23. VIABILITY ANALYSIS FOR ENDANGERED METAPOPULATIONS: A DIFFUSION APPROXIMATION APPROACH

E.E. Holmes and **B. Semmens**

23.1 INTRODUCTION

Population viability analysis (PVA) assesses the rate of population decline and the risks of extinction or quasiextinction over a defined time horizon for a population of concern (Gilpin and Soule, 1986; Boyce, 1992; Morris and Doak, 2002). Although the techniques employed to conduct PVA are varied, they typically involve building quantitative models that are parameterized by demographic and environmental data. PVA was first used in the early 1980s (Shaffer, 1981), and in the past decade it has gained broad acceptance in the conservation community as a useful tool for assessing and managing "at-risk" species (Beissinger, 2002; Morris and Doak, 2002; Reed et al., 2002). This is particularly true for demogaphic PVAs, due in large part to the advancements in Monte Carlo techniques and desktop computers (Beissinger, 2002). The International Union for the Conservation of Nature (IUCN)'s Red List Criteria,

probably the most widely applied set of decision rules for determining the status of at risk species, is partially defined by metrics that require some form of PVA (IUCN, 1994). For instance, under one of the Red List criteria, a taxon may be classified as endangered if a "reduction of at least 50%, projected or suspected to be met within the next ten years or three generations" is predicted.

Although many PVAs are focused on single populations in single sites, there are often needs for spatially explicit PVAs: many populations of conservation concern are distributed across multiple sites and additionally, the primary anthropogenic threats facing at-risk species are habitat destruction and alteration, which are fundamentally spatial processes (Wilcove et al., 1998). Several software packages have been written for spatially explicit PVA, including RAMAS Metapop (Akçakaya, 1997) and RAMAS GIS (Boyce, 1996), ALEX (Possingham and Davies, 1995), and VORTEX (Lacy, 1993). These models incorporate a diversity of demographic and spatial attributes such as distance-dependent migration, allee effects, social population structure, habitat quality and spatial arrangement, and genetic variability. The development of flexible sophisticated PVA software packages such as these has made the construction and simulation of spatially explicit PVA models feasible for those who are not highly skilled programmers and has greatly increased the number of managers and scientists capable of using spatially realistic PVA models.

As the use of PVA has grown in conservation science, so have concerns that PVAs are often overextended given limited data sets (Reed et al., 2002). Beissinger and Westpahl (1998) suggested that PVA should be limited to assessing relative risks over short time frames using the simplest models that can reasonably be justified. For single species with spatially simple structure, data needs can often be met when Beissinger and Westpahl's call for model moderation and simplicity are heeded. When one is faced with species with more complex spatial structure, a much larger amount of data is needed to parameterize the dynamics of individual local populations, the levels and patterns of dispersal, and the spatial pattern of temporal correlations among local populations (e.g., Ralls et al., 2002). Unfortunately, collection of data needed to parameterize a spatial model is rare for species of conservation concern, at least in the United States (Morris et al., 2002), and there is a disconnect between the parameter requirements for spatially explicit PVA models and the willingness and/or ability of management agencies to collect the types of data needed to appropriately apply such tools. Because it is usually impossible to retroactively fulfill data requirements for a spatial PVA and there will always be cases where collection of spatial data is infeasible, managers require PVA tools that can help guide conservation of metapopulations in the absence of large amounts of spatial data.

Diffusion Approximation for Metapopulations

One approach to the problem of limited population data is to find a diffusion approximation that correctly models the long-run statistical properties of a complex population process. This approach has been used successfully for single population models (Karlin and Taylor, 1981; Lande and Orzack, 1988; Lande, 1993; Dennis et al., 1991; Hill et al., 2002; see also Morris and Doak, 2002; Lande et al., 2003) and reduces the problem of parameterizing a large model with many parameters to the much simpler task of parameterizing a
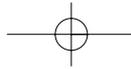
two-parameter diffusion model. One of the main practical implications of the diffusion approximation approach is that it is not necessary to know the multitude of parameters describing the local dynamics, dispersal levels, spatial patterns of dispersal, and spatial synchrony between local populations in order to make basic predictions about the statistical distribution of the long-term metapopulation or local population trajectories. The relevant two parameters for the diffusion approximation can be estimated from a simple time series of counts from the population process.

This chapter uses the diffusion approximation approach to model the long-run behavior of spatially structured populations. Our focus is on stochastic metapopulations characterized by structured population size, density-independent local dynamics, and, in keeping with the assumption of density independence, a metapopulation that is declining as a whole. Local populations are assumed to have patch-specific structured local dynamics and dispersal rates, with spatial structure among local populations in terms of both their local dynamics and dispersal patterns. Description of the long-run statistical behavior of the metapopulation trajectories using a diffusion approximation allows the estimation of PVA risk metrics such as the long-term rate of metapopulation decline and the probability of reaching different threshold declines over different time horizons (i.e., probabilities of extinction or quasiextinction). These methods for estimating metapopulation PVA metrics are illustrated using data from two chinook salmon metapopulations in the U.S. Pacific Northwest.

23.2   A STOCHASTIC METAPOPULATION MODEL

Our focus is on declining metapopulations, and thus what has been termed nonequilibrium metapopulations. We model a collection of local populations connected by dispersal where local populations have density-independent local dynamics, which may be "sources" or "sinks," but the metapopulation as a whole is declining. Dispersal levels could be very low, resulting in basically independent local populations, or extremely high, resulting in essentially one population. From a practical standpoint, this approach is most appropriate when dispersal is not insignificant (e.g., above 2% per year localized dispersal or 0.1% global dispersal), otherwise parameterization of the model requires inordinately long time series. Data from this type of metapopulation would be characterized by fluctuating local population trajectories, but actual extinctions would be unusual until the metapopulation has very few individuals. Our model assumes no density dependence nor carrying capacities within the individual local populations. Such a model is only appropriate in cases where the population is declining and all local populations are well below their carrying capacities. Our example using data on chinook salmon illustrates a situation that is likely to be well modeled as this type of metapopulation.

The following section gives a rather parameter-intensive mathematical description of a stochastic, declining metapopulation. However, the reader should keep in mind that this model will not be parameterized. Rather the asymptotic behavior of this model's trajectories will be derived and that information will be used to develop a diffusion approximation of the process. Time series data will then be used to parameterize the diffusion approximation.
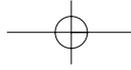
The Model

Consider an individual local population i with stochastic yearly growth and stochastic dispersal to and from other local populations. Such a local population's numbers in year t, $N_i(t)$, could be described as follows:

$$N_i(t) = \text{growth} - \text{dispersal out} + \text{dispersal in}$$
$$= N_i(t-1)e^{z_i(t-1)} - d_i(t-1)N_i(t-1)e^{z_i(t-1)} \qquad (23.1)$$
$$+ \sum_{j \neq i} \alpha_{ji}(t-1)d_j(t-1)N_j(t-1)e^{z_j(t-1)}$$

where $z_i(t)$ is the stochastic growth rate of local population i in year t and is a random variable with some unspecified statistical distribution with mean μ_i and variance σ_i^2. The μ_i term will be referred to as the local population's intrinsic growth rate; it will not be observed, as the local population is subject to immigration and emigration. Some fraction of individuals, $d_i(t)$, leaves local population i at year t and disperses to other local populations, and dispersal into local population i occurs from other local populations. The fraction of dispersers from local population j that go to local population i in year t is $\alpha_{ji}(t)$ and can vary depending on the destination, i, thus allowing for spatially structured dispersal. The dispersal parameters, $d_i(t)$ and $\alpha_{ji}(t)$, are assumed to be temporally random variables from some unspecified statistical distribution.

The Model in Matrix Form

The model for the entire metapopulation can be written using a random transition matrix, $\mathbf{A}(t)$, which encapsulates both dispersal and local growth:

$$\begin{bmatrix} N_1(t+1) \\ N_2(t+1) \\ N_3(t+1) \\ \ldots \\ N_k(t+1) \end{bmatrix} = A(t) \times \begin{bmatrix} N_1(t) \\ N_2(t) \\ N_3(t) \\ \ldots \\ N_k(t) \end{bmatrix} \qquad (23.2)$$

where

$$A(t) = \begin{bmatrix} (1-d_1)e^{z_1} & \alpha_{21}d_2e^{z_2} & \alpha_{31}d_3e^{z_3} & \ldots & \alpha_{k1}d_ke^{z_k} \\ \alpha_{12}d_1e^{z_1} & (1-d_2)e^{z_2} & \alpha_{32}d_3e^{z_3} & \ldots & \alpha_{k2}d_ke^{z_k} \\ \alpha_{13}d_1e^{z_1} & \alpha_{23}d_2e^{z_2} & (1-d_3)e^{z_3} & \ldots & \alpha_{k3}d_ke^{z_k} \\ \ldots & \ldots & \ldots & \ldots & \ldots \\ \alpha_{1k}d_1e^{z_1} & \alpha_{2k}d_2e^{z_2} & \alpha_{3k}d_3e^{z_3} & \ldots & (1-d_k)e^{z_k} \end{bmatrix} \qquad (23.3)$$

The '(t)' on the d's, α's, and z's have been left off to remove clutter. There may be any level or spatial pattern of temporal correlation among the intrinsic local growth rates, z_i's, dispersal rates, d_i's, and dispersal patterns, α_{ji}'s.

In the matrix model, each row represents 1 unit of habitat. Local populations with multiple units of habitat appear as multiple rows with very high dispersal
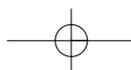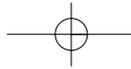
between the units of habitat in that local population. The habitat units within a local population could vary in quality (i.e., habitat within a local population need not be uniform) and different local populations certainly differ in the number of habitat units they contain. The d_i's and α_{ji}'s are assumed to be drawn from some distribution that can be different for each local population or local population pair. Although the d_i's, α_{ji} and z_i's are temporally random variables, they are assumed to be stationary, i.e., that there is no overall change in the mean values over time. For the purposes of this chapter, it will be assumed that the d_i's, α_{ji}'s, and z_i's are all strictly positive, which means that all local populations are connected to each other to some (although possibly very low) degree and that mean yearly geometric growth rates, $\exp(\mu_i)$'s, while possibly very small are not zero. These assumptions imply that the $\mathbf{A}(t)$ describe an ergodic set of matrices (Caswell, 2001). The assumption of strict positivity is not strictly necessary. It is possible for $\mathbf{A}(t)$ to describe an ergodic set if some elements of $\mathbf{A}$ are zero; it depends on the pattern of zeros within $\mathbf{A}$ [cf. Caswell (2001) for a discussion of the conditions under which matrices are ergodic].

The model is very general, allowing some sites to be dispersal sources and others to be dispersal targets, allowing any spatial pattern of dispersal or spatially correlated local growth rates, allowing any pattern of temporal correlation amongst local growth rates, and allowing any combination or pattern of habitat sizes of local sites.

Using Random Theory to Understand the Model's Statistical Behavior

Together, Eqs. (2) and (3) describe a quite generic model of a declining metapopulation with density-independent local dynamics. From a viability analysis perspective, one might ask the question: "Can one predict the viability of the total metapopulation?" In more precise terms, this is asking what are the statistical properties of the metapopulation trajectories of this type of connected collection of local populations [of the form in Eqs. (2) and (3)]? Clearly, the matrix $\mathbf{A}(t)$ has a large number of parameters that would be difficult, if not impossible, to estimate for any given metapopulation of conservation concern. However, using random theory, it can be shown that the long-term dynamics can be described by only two parameters and that it is unnecessary to know the multitude of other parameters for the purpose of projecting long-run dynamics.

To use this random theory, we first need to recognize that this stochastic metapopulation model falls into the class of random processes that involve products of ergodic random matrices, in this case products of $\mathbf{A}(t)$, which can be seen by using Eq. (2) to project the vector of local population sizes forward:

$$\mathbf{N}(1) = \mathbf{A}(0)\mathbf{N}(0)$$
$$\mathbf{N}(2) = \mathbf{A}(0)\mathbf{A}(1)\mathbf{N}(0)$$
$$\dots$$
$$\mathbf{N}(t) = \mathbf{A}(0)\mathbf{A}(1)\mathbf{A}(2) \dots \mathbf{A}(t-1)\mathbf{N}(0) \tag{23.4}$$

where $\mathbf{N}(t)$ is the column vector of N_i values at time t in Eq. (2). Products of random ergodic matrices have a well-established theoretical foundation and have certain well-studied asymptotic statistical properties. A brief review of two of the key results from this theory is provided in Box 23.1 and a simulated

BOX 23.1   Key Results from Random Theory

Two of the fundamental results from the theory of products of random matrices are reviewed and interpreted in the context of our metapopulation model. The reader is referred to Caswell (2001) and Tuljapurkar (1990) for other reviews interpreted in the context of demographic, single population models.

