

Ecological Applications, 16(1), 2006, pp. 62–73
© 2006 by the Ecological Society of America

ANALYSIS OF NEIGHBORHOOD DYNAMICS OF FOREST ECOSYSTEMS USING LIKELIHOOD METHODS AND MODELING

CHARLES D. CANHAM¹ AND MARÍA URIARTE

Institute of Ecosystem Studies, Box AB, Millbrook, New York 12545 USA

Abstract. Advances in computing power in the past 20 years have led to a proliferation of spatially explicit, individual-based models of population and ecosystem dynamics. In forest ecosystems, the individual-based models encapsulate an emerging theory of “neighborhood” dynamics, in which fine-scale spatial interactions regulate the demography of component tree species. The spatial distribution of component species, in turn, regulates spatial variation in a whole host of community and ecosystem properties, with subsequent feedbacks on component species. The development of these models has been facilitated by development of new methods of analysis of field data, in which critical demographic rates and ecosystem processes are analyzed in terms of the spatial distributions of neighboring trees and physical environmental factors. The analyses are based on likelihood methods and information theory, and they allow a tight linkage between the models and explicit parameterization of the models from field data. Maximum likelihood methods have a long history of use for point and interval estimation in statistics. In contrast, likelihood principles have only more gradually emerged in ecology as the foundation for an alternative to traditional hypothesis testing. The alternative framework stresses the process of identifying and selecting among competing models, or in the simplest case, among competing point estimates of a parameter of a model. There are four general steps involved in a likelihood analysis: (1) model specification, (2) parameter estimation using maximum likelihood methods, (3) model comparison, and (4) model evaluation. Our goal in this paper is to review recent developments in the use of likelihood methods and modeling for the analysis of neighborhood processes in forest ecosystems. We will focus on a single class of processes, seed dispersal and seedling dispersion, because recent papers provide compelling evidence of the potential power of the approach, and illustrate some of the statistical challenges in applying the methods.

Key words: *AIC; likelihood methods; model selection; neighborhood processes; parameter estimation; seed dispersal; seedling establishment; spatially explicit models.*

INTRODUCTION

Advances in computing power in the past 20 years have led to a proliferation of spatially explicit, individual-based models of population and ecosystem dynamics (DeAngelis and Gross 1992, Judson 1994, Dunning et al. 1995, Mooij and DeAngelis 1999a). Spatially explicit models of forest dynamics such as ZELIG (Urban et al. 1989) and SORTIE (Pacala et al. 1996) shift the focus from the dynamics of discrete patches or gaps (*sensu* Watt 1947, Levin and Paine 1974) to interactions among individual organisms, and allow a more detailed treatment of environmental heterogeneity at a variety of spatial scales. While these models have shifted focus away from the visualization of forests as mosaics of discrete gaps, there has yet been no

clear, corresponding statement of a theoretical alternative to traditional gap-phase dynamics (but see Dieckman et al. 2000 for a review of recent attempts). By tracking birth, death and recruitment for each individual, spatially explicit, individual-based models have the potential to provide insights into the ecological and evolutionary processes that determine the structure and composition of plant communities. In effect, the individual-based models encapsulate an emerging theory of “neighborhood” dynamics, in which fine-scale spatial interactions regulate the demography of component tree species (Pacala et al. 1996, Law and Dieckmann 2000). The spatial distribution of component species, in turn, regulates spatial variation in a whole host of community and ecosystem properties, with feedbacks on component species.

Studies of neighborhood dynamics appear to be a particularly powerful vehicle for the integration of population dynamics and ecosystem processes. Attempts to “link species and ecosystems” have been a pervasive theme in ecology in the past decade (Jones and

Manuscript received 14 April 2004; revised 8 July 2004; accepted 26 August 2004; final version received 29 November 2004. Corresponding Editor: D. S. Schimel. For reprints of this Invited Feature, see footnote 1, p. 3.

¹ E-mail: ccanham@ecostudies.org

Lawton 1994, Loreau 2000). Canopy tree species exert strong control over many aspects of forest ecosystem processes, which, in turn, have feedbacks that influence canopy tree dynamics (Canham and Pacala 1994, Rothe and Binkley 2001). Studies of forest ecosystem processes have traditionally used plot designs that average across variation in local neighborhood composition and structure. It is very difficult using this approach, however, to predict changes in ecosystem processes as canopy composition changes. This is particularly true when the spatial interactions of different species are nonadditive, and the properties of a plot cannot be predicted from simple data on the relative abundance of species within the plot. The emerging theory of neighborhood dynamics provides an alternative approach that incorporates the mechanistic, spatially explicit interactions between species dynamics and ecosystem processes, including heterogeneity in the physical environment. Spatially explicit models then provide the tool for synthesizing these processes into spatial and temporal predictions of system dynamics (Gratzler et al. 2004).

The shift to spatially explicit, individual-based models has been facilitated by development of new methods of analysis of field data (i.e., Ribbens et al. 1994, Kobe et al. 1995, Canham et al. 2001), in which measurements of environmental factors (e.g., light) and critical demographic rates (e.g., recruitment) and ecosystem processes (e.g., nitrogen mineralization) are analyzed in terms of the spatial distributions of neighboring trees and physical environmental factors. The analyses are based on likelihood methods and information theory (Edwards 1992, Eliason 1993, Hilborn and Mangel 1997, Burnham and Anderson 2002, Hobbs and Hilborn 2006), and allow a tight linkage between the models and explicit parameterization of the models from field data. They also allow quantification of uncertainty in both parameter values and model predictions (Pacala et al. 1996, Mooij and DeAngelis 1999b).

Likelihood methods have been widely used in other disciplines for many years, but have only recently gained a foothold in ecological research (Hobbs and Hilborn 2006). Hilborn and Mangel (1997) provided one of the first general primers on the approach written specifically for ecologists. Burnham and Anderson (2002) provide a more formal treatment of the methods and put them in the broader context of information theory. At the heart of the methods is the explicit interplay between data and models, with “model” used in the sense of a mathematical statement of the quantitative relationships that are assumed to have generated the observed data. Classical hypothesis testing is replaced by the more general process of model selection and comparison, using likelihood and parsimony to compare the strength of evidence for competing models. Likelihood methods provide analogues for many traditional parametric statistical tests, but often without

many of the restrictive assumptions required for parametric statistics.

Our goal in this paper is to review recent developments in the use of likelihood methods and modeling for the analysis of neighborhood processes in forest ecosystems. We will not try to review the wide range of issues that can be addressed with the methods. Instead, we will focus on a single class of processes, seed dispersal and seedling dispersion, because a spate of recent papers on this subject provide both compelling evidence of the potential power of the approach, and illustrate some of the statistical challenges in applying the methods.

