Article:
Ellner, S.P. and Rees, M. (2006) Integral projection models for species with complex demography. American Naturalist, 167 (3). pp. 410-428. ISSN: 0003-0147

https://doi.org/10.1086/499438

Integral Projection Models for Species with Complex Demography

Stephen P. Ellner[1,*] **and Mark Rees**[2,†]

1. Department of Ecology and Evolutionary Biology, Cornell University, Ithaca, New York 24853;
2. Department of Plant and Animal Sciences, University of Sheffield, Sheffield S10 2TN, United Kingdom

Submitted January 21, 2005; Accepted October 27, 2005; Electronically published February 14, 2006

Online enhancements: appendixes, zip archive.

ABSTRACT: Matrix projection models occupy a central role in population and conservation biology. Matrix models divide a population into discrete classes, even if the structuring trait exhibits continuous variation (e.g., body size). The integral projection model (IPM) avoids discrete classes and potential artifacts from arbitrary class divisions, facilitates parsimonious modeling based on smooth relationships between individual state and demographic performance, and can be implemented with standard matrix software. Here, we extend the IPM to species with complex demographic attributes, including dormant and active life stages, cross-classification by several attributes (e.g., size, age, and condition), and changes between discrete and continuous structure over the life cycle. We present a general model encompassing these cases, numerical methods, and theoretical results, including stable population growth and sensitivity/elasticity analysis for density-independent models, local stability analysis in density-dependent models, and optimal/evolutionarily stable strategy life-history analysis. Our presentation centers on an IPM for the thistle *Onopordum illyricum* based on a 6-year field study. Flowering and death probabilities are size and age dependent, and individuals also vary in a latent attribute affecting survival, but a predictively accurate IPM is completely parameterized by fitting a few regression equations. The online edition of the *American Naturalist* includes a zip archive of R scripts illustrating our suggested methods.

Keywords: structured populations, integral model, matrix model, sensitivity analysis, latent variability, thistle.

* Corresponding author; e-mail: spe2@cornell.edu.

† E-mail: m.rees@sheffield.ac.uk.

Am. Nat. 2006. Vol. 167, pp. 410–428.

Matrix projection models are probably the most commonly used approach for modeling structured biological populations (Caswell 2001) and play a central role in population and conservation biology (e.g., Morris and Doak 2002). The popularity of matrix models is easy to understand. They are conceptually the simplest way to represent population structure, can be parameterized directly from observational data on the fate and reproductive output of individuals, and yield a great deal of useful information. The dominant eigenvalue λ of the projection matrix gives the population's projected long-term growth rate; the dominant right and left eigenvectors are, respectively, the stable stage distribution w and relative reproductive value v; and the eigenvectors determine the effect on λ of changes in individual matrix entries, which are often the key quantities for management applications. These and other metrics can be used as response variables to summarize population responses to changes in environmental conditions (Caswell 2001, chap. 10). Density dependence, stochasticity, and spatial structure can all be incorporated, and there is a growing body of theory for these situations (e.g., Tuljapurkar 1990; Cushing 1998; Caswell 2001; Tuljapurkar et al. 2003; Doak et al. 2005).

A matrix model divides the population into a set of classes or "stages," even when individuals are classified using a continuously varying trait such as body size. Indeed, the majority of empirical case studies reviewed by Caswell (2001) use size-based classifications rather than an actual discrete stage of the life cycle. In such cases, the definition of stages is to some degree arbitrary, giving rise to some potential problems, including the following: first, treating a range of heterogeneous individuals as a discrete stage inevitably creates some degree of error. Increasing the number of stages to minimize this problem leads to higher sampling error because fewer data are available on each stage. "Optimal" stage boundaries (Vandermeer 1978; Moloney 1986) minimize but cannot eliminate the resulting errors in population projection and can be difficult to implement (Pfister and Stevens 2003). Optimal boundaries for population projection may be poor for other purposes (Easterling et al. 2000). Second, the sensitivities

and elasticities are very sensitive to stage duration (Enright et al. 1995), affecting comparisons both within and between species. Easterling et al. (2000) proposed that these issues could be avoided by using a continuous individual-level state variable x. The population vector is replaced by a distribution function $n(x, t)$, where $n(x, t)dx$ is the number of individuals with their state variable in the range $[x, x + dx]$. The projection matrix **A** is replaced by a projection kernel $K(y, x) = P(y, x) + F(y, x)$, where P represents survival and growth from state x to state y and F represents the production of state y offspring by state x parents. The population dynamics are then

$$n(y, t + 1) = \int_L^U K(y, x)n(x, t)dx, \qquad (1)$$

where $[L, U]$ is the range of possible states. This is the continuous analogue of the matrix model $n_i(t + 1) = \sum_j a_{ij}n_j(t)$, where a_{ij} is the (i, j)th entry in the projection matrix **A**. Under similar assumptions to matrix models, the integral projection model (IPM) predicts a population growth rate λ with associated eigenvectors and state-dependent sensitivity and elasticity functions (Easterling 1998). An IPM is implemented on the computer as a matrix iteration, but this is just a technique for computing integrals and not a discretization of the life cycle.

Several empirical studies have illustrated how a kernel can be estimated from the same data as a matrix model (Easterling et al. 2000; Rees and Rose 2002; Childs et al. 2003, 2004; Rose et al. 2005). The functions making up the kernel can often be estimated by regression; for example, size-dependent survival and fecundity can be fitted by generalized linear or additive models and growth by parametric or nonparametric regression (Metcalf et al. 2003). Consequently, the appropriate model complexity for the available data can be identified using well-established statistical criteria and software rather than the typically ad hoc process of choosing the number of size classes and their boundaries.

The currently available general theory for IPMs (Easterling 1998; Easterling et al. 2000) applies only to models where individuals are characterized by one continuous quantity, but for many species, demographic rates are affected by multiple attributes. These might be observed attributes such as age and size (Rees et al. 1999; Rose et al. 2002) or age and sex (Coulson et al. 2001) or unobserved variables that reflect individual quality. For example, the probability of survival in the thistle *Onopordum illyricum* depends on size, age, and an unobserved measure of individual quality (Rees et al. 1999). In the kittiwake *Rissa tridactyla*, survival was age dependent, and individuals with high survival probability were also more likely to breed, presumably because of between-individual variation in quality (Cam et al. 2002). Failure to account for such latent between-individual differences can lead to systematic overestimation of population variability and extinction risk (Fox and Kendall 2002; Kendall and Fox 2002, 2003), underestimation of the uncertainty in population forecasts (Clark 2003), very large biases in estimates of demographic rates (Clark et al. 2003, 2004), and incorrect predictions of population responses to demographic perturbations (Benton et al. 2004).

When several variables are needed to predict demographic performance, estimating a matrix model becomes difficult because many between-class transitions have to be estimated (Law 1983; Caswell 2001). For example, "the construction of models using both size and age … may be impractical because of the large numbers of categories required" (Caswell 1988, p. 94). Consequently, matrix models with size × age classification have rarely been used, despite numerous studies documenting size- and age-dependent demography (Werner 1975; Gross 1981; Klinkhamer et al. 1987; van Groenendael and Slim 1988; McGraw 1989; Lei 1999; Rees et al. 1999; Rose et al. 2002). In contrast, a size- and age-dependent IPM for *Carlina vulgaris* required only one extra parameter to describe the effect of age on flowering probability (Childs et al. 2003).

In addition, many species have complex life cycles where individuals should be classified by different attributes at different points in the life cycle. For example, many plant populations have long-lived seed banks, so an additional discrete-state variable is required to keep track of seed numbers.

In this article, we generalize the IPM to accommodate species with complex demography, including complex life cycles and multiple attributes affecting individual performance. We begin by developing an IPM for the thistle *O. illyricum*, which has size- and age-dependent demography as well as substantial latent between-individual variation in survival unrelated to size or age. We then present a general IPM, explain how it can be implemented on a computer, and outline stable population theory for the density-independent model (λ, sensitivity/elasticity analysis, etc.) under assumptions very similar to those for matrix models. For density-dependent models, we give a criterion for local stability of a steady state population. The general theory and numerical methods are illustrated using the *Onopordum* IPM. But for the most part, this article's "Results" are "Methods," with detailed applications appearing elsewhere (e.g., Rees and Rose 2002; Childs et al. 2003, 2004; Rose et al. 2005). Our goal here is a unified development that makes IPMs a practical alternative to deterministic matrix models for structured

Table 1: Statistical models and parameter estimates describing the demography of *Onopordum illyricum*

Demographic process	Model
Growth	$\bar{y} = 3.24(.12) + .056(.02)x$, variance about the growth curve, $\sigma^2 = 42.47 \exp(-.71\bar{y})$, $n = 808$, $P < .0001$
Survival probability	$\text{Logit}(s) = -1.42(.21) + q + 1.08(.12)x - 1.09(.32)a$, $n = 1{,}397$, $P < .0001$
Flowering probability	$\text{Logit}(p_f) = -24.01(3.23) + 2.91(.43)x + .84(.31)a$, $n = 721$, $P < .0001$
Fecundity (seeds per flowering plant)	$f_n = \exp(-11.84(4.43) + 2.27(.60)x)$, $n = 49$, $P < .0001$
Probability of seedling establishment	$p_e = .025$ (density-independent model) or $p_e \propto 1/S_t$ (density-dependent model)
Distribution of seedling size	Gaussian with mean = 1.06, variance = 3.37, truncated at 0, $n = 389$
Distribution of seedling quality	Gaussian with mean 0 and standard deviation $\sigma_s = .82(.37)$

Note: The models are functions of log rosette area x, age a, and individual quality q; values in parentheses are standard errors of parameter estimates. The predicted values are the conditional mean $\bar{y}$ and variance σ^2 of log size next year given current size, the survival probability s, flowering probability p_f, and fecundity f_n.

populations with continuous trait variation. Subsequent articles will consider stochastic integral models and data-based models for species with complex life cycles.

To make the article more accessible, most technical details are in the appendixes. Appendix A covers computational methods and should be read before building an IPM from your own data. The mathematical level is the same as the main text, roughly that of Caswell (2001). Script files in R (R Core Development Team 2005) for the methods described in appendix A and for our *Onopordum* model are provided as a zip archive in the online edition of the *American Naturalist*. Appendixes B and C in the online edition of the *American Naturalist* are mathematical derivations, written for theoreticians.

Modeling Complex Demography in *Onopordum illyricum*

To motivate the general framework, we develop in this section an IPM for the thistle *O. illyricum* derived from a 6-year field study. The field study and data analysis summarized below, including model selection, are described in detail by Rees et al. (1999). This example illustrates how statistical analysis of demographic data translates directly into an integral model that involves far fewer fitted parameters than a conventional matrix model. In later sections, we use the model to demonstrate the mechanics of working with an integral model, with emphases on exploring the effects of latent heterogeneity in ways that would be difficult in a matrix model and on using the model for analyses of life-history evolution.

Field Study

Onopordum illyricum is a monocarpic perennial (reproduction is fatal) across its entire current range (Pettit et al. 1996). Reproduction occurs only by seed; these form a seed bank (up to 190 seeds m^{-2}), with a typical half-life of 2–3 years (Allan and Holst 1996). There were two study sites, but to simplify presentation, we focus on one, Plaine du Crau, an area of sheep-grazed semiarid steppe. Statistical analyses used data from both sites, with site effects fitted when significant.