The Metapopulation and Local Populations Decline at the Same Rate

One of the basic results from Furstenberg and Kesten's "Products of Random Matrices" (1960) is that the product of ergodic random matrices asymptotically goes to an equilibrium. Say that $\mathbf{X_t}$ is an ergodic random "$k \times k$" matrix and that $\mathbf{Y}$ (also a $k \times k$ matrix) denotes the product of n of the $\mathbf{X}$ matrices: $\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3, \ldots \mathbf{X}_n$. Then Furstenberg and Kersten's results say that $\mathbf{Y}$ goes an equilibrium state such that

$$\lim_{t \to \infty} \frac{1}{t} \log \sum_{i \in \alpha} \sum_j \mathbf{Y}_{ij} = \text{a constant which is the same for all } a \qquad \text{(B1)}$$

We can use this result to show that the long-run exponential growth rate of the metapopulation and the local populations will be the same.

$$\mathbf{N}(t) = \mathbf{A}(0)\,\mathbf{A}(1)\,\mathbf{A}(2)\ldots\mathbf{A}(t-1)\mathbf{N}(0) \quad \text{our metapopulation model}$$

$$\text{Let } \mathbf{Y} = \mathbf{A}(0)\,\mathbf{A}(1)\,\mathbf{A}(2)\ldots\mathbf{A}(t-1)$$

$$\text{Then } \log N_i(t) = \log \sum_j \mathbf{Y}_{ij} + \log N_i(0)$$

$$\text{and } \log M(t) = \log \sum_i \sum_j \mathbf{Y}_{ij} + \log M(0)$$

Thus from Eq. (B1),

$$\lim_{t \to \infty} \frac{1}{t} \log \sum_j \mathbf{Y}_{ij} = \lim_{t \to \infty} \frac{1}{t} \log \sum_i \sum_j \mathbf{Y}_{ij} = \text{a constant} = \mu_m$$

The Distribution of Local Population and Metapopulation Sizes is Distributed Lognormally

One of the most powerful results, for our purposes at least, concerns the statistical distribution of the metapopulation and local trajectories. This tells us what distribution of sizes we would see if we ran our model over and over again and allows us to make population viability analyses for metapopulations since we have a prediction about the likelihood of different metapopulation futures. Random theory (Furstenberg and Kersten, 1960; Tuljapurkar and Orzack, 1980) shows that any sum of the $N_i(t)$'s, such as the total metapopulation (all i's), a single local population (one i), or any other subset, goes to the same distribution:

$$\log \frac{\mathbf{c}^{\hat{}}\mathbf{N}(t)}{\mathbf{c}^{\hat{}}\mathbf{N}(0)} \xrightarrow[t \to \infty]{} \text{Normal}(t\mu_m,\ t\sigma_m^2) \qquad \text{(B2)}$$

where the sum of local populations is denoted in matrix terms as $\mathbf{c}^{\hat{}}\mathbf{N}(t)$ and $\mathbf{c}$ is a column vector with 0's and 1's to show which local populations to sum together.
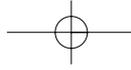
Example

These results are simple to see with simulations. An example of a linear chain of 10 local populations connected via 2% yearly dispersal to their nearest neighbors and 0.2% to nonnearest neighbors is shown. The local dynamics were e^{z_i} where z_i is a normally distributed random variable, Normal(μ_i, σ_i^2). The local growth rates, μ_i's, for local populations 1 to 10 were, respectively, 0.97, 1.00, 0.96, 0.83, 0.88, 1.00, 1.00, 0.89, 0.99, and 0.81. Figure 23.1A shows that the long-run growth rate of the local population and metapopulations is equal to the same constant. Figure 23.1B shows that the distribution of metapopulation size after 100 yr is Normal($100\mu_m$, $100\sigma_m^2$). The expected distribution was estimated using the maximum likelihood (ML) estimates for μ_m and σ_m^2 [Eq. (9)] from a single 1000-yr time series of metapopulation counts. The ML estimate for σ_m^2 relies on an assumption of normality for $t = 1$, although strictly speaking normality only holds for t large. However, it does quite well as can be seen in Fig. 23.1B.

example is shown to illustrate these results. As described in Box 23.1, the theory demonstrates that this stochastic, density-independent metapopulation will have an asymptotic growth rate and that the metapopulation, $M(t) = \Sigma N_i(t)$, the individual $N_i(t)$'s, and sets of $N_i(t)$'s representing the units of habitat comprising a semi-independent local population will be distributed lognormally with the same parameters:

$$\log M(t)/M(0) \xrightarrow[t \to \infty]{} \text{Normal}(t\mu_m, t\sigma_m^2)$$

$$\log N_i(t)/N_i(0) \xrightarrow[t \to \infty]{} \text{Normal}(t\mu_m, t\sigma_m^2)$$

$$\log \sum_{i \in a} N_i(t)/\sum_{i \in a} N_i(0) \xrightarrow[t \to \infty]{} \text{Normal}(t\mu_m, t\sigma_m^2)$$

$$\text{where } a = \left\{a_1, a_2, \ldots, a_m\right\} \tag{23.5}$$

Figure 23.1 shows an example of this behavior. A metapopulation is simulated (described in Box 23.1) and, over time, the metapopulation declines at a constant rate and all $N_i(t)$'s have the same long-term fate. When viewed over short time frames, t small in Fig. 23.1, the local sites show different growth rates with some declining more or less than the long-term rate, but over the long-term their rate of decline is the same.

The model studied here approximates the local dynamics by a simple exponential growth (or decline) model. However, it has been shown that results from random theory (presented in Box 23.1) also apply to a more complicated metapopulation model where local dynamics are described by stochastic age-structured Leslie matrices (Sanz and Bravo de la Parra, 1998). Essentially, this occurs because even when the local dynamics are described by a local matrix model, the system can still be described by products of random matrices.
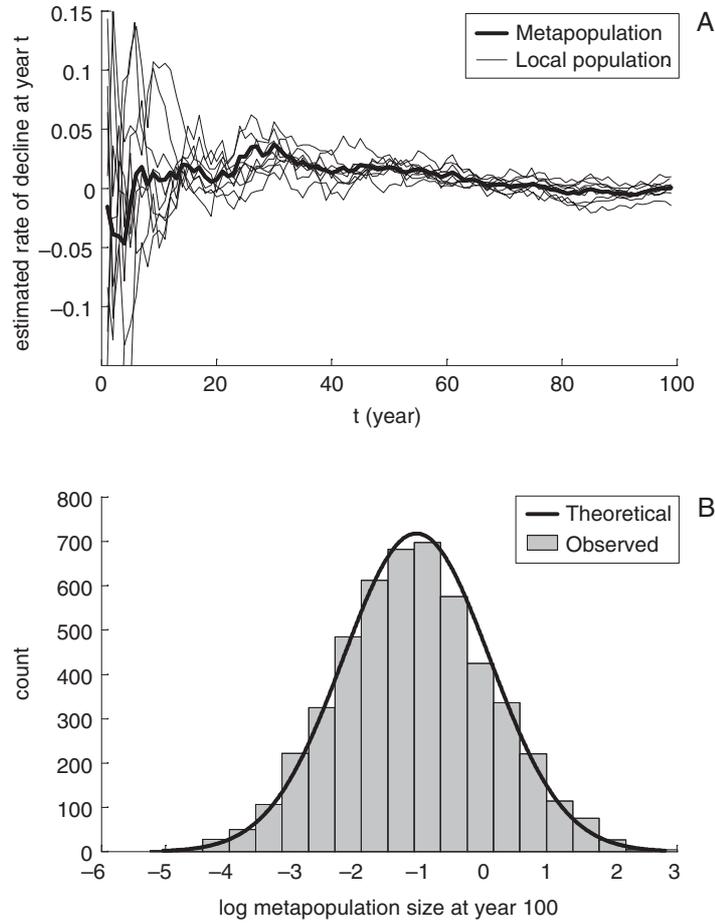
Fig. 23.1 Illustration of two of the main results from random theory. (A) All local populations go toward the same long-term rate of population growth (or decline) as t gets large. (B) The distribution of log $M(t)$ is a normal distribution with mean given by the long-term rate of growth (or decline) multiplied by t and the variance given by t multiplied by the rate that variance increases in an individual trajectory, i.e., $t \times (1/\tau)\log M(t + \tau)/M(t)$ for τ not overly small. Here the variance was estimated from one time series using $\tau = 10$ and this is used to predict the distribution at $t = 100$.

23.3 DIFFUSION APPROXIMATION

The asymptotic distribution of log $M(t)$ in Eq. (5) has the same properties as the distribution of a diffusion process with drift; it is normal and the mean and variance of the distribution of log $M(t)$ increase linearly with time, t. This observation in the context of age-structured matrix population models (Lande and Orzack, 1988; Dennis et al., 1991) led to the use of a diffusion approximation to enable parameterization using simple time series and to enable calculation of extinction probabilities. Diffusion approximation methods for single population populations are an important and established method for approximating stochastic trajectories (Lande and Orzack, 1988; Dennis et al.,
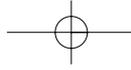
1991; Chapter 3 in Morris and Doak, 2002; Chapter 5 in Lande et al., 2003). Models for single populations are mathematically analogous to the models used here for metapopulations with a stochastic process involving products of random matrices. However, in single population models, the matrix represents a life history matrix rather than a growth and dispersal matrix, and the $\mathbf{N}(t)$ vector [in Eq. (2)] represents different age or stage classes, whereas in the metapopulation matrix, it represents different local sites and populations.

A diffusion approximation with drift is a stochastic process with the following properties (cf. Karlin and Taylor 1981):

$$\left.\begin{array}{l} X(t) - X(0) = \mu_m t + \varepsilon \\ \varepsilon \sim \text{normal}(0, \sigma_m^2 t) \end{array}\right\} \text{for } t = 1, 2, 3, \ldots \qquad (23.6)$$

For any nonoverlapping pair of time periods, $t_1 < t_2$ and $t_3 < t_4$, $X(t_2) - X(t_1)$, and $X(t_4) - X(t_3)$ are independent random variables. $X(t + \tau)$ is a random variable with distribution Normal $(X(t) + \mu_m \tau, \sigma_m^2 \tau)$. Correspondingly, the probability density function for $X(t + \tau)$ given log $X(t)$ is

$$p(X(t + \tau) \mid X(t)) = \frac{1}{\sqrt{2\pi\sigma_m^2\tau}} \exp\left[\frac{-(X(t + \tau) - X(t) - \mu_m\tau)^2}{2\sigma_m^2\tau}\right] \qquad (23.7)$$

Behavior of Metapopulation Trajectories Versus Diffusion Trajectories

Diffusion approximation is based on the behavior of log $M(t)$ as t goes to infinity; however, in PVA settings the time frame of interest is substantially less than infinity and is typically in the range of 25 to 100 yr. How well does the diffusion approximation do over these finite time periods? To explore this, a collection of 50 local populations were simulated that were connected by global dispersal ranging from 0.1 to 5% per year and that had correlated local dynamics, $z_i(t)$, drawn from a Normal(mean $= -0.05$, variance $= 0.09$) and a temporal covariance of $(0.2)(0.09)$ between the $z_i(t)$'s of local populations in any given year.

If the log metapopulation trajectories behave like a diffusion process, and if [AU1] we repeatedly generate a large sample of replicate metapopulation trajectories, the mean and variance of $(1/t)\log M(t)/M(0)$ from those trajectories should be a constants over the time period of interest. Additionally, $(1/t)\log M(t)/M(0)$ should be normally distributed. To examine whether the metapopulation trajectories had these properties, the simulations were started from a distribution of local population sizes selected from the equilibrium set of local population distributions and then run forward for 200 yr. This was repeated (using the same initial distribution of local populations) 1000 times to estimate the distribution of $(1/t)\log M(t)/M(0)$. This process was repeated for four randomly chosen initial distributions of local population sizes. The mean and variance of $(1/t)\log M(t)/M(0)$ are denoted as $\mu_m(t)$ and $\sigma_m^2(t)$, respectively, in Fig. 23.2 and in the discussion given later.

