



Accelerating Bayesian inference for stochastic epidemic models using incidence data

Andrew Golightly¹ · Laura E. Wadkin² · Sam A. Whitaker² · Andrew W. Baggaley² · Nick G. Parker² · Theodore Kypraios³

Received: 27 March 2023 / Accepted: 26 September 2023
© The Author(s) 2023

Abstract

We consider the case of performing Bayesian inference for stochastic epidemic compartment models, using incomplete time course data consisting of incidence counts that are either the number of new infections or removals in time intervals of fixed length. We eschew the most natural Markov jump process representation for reasons of computational efficiency, and focus on a stochastic differential equation representation. This is further approximated to give a tractable Gaussian process, that is, the linear noise approximation (LNA). Unless the observation model linking the LNA to data is both linear and Gaussian, the observed data likelihood remains intractable. It is in this setting that we consider two approaches for marginalising over the latent process: a correlated pseudo-marginal method and analytic marginalisation via a Gaussian approximation of the observation model. We compare and contrast these approaches using synthetic data before applying the best performing method to real data consisting of removal incidence of oak processionary moth nests in Richmond Park, London. Our approach further allows comparison between various competing compartment models.

Keywords Stochastic epidemic model · Bayesian inference · Linear noise approximation · Incidence data · Oak processionary moth

1 Introduction

The starting point for formulating a model of epidemic spread is usually a set of compartments or classes, characterising the individuals participating in the epidemic (see e.g. Jacquez 1972). A transmission model then describes the dynamics of individuals within each class, and this can be combined with models of severity and detection/observation (see e.g. Birrell et al. 2020) to give an overarching model that links epidemic data to latent disease transmission. The transmission model can be deterministic or stochastic, with the former typically taking the form of a coupled, nonlinear ordinary differential equation (ODE) system and the latter a continuous-time

Markov jump process (MJP). The link between deterministic and stochastic models is made clear in Kurtz (1970) (see also Kurtz 1971), with the ODE approach seen as a large population limit of the MJP. However, while fitting ODE models in the presence of a simple observation model is relatively straightforward, they ignore intrinsic stochasticity, which may play an important role in the dynamics of the epidemic, particularly at the start or the end of an outbreak, when numbers of individuals with the disease are likely to be small. Stochastic transmission models have been widely adopted for small-size epidemics (see e.g. Boys and Giles 2007; Stockdale et al. 2017; Funk et al. 2018) whereas Corbella et al. (2022) (see also Birrell et al. 2011) combine a deterministic transmission model with sophisticated stochastic observation models.

Although complex stochastic epidemic models can provide an accurate description of disease transmission, fitting such models to discrete-time data that may be incomplete and subject to error is typically complicated by the intractability of the observed data likelihood. Earlier attempts to address this issue include the use of data augmentation and Metropolis-Hastings steps to integrate over the uncertainty

✉ Andrew Golightly
andrew.golightly@durham.ac.uk

¹ Department of Mathematical Sciences, Durham University, Stockton Road DH1 3LE, UK

² School of Mathematics, Statistics and Physics, Newcastle University, Newcastle Upon Tyne NE1 7RU, UK

³ School of Mathematical Sciences, University Of Nottingham, Nottingham NG7 2RD, UK

associated with unobserved dynamic components (Gibson and Renshaw 1998; O’Neill and Roberts 1999). Recent attention has focused on approaches that only require the ability to generate forward realisations from the transmission model such as approximate Bayesian computation (McKinley et al. 2009; Kypraios et al. 2017) and pseudo-marginal methods (McKinley et al. 2014; Spannaus et al. 2020; Corbella et al. 2022). Nevertheless, these methods typically require many (millions of) model simulations which can be impracticable for large-size epidemics, that is, epidemics in which the population size is greater than a few thousand individuals. Alternative approaches that aim to reduce the computational burden include the use of cheap but approximate “surrogate” methods including the use of Gaussian process emulation (Scarponi et al. 2022; Swallow et al. 2022), direct approximation of the MJP e.g. via a linear noise approximation (LNA, Fintzi et al. 2021) and tractable approximation of the observed data likelihood (see e.g. Whiteley and Rimella 2021).

In this paper, we consider incidence data consisting of the number of new cases (either infections or removals) accumulated in an observation interval. We further assume imperfect observations and consider two approaches commonly used to allow for under reporting or over dispersion. We build on the work of Fintzi et al. (2021) by adopting a linear noise approximation of the latent cumulative incidence process. Whereas Fintzi et al. (2021) integrate over the uncertainty in the latent process via a sampling approach, we introduce a further Gaussian approximation of the observation model, allowing analytic integration of the latent incidence process. Our framework additionally allows for a time-varying infection rate, as is typically required for accounting for seasonality and/or interventions. The infection rate is modelled stochastically as an additional component in the system of stochastic differential equations which the LNA approximates. Hence, our contribution is a fast and efficient sampling-based framework for inferring the parameters of a general class of stochastic epidemic models. We benchmark the performance of our approach against state-of-the-art correlated pseudo-marginal methods (Dahlin et al. 2015; Deligiannidis et al. 2018) in terms of both accuracy and efficiency, using a susceptible-infectious-removed (SIR) model.

Finally, we consider the application of the methodology to the infestation of the oak processionary moth (OPM), *Thaumetopoea processionea*, in Richmond Park, London. Using time course data consisting of the yearly removal incidence of infested trees between the years 2013 and 2020, we compare and contrast an SIR model with a model in which infected trees can re-enter the susceptible class. The assumed initial population size is some 40,000 oak trees with the number of susceptible trees reducing to around 35,000 over the time frame of the data set. The size of the epidemic necessitates analytic integration of the latent incidence process;

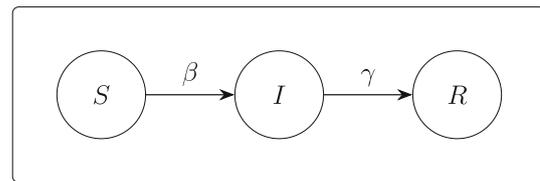


Fig. 1 SIR compartment model

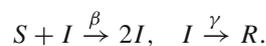
discrete stochastic models combined with exact (simulation-based) inference methods such as data augmentation (see e.g. Jewell et al. 2009) are practically infeasible here; see Stockdale et al. (2021) for a recent discussion.

The remainder of this paper is organised as follows. Stochastic epidemic models, and in particular, the jump process and subsequent LNA representations of the cumulative incidence process are considered in Sect. 2. The inference task is described in Sect. 3, including marginalisation of the incidence process via pseudo-marginal and analytic methods, with the latter using an additional approximation. Applications are considered in Sect. 4 with a discussion of limitations of the proposed approach in Sect. 5.

2 Stochastic epidemic models

For ease of exposition, we consider an SIR epidemic model (Andersson and Britton 2000; Kermack and McKendrick 1927) within which a population of fixed size N_{pop} is classified into compartments consisting of susceptible (S), infectious (I) and removed (R) individuals. However, we note that the framework can be easily extended to allow for additional compartments e.g. SEIR models (Hethcote 2000) which allows for susceptibles entering an exposed class prior to becoming infectious.

The SIR compartment model is shown in Fig. 1. Transitions between compartments can be described by the set of pseudo-reactions given by



The first transition describes contact of an infective individual with a susceptible and with the net effect resulting in an additional infective individual and one fewer susceptible. The final transition accounts for removal (recovered with immunity, quarantined or dead) of an infected individual. The components of $\theta = (\beta, \gamma)'$ denote infection and removal rates.

In what follows, we describe the most natural stochastic SIR model as a Markov jump process before considering a linear noise approximation suitable for incidence data.

2.1 Markov jump process

Let $X_t = (S_t, I_t)'$ denote the numbers in each state at time $t \geq 0$ and note that $R_t = N_{\text{pop}} - S_t - I_t$ for all $t \geq 0$. We denote by $\theta = (\beta, \gamma)'$ the vector of parameters. The dynamics of $\{X_t, t \geq 0\}$ are most naturally described by a Markov jump process (MJP), that is, a continuous time, discrete valued Markov process. Assuming that at most one event can occur over a time interval $(t, t + \Delta t]$ and that the state of the system at time t is $x_t = (s_t, i_t)'$, the MJP representation of the prevalence process is characterised by transition probabilities of the form

$$\begin{aligned} \mathbb{P}(X_{t+\Delta t} = (s_t - 1, i_t + 1)' | x_t, \theta) &= \beta s_t i_t \Delta t + o(\Delta t), \\ \mathbb{P}(X_{t+\Delta t} = (s_t, i_t - 1)' | x_t, \theta) &= \gamma i_t \Delta t + o(\Delta t), \\ \mathbb{P}(X_{t+\Delta t} = (s_t, i_t)' | x_t, \theta) &= 1 - (\beta s_t i_t + \gamma i_t) \Delta t + o(\Delta t), \end{aligned}$$

where $o(\Delta t)/\Delta t \rightarrow 0$ as $\Delta t \rightarrow 0$. Similarly, the cumulative incidence of infection and removal events $\{N_t, t \geq 0\}$ is an MJP governed by the transition probabilities

$$\begin{aligned} \mathbb{P}(N_{t+\Delta t} = (n_{1,t} + 1, n_{2,t})' | n_t, x_t, \theta) &= \beta s_t i_t \Delta t + o(\Delta t), \\ \mathbb{P}(N_{t+\Delta t} = (n_{1,t}, n_{2,t} + 1)' | n_t, x_t, \theta) &= \gamma i_t \Delta t + o(\Delta t), \\ \mathbb{P}(N_{t+\Delta t} = (n_{1,t}, n_{2,t})' | n_t, x_t, \theta) &= 1 - (\beta s_t i_t + \gamma i_t) \Delta t + o(\Delta t), \end{aligned}$$

where $n_t = (n_{1,t}, n_{2,t})'$ denotes the cumulative number of infection and removal events at time t .

The processes describing the prevalence $\{X_t, t \geq 0\}$ and cumulative incidence $\{N_t, t \geq 0\}$ are linked via the equation

$$X_t = x_0 + \sum_{i=1}^2 S^i N_{i,t} \tag{1}$$

where x_0 represents the initial number of susceptible and infected individuals, $N_{i,t}$ denotes the i th component of the incidence process at time t and S^i is the i th column of the stoichiometry matrix

$$S = \begin{pmatrix} -1 & 0 \\ 1 & -1 \end{pmatrix}.$$

Given x_0 and θ , generating exact realisations of the incidence process, and therefore the stochastic SEIR model is straightforward and can be achieved by using well known simulation algorithms from the stochastic kinetic models literature (see e.g. Wilkinson 2018). At this point, it will be

helpful to define the instantaneous rate or hazard function $h(x_t) = (h_1(x_t), h_2(x_t))'$ by

$$\begin{aligned} h(x_t) &= \lim_{\Delta t \rightarrow 0} \mathbb{P}(X_{t+\Delta t} | x_t, \theta) / \Delta t \\ &= (\beta s_t i_t, \gamma i_t)' \end{aligned}$$

for $X_{t+\Delta t}$ resulting from either an infection or removal reaction respectively. Note that θ is suppressed here for notational convenience. Then, given a state x_t at time t , Gillespie's direct method (Gillespie 1977) simulates the time to the next event as an exponential random variable with rate $\sum_{i=1}^2 h_i$. The event that occurs will be of type i (with $i = 1$ infection, $i = 2$ removal) with probability proportional to h_i .

2.2 Linear noise approximation (LNA)

The linear noise approximation (LNA) is most commonly presented as a Gaussian process approximation of the MJP description of the prevalence process $\{X_t, t \geq 0\}$ (see e.g. Ross et al. 2009; Fearnhead et al. 2014; Fuchs 2013 in the epidemic context and Ferm et al. (2008), Komorowski et al. (2009) and Stathopoulos and Girolami (2013) in a wider systems biology context). As in Fintzi et al. (2021), we require an approximation of the cumulative incidence process $\{N_t, t \geq 0\}$. Although this can be derived in a number of more or less formal ways, we give an intuitive derivation along the lines of Wallace (2010).