Sampling ran from August 1987 to August 1992, which included the complete lifetime of the 1987 seedling cohort. Each plant's location and diameters (the longest and its perpendicular) were recorded in August, November, March, and May. We used the log-transformed maximum of the November, March, and May rosette areas as our measure of plant size. Additional visits were made each summer to collect all capitula on flowering plants within study quadrats. All apparently viable seeds were counted and then scattered randomly in their quadrat of origin.

Data Analysis

Results of the data analysis are summarized in table 1. Seedling size was well described by a normal distribution truncated at 0 (fig. 1*a*). Annual changes in plant size were size dependent ($P < .0001$) but not age dependent ($P > .08$) and were fitted by a linear model with size-dependent variance ($P < .0001$; see table 1; fig. 1*b*).

Survival probability was modeled as a mixed logistic regression with size and age as independent variables and a Gaussian distribution of individual intercepts. The main effects of size, age, and site were all highly significant ($P < .0001$), and there was significant between-individual heterogeneity ($P < .02$). Survival probability increases with plant size and decreases with age (fig. 1*c*). The standard deviation of the intercept distribution σ_s quantifies the between-individual variability. This variability may reflect differences in the local competitive environment, abiotic conditions, genetic differences, or other properties that remain constant over an individual's life.

Flowering probability was modeled by standard logistic

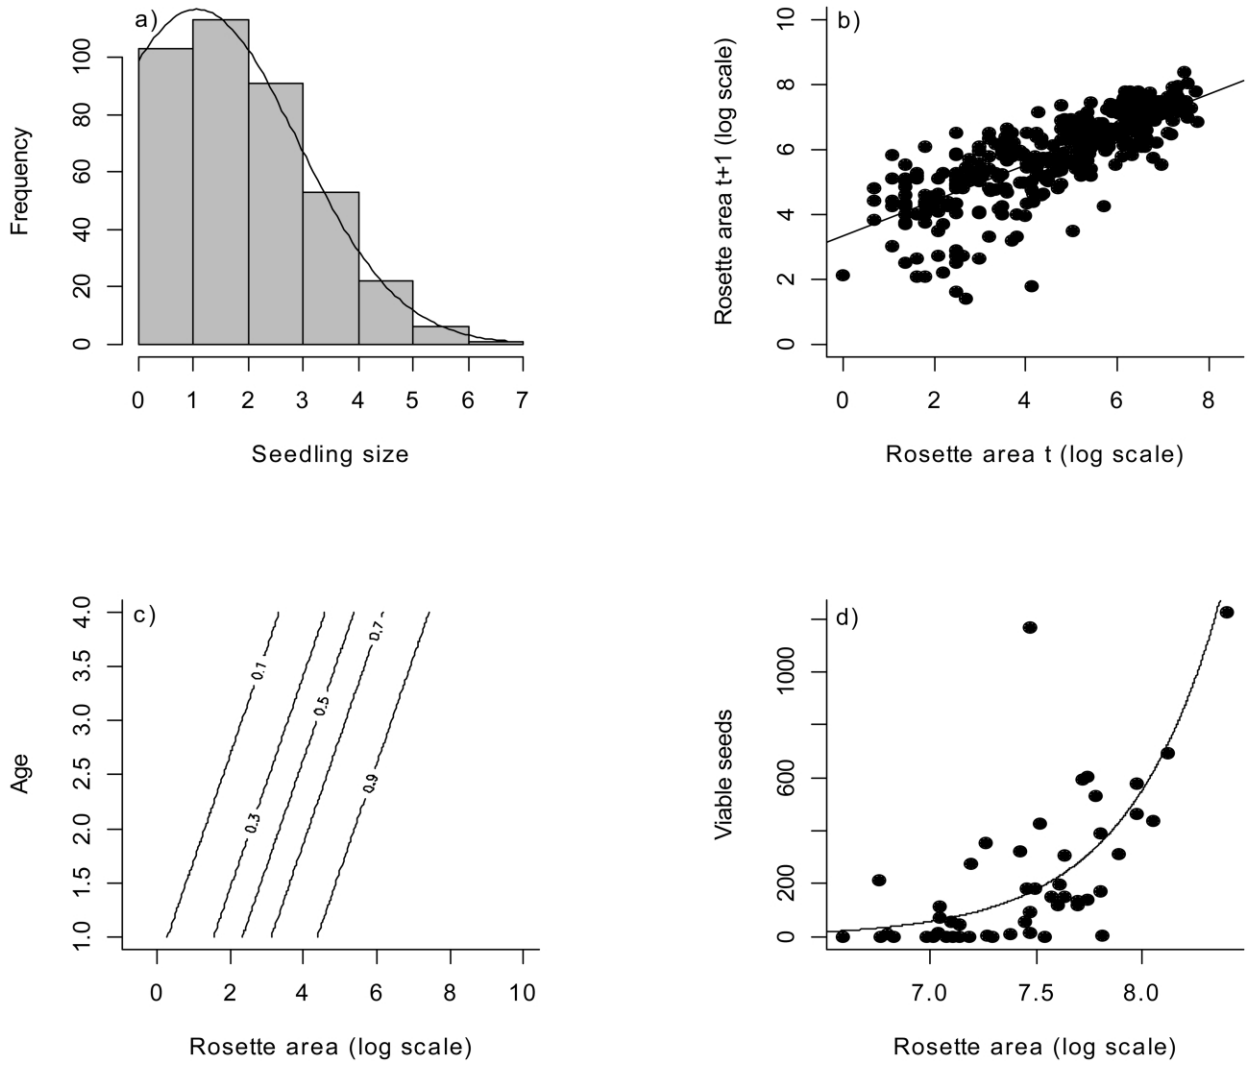

Figure 1: Demographic functions for *Onopordum*. *a*, Distribution of recruit sizes; the curve is the fitted truncated normal distribution. *b*, Growth relationship for plant size in successive years. *c*, Contour plot showing the dependence of survival on plant size and age. *d*, Fecundity (viable seed production) in relation to log-transformed rosette area.

regression because there was no evidence of between-individual variation ($P > .1$). Flowering probability increased with plant size ($P < .0001$) and age ($P < .008$), but site effects were not significant ($P > .1$). Seed production is strongly size dependent ($P < .0002$) and highly variable (fig. 1*d*).

Because of the seed bank, the probability of seedling establishment cannot be estimated by the ratio between recruitment and seed production. We therefore set the probability of seedling establishment in the density-independent IPM to match the observed rate of population increase ($\lambda = (155/140)^{1/4} = 1.026$). We chose not to model the seed bank because estimates of seed germination and death are not available; however, if estimates were available, it would be straightforward to add an additional discrete-state variable representing the number of seeds in the seed bank.

Intraspecific competition with neighbors had very little influence on growth and survival. In contrast, despite seed production being highly variable (0–2,750 per quadrat), the number of recruits was remarkably constant and independent of seed production the previous year ($P > .2$). Our density-dependent model therefore assumes that population growth is limited by microsite availability.

Integral Model for Onopordum

The fate of *Onopordum* plants is influenced by their size x, age a, and quality q (measured by survival intercept). Size is continuous whereas age and quality are discrete in our model. We treat quality as discrete for technical reasons (the conditional probability distribution for $q(t+1)$ given $q(t)$ is singular with respect to Lebesgue measure) that reflect our biological assumption that quality is constant. A dynamic quality variable is treated like size or any other dynamic continuous trait (for an example, see app. A).

To deal with multiple traits affecting demography, we generalize the basic model (1) as follows. First, there is a set of functions $n_{a,k}(x, t)$ that gives the distribution of size for individuals of age a ($a = 0$–7) and quality class k ($k = 1$–Q). Second, there is a set of survival-growth and fecundity kernel components that specifies the fate and fecundity of individuals of each possible age × quality (a, k) combination.

Survival takes individuals from population component (a, k) to $(a + 1, k)$. The survival-growth kernel for (a, k) individuals is derived from the probability of survival and the size distribution of survivors. In the notation of table 1,

$$P_{a,k}(y, x) = s(x, a, q_k)[1 - p_f(x, a)]g(y, x). \qquad (2)$$

That is, to reach size y from size x, the individual must survive, not flower, and make the $x \to y$ size transition. Flowering probability is a factor in equation (2) because flowering is fatal. Formulas for the probabilities of survival and flowering are given in table 1, and the growth kernel $g(y, x)$ is given by the conditional size distribution from table 1, where y is approximately normal, with mean $\bar{y} = 3.24 + 0.56x$ and variance $\sigma_y^2 = 42.74 \exp(-0.71\bar{y})$.

Births go from each component (a, k) to all components $(0, j)$. Seedling size and quality are assigned independently (and independent of parent age, size, and quality) according to the distributions given in table 1. The fecundity kernels are therefore $F_{a,k,j}(y, x) = \alpha_j \varphi_0(y) S(x, a, k)$, where α_j is the fraction of seedlings in quality class j, φ_0 is the probability density of seedling size, and $S(x, a, k)$ is the number of seedlings in year $t + 1$ per parent of age a, size x, and quality class k in year t. Our model breaks down seedling production into survival of the parent, flowering of the parent, seed production, and probability of establishment as a seedling. So, in terms of the demographic models in table 1, we have $S(x, a, k) = p_e f_n(x) p_f(x, a) s(x, a, q_k)$.

Having specified the survival-growth and fecundity kernels, the model is now complete. Note that the kernel of the integral model is implied directly by the statistical analysis of the data. The model expresses the population-level consequences of individual demography without any additional assumptions or approximations.

Why Not Use a Matrix Model?

For illustration, we have chosen a species with several features that the basic IPM cannot accommodate, but the model is still defined by a small number of conventional regression models with a total of 17 fitted parameters. In contrast, a conventional matrix model would have an enormous number of parameters (matrix entries) to estimate. At a bare minimum, we might use four size classes and a maximum age of 4. Then, with Q quality classes, there would be $4Q$ matrices of size 4×4 for survival and growth transitions and fecundity parameters for ages 2–4 (total fecundity and the offspring size × quality distribution). The total parameter count is $65Q + 14$. For accurate solution of the integral model, we have used 40 quality classes, which would entail more than 2,500 parameters for a matrix model, but even with three quality classes, the matrix model would have more than 200 parameters. Moreover, if individual quality is dynamic, the number of parameters in a matrix model would be vastly higher (proportional to Q^2), while an IPM would typically require only a few extra parameters to describe quality dynamics, as in the size × quality model described in appendix A.

This enormous difference in parameter count occurs whenever individual performance is affected by multiple state variables. Even with our large data set (1,402 observations on 1,144 individuals), we could not hope to estimate accurately all 192 survival entries in a projection matrix with three quality classes. For the integral model, we have a generalized linear mixed model with four fitted parameters and the option of fitting a nonlinear model if it were needed (e.g., Wood 2004, 2005). The relative ease and parsimony of model parameterization for IPMs is one of the main points of this empirical application. Continuous regression models can be used to parameterize matrix models (Morris and Doak 2002), followed by some way of averaging across a stage class, for example, to estimate the average fecundity of medium-size individuals. But if the actual response function is smooth, then division into stage classes necessarily distorts the functional relationship—indeed, any supposed average fecundity of medium-size individuals is not well defined because the average depends on the within-stage size distribution, which varies over time. By retaining a continuous-trait variable, the individual-level description in terms of smoothly varying effects of continuous-trait variation is translated exactly into the implied population-level dynamics.

