

although currently some consensus is apparently beginning to emerge (section 5.4).

An interesting thing happened while the regulation debate was raging. First, empirical ecologists began collecting long-term data on population fluctuations of a wide variety of organisms. It is curious that a lot of long-term data sets were started during the 1940s and 1950s (i.e., just when the debate was at one of its peaks!). Next, quantitative ecologists started analyzing these time-series data (Moran 1953; Bulmer 1974; Berryman 1978; Royama 1981; Potts et al. 1984; Turchin 1990) using, in the beginning, such linear approaches as the Box-Jenkins time-series analysis. Then, ecologists (most notably, Robert May) participated in the nonlinear dynamics revolution (Gleick 1988). When physicists invented the new technique of attractor reconstruction in time-delayed coordinates (Takens 1981; Packard et al. 1980), some ecologists began applying it to ecological time series (Schaffer 1985). Classical time-series analyses and nonlinear dynamics approaches were eventually merged in a synthetic approach to the analysis of ecological data (these approaches will be discussed in part II), and applied to issues ranging beyond mere density dependence. Presently, we are seeing how these nonlinear time-series methods are being merged with the theoretical tradition (see chapter 8), and there are also promising beginnings of the synthesis between the population-regulation analyses and experimental approaches (Cappuccino and Harrison 1996).

1.2 GENERAL PHILOSOPHY OF THE APPROACH

Most ecologists do their science without giving much thought to the broad philosophical issues underlying what they do. Among those ecologists who do worry about philosophical foundations, the most vocal, and not afraid of making strong recommendations, are the Popperians (e.g., Chitty 1996; Murray 2000; Lambin et al. 2002). Other ecologists take the view that there are many ways of doing ecology, and one should not be too dogmatic about it (e.g., Fagerstrom 1987; Pickett et al. 1994). I believe that such philosophical discussions are important, because they affect how we do ecology. Furthermore, one

of the broad themes of this book is methodological (see the preface): what are the best approaches to solving the puzzle of population cycles? Thus, I need to describe the philosophical basis of the general approach that I advocate.

While Popper's idea that all theories have to be testable in order even to be called scientific seems quite reasonable to me, I find the rest of his philosophy of science, at least as expounded by his ecological disciples, not to be a very useful way of doing science. I am particularly bothered by the emphasis of Popperians on falsificationism as *the* way of doing science. First, the view that data are "hard facts" is untenable for methodological and psychological reasons (see Fagerstrom 1987 for a very clear discussion of this point). Thus, it is not true that in any contest between theory and data, it is theory that should necessarily lose. Second, I don't think that ecologists are in the business of rejecting theories. "Ecologists, like many others, do not reject theories for the futile reason that they are wrong; theories are retained until better ones emerge" (Fagerstrom 1987). A very good idea of how futile a rejectionist program can be is conveyed by the book of Dennis Chitty (1996), *Do Lemmings Commit Suicide?* There Chitty relates how a consistent application of the rejectionist approach led him to reject all hypotheses that could be tested, leaving him with the explanation that nobody could figure out how to test.

I think that we (ecologists) are, instead, in the business of deciding which of the available alternative theories is the best, or "least wrong" (I shall make this idea more precise later in this section). One thing that any scientist has to come to terms with is that all our theories are, in the final account, wrong (the alternative of not being wrong is to become untestable, that is, nonscientific). The more explicitly we formulate our theories (which, at least in the context of population dynamics, means translating them into mathematical statements) the more wrong they become, simply because our simple theories can never capture all the complexity and detail of nature. So falsifying theories is trivial: just collect detailed data about any aspect of the theory, and you are certain to show that the theory is wrong. If you have not, it simply means that you either collected too few data points or did not measure them carefully enough.

