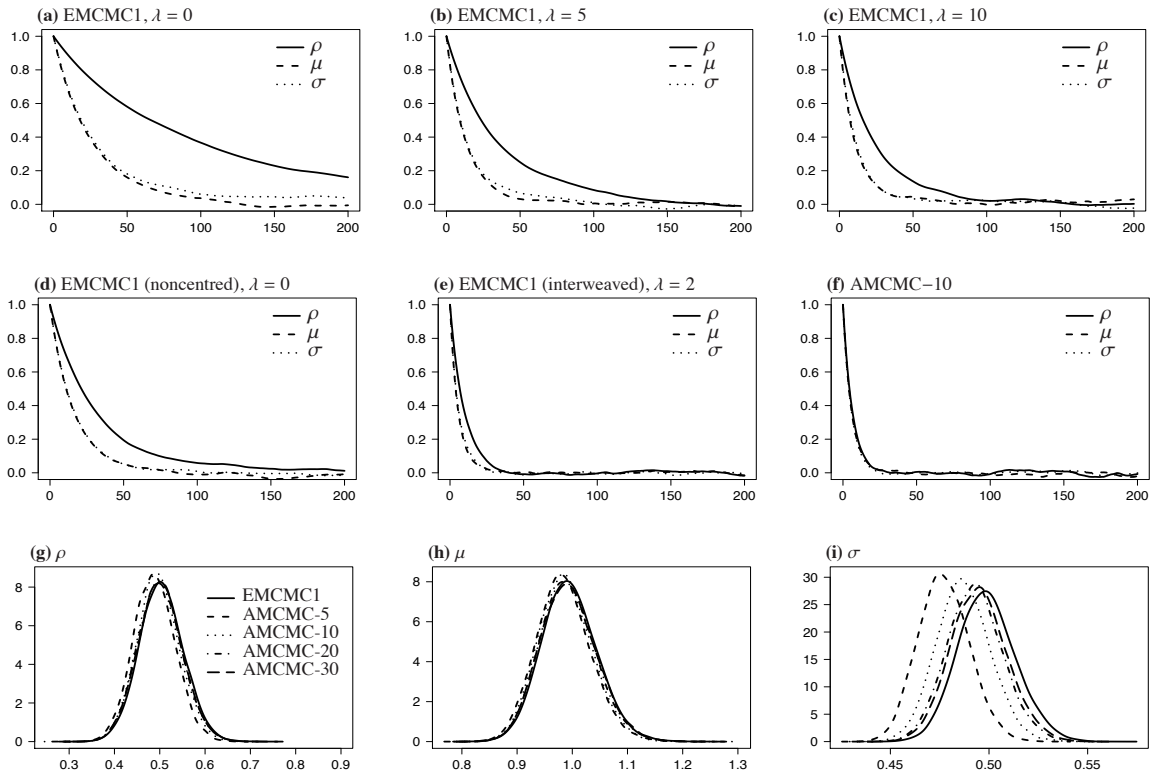


Figure 3: The Pearson diffusion model with $n = 1000$ simulated data points. True values are $(\rho, \mu, \sigma) = (1/2, 1, 1/2)$. Autocorrelations are reported after a burn in of 5000 iterations. Posterior density estimates using EMCMC1 (interweaved) and AMCMC algorithms for (g) ρ , (h) μ and (i) σ .

Method	Par.	λ	M	Mean	SD	ESS _{adj}	ESS	KS	Time	Correlation matrix		
EMCMC1 (noncentred)	ρ	0	2.316	0.505	0.048	6.599	16.789		2.544	1.000	-0.447	0.539
	μ			0.995	0.050	12.432	31.629				1.000	0.008
	σ			0.499	0.015	12.117	30.829					1.000
EMCMC1 (interweaved)	ρ	0	2.314	0.506	0.049	6.686	22.170		3.316	1.000	-0.438	0.540
	μ			0.994	0.050	15.035	49.859				1.000	0.021
	σ			0.499	0.015	14.408	47.779					1.000
EMCMC1 (interweaved)	ρ	2	4.346	0.504	0.048	11.094	53.153		4.791	1.000	-0.416	0.537
	μ			0.995	0.050	18.951	90.792				1.000	0.026
	σ			0.499	0.015	17.022	81.554					1.000
AMCMC-30	ρ		30.000	0.501	0.048	4.916	80.518	0.005	16.380	1.000	-0.439	0.519
	μ			0.993	0.050	5.160	84.516	0.017			1.000	0.021
	σ			0.495	0.014	5.329	87.293	< 0.001				1.000
AMCMC-10 (int-by-parts)	ρ		10.000	0.504	0.049	12.898	81.865	0.286	6.347	1.000	-0.422	0.537
	μ			0.995	0.050	13.878	88.087	0.749			1.000	0.020
	σ			0.499	0.015	13.793	87.545	0.583				1.000