LIKELIHOOD METHODS AND MODEL SELECTION

Maximum likelihood methods have a long history of use for point and interval estimation in statistics. In contrast, likelihood principles (Edwards 1992) have only more gradually emerged in many fields of science as the foundation for an alternative to traditional hypothesis testing (e.g., Royall 1997, Johnson and Om-land 2004, Hobbs and Hilborn 2006). The alternative framework stresses the process of identifying and selecting among competing models (as statements of multiple working hypotheses, *sensu* Chamberlain 1890), or in the simplest case, among competing point estimates of a parameter of a model (i.e., the maximum likelihood estimate [MLE]). This is in stark contrast to the traditional frequentist approach of rejection of a single “null” hypothesis (particularly since the null hypothesis is often of limited utility). In contrast to P values, which do not provide a direct measure of the support in the data for any particular alternative hypothesis, likelihoods (or more commonly, log-likelihoods) can be calculated for an entire set of alternative models and a continuum of parameter values, and provide an explicit measure of the strength of evidence for any particular model or parameter value. Likelihood profiles and “support intervals” (Edwards 1992) provide much more intuitive alternatives to traditional confidence intervals as a means of presenting uncertainty in a particular parameter estimate. Bayesian methods build on the likelihood framework and provide an explicit way of dealing with uncertainty in the parameter estimates. The Bayesian approach is particularly well suited in situations where parameter estimates will be used for policy and decision-making (Ellison 1996, Hobbs and Hilborn 2006).

On its own, the use of likelihood as a basis for model selection will always favor more complex models, without consideration of issues of parsimony (Burnham and Anderson 2002). While there is no guarantee that nature is simple, there are still compelling reasons to favor simple explanations over complex ones. There are a number of methods for model selection that combine the fit of the model (as measured by log-likelihood) and the complexity of the model (as measured

by the number of parameters in the model; Johnson and Omland 2004). The most common methods in use in the ecological literature have been likelihood ratio tests (LRT) and some form of the Akaike information criterion (AIC). Likelihood ratio tests suffer from some of the same limitations as traditional frequentist hypothesis testing, namely, the need to conduct pairwise comparisons of alternate models, using significance tests based on the P value of the LRT statistic. AIC, by comparison, is based on information theory principles, and provides a quantitative measure that can be used to rank all competing models by penalizing more complex models (Burnham and Anderson 2002). In effect, AIC is a measure of the information lost by using a given (simpler) model to approximate reality (Burnham and Anderson 2002). Thus, models with the smallest AIC are preferred. AIC and other information theory based measures also provide a foundation for inference based on model averaging (Burnham and Anderson 2002, Johnson and Omland 2004).

There are four general steps involved in a likelihood analysis: (1) model specification, (2) parameter estimation using maximum likelihood methods, (3) model comparison, and (4) model evaluation. The first step also requires specification of the appropriate underlying statistical distribution of the error term (ϵ). The maximum likelihood estimates of the parameters of each candidate model are then determined using either one of the many standard local optimization methods (i.e., Gauss-Newton or Simplex), or global optimization methods such as simulated annealing (Goffe et al. 1994). Investigators also have the option of using “model averaging” to provide more robust parameter estimates by combining results from alternate models (Neuman 2003, Wintle et al. 2003). This is particularly appropriate when there is no single model with overwhelming support in the data. The models are then typically compared by calculating the (log) likelihood of observing the data, given the functional form (scientific model) and the maximum likelihood estimates of the parameters of the model, and then calculating AIC, given log-likelihood, the number of parameters in the model, and sample size (if correcting AIC due to small numbers of samples relative to the number of parameters) (Burnham and Anderson 2002). Once a “best” model has been identified, the final step involves evaluation of the model in the traditional terms of goodness of fit, bias, and prediction error. Competing assumptions about distribution of model errors (e.g., normal vs. lognormal) can also be compared in a parallel manner.

While likelihood methods provide powerful tools for hypothesis testing via model selection, the effectiveness of the approach ultimately lies in the insight of the investigators in choosing appropriate and interesting scientific models, and their skills in collecting appropriate data. Most of the recent examples of use of

likelihood methods in the ecological literature (and all of the examples presented in this paper) involve analysis of comparative field data; however, the methods are equally powerful in the analysis of experimental data (remember that the “M” in GLM stands for “model”). In particular, by forcing an investigator to focus clearly on both the underlying models being compared and the statistical properties of the data (appropriate distribution of the errors in order to calculate likelihood), likelihood-based analyses typically free the investigator from many of the restrictive assumptions of traditional parametric analyses. From an educational perspective, likelihood methods are slightly more difficult to master initially than traditional frequentist statistics, but students often find likelihood methods more logical, comprehensive, and flexible.

A LIKELIHOOD FRAMEWORK FOR ANALYSIS OF NEIGHBORHOOD PHENOMENA IN FORESTS

Forests are ideal candidates for neighborhood analyses. Canopy trees provide the bulk of both the physical and biological structure—living and dead, above- and belowground—in forest ecosystems. The dominant ecosystem fluxes (e.g., throughfall, litterfall, nutrient uptake, leaching losses) are primarily vertical rather than horizontal, and the identities of canopy trees overhead and nearby have distinctive effects on a whole host of ecosystem processes at the “neighborhood” scale (0–50 m; Binkley and Giardina 1998, Finzi et al. 1998a, b). It is possible to map the distributions of forest trees in large plots (Condit 1998), and readily obtainable measures of plant size such as dbh (diameter of the stem at 1.35 m height) can be used in allometric equations to predict above and belowground biomass (Jenkins et al. 2003) and the physical dimensions of tree crowns (e.g., Canham et al. 1994, Chave et al. 2003).

Consider the state of some ecosystem property y at a point p (i.e., soil nutrient availability, seed rain, growth of a plant at that location, etc.). For practical reasons, $y(p)$ is measured in a sample drawn from a location of finite dimensions (p'). The examples we will discuss in this paper all employ simple, additive models that begin with the general form

$$y(p') = g(p') \sum_{i=1}^n f(x_i) + \epsilon \quad (1)$$

for $i = 1, \dots, n$ “agents” (x_i , typically trees) within some maximum distance (r), and where $g(p')$ is some function of the physical and biological attributes of location p' . Each agent (x_i) has a vector of attributes; for a tree this would typically include species, dbh, and location relative to location p' , and may also include environmental conditions at the location of each tree. In principle, r should be chosen to be large enough that all agents that have measurable impact on the process at location p' are included. In practical terms, this is

often not possible, and the analysis then suffers from a biased sample. This has been a pervasive problem in recent studies using inverse modeling to estimate seed dispersal and seedling dispersion (i.e., Ribbens et al. 1994, Clark et al. 1998, LePage et al. 2000, Muller-Landau et al. 2004). In principle, it is also possible to use the analysis to provide estimates of the effective neighborhood radius for a given process, i.e., the value of $r' < r$ at which the effects of a neighboring agent (e.g., a tree of a given species, size, etc.) become negligible for the process of interest. For instance, in our analyses of neighborhood competition among trees we estimate the maximum radius of the zone of influence of neighbors on a focal tree located at p' (Uriarte et al. 2004, Canham et al. 2004).

