1 Models in Context

1.1 Complexity and Obscurity in Nature and in Models

This book deals with the use of data and models that can enhance understanding and contribute to prediction. These two goals are complementary. Both involve inference, and model analyses can take the form of predictive distributions. For environmental scientists, the challenge stems from the fact that natural and managed systems are high-dimensional, meaning that many interacting forces are at work (Levin 1998; Clark 2005; Clark and Gelfand 2006a). Much of nature is unmeasurable, unobservable, or both. Much cannot be manipulated. Faced with obscure, complex, and uncontrolled processes, environmental scientists have long recognized the need for abstraction (Schaffer 1981; Caswell 1988). Theoreticians and experimentalists attempt to extract the important relationships from nature so they can be studied in a controlled setting. Ecologists write models with only a few variables and parameters. They design experiments with only a few treatments.

The need to simplify on both the theoretical and experimental sides leaves a gap that can isolate those who analyze ecological models from those who collect and draw inference from data. This gap makes it difficult to test theory with data and to model data in appropriate ways (e.g., Oreskes et al. 1994). The goal of this book is to describe methods that can help to bridge the gap, starting from concepts that underlie traditional process and statistical models, and moving toward modern techniques that allow for deeper integration. This introductory chapter starts with some background and motivation.

1.1.1 Why Ecological Models Are Low-Dimensional

Attempts to abstract key features of a process are an important component of all scientific disciplines. From the conceptual (theory) side, this abstraction is accomplished with process models that contain few variables and parameters. High-dimensional process models are intractable; without simplified models, they cannot be analyzed to yield transparent relationships. Complex process

models are difficult to apply beyond the context for which they were developed. If we can abstract the important elements and develop a simple model thereof, analysis might allow us to understand how the process behaves and why.

The simplification needed to describe systems mathematically often requires assumptions that cannot apply to ecological data. Theorists may speak in terms of processes that apply everywhere, in a general sense, but nowhere in particular. In the light of the complexity mismatch between theory and the real world, it is not surprising that mathematical models are often viewed as irrelevant (e.g., Simberloff 1980). In over a decade of teaching mathematical models I do not need one hand to count the number of times that a basic model described in ecological textbooks has been directly applied to a student data set. When it comes to models, irrelevance can be the price of tractability.

1.1.2 Why Statistical Inference Is Low-Dimensional

Traditionally, the statistical analysis needed for inference and prediction was possible only for data collected under a rigid design. Here again, simplification is achieved through use of model systems. Statistical models simplify with assumptions such as each observation is independent of all others; uncertainty is typically allowed only for a response variable; the variables must be observable, in the sense that we can assume that values assigned to predictor variables represent truth. To meet these assumptions, experiments rely on strict design.

By contrast, ecological data are typically complex and interrelated in space and time (Ellner and Turchin 1995; Scheffer and Carpenter 2003). Broad-scale and long-term interactions with superimposed high frequency and difficult-to-measure fluctuations are pervasive. Only small parts of a process readily submit to experimental control and typically only in terms of highly abstracted experimental designs (Carpenter 1996; Brown 1995; Skelly 2002; Hastings 2004; Clark and LaDeau 2005). The small amount of variance in data that is explained by ecological models reflects the high dimensionality of nature, the many interacting processes that affect observations. The fact that simple models typically account for a small fraction of the total variance, even in experimentally controlled settings, leads to obvious questions: Are the important factors included in the analysis? Do the assumptions of independence and uncertainty confined to the response variable impact inference? Can the experimental results be extrapolated to nature?

Because many important processes cannot be studied in controlled experiments, there is a tendency to overlook model assumptions and apply statistical models where assumptions are violated. Spatio-temporal aspects of data and most sources of uncertainty tend to be ignored. Whether nature is abstracted to the point where simple model assumptions can be satisfied or assumptions are violated in the analysis of natural settings, the barriers between data and theory can be large.

The contrast between the high dimensionality of natural systems versus the simplicity that is manageable in experiments and models explains part of the historic gulf between ecological models and data. Throughout, I suggest that simple process models can be powerful. A goal of this book is to demonstrate

that this gulf can often be bridged, but only if process models can remain simple enough to be tractable, and data/parameter models can be sufficiently sophisticated to allow for the high dimensionality of nature. Rather than design away natural influences, I emphasize inference based on relevant scales and settings; the focus is on bringing models to nature, rather than nature to models.

1.2 Making the Connections: Data, Inference, and Decision

The challenges of inference, prediction, and decision making faced by ecologists are shared by many disciplines. I use two examples to highlight some of the specific issues. To emphasize the generality of both the challenges and the emerging tools needed to address them, I draw the first example from medicine. The second example comes from community ecology.

1.2.1 Example: Soft Data, Hard Decisions

Most information does not come from experiments designed by statisticians, and most decisions are subjective. Consider the example of a standard treatment of kidney stones with extracorporeal shockwave lithostripsy (Madigan and York 1995). The stone is located on a real-time ultrasound image by an operator who will focus shockwaves capable of disintegrating it. If the image is of high quality (I), there is an increased chance that it will be properly identified and disintegrated (D). If disintegrated, there is increased chance of clearance through the urinary tract (C). The decision framework involves interpretation of data with a flow outlined in Figure 1.1a. Note that the medical practitioner is concerned with the probability of successful clearance (C) given image quality (I), or $\Pr\{C|I\}$. Between these two events is the probability of disintegration (D). The decision is subjective, in the sense that, faced with the same evidence and choices, two practitioners could arrive at different decisions.

The basic elements of this problem are familiar to environmental scientists, managers, and policy makers. For instance, information is limited, but it can accumulate with time and experience. To many scientists, these data

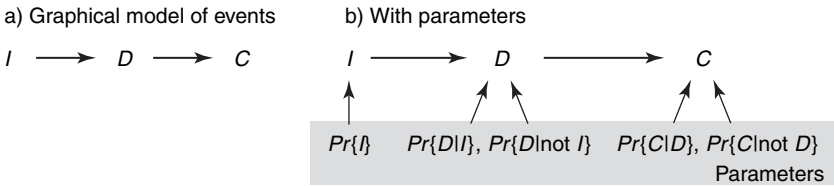

FIGURE 1.1. A graphical model of the kidney stone example (modified from Madigan and York 1995). In (*a*) is the graphical model of events that the image is of high quality (I), the stone is disintegrated (D), and the remnants are cleared through the urinary tract (C). To calculate $\Pr\{C|I\}$ we require five parameters (*b*). By making them stochastic (giving them priors), we can write simple rules for updating knowledge based on experience with new patients.

seem soft. The decision must be based on inadequate knowledge. We may not have the luxury of putting off a decision until hard data can be collected from, say, a series of controlled experiments. However, we would value a means for updating knowledge that can result in better decisions, that is, an adaptive management framework.

How can the practitioner use the model to learn from accumulated experience? The answer is, in large part, technical and a principal motivation of this book. But I provide a partial answer here. The parameters that influence decisions include the probability of C given that disintegration did or did not occur ($\Pr\{C|D\}$ and $\Pr\{C|\text{not } D\}$,[1] respectively), the probability of D given that image quality was good or bad ($\Pr\{D|I\}$ and $\Pr\{D|\text{not } I\}$, respectively), and the probability that the image was of good quality ($\Pr\{I\}$) (Figure 1.1b). The values of these five parameters determine probability of success. Clearly, the more we know about the values of these parameters, the more informed the decision. If we treat these parameters as being fixed values, there is no opportunity for learning. To allow for updating knowledge, the parameters are taken to be random. This random character of parameters allows for regular updating, with current understanding being further refined by each new experience (observation). The posterior knowledge taken from each patient becomes the prior for the next.

