CHAPTER 1
What Is Complexity?

Ideas thus made up of several simple ones put together, I call Complex; such as are Beauty, Gratitude, a Man, an Army, the Universe.

—John Locke, *An Essay Concerning Human Understanding*

> **Brazil: The Amazon rain forest**. Half a million army ants are on the march. No one is in charge of this army; it has no commander. Each individual ant is nearly blind and minimally intelligent, but the marching ants together create a coherent fan-shaped mass of movement that swarms over, kills, and efficiently devours all prey in its path. What cannot be devoured right away is carried with the swarm. After a day of raiding and destroying the edible life over a dense forest the size of a football field, the ants build their nighttime shelter—a chain-mail ball a yard across made up of the workers' linked bodies, sheltering the young larvae and mother queen at the center. When dawn arrives, the living ball melts away ant by ant as the colony members once again take their places for the day's march.

Nigel Franks, a biologist specializing in ant behavior, has written, "The solitary army ant is behaviorally one of the least sophisticated animals imaginable," and, "If 100 army ants are placed on a flat surface, they will walk around and around in never decreasing circles until they die of exhaustion." Yet put half a million of them together, and the group as a whole becomes what some have called a "superorganism" with "collective intelligence."

How does this come about? Although many things are known about ant colony behavior, scientists still do not fully understand all the mechanisms underlying a colony's collective intelligence. As Franks comments further, "I have studied *E. burchelli* [a common species of army ant] for many years, and for me the mysteries of its social organization still multiply faster than the rate at which its social structure can be explored."

The mysteries of army ants are a microcosm for the mysteries of many natural and social systems that we think of as "complex." No one knows exactly how any community of social organisms—ants, termites, humans—come together to collectively build the elaborate structures that increase the survival probability of the community as a whole. Similarly mysterious is how the intricate machinery of the immune system fights disease; how a group of cells organizes itself to be an eye or a brain; how independent members of an economy, each working chiefly for its own gain, produce complex but structured global markets; or, most mysteriously, how the phenomena we call "intelligence" and "consciousness" emerge from nonintelligent, nonconscious material substrates.

Such questions are the topics of *complex systems*, an interdisciplinary field of research that seeks to explain how large numbers of relatively simple entities organize themselves, without the benefit of any central controller, into a collective whole that creates patterns, uses information, and, in some cases, evolves and learns. The word *complex* comes from the Latin root *plectere*: to weave, entwine. In complex systems, many simple parts are irreducibly entwined, and the field of complexity is itself an entwining of many different fields.

Complex systems researchers assert that different complex systems in nature, such as insect colonies, immune systems, brains, and economies, have much in common. Let's look more closely.

Insect Colonies

Colonies of social insects provide some of the richest and most mysterious examples of complex systems in nature. An ant colony, for instance, can consist of hundreds to millions of individual ants, each one a rather simple creature that obeys its genetic imperatives to seek out food, respond in simple ways to the chemical signals of other ants in its colony, fight intruders, and so forth. However, as any casual observer of the outdoors can attest, the ants in a colony, each performing its own relatively simple actions, work together to build astoundingly complex structures that are clearly of great importance for the survival of the colony as a whole. Consider, for example, their use of soil, leaves, and twigs to construct huge nests of great strength and stability, with large networks of underground passages and dry, warm, brooding chambers whose temperatures are carefully controlled by decaying nest materials and the ants' own bodies. Consider also the long bridges certain species of ants build with their own bodies to allow emigration from one nest site to another via tree branches separated by great distances (to an ant, that is) (figure 1.1). Although much is now understood about ants and their social structures, scientists still can fully explain neither their individual nor group behavior: exactly how the individual actions of the ants produce large, complex structures, how the ants signal one another, and how the colony as a whole adapts to changing circumstances (e.g., changing weather or attacks on the colony). And how did biological evolution produce creatures with such an enormous contrast between their individual simplicity and their collective sophistication?

The Brain

The cognitive scientist Douglas Hofstadter, in his book *Gödel, Escher, Bach*, makes an extended analogy between ant colonies and brains, both being complex systems in which relatively simple components with only limited communication among themselves collectively give rise to complicated and sophisticated system-wide ("global") behavior. In the brain, the simple components are cells called *neurons*. The brain is made up of many different types of cells in addition to neurons, but most brain scientists believe that the actions of neurons and the patterns of connections among groups of neurons are what cause perception, thought, feelings, consciousness, and the other important large-scale brain activities.

FIGURE 1.1. Ants build a bridge with their bodies to allow the colony to take the shortest path across a gap. (Photograph courtesy of Carl Rettenmeyer.)

Neurons are pictured in figure 1.2 (top). Neurons consists of three main parts: the cell body (*soma*), the branches that transmit the cell's input from other neurons (*dendrites*), and the single trunk transmitting the cell's output to other neurons (*axon*). Very roughly, a neuron can be either in an active state (*firing*) or an inactive state (*not firing*). A neuron fires when it receives enough signals from other neurons through its dendrites. *Firing* consists of sending an electric pulse through the axon, which is then converted into a chemical signal via chemicals called *neurotransmitters*. This chemical signal in turn activates other neurons through their dendrites. The firing frequency and the resulting chemical output signals of a neuron can vary over time according to both its input and how much it has been firing recently.

These actions recall those of ants in a colony: individuals (neurons or ants) perceive signals from other individuals, and a sufficient summed strength of these signals causes the individuals to act in certain ways that produce additional signals. The overall effects can be very complex. We saw that an explanation of ants and their social structures is still incomplete; similarly, scientists don't yet understand how the actions of individual or dense networks of neurons give rise to the large-scale behavior of the brain (figure 1.2, bottom). They don't understand what the neuronal signals mean, how large numbers of neurons work together to produce global cognitive behavior, or how exactly they cause the brain to think thoughts and learn new things. And again, perhaps most puzzling is how such an elaborate signaling system with such powerful collective abilities ever arose through evolution.

The Immune System

The immune system is another example of a system in which relatively simple components collectively give rise to very complex behavior involving signaling and control, and in which adaptation occurs over time. A photograph illustrating the immune system's complexity is given in figure 1.3.

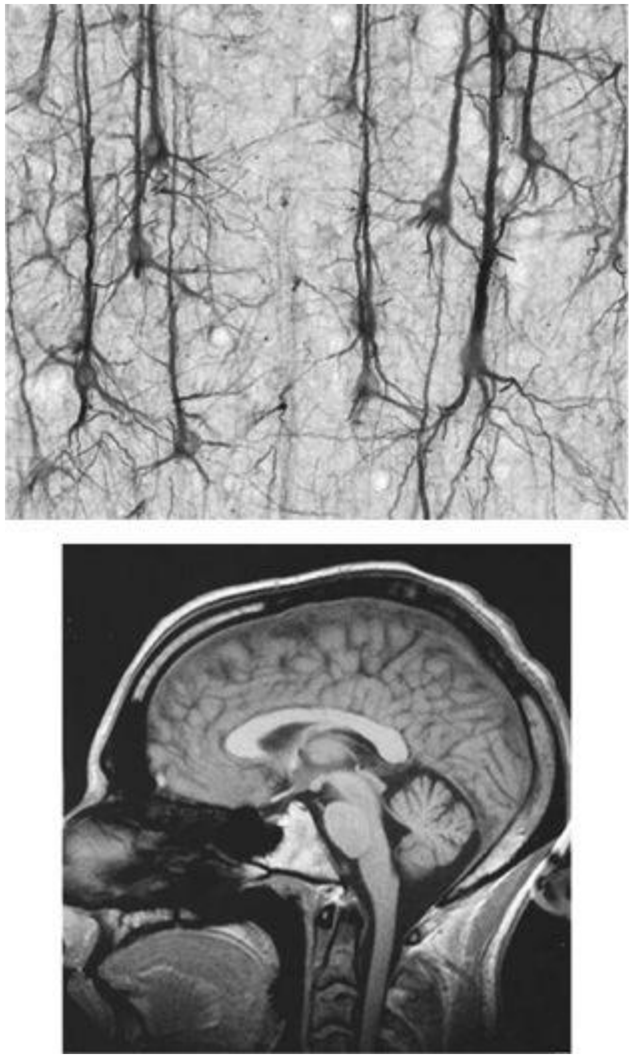

FIGURE 1.2. Top: microscopic view of neurons, visible via staining. Bottom: a human brain. How does the behavior at one level give rise to that of the next level? (Neuron photograph from brainmaps.org [http://brainmaps.org/smi32-pic.jpg], licensed under Creative Commons [http://creativecommons.org/licenses/by/3.0/]. Brain photograph courtesy of Christian R. Linder.)

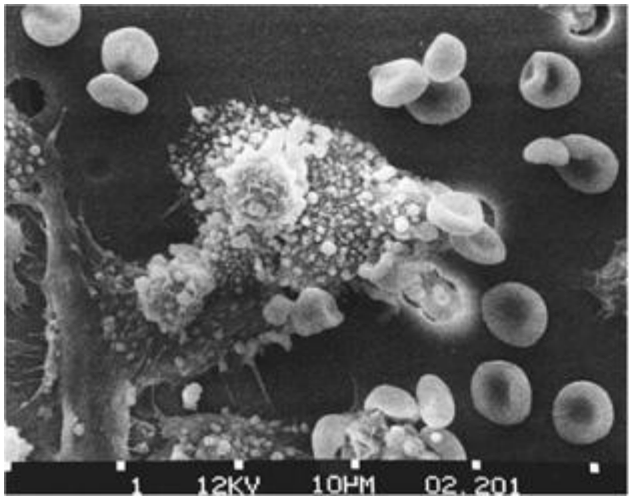

FIGURE 1.3. Immune system cells attacking a cancer cell. (Photograph by Susan Arnold, from National Cancer Institute Visuals Online [http://visualsonline.cancer.gov/details.cfm?imageid=2370].)

The immune system, like the brain, differs in sophistication in different animals, but the overall principles are the same across many species. The immune system consists of many different types of cells distributed over the entire body (in blood, bone marrow, lymph nodes, and other organs). This collection of cells works together in an effective and efficient way without any central control.

The star players of the immune system are white blood cells, otherwise known as *lymphocytes*. Each lymphocyte can recognize, via receptors on its cell body, molecules corresponding to certain possible invaders (e.g., bacteria). Some one trillion of these patrolling sentries circulate in the blood at a given time, each ready to sound the alarm if it is *activated*—that is, if its particular receptors encounter, by chance, a matching invader. When a lymphocyte is activated, it secretes large numbers of molecules—*antibodies*—that can identify similar invaders. These antibodies go out on a seek-and-destroy mission throughout the body. An activated lymphocyte also divides at an increased rate, creating daughter lymphocytes that will help hunt out invaders and secrete antibodies against them. It also creates daughter lymphocytes that will hang around and remember the particular invader that was seen, thus giving the body immunity to pathogens that have been previously encountered.

One class of lymphocytes are called *B cells* (the *B* indicates that they develop in the bone marrow) and have a remarkable property: the better the match between a B cell and an invader, the more antibody-secreting daughter cells the B cell creates. The daughter cells each differ slightly from the mother cell in random ways via mutations, and these daughter cells go on to create their own daughter cells in direct proportion to how well they match the invader. The result is a kind of Darwinian natural selection process, in which the match between B cells and invaders gradually gets better and better, until the antibodies being produced are extremely efficient at seeking and destroying the culprit microorganisms.

Many other types of cells participate in the orchestration of the immune response. *T cells* (which develop in the thymus) play a key role in regulating the response of B cells. *Macrophages* roam around looking for substances that have been tagged by antibodies, and they do the actual work of destroying the invaders. Other types of cells help effect longer-term immunity. Still other parts of the system guard against attacking the cells of one's own body.

Like that of the brain and ant colonies, the immune system's behavior arises from the independent actions of myriad simple players with no one actually in charge. The actions of the simple players—B cells, T cells, macrophages, and the like—can be viewed as a kind of chemical signal-processing network in which the recognition of an invader by one cell triggers a cascade of signals among cells that put into play the elaborate complex response. As yet many crucial aspects of this signal-processing system are not well understood. For example, it is still to be learned what, precisely, are the relevant signals, their specific functions, and how they work together to allow the system as a whole to "learn" what threats are present in the environment and to produce long-term immunity to those threats. We do not yet know precisely how the system avoids attacking the body; or what gives rise to flaws in the system, such as autoimmune diseases, in which the system does attack the body; or the detailed strategies of the human immunodeficiency virus (HIV), which is able to get by the defenses by attacking the immune system itself. Once again, a key question is how such an effective complex system arose in the first place in living creatures through biological evolution.

Economies

Economies are complex systems in which the “simple, microscopic” components consist of people (or companies) buying and selling goods, and the collective behavior is the complex, hard-to-predict behavior of markets as a whole, such as changes in the price of housing in different areas of the country or fluctuations in stock prices (figure 1.4). Economies are thought by some economists to be adaptive on both the microscopic and macroscopic level. At the microscopic level, individuals, companies, and markets try to increase their profitability by learning about the behavior of other individuals and companies. This microscopic self-interest has historically been thought to push markets as a whole—on the macroscopic level—toward an equilibrium state in which the prices of goods are set so there is no way to change production or consumption patterns to make everyone better off. In terms of profitability or consumer satisfaction, if someone is made better off, someone else will be made worse off. The process by which markets obtain this equilibrium is called *market efficiency*. The eighteenth-century economist Adam Smith called this self-organizing behavior of markets the “invisible hand”: it arises from the myriad microscopic actions of individual buyers and sellers.

Economists are interested in how markets become efficient, and conversely, what makes efficiency fail, as it does in real-world markets. More recently, economists involved in the field of complex systems have tried to explain market behavior in terms similar to those used previously in the descriptions of other complex systems: dynamic hard-to-predict patterns in global behavior, such as patterns of market bubbles and crashes; processing of signals and information, such as the decision-making processes of individual buyers and sellers, and the resulting “information processing” ability of the market as a whole to “calculate” efficient prices; and adaptation and learning, such as individual sellers adjusting their production to adapt to changes in buyers’ needs, and the market as a whole adjusting global prices.

The World Wide Web

The World Wide Web came on the world scene in the early 1990s and has experienced exponential growth ever since. Like the systems described above, the Web can be thought of as a self-organizing social system: individuals, with little or no central oversight, perform simple tasks: posting Web pages and linking to other Web pages. However, complex systems scientists have discovered that the network as a whole has many unexpected large-scale properties involving its overall structure, the way in which it grows, how information propagates over its links, and the coevolutionary relationships between the behavior of search engines and the Web's link structure, all of which lead to what could be called "adaptive" behavior for the system as a whole. The complex behavior emerging from simple rules in the World Wide Web is currently a hot area of study in complex systems. Figure 1.5 illustrates the structure of one collection of Web pages and their links. It seems that much of the Web looks very similar; the question is, *why*?

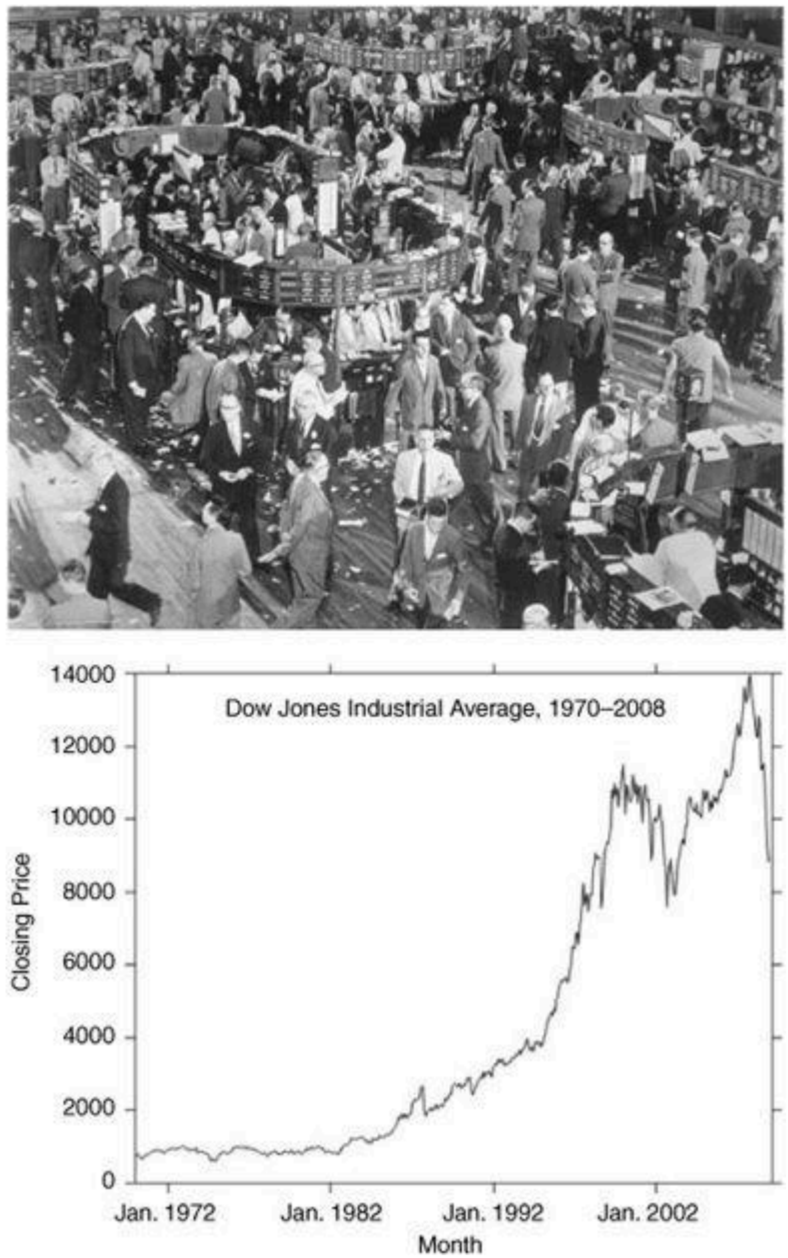

FIGURE 1.4. Individual actions on a trading floor give rise to the hard-to-predict large-scale behavior of financial markets. Top: New York Stock Exchange (photograph from Milstein Division of US History, Local History and Genealogy, The New York Public Library, Astor, Lenox, and Tilden Foundations, used by permission). Bottom: Dow Jones Industrial Average closing price, plotted monthly 1970–2008.