Consider an infinitesimal time interval, $(t, t + dt]$, over which the reaction hazards will remain constant almost surely. Consequently, the counting process over this interval for the i th component, denoted by $dN_{i,t}$, is Poisson distributed with rate $h_i dt$. Stacking these quantities in the vector dN_t it should be clear that

$$E(dN_t) = h(x_t)dt, \quad \text{Var}(dN_t) = \text{diag}\{h(x_t)\}dt,$$

Hence, the Itô Stochastic differential equation (SDE) that best matches the MJP representation of the incidence process is given by

$$dN_t = h(x_t)dt + \text{diag}\{\sqrt{h(x_t)}\} dW_t, \tag{2}$$

where W_t is a length-2 vector of uncorrelated standard Brownian motion processes and $\text{diag}\{\sqrt{h(x_t)}\}$ is a 2×2 diagonal matrix with non-zero entries given by $\sqrt{h_1(x_t)}$ and $\sqrt{h_2(x_t)}$. Note that the RHS of (2) can be written explicitly in terms of n_t via (1) which gives $x_t = (s_0 - n_{t,1}, i_0 + n_{t,1} - n_{t,2})'$. We then further define the hazard function written in terms of n_t as

$$h^*(n_t) = (\beta[s_0 - n_{t,1}][i_0 + n_{t,1} - n_{t,2}], \gamma[i_0 + n_{t,1} - n_{t,2}])'$$

for which (2) becomes

$$dN_t = h^*(n_t)dt + \text{diag}\{\sqrt{h^*(n_t)}\} dW_t. \tag{3}$$

The SDE in (3) can be linearised as follows. Consider a partition of N_t as $N_t = \eta_t + R_t$ where η_t is a deterministic process satisfying the ordinary differential equation (ODE)

$$\frac{d\eta_t}{dt} = h^*(\eta_t) \tag{4}$$

and $R_t = N_t - \eta_t$ is a residual stochastic process satisfying the (typically) intractable SDE

$$dR_t = \{h^*(n_t) - h^*(\eta_t)\} dt + \text{diag}\{\sqrt{h^*(n_t)}\} dW_t. \tag{5}$$

We obtain an approximate, tractable \hat{R}_t by Taylor expanding $h^*(n_t)$ and $\text{diag}\{h^*(n_t)\}$ about η_t . Retaining the first two terms in the expansion of the former and the first term in the expansion of the latter gives

$$d\hat{R}_t = F_t \hat{R}_t dt + \text{diag}\{\sqrt{h^*(\eta_t)}\} dW_t \tag{6}$$

where F_t is the Jacobian matrix with (i, j) th element given by the partial derivative of the i th component of $h^*(\eta_t)$ with respect to the j th component of η_t . Hence, we obtain

$$F_t = \begin{pmatrix} \beta(s_0 - i_0 - 2\eta_{t,1} + \eta_{t,2}) & \beta(\eta_{t,1} - s_0) \\ \gamma & -\gamma \end{pmatrix}.$$

The solution of (6) is straightforward to obtain (see e.g. Fearnhead et al. 2014, among several others). Omitting details, we arrive at

$$N_t | (N_0 = \eta_0 + r_0) \sim N(\eta_t + G_t r_0, V_t) \tag{7}$$

where the fundamental matrix G_t satisfies

$$\frac{dG_t}{dt} = F_t G_t, \quad G_0 = I_2 \tag{8}$$

and V_t satisfies

$$\frac{dV_t}{dt} = V_t F_t' + \text{diag}\{h^*(\eta_t)\} + F_t V_t, \quad V_0 = 0_2. \tag{9}$$

Note that I_2 and 0_2 are the 2×2 identity and zero matrices respectively. The LNA for the cumulative incidence process is then summarised by (7) and (4), (8), (9).

2.3 Time varying infection rate

In practice it may be unreasonable to assume that the infection rate in the SIR model remains constant throughout the epidemic (e.g. due to seasonality and/or interventions). We therefore follow Dureau et al. (2013) and Spannaus et al. (2020) (among others) and describe the contact rate via an

Itô stochastic differential equation (SDE). Let $\{\beta_t, t \geq 0\}$ denote the infection process and consider $N_{3,t} = \log(\beta_t)$, assumed to satisfy a time-homogeneous SDE of the form

$$dN_{3,t} = a(n_{3,t})dt + b(n_{3,t})dW_{3,t} \tag{10}$$

where $\{W_{3,t}, t \geq 0\}$ is a standard Brownian motion process. Combining (10) with (3) gives a coupled SDE for $N_t = (N_{1,t}, N_{2,t}, N_{3,t})'$ of the form

$$dN_t = \{h_1^*(n_t), h_2^*(n_t), a(n_{3,t})\} dt + \text{diag}\left\{\sqrt{h_1^*(n_t)}, \sqrt{h_2^*(n_t)}, b(n_{3,t})\right\} dW_t \tag{11}$$

where W_t is a length-3 vector of uncorrelated standard Brownian motion processes. The LNA of (11) follows in the same way as Sect. 2.2, albeit with the (transpose of the) Jacobian matrix F_t redefined as

$$F_t' = \begin{pmatrix} \exp(\eta_{3,t})(s_0 - i_0 - 2\eta_{t,1} + \eta_{t,2}) & \gamma & 0 \\ \exp(\eta_{3,t})(\eta_{t,1} - s_0) & -\gamma & 0 \\ \exp(\eta_{3,t})(s_0 - \eta_{t,1})(i_0 + \eta_{t,1} - \eta_{t,2}) & 0 & \frac{\partial a(\eta_{3,t})}{\partial \eta_{3,t}} \end{pmatrix}$$

and the RHS of (4) and (9) augmented to include $a(n_{3,t})$ and $b^2(n_{3,t})$.

3 Bayesian inference

In this section, we consider the problem of performing fully Bayesian inference for the parameters (and unobserved dynamic processes) governing the SIR (LNA) model, based on incidence observations that we assume are incomplete and subject to measurement error. We describe the observation model before considering the inference task. For ease of exposition, we assume a constant infection rate, but note that extension of the methodology to a time varying infection rate is straightforward and considered in Sect. 4.

3.1 Observation model

Without loss of generality, consider data $y = (y_1, \dots, y_T)'$ at integer times, where y_t is a (partial) observation on the cumulative incidence $\Delta N_t = N_t - N_{t-1}$ over a time interval $(t - 1, t]$. Commonly used models for incidence data include additive Gaussian noise (Dureau et al. 2013), the Binomial distribution (Cauchemez and Ferguson 2008) and the Negative Binomial distribution (Lloyd-Smith 2007; Fintzi et al. 2021; Spannaus et al. 2020). The latter two models are typically used under the assumption of under reporting and to capture overdispersion, respectively. In large population settings, they may be well approximated by a Gaussian

distribution which may offer computational benefits when combined with a Gaussian description of the underlying epidemic dynamics, such as the LNA described in Sect. 2.2. These models take the form

$$Y_t | (\Delta N_t = \Delta n_t) \sim N(P' \Delta n_t, \sigma^2), \tag{12}$$

$$Y_t | (\Delta N_t = \Delta n_t) \sim \text{Bin}(P' \Delta n_t, \lambda), \tag{13}$$

$$Y_t | (\Delta N_t = \Delta n_t) \sim \text{NegBin}(\mu = \lambda P' \Delta n_t, \sigma^2 = \mu + \phi \mu^2) \tag{14}$$

for $t = 1, \dots, T$. We assume that Y_t is univariate and P is a constant matrix with $P' = (1, 0)$ corresponding to noisy counts of new infections in a given time window and $P' = (0, 1)$ for removals. We further assume that the observations are independent (given the latent process) and we let $\pi(y_t | \Delta n_t, \psi)$ denote the probability mass function linking y_t and $\Delta n_t = n_t - n_{t-1}$, with ψ denoting the parameters governing the observation model. For example, in (14), $\psi = (\lambda, \phi)'$ with λ controlling the average proportion of cases seen and ϕ is the (inverse) overdispersion parameter.

3.2 Inference task

We assume that interest lies in the vector of all static parameters $\theta = (\beta, \gamma, \psi', x_0')$ including the initial state $x_0 = (s_0, i_0)'$ for convenience, and the latent incidence process at the observation times $n = \{n_t, t = 0, \dots, T\}$.

Upon ascribing a prior density $\pi(\theta)$ to θ , Bayesian inference proceeds via the joint posterior

$$\pi(\theta, n | y) \propto \pi(\theta) \pi(n | x_0, \beta, \gamma) \pi(y | \Delta n, \psi) \tag{15}$$

where

$$\pi(n | x_0, \beta, \gamma) = \prod_{t=1}^T \pi(n_t | n_{t-1}, x_0, \beta, \gamma)$$

and $\pi(n_t | n_{t-1}, x_0, \beta, \gamma)$ is the Gaussian transition density obtained from (7) with the ODEs (4) and (9) integrated over $[t - 1, t]$ with $\eta_{t-1} = n_{t-1}$ and $V_{t-1} = 0_2$. Note that in this case the residual $r_{t-1} = n_{t-1} - \eta_{t-1} = 0$ and subsequently the ODE satisfied by G_t in (8) need not be integrated. Finally,

$$\pi(y | \Delta n, \psi) = \prod_{t=1}^T \pi(y_t | \Delta n_t, \psi).$$

Since the joint posterior in (15) will be intractable, we resort to Monte Carlo methods for generating samples of the parameters and latent dynamic process. A Gibbs sampler provides a natural mechanism for sampling (15), whereby one alternates between draws of $\theta | n, y$ and $n | \theta, y$. However, dependence

between n and θ can lead to poor mixing. For this reason, Fintzi et al. (2021) use a non-centred parameterisation whereby standard Gaussian innovations driving the generative form of the LNA are used as the effective components to be conditioned on. In what follows, we take a different approach by marginalising out the latent process, either via (correlated) pseudo-marginal methods (e.g. Andrieu et al. 2010; Deligiannidis et al. 2018) or by further approximating the observation model in the non Gaussian case.

3.3 Marginalisation of the incidence process

The joint posterior density in (15) can be factorised as

$$\pi(\theta, n | y) = \pi(\theta | y) \pi(n | \theta, y) \tag{16}$$

where

$$\pi(\theta | y) \propto \pi(\theta) \pi(y | \theta). \tag{17}$$

The form of (16) suggests a two step approach to inference whereby samples are first drawn from the marginal parameter posterior $\pi(\theta | y)$ in step 1, and then conditioned on in a second step when drawing samples of the latent process from $\pi(n | \theta, y)$. However, unless the observation model takes the linear Gaussian form of (12), neither the observed data likelihood $\pi(y | \theta)$ in (17) nor the constituent densities in (16) will be tractable, despite the linear Gaussian structure of the LNA. The main focus of this paper is exactly this *intractable* scenario, and we now consider two approaches to address it.

3.3.1 Pseudo-marginal methods

Consider first the intractable observed data likelihood $\pi(y | \theta)$ which can be factorised as

$$\pi(y | \theta) = \pi(y_1 | \theta) \prod_{t=2}^T \pi(y_t | y_{1:t-1}, \theta) \tag{18}$$

where $y_{1:t-1} = (y_1, \dots, y_{t-1})'$. The terms in (18) can be recursively estimated using a particle filter (see e.g. Chopin and Papaspiliopoulos 2020, for an overview) in such a way that realisations of a non-negative unbiased estimator of the full likelihood are obtained. We denote this estimator by

$$\hat{\pi}_U(y | \theta) = \hat{\pi}_{U_1}(y_1 | \theta) \prod_{t=2}^T \hat{\pi}_{U_t}(y_t | y_{1:t-1}, \theta)$$

where the flattened vector $U = (U'_1, \dots, U'_T)'$ $\sim g(u)$ denotes all random variables used in the construction of the estimator. Hence, unbiasedness here means that $E_{U \sim g} \{\hat{\pi}_U(y | \theta)\} = \pi(y | \theta)$. Algorithm 1 gives step $t + 1$ of

the particle filter and can be executed for $t = 0, \dots, T - 1$ upon initialising with particles $\{n_0^{(k)} = 0_2, k = 1, \dots, N\}$. Note that component $t + 1$ of U is partitioned as $(\tilde{U}'_{t+1}, \tilde{U}_{t+1})'$ where the k th value of \tilde{U}_{t+1} is $\tilde{U}_{k,t+1} \sim N(0, I_2)$ and used to propagate particles in step 1(b); $\tilde{U}_{t+1} \sim \text{Unif}(0, 1)$ is used in the systematic resampling step 3. As presented, step 1a has that the ODE system governing the LNA is initialised and integrated for each particle $n_t^{(k)}$. This ‘restarting’ approach avoids potential mismatch between the deterministic process η_t and the latent stochastic process n_t (see e.g. Fearnhead et al. 2014; Minas and Rand 2017, for further discussion). Upon iterating Algorithm 1 over t , the product (over observation times) of the average unnormalised weight gives an unbiased estimator of $\pi(y|\theta)$ (Del Moral 2004) and is key to the construction of a pseudo-marginal scheme that we now describe.

Algorithm 1 Step $t + 1$ of the Particle Filter

Input: Parameter θ , auxiliary variable $u_{t+1} = (\tilde{u}'_{t+1}, \tilde{u}_{t+1})'$, next observation y_{t+1} , N particles $\{n_t^{(k)}, k = 1, \dots, N\}$.