Whether the goal is increased understanding (as in, say, inference), prediction, or decision, the model graph provides a road map that facilitates not only modeling, but also computation. It emphasizes the importance of conditional probability, represented by arrows connecting events that are directly linked. In this particular instance, it describes a decision process that involves uncertainty that can be reduced though a prior-update-posterior cycle.

I use graphs like Figure 1.1 to represent models throughout this book. Early in the text, I use the convenient structure that involves data, process, and parameter submodels. This hierachical framework for submodels serves to decompose potentially complex problems into simple elements. By the end of this book I extend this framework to the more general notion of models as networks.

1.2.2 Example: Ecological Model Meets Data

A second example illustrates how the graphical framework of Figure 1.1 extends to ecological processes and brings in some of the challenges that confront ecologists attempting to integrate models and data. Ecological models predict that differences in how species respond to limiting resources can determine whether they can coexist (Levins 1979; Armstrong and McGehee 1980; Tilman 1982; Pacala and Silander 1990; Grover 1991; Pacala and Rees 1998; Murdoch et al. 1998). Is one species better able to exploit a limiting resource than another species? Does the advantage depend on how much resource is present? To evaluate the role of limiting resources, ecologists gather data describing growth responses at different resource levels.

[1] The notation | means "given that." Thus, $\Pr\{C|D\}$ means "the probability of event C given that event D has occurred."

Figure 1.2 shows data on seedling response (height growth) to a resource (light). The degree of scatter in Figure 1.2a and 1.2b is not unusual (Kobe 1999; Finzi and Canham 2000). Growth rate data are obtained from measurements of seedling height, together with estimates of light that penetrates the forest canopy. Full sunlight has a value of one, and complete darkness has a value of zero. Although the raw data do not show obvious differences between the two species, the fitted model says that the differences are highly significant, with 95 percent confidence intervals of each assigning near zero probability to the other (Figure 1.2c). It might seem paradoxical that these broadly overlapping data clouds are represented by significantly different models, yet this common situation is rarely mentioned. In fact, it is central to the interpretation of models.

Here is the standard analysis that led to this result. A model might include parameters that describe a minimum light requirement, a half-saturation constant, and a maximum rate, or asymptote. We could write the model as

$$y_{ijt} = \mu(x_{jt};\theta) + \varepsilon_{ijt}$$

$$\varepsilon_{ijt} \sim N(0,\sigma^2)$$

where x_{jt} and y_{ijt} are the predictor and response variables (light and growth) for the i^{th} individual at location j at time t, and μ is the saturating function of light availability x. θ represents the three parameters that determine μ, representing (1)

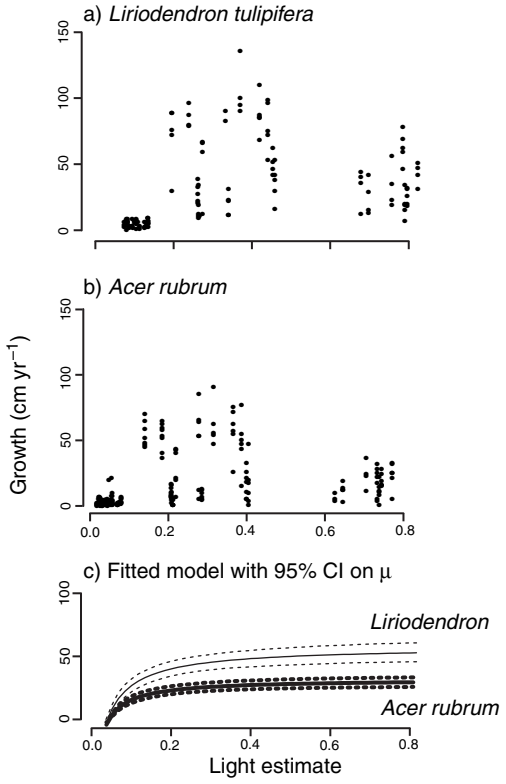

FIGURE 1.2. Growth rate responses to light availability. (*a*) and (*b*) show data obtained as measurements of seedling height growth and light availability. (*c*) shows a traditional model fitted to these data, the solid line being the estimate of μ. The dashed lines are 95 percent confidence intervals for μ. From Clark et al. (2003).

the minimum light requirement, (2) a half-saturation constant, and (3) the asymptotic growth rate. This saturating function does not fully describe the relationship between x and y—there is residual scatter. To accommodate this scatter, this function is embedded within a stochastic shell, which might allow for the fact that the model cannot be correct, and growth rates are not precisely known. Together, deterministic and stochastic elements are termed a *likelihood function*. This model says that y is assumed to have a normal distribution with a mean determined by the function μ and residual variance σ^2. This is the fitted model of Figure 1.2c.

The saturating function μ runs through the data cloud and is described by estimates of three parameters θ, which are represented by error distributions in Figure 1.3. The spread of these error distributions represents the level of confidence in parameter estimates, which depends, in turn, on the number of observations, the scatter, and the model. A 95 percent confidence interval can be viewed as the central 95 percent of the error distribution (although this is not its precise definition, as discussed in Chapter 5). The errors in these estimates are asymptotically zero in the sense that, as sample size increases, the confidence intervals (spread of error distributions) decrease. Hereafter, when I speak of asymptotics, I mean it in this sense of sample size. A predictive interval for the function μ is obtained by propagating the error in parameters θ to error in μ (dashed lines in Figure 1.3).

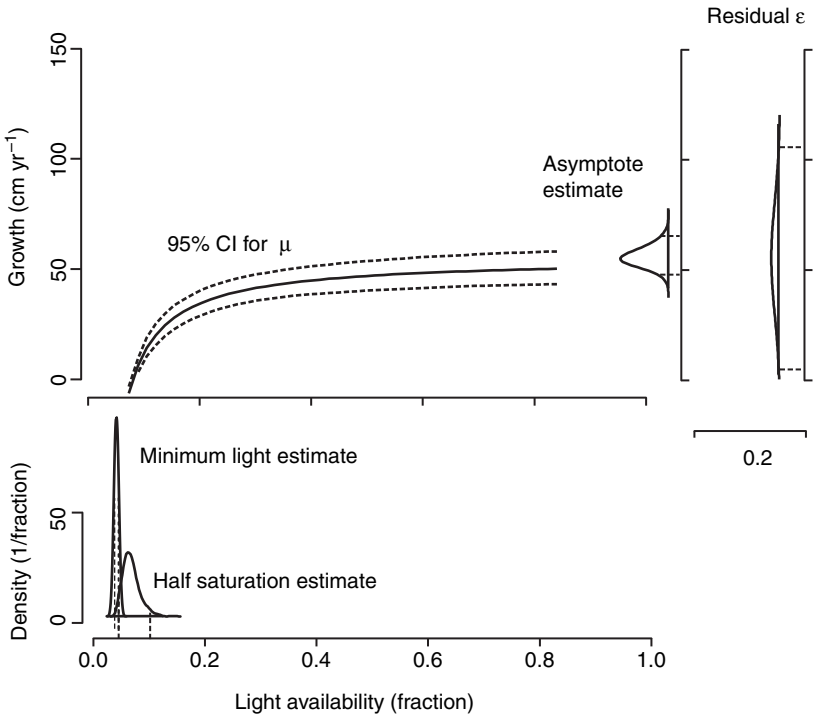

FIGURE 1.3. Elements of the traditional model used to analyze data in Figure 1.1. There are error distributions for three parameters and residual stochasticity in the error term ϵ (From Clark et al. 2003).

