

6

Is Nature Chaotic?




Figure 6.1 The number of pupae of the green bottle (sheep blowfly), in a laboratory population monitored every two days for two years. Data kindly made available to researchers by Robert Smith and colleagues (see <http://mcs.open.ac.uk/drm48/chaos/>).¹

... even if it were the case that the natural laws had no longer any secret for us, we could still only know the initial situation approximately... it may happen that small differences in the initial conditions produce very great ones in the final phenomena. A small error in the former will produce an enormous error in the latter. Prediction becomes impossible...

—Henri Poincaré (1908).²

Centuries before King Harold of England famously received an arrow in the eye (AD 1066), Chinese officials in the T'ang dynasty (AD 618–907) began collecting annual reports on the abundance of migratory locusts.^{3,4} The primary aim of this initiative was to make sense of the changes over time (the dynamics) of this devastating agricultural pest, and thereby predict the timing and intensity of outbreaks. Now, despite a staggering 1,300 years of faithful recording, few patterns are evident and the data look decidedly messy.⁵ Irregular climatic fluctuations, particularly those involved in the drying up of grasslands on river deltas, may explain some of the variability.⁴ However, one might wonder whether some of this 'messiness' was internally driven, caused by some sort of 'feedback' arising within the dynamics themselves. Many long-term data sets on population dynamics have these extremely messy qualities, ranging from the daily number of damselfish reaching maturity on the Great Barrier Reef⁶ to the number of feral sheep on Scottish Islands,⁷ and it is important to know where it all comes from.

The study of 'chaos' (easiest to define negatively as an absence of order, but we will get to a more formal definition later) has its roots in precisely the type of feedback processes referred to above, reflecting what mathematicians call 'non-linearities' (relationships that are not straight lines). Several mathematicians, most notably, the eminent French mathematician Henri Poincaré (1854–1912), had long noted that non-linear systems could generate some extremely unusual dynamics, such that the precise trajectory a system took was highly sensitive to the initial conditions. However, observations such as these were largely overlooked by ecologists until a new generation of researchers, notably Robert May (a physicist turned ecologist, now Lord May of Oxford), began toying with their own simple ecological models and appreciating that the behaviour of these models was not always simple.^{8,9} Until ecologists were made aware of the potential effects of non-linearities in the 1970s, the prevailing view was that complex dynamics must have complex causes. One of the many benefits of the development of chaos theory is that it has led to an appreciation that sometimes extremely complicated dynamics can arise out of the simplest and most innocuous looking of mathematical models, even those without any elements of chance built in.

We begin this chapter by describing one such simple model with potentially complicated dynamics, called the 'discrete-time logistic growth model'. A version of the logistic model was introduced in 1838 by Pierre François Verhulst (and later rediscovered by Raymond Pearl¹⁰) in an attempt to formalize arguments he encountered in Thomas Malthus' *An Essay on the Principle of Population*¹¹ (an essay that also had a famous influence on Charles Darwin's ideas). This mathematical model will help define

what chaos is, and how it is arrived at, before we go on to ask our ultimate question of whether natural populations fluctuate in a chaotic manner. If chaotic dynamics are a common feature of natural population fluctuations, then it has all sorts of implications for conservation biology, disease control, and many other areas of ecology; therefore, we take time to consider what a ‘yes’ answer would mean for ecology. We also ask some related questions, such as whether natural selection tends to produce population fluctuations that lack chaotic dynamics, and whether human intervention can make some non-chaotic populations chaotic and *vice versa*.

The fish pond

Imagine a population of fish in a pond. We census the population each year at the end of the breeding season; let the symbol x_t represent the population size of these fish in generation t , expressed as a fraction of the absolute maximum number of fish that could ever live there (this conveniently helps keep all numbers between 0 and 1). How might x_t vary over consecutive generations? At extremely low densities, each individual would have access to plenty of resources so it is likely that each individual would produce a relatively high number of offspring. In contrast at high density, individuals would be competing over resources, so that each individual would not leave as many surviving offspring. In effect, the population should ‘feedback’ on itself—at low population densities the *per capita* population growth rate would be relatively high, but at high population densities the *per capita* population growth rate would be relatively low. It is a good bet that something like this goes on in many populations—after all, no species on the planet goes through permanently unfettered geometric growth. We know this for sure, because (as Darwin had argued in the case of elephants¹²), were it any different, we would soon be up to our eyeballs in them.

How do we express this type of ‘density-dependent’ feedback mathematically? There are lots of different ways, many of which would yield qualitatively similar results, but one of the simplest is to simply let $x_{t+1} = rx_t(1-x_t)$ where r is a mathematical constant. Although we have largely avoided formal mathematics in this book, in this case it is worth working through the implications of this simple equation because of the insights that it provides. Here we see that when x_t is extremely small, then the index of population density in the next generation (x_{t+1}) is approximately rx_t (since $1-x_t$ is approximately equal to 1). In other words, the *per capita* population growth rate is almost r when the population size is small. Yet, when we increase x_t the feedback term $(1-x_t)$ now becomes increasingly smaller, so the *per capita* population growth rate diminishes. The mathematical function we have assumed might appear somewhat arbitrary, and probably there is not a population on Earth that actually shows precisely this dynamic, but it does the trick of introducing a feedback, and it makes sense to start with a simple rule. We also note in passing that this is a *discrete-time* version of the logistic equation, representing population size in the next generation as a function of population

size in the current generation. As such, the equations used to predict the changes in population size are called ‘difference equations’. However, we could let generation time tend to zero and end up with smoother, completely continuous changes in population size. Under these conditions, we would have a ‘differential equation’ (the way the logistic equation is often presented in ecology textbooks) and in this case the tools of calculus could be used to understand their dynamics.

Let us get back to the discrete-time logistic equation. We can see directly that the feedback involves a non-linearity when we plot x_{t+1} against x_t for a variety of values of x_t between 0 and 1 (Fig. 6.2). Thus, when $x_t = 0$ then the predicted population size in the next generation (x_{t+1}) is 0. Equally, when $x_t = 1$ then $x_{t+1} = 0$ (since $1 - x_t = 0$). Hence, only intermediate values of x_t generate non-zero values for x_{t+1} and the end result is a curve that bends over on itself rather like a hairpin (a ‘fold’). In fact, these particular curves are ‘parabolas’ (yes, the trajectory of a cannon ball we all know and love from high-school mathematics). Interestingly, increasing the value of r increases the intensity of the feedback and hence the severity of the folding (Fig. 6.2).

To see what dynamics are predicted by the model, we can start with a particular index of population density x_0 (e.g. 0.4) and simply update the equation iteratively




Figure 6.2 A ‘map’ of population density x_t against x_{t+1} as assumed by the logistic equation. Increasing the value of r increases the curvature of the relationship because changes are higher the higher the value of r . The graph also shows a line (dotted) in which x_t is plotted against x_{t+1} , allowing us to highlight where potential equilibria occur (for $r = 2$ population densities move towards this equilibria, but for $r = 3.5$ population changes are just too lumpy to allow the equilibrium to be converged upon and population densities vary around this equilibrium instead).

