

Ecological Archives M074-010-A1

James S. Clark, Shannon LaDeau, and Ines Ibanez. 2004. Fecundity of trees and the colonization-competition hypothesis. *Ecological Monographs* 74:415-442.

Appendix A. MCMC algorithm: conditional posteriors and convergence.

The Gibbs sampler for our model is based on iterative sampling from the full conditional density for each parameter (Gelfand and Smith 1990, Gelman et al. 1995). Here we summarize our method.

Fecundity parameters

Regression parameters can be directly sampled based on distributions derived from theory of linear models. Parameters for fixed effects α conditionally depend on all tree fecundity schedules and on a bivariate normal prior,

$$p(\alpha|y, \mathbf{X}, \beta, \Sigma, \dots) \propto \prod_{i=1}^m N_{T_i}(y_i | \mathbf{X}_i \alpha + \mathbf{1} \beta_i, \Sigma_i) N_2(\alpha | \mathbf{c}, \mathbf{V}_a)$$

$$= N(\alpha | \mathbf{V} \mathbf{v}, \mathbf{V})$$

where $\mathbf{V}^{-1} = \sum_{i=1}^m \mathbf{X}_i^T \Sigma_i^{-1} \mathbf{X}_i + \mathbf{V}_a^{-1}$ and $\mathbf{v} = \sum_{i=1}^m \mathbf{X}_i^T \Sigma_i^{-1} (y_i - \mathbf{1} \beta_i) + \mathbf{V}_a^{-1} \mathbf{c}$.

The posterior conditional for the β_i contains only the fecundity schedule for the i^{th} tree,

$$p(\beta_i | y, \mathbf{X}, \alpha, \Sigma, \dots) \propto N_{T_i}(y_i | \mathbf{X}_i \alpha + \mathbf{1} \beta_i, \Sigma_i) N_2(\beta_i | 0, \tau^2)$$

$$= N(\beta_i | \mathbf{V}\mathbf{v}, \mathbf{V})$$

where

$$\mathbf{V}^{-1} = \mathbf{1}^T \Sigma_i^{-1} \mathbf{1} + \tau^{-2} = \frac{2(1-\rho) + (T_i - 2)(1-\rho)^2}{\sigma^2} + \frac{1}{\tau^2}$$

and

$$\begin{aligned} \mathbf{v} &= \mathbf{1}^T \Sigma_i^{-1} (y_i - \mathbf{X}_i \boldsymbol{\alpha}) \\ &= \frac{1}{\sigma^2} \left[(1-\rho)(y_{i1} + y_{iT_i} - 2\alpha_0 - \alpha_1(d_{i1} + d_{iT_i})) + (1-\rho)^2 \sum_{t=2}^{T_i-1} (y_{it} - \alpha_0 - \alpha_1 d_{it}) \right]. \end{aligned}$$

The variance on hyperparameters has conditional posterior

$$\begin{aligned} p(\tau^2 | \dots) &\propto \prod_{i=1}^m N(\beta_i | 0, \tau^2) IG(\tau^2 | a_\tau, b_\tau) \\ &= IG\left(\tau^2 \middle| a_\tau + \frac{m}{2}, b_\tau + \frac{1}{2} \sum_{i=1}^m \beta_i^2\right). \end{aligned}$$

Prior parameter values were $\mathbf{c} = [2, 0.5]^T$, $v_{a_1} = 1000$, $v_{a_2} = 1.0$, and $a_\tau = b_\tau = 0.01$.

Seed production from the i^{th} tree conditionally depends on the likelihood for all seed traps in all years (because there is dispersal and autocorrelation) and on its' own fecundity schedule. For the i th tree we have the full posterior conditional

$$p(y_i | X_i, \boldsymbol{\alpha}, \beta_i, \Sigma_i, \dots) \propto \prod_{t=1}^{T_j} \prod_{j=1}^n \text{Pois}(s_{jt} | A_j g_j(\mathbf{y}_t)) N_{T_i}(y_i | \mathbf{X}_i \boldsymbol{\alpha} + \mathbf{1} \beta_i, \Sigma_i).$$

We use a Metropolis step based on the proposal density $N_{T_i}(y_i, 0.1\Sigma_i)$. Thus, we propose a full time series for each tree, and we accept or reject on the basis of the full series, not on specific years.