Example: seed dispersal and seedling establishment

There have been three general approaches to the characterization of patterns of seed dispersal and seedling dispersion around parent trees (Greene and Calogeropoulos 2002): (1) direct measurement at locations around single, isolated individuals (i.e., Kitajima and Augspurger 1989, Stoyon and Wagner 2001), (2) development of theoretical models based on the specific mode of dispersal (particularly the aerodynamics of wind dispersal) and parameterization from measurements of the properties of both the propagules and the dispersing agents (wind, water, animals, etc.; reviewed in Turchin 1998), and (3) neighborhood analyses using inverse modeling to estimate the parameters of dispersal functions (or “kernels”), based on measurements of seed rain or seedling abundance at a set of sample locations, and a map of the spatial distribution and sizes of potential parent trees in the vicinities of the sample locations (e.g., Ribbens et al. 1994, Clark et al. 1998, 1999, LePage et al. 2000, Muller-Landau et al. 2004). We will outline the general approach used in this third approach, and use it to highlight some of the challenges in the use of likelihood methods.

Model specification and functional forms.—Recently published neighborhood models of seed rain and seedling dispersion (e.g., Ribbens et al. 1994, Clark et al. 1998, 1999, LePage et al. 2000) have all taken the general form of Eq. 1. Models of seed rain generally assume that seed input is independent of the environment at location p' , so $g(p') = 1$. In contrast, seedling establishment is often strongly influenced by local conditions. LePage et al. (2000) treated $g(p')$ as a vector of scalars (ranging from 0 to 1) to take into account variation in the favorability of a set of seedbed substrates within sample quadrats. Most of the attention, however, has been focused on the selection of an appropriate functional form for $f(x_i)$. $F(x_i)$ is typically decomposed into two, multiplicative parts:

$$f_1(x_i) = A \left(\frac{\text{dbh}_i}{C} \right)^\alpha \quad (2)$$

and

$$f_2(x_i) = \frac{1}{\eta} \exp[-B(\text{dist}_i)^\beta] \quad (3)$$

where dbh_i and dist_i are the size of and distance to tree i , respectively; A , α , B , and β are estimated parameters, and η is a normalization constant (i.e., the arcwise integration of the exponential seed dispersal kernel). The first part predicts total number of seeds or seedlings (A) produced by potential parent trees within a specified radius of the sampled point (p') and the second describes the dispersal kernel of seeds or seedlings. The use of the normalization constant produces estimates of A in units of the total number of seeds or seedlings produced by a tree of a standard size (of $\text{dbh } C$) throughout the seed shadow. Clark et al. (1999) treat B as a random variable rather than a fixed parameter in deriving their “2Dt” model. The parameter B determines the speed in the decline of recruitment as the distance to the parent tree increases while β determines the general shape of the dispersal kernel.

All of the recent papers on seed rain and seedling dispersion have assumed that the process is isotropic (i.e., seed input does not vary as a function of the direction to the source tree), but in an analysis of leaf litterfall, Staelens et al. (2003) have introduced a simple modification of the B term to account for anisotropy. We have tested for anisotropy in both seed and seedling distributions on our own data, and find common but not ubiquitous presence of strong directional effects in both seed rain and seedling distributions (C. D. Canham, *unpublished results*). The anisotropy term estimates both the angle and increment in distance attributable to wind. Independent data on prevailing wind direction at the experimental site can provide a reality check for these estimates.

There are a number of critical assumptions embedded in Eqs. 2 and 3, both implicitly and as an overt attempt to simplify the analysis for methodological reasons. Perhaps the most fundamental is the assumption that seed rain is a monotonically declining function of distance from the parent tree (Eq. 3). This appears to be a pervasive assumption in studies of seed dispersal, but Greene and Calogeropoulos (2002) have argued persuasively that there is a high likelihood that the modal dispersal distance is displaced some distance away from the source tree for both wind and animal dispersed species, and that a lognormal function is more appropriate than the exponential function in Eq. 3. Greene et al. (2004) have examined seed and seedling data from a number of temperate tree species at sites in Quebec and find that a lognormal function often has higher likelihood (and lower AIC) than monotonically declining functions, including the “2Dt” function of

Clark et al. (1999). We have reanalyzed the data on seedling dispersion for five of the species presented in LePage et al. (2000), and the lognormal function is a better fit to the data than their original exponential model for four of the five species (Appendix A). This highlights one of the most important principles of a likelihood approach (or indeed in any approach based on model selection): the results are conditioned by the scope of the models under consideration (Beissinger and Snyder 2002). The likelihood analysis will identify the "best" model, given the set of models under consideration. This does not in any way guarantee that the selected model is the best model among all possible models. This leads to Recommendation #1:

It is always preferable to focus on models with functional forms that are motivated by hypotheses about underlying mechanisms, but it is equally important to have an open mind and test competing models that are flexible enough to allow the data to display novel or unexpected patterns.

Parameter estimation.—

1. *Sampling considerations and parameter trade-offs.*—The published studies have generally simplified Eqs. 2 and 3 even further by assuming that the α and β parameters are fixed. In all of the studies, α has been fixed arbitrarily at a value of 2 (e.g., Ribbens et al. 1994, Clark et al. 1998, LePage et al. 2000, Muller-Landau et al. 2004). As a result, total plant fecundity is assumed to scale roughly linearly with plant biomass (since tree biomass is approximately linearly related to dbh^2 ; Jenkins et al. 2003). This appears to be a reasonable assumption, and Greene and Johnson (1994) have integrated results from a wide range of empirical studies on total plant fecundity using a relationship that is linearly related to tree basal area (i.e., proportional to dbh^2) and a power function of seed mass. As Ribbens et al. (1994) and Clark et al. (1998) note, the original motivation for fixing the exponent (rather than estimating the most likely value given the data) was related to difficulties in parameter estimation due to trade-offs between the A and α parameters. The trade-offs will be particularly problematic if there is a limited range of tree diameters in the neighborhoods (for example, if the stands are even-aged, with relatively uniformly sized canopy trees). There is much less a priori justification for arbitrarily fixing the value of β at any given value (studies have used values ranging from 0.5 to 3 to arbitrarily produce various degrees of convexity in the seed dispersal function). There is no compelling mechanistic theory to support any given value of β , and there is much less potential for trade-offs with the B parameter as long as there is reasonable variation in the distribution of distances to potential source trees

Parameter trade-offs, i.e., when the same value of the response variable [$y(p')$] can occur as a result of different combinations of two or more parameters, are

common in models with more than just a few variables. One common consequence of parameter trade-offs is that the likelihood surface is relatively flat in the region of parameter space defined by the trade-off. The particular maximum likelihood parameter estimates produced by the optimization procedure may then vary widely given idiosyncrasies of the data set. This can seriously weaken the generality of the resulting model, and limit the model's value when applied to a new data set.