If all our theories are a priori wrong, what can we do? Well, science is still the search for truth, but any scientific truth that we find is

both approximate and tentative. Approximate, because of the reasons discussed in the paragraph above; tentative, because we have no guarantee that somebody smarter or possessing better data and analytical tools will not come up with a better "truth" sometime in the future. Therefore, we should not be in the business of rejecting theories, as ecological Popperians would have us do, but in the business of contrasting two or more theories with each other, using the data as an arbiter. The corollary of this approach is that our best theory may not explain or predict data very well, but we should still use it until we have something better. Even the theory that explains only 10% variation in the data is useful, because it sets a standard to be bettered.

In the rest of the section, I make this idea precise for the specific context of population dynamics. The basic notions are three: (1) define very carefully what you are trying to explain; (2) translate your verbal theories into explicit mathematical models (note the plural here); and (3) use formal statistical methods to quantify the relative ability of the rival models to predict data. Data may already be available, or they may be specifically collected to distinguish between predictions of the rival hypotheses (the latter constitutes *an experiment*).

1.2.1 Defining the Phenomenon to Be Explained

The broad question that I address in this book is, why do population numbers change with time? Or, to put it more succinctly, "why do populations behave as they do?" (Royama 1992:1). In any particular case study, this broad question can be broken into more specific issues. First, are dynamics of the studied population characterized by a stationary distribution of densities? (This is the issue of population regulation.) If yes, there is some characteristic mean level around which the population fluctuates, and fluctuations are characterized by a certain (finite) variance. What ecological mechanisms are responsible for setting this mean level? (This is the focus of population statics.) What mechanisms set the amplitude of fluctuations? Finally, are there detectable statistical periodicities, and what is the order and trajectory stability characterizing dynamics?

At the most general level, the phenomenon to be explained is quantified by a temporal record of population fluctuations, or time-series

data. Time-series data are often available even before the beginning of the formal inquiry into dynamics of a particular population (although we often have to do with an index of population rather than an absolute measure of population density; examples include fur returns, bag records, and pheromone trap catches). If time-series data are not available, a systematic program for their collection should be initiated immediately. (One should not worry too much about limited usefulness of short time series; after all, it may take many decades to approach the solution, by which time time-series data will be long enough to be useful!)

I will call the density measurements of the "focal species" (the one whose dynamics we are trying to understand), $\{N_t\}$, the **primary data**.¹ We may have time-series data on other aspects of system dynamics available (e.g., temporal changes in mean body mass, fluctuations in the availability of food, and densities of predators or parasitoids). Such **ancillary data** may be extremely useful, but are secondary in the sense that we do not require that our explanation of the focal species dynamics would account for all of them. For example, if we are studying a forest defoliator, then a model based on plant quality does not need to explain why parasitism rates vary (perhaps parasitoids are simply responding to the oscillations of their food supply, without a detectable feedback effect on defoliator densities). Vice versa, a parasitism-based explanation does not need to account for changes in plant quality. Of course, the model based on a particular factor has to be consistent with time-series data for this factor.

A focus on the primary data permits us to use the same metric when comparing hypotheses based on very different factors. One particular metric that I will use extensively is the coefficient of prediction, R^2_{pred} (the proportion of variance in log-transformed density explained by the hypothesis). However, this is not the only metric that can be employed to quantitatively compare the performance of different hypotheses. Another approach is to first quantify the observed dynamical pattern with **probes** such as the period and amplitude of oscillations (and others, see section 6.2.2), and then to determine how well rival models predict the numerical values of probes.

¹Concepts emphasized in boldface type are defined in the glossary.

