

## COEXISTENCE: HOW TO IDENTIFY TROPHIC TRADE-OFFS

JAMES S. CLARK,<sup>1,2</sup> JACQUELINE MOHAN,<sup>2</sup> MICHAEL DIETZE,<sup>2</sup> AND INES IBANEZ<sup>2</sup>

<sup>1</sup>*Department of Biology and Nicholas School of the Environment, Duke University, Durham, North Carolina 27708 USA*

<sup>2</sup>*University Program in Ecology, Duke University, Durham, North Carolina 27708 USA*

**Abstract.** Analyses of growth response to resource availability are the basis for interpreting whether trophic trade-offs contribute to diversity. If different species respond most to resources that are limiting at different times, then those differences may trade off with other trophic or life-history traits that, together, help to maintain diversity. The statistical models used to infer trophic differences do not accommodate uncertainty in resources and variability in how individuals use resources. We provide hierarchical models for resource–growth responses that accommodate stochasticity in parameters and in data, despite the fact that causes are typically unknown. A complex joint posterior distribution taken over  $>10^2$  parameters is readily integrated to provide a comprehensive accounting of uncertainty in the growth response, together with a small number of hyperparameters that summarize the population response. An application involving seedling growth response to light availability shows that large trophic differences among species suggested by traditional models can be an artifact of the assumption that all individuals respond identically. The hierarchical analysis indicates broad trophic overlap, with the implication that slow dynamics play a more important role in preserving diversity than is widely believed.

**Key words:** *coexistence; competition; diversity; growth response; hierarchical Bayes; light; resource–consumer interactions; species interactions; tree seedlings.*

### INTRODUCTION

Theory, experiment, and observation suggest that trade-offs among species involving trophic interactions or life history are the logical and parsimonious explanation for the often rich diversity of plant communities (Tilman 1988, 1994, Pacala et al. 1996, Rees et al. 2001). Trophic trade-offs result from patterns of consumption. Examples include different minimal requirements for different resources (Tilman 1982), light response curves that shift the advantage from early- to late-successional species as canopies close (Bazzaz 1979), and natural enemies that can promote diversity through preferential predation on a superior competitor (Paine 1966, Pacala and Crawley 1992). Life-history trade-offs involve timing of reproductive effort, and life-history and trophic trade-offs often interact through allocation that can affect growth, seed size, fecundity, dispersal, and survivorship (Loehle 1988, Tilman 1988, Clark 1991, Rees et al. 2001).

The apparent agreement of evidence from disparate approaches is compelling. Theory emphasizes trade-

offs, because models predict the extinction of species lacking parameter combinations that are sometimes favored in competition or in potentially rare or transient environments (MacArthur 1972, May 1973, Tilman 1994). Field data can be found to support this view (reviews of Connell and Slatyer 1977, Rees et al. 2001). The alternative, that species are not importantly different and densities therefore “drift” (Hubbell 2001), does not see much coverage in recent reviews of the subject. To some, the hypothesis is made more palatable by the possibility that speciation might offset the inevitable extinction losses.

Here we demonstrate that the motivation for trade-offs is less compelling and finds less support in data than is generally appreciated. The theoretical demand for trade-offs and the empirical support are influenced by assumptions that concern the structure of variability and the degree of uncertainty. Theoretical models assume that differences among species overwhelm variability among individuals, so much so that individual differences can be ignored. Individual differences need not be genetic, and they need not be associated with measurable environmental variation; any individual variation violates the assumptions of most ecological theory and almost all classical statistical models used to test it.

For example, growth responses depend on how organisms integrate variation (Levins 1979). Ecologists aspire to tightly controlled experiments to insure that resource treatments match the consumer's "definition" and the assumptions of statistical models. Such experiments may be feasible at the ecophysiology level and, rarely, at the population level. Examples include those used in comparisons of "A/C<sub>i</sub>" curves (the effect of internal CO<sub>2</sub> concentration on the rate of carbon assimilation; Pearcy et al. 1981, Sharkey et al. 1986), "light response curves" (the effect of light intensity on the photosynthetic rate; Bazzaz 1979, Ellsworth 2000), and population growth in chemostats (e.g., Tilman 1982). In such cases, the resource supply, arguably, can be known with some precision. More commonly, even tightly controlled experiments are characterized by wide fluctuations. The impact of variability is poorly understood for plant responses to light (Pearcy and Yang 1998), soil moisture (Hinckley et al. 1978, Oren and Pataki 2001), nutrients (Chapin et al. 1990, Glimskar and Ericsson 1999), and CO<sub>2</sub> (Ellsworth 2000, Alistair et al. 2001). Consumer responses to hosts and prey depend on many interacting variables (Chan and Godfray 1993, Turchin et al. 1999, 2000, Krebs et al. 2001).

Ignoring sources of stochasticity can result in biased estimates and inaccurate and overconfident predictions. It can foster misleading conclusions that species differ more than they actually do. In view of the fact that many of the dominant sources of variation will typically be unknown (and unknowable), how can trophic relationships be analyzed in a way that leads to useful inference concerning their role in community dynamics?

Our approach provides insight into the types of stochasticity that are important for ecological inference and prediction, and the methods that can be used to accommodate them. We do not establish here whether trophic or life-history trade-offs explain diversity. Rather, we demonstrate how to determine differences among species. We focus on a single trophic interaction, seedling growth responses to light, but the problem and the solution are generalizable. We do not attempt tight control over variables that vary widely in nature, although the issues and methods for addressing them apply to both observational and experimental data. The hierarchical structure of our model admits uncertainties in resource availability and variability in how individuals respond to it. We then compare results with those obtained using a traditional frequentist framework. For our example, we compare *Acer rubrum* and *Liriodendron tulipifera*, two species that coexist. This simple example illustrates the approach and its impact on inference. We take up a full analysis of a large number of species in a separate study (J. Mohan and J. S. Clark, unpublished manuscript).

## THE DATA

We illustrate the hierarchical approach using seedling growth responses to resources. We consider soil moisture, but devote most attention to light, because it is commonly analyzed. Growth data derive from annual measurements of seedling heights on 1-m<sup>2</sup> plots from the Duke Forest, Orange County, North Carolina, USA. We include several data sets, including one for *Acer rubrum* seedlings growing beneath a closed canopy and larger data sets for *Acer rubrum* and *Liriodendron tulipifera* seedlings that are derived from a range of canopy conditions. Species co-occur on most plots used in the hierarchical analysis (parameter estimates are provided for each plot in the online Supplement). The growth of seedling *i* on plot *j* is defined as the increase in height,  $y_{ij}$ , averaged over two years (three censuses). There are  $n_j$  seedlings per plot,  $m$  total plots, and  $M = \sum_{j=1}^m n_j$  total seedlings.

At each location, we measured soil volumetric water content (as a percentage) in the upper 15 cm of soil using a time domain reflectometer, TDR (Tektronix 1502B, Tektronix, Beaverton, Oregon, USA). Two measurements per plot were obtained every two weeks from mid-May until early September. Average values (two replicates by eight measurement dates) are used.

To estimate understory light, we obtained hemispherical canopy photographs during uniform sky conditions (early morning/late afternoon) in midsummer at a height of 1.15 m above each seedling plot. Images were obtained on 400-speed color slide film using a Nikon FM2 camera with a Sigma 8 mm 180° fish-eye lens and leveling tripod. Scanned images were analyzed using HemiView Canopy Analysis Software (Version 2.1, Delta-T Devices, Cambridge, UK). Photo analysis involves a user-defined threshold intensity for each photo that determines whether pixels are classified as open (sky) or obscured (canopy). The Global Site Factor (GSF) represents the proportion of full sunlight penetrating the forest canopy. The GSF combines direct radiation, by calculating the annual solar track, and diffuse radiation, based on a uniform overcast sky model. It does not account for backscatter within the canopy.

