



Parameterizing Lognormal state space models using moment matching

John W. Smith^{1,4} · R. Quinn Thomas^{2,3} · Leah R. Johnson^{3,4}

Received: 19 December 2022 / Revised: 12 May 2023 / Accepted: 19 June 2023
© The Author(s) 2023

Abstract

In ecology, it is common for processes to be bounded based on physical constraints of the system. One common example is the positivity constraint, which applies to phenomena such as duration times, population sizes, and total stock of a system's commodity. In this paper, we propose a novel method for parameterizing Lognormal state space models using an approach based on moment matching. Our method enforces the positivity constraint, allows for arbitrary mean evolution and variance structure, and has a closed-form Markov transition density which allows for more flexibility in fitting techniques. We discuss two existing Lognormal state space models and examine how they differ from the method presented here. We use 180 synthetic datasets to compare the forecasting performance under model misspecification and assess the estimation of precision parameters between our method and existing methods. We find that our models perform well under misspecification, and that fixing the observation variance both helps to improve estimation of the process variance and improves forecast performance. To test our method on a difficult problem, we compare the predictive performance of two Lognormal state space models in predicting the Leaf Area Index over a 151 day horizon by using a process-based ecosystem model to describe the temporal dynamics. We find that our moment matching model performs better than its competitor, and is better suited for intermediate predictive horizons. Overall, our study helps to inform practitioners about the importance of incorporating sensible dynamics when using models of complex systems to predict out-of-sample.

Keywords Bayesian statistics · Forecasting · MCMC · Particle filter · State space model

Handling Editor: Luiz Duczmal.

R. Quinn Thomas and Leah R. Johnson have contributed equally to this work.

Extended author information available on the last page of the article

1 Introduction

Process-based models are mathematical representations of the evolution of biological or physical systems (Buck-Sorlin 2013). These models are often comprised of systems of ordinary or partial differential equations in time and space, or are discretizations of such systems. Because they encapsulate known or hypothesized mechanisms of physical or biological systems, process-based modeling approaches have advantages over empirical or phenomenological modeling approaches, particularly when low data availability limit the ability of empirical models to accurately represent complex processes. As a result, process-based models remain common in ecological forecasting applications (Lewis et al. 2022). However, process-based modeling approaches have their own set of challenges. For example, quantification of uncertainty in both model dynamics and structure is difficult, as is assessing over-parameterization (Luo et al. 2009).

Quantifying uncertainty in process-based models is one of the most challenging tasks when using them for forecasting applications that involve uncertainty propagation. Uncertainty comes from a wide variety of sources, including but not limited to: process, measurement, initial conditions, driver data, and parameter estimation (Dietze et al. 2018). Process uncertainty (or process stochasticity) is particularly important to address, as it acknowledges that the modeling framework may contain unknown errors that are best represented stochastically or contain elements that are known to be non-deterministic and that affect the dynamics of the biological process. For this reason, the state space modeling (SSM) framework (Durbin and Koopman 2012; Petris et al. 2009; Auger-Méthé et al. 2021) has been used frequently in ecological applications (see Thomas et al. 2017; Dowd and Meyer 2003; Dennis et al. 2006, for examples). State space models provide a flexible framework that is able to handle missing data and partition multiple sources of uncertainty (Dietze et al. 2018). Many existing fitting methods for ecosystem process-based models already account for initial condition, parameter, and observation uncertainty. These models can also add an uncertainty structure to describe stochastic elements in the process dynamics. This allows them to be analyzed as state space models.

One of the nuances of including uncertainty/stochasticity in the process model is the trade-off between complexity and ecological coherence. For example, in forest carbon modeling, Gaussian error, with its positive to negative infinity bounds, is commonly assumed for carbon stocks (pools) (Thomas et al. 2017; Jiang et al. 2018), despite the biophysical impossibility of an ecosystem having a negative amount of carbon. In general, biological processes may have well defined lower bounds (and potentially upper bounds) that are not accurately captured by error structures that have support over the entire real line. Consequently, models that have biophysically inadequate error structures can produce nonsensical predictions as the states approach these well defined bounds, or if the observations have large measurement error which can allow for predictions with negative values as the modeled states approach the lower bounds. This is especially relevant in forecasting applications, where the uncertainty compounds as the forecast horizon increases.

A wealth of probability distributions are available for modeling non-negative processes. In particular, the Lognormal distribution has a rich history in ecology (Dennis and Patil 1988), and is frequently used in state space modeling of populations (e.g., Buckland et al. 2004; Dennis et al. 2006; Knape et al. 2011) and species abundance (e.g, Maunder et al. 2015; Mäntyniemi et al. 2015). Two common formulations for Lognormal state space models (LN-SSMs) are the stochastic Gompertz SSM (Gompertz 1825) and the stochastic Moran-Ricker SSM (Ricker 1954; Dennis et al. 2006). In their simplest forms, these models include assumptions that may not accurately describe some ecological processes—namely the assumption of density dependent variance (Dennis et al. 2006) and systematic bias in the process evolution and observation functions. When these assumptions are not appropriate for an application, we need a mechanism to insert more appropriate assumptions about the process evolution, observation, and variance dynamics. To incorporate ecologically coherent error structures into state space models, we propose a novel Lognormal moment matching approach that allows users to specify the mean and variance of their process evolution and observation models.

Here, we fit our Lognormal moment matching models from a Bayesian perspective using Markov Chain Monte Carlo (MCMC) and Particle Markov Chain Monte Carlo (pMCMC; Andrieu et al. 2010). The Bayesian paradigm provides a flexible framework for fitting complex models, and the moment matching approach we introduce here offers a closed form Markov transition density. This closed form transition density provides the option to fit these models using MCMC, while still supporting access to particle filter (Cappé et al. 2007; Doucet and Johansen 2011) methods such as pMCMC. Both MCMC and pMCMC also provide a rigorous framework for quantifying parameter uncertainty and assessing forecast performance. MCMC methods generate samples from the posterior distribution, allowing practitioners to generate parameter estimates and build empirical density functions for their forecasts (Krüger et al. 2021). We can then validate our models by combining out-of-sample forecasts with MCMC output and evaluating them with proper scoring rules (Gneiting and Raftery 2007), such as the Continuous Ranked Probability Score (CRPS; Matheson and Winkler 1976) and the Ignorance Score (IGN; Good 1952; Roulston and Smith 2002).

We create four different models using the Lognormal moment matching technique developed here. The four models are all based off of the Gompertz and Moran-Ricker models (Gompertz 1825; Dennis et al. 2006; Ricker 1954), and are parameterized to have unbiased process evolution. We explore two different variance structures: a density dependent variance (Dennis et al. 2006) for the evolution and observations, and a constant variance for the evolution and observations.

First, we discuss interpretations of the Gompertz and Moran-Ricker SSMs and contrast them with interpretations of the models we present here. Next, we design and conduct simulation studies to compare forecast performance under model misspecification, and assess estimation of precision parameters. In particular, we follow the approach of Auger-Méthé et al. (2016) and investigate how the forecast performance and precision estimates change when the observation precision is fixed versus when it is estimated. Finally, we create a Lognormal State Space model that uses a two-dimensional process-based ecosystem model of Bloom and Williams (2015)

to model the state dynamics, and use it to predict the Leaf Area Index (LAI) at a focal forest site (University of Notre Dame Environmental Research Center; UNDE) in Wisconsin, USA. We parameterize our models in two different ways: once using a parameterization that assumes biased process evolution and observations with a density dependent variance structure, and once using a parameterization obtained using our moment matching method with a density dependent variance structure. We design and perform an analysis to test the viability of using a state space modeling framework to predict out-of-sample LAI, and to identify similarities and differences between our moment matching parameterization and the parameterization that assumes biased process dynamics and observations.

2 Motivating examples

In this section, we motivate the use of biologically coherent error structures in applications by demonstrating problems that may occur when modeling positive processes with error distributions that have mass over the entire real line. In particular we demonstrate how process models produce nonsensical predictions as they approach or exceed well defined biophysical bounds. First, we consider the following toy dynamical system in time:

$$X_t|X_{t-1} \sim \log\mathcal{N}\left(\mu = \log\left(\frac{X_{t-1}^2}{\sqrt{X_{t-1}^2 + \sigma^{2*}}}\right), \sigma^2 = \log\left(1 + \frac{\sigma^{2*}}{X_{t-1}^2}\right)\right), \quad (1)$$

where $\log\mathcal{N}(\mu, \sigma^2)$ represents a Lognormal distribution with parameters μ and σ^2 . Since the stochastic evolution function is Lognormally distributed, the process is positive, i.e. $X_t \in (0, \infty), \forall t = 1, \dots, T$. Further, the conditional mean and variance are given by $\mathbb{E}[X_t|X_{t-1}] = X_{t-1}$ and $\mathbb{V}[X_t|X_{t-1}] = \sigma^{2*}$.

Suppose that the positivity of the dynamical system is ignored, and instead the process is modeled as simple a Gaussian random walk:

$$X_t|X_{t-1} \sim \mathcal{N}(\mu = X_{t-1}, \sigma^2 = \sigma^{2*}), \quad (2)$$

where $\mathcal{N}(\mu, \sigma^2)$ represents a Gaussian distribution with parameters μ and σ^2 . To showcase the differences in forecast performance when the positivity constraint is ignored, we simulated two sample trajectories of length 50 from the Lognormal dynamical system in Eq. 1. Then, we simulate forward in time using the Lognormal model (Eq. 1) and the Gaussian model (Eq. 2) 10,000 times each. We used these 10,000 simulations to compute the median forecasts and 95% credible intervals for each model. When the system starts sufficiently far from zero the two models are indistinguishable from each other, producing nearly identical forecasts in terms of median estimates and 95% credible intervals (Fig. 1, top two panels). When the starting value X_0 is chosen to be close to the biological lower bound, however, the differences in forecasts become more apparent. The Lognormal model forecast intervals are now asymmetric and bounded below at zero, and quickly reach zero as the

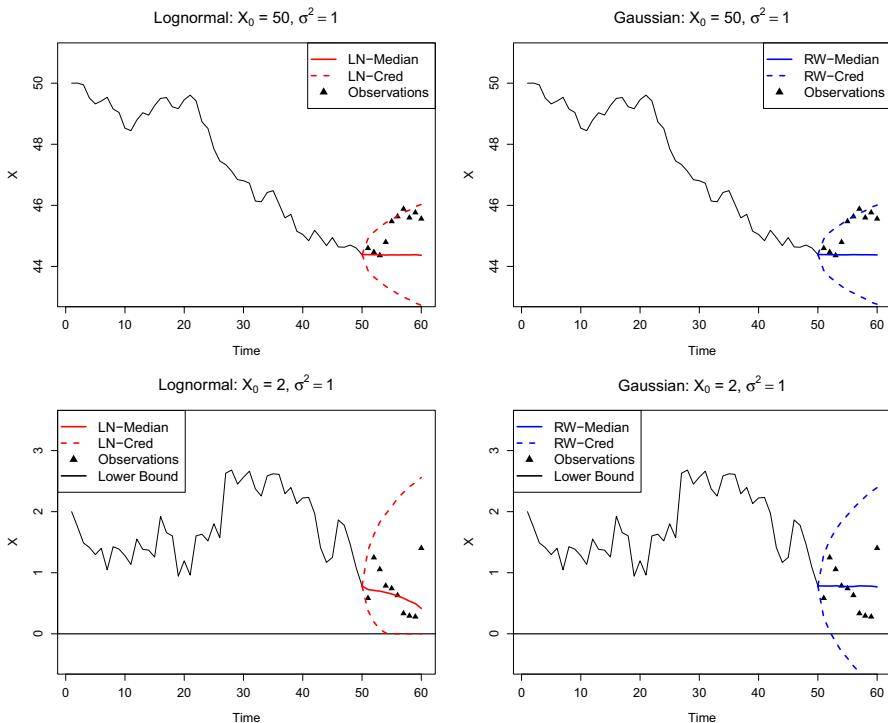


Fig. 1 Two sample trajectories (black lines) from the Lognormal dynamical system in Eq. 1. The top panels start from $X_0 = 50$ and the bottom panels start from $X_0 = 2$. Medians and 95 % credible intervals for 10 day forecast horizons (dotted lines) are created by simulating forward using the Lognormal model (Eq. 1) and the Gaussian model (Eq. 2) 10,000 times

forecast horizon increases. The random walk model forecast intervals retain their symmetry and continue to put considerable forecast mass below zero as the forecast horizon increases (Fig. 1, bottom two panels).

In this simple example, the negative forecasts are not causing any issues in the process model, thus only producing unrealistic forecasts. To illustrate an example where issues arise in the process model, consider the following dynamical system:

$$X_t | X_{t-1} \sim \log \mathcal{N} \left(\mu = \log(|X_{t-1}|) - a, \sigma^2 = \log \left(1 + \frac{\sigma^{2*}}{|X_{t-1}|} \right)^2 \right). \quad (3)$$

We note that, formally, absolute values are unnecessary for the evolution of the true model (since X_t is positive for all t). However, they will become necessary when we attempt to find an analogous Gaussian model with which to forecast.