(i.e. in steps), calculating x_1 and placing it back into the equation to find x_2 and so on. You can try yourself—it requires no mathematics beyond arithmetic—yet had you done these simple calculations in the 1960s, *and realized their significance*, then you would have made an important scientific discovery. It turns out that the type of dynamics one predicts is only dependent on the value of r (and not the value of x in the starting generation). If r is relatively small (e.g. 2), then the fish population size always rises to a single value (the ‘equilibrium’) and stays there indefinitely (see Fig. 6.3a). The more mathematically minded reader might wish to confirm, by setting $x_{t+1} = x_t$, that this equilibrium is $(r - 1)/r$. Equilibria similar to these have a certain appeal and they imply a reassuring sense of stability and order. Indeed, before the 1970s, equilibrium solutions were the type of result most ecologists concentrated on when developing and exploring their models,¹³ almost going out of their way to ignore complications.¹⁴ The oversight comes in part from the absence of fast computers to help visualize the dynamics (the fastest computers in the world in 1970 were several orders of magnitude slower than a good modern desktop¹⁵), and it is no coincidence that the development of ideas about chaos came with the increase of computing power.

As we increase r further, then strange things happen. First, we get regular repeated cycles occurring in which the population overshoots the equilibrium then undershoots it, overshoots, then undershoots (Fig. 6.3b). This can be seen as a simple consequence of the discrete (‘lumpy’) nature of the change—the higher r , the higher the potential size of the changes from generation to generation, and the less fine-scale adjustment is possible (rather like adjusting temperature in a shower, in which the time delay between adjusting the handle and experiencing its effects means you can never get it just right). In Fig. 6.3a, the population compensates for being above or below the equilibrium value, so that each generation is closer to the equilibrium than the last. However, a higher r value makes the system feedbacks larger, and this higher sensitivity tends to lead to overcompensation and so the population never settles down to the equilibrium, but fluctuates around it. Note that although the system shown in Fig. 6.3b does not settle down to a single equilibrium, the dynamics are entirely predictable, so that the size of the population is always exactly as it was four generations previously. Increasing the value of r still further produces dynamics that seem to lack any sort of pattern at all (Fig. 6.3c)—it is no longer a question of consistently overshooting and undershooting but rather irregular behaviour that never quite repeats itself. Welcome to the world of chaos. There are no elements of chance whatsoever built into these dynamics (in the jargon, the model is ‘deterministic’, as opposed to ‘stochastic’)—the apparent noise is solely driven by the high degree of non-linearity in the system (it is not noise at all, but ‘deterministic chaos’).

Beautiful bifurcations

To see more clearly how chaos is arrived at, imagine starting at some arbitrary value ($x_0 = 0.3$ say, the exact value does not matter) and iteratively calculating the population