Defining the problem the way I do here is not the only way to study population cycles. One alternative way, as practiced by the Canadian school (Chitty, Krebs, Boonstra, and others), is to define a "cyclic syndrome," which includes such features of dynamics as rapid changes of population density, systematic variation in body weight with the phase of the cycle, and perhaps certain changes in behavior, such as aggressiveness versus "docility." The problem with this approach is twofold. First, the "cyclic syndrome" is often defined without reference to whether population dynamics are characterized by a periodicity or not. In the cases where population dynamics are not periodic (or when there are no long-term data to determine whether there is periodicity or not), we find ourselves in a situation of studying a population "cycle" that is not a cycle by any formal definition. The second problem is that by including in the definition changes in individual quality and behavior, the Canadians tilt the field in favor of their favorite hypotheses. Suppose, for example, that periodic dynamics in a particular rodent population are driven by an interaction between rodents and their food supply. By measuring such processes as food requirements of rodents and growth dynamics of vegetation after being consumed, we may be able to construct an empirically based model that would predict the cyclic changes in rodent numbers (the primary data) very well. However, since we have not explicitly dealt with the physiological or behavioral responses of individuals to food scarcity or abundance, the model will say nothing about systematic changes of body weights with the cycle phase. Thus, the model will fail to explain the "cyclic syndrome." We could, of course, include such individual responses in the model. But this would be done at the expense of complicating the model structure, with the only yield an explanation of what really are side effects of population cycles—that individuals would be of low weight and fight more when food is scarce and population density is collapsing. This argument suggests to me that we should give logical preeminence to the primary data. I repeat that this does not mean that we should ignore various kinds of ancillary data, but neither should we necessarily aim at a theory that explains every bit of data collected about the focal population.

Before leaving the subject of problem definition, I want to reiterate that in this book I focus exclusively on nonspatial aspects of

population dynamics. I recognize that movement and spatial dynamics are very important. However, one has to start somewhere, and the magnitude of the task—disentangling the mechanistic causes of temporal oscillations in any particular system—is already enormous.

1.2.2 Formalizing Hypotheses as Mathematical Models

Having defined the explanandum (what is to be explained), I now turn to the explanans (the means of explanation). The question is, what ecological mechanisms underlie temporal change in natural populations? This issue is at the core of the book.

But what do I mean by ecological mechanisms? I believe that the most useful approach to understanding population dynamics is the reductionist one. Thus, the mechanistic basis for population ecology should be provided by the properties of entities one hierarchical level lower than populations, that is, by the behavior and physiology of individual organisms (Metz and Diekmann 1986; Caswell et al. 1997): individual consumption, growth, and reproduction rates; the probabilities of being killed by a predator or succumbing to a pathogen; characteristics of individual movement; and so on. I believe that such **methodological individualism** is a valid principle, but in practice it is not always possible, nor desirable, to follow this reductionist program to the logical extreme.

For example, when studying a predator-prey interaction, we need not follow each individual predator while keeping track of its size, sex, hunger level, spatial position, and so on. We might instead summarize this wealth of information with just a few numbers, for example, the number of predators in each size class at any given time, or the density of predators in each patch, or even, most simply, the density of all predators. The mapping here is "many to one," because many potential descriptions in individual terms will map to a single number or set of numbers at the population level. Thus, an understanding of predator-prey dynamics may be approached in two steps. In the first step, the investigator performs a careful study of the individual predation process and attempts to summarize it with simple relationships, such as the functional response curve. In the second

step, functions summarizing behavior and physiology of individuals serve as building blocks in a population dynamics model.

Ecological **mechanisms**, as used in this book, can refer both to detailed descriptions of what individuals do and to functions summarizing salient features of individual behaviors. I agree that it is more satisfying to build fully mechanistic explanations of population dynamics that are firmly based on what individuals do. However, it is not necessary to do it in one step. The history of population ecology shows that such concepts as "population density," "functional response," and "density-dependent population growth rate" turn out to be very useful conceptualizations for connecting population dynamics to individual-based explanations. Thus, we should continue employing these concepts, while keeping in mind their limitations.

The next step is to decide how to connect specific ecological mechanisms to testable predictions. I will require that the answer take the form of a **fully specified model**. The main reason for this requirement is that translating each rival hypothesis into an explicit model will allow us to perform quantitative cross-comparisons between different hypotheses. In other words, we shall be able to say which model explains the data better.

Constructing a fully specified model is done in three steps. First, we choose the mathematical framework and, most important, the **state variables**. Mathematical framework is often suggested by the biology of the system. For example, if we are dealing with a forest defoliator who has one generation a year, then we should probably use the discrete (difference) equations. If, on the other hand, we deal with a large ungulate population, in which the time step at which reproduction occurs (one year) is a small fraction of an average life span, then a continuous differential equations framework provides a good approximation (we might also consider adding seasonality explicitly to the model).