1. Forward propagation. For $k = 1, \dots, N$:
 - (a) Integrate (4) and (9) over $(t, t + 1]$ with initial conditions $\eta_t^{(k)} = n_t^{(k)}$ and $V_t^{(k)} = 0_2$ to give $\eta_{t+1}^{(k)}$ and $V_{t+1}^{(k)}$. Note that (8) need not be integrated.
 - (b) Set $n_{t+1}^{(k)} = \eta_{t+1}^{(k)} + \sqrt{V_{t+1}^{(k)}} \tilde{u}_{k,t+1}$ and $\Delta n_{t+1}^{(k)} = n_{t+1}^{(k)} - n_t^{(k)}$.

2. Compute the weights. For $k = 1, \dots, N$:

$$\tilde{w}_{t+1}^{(k)} = \pi(y_{t+1} | \Delta n_{t+1}^{(k)}, \psi),$$

$$w_{t+1}^{(k)} = \frac{\tilde{w}_{t+1}^{(k)}}{\sum_{j=1}^N \tilde{w}_{t+1}^{(j)}}.$$

3. Resample N particles using systematic resampling with uniform draw \tilde{u}_{t+1} and weights $w_{t+1}^{(k)}, k = 1, \dots, N$.

Output: N particles $\{n_{t+1}^{(k)}, k = 1, \dots, N\}$ to be used in step $t + 1$, an estimate for the current marginal likelihood term $\hat{\pi}_{u_{t+1}}(y_{t+1} | y_{1:t}, \theta) = \frac{1}{N} \sum_{k=1}^N \tilde{w}_{t+1}^{(k)}$.

Pseudo-marginal Metropolis-Hastings methods (PMMH, Andrieu et al. 2009; Andrieu and Roberts 2009) are a class of Metropolis-Hastings (MH) scheme that target the joint density

$$\pi(u, \theta) \propto \pi(\theta)g(u)\hat{\pi}_u(y|\theta)$$

for which it is easily checked that marginalising over U gives the marginal parameter posterior $\pi(\theta|y)$. Hence, an MH scheme with proposal density $q(\theta^*|\theta)g(u^*)$ and acceptance probability

$$\alpha(\{\theta^*, u^*\}|\{\theta, u\}) = \min \left\{ 1, \frac{\pi(\theta^*)\hat{\pi}_{u^*}(y|\theta^*)}{\pi(\theta)\hat{\pi}_u(y|\theta)} \times \frac{q(\theta|\theta^*)}{q(\theta^*|\theta)} \right\}$$

targets the joint density $\pi(u, \theta)$ for with retaining draws of θ gives (dependent) samples from the marginal parameter posterior.

The efficiency of the PMMH scheme can be improved by proposing the auxiliary variable from a g -reversible kernel $f(u^*|u)$ that induces positive correlation between u and u^* , and in turn, $\hat{\pi}_u(y|\theta)$ and $\hat{\pi}_{u^*}(y|\theta^*)$, so that the variance of the acceptance probability is reduced. Suppose that $g(u) = N(u; 0, I_{\dim(u)})$ and note that where necessary, the inverse CDF method can be used to transform the auxiliary variable to be Gaussian (for example, in step 3 of Algorithm 1) where a uniform draw is required). A practical choice of $f(u^*|u)$ is the g -reversible Crank–Nicolson kernel

$$f(u^*|u) = N(u^*; \rho u, (1 - \rho^2) I_{\dim(u)})$$

for which the tuning parameter ρ controls correlation between u and u^* . The resulting correlated PMMH scheme (CPMMH, Dahlin et al. 2015; Deligiannidis et al. 2018) is an MH scheme targeting $\pi(u, \theta)$ with proposal density $q(\theta^*|\theta)f(u^*|u)$ and acceptance probability as above. CPMMH can result in significant gains in computational efficiency over PMMH (see e.g. Golightly et al. 2019, in the context of stochastic kinetic models), provided that the positive correlation between u and u^* induces positive correlation between successive likelihood estimates. To alleviate the issue of resampling in the particle filter potentially eroding this correlation, we follow Choppala et al. (2016) by sorting particles (according to Euclidean distance from the particle with the smallest first component) before propagation.

Finally we note that when interest lies in the posterior for the latent incidence process, samples can be obtained via modification of the (C)PMMH scheme (Andrieu et al. 2010) by drawing a particle path at each iteration of the algorithm. Note that this requires storing the ancestral lineages of the particles in each run of the particle filter.

3.3.2 Analytic method via Gaussian approximation

The LNA, when combined with the linear Gaussian observation model (12), permits analytic calculation of the observed data likelihood $\pi(y|\theta)$ and the conditional posterior $\pi(n|\theta, y)$. Evaluation of the former can be efficiently achieved via a forward filter and draws from the latter via backward sampling. We apply these methods to the Binomial and Negative Binomial observation models in (13) and (14) through suitable Gaussian approximations thereof. For

reasons of brevity, we focus on the Binomial case but note that our approach is easily extended to the Negative Binomial case. Where appropriate, we suppress the parameter vector θ from the notation for simplicity.

To make clear the two approximations to be used in the filtering recursions, consider the LNA written in state-space format over a time interval $(t, t + 1]$, using a Gaussian approximation to the Binomial observation model. We have that

$$N_{t+1}|(N_t = n_t) \sim N(\eta_{t+1} + G_{t+1}(n_t - \eta_t), V_{t+1}), \tag{19}$$

$$Y_{t+1}|(N_{t+1} = n_{t+1}, N_t = n_t) \sim N(\lambda P' \Delta n_{t+1}, \lambda(1 - \lambda)P' \Delta n_{t+1}), \tag{20}$$

where η_{t+1} , G_{t+1} and V_{t+1} are obtained by integrating (4), (8) and (9) over $(t, t + 1]$ with initial conditions of η_t (itself integrated from time 0), I_2 and O_2 . Although (19) is linear in n_t , as noted in Sect. 3.3.1, η_t should be initialised at n_t . ‘Restarting’ the LNA in this way can avoid issues arising from the ODE solution becoming poor over long time intervals (Fearnhead et al. 2014; Minas and Rand 2017). However, we now have that both (19) and (20) involve nonlinear expressions of the latent process. Therefore, to permit the use of standard Kalman-filtering recursions we make further linear approximations. Suppose that the filtering distribution at time t is $N_t|(Y_{1:t} = y_{1:t}) \sim N(a_t, C_t)$. Firstly, we set $\eta_t = a_t$, $V_t = C_t$ and integrate (4) and (9) over $(t, t + 1]$ to obtain η_{t+1} and V_{t+1} . Finally, we replace Δn_{t+1} in the variance of (20) with $\Delta \hat{n}_{t+1} := E(\Delta N_{t+1}) = \eta_{t+1} - a_t$. Note that explicit conditioning of the expectation on $y_{1:t}$ has been suppressed for simplicity. The resulting linear and Gaussian state-space model is

$$N_{t+1}|(N_t = n_t) \sim N(\eta_{t+1}, V_{t+1}), \tag{21}$$

$$Y_{t+1}|(N_{t+1} = n_{t+1}, N_t = n_t) \sim N(\lambda P' \Delta n_{t+1}, \lambda(1 - \lambda)P' \Delta \hat{n}_{t+1}). \tag{22}$$

In the remainder of this section, we derive the filtering recursions based on (21) and (22) to compute the observed data likelihood and conditional posterior of the latent process.

We construct the observed data likelihood contribution $\pi(y_{t+1}|y_{1:t}, \theta)$ as follows. Conditional on $y_{1:t}$, we have that

$$\text{Var}(\Delta N_{t+1}) = V_{t+1} + C_t - C_t G'_{t+1} - G_{t+1} C_t$$

where we have used that $\text{Cov}(N_{t+1}, N_t) = G_{t+1} \text{Var}(N_t)$. Hence, combining with (22) gives

$$\pi(y_{t+1}|y_{1:t}, \theta) = N\left(y_{t+1}; \lambda P' E(\Delta N_{t+1}), \lambda^2 P' \text{Var}(\Delta N_{t+1}) P + \hat{\sigma}^2\right) \tag{23}$$

where $\hat{\sigma}^2 = \lambda(1 - \lambda)P' \Delta \hat{n}_{t+1}$ is the observation variance in (22). To update the filtering distribution, we construct the joint density of N_{t+1} and Y_{t+1} conditional on $Y_{1:t} = y_{1:t}$ as

$$\begin{pmatrix} N_{t+1} \\ Y_{t+1} \end{pmatrix} \sim N \left\{ \begin{pmatrix} \eta_{t+1} \\ \lambda P' E(\Delta N_{t+1}) \end{pmatrix}, \begin{pmatrix} V_{t+1} & \text{Cov}(N_{t+1}, Y_{t+1}) \\ \text{Cov}(Y_{t+1}, N_{t+1}) & \lambda^2 P' \text{Var}(\Delta N_{t+1}) P + \hat{\sigma}^2 \end{pmatrix} \right\}$$

where $\text{Cov}(N_{t+1}, Y_{t+1}) = \lambda(V_{t+1} - G_{t+1} C_t)P$. Hence, conditioning on $Y_{t+1} = y_{t+1}$ gives $N_{t+1}|(Y_{1:t+1} = y_{1:t+1}) \sim N(a_{t+1}, C_{t+1})$ with mean

$$\begin{aligned} a_{t+1} &= \eta_{t+1} + \text{Cov}(N_{t+1}, Y_{t+1}) \\ &\quad \times (\lambda^2 P' \text{Var}(\Delta N_{t+1}) P + \hat{\sigma}^2)^{-1} \\ &\quad \times (y_{t+1} - \lambda P' E(\Delta N_{t+1})) \end{aligned} \tag{24}$$

and variance

$$\begin{aligned} C_{t+1} &= V_{t+1} - \text{Cov}(N_{t+1}, Y_{t+1}) \\ &\quad \times (\lambda^2 P' \text{Var}(\Delta N_{t+1}) P + \hat{\sigma}^2)^{-1} \\ &\quad \times \text{Cov}(Y_{t+1}, N_{t+1}). \end{aligned} \tag{25}$$

Calculation of (23), (24) and (25) constitutes a single step of the forward filter; see Algorithm 2, which can be iterated over t to give an evaluation of the observed data likelihood (under the LNA), $\pi(y|\theta)$. Hence, draws from the marginal parameter posterior $\pi(\theta|y)$, with the LNA as the inferential model, are obtained in a straightforward manner via Metropolis-Hastings e.g. random walk Metropolis (RWM).

Algorithm 2 Step $t + 1$ of the LNA Forward Filter

Input: Parameter θ ; a_t and C_t , the initial conditions of (4) and (9); $\pi(y_{1:t}|\theta)$, the current observed data likelihood; y_{t+1} , the next observation.

1. Prior at $t + 1$. Initialise the LNA with $\eta_t = a_t$, $G_t = I_2$ and $V_t = C_t$. Integrate (4), (8) and (9) forward to $t + 1$ to obtain η_{t+1} , G_{t+1} and V_{t+1} . Thus

$$N_{t+1}|(Y_{1:t} = y_{1:t}) \sim N(\eta_{t+1}, V_{t+1}).$$

2. Likelihood update. Compute

$$\pi(y_{1:t+1}|\theta) = \pi(y_{1:t}|\theta)\pi(y_{t+1}|y_{1:t}, \theta)$$

where $\pi(y_{t+1}|y_{1:t}, \theta)$ is given by (23).

3. Posterior at $t + 1$. Combining the distributions of N_{t+1} and Y_{t+1} (given $y_{1:t}$) and then conditioning on y_{t+1} gives $N_{t+1}|(Y_{1:t+1} = y_{1:t+1}) \sim N(a_{t+1}, C_{t+1})$ where a_{t+1} and C_{t+1} are given by (24) and (25).

Output: $\pi(y_{1:t+1}|\theta)$, a_{t+1} and C_{t+1} .

It remains that we can use the LNA to generate draws from $\pi(n|\theta, y)$. Note the factorisation

$$\pi(n|\theta, y) = \prod_{t=1}^{T-1} \pi(n_t|n_{t+1}, y_{1:t}, \theta)$$

where each constituent term is a Gaussian density. Under the LNA, the joint density of N_t and N_{t+1} conditional on $Y_{1:t} = y_{1:t}$ is

$$\begin{pmatrix} N_t \\ N_{t+1} \end{pmatrix} \sim N \left\{ \begin{pmatrix} a_t \\ \eta_{t+1} \end{pmatrix}, \begin{pmatrix} C_t & C_t G'_{t+1} \\ G_{t+1} C_t & V_{t+1} \end{pmatrix} \right\}.$$

Conditioning on $N_{t+1} = n_{t+1}$ gives $N_t | (N_{t+1} = n_{t+1}, Y_{1:t} = y_{1:t}, \theta) \sim N(\tilde{a}_t, \tilde{C}_t)$ with mean and variance

$$\begin{aligned} \tilde{a}_t &= a_t + C_t G'_{t+1} V_{t+1}^{-1} (n_{t+1} - \eta_{t+1}), \\ \tilde{C}_t &= C_t - C_t G'_{t+1} V_{t+1}^{-1} G_{t+1} C_t. \end{aligned}$$

Hence, the components of the cumulative incidence n can be drawn via backward sampling for $t = T, T - 1, \dots, 1$, given storage of the LNA ODE output and filtering mean/variance from the forward filter. Then, the latent process x can be constructed deterministically from n and the initial values x_0 using (1).