## WHY TRADITIONAL METHODS CAN BE MISLEADING

A traditional analysis involves fitting a saturating function of growth to an estimate of resource variability. Functional forms typically have parameters that describe the asymptotic rate (e.g.,  $A_{\max}$  for photosynthesis), the minimal resource at which growth is possible ("light compensation point" and Tilman's [1982]  $R^*$  are physiological- and population-level examples, respectively), and a half saturation constant (Holling 1959). The Monod function is standard:

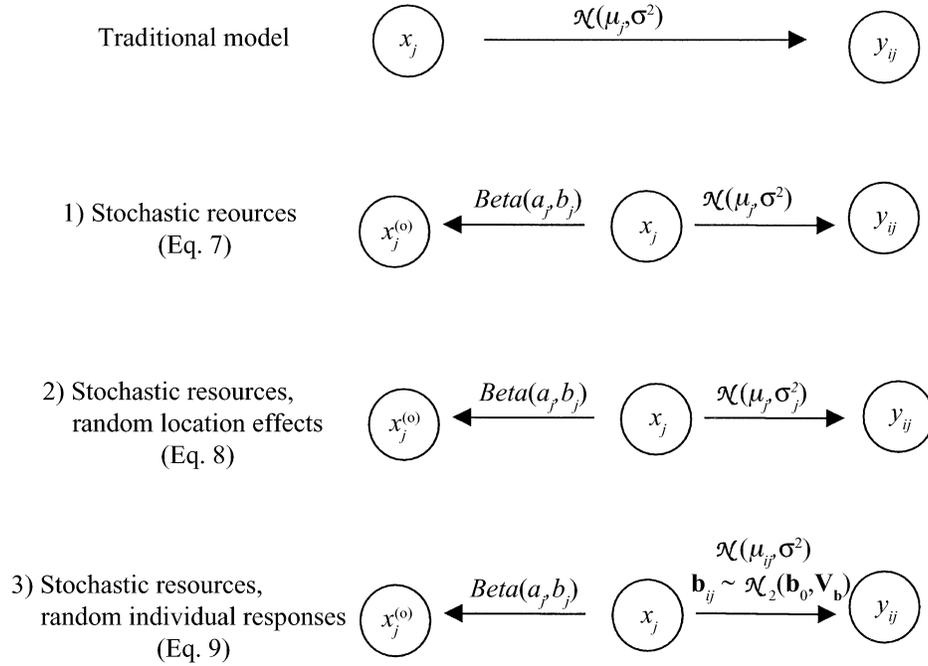


FIG. 1. The traditional model of uncertainty in a resource–growth analysis and three options (Methods 1–3) involving different types of stochasticity in growth.

$$\mu_j = \beta_0 + \beta_1 \left( \frac{x_j}{\theta + x_j} \right) \quad (1)$$

for resource level  $x_j$  at location  $j$ , having asymptote  $G = \beta_0 + \beta_1$ , minimum resource for positive growth  $x_0 = -\beta_0\theta/G$ , and “half saturation”  $\theta$ , the resource level at which growth is halfway between  $\beta_0$  and  $\beta_1$ . For resources that have an upper bound (e.g., the GSF has a maximum value of 1), parameter  $G$  is not a true asymptote. An alternative parameterization,

$$\mu_j = G \left( \frac{x_j - x_0}{\theta + x_j} \right) \quad (2)$$

is obtained by making the appropriate substitutions in Eq. 1. Both Eqs. 1 and 2 are useful. We fit parameters using Eq. 1 because we can write more efficient algorithms that take advantage of the linear parameters  $\beta_0$  and  $\beta_1$ . Parameters of Eq. 2 are more “biological,” so we interpret the reparameterized Eq. 2.

Here is a classical approach. Let  $y_{ij}$  be the  $i$ th measurement of growth at resource level  $x_j$ . A simple model says:

$$y_{ij} = \mu_j + \varepsilon_{ij} \quad (3)$$

The first term is the deterministic model. The only uncertainty we associate with it is that conferred by our degree of confidence in the estimates of the parameters  $\mathbf{b} = [\beta_0 \ \beta_1]^T$  and  $\theta$ . All variability is contained in the second term, which is often assumed to follow a normal distribution,  $\varepsilon_{ij} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2)$ ; “i.i.d.” indicates “independent and identically distributed. To estimate param-

eters, we write a likelihood for the data under the assumption that the response is  $\mu_j$  with error, described by a fitted parameter  $\sigma^2$ ,  $y_{ij} \sim \mathcal{N}(\mu_j, \sigma^2)$ , the “traditional model” (Fig. 1). The likelihood

$$p(\mathbf{y} | \mathbf{x}, \mathbf{b}, \theta, \sigma^2) = \prod_{j=1}^m \prod_{i=1}^{n_j} \mathcal{N}(y_{ij} | \mu_j, \sigma^2) \quad (4)$$

is the basis for maximum likelihood (ML) or Bayesian estimation, the latter requiring a prior for parameters.

#### Confidence (credible) intervals

Confidence intervals are used for inference on parameters and for prediction. Within the classical framework, we assume that there exists a true parameter value. Parameters do not have probability distributions; they are constants. An  $X\%$  frequentist confidence interval should contain this true parameter value in  $X\%$  of repeated trials. Error propagation translates this concept to the response variable; confidence intervals for the predicted response are expected to cover the true response in the specified fraction of repeated trials. Ecologists can view the “response” in several ways, depending on their interpretation of scatter.

A fitted model can be used to predict all or part of Eq. 3, and most analyses predict only the first part. Eq. 3 consists of a deterministic term  $\mu$ , which is usually associated with the process of interest. All stochasticity is contained in the second, “error” term, which is associated with things gone wrong (hence, the name). Of course, the process  $\mu$  also contains data  $x$ , but  $x$  is

assumed to have no estimation error and no stochasticity; once observed (or administered as a treatment),  $x$  is precisely known. Because we are interested in  $\mu$ , the only stochasticity that impacts inference involves estimation (parameter uncertainty). Thus, a confidence interval typically propagates estimation error in  $\mathbf{b}$  and  $\theta$  to the estimate of  $\mu$ , while ignoring the stochasticity associated with  $\varepsilon$ .

Because estimates are subject to asymptotics, so, too, are estimates of  $\mu$ ; the confidence interval on  $\mu$  becomes vanishingly small as the sample size increases. This is true whether the analysis is frequentist or Bayesian (Clark 2003). The disappearing confidence interval reflects the fact that it is based on fixed entities, namely parameters and  $x$ . Thus, an  $X\%$  confidence interval on  $\mu$  does not contain  $X\%$  of the data. A confidence interval calculated in this way can be viewed as information about parameters translated to information about  $\mu$ .

To construct a prediction envelope on data ( $y$ ), we must include the “error term”  $\varepsilon$ , which (in this case) is conditionally normal and marginally Student’s  $t$ . If scatter is large, then this confidence interval is wide. This  $X\%$  confidence interval does include approximately  $X\%$  of the observed data. But, of course, data include the undesirable scatter (error) that obscures the process of interest  $\mu$ . Ecologists would typically use a confidence interval on  $\mu$  to assess whether species differ in their responses to resources. This practice of ignoring scatter in data (a confidence interval on  $\mu$  rather than on  $y$ ) might be justifiable based on the initial model assumptions: stochasticity is “error.”

What if the response itself ( $\mu$ ) is stochastic? This occurs if any part of  $\mu$  is stochastic, including variable  $x$  or variable parameters. In the latter case, “parameters” are not subject to asymptotics. They are more like variables. Within a classical ML framework, such considerations are of passing interest, because there exists no structure that is generalizable to the vast number of ways in which stochasticity can affect data.

The 1990s saw a revolution in computational statistics resulting in a general framework that accommodated nearly all high-dimensional problems. The Gibbs sampler, a type of Markov chain Monte Carlo (MCMC simulation), is based on factorization of a joint (potentially high-dimensional) parameter distribution into low-dimensional (often univariate) distributions that can be simulated in a stepwise fashion (Gelfand and Smith 1990). Parameter variability is accommodated by adding a stage to the model, rather than by restructuring it. Stochasticity in  $x$  is obtained by including a probability for it, again, without modification of the basic structure. More meaningful counterparts for the frequentist confidence interval are available. If  $\mu$  is a latent (unobserved) process, then a posterior distribution defines our Bayesian confidence (“credible”) interval for  $\mu$  (this can be done for each  $\mu_j$ ). Predictive distributions for as yet unobserved data  $y$  are the basis

for the probability statement that (given model and data)  $X\%$  of observations will fall within a prescribed range. In the sections that follow, we illustrate these concepts in the context of the current example.

#### *When parameters are not identifiable*

Here we develop several aspects of the classical approach that limit inference. First, this approach provides no simple way to assimilate different types of information. Growth relationships have been studied many times, yet the traditional approach requires that each new data set be subjected to a new analysis. The *Acer rubrum* data set from closed-canopy conditions is typical (Fig. 2). There is a weak relationship with light, but insufficient information to identify the asymptote or half saturation constant; standard algorithms converge to a straight line. Moreover, the positive intercept does not agree with our expectation that there should be some “minimal” resource level below which plants cannot survive. Although seedlings can survive on stored reserves for extended periods, we do not expect them to persist in complete darkness.

There are several options. Common next steps include collecting additional light response data. Additional growth estimates at high light levels might help to identify an asymptote. It is possible that the broad scatter may result from variables that have not yet been considered. For seedling growth, plant ecologists usually turn next to soil moisture (Fig. 2, inset). This second resource does not explain appreciable variability in the data, nor does it explain variability that remains after accounting for the light response. The many previous studies have no impact on the analysis of this data set.

#### *Uncertain data*

A second limitation of the traditional approach relates to the fact that it ignores most sources of variability and uncertainty. Indices used to summarize resource density do not capture spatial and temporal variability in the same way as a consumer. In the case of light, the GSF represents how the changing sun angle throughout the day and year might be perceived by an understory plant that has variable and uncertain phenology and response times to sunflecks, other resources, and so on. A canopy photo does not uniquely define a GSF; two different photos could yield the same GSF, but there is no guarantee that a plant would perceive them identically. Different operators can obtain different GSF values from the same photo (the operator chooses a threshold parameter that determines how pixels are classified by the algorithm), and the same operator can obtain different values from photos taken on different days or at different times of day. The traditional model (Fig. 1) provides no accommodation for these “errors in variables” that affect perhaps all studies of trophic interactions.

