

tell us much about the dynamics of the system. There is, however, a difference between populations described by eqn 3.6 and those that we studied in Figures 3.1–3.7. In eqn 3.6 the effects of population density act instantaneously on population growth rate. In our discrete-time examples there is an implicit time lag: the level of competition experienced by individuals in the current generation is set by interactions that occurred in the last generation. Such time lags tend to be destabilizing as they delay the onset of a reduction in population growth rates as densities climb, and make it more likely that any equilibrium is overshoot. Indeed, it is mathematically impossible for a population governed by eqn 3.6 to show chaos. We should stress that it is not the difference between continuous- and discrete-time formalisms that lies behind the contrasting stability, but the presence of the time lag. Indeed, in continuous time we can get exactly the same dynamics by explicitly making net population growth rates a function of previous population densities,

$$\frac{dN_t}{dt} = r(N_{t-\tau})N_t \quad (3.7)$$

where  $\tau$  is a time lag of approximately one generation.

Structured population models with density dependence can be built using the same matrix, integral equation, or partial differential equation approaches discussed in the section on density independence. Naturally they are more complex, and often with a greater potential for destabilizing time lags. Relaxing the assumption that all individuals are equal also leads to the possibility of more complicated types of interaction than are possible for unstructured populations. Competition may be asymmetric, typically with smaller individuals suffering disproportionately at the hands (or roots) of larger individuals. Moreover, cannibalism is much more common in the animal kingdom than often realized, and when it occurs it is nearly always size-related, with older larger individuals consuming their smaller conspecifics. Such age-specific interactions have been studied in detail, particularly in insect systems that can be maintained in the laboratory for multiple

generations. Many of these systems show population cycles with periods shorter than those predicted by unstructured models (e.g. Figure 3.7). The details differ with the natural history of the different systems but a common pattern is for an older cohort of individuals to reduce the numbers in a younger cohort by out-competing them for food or through cannibalism. When the depleted younger cohort grow old enough to be dominant competitors or cannibals themselves there are not enough of them to reduce significantly the next cohort coming through. This means that the next group of individuals to mature into the older cohort are very numerous and decimate the current younger cohort, and the cycle begins again.

### 3.2.1 Chaos

The pioneers of modern mathematical dynamics, particularly Poincaré at the end of the nineteenth century, realized that the behaviour of highly nonlinear systems could be very odd, but in the absence of computers to help visualize their dynamics, progress on understanding what was happening was very slow. When computers began to become available in the 1960s workers in fields such as meteorology and ecology were able to see the complex dynamics produced by beguilingly straightforward equations, and this led to a burst of interest in both pure and applied mathematics that laid the foundations of the modern field of chaotic dynamics. In population ecology, the classic paper is May's *Simple mathematical models with very complicated dynamics* (May, 1976a), which not only introduced the notion of chaos to the field but showed that lurking underneath the seeming unpredictability of chaotic dynamics was considerable order and pattern. We shall now explore a population model of exactly the type that May analysed.

Chaos has already been encountered in this chapter as the dynamics that emerge in a simple discrete-time population model as the population growth curve (or map) becomes sufficiently nonlinear (the 'humpiness' of the curves in Figure 3.7). Let us now specify a family of curves that can give rise to the maps in Figures 3.3 and 3.7. For reasons that will be explained in a few paragraphs it does

not particularly matter which family we chose, and we plump for the Ricker equation as it is commonly used in applied population biology, particularly for fisheries (Ricker, 1954).

$$n_{t+1} = n_t \exp[r(1 - n_t)] \quad (3.8)$$

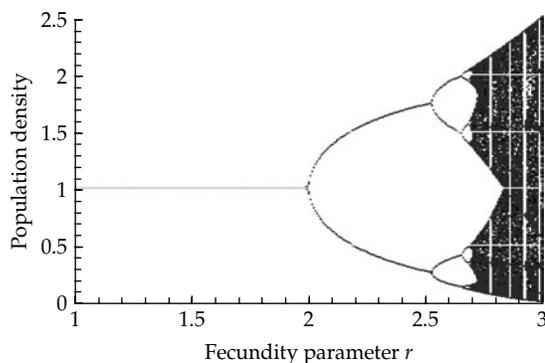
Here  $n_t$  is population density (scaled to equal 1 at equilibrium). When rare the population increases each generation by a factor  $\exp[r]$  but as densities approach 1 the increase slows and above 1 it reverses. If  $r$  is high there is the potential for the population to overshoot the equilibrium.

