

APPENDIX: GIBBS SAMPLER FOR BAYESIAN STATE-SPACE MODELS

Exponential model

The Gibbs sampler is described for ecological examples by Clark et al. (2003, 2004) and Wikle (2003). It involves alternately sampling from each of the conditional posterior densities that are obtained by factoring the joint posterior. Here we outline the basic elements of the Gibbs sampler for the exponential model. Models are implemented in the language *R*. Algebra is reviewed in Clark (2004).

States x_t —The conditional posterior for each of the unobserved log-abundances, x_t , is the product of three normal distributions (eqn 5), which is solved by “completing the square”,

$$\begin{aligned} p(x_t | x_{t-1}, x_{t+1}, y_t, b, \sigma^2, \tau^2) &\propto N(x_t | x_{t-1} + b, \sigma^2) N(x_{t+1} | x_t + b, \sigma^2) N(y_t | x_t, \tau^2) \\ &= N(x_t | Vv_t, V), \end{aligned} \tag{A.1}$$

The resulting normal distribution has mean Vv_t and variance V , given by

$$\begin{aligned} Vv_t &= V \left[\frac{(x_{t-1} + x_{t+1})}{\sigma^2} + \frac{y_t}{\tau^2} \right] \\ V^{-1} &= \frac{2}{\sigma^2} + \frac{1}{\tau^2} \end{aligned} \tag{A.2}$$

The variance is the harmonic mean of variances from the three arrows (A, B, and C) in Figure 1. The mean is the weighted average of contributions from the three arrows, the weights being the inverses of the three variances. Because the first census is conditioned on the ‘prior’ x_0 , it is treated like the other censuses. Had we conditioned instead on the first census, it would conditionally depend only on arrows B and C, in which case we have,

$$V_{V_1} = V \left[\frac{x_2}{\sigma^2} + \frac{y_1}{\tau^2} \right]$$

$$V^{-1} = \frac{1}{\sigma^2} + \frac{1}{\tau^2}$$

For the last census, we have (from arrows A and C) the mean value

$$V_{V_T} = V \left[\frac{x_{T-1}}{\sigma^2} + \frac{y_T}{\tau^2} \right]$$

Process error σ^2 —The conditional posterior for the process error is the product of the normal likelihood and the inverse gamma prior to give the conditional posterior

$$\sigma^2 | \mathbf{x}, b \sim IG \left(\alpha_\sigma + \frac{T}{2}, \beta_\sigma + \frac{1}{2} \sum_{t=1}^T (x_t - x_{t-1} - b)^2 \right) \quad \text{A.3}$$

Observation error τ^2 —If the same distribution applies to all censuses, then the conditional posterior is

$$\tau^2 | \mathbf{x}, \mathbf{y} \sim IG \left(\alpha_\tau + \frac{T}{2}, \beta_\tau + \frac{1}{2} \sum_{t=1}^T (y_t - x_t)^2 \right) \quad \text{A.4}$$

Rate parameter b —For a normal prior, the conditional posterior for the parameter b is also normal. It is proportional to the product

$$\begin{aligned} p(b | \mathbf{x}, \sigma^2) &\propto \prod_{t=1}^T N(x_t | x_{t-1} + b, \sigma^2) N(b | b_0, V_b) \\ &= N(b | V_V, V) \end{aligned} \quad \text{A.5}$$

with mean and variance

$$V_V = V \left(\frac{\sum_{t=1}^T (x_t - x_{t-1})}{\sigma^2} + \frac{b_0}{V_b} \right)$$

$$V^{-1} = \frac{T}{\sigma^2} + \frac{1}{V_b} \quad \text{A.6}$$

For a noninformative prior (large V_b), the posterior conditional mean tends to the mean change in log density per time increment.

Variable observation errors--Variable observation errors mean that each must be sampled from a different distribution. For the conditional posterior on x_t , we replace τ^2 in eqn A.4 with the value for that specific observation, τ_t^2 . The conditional posterior for each variance is then

$$\tau_t^2 | x_t, y_t \sim IG\left(\alpha_t + \frac{1}{2}, \beta_t + \frac{1}{2}(y_t - x_t)^2\right). \quad \text{A.7}$$

These conditionals replace eqn A.4.