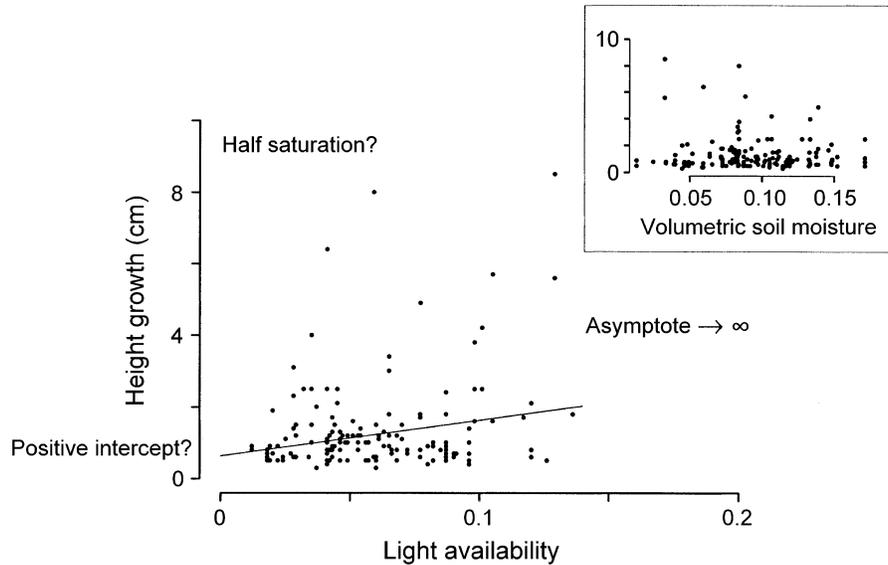


FIG. 2. The relationship between the fraction of light availability (the Global Site Factor, GSF) and height growth for *Acer rubrum* seedlings in a closed canopy. The large scatter is typical for such data sets. The nonlinear regression line does not resolve an asymptote or a half saturation constant, and it suggests the impossible positive intercept. The inset shows the same data plotted against soil moisture.

#### *Individuals differ*

Hierarchical models were developed for responses that are too variable for the assumptions of classical models (Carlin and Louis 2000). In the classical model, we ignore  $\varepsilon$  when making inference on  $\mu$ , because  $\varepsilon$  adds noise but not insight. Residual variation represented by  $\varepsilon$  has nothing to do with light response (although the estimate of  $\sigma^2$  does depend on the estimate of  $\mu$ ). The error  $\varepsilon$  enters this model as an addition to the deterministic response  $\mu$ , and it applies identically to all individuals. In other words, the classical model assumes that individual differences in  $y$  are independent of light availability and independent of the light response  $\mu$ .

If individuals vary in their response to light, say  $\mu_{ij}$ , then the response itself is stochastic. Now  $\mu$  has a mean, and its variance is not asymptotically zero. To determine whether one species has a different response from another, we construct credibility intervals that combine variability in  $\mu$  with the uncertainty in parameter estimates.

#### *The consequences of inappropriate models*

Predictions propagate bias in parameters and they can imply differences that do not exist. Bias can result from allocating scatter to  $\varepsilon$  (described by parameter  $\sigma^2$ ) when it really belongs somewhere else, e.g., in  $\mu$ . The fit is adjusted to put the scatter where the model permits. Confidence intervals for  $\mu$  based on a classical model can be narrow (if sample size is large), despite potentially large variability in the process summarized by  $\mu$ .

Clearly, we cannot identify all of the variables that contribute stochasticity, not even all of the important ones. Nonetheless, we have reason to believe that the response must involve some minimal requirement (trees do not grow in total darkness) and a saturating response (growth cannot increase indefinitely). Two things are obvious: (1) the resources themselves are variable and uncertain, and we have not allowed for their existence, and (2) individuals have a range of responses to light levels, depending on their own perception of it and on other factors that cannot be fully known. We require a method that allows for these sources of stochasticity and that accommodates their influences on inference and prediction.

#### ACCOMMODATING THE STOCHASTICITY

Hierarchical structures allow for context (Ver Hoef 1996), including variability among individuals within populations. Our implementation is Bayesian (Ellison 1996). Parameters have prior distributions that might be informed by assumptions: the asymptote  $G$  should be finite, the intercept should be negative ( $\beta_0 < 0$  or, equivalently,  $x_0 > 0$ ), the half saturation should be positive, and so on. We do not impose strong restrictions here (our priors provide flexibility). However, we note that restrictive priors need be no more subjective than the model selection itself; if we did not believe that growth requires at least some resource or that growth eventually saturates, we would not have chosen model 1–2 (Eqs. 1 and 2). Incorporating priors no more subjective than those that motivated model selection yields parameter estimates for Fig. 1 similar to those

obtained from more extensive data sets. We do not pursue that result here, because a Bayesian motivation more compelling than the capacity to formally incorporate external information is the natural structure that it provides for including stochasticity (Gelman et al. 1995).

We demonstrate a single model structure that admits reasonable assumptions about stochasticity (Fig. 1). We consider first variability and uncertainty in resources, followed by its interaction with variability in growth. Finally, we incorporate variable responses among individuals. In each case, we build on that same structure.

#### Uncertain resources

The observed resource level  $x_j$  is related to, but not the same as, the resource level experienced by the organism. Hereafter, the observed light level is  $x_j^{(o)}$ , and the underlying light level experienced by the plant is  $x_j$ . We use repeated observations (repeated photos) to estimate distributions of  $x_j$ . Because light availability ranges from complete darkness to full sunlight (GSF ranges from 0 to 1), we use a Beta density  $x_j \sim \text{Beta}(a_j, b_j)$  with parameters obtained by moment matching (Appendix A). The likelihood of the model now includes both types of “data” (see Fig. 1):

$$p(\mathbf{x}^{(o)}, \mathbf{y} | \mathbf{x}, \dots) = \prod_{j=1}^m \prod_{i=1}^{n_j} p_1(y_{ij} | x_j, \dots) \prod_{j=1}^m p_2(x_j | x_j^{(o)}, \dots) \quad (5)$$

where ellipses (...) on the left-hand side denote prior parameters. The second product series is the likelihood for the “error in variables,” with the light availability being a latent variable that can be estimated, but not directly observed.

We treat the resource–growth relationship (Eq. 1) as a nonlinear regression, but take advantage of linear parameters. Define the transformed variable for light at plot  $j$  to be  $z_j = x_j / (\theta + x_j)$ . A design matrix consists of a column of ones and a column with each  $z_j$  repeated  $n_j$  times. For the first plot, we have

$$\mathbf{Z}_1 = \begin{bmatrix} 1 & z_1 \\ \vdots & \vdots \\ 1 & z_1 \end{bmatrix}_{n_1 \times 2}$$

The full design matrix is

$$\mathbf{Z}_0 = \begin{bmatrix} \mathbf{Z}_1 \\ \vdots \\ \mathbf{Z}_m \end{bmatrix}_{M \times 2}$$

Further define parameter vector  $\mathbf{b}^T = [\beta_0 \beta_1]$ . Together with the beta distribution for light observations, we have the likelihood

$$p(\mathbf{x}, \mathbf{y} | \dots) = \prod_{j=1}^m \prod_{i=1}^{n_j} \mathcal{N}(y_{ij} | \mu_j, \sigma^2) \prod_{j=1}^m \text{Beta}(x_j | a_j, b_j). \quad (6)$$

The model includes priors on regression parameters,

$\mathcal{N}_2(\mathbf{b} | \mathbf{b}_0, \mathbf{V}_b)$ , with prior means  $\mathbf{b}_0^T = [b_{01}, b_{02}]$  and covariance matrix  $\mathbf{V}_b$ .

We examined three types of stochasticity in growth. As a baseline, we used the standard assumption of an unknown distribution of variances that applies everywhere (Fig. 1, Method 1). The conjugate prior for the residual variance is inverse gamma  $\text{IG}(\sigma^2 | s_1, s_2)$  with parameters  $s_1 = s_2 = 0.1$ . The prior covariance matrix is  $\mathbf{V}_b = v_0 \mathbf{I}_2$ , with  $v_0 = 1000$ , and  $\mathbf{I}_2$  being the rank 2 identity matrix. The complete model is

$$p(\mathbf{b}, \theta, \sigma^2, \mathbf{x} | \mathbf{x}^{(o)}, \mathbf{y}, \dots) = \prod_{j=1}^m \prod_{i=1}^{n_j} \mathcal{N}(y_{ij} | \mu_j, \sigma^2) \prod_{j=1}^m \text{Beta}(x_j | a_j, b_j) \quad (\text{likelihood}) \\ \times \mathcal{N}_2(\mathbf{b} | \mathbf{b}_0, \mathbf{V}_b) \text{Beta}(\theta | a_\theta, b_\theta) \text{IG}(\sigma^2 | s_1, s_2) \quad (\text{priors}). \quad (7)$$

The first two distributions of Eq. 7 are the likelihood (Eq. 6). The remaining three distributions are priors for regression parameters.