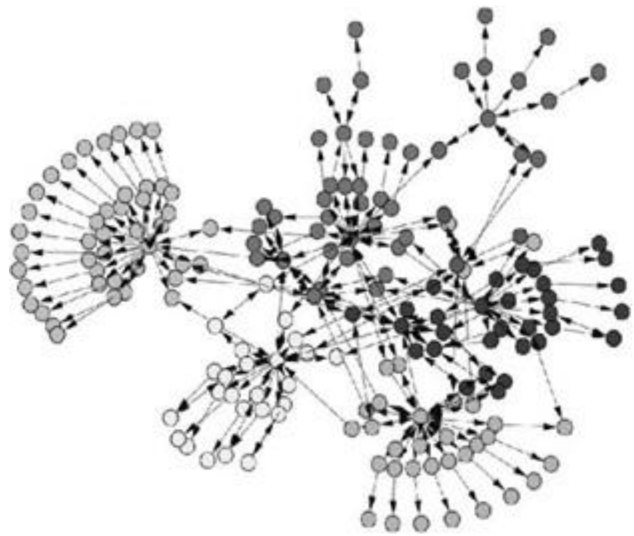

FIGURE 1.5. Network structure of a section of the World Wide Web. (Reprinted with permission from M.E.J. Newman and M. Girvin, *Physical Review Letters E*, 69,026113, 2004. Copyright 2004 by the American Physical Society.)

Common Properties of Complex Systems

When looked at in detail, these various systems are quite different, but viewed at an abstract level they have some intriguing properties in common:

1. **Complex collective behavior**: All the systems I described above consist of large networks of individual components (ants, B cells, neurons, stock-buyers, Website creators), each typically following relatively simple rules with no central control or leader. It is the collective actions of vast numbers of components that give rise to the complex, hard-to-predict, and changing patterns of behavior that fascinate us.

2. **Signaling and information processing**: All these systems produce and use information and signals from both their internal and external environments.
3. **Adaptation**: All these systems adapt—that is, change their behavior to improve their chances of survival or success—through learning or evolutionary processes.

Now I can propose a definition of the term *complex system*: **a system in which large networks of components with no central control and simple rules of operation give rise to complex collective behavior, sophisticated information processing, and adaptation via learning or evolution.** (Sometimes a differentiation is made between *complex adaptive systems*, in which adaptation plays a large role, and nonadaptive complex systems, such as a hurricane or a turbulent rushing river. In this book, as most of the systems I do discuss are adaptive, I do not make this distinction.)

Systems in which organized behavior arises without an internal or external controller or leader are sometimes called *self-organizing*. Since simple rules produce complex behavior in hard-to-predict ways, the macroscopic behavior of such systems is sometimes called *emergent*. Here is an alternative definition of a *complex system*: **a system that exhibits nontrivial emergent and self-organizing behaviors**. The central question of the sciences of complexity is how this emergent self-organized behavior comes about. In this book I try to make sense of these hard-to-pin-down notions in different contexts.

How Can Complexity Be Measured?

In the paragraphs above I have sketched some qualitative common properties of complex systems. But more quantitative questions remain: Just how *complex* is a particular complex system? That is, how do we measure *complexity*? Is there any way to say precisely how much more complex one system is than another?

These are key questions, but they have not yet been answered to anyone's satisfaction and remain the source of many scientific arguments in the field. As I describe in chapter 7, many different measures of complexity have been proposed; however, none has been universally accepted by scientists. Several of these measures and their usefulness are described in various chapters of this book.

But how can there be a science of complexity when there is no agreed-on quantitative definition of complexity?

I have two answers to this question. First, neither a single *science of complexity* nor a single *complexity theory* exists yet, in spite of the many articles and books that have used these terms. Second, as I describe in many parts of this book, an essential feature of forming a new science is a struggle to define its central terms. Examples can be seen in the struggles to define such core concepts as *information*, *computation*, *order*, and *life*. In this book I detail these struggles, both historical and current, and tie them in with our struggles to understand the many facets of complexity. This book is about cutting-edge science, but it is also about the history of core concepts underlying this cutting-edge science. The next four chapters provide this history and background on the concepts that are used throughout the book.

CHAPTER 2
Dynamics, Chaos, and Prediction

It makes me so happy. To be at the beginning again, knowing almost nothing…. The ordinary-sized stuff which is our lives, the things people write poetry about—clouds—daffodils—waterfalls…. these things are full of mystery, as mysterious to us as the heavens were to the Greeks…It's the best possible time to be alive, when almost everything you thought you knew is wrong.

—Tom Stoppard, *Arcadia*

DYNAMICAL SYSTEMS THEORY (or *dynamics*) concerns the description and prediction of systems that exhibit complex *changing* behavior at the macroscopic level, emerging from the collective actions of many interacting components. The word *dynamic* means changing, and dynamical systems are systems that change over time in some way. Some examples of dynamical systems are

The solar system (the planets change position over time)

The heart of a living creature (it beats in a periodic fashion rather than standing still)

The brain of a living creature (neurons are continually firing, neurotransmitters are propelled from one neuron to another, synapse strengths are changing, and generally the whole system is in a continual state of flux)

The stock market

The world's population

The global climate

Dynamical systems include these and most other systems that you probably can think of. Even rocks change over geological time. Dynamical systems theory describes in general terms the ways in which systems can change, what types of macroscopic behavior are possible, and what kinds of predictions about that behavior can be made.

Dynamical systems theory has recently been in vogue in popular science because of the fascinating results coming from one of its intellectual offspring, the study of chaos. However, it has a long history, starting, as many sciences did, with the Greek philosopher Aristotle.

Early Roots of Dynamical Systems Theory

Aristotle was the author of one of the earliest recorded theories of motion, one that was accepted widely for over 1,500 years. His theory rested on two main principles, both of which turned out to be wrong. First, he believed that motion on Earth differs from motion in the heavens. He asserted that on Earth objects move in straight lines and only when something forces them to; when no forces are applied, an object comes to its natural resting state. In the heavens, however, planets and other celestial objects move continuously in perfect circles centered about the Earth. Second, Aristotle believed that earthly objects move in different ways depending on what they are made of. For example, he believed that a rock will fall to Earth because it is mainly composed of the element *earth*, whereas smoke will rise because it is mostly composed of the element *air*. Likewise, heavier objects, presumably containing more earth, will fall faster than lighter objects.

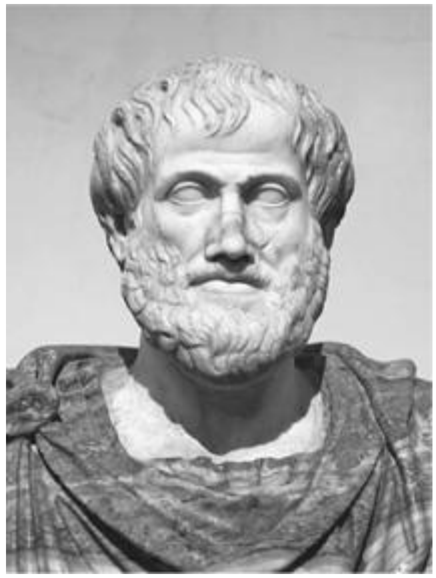

Aristotle, 384–322 B.C.
(Ludovisi Collection)

Clearly Aristotle (like many theorists since) was not one to let experimental results get in the way of his theorizing. His scientific method was to let logic and common sense direct theory; the importance of testing the resulting theories by experiments is a more modern notion. The influence of Aristotle's ideas was strong and continued to hold sway over most of Western science until the sixteenth century—the time of Galileo.

Galileo was a pioneer of experimental, empirical science, along with his predecessor Copernicus and his contemporary Kepler. Copernicus established that the motion of the planets is centered not about the Earth but about the sun. (Galileo got into big trouble with the Catholic Church for promoting this view and was eventually forced to publicly renounce it; only in 1992 did the Church officially admit that Galileo had been unfairly persecuted.) In the early 1600s, Kepler discovered that the motion of the planets is not circular but rather elliptical, and he discovered laws describing this elliptical motion.

Whereas Copernicus and Kepler focused their research on celestial motion, Galileo studied motion not only in the heavens but also here on Earth by experimenting with the objects one now finds in elementary physics courses: pendula, balls rolling down inclined planes, falling objects, light reflected by mirrors. Galileo did not have the sophisticated experimental devices we have today: he is said to have timed the swinging of a pendulum by counting his heartbeats and to have measured the effects of gravity by dropping objects off the leaning tower of Pisa. These now-classic experiments revolutionized ideas about motion. In particular, Galileo's studies directly contradicted Aristotle's long-held principles of motion. Against common sense, rest is *not* the natural state of objects; rather it takes *force* to stop a moving object. Heavy and light objects in a vacuum fall at the same rate. And perhaps most revolutionary of all, laws of motion on the Earth could explain some aspects of motions in the heavens. With Galileo, the scientific revolution, with experimental observations at its core, was definitively launched.

The most important person in the history of dynamics was Isaac Newton. Newton, who was born the year after Galileo died, can be said to have invented, on his own, the science of dynamics. Along the way he also had to invent calculus, the branch of mathematics that describes motion and change.

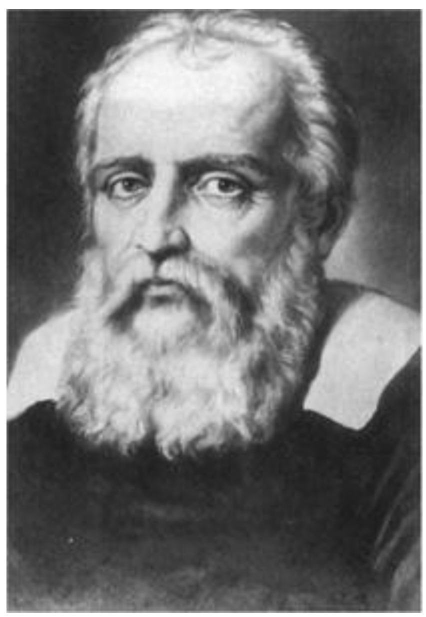

Galileo, 1564–1642 (AIP Emilio Segre Visual Archives, E. Scott Barr Collection)

Isaac Newton, 1643–1727 (Original engraving by unknown artist, courtesy AIP Emilio Segre Visual Archives)

Physicists call the general study of motion *mechanics*. This is a historical term dating from ancient Greece, reflecting the classical view that all motion could be explained in terms of the combined actions of simple "machines" (e.g., lever, pulley, wheel and axle). Newton's work is known today as *classical mechanics*. Mechanics is divided into two areas: kinematics, which describes how things move, and dynamics, which explains why things obey the laws of kinematics. For example, Kepler's laws are kinematic laws—they describe *how* the planets move (in ellipses with the sun at one focus)—but not *why* they move in this particular way. Newton's laws are the foundations of dynamics: they explain the motion of the planets, and everything else, in terms of the basic notions of force and mass.

Newton's famous three laws are as follows:

1. Constant motion: Any object not subject to a force moves with unchanging speed.
2. Inertial mass: When an object is subject to a force, the resulting change in its motion is inversely proportional to its mass.
3. Equal and opposite forces: If object A exerts a force on object B, then object B must exert an equal and opposite force on object A.

One of Newton's greatest accomplishments was to realize that these laws applied not just to earthly objects but to those in the heavens as well. Galileo was the first to state the constant-motion law, but he believed it applied only to objects on Earth. Newton, however, understood that this law should apply to the planets as well, and realized that elliptical orbits, which exhibit a constantly *changing* direction of motion, require explanation in terms of a force, namely gravity. Newton's other major achievement was to state a universal law of gravity: the force of gravity between two objects is proportional to the product of their masses divided by the square of the distance between them. Newton's insight—now the backbone of modern science—was that this law applies everywhere in the universe, to falling apples as well as to planets. As he wrote: "nature is exceedingly simple and conformable to herself. Whatever reasoning holds for greater motions, should hold for lesser ones as well."

Newtonian mechanics produced a picture of a "clockwork universe," one that is wound up with the three laws and then runs its mechanical course. The mathematician Pierre Simon Laplace saw the implication of this clockwork view for prediction: in 1814 he asserted that, given Newton's laws and the current position and velocity of every particle in the universe, it was possible, in principle, to predict everything for all time. With the invention of electronic computers in the 1940s, the "in principle" might have seemed closer to "in practice."

Revised Views of Prediction

However, two major discoveries of the twentieth century showed that Laplace's dream of complete prediction is not possible, even in principle. One discovery was Werner Heisenberg's 1927 "uncertainty principle" in quantum mechanics, which states that one cannot measure the exact values of the position and the momentum (mass times velocity) of a particle at the same time. The more certain one is about where a particle is located at a given time, the less one can know about its momentum, and vice versa. However, effects of Heisenberg's principle exist only in the quantum world of tiny particles, and most people viewed it as an interesting curiosity, but not one that would have much implication for prediction at a larger scale—predicting the weather, say.

It was the understanding of *chaos* that eventually laid to rest the hope of perfect prediction of all complex systems, quantum or otherwise. The defining idea of chaos is that there are some systems—*chaotic* systems—in which even minuscule uncertainties in measurements of initial position and momentum can result in huge errors in long-term predictions of these quantities. This is known as "sensitive dependence on initial conditions."

In parts of the natural world such small uncertainties will not matter. If your initial measurements are fairly but not perfectly precise, your predictions will likewise be close to right if not exactly on target. For example, astronomers can predict eclipses almost perfectly in spite of even relatively large uncertainties in measuring the positions of planets. But sensitive dependence on initial conditions says that in chaotic systems, even the tiniest errors in your initial measurements will eventually produce huge errors in your prediction of the future motion of an object. In such systems (and hurricanes may well be an example) *any* error, no matter how small, will make long-term predictions vastly inaccurate.

This kind of behavior is counterintuitive; in fact, for a long time many scientists denied it was possible. However, chaos in this sense has been observed in cardiac disorders, turbulence in fluids, electronic circuits, dripping faucets, and many other seemingly unrelated phenomena. These days, the existence of chaotic systems is an accepted fact of science.

It is hard to pin down who first realized that such systems might exist. The possibility of sensitive dependence on initial conditions was proposed by a number of people long before quantum mechanics was invented. For example, the physicist James Clerk Maxwell hypothesized in 1873 that there are classes of phenomena affected by "influences whose physical magnitude is too small to be taken account of by a finite being, [but which] may produce results of the highest importance."

Possibly the first clear example of a chaotic system was given in the late nineteenth century by the French mathematician Henri Poincaré. Poincaré was the founder of and probably the most influential contributor to the modern field of dynamical systems theory, which is a major outgrowth of Newton's science of dynamics. Poincaré discovered sensitive dependence on initial conditions when attempting to solve a much simpler problem than predicting the motion of a hurricane. He more modestly tried to tackle the so-called three-body problem: to determine, using Newton's laws, the long-term motions of three masses exerting gravitational forces on one another. Newton solved the *two*-body problem, but the three-body problem turned out to be much harder. Poincaré tackled it in 1887 as part of a mathematics contest held in honor of the king of Sweden. The contest offered a prize of 2,500 Swedish crowns for a solution to the "many body" problem: predicting the future positions of arbitrarily many masses attracting one another under Newton's laws. This problem was inspired by the question of whether or not the solar system is stable: will the planets remain in their current orbits, or will they wander from them? Poincaré started off by seeing whether he could solve it for merely three bodies.

He did not completely succeed—the problem was too hard. But his attempt was so impressive that he was awarded the prize anyway. Like Newton with calculus, Poincaré had to invent a new branch of mathematics, *algebraic topology*, to even tackle the problem. Topology is an extended form of geometry, and it was in looking at the geometric consequences of the three-body problem that he discovered the possibility of sensitive dependence on initial conditions. He summed up his discovery as follows:

> If we knew exactly the laws of nature and the situation of the universe at the initial moment, we could predict exactly the situation of that same universe at a succeeding moment. But 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. If that enabled us to predict the succeeding situation with the same approximation, that is all we require, and we should say that the phenomenon has been predicted, that it is governed by laws. But it is not always so; it may happen that small differences in the initial conditions produce very great ones in the final phenomenon. A small error in the former will produce an enormous error in the latter. Prediction becomes impossible....