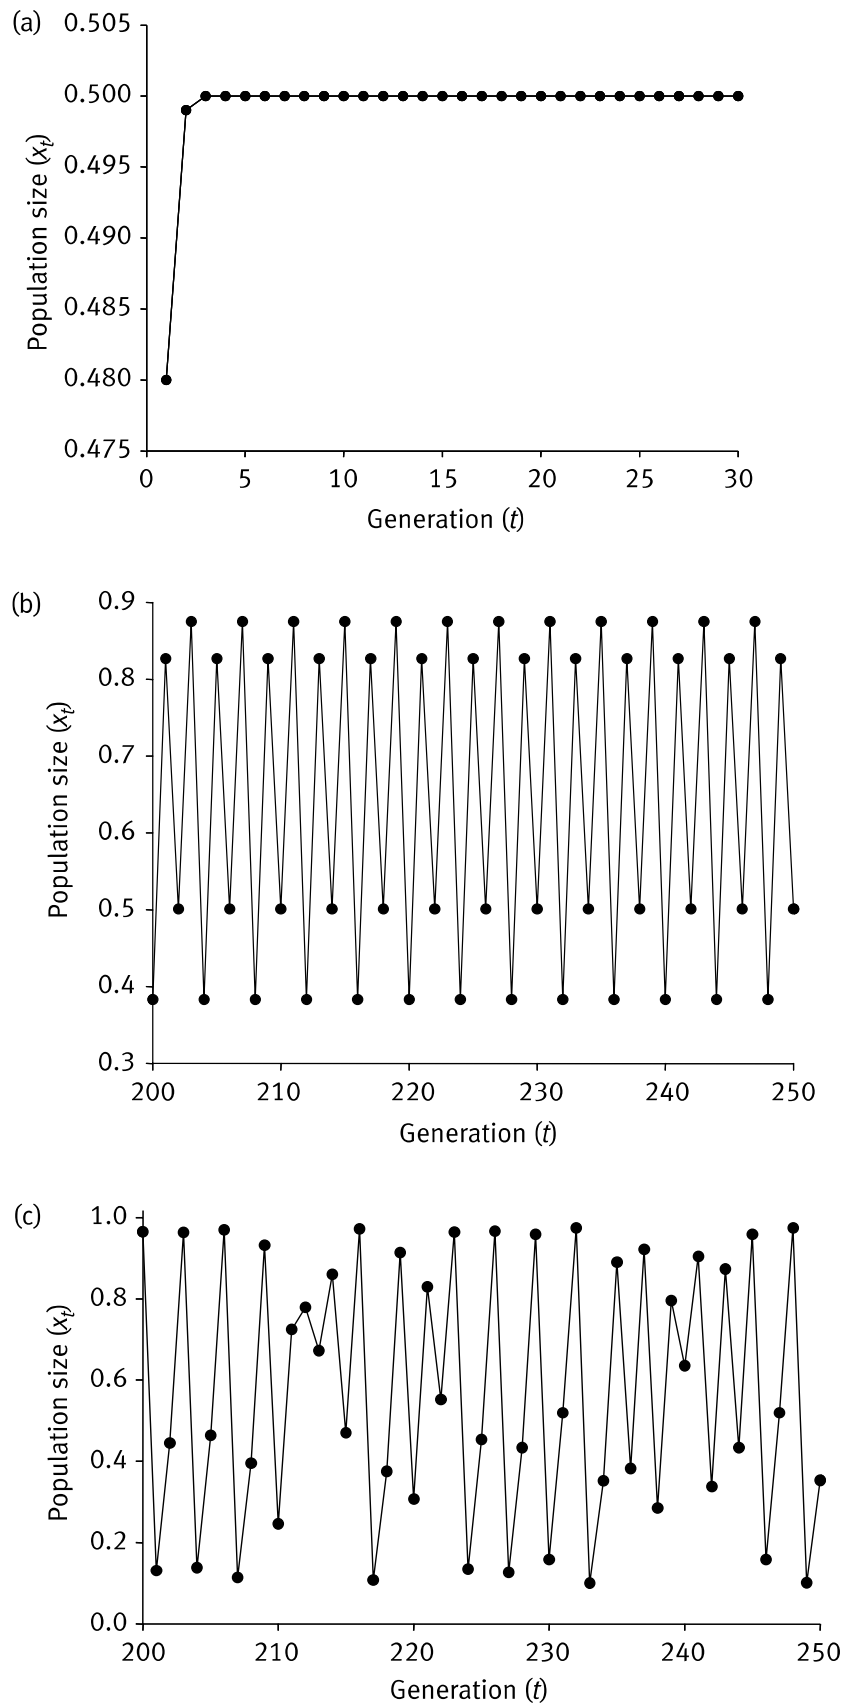


Figure 6.3 Starting with $x_0 = 0.4$, the graphs show the values of x_t iteratively calculated over multiple generations for different values of r . (a) When the parameter r is small (here $r = 2$), we observe a simple rise towards equilibrium. (b) When the parameter r is increased (here $r = 3.5$), we start to see cycles (here a four-point cycle is shown). (c) Increasing r even further (here $r = 3.9$), we start to see chaotic dynamics with no underlying pattern.

dynamics over 10,000 generations—very time-consuming to do ‘by hand’ but quick and easy on a computer. Now let us plot out the value of r on the x -axis, against the population size(s) calculated over the final 9,000 generations of the 10,000 generation iteration (Fig. 6.4). We ignore the first 1,000 generations because they will include the ‘transient’ population sizes that inevitably arise from the particular starting value (x_0). Once we have done this for one value of r , let us repeat the whole exercise for a slightly higher value of r until we have explored the full range of r . When r is low then the final 9,000 population sizes will be exactly the same as one another (the equilibrium) and so they will be represented by single point on the graph. As r increases, the equilibrium value changes (it increases in this case) but it is still a single value (the equilibrium) for a given value of r and so the 9,000 values are again represented by a single point. Nevertheless, as we increase r into the range at which two-point cycles arise, then the population sizes in the final 9,000 generations will fluctuate between two values and two points will start to appear on the graph (Fig. 6.4). You will see that there is a relationship between these two-point cycles and the former equilibrium, with the single line effectively




Figure 6.4 A bifurcation diagram, created here by plotting out the population sizes x_t in the final 9,000 generations of a 10,000 generation iteration of the logistic equation. We omit the early values because they may be ‘transients’, not typical of the steady-state dynamics. For low values of r , only a single value of x is recorded (the equilibrium) for all 9,000 generations. As r increases, then two values are reported (a period-2 cycle), then four values are reported (a period-4 cycle): 8, 16, 32, and so on. The range of values of r with a particular period cycle gets progressively smaller, allowing a point of accumulation beyond which an infinite number of points occur in the cycle and there is no pattern to the dynamics. Even within the chaotic regime, however, we can have ranges of r that give regular predictable cycles and these can be of odd numbers such as period 3.

'bifurcating' into two. As r increases further, we see the amplitude of the oscillations increases (another consequence of the increase in lumpiness) until another bifurcation takes place and a four-point cycle arises. After a smaller range of r then the period 4 cycle bifurcates into a period 8, then (over an increasingly smaller range) period 16, period 32, 64, 128, and so on until a 'point of accumulation' is reached where we break out to an infinite-point cycle. Beyond this point is the chaotic region, the period of the oscillation becomes infinite, and so the dynamics never repeat themselves.

The beauty of bifurcations does not end here. May and Oster¹⁶ proposed that there was something quite predictable to the cascade of period doublings, noting that the ratio of the intervals between successive period doublings was approximately constant, and they did some mathematical work to characterize it. About the same time, Mitchell Feigenbaum took on the challenge of measuring these ratios directly. Using a (now-ancient) Hewlett-Packard HP 65 programmable calculator, he observed that the ratio of the difference between the values at which successive period-doubling bifurcations arise rapidly approached a constant as the number of period doublings increased. This constant was eventually estimated as 4.6692 (to four decimal figures). The fact that the ratios are constant is surprising, but the really surprising thing that Feigenbaum discovered (and mathematicians subsequently helped formally understand) is that a whole range of dynamical equations that likewise have a chaotic region, such as the Ricker equation ($x_{t+1} = x_t \exp[r(1-x_t)]$) used in fisheries research, and the trigonometric mapping $x_{t+1} = k \sin(-x_t)$ used in pure mathematics, all have precisely the same Feigenbaum constant of 4.6692. In other words, the 'scaling ratio' of the bifurcation does not depend on the specific equation. Indeed, Feigenbaum's constant can be used to demonstrate that a model is capable of generating chaos even if it is not directly observed.

So, chaos can be seen as dynamics with an infinite number of points (never repeating) in a cycle—they are 'aperiodic'. Yet peer into the chaotic regime past the point of accumulation and you see that for certain values of r we get regular 3-point cycles. These 3-point cycles bifurcate into 6-point, then 12-point cycles each reaching its own point of accumulation. Elsewhere we have 5-point cycles bifurcating to 10- then 20-point cycles and so on. The bifurcation diagram has what we call fractal structure (more on this later), in that if we focus on smaller and smaller ranges of r and blow them up, we would see the same complex pattern dominated by chaos but with bifurcations once again breaking out. In fact, the first scientific paper to use the word 'chaos' in this context was by Tien-Yien Li and Jim Yorke in 1975, and it highlighted the unusual occurrence of cycles with an odd number of points and explored the implications. The authors entitled their paper 'Period three implies chaos'.¹⁷ Apparently, colleagues had suggested using a rather more sober description, but by using a catchy term the scientists (and many that followed) had an appealing banner under which to sell their work.¹³ Using a colourful label for a scientific idea can be very helpful in attracting attention to it; think of 'selfish genes', 'The Red Queen', or 'Gaia'. As Stephen Jay Gould¹⁸ argued 'phenomena without names... will probably not be recognized at all'.

Defining chaos

One of the problems with using terms with broad appeal is that it can also attract criticism for its lack of precision and scope for misunderstanding. ‘Chaos’ means different things to different people, so we have to be careful to use the term in a strict scientific way. Probably, the easiest definition of chaos is that it is an ‘intrinsically driven’ absence of order—this ‘absence of order’ in turn may be interpreted as dynamics that lack any underlying pattern, so that you cannot predict what is going to happen in the long term. Yet many populations have dynamics that appear to lack any form of pattern, and not all of it may be driven directly by internal feedbacks within the population itself. For example, weather may add what we might think of as extrinsic ‘noise’ to the underlying dynamic (think of an extremely crackly radio reception, which crackles with noise obscuring the ‘signal’ you are trying to listen to), and so might simple measurement ‘error’ (not mistakes *per se*, but chance sampling variation when attempting to estimate population size). So if chaos is ‘internally driven’ unpredictability, then we will need some good mathematical tools for distinguishing intrinsically driven disorder, from extrinsic ‘noise’. In other words, how can we tell whether the population dynamics of fish in our pond (or antelopes on a savannah, say) are truly chaotic? As we will see, thankfully chaos has some rather different properties than a sequence of random numbers.