There are many parameters. Three come from  $\mu_j$ , and one ( $\sigma^2$ ) describes departures from this model. There are  $m$  unknown resource availabilities, for a total of  $m + 4$  parameters ( $m = 37$ ), i.e.,  $\beta_0, \beta_1, \theta, x_1, \dots, x_{37}, \sigma^2$ . Standard approximation methods are impractical; we use a Gibbs sampler (Appendix B).

#### Growth varies by location

Our second method allows that variances in growth might differ among plots (Fig. 1, Method 2). Plot-specific variances are

$$\sigma_j^2 = \frac{1}{n_j} \sum_{i=1}^{n_j} (y_{ij} - \mu_j)^2.$$

The complete model is

$$p(\mathbf{b}, \theta, \sigma_1^2, \dots, \sigma_m^2, \mathbf{x} | \mathbf{x}^{(o)}, \mathbf{y}, \dots) = \prod_{j=1}^m \prod_{i=1}^{n_j} \mathcal{N}(y_{ij} | \mu_j, \sigma_j^2) \prod_{j=1}^m \text{Beta}(x_j | a_j, b_j) \quad (\text{likelihood}) \\ \times \mathcal{N}_2(\mathbf{b} | \mathbf{b}_0, \mathbf{V}_b) \text{Beta}(\theta | a_\theta, b_\theta) \prod_{j=1}^m \text{IG}(\sigma_j^2 | s_1, s_2) \quad (\text{priors}). \quad (8)$$

There is a prior for each variance. The Gibbs sampler for this model is described in Appendix B.

#### The response varies among individuals

Our third method accommodates individual differences in light response in the form of a parameter set that describes the response for the  $i$ th individual on the  $j$ th plot,  $\mu_{ij}$ . Individual parameters are linked to the full population by distributions, defined by hyperparameters. These are “random effects” at the individual level. The full model,

$$\begin{aligned}
p(\mathbf{b}, \mathbf{b}_0, \mathbf{V}_b, \theta, \sigma^2, \mathbf{x} | \mathbf{x}^{(o)}, \mathbf{y}, \dots) \\
= \prod_{j=1}^m \prod_{i=1}^{n_j} \mathcal{N}(y_{ij} | \mu_{ij}, \sigma^2) \prod_{j=1}^m \text{Beta}(x_j | a_j, b_j) \\
\text{(likelihood)} \\
\times \prod_{j=1}^m \prod_{i=1}^{n_j} \mathcal{N}_2(\mathbf{b}_{ij} | \mathbf{b}_0, \mathbf{V}_b) \text{Beta}(\theta | a_\theta, b_\theta) \text{IG}(\sigma^2 | s_1, s_2) \\
\text{(priors)} \\
\times \mathcal{N}_2(\mathbf{b}_0 | \mathbf{h}, \mathbf{D}) W_2(V_b^{-1} | (w\mathbf{R})^{-1}, w) \text{ (hyperpriors)} \\
\text{(9)}
\end{aligned}$$

has a vector of growth parameters for each individual  $\mathbf{b}_{ij} = [\beta_{ij0}, \beta_{ij1}]^T$ , which is drawn from the distribution  $\mathcal{N}_2(\mathbf{b}_{ij} | \mathbf{b}_0, \mathbf{V}_b)$ . We do not have plot-specific growth variances, because we have allowed for variance at the individual level. The parameters have hyperprior density  $\mathcal{N}_2(\mathbf{b}_0 | \mathbf{h}, \mathbf{D})$  and the Wishart  $W()$  parameter covariance matrix, a multivariate generalization of the gamma distribution. The hyperprior for regression parameters is made noninformative by applying small prior precision  $\mathbf{D}^{-1}$ . The prior covariance matrix for hyperparameters  $\mathbf{D} = \text{Diag}(800, 1000)$  is relatively weak (Appendix B). The Wishart hyperparameters consist of weight  $w$  “degrees of freedom” and matrix  $\mathbf{R}$ . A noninformative prior has small  $w$  (but not less than the dimension of  $\mathbf{R}$ , i.e., 2) and  $\mathbf{R} \ll \mathbf{V}_b$ .  $\mathbf{R}$  is roughly the prior mean of  $\mathbf{V}_b$ . We use Carlin and Louis’ (2000) rule of thumb  $\mathbf{R} = \text{Diag}((r_1/8)^2, (r_2/8)^2)$ , where  $r_1 = 20$  and  $r_2 = 80$  are plausible parameter ranges for  $\mathbf{b}_{ij}$ , and  $w = M/20$ . Results here use prior parameter values for  $\mathbf{h} = [-30, 100]^T$ .

Priors based on previous studies could inform our estimates of regression parameters  $\mathbf{b}_{ij}$  (through  $\mathbf{h}$  and  $\mathbf{D}$ ) or their covariances  $\mathbf{V}_b$  (through  $w$  and  $\mathbf{R}$ ). Prior weight could be based on sample sizes used in previous studies and on the degree to which prior evidence is believed to inform this analysis. We defer further interpretation to the *Results*, because they illustrate the contributions of these different elements. The Gibbs sampler is described in Appendix B.

As a rough model selection guide, we used the Deviance Information Criterion (DIC), a generalization of the Bayesian Information Criterion (BIC). Both are based on the deviance statistic,  $D(\phi) = -2\ln L$  for likelihood  $L$  and parameter vector  $\phi$ , with a penalty for model complexity. For BIC, this penalty is the number of parameters,  $p_{\text{tot}}$ , times  $\ln M$ , where  $M$  is the total number of seedlings across all plots. For hierarchical models  $p_{\text{tot}}$  does not have the usual definition. Model complexity (effective number of parameters) is taken as  $p_e = \overline{D(\phi)} - D(\bar{\phi})$ , where  $\overline{D(\phi)}$  is the expected deviance over the posterior distribution of parameter vector  $\phi$  (the deviance evaluated over the full MCMC (Markov chain Monte Carlo) run), and  $D(\bar{\phi})$  is the deviance evaluated at the posterior mean parameter vec-

tor. Then  $\text{DIC} = \overline{D(\phi)} + p_e = 2\overline{D(\phi)} - D(\bar{\phi})$  (Carlin and Louis 2000), with 95% CIs for parameter and model.

Because of large variability, we employ two constraints on the hierarchical model in addition to those for other models. The hierarchical model allows for individual parameter estimates of  $\beta_{ij0} > 0$  and  $\beta_{ij1} < 1$ . This will be the case when individual variation is large. We allow for such estimates in the MCMC run, but exclude them from posteriors, which is equivalent to use of the equivalent prior restriction, truncated at zero (e.g., Gelman et al. 1995). We further constrain  $\theta$  to values close to the MLE for the entire population, using prior Beta values of  $a_\theta = 100$  and  $b_\theta = a_\theta (1/\theta_{\text{ML}} + 1)$ . Otherwise,  $\theta$  can be slow to converge, sometimes requiring  $>10^5$  MCMC steps. Note that this prior constraint emphasizes the contribution of the likelihood. We would not use this approach if our focus were inference on  $\theta$ ; our focus is on  $\mu$  and  $y$ . Our prior has no discernable impact on predictions of  $\mu$  and  $y$ , because other parameters are flexible to values of  $\theta$ . Moreover, by using different priors for each species (each centered on the ML estimate), we avoid any tendency for identical priors to cause species to appear more similar than they actually do. We use exclusively proper priors to insure identifiability (Gelfand and Sahu 1999).

The three approaches are modifications of several in the literature. Linear regression with known covariance is discussed by Carlin and Louis (2000). Our hierarchical structure is similar to that of Gelfand et al. (1990), who used a linear model to examine multiple observations per individual. Our “linearization” is recommended when efficiency is an issue (as it is here). Gamerman (1997) mentions its application in a Bayesian context. Gelman et al. (1995) and Carlin and Louis (2000) discuss the Wishart as a conjugate for the regression covariance matrix.