4 Applications

We consider two applications of the methodology described in Sect. 3. Firstly, using synthetic data generated from the SIR model, we compare the performance of the two marginalisation techniques described in Sects. 3.3.1 and 3.3.2; these are the correlated pseudo-marginal Metropolis-Hastings (CPMMH) and forward filtering Metropolis-Hastings (henceforth FFMH) based inference schemes. We compare the accuracy of posterior output from these schemes with inferences obtained by assuming the most natural Markov jump process as the inferential model. We fit this model using the pseudo-marginal Metropolis-Hastings (PMMH) scheme described in Golightly and Wilkinson (2011). Additionally, we include inferences based on a deterministic ODE model of latent incidence (fit via MH). In the second application, we use FFMH to fit SIR and SIRS models with two different choices of observation model, to a real data set consisting of pest removals from trees in a London park.

All algorithms are coded in R and were run on a desktop computer with an Intel quad-core CPU. Source code is available at <https://github.com/AndyGolightly/LNA-Incidence>. All schemes use random walk proposals with Gaussian innovations for the log-transformed parameters. For CPMMH, we fixed $\rho = 0.99$, which we found to give a good balance

between mixing over the auxiliary variable and parameter chains. We chose the number of particles N by following the practical advice of Deligiannidis et al. (2018). That is, we choose N so that the variance of $\log \hat{\pi}_{u^*}(y|\theta^*) - \log \hat{\pi}_u(y|\theta) \approx 1$. For CPMMH and FFMH, we took the random walk innovation variance to be $\widehat{\text{Var}}(\log \theta|y)$ estimated from a pilot run, and subsequently scaled to meet a desired empirical acceptance rate (see e.g. Schmon et al. (2021) for (C)PMMH and Schmon and Gagnon (2022) for Metropolis-Hastings). CPMMH was run for 50,000 iterations and the remaining schemes were run for 10,000 iterations, which we found gave reasonable mixing efficiency as measured by, for example, effective sample size (see e.g. Plummer et al. 2006).

4.1 Simulation study

We generated three synthetic data sets (denoted $\mathcal{D}_i, i = 1, 2, 3$) from the SIR model, each consisting of the number of new infections in time intervals $(t, t + 10]$ for $t = 10, 20, \dots, 70$. For \mathcal{D}_1 , we used $x_0 = (119, 1)'$ and $(\beta, \gamma)' = (0.00091, 0.082)'$; these choices are consistent with inferences from the well studied Abakaliki small pox data (see e.g. Bailey 1975). For \mathcal{D}_2 , we constructed a larger outbreak by scaling the total population size N_{pop} and removal rate by a factor of 3, resulting in $x_0 = (359, 1)'$ and $(\beta, \gamma)' = (0.00091, 0.246)'$. For \mathcal{D}_3 , we scaled N_{pop} by a factor of 10 (compared to \mathcal{D}_1) and set $x_0 = (1180, 20)'$. We scaled both the infection and removal rates to give $(\beta, \gamma)' = (0.00018, 0.164)'$. Note that all data sets have the same basic reproduction number, $\mathcal{R}_0 = N_{\text{pop}}\beta/\gamma = 1.33$. We corrupted the resulting incidences via the Binomial observation model (13) with $\lambda = 0.8$ in each case. The data sets are shown in Fig. 2 alongside the underlying traces of S_t and I_t (assumed unobserved).

When analysing \mathcal{D}_1 , we adopted an independent prior specification with Gamma and Uniform components by taking $\beta \sim \text{Gamma}(10, 10^4)$, $\gamma \sim \text{Gamma}(10, 10^2)$ and $\lambda \sim \text{Unif}(0, 1)$. The first two choices have been used by Fearnhead and Meligkotsidou (2004) and many others when analysing the Abakaliki small pox data. For \mathcal{D}_2 and \mathcal{D}_3 we adopted a more diffuse prior for the removal rate to better reflect the increase in the ground truth value; specifically, $\gamma \sim \text{Gamma}(10, 30)$, with the prior specification for the remaining components as for \mathcal{D}_1 . We assume that the initial state x_0 is fixed and known but note that inference for x_0 is possible by augmenting θ to include the components of x_0 .

Table 1 and Figs. 3 and 4 summarise the posterior output from each scheme. For data set \mathcal{D}_1 (population size $N_{\text{pop}} = 120$), although all inferential models give posterior output that is consistent with the ground truth parameter values, there are noticeable inconsistencies between LNA- and MJP-based inferences. Using the LNA to model the latent incidence process but with the correct observation model

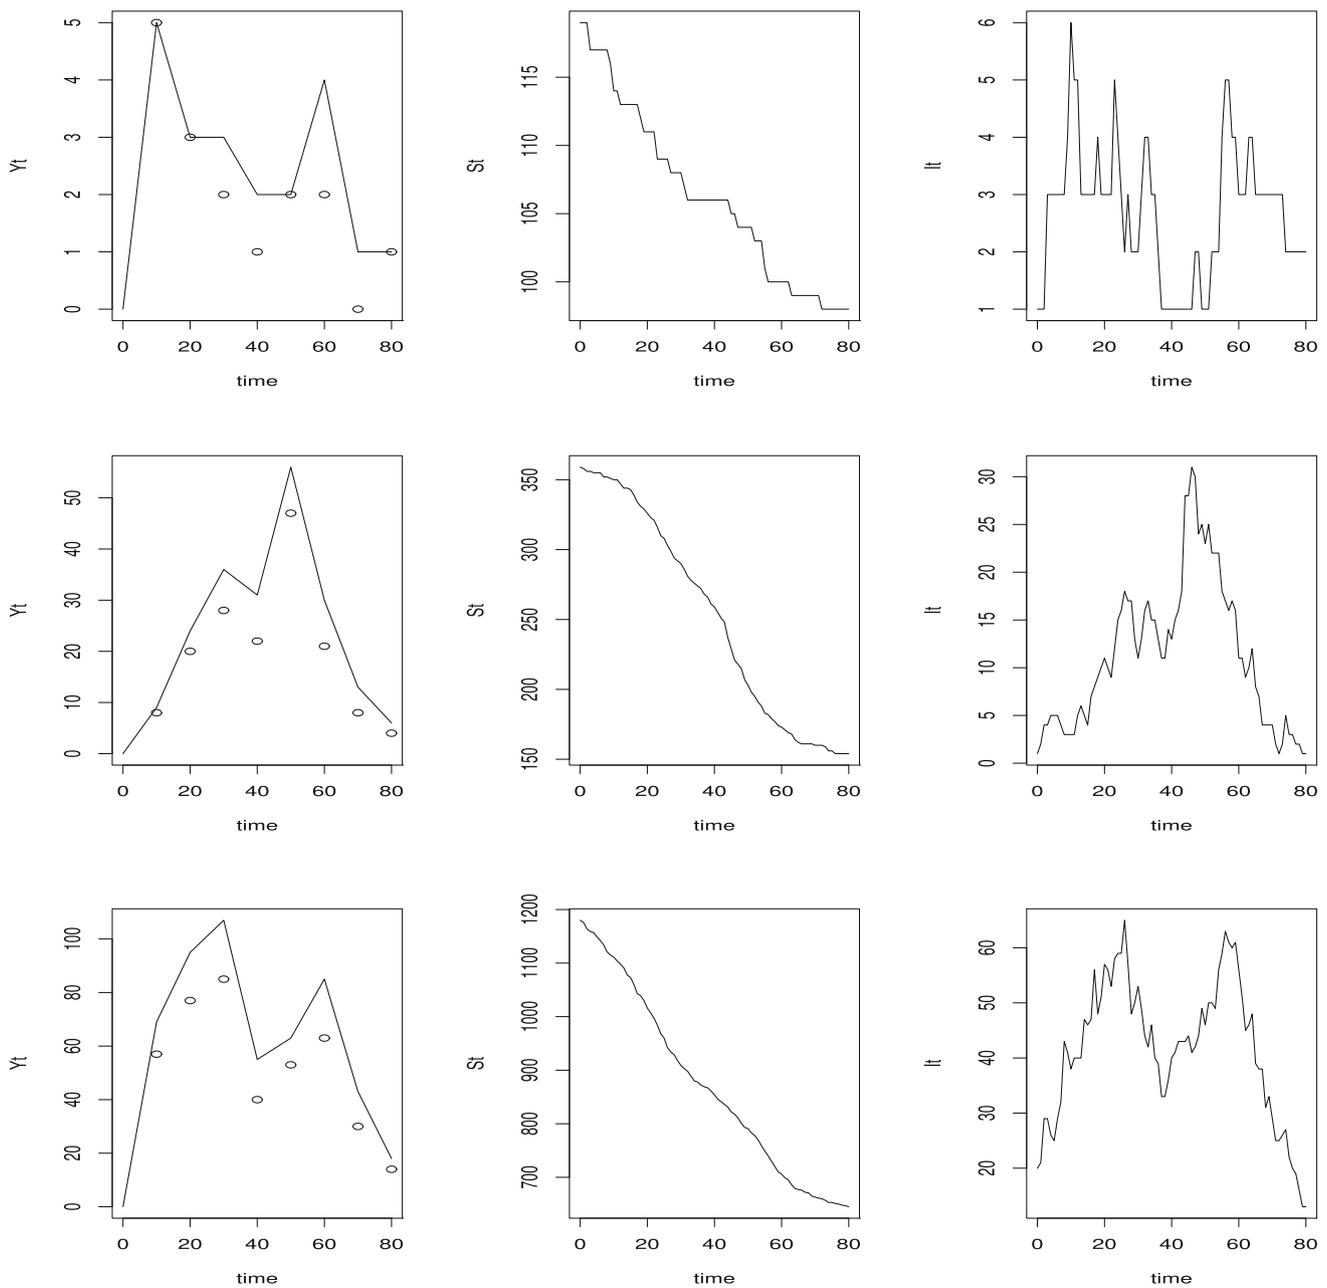


Fig. 2 Synthetic data sets \mathcal{D}_1 (top panel), \mathcal{D}_2 (middle panel) and \mathcal{D}_3 (bottom panel). Left: noisy numbers of new infecteds in a 10 day interval (circles) and latent values (line). Middle and right: corresponding susceptible and infected states

(LNA/(C)PMMH) results in an overestimation of the infection and removal rates, although there is little difference between this approach compared to the MJP when considering the basic reproduction number \mathcal{R}_0 . However, differences are more pronounced when further approximating the observation model as Gaussian (LNA / FFMH), which results in under estimation of \mathcal{R}_0 . Nevertheless, these differences are relatively small, and the advantages (in terms of overall efficiency) of analytically integrating out the latent process (as per LNA/FFMH) are clear.