Figure 23.2 illustrates the results. For dispersal levels 1% or higher, the trajectories behaved like a diffusion process with $\mu_m(t)$ and $\sigma_m^2(t)$ roughly constant and the distributions approximately normal according to a Kolmogorov–Smirnov test
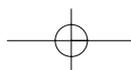
at $P = 0.05$. For low dispersal, 0.1%, the trajectories did not behave like a diffusion process for t less than 200 at least. Neither $\mu_m(t)$ nor $\sigma_m^2(t)$ was constant, except for $t > 150$, and the normality assumption was generally violated except again at large t. This means that when dispersal is very low, diffusion approximations for this metapopulation would be more approximate than for metapopulations with higher dispersal.

Figure 23.2 illustrates results from one particular model. Repeating this process for a number of different models indicated some general behaviors. The higher the dispersal levels, the more trajectories behaved like a diffusion process. Global dispersal levels of at least 2 to 5% were generally high enough to result in diffusion-like behavior within a short time frame. Note that localized dispersal has the effect of lowering the effective dispersal rates. The higher the amount of temporal covariance between local populations in terms of their yearly growth rates, the more the trajectories behaved like a diffusion process. The simulations were done with the local population sizes within the equilibrium set of local population distributions — indeed the theory is predicated on the local populations being near equilibrium. For metapopulations with 2 to 5% dispersal, the local populations equilibrated fairly quickly starting from all local populations with equal numbers. However, at very low dispersal, equilibration took thousands of time steps. This suggests that the assumption of equilibrium should be viewed cautiously for metapopulations that have very low dispersal rates between local populations.
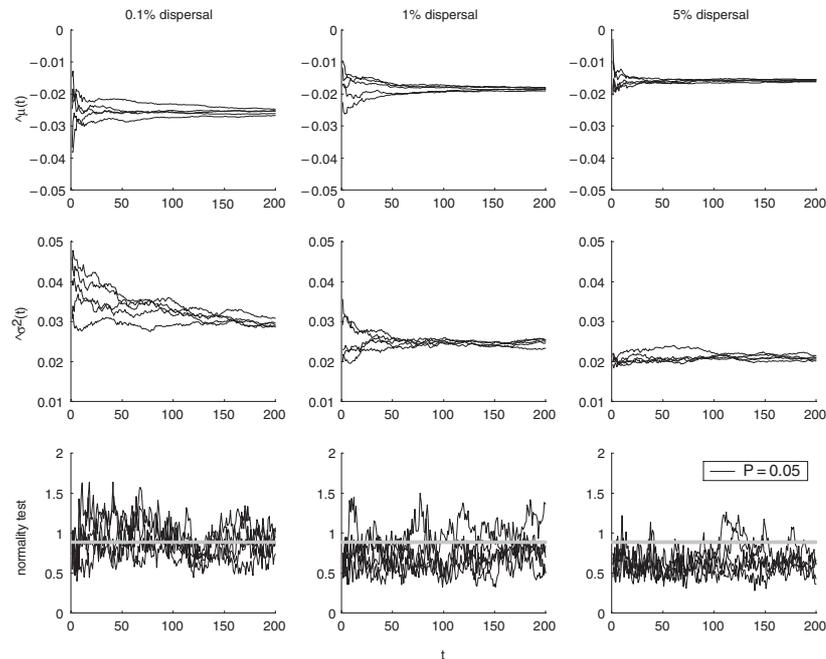
Fig. 23.2 Illustration of the performance of a diffusion approximation for modeling the behavior of a metapopulation with 50 local populations and uniform 0.1, 1, or 5% yearly dispersal. The diffusion approximation performs well for a given time frame when $\mu_m(t) = (1/t)\log M(t)/M(0)$ and $\sigma_m^2(t) = (1/t)$ var $[\log M(t)/M(0)]$ are constants over that time frame and when $\log M(t)/M(0)$ is normal.
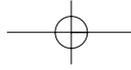
23.4 ESTIMATING THE PARAMETERS

Maximum likelihood estimates of μ_m and σ_m^2 can be calculated using the diffusion approximation for log $M(t)$. Denote the observed time series as $M = M(0)$, $M(1)$, $M(2)$, . . . , $M(n)$. If we approximate log $M(t)$ as a diffusion process, the likelihood function $L(\mu_m, \sigma_m^2/M)$ is given by the product of the probability function distributions for the transitions from log $M(t + 1)$ to log $M(t)$, which is Eq. (7) with $\tau = 1$, over $t = 0, 1, 2, . . . , n$. Thus the log likelihood function is

$$\log L(\mu_m, \sigma_m^2 \mid M) = -(n / 2) \log (2\pi\sigma_m^2)$$
$$-\frac{1}{2\sigma_m^2} \sum_{i=1}^{n} [\log(M(i)/M(i - 1)) - \mu_m]^2 \qquad (23.8)$$

Maximum likelihood estimates are obtained by solving for μ_m and σ_m^2, which maximize Eqn. (8),

$$\hat{\mu}_m = \frac{1}{n} \log \left(\frac{M(n)}{M(0)}\right)$$
$$\hat{\sigma}_m^2 = \frac{1}{n} \sum_{i=1}^{n} \left[\log \left(\frac{M(i)}{M(i - 1)}\right) - \hat{\mu}_m \right]^2 \qquad (23.9)$$

Note that the unbiased estimator for σ_m^2 would use $(n - 1)$ rather than n. The $\hat{\mu}_m$ and $\hat{\sigma}_m^2$ are analogous to the estimates of mean and variance from n samples from a normal distribution, and confidence intervals on μ_m and σ_m^2 are analogous:

$$\left(\hat{\mu}_m - t_{\alpha/2,n-1}\sqrt{\hat{\sigma}_m^2 /n}, \ \hat{\mu}_m + t_{\alpha/2,n-1}\sqrt{\hat{\sigma}_m^2 /n} \right)$$
$$(n\hat{\sigma}_m^2/\chi_{\alpha,n-1}^2, n\hat{\sigma}_m^2/\chi_{1-\alpha,n-1}^2 \qquad (23.10)$$

where t_{α},q is the critical value of a t distribution at $P = \alpha$ and q degrees of freedom and $\chi_{\alpha,q}^2$ is the critical value of a χ^2 distribution at $P = \alpha$ and q degrees of freedom. See Dennis et al. (1991) for a more in-depth discussion of maximum likelihood estimates for diffusion processes. Following Dennis et al.'s monograph, parameter estimation based on the diffusion approximation has been widely used for the analysis of single population trajectories. For a discussion of parameter estimation that is not based on the diffusion approximation, the reader is referred to Heyde and Cohen (1985).

Maximum likelihood estimates assume that the metapopulation has reached a stochastic equilibrium and thus that the diffusion approximation is reasonable. When exploring these methods using simulations, it is important to allow the system to equilibrate, after starting the simulation with something peculiar like all local populations at the same size. Equilibruim can be monitored by waiting for the variance of $(\log(N(t)_i) - \log[\mathrm{mean}(N_i(t))])$ to stabilize. In simulations done for this chapter, the distribution stabilized relatively quickly when dispersal was nonzero. If dispersal is zero, however, the distribution never stabilizes and the variance of $(\log(N_i(t)) - \log[\mathrm{mean}(N_i(t))])$ increases continually. For an actual metapopulation, for which one wants to
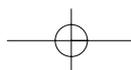
conduct a PVA, it is also critical to test the appropriateness of the diffusion approximation for one's time series data. Dennis et al. (1991) and Morris and Doak (2002) reviewed how to do this, which is based on diagnostic procedures for evaluating the appropriateness of linear models.

Parameter Bias

The estimators are unbiased maximum likelihood estimators for the diffusion approximation, $X(t)$. It is important to understand whether and how these estimates are biased when working with short time series of metapopulation trajectories, $M(t)$, as opposed to an actual diffusion process. In particular, $\hat{\sigma}_m^2$ is certain to be biased to some degree, as it relies on the diffusion approximation holding for $\tau = 1$ in log $M(t + \tau)/M(t)$, regardless of the length of the time series used for estimation. This is not the case for $\hat{\mu}_m$, which is also an unbiased predictor for $M(t)$ given a long time series (Heyde and Cohen, 1985).

To numerically explore parameter bias from short time series, simulations were used to look at the difference between $\hat{\mu}_m$ and $\hat{\sigma}_m^2$ from a 20-yr time series versus their true values μ_m and σ_m^2. An example metapopulation of 50 local sites was simulated with global dispersal and correlated local growth rates, $z_i(t)$, drawn yearly from a normal distribution with mean = μ_i, variance = ϕ, and covariance of $0.2*\phi$ between any two local growth rates Two versions of the simulation were run: one to model uniform site quality (spatially uniform $\mu_i = -0.05$) and one to model highly variable site quality (spatially variable μ_i's). To explore biases over a range of different dispersal and variability levels, models were run with dispersal between 0.1 and 5% per year and local variability, ϕ, between 0.1 and 0.5. These parameters translated to metapopulation level rates, μ_m, in the range of 0.01 to -0.05 and metapopulation level variability, σ_m^2, in [AU2] the range of 0.001 to 0.08. For each dispersal and local variability pair, 1000 replicate metapopulation trajectories were simulated, each with an initial distribution of local population sizes selected randomly from the equilibrium set.

The mean difference between $\hat{\mu}_m$ and μ_m over the dispersal and local variability parameter space was very low, <0.0015, for both uniform and variable μ_i simulations. Overall the lack of bias in $\hat{\mu}_m$ supports metrics that rely primarily on this parameter, such as the metapopulation λ (next section). For most of the parameter space explored, $0 < |\hat{\sigma}_m^2 - \sigma_m^2| < 0.01$, representing a 0 to 20% under- or overestimation of σ_m^2. Larger biases, $|\hat{\sigma}_m^2 - \sigma_m^2| > 0.01$, representing a >20% under- or overestimation, were seen for some parameter combinations. The impact of this bias depended on where $\hat{\sigma}_m^2$ was used. For instance, the effect on estimated confidence intervals on μ_m [Eq. (10)] was minimal with the width of the interval changing by a median 0.002. The effect on estimated passage probabilities was higher, although not dramatic. For example, the estimated probability that the metapopulation will be 10% of current levels at the end of 50 yr was decreased by 0 to 0.04 (on a scale from 0 to 1) for the uniform μ_i simulation and increased by 0 to 0.04 for the variable μ_i simulation. The estimated probability that the metapopulation will pass below 10% of current levels at any point during the next 50 yr was changed by 0 to 0.09. Overall, the effect of $\hat{\sigma}_m^2$ bias was low in these simulations, but this will depend on the particular metapopulation and will need to be investigated for individual cases of interest.
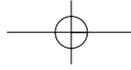
In practical applications, one must contend with other factors that can lead to parameter bias, but which are outside the scope of this chapter. In particular, observation error, nonequilibrium local population distributions, and temporal autocorrelation can lead to parameter bias. Such problems are being studied in the context single population PVA. Much of this work is likely to be relevant for metapopulation PVA. See Morris and Doak (2003) for a review and discussion of current work in this area.

23.5 METAPOPULATION VIABILITY METRICS

One of the most basic viability metrics is the long-term geometric rate of decline (or growth) of a population, termed generally λ in the PVA literature. If λ is less than 1.0, the population ultimately declines to extinction and $100*(1 - \lambda)$ is roughly the average yearly percent decline. The metapopulation λ is $\exp(\mu_m)$ and its estimate is then

$$\hat{\lambda} = \exp(\hat{\mu}_m). \tag{23.11}$$

This definition of λ follows Caswell's use of the symbol λ_s as the long-term average stochastic growth rate: $\lambda_s = [N(t)/N(0)]^{1/t}$ as $t \to \infty$ (Caswell, 2001). This is the long-run geometric growth rate that would be observed in almost every trajectory. Defined this way, if $\lambda < 1$, the population goes extinct with certainty, eventually. This differs from Dennis et al.'s use of the symbol λ where λ is used for $\exp(\mu + \sigma^2/2)$ and the long-term average geometric growth rate is instead denoted by $\alpha = \exp(\mu)$. The maximum likelihood estimate of λ is a biased estimator; because $\hat{\mu}_m$ is normally distributed, the median value of $\exp(\hat{\mu}_m)$ is $\exp(\mu_m)$ but the mean value is not. Dennis et al. (1991) gave an unbiased estimator [mean($\hat{\lambda}$) = λ] based on Shimizu and Iwase (1981), although Dennis and colleagues found negligible differences between biased and unbiased estimators in their examples.

