


# In Which We Provide a Context

A model is any representation or abstraction of a system or process. We build models because they help us to (1) define our problems, (2) organize our thoughts, (3) understand our data, (4) communicate and test that understanding, and (5) make predictions. A model is therefore an intellectual tool.

People's perceptions about models and modeling can be very different depending on the types of problems they usually face and hence the types of tools they commonly use. There is a vast difference, for example, between trying to predict next week's weather and plotting a rocket's trajectory to the moon. These differences are reflected in the models chosen and how they are used.

Holling (1978) has a diagram (Fig. 1.1) that provides a simple and useful classification. The horizontal axis represents how well we understand the problem we are trying to solve; the vertical axis represents the quality and/or quantity of relevant data. Holling divides the quadrant between the two axes into four areas, corresponding to four classes of problems.



**Figure 1.1** Holling's (1978) classification of modeling problems. This book is concerned mainly with regions 2 and 4.

Area 1 is the region of good data but little understanding. This is where statistical techniques are useful; they enable one to analyze the data, search for patterns or relations, construct and test hypotheses, and so on.

Area 3 is the region of good data and good understanding. Many problems in engineering and the physical sciences (for example, the problem of plotting the rocket's trajectory to the moon) belong to this class of problems. This is the area where models are used routinely and with confidence because their effectiveness has been proved repeatedly.

In area 2 there is little in the way of supporting data but there is some understanding of the structure of the problem; in area 4 even the understanding of the problem is tenuous. Many of the problems in the nonphysical sciences belong to either region 2 or 4, and these are the types of problems we will address in this book. They present us with two rather daunting challenges:

1. From the management point of view, decisions may have to be made despite the lack of data and understanding. How do we make good, scientific decisions under these circumstances?
2. How do we go about improving our understanding and collecting the data we need? (In other words, how do we progress from area 4 to area 3 of Holling's diagram?)

We will attempt to show how models help us meet both these challenges.

Some would argue that this attempt is bound to fail. They believe that the first priority is to collect as much data as possible and that model building should be postponed until the data have accumulated and been analyzed statistically. Others, noticing the routine way in which models

are used in area 3 of Holling's diagram, are convinced that their problems are too ill-defined to model in that way.

The latter are correct; one does not use exactly the same tools in the same way in areas 2 and 4 as in area 3. However, they are incorrect in their assumption that the modeling toolkit is designed only for the well-defined problems in area 3. The purpose of this book is to show:

- How models can be built, very tentatively at first
- How the properties of the models can be explored
- How one can speculate, using hypothetical data
- How one can then cautiously reach some conclusions and search for evidence that supports them

Models built this way are bound to be speculative. They will never have the respectability of models built for solving problems in area 3 because it is unlikely that they will be sufficiently accurate or that they can ever be tested conclusively. They should therefore never be used unquestioningly or automatically. The whole process of building and using these models has to be that much more *thoughtful* because we do not really understand the structure of the problem and do not have (and cannot easily get) supporting data.

We therefore build models to *explore the consequences* of what we believe to be true. Those who have a lot of data and little understanding of their problem (area 1) gain understanding by "living with" their data, looking at it in different ways, and searching for patterns and relationships (Tukey, 1977). Because we have so little data in areas 2 and 4, we learn by living with our models, by exercising them, manipulating them, questioning their relevance, and comparing their behavior with what we know (or think we know) about the real world. This process often forces us to reevaluate our beliefs, and that reevaluation in turn leads to new versions of the models. The mere act of assembling the pieces and building a model (however speculative the model might be) usually improves our understanding and enables us to find or use data we had not realized were relevant. That in turn leads us to a better model.

The process is one of *boot-strapping*: If we begin with little data and understanding in the bottom left-hand corner of Holling's diagram, models help us to zigzag upwards and to the right. This is a far healthier approach than one of just collecting data because we improve our understanding as we go along. (Those who collect data without building models run the very real risk of discovering, when they eventually analyze their data, that they have collected the wrong data!)

The approach we are advocating can never be routine. It is a subtle process, which is why we need a whole book and not just an introduction to describe it.

## A TRIVIAL MODEL

To introduce some of the concepts and terminology of modeling, consider a very simple process—converting U.S. dollars into German marks. We can write an equation that models this process:

$$M = kD \quad (1.1)$$

where  $D$  and  $M$  represent the number of dollars and marks respectively and  $k$  is the conversion rate. A number of points are worth noting about even as trivial an example as this.