Obviously, the 95 percent confidence interval on $\mu(x)$ does not bound 95 percent of the data. The uncertainty in parameter estimates, which is used to construct the envelope for μ, is relatively small in comparison with variability in the data. The bulk of the variation is associated with the error term ε_{ijt}. This scatter is represented by a density shown on the far right-hand side of Figure 1.3. We could calculate a predictive interval for y that incorporates this scatter (it integrates ε) (Feiburg and Ellner 2003; Clark 2003), but that is rarely done. We would do this to predict as-yet unobserved growth response data y. But if this scatter is viewed as noise or error, we might not have much use for this predictive interval—it adds clutter, but not insight. Moreover, if we thought the scatter captured by ε_{ijt} was anything other than error, it might be difficult to justify this model in the first place.

So what is this leftover scatter? I began by saying that the scatter sopped up by ε_{ijt} might be associated with observation errors or model misspecification. But measurements of seedling height can be off by a centimeter. They are not off by a meter. So observation error is not the explanation for the broad scatter in Figure 1.2. If the deterministic part of the model μ is inadequate, we might increase its complexity in terms of a more flexible form or by including additional covariates that could explain the scatter. In either case, we require more parameters. In fact, ecologists have studied seedling growth many times, and measurements of more variables often do not explain much additional scatter. In other words, we often cannot account for this variability by increasing the complexity of the deterministic part of the model $\mu(x)$ (Clark et al. 2003a). If the scatter is not observation error, and we cannot accommodate it by incorporating more deterministic complexity, traditional methods do not leave many options.

In fact, there are many ways in which stochasticity might stand in for unobservable aspects of this relationship. For example, the light data x might be variable or imprecisely known. This brings an additional source of stochasticity and is sometimes termed an error-in-variables problem. I distinguish between the light seen by a plant at location j, x_j, and the observation of it using the notation $x_j^{(obs)}$. To allow that observations depend on the true light level, with uncertainty, I include a density for light observations, $x_j^{(obs)} \sim p(x_j,\phi)$, where ϕ represents any parameters that enter the model for observations. We now have the model

$$\begin{aligned} y_{ijt} &= \mu(x_{jt};\theta) + \varepsilon_{ijt} \\ x_{jt}^{obs} &\sim p(x_j,\phi) \qquad \text{uncertainty in } x \\ \varepsilon_{ijt} &\sim N(0,\sigma^2) \qquad \text{error in } y \end{aligned}$$

As in Section 1.2.1, we can represent this model with a graph. The basic model that ignores uncertainty in light is represented by Figure 1.4a. If implemented in a Bayesian framework (Chapter 4), I could refer to this as "simple Bayes." The stochastic element in this graph is represented by the connection between σ and y, indicating the error in y. In part b of this figure, there is an additional source of stochasticity associated with observations of x. The graph has increased in complexity to accommodate this additional relationship.

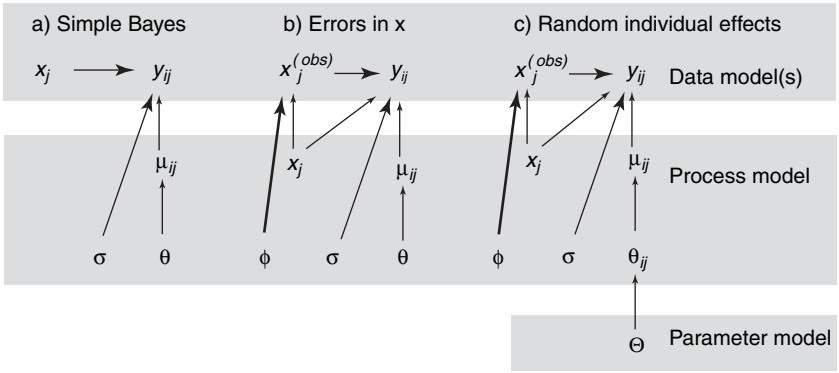

FIGURE 1.4. Four models that can be used to model growth response to light. Complexity increases from a simple model with error in y (*a*), to error in both variables (*b*), to variability among individuals (*c*). Modified from Clark (2005).

But there is still much more going on in Figure 1.2. We expect that individuals will have different responses to light, depending on many factors that cannot be measured. If individual responses result from unmeasured or unmeasurable factors, then we can include random effects in the model. In this case, there is a parameter vector that applies to each individual θ_{ij}. These individual parameters are drawn from a distribution having parameters Θ. This model is:

$$
\begin{aligned}
y_{ijt} &= \mu(x_{jt};\theta_{ij}) + \varepsilon_{ijt} \\
x_{jt}^{(obs)} &\sim p(x_j,\phi) \qquad \text{error in } x \\
\theta_{ij} &\sim p(\Theta) \qquad \text{random individual effects} \\
\varepsilon_{ijt} &\sim N(0,\sigma^2) \qquad \text{error in } y
\end{aligned}
$$

There are now $n+k$ new parameters that must be estimated, one θ_{ij} for each individual, and k parameters in θ, describing population heterogeneity.

The model is getting complex, and the list of potentially important influences is still not exhaustive. For example, we might include random effects associated with location, fixed or random effects for years, and there might be autocorrelation between years or with distance. We might have additional sources of information that are not yet accommodated by the model. Already the model is beginning to look like a network of elements, some of which are known and some unknown. This network perspective is readily depicted with graphs. The modules labeled as "Data," "Process," and "Parameter" on the right-hand side of Figure 1.4 help organize relationships.

With these basic principles in mind, I return to the traditional analysis, which superficially appears to have clarified a relationship involving broad scatter with a tight relationship between resource and growth (Figure 1.2). So where did the scatter go? And did we lose anything important? Using models available in standard software, there is a deterministic function $\mu(x)$ buried

in noise. The function $\mu(x)$ is like the mean response, and it is surrounded by error that is independent of $\mu(x)$. We now have the answer to the first question: the scatter was sopped up by a stochastic shell.

The second question is more difficult to answer. The insights used to extend the simple model came from recognition that observations of height growth are imprecise, light observations or treatments are variable and imprecisely known, individuals have different genotypes, and their responses to light depend on other factors that vary in time and space. Some of these factors can vary on such fine spatial and temporal scales that we could not hope to measure them. Because there are response times associated with many variables, we would have a hard time even deciding on a scale to measure them. If all of these factors contribute to the scatter, then the problem is high-dimensional.

How do the assumptions of the statistical model used to fit Figure 1.2c compare with the factors we identified on the basis of ecological insight? One could argue that at least one of these factors, observation error, is consistent with the statistical model (the deterministic-response-in-a-stochastic-shell approach). This would be the case if there were a deterministic response that applies to all individuals, and it was obscured by errors that result from sampling. The two might be independent if meter stick error does not depend on tree height. But the data suggest that this assumption does not apply here. The measurements could be off by perhaps a centimeter, but not by a meter.

Once we move beyond simple observation error, the deterministic-response-in-a-stochastic-shell model becomes less plausible. The other factors are not well described by this model, because their effects depend on the function $\mu(x)$ itself. Of course, we could make the stochasticity depend on $\mu(x)$ in a specific way by, say, modeling log values. Where such specific transformations can provide a quick fix, use them. For a number of reasons, this quick fix will not suffice here. If the estimates of the light level at which a plant grows are only approximate, then x is stochastic, and, in turn, $\mu(x)$ is stochastic; it is a function of stochastic x. If individuals have different responses to light due to genotype or factors that vary, then there is a function $\mu(x)$ for each individual i.