Total variance σ^2 was sampled directly from the posterior conditional

$$\begin{aligned} p(\sigma^2 | \dots) &\propto \prod_{i=1}^m N_{T_i}(y_i | \mathbf{X}_i \boldsymbol{\alpha} + \mathbf{1} \beta_i, \Sigma_i) \text{Gam}(\sigma^2 | a_\sigma, b_\sigma) \\ &= IG(\sigma^2 | s_1, s_2) \end{aligned}$$

where

$$s_1 = a_\sigma + \frac{1}{2} \sum_{i=1}^m T_i$$

$$s_2 = b_\sigma + \frac{1}{2} \sum_{i=1}^m (y_i - \mathbf{X}_i \boldsymbol{\alpha} - \mathbf{1} \beta_i)^T \mathbf{R}_i^{-1} (y_i - \mathbf{X}_i \boldsymbol{\alpha} - \mathbf{1} \beta_i)$$

$$\text{and } \mathbf{R}_i = \sigma^{-2} \Sigma_i.$$

Prior parameter values were $a_\sigma = b_\sigma = 0.01$.

For ρ we used the symmetric proposal density $Unif(\rho^{(g)} - h, \rho^{(g)} + h)$ and a Metropolis step, where $\rho^{(g)}$ is the current value of ρ , and $h = \min[0.2(1 - |\rho^{(g)}|), 0.05]$ to sample from the conditional posterior

$$p(\rho | \dots) \propto \prod_{i=1}^m N_{T_i}(y_i | \mathbf{X}_i \boldsymbol{\alpha} + \mathbf{1} \beta_i, \Sigma_i) Unif(\rho | -1, 1).$$

Unless near -1 or 1 , the proposal density has width 0.1 . Near either limit, it narrows.

Because correlations are always bounded away from limits, the algorithm does not bog

down there. As with σ^2 , the covariance matrix is assembled as basis for evaluating the conditional posterior.

Dispersal parameter

The dispersal parameter conditionally depends on all of the seed trap data,

$$p(u|\dots) \propto \prod_{t=1}^{T_j} \prod_{j=1}^n \text{Pois}(s_{jt} | A_j g_j(\mathbf{y}_t)) \text{Gam}(u | a_u, b_u).$$

We used a Metropolis-Hastings step with a gamma proposal density. Gamma distributions $\text{Gam}(a, b)$ have mean a/b . Prior parameter values were $a_u = 1$ and $b_u = 0.01$.

Tree-status parameters

Conditional posteriors for tree status parameters contain both full likelihoods, because statuses q_i directly influence all seed traps through expected seed rain (Eqs. 13 and 14).

For a_θ , the conditional is

$$p(a_\theta|\dots) \propto \prod_{t=1}^{T_j} \prod_{j=1}^n \text{Pois}(s_{jt} | A_j g_j(\mathbf{y}_t)) \left[\prod_{i=1}^m \text{Multinom}(q_i | 1, w_i) \right]^k \text{Gam}(a_\theta | a_0, b_0)$$

with that for b_θ having the appropriate prior. For both a_θ prior parameter values were $a_0 = 0.2$ and $b_0 = 0.1$, and b_θ prior parameter values were $a_0 = 0.1$ and $b_0 = 1$. For both we used Gamma proposals and Metropolis-Hastings. Likewise, the female fraction full conditional is

$$p(\phi|\dots) \propto \prod_{t=1}^{T_j} \prod_{j=1}^n \text{Pois}(s_{jt} | A_j g_j(\mathbf{y}_t)) \left[\prod_{i=1}^m \text{Multinom}(q_i | 1, w_i) \right]^k \text{Unif}(\phi | 0, 1).$$

Recognition error was

$$p(v|\dots) \propto \prod_{t=1}^{T_j} \prod_{j=1}^n \text{Pois}(s_{jt} | A_j g_j(\mathbf{y}_t)) \left[\prod_{i=1}^m \text{Multinom}(q_i | 1, w_i) \right]^k \text{Unif}(v | 0, 1).$$

The weighting parameter k was introduced as a means for manipulating the contribution of the two data sets. A range of values was explored. We ultimately used $k = 1$, i.e., both data sets weighted strictly in terms of their representation in the data sets.

Convergence

Despite the high dimensionality of the model, Markov chains converged for all but a few species. We were not satisfied that convergence was reached for several species having few seeds; these species were not analyzed further (see *Results* in text). To initialize seed production rates y_{it} , plot-by-plot maximum likelihood (ML) estimates of α_0 were obtained for each year under the simplifying assumption of $\theta = 1$, and $\alpha_0 = 0.5$. These estimates were used to initialize fecundity for each tree. A burn-in of 2000 to 20,000 iterations was sufficient for all species included in this analysis, with steps to convergence depending in large part on number of trees. Parameters that could be sampled directly converged to values near the target distribution rapidly, with some slow drift as the stickier y_{it} 's converged more slowly.