We have explored the consequences of potential trade-offs between A vs. α and B vs. β in a number of both published and unpublished data sets, and using hypothetical data sets generated using known parameter values in hypothetical stands. The hypothetical data sets clearly demonstrate that it is possible to estimate all four parameters simultaneously as long as there is reasonable variation in both size and distances to parent trees around sample locations (Fig. 1, Appendix A). Our reanalysis of data for seedling dispersion of five species in LePage et al. (2000) (where α was fixed at 2 and β was fixed at 3) suggests that the assumption that fecundity is linearly related to parent plant biomass (i.e., that $\alpha = 2$) is not well supported by the data (Appendix B). We allowed α to vary from 0 to 4, and maximum likelihood estimates for the five different species varied from 0.1 to 3.7, and models in which α was allowed to vary typically had much stronger support in the data than models in which it was fixed at an arbitrary value of 2 (Appendix B). The lowest estimated value for α occurred in a fast-growing, relatively short-lived pioneer tree species (paper birch, *Betula papyrifera*). For this species, fecundity does not appear to vary significantly with tree size once the trees are above the minimum tree size used in the analysis (15 cm dbh). Values greater than 3 occurred for two of the most shade tolerant species (western red cedar, *Thuja occidentalis*; Amabilis fir, *Abies amabilis*). We interpret these to reflect patterns in which the bulk of seed production occurs in very large individuals. We have explored this in subsequent analyses by estimating an additional parameter to specify a minimum size (dbh) below which trees are assumed to be nonreproductive. The estimated minimum reproductive size was in the range from 26 to 29 cm dbh for two tree species in a temperate rain forest in New Zealand (W. Ruscoe, *personal communication*). Finally, it is worth noting that none of the recently published analyses or the tests described here allow for functional relationships that are either asymptotic (i.e., fecundity reaches a plateau at some size), or in which fecundity actually declines at some point with increasing size (senescence).

These results suggest that, even given the importance of avoiding parameter trade-offs, caution should be used in fixing parameters at some arbitrary value (no matter how reasonable the assumption may be). This is a simple corollary to Recommendation #1. Never-




FIG. 1. Dispersal functions estimated using maximum likelihood methods and a 50-m neighborhood radius for six hypothetical data sets (A–F). The data sets were generated as a Poisson process with an expected mean density in each quadrat given by Eqs. 2 and 3, with parameter values as described in Appendix A. The estimated functions [A(est)–F(est)] are compared with the true, underlying functions (A–F) specified by the parameters in Appendix A. The functions are asymptotic, and only the first 15 m are shown.

theless, the potential for parameter trade-offs is likely to be high in neighborhood models. The most effective means of avoiding trade-offs is to make sure that the data set contains a wide range of values for the attributes of the agents in the neighborhood model (i.e., tree size). The goal of the sampling design for a neighborhood analysis is not to provide unbiased estimates of the population of neighborhood represented in the sample, but rather to provide a sufficient range of observations so that the estimation of parameters in the neighborhood model is unbiased. This is a critical distinction: the former goal leads to the traditional dictates of random sampling, etc., while the latter places a premium on distributing sampling effort across a gradient of neighborhood conditions. This leads to Recommendation #2:

Always be alert for potential parameter trade-offs, and distribute sampling effort broadly across gradients in the critical attributes of the agents in the neighborhood models. Where the spatial distribution and properties of the agents are known beforehand (as in the case of studies in existing mapped forest stands), that knowledge should be used to guide the spatial distribution of sampling effort.

2. *Bias due to misidentification of agents.*—The analysis of seed rain highlights another critical issue in sampling design and data collection for neighborhood analyses: the misidentification of the agents (x_i) in Eq. 1. The misidentification typically takes two forms: (1) absence of data on distant agents (i.e., agents

located outside the mapped neighborhood), and (2) inclusion of “inactive” agents (i.e., inclusion of all trees in the neighborhood for analysis of seed rain, when only some unknown subset of the trees has actually produced seed). Both forms of misidentification have potentially serious implications for parameter estimation and model interpretation. For many species of trees, particularly wind-dispersed species, the asymptote of the dispersal function is clearly nonzero, and some fraction of seeds disperses long distances (Nathan and Mueller-Landau 2000). In these cases, a finite mapped stand represents a truncated data set of potential parent trees. In some ways, this problem is analogous to censored data problems in traditional survival analyses.

Data sets with misidentified agents have the clear potential to bias estimates of both fecundity (Eq. 2) and dispersal (Eq. 3). The potential bias in estimating fecundity is immediately evident: the analysis attributes the observed seed rain or seedling density to the agency of the $i = 1, \dots, n$ trees used in the analysis for that sample location. If n does not include all source trees, then the method will consistently overestimate the average per-tree fecundity. It is worth pointing out that this bias in parameter estimation does not necessarily translate into bias in the average predicted seed rain or seedling establishment when the functions are used in a model such as SORTIE (Ribbens et al. 1994, Pacala et al. 1996), as long as the model uses the same neighborhood radius r to predict seed rain or seedling establishment as was used in the analysis of field data.




FIG. 2. Estimated dispersal functions using different neighborhood search distances and functional forms of the dispersal kernel for a hypothetical data set generated from 1-km² area of source trees and 200 1-m² quadrats distributed in two transects through a central 1-ha portion of the source area. Simulated seed rain was generated as a Poisson process, with a lognormal dispersal function (Logn) and parameter values given in Appendix C. The “theoretical” curve is the true, underlying lognormal dispersal function used to generate the hypothetical data set. The Weibull function allowed both B and β in Eq. 3 to vary. The 2Dt curve was estimated using the dispersal function presented by Clark et al. (1999). Distances in the legend refer to neighborhood search distances (radius).

The degree of bias in estimation of parameters related to the dispersal kernel (Eq. 3) caused by finite mapping is not as self-evident. We have explored the nature of the bias using Monte Carlo simulations with hypothetical mapped stands with a variety of distributions and abundances of parent trees to create simulated patterns of seed rain generated from known parameter values (Fig. 2, Appendix C). We then sampled the patterns with varying sizes of mapped stands and spatial distributions of samples within the stands (Appendix C). It's hard to generalize because of the almost infinite variety of potential distributions of source trees and samples, but it appears that there is relatively little bias in parameter estimation for the dispersal kernel with a finite mapping of potential source trees as long as the mapping extends well out past the mean dispersal distance (MDD; i.e., $r > 2 \times \text{MDD}$; Fig. 2). When the neighborhood radius (r) is in the range from 1 to $2 \times \text{MDD}$, the analysis generally produces slight overestimates of dispersal and large overestimates of fecundity. When the neighborhood radius is $< \text{MDD}$, there is considerable bias (and increasing uncertainty) in parameter estimates for both fecundity and the dispersal kernel. The bottom line is that if $f(x_i)$ (Eq. 1) does not go to zero at some finite distance (and in practical terms, within the maximum distance included in the neighborhood mapping), then there will be some degree of bias in the parameter estimates. The bias may be quite small and entirely acceptable, but it should be

acknowledged and reasonable attempts made to assess its magnitude.