We want to picture the dynamics of the whole system for different values of the sole adjustable parameter  $r$ . To do this, imagine iterating the equation by cobwebbing as in Figures 3.3 and 3.7 and then throwing away all the transient dynamics, perhaps the first 50 generations. For Figure 3.3 (corresponding to a value of  $r=1$ ) the non-transient dynamics would not be very interesting: it would simply be a population at stable equilibrium, in this case  $n=1$ . In Figure 3.8 we plot  $r$  along the  $x$  axis and the non-transient dynamics on the  $y$  axis; for  $r=1$  there is a single point at  $n=1$ . The dynamics of the population described by the first panel in Figure 3.7 ( $r=1.9$ ) differ only in their transient behaviour and so it too would be represented by a single point at  $n=1$ . Indeed, for the Ricker equation a stable equilibrium occurs for all persistent populations with  $r < 2$ , which gives the straight line at  $n=1$  in the left-hand part of Figure 3.8.

The persistent dynamics depicted by the middle panel of Figure 3.7 ( $r=2.3$ ) are a two-point limit

cycle: the population oscillates for ever between two densities, one greater and one less than the now unstable equilibrium  $n=1$ . In Figure 3.8 this appears as two points. The value of this representation now becomes clear because instead of having to try to compare a large number of cobweb diagrams we can see at one glance how the cycles appear at  $r=2$  and then increase in amplitude as  $r$  gets bigger. The change of behaviour at  $r=2$  is called for obvious reasons a bifurcation, and the representation itself is a bifurcation diagram. We can also see that at  $r=2.5$  a second bifurcation occurs to give a four-point cycle, and then further bifurcations at increasingly smaller intervals of  $r$  until a limit is reached. "What happens at the point of accumulation [the limit]?" is what May scrawled on a blackboard in the Theoretical Physics Department at Sydney University in the early 1970s.

May showed that what happens is chaos. As the third panel in Figure 3.7 illustrates, the trajectory never converges on a simple cycle but fluctuates aperiodically around very many values of  $n$ , never repeating itself. This is represented in Figure 3.8 by a vertical line containing numerous, in fact an infinite number of, points. Cobwebbing can also be used to demonstrate a cardinal property of chaos: namely sensitivity to initial conditions. Start two trajectories very close together and sooner or later they will diverge. This is not due to a lack of computing power: no matter how close the two initial values they will come to diverge. More accurate estimation of initial values, so that the



**Figure 3.8** The bifurcation diagram for the Ricker population model (see the legend of Figure 3.7).

measured value is close to the 'true' value, can delay the divergence, but not prevent it, and this means that there is an absolute limit to our ability to predict into the future the behaviour of chaotic systems.

The dynamics of a system can be described with a quantity known as the Lyapunov exponent (named after a Russian mathematician whose name is also transliterated Liapunov or Ljapunov). The Lyapunov exponent describes the rate of separation of infinitesimally close trajectories; a positive value means that the trajectories diverge exponentially, and this extreme sensitivity to initial conditions is the hallmark of chaos. Algorithms have been derived to estimate Lyapunov exponents directly from time series (Wolf *et al.*, 1985) and have proved very valuable, especially in the physical sciences where relatively long time series are typically easier to obtain.