General Integral Model

We now describe a general integral model that can accommodate species such as *Onopordum* with complex demography, including species with complex life cycles and multiple traits (discrete or continuous) affecting demographic performance. We then describe the model's general properties and how it can be implemented on a computer, returning to *Onopordum* to illustrate practical application of the model.

The space of individual states $\mathbf{X}$ can include a set of discrete points $\mathbf{D} = \{x_1, \ldots, x_D\}$ and a set of continuous domains $\mathbf{C} = \{\Omega_{D+1}, \Omega_{D+2}, \ldots, \Omega_{D+C}\}$. Each continuous domain is either a closed interval or a closed rectangle in d-dimensional space. Each set in $\mathbf{D}$ or $\mathbf{C}$ will be called a component and denoted as Ω_j, $j = 1, 2, \ldots, N = D + C$. For example, Ω_1 and Ω_2 might represent different genders or a pair of discrete morphs in a species with phenotypic plasticity, within which individuals are classified by size or weight. The state of the population is described by a function $n(x, t) \geq 0$ that gives the distribution of individual states x at time t; note that this function also consists of the following components: discrete values $n_j = n(x_j)$, $j = 1, 2, \ldots, D$ and continuous functions $n_j(x)$, $j = D + 1, 2, \ldots, N$.

To describe transitions within and among components, there is a set of kernel components $K_{ij}(y, x)$, $1 \leq i, j \leq N$, with $K_{ij} \neq 0$ whenever individuals in Ω_j contribute to next year's population in Ω_i. In terms of kernel components, the general deterministic model is

$$n_i(y, t + 1) = \sum_{j=1}^{N} \int_{\Omega_j} K_{ij}(y, x) n_j(x, t) dx. \qquad (3)$$

There are four possible types of kernel components: discrete-to-discrete ($i, j \leq D$): $K_{ij}(y, x) = k_{ij}$, a number (e.g., the number of two-leaf seedlings next year per one-leaf seedlings this year); discrete-to-continuous ($j \leq D$, $i > D$): $K_{ij}(y, x) = k_{ij}(y)$ (e.g., the size distribution of recruits produced by seeds that germinate); continuous-to-discrete ($j > D$, $i \leq D$): $K_{ij}(y, x) = k_{ij}(x)$, the state-dependent contribution to discrete component i (e.g., the number of resting eggs produced by *Daphnia* of age j and size x); and continuous-to-continuous ($i, j > D$): $K_{ij}(y, x)$ is a genuinely bivariate function giving the contribution of state x individuals in component j to state y individuals in component i.

Kernel components can be separated into survival-growth and fecundity contributions, $K_{ij} = P_{ij} + F_{ij}$. We assume that all P_{ij} and F_{ij} are continuous functions so that the kernel is continuous. The model as just described is density independent; the density-dependent model is the same except that K can depend on the population state, $K = K(y, x, n)$, or on some measure $\mathbf{N}$ of total population density, $K = K(y, x, \mathbf{N})$.

In our *Onopordum* model, the domains are all intervals representing the range of possible sizes, one for each age × quality class combination. All kernel components are therefore of the continuous-to-continuous type. To incorporate the seed bank (as in M. Rees, D. Z. Childs, J. C. Metcalf, K. E. Rose, A. Sheppard, and P. J. Grubb, unpublished manuscript), we would add a discrete component for the number of seeds, a series of continuous-to-discrete fecundity kernels for seed production by each age × quality combination, a series of discrete-to-continuous kernels for emergence of seeds into each quality category at age 0, and a discrete-to-discrete kernel (a number) for seeds remaining in dormancy and surviving.

Our only restrictive assumption is that continuous domains are bounded. For simplicity, we have assumed rectangular domains, but the theory applies so long as each continuous domain is a bounded and closed subset of Euclidean space. Unbounded domains, however, lead to technical complications (see app. C) that we believe should be avoided. An integral model with bounded domains is a natural generalization of matrix models and shares many of their properties (the domain of a matrix model, a finite set of points, is always bounded in the sense used here, though it can sometimes represent an unbounded set of individual states, e.g., a single matrix stage class or IPM domain for individuals of age A or higher). An integral model with unbounded domains can behave very differently, for example, allowing a population to spread forever in "trait space" without ever reaching a stable distribution (see app. B).

There are two ways to make a bounded integral model. First, the modeler can specify a finite range of possible values for all individual attributes. For example, each variable can be truncated at several standard deviations beyond the range of observed values, with the kernel set to 0 outside those limits. Suitable limits can also be set by expanding the range until further increases have no impact on model predictions. Individuals far beyond the range of the data are a fiction that results from using an unbounded statistical distribution to model a finite data set. This is harmless for short-term prediction but can lead to unrealistic long-term behavior. Second, unbounded attributes can be transformed onto a bounded domain using, say, the logistic transformation $x \to e^x/(1 + e^x)$. Whether this produces a model satisfying our assumptions depends on how the kernel is defined outside the range of the data (see app. C). Either way, bounded components result when the model is not allowed to produce individuals very different from those actually observed. The first approach is much simpler and is therefore recommended.

Implementing a General Integral Model

Evaluating Integrals

Equation (3) shows that each iteration of an IPM consists of one or more integrals that must be evaluated numerically. We recommend doing this using a simple method called the midpoint rule. To explain the midpoint rule, consider the basic model (1). We define mesh points x_i by dividing the interval $[L, U]$ evenly into m size classes and setting x_i at the midpoint of the ith class:

$$x_i = L + (i - 0.5)h,$$
$$i = 1, 2, \ldots, m, \tag{4}$$

where $h = (U - L)/m$. The midpoint rule approximation to equation (1) is then

$$n(x_j, t + 1) = h\sum_{i=1}^{m} K(x_j, x_i)n(x_i, t). \tag{5}$$

This is a matrix multiplication

$$\mathbf{n}(t + 1) = \mathbf{Kn}(t), \tag{6}$$

where $\mathbf{K}$ is the matrix whose (i, j)th entry is $hK(x_i, x_j)$ and $\mathbf{n}(t)$ is the vector whose ith entry is $n(x_i, t)$. The same idea can be used to approximate higher-dimensional integrals for models with multiple state variables by defining mesh points for each variable. The result is again a matrix iteration for the population distribution at the mesh points.

The accuracy of the midpoint rule depends on the number of mesh points m. Determining m is a trade-off between accuracy and computational cost, and in practice one should explore a range of mesh sizes to ensure that the population growth rate and other quantities of interest are calculated accurately. In appendix A, we suggest methods for implementing integral models when individuals are cross-classified by age and size or by size and quality.

Computing λ, w, *and* v

The usual procedure for a matrix model is to compute the complete set of eigenvalues and eigenvectors and then find the dominant pair. For the large matrices representing a complex integral model, it is much more efficient to compute only the dominant pair by iterating the model. Let $\mathbf{n}(t)$ denote the population state in generation t—either one vector or the set of vectors for each component of $\mathbf{X}$. Choose any nonzero initial distribution $\mathbf{n}(0) = \mathbf{n}_0$, and let $\mathbf{u}_0 = \mathbf{n}_0/\|\mathbf{n}_0\|$, where $\|\mathbf{x}\|$ is the sum of all entries in the vector. The iteration for the population growth rate λ_t and population structure $\mathbf{u}(t) = \mathbf{n}(t)/\|\mathbf{n}(t)\|$ is then

$$\mathbf{u}(t + 1) = \frac{\mathbf{Ku}(t)}{\|\mathbf{Ku}(t)\|},$$
$$\lambda_t = \|\mathbf{Ku}(t)\|, \tag{7}$$

where $\mathbf{K}$ is the matrix used to iterate the model numerically (e.g., eq. [6]). Iterating equation (7) until it converges gives λ and a state vector $\mathbf{u}$ whose entries are proportional to the stable state distribution function $w(x)$ (Isaacson and Keller 1966).

The dominant left eigenvector v, representing state-dependent reproductive value, is the dominant right eigenvalue for the transpose kernel $K^T(y, x) \equiv K(x, y)$ and can be obtained by iterating equation (7) with the transpose kernel. Some suggestions for implementing transpose iteration are in appendix A.

Stable Population Growth: Assumptions and Their Meaning

General Theory

Stable population growth refers to properties centered on the existence of a unique stable population distribution and asymptotic growth rate, to which a density-independent population converges from any initial composition:

$$\lim_{t\to\infty} \frac{n(x, t)}{\lambda^t} = Cw(x). \tag{8}$$

Here, λ and w are the dominant eigenvalue and eigenvector for the kernel, respectively, and C is a constant depending on the initial population. Thus, λ is the long-term population growth rate, and w is the stable state distribution, with $\lambda > 0$ and $w(x) \geq 0$. In this section, we describe two conditions that imply stable population growth for our general integral model, and in "Stable Population Growth: Results" we state the conclusions. Proofs and additional discussion of our assumptions are in appendixes B and C.

For matrix models, equation (8) is guaranteed to occur if the projection matrix $\mathbf{A}$ is power positive: all entries of $\mathbf{A}^m$ are positive for some $m > 0$ (Caswell [2001, sec. 4.5] calls such a matrix "primitive"). Note that power positivity is possible only if postreproductive stages are removed from the model. The analogous condition implying stable population growth in our general IPM is that for some $m > 0$,

$$K^{(m)}(y, x) > 0 \tag{9}$$

for all x, y in $\mathbf{X}$, where $K^{(1)} = K$ and $K^{(m)}$ is the m-step-ahead projection kernel defined by the Chapman-Kolmogorov formula:

$$K^{(t+1)}(y, x) = \int_{\mathbf{X}} K(y, z)K^{(t)}(z, x)dz. \quad (10)$$

Equation (10) is the continuous equivalent of matrix multiplication, so $K^{(m)}$ is analogous to the mth power of a projection matrix. We therefore refer to property (9) as power positivity of the kernel. In a matrix model, all components of **X** are discrete points, so property (9) is exactly equivalent to the projection matrix being power positive.

A second condition that also guarantees stable population growth in our model is mixing at birth, meaning that the relative frequency of offspring states is similar for all parents. This condition is likely to hold in many species. For example, although many plant species have great plasticity in the number of seeds produced, there is much less plasticity in seed size or quality. In addition, maternal environment effects are often small compared with the effect of the environment in which a seedling grows (Weiner et al. 1997), which means that the distribution of offspring size will be similar for all parents, provided the population is censused sometime after recruitment.

Technically, suppose that for a parent with state x, the fecundity kernel satisfies $A(x)\varphi_0(y) \le F(y, x) \le B(x)\varphi_0(y)$, where $A, B > 0$ and φ_0 is a probability distribution. Then, if there is a finite maximum age for reproduction, we can construct a Leslie matrix **L** from the mean age-specific survival and fecundity of a cohort of newborns with state distribution $\varphi_0(y)$. In appendix B, we show that if there is mixing at birth and **L** is power positive, then some iterate of the kernel has a property called *u*-boundedness (Krasnosel'skij et al. 1989), which implies stable population growth (app. C). This continues to hold without a finite maximum age if senescence is not too slow—specifically, if the survivors from a cohort of newborns are eventually outnumbered by descendants of the same cohort at all states in **X** (we call this uniform senescence; see app. B for the formal definition).

What about My Model?