Computing the conditional mean and variance of the Lognormal distribution with the given values of μ and σ^2 and using them as the conditional mean and variance of a Gaussian distribution results in:

$$X_t | X_{t-1} \sim \mathcal{N}\left(\exp\left(\mu + \frac{\sigma^2}{2}\right), (\exp(\sigma^2) - 1) \exp(2\mu + \sigma^2)\right). \quad (4)$$

To compare the differences between our Lognormal and Gaussian models, we simulated two sample trajectories of length 50 from the Lognormal dynamical system in Eq. 3. The first trajectory was simulated with parameters $X_0 = 100$, $\sigma^2 = 5$, $a = 0.01$. The second trajectory was simulated with parameters with $X_0 = 0.8$, $\sigma^2 = 0.1$, $a = 0.2$. To generate forecast summaries, we simulate forward in time using the Lognormal model (Eq. 3) and the Gaussian model (Eq. 4) 10,000 times each. We used these 10,000 simulations to compute the median forecasts and 95% credible intervals for each model. For initial values sufficiently far from zero, the forecast medians and 95% credible intervals from the Lognormal model are nearly identical to the forecast medians and 95% credible intervals from the Gaussian model (Fig. 2, top two panels). When the initial values are close to zero, the Lognormal model (bottom left panel, Fig. 2) enforces the lower bound and produces asymmetric credible intervals. However, the Gaussian forecasts when the system approaches zero (bottom right panel, Fig. 2) produce credible intervals that put mass

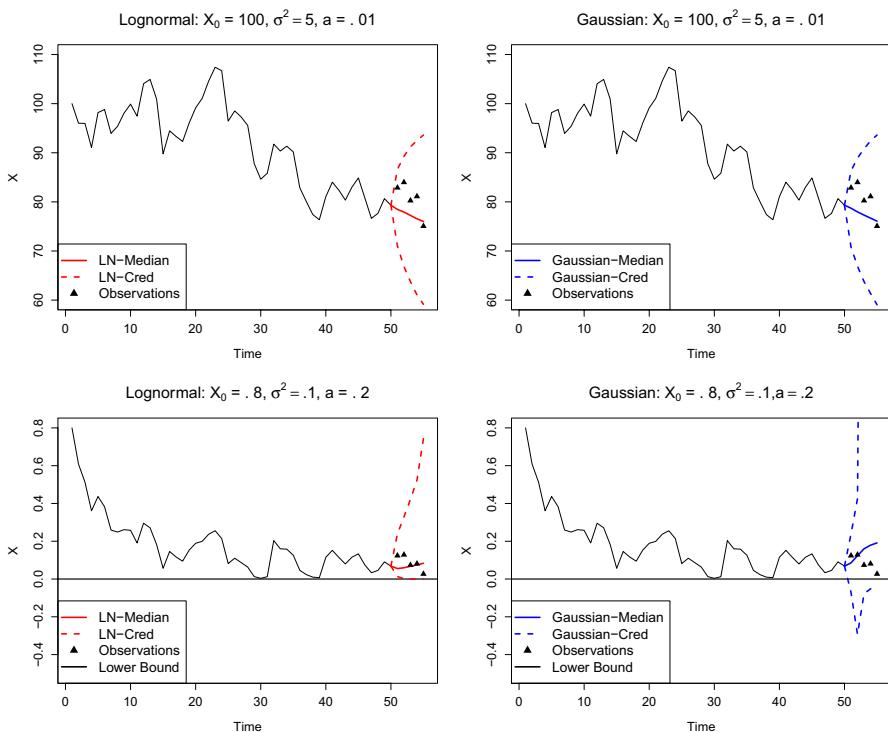


Fig. 2 Two sample trajectories from the Lognormal dynamical system in Eq. 3. The top panels starting from $X_0 = 100$ and the bottom panels starting from $X_0 = 0.8$. Medians and 95% credible intervals for 5 day forecast horizons (dotted lines) are created by simulating forward using the Lognormal model (Eq. 3) and the Gaussian model (Eq. 4) 10,000 times

below the biological lower bound. Moreover, the credible intervals for the Gaussian model are considerably wider than the credible intervals for the Lognormal model. When working on specific applications, predictions like these may be challenging to analyze because it is easy to associate the dynamics with an inadequate process model. In reality, the problem may be a consequence of having an error structure that can put probability mass on regions outside of the well defined upper and lower bounds of our ecological process.

3 Methods

3.1 Lognormal SSMs

State Space Models, sometimes referred to as Hidden Markov Models (HMMs), are a broad class of models used to track the states of a model (some unobserved process $X_{1:T}$) through a set of observations, $Y_{1:T}$ (Durbin and Koopman 2012; Petris et al. 2009). A state space model has three components—the state component, the observation component, and additional parameters. The state component consists of an (unobserved) Markov process, X_t , being moved through time by an evolution function (or process function) $f(X_t|X_{t-1}, \Theta)$. The implication of $X_{1:T}$ following a Markov process is that the past and the future are independent conditional on the state at the current time, X_t (Shumway and Stoffer 2011). The observation component consists of noisy observations of the latent process, $Y_{1:T}$, that are governed by an observation density function, $g(Y_t|X_t, \Theta)$. These observations are assumed to be independent of one another conditional on the latent states $X_{1:T}$. The additional parameter component, Θ , contains the parameters that govern the evolution function and observation function.

The Gompertz (1825) and Moran-Ricker (Ricker 1954) SSMs are simple discrete density-dependent state space models with Lognormally distributed process and measurement error (Dennis et al. 2006). The Gompertz and Moran-Ricker SSMs are popular choices for Lognormal state space models because they can be easily transformed into Normal Dynamic Models (NDMs), and can then take advantage of a suite of well studied fitting methods, including Kalman filtering (Kalman 1960), extended Kalman filtering (Julier and Uhlmann 1997), and Gibbs sampling (Geman and Geman 1984; Carter and Kohn 1994), alleviating computational difficulties associated with fitting SSMs.

The latent process models for the Gompertz model can be written as:

$$X_t = X_{t-1} \exp(a + b \log(X_{t-1}) + \epsilon_t), \quad \epsilon_t \sim \mathcal{N}(0, \phi) \quad (5)$$

The latent process model for the Moran-Ricker model can be written as:

$$X_t = X_{t-1} \exp(a + bX_{t-1} + \epsilon_t), \quad \epsilon_t \sim \mathcal{N}(0, \phi) \quad (6)$$

For both the Gompertz and Moran-Ricker models, the error terms ϵ_t are assumed to be: (a) independent and identically distributed and (b) independent of the initial latent state X_0 (Shumway and Stoffer 2011).

Letting $A = \exp(a)$, we can rewrite Eqs. 5 and 6, the Gompertz and Moran-Ricker process equations (respectively), in the form:

$$X_t = A(X_{t-1})^{b+1} \exp(\epsilon_t), \quad \epsilon_t \sim \mathcal{N}(0, \phi) \quad (7)$$

$$X_t = AX_{t-1} \exp(bX_{t-1}) \exp(\epsilon_t), \quad \epsilon_t \sim \mathcal{N}(0, \phi) \quad (8)$$

More generally, we can think of both the Gompertz and Moran-Ricker process models as belonging to a class of Lognormally distributed process models that we introduce in this work. These models have the form:

$$X_t = f^*(X_{t-1} | \Theta) \exp(\epsilon_t), \quad \epsilon_t \sim \mathcal{N}(0, \phi), \quad f^*(X_{t-1} | \Theta) > 0, \quad (9)$$

where $f^*(X_{t-1} | \theta)$ is the model of choice for the non-negative physical process, and the process error is governed by a multiplicative zero mean Lognormal distribution. The conditional mean ($\mathbb{E}[X_t | X_{t-1}]$), conditional variance ($\mathbb{V}[X_t | X_{t-1}]$) and conditional median ($\mathcal{M}[X_t | X_{t-1}]$) are given by:

$$\mathbb{E}[X_t | X_{t-1}] = f^*(X_{t-1} | \Theta) \exp((2\phi)^{-1}), \quad (10)$$

$$\mathbb{V}[X_t | X_{t-1}] = f^*(X_{t-1} | \Theta)^2 \exp(\phi^{-1}) \left(\exp(\phi^{-1}) - 1 \right) \quad (11)$$

$$\mathcal{M}[X_t | X_{t-1}] = f^*(X_{t-1} | \Theta). \quad (12)$$

Examining the conditional expected value, we see that this class of models is biased in terms of mean process evolution, but unbiased in terms of median process evolution. Further, the variance of the latent state is controlled by the value of the process model $f^*(\cdot)$, with larger values of the process assumed to have larger variation than smaller ones. Therefore, in the class of process models that describe the Gompertz and Moran-Ricker models, the density dependent relationship of the latent process variance is assumed to scale quadratically with the predicted value of the process model.

The generalization of the Gompertz and Moran-Ricker process models also holds true for the observation model. For example, for continuous responses Y_t , the assumed relationship between an arbitrary observation Y_t and the corresponding latent state X_t for the Gompertz model may be given by:

$$Y_t = X_t \exp(\eta_t), \quad \eta_t \sim \mathcal{N}(0, \tau) \quad (13)$$

Similar to the ϵ_t sequence, η_t are independent, identically distributed, and independent of the initial latent state X_0 .

Broadly, the Gompertz and Moran-Ricker observation models also belong to the same class of Lognormally distributed observation models in Eq. 13 that have the form:

$$Y_t = g^*(X_t | \Theta) \exp(\eta_t), \quad \eta_t \sim \mathcal{N}(0, \tau), \quad g^*(X_t | \Theta) > 0, \quad (14)$$

where $g^*(X_t|\Theta)$ is now interpreted as the model used to link our observations to their respective latent states, and the measurement error is dictated by a multiplicative zero mean Lognormal distribution. The conditional mean, conditional variance, and conditional median of this class of observation models are given by:

$$\begin{aligned}\mathbb{E}[Y_t|X_t] &= g^*(X_t|\Theta) \exp((2\tau)^{-1}), \\ \mathbb{V}[Y_t|X_t] &= g^*(X_t|\Theta)^2 \exp(\tau^{-1}) \left(\exp(\tau^{-1}) - 1 \right), \\ \mathcal{M}[Y_t|X_t] &= g^*(X_t|\Theta).\end{aligned}$$

Importantly, while this class of observation models assumes that observations are unbiased in log space, they assume there is systematic observation bias in their measurement process, which is given by $B(Y_t|X_t) = g^*(X_t|\Theta)(1 - \exp((2\tau)^{-1}))$.

Though we have only examined the Gompertz and Moran-Ricker models when discussing the class of Lognormal process and observation models presented here, there are many examples of Lognormal modeling frameworks in the animal population modeling literature (see Buckland et al. 2004; Knape et al. 2011, for examples) and fisheries modeling literature (see Maunder et al. 2015; Mäntyniemi et al. 2015, for examples) that fall into the classes we discuss here. Including a term to correct the bias is a common method used in applications to rectify assumptions of biased process evolution and biased observations (Knape et al. 2011; Maunder et al. 2015; Mäntyniemi et al. 2015) but this bias correction term changes the variance structure, as the Lognormal variance is a function of both μ and σ^2 . In particular this bias correction term can lead to an increase in variance, a consequence of the bias-variance trade-off (Casella and Berger 2002).

3.2 Lognormal moment matching models

Instead of assuming an unbiased median process evolution and a density dependent variance structure, a modeler developing a Lognormal SSM for their application may want to specify the mean evolution and the variance structure of the stochastic Lognormal process model. That is, we desire a Lognormally distributed stochastic evolution function such that $\mathbb{E}[X_t|X_{t-1}] = f^*(X_{t-1}|\Theta)$ and $\mathbb{V}[X_t|X_{t-1}] = \phi_t^{*-1}$. We can create a *process* model with these characteristics by using a moment matching transformation on the mean and precision, specifically:

$$\mu_t = \log \left(\frac{f^*(X_{t-1}|\Theta)^2}{\sqrt{f^*(X_{t-1}|\Theta)^2 + \phi_t^{*-1}}} \right), \quad (15)$$

$$\phi_t = \log \left(1 + (f^*(X_{t-1}|\Theta)^2 \phi_t^*)^{-1} \right)^{-1}. \quad (16)$$

This approach provides a stochastic Lognormal process model with the desired properties (see Appendix 7.1 for derivation):

$$\begin{aligned}
X_t | X_{t-1} &\sim \log \mathcal{N}(\mu = \mu_t, \sigma^2 = \phi_t^{-1}), \\
\mathbb{E}[X_t | X_{t-1}] &= f^*(X_{t-1} | \Theta), \\
\mathbb{V}[X_t | X_{t-1}] &= \phi_t^{*-1}.
\end{aligned}$$

Thus, our framework allows for us to parameterize the process model such that the temporal evolution is unbiased, and that allows for a flexible way to model the variance of the process through time. For example, we can choose $\phi_t = \phi$ to model constant variance through time, $\phi_t = \phi / f^*(X_{t-1} | \Theta)^2$ to model density dependent variance, or $\phi_t = \phi \exp(-t^{-1})$ to model variance that dissipates over time.