Bifurcation diagrams are beautiful objects that contain a wealth of mathematical detail. They have quite literally been the subject of tens or possibly hundreds of mathematics PhD theses. Our focus here is on their relevance to biology and we have space to mention only a very few more technical results. First, May and others showed that the patterns in Figure 3.8 do not just apply to the Ricker equation but to a very broad class of models that all show the same transition through period-doubling from order to chaos (May, 1973a, 1974c; Li and Yorke, 1975). There is a limited number of routes to chaos and one can derive general results that apply to very many systems. For example, the ratio of the interval of  $r$  values in which two-point cycles are found and in which four-point cycles are found is 4.6692. In fact the same ratio is found for every adjacent interval (four-point/eight-point, etc.) not only for the Ricker equation but for *every* map that shows this type of transition from order to chaos (Feigenbaum, 1978). Second, if you look closely at the bifurcation diagram to the right of the accumulation point you see that the region of chaos contains intervals of simpler dynamics, including period-three cycles that undergo their own transition back into chaos. In fact there is an infinite number of narrow, periodic windows. Finally, the bifurcation diagram has fractal structure: enlarge part of the region of chaos and you

will see a complex pattern of bifurcations, aperiodic and period trajectories; chose part of this picture and enlarge yet again and the same patterns appear in miniature, and so on *ad infinitum*.

The beauty of bifurcation diagrams is fragile: add a little stochastic noise—inescapable in real biological systems—and their more rococo patterns disappear. However, the extreme sensitivity to initial conditions, the signature of chaos, remains. So while it is not mathematically true that the Ricker model predicts chaos for all  $r > 2.69$  it might as well be for any biological purposes. Another biologically relevant property of chaos is also shown in Figure 3.8. Although precise prediction is not possible the different population trajectories are bounded, that is they cannot become arbitrarily large or small. A pure random walk would not be bounded (except of course by  $n = 0$ ). Indeed, it is sometimes possible to calculate the probability distribution of different population states. Depending on the system this may be valuable information for ecologists and population managers.

Chaos is not just a property of discrete-time systems and chaos in continuous time systems has also been extensively studied. Consider the non-transient behaviour of a continuous system. If the system is at equilibrium this will be a simple point but if there are persistent cycles or chaos then it will be a continuous line. For single-species populations this line can be plotted in a space where the coordinates are population densities now and at times in the past. For example, on a three-dimensional graph the coordinates might be densities now, 1 month ago, and 2 months ago. In this space a cycle will be a closed loop while a chaotic trajectory will be an object such as that on the left of Figure 3.9. This object looks like a twisted diaphanous sheet and is a fractal: successive magnifications of parts of the sheet show the same self-similar pattern. One point to note is that chaos occurs in simple (ordinary) differential equations only for systems of three or more variables: the dynamics of a two variable-system can be described in a two-dimensional space which does not allow for the twisting and mixing of trajectories that are the hallmarks of chaos.

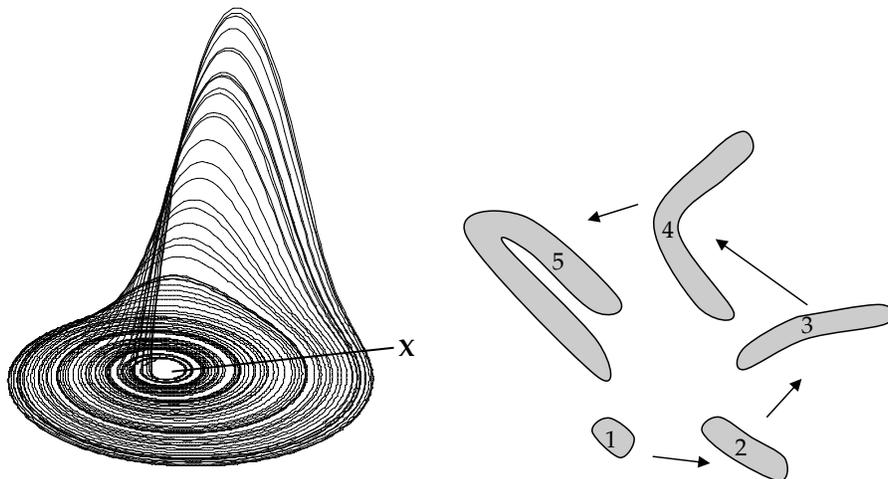
There is a close link between chaos and fractals. Objects that represent the non-transient behaviour

