

Models for Bounded Systems with Continuous Dynamics

Amanda R. Cangelosi and Mevin B. Hooten*

Department of Mathematics and Statistics, Utah State University, 3900 Old Main Hill,
Logan, Utah 84322, U.S.A.

**email*: amanda.c@aggiemail.usu.edu

SUMMARY. Models for natural nonlinear processes, such as population dynamics, have been given much attention in applied mathematics. For example, species competition has been extensively modeled by differential equations. Often, the scientist has preferred to model the underlying dynamical processes (i.e., theoretical mechanisms) in continuous time. It is of both scientific and mathematical interest to implement such models in a statistical framework to quantify uncertainty associated with the models in the presence of observations. That is, given discrete observations arising from the underlying continuous process, the unobserved process can be formally described while accounting for multiple sources of uncertainty (e.g., measurement error, model choice, and inherent stochasticity of process parameters). In addition to continuity, natural processes are often bounded; specifically, they tend to have nonnegative support. Various techniques have been implemented to accommodate nonnegative processes, but such techniques are often limited or overly compromising. This article offers an alternative to common differential modeling practices by using a bias-corrected truncated normal distribution to model the observations and latent process, both having bounded support. Parameters of an underlying continuous process are characterized in a Bayesian hierarchical context, utilizing a fourth-order Runge–Kutta approximation.

KEY WORDS: Bias correction; Differential equations; Hierarchical models; Numerical approximations; Population dynamics; Truncated normal.

1. Introduction

The dynamics of natural systems have been extensively studied by theoretical scientists and applied mathematicians to describe the underlying processes (i.e., mechanisms) assumed to drive the phenomena of interest (e.g., population dynamics). For various scientific and philosophical reasons, the underlying theoretical mechanism is often assumed to be continuous in time, and thus the scientist assigns one of numerous possible continuous process models, usually involving differential equations, to describe the dynamical system of interest (e.g., Turchin, 2003, pp. 93–108). This article addresses the need to quantify uncertainty associated with such mathematical models, highlighting the advantage of a statistical approach. Statistical implementations of mathematical models for describing dynamical ecological systems are appearing more frequently in the literature (e.g., Wikle, 2003; Buckland et al., 2007; Hooten and Wikle, 2007; Gianni, Pasquali, and Ruggeri, 2008; Johnson et al., 2008), though they are often coarsely discretized. While there are situations that naturally warrant the use of discrete models (e.g., matrix models; Caswell, 2001), this article focuses on the implementation of continuous models to describe underlying processes, even when the corresponding observations may be discrete.

In addition to continuity, natural processes are often bounded, meaning that their quantification is restricted to some subset of the real number line. Almost exclusively, natural processes are measured in quantities of mass, thus restricting the continuous process model (as well as the observations) to nonnegative realizations. Numerous techniques have been utilized to accommodate issues of bounded support

(e.g., log transformations), but often these techniques misrepresent system dynamics or fail to accommodate complicated nonlinear models. When data contain zeros as realizations (e.g., species competition, rainfall data, rare species counts), further complications arise. A common but ad hoc approach is to add a small number so that a log transformation can be taken. While these practices are widespread, they often do not represent the dynamics well. To accommodate the occurrence of zeros in data, discrete probability models (e.g., Poisson) and zero-inflated models are commonplace. However, discrete probability models are reasonable only when the support of the data is discrete; this is not often the case for many measurements of mass (e.g., plant basal area quadrat cover). Zero-inflated models are reasonable only when a natural model exists for the zero process; however, in many situations, zero data clearly belong to the same dynamical process as the data with positive support. This article offers an alternative method of accommodating bounded data containing zeros through the incorporation of a bias-corrected truncated normal model.

We focus on a general methodology that one may use to quantify the uncertainty associated with a given continuous model, as well as introduce a bias-corrected truncated normal distribution as a general and robust model for the data and underlying dynamical process. The truncated normal distribution is particularly desirable over the gamma or lognormal distributions, as it allows truncation boundaries to occur anywhere along the real line; specifically, the truncated normal distribution allows zero-valued realizations to have nonzero density. We show that within a hierarchical

framework (Berliner, 1996), a bias-corrected truncated normal data model can be utilized to accommodate observations having bounded support. In addition, a similar bias-corrected truncated normal physical process model is introduced to describe both single-species population growth (logistic growth) and interspecies competition (Lotka–Volterra equations, e.g., Edelstein-Keshet, 1988, p. 224–231). Although the example presented here contains nonnegative integer observations, we implement the truncated normal data model (which has continuous support) to demonstrate its general utility.

We also utilize a fourth-order Runge–Kutta (RK4) approximation to the system of differential equations in conjunction with Markov chain Monte Carlo (MCMC) methods to implement the formal statistical model. While a high-order approximation, such as RK4, is not new to statistics (e.g., Rumelin, 1982; Scipione and Berliner, 1993; Shoji and Ozaki, 1998), its implementation within a hierarchical framework is underutilized. RK4 is a useful approximation, well known for its accuracy, efficiency, robustness, and applicability to dynamical processes of any finite dimension.