One popular way to ascertain whether a *mathematical model* is capable of exhibiting chaos is to examine the nature of the non-linearity involved and examine how particular parameters might affect the extent of the non-linearity, just as we have done for the discrete-time logistic equation. Once characterized in this way, one can explore the impact of the non-linearity by identifying any potential bifurcation points, and the point of accumulation beyond which chaos lies. Bifurcation diagrams are usually straightforward to generate when there is one dynamical variable of interest (such as the population size of one species), but similar techniques can be used with multiple dynamical variables (such as the population sizes of several species simultaneously). Of course, it is hard to do these types of manipulations with real observations of natural populations (although experiments using flour beetles have met with certain success—as we describe later), so other techniques must be used to look for chaos in real data.

The butterfly effect

There is one property of chaotic dynamics that we have not mentioned yet, but it is such an important and universal property of chaotic systems, that it has now become its key defining characteristic.^{19,20} Mathematical models fluctuating chaotically always show extremely sensitive dependence on the initial conditions. This feature has been called the ‘butterfly effect’ following a 1972 talk by meteorologist, and father of modern chaos theory, Edward Lorenz,²¹ who sadly died in early 2008 as we were completing our book. Lorenz’s original insight came in the 1960s when he recognized and documented

the extreme sensitivity to initial conditions exhibited by a simple non-linear model of fluid convection in the atmosphere.²¹ Thus, the story goes that if weather systems were chaotic then an almost negligible change in local wind speed in South America, such as that created by a wing flap of a butterfly, may ultimately mean the difference between having a hurricane in the northern hemisphere and not having one. Of course, this sensitivity has nothing to do with butterflies *per se*, and butterflies do not directly trigger anything—it is simply that with chaos, a small difference will always cascade to produce uncorrelated futures (not necessarily bad ones either). The nursery rhyme ‘For the want of a nail, the shoe was lost; for the want of a shoe the horse was lost...’ captures some of this contingency.

Naturally, the butterfly effect could be called something else, such as the ‘seagull effect’ (Lorenz’s original metaphor). However, the butterfly neatly captures the shape of Lorenz’s strange attractor (Fig. 6.5, see later for a full explanation) and, bizarrely, Ray Bradbury’s 1952 short story *A Sound of Thunder* also uses a butterfly to depict the nature of extreme sensitivity. In this story, a prehistoric butterfly is crushed underfoot by a time-traveller and this perturbation to the world is sufficient to change the outcome




Figure 6.5 The famous Lorenz attractor, derived from a continuous-time model of fluid flow in the atmosphere involving three dynamical variables (x , y , z). Under chaotic conditions the same dynamics never repeat, so the continuous line never intersects itself.

of a presidential election many years later. Apparently, Al Gore hired a mathematician to teach him chaos theory after failing to gain the U.S. presidency in 1988.¹³ In almost poetic irony, given Bradbury's short story, it appears that another butterfly—this time the infamous butterfly ballot paper in Florida—may have cost Gore the 2000 election.

Mathematicians have characterized this extreme sensitivity to initial conditions, noting that if a system has chaotic dynamics, then the difference between the trajectories of two populations that have slightly different initial conditions grows exponentially (geometrically) until this difference is essentially as large as the variation in either trajectory. At this point, the two population trajectories have no relationship to one another, although they may have started out at almost the same densities. The rate at which trajectories from similar, but not identical, starting conditions diverge from one another can be characterized by a quantity (or a series of quantities) known as a 'Lyapunov exponent(s)' (spelt in various ways) after the Russian mathematician, Aleksandr Lyapunov. A positive Lyapunov value means that the trajectories do indeed diverge exponentially from one another. In effect, due to their sensitivity to initial conditions, chaotic systems are 'noise amplifiers' while non-chaotic systems with deterministic rules tend to be 'noise mufflers'.²²

Let us stop to think what this means. If natural populations (or the weather, or atmospheric carbon dioxide levels, or whatever dynamics we are interested in) did fluctuate chaotically, then we could give up on long-term forecasting. We cannot measure the 'start conditions' with infinite precision (imagine trying to record all aspects of the weather simultaneously across the entire globe—or even one small part of it—with complete accuracy), so even if we had the best mathematical model one could ever produce, then the difference between what we thought would happen and what will happen will diverge exponentially. This is just one 'casualty of chaos' and we will return to the full casualty list later, as it helps to demonstrate the great potential significance of chaos.

Fatal attraction

Another important way of determining whether a particular mathematical model or ecological data set exhibits chaos is to present the dynamics in a rather different way, not as population size (or whatever variable you are interested in, such as temperature) against time, but as population sizes against one another. This is most easily seen when there are two or more variables such as densities of a predator and a prey species, or densities of three competing species. Instead of plotting the number of predators, and the number of prey separately against time, we can plot the number of predators at given times against the number of prey at the same times directly, and effectively ignore time. The technical term for displaying dynamics in this way is to show the results in 'phase space'. If predators and prey quickly reach an equilibrium, then this equilibrium will appear as a single point on a graph of predators vs prey, and the dynamics will stay at that point for all the remaining time. In effect, the dynamics will appear as if predators and prey get 'sucked in' to an equilibrium point in phase space, and this

equilibrium point is therefore known as an ‘attractor’. Other forms of attractor are also possible. For example, if the number of predators goes up when there are plenty of prey to eat, but the number of prey goes down when there are many predators, then predators and prey might enter into regular and predictable cycles. If we plot these cycles not as predator vs time, and prey vs time, but as predators vs prey, then again we would see an attractor, but this attractor would be a regular orbit (a closed loop), with predators and prey continually circling around it.

Now let us consider what chaos would look like in phase space. Chaos almost by definition must be bounded—while lacking order, the variable(s) in question should fall in a finite range between extinction and unfettered growth. Yet at the same time, the lack of order means that the same pattern is never repeated (if it did so, then with no built-in elements of chance, the dynamics would simply have to repeat itself). Imagine, therefore, the long-term dynamics of a population as an infinitely long ball of wool. How can you get an infinitely long ball of wool into a finite space without ever crossing over (repeating) itself? The answer is by having peculiar properties of folding and self-similarity that we alluded to earlier when discussing bifurcation diagrams. In other




Figure 6.6 A strange attractor in three-dimensional ‘phase’ space, showing how the population dynamics of a predator and two prey species fluctuate together over time. These chaotic dynamics were discovered by Michael Gilpin²³ in a simple model in which predators and prey continually reproduce and interact with one another.

words, our ‘attractor’ has to be a strange geometrical object, and for this reason chaotic attractors are known as ‘strange attractors’. Figure 6.5 (showing the relationship between dynamical variables in Edward Lorenz’s metrological model²¹) and Fig. 6.6 (showing the relationship between predator and prey density in Michael E. Gilpin’s three-species ecological model²³) each depict strange attractors. They look weird, and indeed they are. Their complex beauty has not only attracted biologists, but also artists, art historians, and poets.²⁴ Choose points on two separate lines and you will see that the trajectories rapidly diverge from one another—the stretching and folding effectively pulls them apart—reflecting the high sensitivity to initial conditions. One quantitative measure of strangeness is a measure of their self-similarity at different scales, an attribute that is measured by their ‘fractal dimension’. It is primarily for this reason that chaos is associated with the world of fractals, although we will not be exploring fractals any further in this chapter.