From the asymptotic distribution of log $M(t)$, Eq. (5), the probability that the metapopulation is below a threshold b at the end of y years can be calculated as

$$P[M(t) \le b \mid M(0)] = \Phi\left(\frac{\log(b/M(0)) - \mu_m t}{\sqrt{\sigma_m^2 t}}\right) \tag{23.12}$$

Although this uses asymptotic distribution, this is mitigated by the fact that it is used for the distribution at the end of y years but not at any time before that. The estimate of $P[M(t) \le b \mid M(0)]$ replaces μ_m and σ_m^2 by their estimates $\hat{\mu}_m$ and $\hat{\sigma}_m^2$. Like the estimate of λ, the median estimate of $P[M(t) \le b \mid M(0)]$ is equal to the true value, but not the mean.

Some metapopulations can have a low long-term risk of being below a threshold due to a λ near 1.0, but high short-term risks of hitting that threshold due to high variability. Such quasiextinction or extinction probabilities are commonly used and very important PVA metrics. The diffusion approximation for log $M(t)$ can be used to estimate these probabilities for the metapopulation. The probability of that the diffusion process, $X(t)$, experiences a decline below
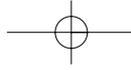
a threshold log b at some time T less than y years is calculated by integrating over the probability density function for first passage times for a diffusion process with drift (Karlin and Taylor, 1981). Lande and Orzack (1988) go through the calculation, which leads to

$$P(T \leq y) = \Phi\left(\frac{-(X(0) - \log b) - \mu_m y}{\sqrt{\sigma_m^2 y}}\right)$$

$$+ \exp(-2(X(0) - \log b)\mu_m/\sigma_m^2)\Phi\left(\frac{-(X(0) - \log b) + \mu_m y}{\sqrt{\sigma_m^2 y}}\right) \quad (23.13)$$

$\Phi()$ is the cumulative distribution function for a standard Normal(mean = 0, variance = 1). The estimate of $P(T \leq y)$ for the metapopulation uses $\hat{\mu}_m$ and $\hat{\sigma}_m^2$ with $X(0) = \log M(0)$. The estimated probability of extinction (to 1 individual) is calculated using Eq. (13) and setting b equal to 1. The reader is cautioned that estimates of extinction are problematic and that estimates of quasiextinction (e.g., some threshold greater than 1 individual) are more robust (cf. Morris and Doak, 2002). Also, Eq. (13) uses diffusion approximation over short time scales, as it calculates the probability of hitting a threshold at any time, including short times, before y years. This makes Eq. (13) more approximate than other metrics.

Other viability metrics based on diffusion approximation, such as the mean time to extinction, and median time to extinction, are discussed in Lande and Orzack (1988) and Dennis et al. (1991).

Risk Metric Uncertainty

The $100(1 - \alpha)\%$ confidence intervals are often used as characterizations of uncertainty. These can be calculated for risk metrics using the estimated distributions of $\hat{\mu}_m$ and $\hat{\sigma}_m^2$. The confidence intervals for $\hat{\lambda}$ are

$$(\exp(\hat{\mu}_m - t_{\alpha/2,n-1}\sqrt{\hat{\sigma}_m^2/n}), \exp(\hat{\mu}_m + t_{\alpha/2,n-1}\sqrt{\hat{\sigma}_m^2/n})). \quad (23.14)$$

where t_{α},q is the critical value of a t distribution at $P = \alpha$ and q degrees of freedom. The corresponding significance level, α, for a hypothesis test, such as "Is $\lambda < b$" is the α such that

$$\frac{\hat{\mu}_m - \log b}{\sqrt{\hat{\sigma}_m^2/n}} = t_{\alpha,n-1}. \quad (23.15)$$

Confidence intervals on $P(T \leq y)$ and $P[M(y) \leq b \mid M(0)]$ can be calculated by parametric bootstrapping from the estimated distributions of $\hat{\mu}_m$ and $\hat{\sigma}_m^2$: Normal($\hat{\mu}_m$, $\hat{\sigma}_m^2/n$) and Gamma(shape = $(n - 1)/2$, scale = $2\hat{\sigma}_m^2/(n - 1)$). A large number of $(\hat{\mu}_b, \hat{\sigma}_b^2)$ pairs are generated randomly by sampling from these distributions and the risk metric Ψ is calculated [Eqs. (13) or (12)] for each pair. The range of Ψ over the $(\hat{\mu}_b, \hat{\sigma}_b^2)$ bootstrapped pairs, for which both parameters are within their respective $100(1-\alpha)\%$ confidence intervals, defines the $100(1 - \alpha)\%$ confidence interval for Ψ. This and other methods for calculating confidence intervals for diffusion approximation risk metrics are discussed in Dennis et al. (1991).
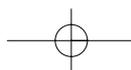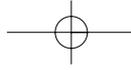
An alternate way to present the level of uncertainty is to estimate the data support for different values of a risk metric. There are both frequentist and Bayesian approaches for this [see Wade (2001) for a review geared toward conservation applications]. Holmes (2003) presented a Bayesian approach, which uses posterior probability distributions to illustrate data support. That method is adapted here for estimating the level of data support for the metapopulation risk metrics. Let Ψ be a risk metric. The probability that Ψ is greater than some threshold φ given the data are

$$P(\Psi > \varphi \mid \hat{\mu}_m, \hat{\sigma}_m^2) = \int_{\substack{\text{all } (\mu_m, \sigma_m^2) \\ \text{for which } \Psi > \varphi}} \frac{L(\mu, \sigma_m^2 \mid \hat{\mu}_m) L(\sigma_m^2 \mid \hat{\sigma}_m^2) \pi(\mu_m) \pi(\sigma_m^2)}{\eta(\hat{\mu}_m, \hat{\sigma}_m^2)} \quad (23.16)$$

where $L(\mu, \sigma_m^2 \mid \hat{\mu}_m)$ is the likelihood function given $\hat{\mu}_m \sim \text{Normal}(\mu_m, \sigma_m^2/n)$, $L(\sigma_m^2 \mid \hat{\sigma}_m^2)$ is the likelihood function given $\hat{\sigma}_m^2 \sim \text{Gamma}((n - 1)/2, 2\sigma_m^2/(n - 1))$, $\pi(\mu_m)$ and $\pi(\sigma_m^2)$ are the priors on μ_m and σ_m^2, and the normalizing constant is

$$\eta(\hat{\mu}_m, \sigma_m^2) = \int_{\mu_m = -\infty}^{\mu_m = +\infty} \int_{\sigma_m^2 = 0}^{\sigma_m^2 = \infty} L(\mu_m, \sigma_m^2 \mid \hat{\mu}_m) L(\sigma_m^2 \mid \hat{\sigma}_m^2) \pi(\mu_m) \pi(\sigma_m^2) d\mu_m d\sigma_m^2 \quad (23.17)$$

The posterior distribution of Ψ is $[P(\Psi < \varphi + d\varphi \mid \hat{\mu}_m, \hat{\sigma}_m^2) - P(\Psi < \varphi \mid \hat{\mu}_m, \hat{\sigma}_m^2)]/d\varphi$ over all φ. Examples of this calculation for λ and the probability of being below thresholds at the end of 25 yr are shown in the salmon examples. Holmes (2003) supplied Splus code for these calculations.

23.6   A SIMULATED EXAMPLE

In this example, a collection of 49 local populations in a 7×7 grid was simulated with neighborhood dispersal. Local populations were specified with variable mean local growth rates; thus, some μ_i values were much larger than others. The local growth rates in any given year were slightly correlated between sites. Thus all sites were more likely than random to have good and bad years together. Dispersal was variable between 5 and 10% from year to year and was mainly to the four nearest neighbors (or two and three for corner or edge sites). In specific terms, $\mathbf{A}(t)$ was specified with $z_i(t)$'s drawn from a normal distribution with mean $= \mu_i$ and a variance of 0.0625. The μ_i were different for each local population and were chosen randomly between -0.22 and -0.01. Each year, new $z_i(t)$'s were selected from the normal distribution for that local population. The $z_i(t)$'s were correlated among the local populations such that the covariance of $z_i(t)$ and $z_j(t)$ was $(0.1)(0.0625)$. The $d_i(t)$ varied from year to year. Each year and separately for each local population, $d_i(t)$ was selected from a uniform random distribution between 0.05 and 0.1; thus the dispersal varied from year to year and between local populations in any given year. Most of this dispersal, 80%, was to nearest neighbors. Thus for nearest neighbors, $\alpha_{ji} = 0.80 \, d_j(t)/nn$, where nn is the number of nearest neighbors, and for nonnearest neighbors, $\alpha_{ji} = 0.2 \, d_j(t)/nnn$; where nnn is the number of nonnearest neighbors.
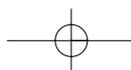
The simulation was started from a set of local population sizes drawn randomly from the stochastic equilibrium, and starting sizes were drawn anew from this distribution for each replicate of the simulation. For each replicate, a 25-yr time series was generated, and from this time series, μ_m and σ_m^2 were estimated using Eq. (9). From the estimates, the probability that the metapopulation would be below different thresholds (50 or 75% of starting levels) at the end of 25 yr was predicted and compared to the actual probabilities obtained by repeatedly (1000 times) running the simulation for 25 yr starting from the point where the initial 25-yr time series stopped. This simulation was replicated 500 times to generate the distribution of estimated probabilities of 50 and 75% decline in 25 yr versus the true probability. Also, from each 25-yr simulation, the metapopulation λ was estimated and compared to the actual value calculated by running a 10000-yr simulation. For each estimated risk metric, confidence intervals were estimated via the methods in Section 23.5.

Figure 23.3 shows the distribution of λ estimates and the estimated probabilities of 50 and 75% decline versus true values. As expected, the median estimate of λ was equal to the true value ($\hat{\mu}_m$ is an unbiased estimator of μ_m).
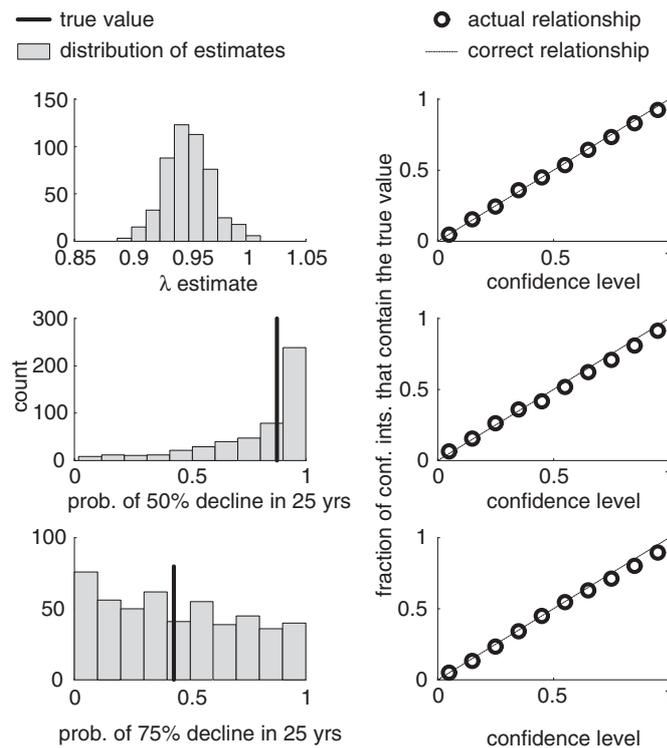
Fig. 23.3 Estimated viability metrics and their estimated confidence intervals versus the true values for a 49 site metapopulation in a 7×7 grid with 5–10% dispersal to the closest neighboring sites. (Left) True metrics compared to the distribution of estimated metrics from 500 simulations starting from the same initial conditions. (Right) Performance of the estimated confidence intervals by looking at the fraction of estimated $100(1-\alpha)$% confidence intervals that contain the true values.
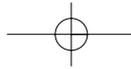
The median estimate of λ was 0.97 compared to the true value of 0.97. The median estimates of 50 and 75% decline were 0.63 and 0.14 compared to the true values of 0.62 and 0.12, respectively. Although the median estimates were very close to the true values, the estimates were variable. The estimates of λ ranged between 0.9 and 1.0. The estimates of declines to thresholds were also variable. The variability depended on the threshold and the time frame. In this example, there was low variability around the estimate of 50% decline in 25 yr, but high variability in the estimate of 75% decline. The true values for each of the metrics are shown by the solid lines in the middle of the distributions.