To illustrate our methods, we use a historical data set (Gause, 1934) containing population measurements of *Paramecium aurelia* and *Paramecium caudatum* grown in nutrient medium both separately and together. Although this particular data set has been used in numerous studies (Leslie, 1957; Edelstein-Keshet, 1988; Pascual and Kareiva, 1996; Lele, Dennis, and Lutscher, 2007), often serving as a textbook example for biological logistic growth models, many previous studies utilize low-order discretizations, wherein population size is exclusively modeled at times coinciding with those of the observed data. In contrast, we augment the process so that it can be modeled at arbitrarily small time increments. It should be noted that the focus here is concerned mainly with the methods, while the data are used primarily for illustration.

In the following section, the truncated normal distribution is briefly described and the necessity of a bias correction within the hierarchical framework is explained. The process model, wherein RK4 approximation is utilized, is outlined and followed by a discussion of the entire hierarchical model. The “Results” section includes details about model fits pertaining to both simulated data and the Gause data. The “Discussion” section provides brief interpretations of results from the classical Gause data, followed by a detailed assessment of the methodology in general, a comparison to other work, and possible extensions.

2. Methods

2.1 The Truncated Normal Distribution

If $X \sim N(\mu, \sigma^2)$ and $I \subset R$ is an interval, then $P(X \leq x | X \in I)$, defined for any real x , is the truncated normal distribution of X (adapted from Rohatgi, 1976, p. 116). If a normal distribution is truncated, the probability mass from the tail(s) is dispersed throughout the remaining portion. Thus, if a normal density is left-truncated at zero, then the new mode resulting from truncation is necessarily located in one of two places: If the original (untruncated) mode was positive, then the mode of the left-truncated density is unchanged (though the expectation is not); if the original mode was negative, then the

mode of the left-truncated density is zero. This characteristic becomes important in justifying our choice of parameter estimation.

The truncated normal (T.N.) distribution can accommodate random variables with nonnegative support, but a more appropriate model can be specified by implementing a bias correction in the truncated normal with a target expectation. To illustrate the suitability of such a bias correction, first suppose $X \sim \text{T.N.}(\mu, \sigma^2)$, where we desire $E(X) = \mu > 0$. The actual expectation is not μ however, but rather

$$E(X) = \mu + \sigma \left(\frac{\phi\left(\frac{b_1 - \mu}{\sigma}\right) - \phi\left(\frac{b_2 - \mu}{\sigma}\right)}{\Phi\left(\frac{b_2 - \mu}{\sigma}\right) - \Phi\left(\frac{b_1 - \mu}{\sigma}\right)} \right), \quad (1)$$

where b_1 and b_2 are the lower and upper bounds of truncation, and ϕ and Φ are the standard normal probability density function and cumulative distributive function, respectively (Johnson, Kotz, and Balakrishnan, 1994; Horrace, 2005). That is, the expectation is not equivalent to the parameter μ , as desired. In particular, when $b_1 = 0$ and $b_2 = \infty$, we have

$$E(X) = \mu + \sigma \left(\frac{\phi\left(\frac{\mu}{\sigma}\right)}{\Phi\left(\frac{\mu}{\sigma}\right)} \right). \quad (2)$$

Though simpler than that in equation (1), the expectation in equation (2) is a complicated nested integral equation that is analytically intractable; however, it can be approximated numerically. If we specify $X \sim \text{T.N.}(g, \sigma^2)$, where g is a function of μ and σ^2 such that $E(X) = \mu$, then we need only find the bias-correcting function g . That is, for $b_1 = 0$ and $b_2 = \infty$, one can approximate a truncated normal distribution with a target expectation μ by solving for the function g such that

$$E(X) = g + \sigma \left(\frac{\phi\left(\frac{g}{\sigma}\right)}{\Phi\left(\frac{g}{\sigma}\right)} \right) = \mu. \quad (3)$$

In addition to obtaining an appropriate distribution within the hierarchical model (i.e., one that obtains the target expectation), the implementation of this bias correction also decreases process uncertainty (i.e., standard deviation of the distribution), particularly near truncation bounds, as often appears in data sets and is depicted in Figure 1.

2.2 Constructing the Process Model

Let $\frac{dx}{dt} = h(t, \mathbf{x})$ be a differential equation (or system of differential equations) that models a continuous dynamical process \mathbf{x} in time t . Such a process can be well approximated by the classical RK4 method (Burden and Faires, 2001):

$$\mathbf{x}_{t+\Delta t} = \mathbf{x}_t + \frac{\Delta t}{6}(a_1 + 2a_2 + 2a_3 + a_4) \quad (4)$$

$$\equiv f(\mathbf{x}_t, \boldsymbol{\theta}), \quad (5)$$

where $a_1 = h(t, \mathbf{x}_t)$, $a_2 = h(t + \frac{\Delta t}{2}, \mathbf{x}_t + \frac{\Delta t}{2}a_1)$, $a_3 = h(t + \frac{\Delta t}{2}, \mathbf{x}_t + \frac{\Delta t}{2}a_2)$, $a_4 = h(t + \Delta t, \mathbf{x}_t + a_3\Delta t)$, and $\boldsymbol{\theta}$ is a vector of parameters. Thus, the process at time $t + \Delta t$ is determined by