1. The model has an *input* and an *output*. We feed in  $D$ , the number of dollars, and it returns  $M$ , the number of marks. We call  $D$  and  $M$  the *variables* of the model. They are what the model is all about—we want to know how  $M$  changes as we vary  $D$ .

2. We call the exchange rate  $k$  a *parameter* of the model. It is not a variable because it is a quantity we have to estimate before we can use the model; it mediates the relationship between the variables. If exchange rates were fixed officially and did not change for years at a time, the parameter  $k$  would be a *constant* of the model.

3. We can distinguish between results that are a consequence of the *structure* of the model and results that depend on specific *data*. The structure of the model is determined by the governing equations. In this case we have only one equation and its significant property is that it is *linear*: If we draw a graph of  $M$  versus  $D$  we get a straight line. Moreover, the line passes through the origin, so if we double the number of dollars we wish to convert, we will get exactly twice the number of marks—i.e., as we increase or decrease the input, so the output increases or decreases in the same proportion.

Results such as this, which follow from the structure of a model, often have important practical implications. If our example were not so trivial, we might well be excited to discover that it makes no difference whether we exchange \$200 in a single transaction or in two separate transactions of, say \$100 each. This result is independent of the exchange rate  $k$ . On the other hand, if we want to know how many marks we will actually receive for our \$200, we have to know the value of  $k$ —i.e., we need data. Results that are a consequence of the structure of the model are thus independent of the data and are general in their scope. Specific instances, or numerical examples, however, require data.

4. If we probe into the details of currency exchange, we will discover that Eq. (1.1) is an oversimplification and that the general result of the previous paragraph is not always valid. For example, the commission charged on very small transactions is likely to be higher (as a percentage of the transaction) than that charged on large transactions. Four hundred transactions of 50 cents each would in practice yield fewer marks than a single exchange of \$200. It is fortunate for our purpose that there is this

dis  
mo  
prc  
rea

ord  
bef  
our  
cor

tati  
not  
und  
only  
trac  
not  
the  
qua

excl  
no p  
con  
excl  
amp  
of, s  
valu  
the f  
tribu

tren  
the s  
cann  
it oft  
shou  
stocl  
a seq  
it. It  
taint  
feren

comp  
be m  
Whe  
with  
sequ

ng, consider  
marks. We

(1.1)

respectively  
noting about

, the number  
D and M the  
—we want to

el. It is not a  
e can use the  
If exchange  
t a time, the

quence of the  
ta. The struc-  
. In this case  
it is linear: If  
over, the line  
llars we wish  
s—i.e., as we  
creases in the

e of a model,  
e were not so  
no difference  
two separate  
f the exchange  
rks we will ac-  
—i.e., we need  
model are thus  
cific instances,

we will discov-  
al result of the  
ne commission  
is a percentage  
. Four hundred  
r marks than a  
hat there is this

discrepancy between our model and the real world because our trivial model is a little too slick. Models are an abstraction, a simplification of a process rather than a replication of the process. They never describe the real world *exactly* and often do not even attempt to do so.

People working in the physical sciences distinguish between first-order and second-order effects; the former explain much of the observed behavior of a system, while the latter can be considered as refinements. In our trivial model, Eq. (1.1) contains the first-order effect; the question of commissions is a second-order effect.

It follows from this discussion that it is always possible to find limitations to a model. As a result, we have to be wary on two counts: first, not to rely on a model when it is stretched to the limit and, second, not to undervalue a model because it has limitations. The fact that a model is only valid within certain limits or under certain conditions does not detract from its usefulness within those limits. *The quality of a model does not depend on how realistic it is, but on how well it performs in relation to the purpose for which it was built.* For most tourists Eq. (1.1) is an adequate and useful model.

5. Equation (1.1) is a *deterministic* model. Given a value for the exchange rate  $k$ , once we have converted, say, \$100 into marks, there is no point in repeating the calculation; the model is entirely predictable and contains no element of uncertainty. Suppose, however, we knew that the exchange rate fluctuated, within certain limits, from day to day. For example, our information might be that  $k$  can vary from 2.70 to 3.20 (in steps of, say, 0.02) and that on a particular day it is as likely to have any one value as any other value within that range; in other words, our data is in the form of a *statistical distribution* for  $k$ . In this case the shape of the distribution is flat or uniform (equal probability) over the range 2.70 to 3.20.