The Lognormal moment matching approach can be similarly applied to the *observation* model to produce a Lognormally distributed observation density with the properties $\mathbb{E}[Y_t | X_t] = g^*(X_t | \Theta)$, $\mathbb{V}[Y_t | X_t] = \tau_t^{*-1}$ by choosing parameters m_t and τ to be:

$$m_t = \log \left(\frac{g^*(X_t | \Theta)^2}{\sqrt{g^*(X_t | \Theta)^2 + \tau_t^{*-1}}} \right), \quad (17)$$

$$\tau_t = \log \left(1 + (g^*(X_t | \Theta)^2 \tau_t^*)^{-1} \right)^{-1}. \quad (18)$$

With the forms of both the process model and observation model fully specified in this manner, it is possible to write a likelihood equation for a model that includes both components, that we call the Lognormal Moment Matching model (LNM3). Suppose that we have observations at a subset of time points $I \subset \{1, \dots, T\}$. We denote these observations by $\mathbf{Y} = (Y_{i_1}, \dots, Y_{i_n})$ such that $(i_1, \dots, i_n) \in I$. Then the likelihood for the LNM3 is given by

$$\begin{aligned}
\mathcal{L}(X_{1:T}, \Theta | \mathbf{Y}) &= \prod_{t=1}^T \log \mathcal{N}(X_t | \mu = \mu_t, \sigma^2 = \phi_t^{-1}) \\
&\times \prod_{i \in I} \log \mathcal{N}(Y_i | \mu = m_i, \sigma^2 = \tau_i^{-1}),
\end{aligned} \quad (19)$$

where $\log \mathcal{N}$ denotes the Lognormal density function. μ_t , ϕ_t , m_i , and τ_i denote the moment matching parameterizations for the process and observations models from Eqs. 15–18.

The moment matching method can be used to parameterize our stochastic Lognormal model to describe any positive process or observation model and any variance structure, and thus helps to maintain biophysical coherence when modeling positive processes. Fitting these models as state space models further allows them to easily handle missing data and partition between process and measurement error. Thus the LNM3 approach provides a framework that is flexible, biophysically adequate, and statistically coherent.

3.3 Model fitting

Our analysis focused on comparing and contrasting the LNM3 approach to the Gompertz and Moran-Ricker approaches, two common Lognormal SSMs. We estimate the latent states, precisions, and model parameters for six different models presented here using Bayesian inference via Markov Chain Monte Carlo (MCMC; Robert and Casella 2005) and Particle Markov Chain Monte Carlo (pMCMC; Andrieu et al. 2010). The six different models include the Gompertz SSM, the Moran-Ricker SSM, and four different SSM formulations using the Lognormal Moment Matching (LNM3) method presented in Eqs. 15–18. Model parameters, latent states, and precisions were estimated via Markov Chain Monte Carlo (MCMC). MCMC is an estimation method that generates samples of parameters from their posterior distributions using Markov chains. Parameter uncertainty can be quantified using posterior samples from these Markov chains. All models were fit using JAGS (Plummer 2003) using the `rjags` package (Plummer 2019) in R version 4.1.0 (R Core Team 2016).

3.3.1 Gompertz SSM fitting

We fit the Gompertz model in log-space, by taking the logarithm of the latent states ($X_{1:T}$) and observations ($Y_{i \in I}$). We define these as $D_t = \log(X_t)$ and $\mathbf{F} = (F_{i_1}, \dots, F_{i_n}) = (\log(Y_{i_1}), \dots, \log(Y_{i_n}))$ such that $(i_1, \dots, i_n) \in I$. Under this transformation, the process model and observation model can be written as a Normal Dynamic Linear Model (NDLM) (West and Harrison 1997):

$$D_t = a + (1 + b)D_{t-1} + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \phi) \quad (20)$$

$$Y_i = D_i + \epsilon_{obs,i}, \quad \epsilon_{obs,i} \sim \mathcal{N}(0, \tau) \quad (21)$$

If the log-observations (\mathbf{F}) are available at a subset of time points, $I \subset \{1, \dots, T\}$, the likelihood for the Gompertz model can be written as:

$$\begin{aligned} \mathcal{L}(D_{1:T}, a, b, \tau, \phi | \mathbf{F}) &\propto \prod_{t=1}^T \sqrt{\phi} \exp\left(-\frac{\phi}{2}(D_t - a - (1 + b)D_{t-1})^2\right) \\ &\quad \times \prod_{i \in I} \sqrt{\tau} \exp\left(-\frac{\tau}{2}(F_i - D_i)^2\right). \end{aligned} \quad (22)$$

Prior choices for the parameters a , b , D_0 , ϕ , and τ used during our MCMC estimation of the Gompertz SSM are detailed in Sect. 7.2 of the Appendix.

The full conditional distributions for $D_{1:T}$ in the Gompertz model are analytically tractable, allowing for Gibbs sampling (Geman and Geman 1984). For interior latent states ($k = 1, 2, \dots, T - 1$) the Gibbs updates are given by:

$$\pi(D_k | \cdot) \sim N\left(\mu^* = \frac{\phi((1+b)D_{k-1} + a + (1+b)(D_{k+1} - a)) + \tau F_k \mathbb{1}_{k \in I}}{\phi(1 + (1+b)^2) + \tau \mathbb{1}_{k \in I}}, \phi^* = \phi(1 + (1+b)^2) + \tau \mathbb{1}_{k \in I}\right), \quad (23)$$

The Gibbs updates for the initial latent state and the final latent state are given by

$$\pi(D_T | \cdot) \sim N\left(\frac{\phi(1 + (1+b)D_{T-1} + a) + \tau F_T \mathbb{1}_{T \in I}}{\phi + \tau \mathbb{1}_{T \in I}}, \phi + \tau \mathbb{1}_{T \in I}\right), \quad (24)$$

$$\pi(D_0 | \cdot) \sim N\left(\frac{\mu_0 \phi_0 + \phi(1+b)(D_1 - a)}{\phi(1 + (1+b)^2) + \phi_0}, \phi(1 + (1+b)^2) + \phi_0\right), \quad (25)$$

where the function $\mathbb{1}_{i \in I}$ is an indicator function that returns $\mathbb{1}_{i \in I} = 1$ if a log-observation F_i is available at time i , and $\mathbb{1}_{i \in I} = 0$ otherwise.

We ran the MCMC for the Gompertz model for a total of 10,000 iterations, with a burn-in of 2000 iterations, and an adaptation period (`n.adapt` in JAGS) of 1000.

3.3.2 Moran-Ricker SSM fitting

We also fit the Moran-Ricker model in log-space by taking the logarithm of the latent states ($X_{1:T}$) and observations ($Y_{i \in I}$), $D_t = \log(X_t)$, $F_i = \log(Y_i)$. Under this log transformation, the process model and observation model can be written as:

$$D_t = a + D_{t-1} + \exp(bD_{t-1}) + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, \phi), \quad (26)$$

$$Y_i = D_i + \epsilon_{obs,i}, \quad \epsilon_{obs,i} \sim \mathcal{N}(0, \tau). \quad (27)$$

Assuming that log observations (\mathbf{F}) are available at a subset of time points, $I \subset \{1, \dots, T\}$, the likelihood for the Moran-Ricker model fit in log-space is given by:

$$\begin{aligned} \mathcal{L}(X_{1:T}, a, b, \tau, \phi | Y_I) \propto & \prod_{t=1}^T \sqrt{\phi} \exp\left(-\frac{\phi}{2}(D_t - a - D_{t-1} - \exp(bD_{t-1}))^2\right), \\ & \prod_{i \in I} \sqrt{\tau} \exp\left(-\frac{\tau}{2}(F_i - D_i)^2\right). \end{aligned} \quad (28)$$

Prior choices for the parameters a , b , D_0 , ϕ , and τ used during our MCMC estimation of the Gompertz SSM are detailed in Sect. 7.3 of the Appendix.

We ran the MC for the Moran-Ricker model for a total of 10,000 iterations, with a burn-in of 2000 iterations and an adaptation period (`n.adapt` in JAGS) of 1000.

3.3.3 Lognormal moment match SSM fitting

We explored four different Lognormal models using the moment matching approach. These models include: the Gompertz process model parameterized for unbiased mean process evolution and unbiased observation model with constant variances (LGC), the Moran-Ricker process model parameterized for unbiased mean process evolution and unbiased observation model with constant variances (LMRC), the Gompertz process model parameterized for unbiased mean process evolution and unbiased observation model with density dependent variances that scale quadratically with the process (LGD), and the Moran-Ricker process model parameterized for unbiased mean process evolution and unbiased observation model with density dependent variances that scale quadratically with the process (LMRD). For the remainder of this paper, we will be referring to the first two models (LGC, LMRC) as the constant variance models. The term density dependent models will be used to describe the second two models (LGD, LMRD) as well as the classical Gompertz and Moran-Ricker models. We carefully chose these four models to describe common assumptions for different SSM formulations. The unbiased mean process evolution, unbiased observation density function, and constant variance models (LGC, LMRC) were chosen to mimic the assumptions of homoskedastic Gaussian SSMs, which are not frequently used in Lognormal SSMs. The parameterizations for the LGD and LMRD models were chosen to mimic the variance structure of the Gompertz and Moran-Ricker models while maintaining an unbiased process evolution function and unbiased observation density function.

The likelihood for the constant variance models is obtained by substituting the values for $f^*(X_{t-1}|\Theta)$, ϕ_t , $g^*(X_t|\Theta)$ and τ_t from Table 1 into the LNM3 likelihood (Eq. 19).

The density dependent moment matching models were fit in log-space, by taking the log of the latent states and observations; $D_t = \log(X_t)$, $F_t = \log(Y_t)$. The model likelihood in log-space for the LGD model can be written as:

Table 1 Process evolution functions, process error structure, observation density function, and observation error structure for the four types of Lognormal moment matching SSMs used

Model name	$f^*(X_{t-1} \Theta)$	ϕ_t	$g^*(X_t \Theta)$	τ_t
LGC	$\exp(a)(X_{t-1})^{b+1}$	ϕ	X_t	τ
LMRC	$X_{t-1} \exp(a + bX_{t-1})$	ϕ	X_t	τ
LGD	$\exp(a)(X_{t-1})^{b+1}$	$\phi(\exp(a)(X_{t-1})^{b+1})^{-2}$	X_t	τX_t^{-2}
LMRD	$X_{t-1} \exp(a + bX_{t-1})$	$\phi(X_{t-1} \exp(a + bX_{t-1}))^{-2}$	X_t	τX_t^{-2}

LGC and LMRC represent the Gompertz and Moran-Ricker process functions and observation functions with constant process and measurement variance. LGD and LMRD represent the Gompertz and Moran-Ricker process functions and observation functions with density dependent process and measurement variance

$$\begin{aligned}
\mathcal{L}(D_{1:T}, a, b, \tau, \phi | \mathbf{F}) \propto \\
\prod_{t=1}^T \frac{\exp\left(-\frac{\log(1+\phi^{-1})^{-1}}{2}(D_t - (a + (1+b)D_{t-1} - .5 \log(1 + \phi^{-1})))^2\right)}{\sqrt{\log(1 + \phi^{-1})}} \\
\prod_{i \in I} \frac{\exp\left(-\frac{\log(1+\tau^{-1})^{-1}}{2}(F_i - (D_i - .5 \log(1 + \tau^{-1})))^2\right)}{\sqrt{\log(1 + \tau^{-1})}}.
\end{aligned} \tag{29}$$

Similarly, the model likelihood in log-space for the LMRD model can be written as:

$$\begin{aligned}
\mathcal{L}(D_{1:T}, a, b, \tau, \phi | \mathbf{F}) \propto \\
\prod_{t=1}^T \frac{\exp\left(-\frac{\log(1+\phi^{-1})^{-1}}{2}(D_t - (a + D_{t-1} + b \exp(D_{t-1}) - .5 \log(1 + \phi^{-1})))^2\right)}{\sqrt{\log(1 + \phi^{-1})}} \\
\prod_{i \in I} \frac{\exp\left(-\frac{\log(1+\tau^{-1})^{-1}}{2}(F_i - (D_i - .5 \log(1 + \tau^{-1})))^2\right)}{\sqrt{\log(1 + \tau^{-1})}}.
\end{aligned} \tag{30}$$

Prior choices for the parameters a , b , D_0 , X_0 , ϕ , and τ used during our MCMC estimation for each of the four moment matching SSMs are detailed in Sect. 7.4 of the Appendix.