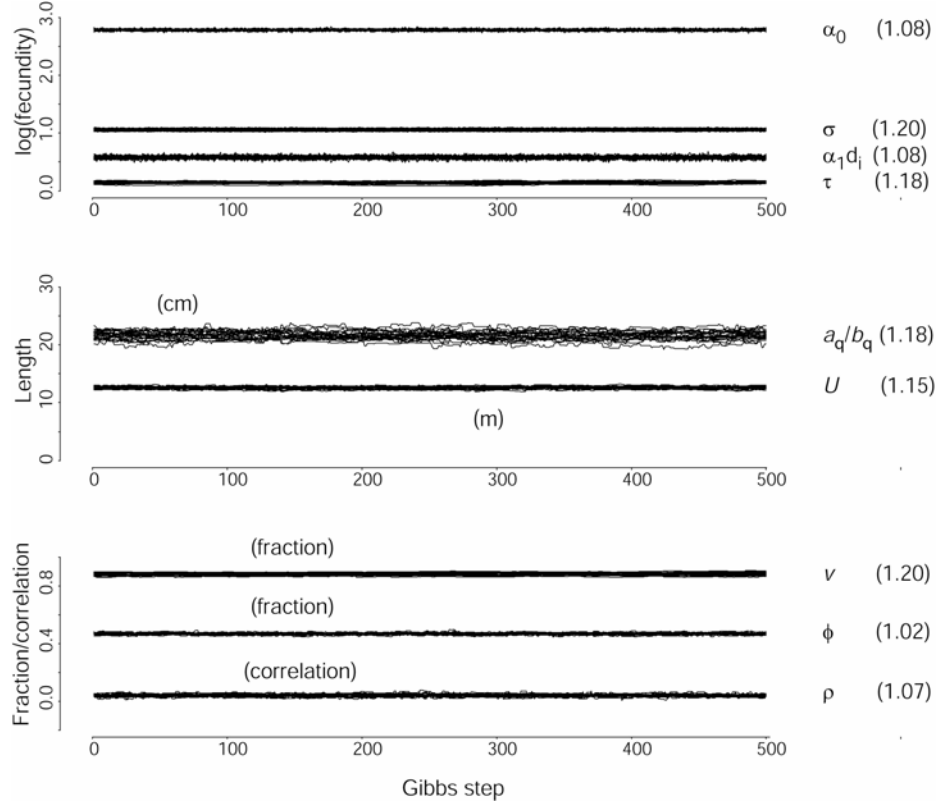


Fig. A1. Markov-chain Monte Carlo (MCMC) output for 20 parallel chains after discarding 20,000 burn-in iterations. Parameters are scaled to familiar or consistent units. Gelman and Rubin's (1992) potential scale reduction values are in parentheses at right. Posterior densities in the original scales are provided in the text figures. The four terms of log fecundity are shown in the upper panel, where the diameter used to scale α_1 is $d_{i1} = 1.46$ (or 28.6 cm), the diameter of the example tree in Fig. A2. Thus, the intercept term (α_0) dominates, followed by process variability (σ), diameter effect ($\alpha_1 d_{ii}$), and individual effects (τ). The middle panel shows the diameter at which the maturation schedule $\theta = 0.5$ (a_θ/b_θ) and the mean dispersal distance ($U \equiv \pi\sqrt{u}/2$). The lower panel shows the recognition error (ν), female fraction (ϕ), and AR(1) parameter (ρ).

An example from *Acer rubrum* demonstrates a “worst-case” scenario. When there are many trees in a stand, there are many potential combinations of fecundity values that might “satisfy” the annual trap data, and there is more “overlap” in seed shadows, making

it difficult to resolve parameters. Moreover, when interannual variability is large, there can be years of low seed recovery for which identification of source trees is especially challenging. *A. rubrum* is the most abundant species, and it is at highest density in stand C1 (720 trees/ha). After 20,000 burn-in iterations on parallel chains, all population level parameters are well identified (Fig. A1). Convergence occurred from multiple initial conditions, and potential scale reduction factors are at acceptable values (Gelman and Rubin 1992).