## RESULTS

### *The traditional model*

The classical method results in low parameter uncertainties (parameter error distributions are in Fig. 3a), in large part due to the fact that standard errors are proportional to  $M^{-1/2}$ . Error propagation to the response  $\mu$  (dashed lines in Fig. 3a) can give a misleading impression, in that  $\mu$  is estimated with great confidence, a result of assumptions that light is known precisely and that all individuals have this identical response to light. Tight confidence intervals on  $\mu$  belie the true scatter; far more than 95% of the data fall outside the 95% propagated parameter error. The distribution of residual variation  $\varepsilon$  accounts for the scatter (right side of Fig. 3a). Regardless of cause, the analysis assigns it to  $\varepsilon$ , because we have provided no other place for it. Of course, we could integrate the stochasticity associated with  $\varepsilon$  (and uncertainty in the estimate  $\sigma^2$ ), but this predictive interval for data, with negative val-

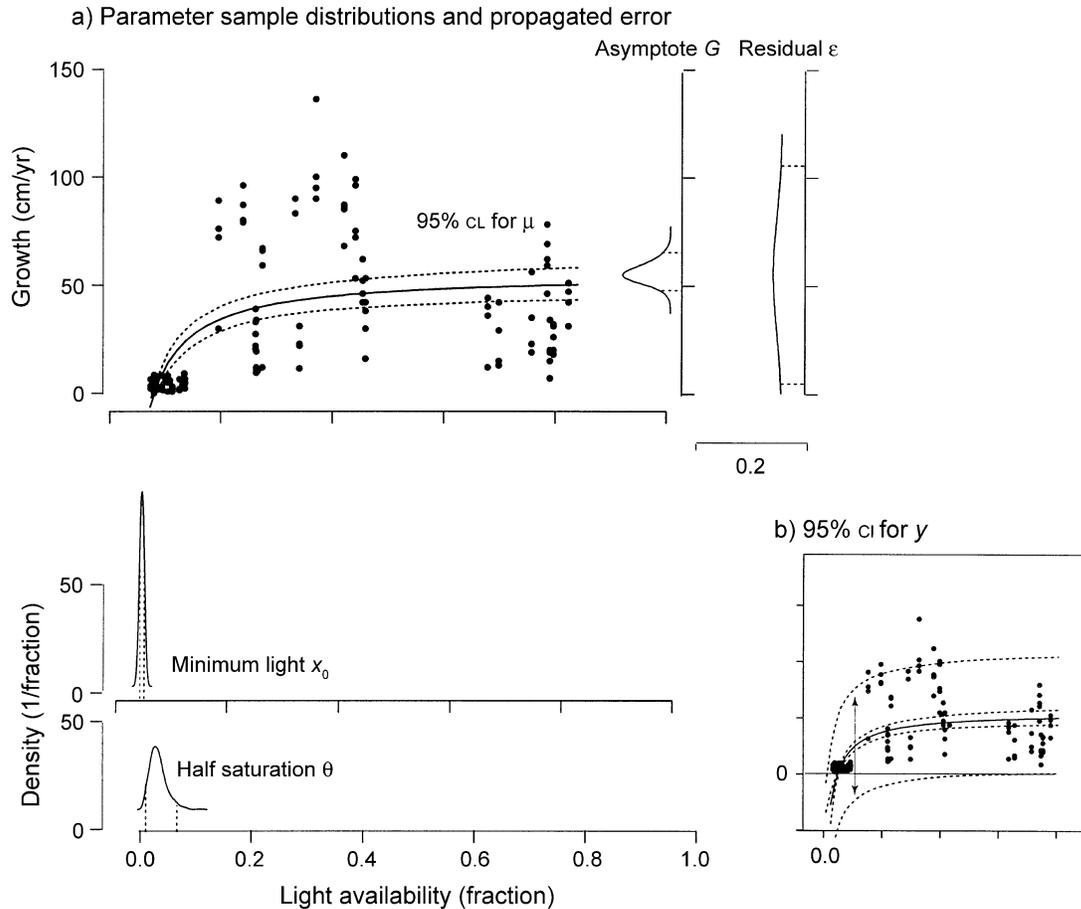


FIG. 3. (a) Maximum likelihood (Traditional model in Fig. 1) fit for the *Liriodendron* data set, showing parameter error distributions and the distribution of residuals  $\varepsilon$  (with variance determined by the MLE for  $\sigma^2$ ). Dashed lines show 95% confidence intervals for parameters and the model  $\mu$ , which is propagated from parameter error distributions using a nonparametric bootstrap. Part (b) shows the additional effect of error represented by the term  $\varepsilon$  (outer predictive intervals). For the  $k$ th bootstrap resample, we estimate a growth rate as a random draw from  $\mathcal{N}(\mu_{jk}, \sigma_k^2)$ . The marginal distribution is Student's  $t$  due to the sampling distribution of  $\sigma^2$ .

ues at the low light levels where light limitation matters most (Fig. 3b), would be of questionable use. Moreover, the assumption that scatter is “error” (the impetus for this classical model) is used to justify ignoring it and, hence, focus on the CI for  $\mu$ .

#### Uncertain resources

Acknowledging uncertainty in light affects the estimate of the response  $\mu$  by structuring the stochasticity in a different way. Each plot has its own distribution of light levels (Fig. 4). Each is conditionally informed (1) by the variances in light measurements, described by the sampling distribution  $x_j | x_j^{(o)} \sim \text{Beta}(a_j, b_j)$  and (2) by the probability that we would observe the sample of growth rates given the observed light level,  $y_{ij} | x_j \sim \mathcal{N}(\mu_j, \sigma^2)$ . The  $m$  densities of light (one for each plot) in Fig. 4 integrate contributions from both. Admitting that light is uncertain has impact on parameter estimates that control the fitted response at low light ( $\theta$ ) and a small effect

on the asymptote  $G$  (Appendix C). The credible interval shown for  $\mu$  is provided for comparison with the 95% CI obtained from the frequentist approach in Fig. 3a. It integrates uncertainty in regression parameters and resource variability and uncertainty. The error term is not included in this CI, so it is not a prediction of  $y$ . The contribution of  $\varepsilon$  is shown separately on the right side of the figure. The differences from traditional ML are not due to prior effects, because priors used for Fig. 4 are swamped by large  $M$ .

Fig. 4 indicates that uncertain light cannot account for a large fraction of the structure in these data. Residual uncertainty (described by  $\varepsilon$ ) is only slightly lower for Method 1 than for the traditional ML approach (Appendix C), because no amount of light uncertainty can bring the high growth rates in line with a model  $\mu$  that would describe all of the data well. The large scatter in growth rates could result from factors that vary from plot to plot or from differences among in-

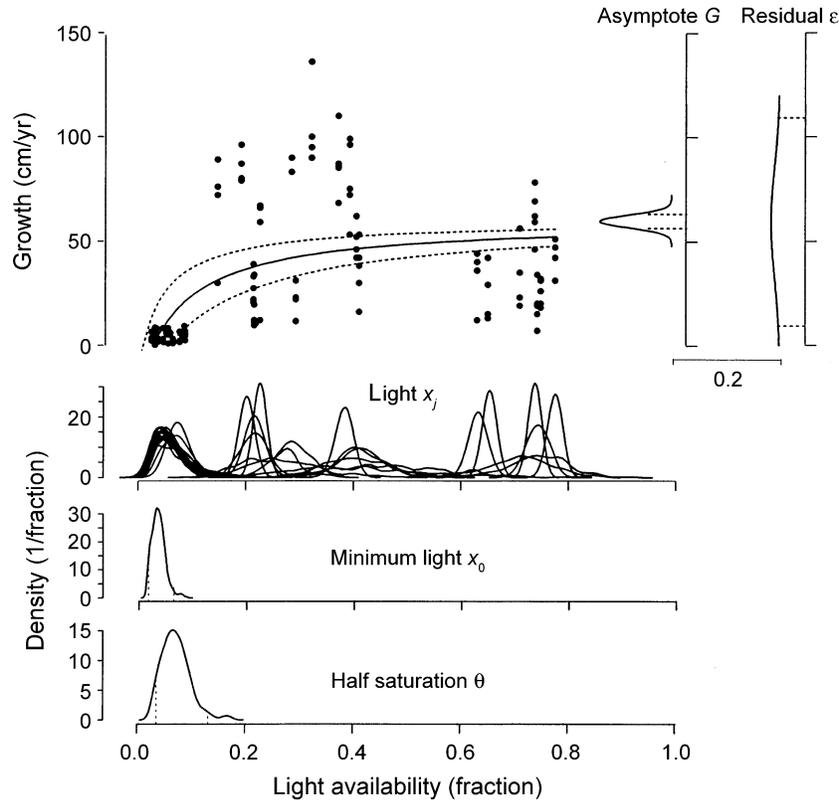


FIG. 4. Method 1 analysis (see Fig. 1) that accommodates variability in light. Bayesian posteriors are shown for all parameters, together with the predictive distribution for the model  $\mu$ . Dashed lines are 95% credible intervals. Those for the response  $\mu$  were constructed from quantiles drawn for each Gibbs step. The term  $\epsilon$  is shown, rather than the parameter  $\sigma^2$  itself, to provide a sense of the variation not “explained” by the model.

dividuals in how they respond to light. The next two methods address these possibilities.

*Growth uncertainty varies by location*

Acknowledgement of plot-to-plot differences in response provides a more realistic classification of uncertainty, but not in the best way. The variances in growth rates change from one plot to the next (right side of Fig. 5) in a way that is not captured well by a single residual variance  $\sigma^2$  (Fig. 4). When light is below the proportion 0.2, the distributions of light levels are tighter in Fig. 5 than they are in Fig. 4, because growth variances on these plots are low. The low variances on observations at low light cause them to dominate the regression, resulting in large estimates of  $\theta$  and a low asymptote  $G$ . To remedy this, we might choose a different parametric form (a function  $\mu$  with more parameters) or change the assumption about how variability operates (next section). This collection of distributions based on unknown factors captures the data structure in a way that seems more “realistic” than previous methods, to the extent that plot-to-plot differences are large, and previous methods would not allow it. The undue weight that it places on subsets of the data may

be undesirable, and the collection of empirical variance estimates is unenlightening.

