

RESEARCH ARTICLE

# Efficient sequential Bayesian inference for state-space epidemic models using ensemble data assimilation

Dhorasso Temfack<sup>1\*</sup>, Jason Wyse

School of Computer Science and Statistics, Trinity College Dublin, Dublin, Ireland

\* [temfackd@tcd.ie](mailto:temfackd@tcd.ie)



## Abstract

Estimating latent epidemic states and model parameters from partially observed, noisy data remains a major challenge in infectious disease modeling. State-space formulations provide a coherent probabilistic framework for such inference, yet fully Bayesian estimation is often computationally prohibitive because evaluating the observed-data likelihood requires integration over a latent trajectory. The Sequential Monte Carlo squared (SMC<sup>2</sup>) algorithm offers a principled approach for joint state and parameter inference, combining an outer SMC sampler over parameters with an inner particle filter that estimates the likelihood up to the current time point. Despite its theoretical appeal, this nested particle filter imposes substantial computational cost, limiting routine use in near-real-time outbreak response. We propose Ensemble SMC<sup>2</sup> (eSMC<sup>2</sup>), a computationally efficient variant that replaces the inner particle filter with an Ensemble Kalman Filter (EnKF) to approximate the incremental likelihood at each observation time. While this substitution introduces bias via a Gaussian approximation, we mitigate finite-sample effects using an unbiased Gaussian density estimator and adapt the EnKF for epidemic data through state-dependent observation variance. This makes our approach particularly suitable for overdispersed incidence data commonly encountered in infectious disease surveillance. Simulation experiments with known ground truth and an application to 2022 United States (U.S.) monkeypox incidence data demonstrate that eSMC<sup>2</sup> achieves substantial computational gains while producing posterior estimates comparable to SMC<sup>2</sup>. The method accurately recovers latent epidemic trajectories and key epidemiological parameters, providing an efficient framework for sequential Bayesian inference from imperfect surveillance data.

## OPEN ACCESS

**Citation:** Temfack D, Wyse J (2026) Efficient sequential Bayesian inference for state-space epidemic models using ensemble data assimilation. PLoS Comput Biol 22(5): e1014301. <https://doi.org/10.1371/journal.pcbi.1014301>

**Editor:** Nicola Perra, Queen Mary University of London, UNITED KINGDOM OF GREAT BRITAIN AND NORTHERN IRELAND

**Received:** December 4, 2025

**Accepted:** May 7, 2026

**Published:** May 18, 2026

**Copyright:** © 2026 Temfack, Wyse. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

**Data availability statement:** The code implementing the methods for Bayesian inference in state-space epidemiological models is available at [https://github.com/Dhorasso/eSMC2\\_Epi\\_Inference](https://github.com/Dhorasso/eSMC2_Epi_Inference). The 2022 U.S. monkeypox incidence data used in this study are publicly accessible

## Author summary

During infectious disease outbreaks, public health officials need timely estimates of how transmission is changing over time to make informed decisions about

from the Centers for Disease Control and Prevention (CDC) at <https://www.cdc.gov/monkeypox/data-research/cases/index.html>.

**Funding:** This work was supported by Taighde Éireann – Research Ireland (Grant number 21/FFP-P/10123 supported both DT and JW). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

**Competing interests:** The authors have declared that no competing interests exist.

interventions. However, existing statistical methods for extracting these insights from noisy surveillance data are computationally demanding, which limits their practical use when rapid answers are needed. In this work, we develop a more efficient computational approach that substantially speeds up these calculations while maintaining reliable inference. Our method combines two established techniques: one that efficiently tracks how disease states evolve over time, and another that systematically updates parameter estimates as new data arrive. We demonstrate that our approach reliably recovers important epidemic features like time-varying transmission rates and the effective reproduction number from incomplete and uncertain case count data. By substantially reducing the computational burden, our method makes real-time monitoring more feasible during ongoing outbreaks, potentially enabling faster and more informed public health responses when decisions must be made quickly.

## 1. Introduction

Accurate real-time inference of epidemic dynamics is essential for situational awareness, forecasting, and evaluating interventions during outbreaks. Public health decisions rely on rapidly updating estimates of quantities such as the time-varying reproduction number, incubation period, and infectious duration. Epidemic surveillance data are often noisy, incomplete, and reported irregularly, making the recovery of latent infection processes a challenging inverse problem: the unobserved transmission process must be inferred from partial, indirect observations of incidence or prevalence [1,2]. This motivates a state-space modeling framework, in which latent epidemic states evolve according to mechanistic transmission dynamics, such as the Susceptible–Exposed–Infectious–Removed (SEIR) model, while observations provide indirect information about the underlying epidemic process.

State-space models (SSMs) have become central to modern infectious disease modeling because they explicitly account for stochastic transmission and observational uncertainty [3,4]. Within this framework, Bayesian inference provides a principled approach to jointly estimate latent epidemic trajectories and epidemiological parameters. However, fully Bayesian inference for SSMs remains computationally demanding: the observed-data likelihood is typically intractable, and standard Monte Carlo methods may be impractical for large-scale or real-time applications.

Particle filtering and related methods have made Bayesian state estimation tractable in epidemic modeling, with applications to influenza, Ebola, COVID-19, and more recently, monkeypox [3,5–7]. Particle Markov Chain Monte Carlo (PMCMC) [8] enables exact Bayesian inference by combining MCMC with unbiased likelihood estimates from a particle filter (PF), but it can be computationally expensive since a full PF must be run at each MCMC iteration. Likelihood estimates may also have high variance, leading to poor mixing and increased Monte Carlo error, especially when the data provide limited information about parameters. Sequential Monte Carlo Squared (SMC<sup>2</sup>) [9] extends PMCMC by maintaining a population of parameter

particles, each with an internal PF for state inference and likelihood evaluation. Parameter particles are reweighted, and a resample–move step with a PMCMC mutation kernel mitigates degeneracy and maintains diversity. This nested structure allows sequential Bayesian updating as new data arrive and has been applied successfully in epidemic modeling [10,11], but computational cost grows with particle number and weight degeneracy remains a concern [12]. Using the Auxiliary Particle Filter (APF) [13] within SMC<sup>2</sup> can improve particle selection and reduce the number of particles required [10], but it adds computational complexity due to more elaborate weighting and proposal steps. Likelihood-free approaches, notably Approximate Bayesian Computation (ABC), offer an alternative when the likelihood is fully intractable [14]. However, ABC typically requires many model simulations and careful choice of summary statistics and thresholds, which can make real-time application challenging.

As a computationally efficient and appealing method, the Ensemble Kalman Filter (EnKF) [15] has demonstrated strong performance in data assimilation for near-Gaussian systems. Compared with PF, the EnKF can achieve reliable results with substantially fewer particles (also referred to as ensemble members), since it updates the ensemble through linear transformations rather than importance resampling, thereby mitigating weight degeneracy. In practice, this often allows the EnKF to outperform particle filters in moderately nonlinear settings where the latter suffer from sample depletion [16]. In epidemiological applications, the EnKF is used to track infection dynamics and estimate transmission parameters [17–19]. Most implementations employ an augmented EnKF to jointly estimate states and parameters [18,20]. While effective, this approach can be less flexible in adapting to changing transmission dynamics, since parameters are treated as additional static states and their artificial evolution (e.g., via random-walk perturbations) may lead to filter degeneracy or drift over time, a well-known issue shared by most Bayesian augmented filtering methods [21].

Several hybrid approaches have been developed to combine the scalability of the EnKF with the inferential rigor of full Bayesian methods. For instance, Stroud et al. [21] proposed a two-layer filter pairing an EnKF for states with a particle filter for parameters, while Katzfuss et al. [22] suggested replacing the particle filter with an EnKF within the particle MCMC framework, an approach that substantially reduces computational cost at the expense of a small approximation bias. Extending this idea, Drovandi et al. [23] introduced the Ensemble MCMC (eMCMC) algorithm, which corrects the bias in the EnKF likelihood using an unbiased Gaussian density estimator together with an early-rejection scheme to improve efficiency. Their results showed that EnKF-based likelihoods can achieve accurate posterior inference in moderately nonlinear systems with orders-of-magnitude speedups compared to particle filters.

In this work, we propose Ensemble SMC<sup>2</sup> (eSMC<sup>2</sup>), a novel algorithm that integrates the EnKF into the inner layer of the SMC<sup>2</sup> framework to enable fully sequential Bayesian inference over both latent states and parameters. This approach substantially reduces computational cost while maintaining accurate posterior inference. To mitigate the bias inherent in the standard EnKF likelihood estimate, we incorporate the unbiased Gaussian likelihood correction employed by Drovandi et al. [23], which reduces sensitivity to the number of ensemble members. While we cannot provide formal bias quantification, prior studies indicate that EnKF-based likelihoods yield stable and accurate inference in moderately nonlinear systems [22–24]. We further employ a variant inspired by the Poisson Kalman Filter [25], making the approach suitable for observations typical in disease surveillance and relaxing the Gaussian noise assumptions of standard EnKF methods. We demonstrate its performance on a diffusion-driven SEIR model with incidence data aggregated over fixed reporting intervals, showing that eSMC<sup>2</sup> delivers an accurate estimate of the time-varying transmission dynamics and key epidemiological parameters using both synthetic and real datasets.

In parallel with our work, Golightly et al. [26] independently proposed an extension of SMC<sup>2</sup> in which the inner particle filter is replaced by an EnKF. Their nested EnKF (NEnKF) maintains parameter diversity via a resample–move strategy, combined with delayed-acceptance steps to improve efficiency. While both approaches embed an EnKF within SMC<sup>2</sup>, they differ in their motivating application and methodological focus. Golightly et al. [26] consider a broad class of state-space models, whereas our method is motivated by epidemic inference and introduces an unbiased Gaussian likelihood correction with

variance adaptation for count-valued data. Nevertheless, both approaches highlight the potential of embedding the EnKF within SMC<sup>2</sup> as an efficient alternative to a nested particle filter for sequential state and parameter inference.

The remainder of this paper is organized as follows. In [Section 2](#), we present the SEIR model within a state-space framework, review the EnKF, and describe the eSMC<sup>2</sup> algorithm. [Section 3](#) presents simulation studies and an application to U.S. monkeypox incidence data. Finally, [Section 4](#) discusses the limitations of our approach, potential improvements, and directions for future research.

## 2. Methodology

### 2.1. SEIR model in a Bayesian state-space framework

We model the spread of infectious diseases using a stochastic SEIR framework that explicitly tracks both latent epidemic states and the time-varying transmission rate. While we present the SEIR model here as a concrete example, the methodology is applicable to a wide range of compartmental epidemic models. Let the latent state at time  $t$  be  $x_t = (S_t, E_t, I_t, R_t, Z_t, \log(\beta_t))^T$ , where  $S_t$ ,  $E_t$ ,  $I_t$ ,  $R_t$  denote the numbers of susceptible, exposed (infected but not yet infectious), infectious, and removed (recovered or deceased) individuals, respectively. The variable  $Z_t$  is not a dynamical compartment but an auxiliary quantity representing the new incidence over each reporting interval, included for convenience to directly match the observation model. The variable  $\log(\beta_t)$  is incorporated into the latent state to allow inference on the time-varying transmission rate. The latent dynamics follow the system

$$\begin{cases} \frac{dS_t}{dt} = -\beta_t \frac{S_t I_t}{N}, & \frac{dE_t}{dt} = \beta_t \frac{S_t I_t}{N} - \alpha E_t, & \frac{dI_t}{dt} = \alpha E_t - \gamma I_t, & \frac{dR_t}{dt} = \gamma I_t, \\ Z_t = \int_{t-1}^t \alpha E_s ds, \\ d \log(\beta_t) = \nu_\beta dB_t. \end{cases} \quad (1)$$

where  $N = S_t + E_t + I_t + R_t$  is the constant population size, and  $1/\alpha$  and  $1/\gamma$  denote the average incubation and infectious periods, respectively. The transmission rate  $\beta_t$  is modeled as a stochastic process on the log scale, following a Brownian motion with volatility  $\nu_\beta$  and driven by a standard Wiener process  $B_t$ . This allows  $\beta_t$  to vary over time in response to behavioral changes, public health interventions, or seasonal effects [3,27]. The parameter  $\nu_\beta$  controls the magnitude of these fluctuations: larger values permit faster changes, while smaller values constrain variation. This formulation enables the model to capture unobserved temporal variability in transmission intensity throughout an epidemic. The effective reproduction number is defined as  $R_{\text{eff}}(t) = \beta_t S_t / (\gamma N)$ . To simulate sample trajectories, the continuous-time system (1) is propagated forward using a discrete-time approximation. Let  $\theta$  denote the vector of unknown model parameters governing both the latent dynamics in (1) and the observation model described below. The latent state is updated according to

$$x_t | x_{t-1}, \theta \sim p(\cdot | x_{t-1}, \theta). \quad (2)$$

The forward process  $p(\cdot | x_{t-1}, \theta)$  updates  $S_t$ ,  $E_t$ ,  $I_t$ ,  $R_t$ , and  $Z_t$  deterministically via a forward Euler scheme, while  $\log(\beta_t)$  evolves stochastically using the Euler–Maruyama method [28]. Formally, the deterministic updates are represented as a Dirac delta distribution within the transition density  $p(x_t | x_{t-1}, \theta)$ , whereas the stochasticity of  $\log(\beta_t)$  introduces process noise. Details on the discretization and transition density are provided in Appendix A in [S1 Text](#).

We assume that the available data  $y_{1:t} = \{y_1, \dots, y_t\}$ , representing daily reported cases, are noisy measurements of the latent incidence  $Z_t$ , linked through the conditional distribution  $p(y_t | x_t, \theta) = p(y_t | Z_t, \theta)$ . A natural starting point is to model the counts as realizations of a Poisson random variable. In practice, however, daily case counts often fluctuate more than a

Poisson model would predict, a phenomenon known as overdispersion. To account for this, we also consider a Negative Binomial formulation.