We also note that the specification of the density dependent variance as $\phi_t = \phi f^*(X_{t-1} | \Theta)^2$ and the log-linear nature of the Gompertz curve lead to closed form full conditional distributions for the latent states in the LGD model. Similar to the Gompertz model, this allows for Gibbs sampling to be used to update the latent states, which helps decrease computation time.

We ran the MC for each of the four moment matching models for a total of 10,000 iterations, with a burn-in period of 2000 iterations, and an adaptation period (n. adapt in JAGS) of 1000.

3.4 Simulation study

We designed a simulation study to assess the forecasting performance of the six models presented here, both under the cases where they are the true generating models and where the models are misspecified. Our three primary objectives for the simulation study were: (1) assess forecasting performance of each model when it is the true generating model and when it is misspecified, when observation precision is being estimated (2) assess forecasting performance of each model when it is the true generating model and when it is misspecified, with fixed observation precision; (3) analyze the estimation of precisions for the models considered in this manuscript.

To assess the forecast performance of our models, we used proper scoring rules (Gneiting and Raftery 2007). Broadly, proper scoring rules use information about the predictive distribution coupled with observations to assign a measure of agreement of the forecast and the observations (Krüger et al. 2021). Specifically, Gneiting

and Raftery (2007) define a scoring rule to be proper if the expected value of the score is maximized by a draw from the true forecast distribution, and show that both the Continuous Ranked Probability Score (CRPS; Matheson and Winkler 1976) and the Ignorance Score (IGN; Good 1952; Roulston and Smith 2002) score are proper scoring rules. The IGN and CRPS are expressed in terms of a forecast probability density function $f(\cdot)$ and a forecast cumulative distribution function $F(\cdot)$. Given an observation y , the IGN and CRPS are defined as:

$$\text{IGN}(y) = -\log(f(y)); \quad (31)$$

$$\text{CRPS}(y) = \int_{\mathbb{R}} (F(z) - \mathbb{1}_{z \geq y})^2 dz. \quad (32)$$

For our simulation studies, we consider both the CRPS and IGN scores, so that we have one scoring rule defined in terms of the probability density function (IGN) and one scoring rule defined in terms of the cumulative distribution function (CRPS). We use the IGN and CRPS scores to quantitatively compare forecasts, with lower scores within a scoring rule indicating a better performance. CRPS has support over the positive real line, $[0, \infty)$, while IGN can take values between $[-\log(f(y^*)), \infty)$, where $y^* = \text{argmax}_{y \in \mathbb{R}} f(y)$.

To analyze the estimation of precision parameters in each of our six models, we used the empirical coverage rate of the 95% highest posterior density credible intervals for the precision parameters. Auger-Méthé et al. (2016) show that SSMs can have difficulty recovering the process and observation precisions even when the models are linear and Gaussian. To perform a thorough analysis on the estimation of the precision parameters, we fit each model under two different scenarios. In the first scenario, we fixed the observation precision and estimated the process precision for each model. In the second scenario, we estimated both the observation precision and the process precision for each model. This approach allowed us to assess the estimation by looking at the increase in empirical coverage rate for the process precision when the observation precision is fixed compared to when the observation precision is estimated. By choosing to use coverage and proper scoring rules to quantify our model performance, we follow a common approach used in the literature—using frequentist concepts to assess Bayesian models (Box 1980; Rubin 1984; Little 2006, 2012).

To quantify our three objectives, we performed the simulation study as follows: for each of the six models, thirty different synthetic datasets of length 575 were generated by simulating from the underlying process model and observation model, with parameter values taken from Table 2. Parameter values for each model were chosen so that the systems had similar mean dynamics for each generating model. Each of the six models discussed here was fit to each synthetic dataset. Models were initially fit with the first 365 days of data, and then forecasts were computed for days 366–372, a seven day forecast horizon. The average IGN and CRPS for the seven day forecasts horizons were computed using the `logs_sample` and `crps_sample` function from the `scoringRules` R package (Jordan et al. 2017). We also saved the highest posterior density credible intervals for the estimated precisions.

Table 2 Parameter values used for the six different model formulations to create thirty synthetic datasets for each generating model

Param.	Gomp	MR	LGC	LMRC	LGD	LMRD
$\exp(a)$	0.82	1.26	1.21	1.11	1.21	1.11
b	−0.658	−0.034	−0.099	−0.014	−0.099	−0.014
ϕ	70.2	51.9	4	4	70.2	70.2
τ	188.7	188.7	4	4	188.7	188.7

We then re-fit the models with the first 372 days of data, and forecasts were computed for days 373–379. This process was repeated until the synthetic data was exhausted. Overall, 12 models were fit to each individual synthetic data set—once by each model with observation precision known and once by each model with observation precision estimated, for a total of 2,160 simulations. By using the same synthetic datasets, we were able to isolate the differences in forecast performance for each precision scenario and quantify the difference using paired Wilcoxon tests (Wilcoxon 1945) with a Holm-Bonferroni adjustment (Holm 1979).

3.5 Application: Leaf Area Index predictions

We examined the differences in predictive performance between two different LN-SSM formulations by modeling Leaf Area Index (LAI) at University of Notre Dame Environmental Research Center (UNDE), a National Ecological Observatory Network (NEON) site. LAI is the total surface of leaves per area of ground and is a metric of the photosynthetic capacity of an ecosystem. We obtained estimates LAI at UNDE from Oak Ridge National Laboratory Distributed Active Archive Center using their fixed subsets feature (DAAC 2018). The dataset contains estimates of LAI at four day intervals, computed using reflectance estimates from the Moderate Resolution Imaging Spectroradiometers (MODIS) satellite as inputs to a phenological model that predicts LAI (Yang et al. 2006). The two statistical models that we considered were a biased LN-SSM, based on Eqs. 9 and 14, and a moment matching LN-SSM, based on Eqs. 15, 16, 17, and 18. Our analysis was designed with two primary questions in mind: (1) can we construct accurate out-of-sample predictions of LAI for intermediate horizons (e.g., multiple seasons) using a Lognormal state space model?; (2) do the moment matching and biased formulations show differences in out-of-sample predictive performance for intermediate horizons, when measured by CRPS and IGN?

To fit the LAI data, we used a reduced version of the Data Assimilation Linked Ecosystem Carbon Model, Version 2 (DALEC2; Bloom and Williams 2015). In DALEC2, the LAI is modeled as a function of the carbon stored in foliage. Rather than fitting the full model (a six dimensional dynamical system with 23 process parameters) we only considered the interaction between foliage carbon (C_f) and the labile carbon (C_{lab}). This reduced the model to a two dimensional dynamical system with eight process parameters. Descriptions and units for DALEC2 process parameters can be found in Table 7. The form of the reduced process model is given as:

$$\mathbf{C}^{(t)} = M_t \mathbf{C}^{(t-1)} + p_t, \text{ where} \quad (33)$$

$$M_t = \begin{bmatrix} 1 - \Phi_f^{(t)} & \Phi_o^{(t)} \\ 0 & 1 - \Phi_o^{(t)} \end{bmatrix}, \mathbf{C}^{(t-1)} = \begin{bmatrix} C_f^{(t-1)} \\ C_{lab}^{(t-1)} \end{bmatrix}, p_t = \begin{bmatrix} G(\mathbf{D}^{(t)}, c_{lma}, c_{eff}) f_f \\ G(\mathbf{D}^{(t)}, c_{lma}, c_{eff}) f_{lab} \end{bmatrix} \quad (34)$$

$$\text{LAI}^{(t)} = \frac{C_f^{(t)}}{c_{lma}}$$

$$\Phi_f(t, d_f, c_r, c_{lf}) = \sqrt{\frac{2}{\pi}} \cdot \frac{-\log(1 - c_{lf})}{c_{rf}} \exp\left(-\left(\sin\left(\frac{t - d_f + \psi_f}{s}\right) \frac{\sqrt{2}s}{c_{rf}}\right)^2\right), \quad (35)$$

$$\Phi_o(t, d_o, c_{ro}) = \sqrt{\frac{2}{\pi}} \left(\frac{6.9088}{c_{ro}}\right) \exp\left(-\left(\sin\left(\frac{t - d_o + .6245c_{ro}}{s}\right) \frac{\sqrt{2}s}{c_{ro}}\right)^2\right), \quad (36)$$

where $s = 365.25/\pi$ and $\psi_f = -\sqrt{W_0\left(\left(2\pi \log(1 - c_{lf})^2\right)^{-1}\right)}/\sqrt{2}$, where W_0 is the principal branch of the Lambert W function (Lambert 1758), and $G(\mathbf{D}^{(t)}, c_{lma}, c_{eff})$ is the output of the Aggregated Canopy Model (ACM) for gross photosynthetic production (Williams et al. 1997). $\mathbf{D}^{(t)}$ represents meteorological driver variables for day t that include: daily minimum and maximum temperatures ($^{\circ}\text{C}$), daily incoming shortwave radiation (g Cm^{-1}), and atmospheric carbon (CO_2 ppm). Minimum temperatures, maximum temperatures, and daily shortwave radiation were obtained from the National Ecological Observatory Network (NEON; National Ecological Observatory Network 2020; National Ecological Observatory Network (NEON) 2022b, a). We imputed any missing NEON observations using a piece-wise linear interpolation. We took monthly measurements of atmospheric carbon from the Scripps Project (Keeling et al. 2005), and interpolated daily measurements by assigning the monthly values to each day within the month.

To fit the reduced DALEC2 as a LN-SSM, we used the process-based model (Eq. 33) as the state component and the LAI equation (Eq. 34) as the observation component. To model the variance structure of the process evolution function and observation function, we used a density dependent variance. We chose a density dependent variance structure (Dennis et al. 2006) because we expect more process variation and measurement error for foliage carbon when it is large during the early spring and summer months, and less when it is low in the winter months. We considered two different LN-SSM formulations with density dependent variance structures. The first model, the biased model, parameterizes the process and observation components in a biased manner using Eqs. 9 and 14. The second model that we consider parameterizes the process and observation components in an unbiased manner using our moment matching approach. The process evolution function and the observation function for the biased model are given by:

$$\mathbf{C}^{(t)} | \mathbf{C}^{(t-1)} \sim \text{MV log } \mathcal{N}(\log(M_t \mathbf{C}^{(t-1)} + p^{(t)}), \Omega), \quad (37)$$

$$\text{LAI}^{(i)} | \mathbf{C}^{(i)} \sim \log \mathcal{N} \left(\log \left(\frac{C_f^{(i)}}{c_{lma}} \right), \tau^{-1} \right), \quad (38)$$

where $\text{MV log } \mathcal{N}$ represents the multivariate Lognormal density, M_t, p_t and $\mathbf{C}^{(t)}$ take on the same values from Eq. 33, $\Omega = \text{Diag}(\omega_f^2, \omega_{lab}^2)$, and the \log in the process evolution mean represents the componentwise logarithm of each element of the vector.

The process evolution function and observation function for the moment matching model are given by:

$$\mathbf{C}^{(t)} | \mathbf{C}^{(t-1)} \sim \text{MV log } \mathcal{N}(\log(M_t \mathbf{C}^{(t-1)} + p^{(t)}) - .5 \log(\mathbf{1}\mathbf{1}^T + \Omega)\mathbf{1}, \log(\mathbf{1}\mathbf{1}^T + \Omega)), \quad (39)$$

$$\text{LAI}^{(i)} | \mathbf{C}^{(i)} \sim \log \mathcal{N} \left(\log \left(\frac{C_f^{(i)}}{c_{lma}} \right) - .5 \log(1 + \tau^{-1}), \log(1 + \tau^{-1}) \right), \quad (40)$$

where $\mathbf{1}$ represents a 2×1 vector of ones, $\Omega = \text{Diag}(\omega_f^2, \omega_{lab}^2)$, and the \log once again represents a componentwise logarithm. Additional information on expected values and variances for the two different DALEC2 LN-SSM formulations can be found in Table 6 in the Appendix.

The equations above that we use to define our LN-SSM also highlight an interesting difficulty of our analysis: we are fitting a two state dynamical system, but we only have observations for one of the states, (C_f). Although we are primarily interested in predicting LAI, which is a function of C_f , we still need to estimate labile carbon so that we can model its contribution to foliage during the period of leaf regrowth. This also highlights an advantage of our state space modeling approach: the uncertainty from our lack of labile carbon measurements is propagated through time, giving us a more complete picture of the uncertainty in the foliage carbon.