One property of a long sequence of random numbers is that the trend of the numbers is predictable even though the individual numbers in the sequence are unpredictable. Think, for example, of throwing a die. We cannot tell what number will come up next, but we know that if we throw it often enough (and if it is unbiased), all the integers from one through six should come up with equal frequency (another uniform distribution). A *stochastic* version of our dollar-to-mark conversion model is one that uses a sequence of random numbers to provide a value for  $k$  each time we need it. It follows that the stochastic model contains an element of uncertainty—if we repeat a calculation we are likely to get a somewhat different answer every time.

How do we obtain an appropriate sequence of random numbers on a computer? We use subroutines called *random number generators* that can be modified to produce a sequence with the required statistical properties. When we use stochastic models we will often want to run the model once with one sequence of random numbers, then run it again with another sequence. The “seed” of a random number generator is a number (any

number) that we feed in at the beginning of a simulation; the subroutine uses it to choose a starting point in the sequence of random numbers. We can produce as many sequences as we need by changing the seed.

The choice of whether to build a deterministic or a stochastic model depends on the purpose of the model. For the average tourist there would be no point in building a stochastic model. Even if he knew that the exchange rate fluctuated he would probably be satisfied with a deterministic model that used the average value of  $k$ .

Suppose, however, that our model is going to be used by a vice-president of a large corporation who knows that he will have to convert \$10 million into marks some time within the next week. He is quoted an exchange rate of 2.98 today; should he convert or should he wait in expectation of getting a better exchange rate later in the week? The deterministic model cannot answer this question, but the vice-president could use a stochastic model (in the absence of any understanding of the relevant market forces) to estimate the probabilities of losing or gaining by postponing the decision. Generally speaking we use stochastic models whenever the *variance* in the behavior of the system is important. Two examples of this will be found in Chaps. 3 and 4.

6. Complete treatises could be written on currency exchange; we have written one simple equation. Should we build a more detailed model? Should we, for instance, try to predict the value of  $k$  on the basis of other information (such as the balance of payments and prime interest rates in the two countries)? These questions relate to the *resolution* of the model, and the discussion of resolution merits a section of its own.

Before moving on to that section, note that the answers to these questions will depend, just as in points 4 and 5, on the *purpose* of the model.

Note also that points 4, 5, and 6 will crop up in various guises throughout this book. We introduce them here only to alert the reader to some of the issues involved in model building.

## THE RESOLUTION OF A MODEL

We will use a number of different analogies to illustrate what we mean by resolution. We use the word in the same sense as the resolution of a microscope or telescope, where we are concerned with the extent to which the optical instrument enables us to distinguish the components of the object we are viewing. Similarly, the resolution of a model tells us which aspects of the subject being modeled are distinguishable or clear and which are hidden, ignored, submerged, or blurred.

Whenever we observe the world around us, we do so selectively; we pay attention to some features and ignore others. Think of driving along a highway. We pay particular attention to the cars immediately in front of and behind us, some attention to road signs, less attention to the rest of

abroutine  
ibers. We  
ed.  
tic model  
ere would  
/ that the  
eterminis-  
  
oy a vice-  
to convert  
quoted an  
wait in ex-  
The deter-  
dent could  
of the rele-  
gaining by  
tic models  
rtant. Two  
  
change; we  
re detailed  
n the basis  
me interest  
ution of the  
s own.  
ers to these  
pose of the  
  
rious guises  
he reader to  
  
we mean by  
tion of a mi-  
ent to which  
nts of the ob-  
us which as-  
ar and which  
  
electively; we  
riving along a  
ly in front of  
to the rest of

the traffic, and only peripheral attention to the scenery. The way in which our attention is divided depends on circumstances. If the route is strange, we will pay more attention to the road signs; if it is scenic, we will take more notice of the scenery. Moreover, the speed of the neighboring cars is their important attribute, not their color, style, or attractiveness of the occupants. If we are unobservant, or look at the wrong things, we will be bad drivers. If we notice too much as we drive, we are also likely to impair rather than improve the quality of our driving.

The same is true of modeling. If our resolution is too coarse, the model will be inadequate. If it is too fine, we will be distracted by irrelevant detail. The balance is important too. The driver who concentrates on the scenery at the expense of watching the cars closest to him is likely to have an accident. A model that is not properly balanced can be equally disastrous. There are thus two aspects to the concept of resolution: first, which components we include in our model and which we leave out, and second, how much detail or emphasis we ascribe to the components we include. By *resolution* we therefore mean a combination of *scope* and *detail*.