Is this simply nit-picking? How much of a difference can it make? Would we not fit roughly the same model, regardless of seemingly arcane decisions about what causes the scatter? Can these details affect the inference on the process of interest, example, competition and coexistence? The answer is, "it depends." We will confront this issue repeatedly. For now simply consider a few more points. First, the range of variability in growth described by the fitted function $\mu(x)$ in Figure 1.2c is a small fraction of the total (Figure 1.2a, b). If pressed to identify the dominant term in our model, we would have to say that it is the part we have relegated to the noise or error term, not to the signal $\mu(x)$.

The first point leads to a second. Inference is based on a model that differs substantially from our view of the relationship. If the scatter were modest we might ignore these factors, but how do we justify it here? If we are satisfied with sweeping under the rug the dominant part of the relationship, is there any

point to statistics? Do we have any more confidence in a predictive interval constructed from the traditional model than we would in a line drawn through the scatter by eye? It is worth mentioning that the prevailing confidence in ecological predictions is not great, in part because it is not clear how they relate to underlying processes and to the data used to fit them. The traditional view that the scatter is error underlies the interpretation that one species grows faster at all light levels (Figure 1.5a). In light of the assumptions, can we be sure that tulip poplar outcompetes red maple?

If we trust the data and we trust the theory (e.g., growth is a saturating function of light), then the statistical analysis is the weak link. Off-the-shelf software does not provide flexible options for most ecological data sets. In other words, it may be a bad idea to let default assumptions from canned software determine everything beyond the basic process model.

As preface to techniques covered in this book, consider what happens if we allow for the sources of variability that are consistent with ecological insight (uncertainties in observations, light, and the growth response $\mu(x)$). An analysis that admits these considerations suggests broad overlap (Figure 1.5b). It is true that we could also construct broadly overlapping prediction intervals for the classical approach in Figure 1.5a. This interval could be constructed by putting back the scatter (the stochastic shell) that we threw away to produce Figure 1.5a. I discuss this in later sections of the book. This is generally not done in ecological analyses, because modeling begins from the premise that everything other than $\mu(x)$ constitutes error and is independent of $\mu(x)$. Moreover, the predictive interval obtained by mixing in ε_{ijt} is not the one we obtain under the assumption that many factors contribute stochasticity, and many of them come through $\mu(x)$ itself.

Throughout this book I focus on the weak interface, the critical connection that is needed to evaluate relationships and to predict. Modern approaches use structure to allow for complexity and stochastic elements to stand in for uncertainty. Getting this connection right helps make theory relevant, and it allows us to exploit data fully.

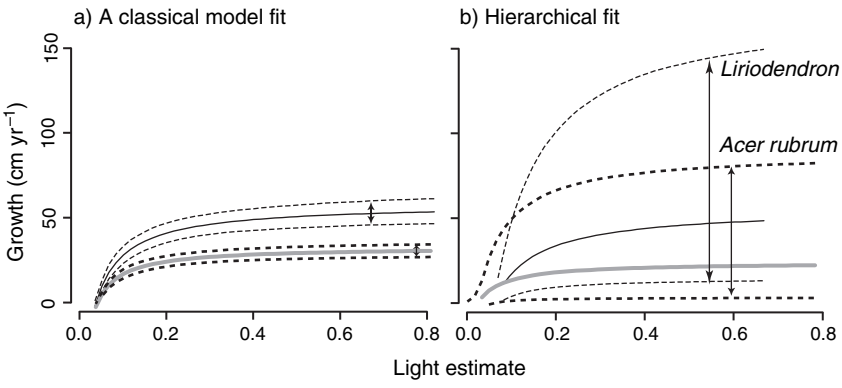

FIGURE 1.5. Comparison of the predictive interval on μ from a classical analysis, as has been standard practice in ecological analyses, with the predictive interval on y from the model in Figure 1.4c.

1.3 Two Elements of Models: Known and Unknown

I emphasize two elements of models, including variables and relationships that are known, either because they can be seen or we can assume them to be true, and those that are unknown, because they are obscure. The former constitutes much of traditional modeling in ecology. Unknowns are treated stochastically. Section 1.3.1 gives a brief overview of process models, followed in Section 1.3.2 by some aspects of unknown model elements, which are taken up in stochastic components.

1.3.1 A Point of Departure: Ecological Process Models

Environmental scientists do not have the luxury of continuing to limit attention to simple models. Low-dimensional models and designed experiments will continue to contribute insight, but the growing demand for relevant inference and prediction calls for the added capacity to address complex interactions. Modeling often begins with a process component to describe how things work. We may seek to determine whether or not relationships exist (e.g., hypothesis tests), to quantify relationships, and to make predictions. Where possible, the modeling strategy will focus on simple process models and, in many cases, allow for complexity on the stochastic side. Still, many processes have multiple interacting elements. Here I mention some of the types of environmental questions that entail integration of data and models, some of which are necessarily high-dimensional. Not all of these examples are included in this book, but they are all amenable to the approaches considered here.

Populations—At a time of intense interest in protecting rare species, the awareness that many nonendangered species undergo extreme fluctuations in density is of great interest. A broad range of questions and processes is addressed using population models and data. What places a population at risk? Will changes in the demography or age structure of a population have consequences in the future? What are the ecological implications of this structure? What kinds of constraints and trade-offs gave rise to observed life history schedules?

Ecologists use models to explore relationships involving individual behavior and dynamics of populations. For example, reproductive episodes can result in oscillations if there are strong feedback effects of density on population growth rate. Although evidence for chaos in population dynamics is weak (Hassell et al. 1976; Ellner and Turchin 1995), feedback involving resources or natural enemies can result in complex dynamics (Elton 1924; Saucy 1994; Ostfeld and Jones 1996; Earn et al. 2000). Oscillations can bring a population seemingly close to extinction, only to be followed by rebound to high abundances. Such fluctuations can occur repeatedly. They can be periodic, and, thus, are potentially predictable. Important strides in the 1970s (May 1974) initiated fertile research that has expanded to include feedbacks and environmental variation (e.g., Bjørnstad et al. 2001). Age represents

one of several kinds of structure in populations. It can have several classes of effects, depending on its contribution to time delays and response to perturbations (Lande et al. 2002; Coulson et al. 2001). The inherent spatial variability in nature (e.g., soils and topography) can explain much pattern in populations, but spatial pattern can also result from processes endogenous to populations themselves. Spatial coupling can result from environmental variation (the Moran effect) (Grenfell et al. 2000), dispersal (Kelly and Sork 2002), or other biotic interactions (Ostfeld and Jones 1996).

The challenges faced by scientists and managers come from the fact that populations are highly variable and subject to many influences, and demographic processes are hard to see. The basic process models used to infer population dynamics are reviewed in a number of recent theoretical ecology texts (Chapter 2). Spatio-temporal variability in demographic rates can complicate models (Wikle 2003a; Clark et al. 2003a) and require many years of data to estimate parameters (Clark et al. 1999b; Coulson et al. 2001). Many of the variables that ecologists would like to infer or predict, such as demographic rates for endangered species and extinction time, occur at low density. Because different processes may dominate at low density, including Allee effects, we cannot simply extrapolate parameter estimates obtained from populations studied at higher densities. Spatio-temporal processes are inherently high-dimensional, in the sense that variables not only interact with one another, but they do so in ways that depend on when and where they occur (Legendre 1993a; Ranta et al. 1995). Natural ecosystems are characterized by both continuous and discrete variation in space (Hanski 1998). Parameter uncertainty may have a large impact on predictions of extinction risk (Ludwig 1999; Ellner and Feiberg 2003). Ecologists have long struggled to characterize land cover in ways that are tractable, yet relevant for population dynamics (Lande 1988; Franklin et al. 2000).