Table 1: The Pearson diffusion model with $n = 1000$ simulated data points. True values are $(\rho, \mu, \sigma) = (1/2, 1, 1/2)$. Statistics are reported after a burn-in period of 5000 iterations. The M column shows the average number of imputed points between consecutive observations. The SD column shows the standard deviation. The ESS column shows the effective sample size per 1000 iterations. The Time column shows the time in seconds required for 1000 iterations of each chain. The adjusted effective sample size is shown in column ESS_{adj}. The KS column shows the p -value for the Kolmogorov-Smirnov test with null hypothesis that draws from the exact and approximate marginals come from the same distribution.

where $\rho > 0, \mu > 0, \sigma > 0$. The parameter vectors are identified as $\theta_1 = (\rho, \mu)^T$ and $\theta_2 = \sigma$. The process is known as the double well potential process (denoted by DWELL hereafter). We simulated $n = 1000$ (excluding the initial point) equidistant observations with $\Delta t_i = 1$ for parameter setting $(\rho, \mu, \sigma) = (0.1, 2, 1/2)$ and $V_0 \sim N(2, 1/4)$ (Figure 2b). Reduction to a unit volatility process is easily achieved using $X_s := V_s/\sigma$, which solves the SDE

$$dX_s = -\rho X_s (\sigma^2 X_s^2 - \mu) ds + dW_s.$$

Simple calculations reveal that the function $f(u; \theta) := \|\alpha(u; \theta)\|^2 + \Delta_x H(u; \theta)/2$ is given by

$$f(u; \theta) = \frac{\rho}{2} \{ \rho \sigma^4 u^6 - 2\rho \mu \sigma^2 u^4 + (\rho \mu^2 - 3\sigma^2) u^2 + \mu \},$$

and that the algorithms are applicable with $l(\theta) = f(u_i; \theta)$, where

$$u_i^2 = \frac{2\rho\mu + \sqrt{\rho(\rho\mu^2 + 9\sigma^2)}}{3\rho\sigma^2}.$$

Finally, for a given realisation of the layer \tilde{L} the upper bound (13) is easy found by noticing that

$$f(u; \theta) \leq \frac{\rho}{2} (\rho \sigma^4 u^6 + \rho \mu^2 u^2 + \mu) =: g(u; \theta),$$

which is convex and has a minimum at $u = 0$, implying that

$$r(\tilde{L}; \theta) = \left[g\{x(\theta_2) - \tilde{L}\delta; \theta\} \vee g\{y(\theta_2) + \tilde{L}\delta; \theta\} \right] - l(\theta). \quad (21)$$

We assign improper prior densities to the parameters with $\pi(\rho) \propto 1$, $\pi(\mu) \propto 1$, and $\pi(\sigma) \propto 1/\sigma$, and run the chains for 10^5 iterations. The performance of the algorithms and posterior density estimates are shown in Figure 4. Notice that the performance of EMCMC3 under the original parametrisation is very poor, due to strong posterior correlation between Poisson points and parameters; a fact attributed to the sensitivity of $r(\tilde{L}; \theta)$ to the parameters (see expression (21)). On the other hand, the noncentred algorithm exhibits a much stronger performance with low serial correlation for each parameter even after 50 lags.

Posterior summaries from the output of the chains and computational performance are gathered in Table 2. In contrast to the EA1 example presented earlier, the interweaved strategy, after accounting for the additional computational cost, does not offer any significant improvement over the noncentred algorithm. Finally, we found that AMCMC with $M = 40$ paired with integration by parts provides a reasonable approximation to the posterior marginal densities and exhibits slightly larger adjusted ESS than the noncentred algorithm with $\lambda = 2$. The results were robust in changes in parameter prior distributions.

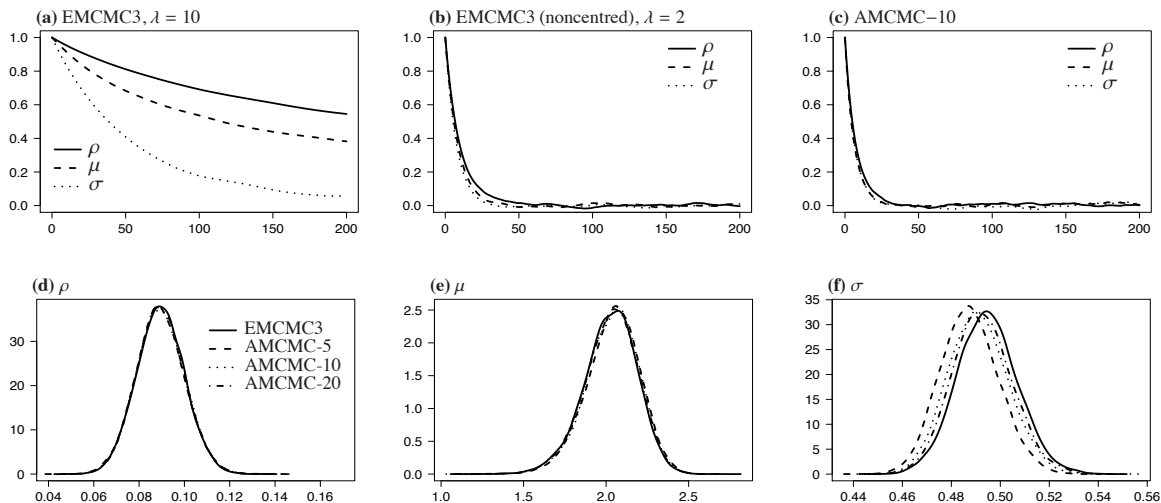


Figure 4: The DWELL diffusion model with $n = 1000$ simulated data points. True values are $(\rho, \mu, \sigma) = (0.1, 2, 1/2)$. Autocorrelations are reported after a burn in of 5000 iterations. Posterior density estimates using EMCMC3 (noncentred) and AMCMC algorithms for (d) ρ , (e) μ and (f) σ .