In Fig. 1.2 we change our analogy to illustrate how the resolution of a model determines what the model can and cannot do. In each diagram we are concerned with two antelope. In Fig. 1.2(a) we are interested in the two antelope within the context of the herd; our model enables us to make statements such as "those are the two, the ones in the middle of the herd" or "they are in a large herd" (or "a small group" or "a dense herd"). We emphasize the context at the expense of being able to say something specific about the two animals themselves; we do not even know whether they are male or female. In Fig. 1.2(b) we ignore the context—i.e., we reduce the scope of the model. The only statements we can make with this model are ones such as "the one is bending down while the other is looking up" or "they are standing close together." Figure 1.2(c) has the same scope as Fig. 1.2(b) but considerably more detail. Here we can identify the one antelope as a female kudu, the other as an oryx. We can comment on the stripe patterns on the kudu or the shape of the oryx's horns. Figure 1.2(d) illustrates a model that is unbalanced. Unless there are good reasons for concentrating on the two rear stripes on the kudu and the upper back legs of the oryx, there is something wrong with the resolution in this diagram.

Comparing Fig. 1.2(a), (b), and (c) highlights the extent to which the resolution of a model really determines what the model is all about. In each case the subject is two antelope, but the three representations of those antelope lead to three models that have very different structures, data requirements, and types of output. It is often useful, when trying to choose a suitable level of resolution for a model, to sketch something similar to these three diagrams and list for each what the model will be able to achieve and what data it will require.

There is good reason for belaboring the concept of resolution. Those




Figure 1.2 Illustrating resolution and balance. Note that (d) is unbalanced.



who build models in the physical sciences scarcely give a thought to the question of resolution. Years of experience have been transmitted from one generation to the next in such a way that a physicist or engineer automatically and almost instinctively chooses the resolution appropriate to a particular problem. This experience has *not* been developed to the same extent in the nonphysical sciences, and perhaps the most important skill to master in disciplines such as ecological modeling is choosing the appropriate level of resolution. The choice depends as much on the purpose of the model as on the structure of the system. It also depends on the *time scale* of the model. (We can ignore changes in the polar icecaps in a model with a time scale of decades, but not if we are trying to predict climatic changes over thousands of years.)

In the next section we introduce time-dependent models and show how questions of time and resolution are interlinked.


TIME-DEPENDENT MODELS

Our trivial currency exchange model is static; time does not feature explicitly in it. However, if we modified it to predict how the exchange rate might change as a function of time, it would be a *time-dependent* model. Most models are time-dependent for the simple reason that we tend to collect data, think, and plan in terms of time. Often our objective is to project into the future. We want to know what might happen if we do this instead of that.

There are different ways we can model the passage of time. Suppose, for example, that we are responsible for the control of mosquitoes in a city. We might want to monitor and model what happens to the mosquito population from one year to the next. We then have the choice of representing the population either as a *continuous* or a *discrete* function of time. If we choose the continuous representation, mathematically we can think of the population as some function  $P$  of the time  $t$ , and in practice we could choose *any* value for  $t$  and our model would give us an estimate of  $P(t)$ . If we plot  $P$  versus  $t$  we will get a continuous graph, although the value of  $P$  could be zero throughout the winter months and suddenly increase at the beginning of summer.

Alternatively, we might choose to look at the mosquito population only once a month or once a year. The appropriate mathematical representation would then be one in which we use subscripts—for example,  $P_t$  might represent the population at time step  $t$ , in which case the population at the next time step would be  $P_{t+1}$ . This is the discrete representation. If we use it, it is not meaningful to ask what the population is at any time other than the specific times  $t, t + 1$ , etc.

Thus, a continuous model is one in which time *flows*, while a discrete model is one in which time *jumps*. Which representation should we choose? The answer depends partly on the format of our data, partly



s unbalanced.

on mathematical convenience, but mainly on the objectives and resolution of our model.

Suppose our objective is to predict how the mosquito population will respond to various control measures. At one level of resolution, those control measures might be described in fairly gross terms, such as "tons of insecticide used per year." In this case our model would be unbalanced if we thought of the number of mosquitoes as a continuous variable. The appropriate model is discrete with a time step of one year.  $P_t$  might then represent the total number of mosquitoes hatched during year  $t$  or the peak value of the population during the year.