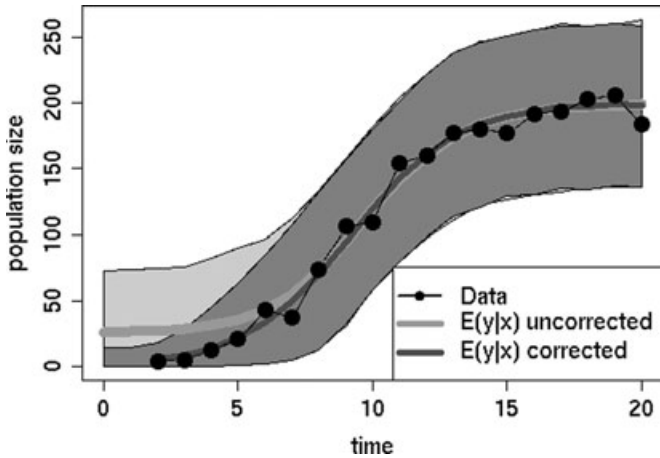


Figure 1. Effect of bias correction on uncertainty. Assume a truncated normal data model, left-truncated at zero. The dark area depicts a probability envelope for a data model containing a bias correction; the light area is that for a model without a bias correction. Note the shift in the mean and reduction of uncertainty in the dark envelope near the truncation boundary.

the process at time t plus the product of the size of the time interval (Δt) and an estimated weighted average of slopes $(\frac{a_1 + 2a_2 + 2a_3 + a_4}{6})$, where a_i is an estimated slope at one of four points within the time interval Δt . Note that the approximation for $\mathbf{x}_{t+\Delta t}$ is computed using four points within the time interval; this accounts for the accuracy of RK4, as the error at each of the four steps is on the order of $(\Delta t)^5$ and has an overall error on the order of $(\Delta t)^4$.

The RK4 approximation, which we have labeled $f(\mathbf{x}_t, \boldsymbol{\theta})$, can then be embedded within a hierarchical probability model at the process level. The process is “augmented” in the following way: If the data consist of an $l \times m \times n$ array (e.g., m variables observed at n points in time, replicated l times), then the process will have dimension $l \times m \times (\frac{n}{\Delta t} + 1)$, where $0 < \Delta t < 1$. Note that, depending on the value of Δt , the process can be approximated to a desired level of precision. If the process is allowed to be stochastic at each discretization time t , while Δt is specified to be small, then a state-space model with Markov dependence and an augmented process component results; i.e., $\mathbf{x}_t | \mathbf{x}_{t-\Delta t}, \boldsymbol{\theta} \sim [\mathbf{x}_t | f(\mathbf{x}_{t-\Delta t}, \boldsymbol{\theta})]$, where temporal resolution of \mathbf{x}_t is finer than that of the data.

In our case, $f(\mathbf{x}_{t-\Delta t}, \boldsymbol{\theta})$ is the RK4 approximation to the Lotka–Volterra model for two-species competition:

$$\frac{dx_1}{dt} = r_1 x_1 \frac{k_1 - x_1 - \beta_1 x_2}{k_1}, \tag{6}$$

$$\frac{dx_2}{dt} = r_2 x_2 \frac{k_2 - x_2 - \beta_2 x_1}{k_2}, \tag{7}$$

where the parameter vector $\boldsymbol{\theta}$ contains $r_1, k_1, \beta_1, r_2, k_2$, and β_2 , the parameters controlling the dynamics of the system, each of which is assumed to exist within the interval $(0, \infty)$. Specifically, for the i th species, x_i is the underlying continuous process of population growth, r_i is the growth rate, k_i is

the carrying capacity, and β_i is the competition parameter. Regarding competition parameters, β_1 denotes the deleterious effect of species 2 on species 1, and vice versa. Note that in the absence of one species, this particular process model defaults to logistic growth of the other species, and the species competition model can be generalized to any number of species (Edelstein-Keshet, 1988, p. 231–236). Further note that there are a wide variety of systems of differential equations in the literature with which one can model a given continuous population process, each having strengths and weaknesses (Turchin, 2003).

2.3 The Hierarchical Representation

To introduce the hierarchical model, we begin with the model for observations (i.e., likelihood). Let observations for the process under study at time $t = 1, \dots, n$ be denoted by \mathbf{y}_t , an $l \times m \times n$ data array, where l is the number of replicates, m is the number of species observed, and n is the number of times each species is observed. Then a probability model can be written for these measurements:

$$\mathbf{y}_t | \mathbf{x}_t, \sigma_y^2 \sim \text{T.N.}(\mathbf{g}_{y_t}, \boldsymbol{\Sigma}_y)_{b_1}^{b_2}, \boldsymbol{\Sigma}_y = \sigma_y^2 \mathbf{I}, \tag{8}$$

where \mathbf{g}_{y_t} is a bias correction such that $E(\mathbf{y}_t) = \mathbf{x}_t > \mathbf{0}$, and b_1, b_2 are the lower and upper truncation boundaries, respectively. We focus on the expectation (rather than the median or mode) for two reasons: First, the mean is a natural parameter for unconstrained models and it is convenient for reparameterization; second, if the uncorrected mode is negative, then the corrected mode is necessarily zero. Note that we assigned the same variance (σ_y^2) to each species at the data level because we expect consistent measurement error between species; we could have easily allowed the variance of each species at the data level to be different.