Specifically, the number of reported cases  $y_t$  in a reporting interval  $(t-1, t]$ ,  $t = 1, \dots, T$ , is modeled as either

$$y_t | x_t, \theta \sim \text{Poisson}(\rho Z_t) \quad \text{or} \quad (3a)$$

$$y_t | x_t, \theta \sim \text{NegBin}(\text{mean} = \rho Z_t, \text{var} = \rho Z_t + \phi(\rho Z_t)^2), \quad (3b)$$

The Negative Binomial is parameterized by its mean  $\rho Z_t$  and variance  $V_t = \text{Var}[y_t | x_t] = \rho Z_t + \phi(\rho Z_t)^2$ , where  $\phi > 0$  is the overdispersion parameter; the Poisson model, with  $V_t = \text{Var}[y_t | x_t] = \rho Z_t$ , is recovered as  $\phi \rightarrow 0$ . This mean-variance parameterization is used directly in the ensemble-based observation variance described in Section 2.2 below, ensuring consistency between the observation model and the EnKF update step.  $\rho \in (0, 1]$  is the reporting fraction, informed by expert knowledge. Specifying  $\rho$  helps mitigate practical non-identifiability issues highlighted in [27], since decreases in observed cases could otherwise be attributed equally to a lower reporting rate, a reduced transmission rate, or a combination of both. Practical considerations regarding the structural identifiability of parameters in such models were also discussed in [11].

Together, Eqs (2)–(3) define our complete state-space model [4]. Let  $x_{0:t} = \{x_0, \dots, x_t\}$  denote the latent states up to time  $t$ . Prior knowledge about  $\theta$  and initial state  $x_0$  is encoded in the distributions  $p(\theta)$  and  $p(x_0 | \theta)$ , respectively. The joint posterior of states and parameters given the observations is then

$$p(x_{0:t}, \theta | y_{1:t}) \propto p(\theta) p(x_0 | \theta) \prod_{s=1}^t [p(x_s | x_{s-1}, \theta) p(y_s | x_s, \theta)]. \quad (4)$$

In this sequential learning setting, the inferential objectives at each time  $t$  are twofold. First, for a given fixed value of  $\theta$ , filtering aims to infer the current state  $x_t$  conditional on all observations up to that time, corresponding to:

$$p(x_t | y_{1:t}, \theta) \propto p(y_t | x_t, \theta) \underbrace{\int p(x_t | x_{t-1}, \theta) p(x_{t-1} | y_{1:t-1}, \theta) dx_{t-1}}_{\text{Forecast density } p(x_t | y_{1:t-1}, \theta)}. \quad (5)$$

Second, parameter estimation focuses on the marginal posterior of  $\theta$ ,

$$\begin{aligned} p(\theta | y_{1:t}) &\propto p(\theta) \prod_{s=1}^t p(y_s | y_{1:s-1}, \theta) \\ &\propto p(\theta) \underbrace{\int \prod_{s=1}^t [p(y_s | x_s, \theta) p(x_s | y_{1:s-1}, \theta)] dx_s}_{\text{marginal likelihood } p(y_{1:t} | \theta)}, \end{aligned} \quad (6)$$

where we adopt the convention that  $p(y_1 | \theta) = p(y_1 | y_{1:0}, \theta)$ . The filtering distribution that accounts for parameter uncertainty is obtained by marginalising the conditional filtering distribution over the parameter posterior, i.e.,

$$p(x_t | y_{1:t}) = \int p(x_t | y_{1:t}, \theta) p(\theta | y_{1:t}) d\theta. \quad (7)$$

For nonlinear state-space models, Eqs (5) and (6) cannot be solved analytically and are generally intractable. Consequently, practical inference relies on numerical approximations such as particle filtering or EnKF-based methods to generate samples from the filtering distribution or to estimate the marginal likelihood. In this work, motivated by the computational efficiency and scalability considerations discussed in the introduction, we focus on EnKF-based inference.

## 2.2. Ensemble Kalman filter

We now describe the ensemble Kalman filter (EnKF) [15], a sequential inference method for state-space models with nonlinear latent dynamics. While the EnKF was originally developed for systems with Gaussian observations, many real-world datasets, including those arising in epidemic surveillance often deviate from this assumption and may exhibit overdispersion. To accommodate such data, we follow the approach of Ebeigbe et al. [25] and extend it by constructing observation noise terms consistent with the observation models in Eq (3).

The EnKF alternates between prediction and update operations at each time step to sequentially refine the latent state estimate based on incoming observations (Fig 1). During the prediction step, the ensemble is propagated forward through the system dynamics, while in the update step, the ensemble is adjusted using the most recent observation.

Let us first suppose that the parameters  $\theta$  and a sequence of observations  $y_{1:T}$  (in our case, the daily incidence) are known. At each time step  $t$ , the EnKF aims to infer the latent state  $x_t$  given the data  $y_{1:t}$ . Let the ensemble at time  $t-1$  be denoted by  $\{x_{t-1}^{(i)}\}_{i=1}^{N_x}$ , representing samples from the conditional  $p(x_{t-1}|y_{1:t-1}, \theta)$ . The forecast ensemble  $\{x_t^{(f,i)}\}_{i=1}^{N_x}$  is generated by propagating each particle forward through the transition dynamics, that is,

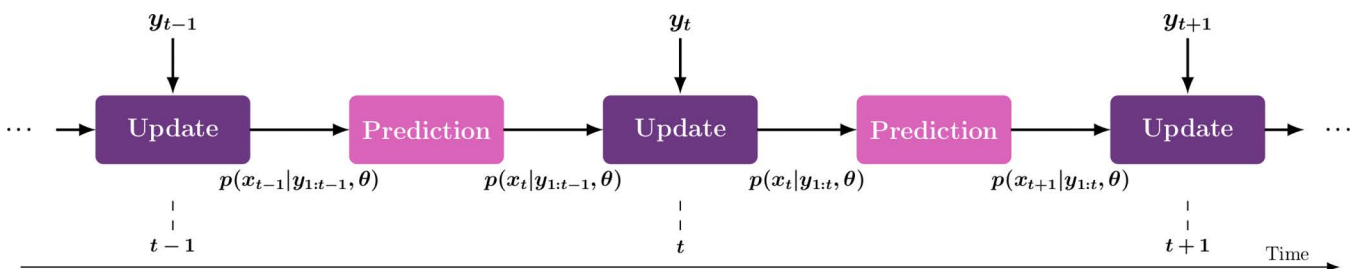
$$x_t^{(f,i)} \sim p(\cdot|x_{t-1}^{(i)}, \theta), \quad i = 1, \dots, N_x. \tag{8}$$

The forecast density  $p(x_t|y_{1:t-1})$  is then approximated by:

$$p(x_t|y_{1:t-1}, \theta) \approx \mathcal{N}(x_t; \hat{\mu}_{t|t-1}, \hat{\Sigma}_{t|t-1}), \tag{9}$$

where  $\hat{\mu}_{t|t-1}$  and  $\hat{\Sigma}_{t|t-1}$  are the sample mean and covariance of the forecast ensemble.

A key adaptation for epidemic data is the specification of the observation error variance  $V_t$ . Classical Kalman filtering assumes Gaussian observations with constant variance (or at least does not depend on the latent state), an assumption often violated for epidemiological models, where the variability of reported incidence typically scales with its magnitude [25]. A principled approach is to set  $V_t$  to match the expected conditional variance of the observation given the latent state  $E[\text{Var}[y_t|x_t]]$ , as implied by the observation model (3). In practice, we approximate the expectation above using the forecast ensemble. We define the ensemble estimate of the observation error variance by



**Fig 1. Schematic illustration of the sequential update–prediction cycle in the EnKF.** Each time step alternates between an update stage, where observations  $y_t$  are assimilated to refine the latent state estimate, and a prediction stage, where the ensemble is propagated forward through the system dynamics.

<https://doi.org/10.1371/journal.pcbi.1014301.g001>

$$\begin{aligned}
 E[\text{Var}[y_t|x_t]] &\approx V_t^{N_x} = \frac{1}{N_x} \sum_{i=1}^{N_x} \text{Var}[y_t|x_t^{(f,i)}], \\
 &= \begin{cases} \frac{1}{N_x} \sum_{i=1}^{N_x} Hx_t^{(f,i)}, & \text{Poisson model,} \\ \frac{1}{N_x} \sum_{i=1}^{N_x} [Hx_t^{(f,i)} + \phi(Hx_t^{(f,i)})^2], & \text{Negative Binomial model,} \end{cases} \quad (10)
 \end{aligned}$$

where  $H = [0, 0, 0, 0, \rho, 0]$  projects the model state onto the observed component, i.e.,  $Hx_t^{(f,i)} = \rho Z_t^{(f,i)}$ ,  $i = 1, \dots, N_x$ . To prevent numerical instability when  $V_t^{N_x}$  is very small (or near zero), we regularize the variance estimate by setting

$$\widehat{V}_t = \max(\eta, V_t^{N_x}), \quad (11)$$

where in our implementation we set  $\eta = 0.1$ , small relative to the typical magnitude of  $V_t^{N_x}$ . We found that results are insensitive to this choice provided  $\eta < 1$ . Appendix A in [S1 Text](#) provides a detailed justification for this specification. This specification allows the EnKF update to capture key features of the observation distribution (3), though it can overestimate uncertainty when case counts are extremely low. Such low-incidence scenarios are not the primary focus of our study, and alternative sequential methods, such as the Lifebelt particle filter [29], may be preferred in those cases. A similar treatment of adaptive variance specification is discussed in (Section 3.2, [30]).

During the analysis (update) step, each forecast ensemble member is adjusted using the Kalman gain and the innovation. In the stochastic (perturbed-observation) EnKF, the update is

$$x_t^{(i)} = x_t^{(f,i)} + K_t (y_t + v_t^{(i)} - Hx_t^{(f,i)}), \quad v_t^{(i)} \sim \mathcal{N}(0, \widehat{V}_t), \quad (12)$$

where the Kalman gain is given by

$$K_t = \widehat{\Sigma}_{t|t-1} H^T (H \widehat{\Sigma}_{t|t-1} H^T + \widehat{V}_t)^{-1}. \quad (13)$$

Using the Gaussian approximation of the forecast density in (9), the one-step predictive likelihood can be expressed as

$$p(y_t|y_{1:t-1}, \theta) \approx \mathcal{N}(y_t; H\widehat{\mu}_{t|t-1}, H\widehat{\Sigma}_{t|t-1}H^T + \widehat{V}_t). \quad (14)$$

Only the key components of the EnKF algorithm for state inference are presented above; for a comprehensive overview, see [22]. For linear–Gaussian SSMs, the likelihood approximation in (14) converges to the true likelihood as  $N_x \rightarrow \infty$ , but for finite ensembles it remains biased due to sampling variability in the forecast mean and covariance [23]. Importantly, even if the forecast ensemble is exactly Gaussian, the estimated density  $\mathcal{N}(y_t; H\widehat{\mu}_{t|t-1}, H\widehat{\Sigma}_{t|t-1}H^T + \widehat{V}_t)$  is not an unbiased estimate of the true likelihood. This occurs because the empirical normal density  $\mathcal{N}(y; \mu_{N_x}, \Sigma_{N_x})$ , computed from the finite-ensemble mean  $\mu_{N_x}$  and covariance  $\Sigma_{N_x}$ , is a biased estimator of the true density  $\mathcal{N}(y; \mu, \Sigma)$  for any finite  $N_x$ .

Nevertheless, empirical results in the literature [21,23,24] show that the EnKF likelihood performs remarkably well for inference in both linear and nonlinear systems, with posterior estimates only weakly dependent on the ensemble size [23]. To mitigate the finite-sample bias, we adopt the unbiased Gaussian density estimator of Ghurye and Olkin [31], later employed by Price et al. [32] and Drovandi et al. [23]. This estimator yields an unbiased estimate of  $\mathcal{N}(y; \mu, \Sigma)$  provided  $N_x > d_y + 3$ , where  $d_y$  is the dimension of  $y$ . It is defined as

$$\widehat{\mathcal{N}}_{\text{unb}}(\mathbf{y}; \mu_{N_x}, \Sigma_{N_x}) = (2\pi)^{-d_y/2} \frac{c(d_y, N_x - 2)}{c(d_y, N_x - 1)(1 - 1/N_x)^{d_y/2}} |M_{N_x}|^{-(N_x - d_y - 2)/2} \times \psi \left( M_{N_x} - \frac{(\mathbf{y} - \mu_{N_x})(\mathbf{y} - \mu_{N_x})^\top}{1 - 1/N_x} \right)^{(N_x - d_y - 3)/2}, \quad (15)$$

where  $M_{N_x} = (N_x - 1)\Sigma_{N_x}$ , and  $\psi(A) = |A|$  if  $A$  is positive definite, and 0 otherwise. Here,  $|A|$  denotes the determinant of the matrix  $A$  and the constants  $c(d, \nu)$  are defined as

$$c(d_y, \nu) = \frac{2^{-d_y/2} \pi^{-d_y(d_y-1)/4}}{\prod_{i=1}^{d_y} \Gamma(\frac{1}{2}(\nu - i + 1))}. \quad (16)$$

We can now replace the standard EnKF likelihood term in (14) with its unbiased counterpart:

$$\widehat{p}_{\text{enkf}}^{N_x}(\mathbf{y}_t | \mathbf{y}_{1:t-1}, \theta) = \widehat{\mathcal{N}}_{\text{unb}} \left( \mathbf{y}_t; H\widehat{\mu}_{t|t-1}, H\widehat{\Sigma}_{t|t-1}H^\top + \widehat{V}_t \right). \quad (17)$$

The approximation of the marginal likelihood up to the time  $t$  is then given by:

$$\widehat{p}_{\text{enkf}}^{N_x}(\mathbf{y}_{1:t} | \theta) = \prod_{s=1}^t \widehat{p}_{\text{enkf}}^{N_x}(\mathbf{y}_s | \mathbf{y}_{1:s-1}, \theta). \quad (18)$$

Unlike in the case of the particle filter [8], the resulting posterior approximation  $\widehat{p}_{\text{enkf}}^{N_x}(\theta | \mathbf{y}_{1:t}) \propto \widehat{p}_{\text{enkf}}^{N_x}(\mathbf{y}_{1:t} | \theta) p(\theta)$ , still does not target the true posterior (due to the Gaussian assumption). However, when the Gaussian assumption holds, the resulting posterior no longer depends on  $N_x$ , eliminating the finite-ensemble bias and yielding a stable target for Bayesian inference [23,31]. The EnKF procedure, including unbiased likelihood estimation, is outlined in Algorithm 1. This estimator can then be incorporated into a Metropolis–Hastings algorithm to target the sequence of posterior distributions  $\{p(\theta), p(\theta | \mathbf{y}_1), p(\theta | \mathbf{y}_{1:2}), \dots, p(\theta | \mathbf{y}_{1:T})\}$ .