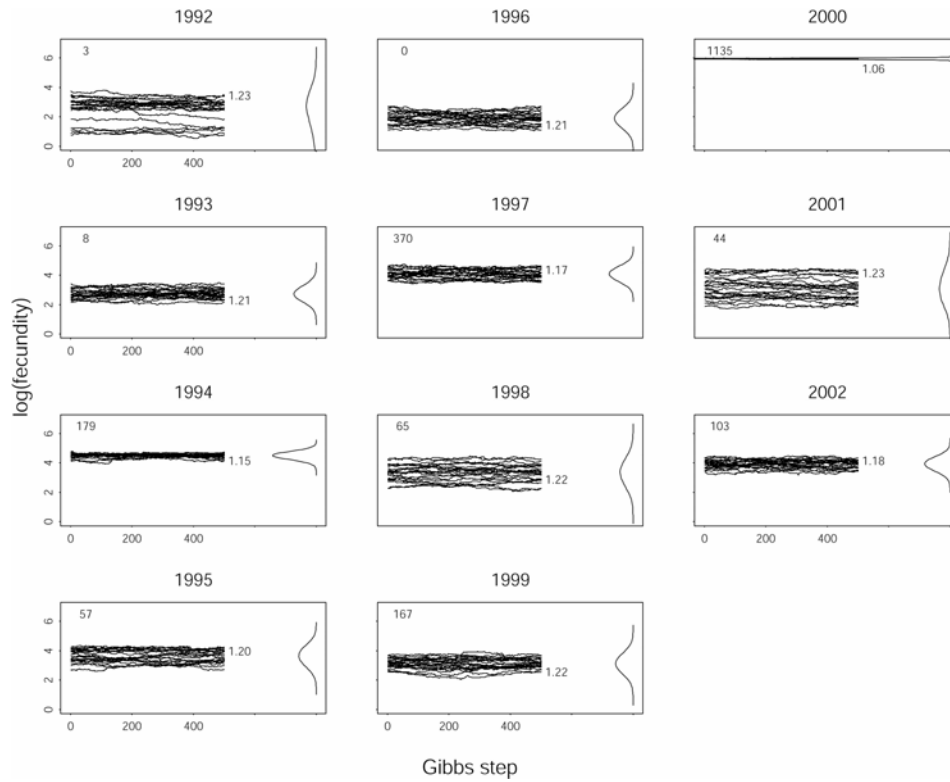


Fig. A2. Twenty parallel MCMC chains for an example *Acer rubrum* tree on plot C1 corresponding to those in Fig. A1. Numbers of seeds collected for each year are indicated in the upper left of each panel. Potential scale-reduction factors and smoothed posterior densities are at right. For years of low seed recovery (< 100 seeds), estimated fecundities are low ($< 10^4$ seeds) and not well identified.

An example of fecundity estimates for a tree from plot C1 show how parameters are well identified for years with adequate seed recovery (say, more than 100 seeds), when trees produce more than 1000 seeds (1994, 1995, 1997, 2000, and 2002 in Fig. A2). In years of low seed recovery, fecundity estimates are low, and poorly identified on the log scale plots of Fig. A2 (1992, 1993, 1996, 1998, 1999, 2001). The years of low fecundity estimates do not contribute substantially to overall fecundity, as is evident from the linear scale on the lower panel of Fig. A3. As this is an example of the most challenging species on the most challenging plot, most simulations converged much more rapidly.




Fig. A3. Annual estimates of fecundity from the example tree in Fig. A2 with 95% CIs. Poorly resolved estimates come from years of especially low fecundity.

Extraction of species effects

The extraction of species effects is summarized here. From the Gibbs output for a data set containing all trees of a genus, we extracted the chain of β_i and y_i estimates for each individual of a given species p . For the i^{th} tree, the estimate of the full (fixed plus random) effect at Gibbs step g is $a_0^{(g)} + \beta_i^{(g)}$. The fixed effect for species p is the mean of random effects taken over all individuals of that species,

$$a_{0,p}^{(g)} = a_0^{(g)} + \frac{1}{m_p} \sum_{i \in P} \beta_i^{(g)} = a_0^{(g)} + \overline{\beta_p^{(g)}}$$

where P is set of m_p individuals belonging to species p . The individual effect for tree i of species p is centered on the species p population effect

$$\beta_{i,p}^{(g)} = a_0^{(g)} + \beta_i^{(g)} - a_{0,p}^{(g)} = \beta_i^{(g)} - \overline{\beta_p^{(g)}}.$$

The g th estimate of τ^2 is

$$(\tau_p^2)^{(g)} = \frac{1}{m_p} \sum_{i \in P} (\beta_i^2)^{(g)}.$$

Fecundity schedules y_i were simply extracted from trees belonging to species p .

Literature cited

Gelfand, A. E., and A. F. M. Smith. 1990. Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association* **85**:398-409.

Gelman, A., J. B. Carlin, H. S. Stern, and D. B. Rubin. 1995. *Bayesian data analysis*. Chapman and Hall, London, UK.

Gelman, A., and D. B. Rubin. 1992. Inference from iterative simulation using multiple sequences. *Statistical Science* **7**:457-511.