Uneven sample intervals—From eqn 9, we have the conditional density for the state at time t

$$\begin{aligned} p(x_t | x_{t-\delta_t}, x_{t+\delta_{t+1}}, y_t, b, \sigma^2, \tau^2) &\propto N(x_t | x_{t-\delta_t} + \delta_t b, \delta_t^2 \sigma^2) N(x_{t+\delta_{t+1}} | x_t + \delta_{t+1} b, \delta_{t+1}^2 \sigma^2) N(y_t | x_t, \tau_t^2) \\ &= N(x_t | V_t v_t, V_t) \end{aligned}$$

This replaces eqn A.1. Completing the square, we obtain mean and variance

$$V_t v_t = V_t \left(\frac{x_{t-\delta_t} + \delta_t b}{\delta_t^2 \sigma^2} + \frac{x_{t+\delta_{t+1}} - \delta_{t+1} b}{\delta_{t+1}^2 \sigma^2} + \frac{y_t}{\tau_t^2} \right)$$

$$V_t^{-1} = \frac{1}{\tau_t^2} + \frac{1}{\delta_t^2 \sigma^2} + \frac{1}{\delta_{t+1}^2 \sigma^2}$$

The process error is sampled from

$$\sigma^2 | \mathbf{x}, b \sim IG\left(\alpha_\sigma + \frac{T}{2}, \beta_\sigma + \frac{1}{2} \sum_{t=1}^T \left(\frac{x_t - f(x_{t-1})}{\delta_t} \right)^2\right) \quad \text{A.8}$$

For the growth rate parameter, the conditional posterior is

$$\begin{aligned}
b|\mathbf{x}, \mathbf{y}, \sigma^2 &\propto \prod_{t=1}^T N(x_t | x_{t-\delta_t} + \delta_t b, \delta_t^2 \sigma^2) N(b | b_0, V_0) \\
&= N(b | V_V, V)
\end{aligned}$$

having mean V_V and variance V

$$V_V = V \left(\frac{1}{\sigma^2} \sum_{t=2}^T (x_t - x_{t-\delta_t}) / \delta_t + \frac{b_0}{V_0} \right)$$

$$V^{-1} = \frac{T}{\sigma^2} + \frac{1}{V_0}$$

Missing values-- For missing values, the conditional posterior involves only arrows A and B in Figure 3,

$$V_{V_t} = V \left[\frac{x_{t-1} + x_{t+1}}{\sigma^2} \right] = \frac{x_{t-1} + x_{t+1}}{2}$$

$$V = \sigma^2 / 2$$

So we have two ways of drawing estimates of x_t , depending on whether or not there is an observation available at census t .

Nonlinear models

For the nonlinear model with unequal sample intervals (eqn 10), the conditional posterior for x_t is proportional to

$$p(x_t | x_{j \neq t}, \mathbf{b}, \sigma^2, \tau^2) \propto N(x_t | f(x_{t-\delta_t}; \mathbf{b}), \delta_t^2 \sigma^2) N(x_{t+\delta_t} | f(x_t; \mathbf{b}), \delta_{t+1}^2 \sigma^2) N(y_t | x_t, \tau^2)$$

The product of A and C can be integrated, being quadratic in x_t , with mean and variance

$$V_t v_t = V_t \left(\frac{f(x_{t-\delta_t})}{\delta_t^2 \sigma^2} + \frac{y_t}{\tau^2} \right)$$

$$V_t^{-1} = \frac{1}{\delta_t^2 \sigma^2} + \frac{1}{\tau^2}$$

The target distribution can now be written as the product of two normal distributions, one of which (B) is nonlinear in x_t ,

$$p(x_t | x_{j \neq t}, \mathbf{b}, \sigma^2, \tau^2) = \overbrace{N(x_t | V_{\mathbf{V}_t}, V)}^{A \times C} \overbrace{N(x_{t+\delta_t} | f(x_t; \mathbf{b}), \delta_{t+1}^2 \sigma^2)}^B$$

To sample from this target distribution we must embed within the Gibbs sampler either a rejection (Carlin et al. 1992) or a Metropolis (Calder et al. 2003) step. Methods for both are described in Clark (2004). The conditionals for the variances are

$$\sigma^2 | \mathbf{x}, \mathbf{y}, \mathbf{b}, x_0 \sim IG \left(\alpha_\sigma + \frac{T}{2}, \beta_\sigma + \frac{1}{2} \sum_{t=1}^T \left(\frac{x_t - f(x_{t-\delta_t})}{\delta_t} \right)^2 \right)$$

$$\tau^2 | \mathbf{x}, \mathbf{y} \sim IG \left(\alpha_\tau + \frac{T}{2}, \beta_\tau + \frac{1}{2} \sum_{t=1}^T (y_t - x_t)^2 \right)$$