of a dynamic system are called *attractors* (because trajectories originating elsewhere in state space are attracted to them). In continuous time, points (stable equilibria) and closed loops (cycles) are examples of normal attractors whereas fractal objects such as that in Figure 3.9 are termed *strange attractors*. All chaotic systems are governed by strange attractors and, as we shall return to shortly, determining that a system's attractor is fractal is one way of identifying chaos in nature. The attractor in Figure 3.9 also provides an insight into why chaos is always associated with extreme sensitivity to initial conditions. The right-hand panel in Figure 3.9 is a cartoon to illustrate the evolution of a set of initially very similar trajectories: the bundle marked 1, which should be imagined as lying flat on the horizontal surface of the attractor in front of the line  $X$ . Flow on the attractor occurs in the counter-clockwise direction and sets of points are first stretched (2, 3) and then folded (4, 5). If you imagine this occurring numerous times it is easy to see how trajectories that start off near each other quickly become separated. The degree of stretching in a system is quantified by the Lyapunov exponent.

Chaos in continuous- and discrete-time systems is intimately related. Consider the section  $X$  (called a Poincaré section) through the attractor

in Figure 3.9. If the position along the section is treated as a variable, and if the position in the current traverse is plotted against that in the previous, one arrives at a map exactly equivalent to the chaotic Ricker map discussed above. Now, however, the  $r$  parameter is not simply a measure of single-species fecundity, but a more complex amalgam of the life histories of all species or development stages that influence the dynamics.

How might one seek to decide whether natural populations are chaotic? Typically this has to be done from time-series data, which at least in comparison with data from the physical sciences are inevitably of relatively short duration. There are two broad approaches. The first is to try to fit a flexible population model to the time-series data and then to determine by iterating the model whether the dynamics are chaotic. The second is to try directly to reconstruct the attractor governing the system and determine whether it is fractal. Both approaches are helped by a very important theorem (Takens, 1981) that states that the attractor of a multi-species or complex single-species interaction can always be reconstructed from single-variable time-series data in a space made up of a sufficient number of time-lagged dimensions (i.e. the coordinates are densities at time  $t$ ,  $t - \tau$ ,  $t - 2\tau \dots$  where  $\tau$  is a lag). The major proviso is



**Figure 3.9** Chaos in continuous time. The object on the left is a strange attractor describing the flow of trajectories of a continuous-time system in three-dimensional space (the Rössler attractor).  $X$  is a Poincaré section discussed in the text. The cartoon on the right describes how bunches of nearby trajectories become stretched and folded as they move around the attractor. See text for further details.

that you have to have sufficient data, which in practice is usually a very demanding requirement.

The first attempt to fit models to data did not use time series but life-history data on fecundity and density-dependent mortality. Hassell *et al.* (1976) fitted a two-parameter model to data from 24 species of insects with reasonably discrete generations and concluded that the vast majority had stable dynamics, indeed not even showing an oscillatory return to equilibrium. Although the authors were at pains to stress the provisional nature of their conclusions, this paper had a very major impact, and to a certain extent inadvertently licensed ecologists to treat chaos as a theoretical curiosity for the next decade.

The next major attempt to search for chaos used model-free approaches and was spurred by the growth of empirical chaos studies in the physical sciences (Schaffer, 1985; Schaffer and Kot, 1985a, 1985b; Olsen and Schaffer, 1990). The basic idea was to reconstruct the attractor by embedding the time series in time-lagged coordinates and then either to take a Poincaré section and look for a one-dimensional chaotic map, or to estimate the *attractor dimension*. In our daily lives we do not normally need tests to tell us whether an object is one-, two-, or three-dimensional but mathematicians who often work in much higher dimensional space have derived algorithms to estimate arbitrary dimensionality. When these are applied to fractal objects they return a non-integer dimension. A non-integer dimension implies a fractal and a fractal implies chaos. Though clearly worth trying, ultimately this research programme was defeated by the quality of the data available. To quote Schaffer (2000), 'Only in the instance of recurrent outbreaks of measles in human populations, was there sufficient data to justify our initial enthusiasm' and, he added, even here the argument chiefly rested on the comparison of time-series data with the output of epidemiological models.