An IPM with stable population growth behaves very much like a power-positive matrix model. If it behaves differently, the cause is probably biological rather than mathematical, such as inclusion of postreproductive individuals or semelparous reproduction at a fixed age. The model's behavior will then be totally different from equation (8), so it is easy to check computationally whether your IPM has stable growth. If the model is implemented as a single large matrix (as we suggest in app. A for populations structured by size and additional continuous traits), the check is to verify that in some high power of the matrix (e.g., representing 100 times the average life span), the columns are all multiples of one another, indicating convergence to that population structure from any initial distribution. High matrix powers can be computed efficiently using $A^4 = A^2 \times A^2$, $A^8 = A^4 \times A^4$, and so forth. If the model is implemented via separate matrices for each kernel component (as we suggest in app. A for populations with age structure), the same check can be done indirectly by computing the population state after many generations from many different initial conditions.

Stable Population Growth: Results

Under the assumptions described in "Stable Population Growth: Assumptions and Their Meaning," the long-term behavior of the integral model is identical to that of a power-positive matrix model. These results are derived in appendix C, which is based on work by Easterling (1998) and the theory of positive linear operators (Krasnosel'skij et al. 1989).

First, for any nonzero initial population $n(x, 0)$, the long-term dynamics are given by equation (8). Corresponding to λ, there is also a left eigenvector v that gives the state dependence of relative reproductive value, that is, the relative contribution to long-term population growth as a function of individual state. The convergence in equation (8) is asymptotically exponential, with the error decreasing in proportion to $(|\lambda_2|/\lambda)^t$ once initial transients have died out, where $|\lambda_2|$ is the maximum magnitude of nondominant eigenvalues.

Second, following the development for matrix models (Cushing 1998, chap. 1; Caswell 2001, sec. 5.3.4), we can calculate the net reproductive rate R_0, which is the long-term generation-to-generation population growth rate. That is, if g_0 is the total current population (generation 0), g_1 is their total number of offspring counted at birth (generation 1), g_2 is the total number of offspring of generation 1 individuals, and so on, then

$$\lim_{k\to\infty} \frac{g_k}{R_0^k} = G, \quad (11)$$

with the value of G depending on the initial population distribution (see app. B). As in matrix models, R_0 and λ are related: $\lambda - 1$ and $R_0 - 1$ have the same sign, and R_0 is also important for analysis of evolutionarily stable strategies (ESS) in density-dependent models. We use it below to identify ESS flowering strategies in the *Onopordum* model.

Sensitivity analysis is also nearly identical to the matrix case, under the appropriate definitions. The question is, How much does λ change when we perturb $K(y_0, x_0)$ for some $x_0, y_0 \in \mathbf{X}$? With continuous-state variables, it is not meaningful to perturb K at a single point. Instead, we

consider a small perturbation near (y_0, x_0), replacing $K(y, x)$ with $K(y, x) + \varepsilon f_r(y, x|y_0, x_0)$, where f_r is an approximate δ function—a smooth nonnegative function that is 0 if the distance from (y_0, x_0) to (y, x) is larger than r and such that $\int_{X\times X} f_r(y, x|y_0, x_0)dydx = 1$. Let $\lambda(\varepsilon|y_0, x_0)$ denote the resulting dominant eigenvalue. The sensitivity function is then

$$s(y_0, x_0) = \lim_{\varepsilon, r\to 0} \frac{\lambda(\varepsilon|y_0, x_0) - \lambda}{\varepsilon}. \quad (12)$$

This satisfies the familiar sensitivity formula

$$s(y, x) = \frac{v(y)w(x)}{\langle v, w\rangle}, \quad (13)$$

where $\langle v, w\rangle$ is the inner product $\langle v, w\rangle = \int_X v(x)w(x)dx$. The elasticity function is

$$e(y_0, x_0) = \frac{K(y_0, x_0)s(y_0, x_0)}{\lambda}. \quad (14)$$

The elasticity function integrates to 1, so $e(y_0, x_0)$ can be interpreted as the proportional contribution of $K(y_0, x_0)$ to population growth, just like elasticities for a matrix model.

There are two approaches to computing the sensitivity and elasticity functions: first, once w and v have been computed, equation (13) gives s at all mesh points and, second, directly use the definition, equation (12), by perturbing K and recomputing λ by iteration. This will be slower and less accurate than equation (13) but may be preferable if the model structure makes it hard to compute v. In appendix A, we suggest numerical methods for computing sensitivities efficiently by perturbation.

Stability Analysis for Density-Dependent Models

In this section, we state a criterion for local stability of equilibria in models with a density-dependent kernel, which is derived in appendix B. For simplicity (following Caswell 2001, chap. 16), we assume that the density-dependent kernel has the form $K(y, x, \mathbf{N})$, where $\mathbf{N}$ is a weighted total population size,

$$\mathbf{N}(t) = \int_X W(x)n(x, t)dx, \quad (15)$$

for some weighting function $W \geq 0$. Let $\bar{n}(x)$ denote an equilibrium population distribution and $z(x)$ a small perturbation. Then, starting from $\bar{n} + z$ in year t, the population in year $t + 1$ is $\bar{n} + z'$, where

$$z'(y) = \int_X J(y, x, \overline{\mathbf{N}})z(x)dx + O(z^2). \quad (16)$$

Here $\overline{\mathbf{N}}$ is given by equation (15), with $n = \bar{n}$, and J is the Jacobian kernel given by

$$J(y, x, \overline{\mathbf{N}}) = K(y, x, \overline{\mathbf{N}}) + Q(y)W(x),$$
$$Q(y) = \int_X \frac{\partial K}{\partial \mathbf{N}}(y, x, \overline{\mathbf{N}})\bar{n}(x)dx, \quad (17)$$

where J is analogous to the Jacobian matrix for density-dependent matrix models (Caswell 2001, sec. 16.4). So long as J is continuous, the conclusion from equation (16) is that $\bar{n}(x)$ is locally stable if the dominant eigenvalue of $J(y, x, \overline{\mathbf{N}})$ is <1 in magnitude.

The stability criterion would generally be applied numerically, by computing J and its dominant eigenvalue. However, in appendix B, we prove that a class of models with density-dependent fecundity, which includes our density-dependent *Onopordum* model, has at most one positive equilibrium, which is locally stable whenever it exists.

Onopordum Model Results

We now apply the general theory to analyze the IPM for *Onopordum*. Some additional results on evolutionary analysis—characterization of optimal and ESS life histories—are also presented and applied.

The model has a finite maximum age, and the distribution of offspring states is independent of parent state, so the density-independent model satisfies the mixing at birth assumption and therefore has a unique dominant eigenvalue and associated eigenvectors. The model's stable distribution provides an accurate description of the bimodal distribution of sizes observed in the population (fig. 2*a*), predicts that the population will be dominated by new recruits (fig. 2*b*), and shows that older individuals have larger survival intercepts, on average, relative to new recruits (fig. 2*c*). This shift reflects the increased survival of individuals with larger survival intercepts. Reproductive value increases monotonically with size, reflecting the fact that the probability of flowering and seed production increases as plants become larger (fig. 2*d*). Reproductive value decreases with age (fig. 2*e*) because mortality increases as plants grow older and increases with individual survival intercept (fig. 2*f*); the age dependence of reproductive value is not an artifact of the finite maximum age imposed in the model, as it is unaffected by increasing the maximum age.

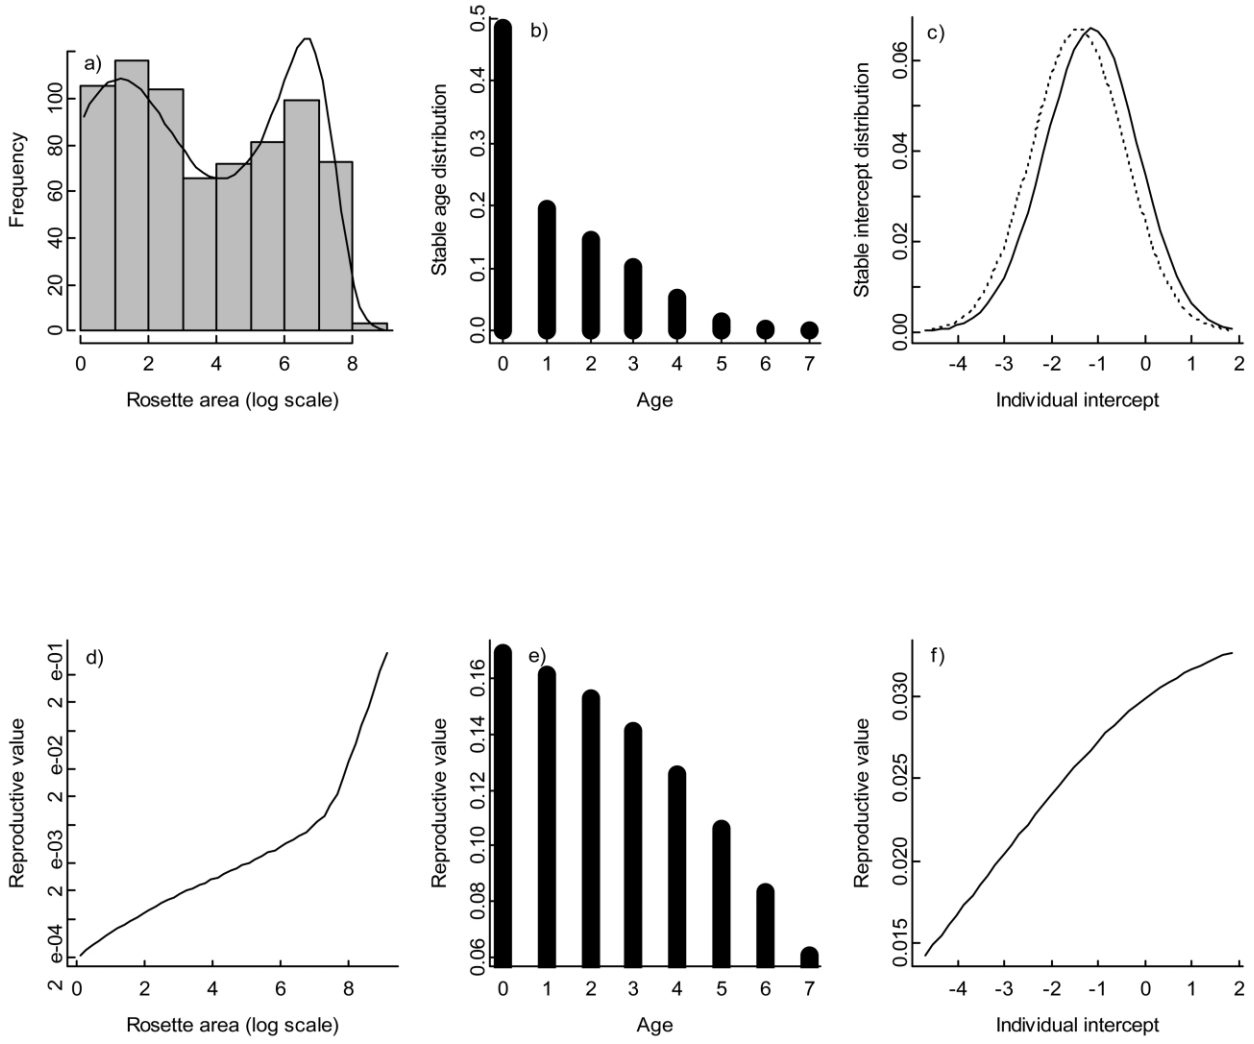

Figure 2: Predicted stable state distribution for (*a*) size (*line* = integral projection model prediction; *histogram* = observed size distribution), (*b*) age, and (*c*) individual survival intercept (*dotted line* = distribution of intercepts in new recruits). Predicted relative total reproductive value as a function of (*d*) size, (*e*) age, and (*f*) individual survival intercept.