It is hard to generalize about the potential bias caused by inclusion of “inactive” agents (i.e., inclusion of nonreproductive trees) in the analysis. If only some fraction of trees produce seed in a given year (as is common for all sexually dimorphic species and for many other tree species [cf. Herrera et al. 1998, Koenig and Knops 2000]), then the estimates of fecundity will clearly be a biased underestimate for the trees that actually reproduce, although the estimate may be unbiased as a population average for all trees in the neighborhood. We have done limited tests with hypothetical data sets, and when the nonreproductive trees (inactive agents) are a random subset of the adult trees, bias is restricted to estimates of fecundity, while estimates of the parameters of the dispersal kernel are generally unaffected. There are, however, almost certainly cases in nature where the spatial distribution of reproductive vs. nonreproductive individuals is nonrandom, and the potential bias in estimates of the dispersal function caused by treating all trees as reproductive will depend on the exact nature of the spatial pattern.

Existing studies have differed in how they address this issue. Ribbens et al. (1994) and LePage et al. (2000) used a fixed radius r representing the most conservative distance within which all potential parent trees were known (i.e., used only the trees within a mapped stand that were within a radius r of each sample




FIG. 3. Dispersal functions estimated using different neighborhood configurations for the hypothetical data set and the underlying “theoretical” parameters for a lognormal dispersal function used in Fig. 2. The “25 m radius” and “50 m radius” curves used a circular neighborhood of all trees within 25 and 50 m, respectively, of each sample location. The “0 m from edge” curve represents a sample layout in which the two transects of adjacent 100 1-m² quadrats were placed along two opposite edges of a 100 × 100 m mapped stand. The “25 m from edge” and “50 m from edge” curves were estimated using the same sample layout as the “0 m from edge” curve, but with a mapped stand that had an additional 25- or 50-m buffer around the central 100 × 100 m stand.

location). The advantage of this approach is that each sample has the same “censoring” (i.e., omitting all trees beyond distance r). Another alternative is to use all trees within a given stand in the analysis for all seed traps in the stand (e.g., Clark et al. 1998). Since sample locations differ in their position relative to the nearest edge of the plot, the nature of the censoring will be different for each observation. A third approach uses data from within the mapped stands to create a hypothetical “buffer” zone around the mapped stand to account for the contribution from parent trees outside the mapped stand (Muller-Landau et al. 2004). The buffer zone could have a uniform density of adult trees equivalent to that found in the plot, or a random distribution of parent trees drawn from size distributions of parent trees inside the plot. We examined the consequences of the first two options (using a conservative fixed minimum radius for the neighborhood around each sample, vs. using all trees in the plot for all samples) using hypothetical data sets (Fig. 3, Appendix C). Our results suggest that the first option (used by Ribbens et al. 1994 and LePage et al. 2000) is unnecessarily conservative. As long as there is a reasonable minimum radius mapped around each sample (\approx MDD), then using the additional but censored data for trees beyond that minimum radius appears to reduce potential bias (compare the “25 m radius” curve with the “25 m from edge” curve in Fig. 3). If many of the quadrats are very near the edge of the plot, however, even using all of the

trees in the plot can result in significant bias, particularly in the estimation of fecundity (Fig. 3, “0 m from edge”).

These simple simulations are hardly definitive, but they highlight the benefits of large mapped stands for studies of forest neighborhood dynamics (e.g., Uriarte et al. 2004). They also lead to Recommendation #3:

Monte Carlo simulations with either hypothetical neighborhoods or actual mapped forest stands can and should be used to evaluate potential bias in parameter estimation given assumptions about parameter values and the size of the neighborhood included in the analysis. In the absence of prior information on the neighborhood size needed to reduce bias to acceptable levels, be prepared to sample in very large neighborhoods to reduce the potential for bias.

Model comparison: selecting the best model.—There are a number of methods available for model comparison, i.e., the assessment of the strength of support in the data for alternate models (reviewed in Burnham and Anderson 2002, Johnson and Omland 2004, Hobbs and Hilborn 2006). Measures such as AIC combine the likelihood of the model with a penalty for increasing model complexity and in some cases a penalty for small sample size relative to the number of parameters in the model (i.e., AIC_{corr}). AIC has the benefit of a foundation in information theory as a basis for balancing goodness of fit and model complexity: it is proportional to the

expected Kullback-Leibler information lost through use of the model to approximate reality (Burnham and Anderson 2002). Thus, among a set of models, the model with the lowest AIC (or AIC_{corr} for small sample sizes) is selected. The differences in AIC between alternate models provide a measure of the relative difference in the strength of evidence for the different models, relative to the complexity of the model. Akaike weights (W_i) provide a means to normalize the strength of evidence for one of a fixed set of alternate models (AIC_i) relative to the best model (AIC_{min}):

$$W_i = \frac{\exp\left[-\frac{1}{2}(AIC_i - AIC_{min})\right]}{\sum_{j=1}^n \exp\left[-\frac{1}{2}(AIC_j - AIC_{min})\right]} \quad (4)$$

The Akaike weights have two noteworthy uses: they can be summed across all candidate models containing a particular parameter to assess the relative importance of that parameter, and the weights can be used in model averaging to combine parameter estimates from across the set of candidate models to produce more robust estimates (Burnham and Anderson 2002, Hobbs and Hilborn 2006).

The issue of parsimony lies at the heart of model selection. While AIC provides an objective basis for penalizing more complex models (on the basis of information theory), there is still room for judgment in how far an investigator should go in pursuit of the simplest model that is consistent with the data. Consider the analysis of seedling dispersion by LePage et al. (2000). Field observations suggested that variation in seedbed substrate had a strong effect on seedling establishment and subsequent density (LePage et al. 2000). They included a vector of estimated favorabilities for eight distinct substrate types in the $g(p')$ term in Eq. 1. Asymptotic support intervals for the parameter estimates overlapped for many of the substrate favorability terms. It is likely that there were simpler models that could have been produced by combining similar substrate favorability parameters, and that these simpler models would have had lower AIC scores than the full model. However, the number of plausible groupings, even based on a posteriori analysis of the estimated differences among substrate types, is very large. We would argue that the most sensible application of the principle of parsimony in this case would be to compare the full model with a simpler model that omitted all consideration of substrate differences. If that simpler model is demonstrably worse (i.e., has a higher AIC), and there is an a priori interest in the parameter estimates for the different substrates, then we would argue that there is little benefit in exhaustively testing alternate models to identify the most parsimonious groupings of substrate types. On the other hand, we can imagine cases where an investigator has a priori

interest in a categorical variable with m possible values, but where the objective is to specifically test whether there is support in the data for distinguishing between different subsets of the m values. In this case, thorough testing of alternate groupings would be justified (although increasingly tedious as m increases).