In other words, even if we know the laws of motion perfectly, two different sets of initial conditions (here, initial positions, masses, and velocities for objects), even if they differ in a minuscule way, can sometimes produce greatly different results in the subsequent motion of the system. Poincaré found an example of this in the three-body problem.

Henri Poincaré, 1854–1912 (AIP Emilio Segre Visual Archives)

It was not until the invention of the electronic computer that the scientific world began to see this phenomenon as significant. Poincaré, way ahead of his time, had guessed that sensitive dependence on initial conditions would stymie attempts at long-term weather prediction. His early hunch gained some evidence when, in 1963, the meteorologist Edward Lorenz found that even simple computer models of weather phenomena were subject to sensitive dependence on initial conditions. Even with today's modern, highly complex meteorological computer models, weather predictions are at best reasonably accurate only to about one week in the future. It is not yet known whether this limit is due to fundamental chaos in the weather, or how much this limit can be extended by collecting more data and building even better models.

Linear versus Nonlinear Rabbits

Let's now look more closely at sensitive dependence on initial conditions. How, precisely, does the huge magnification of initial uncertainties come about in chaotic systems? The key property is *nonlinearity*. A linear system is one you can understand by understanding its parts individually and then putting them together. When my two sons and I cook together, they like to take turns adding ingredients. Jake puts in two cups of flour. Then Nicky puts in a cup of sugar. The result? Three cups of flour/sugar mix. The whole is equal to the sum of the parts.

A nonlinear system is one in which the whole is different from the sum of the parts. Jake puts in two cups of baking soda. Nicky puts in a cup of vinegar. The whole thing explodes. (You can try this at home.) The result? *More* than three cups of vinegar-and-baking-soda-and-carbon-dioxide fizz.

The difference between the two examples is that in the first, the flour and sugar don't really interact to create something new, whereas in the second, the vinegar and baking soda interact (rather violently) to create a lot of carbon dioxide.

Linearity is a reductionist's dream, and nonlinearity can sometimes be a reductionist's nightmare. Understanding the distinction between linearity and nonlinearity is very important and worthwhile. To get a better handle on this distinction, as well as on the phenomenon of chaos, let's do a bit of very simple mathematical exploration, using a classic illustration of linear and nonlinear systems from the field of biological population dynamics.

Suppose you have a population of breeding rabbits in which every year all the rabbits pair up to mate, and each pair of rabbit parents has exactly four offspring and then dies. The population growth, starting from two rabbits, is illustrated in figure 2.1.

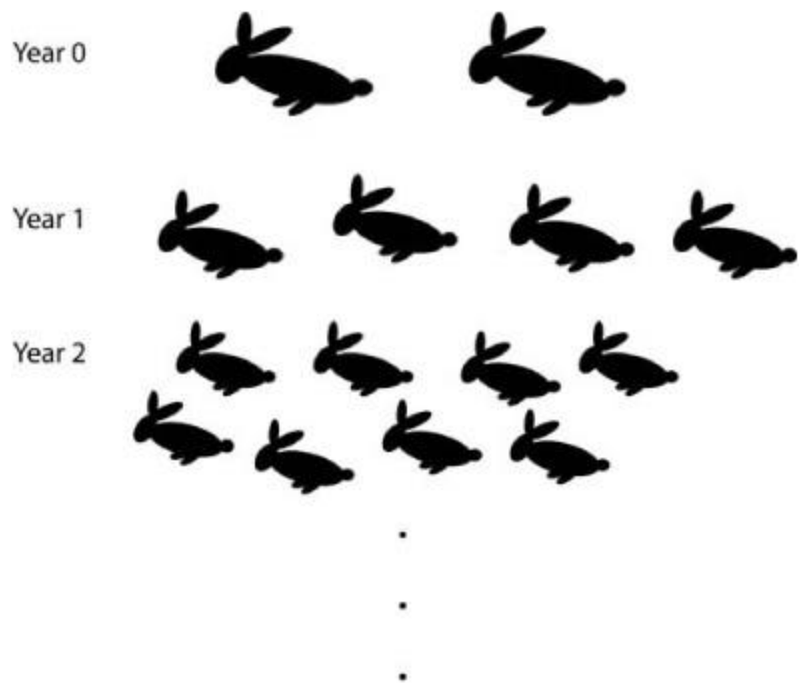

FIGURE 2.1. Rabbits with doubling population.

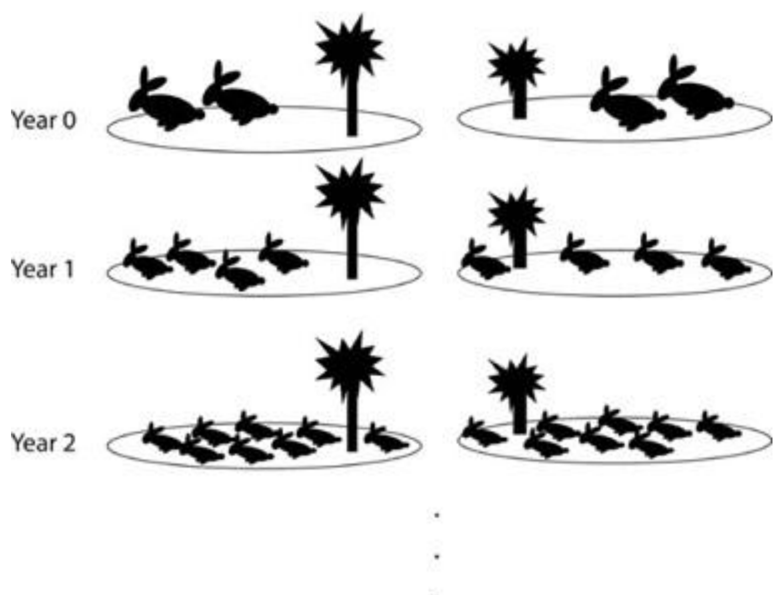

FIGURE 2.2. Rabbits with doubling population, split on two islands.

It is easy to see that the population doubles every year without limit (which means the rabbits would quickly take over the planet, solar system, and universe, but we won't worry about that for now).

This is a linear system: the whole is equal to the sum of the parts. What do I mean by this? Let's take a population of four rabbits and split them between two separate islands, two rabbits on each island. Then let the rabbits proceed with their reproduction. The population growth over two years is illustrated in figure 2.2.

Each of the two populations doubles each year. At each year, if you add the populations on the two islands together, you'll get the same number of rabbits that you would have gotten had there been no separation—that is, had they all lived on one island.

If you make a plot with the current year's population size on the horizontal axis and the next-year's population size on the vertical axis, you get a straight line (figure 2.3). This is where the term *linear system* comes from.

But what happens when, more realistically, we consider limits to population growth? This requires us to make the growth rule nonlinear. Suppose that, as before, each year every pair of rabbits has four offspring and then dies. But now suppose that some of the offspring die before they reproduce because of overcrowding. Population biologists sometimes use an equation called the *logistic model* as a description of population growth in the presence of overcrowding. This sense of the word *model* means a mathematical formula that describes population growth in a simplified way.

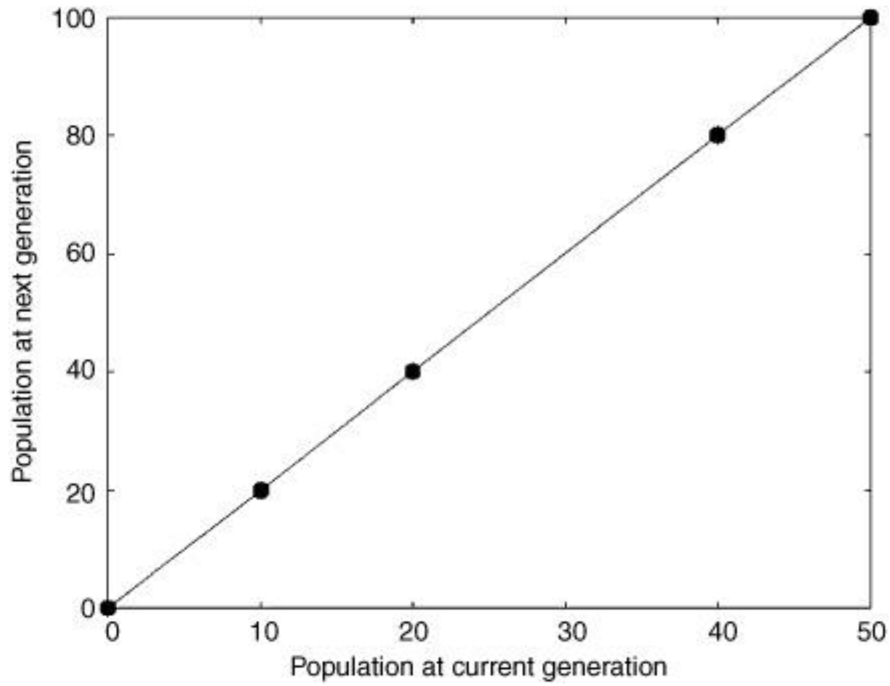

FIGURE 2.3. A plot of how the population size next year depends on the population size this year for the linear model.

In order to use the logistic model to calculate the size of the next generation's population, you need to input to the logistic model the current generation's population size, the *birth rate*, the *death rate* (the probability of an individual will die due to overcrowding), and the maximum *carrying capacity* (the strict upper limit of the population that the habitat will support.)

I won't give the actual equation for the logistic model here (it is given in the notes), but you can see its behavior in figure 2.4.

As a simple example, let's set *birth rate* = 2 and *death rate* = 0.4, assume the carrying capacity is thirty-two, and start with a population of twenty rabbits in the first generation. Using the logistic model, I calculate that the number of surviving offspring in the second generation is twelve. I then plug this new population size into the model, and find that there are still exactly twelve surviving rabbits in the third generation. The population will stay at twelve for all subsequent years.

If I reduce the death rate to 0.1 (keeping everything else the same), things get a little more interesting. From the model I calculate that the second generation has 14.25 rabbits and the third generation has 15.01816.

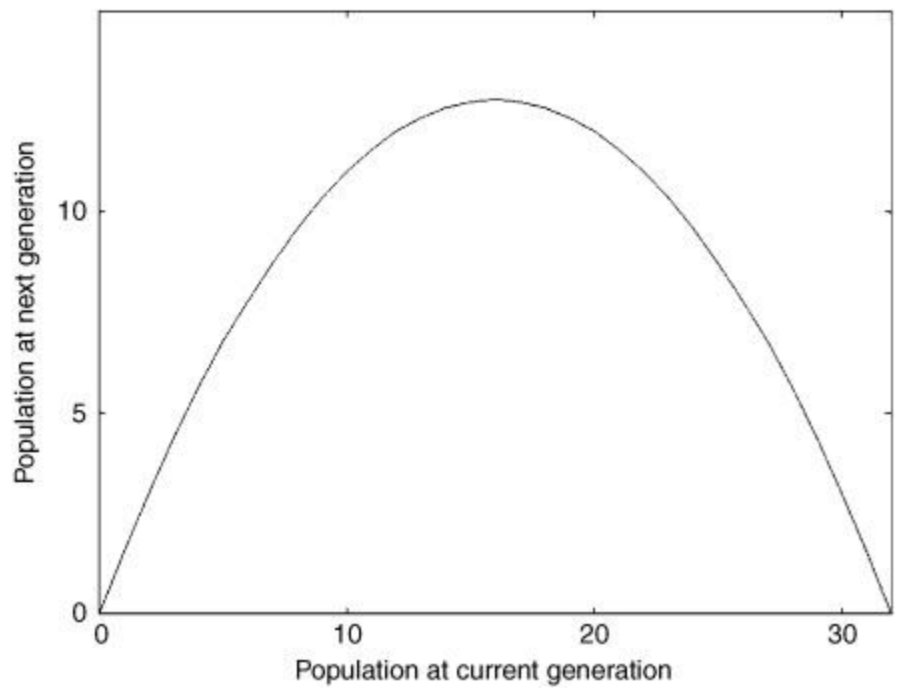

FIGURE 2.4. A plot of how the population size next year depends on the population size this year under the logistic model, with birth rate equal to 2, death rate equal to 0.4, and carrying capacity equal to 32. The plot will also be a parabola for other values of these parameters.

Wait a minute! How can we have 0.25 of a rabbit, much less 0.01816 of a rabbit? Obviously in real life we cannot, but this is a mathematical model, and it allows for fractional rabbits. This makes it easier to do the math, and can still give reasonable predictions of the actual rabbit population. So let’s not worry about that for now.

This process of calculating the size of the next population again and again, starting each time with the immediately previous population, is called “iterating the model.”

What happens if the death rate is set back to 0.4 and carrying capacity is doubled to sixty-four? The model tells me that, starting with twenty rabbits, by year nine the population reaches a value close to twenty-four and stays there.

You probably noticed from these examples that the behavior is more complicated than when we simply doubled the population each year. That's because the logistic model is nonlinear, due to its inclusion of death by overcrowding. Its plot is a parabola instead of a line (figure 2.4). The logistic population growth is not simply equal to the sum of its parts. To show this, let's see what happens if we take a population of twenty rabbits and segregate it into populations of ten rabbits each, and iterate the model for each population (with *birth rate* = 2 and *death rate* = .4, as in the first example above). The result is illustrated in figure 2.5.

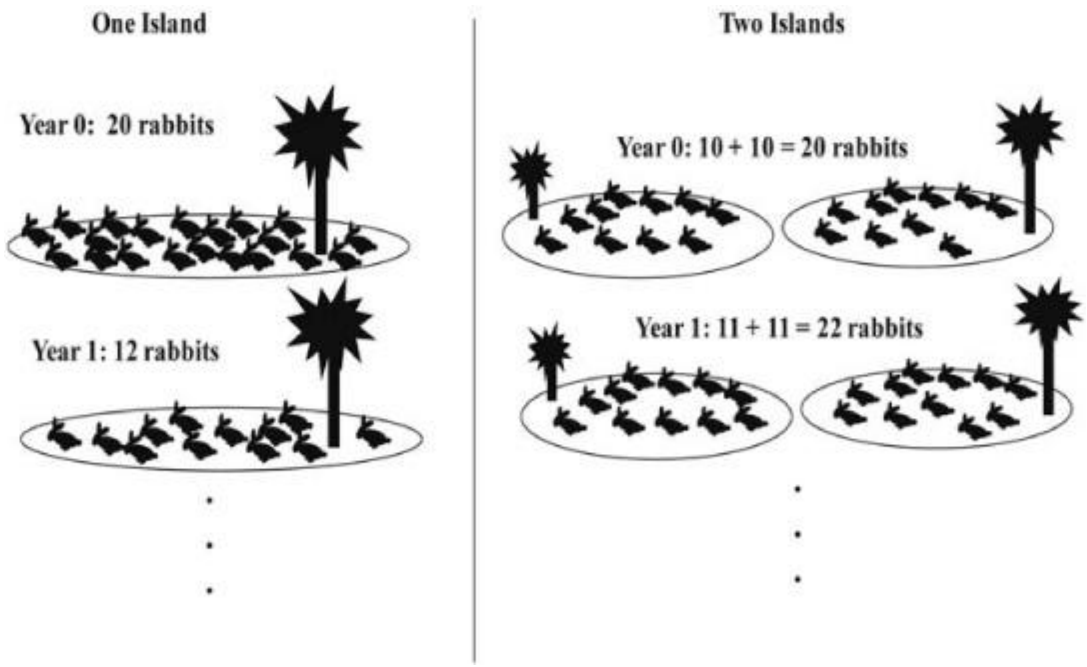

FIGURE 2.5. Rabbit population split on two islands, following the logistic model.

At year one, the original twenty-rabbit population has been cut down to twelve rabbits, but each of the original ten-rabbit populations now has eleven rabbits, for a total of twenty-two rabbits. The behavior of the whole is clearly not equal to the sum of the behavior of the parts.

The Logistic Map

Many scientists and mathematicians who study this sort of thing have used a simpler form of the logistic model called the *logistic map*, which is perhaps the most famous equation in the science of dynamical systems and chaos. The logistic model is simplified by combining the effects of birth rate and death rate into one number, called *R*. *Population size* is replaced by a related concept called "fraction of carrying capacity," called *x*. Given this simplified model, scientists and mathematicians promptly forget all about population growth, carrying capacity, and anything else connected to the real world, and simply get lost in the astounding behavior of the equation itself. We will do the same.

Here is the equation, where x_t is the current value of x and x_{t+1} is its value at the next time step:[1]

$$x_{t+1} = R\, x_t\, (1 - x_t).$$

I give the equation for the logistic map to show you how simple it is. In fact, it is one of the simplest systems to capture the essence of chaos: sensitive dependence on initial conditions. The logistic map was brought to the attention of population biologists in a 1971 article by the mathematical biologist Robert May in the prestigious journal *Nature*. It had been previously analyzed in detail by several mathematicians, including Stanislaw Ulam, John von Neumann, Nicholas Metropolis, Paul Stein, and Myron Stein. But it really achieved fame in the 1980s when the physicist Mitchell Feigenbaum used it to demonstrate *universal* properties common to a very large class of chaotic systems. Because of its apparent simplicity and rich history, it is a perfect vehicle to introduce some of the major concepts of dynamical systems theory and chaos.

