APPENDIX- A GIBBS SAMPLER FOR THE HIERARCHICAL MODEL

A full description of Bayesian computation is beyond the scope of this paper, but is the subject of several recent overviews (Gelman et al. 1995, Carlin and Louis 2000). Markov chain Monte Carlo (MCMC) methods simulate a posterior distribution by executing a random walk (Markov process) having a stationary density that is equivalent to the target (joint posterior) distribution. For example, the joint posterior for fecundity parameters is

$$p(\mathbf{b},\alpha_b,\beta_b|\mathbf{y}) = \prod_{i=1}^{n} f(y_i|\mathbf{b}_i)\prod_{i=1}^{n}\pi(\mathbf{b}_i|\alpha_b,\beta_b)u(\alpha_b)v(\beta_b). \quad \text{A.1}$$

This joint posterior contains parameters b_i, $i= 1, \ldots, n$, α_b, and β_b and is determined by the likelihood $f(\ )$, parameter density $\pi(\ )$, and priors $u(\ )$ and $v(\ )$.

The Gibbs sampler (Gelfand and Smith 1990) is a MCMC algorithm that makes use of the fact that a joint distribution of Q random variables is uniquely determined by the Q conditional distributions. The algorithm proceeds by alternately sampling from each of the full conditional distributions. One iteration of the algorithm consists of sequential sampling from each conditional posterior, updating each parameter value before proceeding to the next. The Gibbs sampler is described here.

Fecundity

The three levels for the fecundity model in Eq. A.1 include the Poisson likelihood, $f(\mathbf{y}|\mathbf{b}) = \prod_{i=1}^{n} Pois(y_i|,b_i)$, a gamma distribution ('prior') for the Poisson parameters $\pi(b_i|\alpha_b,\beta_b) = Gam(b_i|\alpha_b,\beta_b)$, and hyperpriors $u(\alpha_b) = Exp(\alpha_b|\rho_0)$ and $v(\beta_b) = Gam(\beta_b|c_0,d_0)$. From Eq. A.1 we have the joint posterior

$$p(\mathbf{b},\alpha_b,\beta_b|\mathbf{y}) \propto \prod_{i=1}^{n} Pois(y_i|b_i)\prod_{i=1}^{n} Gam(b_i|\alpha_b,\beta_b)Exp(\alpha_b|\rho_0)Gam(\beta_b|c_0,d_0)$$

For the parameters b_j $j = 1, 2, \ldots, n$, α_b, and β_b, the i^{th} step of the algorithm is as follows:

1. The Gibbs step begins with draws from the conditional posteriors for fecundity parameters

   $b_i | \alpha_b, \beta_b, y_j \sim Gam(\alpha_b + y_i, \beta_b + n_i)$.

2. The gamma hyperprior on β_b allows for direct sampling from the conjugate conditional posterior $\beta_b | \mathbf{b}, \alpha_b \sim Gam\left(c_0 n + \alpha_b, d_0 + \sum_{i=1}^{n} b_i\right)$.

3. A Metropolis-Hastings step was used for α_b in a manner similar to that described for survival (see below). A Gamma proposal distribution was used to generate candidate values α^* with acceptance probabilities determined from the conditional posterior $p(\alpha^* | \mathbf{b}, \beta) \propto \prod_{i=1}^{n} Gam(b_i | \alpha^*, \beta) * Exp(\alpha^* | \rho_0)$.

Prior parameter estimates used in figures were taken to be noninformative (Table 3).

Survival

Assuming a binomial likelihood with beta distributed probability s and gamma distributed priors we have three stages:

$$f(\mathbf{y}|\mathbf{s}) = \prod_{i=1}^{n} Bin(y_i | n_i, s_i) \qquad \text{(likelihood)}$$

$$\pi(s_i | \alpha_s, \beta_s) = Beta(s_i | \alpha_s, \beta_s) \qquad \text{(prior)}$$

$$u(\alpha_s) = Gam(\alpha_s | a_0, b_0) \text{ and } v(\beta_s) = Gam(\beta_s | a_0, b_0) \qquad \text{(hyperpriors)}$$

For individual variation, we have $Bernoulli(y|s) \equiv Bin(y|1, s)$ and indicator variable y, assuming values of zero or 1. From Eq. A.1, the joint posterior is

$$p(\mathbf{s}, \alpha_s, \beta_s | \mathbf{y}) \propto \prod_{i=1}^{n} Bin(y_i | n_i s_i) \prod_{i=1}^{n} Beta(s_i | \alpha_s, \beta_s) Gam(\alpha_s | a_0, b_0) Gam(\beta_s | a_0, b_0).$$

1. The s_i are each drawn from the conditional beta posteriors

   $$s_i | y_i, \alpha_s, \beta_s \sim Beta(\alpha_s + y_i, \beta_s + n_i - y_i).$$

2. The hyperparameters are not conjugate. I used a Metropolis-Hastings step.

   Sampling was accomplished in the same way for both survival parameters. The conditional posterior for α_s is $p(\alpha_s | \mathbf{s}, \mathbf{y}, \beta_s) \propto \prod_{i=1}^{n} Beta(s_i | \alpha_s, \beta_s) Gam(\alpha_s | a_0, b_0)$.

   A Gamma proposal distribution with a shape parameter of 5 and mean equal to the current parameter value α_j was used to draw candidate values $J(\alpha^* | \alpha_j) = Gamma(\alpha^* | 5, 5/\alpha_j)$. The importance ratio

   $$r = \frac{p(\alpha^* | \mathbf{y}) \big/ J(\alpha^* | \alpha_j)}{p(\alpha_i | \mathbf{y}) \big/ J(\alpha_j | \alpha^*)}$$

determines acceptance of the candidate value with probability min(r,1). The Gibbs sampler was run for 5000 iterations. The first 1000 iterations were discarded. Posterior quantiles and kernel estimates were determined on the remainder.

Literature cited

Carlin, B. P., and T. A. Louis. 2000. Bayes and empirical Bayes methods for data analysis. Chapman and Hall, Boca Raton, Florida, USA.

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 Hall, London, UK.