We treated two model parameters as fixed: the Leaf Mass per Area (c_{lma}) and the density dependent observation precision parameter (τ). We fixed c_{lma} to avoid identifiability issues with the canopy efficiency parameter (c_{eff}), as the two parameters appear exclusively together in the ACM (Williams et al. 1997). For both models, we fixed c_{lma} to a value of 75, based on empirical results from Serbin et al. (2019) that were calibrated specifically at UNDE. We estimated the density dependent observation precision parameter using historical data from MODIS, which reports both the standard deviation and the mean of the LAI estimates. For the moment matching model, we took the median value of the means divided by the standard deviations. This gave us a value of $\hat{\tau} \approx 4$. For the standard Lognormal SSM formulation, we used the form of the variance for the Lognormal distribution (Eq. 10) to obtain $\hat{\tau} \approx 4.18$. For the DALEC2 process parameters, we used uniform prior distributions over the range of acceptable values taken from Bloom and Williams (2015). For the density dependent process variance components, we used a Uniform(0, 1) distribution as the prior. We chose this because

the ω parameters for the biased and moment matching parameters have the interpretation that they roughly control the average proportion of process error at each time point. All prior distributions for parameters were identical for the moment matching model and the biased model. A full table detailing prior distributions and fixed parameters can be found in Appendix 7.6.

We fit both of the models using particle MCMC (pMCMC; Andrieu et al. 2010). We ran each model for a total of 881 days, from September 5th, 2019 to February 2nd, 2022. We used four day MODIS LAI measurements from September 5th, 2019 to September 6th, 2021 to fit each model, and we used the remaining MODIS LAI measurements from September 7th, 2021 to February 2nd, 2022 to assess out-of-sample prediction. To assess out-of-sample predictions, we used the pMCMC samples of the latent states to generate samples from the posterior predictive distribution for the observations. We then used these samples to validate against the out-of-sample MODIS LAI measurements using CRPS and IGN using the `scoringRules` package (Jordan et al. 2017) in the R programming language (R Core Team 2016). By doing this, we are scoring on $Y|X$ rather than directly on the observation Y , and acknowledge that the LAI observations that we are using for validation have measurement error (Ferro 2017; Bessac and Naveau 2021). We implemented pMCMC using the `pmcmc` function in the R package `pomp` (King et al. 2016). The `pmcmc` function implements the Particle Marginal Metropolis Hastings (PMMH) algorithm of Andrieu et al. (2010), using a bootstrap particle filter (Gordon et al. 1993). We ran each model for a total of 100,000 iterations, with a burn in of 50,000 iterations, 500 particles, and an adaptive multivariate normal proposal distribution (Andrieu and Thoms 2008; Rosenthal 2009) that began using a scaled empirical covariance matrix after 1000 samples are accepted. To obtain initial parameter estimates that start in a region of high posterior density, we used a Gaussian process surrogate model optimization (Gramacy 2020) using `TRBO` (Eriksson et al. 2019), with the particle filter marginal log-likelihood as the objective function.

LAI is ubiquitous for forecasting changes in carbon stored in different components of a terrestrial ecosystem under different projections of climate. For our analysis, we chose to predict over multiple seasons, in an attempt to emulate a forecasting scenario of LAI response to predicted changes in climate, such as a particularly hot or cold winter. Our analysis also serves as an efficacy test for predictive modeling of LAI using a state space framework. While much work has been done on LAI prediction using ecosystem process-based models (see Mahowald et al. 2016; Ercanli et al. 2018, for examples), to our knowledge there has been little work done on predicting LAI by using a mechanistic process-based ecosystem model as the process component of a statistical state space model.

4 Results

4.1 Simulation study

The first objective of our simulation study was to assess the forecasting performance of each model under model misspecification when observation precision is

estimated. We expected that the true generating model would perform best on average for its thirty synthetic datasets, using the average IGN and CRPS scores over each seven day forecast horizon. For the case where both the process precision, ϕ , and the observation precision, τ , were estimated, the Gompertz model (Gomp in Table 3) had the best forecasting performance among the density dependent models (MR, Gomp, LMRD, LGD) for both CRPS and IGN. The unbiased moment matching analog to the Gompertz, LGD, also had a strong performance, scoring second highest for IGN on all four density dependent variance models, and scoring second highest for CRPS on three out of the four. For the constant variance models (LMRC, LGC), the constant variance Moran-Ricker model (LMRC) had the best forecasting performance for both CRPS and IGN.

The second objective of our simulation study was to assess the forecasting performance of each model under model misspecification when observation precision was fixed. For the simulations where the observation precision, τ , was fixed and the process precision, ϕ , was estimated, we found higher consistency between our scoring rules and the generating model. For the simulations that estimated both precision parameters (Table 3), the Gomp and LGD models consistently outperformed the other density dependent variance models for both scoring rules. In contrast, the simulations where τ was fixed and ϕ was estimated (Table 4) showed a more equal representation among the density dependent models. For the four density dependent generating models, the Moran-Ricker and its unbiased moment matching

Table 3 Average CRPS and IGN scores for the simulations where both τ and ϕ are estimated

Generating model							
Model	MR	Gomp	LMRC	LGC	LMRD	LGD	
Average CRPS							
MR	0.8587	0.8027	0.6496	0.6362	0.8752	0.8970	
Gomp	0.8344	0.7773	0.5997	0.5952	0.8279	0.8500	
LMRC	0.8356	0.7810	0.5913	0.5889	0.8411	0.8672	
LGC	0.8375	0.7809	0.5927	0.5900	0.8440	0.8672	
LMRD	0.8380	0.7872	0.6181	0.6100	0.8357	0.8585	
LGD	0.8357	0.7789	0.6080	0.6016	0.8310	0.8517	
Average IGN Score							
MR	1.816	1.744	1.595	1.557	1.806	1.830	
Gomp	1.785	1.710	1.514	1.494	1.751	1.774	
LMRC	1.812	1.757	1.456	1.452	1.858	1.915	
LGC	1.806	1.750	<i>1.458</i>	<i>1.454</i>	1.840	1.882	
LMRD	1.793	1.723	1.537	1.507	1.762	1.790	
LGD	<i>1.792</i>	<i>1.714</i>	1.530	1.500	<i>1.755</i>	<i>1.778</i>	

Columns represent the generating model for the synthetic datasets and rows represent the models used to fit the datasets. Scores are averaged over thirty different synthetic datasets and thirty different 7 day forecast horizons for each combination of generating model and model used to fit the data. Bold entries represent the lowest score for a given generating model, and italicized entries represent the second lowest score

counterpart scored the best for CRPS, with each model producing the lowest score twice. When measuring performance with IGN, the Moran-Ricker and Gompertz models each scored lowest twice. The density dependent models all scored similarly, regardless of the choice of true generating model, both for CRPS (max difference ≈ 0.0052) and for IGN (max difference ≈ 0.006).

Our third objective was to analyze the estimation of precisions for each of the six models we used. Our first investigation for the estimation of the precision parameters was to quantify the differences in IGN and CRPS scores between the simulations where the observation, τ was fixed and the simulations where τ was estimated. This investigation is also intimately related to the performance of the models under model misspecification, and therefore provides insight for all three of our objectives. To evaluate the impacts statistically, we used a paired Wilcoxon test (Wilcoxon 1945). To do this, we took the average IGN and CRPS scores over each seven day forecast horizon for the two different scenarios, treating them as "before" and "after". We justify this by noting that each precision scenario was fit using identical synthetic datasets, with the only difference being whether τ was estimated or not. A paired Wilcoxon test was chosen over a paired t-test after finding that the normality assumption of the paired t-test was not satisfied. Unsurprisingly, we found that fixing the observation precision helped to improve forecasting performance for nearly all models. The LGD model performed significantly better in terms of IGN when τ was fixed (LGD: $p = 0.023$), but did not perform significantly better in terms of

Table 4 Average CRPS and IGN scores for the simulations where ϕ is estimated and τ is known

Generating model						
Model	MR	Gomp	LMRC	LGC	LMRD	LGD
Average CRPS						
MR	0.8332	0.7770	0.5968	0.5945	0.8230	0.8492
Gomp	0.8358	0.7774	0.6000	0.5954	0.8276	0.8508
LMRC	0.8357	0.7813	0.5911	0.5891	0.8337	0.8619
LGC	0.8362	0.7805	0.5929	0.5897	0.8349	0.8603
LMRD	0.8332	0.7766	0.5971	0.5937	0.8233	0.8495
LGD	0.8359	0.7777	0.5999	0.5959	0.8282	0.8510
Average IGN Score						
MR	1.785	1.711	1.502	1.487	1.748	1.775
Gomp	1.788	1.710	1.515	1.494	1.752	1.774
LMRC	1.819	1.764	1.455	1.453	1.848	1.905
LGC	1.813	1.750	<i>1.458</i>	<i>1.453</i>	1.834	1.887
LMRD	<i>1.786</i>	1.712	1.505	1.486	<i>1.752</i>	1.780
LGD	1.788	<i>1.710</i>	1.513	1.494	1.752	1.776

Columns represent the generating model for the synthetic datasets and rows represent the models used to fit the datasets. Scores are averaged over thirty different synthetic datasets and thirty different 7 day forecast horizon for each combination of generating model and model used to fit the data. Bold entries represent the lowest score for a given generating model, and italicized entries represent the second lowest score

CRPS (LGD: $p = 0.14$). The Moran-Ricker, LMRD, LGC, and LMRC models all had statistically significant decreases in IGN (MR: $p < 2e-16$; LMRD: $p = 1.6e-07$; LGC: $4.3e-07$; LMRC: $p = 7.3e-08$) as well as CRPS (MR: $p < 2e-16$; LMRD: $p = 1.6e-07$; LGC: $1.8e-08$; LMRC: $p = 1.1e-05$) when τ was fixed. Fixing the observation precision did not significantly impact the performance of the Gompertz model for IGN ($p = 0.74$) or CRPS ($p = 0.75$).

Our second investigation for the estimation of the precision parameters was based on the empirical coverage rate of the HPD intervals. The empirical coverage of the highest posterior density (HPD) credible intervals for ϕ increased for every simulation where τ was fixed, and all six models had empirical coverage that fell within 1.5% of the nominal rate. For the simulations where both ϕ and τ are estimated, Table 5 shows that the Gompertz model and LGD models produce the best empirical coverage for the precision parameters. This is likely to be related to their excellent performance in these simulations, where other models struggled to consistently produce precision estimates that contained the ground truth in their 95% HPD intervals. The empirical undercoverage of the Moran-Ricker (MR) model for both ϕ and τ may also explain its poor performance in the forecast results from Table 3, where it came in last place for CRPS for all six generating models and last place for IGN for three out of six generating models, including the case where it was itself the true generating model.

4.2 Leaf Area Index predictions at UNDE

We found that both the moment matching LN-SSM and the biased LN-SSM produced predictive distributions that captured the dynamics of both the in-sample and the out-of-sample LAI observations. Both models showed similar fits for the in-sample LAI observations when looking at medians (Fig. 3) and 90% highest posterior density intervals (Fig. 3). This was not surprising to us, as both models used an identical process model and prior distributions, and only differed slightly in the formulation of process evolution and observation functions.

For the out-of-sample LAI predictive distributions, the models behaved differently. The biased model (Fig. 3, top panel) had lower variance at the start of the predictive horizon, and then began to tail off at the end of the horizon. The moment matching model (Fig. 3, bottom panel) had larger predictive variance at first, but

Table 5 Average empirical HPD credible interval coverage for the precisions of each generating model, under the scenarios where both ϕ and τ are estimated and when τ is fixed

Generating model						
Parameter	MR	Gomp	LMRC	LGC	LMRD	LGD
Scenario 1: Both ϕ and τ estimated						
ϕ	72.3%	89.1%	89.2%	86.9%	85.8%	89%
τ	74.7%	96.4%	92.1%	91.8%	91.7%	95.3%
Scenario 2: ϕ Estimated, τ Fixed						
ϕ	95.2%	94.8%	94.1%	93.6%	94.2%	94.9%

The nominal expected coverage rate is 95%. Coverage rates are averaged over all synthetic datasets and forecast horizons, for a total of 900 samples

then leveled off and accumulated slowly at the end of the horizon. The median predicted values for the moment matching model are slightly larger than the median predicted values for the biased model over the entire horizon. This is an interesting result: for identical process parameter values, the median of the moment matching models should be strictly smaller than the median of the biased model. This tells us that there is some mismatch between the parameters being estimated for the biased model and the moment matching model. This is something that we must be cautious about when we are using parameters that have physical interpretations. Modeling the ecosystem dynamics using a biased state process evolution may lead to biased parameter estimates and incorrect inferences for our interpretable parameters.

We also found differences in performance between the two models when measured by IGN and CRPS. When measured by mean IGN across the 32 out-of-sample LAI measurements, the two models had similar performance, with the moment matching model performing slightly better (mean $IGN_{mm} = 0.3566$; mean $IGN_{biased} = 0.3850$). The moment matching model had marginally better IGN scores for the observations where neither model performed well (e.g. out-of-sample measurements two, three, and four, Fig. 4, top panel), and the models had similar IGN score performance otherwise. When measured by mean CRPS across the 32 out-of-sample LAI measurements, the moment matching model showed a much better performance (mean $CRPS_{mm} = 0.2389$; mean $CRPS_{biased} = 179.93$). Towards the end of the prediction horizon the CRPS for the biased model quickly increases, while the CRPS for the moment matching model stays comparatively small. We believe that this happens because of an accumulation of bias in the biased model combined with the heavy tails of the Lognormal distribution.

