

CHAPTER ONE

**WHAT IS POPULATION VIABILITY ANALYSIS,
AND HOW CAN IT BE USED IN CONSERVATION DECISION MAKING?**

What is Population Viability Analysis?..... 2

Potential products and uses of PVA..... 2

 PVA Use #1: Assessing the extinction risk of a single population..... 3

 PVA Use #2: Comparing relative risks of two or more populations 4

 PVA Use #3: Analyzing and synthesizing monitoring data 5

 PVA Use #4: Identifying key life stages or demographic processes as management targets..... 6

 PVA Use #5: Determining how large a reserve needs to be to achieve a desired level of protection from extinction 7

 PVA Use #6: Determining how many individuals to release to establish a new population..... 7

 PVA Use #7: Setting limits on the harvest or “take” from a population that are compatible with its continued existence 8

 PVA Use #8: Determining how many (and which) populations are needed to achieve a desired overall likelihood of species persistence 8

Types of Population Viability Analysis 9

A “roadmap” to this book 12

Our modeling philosophy: keep it simple..... 14

Tables 17

WHAT IS POPULATION VIABILITY ANALYSIS?

Broadly defined, population viability analysis (PVA) is the use of quantitative methods to predict the likely future status of a population or collection of populations of conservation concern. By “future status”, we typically mean the likelihood that the population (or the total number of individuals across all populations) will be above some minimum size at a given future time. That size may be zero if we simply want to know if the population will avoid outright extinction, or it could be a larger number if we want to have warning of impending disaster in time to take aggressive action to try to save the population. More indirect indices of future status, such as simply whether the population is expected to grow or decline in the coming years, are also valuable to conservation planners. Assigning concrete numbers to these measures of future status is the aim of population viability analysis. As we use the term in this book, PVA refers to a wide range of methods to analyze data and link the results to models of population growth and decline.

Having read the foregoing description of PVA, your first impression may be similar to the one that most people have of the field of accounting, i.e. that it sounds rather dry and esoteric. But as with accounting, a little more thought should lead you to acknowledge that PVA can also be incredibly useful¹. To help convince you of this utility, before we say anything about how PVA is actually performed, we begin by listing specific conservation goals that a viability assessment can help to achieve. That is, we first address the important question: for what reasons might we wish to make quantitative predictions about future population status?

POTENTIAL PRODUCTS AND USES OF PVA

There are many uses for projections of future population state. In Table 1.1, we group eight of these uses under the headings of assessment (in which our goal is simply to ask how well the population is doing) and management (in which our goal is to determine what interventions will reduce the population’s likelihood of extinction). A recurring theme of this book is the reciprocal interplay between these two major uses of PVA. To illustrate these uses, we briefly

¹ And, we believe, more interesting and valuable than accounting!

describe specific examples in which PVA has been applied to populations of threatened and endangered animals and plants in the real world.

PVA Use #1: Assessing the extinction risk of a single population

If there is only a single remaining population of a threatened or endangered taxon, we would obviously want to know whether that population is likely to decline, and thus whether we should intervene to save it. We'd also like to know how soon extinction is likely to occur, as it would help us to determine the necessity of immediate intervention. Even if there are many extant populations of a rare species, conservation biologists charged with preserving one of those populations might profit from a quantitative assessment of its future state. A good example of such a population is the grizzly bears in the Greater Yellowstone Ecosystem (Yellowstone National Park, USA, and the surrounding National Forest lands). Grizzlies are common in Alaska and Canada, but because of its unique status as one of the few remaining populations in the lower 48 United States, the Yellowstone population is currently listed as threatened under the U.S. Endangered Species Act.

In one of the first quantitative assessments of population viability, Mark Shaffer (1978, 1981, Shaffer and Sampson 1985) used computer simulations to ask whether the Yellowstone grizzly bear population was large enough to have at least a 95% chance of surviving for different periods of time into the future. He found that, given the number of bears in the population at the time of his analysis, persistence for 100 years was likely, but that the chance of persisting for 300 years was far more bleak. A larger population, and hence a larger area, would be needed to achieve a reasonable likelihood of long-term persistence. This analysis, and a bevy of subsequent PVAs of the same population (Knight and Eberhardt 1985, Eberhardt et al. 1994, Doak 1995, Pease and Mattson 1999) have, via law suits as well as less confrontational means, had a significant impact on management policy (for example, they resulted in reduced clear-cutting and mining development on the Forest Service lands that comprise the majority of the Greater Yellowstone Ecosystem) and have helped to prevent the premature removal of the Yellowstone grizzly from the ESA's "threatened" list. Shaffer's quantitative approach to evaluating whether a single population was large enough to make near-term extinction unlikely also inspired many

similar analyses for other taxa, and spurred development of the new field of population viability analysis.