The data model in equation (8) is conditional on the second level in the hierarchy, the process \mathbf{x}_t . Let us adopt the formulation described in Section 2.2 as a model for the process:

$$\mathbf{x}_t | \mathbf{x}_{t-\Delta t} \sim \text{T.N.}(\mathbf{g}_{x_t}, \boldsymbol{\Sigma}_x)_{b_1}^{b_2}, \tag{9}$$

$$\boldsymbol{\Sigma}_x = \begin{pmatrix} \sigma_{x_1}^2 & 0 \\ 0 & \sigma_{x_2}^2 \end{pmatrix},$$

where $\boldsymbol{\Sigma}_x$ is a diagonal covariance matrix, and \mathbf{g}_{x_t} is a bias-correcting function (dependent upon $f(\mathbf{x}_{t-\Delta t}, \boldsymbol{\theta}), \sigma_x^2, b_1$, and b_2) such that $E(\mathbf{x}_t) = f(\mathbf{x}_{t-\Delta t}, \boldsymbol{\theta}) > \mathbf{0}$, the RK4 approximation to the Lotka–Volterra species competition model at time t .

The latent variables \mathbf{x}_t can be thought of as unobserved state vectors and may occur at a finer temporal resolution than the data, \mathbf{y}_t . In other words, if $\boldsymbol{\tau}_x, \boldsymbol{\tau}_y$ are finite sets of times, then $|\{\mathbf{x}_t, \forall t \in \boldsymbol{\tau}_x\}| > |\{\mathbf{y}_t, \forall t \in \boldsymbol{\tau}_y\}|$, where “|·|” represents the size of the set. This point is critical for obtaining a precise approximation to the motivating system of differential equations, because in practice we could choose the set $\boldsymbol{\tau}_x$ to be large enough to attain any desired precision in the approximation.

The parameter model makes up the third and final level of the hierarchy, controlling the dynamics of the system (whence, $\boldsymbol{\theta}$) as well as additional stochasticity (whence, σ_x^2 ,

σ_y^2 , the uncertainty associated with the process and data models, respectively).

$$\begin{aligned}\boldsymbol{\theta} &\sim \text{T.N.}(\boldsymbol{\mu}_\theta, \boldsymbol{\Sigma}_\theta)_{d_1}^{d_2}, \\ \mathbf{x}_0 &\sim \text{T.N.}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)_{b_1}^{b_2}, \\ \sigma_y^2 &\sim \text{Inverse gamma}(r_y, q_y), \\ \sigma_x^2 &\sim \text{Inverse gamma}(r_x, q_x).\end{aligned}$$

The truncated normal distributions for $\boldsymbol{\theta}$ and \mathbf{x}_0 are desirable because the distributions are both flexible and proper, guaranteeing a proper posterior. The initial state vector \mathbf{x}_0 accommodates uncertainty in the process at the time period before data are available. In the case where Δt is small, this initial random state vector resembles the process at the time when data are first available; thus we specify its prior distribution to have a location near that of the data where $t = 1$ and a diagonal covariance matrix with variance components larger than the variance in the data at any time point. Locational hyperparameters pertaining to the dynamic nature of the process (i.e., $\boldsymbol{\mu}_\theta$) can be specified to provide reasonable dynamics while still allowing for a wide range of behavior. Specifically, it is possible to fully specify the covariance matrices $\boldsymbol{\Sigma}_\theta$ and $\boldsymbol{\Sigma}_0$ to allow for prior dependence between parameters, but in the absence of specific knowledge about the form of dependence, we allowed them to be diagonal with standard deviations proportional to the magnitudes of the location parameters (e.g., larger prior variability for larger parameters such as k_1 and k_2).

The inverse-gamma priors were specified to be conservative in location and wide in variability. Variance components can be modeled in different ways (Gelman, 2006), though the specification used here yielded a stable algorithm and performed well. In hierarchical models without significant structure at the process level, identifiability issues may arise with the variance components (Dennis et al., 2006). In this specification, substantial structure is provided by the latent dynamical system, thus helping partition the sources of uncertainty.

2.4 Model Implementation and Inference

The hierarchical model described above can be implemented in an MCMC setting in the usual manner by analytically identifying full-conditional distributions for parameters and latent state vectors, and then sampling from them sequentially. In this case, the truncated distributions imply nonconjugacy in all parameters and a Metropolis approach must be taken in the MCMC algorithm. Additionally, although the bias corrections are functions of $\{\mathbf{x}_t\}$, they are analytically intractable and thus must be approximated numerically, using one of several possible numerical optimization routines to find the solution.

To illustrate our numerical method of choice for the univariate case, suppose we desire a truncated normal random variable y to have expectation μ . Recall that, without a bias correction, the expectation would be as specified in equation (1). Thus, we need to approximate the function g such that $E(y)$ appears as in equation (3). Then each time that a bias-corrected truncated normal density needs to be evaluated in the MCMC algorithm, we can numerically find g conditional on μ , σ^2 , b_1 , b_2 . This bias correction is, in effect, a reparameterization that results in a more appropriate probability model. As described below, this method is generalized for the multivariate case.

parameterization that results in a more appropriate probability model. As described below, this method is generalized for the multivariate case.

