Appendix S1: Details of model fitting

The Bayesian approach used here for model fitting involves simulating from the distribution pr($\theta \mid M$), the probability of the vector of model parameters θ conditional on the data. For the models in this paper, θ is the vector of all the regression coefficients (α_k and β_k for Eq. 1; γ_k and δ_k for Eq. 2), and M is the vector of all the species deaths, m_{jk}. Often the analytical form of the conditional distribution is intractable, and Markov chain Monte Carlo (see Gelman et al. 2004, Carlin et al. 2006) is a method from drawing simulations of θ from this conditional distribution when the density function is known only up to a constant of proportionality. For the model in Eq. 1, the distribution function of interest is

$$P(\beta_k, \alpha_k \mid m_{jk}; N_{jk}, g_{jk}) \propto \prod_{j=1}^{n}\prod_{k=1}^{4} \text{Binomial}(m_{jk} \mid \beta_k, \alpha_k; N_{jk}, g_{jk}) \prod_{k=1}^{4} P(\beta_k)\, P(\alpha_k) \quad \text{(S3)}$$

where j indexes species, and k indexes soil types (note that not all species occur on all soil types). m_{jk} is the number of trees observed dead by 1997, N_{jk} is the number of individual trees observed alive in 1992, and g_{jk} is the observed mean growth rate for species j on soil k for all species with $n \geq 10$ on a soil. To the right of the proportionality symbol are the likelihood or distributional model for the observed data and the prior probabilities for model parameters. Uninformative prior probability densities for β_k and α_k were chosen to allow the data to dominate inference. The normal distribution was parameterized based on precision, where precision is the inverse of the variance:

$P(\beta_k) \sim \text{Normal}(0, 1e^6)$

$P(\alpha_k) \sim \text{Normal}(0, 1e^6)$

For models in Eq. 2, the joint posterior probability distribution of parameters is the same, except: (1) the variable name changes as noted in the "Methods" section in the main text, (2) the mortality terms (q_{jl}) were estimated based on a species' mortality rate on sandy loam; and (3) the g_{jl} were estimated based on that species' growth rate on one of the other three richer soils.

Markov Chain Monte Carlo (MCMC) simulation (as executed in WinBugs v.1.4.1; Spiegelhalter et al. 2003) was used to simulate a large number of random samples from the marginal posterior distributions of parameters. WinBugs adaptively selects the appropriate sampling algorithm based on the model and can use a combination of different algorithms, including the Gibbs sampler and Metropolis algorithm. The Gibbs sampler generates samples from the marginal probability distribution of each parameter, conditional on the current values of all other parameters in the model (Gelman et al. 2004). These random samples were used to make inferences on parameters of interest in the process model.

For each Markov chain, we ran 50,000 iterations and discarded the early iterations ("burn-in," generally the first 20,000 iterations). We thinned post-burn-in chains by retaining only every 30th sample, which increases the independence of samples in the Markov chain (Gelman et al. 2004). We ran three chains to facilitate monitoring of convergence of chains to the target distribution based on (1) visual inspection of plots of the parameter space that each chain was sampling at each iteration and (2) within- versus between-chain variance components (the potential scale reduction factor for each parameter, $\hat{R}$; Gelman et al. 2004). It is advised that $\hat{R}$ should not exceed 1.1 (Gelman et al. 2004). No $\hat{R}$ for any parameter in our models exceeded this value, and based on visual inspection no chains showed any evidence suggesting lack of convergence. Inference was based on at least 2000 post-burn-in, thinned samples.

Convergence diagnostics were assessed using the R contributed package, coda (Plummer et al. 2005) and standard output from WinBugs v.1.4.1.

Literature cited

Carlin, B.P., Clark, J.S. & Gelfand, A.E. 2006. Elements of hierarchical Bayesian inference. Pages 3-24 *in* A. Gelfand and J. S. Clark, editors. Hierarchical Modelling for the Environmental Sciences. Oxford University Press, Oxford.

Gelman, A., Carlin, J.B., Stern, H.S. & Rubin, D.B. 2004. Bayesian Data Analysis, 2nd edition. Chapman and Hall/CRC, New York.

Plummer, M., Best, N.G., Cowles, M.K. & Vines, S.K. 2005. coda: Output analysis and diagnostics for MCMC. R package version 0.9-5.