What chaos is, and is not

We now get to a workable definition of chaos and clear up a few misconceptions. A recent definition was proposed by Cushing and colleagues in their book *Chaos in Ecology*.²⁵ In their definition, which we will use ourselves, they combine elements of disorder, ‘boundedness’ and sensitivity to initial conditions all in one: ‘a trajectory is chaotic if it is bounded in magnitude, is neither periodic nor approaches a periodic state, and is sensitive to initial conditions’. So, it is the sensitivity to initial conditions that provides a key clue to chaotic dynamics.

The first potential misconception is easily cleared up by pointing out that chaos is not only a property of mathematical models expressed in terms of difference equations. We introduced chaos through a simple difference equation, but models based on continuous changes can also exhibit chaos—indeed the two strange attractors in Figs. 6.5 and 6.6 were generated by models with continuous rates of change involving three dynamical variables. It turns out that chaos only occurs in simple differential equation systems involving three or more variables,²⁰ but the possibilities for chaos get richer as we increase the number of variables.²⁶ As mathematician Mark Kot²⁷ noted, ‘As soon as you move to three or more species, there are hundreds of ways to get chaos’. Second, while early researchers were taken aback by the complex dynamics predicted by simple sets of equations with no elements of chance involved (so-called deterministic equations), and many investigators continue to emphasize chaos as a primarily deterministic phenomenon, work has also been done to understand the role of small random elements (noise) in these chaotic systems.²² For example, small amounts of noise added to the dynamic can make something of a mess of bifurcation diagrams we described earlier, but the underlying bifurcations are still evident and the extreme sensitivity to initial conditions remains.²⁸ Despite this, depending on one’s specific definitions, noise may have the potential to turn non-chaotic systems intrinsically chaotic,²⁹ thereby creating much *more* unpredictability than one would expect from the random elements alone.

The role of noise is currently under debate^{30–32} but it is clear that noise may do much more than provide a fuzzy cloud around a deterministic skeleton. We leave further consideration of the influence of such ‘stochasticities’, particularly in connection with cycling populations, until later in this chapter.

The casualties of chaos

Now that we know what chaos is, we briefly ask what are its implications if it turns out that many populations do indeed exhibit chaotic dynamics. In other words, is it worth finding out whether natural populations are chaotic? We have already pointed to several potential benefits of this branch of research. In particular, if ecologists observe a fluctuating population, it is only natural to wonder whether the fluctuations are caused by external environmental events such as temperature and rainfall, or whether they are caused by internal feedbacks within the population itself. By carefully analysing the data and looking for the signatures of chaos, we can hope to find out.

In some ways, just asking the question moves the debate forward. For example, for several decades in the past century there was a heated debate over whether natural populations were regulated by internal density-dependent mechanisms (such as competition for resources) or external density-independent mechanisms (such as periods of bad weather). At first glance, one might assume (as many population biologists did) that density-dependent mechanisms would tend to produce stable dynamics, while density-independent mechanisms would tend to produce erratic fluctuations. Yet, armed with an understanding of chaos, all this is turned on its head—too strong a density-dependent feedback, and one could end up with highly erratic fluctuations. Not only does this suggest that standard tests for density dependence in time-series data are invalid for chaotic systems, but it also means that just because you see unusual fluctuations does not mean that there is no density-dependence operating. This has implications for global ecology, not just population dynamics—for example, asking questions about the potential regulation of carbon dioxide or oxygen on Earth over geological time.

Perhaps the single most important reason why it is helpful to know whether populations are chaotic relates to the sensitivity with respect to initial conditions. Spontaneous, unpredictable events are a central element of quantum theory, but the thought that natural populations could also show extreme unpredictability due to the sensitivity of their dynamics must have come as a shock to many ecologists. As already noted, if a high proportion of natural populations fluctuated chaotically due to intrinsic non-linear feedbacks, then we could hang on to very short-term forecasting, but kiss goodbye to the goal of long-term forecasting (this appears to be the case for local weather forecasting). We might be able to say statistically what the mean and likely range of population sizes were (in the same way that we can predict climate into the future, even if local weather forecasts are much more constrained), but beyond that we can simply give up. Of course most of us take long-term forecasts of any complex system, be it the weather,

the economy, or a natural population, with a pinch of salt, but if they were chaotic, we know that our predictive power has built-in limitations, and we might usefully be able to quantify these limitations. For example, if global weather patterns are chaotic then we might be able to describe how the accuracy of predictions might be expected to decay over longer and longer prediction intervals.

Is mother nature a strange attractor?

Therefore, the crunch question is: do natural populations fluctuate chaotically? To answer this, we must quantitatively examine real data on the estimated size of some specified population over many generations. Early approaches to address the question involved assuming that a particular mathematical model (which was capable of exhibiting chaos under some conditions) was an accurate descriptor of the underlying dynamics. The model was then fitted to the observational data and the parameters were estimated (such as the parameter r in the discrete-time logistic). If the estimated parameter values were such that they would generate chaotic dynamics in the model, then one might be tempted to believe that the dynamics being investigated were also chaotic.

Population biologist Mike Hassell and colleagues³³ took just the above approach in 1976 when they fitted a general discrete-time population model involving three parameters (α , λ , and β) to 28 different data sets on the dynamics of insects (24 from field situations, and four from laboratory studies). It turns out that in their particular model, not one but two parameters— λ (growth rate) and β (a competition coefficient)—influence the degree of non-linearity, and that only high combinations of both λ and β would generate conditions sufficient to produce chaos. After fitting the model, the authors cautiously concluded that the vast majority of insect data sets had parameter combinations that would put the dynamics into a simple equilibrium, and only one case—the classical laboratory study of blowflies conducted by Nicholson³⁴—had parameter combinations that would put the dynamics into the chaotic regime. However, as the authors pointed out, even this case may have simply arisen as a laboratory artefact—the flies were not subject to many natural mortality factors such as parasitic wasps, which may have exaggerated the non-linear qualities of the dynamics.