*The response varies among individuals*

The model has not yet “captured” the structure of this relationship in the most useful form. The error term accommodates stochasticity that is unrelated to  $\mu$ . In the present context, this might typically be associated with measurement errors. Although growth data show broad scatter, height can, in most cases, be measured with relatively small error. Variability results from die-back and browse, but these effects can often be recognized. Poor meter-stick technique should not be the cause of order-of-magnitude error. In other words, the scatter contains a small contribution from measurement error.

Light uncertainty and variability can explain some, but not most, of this variability. If uncertain light accounts for the scatter, then data points well above the fitted model  $\mu$  experience true light levels that are much higher than those that we estimated from canopy photos ( $x \gg x^{(o)}$ ). However, even if we allowed the possibility that a plant at a measured light level of  $x^{(o)} = 0.1$  might, in fact, be receiving full sunlight ( $x$  near 1), we could

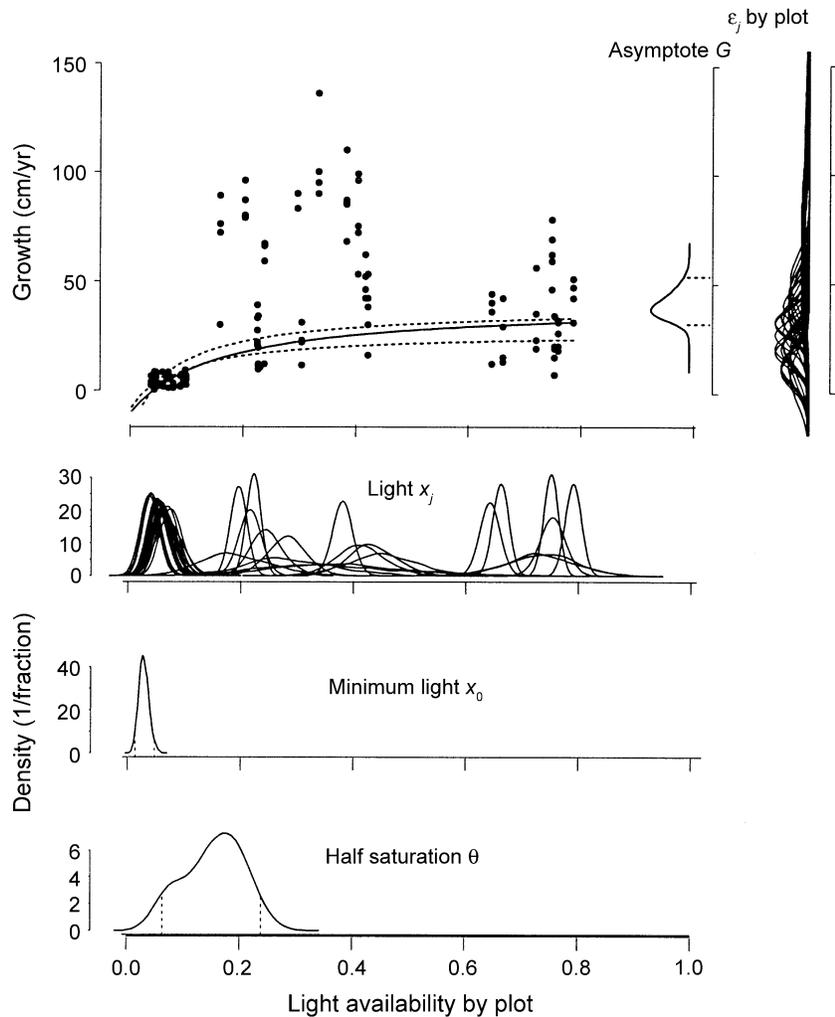


FIG. 5. Method 2 analysis (see Fig. 1) with plot-to-plot variability in growth (right-hand side). Bayesian posteriors are shown, together with the predictive distribution for the model. Symbols follow Fig. 4.

not accommodate much of the scatter in these data sets. There may be differences among individuals, related to inherent or local conditions, that cause each to respond to light somewhat differently. Although we cannot know all of these potential influences, we can allow for the variability that they confer.

The hierarchical model allows for this possibility, and it predicts that the individual differences are large. We have not changed our view that there is a relationship, like Eq. 2, that applies to all individuals in the population. By contrast, we also do not assume that they are completely independent. Each has its own parameter estimates (“individual”  $G$  and  $x_0$  posteriors in Fig. 6). The hierarchical model provides a summary of the population in the form of estimated hyperparameters. The structure of variability and uncertainty here is more plausible (and useful) than in the previous cases, because it is classified in the way in which ecologists believe that variability operates.

We have retained the assumption that true light levels are unknown. The parameters for any given individual can be viewed as a combination of the contributions of the data for that individual (light and growth), which are weighted by the precision  $\sigma^{-2}$ , and that for the population as a whole. The latter is the hyperparameter vector  $\mathbf{b}_0$ , which is weighted by the inverse of the parameter covariance  $\mathbf{V}_b$ . There is thus a “balance” between contributions from the individual and the population, maintained by the degree to which each contributes to the overall fit. The population provides the “glue” that precludes the overfitting of independent parameters for every individual. The 95% credible interval for  $\mu$  (dashed line in Fig. 6) contains  $\sim 95\%$  of the data points, because the residual variation is small (right side of Fig. 6). This credible interval for the population response is consistent with our view that individuals do not have negative growth rates (e.g., Fig. 3b), although they can lose height for

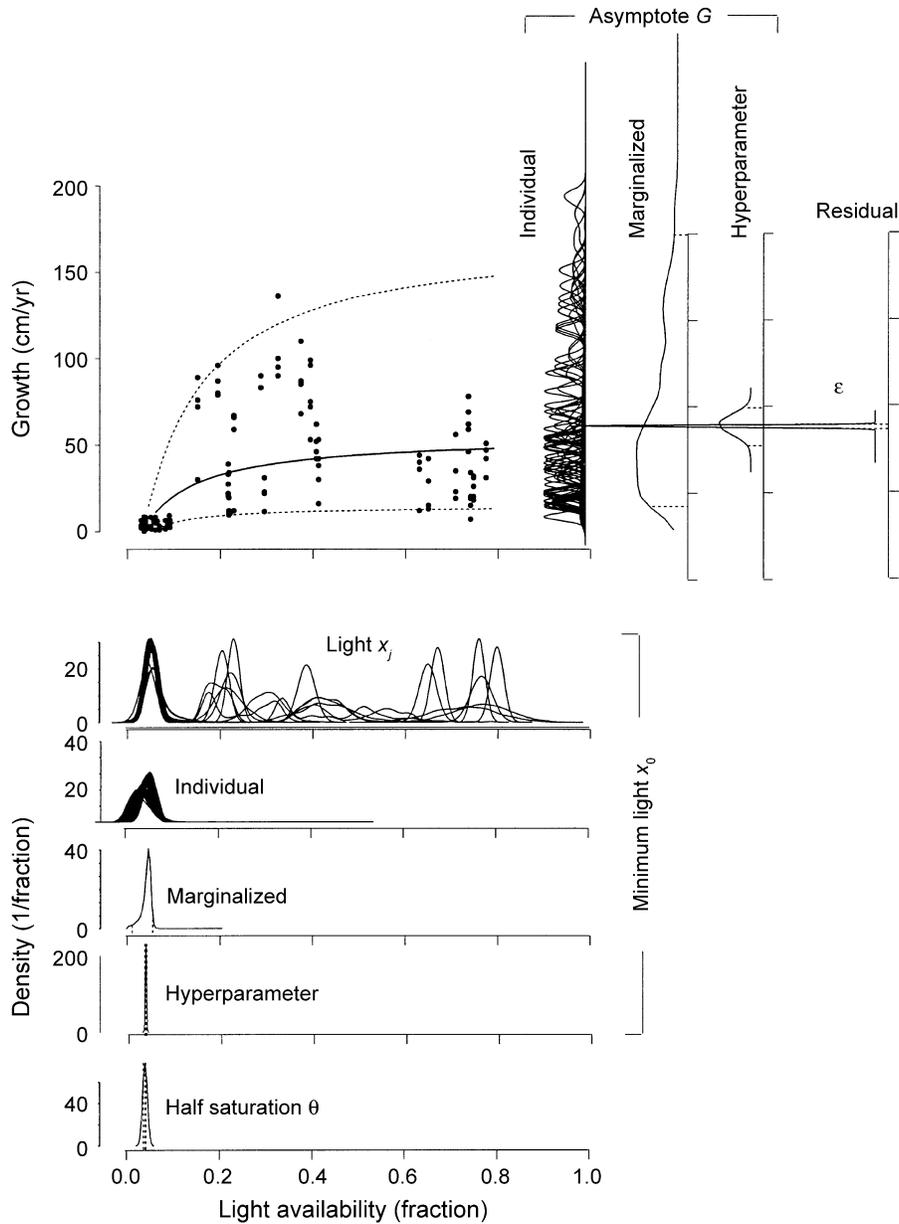


FIG. 6. The hierarchical analysis (Method 3 in Fig. 1), which assumes that each individual may respond differently to light (“individual” parameters), with individual distributions summarized by hyperparameters. Bayesian posteriors are shown for all parameters except  $\sigma^2$ , which is represented instead by the distribution for  $\epsilon$ , together with the predictive distribution for the model. The two hierarchical parameters are shown as individual posteriors, marginalized over all individuals, and hyperparameters.

reasons that would typically be unrelated to light. A low DIC indicates that the hierarchical model can explain much of the data structure; despite the large number of parameters, the number of “effective” parameters is relatively small.