In the last 15 years, interest has grown again in the challenge of detecting chaos from time series. Sugihara and May (1990) developed a technique called nonlinear forecasting which measures the extent to which predictability decays with time. In chaotic systems this occurs in a characteristic way determined by the magnitude of the Lyapunov

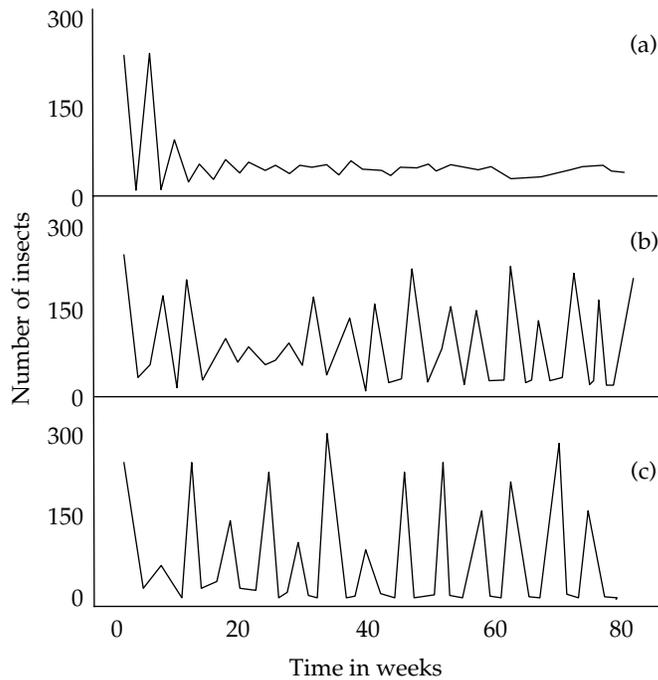
exponent. This method has since found wide application beyond biology in econometrics. Model-based approaches have also enjoyed renewed attention. One strand has sought to develop more accurate mechanistic population models, capitalizing on both the more powerful computing tools now available and statistical advances in extracting parameter values from data. A different strand, with similarities to Sugihara and May's approach, fits very flexible non-mechanistic population models to time-series data typically using response surfaces that are optimized either by traditional least-squares methods or more exotic techniques such as thin-plate splines or neural nets (Ellner and Turchin, 1995). The magnitude of the dominant Lyapunov exponent is calculated directly from the fitted model. It is still too early to judge the long-term value of these methods, although they have revealed a number of systems with apparent chaotic dynamics, in particularly involving human-disease and predator-prey interactions.

For single-species interactions, the best examples of possible chaos involve laboratory systems, including Nicholson's famous long-term blow fly experiment. A very nice experimental example is the work of Costantino *et al.* (1997) on the flour beetle, *Tribolium castaneum*. Recall we mentioned above that strong interactions between different life-history stages can give rise to complex dynamics. In *Tribolium*, adults and larvae cannibalize eggs while adults also eat pupae. A population model showed that by varying a single parameter (pupal mortality) the dynamics of the system moved from stability to chaos and then to a three-point cycle. Figure 3.10 shows that experimentally manipulating pupal mortality leads to dynamics that look very like those predicted. It is true that this is a highly artificial system, yet it is an impressive demonstration that the dynamics of these insects have been understood.

### 3.3 Randomness

#### 3.3.1 Types of random effect

Real animals, plants, and micro-organisms are continually buffeted by the effects of random



**Figure 3.10** Time series of the number of larval beetles in laboratory populations for different rates of pupal mortality which were artificially manipulated. Theoretical models predict that the population in panel a should have a stable equilibrium, panel b should be chaotic, and panel c should have a three-point cycle. The experimental data show good agreement with the predictions (after Costantino *et al.*, 1997).

processes and a critical question in population biology is the extent to which insights gained from the analysis of deterministic models survive the insults thrown at them by stochastic nature.