### Algorithm 1 Ensemble Kalman filter

Operations involving index  $i$  must be performed for  $i = 1, \dots, N_x$ .

**Inputs:** Observation  $\mathbf{y}_{1:T}$ , Number of ensemble members  $N_x$ , Parameter vector  $\theta$ .

**Output:** Filtering ensembles  $\{\mathbf{x}_{0:T}^{(i)}\}_{i=1}^{N_x}$  and likelihood estimate  $\widehat{p}_{\text{enkf}}^{N_x}(\mathbf{y}_{1:T} | \theta)$

- 1: Sample initial particles:  $\mathbf{x}_0^{(i)} \sim p(\mathbf{x}_0)$
- 2: **for**  $t=1$  to  $T$  **do**
- 3:   Sample forecast ensemble as:  $\mathbf{x}_t^{(f,i)} \sim p(\cdot | \mathbf{x}_{t-1}^{(i)}, \theta)$
- 4:   Compute the observation variance, ensemble forecast mean, and variance:  $\widehat{V}_t, \widehat{\mu}_{t|t-1}$ , and  $\widehat{\Sigma}_{t|t-1}$ .
- 5:   Compute the incremental Likelihood:  $\widehat{p}_{\text{enkf}}^{N_x}(\mathbf{y}_t | \mathbf{y}_{1:t-1}, \theta) = \widehat{\mathcal{N}}_{\text{unb}} \left( \mathbf{y}_t; H\widehat{\mu}_{t|t-1}, H\widehat{\Sigma}_{t|t-1}H^\top + \widehat{V}_t \right)$ .
- 6:   Shift each ensemble member as:  $\mathbf{x}_t^{(i)} = \mathbf{x}_t^{(f,i)} + K_t(\mathbf{y}_t + \mathbf{v}_t^{(i)} - H\mathbf{x}_t^{(f,i)}), \quad \mathbf{v}_t^{(i)} \sim \mathcal{N}(0, \widehat{V}_t)$
- 7: **end for**
- 8: Compute overall likelihood:  $\widehat{p}_{\text{enkf}}^{N_x}(\mathbf{y}_{1:T} | \theta) = \prod_{t=1}^T \widehat{p}_{\text{enkf}}^{N_x}(\mathbf{y}_t | \mathbf{y}_{1:t-1}, \theta)$ .

### 2.3. Ensemble SMC<sup>2</sup>

The EnKF algorithm described previously assumes that the model parameters  $\theta$  are known. In practice, these parameters are rarely known, and it is preferable to assign a distribution  $p(\theta)$  that reflects prior knowledge. To account for this

uncertainty, inference on  $\theta$  can be performed by embedding the EnKF within an SMC sampler over the parameter space [33]. This approach follows the general principle of SMC<sup>2</sup> [9], which traditionally employs a particle filter to estimate the likelihood. At each observation time  $t$ , the posterior distribution of the parameters given data  $y_{1:t}$  satisfies the recursive relation

$$p(\theta|y_{1:t}) = \frac{p(y_t|y_{1:t-1}, \theta) p(\theta|y_{1:t-1})}{\int p(y_t|y_{1:t-1}, \theta) p(\theta|y_{1:t-1}) d\theta}. \quad (19)$$

A practical way to approximate this posterior is to represent it with a finite set of weighted parameter particles. Let's assume that  $\{\theta^m, \omega_t^m\}_{m=1}^{N_\theta}$  is particle approximation from the previous posterior  $p(\theta|y_{1:t-1})$ . The updated posterior can then be expressed as:

$$p(\theta|y_{1:t}) \approx \sum_{m=1}^{N_\theta} \omega_t^m \delta_{\theta^m}(\theta), \quad (20)$$

where  $\omega_t^m$ , for  $m = 1, \dots, N_\theta$  are the normalized weights given by

$$\omega_t^m = \frac{\omega_{t-1}^m p(y_t|y_{1:t-1}, \theta^m)}{\sum_{j=1}^{N_\theta} \omega_{t-1}^j p(y_t|y_{1:t-1}, \theta^j)}. \quad (21)$$

As the incremental likelihood  $p(y_t|y_{1:t-1}, \theta)$  is intractable in a nonlinear system, at time  $t$ , the weight of the  $\theta$ -particle  $\theta^m$  is updated using an estimator  $\hat{p}_{\text{enkf}}^{N_x}(y_t|y_{1:t-1}, \theta^m)$  obtained from the associated EnKF. This differs from the related work of Wu et al. [34], where the EnKF is used to construct proposal kernels in the SMC sampler but not as a likelihood estimator. As in standard SMC<sup>2</sup>, particle degeneracy may occur over time. To mitigate this, resampling (described in Appendix B in S1 Text) is triggered whenever the effective sample size (ESS) drops below  $N_\theta/2$  [9,35]. This is followed by a jittering mechanism that rejuvenates the particles by proposing new candidates, which are then accepted or rejected through a few iterations of Particle Marginal Metropolis–Hastings (PMMH). Since the weighted particles already approximate  $p(\theta|y_{1:t})$  (under the Gaussian assumption), applying a Metropolis–Hastings update to each one can only enhance this approximation (see Algorithm 2).

### Algorithm 2 PMMH mutation

**Inputs:** Observations  $y_{1:t}$ , Current parameter  $\theta$ , current likelihood estimate:  $\hat{p}_{\text{enkf}}^{N_x}(y_{1:t}|\theta)$ , Number of ensemble members  $N_x$ , Number of PMMH moves  $R$

**Output:** New parameter  $\theta'$ , New likelihood:  $\hat{p}_{\text{enkf}}^{N_x}(y_{1:t}|\theta')$

1: **for**  $r=1$  to  $R$  **do**

2: Propose  $\theta^* \sim q(\theta^*|\theta)$

3: Run a new EnKF (Algorithm 1) with  $\theta^*$  and compute  $\hat{p}_{\text{enkf}}^{N_x}(y_{1:t}|\theta^*)$

4: Calculate acceptance ratio:

$$\vartheta(\theta^*, \theta) = \min \left\{ 1, \frac{\hat{p}_{\text{enkf}}^{N_x}(y_{1:t}|\theta^*) p(\theta^*)}{\hat{p}_{\text{enkf}}^{N_x}(y_{1:t}|\theta) p(\theta)} \times \frac{q(\theta|\theta^*)}{q(\theta^*|\theta)} \right\}$$

5: Accept  $\theta^*$  with probability  $\vartheta(\theta^*, \theta)$ , otherwise retain  $\theta$

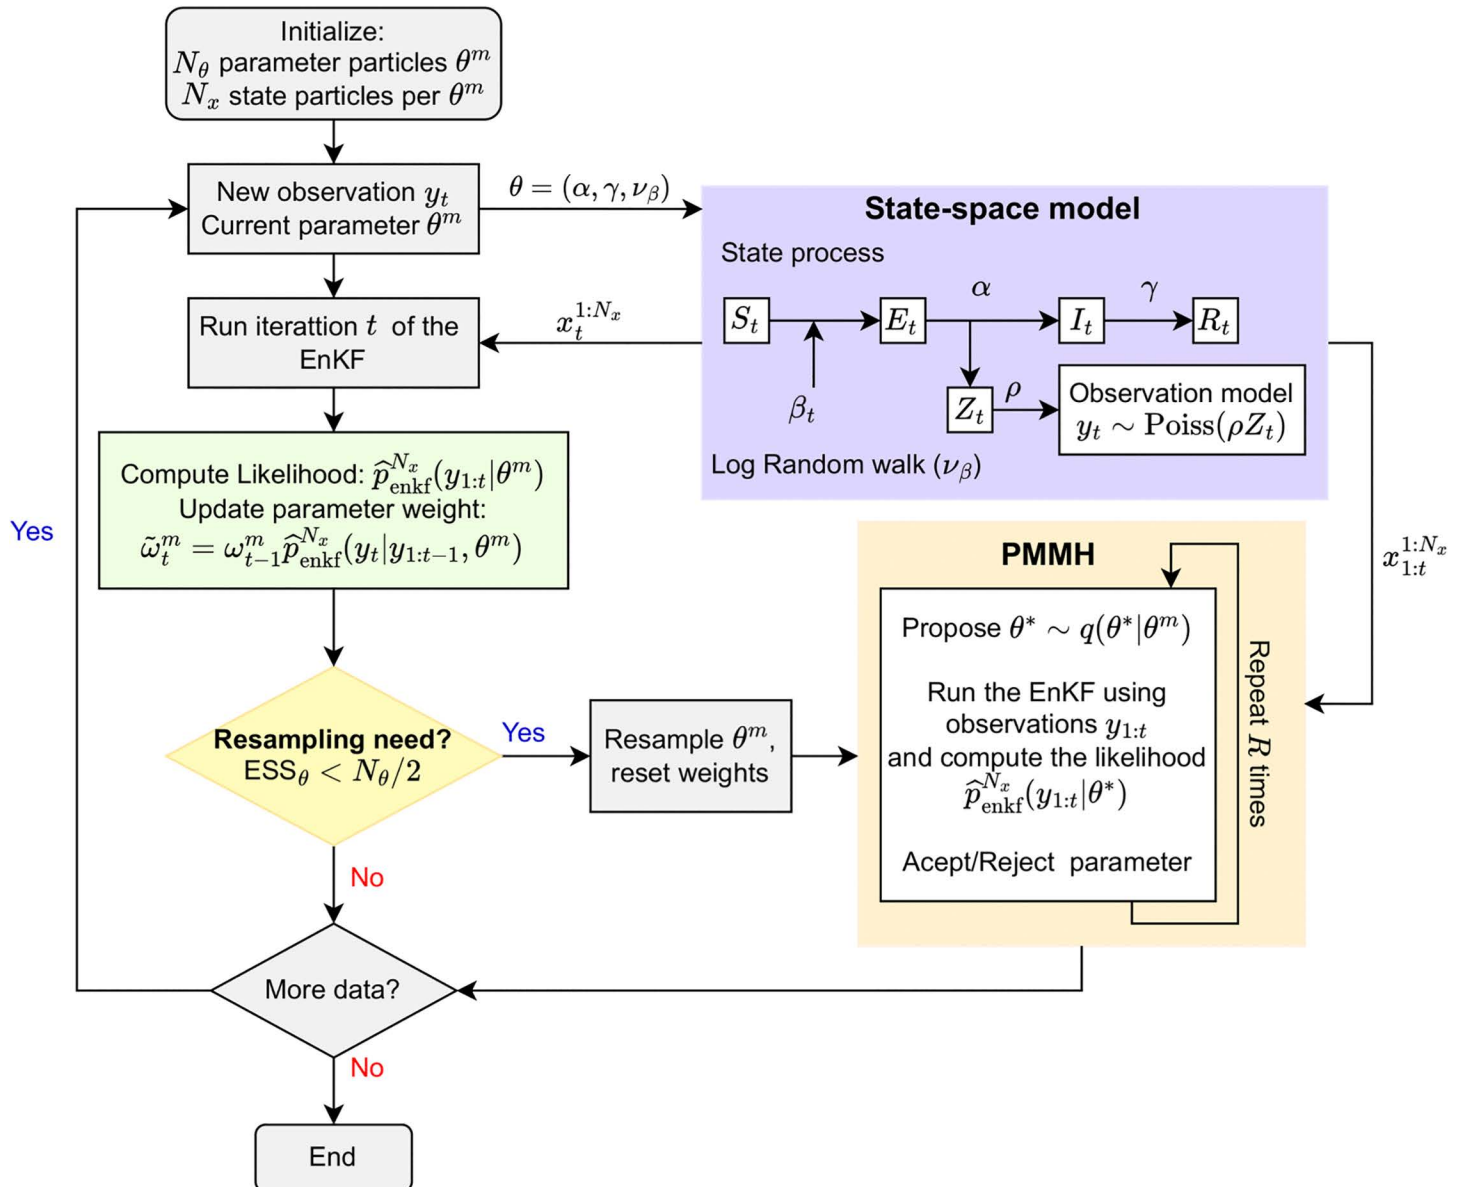
6: **end for**

In the numerical experiments presented in Section 3, we use an independent proposal distribution,  $q(\theta^*|\theta) = q(\theta^*)$ . The proposal is constructed directly on the parameter space using a multivariate normal distribution. Specifically, new parameter particles  $\theta^*$  are drawn from

$$q(\theta^*) = \mathcal{N}(\theta^*; \hat{E}[\theta|y_{1:t}], c \widehat{\text{Var}}[\theta|y_{1:t}]), \quad (22)$$

where  $\widehat{E}[\theta|y_{1:t}]$  and  $\widehat{\text{Var}}[\theta|y_{1:t}]$  denote the empirical mean and covariance of the current particle population  $\{\theta^m, \omega_t\}_{m=1}^{N_\theta}$ , and  $c$  is a scaling constant that controls the covariance matrix.