*Comparing trophic responses*

A comparison of two species shows that a more realistic structure blurs species differences at all light levels. A classical fit predicts large differences in  $\mu$

between *Liriodendron* and *Acer rubrum* (Fig. 7b). Sample sizes are large and, thus, confidence intervals are narrow, so narrow that each  $\mu$  estimate assigns essentially zero probability to the other response. The models for the two species are statistically different.

The resource and individual variability result in broadly overlapping confidence intervals (Fig. 7a). Although the classical model predicts distinct differences, a fuller accounting of stochasticity predicts broad overlap. A full accounting of the uncertainty in the classical

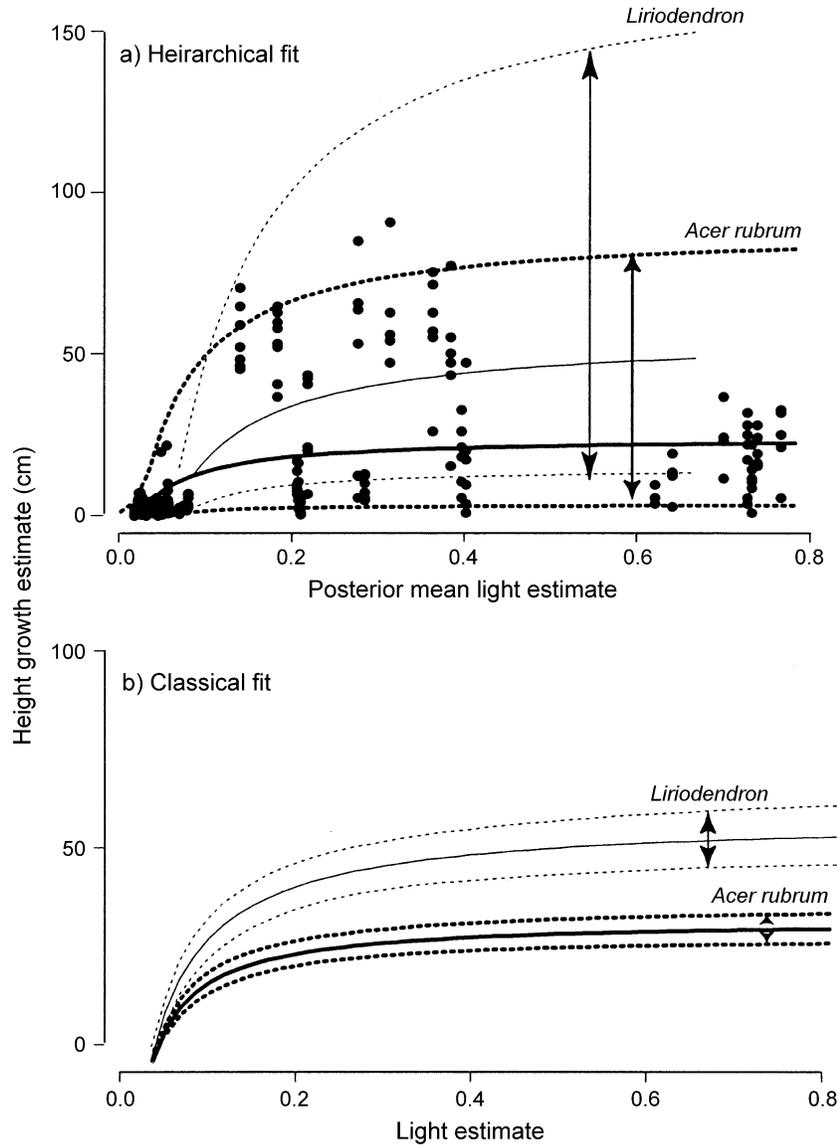


FIG. 7. A comparison of (a) hierarchical and (b) classical models for two species. Data are for *Acer rubrum* (data for *Liriodendron* are shown in previous figures). The figure shows 95% confidence intervals as they are typically interpreted. For the classical model, they are for the species-specific response  $\mu$ . For the hierarchical Bayesian model, they are credible intervals on the population response.

fit (i.e., propagating the effect of  $\sigma^2$ ) would also produce broad overlap (e.g., Fig. 3b), provided we were willing to accept the scatter as something other than "error." However, if we believed this, we could not justify the classical assumption that the scatter is "noise."

#### DISCUSSION

Inference from analyses such as Fig. 3 are used to identify species differences that might explain diversity. The traditional model assumes that resources are known precisely and that individuals respond identically. Species interactions parameterized in this way

challenge ecologists to identify trade-offs that permit coexistence. If the interaction depends on light interaction alone, Fig. 7b does not admit coexistence (e.g., Tilman 1982). When embedded in a dynamic model, Fig. 7b predicts rapid extinction of the poorer competitor. The fact that *Acer rubrum* is not extinct demands a trade-off, some combination of circumstances within which *Acer rubrum* can dominate (and thus, persist). The search for such trade-offs is a tradition in ecology. Our analysis from one site does not mean that this relationship will apply everywhere. The important lesson derives from the impact of admitting stochasticity in a realistic way.

Our approach is consistent with the way in which ecologists view sources of stochasticity, individually and in combination. Light is among the most important controls on plant growth. Light is highly variable, as are responses (consumption rates). Others (Pacala et al. 1994, Kobe 1999, Finzi and Canham 2000) find comparable levels of variability in growth responses to light. The hierarchical model is consistent with the view that individuals differ in their response, and that resources cannot be precisely known. It describes these contributions in ways that are readily applied to predictions of trophic interactions and to community dynamics. Parameters describe individual responses, but responses are interdependent and linked by hyperparameters that summarize the population. Confidence intervals are consistent with data: a  $X\%$  CI bounds approximately  $X\%$  of the data (Fig. 6). Because we begin with the assumption that these differences exist, we do not focus on a deterministic species-specific response that can misrepresent the degree of overlap (Fig. 7).

A sensible accounting of uncertainties need not motivate the search for trade-offs needed to “explain” coexistence. Trophic relationships are variable and broadly overlapping. Confidence envelopes do not predict large differences among species (Fig. 7a). Broad overlap does not mean that there are no biologically important differences. A number of recent reviews emphasize the potential importance of trade-offs, and there is a legacy of research to indicate that such trade-offs can promote coexistence. For example, the traditional classification of traits for early- and late-successional tree species undoubtedly contributes to patterns of diversity (e.g., Rees et al. 2001). This analysis, however, demonstrates that outcomes of trophic interactions will be highly variable. The rapid extinction implied by traditional data modeling (Fig. 7b) is not implied by the model that admits variability in a plausible way (Fig. 7a). Instead of a deterministic winner, either species may have the advantage at any light level. Regardless of whether “trade-offs” exist, their contribution is less efficacious than advocated by most current theory and as implied by classical data modeling.

#### *Alternative structures*

This analysis demonstrates a strategy for computational statistics that admits a broad range of assumptions within a common structure (Gelfand and Smith 1990). This example is not exhaustive. We explored four sets of assumptions by simply extending the original model (Fig. 1). The specific forms that we explored were motivated by insight concerning how variability might operate. Others can disagree with our specific assumptions, but it would be difficult to argue that they are inferior to those implicit to a classical approach. In competition with a model lacking individual variability, results indicate that individual differences matter. Obviously, no model is “correct,” and this flexible

framework serves to emphasize the importance of insight. The comparison of assumptions that can be explored within this framework further serves to emphasize the inadequacy of a classical approach when stochasticity can enter a model in many ways. Classical methods will continue to be valuable, provided their assumptions are explicit.

We do not explore all options, but we mention several that could be pursued further. A few extensions include different classifications of error structures and more detailed specification of covariate effects. For example, we used two hierarchical regression parameters, but assumed a single half saturation constant for the population. This was a practical choice, because (1) this approach seemed sufficient to capture the random effects of individuals, and (2)  $\theta$  is the “sticky” parameter that is slow to converge.  $M$  different values of  $\theta_i$  would be substantially slower. Because we are interested in species differences, rather than parameter inference, the regression covariances that result from this assumption are not of great concern. Nonetheless, this is an option that could be pursued.

The approach readily admits alternative or further subclassifications of variability, e.g., among subpopulations that might occupy different environmental settings. In principle, random effects on growth might be explored at a range of levels.

By omitting priors on the parameters for light availability, we did not admit as much variability as we might. We reasoned that the variability estimated from repeated photos was already large, and that the nature of the scatter in the data did not suggest that extreme variability in light was the cause. Nonetheless, this option is easily implemented and would not dramatically slow convergence.