Elasticity Analysis

Using equation (14), we can compute the elasticities of λ, and these can be partitioned into contributions from the survival-growth and reproduction components of the kernel (though these are interdependent because the survival and flowering functions contribute to both kernel components). In this partitioning, the survival-growth component has a larger influence than the reproduction component (0.75 vs. 0.25, respectively).

The elasticities of the survival-growth component of the kernel can be partitioned into contributions from plants of different ages, sizes, and survival intercepts (fig. 3). Younger plants make the largest contribution to λ (fig. 3*a*) because they represent a larger proportion of the stable age distribution (fig. 2*a*). The distribution of elasticities for the survival intercepts is shifted toward larger intercepts, relative to distribution of intercepts in recruits (fig. 3*b*). Again, this is a consequence of individuals with larger survival intercepts being overrepresented in the stable distribution (fig. 2*b*). This effect is tempered by the fact that few individuals with larger survival are produced (fig. 2*c*). The contribution of growth to λ is dominated by transitions into the larger size range where reproduction occurs (fig. 3*c*).

For the reproduction component of the kernel, age 3 plants make the greatest contribution to λ (fig. 4*a*). This

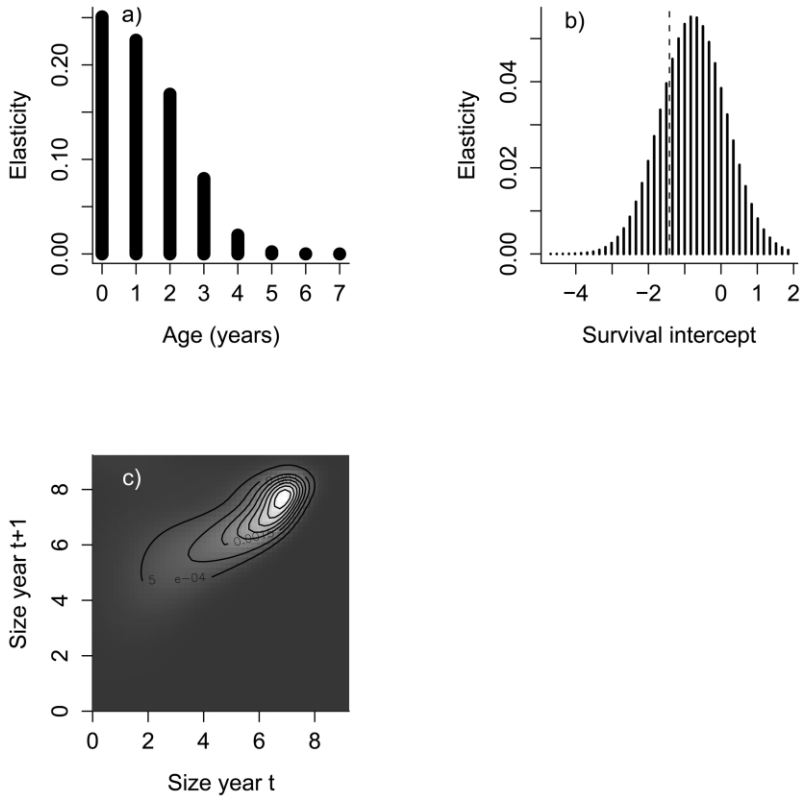

Figure 3: Survival-growth elasticities for the *Onopordum* integral projection model partitioned according to (*a*) age, (*b*) survival intercept (*dashed vertical line* = mean survival intercept in recruits), and (*c*) size transitions from one year to the next.

is because plants are, on average, larger as they get older and because fecundity increases with size; however, this effect is counteracted by older plants that make up a smaller proportion of the stable age distribution (fig. 2*b*). The survival intercept elasticities, for both adults and recruits, are identical and shifted toward larger intercepts (fig. 4*b*, 4*d*). These elasticities are identical because the distribution of survival intercepts in recruits is independent of adult survival intercept. The contributions of different size transitions to λ are dominated by movement of individuals from large sizes to recruits (fig. 4*c*), which is a consequence of larger individuals having higher probabilities of flowering and producing more seeds.

To model density-dependent recruitment, we assumed that in each year, the probability of establishment was equal to the observed average number of recruits (2.25 m^{-2}) divided by total seed production (m^{-2}). Iteration of the resulting model shows smooth convergence to a stable equilibrium density of 4.8 plants m^{-2}, in good agreement with the average density recorded in the field, 4.4 plants m^{-2} (Rees et al. 1999). The density-dependent kernel has the form studied analytically in appendix B (eq. [B12]), so our analytical results confirm the numerical observation of a unique stable equilibrium.

The *Onopordum* model illustrates how latent variability between individuals can be parsimoniously modeled using an IPM. To explore this further, we looked at the effects of ignoring individual variability when fitting the survival model and of varying the distribution of survival intercepts while leaving the other parameters of the mixed model fixed. Varying σ_s, the standard deviation of survival intercepts, results in a rapid increase in λ at low levels of variability, with the relationship reaching an asymptote at $\sigma_s \approx 1$ (fig. 5*a*). Ignoring individual effects when fitting the survival model results in λ being slightly overestimated (1.036 vs. 1.026). The effects of varying σ_s on the equilibrium population size are relatively small except at very low levels of variability when the population is no longer persistent ($\lambda < 1$), and in this case, the error resulting from ignoring individual effects when fitting the survival model is slight (fig. 5*b*). In contrast, changes in mean offspring quality have dramatic effects (fig. 5*c*, 5*d*), and a relatively small decrease in mean quality is predicted to render the population unviable. On the likely assumption that quality

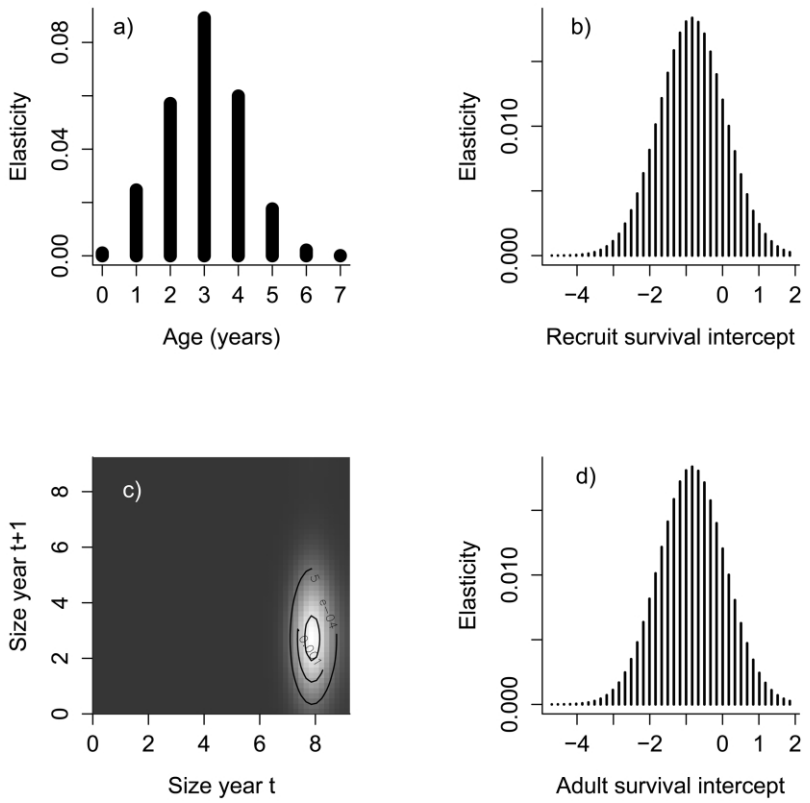

Figure 4: Fecundity elasticities for the *Onopordum* integral projection model partitioned according to (*a*) age, (*b*) recruit survival intercept, (*c*) size transitions from one year to the next, and (*d*) adult survival intercept.

variation is largely a result of microsite variability, these calculations address the impact of habitat modifications that would affect individuals' complete life histories rather than a specific set of transitions. In particular, they suggest a large impact of relatively small changes in the mean quality of available microsites for seedling establishment.

Evolutionary Analysis

The field data on *Onopordum* were originally used to construct an individual-based model in order to study the evolution of the flowering strategy as a function of age and size (Rees et al. 1999). At that time, no analytic framework was available that could incorporate demography depending on size, age, and individual quality. Consequently, ESS flowering strategies had to be identified using an individual-based model simulating competition among genotypes.

With an IPM, we can instead use familiar analytic and numerical methods for structured population models. In a density-independent model, the optimal life history maximizes λ and so can be found by numerically optimizing λ as a function of the model parameters. In our density-dependent model, density dependence acts only on seedling establishment. ESS life histories are therefore characterized by maximization of the net reproductive rate R_0 when offspring are counted before the impact of density dependence (app. B). Because of the mixing at birth in our model, R_0 equals the average per capita lifetime seed production of a cohort of newborns (app. B). We compute this by starting the population as a cohort of newborns and summing up their seed production until all in the founding cohort have died.

ESS life histories in the *Onopordum* model can then be computed by numerical optimization of R_0, which is far quicker than individual-based simulation of the evolutionary process. This is not just a convenience—it allows far more intensive study of the model.

As an example, we revisit the comparison by Rees et al. (1999) between observed and predicted ESS flowering strategies. The flowering strategy is defined by three parameters: intercept, slope for age effect, and slope for size effect. As in the work by Rees et al. (1999), when all three parameters are allowed to vary, the predicted strategy is a

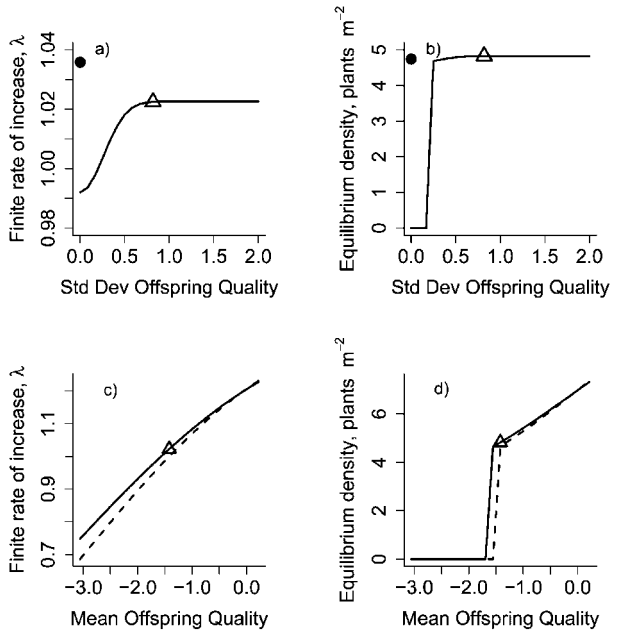

Figure 5: Predicted effects of changing the distribution of offspring quality on the finite rate of increase λ and equilibrium population size. *a*, *b*, Effects of changing the standard deviation of the distribution of survival intercepts σ_s. Solid dots show predictions obtained by ignoring individual effects when fitting the survival model. *c*, *d*, Effects of changing the mean of the distribution of survival intercepts by plus or minus twice the standard deviation of survival intercepts. The solid line shows the estimated value of σ_s, the dashed line shows σ_s reduced by a factor of 4; increases in σ_s above the estimated value have no visible effect. In all plots, the triangles are plotted at the estimated parameters of the offspring distribution.

sharp threshold: all plants of a given age should flower with probability 0 or 1, depending on whether their size is below or above a threshold that decreases with age. The observed size dependence of flowering is gradual, representing possibly a constraint or else a decision that depends on plant size at some time between censuses (Rees et al. 1999). We therefore imposed gradual size dependence by holding the size slope fixed at its estimated value and optimizing the intercept and age slope (see Childs et al. 2003).