Another highly influential analysis of a single population was Russell Lande's (1988a) assessment of the Northern Spotted Owl (*Strix occidentalis caurina*), then under consideration for listing as threatened under the U.S. Endangered Species Act due to logging of the old-growth forests on which the owl depends. Treating the entire subspecies as a single population, Lande used a relatively simple model with age structure and implicit spatial dynamics (as described in Chapters 6-9 of this book) to ask whether the population was declining in the face of logging pressure. Although Lande's best estimate of the rate of population growth suggested that the population was in fact in decline, the data were insufficient to rule out the possibility of a static or even a slowly growing population. Lande's finding that the population was most likely to be declining played a role in the eventual listing of the subspecies as threatened, with a consequent reduction in logging activity on federal lands within the owl's range. The uncertainty around Lande's estimate of the population growth rate also served a useful function in indicating that more data needed to be collected to more accurately assess the population's viability. In fact, given the typical scarcity of data for the rare species that are most in need of PVAs, one of the benefits of population models is that they can help us to determine which types of data would be most useful to collect in the future.

PVA Use #2: Comparing relative risks of two or more populations

Perhaps of even greater utility than a single measure of risk for one population is the ability to compare the risk of extinction among multiple populations of the same taxon. Such comparative estimates are useful in at least three ways. First, if we want to know if the entire taxon is at risk of extinction, we can ask whether all or most of its populations are declining. For example Forsman et al. (1996), motivated by Lande's (1988a) analysis of the entire subspecies, estimated that 10 out of 11 local populations of the Northern Spotted Owl they studied were declining significantly, strongly supporting the need to address the impact of logging on the subspecies. Second, if we seek to save all existing populations of a taxon (for example, because each harbors unique alleles) but have limited funding, measures of relative risk would allow us to prioritize which populations should receive management attention first. Third, if we have

determined that we need not or cannot save all populations, metrics of relative viability could form the basis for “triage” to select the most promising populations to preserve. Allendorf et al. (1997) advocated this use of PVA to determine which populations of Pacific salmon are most worthy of preservation. Having a tool to identify the healthiest populations may be especially useful to organizations such as The Nature Conservancy, that has the goal of preserving all native species within each of its large “ecoregions” (The Nature Conservancy 1997) but cannot afford to purchase lands harboring every extant population of every rare species.

PVA Use #3: Analyzing and synthesizing monitoring data

Monitoring information is frequently gathered on rare or potentially threatened species. For example, Recovery Plans for species listed under the U.S. Endangered Species Act (ESA) often specify that important population processes, such as survival, growth, and reproduction, be monitored to assess population health (Morris et al. 2002). Similarly, Habitat Conservation Plans, also implemented under the ESA, are required to contain a monitoring component, which ranges from estimates of relative population numbers to gathering of full demographic data (Harding et al. 2001a). Thus calls to collect monitoring data are common; unfortunately, carefully-considered plans to analyze these data in such a way as to inform management decisions are far less frequent. For example, how will we decide whether evidence of improvement or decline in monitoring data is sufficient to indicate a real change in imperilment? Moreover, what if monitoring indicates that the survival rates in a focal population are high relative to other known populations or species but that rates of reproduction are low? The population models that form the basis for PVA provide a way to combine such seemingly contradictory monitoring data into a single assessment of population health and to gauge how a change in a monitored population translates into a change in our assessment of its viability.

In a simple but very useful application of a PVA model to monitoring data, Leah Gerber and her co-authors (1999) asked how many years of population census data were needed to conclude with reasonable certainty that the eastern North Pacific gray whale population met criteria for delisting under the ESA. This formerly very rare population had increased rapidly from 1965 to the mid 1990s. But while delisting occurred in 1994, at that point there had been no quantitative assessment that the population was truly out of danger of near-term extinction.

Gerber et al. used the simplest of PVA methods (presented in Chapter 3) to show not only that the decision to delist was warranted, but also that it could have been made safely several years earlier.