Ecological Communities—Models are used to understand trophic interactions among species (Hutchinson 1959; MacArthur 1972; Tilman 1982, 1988). For example, why are there so many species? Why are some abundant and others rare? Trophic interactions and environmental controls contribute to these patterns. Interactions among species on the same and on different trophic levels must somehow explain much of the pattern (or lack thereof) observed in nature. Why don't the best competitors for the few resources that appear to limit growth drive all others to extinction (Tilman 1988)? How do species interact by way of shared natural enemies (Holt 1977) or other types of indirect effects (Wootton 1993)?

Early ecological models were deterministic, but incorporation of temporal stochasticity has become increasingly popular. Ecologists early suspected that variability in time could have some obvious as well as some mysterious implications for fundamental ecological processes (e.g., Elton 1924; Hutchinson 1959). Fluctuations in the environment coupled with temporal integration by organisms (e.g., long life, dormancy) can contribute to coexistence of species, and it may have evolutionary consequences (Chesson 2000). I consider different approaches for accommodating temporal heterogeneity in Chapter 9.

Spatial relationships complicate modeling, including movement and re-source patchiness. Recent theoretical models have increasingly focused on how predictions of spatial models differ from those of nonspatial ones (e.g., Lewis 1997; Neuhauser 2001; Bolker et al. 2003). A rich literature is developing that emphasizes problems related to spatial covariance and scale (Levin 1992). Theory has facilitated understanding of how aggregation of individuals in space affects interactions such as competition and predation. New statistical approaches allow us to consider processes where spatial relationships are not spatially coherent and change over time (Chapter 10).

Recent models have contributed to our understanding of how disease can affect population dynamics. For example, models have proved invaluable for exploring the spread of AIDS (Anderson and May 1991), foot-and-mouth (Keeling et al. 2001), and the temporal trends in measles (Bjørnstad et al. 2002) and whooping cough (Rohani et al. 1999). Examples of spatio-temporal analyses based on epidemiological data arise in Chapters 9 and 10.

Ecosystem Function—Biogeochemistry involves models and data at a range of spatial and temporal scales. While the 7 percent perturbation of carbon exchange between plants and the atmosphere has been enough to awaken global concerns of climate change, it is one of many interacting cycles that are absorbing the dramatic alterations by humans (Schlesinger 2004). Freshwater shortages, eutrophication from mobilized reactive N and P, acidified precipitation from mobilized S, and pollution of rivers and coastal oceans are just a few of the now-recognized transformations in the chemical environmental upon which life depends. Fertilizer applied in the Upper Midwest impacts fisheries in the Gulf of Mexico (Galloway 2004). Growing awareness that important human perturbations transcend the periodic table (e.g., recent concerns for Pb, Fe, Cl, and B, to name a few) underscores the need to understand linkages. Biology is central to such linkages. Stochiometric relationships demanded by organisms mean that change in the supply of one constituent can cascade through many others (Elser and Sterner 2002). Stochiometric relationships put biology in the driver's seat, as a regulator and a place where nonlinear feedbacks can reside, while simultaneously lending a degree of predictability to element interactions.

With a pending global energy crisis and its broad impact through climate, carbon will keep the attention of biogeochemists for some time (Jackson and Schlesinger 2004). Environmental scientists must anticipate not only the future impact of continuing human perturbations, but they also must weigh in on potential engineering fixes, such as the large-scale N fertilization of soils and Fe fertilization of oceans to stimulate uptake of atmospheric CO_2.

The modeling challenges are great, involving physical processes in soils, waters, and the atmosphere at a range of scales. I mention several continental-scale analyses involving air quality in Chapter 10.

Biodiversity Feedbacks on Ecosystems—Through integration of data and models, the melding of the classically distinct disciplines of community and eco-system ecology has shown that nutrient addition can lead to reduced diversity (Wedin and Tilman 1996), while diversity can play an important role

in nutrient supply (Vitousek et al. 1996) and retention (Tilman et al. 1996). Reorganization of food webs comes with changes in nutrient loading that propagate through primary producers to herbivores (Rosenzweig 1971; Carpenter et al. 2001). Nonindigenous species not only respond to land use changes, native biodiversity, and climate change, but they can also change native ecosystems, altering fire regimes (D'Antonio et al. 2000), promoting spread of infectious disease (some of which are themselves caused by nonindigenous pathogens) (Daszak et al. 2000), and affecting nutrient cycling (Vitousek et al. 1996). How does the loss or addition of species affect ecosystem functioning? Is there a minimal number of species needed to maintain ecosystem function in a given setting? Is redundancy needed as a buffer against change (e.g., Loreau et al. 2001)? Does variability at the level of individual populations propagate broadly? Or are species with similar or complementary function and niche requirements roughly equivalent (Doak et al. 1998; Reich et al. 2003)?

The many scale-dependent issues that arise in the context of ecosystem function can be modeled using the graphical framework outlined in Section 1.2. This framework is applied throughout this book.

Human Dimensions—Models and data are used to provide guidance on potential impacts of climate and land-cover change. The pace of contemporary climate change threatens reorganization of food webs, as species outrun their natural enemies and hosts, forming new networks of interaction, while others are left behind. Anthropogenic climate change has already affected many populations, but to what degree, and how has it propagated through communities? Lags, transient effects, and poorly understood changes in phenology are already combining to produce unanticipated effects on migratory birds, pollinators, and plant reproduction (Parmesan and Yohe 2003). Migration potential of plants became a research priority when general circulation models (GCMs) of the atmosphere suggested that contemporary migration rates would need to exceed apparent dispersal capacities of many species (Pitelka et al. 1997; Clark et al. 1998, 2003b). A growing number of examples from contemporary invasions demonstrate the extensive impacts that can follow from a small change in species richness (Mack et al. 2000; Callaway and Manon 2006). Ecologists are increasingly applying models to help understand which species, communities, ecosystems, and habitats are susceptible to climate change impacts.

Models and data are central to CO_2 fertilization study. Where can we expect CO_2-enhanced growth responses, for which species, and how will changing competitive relationships affect diversity (Hattenschwiller and Korner 1996; Mohan et al. 2005)? How do climate and CO_2-induced shifts in the length and timing of the growing season affect species distributions and interactions, and ultimately ecosystem-level processes (Farnsworth et al. 1996; DeLucia et al. 1999; Parmisan and Yohe 2003; Root et al. 2003; Ibáñez et al. 2006)? How do the combined effects of climate, CO_2, and habitat fragmentation affect migration potential?

These examples, drawn from a range of ecological concerns, have been traditionally modeled deterministically, in the sense that unknowns are ignored

or treated in unrealistic ways. There is a growing number of important exceptions that constitute many of the examples used in this book. Because process models do not include everything that goes on in nature, we need to consider what to do with the things that are left out. As suggested in Section 1.2, these leftovers are treated stochastically.

1.3.2 Stochasticity and a Structure for Complexity

The foregoing processes can be complex and obscure, and involve the types of unknowns described in Section 1.2.2. As discussed for that simple example, statistical inference is used to integrate data with the phenomena that are formulated as process models. That integration involves stochasticity. I began this chapter by saying that a statistical model may be as simple as a process model wrapped in stochasticity, a place to park the scatter. In fact, stochasticity is not confined to places where data enter the model. In Section 1.2.2 it stands in for unknowns, such as differences among individuals that cannot be related to specific variables or processes that are only partially understood. As background and motivation for the methods that follow, I consider the concept of stochasticity a bit further.