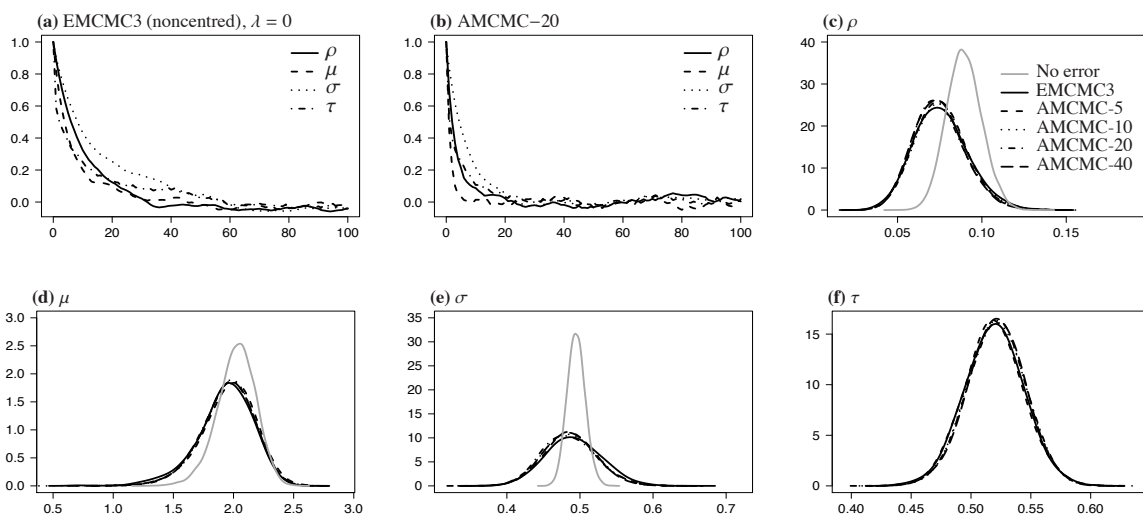


Figure 5: The DWELL diffusion model observed with error with $n = 1000$ simulated data points. True values are $(\rho, \mu, \sigma, \tau) = (1/2, 1, 1/2, 1/2)$. The outputs of the MCMC chains are subsampled every 10 iterations. Autocorrelations are reported after a burn in of 5000 iterations. Posterior density estimates using EMCMC3 (noncentred) and AMCMC algorithms for (c) ρ , (d) μ , (e) σ , and (f) τ . In (c)-(e) we superimpose the posterior density obtained when the same process is directly observed.

6.3 Noisy observations

We now illustrate the performance of EMCMC by adding to the observations of the previous example a Gaussian error with mean 0 and variance $\tau^2 = 1/4$ (Figure 2c). We have assigned the same priors for the diffusion parameters as before, and an improper prior proportional to $1/\tau$ for τ . We employ EMCMC3 under the noncentred parametrisation and the HFA scheme with increasing values of $M = \{5, 10, 20, 40\}$. The algorithms are run for 10^5 iterations and the MCMC outputs are thinned every 10 iterations. Figure 5 presents autocorrelation plots for the parameters along with the marginal posterior density estimates. It is interesting to notice that the presence of noise in the observations seems to aid the approximation of the AMCMC methods, since the posterior densities do not change significantly as M increases, and seem to provide a reasonable approximation to the exact ones.

Method	Par.	λ	M	Mean	SD	ESS _{adj}	ESS	KS	Time	Correlation matrix		
EMCMC3 (noncentred)	ρ	2	8.466	0.089	0.010	2.820	52.718		18.695	1.000	0.464	0.370
	μ			2.026	0.161	3.130	58.521				1.000	-0.004
	σ			0.495	0.012	3.510	65.616				1.000	
EMCMC3 (interweaved)	ρ	2	8.313	0.089	0.010	2.769	54.197		19.574	1.000	0.465	0.384
	μ			2.027	0.160	3.189	62.422				1.000	0.012
	σ			0.495	0.012	3.655	71.534				1.000	
AMCMC-40	ρ	40.000	40.000	0.090	0.011	3.337	73.816	0.059	22.123	1.000	0.472	0.386
	μ			2.029	0.161	3.631	80.322	0.162			1.000	0.026
	σ			0.494	0.012	3.594	79.502	< 0.001			1.000	
AMCMC-40 (int-by-parts)	ρ	40.000	40.000	0.089	0.010	3.409	74.472	0.082	21.848	1.000	0.467	0.376
	μ			2.023	0.160	3.916	85.557	0.704			1.000	0.006
	σ			0.495	0.012	3.796	82.930	0.293			1.000	

Table 2: The DWELL diffusion model with $n = 1000$ simulated data points. True values are $(\rho, \mu, \sigma) = (0.1, 2, 1/2)$. Statistics are reported after a burn-in period of 5000 iterations. Column details as in Table 1.