We employ a linear interpolation method (see Web Appendix A for details and an example) in order to exploit efficient matrix operations in the R programming environment (R Core Development Team, 2007) and find that it is precise in approximating the well-behaved function g for our purposes here. That is, in general, if we desire a truncated normal random variable \mathbf{y} to have expectation $\boldsymbol{\mu}$, then we can find a reparameterization $g(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, using the method described in Web Appendix A, such that $E(\mathbf{y}) = \boldsymbol{\mu}$, for $\mathbf{y} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma} \sim \text{T.N.}(g(\boldsymbol{\mu}, \boldsymbol{\Sigma}), \boldsymbol{\Sigma})$, where $\boldsymbol{\Sigma}$ is a diagonal matrix. Performance of numerical algorithms vary by situation and thus should be experimented with for best results in terms of precision and efficiency.

3. Results

In fitting the model described in the previous section to data, we actually fit three separate models: one univariate logistic growth model to each of the single-species data, followed by a competition model using the two-species data. The single-species models were implemented by setting competition parameters equal to zero, in which case the Lotka–Volterra equations default to logistic growth. In this case, because populations were observed both in isolation and in competition (independently), the advantage is that posterior distributions from the single-species models can be used to inform prior distributions for germane parameters (i.e., r_1 , r_2 , k_1 , k_2) in the two-species model. Thus, single-species posterior means of r_1 , k_1 , r_2 , and k_2 were used within the competition model to then obtain estimates of the competition parameters, β_1 and β_2 . This use of independent data to gain a priori scientific knowledge about the process is a strength of the Bayesian approach and allows us to focus all available power on the estimation of the additional competition parameters (i.e., β_1 , β_2). We set $\Delta t = \frac{1}{4}$, yielding an augmented process with four times the temporal resolution of the data and a very precise approximation to the differential equations. The augmentation level of $\frac{1}{4}$ was chosen because it represented the process well for this scale in the simulated data; it was large enough for computational efficiency, while providing an accurate approximation. Higher-order systems may require a smaller Δt depending on the dynamics. The MCMC algorithms were run for 20,000 iterations with a burn-in of 2000 chosen by study of trace plots (see Web Appendix B).

To assess the model's capability in situations similar to the Gause data, the model was evaluated using simulated data first; the results of this are summarized in Table 1. Figures 2–5 depict results of the model, applied to Gause's *Paramecium* data. Figure 2 shows the data for single-species logistic growth, overlaid with the augmented process 95% credible intervals. With regard to the model fit using the Gause data, Figure 3 shows posterior distributions for single-species logistic growth rates and carrying capacities. Posterior mean single-species growth rates (r_1 and r_2) were 0.6762 and 1.068 with standard deviations 0.1684 and 0.5109, for *P. aurelia* and *P. caudatum*, respectively. Here, the prior distributions were $[r_1] = [r_2] = \text{T.N.}(0.5, 100)_{d_{r_1}}^{d_{r_2}}$, having

Table 1
Results from simulated data

Parameter	Truth	Posterior mean	Posterior S.D.	Prior mean	Prior S.D.
Single species 1					
r_1	0.5	0.5313	0.1473	0.5	10
k_1	600	598.8	45.45	650	316.2
Single species 2					
r_2	0.3	0.3547	0.1360	0.5	10
k_2	400	419.7	41.67	350	316.2
Species competition					
r_1	0.5	0.4873	0.0821	0.5313	0.1473
k_1	600	604.6	47.59	598.8	45.45
β_1	2	1.989	0.5363	1	3.162
r_2	0.3	0.2945	0.1260	0.3547	0.1360
k_2	400	410.8	42.42	419.7	41.67
β_2	0.5	0.5870	0.3336	1	3.162

truncation bounds $d_{r_1} = 0$ and $d_{r_2} = \infty$. Posterior mean single-species carrying capacities (k_1 and k_2) were 564.4 and 202.2, with standard deviations 26.12 and 24.15, for *P. aurelia* and *P. caudatum*, respectively. The prior distributions for the carrying capacities were $[k_1] = \text{T.N.}(600, 10^5)_{d_{k_1}}$ and $[k_2] = \text{T.N.}(200, 10^5)_{d_{k_2}}$, each with truncation bounds $d_{k_1} = 0$, and $d_{k_2} = \infty$.

Figure 4 shows the data for two-species competition, overlaid with the augmented process credible intervals, after using the single-species growth rate and carrying capacity parameters to inform the two-species priors. Figure 5 shows posterior distributions for competition parameters (β_1 and β_2). Posterior means of the competition parameters were 2.553 and 0.4742 with standard deviations 1.152 and 0.2916, for *P. aurelia* and *P. caudatum*, respectively. Here, the prior distributions were $[\beta_1] = [\beta_2] = \text{T.N.}(1, 10)_{d_{\beta_1}}$, having truncation bounds $d_{\beta_1} = 0$ and $d_{\beta_2} = \infty$. The hyperparameters for both single-species and competition model variance components were $r_y = r_x = 10^{-3}$ and $q_y = q_x = 2.01$, yielding $E(\sigma_y^2) = E(\sigma_x^2) = 1000$ and $\text{Var}(\sigma_y^2) = \text{Var}(\sigma_x^2) = 10^8$.