The variability of the estimates is due to the stochastic nature of the process and is not a fault of the estimation methods per se; by chance, short trajectories will appear to have underlying parameters that are different than the true underlying parameters, which leads to variability in the estimated viability metrics. When estimates are inherently variable, it is critical that the confidence intervals for the estimates be estimated correctly. Figure 23.3 (right) confirms that the estimated confidence intervals properly characterize the uncertainty for the estimate risk metrics: e.g., $100(1-\alpha)\%$ of the time the $100(1-\alpha)\%$ confidence intervals contain the true values.

23.7 SALMON AS METAPOPULATIONS

Salmonid populations (*Oncorhynchus* spp.) show strong spatial structuring and they have often been referred to as metapopulations (Reiman and McIntyre, 1995; Policansky and Magnuson, 1998; Cooper and Mangel, 1999; Hill et al., 2002). Spawning and rearing habitats of different salmon stocks occur on discrete and physically separated river or stream sections. Salmon have a well-known and strong tendency to return to their natal streams with a low (1 to 20%) dispersal to other stocks (Fulton and Pearson, 1981; Mathews and Waples, 1991; Quinn, 1993). Within the U.S. Pacific Northwest, collections of anadromous salmon stocks have been divided into "evolutionary significant units" (ESUs) (Waples, 1991), which represent substantially reproductively isolated conspecific groups that can be distinguished based on their coherence on a genetic level and known dispersal between the stocks. Salmon within a stock spawn on individual streams or river sections and the majority of offspring return to spawn in their natal stream or river. Straying of returning adults to nonnatal streams is spatially structured and occurs more frequently within subbasins. Stocks within an ESU have some level of synchrony due to exposure to common migratory corridors between the ocean and the natal stream and also due to exposure to similar large-scale ocean dynamics (Pearcy, 1992; Ware, 1995; Mantua et al., 1997). However, stocks also show a great deal of asynchrony due to exposure to their independent spawning and juvenile rearing habitats and variability in migration timing between stocks (e.g., PSTRT, 2001). Throughout the Pacific Northwest, most salmonid populations show regional decline with the majority of individual stocks showing steady declines with densities well below historical levels (Rieman and Dunham, 2000; McClure et al., 2003).

23.8   SNAKE RIVER SPRING/SUMMER CHINOOK ESU

The Snake river spring/summer chinook ESU (Fig. 23.4) includes all spring and summer chinook spawning within the subbasins of the Tucannon river, Grande Ronde river, and the south, middle, and east fork Salmon rivers, which flow into the Snake river below the Hells Canyon dam (Mathews and Waples, 1991). Juvenile fish rear in the mountain streams and then migrate down the Snake and Columbia rivers to the ocean. After maturing in the ocean, adult fish return to spawn at variable ages between 3 and 5 yr (mean = 4.5 yr). Tagging experiments in the Columbia river basin (which the Snake river basin is a part of) have found that the proportion of individuals that disperse and spawn away from their natal sites is on the order of 1–3% for wild-born individuals (Quinn, 1993).

The Snake river spring/summer chinook ESU was listed as threatened under the U.S. Endangered Species Act in 1992. Stocks within this large and complex basin, like salmon stocks throughout the Pacific Northwest, are impacted negatively by a variety of factors (Wissman et al., 1994) and many have experienced substantial declines (Myers et al., 1998; McClure et al., 2003). There is habitat degradation in many areas related to forestry, grazing, mining, and irrigation practices, resulting in lack of pools, high temperatures, low flows, poor overwintering conditions, and high sediment loads in many areas. At the same time, a substantial portion of the ESU is protected as part of federally designated wilderness (Mathews and Waples,
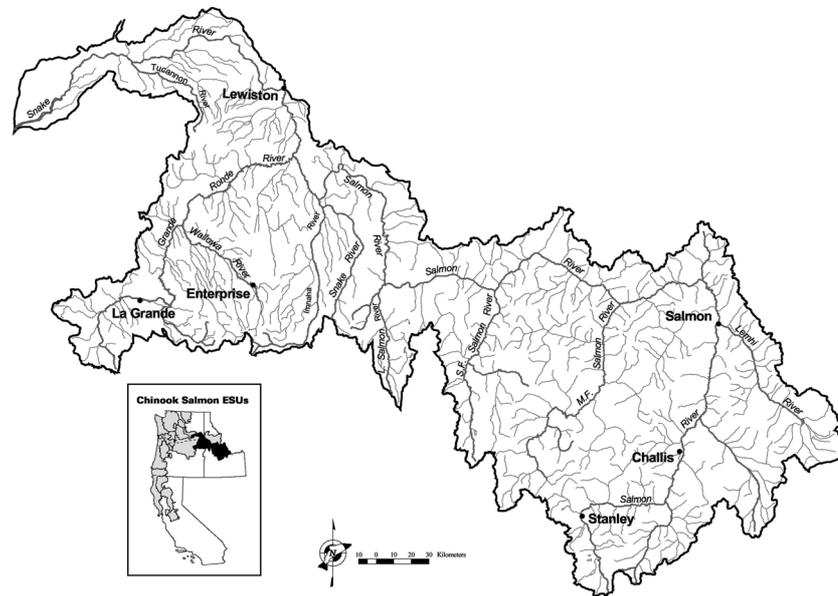
Fig. 23.4   Map of the Snake river spring/summer chinook ESU. The ESU includes stocks from the Snake river and its tributaries between Ice Harbor and Hells Canyon dams. The Hells Canyon hydropower dam has no passage facilities and blocks the migration of salmon into their historical habitat in the upper Snake river basin.
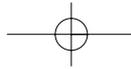
1991). The official ESU designation does not include salmon in the Clearwater basin, as chinook in this subbasin originate from hatchery fish that were stocked in the subbasin after the original natural fish were extirpated in the 1940s. However, from a metapopulation dynamic perspective, current stocks in the Clearwater river basin interact with stocks within other subbasins. Thus, in this analysis, all stocks in the entire Snake river basin were analyzed together.

A total metapopulation level time series was available for this ESU from counts of the total number of wild-born spawners returning through the Ice Harbor dam at the downstream end of the ESU (Fig. 23.4). Returning spawners can be either wild born or hatchery born as hatcheries have been operating in the basin since the early 1970s. McClure et al. (2003) discussed the effects of hatchery production on viability analyses. By focusing on the wild-born spawner time series and not incorporating a correction for hatchery production, the in-stream viability metrics assume that hatchery-born fish all return to the hatchery and do not spawn in stream (which would produce wild-born offspring). As discussed by McClure et al. (2003), this means our viability metrics are optimistic upper bounds, as some unknown fraction of hatchery fish do stray to the wild spawning grounds and potentially reproduce.

In addition to the metapopulation level dam count, time series of redds per mile (rpm), which are indices of the density of gravel egg nests made by spawning females, were available for the majority of stocks within the Snake river basin. Redds per mile are an index of the redds (and consequently returning spawners) trend within a stock, but the total redds are unknown, as the total spawning habitat is not surveyed. The majority of rpm and dam data are available in the digital appendices of McClure et al. (2003).

Parameter Estimation

Our Ice Harbor dam time series starts in 1962 and ends in 1999. The wild-born component of the dam count is denoted $M(0), M(1), M(2), \ldots M(37)$, where $M(0)$ is the 1962 count and $M(37)$ is the 1999 count. The maximum likelihood estimates presented in Eq. (9) assume that data do not contain sampling error or other nonprocess error; however, salmon data typically have high levels of sampling error and boom–bust cycles that confound estimation of μ_m and especially σ_m^2 (Holmes, 2001). An alternate approach uses data filtering and examination of the rate at which variance increases within the time series to improve parameter estimation and separate out sampling error variance from the time series (Holmes, 2001; Holmes and Fagan, 2002; cf also Morris and Doak, 2002). These methods have been cross-validated extensively with salmon data (Holmes and Fagan, 2002; Fagan et al., 2003) and are used here to estimate parameters. First, data are transformed using a running sum:

$$\widetilde{M}(t) = \frac{1}{4} \sum_{j=0}^{3} M(t+j) \text{ for } t = 0 \text{ to } 34 \tag{23.18}$$
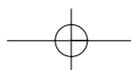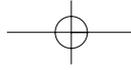
The estimates of μ_m and σ_m^2 are then

$$\hat{\sigma}_m^2 = \frac{1}{3}\left(\mathrm{var}\left(\log \widetilde{M}(t+4)\,/\,\widetilde{M}(t)\right) - \mathrm{var}\left(\log \widetilde{M}(t+1)\,/\,\widetilde{M}(t)\right)\right) = 0.0353$$

$$\hat{\mu}_m = \frac{1}{34}\sum_{t=0}^{34}\log \widetilde{M}(t+1)/\widetilde{M}(t) = -0.0561 \tag{23.19}$$

The estimate of σ_m^2 uses the property that the variance of the underlying stochastic process should increase linearly with time: $E[\mathrm{var}(\log M(t)/M(0))] = \sigma_m^2 t$. The confidence intervals for $\hat{\mu}_m$ and $\hat{\sigma}_m^2$ using $\widetilde{M}(t)$ are slightly different than Eq. (10) (Holmes and Fagan, 2002):

$$(\hat{\mu}_m - t_{\alpha/2,df}\,\sqrt{\hat{\sigma}_m^2\,/(t - L + 1)},\ \hat{\mu}_m + t_{\alpha/2,df}\sqrt{\hat{\sigma}_m^2\,/(t - L + 1)})$$

$$(df\,\hat{\sigma}_m^2/x_{\alpha,\,df}^2,\ df\,\hat{\sigma}_m^2/\chi_{1-\alpha,\,df}^2) \tag{23.20}$$

where L is the number of counts summed together for the running sum and $df = 0.333 + 0.212\,(n+1) - 0.387\,L = 6.84$ ($L = 4$ and $n = 38$ here). The estimated 95% confidence intervals on μ_m and σ_m^2 are $(-0.133, 0.020)$ and $(0.017, 0.111)$, respectively.

Metapopulation Viability Metrics

The estimate of λ for the Snake river spring/summer chinook ESU is $\hat{\lambda} = \exp(\hat{\mu}_m) = 0.94$. To the extent that long-term trends continue, the expected population size in 25 yr is 21% of current levels $(= \hat{\lambda}^{25})$. The point estimate of the probability of that the ESU drops below 10% of current levels at any time over the next 25 yr is

$$\exp(-2\,\hat{\mu}_m\log(10/1)/\hat{\sigma}_m^2) \times \left[\,\Phi\!\left(\frac{-\log(10/1)\ +\ |\hat{\mu}_m|\,25}{\sqrt{\hat{\sigma}_m^2\,25}}\right) + \right.$$

$$\left. \exp(2\log(10/1)|\hat{\mu}_m|/\hat{\sigma}_m^2)\,\Phi\!\left(\frac{-\log(10/1)\ +\ |\hat{\mu}_m|25}{\sqrt{\sigma_m^2\,25}}\right)\right] \tag{23.21}$$

$$= 0.23$$

The corresponding estimate of 90% decline over the next 50 yr is 0.74. The probability of extinction was not estimated, as this requires an estimate of the total population size. The number of returning spawners is not the total population size, as nonmature fish remain in the ocean. However, if the true λ of the metapopulation is less than 1.0, the population will eventually go extinct.

The posterior probability density functions [Eq. (16)] for the estimated metrics are shown in Fig. 23.5. The posterior probability distributions give an indication of the degree to which data support different risk levels. The distribution for λ shows considerable data support for a $\lambda < 1$, indicating a declining metapopulation. There is also strong data support for a high risk of 90% decline over the next 50 yrs; however, the estimate of 90% decline over 25 yr is very uncertain. The mean value is 0.23, but the probability distribution is
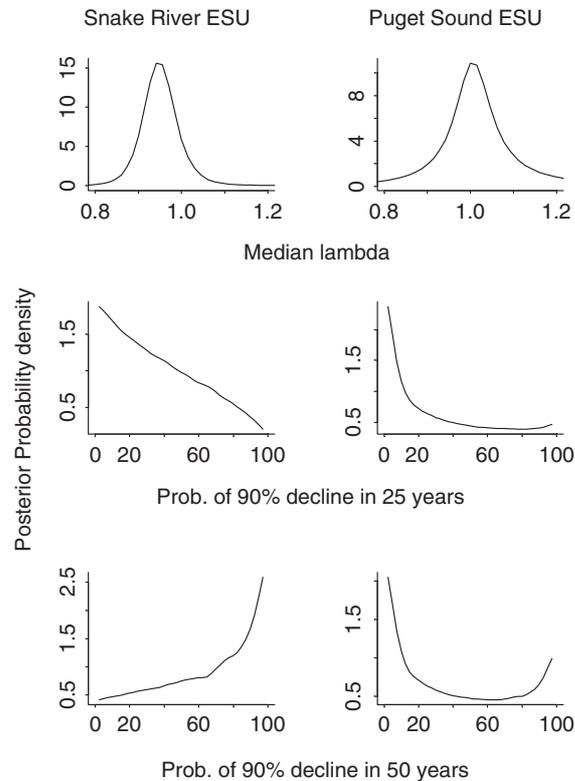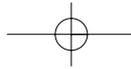
Fig. 23.5 Estimated posterior probability distributions for λ and the probability of 90% decline in 25 and 50 yr. Posterior probability distributions, which were calculated using uniform priors on μ_m and σ_m^2, indicate the relative levels of data support for different risk metric values. Distributions for Snake river spring/summer chinook (left) and Puget Sound ESUs (right).

very broad over the 0 to 1 range. This illustrates that uncertainty in estimates of probabilities of quasiextinction can vary widely depending on the time frame over which one is interested.