Of course, the entire model-fitting approach is fraught with problems,^{35,36} not the least of which is that one must be extremely confident that the model you have fitted does indeed represent the underlying dynamics. Another related objection is that the dynamics of natural populations are often dependent on the influence of many other species, so fitting such a simple model is inappropriate—although one might argue (with a degree of mathematical justification) that the fitted model could be considered a representation of the outcome of all relevant species interactions.³⁷ Nevertheless, Hassell's approach remained an obvious and sensible way to treat the data, especially since it helped readers see the underlying chaos (still a novel concept at that time) in the mathematical model, before the model was fitted to the data.

Thomas and colleagues³⁸ followed suit in 1980, this time fitting a ‘ θ -Ricker’ model capable of exhibiting chaos, to the dynamics of 27 species of fruit fly in the laboratory. They expected the functions to be highly folded because they maintained very high population growth rates by changing the flies’ food regularly. Somewhat surprisingly however, when they came to analyse their data they found that the estimated parameters were not sufficient to put any of the 27 species into the chaotic regime. A similar result was obtained when Mueller and Ayala³⁹ examined the dynamics of 25 genetically distinct populations of the fruit fly *Drosophila melanogaster* and found no evidence that the estimated population parameters were sufficiently large to push the dynamic into the chaotic regime. Collectively these studies, on an impressively large number of species populations, were sufficient to convince many ecologists that chaos was relatively unimportant in natural populations, and therefore simply a ‘plaything’ for theorists. More recent studies applying much the same models, such as the Hassell model to the population dynamics of a number of species of weeds, have likewise come to the conclusion that the dynamics were not chaotic.⁴⁰

Nevertheless it has not all been one-way traffic. In the mid-1980s, Schaffer and Kot⁴¹ began looking at the dynamics of measles cases in New York City and Baltimore, reported monthly from 1928 to 1963. Before the widespread employment of vaccines, measles epidemics arose almost every year in large American and European cities, but major peaks were unpredictable, occurring every second or third year in New York and less frequently in Baltimore. The combination of seasonal ‘forcing’ (contact rates among school children are higher in the winter when schools are in session, compared to the summer) and feedbacks via gaining immunity made childhood infections very plausible candidates for chaotic dynamics. Rather than fitting a model, the authors attempted to reconstruct the attractor and test whether it had sufficient strangeness (stretching and folding) to qualify as a strange attractor. Of course with only a single dynamical variable, it is hard to produce a strange attractor directly, but thanks to a neat solution proposed by the physicist Floris Takens,³⁶ it is possible to plot the number of measles cases at time t against the number of cases at time $t+\tau$ and the number of cases against time $t+2\tau$ with τ a variable time difference and thereby (assuming you have enough data) build up an equivalent picture of the underlying dynamic. Although chance may have played some role in generating the measles unpredictability (e.g. some cases will go unreported, and reports will be lost), by reconstructing and analysing the attractor that the epidemics represented (followed up by an estimate of the Lyapunov exponent) the authors argued that there was a strong deterministic component to this unpredictability (chaos) in both the New York and Baltimore data sets. More recent analyses have generally supported these conclusions,⁴² including an analysis based on time-series analysis.⁴³ However, doubts still remain,⁴⁴ most notably because underlying factors such as birth rates have changed over time, and because the amount of seasonal forcing required to generate chaos in mathematical models of measles epidemics is considerably more than actually observed.⁴⁵

The year before (in 1984) Schaffer⁴⁶ had analysed the oscillatory dynamics of the Canadian lynx as recorded by the numbers of skins shipped yearly by the Hudson’s

Bay Company over the 1800s and 1900s—these data have been a classic of ecology textbooks since the 1920s.^{47,48} As the author himself has noted, these data made a somewhat less convincing case, but they were again suggestive of chaotic signal (and indeed more recent analyses provide additional support for this⁴⁹).

By the mid-1980s, Schaffer was beginning to see sufficient signs of chaos, that he issued (with Kot) a call to arms, arguing that ecologists were ignoring the very real possibility that chaos could be an important component of ecological systems and likened chaos to ‘the coals that Newcastle forgot’⁵⁰ (the implication being that this rich vein of science was under the noses of ecologists and they did not realize its potential). With titles like that, coupled with the bestseller (and all-round wonderful read) *Chaos* by James Gleick,⁵¹ scientists were well and truly waking up to the possibility of chaos.

In 1991, Tilman and Wedin provided experimental field evidence of the signature of chaos in dynamics of the perennial grass, *Agrostis scabra*, grown at two different initial densities on 10 different soil mixtures. For progressively richer soils, the dynamics evaluated over 5 years tended to exhibit higher amplitude oscillations with the richest (highest nitrogen) soil exhibiting dynamics the authors described as chaotic. Of course, with only a 5-year data set this interpretation is at best speculative (once again model-fitting methods and parameter estimation were used), but the inherent time scale of the annual dynamics clearly poses experimental challenges. One potential source for the significant non-linearity was the accumulation of leaf litter. Thus, in high-density years the accumulation of leaf litter (dead plant material at the end of the growing season) may inhibit growth of the following year. More recent work on the dynamics of another plant species, an annual greenhouse weed *Cardamine pensylvanica* likewise found evidence of oscillatory dynamics over 15 years, but in this instance found no evidence of chaos.⁵²

Chaos in small mammals?

The regular oscillations of small mammal populations such as voles and lemmings have given population ecologists plenty of data (they are often pests of forestry plantations and leave signs—such as grass clippings and bark scrapings—which can be used to estimate their densities).⁵³ No, lemmings do not jump off cliffs into the sea on a ‘suicide drive’ as Disney’s 1958 documentary *White Wilderness* would have us believe (indeed, the shot of lemmings jumping was entirely contrived—not only were they pushed, but also the sequence was filmed in Alberta, Canada, which has neither lemmings nor sea). However, many populations of small mammals exhibit remarkable high-amplitude 3–5 year oscillations in population size. These dynamics appear rather different in form in different regions. In particular, in southern Fennoscandia (including the Scandinavian Peninsula, Finland, and Denmark) and central Europe, populations seem to exhibit far lower amplitude fluctuations than in northern Fennoscandia. A possible reason for this is that the density of generalist predators is low in the north, and here specialist predators, notably the weasel, drive the dynamics.⁵⁴

It has been proposed that the shift from south to north in small mammal dynamics is not from equilibrium to an entirely regular cycle, but rather from stability to chaos. Indeed, Hanski and colleagues⁵⁵ analysed data on the population sizes of *Microtus* voles in western Finland and, on the basis of time-series analysis that revealed positive Lyapunov exponents, the authors argued that the observed dynamics in these populations were chaotic (with the exception of the most southerly population), albeit with a significant periodic component. In essence, chaos may be superimposed on top of a more regular signal. They supported their interpretation with a predator–prey model involving seasonality that readily generates the type of chaos they had revealed in the data. Nevertheless, it is fair to say that not all researchers are fully convinced, and there has been considerable debate over the issue, centring on how the Lyapunov exponent (and its likely range) is best estimated in systems involving noise.^{56–58} Of course, added noise is inevitable if you are trying to evaluate vole density across a large part of Finland using indirect methods of estimation.