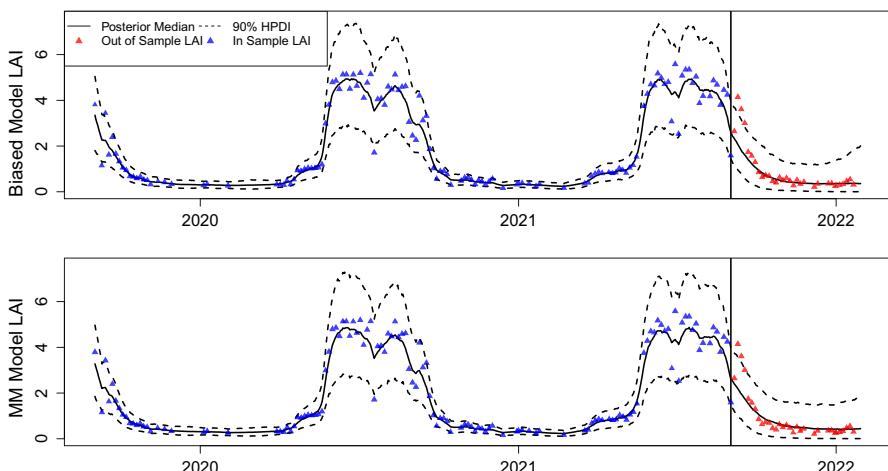


Fig. 3 Model fits for the biased model (top panel) and moment matching model (bottom panel). Medians (solid line) and 90% highest posterior density intervals (dashed lines) were computed using 50,000 post burn-in samples of the latent states generated by pMCMC. Blue triangles denote LAI measurements that were used to train the model, and red triangles denote LAI measurements that were not seen by the model and used only for validation purposes. The vertical black represents the time value where the model began predicting out-of-sample

5 Discussion

Although the Gaussian distribution is a common choice in modeling applications, many ecological processes have strict lower bounds that are not adequately captured by Gaussian models. This mismatch becomes especially problematic in forecasting applications, where uncertainty grows as the forecast horizon increases. However, incorporating assumptions about the evolution of the latent process and variance dynamics into non-Gaussian SSM frameworks can be challenging. To remedy this, we proposed a method for incorporating non-negative process and observation models with arbitrary variance structures into Lognormal SSMs using moment matching (LNM3). The primary advantage of our method is flexibility: it allows practitioners to create stochastic processes and measurement components that are unbiased in terms of their mean evolution, have a flexible variance that can change through time, and offer a closed form Markov transition density that allows models to be fit with standard MCMC software such as JAGS (Plummer 2003).

We used Monte Carlo simulations to assess the forecasting performance of the six models discussed here, using a total of 180 synthetic datasets that were fit twelve times each: once by each model with the observation precision estimated, and once by each model with the observation precision fixed. We found that the forecasting performance of our models under misspecification was heavily dependent on whether or not the observation precision was fixed, and also dependent on the metric used for evaluation: CRPS or IGN. With the observation precision estimated, the Gompertz model had the best average CRPS and IGN scores across all of the synthetic datasets for four out of the six generating models. With the observation precisions fixed at the true values, we found that every model except for the Gompertz model had a significant increase in forecast performance when measured by average

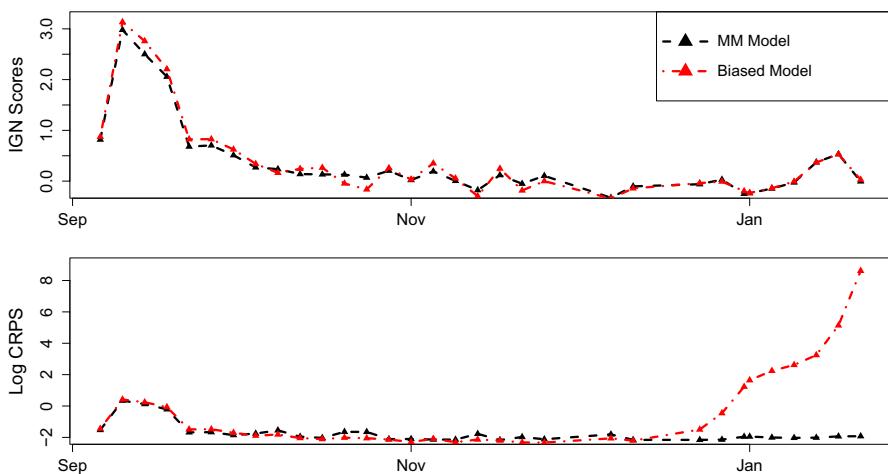


Fig. 4 Visualizations of the IGN scores (top panel) and the log CRPS (bottom panel). The dashed black line denotes the moment matching model, and the dashed and dotted red line denotes the biased model. Triangles along these lines represent points where there were out-of-sample LAI measurements for validation. For visualization purposes, we used a linear interpolation between measurements

CRPS or average IGN. For these simulations, no model dominated the others in terms of forecast performance for either metric. The Gompertz model outperforming the true generating models identifies one of the difficulties of using proper scoring rules to evaluate forecast performance, especially if using forecasting performance as a way to guide model choice. Although it is usually expected that the CRPS and IGN scores should favor the generating model, we are unlikely to have access to the true parameter values that underlie the generating model and instead have to rely on estimates of parameter values from our MCMC.

Even simple linear Gaussian SSMs can be prone to estimation problems, especially for parameters that govern the variance structure of the process and observations (Auger-Méthé et al. 2016). We found that our models were no exception to this: when both the process precision and observation precision were estimated, no model came within 5% of the nominal 95% coverage rate for the process precision HPDs. We also found that estimation of the precision parameters was closely related to forecasting performance. For the simulation where both observation and process precision parameters were estimated, the two models with the best performance (Gompertz and LGD) were also the models that had empirical coverage rates closest to 95% for the precision parameters. Similarly, the Moran-Ricker model was 20% below the nominal coverage rate for both the observation and process precision parameters, and had the worst average CRPS scores for every generating model, including itself. For the simulation studies where the observation precision was fixed, we found that coverage rates for each of the six models were close to the nominal 95% coverage rate, with the empirical coverage rates ranging from 93.6% to 95.2%. This supports the findings from Auger-Méthé et al. (2016), who show that fixing the measurement error in linear Gaussian SSMs can help to alleviate estimation problems.

We tested the efficacy of our method when applied to a challenging problem by using a two-dimensional process-based ecosystem model to describe the state process dynamics in a LN-SSM and using it to predict the Leaf Area Index (LAI). Overall, we found that both models performed well in reconstructing latent states that had good agreement with the LAI measurements while adequately capturing the dynamics of the out-of-sample measurements that we used for validation. Both models showed similar fits for the in-sample LAI measurements, but showed differences in out-of-sample predictive performance. The moment matching model, developed using the methodology we present here, had a superior performance for the out-of-sample LAI when assessed using both IGN scores and CRPS. This supports our idea that having a flexible mechanism for adjusting the process evolution, observation function, and variance structure can help to better capture out-of-sample dynamics and improve predictive performance. Towards the end of the predictive horizon, the biased model begins to perform very poorly when measured by CRPS, and the moment matching model continues to perform well. This indicates that the moment matching framework used here may be better for applications that have predictive horizons of intermediate length.

Though we saw better performance using our moment matching technique, the analysis that we did here serves mainly as proof-of-concept for state space modeling of LAI, and there is much room for improvement in future work. The largest improvement would be to integrate additional data streams, and to include weather

drivers that are forecast. For example, including data on litterfall accumulation from the National Ecological Observatory Network (NEON) would help to further constrain process parameters, process variance, and potentially improve out-of-sample prediction performance, and including forecasted weather drivers adds an additional level of uncertainty to our model predictions. Similarly, in future work we can consider additional methods for constraining parameters in the absence of direct observations, such as the ecological data constraints described in Bloom and Williams (2015). Finally, in future work we can study the performance of the biased and moment matching DALEC2 model formulations over longer predictive horizons, such as multi-year projections of LAI under different climate scenarios.

Though the methods presented here use the Lognormal distribution to represent stochasticity, the moment matching approach broadly applies to other distributions as well, and provides opportunities for future directions. For example, the gamma distribution has been considered for state space modeling (Smith and Miller 1986) and stochastic differential equation modeling (Dennis and Patil 1984) applications, has non-negative support, can be parameterized in terms of its mean and variance using a moment matching approach, and has lighter tail behavior than the Lognormal distribution. The beta distribution, which has previously been used in SSMs (see Osthust et al. 2017; Deo and Grover 2021, for examples), has adequate support for modeling proportions and can also be parameterized in terms of its mean and variance to allow for moment matching approaches.

In conclusion, to address biological non-coherence in models of physical systems, we proposed a novel Lognormal state space modeling framework that preserves the positivity of the latent process and observations. The methods presented here allow practitioners to incorporate complex process models and error dynamics into state space models while ensuring that the forecasts produced by the model agree with the constraints of the system. The flexibility of the moment matching method for representing complex systems along with the variance partitioning of the state space model provide a coherent statistical framework for forecasting, in terms of biophysical adequacy, forecast assessment, and uncertainty quantification.

Appendix

Lognormal moment matching

In this section, we derive the moment matching transformation that we use throughout this work to specify the dynamics of the process and observation models. Suppose that we are interested in finding a transformation for random variable $X \sim \text{Lognormal}(\mu, \sigma^2)$, such that $\mathbb{E}[X] = \mu^*$, $\mathbb{V}[X] = \sigma^{2*}$. The Lognormal probability density function is given by:

$$f(x|\mu, \sigma) = \frac{1}{\mu\sigma\sqrt{2\pi}} \exp\left(-\frac{(\log(x) - \mu)^2}{2\sigma^2}\right), \mu \in \mathbb{R}, \sigma > 0, x \in (0, \infty)$$

Then, we have the following equations for the mean and variance:

$$\mathbb{E}[X] = \mu^* = \exp\left(\mu + \frac{\sigma^2}{2}\right) \quad (41)$$

$$\mathbb{V}[X] = \sigma^{2*} = (\exp(\sigma^2) - 1) \exp(2\mu + \sigma^2) \quad (42)$$

By taking the logarithm on both sides, we can rewrite Eq. 41 as: $\sigma^2 = 2(\log(\mu^*) - \mu)$. We can substitute this into Eq. 42 to get μ^* written in terms of known constants,

$$\begin{aligned} \sigma^{2*} &= \left(\exp(2\log(\mu^*) - 2\mu) - 1\right) \exp(2\mu + 2\log(\mu^*) - 2\mu) \\ &= \left(\exp(2\log(\mu^*) - 2\mu) - 1\right) \exp(2\log(\mu^*)) \\ &= (\mu^{*2} \exp(-2\mu) - 1) \mu^{*2} \end{aligned} \quad (43)$$

Then, rewriting Eq. 43 in terms of $\exp(-2\mu^*)$, we have:

$$\exp(-2\mu) = \frac{\sigma^{2*} + \mu^{*2}}{\mu^{*4}} \quad (44)$$

We can then solve Eq. 44 for μ . This yields:

$$\begin{aligned} -2\mu &= \log\left(\frac{\sigma^{2*} + \mu^{*2}}{\mu^{*4}}\right) \\ \mu &= -\frac{1}{2} \log\left(\frac{\sigma^{2*} + \mu^{*2}}{\mu^{*4}}\right) \\ \mu &= \log\left(\frac{\mu^{*2}}{\sqrt{\mu^{*2} + \sigma^{2*}}}\right), \end{aligned} \quad (45)$$

where the final step comes from the property that $a \log(b) = \log(b^a)$. Then, we can substitute our expression for μ back into the Equation $\sigma^2 = 2(\log(\mu^*) - \mu)$. This gives us:

$$\begin{aligned} \sigma^2 &= 2(\log(\mu^*) - \mu) \\ &= 2\left(\log(\mu^*) - \log\left(\frac{\mu^{*2}}{\sqrt{\mu^{*2} + \sigma^{2*}}}\right)\right) \\ &= 2 \log\left(\frac{\sqrt{\mu^{*2} + \sigma^{2*}}}{\mu^*}\right) \\ &= \log\left(\frac{\mu^{*2} + \sigma^{2*}}{\mu^{*2}}\right) \\ &= \log\left(1 + \frac{\sigma^{2*}}{\mu^{*2}}\right) \end{aligned} \quad (46)$$

Thus our desired transformations for our moment matching approach are given by $\mu = \log\left(\frac{\mu^{*2}}{\sqrt{\mu^{*2} + \sigma^{2*}}}\right)$ and $\sigma^2 = \log\left(1 + \frac{\sigma^{2*}}{\mu^{*2}}\right)$.