4. Discussion

Overall, model fits suggested that the data contain significant information about the process and parameters. More specifically, regarding the simulated data summarized in Table 1, note that the posterior variance is reduced for the growth rates in the competition model, as compared to the single-species cases. This suggests that the competition data hold additional information about the single-species growth parameters. Further, note that each of the parameters used for simulation is indeed captured by the model, suggesting that the model describes the simulated process accurately. Regarding the Gause data, the overlap in posterior growth rates for single-species (Figure 3a) provides little evidence of a significant difference, although the variability differs. The fact that the posterior variability of r_2 is greater than that of r_1 may seem surprising because the *P. caudatum* observations appear to have less spread than the *P. aurelia* observations. Perhaps the difference in variability is due to the fact that *P. caudatum* initially experiences a delayed growth relative to *P. aurelia*, yet *P.*

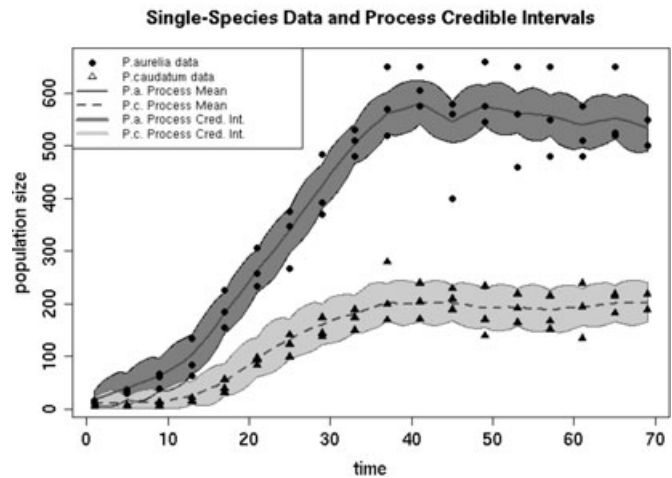


Figure 2. *P. aurelia* and *P. caudatum* observations with augmented process 95% credible interval overlaid. Each species was grown in its own medium, independent of the other species (i.e., noncompetition).

caudatum appears to meet its carrying capacity earlier than *P. aurelia*. Figure 3b suggests that *P. aurelia* is naturally capable of attaining a larger population size than *P. caudatum* in the absence of the other species. Figure 5 evinces a substantial difference among competition parameters; perhaps *P. caudatum* is more affected as a percentage of its carrying capacity. One important aspect with regard to the augmented process is that, similar to prediction uncertainty in all statistical models, the credible intervals of the augmented process indicate increased uncertainty at time points farther away from data, as seen in Figures 2 and 4.

Regarding the methodology of our model, the proposed bias-corrected truncated normal models alleviate the need for log transformations and allow us to model a nonnegative process. It should be noted that both the process and data models have bounded support, whereas other studies (e.g., Stein, 1992) have modeled bounded data which were assumed to arise from a measurement model conditioned on a process with real support, rendering no need for a bias correction.