6.4 A bivariate double well potential

We consider a double well potential process in two dimensions (denote by MVWELL hereafter), solution to

$$dV_s = -\frac{\sigma^2}{2}\nabla_v G(V_s)ds + \sigma dW_s, \quad \text{where } G(v) = \rho_1 \left[\left(v^{(2)} \right)^2 - \mu_1 \right]^2 + \rho_2 \left(v^{(2)} - \mu_2 v^{(1)} \right)^2, \quad (22)$$

and $\rho_1, \rho_2, \mu_1, \mu_2, \sigma > 0$. The parameter vectors are identified as $\theta_1 = (\rho_1, \mu_1, \rho_2, \mu_2)^T$ and $\theta_2 = \sigma$. The invariant density of the process is proportional to $\exp\{-G(v)\}$ and has two modes, at $(\sqrt{\mu_1}/\mu_2, \sqrt{\mu_1})$ and $(-\sqrt{\mu_1}/\mu_2, -\sqrt{\mu_1})$. This model belongs in the EA3 class and reduction to a unit volatility process is achieved as $X_s := V_s/\sigma$. The lower bound $l(\theta)$ and Poisson intensity $r(\bar{L}; \theta)$ are given in the Appendix.

We simulate $n = 1000$ equidistant observations with $\Delta t_i = 1$, $V_0 = (0, 0)^T$ and parameters $(\rho_1, \mu_1, \rho_2, \mu_2, \sigma) = (1/2, 2, 1/2, 1, 1/2)$. The simulated dataset is shown in Figures 2d and 2e. We assign improper priors to the parameters with $\pi(\rho_i) \propto 1, \pi(\mu_i) \propto 1, \pi(\sigma) \propto 1/\sigma, i = 1, 2$. All MCMC chains were run for 10^5 iterations and the performance of the algorithms is shown in Figure 6. For $\lambda = 5$, the EMCMC3 algorithm under the original parametrisation performs poorly, whereas the noncentred exhibits a much more rapid mixing. On the other hand, the interweaved strategy with $\lambda = 2$ shows a comparable performance to that of the noncentred. Posterior summary statistics from the algorithms are shown in Table 3. From the approximate methods, we found that AMCMC-60 paired with integration by parts was the most efficient algorithm which provided a sufficiently accurate approximation to the posterior marginal distributions.

Finally, as we pointed out in Section 2.2.2, the EA3 is inflicted by an additional computational cost due to the rejection sampler for the layered Brownian bridge, which is clearly reflected in the CPU time of the interweaved algorithm. In particular, a computational profiling of the algorithm showed that approximately 91% of the total time was used by EA3, out of which 85% was due to the simulation of layered Brownian bridges. This suggests that an alternative more efficient design of the layered Brownian bridge simulation would boost substantially the performance of EMCMC3.

7 Discussion

We have developed exact data augmentation methods for discretely directly and indirectly observed diffusions. We have established the precise connection between this paradigm and the best existing alternative method when the variance-stabilising transformation can be performed, the HFA. The empirical comparison of the two methods showed that in univariate processes EMCMC can perform at least as well as a sufficiently accurate AMCMC, even when ignoring the additional computational cost needed by the latter to determine a good value of M through experimentation. For the considered bivariate example, EMCMC is outperformed by AMCMC since the cost of the former is dominated by the simulation of the layered Brownian bridges, and thus could be improved by considering alternative designs for this simulation.

We have also pointed out an intriguing connection between exact and approximate methods: the degree of freedom rendered by the Poisson sampling rate. On going work involves the rigorous proof of the effect of the auxiliary Poisson sampling. The auxiliary sampling can be seen as a variance reduction scheme. In general,