This leads to Recommendation #4, which echoes other recent discussions of model selection methods (e.g., Burnham and Anderson 2002, Johnson and Omland 2004):

Application of the principle of parsimony in model comparison does not require a systematic search for the simplest possible model. Rather, a priori interests of the investigators in testing a set of alternate modes (representing multiple working hypotheses, sensu Chamberlain 1890) should guide the model comparison process.

Model evaluation: How good is the best model?— Once the best model among the set of candidates is selected, it is still intuitive to ask: "How good is it?" In a likelihood framework, AIC provides an answer, but only within the context of comparison with alternate models evaluated using the same data set. Many of the standard diagnostic tools for evaluating regression models, however, are also appropriate for evaluating neighborhood models in a likelihood framework. Perhaps the most important are measures of prediction bias (typically assessed using the slope of the relationship between predicted and observed values) and prediction error (using RMSE or comparable metrics). We suspect that most ecologists using traditional regression methods are much more interested in general measures of goodness of fit such as R^2 than in the significance of test statistics comparing their model with a null hypothesis. We also suspect that statisticians find it tedious that practitioners are so preoccupied with goodness of fit. To most investigators, a high R^2 value has meaning that transcends issues of parameter estimation and hypothesis testing, and lends confidence that the processes embodied in the model are an apt (and close to sufficient) description of the system. In a sense, goodness of fit is a measure of the skill (or luck) of the investigator in properly specifying the model and collecting appropriate data with minimal error. This is not something that statisticians can generally help us with, so it's perhaps not surprising that the subjects of goodness of fit and model evaluation receive relatively little attention in either standard textbooks on statistics or in recent treatments of likelihood methods (e.g., Burnham and Anderson 2002).

While ecologists are probably most familiar with R^2 as an overall measure of goodness of fit ($= 1 - SSR/SST$, where SSR is the sum of the squared residuals and SST is the total sum of squares), there are a number of alternative measures of the overall goodness of fit of a model, including χ^2 and measures of deviance. It

is likely that R^2 is favored by many investigators simply because it has such an intuitive interpretation and a scale that is invariant with respect to sample size or the measurement scale of the response variable. For models with normally distributed errors and where the mean and the variance are not correlated, R^2 can indeed provide a useful measure of goodness of fit. There are many cases, however, where either process or measurement error is not normally distributed, and is better described by a lognormal or the gamma distribution where the variance is an explicit function of the mean. This is particularly true for studies of seed rain and seedling dispersion, where the likelihood functions typically have either a Poisson or negative binomial distribution (e.g., Clark et al. 1998). For either distribution, the variance of the process is proportional to the mean (and literally equal to the mean for a Poisson process). Thus, a measure such as R^2 will vary as a function of the mean of the expected values of the model. An investigator can then increase R^2 simply by sampling in such a way that the expected values are small (i.e., using small quadrats or short census intervals).

Clark (2003) addresses a far more profound issue regarding the assessment and interpretation of uncertainty in model predictions. Students are routinely taught to assess the adequacy of sampling effort in terms of degree of confidence in a parameter estimate. In contrast, there are both practical and theoretical issues for which assessment of the variability in a parameter or a prediction is of as much interest as the assessment of its central tendency (Clark et al. 2003). Investigators typically consider uncertainty to be a function of measurement error and sampling effort, and something that could be reduced through more diligent effort. It is worth restating the obvious, however, that any analyses involving organisms with either genotypic or phenotypic variability will encounter “process” error as well as “measurement” error. Hierarchical Bayesian models (Clark 2003, Clark et al. 2003, Latimer et al. 2006) offer a powerful method of explicitly incorporating much more complex (and biologically reasonable) error structures than the traditional approach used in Eq. 1, including explicit incorporation of variability in parameters. Clark (2003) illustrates the benefits of incorporating more realistic error structures in models of population dynamics. We expect that those benefits also pertain to a much wider range of studies (e.g., Hoeting et al. 2006). This leads to Recommendation #5:

Variability is a fact of life for biologists. Rather than treat error as a nuisance, there are good reasons to seek more explicit and biologically realistic incorporation of error terms and uncertainty in neighborhood models. Hierarchical Bayesian models appear to offer a promising approach.

FUTURE CHALLENGES IN NEIGHBORHOOD ANALYSES

To date, neighborhood models of the type we have described have largely been in the domain of studies of population dynamics and resource competition for sessile organisms (particularly plants). We believe that there is much wider potential application of the methods for any ecological process where the spatial configuration of component organisms plays a central role. Studies of leaf litterfall (Ferrari and Sugita 1996, Staelens et al. 2003) represent one obvious (but perhaps trivial) example. The approach is likely to be seen as brute phenomenology by investigators who seek mechanistic explanation of ecosystem processes at much more reductionist levels. For studies of forest ecosystem dynamics, however, we would suggest that there is already overwhelming empirical evidence that tree species differ significantly in their effects on a whole host of ecosystem processes, and a growing literature on the exact mechanisms that underlie these differences. There have been a number of attempts to link models of tree population dynamics with models of forest ecosystem dynamics, but other than the pioneering work with LINKAGES (Post and Pastor 1996), none to our knowledge have yet explicitly incorporated a neighborhood approach.

Neighborhood models will be critical where there are nonadditive effects of mixtures of the influence of different species (Finzi and Canham 1998, Gartner and Cardon 2004). In our framework, nonadditive effects would be accommodated using higher-order (nonlinear) terms in Eq. 1. If there are strong nonlinearities, nonspatial models will give biased predictions of the effects of the mixtures, as a result of Jensen's inequality (Pastor et al. 2000, Duursma and Robinson 2003). For example, both antagonistic and synergistic effects of species mixtures are common in studies of leaf litter decomposition (Gartner and Cardon 2004). There has been some recent progress in predicting at least the direction of the effect (antagonistic vs. synergistic) of different mixtures of species from the traits of component species (Wardle et al. 1998, Eviner and Chapin 2003), but the wide range of empirical results found in recent studies suggests that a great deal more field research will be needed to predict the magnitude of nonadditive effects. The challenge for a neighborhood approach is to use the results of recent experiments and process studies to construct relatively simple models that capture the mechanisms that underlie these nonadditive effects, without resorting to the brute force (and extremely large number of resulting parameters) required to describe any potential nonadditive effects among all possible subsets of n species (Finzi and Canham 1998).