The process model  $\mu$  can be further embellished to admit declines in growth rates that might occur at high light and to accommodate effects of covariates. Given the inevitable trade-off between detailed process and a rich error structure, including more covariates should permit a simplified treatment of error.

#### CONCLUSIONS

Our results do not imply that all species are trophically identical or that trophic trade-offs do not exist. There are many potential trade-offs that can affect interactions, even between the two species that we analyze here. Rather, these results demonstrate that traditional methods find differences where they are too weak to matter. Proper treatment of uncertainty and variability has a large impact on ecological inference and prediction. A hierarchical Bayesian framework readily accommodates resource levels that cannot be precisely known (organisms do not assess them in the same way as ecologists), and responses that vary among individuals. Failing to allow for these uncertainties results in biased estimates and inaccurate confidence envelopes.

Finally, general ecological principles derive from simple process models. Simplicity is justified by the potential for broad application. Models are extended to specific situations by adding "effects" to produce more complex process models. Most of the effects cannot be parameterized. Complex models can fail, in part, because they demand more information than data can provide. Application typically involves guesswork, and probability statements are rarely plausible.

Rich "error" structure can step in where complex process models fail or are undesirable. Modern statistical computation provides a way forward with a basic structure that allows application of simple models within realistic context. Hierarchical approaches sidestep deterministic complexity by admitting variability and uncertainty.

#### ACKNOWLEDGMENTS

For comments on the manuscript and discussion we thank K. Calder, A. Ellison, K. Gross, T. Ives, J. Hille Ris Lambers, M. Huston, S. LaDeau, J. MacLachlan, C. Wike, M. Wolosin, and two anonymous reviewers. Research support was provided by NSF grants DEB-9632854 and DEB-9981392 and DOE grant 98999.

#### LITERATURE CITED

- Alistair, R., D. S. Ellsworth, and S. W. Humphries. 2001. Possible explanation of the disparity between the in vitro and in vivo measurements of Rubisco activity: a study in loblolly pine grown in elevated pCO<sub>2</sub>. *Journal of Experimental Botany* **52**:1555–1561.
- Bazzaz, F. A. 1979. The physiological ecology of plant succession. *Annual Review of Ecology and Systematics* **10**:351–371.
- Carlin, B. P., and T. A. Louis. 2000. *Bayes and empirical Bayes methods for data analysis*. Chapman and Hall, Boca Raton, Florida, USA.
- Chan, M. S. and H. C. J. Godfray. 1993. Host-feeding strategies of parasitoid wasps. *Evolutionary Ecology* **7**:593–604.
- Chapin, F. S., E. D. Schulze, and H. A. Mooney. 1990. The ecology and economics of storage in plants. *Annual Review of Ecology and Systematics* **21**:423–447.
- Clark, J. S. 1991. Disturbance and tree life history on the shifting mosaic landscape. *Ecology* **72**:1102–1118.
- Clark, J. S. 2003. Uncertainty and variability in demography and population growth: a hierarchical approach. *Ecology*, **84**, in press.
- Connell, J. H., and R. O. Slatyer. 1977. Mechanisms of succession in natural communities and their role in community stability and organization. *American Naturalist* **111**:1119–1144.
- Ellison, A. M. 1996. An introduction to Bayesian inference for ecological research and environmental decision-making. *Ecological Applications* **6**:1036–1046.
- Ellsworth, D. S. 2000. Seasonal CO<sub>2</sub> assimilation and stomatal limitations in a *Pinus taeda* canopy. *Tree Physiology* **20**:435–445.
- Finzi, A. C. and C. D. Canham. 2000. Sapling growth in response to light and nitrogen availability in a southern New England forest. *Forest Ecology and Management* **131**:153–165.
- Gamerman, D. 1997. *Markov chain Monte Carlo: stochastic simulation for Bayesian inference*. Chapman and Hall, London, UK.
- Gelfand, A. E., S. E. Hills, A. Racine-Poon, and A. F. M. Smith. 1990. Illustration of Bayesian inference in normal data models using Gibbs sampling. *Journal of the American Statistical Association* **85**:972–985.
- Gelfand, A. E., and S. K. Sahu. 1999. Identifiability, improper priors, and Gibbs sampling for generalized linear models. *Journal of the American Statistical Association* **94**:247–253.
- 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.
- Glimskar, A., and T. Ericsson. 1999. Relative nitrogen limitation at steady-state nutrition as a determinant of plasticity in five grassland plant species. *Annals of Botany* **84**:413–420.
- Hinckley, T. M., J. P. Lassoie, and S. W. Running. 1978. Temporal and spatial variations in the water status of forest trees. *Forest Science Monograph* **20**:1–72.
- Holling, C. S. 1959. The components of predation as revealed by a study of small mammal predation of the European pine sawfly. *Canadian Entomologist* **91**:293–320.
- Hubbell, S. P. 2001. *The unified neutral theory of biodiversity and biogeography*. Princeton University Press, Princeton, New Jersey, USA.
- Kobe, R. K. 1999. Light gradient partitioning among tropical tree species through differential seedling mortality and growth. *Ecology* **80**:187–201.
- Krebs, C. J., R. Boonstra, S. Boutin, and A. R. E. Sinclair. 2001. What drives the 10-year cycle of snowshoe hares? *BioScience* **51**:25–35.
- Levins, R. 1979. Coexistence in a variable environment. *American Naturalist* **114**:765–783.
- Loehle, C. 1988. Tree life history strategies: the role of defenses. *Canadian Journal of Forest Research* **18**:209–222.
- MacArthur, R. H. 1972. *Geographical ecology*. Princeton University Press, Princeton, New Jersey, USA.
- May, R. M. 1973. *Stability and complexity in model ecosystems*. Princeton University Press, Princeton, New Jersey, USA.
- Oren, R., and D. Pataki. 2001. Transpiration in response to variation in microclimate and soil moisture in southeast deciduous forests. *Oecologia* **127**:549–559.
- Pacala, S. W., C. D. Canham, J. Saponara, J. A. Silander, R. K. Kobe, and E. Ribbens. 1996. Forest models defined by field measurements: estimation, error analysis, and dynamics. *Ecological Monographs* **66**:1–44.
- Pacala, S. W., C. D. Canham, J. A. Silander, Jr., and R. Kobe. 1994. Sapling growth as a function of resources in a north temperate forest. *Canadian Journal of Forest Research* **24**:2172–2183.
- Pacala, S. W., and M. J. Crawley. 1992. Herbivores and plant diversity. *American Naturalist* **140**:243–260.
- Paine, R. T. 1966. Food web complexity and species diversity. *American Naturalist* **100**:65–75.
- Pearcy, R. W., N. Tumosa, and K. Williams. 1981. Relationships between growth, photosynthesis and competitive interactions for a C<sub>3</sub> and a C<sub>4</sub> plant. *Oecologia* **48**:371–376.
- Pearcy, R. W., and W. Yang. 1998. The functional morphology of light capture and carbon gain in the Redwood forest understorey plant *Adenocaulon bicolor* Hook. *Functional Ecology* **12**:543–552.
- Rees, M., R. Condit, M. Crawley, S. Pacala, and D. Tilman. 2001. Long-term studies of vegetation dynamics. *Science* **293**:650–655.
- Sharkey, T. D., J. D. Seemann, and R. W. Pearcy. 1986. Contribution of metabolites of photosynthesis to postillumination carbon dioxide assimilation in response to light-flecks. *Plant Physiology* **82**:1063–1068.
- Tilman, D. 1982. Resource competition and community

- structure. Princeton University Press, Princeton, New Jersey, USA.
- Tilman, D. 1988. Plant strategies and the dynamics and structure of plant communities. Princeton University Press, Princeton, New Jersey, USA.
- Tilman, D. 1994. Competition and biodiversity in spatially structured habitats. *Ecology* **75**:2–16.
- Turchin, P., L. Oksanen, P. Ekerholm, T. Oksanen, and H. Henttonen. 2000. Are lemmings prey or predators? *Nature* **405**:562–565.
- Turchin, P., A. D. Taylor, and J. D. Reeve. 1999. Dynamical role of predators in population cycles of a forest insect: an experimental test. *Science* **285**:1068–1071.
- Ver Hoef, J. M. 1996. Parametric empirical Bayes methods for ecological applications. *Ecological Applications* **6**: 1047–1055.

#### APPENDIX A

Light parameter equations are available in ESA's Electronic Data Archive: *Ecological Archives* E084-001-A1.

#### APPENDIX B

Algorithms for Gibbs sampling for a growth model are available in ESA's Electronic Data Archive: *Ecological Archives* E084-001-A2.

#### APPENDIX C

Tables of parameter estimates, posterior means, standard errors, and 95% credible intervals for light availability and individual plants are available in ESA's Electronic Data Archive: *Ecological Archives* E084-001-A3.

#### SUPPLEMENT

A supplement containing data files is available in ESA's Electronic Data Archive: *Ecological Archives* E084-001-S1.