PVA Use #4: Identifying key life stages or demographic processes as management targets

In addition to facilitating assessment of likely future population status, PVA is also a potentially powerful tool to identify effective management actions. For example, when different population processes or different life stages contribute differentially to population growth, conservationists will want to know which processes or stages would have the greatest impact on viability if targeted for management. Similarly, we need a means to assess whether a management strategy that affects multiple processes and stages (sometimes in opposing ways) is having the intended effect. Once again, population models can provide answers to these questions.

A now-classic example of the use of PVA to guide management decisions is the work of Deborah Crouse, Larry Crowder, and their associates on the loggerhead sea turtle (*Caretta caretta*), a threatened species that nests on beaches along the southeastern coast of the United States (Crouse et al. 1987, Crowder et al. 1994). Two major threats, trampling of eggs and hatchlings on beaches and drowning of older-aged turtles in fishing nets, were hypothesized to underlie declining numbers of loggerheads. PVA models showed that management efforts aimed solely at protecting eggs and hatchlings would not reverse declines, even if they were 100% effective. In contrast, installing TEDs (“turtle excluder devices”) into fishing nets to prevent drownings could result in a growing population even if the devices do not eliminate all fishing-related mortality. The ability to provide a quantitative projection of population growth under different management scenarios played an important supporting role in the decision to require the year-round use of TEDs by the shrimp-trawling fleet in the USA (Crowder et al. 1994).

PVA Use #5: Determining how large a reserve needs to be to achieve a desired level of protection from extinction

As Shaffer found in his analysis of the Yellowstone grizzly bear population, and as we will see repeatedly throughout this book, the current size of a population has a strong effect on its risk of extinction -- bigger is better. Of course, larger populations require more space. If we can estimate the amount of space each individual needs, we can use PVA models to ask how large a reserve needs to be to achieve a low probability of extinction for a particular population inhabiting the reserve². Armbruster and Lande (1993) exploited this aspect of PVA to ask how large national parks in semi-arid regions of Africa must be to buffer populations of African elephants (*Loxodonta africana*) from extinction in the face of droughts of varying frequency and severity. After considering uncertainty in some of their parameters, they conclude that at least 2500 square kilometers are needed for safe management of elephant populations.

PVA Use #6: Determining how many individuals to release to establish a new population

Closely related to the preceding use of PVA is the problem of determining the optimal number of individuals to release when attempting to establish new populations. Translocation and release of wild-caught or captive-raised individuals to sites where the species was once present, but has since been extirpated, is often used to initiate new populations and hence to reduce species-wide extinction risk. Such translocation programs face an inevitable tradeoff: given limitations on the number of individuals that can be generated by captive breeding programs or be safely removed from existing populations, releasing more of those individuals at one site reduces the number of new populations that can be established. Thus it is advantageous to have a way to ask how small a new population “inoculum” can be and still have a reasonable chance of successful establishment (measured as the probability that the new population will still be in existence at a future time horizon). PVA models offer a tool to quantify how establishment success might scale with the number of individuals released, and have been used to assess the likelihood of success for reintroductions of European beaver, wild boar, and capercaillie into

² Of course, it is important to preserve both the *quality* and quantity of habitat.

Scotland (Howells and Edwards Jones 1997, Marshall and Edwards Jones 1998, South et al. 2000) and release of bearded vultures in the Alps (Bustamante 1996).

PVA Use #7: Setting limits on the harvest or “take” from a population that are compatible with its continued existence

Populations of some species (e.g. ginseng) become a concern for conservationists when direct harvesting by humans depresses their numbers to low levels. Other species suffer from indirect human-caused sources of mortality, such as inadvertent capture in fishing nets (loggerhead sea turtles, short-tailed albatross, harbor porpoise), collisions with motorized vehicles (grizzly bears, desert tortoise, Florida manatee), and destruction of required habitat due to resource extraction (Northern spotted owl, red-cockaded woodpecker, Leadbeater’s possum), agriculture (cheetah, San Joaquin kit fox, blunt-nosed leopard lizard), and urban development (cougars, Santa Cruz long-toed salamander). The question often arises as to how much human-added mortality is compatible with the continued existence of a species or population. For example, the U.S. Endangered Species Act terms all human-related causes of mortality for listed species as “take,” and allows the taking of listed species under some circumstances. The quantitative models on which PVAs are based can be used to translate take into elevated extinction risk; this relationship is usually not straightforward. This approach has been used to identify acceptable levels for intentional harvest of both plants (Nantel et al. 1996, Ratsirarson et al. 1996) and animals (Powell et al. 1996, Heppell and Crowder 1996, Tufto et al. 1999), as well as inadvertent mortality caused by other human activities (Caswell et al. 1998).