In this book, I use the term *stochasticity* to refer to things that are unknown. Although we could speak of a stochastic process, it can be most productive to think of processes themselves as deterministic. *Models* are stochastic, because we cannot measure or specify everything that is relevant. We choose what to treat deterministically and what must be treated stochastically. This pragmatic position is encountered in traditional mathematical treatments (e.g., Taylor and Karlin 1994) and in computer sciences applications such as machine learning (Mitchell 1997; Pearl 2002), and it is unavoidable in modern statistics (e.g., Dawid 2004). Nonetheless, this view of stochasticity may appear to conflict with terminology in the ecological literature, in that it does not suggest that processes are themselves stochastic. When pressed, those who argue that, say, population dynamics are inherently stochastic, may appeal to Heisenberg's uncertainty principle, which says that we cannot simultaneously know a particle's position and momentum. The uncertainty it describes has no demonstrated (or hypothesized) relevance at the level of observable phenomena. There is no obvious answer to the question of whether processes are really stochastic, but there is a practical approach—deterministic relationships are used for relationships that are known; stochasticity stands in for the unknowns. For example, Palmer et al. (2005) discuss the classic Lorenz (1963) attractor, a type of chaotic behavior, produced by a set of three coupled differential equations. Standard model selection arguments (Chapter 6) might tell us to truncate the model to two equations, because the third equation accounts for only 4 percent of the total variance. It is well known that this truncation completely changes the behaviour—it is not even chaotic. So what is the statistician to do with so little information on one-third of the state variables? One answer is a stochastic autoregressive term, which recovers the chaotic behavior, including the Lorenz attractor with the two-equation truncated version. In other words, judicious application of stochasticity stands in for

additional complexity, which may be unknown, unobserved, or both. For an ecological example, births and deaths are modeled as demographic stochasticity in models that treat each individual as a black box; we do not care to model the physiology of individual bodies. But a physiologist might do precisely that and would not view birth and death as stochastic, but rather as explicit processes about which much is known. We could have a philosophical debate about this; thus, I simply adopt this perspective on pragmatic grounds, several of which are outlined here:

1. It helps clarify the role of stochasticity in models—it stands in for things that are unknown.
2. It helps clarify the trade-off between determinism and stochasticity in models. If we could (or desired to) know everything, we would have a high-dimensional, deterministic model. If little is known, we may require a sophisticated treatment of stochasticity. Or we might choose a stochastic representation simply to reduce the complexity of a model.
3. It emphasizes that the terms *noise* and *error* are not always the best way to think about the role of stochasticity. In many cases, variability emanates from many unidentified sources, including from the process itself and from the context of the process. Sources of variability may be unknown, yet still demand careful treatment.

For pragmatic reasons I take the view that there are real processes behind the widely ranging growth rates in Figure 1.2. Stochastic representation is used when we do not know the underlying processes, they cannot be measured, or we do not know how to describe them in a deterministic model. The many components may operate on different scales. Simply collecting them all in one stochastic shell seems to agree with the concept of noise, but the processes they represent should impact the model in diverse ways. Figure 1.4c showed three different ways in which stochasticity stands in for uncertainty in x, in y, and in θ. This example illustrates the pragmatic treatment of stochasticity that is essential to modern data modeling. I began with a deterministic process model that might be motivated by previous understanding (saturating light-growth response). I then identified aspects of the problem that are unknown, including the resource, the differences among individuals, and the process model itself. Stochasticity is introduced for each of these unknowns.

Modern statistical methods provide a framework for complexity that allows for the comprehensive treatment of unknowns that contribute to Figure 1.2d. Assuming the wrong structure for stochastic elements might be just as naïve as an inappropriate function for the deterministic process. The basic approach involves a decomposition of complex problems into those that control the process of interest, those that generate data, and the underlying parameters. This structure admits influences that may impinge in diverse ways, without necessarily requiring that they can be observed or even identified.

The basic structure I follow throughout is outlined in Figure 1.6. This particular decomposition comes from Berliner (1996) and has been used in several recent ecological examples (e.g., Wikle 2003a; Clark et al. 2003b, 2004). Like

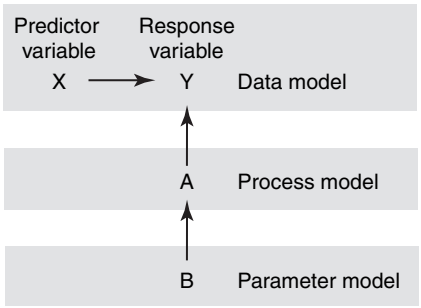

FIGURE 1.6. A graphical structure for an inferential model. Arrows indicate direct relationships that are specified explicitly as conditional distributions.

Figures 1.1 and 1.4, it is a graphical representation that provides a road map for the model. Attention may initially focus on the process level, where we allow for stochasticity in terms of process error or model misspecification. The connection to data accommodates sampling and the stochasticity that results from observation errors, sampling, missing data, and so forth. The parameter model can accommodate structure or relationships among sample units that might result from location, aggregation of individuals or traits, and so on.

By decomposing the problem, we can work directly with a process model that may be simple (and general), while allowing for the complexity that applies to a specific system (parameter models can absorb structure among variables) and data set or experimental design (data model) at other levels. This structure also admits data assimilation, meta-analysis, prediction, and decision analysis. To emphasize this connection between models and data it can be useful to think of inference and forward simulation as two sides of the same coin, one being the inverse of the other.

1.4 Learning with Models: Hypotheses and Quantification

Broadly speaking, the principal products of an analysis consist of (1) quantifying relationships and states (estimates and confidence intervals), (2) testing hypotheses (including model selection), (3) prediction, and (4) decision. Of these four activities, ecological analysis has tended to focus on hypothesis testing. I will not dwell on philosophy, but my inclination to soft-peddle hypothesis testing throughout this book needs some explanation. I discuss hypothesis testing many times, but it will often not be a prime motivation.

Before discussing why I place less emphasis on hypothesis testing than is customary, it is important to first say that hypotheses are critical. Any hope of success requires experiments and observations that are motivated by a clear set of questions with careful thought as to how the range of potential outcomes can lead to progress in understanding. A focus on multiple hypotheses, as opposed to a null hypothesis (e.g., Hilborn and Mangel 1997), can sharpen the approach. The lack of consensus on how to test hypotheses and to interpret them plays into a broader misunderstanding of what models are and how they can be effectively exploited. I summarize some of the issues that arise

with hypothesis testing followed by general considerations concerning the inevitable subjectivity of statistics.

1.4.1 Hypothesis Tests

Consider the (null) hypothesis H_0: $\theta = q$. We collect data and obtain an estimate of θ. Can we make a probability statement about the result? Something about our confidence that reality is different from or the same as q?

Three approaches to hypothesis testing are reviewed by Berger (2003). Fisher significance testing involves a test statistic S that increases in value the farther the estimate is from the hypothesized value q. The larger the value of S, the greater the evidence against the null hypothesis. A P value is associated with S, but it is not the error probability (i.e., the probability of being wrong). It is calculated as the area in the tail of the distribution of S, the tail being everything more extreme than S. If treated as an error probability (as is often done), the evidence against the null hypothesis is overstated. So P values are not error probabilities (they violate the frequentist principle—see below); they are based on values of S more extreme than actually observed.