State variables are typically determined by the verbal hypothesis on which the model is based. For example, if we are modeling the interplay between the individual quality and dynamics, then a minimal model would have two state variables: population density and average individual quality. If we think that taking an average of quality is too restrictive, then we might explicitly model discrete quality classes (e.g., the numbers of "poor"-quality and "high"-quality individuals).

Alternatively, we might employ a partial differential equations framework, and model variation in individual quality smoothly. The choices of mathematical framework and state variables are not independent of each other.

The second step is to choose **functional forms**. These are specific functions that relate state variables and their rates of change to each other. One example is the functional form of the self-limitation term—we could choose to model it using the logistic model, or theta-logistic, and so forth. Another example is the functional response: depending on what we know about the modeled system, we may choose Type I, II, III, or ratio dependence.

The third step is determine the values of **parameters**. Examples are the intrinsic rate of population increase, the carrying capacity, the searching rate, the handling time, and so forth. This task can be accomplished in three basic ways. One is to use the information about the natural history of organisms to deduce the parameter values or, more likely, to deduce the interval where plausible values should be found. The second approach of obtaining parameter values is by fitting models to time-series data (see chapter 8). The disadvantage of this method is that if we wish to use the time-series data to test model predictions, such a test would not be as rigorous, since a degree of circularity is involved. The third way is to design a short-term experiment and directly measure the parameter. This is the best way, but the most laborious one. A short-term experiment may also be designed to measure a whole function, thus providing the empirical foundation for the functional form choice.

These three steps take the model builder progressively from general to specific issues. With each successive step the freedom of choice (or the degree of arbitrariness) increases. The choice of state variables is largely determined by the nature of the hypothesis and the mathematical framework. For functional forms, we usually have a greater latitude, but we usually are limited to discrete choices (e.g., should we use Type II or Type III functional response?). Of course, one could use a qualitative approach; that is, instead of choosing a specific function, one could just say that a function should be monotonically increasing. Such approaches, however, are more useful in building general theory (examples: Rosenzweig 1969; Oksanen et al. 1981) than in the analyses of specific case studies (but see Ellner

et al. 1998 for *semimechanistic* approaches). Finally, parameter values typically vary continuously, and are least constrained by a priori considerations.

For all the above reasons, any single hypothesis can in principle be translated into an infinite number of fully specified models, depending on the choices we make at each stage. This means that rejecting a specific model in favor of another model based on a rival hypothesis may indeed indicate that the rival hypothesis is closer to the truth, but it may also indicate that we did not use the correct functional form in the first model, or perhaps misestimated a key parameter. This is not a lethal problem, since all scientific knowledge is approximate and tentative, but we should keep this caveat in mind.

1.2.3 *Contrasting Models with Data*

Once we have translated a set of competing hypotheses into models, we are ready to start the process of reducing this set to fewer (ideally, one) "winners." For example, suppose we are trying to understand why the population system we are studying exhibits a periodic second-order oscillation (dynamical classes are explained in chapter 5). The first basic test that each model has to pass is the ability to generate the qualitative type of dynamics characterizing the system, a periodic second-order oscillation in our example. Some models simply cannot generate second-order cycles (e.g., one-dimensional differential equation models cannot exhibit cyclic behaviors no matter what functional forms and parameters we use). We immediately eliminate such models, and by implication the hypotheses on which they are based, from the set of plausible explanations of the system's dynamics. The elimination of the verbal hypothesis is somewhat tentative, because it still may be possible to translate the hypothesis into a model (perhaps using a different mathematical framework) that would be able to generate the required qualitative type of dynamics. In any case, no rejection is final (just as no confirmation is final). However, if we do our best and still cannot translate the hypothesis into a model that generates cyclic dynamics, then we shall succeed in throwing a very grave shadow of doubt on the hypothesis. It will now be up to the advocates of the hypothesis

(if there are such) to show that the hypothesis can be translated into a model that generates cycles, and that this translation can be done in a biologically reasonable way.