PVA Use #8: Determining how many (and which) populations are needed to achieve a desired overall likelihood of species persistence

Finally, estimates of risk for separate populations can often be combined to calculate the probability that at least one population will still be in existence at a given future time. In this way, we can gauge how much safety from extinction an entire species would gain if we preserved particular numbers and combinations of extant populations. Such analyses are important both to sets of isolated populations and for metapopulations (archipelagoes of habitat

islands with populations linked by dispersal). For example, David Lindenmayer and Hugh Possingham (1996) analyzed the number and spatial arrangement of fragments of old growth montane ash forest in southeastern Australia that would have to be maintained uncut to ensure the continued existence of Leadbeater's possum. In a quite different circumstance, Eric Menges (1990) performed a simple version of a multi-population analysis for 15 populations of the endangered Furbish's lousewort (*Pedicularis furbishiae*), a plant that is restricted to the banks of a single river in Maine, USA. Populations in moist sites close to the river have a high risk of extinction due to catastrophic events such as ice scouring in winter and burial beneath collapsing river banks, whereas those in drier sites are likely to be overcome by the inexorable encroachment of other, competitively superior plant species. Thus virtually all local populations of Furbish's lousewort are predicted to go extinct in relatively short order, and the species cannot be preserved by protecting extant populations only. Rather, management of Furbish's lousewort must focus on assuring that the establishment of new populations is sufficient to compensate for local extinctions.

Just as frequently as they have been used to suggest preservation strategies, analyses of this type have been used to design strategies for restoring or reconstructing networks of critical habitat patches. Again, the northern spotted owl provides an example. Several different PVA analyses indicated that the U.S. Forest Service's original plan to regrow patches of the owl's old-growth forest habitat wouldn't work (Thomas et al. 1990, Doak 1989, Lamberson et al. 1992, 1994). The plan was to allow small, scattered stands of younger trees to regrow to old growth forest. The models all agreed that much larger patches of habitat, even if they were farther from one another on average, would result in much higher chances of the owl's persistence, a finding that resulted in a dramatic reworking of the final plan to recreate useable habitat for the species.

TYPES OF POPULATION VIABILITY ANALYSIS

The foregoing examples illustrate some of the many ways that quantitative models of population growth, which lie at the heart of PVA, can help to illuminate important questions in conservation biology. Although the acronym PVA is commonly used as though it signified a single method or analytical tool, PVAs are in fact based upon a range of data analysis and

modeling methods that vary widely both in their complexity and in the kinds and amount of data they require. We now describe the types of PVA models covered in this book.

At the simpler end of the spectrum lie models that predict only the total number of individuals in a single population. The data needed to fit these models to a particular population consist of either exhaustive counts or estimates of the total number of individuals in the population or in an easily-recognized subset of the population (such as territory-holding males or mothers with dependent offspring). These counts are obtained from a series of (typically annual) censuses. Throughout the book, we refer to analyses based only on the total number of individuals as “**Count-Based PVAs**”.

Count-based PVAs treat all individuals in the population as though they were identical. Yet we know that for species such as the loggerhead sea turtle, some individuals (e.g. mature breeders) are likely to contribute far more to future population growth than others (e.g. hatchlings, who must survive for many years in order to become reproductive), and that very different fractions of the population may be in categories with high versus low survival or reproduction (e.g., mature trees versus seedlings). Thus, we may be able to obtain a more accurate assessment of the viability of such a population if we account for these differences in individual contributions and for the fractions of individuals of each type currently in the population. Species in which individuals differ substantially in age, size, developmental stage, social status, or any other attribute that affects their contributions to population growth are said to have “structured populations”. Building a PVA for a structured population requires more information than can be obtained from simple counts of the total number of individuals in the population. Specifically, we must estimate the rates of important demographic processes (survival, growth, and reproduction) separately for each type of individual in the population. Estimating these rates typically requires us to conduct a demographic study in which we mark individuals of each type and follow them for several years, each year recording whether they survived, their type (e.g. age or size), and the number of offspring they produced. Thus the data needed to perform a structured PVA are more expensive and labor-intensive to collect, but they allow us to explore how well management techniques aimed at different types of individuals will fare, as we have seen for the loggerhead sea turtle. In this book, we refer to PVAs that require data from a demographic study as “**Demographic PVAs**”. While demographic PVAs can yield more informative predictions for most species, this is only true when we have enough data to

estimate the many parameters these models require. With less information, the simpler count-cased approach can give much more reliable results, even for species with highly structured populations.