We measure overall efficiency using minimum (over each parameter chain) effective sample size (ESS) per second (mESS/s). Given the small population size for \mathcal{D}_1 , using the MJP inside a PMMH scheme is computationally more efficient than using the LNA (which, as implemented, requires numerical integration of 5 coupled ODEs per particle per iteration). Correlating successive likelihood estimates (LNA/CPMMH vs PMMH) increases overall efficiency by a factor of 2, however, the largest gains in overall efficiency are obtained by LNA/FFMH, which improves on MJP/PMMH

Table 1 Synthetic data application

Model/scheme	ρ	N	mESS/s	Mean (Std. Dev.)			
				β	γ	λ	\mathcal{R}_0
Data set \mathcal{D}_1							
				0.00091	0.082	0.8	1.33
MJP/PMMH	0.00	30	0.682	0.00091 (0.00024)	0.088 (0.022)	0.62 (0.22)	1.29 (0.35)
LNA/PMMH	0.00	25	0.039	0.00107 (0.00027)	0.094 (0.025)	0.61 (0.24)	1.37 (0.41)
LNA/CPMMH	0.99	15	0.064	0.00111 (0.00024)	0.101 (0.027)	0.64 (0.22)	1.36 (0.40)
LNA/FFMH	–	–	3.075	0.00102 (0.00021)	0.117 (0.029)	0.82 (0.15)	1.09 (0.26)
ODE/MH	–	–	12.091	0.00203 (0.00024)	0.192 (0.042)	0.48 (0.16)	1.35 (0.42)
Data set \mathcal{D}_2							
				0.00091	0.246	0.8	1.33
MJP/PMMH	0.00	125	0.091	0.00087 (0.00016)	0.225 (0.041)	0.75 (0.12)	1.35 (0.19)
LNA/PMMH	0.00	120	0.018	0.00092 (0.00017)	0.231 (0.050)	0.77 (0.12)	1.51 (0.23)
LNA/CPMMH	0.99	60	0.030	0.00089 (0.00016)	0.228 (0.056)	0.77 (0.12)	1.44 (0.25)
LNA/FFMH	–	–	3.680	0.00094 (0.00021)	0.234 (0.060)	0.78 (0.12)	1.49 (0.22)
ODE/MH	–	–	13.440	0.00087 (0.00005)	0.188 (0.021)	0.65 (0.04)	1.68 (0.10)
Data set \mathcal{D}_3							
				0.00018	0.164	0.8	1.33
MJP/PMMH	0.00	120	0.045	0.00032 (0.00006)	0.290 (0.059)	0.82 (0.09)	1.33 (0.10)
LNA/PMMH	0.00	120	0.012	0.00034 (0.00007)	0.318 (0.070)	0.83 (0.10)	1.30 (0.11)
LNA/CPMMH	0.99	60	0.019	0.00034 (0.00006)	0.308 (0.062)	0.81 (0.12)	1.32 (0.12)
LNA/FFMH	–	–	3.533	0.00023 (0.00005)	0.217 (0.049)	0.85 (0.10)	1.29 (0.11)
ODE/MH	–	–	12.984	0.00019 (0.00001)	0.159 (0.008)	0.63 (0.01)	1.45 (0.04)

Inferential model/scheme, correlation parameter, number of particles, minimum effective sample size per second and marginal parameter posterior summaries. The ground truth parameter values are indicated for each data set

by a factor of 4 and on LNA/CPMMH by a factor of almost 50.

For data set \mathcal{D}_2 (population size $N_{\text{pop}} = 360$), the pseudo-marginal schemes require more particles, due to the intrinsic stochasticity of realisations of the latent process generated inside the particle filter, which is large compared to observation noise. The increased population size (and corresponding parameter values that generated the data) leads to many more reaction occurrences between observation instants (compared to \mathcal{D}_1) reducing the relative efficiency of MJP/PMMH versus LNA/(C)PMMH and LNA/FFMH. We note that for this data set, using the LNA to model the latent process and additionally taking a linear Gaussian approximation of the observation model leads to an inference scheme that is both efficient (with an mESS/s 40 times larger than that of the next best performing scheme) and accurate (see Fig. 3, bottom panel).

The magnitude of typical observations in data set \mathcal{D}_3 (population size $N_{\text{pop}} = 1200$) is broadly consistent with that of \mathcal{D}_2 and we find that the pseudo-marginal schemes require similar particle numbers. Using the LNA with the correct observation model gives parameter inferences that are consistent with the MJP. Using LNA/FFMH appears

to result in underestimates of the infection and removal rates although the basic reproduction number appears to be accurately estimated. In terms of overall efficiency, the advantage of LNA/FFMH over competing schemes is clear, with an mESS/s that is approximately 80 times larger than MJP/PMMH. There is relatively little difference between the performance of LNA/CPMMH and MJP/PMMH.

Table 1 also includes summarised posterior output when using a deterministic ODE model of latent incidence, combined with a Gaussian approximation to the Binomial observation model (ODE/MH). This approach requires only the solution of the ODE system in (4), as opposed to (4), (8) and (9) when using the LNA. Consequently, an approximate 4-fold increase in overall efficiency is achieved for ODE/MH compared to LNA/FFMH. However, ignoring intrinsic stochasticity leads to a clear loss of inferential accuracy. In particular, the reporting rate is underestimated (which is unsurprising, as this leads to a larger observation variance, which can somewhat offset the inability of the latent ODE model to capture intrinsic stochasticity) and the basic reproduction number is typically overestimated. As N_{pop} increases, posterior uncertainty for all static parameters is

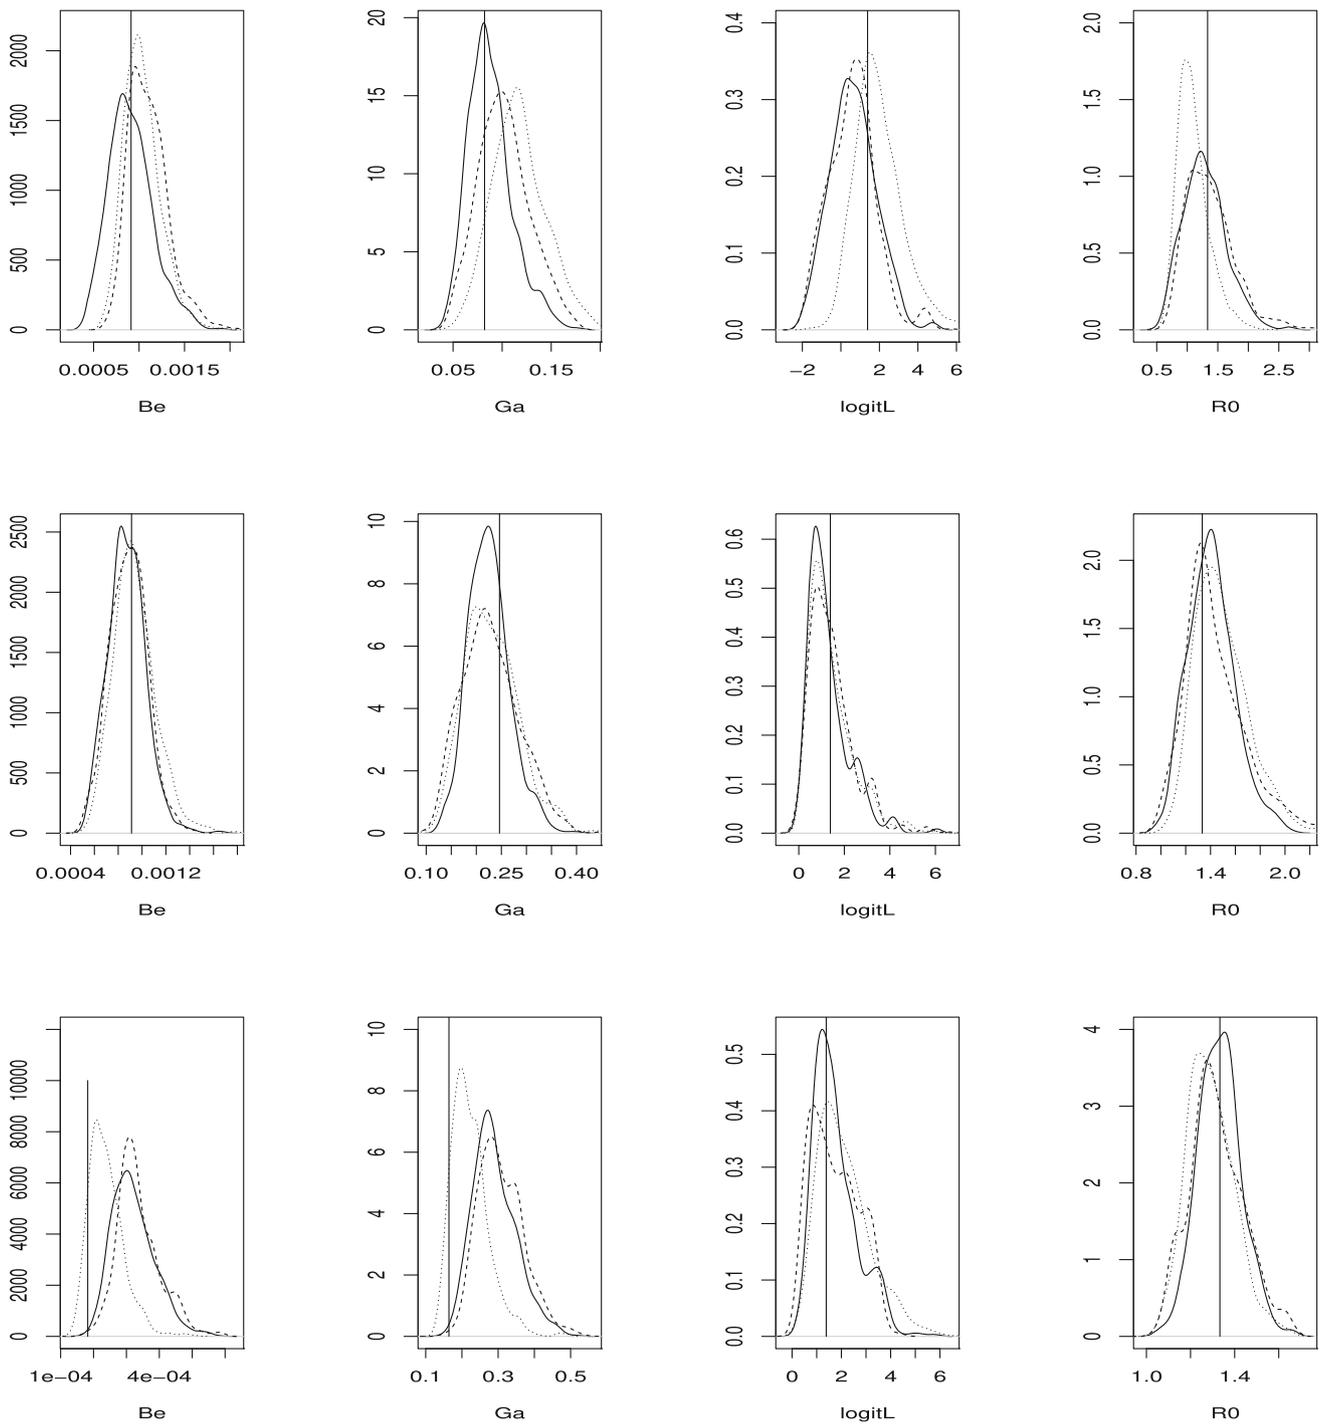


Fig. 3 Synthetic data application. Marginal posterior densities based on \mathcal{D}_1 (top panel), \mathcal{D}_2 (middle panel) and \mathcal{D}_3 (bottom panel), and using the output of MJP/PMMH (solid line), LNA/CPMMH (dashed line), LNA/FFMH (dotted line)