23.9 PUGET SOUND CHINOOK ESU

The Puget Sound ESU is a subset of the major chinook salmon group in Washington's northern coastal basins and Puget Sound. The ESU (Fig. 23.6) includes all spring, summer, and fall runs in the Puget Sound region from the north fork Nooksack river to the Elwha river on the Olympic peninsula (Myers et al., 1998). The Elwha and Dungeness coastal basins of the Strait of Juan de Fuca, Hood Canal, and the Puget Sound area north to the northern Nooksack river basin and the U.S. Canadian border are all a part of the Puget Sound ESU. Basin-to-basin dispersal rates have been observed at between 0.1 and 6% based on recoveries of tagged juveniles returning as adults (PSTRT, 2001). Fish in this ESU typically mature at ages 3 and 4 and are coastally oriented during the ocean phase of their life history. The Puget Sound ESU does not include Canadian or
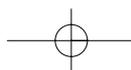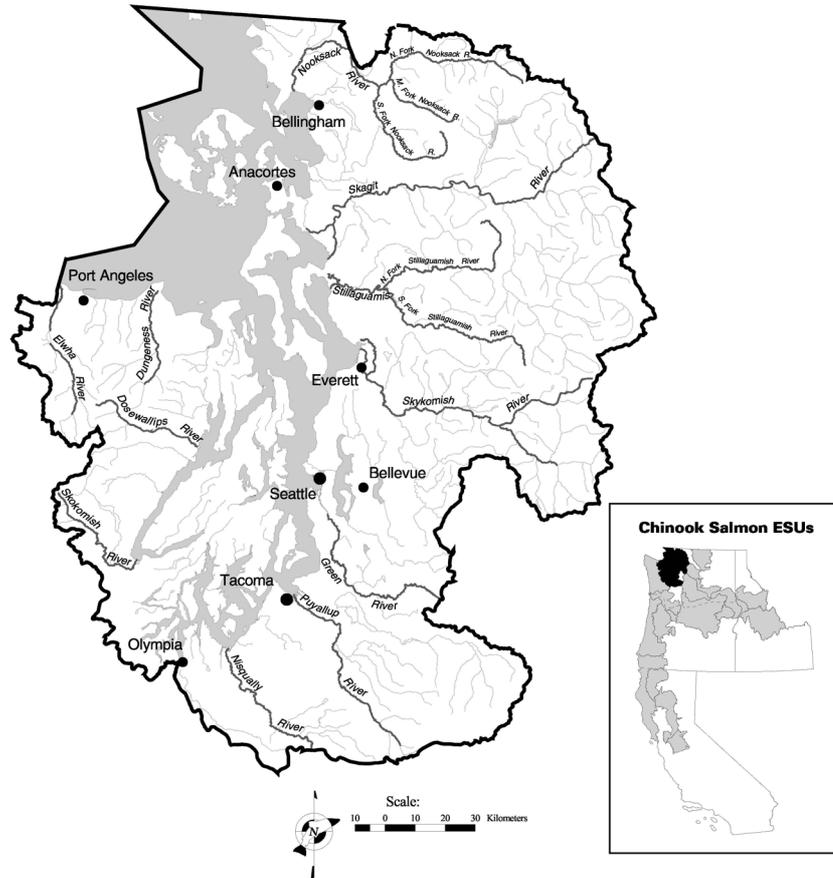
Fig. 23.6 Map of the Puget Sound chinook ESU.

coastal Washington populations. The Puget Sound ESU was listed as threatened under the Endangered Species Act in March of 1999. Trends in abundance throughout the ESU are predominantly downward, with several populations exhibiting severe short-term declines. Degraded spawning and rearing habitats, as well as access restrictions to spawning grounds and migration routes, have all likely contributed to population declines. Salmon in this ESU do not migrate through a hydropower system as the Columbia river ESUs do.

Data for this ESU consist of yearly estimates of the total returning spawners (wild-plus hatchery-born) to the 44 separate river and creek systems feeding into the Puget Sound (Fig. 23.6). These time series were compiled by the National Marine Fisheries Service (Seattle, WA) based on a variety of data: redd counts, carcass counts, in-stream harvest records, weir counts, and hatchery return counts. An independent metapopulation level count was not available; unlike spawners returning to the Columbia river basin, spawners here do not pass through a hydropower system where they can be enumerated. Instead, a 1979–1997 index of the metapopulation was constructed by added together the 29 time series for the local populations with data over the
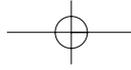
1979–1997 period. As for Snake river analyses, our viability metrics implicitly assumes that hatchery fish have not been reproducing and will be optimistic if some hatchery fish do not return to the hatchery and instead spawn successfully in the wild.

Metapopulation Viability Metrics

Parameters were estimated as for the Snake river. The parameter estimates are $\hat{\mu}_m = 0.0036$, and $\hat{\sigma}_m^2 = 0.012$. The estimate of λ for the Puget Sound chinook ESU is $\hat{\lambda} = \exp(\hat{\mu}_m) = 1.003$. The point estimate of the probability that the ESU drops below 10% of current levels at any time over the next 25 yr is 0.000 and over the next 50 yr is 0.001.

The posterior probability distributions (Fig. 23.5, right) illustrate the high uncertainty, given the data, as to whether this ESU is declining, stable, or increasing. The most that can be said from these data is that there is low data support for a severely declining ($\lambda < 0.9$) or increasing ($\lambda > 1.1$) metapopulation. Interestingly, the low support for small λ values translates into high data support for a low risk of 90% decline in the short term (over 25 yr). Over the longer term, however, the uncertainty as to whether the metapopulation is declining or increasing gives rise to a U-shaped distribution, meaning that data give the most support to a probability of 0 or 1, reflecting that λ could be either less than or greater than 1.0. This example illustrates that while data may be equivocal on some questions of conservation concern, such as "is $\lambda < 1$?", data may still give information on other questions, such as "is the short-term risk of severe decline high?"

23.10   USING THE STOCHASTIC METAPOPULATION MODEL TO INVESTIGATE EFFECTS OF MANAGEMENT

Determining how to distribute effort in order to recover an at-risk species is a routine, and challenging, task of conservation managers. For salmon, management actions tend as a generality to affect an entire ESU or multiple ESUs or to affect individual stocks. Management actions such as harvest reductions or increases to survival during migration (between spawning areas and the ocean) or improvements to estuarine environments are examples of actions that will tend to improve conditions for all stocks within an ESU or multiple ESUs. Habitat improvements or protections that affect spawning areas and management of in-stream water levels are examples of actions that tend to affect individual stocks. Without knowing the local stock dynamics or dispersal rates, one can still give certain types of guidance about how much effort is required for recovery of a declining metapopulation and about how effort should be distributed across all local populations.

Metapopulation Level Actions

When management actions affect all local populations roughly equally, it can be estimated how change would change the metapopulation λ. Mathematically, this means that all μ_i values increase by some $d\mu$.

An absolute $d\mu$ change in all μ_i values is equivalent to multiplying all elements in $\mathbf{A}(t)$ by a constant $= \exp(d\mu)$. The mean of the distribution of $\log M(t)/M(0)$ becomes $(\mu_m + d\mu)t = (\log\lambda_{new})t$. Thus $\log(\lambda_{new}/\lambda_{old}) = d\mu$. The change, $d\mu$, can be translated into currency that is more meaningful from a management standpoint by using the relationship $\lambda = R_0^{1/T}$, between λ, the net reproductive rate, R_0, and the mean generation time, T (Caswell, 2001). This is illustrated here for harvest and hydropower effects on salmon in the Snake river spring/summer chinook ESU (cf. McClure et al., 2003).

Harvest

In the Pacific Northwest, harvest rates for salmon are generally expressed in terms of the fraction of spawners that did not return to the spawning grounds but that would have without harvest, e.g., a harvest rate of 0.8 indicates that the actual number of returning spawners is 20% of what it would have been if there had been no harvest. Harvest rates are expressed in this way so that harvest that occurs in the stream versus in the ocean can be compared via a common currency. We can write the net reproductive rate using fecundity and age-specific survival (cf. Caswell 2001) as

$$R_0 = s_1 F_1 (1 - h)f + s_1(1 - F_1)s_2 F_2 (1 - h)f + s_1(1 - F_1)s_2$$
$$(1 - F_2)s_3 F_3 (1 - h)f \ldots \qquad (23.22)$$

where h is the harvest rate, s_i is the survival from age $i - 1$ to i, F_i is the fraction of spawners that return at age i, and f is the mean offspring per spawner. Using Eq. (22), the change in λ from a change in h alone is

$$\frac{\lambda_{new}}{\lambda_{old}} = \left(\frac{R_{0,new}}{R_{0,old}}\right)^{1/T} = \left(\frac{1 - h_{new}}{1 - h_{old}}\right)^{1/T} \qquad (23.23)$$

Hydropower

Juvenile salmon from the Snake river basin must migrate through the mainstem of the Snake river, enter the Columbia river, and descend down the Columbia river on their journey to the ocean. This migration, and the return migration of spawning adults, involves passage through four large hydropower dams on the Columbia river and four Snake river hydropower dams. Improving the survival of both juvenile and adult fish migrating through the Columbia and Snake river hydropower systems has been the focus of much effort and is one of the human impacts that has been relatively well quantified.

Following a strategy similar to that used for harvest, the effect of changes in survival through the hydropower system on the rate of decline at the ESU level can be estimated. Denoting by c_d and c_u the proportional increase in down- and upstream passage survival due to improvement in the hydropower system, the improved net reproductive rate is

$$R_{0,new} = c_d c_u (s_1 F_1 f + s_1(1 - F_1)s_2 F_2 f + s_1 (1 - F_1)s_2(1 - F_2)s_3 F_3 f \ldots). \qquad (23.24)$$
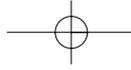
Thus, for assessing the impacts of increased survival through the hydropower system:

$$\frac{\lambda_{new}}{\lambda_{old}} = \left(\frac{R_{0,new}}{R_{0,old}}\right)^{1/T} = (c_d c_u)^{1/T} \qquad (23.25)$$

Estimates of the Impacts of Harvest and Hydropower Changes to the Snake River ESU

The mean ocean and in-river 1980–1999 harvest rate for the Snake river spring/summer chinook ESU was $h = 0.08$ (McClure et al., 2003). By setting $h_{new} = 0$, we can examine the effect of successful selective harvest management that would substantially eliminate harvest impacts on salmon in this ESU. Using Eq. (23) and a mean generation time of 4.5 yr, the estimated increase in λ with $h_{new} = 0$ is roughly 2%. NMFS (2000) has required that agencies operating the Federal Columbia river power system implement a variety of activities, including increased spill, improved passage facilities, and increased barging of salmon around the dams as a means of improving survival through the system. The estimated improvement in passage survival from the improvements proposed by NMFS are on the order of 5–6% (i.e., $c_d c_u = 1.05$–1.06) for the Snake river spring/summer chinook (McClure et al., 2003). This translates into a 1% improvement in λ for this particular ESU using Eq. (25). Thus if the combined effects of substantially reduced harvest and the proposed passage improvements are additive, then roughly a 3% increase in λ is estimated for these actions. If the true λ is less than 0.97, a 3% increase would not be sufficient to achieve $\lambda > 1$. Figure 23.5 indicates that data cannot rule out that the λ in this is ESU is greater than 0.97, but data certainly give more support to a lower λ. This suggests that other recovery actions, such as improvements at the stock level, will also be necessary.