ACKNOWLEDGMENTS

We thank Tom Hobbs, Saran Twombly, and Dave Schimel for organizing a stimulating workshop, and for their feedback on the ideas in this paper. This work was supported by grants

from the National Science Foundation (DEB-0087214) and the A. W. Mellon Foundation. This paper is a contribution to the program of the Institute of Ecosystem Studies.

LITERATURE CITED

- Beissinger, S. R., and N. F. R. Snyder. 2002. Water levels affect nest success of the Snail Kite in Florida: AIC and the omission of relevant candidate models. *Condor* **104**: 208–215.
- Binkley, D., and C. Giardina. 1998. Why trees affect soils in temperate and tropical forests: the warp and woof of tree/soil interactions. *Biogeochemistry* **42**:89–106.
- Burnham, K. P., and D. R. Anderson. 2002. Model selection and multimodel inference: a practical information-theoretic approach. Second edition. Springer-Verlag, New York, New York, USA.
- Canham, C. D., A. C. Finzi, S. W. Pacala, and D. H. Burbank. 1994. Causes and consequences of resource heterogeneity in forests—interspecific variation in light transmission by canopy trees. *Canadian Journal of Forest Research* **24**:337–349.
- Canham, C. D., P. T. LePage, and K. D. Coates. 2004. A neighborhood analysis of canopy tree competition: effects of shading versus crowding. *Canadian Journal of Forest Research* **34**:778–787.
- Canham, C. D., and S. W. Pacala. 1994. Linking tree population dynamics and forest ecosystem processes. Pages 84–93 in C. G. Jones and J. H. Lawton, editors. *Linking species and ecosystems*. Chapman and Hall, New York, New York, USA.
- Canham, C. D., M. J. Papaik, and E. F. Latty. 2001. Interspecific variation in susceptibility to windthrow as a function of tree size and storm severity for northern temperate tree species. *Canadian Journal of Forest Research* **31**:1–10.
- Chamberlain, T. C. 1890. The method of multiple working hypotheses. *Science* **15**:92–96.
- Chave, J., R. Condit, S. Lao, J. P. Caspersen, R. B. Foster, and S. P. Hubbell. 2003. Spatial and temporal variation of biomass in a tropical forest: results from a large census plot in Panama. *Journal of Ecology* **91**(2):240–252.
- Clark, J. S. 2003. Uncertainty and variability in demography and population growth: a hierarchical approach. *Ecology* **84**:1370–1381.
- Clark, J. S., E. Macklin, and L. Wood. 1998. Stages and spatial scales of recruitment limitation in southern Appalachian forests. *Ecological Monographs* **68**:213–235.
- Clark, J. S., J. Mohan, M. Dietze, and I. Ibanez. 2003. Coexistence: how to identify trophic trade-offs. *Ecology* **84**: 17–31.
- Clark, J. S., M. Silman, R. Kern, E. Macklin, and J. HilleRisLambers. 1999. Seed dispersal near and far: patterns across temperate and tropical forests. *Ecology* **80**: 1475–1494.
- Condit, R. 1998. Tropical forest census plots: methods and results from Barro Colorado Island, Panama and a comparison with other plots. Springer-Verlag, Berlin, Germany.
- DeAngelis, D. L., and L. J. Gross, editors. 1992. *Individual-based models and approaches in ecology*. Routledge, Chapman and Hall, New York, New York, USA.
- Dieckmann, U., R. Law, and J. A. J. Metz, editors. 2000. *The geometry of ecological interactions: simplifying spatial complexity*. Cambridge University Press, Cambridge, UK.
- Dunning, J. B., D. J. Stewart, B. J. Danielson, B. R. Noon, T. L. Root, R. H. Lamberson, and E. E. Stevens. 1995. Spatially-explicit population models: current forms and future uses. *Ecological Applications* **5**:3–11.
- Duursma, R. A., and A. P. Robinson. 2003. Bias in the mean tree model as a consequence of Jensen's inequality. *Forest Ecology and Management* **186**:373–380.
- Edwards, A. W. F. 1992. *Likelihood—expanded edition*. Johns Hopkins University Press, Baltimore, Maryland, USA.
- Eliason, S. R. 1993. *Maximum likelihood estimation: logic and practice*. Sage University paper series on quantitative applications in the social sciences, 07–096. Sage University, Newbury Park, California, USA.
- Ellison, A. M. 1996. An introduction to Bayesian inference for ecological research and environmental decision-making. *Ecological Applications* **6**:1036–1046.
- Eviner, V. T., and F. S. Chapin. 2003. Functional matrix: a conceptual framework for predicting multiple plant effects on ecosystem processes. *Annual Review of Ecology Evolution and Systematics* **34**:455–485.
- Ferrari, J. B., and S. Sugita. 1996. A spatially explicit model of leaf litter fall in hemlock–hardwood forests. *Canadian Journal of Forest Research* **26**:1905–1913.
- Finzi, A. C., and C. D. Canham. 1998. Non-additive effects of litter mixtures on net n mineralization in a southern New England forest. *Forest Ecology and Management* **105**:129–136.
- Finzi, A. C., C. D. Canham, and N. Van Breemen. 1998a. Canopy tree soil interactions within temperate forests: species effects on pH and cations. *Ecological Applications* **8**: 447–454.
- Finzi, A. C., N. Van Breemen, and C. D. Canham. 1998b. Canopy tree soil interactions within temperate forests: species effects on soil carbon and nitrogen. *Ecological Applications* **8**:440–446.
- Gartner, T. B., and Z. G. Cardon. 2004. Decomposition dynamics in mixed-species leaf litter. *Oikos* **104**:230–246.
- Goffe, W. L., G. D. Ferrier, and J. Rogers. 1994. Global optimization of statistical functions with simulated annealing. *Journal of Econometrics* **60**:65–99.
- Gratzner, G., C. Canham, U. Dieckmann, A. Fischer, Y. Iwasa, R. Law, M. J. Lexer, H. Sandmann, T. Spies, B. Splechtna, and J. Szwagrzyk. 2004. Spatio-temporal development of forests—current trends in field studies and models. *Oikos* **107**:3–15.
- Greene, D. F., and C. Calogeropoulos. 2002. Dispersal of seeds by animals and wind. Pages 3–23 in J. Bullock, R. Kenward, and R. Hails, editors. *Dispersal ecology*. Blackwell Press, Oxford, UK.
- Greene, D. F., C. D. Canham, K. D. Coates, and P. T. LePage. 2004. An evaluation of alternative seed dispersal functions for trees. *Journal of Ecology* **92**:758–766.
- Greene, D. F., and E. A. Johnson. 1994. Estimating the mean annual seed production of trees. *Ecology* **75**:642–647.
- Herrera, C. M., P. Jordano, J. Gitian, and A. Traveset. 1998. Annual variability in seed production by woody plants and the masting concept: reassessment of principles and relationship to pollination and seed dispersal. *American Naturalist* **152**:576–594.
- Hilborn, R., and M. Mangel. 1997. *The ecological detective: confronting models with data*. Princeton University Press, Princeton, New Jersey, USA.
- Hobbs, N. T., and R. Hilborn. 2006. Alternatives to statistical hypothesis testing in ecology: a guide to self teaching. *Ecological Applications* **16**:5–19.
- Hoeting, J. A., R. A. Davis, A. Merton, and S. E. Thompson. 2006. Model selection for geostatistical models. *Ecological Applications* **16**:87–98.
- Jenkins, J. C., D. C. Chojnacky, L. S. Heath, and R. A. Birdsey. 2003. National-scale biomass estimators for United States tree species. *Forest Science* **49**:12–35.
- Johnson, J. B., and K. S. Omland. 2004. Model selection in ecology and evolution. *Trends in Ecology and Evolution* **19**:101–108.