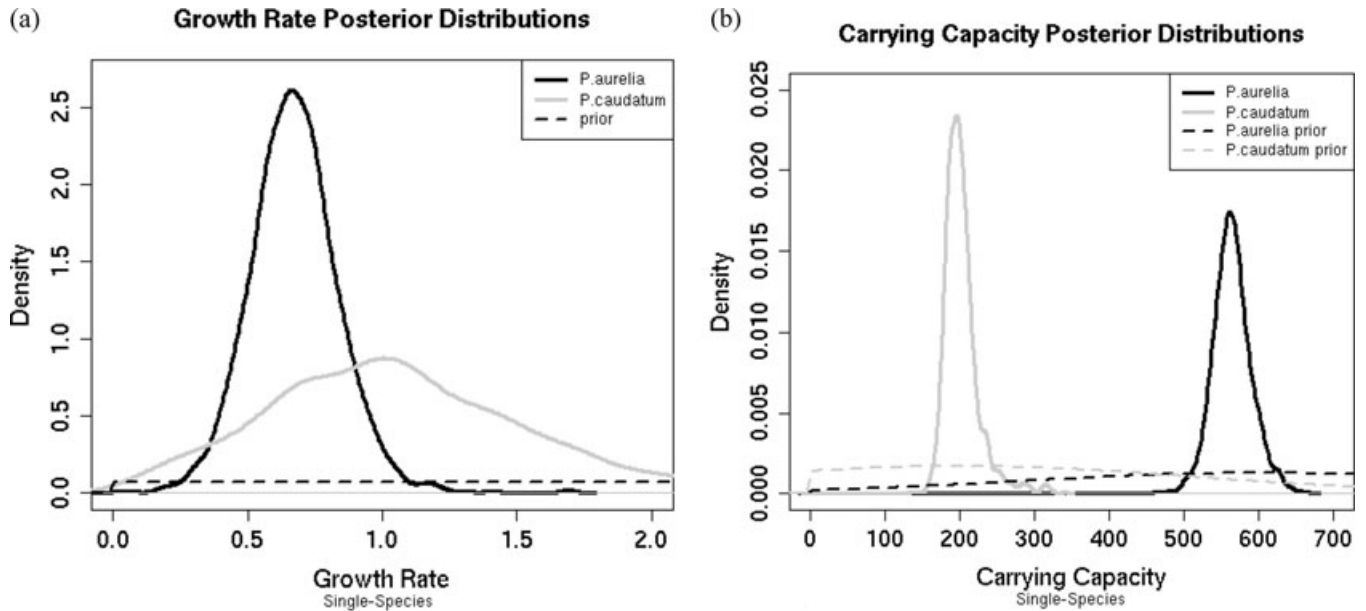


Figure 3. Posterior distributions for single-species logistic growth rates (a) and carrying capacities (b), with prior distribution overlaid.

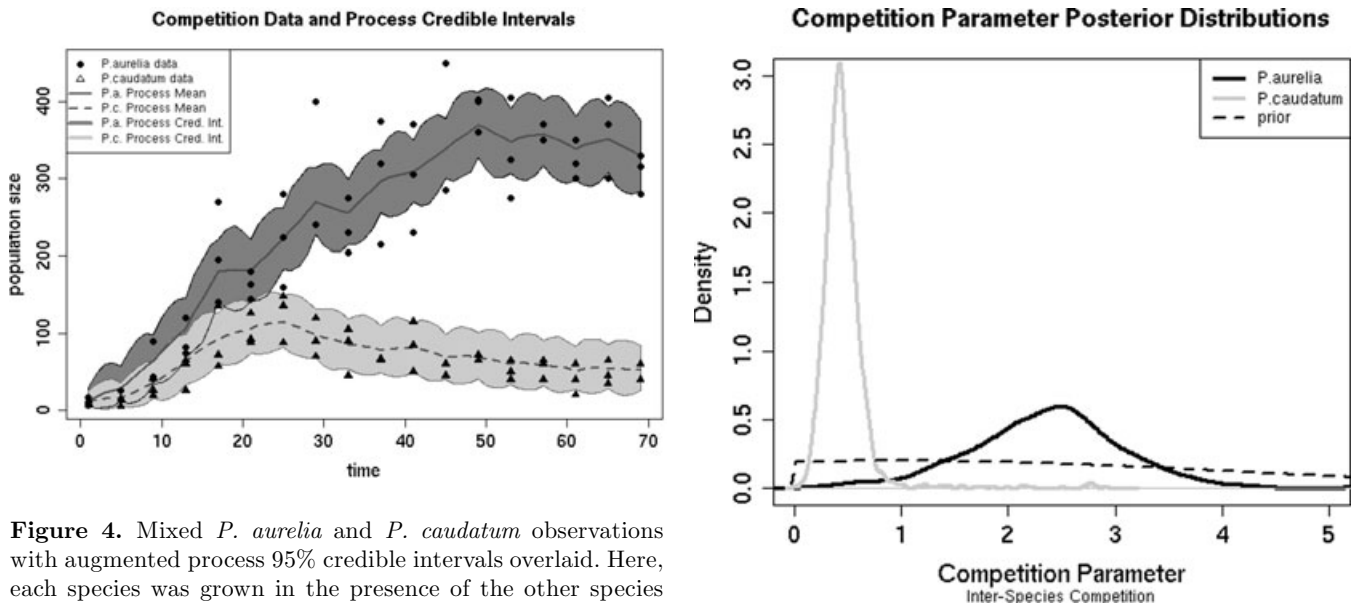


Figure 4. Mixed *P. aurelia* and *P. caudatum* observations with augmented process 95% credible intervals overlaid. Here, each species was grown in the presence of the other species (i.e., interspecies competition).

Figure 5. Posterior distributions for competition parameters, with prior distribution overlaid.

The implementation of a bias correction allows us to specify an appropriate physical process model with positive support. Furthermore, utilizing a finer temporal discretization in the process ($\Delta t < 1$) improves stability of the dynamical system and avoids drawbacks pertaining to the representation of dynamics that accompany many analytically discretized models (e.g., Ricker growth; Turchin, 2003). The process augmentation, resulting from $\Delta t < 1$, also allows our approximation to be faithful to the motivating continuous dynamical model; that is, as $\Delta t \rightarrow 0$, our model preserves the dynamical properties of the differential equations. Additionally, in the presence

of stochasticity, as $\Delta t \rightarrow 0$, our model converges to an underlying stochastic differential process. The incorporation of RK4 for continuous model approximation is useful for emphasizing parameter estimation and ensuring that process uncertainty is focused on model choice rather than model approximation.