Gompertz SSM priors

In this section, we detail and justify our choice of prior distributions for the Gompertz SSM. Based on the likelihood for the Gompertz SSM in Eq. 22, we require prior distributions for a, b, ϕ, τ , and the initial log-latent state D_0 . Given the interpretation of a and b as the multiplicative constant and the growth rate in the Gompertz model (Eq. 7), we used Uniform($-10, 10$) priors for both a and b . There are important considerations when choosing the prior distributions for ϕ and τ .

Though the likelihood in Eq. 22 can exploit conjugacy and use improper Jeffreys priors (Jeffreys 1946) for both ϕ and τ , Gelman (2006) shows that the posterior is sensitive to the choice of ϵ when using a $\text{Gamma}(\epsilon, \epsilon)$ prior in JAGS to emulate the improper Jeffreys prior. Following the advice of Gelman (2006) and Polson and Scott (2012), we use a central Half-Cauchy prior distribution for the precision parameters. We note that though these studies recommend central Half-Cauchy priors on variance parameters, under this choice of prior the inverse variance (precision) is also implied to have a central Half-Cauchy prior (see Appendix 7.5 for a derivation). We chose the initial condition prior, $\pi(D_0)$, to be normally distributed, with an initial mean of μ_0 , and an initial precision of ϕ_0 . Thus the priors are given by:

$$a \sim \text{Uniform}(-10, 10), \quad (47)$$

$$b \sim \text{Uniform}(-10, 10), \quad (48)$$

$$\phi \sim \text{Half-Cauchy}(\gamma = 100), \quad (49)$$

$$\tau \sim \text{Half-Cauchy}(\gamma = 100), \quad (50)$$

$$D_0 \sim \mathcal{N}(\mu = \mu_0, \sigma^2 = \phi_0^{-1}), \quad (51)$$

with $\mu_0 = 4$, and $\phi_0 = 100$ both fixed.

Moran-Ricker SSM

In this section, we detail and justify our choice of prior distributions for the Moran-Ricker SSM. In particular, we must specify priors for parameters a, b, ϕ, τ , and the initial log-latent state D_0 . Prior distributions for the Moran-Ricker are identical to those chosen for the Gompertz model (Eqs. 47–51). a and b were given Uniform($-10, 10$), ϕ and τ were given diffuse Half-Cauchy priors following recommendations by Gelman (2006) and Polson and Scott (2012), and the initial

latent state D_0 was given a normal prior centered at an initial mean μ_0 with initial precision ϕ_0 . Altogether, the prior equations for the Moran-Ricker SSM are:

$$\begin{aligned} a &\sim \text{Uniform}(-10, 10), \\ b &\sim \text{Uniform}(-10, 10), \\ \phi &\sim \text{Half-Cauchy}(\gamma = 100), \\ \tau &\sim \text{Half-Cauchy}(\gamma = 100), \\ D_0 &\sim \mathcal{N}(\mu = \mu_0, \sigma^2 = \phi_0^{-1}), \end{aligned}$$

with $\mu_0 = 4$, and $\phi_0 = 100$ both fixed.

Lognormal moment matching SSM

In this section, we detail and justify our choice of prior distributions for our Lognormal Moment Matching SSMs. We considered four different SSMs based on the Gompertz and Moran-Ricker models. These included two models with unbiased process evolution and constant variance (LGC and LMRC) and two models with unbiased process evolution and density dependent variance (LGD and LMRD).

For the constant variance models, two of the prior choices were modified from the priors used for the Gompertz and Moran-Ricker models. First, the prior distribution on a was modified to a strictly positive uniform distribution, to reflect the non-negativity of the latent process $X_{1:T}$. Second, the prior on the initial condition, X_0 , was changed from a $\mathcal{N}(\mu = \mu_0, \sigma^2 = \phi_0^{-1})$ to $\log \mathcal{N}(\mu = \mu_0, \sigma^2 = \phi_0^{-1})$ since the constant variance models are not being fit in log-space. Priors for b, ϕ , and τ (Eqs. 48–50) were not changed. Altogether, the priors for the constant variance models are

$$a \sim \text{Uniform}(0, 10), \quad (52)$$

$$b \sim \text{Uniform}(-10, 10), \quad (53)$$

$$\phi \sim \text{Half-Cauchy}(\gamma = 100), \quad (54)$$

$$\tau \sim \text{Half-Cauchy}(\gamma = 100), \quad (55)$$

$$X_0 \sim \log \mathcal{N}(\mu = \mu_0, \sigma^2 = \phi_0^{-1}), \quad (56)$$

with $\mu_0 = \log 4 \approx 1.38$, and $\phi_0 = 100$ both fixed.

Priors for these two density dependent moment matching models were chosen to be identical to those chosen for the Gompertz and Moran-Ricker models (Eqs. 47–51). We justify this by noting that the interpretations of the parameters for the LGD and Gompertz model and the LMRD and Moran-Ricker model are identical for a, b , and X_0 . ϕ and τ were given Half-Cauchy priors based on previous work done by Gelman (2006) and Polson and Scott (2012).

Half-Cauchy precision prior

In this section, we justify our choice of prior distribution on the precision parameters for the Gompertz, Moran-Ricker, and LNM3 models. Gelman (2006) and Polson and Scott (2012) both recommend a central Half-Cauchy for variance parameters in hierarchical models. Then, we would like to investigate the implied prior on $\phi = \sigma^{-2}$.

Suppose that $\sigma^2 \sim \text{Half-Cauchy}(\mu = 0, a = \gamma)$. Then, the probability density function for our prior distribution is given by:

$$\pi(\sigma^2) = \frac{2}{\pi\gamma(1 + (\frac{\sigma^2}{\gamma})^2)}, \sigma^2 > 0 \quad (57)$$

Now we let $\phi = \sigma^{-2}$. The Jacobian of our transformation of variables is given by $|J| = \phi^{-2}$. Using this, we can write the probability density function for the implied prior on ϕ :

$$\pi(\phi) = \frac{2}{\pi\gamma(1 + (\frac{1}{\phi\gamma})^2)} \frac{1}{\phi^2}, \phi > 0 \quad (58)$$

$$= \frac{2}{\pi\gamma(\phi^2 + (\frac{1}{\gamma})^2)} \phi > 0 \quad (59)$$

$$= \frac{2}{\pi\frac{1}{\gamma}(1 + (\phi\gamma)^2)} \phi > 0 \quad (60)$$

This is the probability density function for another central Half-Cauchy distribution, with $a = \gamma^{-1}$. Thus if the prior for σ^2 is $\sigma^2 \sim \text{Half-Cauchy}(\mu = 0, a = \gamma)$, then the implied prior on ϕ is $\phi \sim \text{Half-Cauchy}(\mu = 0, a = \gamma^{-1})$.

Reduced DALEC2 details

In this section we provide additional details on the moments and prior distributions for the reduced DALEC2 model that we use based on the DALEC2 model of Bloom and Williams (2015). Table 6 contains the conditional means and variance for the process evolution and observation components of the LN-SSMs that we considered. These models were: a biased formulation using the reduced DALEC2 model (Column Biased) and an unbiased formulation using the reduced DALEC2 model obtained using our moment matching framework (Column MM).

Table 7 contains information on parameter descriptions, units, and the prior distribution used for model parameters estimated during our analysis. Information for parameter descriptions and units are taken from Bloom and Williams (2015). Prior distributions for process parameters (f_{lab} through c_{rf}) used Uniform priors over the range of acceptable values from Table 1 of Bloom and Williams (2015). ω_f and ω_{lab} ,

Table 6 Expected values and variances for the process evolution functions and observation functions for the moment matching LN-SSM and biased LN-SSM used to model LAI data at UNDE

Model	MM	Biased
$\mathbb{E}[\mathbf{C}^{(t)} \mathbf{C}^{(t-1)}, \Theta]$	$M_t \mathbf{C}^{(t-1)} + p_t$	$(M_t \mathbf{C}^{(t-1)} + p_t) \exp((2\phi)^{-1})$
$\mathbb{V}[\mathbf{C}^{(t)} \mathbf{C}^{(t-1)}, \Theta]$	$\phi^{-1}(M_t \mathbf{C}^{(t-1)} + p_t)$	ϕ
$\mathbb{E}[\mathbf{C}_{obs}^{(t)} \mathbf{C}^{(t)}, \Theta]$	$C_f^{(t)} c_{lma}^{-1}$	$C_f^{(t)} c_{lma}^{-1} \exp((2\tau)^{-1})$
$\mathbb{V}[\mathbf{C}_{obs}^{(t)} \mathbf{C}^{(t)}, \Theta]$	$\tau^{-1}(C_f^{(t)} c_{lma}^{-1})$	$\phi(X_{t-1} \exp(a + bX_{t-1}))^{-2}$

Table 7 Information on the 10 parameters estimated in our reduced DALEC2 model (Bloom and Williams 2015)

Param	Description	Units	Prior
f_{lab}	Proportion of GPP allocated to labile carbon	Unitless	Unif(0.01, 0.5)
f_f	Proportion of GPP allocated to foliage carbon	Unitless	Unif(0.01, 0.5)
d_o	Start day of leaf regrowth onset	Unitless	Unif(1, 365)
d_f	Start day of leaf fall	Unitless	Unif(1, 365)
c_{eff}	Canopy efficiency	Unitless	Unif(10, 100)
c_{lf}	Proportion of leaves lost annually	Unitless	Unif(0.125, 1)
c_{ro}	Length of labile carbon release period	Day	Unif(10, 100)
c_{rf}	Length of leaf fall period	Day	Unif(20, 150)
ω_f	Density dependent variance parameter	$g^5 C^5 m^{-1}$	Unif(0, 1)
ω_{lab}	Density dependent variance parameter	$g^5 C^5 m^{-1}$	Unif(0, 1)

This information includes notation, interpretations of the parameters, units, and the prior distribution used in our Bayesian Lognormal SSM. Upper and lower bounds for process parameters (f_{lab} through c_{rf}) are taken from the table of upper and lower bounds in Bloom and Williams (2015)

the parameters that govern the percentage of density dependent process variance in our reduced DALEC2 LN-SSMs, were given Uniform(0,1) priors.

Author contributions JWS wrote the main manuscript text and prepared figures. LRJ and RQT supervised research and contributed ideas to improve the manuscript. 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/>.

References

Andrieu C, Thoms J (2008) A tutorial on adaptive MCMC. *Stat Comput* 18:343–373. <https://doi.org/10.1007/s11222-008-9110-y>

Andrieu C, Doucet A, Holenstein R (2010) Particle Markov chain Monte Carlo methods. *J R Stat Soc: B* 72(3):269–342. <https://doi.org/10.1111/j.1467-9868.2009.00736.x>

Auger-Méthé M, Field C, Albertsen C et al (2016) State-space models' dirty little secrets: even simple linear gaussian models can have estimation problems. *Sci Rep* 6:266–277. <https://doi.org/10.1038/srep26677>

Auger-Méthé M, Newman K, Cole D et al (2021) A guide to state-space modeling of ecological time series. *Ecol Monogr* 91(4):e014,70. <https://doi.org/10.1002/ecm.1470>

Bessac J, Naveau P (2021) Forecast score distributions with imperfect observations. *Adv Stat Climatol Meteorol Oceanogr* 7(2):53–71. <https://doi.org/10.5194/ascmo-7-53-2021>

Bloom A, Williams M (2015) Constraining ecosystem carbon dynamics in a data-limited world: integrating ecological common sense in a model-data fusion framework. *Biogeosciences* 12(5):1299–1315. <https://doi.org/10.5194/bg-12-1299-2015>

Box GEP (1980) Sampling and Bayes' inference in scientific modelling and robustness. *J R Stat Soc A* 143(4):383–430

Buck-Sorlin G (2013) Process-based model. In: Dubitzky W, Wolkenhauer O, Cho KH et al (eds) Encyclopedia of systems biology. Springer, New York, p 1755. https://doi.org/10.1007/978-1-4419-9863-7_1545

Buckland S, Newman K, Thomas L et al (2004) State-space models for the dynamics of wild animal populations. *Ecol Model* 171(1):157–175. <https://doi.org/10.1016/j.ecolmodel.2003.08.002>

Petrис G, Petrone S, Campagnoli P (2009) Dynamic linear models with R, 1st edn. Use R! Springer New York, NY. https://doi.org/10.1007/b135794_2

Cappe O, Godsill SJ, Moulines E (2007) An overview of existing methods and recent advances in sequential Monte Carlo. *Proc IEEE* 95(5):899–924. <https://doi.org/10.1109/JPROC.2007.893250>

Carter CK, Kohn R (1994) On Gibbs sampling for state space models. *Biometrika* 81(3):541–553. <https://doi.org/10.1093/biomet/81.3.541>