However, at a different level of resolution we might want to look at how and when we actually apply the insecticide. We will then need to model changes in the mosquito population during the year. If we believe that mosquitoes hatch in an asynchronous manner, a continuous model is appropriate (although we may prefer to use a discrete model with a time step related to some characteristic time in the life cycle of a mosquito). On the other hand, if eruptions in the mosquito population are triggered by events such as rain storms, a discrete model would be mandatory. In the latter case the time step of the discrete model would not be constant; the interval between  $P_t$  and  $P_{t+1}$  could be quite different from the interval between  $P_{t+1}$  and  $P_{t+2}$ . We distinguish between *time-driven* discrete models (where time jumps regularly) and *event-driven* models (where time jumps forward only when something important occurs).

Most models are mathematical, if only in the sense that they use the language and notation of mathematics. However, in some cases we may design the model with the intention of using mathematical theorems and operations. We will call these *analytical models*, and several examples are given in Chap. 6. In other cases, the model may be designed as an *algorithm* or set of computational rules (as in a flow chart) without any attempt at formal analysis. We call these *simulation models*.

Since the mathematics of continuous functions has been studied more assiduously than that of discrete functions, it follows that if we are trying to build an analytical model, and if there is no compelling reason for choosing a discrete representation, we may prefer to make our model continuous. On the other hand, discrete models are easy to implement on a computer. So if we are building a simulation model, and if there is no compelling reason for a continuous representation, we may prefer to make our model discrete. We do, however, have to be cautious, because there are differences between a discrete and continuous representation of the same system; some of these are highlighted in Chap. 6.

## PROVIDING A CONTEXT

The important issues we have raised so far relate to such questions as:

1. The purpose of the model, its expectations, and its limitations

so w  
We  
mod  
to b  
we  
pres

Afric

"In r  
can p  
ing a  
He g  
tongu

them  
contai  
tah, ar  
with fl  
dense  
Rainf  
about  
where  
vary to  
7

resolution

ulation will  
ion, those  
h as "tons  
unbalanced  
riable. The  
might then  
ar t or the

it to look at  
en need to  
we believe  
ous model is  
with a time  
osquito). On  
triggered by  
atory. In the  
constant; the  
e interval be-  
crete models  
e time jumps

t they use the  
ases we may  
theorems and  
examples are  
signed as an  
without any at-

been studied  
that if we are  
ling reason for  
our model com-  
plement on a  
here is no com-  
fer to make our  
cause there are  
entation of the

h questions as:  
its limitations

2. How to choose the appropriate level of resolution.
3. Whether the problem lends itself to simulation or an analytical approach
4. Whether to use a discrete or continuous representation of time
5. Whether to build a stochastic or deterministic model

The answers to these questions depend on the purpose of the model, so we begin to see that it is futile to talk about modeling outside a context. We cannot criticize or evaluate a model unless we know the problem the model was designed to address. It follows that a book that describes how to build and use models must rely heavily on case histories. The examples we will use all relate to the management of game parks (or game preserves; we will use the words synonymously) in African savannas.

There are several reasons for choosing this theme:

1. The most important is that the authors have the experience of building models in this context.
2. Next is that the context highlights most of the difficulties one encounters trying to build models in the nonphysical sciences.
3. Finally, it can be advantageous, when writing about models, to use examples that are somewhat removed from the experience of the reader. This enables the reader to approach the model with an open mind and to concentrate on the arguments rather than worry about some of the finer details. Most of the readers of this book will not have experienced the difficulties of making management decisions in an African game park. All will, we hope, find the context fascinating and will recognize in it themes that parallel their own experiences.

### African Savanna and Game Parks

"In recent years the term 'savanna' has become synonymous with African plainslands—grasslands studded with flat-crowned acacias and carrying a profusion of wild ungulates." This is a quote from Huntley (1982). He goes on to give a wider definition, but this one, though somewhat tongue-in-cheek, will suffice for our purposes.

Our experience is based on various African game parks, some of them extensive ( $2 \times 10^6$  ha) and some relatively small ( $2 \times 10^4$  ha), some containing large predators (such as lion, Cape hunting dog, leopard, cheetah, and spotted hyena), and others not. In addition to "grassland studded with flat-crowned acacias," there is a mix of open grassland and areas of denser shrubs or woodlands, and the proportions vary from park to park. Rainfall and climate vary too. There are parks of arid savanna receiving about 200 mm of rainfall per year, and others in more moist savanna where the annual rainfall exceeds 600 mm. Topography and soil types vary too, not only from one park to another, but also within a park.