The logistic map gets very interesting as we vary the value of R. Let's start with $R = 2$. We need to also start out with some value between 0 and 1 for x_0, say 0.5. If you plug those numbers into the logistic map, the answer for x_1 is 0.5. Likewise, $x_2 = 0.5$, and so on. Thus, if $R = 2$ and the population starts out at half the maximum size, it will stay there forever.

Now let's try $x_0 = 0.2$. You can use your calculator to compute this one. (I'm using one that reads off at most seven decimal places.) The results are more interesting:

$$\begin{aligned}
x_0 &= 0.2 \\
x_1 &= 0.32 \\
x_2 &= 0.4352 \\
x_3 &= 0.4916019 \\
x_4 &= 0.4998589 \\
x_5 &= 0.5 \\
x_6 &= 0.5 \\
&\vdots
\end{aligned}$$

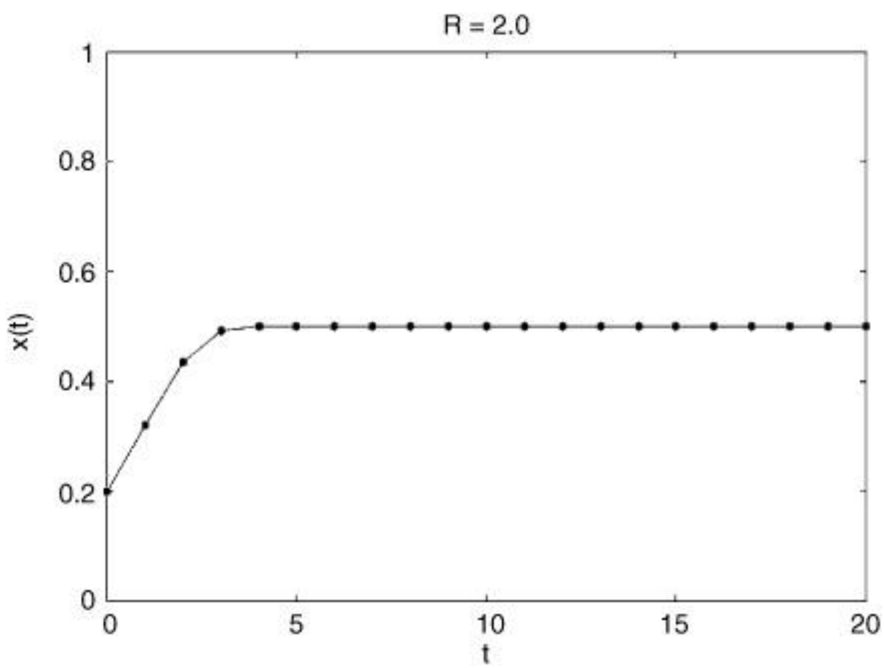

FIGURE 2.6. Behavior of the logistic map for $R = 2$ and $x_0 = 0.2$.

The same eventual result ($x_t = 0.5$ forever) occurs but here it takes five iterations to get there.

It helps to see these results visually. A plot of the value of x_t at each time t for 20 time steps is shown in figure 2.6. I've connected the points by lines to better show how as time increases, x quickly converges to 0.5.

What happens if x_0 is large, say, 0.99? Figure 2.7 shows a plot of the results.

Again the same ultimate result occurs, but with a longer and more dramatic path to get there.

You may have guessed it already: if $R = 2$ then x_t eventually always gets to 0.5 and stays there. The value 0.5 is called a *fixed point*: how long it takes to get there depends on where you start, but once you are there, you are fixed.

If you like, you can do a similar set of calculations for $R = 2.5$, and you will find that the system also always goes to a fixed point, but this time the fixed point is 0.6.

For even more fun, let $R = 3.1$. The behavior of the logistic map now gets more complicated. Let $x_0 = 0.2$. The plot is shown in figure 2.8.

In this case x never settles down to a fixed point; instead it eventually settles into an oscillation between two values, which happen to be 0.5580141 and 0.7645665. If the former is plugged into the formula the latter is produced, and vice versa, so this oscillation will continue forever. This oscillation will be reached eventually no matter what value is given for x_0. This kind of regular final behavior (either fixed point or oscillation) is called an "attractor," since, loosely speaking, any initial condition will eventually be "attracted to it."

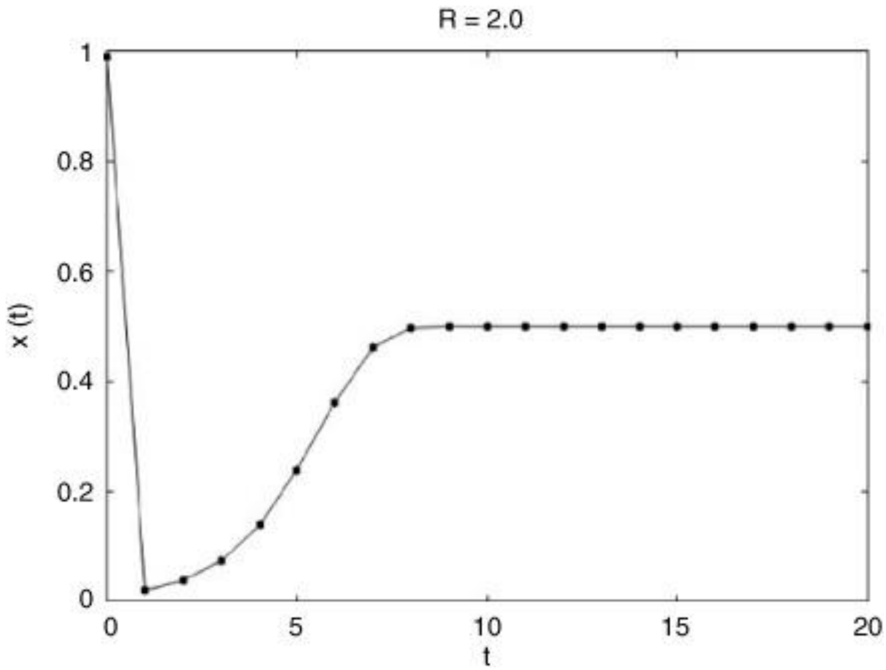

FIGURE 2.7. Behavior of the logistic map for $R = 2$ and $x_0 = 0.99$.

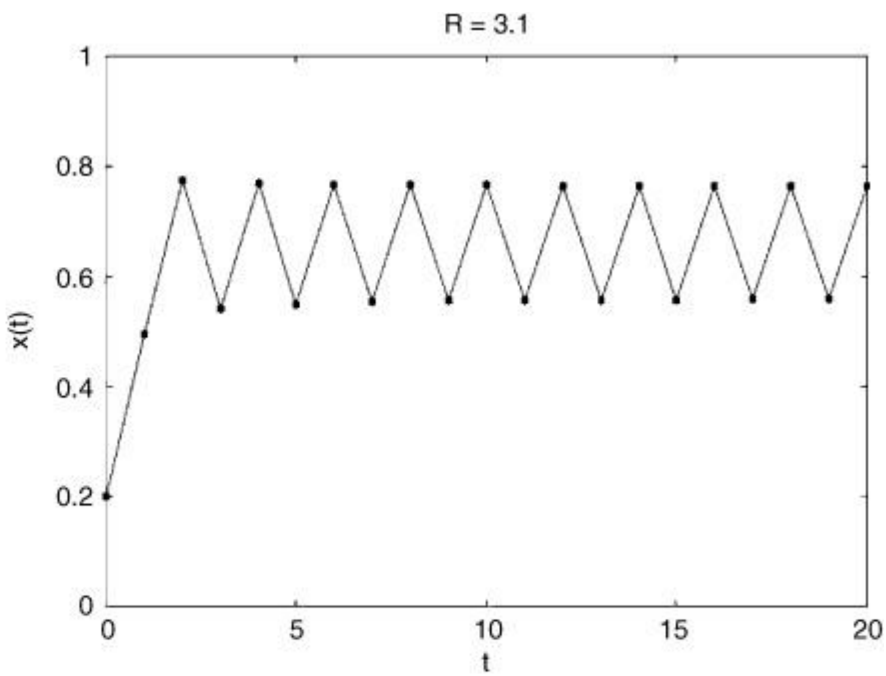

FIGURE 2.8. Behavior of the logistic map for $R = 3.1$ and $x_0 = 0.2$.

For values of R up to around 3.4 the logistic map will have similar behavior: after a certain number of iterations, the system will oscillate between two different values. (The final pair of values will be different for each value of R.) Because it oscillates between two values, the system is said to have *period* equal to 2.

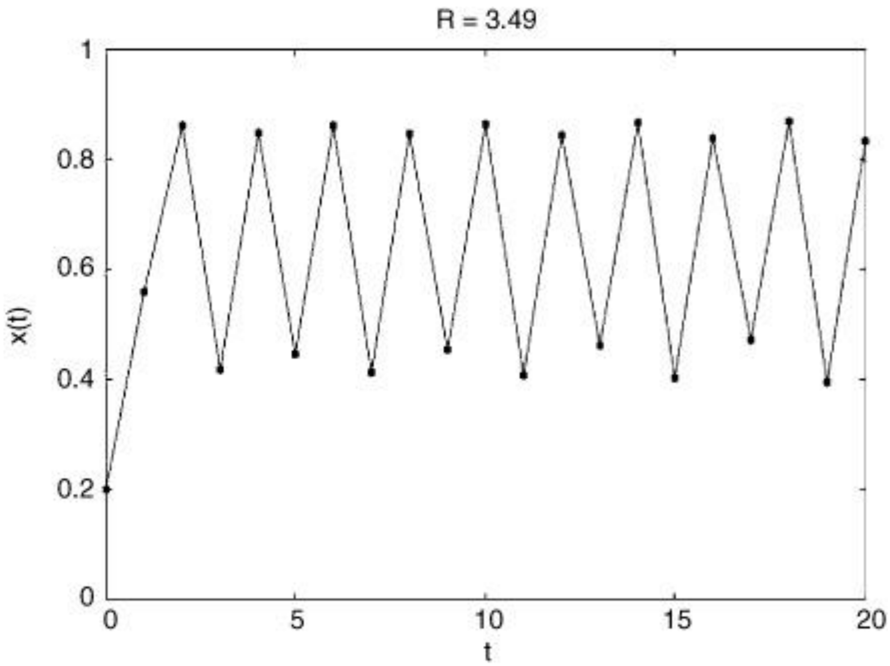

FIGURE 2.9. Behavior of the logistic map for $R = 3.49$ and $x_0 = 0.2$.

But at a value between $R = 3.4$ and $R = 3.5$ an abrupt change occurs. Given any value of x_0, the system will eventually reach an oscillation among *four* distinct values instead of two. For example, if we set $R = 3.49$, $x_0 = 0.2$, we see the results in figure 2.9.

Indeed, the values of x fairly quickly reach an oscillation among four different values (which happen to be approximately 0.872, 0.389, 0.829, and 0.494, if you're interested). That is, at some R between 3.4 and 3.5, the period of the final oscillation has abruptly doubled from 2 to 4.

Somewhere between $R = 3.54$ and $R = 3.55$ the period abruptly doubles again, jumping to 8. Somewhere between 3.564 and 3.565 the period jumps to 16. Somewhere between 3.5687 and 3.5688 the period jumps to 32. The period doubles again and again after smaller and smaller increases in R until, in short order, the period becomes effectively infinite, at an R value of approximately 3.569946. Before this point, the behavior of the logistic map was roughly predictable. If you gave me the value for R, I could tell you the ultimate long-term behavior from any starting point x_0: fixed points are reached when R is less than about 3.1, period-two oscillations are reached when R is between 3.1 and 3.4, and so on.

When R is approximately 3.569946, the values of x no longer settle into an oscillation; rather, they become chaotic. Here's what this means. Let's call the series of values x_0, x_1, x_2, and so on the *trajectory* of x. At values of R that yield chaos, two trajectories starting from very similar values of x_0, rather than converging to the same fixed point or oscillation, will instead progressively diverge from each other. At R = 3.569946 this divergence occurs very slowly, but we can see a more dramatic sensitive dependence on x_0 if we set R = 4.0. First I set x_0 = 0.2 and iterate the logistic map to obtain a trajectory. Then I restarted with a new x_0, increased slightly by putting a 1 in the tenth decimal place, x_0 = 0.2000000001, and iterated the map again to obtain a second trajectory. In figure 2.10 the first trajectory is the dark curve with black circles, and the second trajectory is the light line with open circles.

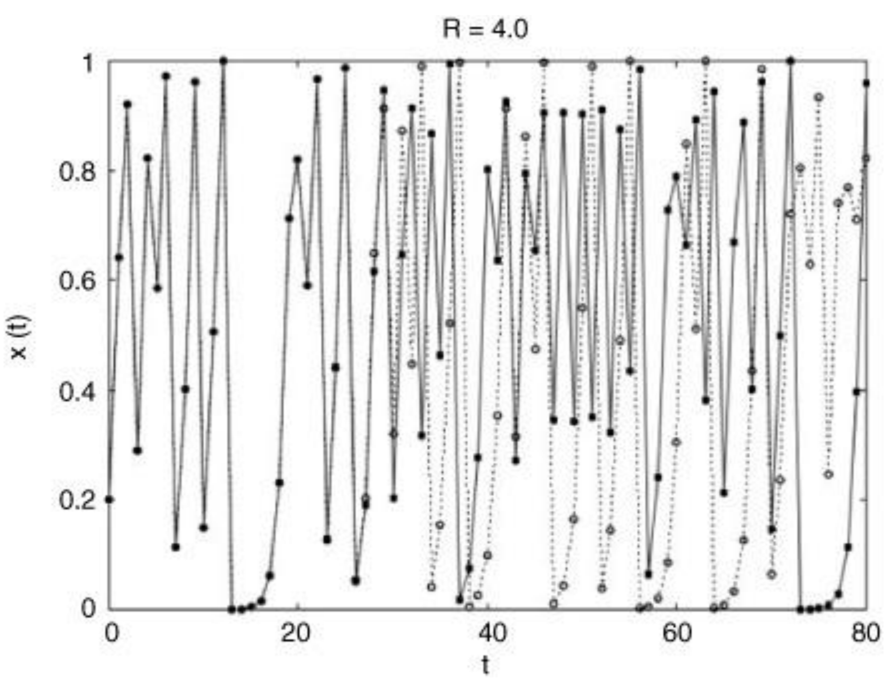

FIGURE 2.10. Two trajectories of the logistic map for $R = 4.0$: $x_0 = 0.2$ and $x_0 = 0.2000000001$.

The two trajectories start off very close to one another (so close that the first, solid-line trajectory blocks our view of the second, dashed-line trajectory), but after 30 or so iterations they start to diverge significantly, and soon after there is no correlation between them. This is what is meant by "sensitive dependence on initial conditions."

So far we have seen three different classes of final behavior (attractors): fixed-point, periodic, and chaotic. (Chaotic attractors are also sometimes called "strange attractors.") *Type of attractor* is one way in which dynamical systems theory characterizes the behavior of a system.

Let's pause a minute to consider how remarkable the chaotic behavior really is. The logistic map is an extremely simple equation and is completely deterministic: every x_t maps onto one and only one value of x_{t+1}. And yet the chaotic trajectories obtained from this map, at certain values of R, look very random—enough so that the logistic map has been used as a basis for generating pseudo-random numbers on a computer. Thus apparent randomness can arise from very simple deterministic systems.

Moreover, for the values of R that produce chaos, if there is any uncertainty in the initial condition x_0, there exists a time beyond which the future value cannot be predicted. This was demonstrated above with $R = 4$. If we don't know the value of the tenth and higher decimal places of x_0—a quite likely limitation for many experimental observations—then by $t = 30$ or so the value of x_t is unpredictable. For any value of R that yields chaos, uncertainty in any decimal place of x_0, however far out in the decimal expansion, will result in unpredictability at some value of t.

Robert May, the mathematical biologist, summed up these rather surprising properties, echoing Poincaré:

> The fact that the simple and deterministic equation (1) [i.e., the logistic map] can possess dynamical trajectories which look like some sort of random noise has disturbing practical implications. It means, for example, that apparently erratic fluctuations in the

> census data for an animal population need not necessarily betoken either the vagaries of an unpredictable environment or sampling errors: they may simply derive from a rigidly deterministic population growth relationship such as equation (1).... Alternatively, it may be observed that in the chaotic regime arbitrarily close initial conditions can lead to trajectories which, after a sufficiently long time, diverge widely. This means that, even if we have a simple model in which all the parameters are determined exactly, long-term prediction is nevertheless impossible.

In short, the presence of chaos in a system implies that perfect prediction *à la* Laplace is impossible not only in practice but also *in principle*, since we can never know x_0 to infinitely many decimal places. This is a profound negative result that, along with quantum mechanics, helped wipe out the optimistic nineteenth-century view of a clockwork Newtonian universe that ticked along its predictable path.

But is there a more positive lesson to be learned from studies of the logistic map? Can it help the goal of dynamical systems theory, which attempts to discover general principles concerning systems that change over time? In fact, deeper studies of the logistic map and related maps have resulted in an equally surprising and profound positive result—the discovery of universal characteristics of chaotic systems.