There are a variety of different ways in which random or stochastic effects can influence population dynamics (May, 1973a). Perhaps the most straightforward is *environmental stochasticity*, where the value of a demographic parameter changes over time. Recall the density-independent, discrete-time model  $N_{t+1} = N_t \lambda$  where  $\lambda$  is the annual population growth rate. This model implicitly assumes that the value of  $\lambda$  is constant, but in fact it will almost certainly vary from generation to generation; we might better write the equation  $N_{t+1} = N_t \lambda_t$  to emphasize this fact. Note that environmental stochasticity affects the demographic rates of all individuals in a population in the same way, and that this effect is independent of population size (Lande *et al.*, 2003). Much research in identifying factors generating environmental stochasticity has focused on climate (Stenseth *et al.*, 2002), although in principle any other factor with unpredictable effects on population parameters can contribute to this process.

Let us return to the discrete-time model of a population with non-overlapping generations,  $N_{t+1} = N_t \lambda$ , and for the sake of argument assume that the value of  $\lambda$  is actually constant over time. But this does not mean that every single individual in the population will produce exactly  $\lambda$  female offspring. In the real world there will always be some between-individual variation or *demographic stochasticity*. For example, consider a parasite that searches randomly for hosts into which it lays a single egg; if the average parasite lays  $\lambda$  female eggs then some will by chance discover more hosts and some by chance fewer. This is a Poisson process where the variance is the same as the mean. One can imagine other natural histories where the variance is much less than a Poisson process (vertebrates that normally produce one offspring a year) and others where the variance is much greater (organisms living in a highly heterogeneous environment). Now suppose the population is small: by chance all individuals in one generation may experience low reproduction and so the following year the population size would be significantly less than the expected  $N_t \lambda$ . Of course, the probability of simultaneous episodes

of good or bad luck become progressively more unlikely in larger populations and hence demographic stochasticity is most important in small populations. In many ways, its action is similar to drift in population genetics.

A further random process that is sometimes distinguished is catastrophic stochasticity: random events that destroy the whole population irrespective of its size or current demographic parameters. We shall not discuss this type of randomness further here, although it is particularly relevant to studies of metapopulations (see Chapter 4 in this volume) and also in conservation biology where populations may be wiped out by human action that can at least be approximated as a random process.

### 3.3.2 Density-independent populations

Let us now see how stochasticity affects population growth rate and population projection. For ease of explanation we shall stick to discrete-time models although the same principles apply to populations that reproduce in continuous time. Return once again to the model  $N_{t+1} = N_t \lambda_t$  where the subscript to the population growth rate emphasizes that it varies between generations, specifically with mean  $\bar{\lambda}$  and variance  $\sigma_{\lambda}$ . This is the way that randomness is most frequently dealt with in population models, and has been referred to as the equilibrium treatment of noise (Coulson *et al.*, 2004). If we take logarithms then we can write

$$\text{Log}[N_t] = \text{log}[N_0] + \sum_{x=0}^{t-1} \text{log}[\lambda_x]. \quad (3.9)$$

If the values of  $\lambda$  vary independently over time, then the right-hand term is the sum of independent random variables, which the Central Limit Theorem tells us is asymptotically normally distributed. This implies that population size itself is lognormally distributed. There are some complexities in calculating long-term population growth rates in this case (Lewontin and Cohen, 1969). An intuitive procedure might be to see how expected population size grows with time. A simple calculation reveals it increases exponentially at a rate determined by  $\bar{\lambda}$ . But the expected

population size is dominated by very rare, huge population sizes in the upper tail of the distribution. In fact the modal population size, the population size that will actually be observed in the field, grows not at a rate determined by the simple arithmetic mean,  $\bar{\lambda}$ , but the geometric mean  $(\lambda_0 \cdot \lambda_1 \cdot \lambda_2 \cdots \lambda_{t-1})^{1/t}$ .