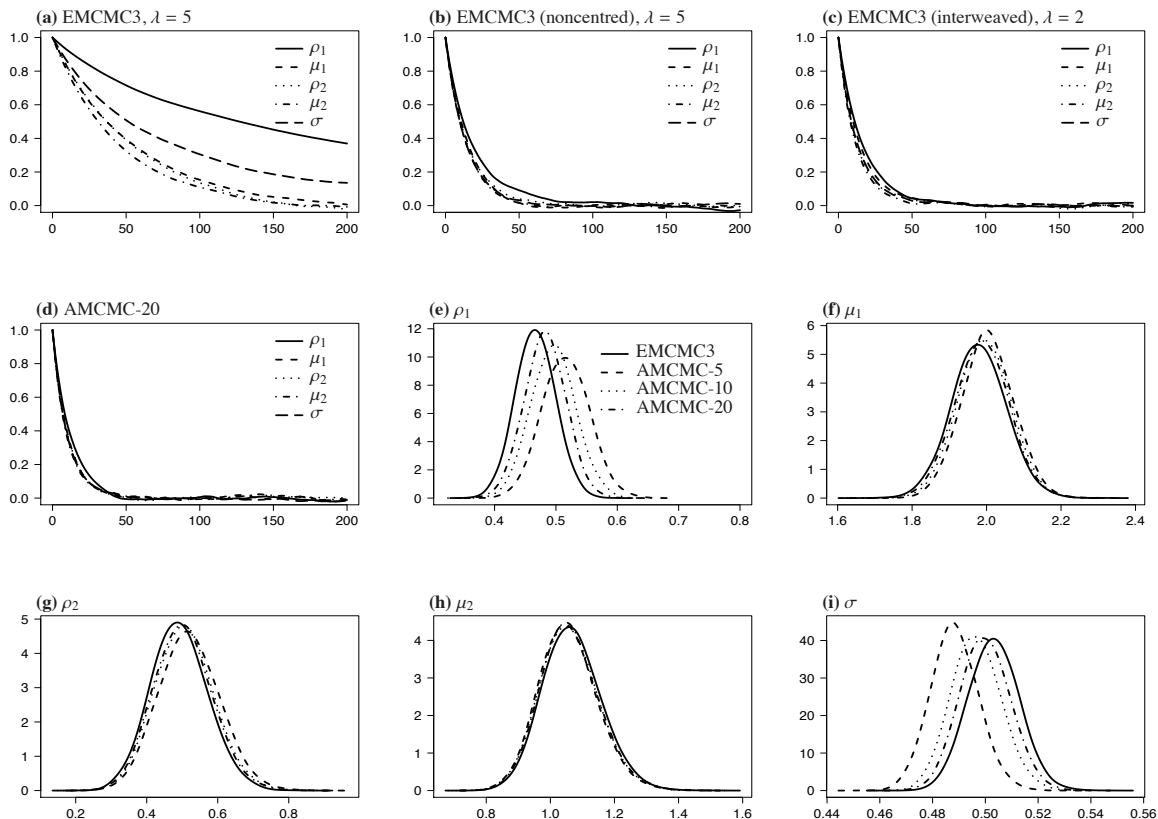


Figure 6: The MVWELL diffusion model with $n = 1000$ simulated data points. True values are $(\rho_1, \mu_1, \rho_2, \mu_2, \sigma) = (1/2, 2, 1/2, 1, 1/2)$. Autocorrelations are reported after a burn in of 5000 iterations. Posterior density estimates using EMCMC3 (interweaved) and AMCMC algorithms for (e) ρ_1 , (f) μ_1 , (g) ρ_2 , (h) μ_2 and (i) σ .

Method	Par.	λ	M	Mean	SD	ESS_{adj}	ESS	KS	Time	Correlation matrix				
EMCMC3 (noncentred)	ρ_1	5	10.455	0.468	0.032	0.482	29.326		60.819	1.000	0.208	0.056	-0.032	-0.069
	μ_1			1.980	0.072	0.547	33.281				1.000	0.492	-0.591	0.016
	ρ_2			0.494	0.079	0.558	33.947					1.000	-0.814	-0.032
	μ_2			1.065	0.090	0.584	35.517						1.000	0.019
	σ			0.503	0.009	0.579	35.224							1.000
EMCMC3 (interweaved)	ρ_1	2	7.233	0.467	0.032	0.674	31.721		47.057	1.000	0.192	0.018	-0.019	-0.062
	μ_1			1.980	0.073	0.751	35.345				1.000	0.482	-0.582	0.020
	ρ_2			0.492	0.078	0.872	41.032					1.000	-0.816	-0.018
	μ_2			1.068	0.089	0.936	44.029						1.000	-0.001
	σ			0.503	0.009	0.843	39.673							1.000
AMCMC-60 (int-by-parts)	ρ_1		60.000	0.467	0.032	1.447	44.949	0.100	31.065	1.000	0.184	0.022	-0.008	-0.053
	μ_1			1.979	0.072	1.518	47.154	0.831			1.000	0.490	-0.594	-0.008
	ρ_2			0.494	0.079	1.594	49.515	0.361				1.000	-0.819	-0.031
	μ_2			1.066	0.090	1.734	53.862	0.268					1.000	0.027
	σ			0.503	0.009	1.621	50.371	0.932						1.000

Table 3: The MVWELL diffusion model with $n = 1000$ simulated data points. True values are $(\rho_1, \mu_1, \rho_2, \mu_2, \sigma) = (1/2, 2, 1/2, 1, 1/2)$. Statistics are reported after a burn-in period of 5000 iterations. Column details as in Table 1.

there is a large scope for investigating other such schemes both for EDA and HFA. In this article we have already demonstrated the effect of performing integration by parts where possible to the efficiency of the algorithms.

The extension of these methods outside the class of processes prescribed by the current version of EA3 is definitely an exciting direction. Another direction of interest for future research is to explore further the connection between unbiased estimation of transition density and MCMC. There is a large and growing literature which develops MCMC algorithms for models with intractable likelihoods using unbiased estimators thereof; see for example Andrieu and Roberts (2009); Andrieu *et al.* (2010). The class of diffusions where such estimators can be

obtained is much larger than that simulated by EA3, see for example Section 4.6 of Papaspiliopoulos (2011).

There exists available software for implementing all the methods in this paper, which is available on request by the authors.