Local Population Level Actions

The effects of changes to individual units of habitat are harder to quantify than the effects of metapopulaion level changes. The change in λ achieved by a change at the level of a specific unit of habitat depends on the level of dispersal, the spatial pattern of dispersal, whether that habitat is connected to source or sink habitat, the level and pattern of synchrony between sites, and so on. In other words, it depends on the type of detailed information that has traditionally been difficult to obtain for metapopulations of conservation concern. Interestingly, although it is difficult to determine how much change in λ can be achieved, it appears possible to estimate where the largest $d\lambda$ from a given $d\mu$ change (per unit of habitat) in the local growth rate is achieved, even though the size of the resultant $d\lambda$ cannot be determined.

Recall that each row of **A** represents a unit of habitat and that a local population is composed of some set of units of habitat with high connectivity. When the intrinsic growth rate, μ_j, in a unit of habitat j is changed by $d\mu$, to
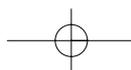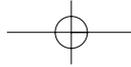
$\exp(\mu_j + d\mu)$, all the a_{ij} elements of column j in matrix $\mathbf{A}(t)$ are multiplied by $\exp(d\mu)$. The goal is to calculate the total change in λ from this $d\mu$ change to all elements in column j by summing over rows i:

$$d\lambda = \sum_i \frac{\partial\lambda}{\partial\mu_{ij}}d\mu = \sum_i \frac{1}{\lambda}\frac{\partial\log\lambda}{\partial\log a_{ij}}d\mu \qquad (23.26)$$

The term $\partial\log\lambda/\partial\log a_{ij}$ is the elasticity of λ. Caswell (2001) presented the calculation for the elasticity of λ for products of stochastic matrices:

$$\frac{\partial\log\lambda}{\partial\log a_{ij}} = \lim_{t\to\infty}\frac{1}{n}\sum_{t=0}^{n-1}\frac{a_{ij}v_i(t+1)w_j(t)}{R(t)\mathbf{v}^\mathrm{T}(t+1)\mathbf{w}(t+1)} \qquad (23.27)$$

where $R(t)$ is the relationship between the right eigenvector and $\mathbf{A}(t)$, $R(t)\mathbf{w}(t+1) = \mathbf{A}(t)\mathbf{w}(t)$. Thus, the $d\lambda$ from a $d\mu$ change in a unit of habitat j can be solved for by summing Eq. (27) over i:

$$d\lambda_j = \frac{d\mu}{\lambda}\lim_{n\to\infty}\frac{1}{n}\sum_{t=0}^{n-1}\sum_i \frac{a_{ij}v_i(t+1)w_j(t)}{R(t)\mathbf{v}^\mathrm{T}(t+1)\mathbf{w}(t+1)} \qquad (23.28)$$

This can be translated into matrix form:

$$[d\lambda_1 \quad d\lambda_2 \ldots d\lambda_k] = \frac{d\mu}{\lambda}\lim_{n\to\infty}\frac{1}{n}\sum_{t=0}^{n-1}\frac{\mathbf{w}^\mathrm{T}(t)\circ\mathbf{v}^\mathrm{T}(t+1)\mathbf{A}(t)}{R(t)\mathbf{v}^\mathrm{T}(t+1)\mathbf{w}(t+1)} \qquad (23.29)$$

where "$\circ$" denotes the scalar product. Using the relationship between the left eigenvector and $\mathbf{A}(t)$, $Q(t+1)\mathbf{v}^\mathrm{T}(t) = \mathbf{v}^\mathrm{T}(t+1)\mathbf{A}(t)$,

$$\begin{bmatrix}d\lambda_1 & d\lambda_2 \ldots d\lambda_k\end{bmatrix} = \frac{d\mu}{\lambda}\lim_{n\to\infty}\frac{1}{n}\sum_{t=0}^{n-1}\frac{Q(t+1)\mathbf{w}^\mathrm{T}(t)\circ\mathbf{v}^\mathrm{T}(t)}{R(t)\mathbf{v}^\mathrm{T}(t+1)\mathbf{w}(t+1)} \qquad (23.30)$$

Thus $d\lambda$ from a change in a unit of habitat j is a weighted temporal average of the reproductive value of local population j times its density:

$$d\lambda_j = \frac{1}{n}\sum_{t=0}^{n-1}c(t)w_j(t)v_j(t) \qquad (23.31)$$

where $c(t)$ is a constant that depends on t but not j. A local population a is composed of units of habitat in the set $a = \{a_1, a_2, a_3, \ldots, a_m\}$, where $\{a_1, a_2, a_3, \ldots, a_m\}$ denotes which rows of $\mathbf{A}$ correspond to the units of habitat in local population a. The $d\lambda$ per $d\mu$ per unit habitat for a particular local population a is $d\lambda_a = (1/m)\sum_{j\in a}d\lambda_j$, where m is the number of units of habitat in local population a. In words, this means the change in λ is proportional to the product of the "average" density of individuals in a particular local population times the "average" reproductive value of its units of habitat.

Although reproductive values are unknown, there are many cases where the product v_jw_j is a positive function of v_j as long as dispersal is not too
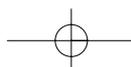
unidirectional (meaning, dispersal from A to B but not B to A). This can be shown analytically in three extreme cases: (a) 100% uniform and equal dispersal, (b) all μ_j values equal, or (c) dispersal extremely low. In cases (a) and (b), the reproductive values are all equal and $v_j w_j = $ (a constant) $\times v_j$. In case (c), $w_j \approx v_j$ and $v_j w_j \approx (v_j)^2$. However, this positive relationship was also found in simulations with variable local growth rates, neighborhood dispersal, and dispersal sources and targets. An obvious exception to this positive relationship is if dispersal is unidirectional, for example, a linear chain of local populations with dispersal via a steady directional wind or ocean current. However, as the following simulations illustrate, the general relationship can still hold even when dispersal is strongly, although not strictly, directional.

Density and λ Sensitivity

Three different types of metapopulation models were used to look at the relationship between average densities in units of habitat versus the $d\lambda$ from a small increase in the local growth rate in each unit of habitat. In each model, dispersal was nonuniform among the local populations so that some sites were dispersal sources (more dispersal out than in) and others dispersal targets (more dispersal in than out). In the first model, local growth rates were equal among all sites and dispersers were distributed globally among all sites. In the second model, local growth rates were variable so that some sites had much higher local growth rates than others and dispersal was mainly to nearest neighbors. In the third model, local growth rates were again variable and dispersal mainly to the two south and east neighbors; however, a small proportion of dispersers were distributed globally. Thus the three examples illustrate global, local, and directional dispersal.

A hundred randomly generated matrix models in each of these three categories were made and $d\lambda_j$ calculated via Eq. (30). Figure 23.7 shows the relationship between the average density of a local population and the $d\lambda$ from increasing the local growth rate in that unit of habitat. The x axis ranks the $d\lambda_j$, thus "1" indicates the local population with the highest $d\lambda_j$ in any simulation and "49" the lowest. The y axis shows the corresponding mean density rank of that local population; "1" indicates the population had the highest density among the 49 sites and "49" the lowest. Results from the 100 randomly generated models are summarized by showing a box plot, which shows the median and range of all density ranks for the sites with a given $d\lambda_j/d\mu$ rank. Thus, the box plot at the x axis position "1" shows the range of density ranks for the units of habitat with the highest $d\lambda/d\mu$ in each model. Model results show a strong positive relationship between the relative density rank within a unit of habitat and which unit of habitat produced the largest increases in the metapopulation λ for a given $d\mu$. The two to three units of habitat with the highest average densities were consistently the units that produced the largest $d\lambda$ for a given $d\mu$. This suggests that plotting the distribution of the relative densities within local populations in a metapopulation could give a rapid indication of the sensitivity of the metapopulation to changes to individual local populations.

Uniform Local Growth Rates
Dispersal Sources and Targets with Global Dispersal

Variable Local Growth Rates
Dispersal Sources and Targets with Neighborhood Dispersal

Variable Local Growth Rates
Dispersal Sources and Targets with Directional Dispersal
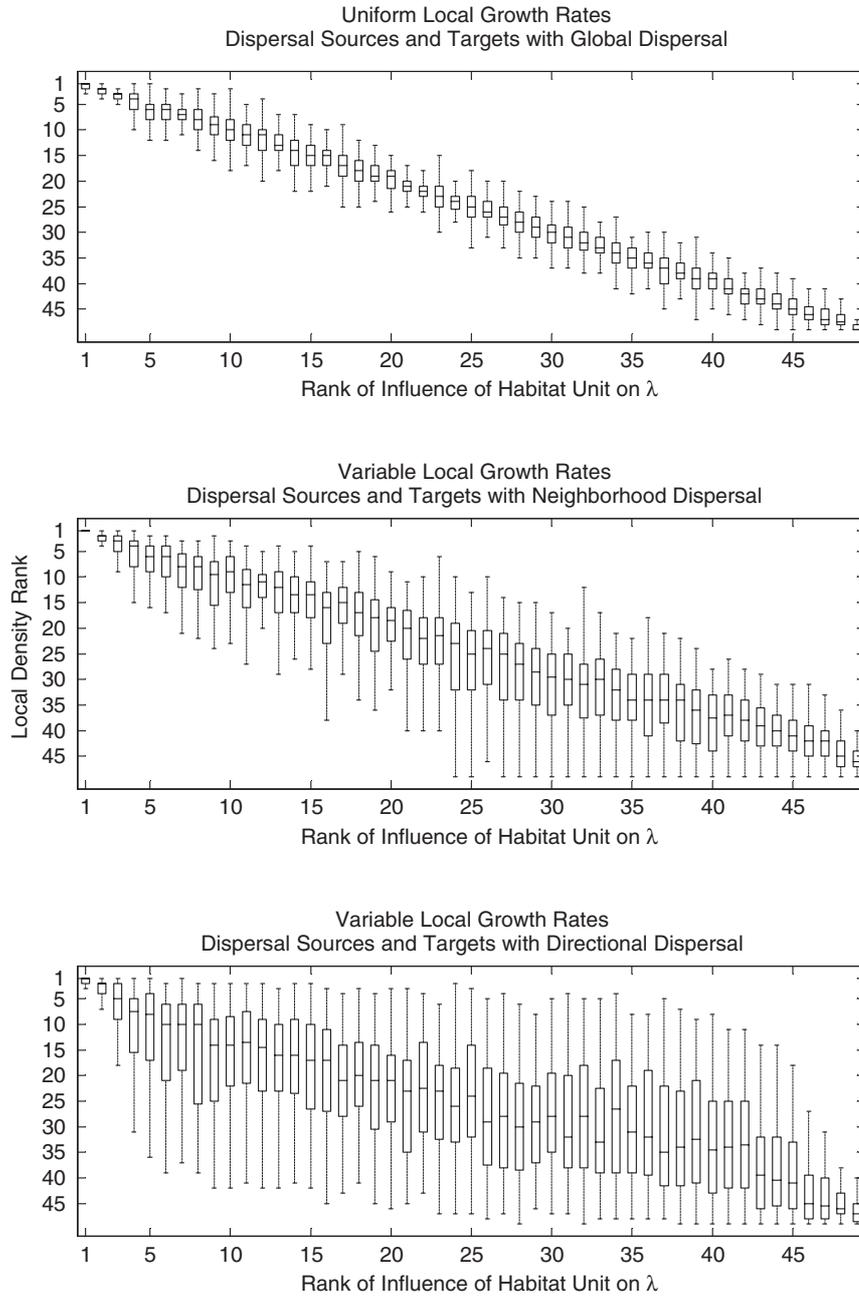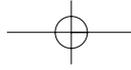
Fig. 23.7   Relationship between the influence of a given habitat unit on the metapopulation λ and the average density in that habitat unit. One hundred 7×7 metapopulations with spatially variable dispersal rates (some sites dispersal sinks and others targets) were generated randomly in each of three classes: (1) spatially uniform growth rates and global dispersal, (2) spatially variable growth rates with neighborhood dispersal, and (3) spatially variable growth rates with directional neighborhood dispersal to the S and E two neighbors only. The *x* axis shows the rank in terms of $d\lambda/d\mu$, and the *y* axis shows a box plot of the distribution of density ranks for sites with a given $d\lambda/d\mu$ rank across all 100 models in each class. Thus the box plot at $x = 1$ shows the distribution of ranks for the sites with the highest $d\lambda/d\mu$ in each model. The line in each box shows the median density rank for the sites with a given $d\lambda/d\mu$ rank, the box encloses 50% of the ranks, and the whiskers show the range from all 100 randomly generated models.
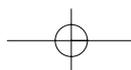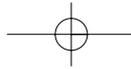
One application of this would be to estimate where negative impacts would lead to the greatest decrease in λ, thus suggesting where protection in most critical. It would also suggest where local improvements would be most effective for a given increase in the local growth rate. However, in actual management situations where improvements are being sited, one is generally trying to maximize the "bang per buck," $d\lambda_j/d\$ = d\lambda_j/d\mu_j \times d\mu_j/d\$$. The cost, $d\$$, is the actual monetary cost or some combination of monetary, logistical, and political costs and $d\mu_j/d\$$ is the cost of a unit improvement to a unit of habitat j. Thus $d\lambda_j/d\mu$ is one part of the equation, and the other part, the cost of a unit improvement in different habitats, would have to come from a specific analysis of the costs and estimated effects of management actions on different local populations.