- Jones, C. G., and J. H. Lawton, editors. 1994. Linking species and ecosystems. Chapman and Hall. New York, New York, USA.
- Judson, O. 1994. The rise of the individual-based model in ecology. *Trends in Ecology and Evolution* **9**:9–14.
- Kitajima, K., and C. K. Augspurger. 1989. Seed and seedling ecology of a monocarpic tropical tree, *Tachigalia versicolor*. *Ecology* **70**:1102–1114.
- Kobe, R. K., S. W. Pacala, J. A. Silander, and C. D. Canham. 1995. Juvenile tree survivorship as a component of shade tolerance. *Ecological Applications* **5**:517–532.
- Koenig, W. D., and J. M. H. Knops. 2000. Patterns of annual seed production by northern hemisphere trees: a global perspective. *American Naturalist* **155**:59–69.
- Latimer, A. M., S. Wu, A. E. Gelfand, and J. A. Silander, Jr. 2006. Building statistical models to analyze species distributions. *Ecological Applications* **16**:33–50.
- Law, R., and U. Dieckmann. 2000. A dynamical system for neighborhoods in plant communities. *Ecology* **81**:2137–2148.
- LePage, P. T., C. D. Canham, K. D. Coates, and P. Bartemucci. 2000. Seed abundance versus substrate limitation of seedling recruitment in northern temperate forests of British Columbia. *Canadian Journal of Forest Research* **30**:415–427.
- Levin, S. A., and R. T. Paine. 1974. Disturbance, patch formation, and community structure. *Proceedings of the National Academy of Sciences (USA)* **71**:2744–2747.
- Loreau, M. 2000. Biodiversity and ecosystem functioning: recent theoretical advances. *Oikos* **91**:3–17.
- Mooij, W. M., and D. L. DeAngelis. 1999a. Individual-based modelling as an integrative approach in theoretical and applied population dynamics and food web studies. Pages 551–575 in H. Olf, V. K. Brown, and R. H. Drent, editors. *Herbivores between plants and predators*. Blackwell, Oxford, UK.
- Mooij, W. M., and D. L. DeAngelis. 1999b. Error propagation in spatially explicit population models: a reassessment. *Conservation Biology* **13**:930–933.
- Muller-Landau, H. C., J. W. Dalling, K. E. Harms, S. J. Wright, R. Condit, S. P. Hubbell, and R. B. Foster. 2004. Seed dispersal and density-dependent seed and seedling mortality in *Trichilia tuberculata* and *Miconia argentea*. Pages 340–362 in E. C. Losos and E. G. Leigh, editors. *Diversity and dynamism: findings from a network of large-scale tropical forest plots*. Chicago University Press, Chicago, Illinois, USA.
- Nathan, R., and H. C. Mueller-Landau. 2000. Spatial patterns of seed dispersal, their determinants and consequences for recruitment. *Trends in Ecology and Evolution* **15**:278–285.
- Neuman, S. P. 2003. Maximum likelihood Bayesian averaging of uncertain model predictions. *Stochastic Environmental Research and Risk Assessment* **17**:291–305.
- 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**:1–43.
- Pastor, L., E. Kisdi, and G. Meszena. 2000. Jensen's inequality and optimal life history strategies in stochastic environments. *Trends in Ecology and Evolution* **15**:117–118.
- Post, W. M., and J. Pastor. 1996. Linkages—an individual-based forest ecosystem model. *Climatic Change* **34**:253–261.
- Ribbens, E., J. A. Silander, and S. W. Pacala. 1994. Seedling recruitment in forests—calibrating models to predict patterns of tree seedling dispersion. *Ecology* **75**:1794–1806.
- Rothe, A., and D. Binkley. 2001. Nutritional interactions in mixed species forests: a synthesis. *Canadian Journal of Forest Research* **31**:1855–1870.
- Royall, R. 1997. *Statistical evidence: a likelihood paradigm*. Monographs on Statistics and Applied Probability #71. Chapman and Hall, New York, New York, USA.
- Staelens, J., L. Nachtergale, S. Luysaert, and N. Lust. 2003. A model of wind-influenced leaf litterfall in a mixed hardwood forest. *Canadian Journal of Forest Research* **33**:201–209.
- Stoyon, D., and S. Wagner. 2001. Estimating the fruit dispersion of anemochorous trees. *Ecological Modelling* **145**:35–47.
- Turchin, P. 1998. *Quantitative analyses of movement*. Sinauer, Sunderland, Massachusetts, USA.
- Urban, D. L., H. H. Shugart, T. M. Smith, and G. B. Bonan. 1989. Spatial applications of gap models. *Forest Ecology and Management* **42**:95–110.
- Uriarte, M., R. Condit, C. D. Canham, and S. P. Hubbell. 2004. A spatially explicit model of sapling growth in a tropical forest: does the identity of neighbours matter? *Journal of Ecology* **92**:348–360.
- Wardle, D. A., G. M. Barker, K. I. Bonner, and K. S. Nicholson. 1998. Can comparative approaches based on plant ecophysiological traits predict the nature of biotic interactions and individual plant species effects in ecosystems? *Journal of Ecology* **86**:405–420.
- Watt, A. S. 1947. Pattern and process in the plant community. *Journal of Ecology* **35**:1–22.
- Wintle, B. A., M. A. McCarthy, C. T. Volinsky, and R. P. Kavanagh. 2003. The use of Bayesian model averaging to better represent uncertainty in ecological models. *Conservation Biology* **17**:1579–1590.

APPENDIX A

A table of the parameter values used to generate hypothetical data sets A–F (*Ecological Archives* A016-005-A1).

APPENDIX B

Results of a reanalysis of patterns of seedling dispersion for five tree species originally reported by LePage et al. (2000) (*Ecological Archives* A016-005-A2).

APPENDIX C

The estimated parameters for dispersal functions using different neighborhood search distances and functional forms of the dispersal kernel (*Ecological Archives* A016-005-A3).