The various parks also have certain features in common. All are

partly if not totally fenced in. Migrations of large herds of ungulates were once a feature of African savanna; the fences either restrict migrations or prevent them altogether. The rainfall, whether it is high or low on average, tends to be variable. There are years of relatively abundant rainfall and years of drought. In some areas there is tentative evidence of a periodicity in the rainfall, cycles with a period of about 20 years during which half the cycle is relatively dry and the other half relatively wet. Particularly in the more arid parks, the bulk of the rainfall occurs during the summer months and the ungulates can be stressed during the long dry periods, especially if summer rains are late. Fire, either natural or man-induced, is also a common feature during the dry months and can influence the regeneration of trees, the balance between bush and grass, and the grazing patterns of the ungulates. These features are all important in relation to the management of a game park.

### Management Problems and Options

There is a lively and ongoing debate about the role of management in a game preserve. One can argue convincingly for as little management as possible in a game park that is very large and very far from the pressures of civilization. Unfortunately, such a park is hypothetical. African game parks have lost their innocence—either they are too small and surrounded by fences, or they are subject to increasing pressures from human population on their borders. Under these circumstances, a game park is not a pristine, natural ecosystem but a human artifice, and the decision not to manage is just one of many other management options; it also needs to be justified.

How much management, then, is necessary? From the ecological viewpoint, much depends on the size of the park (in relation to the size of the home ranges of the larger mammals) and the diversity of the vegetation. Over and above that, the perceived purpose and objectives of the game park are of paramount importance, and these tend to change with time and the pressures of human populations.

It is not our purpose to get involved in this debate. The interested reader is referred to Jewell and Holt (1981) and Owen-Smith (1983) for an introduction to some of the issues and positions. For our purposes we will accept stated management goals and use them as a foil for showing how various modeling techniques can help to refine or meet those goals. Similar techniques would still be useful if the goals were different.

The following are some of the actions that management may take or contemplate in a game park:

1. Not to intervene (as pointed out, this is as much an action as any other action).
2. To build dams or provide water artificially (to compensate for water

ates were  
nigrations  
or low on  
abundant  
evidence  
t 20 years  
f relatively  
fall occurs  
during the  
her natural  
ths and can  
and grass,  
e all impor-

- losses to irrigation upstream or for the fact that the animals can no longer move out of the park in search of water during a drought).
3. To crop (cull or remove) some of the animals. This may be done for a number of reasons:
  - a. To prevent predators from moving out of a park.
  - b. To protect a rare animal or plant species.
  - c. To prevent an animal species that has been too successfully protected from damaging its environment.
  - d. To limit the number of herbivores in a park where natural predators have been eliminated.
  - e. To control disease.
  - f. To finance the maintenance of the park.
4. To reintroduce and foster species that are either extinct or rare within the park.
5. To introduce fire-breaks and other fire control measures.
6. To deliberately burn sections of the park. Possible reasons for burning are:
  - a. To compensate for natural fires that would have occurred if there had been no fire control in the park.
  - b. To alter the pattern of the vegetation.
  - c. To prevent bush encroachment.
  - d. To produce a flush of fresh, nutritious grass for the herbivores.

The major problem management faces is that it cannot be sure these actions will actually achieve the objectives that prompted them; nor can it be sure that an action initiated to solve one problem will not generate other problems or undesirable side-effects. No two situations are ever quite the same, so it is not easy to draw on past experience or the experience of others. Assumptions have to be made; for example, are the next few years likely to be unusually dry or unusually wet? Frequently, decisions have to be made during a crisis, such as a severe drought, when there is no time to collect data or perform an experiment. Even when data are available, they may turn out to be the wrong kind or at the wrong level of resolution. Experiments can be suspect or open to more than one interpretation because it is often difficult to design a control experiment or to evaluate results in the light of environmental factors that changed during the experiment.

It is against this background that one has to evaluate whether models can be useful and, if so, what types of models to build and how to use them. It is a background that forces one to be pragmatic; if the model is not ready in time, decisions will have to be made without it, and if data are not available for the model, the model may have to be built without them.