Example Using the Snake River ESU

The overall level of salmon dispersal between and among stocks within this ESU is known to be fairly low and spatially localized (Mathews and Waples, 1991; Quinn, 1993). In addition, there is high variability in the habitat quality between stocks, with some stocks relatively pristine and protected within wilderness areas, whereas others are exposed to high and multiple impacts (such as stream degradation and disturbance, pollution, in-stream harvest, and irrigation impacts). Figure 23.8 (top) shows the distribution of average normed redds per mile for 50 Snake river spring/summer chinook stocks. For each year between 1980 and 1995, the redds-per-mile count for each stock was divided by the maximum count among the 50 stocks in that year. The average over the 16 yr was then used as an estimate of the average normed redds per mile. The long-tailed distribution is the expectation from theory given low dispersal and high variability in stock habitat quality.

Estimation of the average normed redds per mile was repeated using a variety of different time periods. Regardless of the time period or number of years used for averaging, six stocks consistently appeared among the top five stocks with the highest density of redds: Johnson Cr., Poverty Cr., and Secesh R. in the south fork of the Salmon R., the Lostine R. in the Grande Ronde subbasin, Marsh Cr. in the middle fork of the Salmon R., and the Imnaha R. Perhaps not surprisingly, all of these are in relatively low impacted regions of the ESU. At a subbasin level, the overall highest redd density was in the south fork of the Salmon river where summer-run chinook primarily occur (Fig. 23.8, bottom). The other regions are primarily spring-run chinook. The south fork of the Salmon river is relatively pristine and few hatchery fish have been released into this subbasin; the stocks presumably have experienced relatively low interbreeding with hatchery-reared stocks. In addition, the later run timing may somehow be associated with less straying, lower harvest, or lower hydropower impacts.

This analysis predicts that the λ of the Snake river spring/summer chinook ESU would be most sensitive to changes to the summer-run stocks in the south fork of the Salmon river and to the spring-run stocks, the Lostine R., Imnaha River, and Marsh Creek and should be protected preferentially from impacts. This can be counterintuitive in some situations. For example, imagine making choices about where to allow a limited catch-and-release sport

Snake River spr/sum chinook
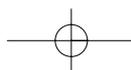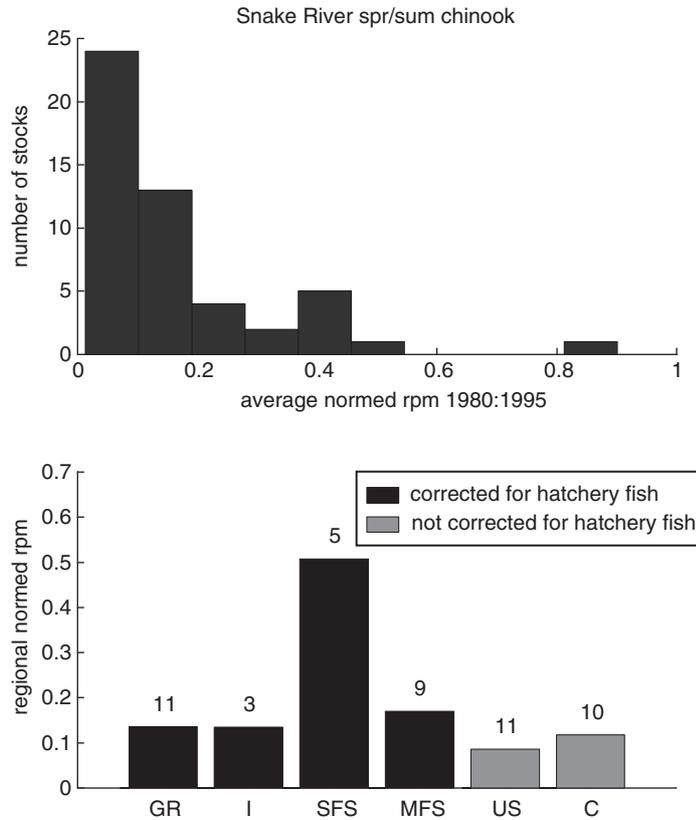
Fig. 23.8   Distribution of densities of redds in the Snake river spring/summer chinook ESU at a stock and subbasin level. The average normed redds densities (top) are shown for the 50 stocks with 1980–1995 data (the years were chosen to maximum the number of stocks with data). For each stock the normed redd density was averaged over the 16 yr to get an estimate of the normed average density. In the lower plot, relative average densities over all stocks within different basins are shown (with the number of stocks in each basin shown above the bars). The basin designations are GR, Grande Ronde; I, Imnaha; SFS, south fork salmon; MFS, middle fork salmon; US, upper salmon; C, Clearwater. Redds due to hatchery fish released into stocks were removed before doing these analyses, as the density will be artificially high simply due to hatchery fish releases. This correction could not be done for the upper salmon or Clearwater regions because the fraction of spawners that are hatchery strays were unknown; however, the hatchery releases are very high in these basins and thus the corrected relative densities would be much lower than shown.

fishery. Sites with the highest density would seem to be the prime candidates, whereas the analysis of $d\lambda/d\mu$ indicates just the opposite. In terms of determining where to direct improvements, the $d\lambda/d\mu$ suggests that these pristine sites are where a given $d\mu$ would produce the greatest metapopulation λ; however, the regions where $d\lambda/d\mu$ is the highest are not necessarily the regions where μ is improved most easily. Indeed a given unit of improvement may be more difficult in pristine sites. Choosing where to direct stock improvements requires consideration of the cost and difficulty of a given $d\mu$ for different stocks in combination with the estimate of the sensitivity of λ to local changes.
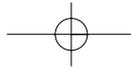
23.11 POPULATION VIABILITY ANALYSIS IN PRACTICE

The purpose of this chapter is to present a theoretical framework for metapopulation PVA using time series data and diffusion approximations. These methods are then illustrated using data from two salmon metapopulations. The salmon analyses are intended as an example of how to calculate the diffusion parameters and metrics. An actual PVA must grapple with other important issues that are outside the scope of this chapter, but which anyone contemplating an actual PVA must be aware. Morris and Doak (2002) gave a review of the criticisms and caveats surrounding the use of PVA and outlined general recommendations and cautions when conducting a PVA. In the context of diffusion approximation methods in particular, Holmes (2003) outlined an approach using matrix models to conduct sensitivity analyses in order to choose among different parameterization methods and metrics for a specific PVA application.

One of the issues that is especially pertinent for our chapter is the issue of variability in estimated risk metrics. A number of recent PVA cross-validations using actual data on a large number of different populations have shown that careful PVA analyses give unbiased risk estimates (Brook et al., 2000; Holmes and Fagan, 2002; Fagan et al., 2003). Although this is very encouraging, a difficult issue is the high inherent variability associated with estimated probabilities (such as probability of extinction), even though they may be unbiased (Ludwig, 1996, 1999; Fieberg and Ellner, 2000; Holmes, 2001; Ellner et al., 2002). How to properly use risk metrics that have high variability is currently being debated within the field with arguments ranging from "don't use them" (Ludwig, 1996, 1999; Fieberg and Ellner, 2000), to "use to estimate risks within collections of populations" (Fagan et al., 2001; Holmes and Fagan, 2002), to "use where data are extensive and high quality" (Coulson et al., 2001), to "PVA metrics based on data, even if variable, are better than the alternatives" (Brook et al., 2002). An encouraging aspect of diffusion approximation methods is that cross-validations using real time series data have indicated that the uncertainty in the estimated metrics appears to be characterized properly (Holmes and Fagan, 2002). Nonetheless, how to use and present metrics with high variability, albeit well characterized, is not an easy question to answer. Presentation of $100(1-\alpha)\%$ is an oft-used approach, but experience in the forum of salmon recovery planning in the Pacific Northwest has shown that it is easy to misinterpret confidence intervals. For example, it is easy to interpret 95% confidence intervals for λ that overlap 1.0 as an indication that data are equivocal as to whether the population is declining or increasing, whereas there may be considerable data support for a declining population. Graphic presentations of data support for different risk levels have been more compelling and informative, although translating levels of data support into numbers that policy makers can use to take uncertainty into account in policy decisions has been challenging.

23.12 DISCUSSION

This chapter focused on the calculation of metapopulation PVA metrics; however, there are other more general PVA insights from an examination of stochastic metapopulations and of this specific class of declining density-independent

metapopulations. First, by definition the trajectory of a stochastic metapopulation is subject to random processes and thus the metapopulation trajectory observed in any one snippet of time is unlikely to capture the long-term dynamics. The shorter the time frame, the farther the observed trend is likely to be from the long-term trend. Thus the trends in any two adjacent time periods are unlikely to be identical, and the difference indicates not necessarily a change [AU3] in the underlying rate of decline, can be due simply to chance. The variability of observed rates of decline can be estimated from the level of the variability driving the long-term dynamics, and thus statistical tests performed to determine the likelihood that an apparent change in trend occurred due to the stochastic nature of the process rather than an underlying change in conditions.

Second, the local populations within a metapopulation are linked and experience the same long-term growth rates, regardless of the underlying difference in local population conditions (i.e., whether they are "sources" or "sinks"). However, measured over a short time period, there will be differences in the observed local population trends due to chance and local conditions. This means that over a given time period, local populations will appear to be declining at different rates, but this is not an indication the long-term trends and not necessarily related local conditions being better or worse than other areas. That the long-term trends of the individual local populations are the same as the metapopulation has a direct impact on PVA for local populations within a metapopulation. The rate of decline observed among the different local populations will differ, as will the apparent level of variability in the local time series. Thus if an individual viability analysis is done using parameters estimated from local population time series alone, it will appear that there is tremendous variability among the local populations risk levels when in fact their long-term risks are similar. When looking at the long-term risks, use of metapopulation level parameters leads to better estimates of the long-term local population risks. Short-term risks, however, are still strongly influenced by local conditions. Clearly estimates of both short-term and long-term risks are needed to capture the whole viability picture for a metapopulation. Although local populations within the type of metapopulations modeled here will be eventually repopulated by dispersal if they undergo extreme declines, the resulting loss of genetic diversity leads to a gradual erosion of the genetic health of the metapopulation. Indeed this has happened for salmon species throughout the Pacific Northwest.

Recovery planning for endangered and threatened species typically requires determining where to put the most effort. Rarely is it the case that maximal effort can be applied everywhere. Using the stochastic metapopulation model, a sensitivity analysis was used to look for local characteristics that predict where local changes would produce the biggest change in the metapopulation growth rate. Interestingly, local density (not absolute numbers) was a strong predictor of where a unit change in local growth rates led to the largest metapopulation growth rate. This relationship was observed even in simulations with dispersal sources and targets and strongly directional dispersal, although it will break down when dispersal is strictly unidirectional. Determining which local populations are best suited for restoration efforts also requires assessing the feasibility, cost, and acceptance of restoration efforts. Indeed when it comes to actually implementing recovery actions, optimizing
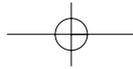
the efficiency of effort in terms affecting recovery requires solving a complex function of biological, economic, and political information. However, understanding the population dynamics of the species of concern and gaining insight regarding how the demography of the species will respond to alternative management actions are fundamental and primary components of this conservation equation.
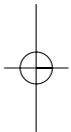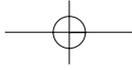
[AU1] Sentence OK?
[AU2] (−0.05 OK)
[AU3] Sentence OK?