Rees et al. (1999) compared estimated and ESS flowering strategies and found that plants following the ESS strategy would flower later in life and at larger sizes (on average) than real plants. In contrast, the ESS flowering strategy for a stochastic model with year-to-year variation in model parameters was very close to that observed. Rees et al. (1999) concluded that the observed flowering strategy is shaped by environmental variability. However, the comparison for the deterministic model ignored the considerable uncertainty in both the empirical estimates and the ESS. The estimated flowering parameters have substantial standard errors (14% and 37% of the point estimate for age and size slopes, respectively), and the predicted ESS depends on the estimated values of all other model parameters.

The characterization of ESSs in terms of R_0 makes it straightforward to quantify the uncertainty in the ESS. We first bootstrapped the original data to compute 1,000 bootstrap estimates of all model parameters (for the survival model, we bootstrapped at the level of individual; for other model components, there is no evidence of heterogeneity between individuals, so we bootstrapped at the level of observations). Then, to find the corresponding ESSs, we computed R_0 and used the Nelder-Mead simplex algorithm to find the ESS flowering parameters for each of the bootstrap parameter sets.

The results (fig. 6) show that there is a considerable difference between the observed flowering strategy and the predicted ESS, with the ESS having a lower flowering probability for all plant ages and sizes observed in the field. The delay in flowering by the ESS results in a difference of about 9 months in the predicted mean age of flowering plants. This difference has relatively little effect on the mean size of flowering plants because the actual mean size is near the fixed point of the mean growth curve (fig. 1*b*).

To test whether these differences are statistically signif-

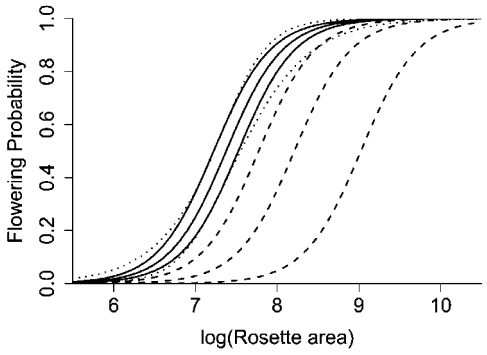

Figure 6: Uncertainty distributions for actual and predicted evolutionarily stable strategy (ESS) relationship between rosette area (on log scale) and flowering probability. The solid and dotted curves show the estimated empirical relationship between size and flowering probability for an individual of age 3, surrounded by pointwise 2.5 and 97.5 percentiles. The solid percentile curves are based on 1,000 bootstrap parameter vectors, with the slope parameter for the effect of size on flowering held constant. The dotted percentile curves are based on 5,000 draws of flowering model parameters from the multivariate normal distribution implied by the variance-covariance matrix of parameter estimates for the fitted model (given in table 1) without the constraint on the slope parameter for size. The dashed curves show the estimated ESS flowering probability, surrounded by 2.5 and 97.5 pointwise percentile curves based on the 1,000 bootstrap parameter vectors. The observed mean age of flowering plants at the study site is 2.55 years (Rees et al. 1999).

icant, we used as test statistics the mean age at flowering and the difference between estimated and ESS flowering probability for an individual with the observed mean age and size at flowering (Rees et al. 1999). The 99% bootstrap confidence intervals (BC method) on these quantities did not overlap 0. The difference in strategies is biologically significant as well: in a population dominated by the observed flowering strategy, the ESS would have approximately 20% higher lifetime seed production and would increase in frequency by 4% per year (invasion $\lambda = 1.039$).

These results support the interpretation by Rees et al. (1999) that the size and age dependence of flowering in *Onopordum* cannot be explained by a deterministic model but rather has adapted in response to environmental variability. ESSs under environmental variability can be predicted using stochastic integral models. Childs et al.'s (2004) work is an example of a stochastic age × size model with offspring state independent of parent state; theory for general stochastic integral models will be presented elsewhere (S. P. Ellner and M. Rees, unpublished manuscript). More generally, the feasibility of using bootstrapping to quantify prediction uncertainty greatly facilitates the use of IPMs for ecological forecasting.

Discussion

In this article, we have shown how IPMs can be applied to species where individual demography is affected by multiple attributes that vary over the life cycle and can overcome some of the practical limitations of matrix models in such cases. Modern computing power means that the IPM is now practical for empirical applications. As recently as a decade ago, desktop computers would have been inadequate for our *Onopordum* model—even using just 50 mesh points for size and 20 for quality, there are 1 million fecundity elasticities for each age at which reproduction is possible.

The existence of a unique stable distribution and asymptotic growth rate for the integral model rest on two kernel properties: either power positivity or *u*-boundedness. Power positivity is the analogue of assuming that a projection matrix is primitive and can be tested computationally. The *u*-boundedness condition is more abstract, but it will generally be satisfied by models with mixing at birth. The mixing at birth condition is likely to hold so long as the range of possible offspring states is the same for all parents, and it can usually be checked in specific models because it involves only properties of newborns.

Our application to *Onopordum* illustrates that when expected growth, survival, and birth rates are a smooth function of continuously varying traits, an integral model is a direct translation of the statistical analysis of individual-level demographic data. This tight connection is particularly important when individuals exhibit substantial variability in multiple traits affecting vital rates, which is often the case. Considerable sophistication is now possible in individual-level demographic modeling because of developments during the past decade in statistical theory and software. For example, hierarchical or mixed models (Pinheiro and Bates 2000; Clark 2003) can be fitted where individual-specific parameters are drawn from a distribution. Arbitrary distributions and correlation structures for latent trait variability can be fitted via Markov chain Monte Carlo in either Bayesian (Gelman et al. 2004) or frequentist (de Valpine 2004) paradigms. Generalized least squares allow models to be fitted where the error structure is heteroscedastic or correlated, as in the growth model of Pfister and Stevens (2002, 2003). Our *Onopordum* data were fitted well by linear models, but nonlinear and nonparametric mixed models can be used when needed (e.g., Wood 2004, 2005). Model-averaging approaches can be represented in the IPM simply by weighted averaging of the kernels implied by each demographic model under consideration.

Elasticities are widely used in comparative studies to partition the contributions of different demographic processes to λ at both the species and population levels (Silvertown et al. 1993; Silvertown and Dodd 1996). Because *Onopordum* individuals are characterized by size, age, and survival intercept, we can partition elasticities according to each of these attributes and so obtain a very detailed understanding of the contributions to λ from individuals in different states. These elasticities are in some ways easier to interpret than those of matrix models, whose values depend on the number of stages (Enright et al. 1995). For *Onopordum*, survival and growth contributed more to λ than did reproduction (75% vs. 25% of the total elasticity, respectively), which is remarkably similar to the results obtained by Childs et al. (2003) using an age- and size-structured IPM for the monocarpic thistle *Carlina vulgaris* (66% vs. 34%). The patterns of age specificity of both the survival-growth and fecundity elasticities are also very similar, with the survival-growth elasticities decreasing with age and the fecundity elasticities having a maximum value for intermediate ages (3 years in *Onopordum* and 2 in *Carlina*).

The use of matrix models in applied situations is often hampered by limited data, and this is often cited as a rationale for constructing a low-dimensional matrix model. This is why Caswell (2001, sec. 3.3) and Morris and Doak (2002, p. 192) discuss ways of choosing among different state variables rather than incorporating multiple state variables and why most of the models reviewed by Caswell (2001) classify individuals by a single state variable. However, an integral model will often make more

parsimonious use of limited data. In our *Onopordum* model, effects of age and individual quality required only three additional parameters: two regression slopes for age effects and survival intercept variance. A matrix model with individuals cross-classified by age, size, and quality would have an enormous number of possible transitions and would require data on multiple individuals in each age × size × quality category to estimate all entries in the projection matrix. In an attempt to minimize this kind of problem for estimating survival rates, Morris and Doak (2002) suggest a two-stage procedure: first, do logistic regression of survival against size or age, and then use the fitted regression to estimate survival in each class. The second step can be done by using the class midpoints or the median observed size within each class or by averaging over the observed size distribution within each class. In contrast to these ad hoc procedures—which still produce a model where survival is a discontinuous function of size—the fitted regression model (e.g., fig. 1*c*) can be expressed exactly as an integral model.

It should also be noted that if one fits an entire matrix model by regression procedures like those advocated by Morris and Doak (2002) for survival rates, then an IPM is obtained simply by increasing the number of size classes. Morris and Doak (2002) suggest that characterizing growth by regression models may be difficult. But given the wide range of regression models that can be used to fit trends in mean and variance, most patterns of growth can be characterized using standard statistical packages, and we have successfully applied this approach to several plant species (e.g., *Onopordum* in this article; *Oenothera glazioviana* in Rees and Rose 2002; *Cirsium canescens* in Rose et al. 2005; *Carlina vulgaris* in Childs et al. 2003, 2004; *Aconitum noveboracense* in Easterling et al. 2000) and to Soay sheep (T. Coulson, personal communication); for many other examples, see the article by Metcalf et al. (2003). The size × quality model (app. A), or similar models classifying individuals by size and energy reserves or by aboveground and belowground sizes, illustrates how temporal correlations in growth, either positive or negative, can be modeled using IPMs.

Recent work on model-based measures of individual fitness (Cam et al. 2002; Link et al. 2002) uses modeling approaches to get around the problem that observed values of realized individual fitness are based on samples of size 1 and can be seriously biased (Link et al. 2002). Instead, Link et al. (2002) propose estimating the distribution of individual fitness by combining age-structured projection matrices with mixed models for individual quality variation. The same approach can be used in integral models, which greatly increases the range of life histories to which these approaches can be applied. More speculatively, by using breeding value for heritable traits as one of the state variables, it should be possible to include quantitative trait dynamics in an integral model. Also, as suggested by Caswell (2001, sec. 8.3), integral models may be useful for describing the spatial dynamics of continuously structured populations.

Acknowledgments

For comments on the manuscript at various stages, we thank E. Crone, W. S. C. Gurney, M. Kot and members of his research group, M. Mangel, A. Patil, C. Pfister and members of her research group, and three anonymous reviewers whose comments were especially helpful. Research support was provided by National Science Foundation (NSF) grant OCE 0326705 in the NSF/National Institutes of Health Ecology of Infectious Diseases program, the Andrew W. Mellon Foundation and the College of Arts and Sciences at Cornell (S.P.E.), and Natural Environment Research Council grant NER/A/S/2002/00940 (M.R.).