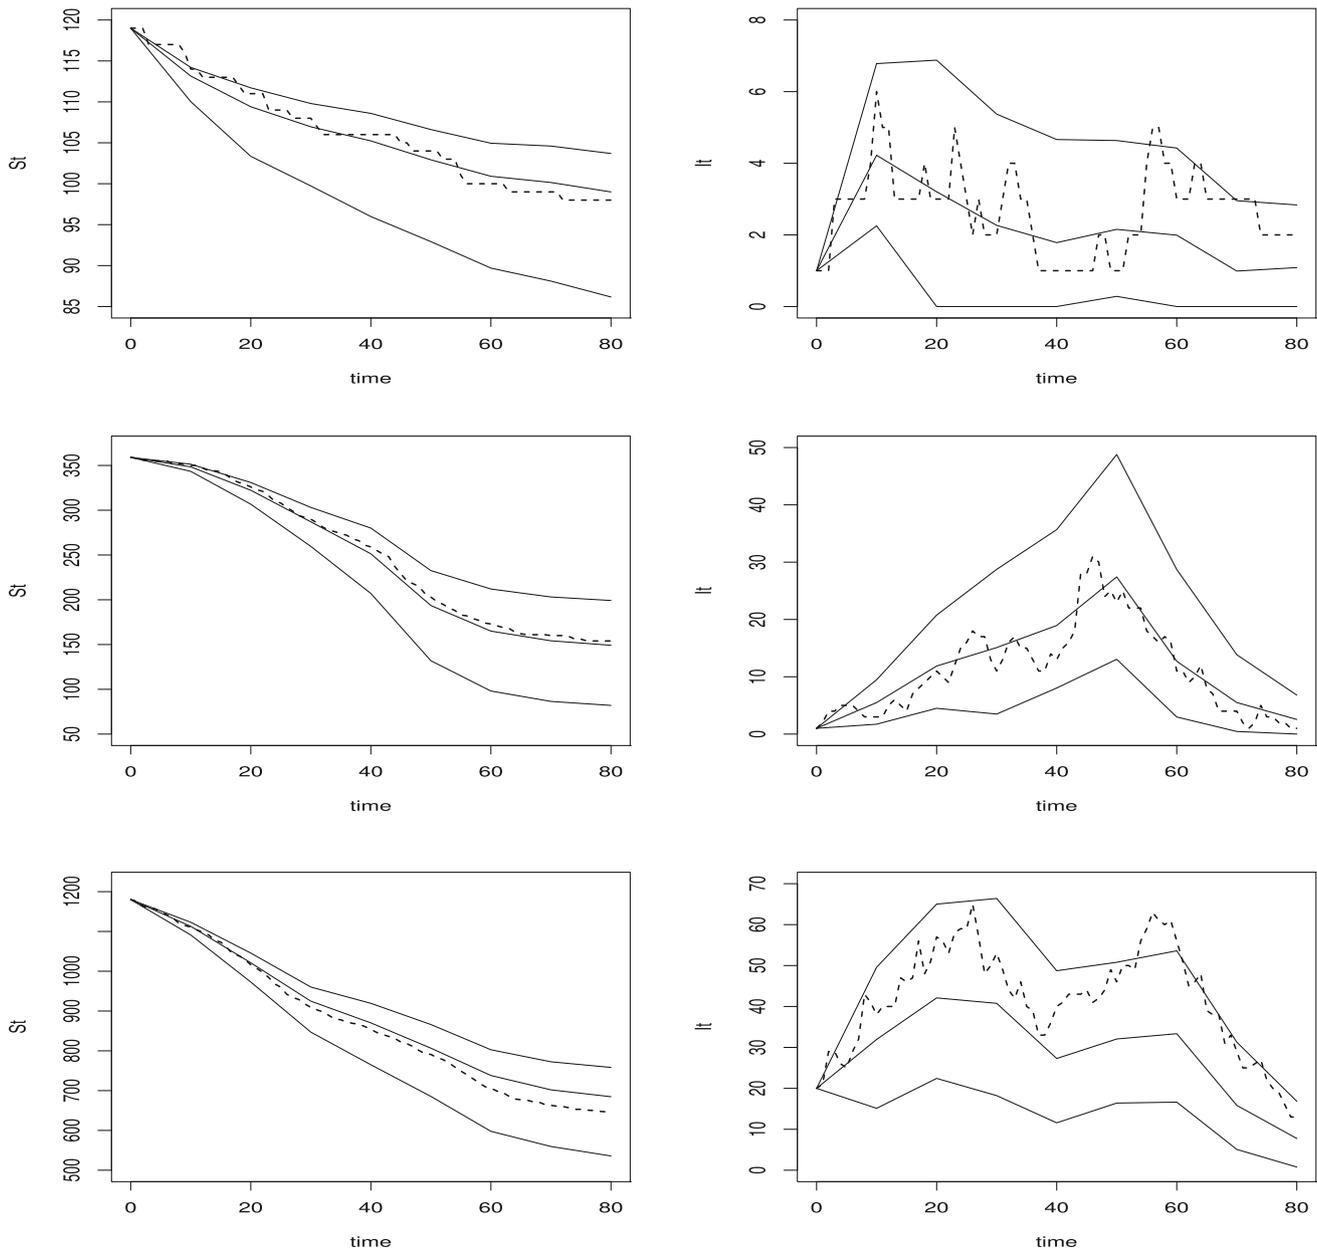


Fig. 4 LNA within-sample predictive distributions (mean and 95% credible intervals) for S_t (left) and I_t (right) based on synthetic data sets \mathcal{D}_1 (top panel), \mathcal{D}_2 (middle panel) and \mathcal{D}_3 (bottom panel). Ground truth latent process trajectories are shown as dashed lines

underestimated (relative to the gold standard MJP approach) irrespective of each data set.

Figure 4 shows within-sample predictive summaries (averaged over parameter uncertainty) for the susceptible and infective states under the LNA. That is, for each of n_{iters} parameter draws from the marginal posterior under the LNA, backward sampling (see Sect. 3.3.2) was run to obtain samples $(S_t^{(k)}, I_t^{(k)})'$ from the conditional posterior $\pi(x_t|\theta^{(k)}, \mathcal{D})$ for $t = 10, 20, \dots, 80$ and $k = 1, 2, \dots, n_{iters}$. We see that although the LNA (unsurprisingly) fails to capture the discrete nature of the infective process (most evident for data set

\mathcal{D}_1), within-sample predictive draws are generally consistent with the ground truth traces generated by the jump process, although there is some suggestion for data sets \mathcal{D}_1 and \mathcal{D}_3 that the LNA is least accurate at the end of the epidemic.

4.2 Oak processionary moth in Richmond park, London

In this section we consider the application of the methodology to the infestation of the oak processionary moth (OPM), *Thaumetopoea processionea*, in Richmond Park, London.

Table 2 OPM data

Year	2013	2014	2015	2016	2017	2018	2019	2020
No. removals	1024	1414	958	540	594	557	587	1029

Number of “removed trees” in a given year, Richmond park, London, 2013–2020

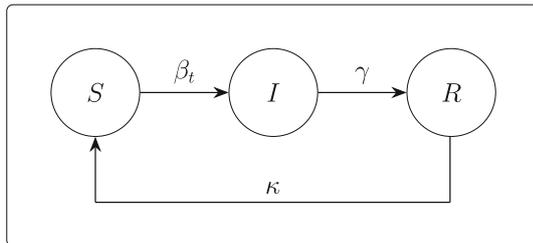


Fig. 5 SIRS compartment model

OPM is an invasive pest, destructive to oak trees and toxic to humans and animals (Maier et al. 2003, 2004; Gottschling and Meyer 2006; Rahlenbeck and Utikal 2015). The moth was first established in the UK in 2006 and despite efforts to initially eradicate, and then contain the infestation, OPM has continued to spread (Suprunenko et al. 2021; Wadkin et al. 2022).

Surveys and control strategies for Richmond Park are carried out by The Royal Parks charity, and this data is then shared with the governmental Oak Processionary Moth Control Programme (Mainprize and Straw 2021). The data records the numbers of OPM nests removed from trees (with recorded locations) between the years 2013 and 2020, allowing the formation of a time series for the yearly removal incidence of infested trees; see Table 2. The removal prevalence of the same set of trees (constructed under the assumption of known initial conditions) was considered in an SIR model in Wadkin et al. (2022). However, upon the manual removal of the OPM nests, it is possible for the trees to become susceptible to re-infestation, and thus we additionally consider the SIRS model below.

4.2.1 Model and prior distribution

To allow for removed trees re-entering the susceptible class, we consider the SIRS compartment model shown graphically in Fig. 5.

Transitions between compartments can be described by the set of pseudo-reactions given by

$$S + I \xrightarrow{\beta_t} 2I, \quad I \xrightarrow{\gamma} R, \quad R \xrightarrow{\kappa} S$$

where β_t is a time varying infection rate whose natural logarithm is described by the SDE

$$d \log \beta_t = \sigma_\beta dW_t.$$

That is, the log infection rate is a scaled Brownian motion process. We note that setting $\kappa = 0$ gives the SIR model and in what follows fit both SIR and SIRS models under the assumption the latent incidence process is well described by the linear noise approximation. Further details of the LNA for the SIRS compartment model are given in Appendix A. We additionally consider two observation models; these are the Binomial and Negative Binomial models given by (13) and (14). This leads to 4 competing models which we compare using the deviance information criterion DIC, (see e.g. Gibson et al. 2018, for a discussion of DIC in the epidemic context) given by

$$DIC = -2E_\theta\{\log \pi(y|\theta)|y\} + p_D$$

where $p_D = -2E_\theta\{\log \pi(y|\theta)|y\} + 2 \log \pi(y|\bar{\theta})$ measures the effective number of parameters in the model. Note that the observed data likelihood $\pi(y|\theta)$ is tractable under the Gaussian approximation approach to inference described in Sect. 3.3.2, which we employ here. Hence, DIC is easily calculated and the model with the smallest DIC is preferred.

We follow Wadkin et al. (2022) by fixing $N_{pop} = 40,000$, $x_0 = (38,600, 1400)'$ and $\log \beta_0 = -10$. We adopt an independent prior specification by taking $\log \gamma \sim N(0, 0.5^2)$, $\log \kappa \sim N(0, 1)$, $\log \sigma_\beta \sim N(1, 1)$, $\lambda \sim \text{Unif}(0, 1)$ and, when using a Negative Binomial observation model, $\log \phi \sim N(0, 1)$. Note that the choice of β_0 and prior for the removal rate γ induces a prior on the basic reproduction number $\mathcal{R}_0 = N_{pop}\beta_0/\gamma$ at time 0 as lognormal $\text{logN}(0.6, 0.5^2)$. This gives a 95% equitailed credible interval of (0.7, 4.8) for \mathcal{R}_0 , which reflects our belief that OPM spread is likely to persist, without precluding $\mathcal{R}_0 < 1$. We note also that the prior for κ gives a 95% credible interval of (0.14, 7.10) years, reflecting vague prior beliefs on the time taken for a removed tree to re-enter the susceptible class.

4.2.2 Results

We ran the marginal Metropolis–Hastings scheme (as described in Sect. 3.3.2) for 50,000 iterations, with the resulting parameter chains suggesting adequate mixing. Tables 3 and 4 and Figs. 6, 7 and 8 summarise the posterior output under each competing model.

Table 3 shows estimated DIC for the SIR and SIRS models, under the assumption of either a Binomial or Negative Binomial observation model. A Binomial observation model is preferred irrespective of the assumed underlying compart-

Table 3 OPM data application

Model	SIR (Bin)	SIR (Neg Bin)	SIRS (Bin)	SIRS (Neg Bin)
DIC	115.4	120.4	113.4	119.8

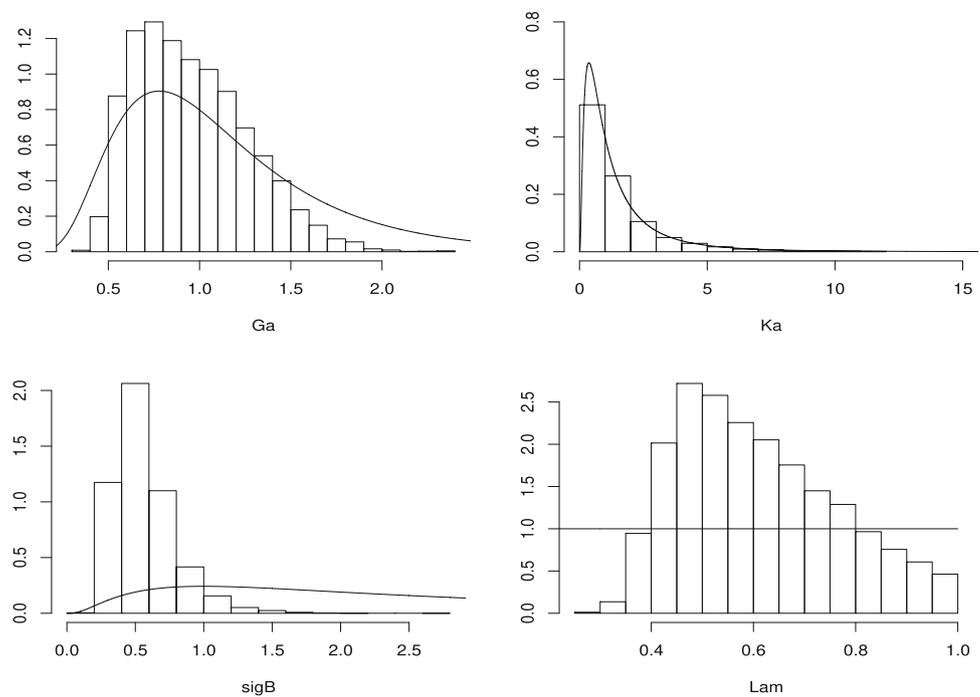
Estimated DIC for the SIR and SIRS models, assuming either a Binomial (Bin) or Negative Binomial (Neg Bin) observation model

Table 4 OPM data application

Parameter	Mean (Standard Deviation)			
	SIR (Bin)	SIR (Neg Bin)	SIRS (Bin)	SIRS (Neg Bin)
γ	0.90 (0.30)	1.44 (0.44)	0.96 (0.31)	1.45 (0.47)
κ	–	–	1.57 (2.02)	1.26 (1.72)
σ_β	0.64 (0.24)	0.58 (0.58)	0.56 (0.22)	0.51 (0.39)
λ	0.64 (0.16)	0.57 (0.20)	0.61 (0.16)	0.62 (0.20)
ϕ	–	0.30 (0.38)	–	0.25 (0.32)

Marginal parameter posterior summaries

Fig. 6 OPM data application. Marginal posterior densities (histograms) and prior (solid line), of the parameters in the SIRS model assuming binomial observations



ment model. This is consistent with the inferred values of the (inverse) dispersion parameter ϕ , which are typically small; see Table 4. Hence, it appears that the true but unobserved removal incidence is much larger than the observed incidence, which is unsurprising given surveyed areas of Richmond park in each year, which typically constitute a small fraction of the total area. Although our findings support the hypothesis that trees can become susceptible to re-infestation over the time scales of the data set considered, we note that the analysis has not been particularly informative for the parameter κ , governing the rate of R to S transitions; see Fig. 6 showing marginal posterior densities for parameters in the SIRS model and the prior specification. Since removed trees are treated with insecticide, this param-

eter is likely to be of interest to practitioners. Nevertheless, improved data collection protocols and a longer study period may provide a partial record of the number of R to I transitions, which would greatly improve inferences on κ .