Several biologically interesting results follow from this. First, as long as there is some variance in  $\lambda$  the geometric mean will always be lower than the arithmetic mean: poor years have a greater negative effect on population growth than the positive effect of good years. Second, a single year with zero net reproduction ( $\lambda = 0$ ) renders the long-term growth rate 0. This makes intuitive sense as the population goes extinct, but note that this is not what a calculation based on the arithmetic mean would suggest. Finally, recall that in the deterministic case persistence was very straightforward: a population would increase if  $\lambda > 1$  and decrease if  $\lambda < 1$ . The situation is now more complicated: populations with geometric mean growth rates less than one will always ultimately go extinct, but some may persist for a long period of time if by luck they experience a chain of propitious years. Similarly, although populations with geometric growth rates greater than 1 will tend to persist, some will by bad luck go extinct. In fact populations which will, on average, grow to infinity also have a probability of extinction of 1 for very long periods of time. This can be seen very simply: if  $E(N_t) = T^2$ , where  $t$  represents time and  $T$  is the length of time since the simulation began, the probability of extinction can be written as  $1 - 1/T$ . When  $T$  gets very large, the expected population size tends to infinity and the probability of extinction tends to unity as  $1/T$  approaches 0. It is possible to calculate the distribution of persistence times of populations governed by different distributions of growth rates, and this may be helpful in population management.

In many ways the population effects of demographic stochasticity are similar to its environmental counterpart. It will increase the variance in  $\lambda$  and so tend to reduce long-term growth rates, and increase the probability of extinction by bad luck. The major difference is that its effects become very weak as population size increases. Indeed, the

total variance in reproductive rates can be thought of as the sum of two components,  $V_E$  (environmental stochasticity) and  $V_D/N$  (demographic stochasticity divided by population size). A reasonable rule of thumb is that demographic stochasticity can be ignored for populations with more than 50 or so female breeders, though note that the population size of large carnivores, even in extensive nature reserves, can often be below this threshold.

We stated above that we were assuming that stochastic effects were uncorrelated over time. Often this will not be the case, especially for short-lived organisms that might, for example, have several generations in a single summer. Quite frequently there will be a positive correlation between the random component of population growth rates in successive seasons (the term *red noise* is sometimes used for these positively correlated random effects). The most important effect of correlated stochasticity is to increase the severity of poor breeding seasons that now tend to follow one another. We note in passing that correlated red noise may lead to patterns in population dynamics that may be very hard to distinguish from an underlying deterministic cause, especially in structured populations. There can also be correlations between environmental and demographic stochasticity, in particular the effect of demographic stochasticity on population growth may be higher in years when the consequences of environmental stochasticity are most severe, a clear concern in conservation biology.

The arguments above apply also to structured populations, though with some complications. First, there is no longer a simple relationship between arithmetic and geometric population growth rates, but a stochastic equivalent to the deterministic growth rate can be calculated (Tuljapurkar, 1982). As with the unstructured population, adding stochastic effects always reduces long-term growth rates. Second, certain age or stage classes may be much more susceptible to stochastic perturbation than others. Random effects may thus lead to perturbations that disrupt the age-structure of the population (*structural variance*; Coulson *et al.*, 2001; Lande *et al.*, 2002). Here, stochasticity influences the population

dynamics via two routes. First, stochasticity has a direct effect on the size and structure of the current population. Second, these changes influence the future trajectory of the population. This interaction between stochasticity and the deterministic skeleton is sometimes referred to as the active treatment of noise, and is currently an area of considerable interest in population biology research. Such effects always reduce the tendency of the population to reach a stable age distribution and, in anticipation of the next section, can also have important consequences on population regulation if the strength and action of density dependence is also influenced by population structure.

### 3.3.3 Density-dependent populations

In a real stochastic environment a population is highly unlikely to remain at the exact same equilibrium value from one generation to the next. But it is still reasonable to talk about an equilibrium if populations above a certain value tend to decline in numbers, and those below the same value tend to increase. Conceptually we can think of an equilibrium not as a fixed population density, but as a probability distribution that remains the same over time and which determines the likelihood of observing the population at any particular level of abundance (Turchin, 2003). Of course, we should also consider the possibility that a population, even one that tends to increase when rare, goes extinct through a run of bad breeding seasons.