The eSMC<sup>2</sup> algorithm can be viewed as a variant of the standard SMC<sup>2</sup>, where the particle filter-based likelihood is replaced by an EnKF approximation in (18). From a practical perspective, this modification is expected to have only a slight effect on the posterior estimates of  $\theta$  and  $x_{1:t}$ . The complete procedure is detailed in Algorithm 3, with a schematic overview in Fig 2.



**Fig 2. Flowchart of the eSMC<sup>2</sup> algorithm.** Each parameter particle carries an ensemble of state particles, which are propagated using the EnKF. Weights are updated based on an EnKF-based likelihood approximation. When the ESS falls below a threshold, parameter particles are resampled and rejuvenated via a PMMH step. This procedure is repeated for every data point, allowing sequential Bayesian learning.

<https://doi.org/10.1371/journal.pcbi.1014301.g002>

To propagate parameter uncertainty in the hidden state estimation, we approximate the marginal posterior distribution of the latent states at time  $t$ ,  $p(x_t|y_{1:t})$ , by integrating over a subset of parameter particles,  $\{\theta^j\}_{j=1}^{N_c} \sim p(\theta|y_{1:t})$ , drawn from Algorithm 3 output [36]. In this work, we set  $N_c = 100$ , which in our numerical examples (Section 3) provides satisfactory coverage of the posterior. For each sampled  $\theta^j$ , an EnKF with  $N_x$  ensemble members is run to approximate the filtering distribution  $p(x_t|y_{1:t}, \theta^j)$ , and the resulting trajectories are combined across all parameter samples to form an empirical approximation of the state posterior:

$$p(x_t|y_{1:t}) \approx \frac{1}{N_c} \sum_{j=1}^{N_c} p(x_t|y_{1:t}, \theta^j), \tag{23}$$

resulting in a total of  $N_c N_x$  trajectories that approximate the filtering distribution.

### Algorithm 3 Ensemble Sequential Monte Carlo Squared (eSMC<sup>2</sup>)

Operations involving index  $m$  must be performed for  $m=1, \dots, N_\theta$ .

**Inputs:** Observations  $y_{1:T}$ , prior  $p(\theta)$ , Number of ensemble members  $N_x$ , Number of parameter particles:  $N_\theta$ , Number of PMMH moves  $R$

**Output:** Parameter particles:  $\{\theta_{0:t}^m\}_{m=1}^{N_\theta}$

- 1: Initialize  $\theta_0^m \sim p(\theta)$ , weights  $\omega_0^m = 1/N_\theta$ , for  $m=1, \dots, N_\theta$
- 2: **for**  $t=1$  to  $T$  **do**
- 3:   Perform iteration  $t$  of EnKF (Algorithm 1) to estimate  $\hat{p}_{\text{enkf}}^{N_x}(y_t|y_{1:t-1}, \theta_t^m)$
- 4:   Update weights:
 
$$\tilde{\omega}_t^m = \omega_{t-1}^m \times \hat{p}_{\text{enkf}}^{N_x}(y_t|y_{1:t-1}, \theta_t^m)$$
- 5:   Normalize weights:  $\omega_t^m = \tilde{\omega}_t^m / \sum_{j=1}^{N_\theta} \tilde{\omega}_t^j$  and compute  $\text{ESS}_\theta = 1 / \sum_m^{N_\theta} (\omega_t^m)^2$
- 6:   **if**  $\text{ESS}_\theta < N_\theta / 2$  **then**
- 7:     Resample parameters  $\{\theta_t^m\}$  with replacement according to  $\{\omega_t^m\}$
- 8:     Reset weights  $\omega_t^m = 1/N_\theta$
- 9:     Perform PMMH move (Algorithm 2) for each  $\theta_t^m$
- 10:   **end if**
- 11: **end for**

It is important to note a key theoretical distinction between SMC<sup>2</sup> and eSMC<sup>2</sup>. Standard SMC<sup>2</sup> is “exact”: as the number of state and parameter particles  $N_x, N_\theta \rightarrow \infty$ , it converges to the true posterior  $p(\theta|y_{1:t})$ . In contrast, eSMC<sup>2</sup> is generally biased because the EnKF approximates the incremental likelihood  $p(y_t|y_{1:t-1}, \theta)$  via a Gaussian assumption. Nevertheless, the EnKF likelihood typically has low variance [22], and unbiased corrections are possible under mild conditions [23]. Empirical studies show that EnKF-based likelihoods yield accurate and stable inference for both linear and moderately nonlinear SSMS [23,24]. Thus, while not formally exact, eSMC<sup>2</sup> provides practically accurate posterior estimates at much lower computational cost.

## 3. Results

### 3.1. Inference on a simulated epidemic

To assess the performance of the proposed eSMC<sup>2</sup> algorithm, we first conducted three simulation studies using the SEIR model from Section 2. All algorithms were implemented in Python and executed on a desktop computer equipped with an Intel Core i7-1300H processor (3.40 GHz). Unless otherwise stated, simulations assumed a population size of  $N=500,000$  with initial conditions  $I_0 = 10$ ,  $S_0 = N - I_0$ , and  $E_0 = R_0 = 0$ . Daily incidence  $y_t$  was generated according to the Poisson observation model in Eq (3), assuming perfect case reporting ( $\rho = 1$ ).

The three examples differ primarily in the temporal structure of the transmission rate  $\beta_t$  and the number of observation points. Example 1 depicts a moderately time-varying epidemic with an initial surge followed by a mild resurgence, capturing short-term fluctuations in transmission. Example 2, by contrast, exhibits a smoother and more prolonged pattern

that sustains infections for longer and decays more gradually. The latter setting was designed to emulate the dynamic observed in the 2022 U.S. mpox epidemic, where behavioral changes and intermittent control measures led to prolonged and uneven declines in case counts. Example 3 is motivated by epidemics such as COVID-19 or seasonal influenza that exhibit secondary infection waves, arising from imperfect immunity or the activation and deactivation of control policies. To capture this behavior, we employ an oscillatory transmission rate that generates recurrent epidemic waves of declining amplitude, providing a more challenging multi-wave setting in which to evaluate algorithmic performance.

**Example 1.** Synthetic data were generated over  $T=60$  time points with time-varying transmission rate  $\beta_t = 0.3 \cdot \exp(\sin(2\pi t/55) - t/80)$ . The incubation and infectious periods were set to  $\alpha = 1/2$  and  $\gamma = 1/7$ , respectively. Each state ensemble was initialized using truncated normal distributions:  $S_0 \sim \mathcal{TN}_{[0, \infty]}(5 \times 10^5, 0.2^2)$ ,  $I_0 \sim \mathcal{TN}_{[0, \infty]}(10, 0.2^2)$  and  $E_0 = R_0 = 0$ . The priors were specified as  $\beta_0 \sim \mathcal{N}(0.3, 0.01^2)$ ,  $\alpha \sim \mathcal{TN}_{[0, \infty]}(0.6, 0.3^2)$ ,  $\gamma \sim \mathcal{TN}_{[0, \infty]}(0.2, 0.1^2)$ ,  $\nu_\beta \sim \mathcal{U}(0, 0.5)$ .


**Example 2.** Synthetic data were generated over  $T=100$  time points with a time-varying transmission rate  $\beta_t = 0.5 \cdot \exp(- (t-15)^2/20^2) + 0.065$ , and fixed parameters  $\alpha = 1/3$  and  $\gamma = 1/8$ . The priors were specified as  $\beta_0 \sim \mathcal{N}(0.35, 0.01^2)$ ,  $\alpha \sim \mathcal{TN}_{[0, \infty]}(0.4, 0.2^2)$ ,  $\gamma \sim \mathcal{TN}_{[0, \infty]}(0.12, 0.2^2)$ ,  $\nu_\beta \sim \mathcal{U}(0, 0.3)$ . Initial state ensemble members were identical to Example 1.

**Example 3.** Synthetic data were generated over  $T=100$  time points with an oscillatory, multi-wave transmission rate  $\beta_t = 0.3 \cdot \exp(\cos(2\pi t/70) - t/100)$ , and fixed parameters  $\alpha = 1/2$  and  $\gamma = 1/6$ . The priors were specified as  $\beta_0 \sim \mathcal{N}(0.8, 0.01^2)$ ,  $\alpha \sim \mathcal{TN}_{[0, \infty]}(0.6, 0.3^2)$ ,  $\gamma \sim \mathcal{TN}_{[0, \infty]}(0.17, 0.05^2)$ ,  $\nu_\beta \sim \mathcal{U}(0, 1)$ . A more informative prior was placed on  $\gamma$  relative to Examples 1 and 2, as the oscillatory structure of  $\beta_t$  introduces identifiability challenges that make the recovery rate harder to pin down from incidence data alone. Initial state ensemble members were identical to Examples 1 and 2.

For each simulation study, we compared the proposed eSMC<sup>2</sup> algorithm with the standard SMC<sup>2</sup> implementation that employs a bootstrap particle filter (BPF) in its inner layer (see Appendix A in [S1 Text](#)). In all cases, we used  $N_\theta = 1000$  parameter particles and  $N_x = 200$  ensemble members (or particles). For simplicity and comparability, the value of  $N_x$  was kept fixed throughout the timeline. Parameter rejuvenation was triggered when the effective sample size of the parameter particles dropped below 50% of  $N_\theta$ , followed by five iterations of the PMMH kernel.


Filtered estimates in [Figs 3–5](#) correspond to a single representative run of each method. The results illustrate that both eSMC<sup>2</sup> (blue) and SMC<sup>2</sup> (orange) accurately capture the epidemic dynamics in Examples 1, 2, and 3. For both methods, the predicted incidence closely follows the observed data, with 95% credible intervals consistently encompassing the actual case counts, indicating reliable uncertainty quantification. Estimates of the time-varying transmission rate  $\beta_t$  and the effective reproduction number  $R_{\text{eff}}(t)$  are broadly consistent across methods and closely align with the ground truth, demonstrating that both approaches yield comparable assessments of the evolving transmission dynamics. The wider credible intervals for  $R_{\text{eff}}(t)$  relative to the filtered incidence reflect two complementary sources of uncertainty: the roughness of the random walk prior on  $\beta_t$ , and parameter non-identifiability, particularly during the early epidemic phase when different combinations of  $\beta_t$ ,  $\alpha$ , and  $\gamma$  can yield indistinguishable incidence curves. This wider uncertainty is therefore an accurate signal that incidence data alone cannot sharply identify instantaneous transmission intensity. Smoother priors on  $\beta_t$ , such as integrated Brownian motion or higher-order random walks, could potentially mitigate this variability but were not considered here and remain an avenue for future work. The filtering estimates of the unobserved latent state also closely track the true latent trajectories, with their credible intervals providing appropriate coverage (see Appendix C in [S1 Text](#)).

[Figs 6–8](#) illustrate the temporal evolution of the posterior means of the inferred parameters, along with their marginal posterior densities at the final time step, each summarizing five independent runs. Across runs, both SMC<sup>2</sup> and eSMC<sup>2</sup> accurately recover the ground truth. For the latency and recovery rates ( $\alpha$  and  $\gamma$ ), the posterior means obtained from both algorithms are closely aligned, with credible intervals narrowing steadily as more data are assimilated, indicating



**Fig 3. Example 1: Filtered estimates of simulated incidence, transmission rate, and effective reproduction number.** Solid orange lines (SMC<sup>2</sup>) and dashed blue lines (eSMC<sup>2</sup>) show posterior means; shaded areas indicate 95% credible intervals. Observed incidence is shown as black dots, and the black dashed line represents the ground-truth latent quantity.


<https://doi.org/10.1371/journal.pcbi.1014301.g003>



**Fig 4. Example 2: Filtered estimates of simulated incidence, transmission rate, and effective reproduction number.** Solid orange lines (SMC<sup>2</sup>) and dashed blue lines (eSMC<sup>2</sup>) show posterior means; shaded areas indicate 95% credible intervals. Observed incidence is shown as black dots, and the black dashed line represents the ground-truth latent quantity.