Acknowledgements

O. Papaspiliopoulos would like to acknowledge financial support by the Spanish government through a “Ramon y Cajal” fellowship and grant MTM2009-09063. G. Sermaidis was funded by the Greek State Scholarships Foundation. G. Roberts acknowledges CRiSM and EPSRC.

References

- Aït-Sahalia, Y. (2008) Closed-form likelihood expansions for multivariate diffusions. *Ann. Statist.*, **36**, 906–937.
- Aït-Sahalia, Y. and Kimmel, R. (2007) Maximum likelihood estimation of stochastic volatility models. *Journal of Financial Economics*, **83**, 413–452.
- Andrieu, C., Doucet, A. and Holenstein, R. (2010) Particle Markov chain Monte Carlo. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, **3**, 269–342.
- Andrieu, C. and Roberts, G. O. (2009) The pseudo-marginal approach for efficient Monte Carlo computations. *Ann. Statist.*, **37**, 697–725.
- Beskos, A., Papaspiliopoulos, O. and Roberts, G. O. (2006a) Retrospective exact simulation of diffusion sample paths with applications. *Bernoulli*, **12**, 1077–1098.
- Beskos, A., Papaspiliopoulos, O. and Roberts, G. O. (2008) A factorisation of diffusion measure and finite sample path constructions. *Methodol. Comput. Appl. Probab.*, **10**, 85–104.
- Beskos, A., Papaspiliopoulos, O., Roberts, G. O. and Fearnhead, P. (2006b) Exact and computationally efficient likelihood-based estimation for discretely observed diffusion processes. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, **68**, 333–382. With discussion and a reply by the authors.
- Brown, P. E., Kåresen, K. F., Roberts, G. O. and Tonellato, S. (2000) Blur-generated non-separable space-time models. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, **62**, 847–860.
- Cotter, S. L., Roberts, G. O., Stuart, A. M. and White, D. (2012) MCMC methods for functions: modifying old algorithms to make them faster. Submitted.
- Elerian, O., Chib, S. and Shephard, N. (2001) Likelihood inference for discretely observed nonlinear diffusions. *Econometrica*, **69**, 959–993.
- Eraker, B. (2001) MCMC analysis of diffusion models with application to finance. *J. Bus. Econom. Statist.*, **19**, 177–191.
- Étoré, P. and Martinez, M. (2011) Exact simulation of one-dimensional stochastic differential equations involving the local time at zero of the unknown process. Tech. rep. Available online from <http://arxiv.org/abs/1102.2565>.
- Forman, J. L. and Sørensen, M. (2008) The Pearson diffusions: a class of statistically tractable diffusion processes. *Scand. J. Statist.*, **35**, 438–465.
- Girolami, M. and Calderhead, B. (2011) Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, **73**, 123–214. With discussion and a reply by the authors.
- Golightly, A. and Wilkinson, D. J. (2006) Bayesian sequential inference for stochastic kinetic biochemical network models. *J. Comput. Biol.*, **13**, 838–851.
- Golightly, A. and Wilkinson, D. J. (2008) Bayesian inference for nonlinear multivariate diffusion models observed with error. *Comput. Statist. Data Anal.*, **52**, 1674–1693.

- Gonçalves, F. and Roberts, G. (2012) Exact simulation problems for jump-diffusions. Submitted.
- Horenko, I. and Schütte, C. (2008) Likelihood-based estimation of multidimensional Langevin models and its application to biomolecular dynamics. *Multiscale Model. Simul.*, **7**, 731–773.
- Kalogeropoulos, K., Roberts, G. and Dellaportas, P. (2010) Inference for stochastic volatility models using time change transformations. *Ann. Statist.*, **38**, 784–807.
- Metzner, P., Schütte, C. and Vanden-Eijnden, E. (2006) Illustration of transition path theory on a collection of simple examples. *The Journal of Chemical Physics*, **125**, 084110.
- Øksendal, B. (2003) *Stochastic differential equations. An introduction with applications*. Universitext. Berlin: Springer-Verlag, 6th edn.
- Papaspiliopoulos, O. (2011) A methodological framework for Monte Carlo probabilistic inference for diffusion processes. In *Bayesian Time Series Models*. Cambridge University Press.
- Papaspiliopoulos, O. and Roberts, G. O. (2012) Importance sampling techniques for estimation of diffusion models. In *Statistical Methods for Stochastic Differential Equations*, 311–337. Monographs on Statistics and Applied Probability, Chapman and Hall.
- Papaspiliopoulos, O., Roberts, G. O. and Sköld, M. (2007) A general framework for the parametrization of hierarchical models. *Statistical Science*, **22**, 59–73.
- Peluchetti, S. and Roberts, G. O. (2008) An empirical study of the efficiency of the EA for diffusion simulation. Tech. rep., University of Warwick.
- Picchini, U., Gaetano, A. and Ditlevsen, S. (2010) Stochastic Differential Mixed-Effects Models. *Scandinavian Journal of Statistics*, **37**, 67–90.
- Pitt, M. and Shephard, N. (1999) Analytic convergence rates and parameterization issues for the Gibbs sampler applied to state space models. *Journal of Time Series Analysis*, **20**, 63–85.
- Plummer, M., Best, N., Cowles, K. and Vines, K. (2010) *coda: Output analysis and diagnostics for MCMC*. URL <http://CRAN.R-project.org/package=coda>. R package version 0.13-5.
- R Development Core Team (2010) *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <http://www.R-project.org>.
- Roberts, G. O., Papaspiliopoulos, O. and Dellaportas, P. (2004) Bayesian inference for non-Gaussian Ornstein-Uhlenbeck stochastic volatility processes. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, **66**, 369–393.
- Roberts, G. O. and Stramer, O. (2001) On inference for partially observed nonlinear diffusion models using the Metropolis-Hastings algorithm. *Biometrika*, **88**, 603–621.
- Rydborg, T. H. (1997) A note on the existence of unique equivalent martingale measures in a Markovian setting. *Finance and Stochastics*, **1**, 251–257.
- Sekhon, J. S. (2007) Multivariate and propensity score matching software with automated balance optimization: The matching package for R. *Journal of Statistical Software*.
- Sermaidis, G. (2010) *Likelihood-based inference for discretely observed diffusions*. Ph.D. thesis, Department of Statistics, University of Warwick.
- Yu, Y. and Meng, X. (2011) To center or not to center, that is not the question: An ancillarity-sufficiency interweaving strategy (ASIS) for boosting MCMC efficiency. *J. Comput. Graph. Statist.*, **20**, 531–570.