More generally, stochastic effects can cause a population to shift from one type of dynamic behaviour to another. Figure 3.5 depicted the dynamics of a species with two locally stable equilibria; it is possible that a sufficiently large random perturbation can move the population from the domain of attraction of one equilibrium to that of the other. Similarly, where there is an Allee effect a species is unable to increase in density when rare so zero population density is locally stable; random effects can push a species density below the critical threshold that leads to extinction. It is also possible that a species that for some reason has fallen below the threshold can be rescued by a random set of good breeding seasons. Of course, even when a species can increase when rare, stochastic extinction

is permanent if there are no sources of migrants to rescue the population. This treatment of stochasticity in population models has been called the passive treatment of noise.

The shape of the equilibrium probability distribution of abundances will obviously be determined by the magnitude and direction of the stochastic perturbations to the demographic parameters, but also by the dynamic consequences of the perturbations; that is, the interaction of the noise with the deterministic dynamics. Consider unstructured populations with deterministically stable equilibria which are approached either smoothly (Figure 3.3) or by damped oscillations (Figure 3.7a). It is very likely that the first population will tend to return towards the equilibrium faster than the population with damped oscillations, and for the same amount of environmental stochasticity will have a lower variance equilibrium population density. A population with an oscillatory approach to a stable equilibrium can more easily be prevented from reaching that equilibrium and thus appear to the observer to be persistently cyclic. This type of dynamic behaviour has been termed *quasicyclic* (Nisbet and Gurney, 1976) and has been seen in several experimental systems, including the flour beetle study described above as an example of chaos (Costantino *et al.*, 1997).

Consider an unstructured dynamic system that is at the edge of chaos, perhaps showing persistent cycles. If one or more parameters were changed slightly, it would move from persistent cycles into the region of chaos where its dynamics would be governed by a strange attractor. Near this threshold, the transient behaviour of the population before it settles into persistent cycles can be very complex. Although in this region there is not a strange attractor, dynamics may be influenced by an object called a *strange repeller* (Rand and Wilson, 1995), which like a strange attractor is a fractal, but repels rather attracts dynamic trajectories. One can think of the system like the ball in a pinball machine, careering from buffer to buffer, perhaps

for a significant period of time. Indeed, this behaviour may go on for ever if stochastic perturbations are large enough to prevent the system ever from settling on the stable cycles. The time series produced by such a process can be indistinguishable from chaos: it can show exactly the same extreme sensitivity to initial conditions, and attempts to reconstruct the attractor would suggest that it had a non-integer number of dimensions.

In discussing the bifurcation diagram in Figure 3.8 we already noted how random effects would interact with the deterministic component of the dynamics to give chaotic population behaviour throughout the region beyond the 'point of accumulation', even though here there are narrow windows of cyclic behaviour. As with chaotic repellers this is another example of the impossibility of separating the deterministic and stochastic aspects of population dynamics in general and chaos in particular.

Although it may seem unarguable that we should seek to develop models with both stochastic and deterministic components, exactly how to do this is not always obvious. For example, adding one type of noise to a model with a deterministically stable equilibrium and a different type of noise to a model governed by a chaotic attractor can produce dynamics that equally well match the type of data that ecological field studies produce. Also it is often not clear how stochasticity should be introduced into the model, onto which demographic parameters, and with what correlation structure. Nevertheless, we are optimistic about the future. For the analysis of time series and other observational data there are a variety of new statistical methods and techniques that will help identify the major stochastic drivers, and reveal how they interact with the underlying biology of the species (Coulson *et al.*, 2001; Lande *et al.*, 2003; Turchin, 2003; Stenseth *et al.*, 2004). There is also an increasing willingness of ecologists to experiment, both in the laboratory and the field, and to integrate modelling with experimental design and analysis.