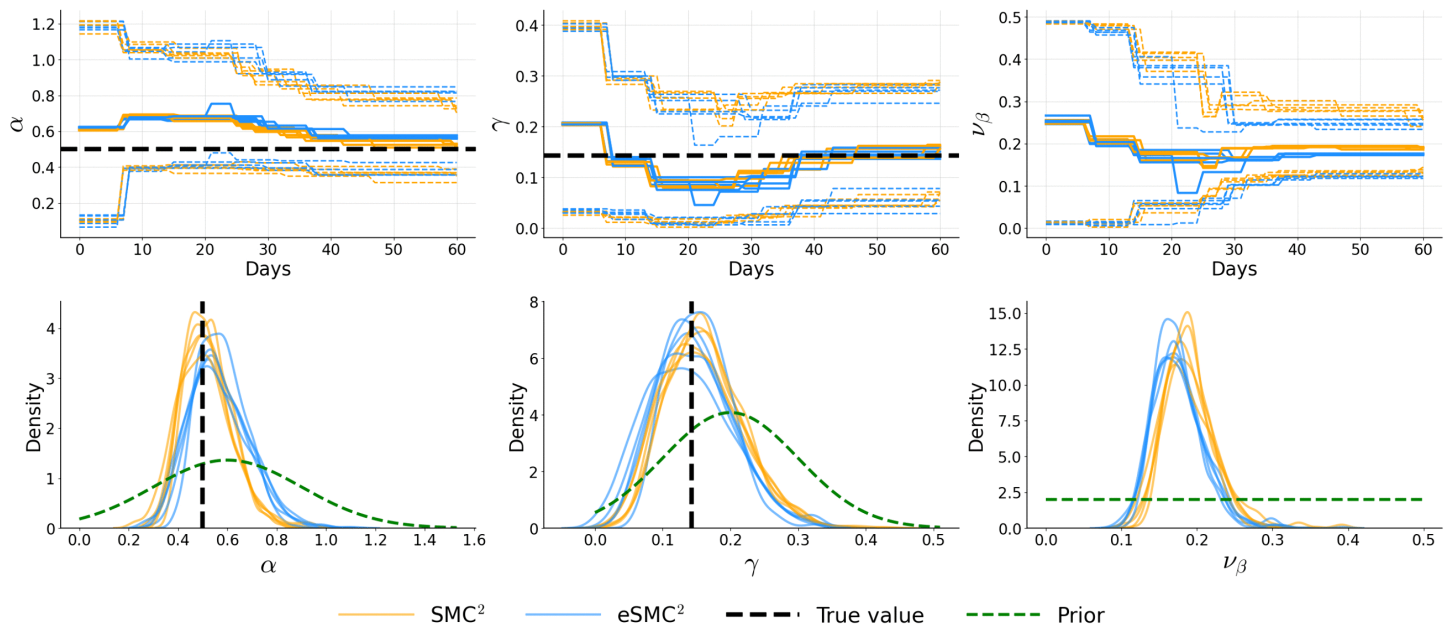
<https://doi.org/10.1371/journal.pcbi.1014301.g004>

progressive information gain. The only notable exception is the diffusion parameter  $\nu_\beta$  in Examples 2 and 3, where a slightly larger discrepancy is observed. This difference likely arises from the stochastic nature of the log-transmission process combined with the approximation introduced by the EnKF-based likelihood in eSMC<sup>2</sup>. Importantly, these differences do not appear to introduce substantial systematic bias or distort the inferred  $\beta_t$  trajectory. The joint posterior contour plots (see Appendix C in [S1 Text](#)) further reveal strong agreement in posterior geometry and dependence structure: the contours for SMC<sup>2</sup> and eSMC<sup>2</sup> largely overlap, are centered near the true values, and display similar correlation patterns among parameters. The only visible difference appears along the  $\nu_\beta$  dimension. Additional sensitivity analyses (Appendix D in [S1 Text](#)) assessing the effect of the number of state particles ( $N_\chi$ ) and the reporting fraction ( $\rho$ ) similarly demonstrate that posterior inferences are robust to ensemble size and exhibit only moderate variation when  $\rho$  is correctly specified.



**Fig 5. Example 3: Filtered estimates of simulated incidence, transmission rate, and effective reproduction number.** Solid orange lines (SMC<sup>2</sup>) and dashed blue lines (eSMC<sup>2</sup>) show posterior means; shaded areas indicate 95% credible intervals. Observed incidence is shown as black dots, and the black dashed line represents the ground-truth latent quantity.

<https://doi.org/10.1371/journal.pcbi.1014301.g005>

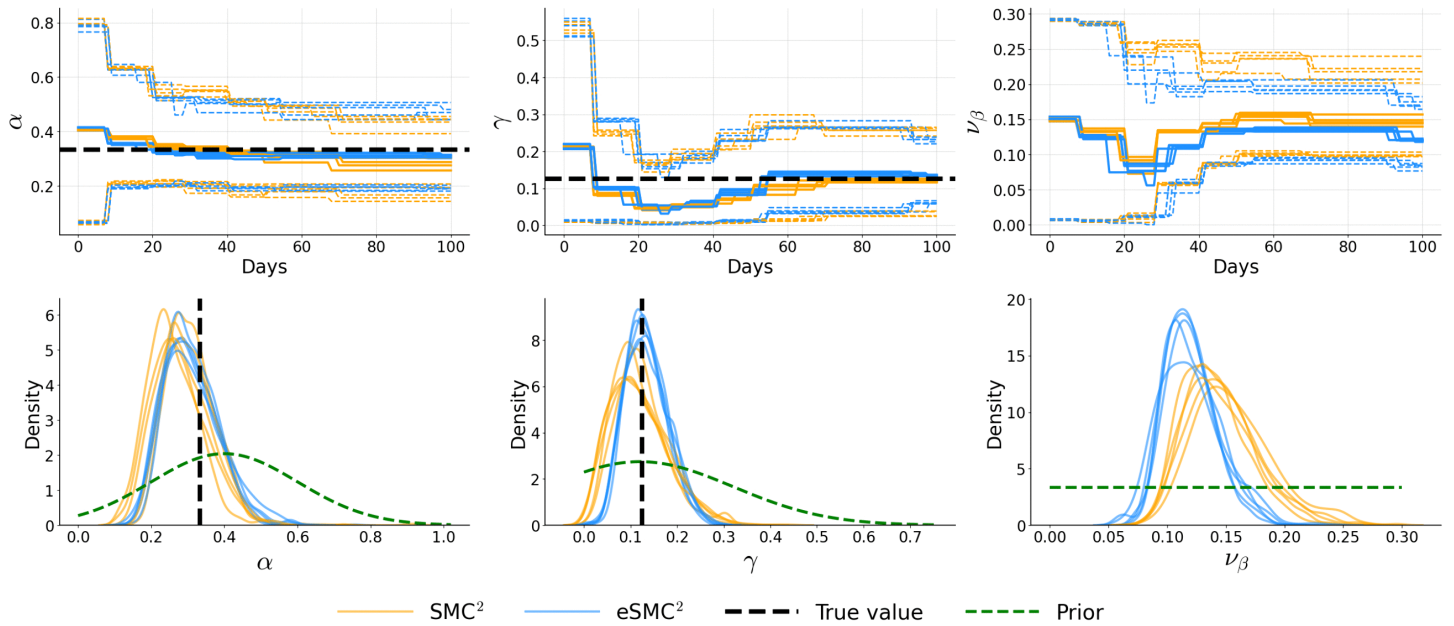


**Fig 6. Example 1: Posterior distributions of  $\alpha$ ,  $\gamma$ , and  $\nu_\beta$  for five independent runs.** Top row shows filtered means with 95% credible intervals. Bottom row shows marginal posterior distributions at  $T=60$ , with prior distributions overlaid (green dashed lines). Black dashed lines indicate true parameter values.

<https://doi.org/10.1371/journal.pcbi.1014301.g006>

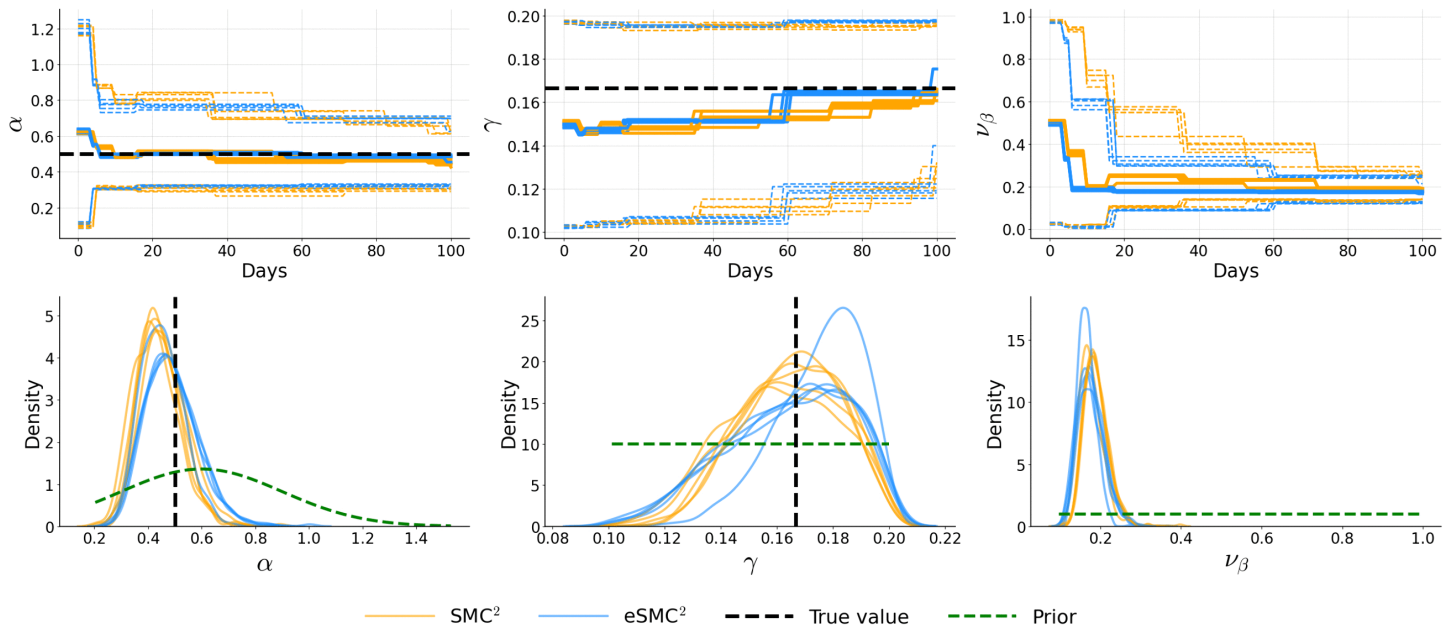
The analyses also investigate the impact of non-informative (flat) priors, demonstrating that posterior estimates remain stable under weakly informative prior specifications.

Results in [Table 1](#) summarize the accuracy and computational cost of each method, averaged over five independent runs. The mean absolute error (MAE) and root mean squared error (RMSE) quantify discrepancies between the filtering



**Fig 7. Example 2: Posterior distributions of  $\alpha$ ,  $\gamma$ , and  $\nu_\beta$  for five independent runs.** The top row shows filtered means with 95% credible intervals. Bottom row shows marginal posterior distributions at  $T=100$ , with prior distributions overlaid (green dashed lines). Black dashed lines indicate true parameter values.

<https://doi.org/10.1371/journal.pcbi.1014301.g007>



**Fig 8. Example 3: Posterior distributions of  $\alpha$ ,  $\gamma$ , and  $\nu_\beta$  for five independent runs.** The top row shows filtered means with 95% credible intervals. Bottom row shows marginal posterior distributions at  $T=100$ , with prior distributions overlaid (green dashed lines). Black dashed lines indicate true parameter values.

<https://doi.org/10.1371/journal.pcbi.1014301.g008>

**Table 1. Comparison of goodness-of-fit and parameter estimates. Each method is evaluated based on CPU time in seconds (with standard deviation) and MAE (RMSE in parentheses) for daily incidence, effective reproduction number, transmission rate and model static parameters. Results are shown for Examples 1, 2, and 3.**

Method	CPU (sd)	Incidence ( $\rho Z_t$ )	$R_{\text{eff}}(t)$	$\beta_t$	$\hat{E}[\alpha y_{1:T}]$	$\hat{E}[\gamma y_{1:T}]$
<b>Example 1</b>						
SMC <sup>2</sup>	2277±145	3.257 (4.691)	0.297 (0.455)	0.044 (0.063)	0.080 (0.101)	0.047 (0.060)
eSMC <sup>2</sup>	467±44	3.067 (4.156)	0.309 (0.444)	0.043 (0.062)	0.103 (0.134)	0.046 (0.057)
<b>Example 2</b>						
SMC <sup>2</sup>	3929±433	4.652 (6.536)	0.266 (0.450)	0.032 (0.050)	0.072 (0.087)	0.049 (0.061)
eSMC <sup>2</sup>	639±48	4.780 (6.697)	0.233 (0.386)	0.033 (0.053)	0.064 (0.078)	0.036 (0.046)
<b>Example 3</b>						
SMC <sup>2</sup>	3766±464	2.610 (3.752)	0.308 (0.467)	0.049 (0.073)	0.055 (0.099)	0.003 (0.019)
eSMC <sup>2</sup>	543±33	2.444 (3.387)	0.280 (0.393)	0.047 (0.065)	0.020 (0.098)	0.004 (0.021)

<https://doi.org/10.1371/journal.pcbi.1014301.t001>

means and the ground truth for both selected latent state ( $Z_t, \beta_t$ ) and static parameters ( $\alpha, \gamma$ ). Across both examples, eSMC<sup>2</sup> achieves error magnitudes closely matching those of SMC<sup>2</sup>, indicating that the ensemble approximation preserves inferential accuracy. In terms of computational effort, eSMC<sup>2</sup> offers a substantial advantage. In Example 1, the average CPU time is approximately 467 seconds for eSMC<sup>2</sup>, compared with 2277 seconds for SMC<sup>2</sup>, corresponding to a reduction by a factor of about 4.87. Similarly, in Example 2, eSMC<sup>2</sup> reduces average runtime to 639 seconds versus 3929 seconds for SMC<sup>2</sup>, a factor of roughly 6.1. In Example 3, eSMC<sup>2</sup> completes in approximately 543 seconds on average, compared with 3766 seconds for SMC<sup>2</sup>, corresponding to a speedup factor of roughly 6.9. The variability in the CPU across runs in both methods reflects the stochastic occurrence of rejuvenation steps, which introduce random fluctuations in computing time.

As an additional benchmark, we evaluated eSMC<sup>2</sup> against the Liu and West filter [37] (using 10<sup>4</sup> particles, runtime ≈9 minutes) and a long PMCMC run (10<sup>4</sup> iterations, runtime ≈3 hours) as a gold-standard reference, both serving as alternative approaches for joint sequential state-parameter estimation. Results provided in Appendix E in S1 Text demonstrate that eSMC<sup>2</sup> achieves superior parameter accuracy and produces narrower, more stable credible intervals compared to the Liu and West filter, while remaining in close agreement with the PMCMC posterior distributions evaluated at the last time point.