Neyman-Pearson hypothesis testing has an explicit alternative hypothesis and appeals to the frequentist principle. This principle can be stated in different ways. Berger's (2003) practical version says that if we could repeat the experiment many times, then the calculated Type I and Type II error probabilities (falsely reject the null, falsely accept the null) would agree with the fraction of Type I and Type II errors from the repeated experiments. Whereas the Fisher alternative is vague (H_0: $\theta = q$, H_1: $\theta \neq q$), the Neyman-Pearson alternative is explicit (H_0: $\theta = q_0$, H_1: $\theta = q_1$). As with a Fisher hypothesis test, there is a test statistic S that increases in value with evidence against the null. The method requires a predefined critical value of S for accepting versus rejecting H_0. The outcome does not discriminate between outcomes (values of S) that may be barely significant versus very significant, because the critical value (e.g., $\alpha = 0.05$) only has meaning if it is designated in advance. Unlike the Fisher method, we do not conclude that S had a P value of, say, 0.03, unless this happened to be the preselected critical value. S is either less than the critical value or not, and the conclusion is the same regardless of whether S exceeds the critical value by a little or by a lot. In other words, data sets that lend very different weights to the null hypothesis in the Fisher sense might lead to the same inference in a Neyman-Pearson sense.

Jeffreys hypothesis testing also involves an explicit alternative hypothesis. The null hypothesis can be accepted if it results in a better fit (a greater likelihood) with a probability determined under the assumption that prior probabilities for null and alternative $= \frac{1}{2}$. One criticism of this approach has been the need to specify a prior probability (Fisher 1935).

The point of this summary is to emphasize that thoughtful experts do not agree on how to test a hypothesis. And the challenges mount as we move to high-dimensional hierarchical models, where the concepts of sample size and number of parameters do not have obvious interpretations (Gelfand and Ghosh 1998; Spiegelhalter et al. 2002). The desire to assign a probability statement to

a study is a laudable goal. It is often unrealistic in more than a highly qualified fashion. Ecologists can take sides in the debates and hope to maintain the "illusion of objectivity" (Berger and Berry 1988; Clark 2005). This may be the only option for those wishing to justify heavy reliance on hypothesis tests. At best, hypothesis tests are guidelines.

Even where we can agree on a hypothesis test, designs based on rejecting the null may not yield much. A study deemed unpublishable, because it fails to reject the null, may have been vacuous from the start. If showing an effect is viewed as the only informative outcome, there probably is no informative outcome. Studies founded on questions like "Does X affect Y?" (H_0: $\theta = q$) rarely provide much guidance, regardless of the P value. Such approaches often devolve to sample size (Spiegelhalter et al. 2003). Alternatively, if designed to parameterize a relationship, as opposed to rejecting a null hypothesis, that relationship may be the important product of the study.

1.4.2 The "Illusion of Objectivity"

The foregoing concerns bear on broader issues regarding what models represent, how they relate to theories and hypotheses, how they are used, and what can be inferred when they are combined with data. Models area not truth, just tools.

I will not advocate a distinction between scientific versus statistical models or between empirical versus mechanistic models. Such terms have had a place in the context of classical data modeling, because theory and data can be difficult to combine. For instance, it is common to fit data to a model like that shown in Figure 1.2, and then to use estimates that fit in a more complex model, for example, one that is dynamic with additional assumptions about other variables, perhaps even at different spatial and temporal scales. More desirable would be to fit the data directly to the model that will ultimately be used for prediction (Pacala et al. 1996). This direct connection is more desirable, because predictive intervals depend on the relationships that are lost when models are fitted piecemeal. The full model is harder to link directly with data, because it is complex.

By admitting complexity, modern methods allow for more direct connections between theory and data. The emerging techniques increasingly make distinctions between scientific and statistical models unnecessary. A motivation for setting such terms aside is the sense they can foster that there is a correct model, a best type of data, and an objective design; that is, there should be some formal and objective statistical test that will tell us what is best and correct. Models differ in complexity for many reasons, including not only the complexity of the process, but also how much we can know about it. Models are caricatures of reality, constructed to facilitate learning. In a world where data are accumulating at unprecedented rates, far more rapidly than we can assimilate them in models (e.g., remote sensing, climate variables, molecular data), we need ways to combine the information coming from multiple sources. Much of this book addresses formal structures for integrating multiple data sources within complementary models.

In most cases, there will be more than one way to model the same process. How well the model fits a data set is one of several considerations in deciding among candidate models. In several sections, I discuss model selection as a basis for helping to identify models that can be of most use. I follow the standard practice of parsimony: increasing complexity in a model requires justification. Still, I place less emphasis on model selection than is typically done within a classical approach focused on rejecting a null hypothesis. This diminished emphasis on model selection stems not only from the view that there is rarely a correct model, but also on several characteristics of ecological phenomena. Ecological processes are inherently spatio-temporal, and the best model can vary from place to place and change over time (e.g., West and Harrison 1997). The frequentist concepts based on the idealized notion of resampling an identical population often does not directly translate to environmental data sets. This is inevitable in high-dimensional systems that are subject to processes that operate at a range of scales. This view can shift the emphasis from that of identifying the correct model to one of identifying models that can be useful in different settings and the need to consider model uncertainty as integral to inference. In some cases it can motivate combining models or model averaging.

Finally, formal statistics should not trump all other considerations, especially when it comes to model selection. The fit to data is not the sole basis for model selection. General circulation models of the atmosphere (GCMs) were constructed to embrace certain physical relationships. Although weather and climate models can benefit from better integration with data (e.g., Bengtsson et al. 2003), a hypothesis-testing framework was not the basis for their formulation and does not play a role in continuing model development.

1.4.3 Quantifying Relationships

Unlike hypothesis tests, confidence envelopes are less controversial than many ecologists think. Confidence envelopes estimated for a given model by different methods and stemming from different perspectives often yield similar results. The processes under investigation are often known to exist. There may be no point in rejecting the existence of an effect. There may be value in quantifying it. For example, the effect of interest may contribute a relatively small fraction of the variance to a data set, because there are other large effects. If so, model selection (a hypothesis test) may reject the added complexity of a model containing the variable of interest. This should not necessarily deter us from examining those variables. Lange et al. (1992) found slight improvement in the fit of a change-point model of treatment effects on CD4 lymphocytes, a surrogate for AIDs infection, in patients with advanced HIV infection (models of this type are described in Chapter 9). The more complex model that included the variable of interest was still valuable, despite the fact that it did not explain nearly as much of the total variance as did other factors. Clark et al. (2004) found that temporal effects on tree fecundity overwhelmed individual effects, yet individual differences were

still large. The more complex model that contained individual effects allowed for quantification of individual effects, *in spite of* large stochasticity standing in for other effects. A posterior density centered near zero might suggest use of a simpler model, yet this parameter might be retained simply as insurance against effects on estimates of other parameters.

1.4.4 Learning from Models

Learning is progressive. It entails continual updating or assimilating information as it accumulates (Section 1.2.1). Models have a central role in this process. This process might be represented as

$$\text{updated knowledge} = f(\text{data, previous knowledge})$$

Bayesian methods do this formally as

$$posterior \propto likelihood \times prior$$

Learning can occur sequentially, with the posterior or predictive distribution that derives from information available before serving as a prior to be combined with data that comes available in the future (West and Harrison 1997). Whether or not we do this formally, the process of successively updating understanding as information accumulates is standard practice. The Bayesian framework for learning plays a large role in this book. The likelihood in this expression admits data. For processes that evolve over time, we might term this the *update/forecast cycle*, with the goal being to ingest new information continually.

1.5 Estimation versus Forward Simulation

Model analysis and prediction are ideally based on the same process model as that used for estimation (Figure 1.7). This connection is critical if we are to construct prediction envelopes that accurately reflect uncertainty in data and models. For a number of reasons, few ecological analyses proceed this way. Predictions are typically made from models that are not the same as those used to obtain parameter estimates. And ecological predictions typically come from process models only after discarding the stochasticity (Clark 2003). It is not necessarily wrong to analyze models that are not fitted to data. But arbitrary treatment of uncertainty can make predictions misleading or even meaningless. This book is partly motivated by the fact that there are fewer obstacles in the way of integrating theory and data than there were just a decade ago. New approaches provide a natural framework for data assimilation and for decision (Figure 1.7).