APPENDIX A

Numerical Methods

This appendix provides more information about numerical methods for implementing integral models. We have tried to suggest methods that are relatively simple to implement but efficient enough for practical use on a current desktop computer. We describe ways to implement age × size and size × quality integral projection models (IPMs) and then discuss transpose iteration (to compute the reproductive value v) and the calculation of sensitivities and elasticities by perturbation. We also explain how we implemented the *Onopordum* model involving age × size × quality variation. Anyone interested in building and using IPMs should read this appendix. The mathematical level is the same as the main text, requiring multivariate calculus and some familiarity with matrix models. Script files illustrating the methods in R (R Core Development Team 2005) are provided as a zip archive in the online edition of the *American Naturalist.*

IPM with Age × Size Classification

In many populations, demographic rates are influenced by both size and age, as discussed in the main text, necessitating an easily parameterized modeling framework. We define $n_a(x, t)$ to be the distribution of size x for age a individuals in year t. The integral model is then

$$n_0(y, t+1) = \sum_{a=0}^{M} \int_{\Omega} f_a(y, x) n_a(x, t)\, dx,$$

$$n_a(y, t+1) = \int_{\Omega} p_{a-1}(y, x) n_{a-1}(x, t)\, dx, \qquad \text{(A1)}$$

where $a = 1, 2, \ldots, M$, $f_a(y, x)$ is the fecundity kernel, $p_a(y, x)$ is the survival-growth kernel representing age a individuals of size x living to age $a + 1$ and growing to size y, and M is the maximum age excluding postreproductive ages (Childs et al. 2003). As in age-structured matrix models, if there is no a priori maximum age, an absorbing state can be used to represent individuals of ages M or older; then, for $a = M$ in equation (A1), we have

$$n_M(y, t+1) = \int_{\Omega} [p_{M-1}(y, x) n_{M-1}(x, t) + p_M(y, x) n_M(x, t)]\, dx.$$

This is like a Leslie matrix with an added nonzero entry in the bottom right corner, representing survival of "old" individuals who will still be "old" next year if they survive.

The model (A1) is a series of one-dimensional integrals. For simplicity, suppose that the size interval for all ages is the same, $\Omega = [L, U]$. Mesh points x_i for the midpoint rule are defined as in the main text (eq. [4]). The survival-growth transitions are then approximated as

$$n_{a+1}(x_j, t+1) = h \sum_{i=1}^{m} p_a(x_j, x_i) n_a(x_i, t). \qquad \text{(A2)}$$

This is a matrix multiplication

$$\mathbf{n}_{a+1}(t+1) = \mathbf{P}_a \mathbf{n}_a(t), \qquad \text{(A3)}$$

where $\mathbf{P}_a$ is the matrix whose (i, j)th entry is $hp_a(x_i, x_j)$ and $\mathbf{n}_a(t)$ is the vector whose ith entry is $n_a(x_i, t)$. For an absorbing state M, we have $\mathbf{n}_M(t+1) = \mathbf{P}_{M-1}\mathbf{n}_{M-1}(t) + \mathbf{P}_M\mathbf{n}_M(t)$. Similarly, the integrals for fecundity are approximated as

$$\mathbf{n}_0(t+1) = \sum_{a=0}^{M} \mathbf{F}_a \mathbf{n}_a(t), \qquad \text{(A4)}$$

where $\mathbf{F}_a$ is the matrix whose (i, j)th entry is $hf_a(x_i, x_j)$.

IPM with Size × Quality Classification

Attributes representing individual quality will often change slowly over an individual's lifetime. This leads to positive autocorrelation in growth or reproduction (Pfister and Stevens 2002, 2003): individuals with rapid growth or high fecundity now are likely to be high-quality individuals who will continue to perform well in the future. Negative autocorrelations can occur if current growth or reproduction depletes storage reserves and limits future growth or reproduction (Ehrlen 2000). Pfister and Stevens (2002) documented within-year positive correlation in growth that degraded the forecasting accuracy of a size-classified matrix model but could be accommodated in an individual-based simulation model (Pfister and Stevens 2003). An IPM with size and quality variables provides an alternative model for these situations, whose properties can be determined without recourse to individual-based simulations.

Let x and q denote individual size and quality, respectively. These will typically be modeled as continuous variables jointly distributed on some rectangle $\Omega = [L_1, U_1] \times [L_2, U_2]$. An IPM in these variables is

$$n(x', q', t+1) = \int_{L_2}^{U_2} \int_{L_1}^{U_1} K(x', q' | x, q) n(x, q, t)\, dx\, dq. \qquad \text{(A5)}$$

To evaluate equation (A5) by midpoint rule, mesh points for both variables are defined:

$$x_i = L_1 + (i - 0.5)h_1,$$
$$q_i = L_2 + (i - 0.5)h_2, \qquad \text{(A6)}$$

where $h_j = (U_j - L_j)m_j$ for $j = 1, 2$. The midpoint rule approximation to equation (A5) is

$$n(x_k, q_l, t+1) = h_1 h_2 \sum_{i=1}^{m_1} \sum_{j=1}^{m_2} K(x_k, q_l | x_i, q_j) n(x_i, q_j, t). \qquad \text{(A7)}$$

In models with three continuous-state variables, equation (A5) is replaced by a triple integral and equation (A7) is replaced by a triple sum over all mesh point combinations for the three variables, and so on, for higher-dimensional models.

As an example, Pfister and Stevens (2002, 2003) assumed linear dynamics for the size and quality of each individual i,

$$x_{i,t+1} = a + bx_{i,t} + q_{i,t} + e_{i,t},$$
$$q_{i,t+1} = \rho q_{i,t} + v_{i,t},$$

where $e_{i,t}$ and $v_{i,t}$ are independent normal random variables with mean 0 and variances σ_X^2 and σ_Q^2, respectively, and a, b, and ρ are constants, with $|\rho| < 1$. The survival-growth component of the kernel is then

$$P(x', q'|x, q) = \varphi(x'; a + bx + q, \sigma_X^2)\varphi(q'; \rho q, \sigma_Q^2),$$

where $\varphi(z; \mu, \sigma^2)$ is the probability density for a normal distribution with mean μ and variance σ^2. This allows unbounded variation in size and quality, so finite limits would be set that extend well beyond the range of observed values. The fecundity component might have the form

$$F(x', q'|x, q) = g_X(x')g_Q(q')B(x),$$

where B is the size-dependent fecundity and g_X and g_Q are the probability densities for offspring size and quality. This assumes that offspring size and quality are independent and independent of parental traits, which might hold when q reflects site quality and offspring settle at random.

It is intuitive to think about the population state as a matrix whose i,jth entry is $n(x_i, q_j)$. However, in matrix languages, it will usually be more efficient to code the model as a single large matrix multiplying a state vector. This contrasts with our advice for age × size models because the block structure implied by stepwise age transitions is not generally present in a size × quality model. To create the state vector, take the state matrix and stack its columns into a single vector: column 1 on top, column 2 underneath column 1, and so on, down to column m_2 at the bottom. This gives a state vector $\mathbf{n}$ of length m_1m_2, whose entry in location $\eta(i, j) \equiv (j - 1)m_1 + i$ corresponds to $n(x_i, q_j)$. Let $\mathbf{A}$ be the $m_1m_2 \times m_1m_2$ matrix whose entry in location $(\eta(k, l), \eta(i, j))$ is $h_1h_2K(x_k,q_l|x_i,q_j)$. The midpoint rule approximation then corresponds to $\mathbf{n}(t + 1) = \mathbf{A}\mathbf{n}(t)$. With 50 mesh points each for size and quality, $\mathbf{A}$ has 6,250,000 entries—larger than the usual matrix model, but a matrix-vector multiplication of this size takes roughly 0.03 s on a current PC (in R, ver. 2.1.1, for Windows, using the matrix library or Matlab 7.0.1). If quality changes slowly, then most entries in $\mathbf{A}$ will be 0, and some speedup may be possible with sparse matrix routines (e.g., roughly twofold for a 2,500 × 2,500 matrix with 95% 0 entries in Matlab 7.01).

Transpose Iteration

The transpose iteration is

$$v(x, t + 1) = \int_\Omega v(y, t)K(y, x)dy, \tag{A8}$$

which is equivalent to forward iteration of the transpose kernel $K^T(y, x) = K(x, y)$. The reproductive value function v (the dominant left eigenvalue of the kernel) can therefore be computed by forward iteration using K^T. A useful numerical check whether transpose iteration has been implemented successfully is to verify that the computed dominant eigenvalues from forward and transpose iteration are identical.

If forward iteration is implemented by a single matrix $\mathbf{A}$, then equation (A8) is the same using $\mathbf{A}^T$. If forward iteration is done using component matrices, let $\mathbf{A}_{ij}$ be the matrix representing kernel component K_{ij}. The forward iteration is $\mathbf{n}_i(t + 1) = \sum_j \mathbf{A}_{ij}\mathbf{n}_j(t)$, with $\mathbf{n}_j(t)$ denoting the vector of population distribution values at mesh points for component j. The transpose iteration is then implemented as

$$\mathbf{v}_j(t + 1) = \sum_i \mathbf{A}_{ij}^{\mathbf{T}}\mathbf{v}_i(t).$$

For example, in an age × size model with component matrices (A3) and (A4), the transpose iteration is analogous to iterating a transposed Leslie matrix. The indexes where $\mathbf{A}_{ij}^{\mathbf{T}} \neq 0$ are $(i, j) = (a + 1, a)$ or $(0, a)$, corresponding to survival and fecundity, respectively, so the transpose iteration is the system

$$\mathbf{v}_a(t + 1) = \mathbf{P}_a^T\mathbf{v}_{a+1}(t) + \mathbf{F}_a^T\mathbf{v}_0(t),$$

with $\mathbf{v}_{M+1} \equiv 0$, where M is the maximum possible age.

Implementing the Onopordum *Model*

Because we assume that quality is constant over the lifetime in *Onopordum*, survival-growth transitions are block structured: individuals in component $\Omega_{a,k}$ (age a and quality class k) either die or else move to $\Omega_{a+1,k}$. We therefore used component matrices for survival-growth transitions, stacking them all into a four-dimensional array $\mathbf{P}$, whose entry in location i, j, a, k is $hP_{a,k}(x_i, x_j)$, where the x's are mesh points for size. The current population state is stored in a three-dimensional array $\mathbf{N}$, where $\mathbf{N}[j, a, k] = n_{a,k}(x_j, t)$. The midpoint rule for survival-growth transitions is then performed by computing for each age × quality combination the matrix-vector product $\mathbf{P}[\cdot, \cdot, a, k]\mathbf{N}[\cdot, a, k]$ that gives $\mathbf{N}[\cdot, a + 1, k]$ at time $t + 1$.

Births are less structured because offspring size and quality are independent of parent age, size, and quality.

We stored values of per capita total seedling production $p_e f_n(x_j) p_f(x_j, a) s(x_j, a, q_k)$ in an array **B** with the same structure as **N**. The total seedling production is computed as $S(t) = h \times \text{sum}(\mathbf{B} \cdot \mathbf{N})$, where $\mathbf{B} \cdot \mathbf{N}$ is the element-by-element product of **B** and **N**. The size × quality distribution for seedlings is then given by $S(t)\mathbf{J}$, where the (j, k)th entry in the matrix **J** is the fraction of seedlings of quality class k and size x_j, calculated from the distributions in table 1. This gives all the age 0 entries of **N** at time $t + 1$, $\mathbf{N}[\cdot, 0, \cdot] = S(t)\mathbf{J}$.

For sensitivity analysis, we used transpose iteration with component matrices to compute the reproductive value v and then applied the sensitivity formula. In all calculations, we used 50 mesh points for size and 40 for quality, and we imposed a maximum possible age of 7 years. The maximum age observed in the field was 5 years, but the study was terminated after the last individual from the initial cohort died, and some plants might live longer. Increasing the maximum age had virtually no effect on model output, as very few individuals survive past age 5 in the model.