### 3.2. Application to the 2022 monkeypox outbreak in the USA

We applied the proposed methodology to infer the transmission dynamics of the 2022 U.S. monkeypox (mpox) outbreak, caused by the West African clade II of the monkeypox virus. This outbreak represented an unprecedented global event, with the United States reporting among the highest case counts worldwide. The typical incubation period ranges from 3–21 days, and infectiousness generally lasts 2–4 weeks. The dataset consists of daily confirmed mpox case counts reported by the U.S. Centers for Disease Control and Prevention (CDC) from 6 May to 31 December 2022. Recent studies have applied Bayesian filtering to mpox data. For instance, Saldaña et al. [38] estimated the instantaneous reproduction number and growth rates using sequential Bayesian updating, while Papageorgiou and Koliass [7] proposed a particle filter incorporating penalty factors specific to mpox transmission. A related stochastic Kalman-based approach, validated on the 2022 Czech Republic mpox data, further discusses the challenges of Bayesian filtering for epidemic modeling with dynamic parameters [39]. Our framework offers a complementary perspective by enabling scalable joint inference of latent epidemic states and model parameters, allowing uncertainty quantification directly from noisy daily surveillance data.

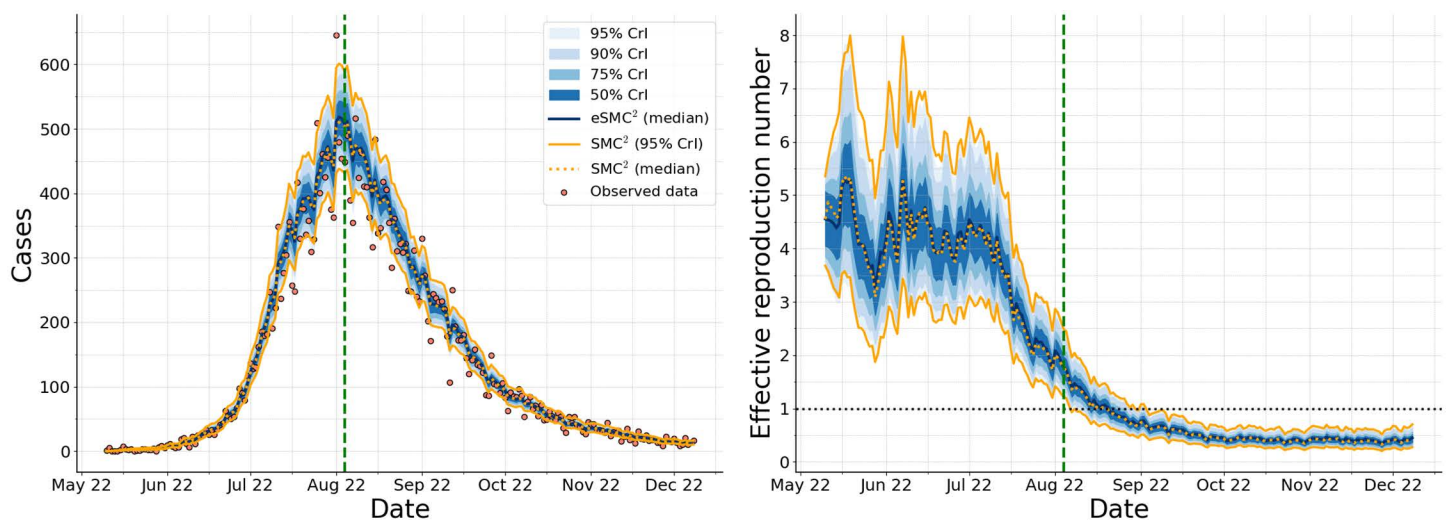
Given the pronounced day-to-day variability in the reported case counts, we adopted a Negative Binomial observation model as the primary specification. In contrast, preliminary analyses under a Poisson observation model led to unstable

posterior trajectories and parameter degeneracy, indicating a poor fit to the observed overdispersion. The Poisson assumption becomes more appropriate when incidence is aggregated over 7-day intervals, which reduces temporal variability. Supporting results for the aggregated analysis are provided in Appendix F in [S1 Text](#).

The algorithm was implemented with  $N_x = 200$  ensemble members and  $N_\theta = 1000$  parameter particles, providing a balance between stability and computational efficiency. All primary results presented below are obtained using eSMC<sup>2</sup>, which constitutes the main inference framework of this study.

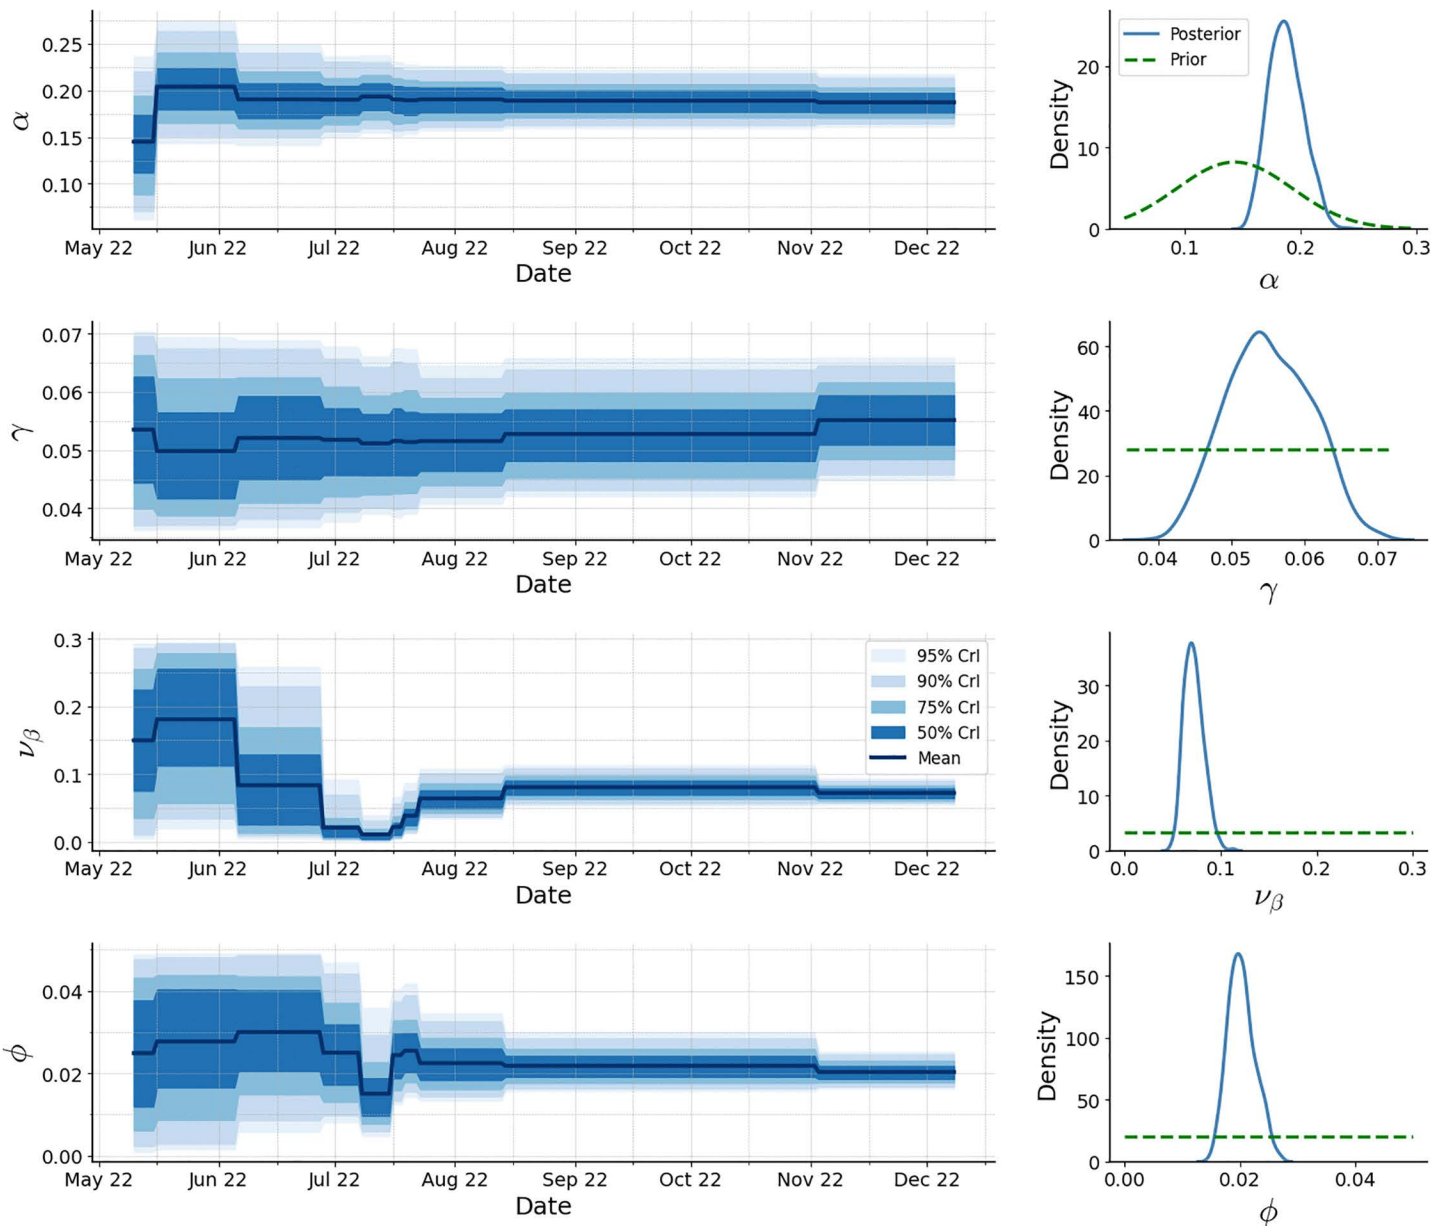
[Fig 9](#) shows the estimated daily incidence and  $R_{\text{eff}}(t)$  over time. The filtered trajectories closely follow the observed case counts, while the credible bands reflect uncertainty arising from both stochastic epidemic dynamics and observation noise. The estimated  $R_{\text{eff}}(t)$  indicates sustained transmission above one during the early phase of the outbreak, followed by a decline after the declaration of the state of emergency, and eventual stabilization below one as the outbreak subsides. For benchmarking purposes, we additionally report results obtained using the standard SMC<sup>2</sup> algorithm under identical model and prior specifications. In [Fig 9](#), SMC<sup>2</sup> results are shown using 95% credible interval bounds (solid orange lines) together with the posterior median trajectory (dotted orange line), enabling a direct comparison of both central estimates and uncertainty quantification between the two methods.

[Fig 10](#) presents the posterior trajectories of key epidemiological parameters, with corresponding priors and posterior summaries reported in [Table 2](#). As more data are assimilated, these parameters become increasingly well-identified, as reflected by the progressive narrowing of the posterior intervals. The analysis indicates a mean incubation period of approximately 5.3 days (95% CrI: 4.6–6.2 days), consistent with early empirical estimates reported by Madewell et al. [40], and a mean infectious period of approximately 18.2 days (95% CrI: 15.2–22.2 days). The overdispersion parameter  $\phi$  exhibits posterior mass distinctly separated from zero, confirming substantial extra-Poisson variability. For comparison, we also implemented the standard SMC<sup>2</sup> algorithm under identical model and prior settings. Results reported in Appendix F in [S1 Text](#) show that both methods produce a very similar evolution of the filtering distributions for the parameters,



**Fig 9. Inference of daily incidence and effective reproduction number for the 2022 U.S. monkeypox outbreak.** Left: filtered estimates of daily incidence. The solid blue line denotes the posterior median obtained using eSMC<sup>2</sup>, with reported case counts shown as red dots. Right: inferred effective reproduction number  $R_{\text{eff}}(t)$ . For eSMC<sup>2</sup>, shaded regions represent 50%, 75%, 90%, and 95% credible intervals. The vertical dashed line indicates the declaration of the national state of emergency. Results from SMC<sup>2</sup> are shown in orange, with solid lines indicating the 95% credible interval bounds and the dotted line representing the posterior median.

<https://doi.org/10.1371/journal.pcbi.1014301.g009>



**Fig 10. Posterior distributions of key epidemiological parameters.** The left column shows filtered means with credible intervals; the right column displays marginal posterior distributions at the last time step.

<https://doi.org/10.1371/journal.pcbi.1014301.g010>

indicating that the ensemble-based approximation in eSMC<sup>2</sup> accurately captures posterior uncertainty. Notably, a single run with these settings required approximately 16 minutes of CPU time, compared with nearly 2.5 hours for the equivalent SMC<sup>2</sup> implementation, demonstrating the efficiency gains achieved by the EnKF-based approach.