At their simplest, both count-based and structured PVAs describe single populations, and they do not keep track of the actual spatial locations of individuals or of suitable habitat. Yet for species such as Furbish's lousewort that exist as metapopulations, assessing the fate of a single population will tell us little about the species' vulnerability to regional or global extinction. Moreover, for organisms such as the northern spotted owl and Leadbeater's possum that are threatened by habitat fragmentation, the spatial arrangement of remaining fragments, and hence the ease with which individuals can disperse between them, can strongly dictate population viability. Thus there is often a need for PVA models to explicitly include more than one local population or patch of suitable habitat. In this book, we refer to such models as "**Multi-site PVAs**". This category encompasses a broad range of model types. "Economy-class" multi-site models simply represent a set of isolated populations of a single species, with no movement of individuals between them, and are useful for calculating the probability that at least one population will persist for a given length of time. At the extreme opposite end of the spectrum from the simplest count-based PVAs lie so-called "spatially-explicit individually-based" (or SEIB) models, which track the actual positions of all individuals as they are born, move, reproduce, and die on a detailed habitat landscape (sometimes constructed using geographical information systems (GIS) databases). Not surprisingly, these highly complex multi-site PVAs have the most demanding data requirements of all PVA models. Not only does one need to estimate the contributions of each type of individual to population growth, one also needs to quantify their detailed movement behaviors, as well as the locations of different types of habitat on the landscape. Consequently, highly detailed multi-site PVAs will only be made possible by extraordinary data collection efforts. As such efforts will not be a feasible option for most rare species, we do not devote much space in this book to SEIB models. We do, however, describe informative analyses that can be performed with a more easily attainable quantity of data using simpler multi-site models.

A “ROADMAP” TO THIS BOOK

This book follows a sequence from simpler to more complex PVA models. Chapter 2 sets the stage for all the chapters that follow by describing several metrics that all types of PVA models use to gauge population viability. Chapter 2 also reviews general factors, both intrinsic to a population and imposed upon it by environmental forces, that influence any population’s risk of extinction. Each of the models we describe in subsequent chapters will attempt to account for some or all of these factors. Hence we recommend that readers unfamiliar with the causes of extinction and the ways that extinction risk is measured read Chapter 2 before proceeding to the more methodological chapters that follow. Because readers familiar with these concepts may choose to jump straight to the following chapter, we recapitulate important theoretical results at the beginning of Chapter 3.

In Chapter 3, we describe procedures to perform the simplest type of count-based PVA, specifically ones in which the change in the number of individuals in the population from one year to the next is assumed to be unaffected by the current population size (so-called “density-independent” models), and the impact of environmental conditions on the change in numbers is assumed to be relatively small. In Chapter 4, we broaden the scope of count-based PVAs by reviewing how several factors omitted in Chapter 3 (i.e., density dependence, correlations in environmental conditions between years, large environmental perturbations, and within-year variation among individuals in the contributions they make to population growth) can be included in the models. The material in Chapter 4 is somewhat more advanced than that in the adjacent chapters, and the data needed to incorporate these additional factors into count-based PVAs will often be lacking. Indeed, of the count-based PVAs done for 70 different species, all used the simpler approach presented in Chapter 3 (Elder et al. 2002). Therefore, we suggest that readers mainly seeking to understand the basics of common PVA methods may wish to skip Chapter 4 on first reading the book.

Chapter 5 takes up the important issue of observation error and how to account for its effects on measures of population viability for count-based models. Whenever population size is determined by a sampling procedure (e.g. aerial surveys, mark-recapture methods, or extrapolation from quadrat samples) rather than exhaustive enumeration, vagaries of the sampling process will introduce variation into the counts. That is, the counts will only be

estimates of the true population size, and some of the variation in the counts from year to year will reflect variation in the sampling process rather than biologically meaningful variation that actually influences population viability. In addition to inflating estimates of variability, observation error introduces uncertainty into our estimated viability measures, and Chapter 5 reviews how to account for this uncertainty.