Computing Sensitivities by Perturbation

For numerical accuracy, it will be better to use centered differences,

$$s(y_0, x_0) \doteq \frac{\lambda(\varepsilon | y_0, x_0) - \lambda(-\varepsilon | y_0, x_0)}{2\varepsilon}. \qquad \text{(A9)}$$

It is important to remember that ε is the integral of the perturbation to the kernel. For example, in an age × size model, the kernel components are all continuous-to-continuous: $K_{ij}(y, x)$, with $x \in \Omega_j$, $y \in \Omega_i$, and $i = j + 1$ (survival) or 0 (reproduction). To evaluate sensitivity at mesh point $(x_{i,l}, x_{j,k})$ by iteration, the perturbation function f_r is taken to be 0 outside the grid cell centered on $(x_{i,l}, x_{j,k})$, and inside that cell $f_r(y, x) \equiv \varepsilon/(h_i h_j)$, where the h's are defined by equation (4). Thus, $\lambda(\pm\varepsilon | x_{i,l}, x_{j,k})$ is computed by adding $\pm\varepsilon/(h_i h_j)$ to $K_{ij}(x_{i,l}, x_{j,k})$ and reestimating λ.

Three speedups are essential when computing sensitivities by iteration: (1) the sensitivity formula (13) implies that $s(y, x) = s(y, x^*)[w(x)/w(x^*)]$ for any x^*. So, for any given y, once a single sensitivity $s(y, x^*)$ is computed by perturbation, all other $s(y, x)$ can be inferred. As a result, the number of sensitivities that must be computed by perturbation is at most the total number of mesh points for all components of **X**, rather than the number of mesh points squared. (2) If elasticities are the goal, sensitivities do not have to be computed for mesh points where the kernel is 0. (3) The stable distribution for the unperturbed kernel is a good approximation to that for the perturbed kernel. So if w is used as the initial population structure, iteration (7) will converge quickly.

To be sure that the perturbed kernel has a dominant eigenvalue, ε should be small enough that the perturbed kernel value is still positive. However, ε cannot be too small or the effect on λ will be too small to estimate numerically. So, for each value of y, the x^* in (1) should be chosen to maximize the impact on λ of a 100% reduction in the kernel near (y, x^*). This is roughly equivalent to maximizing $e(y, x)$ as a function of x, which is equivalent (by the sensitivity formula) to maximizing $K(y, x)w(x)$ as a function of x.

Using these methods, the elasticity surface for a basic size-structured model with size-dependent growth, survival, and fecundity was computed by perturbation in approximately 15 s using 250 grid points. Over the 62,500 resulting grid points for the sensitivity and elasticity surfaces, the maximum relative error in elasticity was 0.06% relative to the elasticity surface computed from the analytic formula (13).

Literature Cited

Allan, C. J., and P. J. Holst. 1996. Longevity of soil based seeds of *Onopordum illyricum*. Plant Protection Quarterly 11:242.

Benton, T. G., T. C. Cameron, and A. Grant. 2004. Population responses to perturbations: predictions and responses from laboratory mite populations. Journal of Animal Ecology 73:983–995.

Cam, E., W. A. Link, E. G. Cooch, J. Y. Monnat, and E. Danchin. 2002. Individual covariation in life-history traits: seeing the trees despite the forest. American Naturalist 159:96–105.

Caswell, H. 1988. Approaching size and age in matrix population models. Pages 85–105 *in* B. Ebenman and L. Persson, eds. Size-structured populations. Springer, London.

———. 2001. Matrix population models: construction, analysis, and interpretation. Sinauer, Sunderland, MA.

Childs, D. Z., M. Rees, K. E. Rose, P. J. Grubb, and S. P. Ellner. 2003. Evolution of complex flowering strategies: an age and size-structured integral projection model approach. Proceedings of the Royal Society of London B 270:1829–1838.

———. 2004. Evolution of size dependent flowering in a variable environment: construction and analysis of a stochastic integral projection model. Proceedings of the Royal Society of London B 271:425–434.

Clark, J. S. 2003. Uncertainty in population growth rates calculated from demography: the hierarchical approach. Ecology 84:1370–1381.

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., S. LaDeau, and I. Ibanez. 2004. Fecundity of trees and the colonization-competition hypothesis. Ecological Monographs 74:415–442.

Coulson, T., E. A. Catchpole, S. D. Albon, B. J. T. Morgan, J. M. Pemberton, T. H. Clutton-Brock, M. J. Crawley, and B. T. Grenfell. 2001. Age, sex, density, winter weather, and population crashes in Soay sheep. Science 292:1528–1531.

Cushing, J. M. 1998. An introduction to structured population mod-

eling. Society for Industrial and Applied Mathematics, Philadelphia.
de Valpine, P. 2004. Monte Carlo state-space likelihoods by weighted posterior kernel density estimation. Journal of the American Statistical Association 99:523–536.
Doak, D. F., W. F. Morris, C. Pfister, B. E. Kendall, and E. M. Bruna. 2005. Correctly estimating how environmental stochasticity influences fitness and population growth. American Naturalist 166: E14–E21.
Easterling, M. R. 1998. The integral projection model: theory, analysis and application. PhD diss. North Carolina State University, Raleigh.
Easterling, M. R., S. P. Ellner, and P. M. Dixon. 2000. Size-specific sensitivity: applying a new structured population model. Ecology 81:694–708.
Ehrlen, J. 2000. The dynamics of plant populations: does the history of individuals matter? Ecology 81:1675–1684.
Enright, N. J., M. Franco, and J. Silvertown. 1995. Comparing plant life-histories using elasticity analysis: the importance of life-span and the number of life-cycle stages. Oecologia (Berlin) 104:79–84.
Fox, G. A., and B. E. Kendall. 2002. Demographic stochasticity and the variance reduction effect. Ecology 83:1928–1934.
Gelman, A., J. B. Carlin, H. S. Stern, and D. B. Rubin. 2004. Bayesian data analysis. Chapman & Hall, London.
Gross, K. L. 1981. Predictions of fate from rosette size in four "biennial" plant species: *Verbascum thapsus*, *Oenothera biennis*, *Daucus carota*, and *Tragopogon dubius*. Oecologia (Berlin) 48:209–213.
Isaacson, E., and H. B. Keller. 1966. Analysis of numerical methods. Wiley, New York.
Kendall, B. E., and G. A. Fox. 2002. Variation among individuals and reduced demographic stochasticity. Conservation Biology 16:109–116.
———. 2003. Unstructured individual variation and demographic stochasticity. Conservation Biology 17:1170–1172.
Klinkhamer, P. G. L., T. J. de Jong, and E. Meelis. 1987. Delay of flowering in the biennial *Cirsium vulgare*: size effects and devernalization. Oikos 49:303–308.
Kot, M., M. A. Lewis, and P. van den Driessche. 1996. Dispersal data and the spread of invading organisms. Ecology 77:2027–2042.
Krasnosel'skij, M. A., J. A. Lifshits, and A.V. Sobolev. 1989. Positive linear systems: the method of positive operators. Heldermann, Berlin.
Law, R. 1983. A model for the dynamics of a plant-population containing individuals classified by age and size. Ecology 64:224–230.
Lei, S. A. 1999. Age, size and water status of *Acacia gregii* influencing the infection and reproductive success of *Phoradendron californicum*. American Midland Naturalist 141:358–365.
Link, W. A., E. G. Cooch, and E. Cam. 2002. Model-based estimation of individual fitness. Journal of Applied Statistics 29:207–224.
McGraw, J. B. 1989. Effects of age and size on life histories and population-growth of *Rhododendron maximum* shoots. American Journal of Botany 76:113–123.
Metcalf, J. C., K. E. Rose, and M. Rees. 2003. Evolutionary demography of monocarpic perennials. Trends in Ecology & Evolution 18:471–480.
Moloney, K. A. 1986. A generalized algorithm for determining category size. Oecologia (Berlin) 69:176–180.
Morris, W. F., and D. F. Doak. 2002. Quantitative conservation biology: theory and practice of population viability analysis. Sinauer, Sunderland, MA.
Pettit, W. J., D. T. Briese, and A. Walker. 1996. Aspects of thistle population dynamics with reference to *Onopordum*. Plant Protection Quarterly 11:232–235.
Pfister, C. A., and F. R. Stevens. 2002. The genesis of size variability in plants and animals. Ecology 83:59–72.
———. 2003. Individual variation and environmental stochasticity: implications for matrix model predictions. Ecology 84:496–510.
Pinheiro, J. C., and D. M. Bates. 2000. Mixed-effects models in S and S-Plus: statistics and computing. Springer, London.
R Development Core Team. 2005. R: a language and environment for statistical computing. R Foundation for Statistical Computing, Vienna. http://www.R-project.org.
Rees, M., and K. E. Rose. 2002. Evolution of flowering strategies in *Oenothera glazioviana*: an integral projection model approach. Proceedings of the Royal Society of London B 269:1509–1515.
Rees, M., A. Sheppard, D. Briese, and M. Mangel. 1999. Evolution of size-dependent flowering in *Onopordum illyricim*: a quantitative assessment of the role of stochastic selection pressures. American Naturalist 154:628–651.
Rose, K. E., M. Rees, and P. J. Grubb. 2002. Evolution in the real world: stochastic variation and the determinants of fitness in *Carlina vulgaris*. Evolution 56:1416–1430.
Rose, K. E., S. M. Louda, and M. Rees. 2005. Demographic and evolutionary impacts of native and invasive insect herbivores: a case study with Platte thistle, *Cirsium canescens*. Ecology 86:453–465.
Silvertown, J., and M. Dodd. 1996. Comparing plants and connecting traits. Philosophical Transactions of the Royal Society of London B 351:1233–1239.
Silvertown, J., M. Franco, I. Pisanty, and A. Mendoza. 1993. Comparative plant demography: relative importance of life-cycle components to the finite rate of increase in woody and herbaceous perennials. Journal of Ecology 81:465–476.
Tuljapurkar, S. 1990. Population dynamics in variable environments. Lecture Notes in Biomathematics 85. Springer, New York.
Tuljapurkar, S., C. C. Horvitz, and J. B. Pascarella. 2003. The many growth rates and elasticities of populations in random environments. American Naturalist 162:489–502
Vandermeer, J. 1978. Choosing category size in a stage projection matrix. Oecologia (Berlin) 32:79–84.
van Groenendael, J. M., and P. Slim. 1988. The contrasting dynamics of two populations of *Plantago lanceolata* classified by age and size. Journal of Ecology 76:585–599.
Weiner, J., S. Martinez, H. Muller-Scharer, P. Stoll, and B. Schmid. 1997. How important are environmental maternal effects in plants? a study with *Centaurea maculosa*. Journal of Ecology 85:133–142.
Werner, P. A. 1975. Predictions of fate from rosette size in teasel (*Dipsacus fullonum* L.). Oecologia (Berlin) 20:197–201.
Wood, S. N. 2004. Low rank scale invariant tensor product smooths for generalized additive mixed models. Technical Report 04-13. Department of Statistics, University of Glasgow.
———. 2005. mgcv: GAMs with GCV smoothness estimation and GAMMs by REML/PQL. R package version 1.2–3. R Foundation for Statistical Computing, Vienna. http://www.R-project.org.

Associate Editor: William F. Morris
Editor: Donald L. DeAngelis