Chaos in the laboratory

Perhaps the most ambitious set of experiments to investigate chaos was conducted in the laboratory on a species of flour beetle *Tribolium castaneum*.^{59,60} *Tribolium* is cannibalistic, with older individuals eating smaller ones, so if the population is at high density then many small larvae will get eaten by the older individuals, reducing recruitment into the next generation. After modelling *Tribolium* dynamics using an age-structured population consisting of larvae, pupae, and adults, the authors concluded that the cannibalistic feedback was capable of generating chaos, as well as unusual dynamics that never quite repeats itself but does not show the sensitivity to initial conditions ('quasi-periodicity'). Having a theoretical model to play with is a helpful way to judge when and where interesting things might happen, and understand why. More importantly, the authors combined this modelling approach with a replicated experimental study in which they artificially manipulated the adult mortality rate⁵⁹ and recruitment rates of pupae to adult stage,⁶⁰ and in each case they found good evidence of the predicted shifts in the dynamics (from stable points, to cycles and quasi-periodicity or stable points through a range of dynamical behaviours ending with chaos). One might argue that by manipulating the ecology, the authors have forced the system to match the model rather than the other way around, but this study remains convincing evidence that populations are at least capable of exhibiting chaos.

Even more recently, Becks and colleagues⁶¹ have manipulated the dynamics of a bacteria-eating ciliate predator and two species of bacteria (rod-shaped and coccus) in a chemostat: a rearing facility ensuring approximately constant environmental conditions. By experimentally manipulating the rate of delivery of the organic food source for the bacteria to the chemostat, the authors found that they could change the underlying dynamic between equilibria, stable cycles, and chaos (based on Lyapunov exponents). Precisely why chaos was generated is unclear, but the system has parallels to

Gilpin's model²³ of a one predator and two prey system. Moreover, this is the first case we know of in which chaos has been demonstrated in a microbial system. A second related example followed in 2008, and involved culturing a functioning planktonic food web isolated from the Baltic Sea under standardized laboratory conditions.⁶² Despite constant external conditions, this microscopic community, which consisted of bacteria, several phytoplankton species, herbivorous and predatory zooplankton, and detritivores, showed marked fluctuations in abundance over the 2,319-day experimental period and yet the populations still persisted intact. Moreover, the dynamics had all the hallmarks of chaos, including positive Lyapunov exponents for each species.⁶² Collectively, these studies indicate that chaotic dynamics can and do arise in complex microbial communities.

The bottom line

Our review of the presence or absence of chaos in populations is not intended to be exhaustive. For example, there are scattered accounts of tests for chaos in the dynamics of bobwhite quail⁶³ (no evidence), water fleas⁶⁴ (no evidence), aphids^{65,66} (no evidence), and moths (no evidence)^{65,66} and no doubt many more. Interestingly, in a recent review of chaos in real data sets,⁶⁷ several of the data sets (including blowflies¹ and flour moths⁶⁸) had dynamics 'on the edge of chaos'; that is, oscillations that, with a little more feedback, would have been chaotic. To this we can add a recent analysis of certain populations of Fennoscandian voles.⁶⁹ Whether this condition is common, or whether it is an artefact of the underlying statistical methodology, is currently unclear. However, it is now known⁷⁰ that noise superimposed on a regular periodic cycle can generate dynamics with no change in period but a change in amplitude—these quasi-periodic dynamics are just the sort of dynamics that give rise to zero Lyapunov exponents and dynamics at the 'edge of chaos'. So, perhaps some of these populations at the edge of chaos are simply ones that have an underlying tendency to show regular cycles in abundance, while being influenced by external noise.

Let us return to the original question we set ourselves. Ecologists have long realized that the systems that they are dealing with are non-linear, but are they sufficiently non-linear to drive chaos? In 1993, an excellent review of chaos in ecology was published using the subtitle 'is mother nature a strange attractor?'²⁰—one we have borrowed for one of our own section headings. The authors knowingly avoided answering their own question directly, preferring instead to suggest that 'chaos is quite likely, but much more work is needed to obtain a fuller answer to the question'. Now 15 years later, ecologists are expressing doubts. More recent opinions have varied from 'the jury is still out'⁶⁷ to 'chaos is rare'.⁷¹ In 1999, science journalist Carl Zimmer⁷² wrote about 'life after chaos'. Our own survey leads us to conclude that there is very little good evidence for chaos in natural populations. We have to be cautious, however, because part of the problem may be that ecological population data are by their nature relatively short term and noisy, making unequivocal proof of the existence of chaos challenging at best. Perhaps this is one reason

why ecologists have recently been more successful in detecting chaos in microbial systems which can be monitored for many more generations. The shortage of good data sets for multicellular organisms has occasionally led to controversy. For example, ecologists have sometimes ended up arguing over the same data sets: as Schaffer has quipped (no slur seems intended): ‘novel claims conjoined with a paucity of data inevitably attract the attention of statistics, much in the manner that offal attracts flies’.³⁷

We cannot rule out the possibility that mother nature is, in general, a strange attractor, but we have to say that the case is looking increasingly shaky, at least for multicellular organisms. If ecological systems are not chaotic then, given that it is a widespread property of many population models, we need to ask the reverse question posed early in the debate by Berryman and Millstein⁷³ in 1989—‘if not, why not?’

If not, why not?

Jeff Goldblum, playing that self-confident ‘chaotician’ in *Jurassic Park* who eventually met the end we could all see coming (‘When you gotta go, you gotta go’), remarked before his demise ‘Life will find a way’. Perhaps natural populations are not chaotic because natural selection somehow finds a way of pushing population parameters towards levels where they would not exhibit chaotic properties. Both Thomas and colleagues³⁸ and Berryman and Millstein⁷³ thought this might be the case, noting that in the chaotic region populations tend to fluctuate wildly yet spend a high proportion of their time at relatively low densities where extinctions are more likely to happen. Their argument was explicitly ‘group selectionist’: ‘it seems reasonable that natural selection might favor parameter values that minimize the likelihood of extinction and, consequently, chaotic dynamics’.⁷³ However, there are several problems with this argument. First, not all chaotic dynamics suffer from a high probability of extinction—some chaotic dynamics are tightly bound well away from zero, and chaos can in some cases reduce the likelihood of species extinctions.⁷⁴ So, despite the biblical impression that chaos is all about doom and destruction, it is not necessarily the case in the ecological sense of the word. Second, it ignores the problem of individual selection for cheats that favour their own reproductive success, even if it ultimately leads to the demise of the group. Chaos is about long-term dynamical behaviour, but natural selection is driven by what genetic variants perform best right now. There are cases where group selection effectively overpowers individual selection, but we generally need rather extreme assumptions³⁸ (see also Chapters 1 and 2).