The next hurdle that each model in the set has to clear is the ability to predict the correct type of dynamics for biologically reasonable parameter values. Lacking good information, we may have to use a wide range for plausible values of some parameters. In such cases, the fact that the model passes this test does not count very heavily in its favor. However, we may find (by numerical explorations) that a certain parameter or parameters are critical for the ability of the model to generate the right dynamics. This means that we have a model prediction that may be tested with an experiment.

Once we are finished with these qualitative tests, we may find ourselves in a situation that none of the models managed to pass them. This means that we have to go back to the drawing board and exercise our creativity again. No cut-and-dried guidelines for generating new hypotheses exist (except, perhaps, Edison's famous dictum about 10% inspiration and 90% perspiration), which is what makes science interesting! Alternatively, there may be only one model still standing. This is a rather happy outcome, since it means that we are essentially done. Not everybody is likely to be satisfied with the conclusion, but it is no longer sufficient simply to advance a verbal hypothesis as an alternative explanation. Having a fully specified model that predicts the correct qualitative dynamics with biologically plausible parameter values substantially raises the stakes for any potential challenger—any alternative hypothesis will have to do at least as well.

A more likely outcome is that two or more models will be able to pass the qualitative tests. This means that we need to subject the remaining hypotheses to quantitative and, ultimately, experimental tests. The most rigorous and objective approach to quantitative testing is to construct the fully specified models using only ancillary data (ideally, by performing focused short-term experiments to quantify functional forms and parameter values), and then use each model to predict the primary data. Models can be compared by (1) how well they predict actual population densities and (2) how well they predict quantitative measures of population dynamics (the probes). Additionally, (3) models must describe the dynamics of other variables on which they are based, and (4) their parameters and functional

forms need to be consistent with what we know about the system biology. Finally, (5) since a simpler explanation is always preferable to a more complex one, models with fewer parameters are given more weight than complicated models. Comparisons (1) and (2) can be translated into a single metric, allowing us to establish a ranking order for the models. Issue (5), although seemingly dependent on investigator judgment, can actually also be incorporated into the overall measure, using approaches based on information criteria, such as AIC (Burnham and Anderson 1998). Issues (3) and (4), by contrast, are difficult to translate into a common quantitative metric. For example, a parasitoid-host model for a forest defoliator may predict parasitism rates better than a food quality model for the same system predicts the changes in food quality. But this does not mean that these are grounds on which to prefer the parasitism model. Perhaps the food quality data are characterized by a higher measurement error. Similarly, one model may be able to predict the dynamics best for a rather marginal value of one of the parameters. But again, we have no common metric to downgrade this model in relation to others. This means that not everything can be formalized, and some aspects of model performance will have to be left to the judgment of individual ecologists.

CHAPTER 2

Population Dynamics from First Principles

2.1 INTRODUCTION

Ecologists rarely discuss the philosophical foundations of research into population dynamics. Foundational issues tend to work in the background shaping inquiry, and are rarely hauled out into the daylight to be closely examined (Cooper 2001). Several controversies in population ecology can be traced to a misunderstanding or a disagreement at the very basic level, for example, the density dependence debate (section 5.4).

My goal in this chapter is to explicitly discuss the logical foundations of population dynamics. The starting point is provided by a set of "self-evident truths" or postulates, from which much of the logical structure of the theory of population ecology can be derived. I will use these postulates to construct several foundational principles of ecological theory that we shall find particularly useful in explaining population oscillations. Furthermore, I will argue that one of the implications following from the principles is that there are three fundamental classes of population dynamics. This typology turns out to be of great practical utility in studying dynamics of specific populations. In particular, it provides a basis for classifying (1) qualitative kinds of dynamical behaviors and (2) the dynamical role of various ecological mechanisms. Additionally, the typology provides (3) the appropriate null hypotheses for testing ideas about population dynamics. In sum, the foundational principles and dynamical classes together