Universals in Chaos

The term *chaos*, as used to describe dynamical systems with sensitive dependence on initial conditions, was first coined by physicists T. Y. Li and James Yorke. The term seems apt: the colloquial sense of the word "chaos" implies randomness and unpredictability, qualities we have seen in the chaotic version of logistic map. However, unlike colloquial chaos, there turns out to be substantial order in mathematical chaos in the form of so-called *universal* features that are common to a wide range of chaotic systems.

THE FIRST UNIVERSAL FEATURE: THE PERIOD-DOUBLING ROUTE TO CHAOS

In the mathematical explorations we performed above, we saw that as *R* was increased from 2.0 to 4.0, iterating the logistic map for a given value of *R* first yielded a fixed point, then a period-two oscillation, then period four, then eight, and so on, until chaos was reached. In dynamical systems theory, each of these abrupt period doublings is called a *bifurcation*. This succession of bifurcations culminating in chaos has been called the "period doubling route to chaos."

These bifurcations are often summarized in a so-called bifurcation diagram that plots the attractor the system ends up in as a function of the value of a "control parameter" such as *R*. Figure 2.11 gives such a bifurcation diagram for the logistic map. The horizontal axis gives *R*. For each value of *R*, the final (attractor) values of *x* are plotted. For example, for $R = 2.9$, *x* reaches a fixed-point attractor of $x = 0.655$. At $R = 3.0$, *x* reaches a period-two attractor. This can be seen as the first branch point in the diagram, when the fixed-point attractors give way to the period-two attractors. For *R* somewhere between 3.4 and 3.5, the diagram shows a bifurcation to a period-four attractor, and so on, with further period doublings, until the onset of chaos at *R* approximately equal to 3.569946.

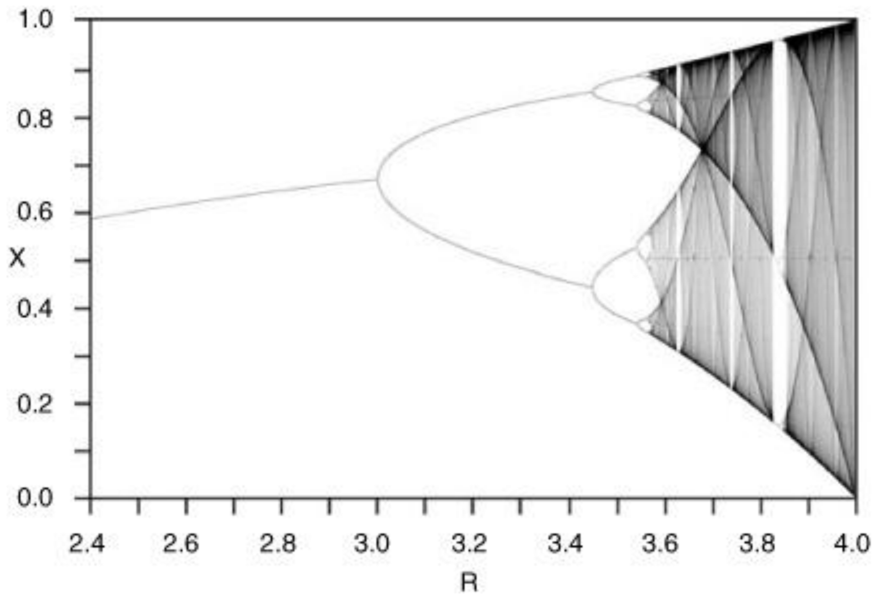

FIGURE 2.11. Bifurcation diagram for the logistic map, with attractor plotted as a function of *R*.

The period-doubling route to chaos has a rich history. Period doubling bifurcations had been observed in mathematical equations as early as the 1920s, and a similar cascade of bifurcations was described by P. J. Myrberg, a Finnish mathematician, in the 1950s. Nicholas Metropolis, Myron Stein, and Paul Stein, working at Los Alamos National Laboratory, showed that not just the logistic map but *any* map whose graph is parabola-shaped will follow a similar period-doubling route. Here, "parabola-shaped" means that plot of the map has just one hump—in mathematical terms, it is "unimodal."

THE SECOND UNIVERSAL FEATURE: FEIGENBAUM'S CONSTANT

The discovery that gave the period-doubling route its renowned place among mathematical universals was made in the 1970s by the physicist Mitchell Feigenbaum. Feigenbaum, using only a programmable desktop calculator, made a list of the R values at which the period-doubling bifurcations occur (where ≈ means "approximately equal to"):

$R_1 \approx 3.0$

$R_2 \approx 3.44949$

$R_3 \approx 3.54409$

$R_4 \approx 3.564407$

$R_5 \approx 3.568759$

$R_6 \approx 3.569692$

$R_7 \approx 3.569891$

$R_8 \approx 3.569934$

.
.
.

$R_\infty \approx 3.569946$

Here, R_1 corresponds to period 2^1 (= 2), R_2 corresponds to period 2^2 (= 4), and in general, R_n corresponds to period 2^n. The symbol ∞

(“infinity”) is used to denote the onset of chaos—a trajectory with an infinite period.

Feigenbaum noticed that as the period increases, the *R* values get closer and closer together. This means that for each bifurcation, *R* has to be increased less than it had before to get to the next bifurcation. You can see this in the bifurcation diagram of Figure 2.11: as *R* increases, the bifurcations get closer and closer together. Using these numbers, Feigenbaum measured the *rate* at which the bifurcations get closer and closer; that is, the rate at which the *R* values *converge*. He discovered that the rate is (approximately) the constant value 4.6692016. What this means is that as *R* increases, each new period doubling occurs about 4.6692016 times faster than the previous one.

This fact was interesting but not earth-shaking. Things started to get a lot more interesting when Feigenbaum looked at some other maps—the logistic map is just one of many that have been studied. As I mentioned above, a few years before Feigenbaum made these calculations, his colleagues at Los Alamos, Metropolis, Stein, and Stein, had shown that any unimodal map will follow a similar period-doubling cascade. Feigenbaum’s next step was to calculate the rate of convergence for some other unimodal maps. He started with the so-called sine map, an equation similar to the logistic map but which uses the trigonometric sine function.

Feigenbaum repeated the steps I sketched above: he calculated the values of *R* at the period-doubling bifurcations in the sine map, and then calculated the rate at which these values converged. He found that the rate of convergence was 4.6692016.

Feigenbaum was amazed. The rate was the same. He tried it for other unimodal maps. It was still the same. No one, including Feigenbaum, had expected this at all. But once the discovery had been made, Feigenbaum went on to develop a mathematical theory that explained why the common value of 4.6692016, now called *Feigenbaum’s constant*, is universal—which here means the same for all unimodal maps. The theory used a sophisticated mathematical technique called *renormalization* that had been developed originally in the area of quantum field theory and later imported to another field of physics: the study of phase transitions and other “critical phenomena.” Feigenbaum adapted it for dynamical systems theory, and it has become a cornerstone in the understanding of chaos.

It turned out that this is not just a mathematical curiosity. In the years since Feigenbaum's discovery, his theory has been verified in several laboratory experiments on physical dynamical systems, including fluid flow, electronic circuits, lasers, and chemical reactions. Period-doubling cascades have been observed in these systems, and values of Feigenbaum's constant have been calculated in steps similar to those we saw above. It is often quite difficult to get accurate measurements of, say, what corresponds to *R* values in such experiments, but even so, the values of Feigenbaum's constant found by the experimenters agree well within the margin of error to Feigenbaum's value of approximately 4.6692016. This is impressive, since Feigenbaum's theory, which yields this number, involves only abstract math, no physics. As Feigenbaum's colleague Leo Kadanoff said, this is "the best thing that can happen to a scientist, realizing that something that's happened in his or her mind exactly corresponds to something that happens in nature."

Mitchell Feigenbaum (AIP Emilio Segre Visual Archives, *Physics Today* Collection)

Large-scale systems such as the weather are, as yet, too hard to experiment with directly, so no one has *directly* observed period doubling or chaos in their behavior. However, certain computer models of weather have displayed the period-doubling route to chaos, as have computer models of electrical power systems, the heart, solar variability, and many other systems.

There is one more remarkable fact to mention about this story. Similar to many important scientific discoveries, Feigenbaum's discoveries were also made, independently and at almost the same time, by another research team. This team consisted of the French scientists Pierre Coullet and Charles Tresser, who also used the technique of renormalization to study the period-doubling cascade and discovered the universality of 4.6692016 for unimodal maps. Feigenbaum may actually have been the first to make the discovery and was also able to more widely and clearly disseminate the result among the international scientific community, which is why he has received most of the credit for this work. However, in many technical papers, the theory is referred to as the "Feigenbaum-Coullet-Tresser theory" and Feigenbaum's constant as the "Feigenbaum-Coullet-Tresser constant." In the course of this book I point out several other examples of independent, simultaneous discoveries using ideas that are "in the air" at a given time.

Revolutionary Ideas from Chaos

The discovery and understanding of chaos, as illustrated in this chapter, has produced a rethinking of many core tenets of science. Here I summarize some of these new ideas, which few nineteenth-century scientists would have believed.

- Seemingly random behavior can emerge from deterministic systems, with no external source of randomness.
- The behavior of some simple, deterministic systems can be impossible, *even in principle*, to predict in the long term, due to sensitive dependence on initial conditions.
- Although the detailed behavior of a chaotic system cannot be predicted, there is some "order in chaos" seen in universal

properties common to large sets of chaotic systems, such as the period-doubling route to chaos and Feigenbaum's constant. Thus even though "prediction becomes impossible" at the detailed level, there are some higher-level aspects of chaotic systems that are indeed predictable.

In summary, changing, hard-to-predict macroscopic behavior is a hallmark of complex systems. Dynamical systems theory provides a mathematical vocabulary for characterizing such behavior in terms of bifurcations, attractors, and universal properties of the ways systems can change. This vocabulary is used extensively by complex systems researchers.

The logistic map is a simplified model of population growth, but the detailed study of it and similar model systems resulted in a major revamping of the scientific understanding of order, randomness, and predictability. This illustrates the power of *idea models*—models that are simple enough to study via mathematics or computers but that nonetheless capture fundamental properties of natural complex systems. Idea models play a central role in this book, as they do in the sciences of complex systems.

Characterizing the dynamics of a complex system is only one step in understanding it. We also need to understand how these dynamics are used in living systems to process information and adapt to changing environments. The next three chapters give some background on these subjects, and later in the book we see how ideas from dynamics are being combined with ideas from information theory, computation, and evolution.

CHAPTER 3
Information

The law that entropy increases—the Second Law of Thermodynamics —holds, I think, the supreme position among the laws of Nature… [*I*] *f your theory is found to be against the Second Law of Thermodynamics I can give you no hope; there is nothing for it but to collapse in deepest humiliation*.

—Sir Arthur Eddington, *The Nature of the Physical World*

COMPLEX SYSTEMS ARE OFTEN said to be "self-organizing": consider, for example, the strong, structured bridges made by army ants; the synchronous flashing of fireflies; the mutually sustaining markets of an economy; and the development of specialized organs by stem cells—all are examples of self-organization. Order is created out of disorder, upending the usual turn of events in which order decays and disorder (or *entropy*) wins out.

A complete account of how such entropy-defying self-organization takes place is the holy grail of complex systems science. But before this can be tackled, we need to understand what is meant by "order" and "disorder" and how people have thought about measuring such abstract qualities.

Many complex systems scientists use the concept of *information* to characterize and measure order and disorder, complexity and simplicity. The immunologist Irun Cohen states that "complex systems sense, store, and deploy more information than do simple systems." The economist Eric Beinhocker writes that "evolution can perform its tricks not just in the 'substrate' of DNA but in any system that has the right information processing and information storage characteristics."

The physicist Murray Gell-Mann said of complex adaptive systems that "Although they differ widely in their physical attributes, they resemble one another in the way they handle information. That common feature is perhaps the best starting point for exploring how they operate."

But just what is meant by "information"?

What Is Information?

You see the word "information" all over the place these days: the "information revolution," the "information age," "information technology" (often simply "IT"), the "information superhighway," and so forth. "Information" is used colloquially to refer to any medium that presents knowledge or facts: newspapers, books, my mother on the phone gossiping about relatives, and, most prominently these days, the Internet. More technically, it is used to describe a vast array of phenomena ranging from the fiber-optic transmissions that constitute signals from one computer to another on the Internet to the tiny molecules that neurons use to communicate with one another in the brain.

The different examples of complex systems I described in chapter 1 are all centrally concerned with the communication and processing of information in various forms. Since the beginning of the computer age, computer scientists have thought of information transmission and computation as something that takes place not only in electronic circuits but also in living systems.

In order to understand the information and computation in these systems, the first step, of course, is to have a precise definition of what is meant by the terms *information* and *computation*. These terms have been mathematically defined only in the twentieth century. Unexpectedly, it all began with a late nineteenth-century puzzle in physics involving a very smart "demon" who seemed to get a lot done without expending any energy. This little puzzle got many physicists quite worried that one of their fundamental laws might be wrong. How did the concept of *information* save the day? Before getting there, we need a little bit of background on the physics notions of *energy*, *work*, and *entropy*.

Energy, Work, and Entropy

The scientific study of information really begins with the science of thermodynamics, which describes energy and its interactions with matter. Physicists of the nineteenth century considered the universe to consist of two different types of entities: *matter* (e.g., solids, liquids, and vapors) and *energy* (e.g., heat, light, and sound).

Energy is roughly defined as a system's potential to "do work," which correlates well with our intuitive notion of energy, especially in this age of high-energy workaholics. The origin of the term is the Greek word, *energia*, which literally means "to work." However, physicists have a specific meaning of "work" done by an object: the amount of force applied to the object multiplied by the distance traveled by the object in the direction that force was applied.

For example, suppose your car breaks down on a flat road and you have to push it for a quarter of a mile to the nearest gas station. In physics terms, the amount of work that you expend is the amount of force with which you push the car multiplied by the distance to the gas station. In pushing the car, you transform energy stored in your body into the kinetic energy (i.e., movement) of the car, and the amount of energy that is transformed is equal to the amount of work that is done plus whatever energy is converted to heat, say, by the friction of the wheels on the road, or by your own body warming up. This so-called heat loss is measured by a quantity called *entropy*. Entropy is a measure of the energy that cannot be converted into additional work. The term "entropy" comes from another Greek word—"trope"—meaning "turning into" or "transformation."

By the end of the nineteenth century two fundamental laws concerning energy had been discovered, the so-called *laws of thermodynamics*. These laws apply to "isolated systems"—ones that do not exchange energy with any outside entity.

First law: *Energy is conserved*. The total amount of energy in the universe is constant. Energy can be transformed from one form to another, such as the transformation of stored body energy to kinetic energy of a pushed car plus the heat generated by this action. However, energy can never be created or destroyed. Thus it is said to be "conserved."

Second law: *Entropy always increases until it reaches a maximum value*. The total entropy of a system will always increase until it reaches its maximum possible value; it will never decrease on its own unless an outside agent works to decrease it.

As you've probably noticed, a room does not clean itself up, and Cheerios spilled on the floor, left to their own devices, will never find their way back into the cereal box. Someone or something has to do work to turn disorder into order.

Furthermore, transformations of energy, such as the car-pushing example above, will always produce some heat that cannot be put to work. This is why, for example, no one has found a way to take the heat generated by the back of your refrigerator and use it to produce new power for cooling the inside of the refrigerator so that it will be able to power itself. This explains why the proverbial "perpetual motion machine" is a myth.

The second law of thermodynamics is said to define the "arrow of time," in that it proves there are processes that cannot be reversed in time (e.g., heat spontaneously returning to your refrigerator and converting to electrical energy to cool the inside). The "future" is defined as the direction of time in which entropy increases. Interestingly, the second law is the only fundamental law of physics that distinguishes between past and future. All other laws are reversible in time. For example, consider filming an interaction between elementary particles such as electrons, and then showing this movie to a physicist. Now run the movie backward, and ask the physicist which version was the "real" version. The physicist won't be able to guess, since the forward and backward interactions both obey the laws of physics. This is what *reversible* means. In contrast, if you make an infrared film of heat being produced by your refrigerator, and show it forward and backward, any physicist will identify the forward direction as "correct" since it obeys the second law, whereas the backward version does not.

This is what *irreversible* means. Why is the second law different from all other physical laws? This is a profound question. As the physicist Tony Rothman points out, "Why the second law should distinguish between past and future while all the other laws of nature do not is perhaps the greatest mystery in physics."

Maxwell's Demon

The British physicist James Clerk Maxwell is most famous for his discovery of what are now called Maxwell's Equations: compact expressions of Maxwell's theory that unified electricity and magnetism. During his lifetime, he was one of the world's most highly regarded scientists, and today would be on any top fifty list of all-time greats of science.

In his 1871 book, *Theory of Heat*, Maxwell posed a puzzle under the heading "Limitation of the Second Law of Thermodynamics." Maxwell proposed a box that is divided into two halves by a wall with a hinged door, as illustrated in figure 3.1. The door is controlled by a "demon," a very small being who measures the velocity of air molecules as they whiz past him. He opens the door to let the fast ones go from the right side to the left side, and closes it when slow ones approach it from the right. Likewise, he opens the door for slow molecules moving from left to right and closes it when fast molecules approach it from the left. After some time, the box will be well organized, with all the fast molecules on the left and all the slow ones on the right. Thus entropy will have been decreased.