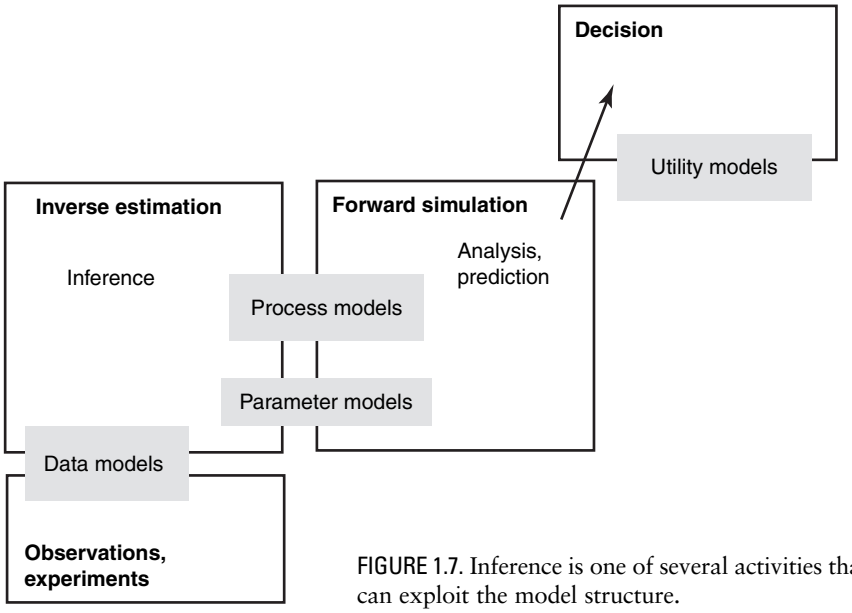

FIGURE 1.7. Inference is one of several activities that can exploit the model structure.

1.6 Statistical Pragmatism

No practicing ecologist is ignorant of the debates over frequentist and Bayesian approaches (e.g., Dixon and Ellison 1996). Already, philosophy has come up in reference to stochasticity (Section 1.3.2) and hypothesis testing (Section 1.4), but it does not play a large role in this book. In light of the philosophical emphasis of ecological writings about Bayes, I should provide some justification for my largely nonphilosophical treatment in this text. Here I summarize several issues that I have written about elsewhere (Clark 2005).

First, the points of controversy have traditionally revolved around *frequency* versus *subjective* concepts of probability (Chapters 4 and 5). The different concepts are especially important when one tries to make sense of classical hypothesis testing, but they also arise within the Bayesian community, in the context of how strictly prior specification must represent current understanding. Ecological writings on Bayes have focused on philosophical differences, advising that an analysis begins with one's philosophical stance, which then prescribes a classical versus a Bayesian approach. This prescription follows past debates in the statistical literature that have become less common and less polarized in recent years. I have expressed my view that the dissipation of such debates and the emergence of modern Bayes in applied fields have less to do with philosophy than pragmatism (Clark 2005). In fact, the philosophy remains important, being central to the challenge of doing objective science using the subjective tools of statistics (e.g., Berger and Berry 1988; Dawid 2004). The expansion of modern Bayes owes much to developing machinery.

Bayesians can now take on high-dimensional problems that were not accessible in the past.

In my view, the persistent focus on philosophy in ecological writings has become counterproductive (Clark 2005). It should not be news to any scientist that strong priors can affect estimates. Weak priors generally do not. The nonpractitioner must wonder at the standard examples provided in many ecological demonstrations of Bayes. Typically a simple Bayesian analysis requires far more work to arrive at essentially the same confidence envelope that could have been obtained with standard software in a classical framework. Given that most ecologists have limited philosophical baggage in this respect, why complicate the analysis? Moreover, the focus on philosophical differences can confuse the issues. For example, although Bayesians refer to parameters as "random," and frequentists do not, in both cases, parameters are fixed. Going further, both approaches view parameters as uncertain. Ecologists have long confused references to random parameters with the idea that the underlying value of a parameter fluctuates. Classical confidence intervals and Bayesian credible intervals both express uncertainty about the underlying true value, which is unknown.

I say more about the underlying assumptions of classical and Bayes models as we begin to implement them. For now, I simply forewarn that both are important in this book, with applications being pragmatic rather than philosophical. By pragmatic, I mean that I avoid unnecessary complexity. The initial models I discuss in Chapter 3 are classical. Bayesian methods require evaluation of integrals, analytically or through simulation. For simple problems this can mean that Bayes requires more effort. I also discuss Bayes for relatively simple problems, but it becomes the dominant approach in later chapters to allow for external sources of information and uncertainty. As models become more complex, the level of difficulty associated with classical and Bayes can reverse. Not only will the Bayesian framework bring flexibility to address complex problems (even where classical methods are still an option), Bayesian approaches can be *easier*, facilitated by new computational techniques.

A pragmatic approach need not be controversial. As suggested above, primary focus on quantification, rather than hypothesis testing, means that the products of a classical and Bayes analysis (e.g., confidence and credible intervals) may often be similar (e.g., Cousins 1995; Clark and Lavine 2001). Controversy can still arise over the role of priors in a Bayesian framework. I will say that informative priors are a good idea, when external information is available. Noninformative priors are useful when information is not available, or when we want to isolate the contribution of a particular data set. Those just testing the Bayesian waters may be motivated by machinery. Those who initially balk at informed priors may come to recognize each new data set as a way of updating what is already known. Priors can be an efficient way to introduce partially known relationships, in contrast to the traditional practice of assuming such relationships are completely known. I return to some of these issues in the final chapter.

Even as models become complex and (in this case) Bayesian, many of the tools will be familiar from classical statistics. Bayesian models will have embedded within them many traditional structures. The power of modern Bayes comes

from the capacity to integrate traditional techniques within high-dimensional models.

Although I embrace both frequentist and Bayesian approaches, there is much that I avoid. In both cases, there is an established framework that underpins the analysis. If I say that I have constructed a confidence envelope based on Fisher Information or a Bayesian posterior, you will have a good idea of what it means (Chapter 5). Although challenging problems can foster excursions into ad hockery (a much bigger temptation for the inflexible classical methods), the formalism of these two approaches is well established.

In this regard, the expanding ecological literature on fitting models with ad hoc approaches that are neither frequentist nor Bayesian is avoided in this book. Many such approaches are highly creative. But there can be numerous pitfalls in terms of inference, prediction, and communication to others. The goal of inference is a probability statement about unobservables—quantities we wish to estimate. A probability statement relies on a framework that is consistent throughout. Ad hockery tends to break down the connection between confidence envelopes and the data used to construct them. Moreover, ad hockery is often unnecessary. When things get complicated, we will go Bayesian. It admits complexity within a consistent, reproducible framework. Although there remains much to understand in the Bayesian world (e.g., selection among hierarchical models in Chapter 8 and complex spatio-temporal models in Chapter 10), there is rapid progress and it builds on a firm foundation. Because we have the Bayesian alternative, there are many recent papers on how to run lines through data points that will not be covered here.

Finally, although I cover basic concepts in classical statistics, I make no effort at comprehensive coverage of classical designs and hypothesis tests. Classical statistics may seem to burden the scientist with a different design and test for each new application. There is substantial jargon that continues to find its way into ecological statistics books, much of which is largely historical and rarely encountered in the modern statistics literature. I focus on general approaches and attempt to minimize jargon.