[Fig 11](#) presents 14-day-ahead posterior-predictive forecasts initiated at multiple time points during the outbreak, assuming the transmission rate  $\beta_t$  remains fixed at its last estimated value. Forecasts are obtained by propagating the latent state forward through the state-space dynamics and sampling synthetic observations from the observation

**Table 2. Description of the different parameters, priors, and posterior estimates at the final time step. Upper and/or lower bounds have been imposed by the observations.  $\mathcal{TN}$  denotes a truncated normal distribution  $\mathcal{TN}_{[\text{inf}, \text{sup}]}(\text{mean}, \text{std}^2)$ .**

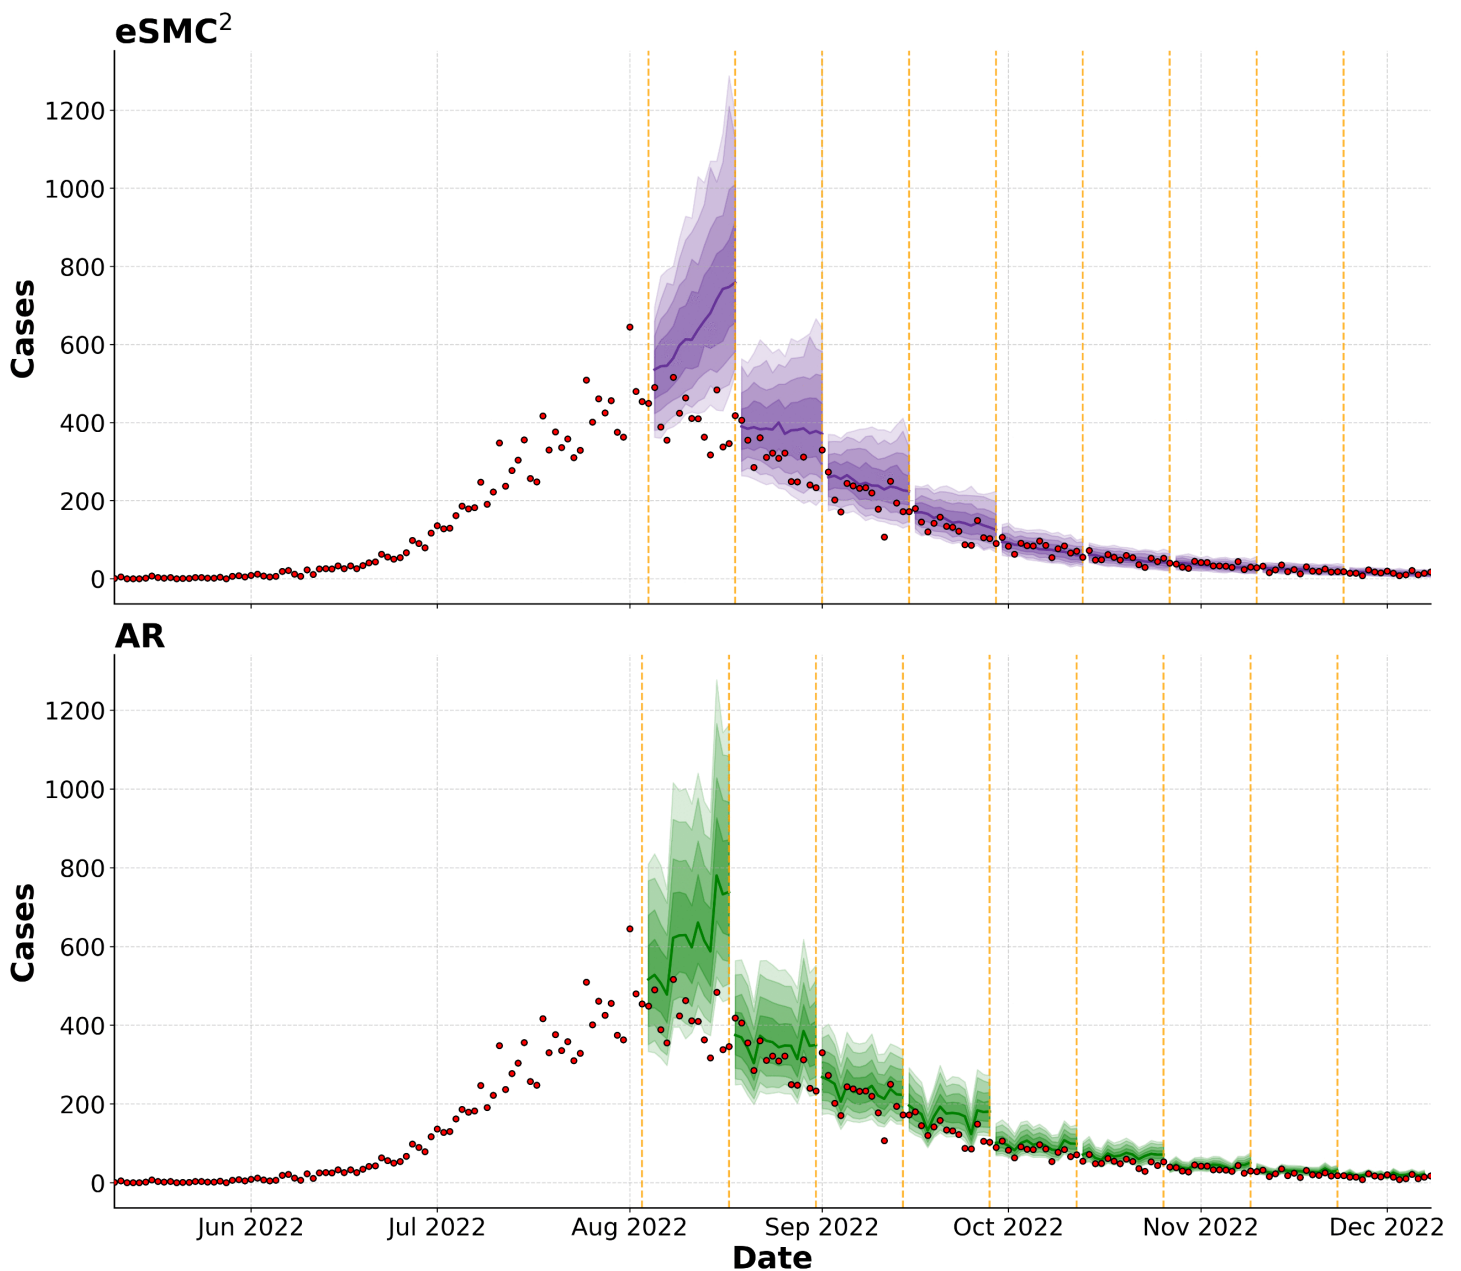
Parameter	Description	Prior/Value	Posterior mean (95% CrI)
$\beta_0$	Initial condition transmission rate	$\mathcal{U}(0.2, 0.3)$	–
$\alpha$	Latency rate ( $1/\alpha$ incubation period)	$\mathcal{TN}_{[1/21, 1/3]}(1/7, 0.05^2)$	0.187 (0.161, 0.218)
$\gamma$	Recovery rate ( $1/\gamma$ infectious period)	$\mathcal{U}(1/28, 1/14)$	0.055 (0.045, 0.066)
$\nu_\beta$	Volatility of the Brownian process	$\mathcal{U}(0, 0.3)$	0.072 (0.055, 0.094)
$\phi$	Overdispersion	$\mathcal{U}(0, 0.05)$	0.020 (0.016, 0.025)

<https://doi.org/10.1371/journal.pcbi.1014301.t002>

model. The purpose of this analysis is to evaluate the short-term predictive performance of the proposed eSMC<sup>2</sup> framework. We compare the results with an autoregressive (AR) model (Appendix F in [S1 Text](#)). The top panel of [Fig 11](#) shows forecasts from eSMC<sup>2</sup> and the bottom panel shows those from the AR model. Quantitative performance is summarised in [Table 3](#), which reports mean absolute error (MAE), weighted interval score (WIS; based on the 50%, 75%, 90%, and 95% intervals), and empirical 95% prediction interval (PI) coverage for each forecast start date. Both models struggle to anticipate the turning point in early August, and around this period the autoregressive model achieves a lower forecast error. In particular, around the epidemic peak, the AR model consistently outperforms eSMC<sup>2</sup> across all evaluation metrics. This likely reflects the intrinsic difficulty of forecasting near the peak, where rapid changes in transmission dynamics and the assumption of a fixed  $\beta_t$  limit the responsiveness of the mechanistic model. In contrast, the AR model, being purely data-driven, can more readily adapt to short-term trends during this highly dynamic phase. Similar behavior has been reported in related work, where data-driven Hawkes-type processes were found to outperform stochastic SEIR models during abrupt regime shifts [\[41\]](#). However, from mid-September onward, eSMC<sup>2</sup> achieves lower MAE and WIS in most forecast windows and maintains coverage close to the nominal 95% level, whereas the autoregressive model exhibits declining coverage in later periods ([Table 3](#)). These results highlight a limitation of the current modeling setup in capturing abrupt turning points, mainly due to the inertia of the mechanistic dynamics, rather than an inherent limitation of SMC-based inference. At the same time, they confirm the strength of the approach in providing well-calibrated probabilistic forecasts once the epidemic enters a more stable phase. Long-term forecasts for the final epidemic phase are presented in Fig. F4 in [S1 Text](#). The model reproduces the overall decline and achieves full 95% predicted interval coverage.

#### 4. Discussion and limitations

In this work, we investigated sequential Bayesian inference for latent states and parameters in compartmental models observed through incidence data. We proposed the eSMC<sup>2</sup>, an extension of the SMC<sup>2</sup> framework that replaces the particle filter with an EnKF-derived likelihood, enabling computationally efficient sequential updates of both latent epidemic trajectories and model parameters. This approach builds on the eMCMC method of Drovandi et al. [\[23\]](#) but extends it to a fully sequential setting, allowing time-resolved inference in dynamic epidemic systems. By leveraging the ensemble approximation of the likelihood, eSMC<sup>2</sup> reduces the computational burden while retaining the ability to capture nonlinear state evolution. Within the broader landscape of Bayesian calibration approaches, including likelihood-free methods such as ABC [\[14\]](#), eSMC<sup>2</sup> represents a complementary strategy. Likelihood-free approaches are particularly useful when the likelihood is intractable or prohibitively expensive to evaluate. In contrast, eSMC<sup>2</sup> relies on an explicit state-space formulation and an approximate likelihood, enabling principled sequential updating of both latent states and parameters as new data arrive. Consequently, eSMC<sup>2</sup> is especially suited to time-series epidemic models where inference is performed sequentially, whereas likelihood-free methods may be preferable for models that do not admit a tractable state-space representation.



**Fig 11. Posterior-predictive forecasts at different starting dates.** The top panel shows eSMC<sup>2</sup> forecasts, and the bottom panel shows AR forecasts. Median forecasts are displayed in purple (eSMC<sup>2</sup>) and green (AR), with corresponding 50%, 75%, 90%, and 95% predictive intervals. Observed case counts are shown in red. Vertical orange dashed lines indicate the forecast start dates.

<https://doi.org/10.1371/journal.pcbi.1014301.g011>

Methodologically, the EnKF provides a scalable alternative to standard particle filtering, particularly in moderately high-dimensional or computationally constrained settings. While the EnKF likelihood is only a Gaussian approximation, we employ an unbiased correction, following Drovandi et al. [23], which reduces sensitivity to the ensemble size. We further allow the observation noise variance to evolve over time by updating it at each filtering step using ensemble-based state predictions. This adaptive treatment of observation uncertainty allows the filter to better reflect time-varying variability in

**Table 3. Evaluation of the 2 weeks-ahead forecasts of reported mpox cases during the 2022 U.S. outbreak. Bold indicates better-performing model.**

Forecast date	eSMC <sup>2</sup>			AR		
	95% PI coverage	MAE	WIS	95% PI coverage	MAE	WIS
2022-08-04	0.43	230.6	269.5	<b>0.57</b>	<b>204.2</b>	<b>251.6</b>
2022-08-18	<b>1.00</b>	78.5	121.3	<b>1.00</b>	<b>52.1</b>	<b>94.2</b>
2022-09-01	0.86	41.2	74.5	<b>0.93</b>	<b>30.8</b>	<b>60.3</b>
2022-09-15	<b>1.00</b>	<b>25.4</b>	<b>43.9</b>	0.79	41.2	55.5
2022-09-29	<b>1.00</b>	<b>10.7</b>	<b>22.6</b>	0.86	15.2	25.2
2022-10-13	<b>1.00</b>	<b>7.8</b>	<b>15.2</b>	0.79	18.3	22.8
2022-10-27	<b>1.00</b>	<b>5.2</b>	<b>11.2</b>	0.86	7.5	11.7
2022-11-10	<b>1.00</b>	<b>4.6</b>	<b>8.5</b>	0.64	7.3	9.6
2022-11-24	<b>0.93</b>	<b>3.9</b>	<b>6.5</b>	0.57	5.9	7.7

<https://doi.org/10.1371/journal.pcbi.1014301.t003>

incidence data. Although we do not provide formal bias quantification, previous studies indicate that EnKF-based likelihoods yield stable and accurate inference in moderately nonlinear systems when combined with MCMC [23,24]. Our simulation studies confirm this behavior, showing that the approximation error remains small and posterior means closely match those obtained with standard SMC<sup>2</sup>. Importantly, the diffusion-driven SEIR formulation combined with Bayesian filtering provides a flexible mechanism for detecting temporal changes in transmission, such as those resulting from behavioral shifts or intervention measures, without requiring explicit change-point modeling. In our simulations, the method accurately recovers both the transmission rate and the effective reproduction number, even without knowledge of their true trajectories. The application to the 2022 U.S. mpox outbreak further demonstrates that eSMC<sup>2</sup> delivers well-calibrated, near-real-time estimates of transmission dynamics and the effective reproduction number, supporting timely situational awareness and public health decision-making. It should be noted that this forecasting analysis is retrospective and does not account for real-time complexities such as data revisions and reporting delays, so predictive performance may be optimistic relative to an operational setting [42].