Figure 7 summarises the within-sample predictive distributions for the susceptible and infective prevalence processes (which are easily reconstructed from the predicted incidences, not shown) and the log infection process, $\log \beta_t$. These results are broadly consistent with those of Wadkin et al. (2022), which suggest a plausibly constant infection rate and an uptick in infected trees from 2018. Figure 8 summarises the marginal posterior distribution of the basic reproduction number \mathcal{R}_0 against year. Sampled values of \mathcal{R}_0 appear to be largely consistent across years, however,

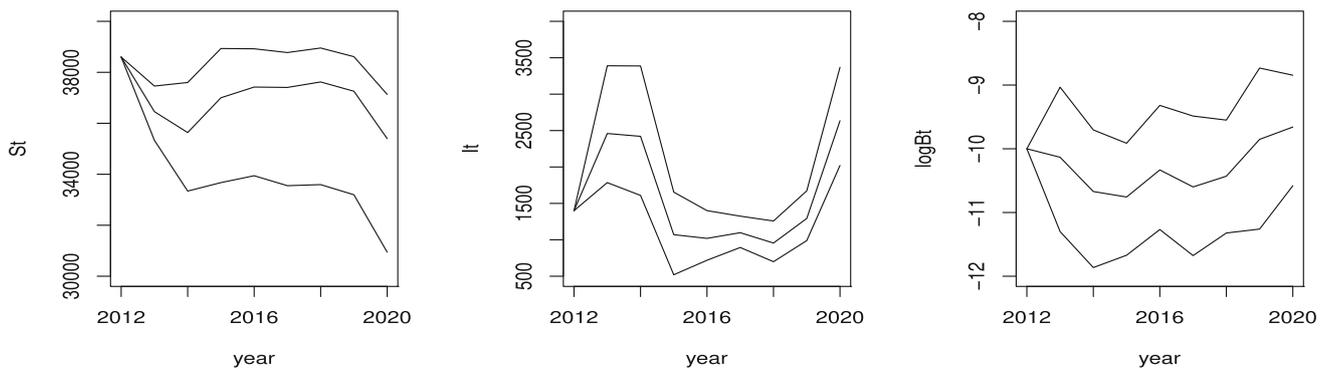
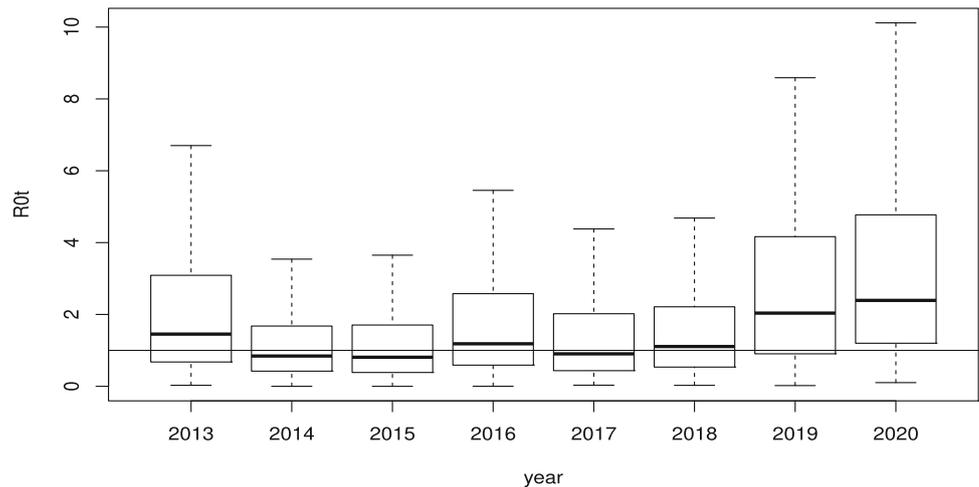


Fig. 7 OPM data application. Within-sample predictive distributions (mean and 95% credible intervals) for S_t (left) and I_t (middle) and $\log \beta_t$ (right)

Fig. 8 OPM data application. Boxplots summarising the marginal posterior distribution of the basic reproduction number \mathcal{R}_0 against year



the marginal posterior distributions in years 2018, 2019 and 2020 have greatest support for $\mathcal{R}_0 > 1$ suggesting that OPM will continue to propagate in Richmond park.

5 Discussion

The construction of efficient and widely applicable approaches to inference for stochastic epidemic models remains a key challenge (Swallow et al. 2022.) In this paper, we have proposed a fast and efficient method for inferring the parameters governing the linear noise approximation (LNA) of a stochastic epidemic model, using incidence data consisting of the cumulative number of new infections (or removals) in fixed-length windows. This setting is considered in Fintzi et al. (2021) who combine the LNA of the incidence process with a Negative Binomial observation model and develop an efficient MCMC scheme targeting the joint density of the parameters and latent incidence process. Our contribution, on the other hand, is a framework for marginalising out the latent incidence process: either by exactly targeting the marginal parameter posterior via a (cor-

related) pseudo-marginal method, or analytically through a Gaussian approximation of the observation process. We additionally allow for a flexible, time-varying stochastic infection rate, which is naturally handled within the LNA framework. Our experiments demonstrated that use of the LNA and a further Gaussian approximation of an observation model can be both accurate and efficient. Using parameter values inspired by the Abakaliki smallpox outbreak, we investigated the accuracy and efficiency of the analytically marginalised LNA as the population size increases (and with the parameters scaled appropriately). In the ‘large epidemic’ setting ($N_{pop} = 1200$), the analytic marginalisation scheme outperforming the next best performing scheme by about a factor of 80. In this scenario, use of the most natural Markov jump process representation of the epidemic is computationally prohibitive.

We further illustrated our approach via an application with real data consisting of numbers of trees infested with oak processionary moth (OPM) nests in Richmond Park, London. Typical observations consist of around 500–1500 removals in a given year, with a total population size of around 40,000 trees, thus necessitating the efficient infer-

ence methods developed here. As well as inferring key quantities of interest, such as the basic reproduction number and latent susceptible and infective trajectories, our approach allows for easy computation of the observed data likelihood which can be used, for example, to compute a deviance information criterion (DIC). We used DIC to compare two different compartment models (SIR versus SIRS) and two different observation models (Binomial versus Negative Binomial). Our analysis suggests the SIRS model as the best fitting compartmentment model, suggesting that trees can re-enter the susceptible class following removal (via treatment). Although improved data collection protocols which include observation of the number of removed trees which subsequently become infected will greatly improve predictive power, our approach demonstrates that meaningful conclusions on the spread of OPM can be drawn, despite a data poor scenario.

5.1 Limitations and extensions

Within the stochastic kinetic models context, the LNA can be derived directly from the most natural Markov jump process (MJP) representation (Kurtz 1970, 1972) but is perhaps most intuitively viewed as a tractable Gaussian process approximation of the Itô Stochastic differential equation (SDE) that best matches the MJP representation Ferm et al. (2008); Fearnhead et al. (2014). As advocated by Fuchs (2013) among others, judging the validity of these continuous-valued approximations should involve comparison with the MJP (e.g. via simulations) for the specific system considered. Nevertheless, we expect, in general, that the best matching SDE and LNA approaches are likely to approximate the MJP particularly poorly when specie numbers are comparatively small (e.g. in the few tens). In such situations, we envisage that our approach is likely to be of most practical benefit in providing initial values and tuning choices for simulation-based inference schemes that target the posterior under the MJP. For inherently multi-scale epidemics, it may be possible to leverage hybrid simulation techniques (see e.g. Sherlock et al. 2014) whereby the LNA is used to model species which frequently change state, coupled with a discrete stochastic updating procedure for species which change state less often.

This work can be further extended in a number of ways. Of particular interest to us is the use of the proposed approach within a spatio-temporal setting, and with application to OPM spread, for example, by allowing importation of pests from nearby locations. Extension of the methodology to allow incorporation of multiple data streams (see e.g. Corbella et al. 2022) also merits further attention.

Acknowledgements We thank Julia Branson from the University of Southampton and Andrew Hoppit from Forestry Commission England for useful discussions regarding the OPM application. This research was supported by: EPSRC New Horizons Grant EP/V048511/1 (AB,

AG, NP, and LW) and NERC Knowledge Exchange Fellows Grant NE/X000478/1 (LW).

Author Contributions Conceptualisation, AG, LW, AB and NP; methodology, AG, LW, SW, AB, NP and TK; software, AG, LW, SW; validation, AG and LW; formal analysis, AG and LW; writing original draft preparation, AG and LW; writing review and editing, AG, LW, SW, AB, NP, TK visualisation, AG; funding acquisition, AG, LW, AB, NP All authors reviewed the manuscript.

Declarations

Conflict of interest The authors declare no competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix A: SIRS model details

Let $X_t = (S_t, I_t)'$ denote the numbers of susceptibles and infectives at time t . Similarly, let $n_t = (n_{1,t}, n_{2,t}, n_{3,t})'$ denote the cumulative number of infection, removal and loss of immunity (that is, removal to susceptible) events at time t . Let β_i, γ and κ denote the corresponding event rates. The cumulative incidence $\{N_t, t \geq 0\}$ is an MJP governed by the transition probabilities

$$\begin{aligned} \mathbb{P}(N_{t+\Delta t} = (n_{1,t} + 1, n_{2,t}, n_{3,t})' | n_t, x_t, \theta) &= \beta_i s_t i_t \Delta t + o(\Delta t), \\ \mathbb{P}(N_{t+\Delta t} = (n_{1,t}, n_{2,t} + 1, n_{3,t})' | n_t, x_t, \theta) &= \gamma i_t \Delta t + o(\Delta t), \\ \mathbb{P}(N_{t+\Delta t} = (n_{1,t}, n_{2,t}, n_{3,t} + 1)' | n_t, x_t, \theta) &= \kappa (N_{pop} - s_t - i_t) \Delta t + o(\Delta t), \\ \mathbb{P}(N_{t+\Delta t} = (n_{1,t}, n_{2,t}, n_{3,t})' | n_t, x_t, \theta) &= 1 - (\beta_i s_t i_t + \gamma i_t + \kappa [N_{pop} - s_t - i_t]) \Delta t + o(\Delta t) \end{aligned}$$

and recall that N_{pop} is the total population size (assumed fixed and known). The stoichiometry matrix is given by

$$S = \begin{pmatrix} -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix}$$

and the hazard function is

$$h(x_t) = (\beta_t s_t i_t, \gamma i_t, \kappa [N_{pop} - s_t - i_t])'$$

Using (1), we may write the hazard function in terms of the incidence process as

$$h^*(n_t) = (\beta[s_0 - n_{t,1} + n_{t,3}][i_0 + n_{t,1} - n_{t,2}], \gamma[i_0 + n_{t,1} - n_{t,2}], \kappa[N_{pop} - s_0 - i_0 - n_{t,3} + n_{t,2}])'$$

Now define $N_{4,t} = \log \beta_t$ as a Brownian motion process scaled by σ_β . This leads to the CLE for the SIRS model (with time varying infection rate) as

$$dN_t = \{h_1^*(n_t), h_2^*(n_t), h_3^*(n_t), 0\} dt + \text{diag} \left\{ \sqrt{h_1^*(n_t)}, \sqrt{h_2^*(n_t)}, \sqrt{h_3^*(n_t)}, \sigma_\beta \right\} dW_t$$

where W_t is a length-4 vector of uncorrelated Brownian motion processes. The LNA then follows from equations (4), (8) and (9), with the (transpose of the) Jacobian matrix F_t' given by

$$F_t' = \begin{pmatrix} \exp(\eta_{4,t})(s_0 - i_0 - 2\eta_{t,1} + \eta_{t,2} + \eta_{3,t}) & \gamma & 0 & 0 \\ \exp(\eta_{4,t})(\eta_{t,1} - \eta_{t,3} - s_0) & -\gamma & \kappa & 0 \\ \exp(\eta_{4,t})(i_0 + n_{t,1} - n_{t,2}) & 0 & -\kappa & 0 \\ \exp(\eta_{4,t})(s_0 - \eta_{t,1} + \eta_{t,3})(i_0 + \eta_{t,1} - \eta_{t,2}) & 0 & 0 & 0 \end{pmatrix}$$