If we look back at Holling's diagram (Fig. 1.1) we see that these considerations place the problems to be solved squarely in area 2 or 4. Our context thus illustrates the class of modeling problems we have chosen to address in this book. If these considerations add to the difficulties of building models, they also add to the interest and importance of doing so.

gement in a  
nagement as  
he pressures  
African game  
d surrounded  
man popula-  
park is not a  
ecision not to  
so needs to be

the ecological  
n to the size of  
of the vegeta-  
jectives of the  
to change with

The interested  
mith (1983) for  
ur purposes we  
oil for showing  
et those goals.  
different.

ient may take or

tion as any other  
ensate for water


Apart from ecological management, there are many other fields in which there is a need to build models pragmatically.

### LEARNING TO COMPROMISE

As a first step toward building pragmatic models, we must learn how to compromise. Figure 1.3 represents an ecosystem at different levels of resolution. The circles in each case represent the components of the ecosystem and the arrows indicate that one component has a direct influence on another. Double-headed arrows indicate a mutual interaction, while the thick arrows represent external influences (the "driving forces" on the model) such as sunshine, rainfall, and management actions.

We will use diagrams of this kind as a sort of logo at the beginning of each chapter to depict the contents of the chapter. In a model, the circles in the diagrams would be associated with variables (a circle might represent one variable or a group of variables) and the arrows with equations. An arrow that starts and ends on the same circle implies that the component has an influence on itself. (Examples are births within a population or density-dependent effects.)

Many people think of ecological modeling in terms of diagrams such as Fig. 1.3(a). Their preconception is that ecosystems are made up of components that interact in a complex way and that models should be built to represent that complexity. However, as we shall see in Chap. 5, it is not easy to build a model as complex as that shown in Fig. 1.3(a), and often



**Figure 1.3** Representing an ecosystem at three different levels of resolution: (a) a detailed system model, (b) isolating a part of the system, (c) a less detailed ("lumped") system model.

lds in which

learn how to  
levels of res-  
of the ecosys-  
t influence on  
ion, while the  
forces" on the

ie beginning of  
del, the circles  
le might repre-  
with equations.  
nat the compo-  
a population or

igrams such as  
ade up of com-  
ould be built to  
hap. 5, it is not  
.3(a), and often

the usefulness of such models, once they have been built, is disappointing. Our first compromise is therefore one of *simplification*. The way to accomplish this is to start with the management problem itself (rather than with a mental picture of the ecosystem), then to find the abstraction of the ecosystem that enables us to contribute effectively to the solution of the problem. This approach usually leads to models at a completely different level of resolution, as illustrated in Figs. 1.3(b) and (c).

In Fig. 1.3(b) we have a management problem that relates to only a part of the ecosystem (the three circles in Fig. 1.3(a) with narrow vertical lines). We therefore focus on that part and try to separate the relevant components and their interactions (the first-order effects) from the rest of the system. Obviously the context within which the chosen components interact cannot be ignored entirely. It is often useful to think of and represent the effect of the rest of the system on the subsystem as an artificial driving force (thick arrow). We will introduce various ways of doing this in Chaps. 2, 3, 4, and 8.

In Fig. 1.3(c) the management problem relates to the system as a whole, but to say something useful with our model we have represented the system at a coarser level of resolution. Here we have combined (or, to use modeling jargon, "lumped") all similarly shaded variables in Fig. 1.3(a) into grosser components and we have concentrated on only the key (or first-order) interactions in the system. We will introduce examples of this in Chaps. 5 and 6.

Obviously the appropriate level of resolution must depend on the management problem to be solved, but in choosing that level we must pay as much attention to what we are likely to achieve with the model as to how well the model represents the problem we are trying to solve. Choosing the appropriate level is thus a pragmatic compromise between the complexity of ecosystems on the one hand, and the need to solve a problem, with limited data and in a reasonable amount of time, on the other. Much of this book will be concerned with learning how to make the most of that compromise.

To quote Bernard Berenson: "Representation is a compromise with chaos."

## FURTHER READING

An appealing definition of a model can be found on pages 7 and 8 of Hall and Day (1977). Approaches to modeling and management that overlap with those described in this chapter can be found in Silvert (1981), in the last chapter of Mann (1982), in a recent book by Walters (in press), and in a paper by Overton (1977). The latter contains a good discussion of resolution. Although the book by Tukey (1977) addresses problems of data analysis rather than modeling, it espouses a similar philosophy. Simon (1982) provides a philosophical and practical approach to complexity in general.

ls of resolution: (a) a  
, (c) a less detailed