Despite its practical and computational advantages, several methodological and theoretical limitations should be noted. Unlike standard SMC<sup>2</sup>, which is theoretically exact in the limit of infinite state and parameter particles, eSMC<sup>2</sup> inherits bias from the EnKF likelihood approximation. This bias may become more pronounced in strongly nonlinear or multimodal epidemic systems, and rigorous convergence guarantees for nonlinear, non-Gaussian settings remain an open challenge. While convergence of the EnKF likelihood has been established in linear-Gaussian systems with rate  $\mathcal{O}(N_x^{-1/2})$  [43], extending such results to nonlinear systems is nontrivial. Although the proposed method is sequential, it is not fully “online”, as each call of the PMMH kernel still involves running the EnKF on all data available up to that point. One potential avenue to further improve computational efficiency would be to develop a fully recursive implementation analogous to a Nested Particle Filter (NPF) [44], in which the inner particle filter could be replaced by an EnKF. From a modeling standpoint, our SEIR application relies on a single data stream (daily incidence). Incorporating additional observations such as hospitalizations, deaths, or serological survey data could improve parameter identifiability, particularly in disentangling reporting rates from transmission parameters [45]. More generally, epidemiological models often include parameters constrained to the unit interval (e.g., reporting or mortality ratios). In such cases, one may either assign a *Beta* prior, which naturally enforces support on (0,1), or apply a *logit* transformation and conduct inference on the unconstrained scale. Finally, while the framework is broadly applicable to other state-space models beyond epidemiology, its performance may vary depending on the degree of nonlinearity, the suitability of the Gaussian approximation for the forecast ensemble and the choice of observation noise. Future research will explore these limitations in more detail.

## Supporting information

**S1 Text. Additional methodological details, algorithms, and extended results** Contains detailed derivation of the SEIR transition density and observation error variance specification (Appendix A), pseudocode for the Bootstrap Particle Filter and stratified resampling procedure (Appendix B), additional simulation results showing unobserved SEIR compartment trajectories and posterior parameter contours (Appendix C), sensitivity analyses examining the impact of ensemble size and prior specifications on parameter recovery (Appendix D), comparison with the Liu and West filter and PMCMC for parameter estimation (Appendix E), and supplementary results from the monkeypox application, including Poisson observation model fits and filtered latent state trajectories (Appendix F).

(PDF)

## Author contributions

**Conceptualization:** Dhorasso Temfack, Jason Wyse.

**Formal analysis:** Dhorasso Temfack, Jason Wyse.

**Funding acquisition:** Jason Wyse.

**Methodology:** Dhorasso Temfack, Jason Wyse.

**Software:** Dhorasso Temfack.

**Supervision:** Jason Wyse.

**Validation:** Dhorasso Temfack, Jason Wyse.

**Writing – original draft:** Dhorasso Temfack, Jason Wyse.

**Writing – review & editing:** Dhorasso Temfack, Jason Wyse.

## References

- Golightly A, Wadkin LE, Whitaker SA, Baggaley AW, Parker NG, Kypraios T. Accelerating Bayesian inference for stochastic epidemic models using incidence data. *Stat Comput.* 2023;33(6). <https://doi.org/10.1007/s11222-023-10311-6>
- Whitaker SA, Golightly A, Gillespie CS, Kypraios T. Sequential Bayesian Inference for Stochastic Epidemic Models of Cumulative Incidence. *Bayesian Anal.* 2025;1(1). <https://doi.org/10.1214/25-ba1547>
- Dureau J, Kalogeropoulos K, Baguelin M. Capturing the time-varying drivers of an epidemic using stochastic dynamical systems. *Biostatistics.* 2013;14(3):541–55. <https://doi.org/10.1093/biostatistics/kxs052> PMID: 23292757
- Birrell PJ, De Angelis D, Presanis AM. Evidence Synthesis for Stochastic Epidemic Models. *Stat Sci.* 2018;33(1):34–43. <https://doi.org/10.1214/17-STS631> PMID: 31975746
- Funk S, Camacho A, Kucharski AJ, Eggo RM, Edmunds WJ. Real-time forecasting of infectious disease dynamics with a stochastic semi-mechanistic model. *Epidemics.* 2018;22:56–61. <https://doi.org/10.1016/j.epidem.2016.11.003> PMID: 28038870
- Ghosh S, Birrell PJ, De Angelis D. An approximate diffusion process for environmental stochasticity in infectious disease transmission modelling. *PLoS Comput Biol.* 2023;19(5):e1011088. <https://doi.org/10.1371/journal.pcbi.1011088> PMID: 37200386
- Papageorgiou VE, Koliass P. A novel epidemiologically informed particle filter for assessing epidemic phenomena. Application to the monkeypox outbreak of 2022. *Inverse Problems.* 2024;40(3):035006. <https://doi.org/10.1088/1361-6420/ad1e2f>
- Andrieu C, Doucet A, Holenstein R. Particle Markov Chain Monte Carlo Methods. *Journal of the Royal Statistical Society Series B: Statistical Methodology.* 2010;72(3):269–342. <https://doi.org/10.1111/j.1467-9868.2009.00736.x>
- Chopin N, Jacob PE, Papaspiliopoulos O. SMC2: An Efficient Algorithm for Sequential Analysis of State Space Models. *Journal of the Royal Statistical Society Series B: Statistical Methodology.* 2012;75(3):397–426. <https://doi.org/10.1111/j.1467-9868.2012.01046.x>
- Golightly A, Kypraios T. Efficient SMC2 schemes for stochastic kinetic models. *Statistics and Computing.* 2018;28(6):1215–30.
- Temfack D, Wyse J. Sequential Monte Carlo Squared for online inference in stochastic epidemic models. *Epidemics.* 2025;52:100847. <https://doi.org/10.1016/j.epidem.2025.100847> PMID: 40850185
- Snyder C, Bengtsson T, Bickel P, Anderson J. Obstacles to High-Dimensional Particle Filtering. *Monthly Weather Review.* 2008;136(12):4629–40. <https://doi.org/10.1175/2008mwr2529.1>

13. Pitt MK, Shephard N. Filtering via Simulation: Auxiliary Particle Filters. *Journal of the American Statistical Association*. 1999;94(446):590–9. <https://doi.org/10.1080/01621459.1999.10474153>
14. Li X, Chadwick F, Swallow B. Advances in approximate Bayesian inference for models in epidemiology. *Epidemics*. 2025;53:100855. <https://doi.org/10.1016/j.epidem.2025.100855> PMID: 40992245
15. Evensen G. Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *J Geophys Res*. 1994;99(C5):10143–62. <https://doi.org/10.1029/94jc00572>
16. Morzfeld M, Hodyss D, Snyder C. What the collapse of the ensemble Kalman filter tells us about particle filters. *Tellus A: Dynamic Meteorology and Oceanography*. 2017;69(1):1283809. <https://doi.org/10.1080/16000870.2017.1283809>
17. Mitchell L, Arnold A. Analyzing the effects of observation function selection in ensemble Kalman filtering for epidemic models. *Math Biosci*. 2021;339:108655. <https://doi.org/10.1016/j.mbs.2021.108655> PMID: 34186054
18. Lal R, Huang W, Li Z. An application of the ensemble Kalman filter in epidemiological modelling. *PLoS One*. 2021;16(8):e0256227. <https://doi.org/10.1371/journal.pone.0256227> PMID: 34411132
19. Sun Q, Miyoshi T, Richard S. Analysis of COVID-19 in Japan with extended SEIR model and ensemble Kalman filter. *J Comput Appl Math*. 2023;419:114772. <https://doi.org/10.1016/j.cam.2022.114772> PMID: 36061090
20. Abbas W, Lee S, Kim S. Joint estimation of hand-foot-mouth disease model and prediction in Korea using the ensemble Kalman filter. *PLoS Comput Biol*. 2025;21(4):e1012996. <https://doi.org/10.1371/journal.pcbi.1012996> PMID: 40245103
21. Stroud JR, Katzfuss M, Wikle CK. A Bayesian Adaptive Ensemble Kalman Filter for Sequential State and Parameter Estimation. *Monthly Weather Review*. 2018;146(1):373–86. <https://doi.org/10.1175/mwr-d-16-0427.1>
22. Katzfuss M, Stroud JR, Wikle CK. Ensemble Kalman Methods for High-Dimensional Hierarchical Dynamic Space-Time Models. *Journal of the American Statistical Association*. 2019;115(530):866–85. <https://doi.org/10.1080/01621459.2019.1592753>
23. Drovandi C, Everitt RG, Golightly A, Prangle D. Ensemble MCMC: Accelerating Pseudo-Marginal MCMC for State Space Models using the Ensemble Kalman Filter. *Bayesian Anal*. 2022;17(1). <https://doi.org/10.1214/20-ba1251>
24. Khalil M, Sarkar A, Adhikari S, Poirel D. The estimation of time-invariant parameters of noisy nonlinear oscillatory systems. *Journal of Sound and Vibration*. 2015;344:81–100. <https://doi.org/10.1016/j.jsv.2014.10.002>
25. Ebeigbe D, Berry T, Schiff SJ, Sauer T. Poisson Kalman filter for disease surveillance. *Phys Rev Res*. 2020;2(4):043028. <https://doi.org/10.1103/physrevresearch.2.043028> PMID: 39211287
26. Golightly A, Heaps SE, Sherlock C, Wadkin LE, Wilkinson DJ. Nested ensemble Kalman filter for static parameter inference in nonlinear state-space models. In: 2025. <https://arxiv.org/abs/251121497>
27. Cazelles B, Champagne C, Dureau J. Accounting for non-stationarity in epidemiology by embedding time-varying parameters in stochastic models. *PLoS Comput Biol*. 2018;14(8):e1006211. <https://doi.org/10.1371/journal.pcbi.1006211> PMID: 30110322
28. Kloeden PE. Stochastic Differential Equations. *International Encyclopedia of Statistical Science*. Springer. 2011. p. 1520–1.
29. Corbella A, McKinley TJ, Birrell PJ, De Angelis D, Presanis AM, Roberts GO, et al. The lifebelt particle filter for robust estimation from low-valued count data. *FoDS*. 2025;7(4):1062–88. <https://doi.org/10.3934/fods.2024052>
30. Asfaw K, Park J, King AA, Ionides EL. A tutorial on spatiotemporal partially observed Markov process models via the R package spatPomp. *arXiv preprint*. 2024.
31. Ghurye SG, Olkin I. Unbiased Estimation of Some Multivariate Probability Densities and Related Functions. *Ann Math Statist*. 1969;40(4):1261–71. <https://doi.org/10.1214/aoms/1177697501>
32. Price LF, Drovandi CC, Lee A, Nott DJ. Bayesian Synthetic Likelihood. *Journal of Computational and Graphical Statistics*. 2017;27(1):1–11. <https://doi.org/10.1080/10618600.2017.1302882>
33. Del Moral P, Doucet A, Jasra A. Sequential Monte Carlo Samplers. *Journal of the Royal Statistical Society Series B: Statistical Methodology*. 2006;68(3):411–36. <https://doi.org/10.1111/j.1467-9868.2006.00553.x>
34. Wu J, Wen L, Green PL, Li J, Maskell S. Ensemble Kalman filter based sequential Monte Carlo sampler for sequential Bayesian inference. *Stat Comput*. 2022;32(1). <https://doi.org/10.1007/s11222-021-10075-x>
35. Chopin N, Papaspiliopoulos O, et al. An introduction to sequential Monte Carlo. vol. 4. Springer; 2020.
36. Steyn N, Parag KV, Thompson RN, Donnelly CA. A Primer on Inference and Prediction With Epidemic Renewal Models and Sequential Monte Carlo. *Stat Med*. 2025;44(18–19):e70204. <https://doi.org/10.1002/sim.70204> PMID: 40772730
37. Liu J, West M. Combined Parameter and State Estimation in Simulation-Based Filtering. *Sequential Monte Carlo Methods in Practice*. Springer New York. 2001. p. 197–223. [https://doi.org/10.1007/978-1-4757-3437-9\\_10](https://doi.org/10.1007/978-1-4757-3437-9_10)
38. Saldaña F, Daza-Torres ML, Aguiar M. Data-driven estimation of the instantaneous reproduction number and growth rates for the 2022 monkeypox outbreak in Europe. *PLoS One*. 2023;18(9):e0290387. <https://doi.org/10.1371/journal.pone.0290387> PMID: 37703247
39. Papageorgiou VE, Vasiliadis G, Tsaklidis G. A new method for the estimation of stochastic epidemic descriptors reinforced by Kalman-based dynamic parameter estimation. Application to mpox data. *Math Biosci*. 2025;381:109365. <https://doi.org/10.1016/j.mbs.2024.109365> PMID: 39672537

40. Madewell ZJ, Charniga K, Masters NB, Asher J, Fahrenwald L, Still W, et al. Serial interval and incubation period estimates of monkeypox virus infection in 12 jurisdictions, United States, May–August 2022. *Emerging Infectious Diseases*. 2023;29(4):818.
41. Temfack D, Wyse J. A sequential ensemble approach to epidemic modeling: combining Hawkes and SEIR models using SMC2. In: 2025. <https://doi.org/doi.org/250615511>
42. Charniga K, Madewell ZJ, Masters NB, Asher J, Nakazawa Y, Spicknall IH. Nowcasting and forecasting the 2022 U.S. mpox outbreak: Support for public health decision making and lessons learned. *Epidemics*. 2024;47:100755. <https://doi.org/10.1016/j.epidem.2024.100755> PMID: [38452454](https://pubmed.ncbi.nlm.nih.gov/38452454/)
43. Chen Y, Sanz-Alonso D, Willett R. Autodifferentiable Ensemble Kalman Filters. *SIAM Journal on Mathematics of Data Science*. 2022;4(2):801–33. <https://doi.org/10.1137/21m1434477>
44. Crisan D, Míguez J. Nested particle filters for online parameter estimation in discrete-time state-space Markov models. *Bernoulli*. 2018;24(4A). <https://doi.org/10.3150/17-bej954>
45. Swallow B, Birrell P, Blake J, Burgman M, Challenor P, Coffeng LE, et al. Challenges in estimation, uncertainty quantification and elicitation for pandemic modelling. *Epidemics*. 2022;38:100547. <https://doi.org/10.1016/j.epidem.2022.100547> PMID: [35180542](https://pubmed.ncbi.nlm.nih.gov/35180542/)