Chapters 6 through 9 review techniques for performing demographic PVAs. These analyses are based on models known as population projection matrices or demographic matrix models. Most published PVAs fall into this category, and many authors still use ‘PVA’ to strictly mean this type of quantitative analysis. Chapter 6 shows how to use data from a study in which marked individuals were followed through time to construct a projection matrix model. In Chapters 7 and 8, we demonstrate how to calculate measures of growth and viability for a structured population using a projection matrix model. In Chapter 9, we examine ways to use projection matrices to explore management options for structured populations, as was done, for example, for the loggerhead sea turtle. We note that in a general book on PVA methods, we cannot possibly cover the topic of projection matrix models as thoroughly as does Caswell (2001) in his comprehensive book *Matrix Population Models: Construction, Analysis, and Interpretation*, to which we refer readers interested in learning more about these models. Nonetheless, our book gives a complete introduction in how to construct matrix models and perform the most useful analyses needed for a demographic PVA.

Chapters 10 and 11 focus on multi-site PVAs. In Chapter 10, we describe and we show how to measure two additional factors, dispersal of individuals and spatial environmental correlations, with which we do not concern ourselves in the preceding chapters but that may need to be estimated and incorporated when the joint viability of more than one population is being assessed. Chapter 11 covers techniques for actually performing a multi-site PVA. As noted above, we concentrate on analytical methods most likely to be useable in many circumstances, eschewing the many spatial models that require extraordinary amounts of data to be useful.

All conservation biologists are familiar with the following conundrum. It is difficult to make an accurate prediction of the future status of a population with only a limited amount of data. Yet it is precisely the rarity of threatened and endangered species that makes it both difficult to obtain a large quantity of population data and especially desirable to have some sort

of population assessment. As a result, practitioners of population viability analysis will be perennially forced to operate in an arena in which data scarcity is the rule rather than the exception. In recognition of this reality, we will repeatedly argue in this book that measures of population viability should be viewed not as iron-clad predictions of population fate but as “works in progress” subject to updating as more data become available. But when data on a particular species are truly scarce, performing a PVA may do more harm than good by engendering a false sense of rigor. In such cases, basing conservation decisions on other methods makes far better sense. In the final chapter of this book, we discuss cautions and caveats that help us to decide when it is useful, and more importantly when it is NOT useful, to perform a PVA. Thus we emphasize from the outset that while we view PVA methods as a potentially useful set of conservation tools, we do not see PVA as a panacea for all conservation problems.

OUR MODELING PHILOSOPHY: KEEP IT SIMPLE

Population biologists have developed a vast array of complex and mathematically sophisticated models, many of which can be adapted to predict the likelihood of population extinction. This book does not attempt to review all of those models. Instead, we focus on the subset of all available PVA methods that we deemed to be the most practical given the types of data typically available for species of concern to conservation biologists. In making the decision of which methods to include, we used two simple and related rules of thumb that we think could be said to capture our general philosophy regarding the use of population models in conservation biology.

The first rule is: “let the available data tell you which type of PVA to perform”. More specifically, we should not seek to build a PVA model that is more complex than the data warrant. It is our view that when data are limited (as they almost always will be when we are dealing with the rare, seldom-studied species that are the typical concern of conservation planners) the benefits of using complex models to perform population viability analyses will often be illusory. That is, while more complex models may promise to yield more accurate estimates of population viability because they include more biological detail (such as migration among semi-isolated populations, the effects of spatial arrangement of habitat patches, and the

nuances of genetic processes such as inbreeding depression, gene flow and genetic drift), this gain in accuracy will be undermined if the use of a more complex model requires us to guess at critical components about which we have no data. Instead, our philosophy is that the choice of models and methods in PVA should be determined primarily by the type and quantity of data that are available, and not the desire to include all interesting and possibly important processes and rates. It is better to use a simple approach (keeping the simplifications in mind) than to construct a complex house of cards that relies on numbers with no empirical justification.