References

Andersson, H.K., Britton, T.: Stochastic Epidemic Models and Their Statistical Analysis. Springer, New York (2000). <https://doi.org/10.1007/978-1-4612-1158-7>

Andrieu, C., Doucet, A., Holenstein, R.: Particle Markov chain Monte Carlo for efficient numerical simulation. In: L'Ecuyer, P., Owen, A.B. (eds.) Monte Carlo and quasi-Monte Carlo methods 2008, pp. 45–60. Springer, Berlin (2009)

Andrieu, C., Doucet, A., Holenstein, R.: Particle Markov chain Monte Carlo methods (with discussion). *J. R. Stat. Soc.: Ser. B (Stat. Methodol.)* **72**(3), 269–342 (2010)

Andrieu, C., Roberts, G.O.: The pseudo-marginal approach for efficient computation. *Ann. Stat.* **37**, 697–725 (2009)

Bailey, N.T.J.: The Mathematical Theory of Infectious Diseases and Its Applications, 2nd edn. Hafner Press [Macmillan Publishing Co., Inc.], New York (1975)

Birrell, P.J., Ketsetzis, G., Gay, N.J., Cooper, B.S., Presanis, A.M., Harris, R.J., Angelis, D.D.: Bayesian modeling to unmask and predict influenza a/h1n1pdm dynamics in London. *Proc. Natl. Acad. Sci.* **108**(45), 18238–18243 (2011)

Birrell, P.J., Wernisch, L., Tom, B.D.M., Held, L., Roberts, G.O., Pebody, R.G., De Angelis, D.: Efficient real-time monitoring of an emerging influenza pandemic: how feasible? *Ann. Appl. Stat.* **12**, 74–93 (2020)

Boys, R.J., Giles, P.R.: Bayesian inference for stochastic epidemic models with time-inhomogeneous removal rates. *J. Math. Biol.* **55**, 223–247 (2007)

Cauchemez, S., Ferguson, N.M.: Likelihood-based estimation of continuous-time epidemic models from time-series data: application to measles transmission in London. *J. R. Soc. Interface* **5**, 885–897 (2008)

Chopin, N., Papaspiliopoulos, O.: An Introduction to Sequential Monte Carlo. Springer, New York (2020)

Choppala, P., Gunawan, D., Chen, J., Tran, M.N., Kohn, R.: Bayesian inference for state space models using block and correlated pseudo marginal methods (2016). Available from [arXiv:1311.3606](https://arxiv.org/abs/1311.3606)

Corbella, A., Presanis, A.M., Birrell, P.J., De Angelis, D.: Inferring epidemics from multiple dependent data via pseudo-marginal methods. Available from [arXiv:2204.08901](https://arxiv.org/abs/2204.08901) (2022)

Dahlin, J., Lindsten, F., Kronander, J., Schon, T.B. (2015). Accelerating pseudo-marginal Metropolis–Hastings by correlating auxiliary variables. Available from [arXiv:1511.05483](https://arxiv.org/abs/1511.05483)

Deligiannidis, G., Doucet, A., Pitt, M.K.: The correlated pseudo-marginal method. *J. R. Soc.: Ser. B (Stat. Methodol.)* **80**(5), 839–870 (2018)

Del Moral, P.: Feynman–Kac Formulae: Genealogical and Interacting Particle Systems with Applications. Springer, New York (2004)

Dureau, J., Kalogeropoulos, K., Baguelin, M.: Capturing the time-varying drivers of an epidemic using stochastic dynamical systems. *Biostatistics* **13**, 541–555 (2013)

Fearnhead, P., Giagos, V., Sherlock, C.: Inference for reaction networks using the Linear Noise Approximation. *Biometrics* **70**, 457–456 (2014)

Fearnhead, P., Meligkotsidou, L.: Exact filtering for partially observed continuous time models. *J. R. Stat. Soc.: Ser. B (Stat. Methodol.)* **66**(3), 771–789 (2004)

Ferm, L., Lötstedt, P., Hellander, A.: A hierarchy of approximations of the master equation scaled by a size parameter. *J. Sci. Comput.* **34**(2), 127–151 (2008)

Fintzi, J., Wakefield, J., Minin, V.N.: A linear noise approximation for stochastic epidemic models fit to partially observed incidence counts. *Biometrics* **78**(4), 1530–1541 (2021)

Fuchs, C.: Inference for Diffusion Processes with Applications in Life Sciences. Springer, Heidelberg (2013)

Funk, S., Camacho, A., Kucharski, A.J., Eggo, R.M., Edmunds, W.J.: Real-time forecasting of infectious disease dynamics with a stochastic semi-mechanistic model. *Epidemics* **22**, 56–61 (2018). **(The RAPIDD Ebola Forecasting Challenge)**

Gibson, G.J., Renshaw, E.: Estimating parameters in stochastic compartmental models using Markov chain methods. *IMA J. Math. Appl. Med. Biol.* **15**, 19–40 (1998)

Gibson, G.J., Streftaris, G., Thong, D.: Comparison and assessment of epidemic models. *Stat. Sci.* **33**, 19–33 (2018)

Gillespie, D.T.: Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.* **81**(25), 2340–2361 (1977)

Golightly, A., Bradley, E., Lowe, T., Gillespie, C.S.: Correlated pseudo-marginal schemes for time-discretised stochastic kinetic models. *Comput. Stat. Data Anal.* **136**, 92–107 (2019)

Golightly, A., Wilkinson, D.J.: Bayesian parameter inference for stochastic biochemical network models using particle Markov chain Monte Carlo. *Interface Focus* **1**(6), 807–820 (2011)

Gottschling, S., Meyer, S.: An epidemic airborne disease caused by the oak processionary caterpillar. *Pediatr. Dermatol.* **23**(1), 64–66 (2006)

Hethcote, W.: The mathematics of infectious diseases. *SIAM Rev.* **42**, 599–653 (2000)

Jacquez, J.: Compartmental Analysis in Biology and Medicine. Elsevier, Amsterdam (1972)

Jewell, C.P., Kypraios, T., Neal, P., Roberts, G.O.: Bayesian analysis for emerging infectious diseases. *Bayesian Anal.* **4**, 465–496 (2009)

- Kermack, W., McKendrick, A.: A contribution to the mathematical theory of epidemics. *Proc. R. Soc. Ser. A* **115**, 700–721 (1927)
- Komorowski, M., Finkenstadt, B., Harper, C., Rand, D.: Bayesian inference of biochemical kinetic parameters using the linear noise approximation. *BMC Bioinform.* **10**(1), 343 (2009)
- Kurtz, T.G.: Solutions of ordinary differential equations as limits of pure jump Markov processes. *J. Appl. Probab.* **7**, 49–58 (1970)
- Kurtz, T.G.: Limit theorems for sequences of jump Markov processes approximating ordinary differential processes. *J. Appl. Probab.* **8**, 334–356 (1971)
- Kurtz, T.G.: The relationship between stochastic and deterministic models for chemical reactions. *J. Chem. Phys.* **57**, 2976–2978 (1972)
- Kypraios, T., Neal, P., Prangle, D.: A tutorial introduction to Bayesian inference for stochastic epidemic models using approximate Bayesian computation. *Math. Biosci.* **287**, 42–53 (2017). (**50th Anniversary Issue**)
- Lloyd-Smith, J.O.: Maximum likelihood estimation of the negative binomial dispersion parameter for highly overdispersed data, with applications to infectious diseases. *PLoS ONE* **2**(2), e180 (2007)
- Maier, H., Spiegel, W., Kinaciyani, T., Hönigsmann, H.: Caterpillar dermatitis in two siblings due to the larvae of *Thaumetopoea processionea* L., the oak processionary caterpillar. *Dermatology* **208**(1), 70–73 (2004)
- Maier, H., Spiegel, W., Kinaciyani, T., Krehan, H., Cabaj, A., Schopf, A., Hönigsmann, H.: The oak processionary caterpillar as the cause of an epidemic airborne disease: survey and analysis. *Br. J. Dermatol.* **149**(5), 990–997 (2003)
- Mainprize, N., Straw, N.: Forestry Commission, Oak Processionary Moth (*Thaumetopoea processionea*) Contingency Plan. Forestry Commission. [Online; accessed 28-October-2021] (2021)
- McKinley, T.J., Cook, A.R., Deardon, R.: Inference in epidemic models without likelihoods. *Int. J. Biostat.* **5**(1), Article 24 (2009)
- McKinley, T.J., Ross, J.V., Deardon, R., Cook, A.R.: Simulation-based Bayesian inference for epidemic models. *Comput. Stat. Data Anal.* **71**, 434–447 (2014)
- Minas, G., Rand, D.A.: Long-time analytic approximation of large stochastic oscillators: simulation, analysis and inference. *PLoS Comput. Biol.* **13**, e1005676 (2017)
- O'Neill, P.D., Roberts, G.O.: Bayesian inference for partially observed stochastic epidemics. *J. R. Stat. Soc. A. Stat. Soc.* **162**, 121–129 (1999)
- Plummer, M., Best, N., Cowles, K., Vines, K.: CODA: convergence diagnosis and output analysis for MCMC. *R News* **6**(1), 7–11 (2006)
- Rahlenbeck, S., Utikal, J.: The oak processionary moth: a new health hazard? *Br. J. Gen. Pract.* **65**(637), 435–436 (2015)
- Ross, J.V., Pagendam, D.E., Pollett, P.K.: On parameter estimation in population models II: multi-dimensional processes and transient dynamics. *Theor. Popul. Biol.* **75**, 123–132 (2009)
- Scarponi, D., Iskauskas, A., Clark, R.A., Vernon, I., McKinley, T.J., Goldstein, M., McCreesh, N.: Demonstrating multi-country calibration of a tuberculosis model using new history matching and emulation package - HMER. Available from <https://www.medrxiv.org/content/early/2022/05/17/2022.05.13.22275052> (2022)
- Schmon, S.M., Deligiannidis, G., Doucet, A., Pitt, M.K.: Large sample asymptotics of the pseudo-marginal method. *Biometrika* **108**, 37–51 (2021)
- Schmon, S.M., Gagnon, P.: Optimal scaling of random walk Metropolis algorithms using Bayesian large-sample asymptotics. *Stat. Comput.* **32**, 1–16 (2022)
- Sherlock, C., Golightly, A., Gillespie, C.S.: Bayesian inference for hybrid discrete-continuous systems biology models. *Inverse Prob.* **30**, 114005 (2014)
- Spannaus, A., Papamarkou, T., Erwin, S., Christian, J.B.: Bayesian state space modelling for COVID-19: with Tennessee and New York case studies. Available from [arXiv:2012.15306](https://arxiv.org/abs/2012.15306) (2020)
- Stathopoulos, V., Girolami, M.A.: Markov chain Monte Carlo inference for Markov jump processes via the linear noise approximation. *Philos. Trans. R. Soc.* **A371**, 20110541 (2013)
- Stockdale, J.E., Kypraios, T., O'Neill, P.D.: Modelling and Bayesian analysis of the Abakaliki smallpox data. *Epidemics* **19**, 13–23 (2017)
- Stockdale, J.E., Kypraios, T., O'Neill, P.D.: Pair-based likelihood approximations for stochastic epidemic models. *Biostatistics* **22**, 575–597 (2021)
- Suprunenko, Y.F., Castle, M.D., Webb, C.R., Branson, J., Hoppit, A., Gilligan, C.A.: Estimating expansion of the range of oak processionary moth (*Thaumetopoea processionea*) in the UK from 2006 to 2019. *Agricultural and Forest Entomology* (2021)
- Swallow, B., Birrell, P., Blake, J., Burgman, M., Challenor, P., Coffeng, L.E., Vernon, I.: Challenges in estimation, uncertainty quantification and elicitation for pandemic modelling. *Epidemics* **38**, 100547 (2022)
- Wadkin, L.E., Branson, J., Hoppit, A., Parker, N.G., Golightly, A., Baggaley, A.W.: Inference for epidemic models with time-varying infection rates: tracking the dynamics of oak processionary moth in the uk. *Ecol. Evol.* **12**(5), e8871 (2022)
- Wallace, E.W.J.: A simplified derivation of the linear noise approximation (2010). Available from [arXiv:1004.4280](https://arxiv.org/abs/1004.4280)
- Whiteley, N., Rimella, L.: Inference in stochastic epidemic models via multinomial approximations. In: *International Conference on Artificial Intelligence and Statistics*, vol. 130, pp. 1297–1305 (2021)
- Wilkinson, D.J.: *Stochastic Modelling for Systems Biology*. CRC Press, Boca Raton (2018)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.