It should be noted that, as Gause's *Paramecium* data are counts, a Poisson data model would also be reasonable. We have found that the implementation of such a model (not presented here) yields similar results in terms of the posterior process and parameters. Although the Poisson specification

addresses the discrete nature of the data in this case, it imposes a distinct mean–variance relationship. An overdispersed Poisson model could also be used, though we focus on the more general truncated normal model, which also implies a mean–variance relationship, yet it is more flexible than the Poisson due to its two-parameter model formulation. Furthermore, the model we use here is also more robust in that it allows for linear transformations of the data (e.g., scaling) as well as various types of data; for example, the proposed model is applicable to data having bounded continuous support (e.g., rainfall data; percentage quadrat cover; plant basal area). Application to such examples is the subject of ongoing research.

5. Supplementary Materials

Web Appendices referenced in Sections 2.4 and 3 are available under the Paper Information link at the *Biometrics* website <http://www.biometrics.tibs.org>.

ACKNOWLEDGEMENTS

We thank Chris Wikle and Peter Adler for many helpful comments and suggestions. We also thank Dave Brown for technical assistance, Brynja Kohler and Brian Yurk for mathematical insights and resources, and the reviewers of this manuscript for their suggestions. This research was supported by Utah State University and the USU College of Science.

REFERENCES

- Berliner, L. M. (1996). Hierarchical Bayesian time-series models. In *Maximum Entropy and Bayesian Methods*, 15–22. Dordrecht, The Netherlands: Kluwer Academic Publishers.
- Buckland, S. T., Newman, K. B., Fernandez, C., Thomas, L., and Harwood, J. (2007). Embedding population dynamics models in inference. *Statistical Science* **22**, 44–58.
- Burden, R. L. and Faires, J. D. (2001). *Numerical Analysis*. Pacific Grove, California: Wadsworth Group.
- Caswell, H. (2001). *Matrix Population Models: Construction, Analysis, and Interpretation*. Sunderland, Massachusetts: Sinauer Associates.
- Dennis, B., Ponciano, J. M., Lele, S. R., Taper, M. L., and Staples, D. F. (2006). Estimating density dependence, process noise, and observation error. *Ecological Monographs* **76**, 323–341.
- Edelstein-Keshet, L. (1988). *Mathematical Models in Biology*. New York: McGraw-Hill, Inc.
- Gause, G. F. (1934). *The Struggle for Existence*. Baltimore, Maryland: The Williams and Wilkins Company.
- Gelman, A. (2006). Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper). *Bayesian Analysis* **1**, 515–534.
- Gianni, G., Pasquali, S., and Ruggeri, F. (2008). Bayesian inference for functional response in a stochastic predator-prey system. *Bulletin for Mathematical Biology* **70**, 358–381.
- Hooten, M. B. and Wikle, C. K. (2007). A hierarchical Bayesian nonlinear spatio-temporal model for the spread of invasive species with application to the Eurasian Collared-Dove. *Environmental and Ecological Statistics* **15**, 59–70.
- Horrace, W. C. (2005). Some results on the multivariate truncated normal distribution. *Journal of Multivariate Analysis* **94**, 209–221.
- Johnson, D. S., London J. M., Lea, M., and Durban, J. W. (2008). Continuous-time correlated random walk model for animal telemetry data. *Ecology* **89**, 1208–1215.
- Johnson, N. L., Kotz, S., and Balakrishnan, N. (1994). *Continuous Univariate Distributions*. New York: John Wiley and Sons.
- Lele, S. R., Dennis, B., and Lutscher, F. (2007). Data cloning: Easy maximum likelihood estimation for complex ecological models using Bayesian Markov chain Monte Carlo methods. *Ecology Letters* **10**, 551–563.
- Leslie, P. H. (1957). An analysis of the data for some experiments carried out by Gause with populations of the protozoa, *Paramecium aurelia* and *Paramecium caudatum*. *Biometrika* **44**, 314–327.
- Pascual, M. A. and Kareiva, P. (1996). Predicting the outcome of competition using experimental data: Maximum likelihood and Bayesian approaches. *Ecology* **77**, 337–349.
- R Core Development Team. (2007). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing, ISBN: 3-900051-07-0. Available at: <http://www.R-project.org>.
- Rohatgi, V. K. (1976). *An Introduction to Probability Theory and Mathematical Statistics*. New York: John Wiley and Sons.
- Rumelin, W. (1982). Numerical treatment of stochastic differential equations. *SIAM Journal on Numerical Analysis* **19**, 604–613.
- Scipione, C. M. and Berliner, L. M. (1993). Bayesian inference in nonlinear dynamical systems. In *Proceedings of the Section on Bayesian Statistical Sciences*. Washington, District of Columbia: American Statistical Association.
- Shoji, I. and Ozaki, T. (1998). A statistical method of estimation and simulation for systems of stochastic differential equations. *Biometrika* **85**, 240–243.
- Stein, M. L. (1992). Prediction and inference for truncated spatial data. *American Statistical Association, Institute of Mathematical Statistics, and Interface Foundation of North America* **1**, 91–110.
- Turchin, P. (2003). *Complex Population Dynamics: A Theoretical/Empirical Synthesis*. Princeton, New Jersey: Princeton University Press.
- Wikle, C. K. (2003). Hierarchical Bayesian models for predicting the spread of ecological processes. *Ecology* **84**, 1382–1394.

Received January 2008. Revised June 2008.

Accepted June 2008.