The second rule is: “make sure you know what your model is doing”. This rule will be easier to follow if we also heed the first rule, that is, if we keep our models simple. This second rule is germane to the question of whether one should build one’s own PVA model or use one of the software packages designed specifically for PVA that are now widely available, such as ALEX (Possingham and Davies 1995), GAPPS (Harris et al. 1986, Downer 1993), INMAT (Mills and Smouse 1994), RAMAS (Ferson 1994, Akçakaya and Ferson 1992, Akçakaya 1997), ULM (Ferrière et al. 1996), or VORTEX (Lacy et al. 1995)³. Careful use of these programs can certainly lead to a defensible viability assessment. However, naïve users of the programs run the risk of making two errors. First, without fully understanding the underlying models used by the programs, users may build into the program incorrect assumptions about the biology of the species or population under study, leading to incorrect estimates of population viability. For example, some of these programs incorporate density dependence in very specific ways that may not be appropriate for the organism under study (Mills et al. 1996). Second, inexperienced users may be lured by the array of options proffered by the software into including risk factors about which no data exist, thus violating our first rule. We hold that, even if one ultimately plans to use one of the software packages to perform PVAs, it is essential that one begins by learning how to build one’s own model from scratch, so as to fully understand what the packages are doing. In this book, we demonstrate through worked examples how to go from raw data to a fully parameterized population model, and then how to use the model for assessment and management. We illustrate this process by showing actual statistical analysis of the raw data, and by providing computer code that actually runs the models. Our computer programs are

³ For comparisons involving some of these programs, see Lindenmayer et al. 1995, Mills et al. 1996, and Brook et al. 2000b.

written in MATLAB⁴, a matrix language that is relatively easy to learn and that is particularly well suited to demographic PVA models. Readers who are not familiar with MATLAB and would like a quick introduction to its use in an ecological context should consult the appendix in Roughgarden (1997). We strongly urge readers to try to understand the structure of the programs we present as a way of truly knowing how the underlying models work. To make the programs as accessible as possible, we pepper them with extensive comments that indicate what the code is doing at each point. Readers who wish to write their own programs in another language can treat the comments as a form of "pseudocode". These programs are not simply another form of "canned" model in which all one needs to do is plug in some numbers. Rather, each program performs only a subset of the tasks involved in a complete PVA. In most cases, you will need to combine them, and even modify them, to suit the needs of your particular PVA. Indeed, a strength of building your own PVA model from scratch is that it can be far more flexible than any pre-packaged software would allow. Nevertheless, the shards of MATLAB code presented in this book should allow you to rapidly do all the operations necessary for most PVAs.

We end by emphasizing that the worked examples in the subsequent chapters employ data from actual populations of rare or declining species. That is, we do not use common species for which data are abundant simply to "beautify" the outcome of our analyses. In this way, we aim to illustrate the constraints of assessing population status with limited data, constraints that are a realistic and unavoidable feature of population viability analysis.

⁴ MATLAB is a widely-available commercial product. It is relatively inexpensive for students and relatively expensive for others (see <http://www.mathworks.com/> for current prices and licensing arrangements). We use it because of its strong advantages for matrix manipulations and, indeed, for compactly programming many other population models. It has also increasingly been adopted by many theoretical ecologists as a lingua franca. There are several less expensive options that still allow you to reap the advantages of MATLAB. For example, SCILAB, a shareware program that uses MATLAB-like syntax, can be downloaded at <http://www-rocq.inria.fr/scilab/>.

TABLES

Table 1.1: Potential uses of PVA “products”

Category of use	Specific use	Examples
Assessment of extinction risk	Assessing the extinction risk of a single population	Shaffer 1981, Shaffer and Samson 1985, Lande 1988
	Comparing relative risks of two or more populations	Forsman et al. 1996, Menges 1990, Allendorf et al. 1997
	Analyzing and synthesizing monitoring data	Menges and Gordon 1996, Gerber et al. 1999
Guiding management	Identifying key life stages or demographic processes as management targets	Crouse et al. 1987
	Determining how large a reserve needs to be to gain a desired level of protection from extinction	Shaffer 1981, Armbruster and Lande 1993
	Determining how many individuals to release to establish a new population	Bustamante 1996, Howells and Edwards Jones 1997, Marshall and Edwards Jones 1998, South et al. 2000
	Setting limits on the harvest or “take” from a population that are compatible with its continued existence	Nantel et al. 1996, Ratsirarson et al. 1996, Tufto et al. 1999, Caswell et al. 1998
	Deciding how many populations are needed to protect a species from regional or global extinction	Menges 1990, Lindenmayer and Possingham 1996