8 Appendix

Proof of Lemma 2

Proof. Expression (15) is derived by writing the density of the accepted random variables $(\tilde{L}, \tilde{X}, \Phi)$ with respect to the law of the proposed

$$\pi(\tilde{L}, \tilde{X}, \Phi | v, w, \theta) = \frac{1}{a(x, y, t, \theta)} \prod_{j=1}^{\kappa} \mathbb{I} \left[\frac{1}{r(\tilde{L}; \theta)} \phi \left\{ \tilde{X}_{\psi_j} + \left(1 - \frac{\psi_j}{t}\right) x(\theta_2) + \frac{\psi_j}{t} y(\theta_2); \theta \right\} < u_j \right],$$

and by invoking a change of measure from the law of a Poisson process of intensity $r(\tilde{L}; \theta)$ to the law of one of unit intensity, thus ensuring a parameter-independent dominating measure. Finally, integrating out the marks $\Upsilon = \{u_j, 1 \leq j \leq \kappa\}$, we obtain expression (15). \square

Proof of Theorem 1

Proof. The factorisation of the density in the three terms is elementary. For $V_0 = v$, $V_t = w$ and $x = \eta(v; \theta_2)$, $y = \eta(w; \theta_2)$, taking expectations on both sides of (5) with respect to $\mathbb{W}_\theta^{(t,x,y)}$ we derive the fundamental identity

$$\tilde{p}_t(x, y; \theta) = \mathcal{N}_t(y - x) \exp \{H(y; \theta) - H(x; \theta) - l(\theta)t\} a(x, y, t, \theta),$$

which combined with (4) gives $a(x, y, t, \theta)$ as a function of $p_t(v, w; \theta)$. Combining this with Lemma 2 yields the expression.

It remains to show that (2) can be obtained by integrating out the auxiliary variables and conditioning on the data. However, this is trivial since by taking expectations with respect to the dominating measure $\otimes_{i=1}^n (\mathbb{M}^{(\Delta t_i)} \times \mathbb{P}^{(\Delta t_i)})$ we obtain the marginal $\pi(\theta) \prod_{i=1}^n p_{\Delta t_i}(V_{t_{i-1}}, V_{t_i}; \theta)$ from which (2) follows as a conditional. \square

Proof of Theorem 2

Proof. For notational simplicity, we define $K(Y, \theta)$ to be the following deterministic function of observations Y and θ :

$$\exp \{H\{x_n(\theta_2); \theta\} - H\{x_0(\theta_2); \theta\} - l(\theta)(t_n - t_0)\} \prod_{i=1}^n D(V_{t_i}; \theta_2) \mathcal{N}_{\Delta t_i} \{x_i(\theta_2) - x_{i-1}(\theta_2)\}.$$

The joint posterior density of θ and imputed data $\{\tilde{L}_i, \tilde{X}_i, \Psi_i, 1 \leq i \leq n\}$ is given (up to a constant) in expression (16). Integrating out the Poisson processes is easily done by first integrating out the coordinates and then the number of Poisson points. Therefore, integrating out $\Psi_i = \{\psi_{i,j}, 1 \leq j \leq \kappa_i\}, 1 \leq i \leq n$, we obtain

$$\pi(\theta) K(Y, \theta) \exp \left(- \sum_{i=1}^n [r(\tilde{L}_i; \theta) - 1] \Delta t_i \right) \prod_{i=1}^n \left\{ \frac{1}{\Delta t_i} \int_0^{\Delta t_i} [r(\tilde{L}_i; \theta) - \phi \{ \tilde{X}_{i,s} + \mu_{i,s}(\theta); \theta \}] ds \right\}^{\kappa_i}.$$