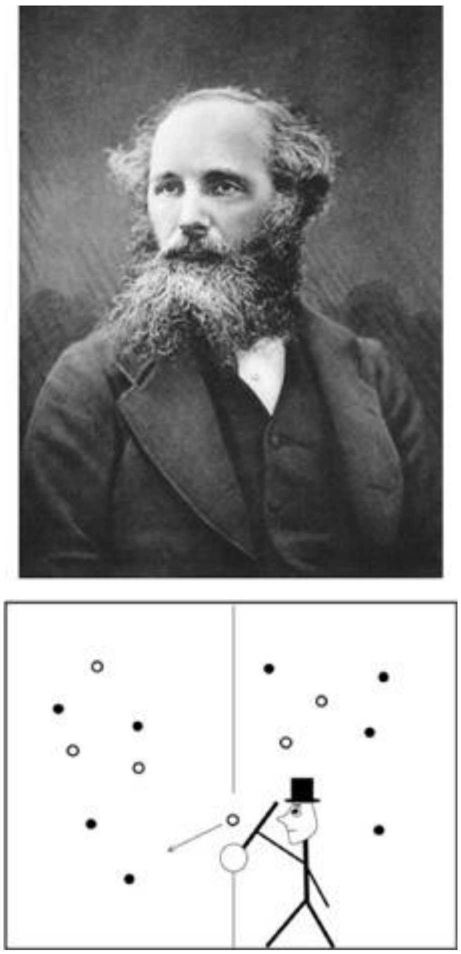

FIGURE 3.1. Top: James Clerk Maxwell, 1831–1879 (AIP Emilio Segre Visual Archives) Bottom: Maxwell's Demon, who opens the door for fast (white) particles moving to the left and for slow (black) particles moving to the right.

According to the second law, work has to be done to decrease entropy. What work has been done by the demon? To be sure, he has opened and closed the door many times. However, Maxwell assumed that a massless and frictionless "slide" could be used as a door by the demon, so that opening and closing it would require negligible work, which we can ignore. (Feasible designs for such a door have been proposed.) Has any other work been done by the demon?

Maxwell's answer was **no**: "the hot system [the left side] has gotten hotter and the cold [right side] colder and yet no work has been done, only the intelligence of a very observant and neat-fingered being has been employed."

How did entropy decrease with little or no work being done? Doesn't this directly violate the second law of thermodynamics? Maxwell's demon puzzled many of the great minds of the late nineteenth and early twentieth centuries. Maxwell's own answer to his puzzle was that the second law (the increase of entropy over time) is not really a law at all, but rather a statistical effect that holds for large collections of molecules, like the objects we encounter in day-to-day life, but does not necessarily hold at the scale of individual molecules.

However, many physicists of his day and long after vehemently disagreed. They believed that the second law has to remain inviolate; instead there must be something fishy about the demon. For entropy to decrease, work must actually have been done in some subtle, nonapparent way.

Many people tried to resolve the paradox, but no one was able to offer a satisfactory solution for nearly sixty years. In 1929, a breakthrough came: the great Hungarian physicist Leo Szilard (pronounced "ziLARD") proposed that it is the "intelligence" of the demon, or more precisely, the act of obtaining information through measurement, that constitutes the missing work.

Szilard was the first to make a link between entropy and *information*, a link that later became the foundation of information theory and a key idea in complex systems. In a famous paper entitled "On the Decrease of Entropy in a Thermodynamic System by the Intervention of Intelligent Beings," Szilard argued that the measurement process, in which the demon acquires a single "bit" of information (i.e., the information as to whether an approaching molecule is a slow one or a fast one) requires energy and must produce at least as much entropy as is decreased by the sorting of that molecule into the left or right side of the box. Thus the entire system, comprising the box, the molecules, and the demon, obeys the second law of thermodynamics.

In coming up with his solution, Szilard was perhaps the first to define the notion of a *bit of information*—the information obtained from the answer to a yes/no (or, in the demon's case, "fast/slow") question.

Leo Szilard, 1898–1964 (AIP Emilio Segre Visual Archives)

From our twenty-first-century vantage, it may seem obvious (or at least unsurprising) that the acquisition of information requires expenditure of work. But at the time of Maxwell, and even sixty years later when Szilard wrote his famous paper, there was still a strong tendency in people's minds to view physical and mental processes as completely separate. This highly ingrained intuition may be why Maxwell, as astute as he was, did not see the "intelligence" or "observing powers" of the demon as relating to the thermodynamics of the box-molecules-demon system. Such relationships between information and physics became clear only in the twentieth century, beginning with the discovery that the "observer" plays a key role in quantum mechanics.

Szilard's theory was later extended and generalized by the French physicists Leon Brillouin and Denis Gabor. Many scientists of the 1950s and later believed that Brillouin's theory in particular had definitively finished off the demon by demonstrating in detail how making a measurement entails an increase of entropy.

However, it wasn't over yet. Fifty years after Szilard's paper, it was discovered that there were some holes in Szilard's and Brillouin's solutions as well. In the 1980s, the mathematician Charles Bennett showed that there are very clever ways to observe and remember information—in the demon's case, whether an air molecule is fast or slow—*without* increasing entropy. Bennett's remarkable demonstration of this formed the basis for *reversible computing*, which says that, in theory, any computation can be done without expending energy. Bennett's discoveries might seem to imply that we are back at square one with the demon, since measurement can, in fact, be done without increasing entropy. However, Bennett noted that the second law of thermodynamics was saved again by an earlier discovery made in the 1960s by physicist Rolf Landauer: it is not the act of measurement, but rather the act of *erasing* memory that necessarily increases entropy. Erasing memory is not reversible; if there is true erasure, then once the information is gone, it cannot be restored without additional measurement. Bennett showed that for the demon to work, its memory must be erased at some point, and when it is, the physical act of this erasure will produce heat, thus increasing entropy by an amount exactly equal to the amount entropy was decreased by the demon's sorting actions.

Landauer and Bennett's solution to the paradox of Maxwell's demon fixed holes in Szilard's solution, but it was in the same spirit: the demon's act of measurement and decision making, which requires erasure, will inevitably increase entropy, and the second law is saved. (I should say here that there are still some physicists who don't buy the Landauer and Bennett solution; the demon remains controversial to this day.)

Maxwell invented his demon as a simple thought experiment to demonstrate his view that the second law of thermodynamics was not a law but a statistical effect. However, like many of the best thought-experiments in science, the demon's influence was much broader: resolutions to the demon paradox became the foundations of two new fields: information theory and the physics of information.

Statistical Mechanics in a Nutshell

In an earlier section, I defined "entropy" as a measure of the energy that cannot be converted into additional work but is instead transformed into heat. This notion of entropy was originally defined by Rudolph

Clausius in 1865. At the time of Clausius, heat was believed to be a kind of fluid that could move from one system to another, and temperature was a property of a system that affected the flow of heat.

In the next few decades, a different view of heat emerged in the scientific community: systems are made up of molecules, and heat is a result of the motion, or kinetic energy, of those molecules. This new view was largely a result of the work of Ludwig Boltzmann, who developed what is now called *statistical mechanics*.

Ludwig Boltzmann, 1844–1906 (AIP Emilio Segre Visual Archives, Segre Collection)

Statistical mechanics proposes that large-scale properties (e.g., heat) emerge from microscopic properties (e.g., the motions of trillions of molecules). For example, think about a room full of moving air molecules. A *classical* mechanics analysis would determine the position and velocity of each molecule, as well as all the forces acting on that molecule, and would use this information to determine the future position and velocity of that molecule. Of course, if there are fifty quadrillion molecules, this approach would take rather a long time—in fact it always would be impossible, both in practice and, as quantum mechanics has shown, in principle. A *statistical* mechanics approach gives up on determining the exact position, velocity, and future behavior of each molecule and instead tries to predict the *average* positions and velocities of large *ensembles* of molecules.

In short, classical mechanics attempts to say something about every single microscopic entity (e.g., molecule) by using Newton's laws. Thermodynamics gives laws of macroscopic entities—heat, energy, and entropy—without acknowledging that any microscopic molecules are the source of these macroscopic entities. Statistical mechanics is a bridge between these two extremes, in that it explains how the behavior of the macroscopic entities arise from *statistics* of large ensembles of microscopic entities.

There is one problem with the statistical approach—it gives only the *probable* behavior of the system. For example, if all the air molecules in a room are flying around randomly, they are most likely to be spread out all over the room, and all of us will have enough air to breathe. This is what we predict and depend on, and it has never failed us yet. However, according to statistical mechanics, since the molecules are flying around randomly, there is some very small chance that at some point they will all fly over to the same corner at the same time. Then any person who happened to be in that corner would be crushed by the huge air pressure, and the rest of us would suffocate from lack of air. As far as I know, such an event has never happened in any room anywhere. However, there is nothing in Newton's laws that says it can't happen; it's just incredibly unlikely. Boltzmann reasoned that if there are enough microscopic entities to average over, his statistical approach will give the right answer virtually all the time, and indeed, in practice it does so. But at the time Boltzmann was formulating his new science, the suggestion that a physical law could apply only "virtually all of the time" rather than *exactly* all of the time was repellent to many other scientists. Furthermore, Boltzmann's insistence on the reality of microscopic entities such as molecules and atoms was also at odds with his colleagues. Some have speculated that the rejection of his ideas by most of his fellow scientists contributed to his suicide in 1906, at the age of 62. Only years after his death were his ideas generally accepted; he is now considered to be one of the most important scientists in history.

Microstates and Macrostates

Given a room full of air, at a given instant in time each molecule has a certain position and velocity, even if it is impossible to actually measure all of them. In statistical mechanics terminology, the particular collection of exact molecule positions and velocities at a given instant is called the *microstate* of the whole room at that instant. For a room full of air molecules randomly flying around, the most probable type of microstate at a given time is that the air molecules are spread uniformly around the room. The least probable type of microstate is that the air molecules are all clumped together as closely as possible in a single location, for example, the corner of the room. This seems simply obvious, but Boltzmann noted that the reason for this is that there are many more possible microstates of the system in which the air molecules are spread around uniformly than there are microstates in which they all are clumped together.

The situation is analogous to a slot machine with three rotating pictures (figure 3.2). Suppose each of the three pictures can come up "apple," "orange," "cherry," "pear," or "lemon." Imagine you put in a quarter, and pull the handle to spin the pictures. It is much more likely that the pictures will all be different (i.e., you lose your money) than that the pictures will all be the same (i.e., you win a jackpot). Now imagine such a slot machine with fifty quadrillion pictures, and you can see that the probability of all coming up the same is very close to zero, just like the probability of the air molecules ending up all clumped together in the same location.

FIGURE 3.2. Slot machine with three rotating fruit pictures, illustrating the concepts *microstate* and *macrostate*. (Drawing by David Moser.)

A *type of microstate*, for example, “pictures all the same—you win” versus “pictures not all the same—you lose” or “molecules clumped together—we can’t breathe” versus “molecules uniformly spread out—we can breathe,” is called a *macrostate* of the system. A macrostate can correspond to many different microstates. In the slot machine, there are many different microstates consisting of three nonidentical pictures, each of which corresponds to the single “you lose” macrostate, and only a few microstates that correspond to the “you win” macrostate. This is how casinos are sure to make money. *Temperature* is a macrostate—it corresponds to many different possible microstates of molecules at different velocities that happen to average to the same temperature.

Using these ideas, Boltzmann interpreted the second law of thermodynamics as simply saying that an isolated system will more likely be in a more probable macrostate than in a less probable one. To our ears this sounds like a tautology but it was a rather revolutionary way of thinking about the point back then, since it included the notion of probability. Boltzmann defined the *entropy* of a macrostate as a function of the number of microstates that could give rise to that macrostate. For example, on the slot machine of figure 3.2, where each picture can come up “apple,” “orange,” “cherry,” “pear,” or “lemon,” it turns out that there are a total of 125 possible combinations (microstates), out of which five correspond to the macrostate “pictures all the same—you win” and 120 correspond to the macrostate “pictures not all the same—you lose.” The latter macrostate clearly has a higher Boltzmann entropy than the former.

FIGURE 3.3. Boltzmann's tombstone, in Vienna. (Photograph courtesy of Martin Roell.)

Boltzmann's entropy obeys the second law of thermodynamics. Unless *work* is done, Boltzmann's entropy will always increase until it gets to a macrostate with highest possible entropy. Boltzmann was able to show that, under many conditions, his simple and intuitive definition of entropy is equivalent to the original definition of Clausius.

The actual equation for Boltzmann's entropy, now so fundamental to physics, appears on Boltzmann's tombstone in Vienna (figure 3.3).

Shannon Information

Many of the most basic scientific ideas are spurred by advances in technology. The nineteenth-century studies of thermodynamics were inspired and driven by the challenge of improving steam engines. The studies of information by mathematician Claude Shannon were likewise driven by the twentieth-century revolution in communications—particularly the development of the telegraph and telephone. In the 1940s, Shannon adapted Boltzmann's ideas to the more abstract realm of communications. Shannon worked at Bell Labs, a part of the American Telephone and Telegraph Company (AT&T). One of the most important problems for AT&T was to figure out how to transmit signals more quickly and reliably over telegraph and telephone wires.

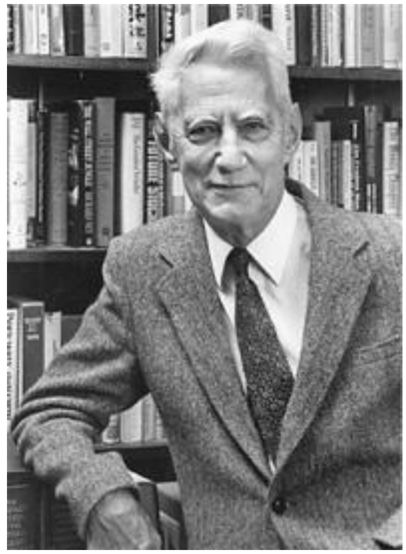

Claude Shannon, 1916–2001. (Reprinted with permission of Lucent Technologies Inc./Bell Labs.)

Shannon's mathematical solution to this problem was the beginning of what is now called *information theory*. In his 1948 paper "A Mathematical Theory of Communication," Shannon gave a narrow definition of information and proved a very important theorem, which gave the maximum possible transmission rate of information over a given channel (wire or other medium), even if there are errors in transmission caused by noise on the channel. This maximum transmission rate is called the *channel capacity*.

Shannon's definition of information involves a *source* that sends *messages* to a *receiver*. For example, figure 3.4 shows two examples of a source talking to a receiver on the phone. Each word the source says can be considered a *message* in the Shannon sense. Just as the telephone doesn't understand the words being said on it but only transmits the electrical pulses used to encode the voice, Shannon's definition of information completely ignores the *meaning* of the messages and takes into account only how often the source sends each of the possible different messages to the receiver.

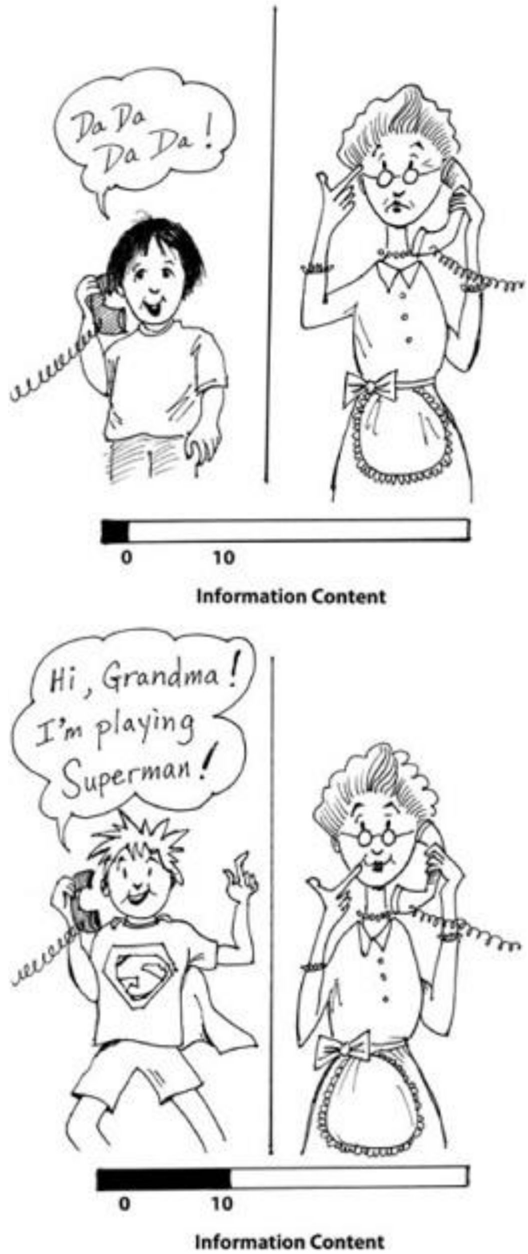

FIGURE 3.4. Top: Information content (zero) of Nicky's conversation with Grandma. Bottom: Higher information content of Jake's conversation with Grandma. (Drawings by David Moser.)