Casella G, Berger R (2002) Statistical inference. Duxbury Press, Pacific Grove

DAAC O (2018) Fixed sites subsetting and visualization tool. oRNL DAAC, Oak Ridge, Tennessee, USA. Subset obtained for MCD15A3Hvproduct at site id us_wisconsin_neon_unde. <https://doi.org/10.3334/ORNLDaac/1567>. Accessed 23 Mar 2022

Dennis B, Patil G (1984) The gamma distribution and weighted multimodal gamma distributions as models of population abundance. *Math Biosci* 68(2):187–212. [https://doi.org/10.1016/0025-5564\(84\)90031-2](https://doi.org/10.1016/0025-5564(84)90031-2)

Dennis B, Patil GP (1988) Applications in ecology. In: Crow EL, Shimizu K (eds) Lognormal distributions. Routledge, Milton Park, pp 303–330

Dennis B, Ponciano JM, Lele SR et al (2006) Estimating density dependence, process noise, and observation error. *Ecol Monogr* 76(3):323–341. [https://doi.org/10.1890/0012-9615\(2006\)76\[323:EDDPNAJ2.0.CO;2](https://doi.org/10.1890/0012-9615(2006)76[323:EDDPNAJ2.0.CO;2)

Deo V, Grover G (2021) A new extension of state-space sir model to account for underreporting—an application to the covid-19 transmission in california and florida. *Results Phys* 24(104):182. <https://doi.org/10.1016/j.rinp.2021.104182>

Dietze MC, Fox A, Beck-Johnson LM et al (2018) Iterative near-term ecological forecasting: needs, opportunities, and challenges. *Proc Natl Acad Sci* 115(7):1424–1432. <https://doi.org/10.1073/pnas.1710231115>

Doucet A, Johansen A (2011) A tutorial on particle filtering and smoothing: fifteen years later. In: Crisan D, Rozovskii B (eds) The Oxford handbook of nonlinear filtering. Oxford University Press, Oxford; N.Y., pp 656–705

Dowd M, Meyer R (2003) A Bayesian approach to the ecosystem inverse problem. *Ecol Model* 168(1):39–55. [https://doi.org/10.1016/S0304-3800\(03\)00186-8](https://doi.org/10.1016/S0304-3800(03)00186-8)

Durbin J, Koopman S (2012) Time series analysis by state space methods, 2nd edn. Oxford University Press, Oxford

Ercanlı I, Günlü A, Şenyurt M et al (2018) Artificial neural network models predicting the leaf area index: a case study in pure even-aged Crimean pine forests from turkey. *For Ecosyst*. <https://doi.org/10.1186/s40663-018-0149-8>

Eriksson D, Pearce M, Gardner J et al (2019) Scalable global optimization via local Bayesian optimization. *Adv Neural Inf Process Syst* 32:5497–5508

Ferro CAT (2017) Measuring forecast performance in the presence of observation error. *Q J R Meteorol Soc* 143(708):2665–2676. <https://doi.org/10.1002/qj.3115>

Gelman A (2006) Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper). *Bayesian Anal* 1(3):515–534. <https://doi.org/10.1214/06-BA117A>

Geman S, Geman D (1984) Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans Pattern Anal Mach Intell PAMI* 6(6):721–741

Gneiting T, Raftery AE (2007) Strictly proper scoring rules, prediction, and estimation. *J Am Stat Assoc* 102(477):359–378. <https://doi.org/10.1198/016214506000001437>

Gompertz B (1825) Xxiv. on the nature of the function expressive of the law of human mortality, and on a new mode of determining the value of life contingencies. in a letter to francis baily, esq. f. r. s. *Philos Trans R Soc Lond* 115:513–583. <https://doi.org/10.1098/rstl.1825.0026>

Good IJ (1952) Rational decisions. *J R Stat Soc B* 14(1):107–114

Gordon NJ, Salmond DJ, Smith AFM (1993) Novel approach to nonlinear/non-gaussian Bayesian state estimation. *IEE Proc F* 140(2):107–113

Gramacy RB (2020) Surrogates: Gaussian process modeling, design and optimization for the applied sciences. CRC, Boca Raton. <http://bobby.gramacy.com/surrogates/>. Accessed 20 Feb 2022

Holm S (1979) A simple sequentially rejective multiple test procedure. *Scand J Stat* 6:65–70

Jeffreys H (1946) An invariant form for the prior probability in estimation problems. *Proc R Soc Lond A* 186(1007):453–461. <https://doi.org/10.1098/rspa.1946.0056>

Jiang J, Huang Y, Ma S et al (2018) Forecasting responses of a northern peatland carbon cycle to elevated CO_2 and a gradient of experimental warming. *J Geophys Res: Biogeosci* 123(3):1057–1071. <https://doi.org/10.1002/2017JG004040>

Jordan A, Krueger F, Lerch S (2017) Evaluating probabilistic forecasts with the r package Scoringrules. *J Stat Softw* 90:1–37. <https://doi.org/10.18637/jss.v090.i12>

Julier SJ, Uhlmann JK (1997) A new extension of the Kalman filter to nonlinear systems. In: The 11th international symposium of aerospace/defense sensing, simulation and controls, multisensor fusion, tracking and resource management II, Orlando, 20–25 April 1997, pp 182–193

Kalman RE (1960) A new approach to linear filtering and prediction problems. *Trans ASME-J Basic Eng* 82(Series D):35–45

Keeling CD, Piper SC, Bacastow RB et al (2005) Atmospheric CO_2 and 13CO_2 exchange with the terrestrial biosphere and oceans from 1978 to 2000: observations and carbon cycle implications. Springer, New York, pp 83–113. https://doi.org/10.1007/0-387-27048-5_5

King AA, Nguyen D, Ionides EL (2016) Statistical inference for partially observed Markov processes via the r package pomp. *J Stat Softw* 69(12):1–43. <https://doi.org/10.18637/jss.v069.i12>

Knape J, Jonzén N, Sköld M (2011) On observation distributions for state space models of population survey data. *J Anim Ecol* 80:1269–77. <https://doi.org/10.1111/j.1365-2656.2011.01868.x>

Krüger F, Lerch S, Thorarinsdottir T et al (2021) Predictive inference based on Markov chain Monte Carlo output. *Int Stat Rev* 89(2):274–301. <https://doi.org/10.1111/insr.12405>

Lambert JH (1758) Observations variae in mathesin puram. *Acta Helvetica, Physico-Mathematico-Anatomico-Botanico-Medica* 3:128–168

Lewis ASL, Woelmer WM, Wander HL et al (2022) Increased adoption of best practices in ecological forecasting enables comparisons of forecastability. *Ecol Appl* 32(2):e02,500. <https://doi.org/10.1002/ea.2500>

Little R (2012) Calibrated Bayes, an alternative inferential paradigm for official statistics. *J Off Stat* 28:309–334

Little RJ (2006) Calibrated Bayes. *Am Stat* 60(3):213–223. <https://doi.org/10.1198/000313006X117837>

Luo Y, Weng E, Wu X et al (2009) Parameter identifiability, constraint, and equifinality in data assimilation with ecosystem models. *Ecol Appl* 19(3):571–574. <https://doi.org/10.1890/08-0561.1>

Mahowald N, Lo F, Zheng Y et al (2016) Projections of leaf area index in earth system models. *Earth Syst Dyn* 7(1):211–229. <https://doi.org/10.5194/esd-7-211-2016>

Matheson JE, Winkler RL (1976) Scoring rules for continuous probability distributions. *Manage Sci* 22(10):1087–1096

Maunder MN, Deriso RB, Hanson CH (2015) Use of state-space population dynamics models in hypothesis testing: advantages over simple log-linear regressions for modeling survival, illustrated with application to longfin smelt (*Spirinchus thaleichthys*). *Fish Res* 164:102–111. <https://doi.org/10.1016/j.fishres.2014.10.017>

Mäntyniemi SHP, Whitlock RE, Perälä TA et al (2015) General state-space population dynamics model for Bayesian stock assessment. *ICES J Mar Sci* 72(8):2209–2222. <https://doi.org/10.1093/icesjms/fsv117>

National Ecological Observatory Network (2020) Woody plant vegetation structure, Data Product DP1.10098.001, Provisional data downloaded from <http://data.neonscience.org>. Accessed 21 Apr 2020

National Ecological Observatory Network (NEON) (2022a) Shortwave radiation (primary pyranometer) (dp1.00022.001). <https://data.neonscience.org/data-products/DP1.00022.001>. Accessed 20 Feb 2022

National Ecological Observatory Network (NEON) (2022b) Triple aspirated air temperature (dp1.00003.001). <https://data.neonscience.org/data-products/DP1.00003.001>. Accessed 20 Feb 2022

Osthus D, Hickmann KS, Caragea PC et al (2017) Forecasting seasonal influenza with a state-space sir model. *Ann Appl Stat*. <https://doi.org/10.1214/16-AOAS1000>

Plummer M (2003) JAGS: a program for analysis of Bayesian graphical models using Gibbs sampling. In: Proceedings of the 3rd international workshop on distributed statistical computing (DSC 2003), Vienna, 20–22 March 2003, pp 1–10

Plummer M (2019) Rjags: Bayesian graphical models using MCMC. R package version 4-10. <https://CRAN.R-project.org/package=rjags>. Accessed 20 Feb 2022

Polson NG, Scott JG (2012) On the Half-Cauchy prior for a global scale parameter. *Bayesian Anal* 7(4):887–902. <https://doi.org/10.1214/12-BA730>

R Core Team (2016) R: a language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org/>

Ricker WE (1954) Stock and recruitment. *J Fish Res Board Can* 11(5):559–623. <https://doi.org/10.1139/f54-039>

Robert CP, Casella G (2005) Monte Carlo statistical methods (Springer texts in statistics). Springer, Berlin

Rosenthal S (2009) Optimal proposal distributions and adaptive MCMC. In: Brooks S, Gelman A, Jones GL, Meng X-L (eds) *Handbook of Markov Chain Monte Carlo*, Chap 4. Chapman & Hall/CRC, pp 91–110

Roulston MS, Smith LA (2002) Evaluating probabilistic forecasts using information theory. *Mon Weather Rev* 130(6):1653–1660. [https://doi.org/10.1175/1520-0493\(2002\)130<1653:EPFUIT>2.0.CO;2](https://doi.org/10.1175/1520-0493(2002)130<1653:EPFUIT>2.0.CO;2)

Rubin DB (1984) Bayesianly justifiable and relevant frequency calculations for the applied statistician. *Ann Stat* 12(4):1151–1172. <https://doi.org/10.1214/aos/1176346785>

Serbin SP, Wu J, Ely KS et al (2019) From the arctic to the tropics: multibiome prediction of leaf mass per area using leaf reflectance. *New Phytol* 224(4):1557–1568. <https://doi.org/10.1111/nph.16123>

Shumway R, Stoffer D (2011) Time series analysis and its applications with R examples, vol 9. Springer, pp 289–384. <https://doi.org/10.1007/978-1-4419-7865-3>

Smith RL, Miller JE (1986) A non-gaussian state space model and application to prediction of records. *J R Stat Soc B* 48(1):79–88

Thomas Q, Brooks E, Jersild A et al (2017) Leveraging 35 years of *pinus taeda* research in the southeastern us to constrain forest carbon cycle predictions: regional data assimilation using ecosystem experiments. *Biogeosciences* 14:3525–3547. <https://doi.org/10.5194/bg-14-3525-2017>

West M, Harrison J (1997) Bayesian forecasting and dynamic models, 2nd edn. Springer, Berlin

Wilcoxon F (1945) Individual comparisons by ranking methods. *Biometrics* 1:196–202

Williams M, Rastetter EB, Fernandes DN et al (1997) Predicting gross primary productivity in terrestrial ecosystems. *Ecol Appl* 7(3):882–894. [https://doi.org/10.1890/1051-0761\(1997\)007\[0882:PGPPIT\]2.0.CO;2](https://doi.org/10.1890/1051-0761(1997)007[0882:PGPPIT]2.0.CO;2)

Yang W, Tan B, Huang D et al (2006) Modis leaf area index products: from validation to algorithm improvement. *IEEE Trans Geosci Remote Sens* 44(7):1885–1898. <https://doi.org/10.1109/TGRS.2006.871215>

Authors and Affiliations

John W. Smith^{1,4} · R. Quinn Thomas^{2,3} · Leah R. Johnson^{3,4}

✉ John W. Smith
john.smith20@montana.edu

R. Quinn Thomas
rqthomas@vt.edu

Leah R. Johnson
lrjohn@vt.edu

¹ Department of Mathematical Sciences, Montana State University, 974 West Cleveland Street, Bozeman, MT 59715, USA

² Department of Forest Resources and Environmental Conservation, Virginia Tech, 310 West Campus Drive, Blacksburg, VA 24060, USA

³ Department of Biological Sciences, Virginia Tech, 926 West Campus Drive, Blacksburg, VA 24060, USA

⁴ Department of Statistics, Virginia Tech, 250 Drillfield Drive, Blacksburg, VA 24060, USA