Integrating out the Poisson points yields the posterior density of θ and $\{\tilde{L}_i, \tilde{X}_i, 1 \leq i \leq n\}$ with respect to $\text{Leb}^p \otimes_{i=1}^n \mathbb{M}^{(\Delta t_i)}$, as

$$\begin{aligned} & \pi(\theta) K(Y, \theta) \exp \left(- \sum_{i=1}^n [r(\tilde{L}_i; \theta) - 1] \Delta t_i \right) \prod_{i=1}^n \exp \left\{ \Delta t_i \left[\frac{1}{\Delta t_i} \int_0^{\Delta t_i} [r(\tilde{L}_i; \theta) - \phi \{ \tilde{X}_{i,s} + \mu_{i,s}(\theta); \theta \}] ds - 1 \right] \right\} \\ & = \pi(\theta) K(Y, \theta) \exp \left\{ - \sum_{i=1}^n \int_0^{\Delta t_i} \phi \{ \tilde{X}_{i,s} + \mu_{i,s}(\theta); \theta \} ds \right\}. \end{aligned}$$

Given the construction of $\mathbb{M}^{(\Delta t_i)}$, by integrating out the layers we obtain the joint density of θ and $\{\tilde{X}_i, 1 \leq i \leq n\}$ with respect to $\text{Leb}^p \otimes_{i=1}^n \mathbb{W}^{(\Delta t_i, 0, 0)}$ which coincides with the posterior density (11) targeted by PA. \square

Proof of Theorem 3

Proof. For a pair of observations $(V_{t_{i-1}}, V_{t_i})$, the joint density of the accepted elements of EA3 $(\tilde{L}_i, \tilde{X}_i, \tilde{\Psi}_i)$ conditionally on $V_{t_{i-1}}, V_{t_i}, \theta$ is given by

$$\frac{\prod_{j=1}^{\infty} \left\{ 1 - \mathbb{I} \left[\tilde{\xi}_{i,j} < r(\tilde{L}_i; \theta) \right] \phi \left\{ \tilde{X}_{i,\tilde{\psi}_{i,j}} + \mu_{i,\tilde{\psi}_{i,j}}(\theta); \theta \right\} / r(\tilde{L}_i; \theta) \right\}}{a(x_{i-1}(\theta_2), x_i(\theta_2), \Delta t_i, \theta)},$$

with respect to the product measure $\mathbb{M}^{(\Delta t_i)} \times \tilde{\mathbb{L}}^{(\Delta t_i)}$, where $\tilde{\mathbb{L}}^{(t)}$ is the measure of a unit rate Poisson process on $[0, t] \times (0, \infty)$. Using the conditional independence of the latent data given Y and θ ,

$$\pi_{nc}(\{\tilde{L}_i, \tilde{X}_i, \tilde{\Psi}_i, 1 \leq i \leq n\} | Y, \theta) = \prod_{i=1}^n \pi_{nc}(\tilde{L}_i, \tilde{X}_i, \tilde{\Psi}_i | V_{t_{i-1}}, V_{t_i}, \theta),$$

and the proof follows along the same lines as Theorem 1. \square

Functions related to the MVWELL model

Applying the Lamperti transformation to (22), we obtain the drift of the transformed process as

$$\alpha(x; \theta) = \nabla_x H(x; \theta), \text{ where } H(x; \theta) = -\frac{1}{2} G(\sigma x).$$

Using second partial derivative tests it is straightforward to verify that the function $f(u; \theta) := \|\alpha(u; \theta)\|^2 + \Delta_x H(u; \theta)/2$ has a global minimum

$$l(\theta) = -\frac{p_1 + p_2}{54\mu_2^2(1 + \mu_2^2)},$$

where

$$p_1 := 2\mu_2\sigma^2 \sqrt{2\rho_1} \left\{ 9 + \mu_2^2 (9 + 2\rho_1\mu_1^2) \right\}^{3/2}, \quad p_2 := \mu_2^2\sigma^2 \left\{ 54\rho_1\mu_1(1 + \mu_2^2) - 8\rho_1^2\mu_1^3\mu_2^2 + 27\rho_2(1 + \mu_2^2)^2 \right\}.$$

The function $f(u; \theta)$ has no local maxima, implying that for a given realisation of the layer \tilde{L} the Poisson rate is

$$r(\tilde{L}; \theta) = \left[\bigvee_{i=1}^4 f(u_i; \theta) \right] - l(\theta),$$

where

$$\begin{aligned} u_1 &= \left(x^{(1)}(\theta_2) - \tilde{L}^{(1)}\delta, x^{(2)}(\theta_2) - \tilde{L}^{(2)}\delta \right)^T, & u_2 &= \left(x^{(1)}(\theta_2) - \tilde{L}^{(1)}\delta, y^{(2)}(\theta_2) + \tilde{L}^{(2)}\delta \right)^T, \\ u_3 &= \left(y^{(1)}(\theta_2) + \tilde{L}^{(1)}\delta, x^{(2)}(\theta_2) - \tilde{L}^{(2)}\delta \right)^T, & u_4 &= \left(y^{(1)}(\theta_2) + \tilde{L}^{(1)}\delta, y^{(2)}(\theta_2) + \tilde{L}^{(2)}\delta \right)^T. \end{aligned}$$