Shannon asked, "How much information is transmitted by a source sending messages to a receiver?" In analogy with Boltzmann's ideas, Shannon defined the information of a macrostate (here, a source) as a function of the number of possible microstates (here, ensembles of possible messages) that could be sent by that source. When my son Nicky was barely a toddler, I would put him on the phone to talk with Grandma. He loved to talk on the phone, but could say only one word—"da." His messages to Grandma were "da da da da da...." In other words, the Nicky-macrostate had only one possible microstate (sequences of "da"s), and although the macrostate was cute, the information content was, well, zero. Grandma knew just what to expect. My son Jake, two years older, also loved to talk on the phone but had a much bigger vocabulary and would tell Grandma all about his activities, projects, and adventures, constantly surprising her with his command of language. Clearly the information content of the Jake-source was much higher, since so many microstates—i.e., more different collections of messages—could be produced.

Shannon's definition of information content was nearly identical to Boltzmann's more general definition of entropy. In his classic 1948 paper, Shannon defined the information content in terms of the *entropy* of the message source. (This notion of entropy is often called *Shannon entropy* to distinguish it from the related definition of entropy given by Boltzmann.)

People have sometimes characterized Shannon's definition of information content as the "average amount of surprise" a receiver experiences on receiving a message, in which "surprise" means something like the "degree of uncertainty" the receiver had about what the source would send next. Grandma is clearly more surprised at each word Jake says than at each word Nicky says, since she already knows exactly what Nicky will say next but can't as easily predict what Jake will say next. Thus each word Jake says gives her a higher average "information content" than each word Nicky says.

In general, in Shannon's theory, a message can be any unit of communication, be it a letter, a word, a sentence, or even a single bit (a zero or a one). Once again, the entropy (and thus information content) of a source is defined in terms of message probabilities and is not concerned with the "meaning" of a message.

Shannon's results set the stage for applications in many different fields. The best-known applications are in the field of coding theory, which deals with both data compression and the way codes need to be structured to be reliably transmitted. Coding theory affects nearly all of our electronic communications; cell phones, computer networks, and the worldwide global positioning system are a few examples.

Information theory is also central in cryptography and in the relatively new field of bioinformatics, in which entropy and other information theory measures are used to analyze patterns in gene sequences. It has also been applied to analysis of language and music and in psychology, statistical inference, and artificial intelligence, among many other fields. Although information theory was inspired by notions of entropy in thermodynamics and statistical mechanics, it is controversial whether or not information theory has had much of a reverse impact on those and other fields of physics. In 1961, communications engineer and writer John Pierce quipped that "efforts to marry communication theory and physics have been more interesting than fruitful." Some physicists would still agree with him. However, there are a number of new approaches to physics based on concepts related to Shannon's information theory (e.g., quantum information theory and the physics of information) that are beginning to be fruitful as well as interesting.

As you will see in subsequent chapters, information theoretic notions such as entropy, information content, mutual information, information dynamics, and others have played central though controversial roles in attempts to define the notion of complexity and in characterizing different types of complex systems.

CHAPTER 7
Defining and Measuring Complexity

THIS BOOK IS ABOUT COMPLEXITY, but so far I haven't defined this term rigorously or given any clear way to answer questions such as these: Is a human brain more complex than an ant brain? Is the human genome more complex than the genome of yeast? Did complexity in biological organisms increase over the last four billion years of evolution? Intuitively, the answer to these questions would seem to be "of course." However, it has been surprisingly difficult to come up with a universally accepted definition of complexity that can help answer these kinds of questions.

In 2004 I organized a panel discussion on complexity at the Santa Fe Institute's annual Complex Systems Summer School. It was a special year: 2004 marked the twentieth anniversary of the founding of the institute. The panel consisted of some of the most prominent members of the SFI faculty, including Doyne Farmer, Jim Crutchfield, Stephanie Forrest, Eric Smith, John Miller, Alfred Hübler, and Bob Eisenstein—all well-known scientists in fields such as physics, computer science, biology, economics, and decision theory. The students at the school—young scientists at the graduate or postdoctoral level—were given the opportunity to ask any question of the panel. The first question was, "How do you define *complexity*?" Everyone on the panel laughed, because the question was at once so straightforward, so expected, and yet so difficult to answer. Each panel member then proceeded to give a different definition of the term. A few arguments even broke out between members of the faculty over their respective definitions. The students were a bit shocked and frustrated. If the faculty of the Santa Fe Institute—the most famous institution in the world devoted to research on complex systems—could not agree on what was meant by *complexity*, then how can there even begin to be a science of complexity?

The answer is that there is not yet a single science of complexity but rather several different sciences of complexity with different notions of what complexity means. Some of these notions are quite formal, and some are still very informal. If the *sciences* of complexity are to become a unified *science* of complexity, then people are going to have to figure out how these diverse notions—formal and informal—are related to one another, and how to most usefully refine the overly complex notion of *complexity*. This is work that largely remains to be done, perhaps by those shocked and frustrated students as they take over from the older generation of scientists.

I don't think the students should have been shocked and frustrated. Any perusal of the history of science will show that the lack of a universally accepted definition of a central term is more common than not. Isaac Newton did not have a good definition of force, and in fact, was not happy about the concept since it seemed to require a kind of magical "action at a distance," which was not allowed in mechanistic explanations of nature. While genetics is one of the largest and fastest growing fields of biology, geneticists still do not agree on precisely what the term *gene* refers to at the molecular level. Astronomers have discovered that about 95% of the universe is made up of "dark matter" and "dark energy" but have no clear idea what these two things actually consist of. Psychologists don't have precise definitions for *idea* or *concept*, or know what these correspond to in the brain. These are just a few examples. Science often makes progress by inventing new terms to describe incompletely understood phenomena; these terms are gradually refined as the science matures and the phenomena become more completely understood. For example, physicists now understand all forces in nature to be combinations of four different kinds of fundamental forces: electromagnetic, strong, weak, and gravitational. Physicists have also theorized that the seeming "action at a distance" arises from the interaction of elementary particles. Developing a single theory that describes these four fundamental forces in terms of quantum mechanics remains one of the biggest open problems in all of physics. Perhaps in the future we will be able to isolate the different fundamental aspects of "complexity" and eventually unify all these aspects in some overall understanding of what we now call complex phenomena.

The physicist Seth Lloyd published a paper in 2001 proposing three different dimensions along which to measure the complexity of an object or process:

How hard is it to describe?

How hard is it to create?

What is its degree of organization?

Lloyd then listed about forty measures of complexity that had been proposed by different people, each of which addressed one or more of these three questions using concepts from dynamical systems, thermodynamics, information theory, and computation. Now that we have covered the background for these concepts, I can sketch some of these proposed definitions.

To illustrate these definitions, let's use the example of comparing the complexity of the human genome with the yeast genome. The human genome contains approximately three billion base pairs (i.e., pairs of nucleotides). It has been estimated that humans have about 25,000 genes—that is, regions that code for proteins. Surprisingly, only about 2% of base pairs are actually parts of genes; the nongene parts of the genome are called *noncoding regions*. The noncoding regions have several functions: some of them help keep their chromosomes from falling apart; some help control the workings of actual genes; some may just be "junk" that doesn't really serve any purpose, or has some function yet to be discovered.

I'm sure you've heard of the Human Genome project, but you may not know that there was also a Yeast Genome Project, in which the complete DNA sequences of several varieties of yeast were determined. The first variety that was sequenced turned out to have approximately twelve million base pairs and six thousand genes.

Complexity as Size

One simple measure of complexity is size. By this measure, humans are about 250 times as complex as yeast if we compare the number of base pairs, but only about four times as complex if we count genes.

Since 250 is a pretty big number, you may now be feeling rather complex, at least as compared with yeast. However, disappointingly, it turns out that the amoeba, another type of single-celled microorganism, has about 225 times as many base pairs as humans do, and a mustard plant called *Arabidopsis* has about the same number of genes that we do.

Humans are obviously more complex than amoebae or mustard plants, or at least I would like to think so. This means that genome size is not a very good measure of complexity; our complexity must come from something deeper than our absolute number of base pairs or genes (See figure 7.1).

Complexity as Entropy

Another proposed measure of the complexity of an object is simply its Shannon entropy, defined in chapter 3 to be the average information content or "amount of surprise" a message source has for a receiver. In our example, we could define a *message* to be one of the symbols A, C, G, or T. A highly ordered and very easy-to-describe sequence such as "A A A A A A A... A" has entropy equal to zero. A completely random sequence has the maximum possible entropy.

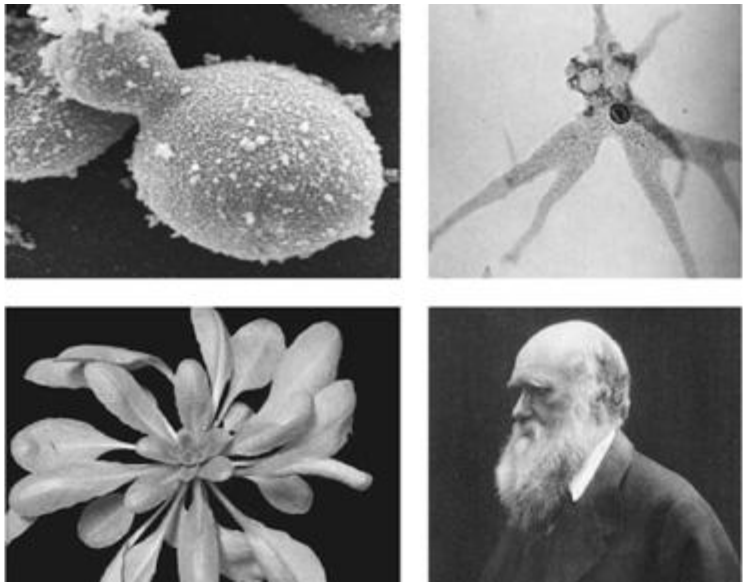

FIGURE 7.1. Clockwise from top left: Yeast, an amoeba, a human, and *Arabidopsis*. Which is the most complex? If you used genome length as the measure of complexity, then the amoeba would win hands down (if only it had hands). (Yeast photograph from NASA, [http://www.nasa.gov/mission_pages/station/science/experiments/Yeast-GAP.html]; amoeba photograph from NASA [http://ares.jsc.nasa.gov/astrobiology/biomarkers/_images/amoeba.jpg]; *Arabidopsis* photograph courtesy of Kirsten Bomblies; Darwin photograph reproduced with permission from John van Wyhe, ed., The Complete Work of Charles Darwin Online [http://darwin-online.org.uk/].)

There are a few problems with using Shannon entropy as a measure of complexity. First, the object or process in question has to be put in the form of "messages" of some kind, as we did above. This isn't always easy or straightforward—how, for example, would we measure the entropy of the human brain? Second, the highest entropy is achieved by a random set of messages. We could make up an artificial genome by choosing a bunch of random As, Cs, Gs, and Ts. Using entropy as the measure of complexity, this random, almost certainly nonfunctional genome would be considered more complex than the human genome. Of course one of the things that makes humans complex, in the intuitive sense, is precisely that our genomes aren't random but have been evolved over long periods to encode genes *useful* to our survival, such as the ones that control the development of eyes and muscles. The most complex entities are not the most ordered or random ones but somewhere in between. Simple Shannon entropy doesn't capture our intuitive concept of complexity.

Complexity as Algorithmic Information Content

Many people have proposed alternatives to simple entropy as a measure of complexity. Most notably Andrey Kolmogorov, and independently both Gregory Chaitin and Ray Solomonoff, proposed that the complexity of an object is the size of the shortest computer program that could generate a complete description of the object. This is called the *algorithmic information content* of the object. For example, think of a very short (artificial) string of DNA:

ACACACACACACACACACAC (string 1).

A very short computer program, "Print A C ten times," would spit out this pattern. Thus the string has low algorithmic information content. In contrast, here is a string I generated using a pseudo-random number generator:

ATCTGTCAAGACGGAACAT (string 2)

Assuming my random number generator is a good one, this string has no discernible overall pattern to it, and would require a longer program, namely "Print the exact string A T C T G T C A A A A C G G A A C A T." The idea is that string 1 is compressible, but string 2 is not, so contains more algorithmic information. Like entropy, algorithmic information content assigns higher information content to random objects than ones we would intuitively consider to be complex.

The physicist Murray Gell-Mann proposed a related measure he called "effective complexity" that accords better with our intuitions about complexity. Gell-Mann proposed that any given entity is composed of a combination of regularity and randomness. For example, string 1 above has a very simple regularity: the repeating A C motif. String 2 has no regularities, since it was generated at random. In contrast, the DNA of a living organism has some regularities (e.g., important correlations among different parts of the genome) probably combined with some randomness (e.g., true junk DNA).

To calculate the effective complexity, first one figures out the best description of the regularities of the entity; the effective complexity is defined as the amount of information contained in that description, or equivalently, the algorithmic information content of the set of regularities.

String 1 above has the regularity that it is A C repeated over and over. The amount of information needed to describe this regularity is the algorithmic information content of this regularity: the length of the program "Print A C some number of times." Thus, entities with very predictable structure have low effective complexity.

In the other extreme, string 2, being random, has no regularities. Thus there is no information needed to describe its regularities, and while the algorithmic information content of the string itself is maximal, the algorithmic information content of the string's *regularities*—its effective complexity—is zero. In short, as we would wish, both very ordered and very random entities have low effective complexity.

The DNA of a viable organism, having many independent and interdependent regularities, would have high effective complexity because its regularities presumably require considerable information to describe.

The problem here, of course, is how do we figure out what the regularities are? And what happens if, for a given system, various observers do not agree on what the regularities are?

Gell-Mann makes an analogy with scientific theory formation, which is, in fact, a process of finding regularities about natural phenomena. For any given phenomenon, there are many possible theories that express its regularities, but clearly some theories—the simpler and more elegant ones—are better than others. Gell-Mann knows a lot about this kind of thing—he shared the 1969 Nobel prize in Physics for his wonderfully elegant theory that finally made sense of the (then) confusing mess of elementary particle types and their interactions.

In a similar way, given different proposed sets of regularities that fit an entity, we can determine which is best by using the test called Occam's Razor. The best set of regularities is the smallest one that describes the entity in question and at the same time minimizes the remaining random component of that entity. For example, biologists today have found many regularities in the human genome, such as genes, regulatory interactions among genes, and so on, but these regularities still leave a lot of seemingly random aspects that don't obey any regularities—namely, all that so-called junk DNA. If the Murray Gell-Mann of biology were to come along, he or she might find a better set of regularities that is simpler than that which biologists have so far identified and that is obeyed by more of the genome.

Effective complexity is a compelling idea, though like most of the proposed measures of complexity, it is hard to actually measure. Critics also have pointed out that the subjectivity of its definition remains a problem.

Complexity as Logical Depth

In order to get closer to our intuitions about complexity, in the early 1980s the mathematician Charles Bennett proposed the notion of *logical depth*. The logical depth of an object is a measure of how difficult that object is to construct. A highly ordered sequence of A, C, G, T (e.g., string 1, mentioned previously) is obviously easy to construct. Likewise, if I asked you to give me a random sequence of A, C, G, and T, that would be pretty easy for you to do, especially with the help of a coin you could flip or dice you could roll. But if I asked you to give me a DNA sequence that would produce a viable organism, you (or any biologist) would be very hard-pressed to do so without cheating by looking up already-sequenced genomes.

In Bennett's words, "Logically deep objects… contain internal evidence of having been the result of a long computation or slow-to-simulate dynamical process, and could not plausibly have originated otherwise." Or as Seth Lloyd says, "It is an appealing idea to identify the complexity of a thing with the amount of information processed in the most plausible method of its creation."

To define logical depth more precisely, Bennett equated the *construction of an object* with the computation of a string of 0s and 1s encoding that object. For our example, we could assign to each nucleotide letter a two-digit code: *A* = 00, *C* = 01, *G* = 10, and *T* = 11. Using this code, we could turn any sequence of *A*, *C*, *G*, and T into a string of 0s and 1s. The logical depth is then defined as the number of steps that it would take for a properly programmed Turing machine, starting from a blank tape, to construct the desired sequence as its output.

Since, in general, there are different "properly programmed" Turing machines that could all produce the desired sequence in different amounts of time, Bennett had to specify which Turing machine should be used. He proposed that the shortest of these (i.e., the one with the least number of states and rules) should be chosen, in accordance with the above-mentioned Occam's Razor.

Logical depth has very nice theoretical properties that match our intuitions, but it does not give a practical way of measuring the complexity of any natural object of interest, since there is typically no practical way of finding the smallest Turing machine that could have generated a given object, not to mention determining how long that machine would take to generate it. And this doesn't even take into account the difficulty, in general, of describing a given object as a string of 0s and 1s.

Complexity as Thermodynamic Depth

In the late 1980s, Seth Lloyd and Heinz Pagels proposed a new measure of complexity, *thermodynamic depth*. Lloyd and Pagels' intuition was similar to Bennett's: more complex objects are harder to construct. However, instead of measuring the number of steps of the Turing machine needed to construct the description of an object, thermodynamic depth starts by determining "the most plausible scientifically determined sequence of events that lead to the thing itself," and measures "the total amount of thermodynamic and informational resources required by the physical construction process."