Perhaps natural selection on individuals, rather than groups, can favour non-chaotic dynamics. The role of natural selection in influencing population behaviour, even in short-term laboratory experiments, is now widely recognized. For example, Yoshida and colleagues⁷⁵ recently successfully produced predator–prey cycles in a laboratory microcosm involving a rotifer feeding on a green alga. However, the cycle periods were far longer than predicted, and the observed predator and prey cycles were almost exactly out of phase, which is not what one would anticipate. Only by accounting for

(and testing for) the possibility of on-going natural selection in their system—in which rotifers effectively traded competitive ability for the ability to defend against predation when predation rates were high—were the authors able to reconcile their experimental and theoretical results. Thus, it seems the prey population was actively evolving at the same time it was undergoing fluctuations in density.

Alexander Nicholson's 'blowflies' represent one of the most celebrated and analysed data sets in the history of ecology^{34,76} (see also Fig. 6.1). Seeking to understand how and why populations fluctuate led him to begin an intensive series of experiments in the 1950s with caged Australian sheep-blowflies. Maintained in the laboratory for several hundred days, the blowflies exhibited characteristically 'double-peaked' oscillatory dynamics. However, in some of his longer-term experiments (lasting over 700 days), the dynamics became rather irregular after about 400 days, and at the same time the period of their oscillations also dramatically halved (to a mean of approximately 38 days). Nicholson himself recognized these patterns in his data and proposed that natural selection was acting in the course of his experiment. George Oster⁷⁷ went further, proposing that natural selection had a destabilizing influence, carrying population parameters into the chaotic regime (thereby neatly explaining Hassell and colleagues³³ earlier observations—see earlier). Yet a more detailed analysis has subsequently revealed the opposite⁷⁸—over the course of the experiment it appears as if there was a reduction in the maximum possible fecundity of adult females, moving the dynamics from unstable to more stable dynamics, tracking the regular addition of protein food supply. So, in the case of Nicholson's blowflies we may have evolution towards stable dynamics. However, we would do well to remember that selection pressures in a jar in a laboratory are likely to be quite different from those found in the wild.

Over the years, a general consensus has been building (with a few notable exceptions⁷⁹) that there may be selection on individuals that happens to take their populations away from the realm of chaos. For example, whether a population will evolve towards stability or towards chaos appears model-dependent, but a fluctuating population with constant carrying capacity (such as that represented via a logistic) will tend to experience selection that results in the population evolving towards population stability (in effect, parents go for offspring quality rather than quantity). Likewise, a suite of general population models,⁸⁰ models of competition,⁸¹ and those involving stage structure⁸² have all been reported to involve natural selection that indirectly promotes population stability. Experimental evolution in fruit fly populations had earlier suggested little or no evolution of parameters affecting population stability,⁸³ but a recent study under rather different conditions did indicate that populations evolve towards stability⁸⁴ (as a consequence of individuals reducing their fecundity to develop more rapidly). Once again, we see a close interrelationship between evolution and ecology, here with natural selection generating demographic parameters that happen to facilitate population stability.

Another set of reasons why populations may fail to exhibit chaos while mathematical models readily exhibit it may have something to do with the particular type of

mathematical models that have been explored. In a recent review, Scheuring⁸⁵ made the case that several biologically relevant details, such as sexual reproduction, population structure, and dispersal tend to be overlooked in simple population dynamical models, yet incorporation of these details into mathematical models generally favour dynamical stability. For example, certain population models that include sexual reproduction show a reduced propensity to exhibit chaos.⁸⁶ Similarly, when you allow small amounts of dispersal between several otherwise chaotically fluctuating populations, the resulting dynamic becomes more stable.^{87–89} Of course, it is very difficult to be general, but reversing the usual statement about chaos, it seems that complicated models with realistic features can generate simple dynamics.

Conclusion

The discovery that simple non-linear relationships, common in ecological systems, could generate extremely complicated dynamics was nothing short of a revelation. The associated finding that these complicated dynamics exhibited extreme sensitivity to initial conditions carries with it implications for all of ecology. In the intervening three decades since these discoveries were made, ecologists have worked hard to find evidence for this chaos in natural populations. Chaos has been formally defined, and methods have been developed to help test for it in the short and noisy data sets that ecologists are forced to deal with. We now know that populations can indeed be manipulated to generate all the features of chaos seen in mathematical models, and there is reasonable evidence of chaos arising in certain cases, such as childhood measles and some microbial systems. So, chaos *can* occur. Nevertheless, the majority of attempts to find chaos in natural populations have either drawn a blank or remain controversial.


Early in the ecological study of chaos, Schaffer and Kot⁵⁰ likened chaos to ‘the coals that Newcastle forgot’. With painful irony, their paper was published shortly after the UK national miners’ strike and all of the coal pits in the Newcastle area are now closed (as from 2007, only six pits remain in operation in the entire United Kingdom). Surveying the literature here leads one to suggest that many of the richest seams (to stretch the coal metaphor) of available ecological data have now also been explored, and few have provided much return. There may be good reasons why natural populations do not exhibit chaos, but only time will tell whether chaos is indeed rare.

Given the wonderful diversity of the natural world and knowledge that many systems have the *propensity* to exhibit chaos, perhaps a better question to have asked is ‘when and how often are natural systems chaotic?’ rather than ‘is this system chaotic?’. We have seen already that voles can exhibit very different dynamics in different populations, and both the blowfly and the lynx data are suggestive of a marked change in dynamics at some point in their history. Likewise, flour beetles can exhibit a range of different dynamics dependent on underlying experimental conditions. So, it is perhaps naive to characterize a population as ‘non-chaotic’ or ‘chaotic’, because dynamics can change according to the prevailing conditions. Human activities could yet turn non-chaotic dynamics into chaotic dynamics by increasing the degree of non-linearity

involved—as has been suggested for some fisheries.⁹⁰ This may particularly be the case for insect pests, such as the migratory locust mentioned in our opening paragraph, if drastic control measures are only implemented if pest density reaches a high level.⁷³

Chaos theory continues to grow and develop in a variety of scientific fields where it has found wide application. As May noted in one of his early seminal papers,⁹ ‘Not only in research, but also in the everyday world of politics and economics, we would all be better off if more people realized that simple non-linear systems do not necessarily possess simple dynamic properties’.

Ecologists are now much more aware of the subtle effects of non-linearities, and appreciate the wide variety of dynamical behaviours they can generate. Yet the truly surprising thing in all this is how long it took scientists to discover chaos. As James Yorke, one of the early pioneers, has recently said ‘I continue to wonder, if nearly all scientists missed this pervasive phenomenon, what other obvious phenomenon might we all be missing now?’⁹¹



Named after the mathematician Benoit Mandelbrot, the Mandelbrot set has become one of the icons of chaos theory. This rather complex object is generated by a relatively simple set of rules and has self-similarity at different scales (hence fractal dimensions), such that zooming in one sees the same patterns at an increasingly finer scale. Due to their aesthetic appeal, fractals have long attracted the interest of graphic designers and artists. Image by TNS.

This page intentionally left blank