There are two regression parameters. Using conjugate normal priors, we can sample directly. For the intercept, we have the normal mean and variance

$$V_{\mathbf{V}} = V \left(\frac{1}{\sigma^2} \sum_{t=1}^T \frac{x_t - x_{t-\delta_t} - \delta_t b_1 e^{x_{t-\delta_t}}}{\delta_t} + \frac{B_0}{V_0} \right)$$

$$V^{-1} = \frac{T}{\sigma^2} + \frac{1}{V_0}$$

For the slope, we have mean and variance

$$V_{\mathbf{V}} = V \left(\frac{1}{\sigma^2} \sum_{t=1}^T \frac{e^{x_{t-\delta_t}} (x_t - x_{t-\delta_t} - \delta_t b_0)}{\delta_t} + \frac{B_1}{V_1} \right)$$

$$V^{-1} = \frac{1}{\sigma^2} \sum_{t=1}^T e^{2x_{t-1}} + \frac{1}{V_1}$$

Lagged effects

From the model in eqn 11, assuming log-normal process error, we have the likelihood

$$N(\Delta|\mathbf{X}\mathbf{b}, \sigma^2\mathbf{I})N(\mathbf{y}|\mathbf{x}, \tau^2\mathbf{I}),$$

where \mathbf{I} is the rank T identity matrix. The conditional posterior for state x_t is the product of normal densities for each of $(x_{t-k}, \dots, x_t, \dots, x_{t+k})$. These can be sampled by straightforward extension of the approach discussed above. The parameter vector \mathbf{b} is sampled directly from the multivariate normal

$$N(\mathbf{V}\mathbf{v}, \mathbf{V})$$

where

$$\mathbf{V}\mathbf{v} = \mathbf{V}(\sigma^{-2}\mathbf{X}^T\Delta + \mathbf{V}_b^{-1}\mathbf{b}_0)$$

$$\mathbf{V}^{-1} = \sigma^{-2}\mathbf{X}^T\mathbf{X} + \mathbf{V}_b^{-1}$$

Variances are sampled as before (eqn A.8, with $\delta = 1$).

Model selection

Our application of Gelfand and Ghosh's (1998) predictive loss entails determination of expected values and second moments from MCMC followed by determination of a goodness of fit G_m and penalty term P_m . At each iteration g of the Gibbs sampler we predict states from currently imputed parameter values for model m ,

$E[x_t^{(g)}] = f_m(x_{t-1}^{(g)}; b^{(g)})$. The mean for each predicted state is taken over the converged sequence and used to calculate

$$G_m = \sum_{t=1}^T (y_t - E[x_t])^2$$

The second moment is calculated as the predicted sum of first moment squared and the

variance, $E\left[\left(x_t^{(g)}\right)^2\right] = \left[f_m\left(x_{t-1}^{(g)}; b^{(g)}\right)\right]^2 + \left(\sigma^2\right)^{(g)} + \left(\tau^2\right)^{(g)}$. The penalty term is

determined as

$$P_m = \sum_{t=1}^T \text{var}[x_t]$$

where the predictive variance is taken as the difference between second moment and the squared first moment. We choose the model with the lowest value of $D_m = G_m + P_m$.

TSIR model

The posterior conditionals required for Gibbs sampling are as follows:

Infected states I_t —State I_t conditionally depends on past and future states and on the observed individuals at time t ,

$$p\left(I_t \mid I_{t-1}, I_{t+1}, I_t^{(o)}, S_{t-1}, \beta_w, \rho, \dots\right) \propto \\ \text{Bin}\left(I_t \mid S_{t-1}, \varphi\left(I_{t-1}, S_{t-1}, \beta_w\right)\right) \text{Bin}\left(I_{t+1} \mid S_t, \varphi\left(I_t, S_t, \beta_{w+1}\right)\right) \text{Bin}\left(I_t^{(o)} \mid I_t, \rho\right)$$

This is sampled with a Metropolis step.

Transmission rates β_w — The β_w each conditionally depend only on the transitions that correspond to the 2-wk period w ,

$$p\left(\beta_w \mid I_{t \in w}, S_{(t-1) \in w}, \dots\right) \propto \prod_{t \in w} \text{Bin}\left(I_t \mid S_{t-1}, \varphi\left(I_{t-1}, S_{t-1}, \beta_w\right)\right) \text{Unif}\left(\beta_w \mid 0, 1000\right)$$

We use a Metropolis step, with the normal proposal density.

The Gibbs sampler begins by drawing 26 values for β_w followed by a single value for ρ , in both cases, based on the currently imputed values for all other parameters. Infected individuals are sampled chronologically, with susceptibles being updated according

to(12). We condition on an initial susceptible value of $S_0 = 120,000$. The last infected value is not constrained by a subsequent observation and, thus, is sampled from

$$Bin(I_t | S_{t-1}, \varphi(I_{t-1}, S_{t-1}, \beta_w)) Bin(I_t^{(o)} | I_t, \rho).$$