For example, to determine the thermodynamic depth of the human genome, we might start with the genome of the very first creature that ever lived and list all the evolutionary genetic events (random mutations, recombinations, gene duplications, etc.) that led to modern humans. Presumably, since humans evolved billions of years later than amoebas, their thermodynamic depth is much greater.

Like logical depth, thermodynamic depth is appealing in theory, but in practice has some problems as a method for measuring complexity. First, there is the assumption that we can, in practice, list all the events that lead to the creation of a particular object. Second, as pointed out by some critics, it's not clear from Seth Lloyd and Heinz Pagels' definition just how to define "an event." Should a genetic mutation be considered a single event or a group of millions of events involving all the interactions between atoms and subatomic particles that cause the molecular-level event to occur? Should a genetic recombination between two ancestor organisms be considered a single event, or should we include all the microscale events that cause the two organisms to end up meeting, mating, and forming offspring? In more technical language, it's not clear how to "coarse-grain" the states of the system—that is, how to determine what are the relevant *macrostates* when listing events.

Complexity as Computational Capacity

If complex systems—both natural and human-constructed—can perform computation, then we might want to measure their complexity in terms of the sophistication of what they can compute. The physicist Stephen Wolfram, for example, has proposed that systems are complex if their computational abilities are equivalent to those of a universal Turing machine. However, as Charles Bennett and others have argued, the ability to perform universal computation doesn't mean that a system by itself is complex; rather, we should measure the complexity of the behavior of the system coupled with its inputs. For example, a universal

Turing machine alone isn't complex, but together with a machine code and input that produces a sophisticated computation, it creates complex behavior.

Statistical Complexity

Physicists Jim Crutchfield and Karl Young defined a different quantity, called *statistical complexity*, which measures the minimum amount of information about the past behavior of a system that is needed to optimally predict the *statistical behavior* of the system in the future. (The physicist Peter Grassberger independently defined a closely related concept called *effective measure complexity*.) *Statistical complexity* is related to Shannon's entropy in that a system is thought of as a "message source" and its behavior is somehow quantified as discrete "messages." Here, predicting the statistical behavior consists of constructing a model of the system, based on observations of the messages the system produces, such that the model's behavior is *statistically* indistinguishable from the behavior of the system itself.

For example, a model of the message source of string 1 above could be very simple: "repeat A C"; thus its statistical complexity is low. However, in contrast to what could be done with entropy or algorithmic information content, a simple model could also be built of the message source that generates string 2: "choose at random from A, C, G, or T." The latter is possible because models of statistical complexity are permitted to include random choices. The quantitative value of statistical complexity is the information content of the simplest such model that predicts the system's behavior. Thus, like effective complexity, statistical complexity is low for both highly ordered and random systems, and is high for systems in between—those that we would intuitively consider to be complex.

Like the other measures described above, it is typically not easy to measure statistical complexity if the system in question does not have a ready interpretation as a message source. However, Crutchfield, Young, and their colleagues have actually measured the statistical complexity of a number of real-world phenomena, such as the atomic structure of complicated crystals and the firing patterns of neurons.

Complexity as Fractal Dimension

So far all the complexity measures I have discussed have been based on information or computation theoretic concepts. However, these are not the only possible sources of measures of complexity. Other people have proposed concepts from dynamical systems theory to measure the complexity of an object or process. One such measure is the *fractal dimension* of an object. To explain this measure, I must first explain what a *fractal* is.

The classic example of a fractal is a coastline. If you view a coastline from an airplane, it typically looks rugged rather than straight, with many inlets, bays, prominences, and peninsulas (Figure 7.2, top). If you then view the same coastline from your car on the coast highway, it still appears to have the exact same kind of ruggedness, but on a smaller scale (Figure 7.2, bottom). Ditto for the close-up view when you stand on the beach and even for the ultra close-up view of a snail as it crawls on individual rocks. The similarity of the shape of the coastline at different scales is called "self-similarity."

The term *fractal* was coined by the French mathematician Benoit Mandelbrot, who was one of the first people to point out that the world is full of fractals—that is, many real-world objects have a rugged self-similar structure. Coastlines, mountain ranges, snowflakes, and trees are often-cited examples. Mandelbrot even proposed that the universe is fractal-like in terms of the distribution of galaxies, clusters of galaxies, clusters of clusters, et cetera. Figure 7.3 illustrates some examples of self-similarity in nature.

Although the term *fractal* is sometimes used to mean different things by different people, in general a fractal is a geometric shape that has "fine structure at every scale." Many fractals of interest have the self-similarity property seen in the coastline example given above. The logistic-map bifurcation diagram from chapter 2 (figure 2.6) also has some degree of self-similarity; in fact the chaotic region of this (*R* greater than 3.57 or so) and many other systems are sometimes called *fractal attractors*.

Mandelbrot and other mathematicians have designed many different mathematical models of fractals in nature. One famous model is the so-called Koch curve (Koch, pronounced "Coke," is the name of the Swedish mathematician who proposed this fractal). The Koch curve is created by repeated application of a rule, as follows.

FIGURE 7.2. Top: Large-scale aerial view of Ireland, whose coastline has self-similar (fractal) properties. Bottom: Smaller-scale view of part of the Irish coastline. Its rugged structure at this scale resembles the rugged structure at the larger scale. (Top photograph from NASA Visible Earth [http://visibleearth.nasa.gov/]. Bottom photograph by

Andreas Borchet, licensed under Creative Commons [http://creativecommons.org/licenses/by/3.0/].)

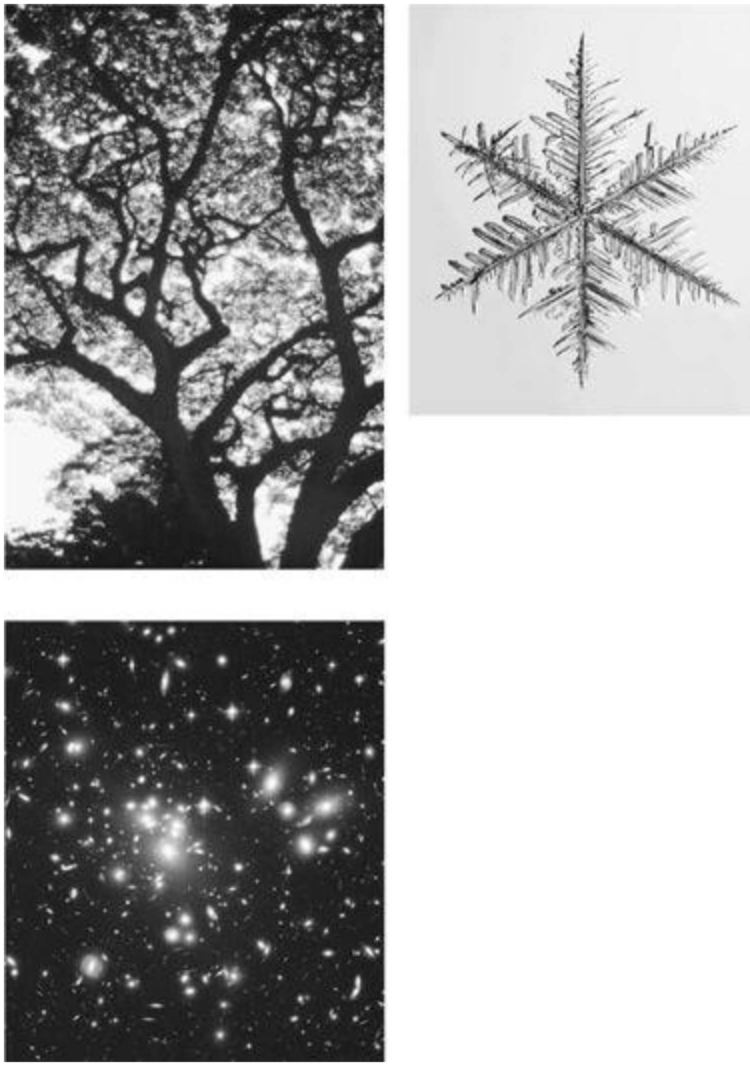

FIGURE 7.3. Other examples of fractal-like structures in nature: A tree, a snowflake (microscopically enlarged), a cluster of galaxies. (Tree photograph from the National Oceanic and Atmospheric Administration Photo Library. Snowflake photograph from [http://www.SnowCrystals.com], courtesy of Kenneth Libbrecht. Galaxy cluster photograph from NASA Space Telescope Science Institute.)

1. Start with a single line.

2. Apply the Koch curve rule: "For each line segment, replace its middle third by two sides of a triangle, each of length 1 ⁄ 3 of the original segment." Here there is only one line segment; applying the rule to it yields:

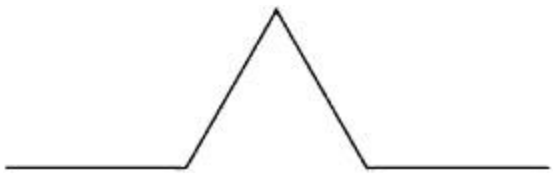

3. Apply the Koch curve rule to the resulting figure. Keep doing this forever. For example, here are the results from a second, third, and fourth application of the rule:

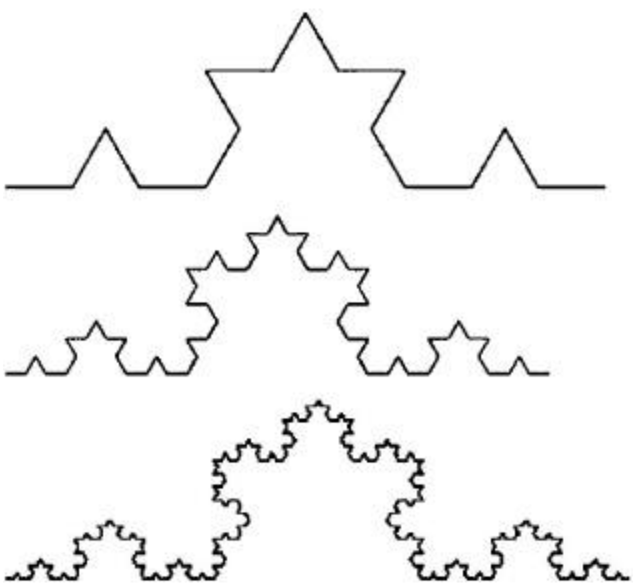

This last figure looks a bit like an idealized coastline. (In fact, if you turn the page 90 degrees to the left and squint really hard, it looks just like the west coast of Alaska.) Notice that it has true self-similarity: all of the subshapes, and their subshapes, and so on, have the same shape as the overall curve. If we applied the Koch curve rule an infinite number of times, the figure would be self-similar at an infinite number of scales—a perfect fractal. A real coastline of course does not have true self-similarity. If you look at a small section of the coastline, it does not have exactly the same shape as the entire coastline, but is visually similar in many ways (e.g., curved and rugged). Furthermore, in real-world objects, self-similarity does not go all the way to infinitely small scales. Real-world structures such as coastlines are often called “fractal” as a shorthand, but it is more accurate to call them “fractal-like,” especially if a mathematician is in hearing range.

Fractals wreak havoc with our familiar notion of spatial dimension. A line is one-dimensional, a surface is two-dimensional, and a solid is three-dimensional. What about the Koch curve?

First, let’s look at what exactly *dimension* means for regular geometric objects such as lines, squares, and cubes.

Start with our familiar line segment. Bisect it (i.e., cut it in half). Then bisect the resulting line segments, continuing at each level to bisect each line segment:

Each level is made up of two half-sized copies of the previous level.

Now start with a square. Bisect each side. Then bisect the sides of the resulting squares, continuing at each level to bisect every side:

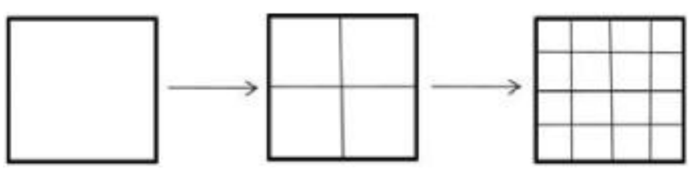

Each level is made up of four one-quarter-sized copies of the previous level.

Now, you guessed it, take a cube and bisect all the sides. Keep bisecting the sides of the resulting cubes:

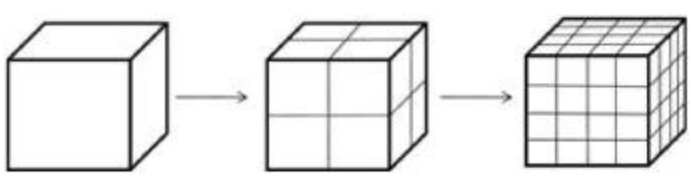

Each level is made up of eight one-eighth-sized copies of the previous level.

This sequence gives a meaning of the term *dimension*. In general, each level is made up of smaller copies of the previous level, where the number of copies is 2 raised to the power of the dimension ($2^{dimension}$). For the line, we get $2^1 = 2$ copies at each level; for the square we get $2^2 = 4$ copies at each level, and for the cube we get $2^3 = 8$ copies at each level. Similarly, if you trisect instead of bisect the lengths of the line segments at each level, then each level is made up of $3^{dimension}$ copies of the previous level. I'll state this as a general formula:

> Create a geometric structure from an original object by repeatedly dividing the length of its sides by a number *x*. Then each level is made up of $x^{dimension}$ copies of the previous level.

Indeed, according to this definition of dimension, a line is one-dimensional, a square two-dimensional and a cube three-dimensional. All good and well.

Let's apply an analogous definition to the object created by the Koch rule. At each level, the line segments of the object are three times smaller than before, and each level consists of four copies of the previous level. By our definition above, it must be true that $3^{dimension}$ is equal to 4. What is the dimension? To figure it out, I'll do a calculation out of your sight (but detailed in the notes), and attest that according to our formula, the dimension is approximately 1.26. That is, the Koch curve is neither one- nor two-dimensional, but in between. Amazingly enough, fractal dimensions are not integers. That's what makes fractals so strange.

In short, the fractal dimension quantifies the number of copies of a self-similar object at each level of magnification of that object. Equivalently, fractal dimension quantifies how the total size (or area, or volume) of an object will change as the magnification level changes. For example, if you measure the total length of the Koch curve each time the rule is applied, you will find that each time the length has increased by 4/3. Only perfect fractals—those whose levels of magnification extend to infinity—have precise fractal dimension. For real-world finite fractal-like objects such as coastlines, we can measure only an approximate fractal dimension.

I have seen many attempts at intuitive descriptions of what fractal dimension means. For example, it has been said that fractal dimension represents the "roughness," "ruggedness," "jaggedness," or “complicatedness” of an object; an object’s degree of “fragmentation”; and how “dense the structure” of the object is. As an example, compare the coastline of Ireland (figure 7.2) with that of South Africa (figure 7.4). The former has higher fractal dimension than the latter.

One description I like a lot is the rather poetic notion that fractal dimension “quantifies the cascade of detail” in an object. That is, it quantifies how much detail you see at all scales as you dive deeper and deeper into the infinite cascade of self-similarity. For structures that aren’t fractals, such as a smooth round marble, if you keep looking at the structure with increasing magnification, eventually there is a level with no interesting details. Fractals, on the other hand, have interesting details at all levels, and fractal dimension in some sense quantifies how interesting that detail is as a function of how much magnification you have to do at each level to see it.

This is why people have been attracted to fractal dimension as a way of measuring complexity, and many scientists have applied this measure to real-world phenomena. However, ruggedness or cascades of detail are far from the only kind of complexity we would like to measure.

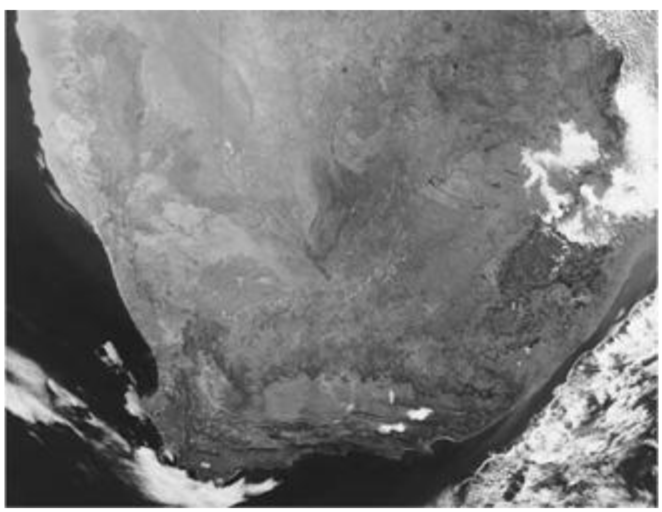

FIGURE 7.4. Coastline of South Africa. (Photograph from NASA Visible Earth [http://visibleearth.nasa.gov].)

Complexity as Degree of Hierarchy

In Herbert Simon's famous 1962 paper "The Architecture of Complexity" Simon proposed that the complexity of a system can be characterized in terms of its degree of *hierarchy*: "the complex system being composed of subsystems that, in turn, have their own subsystems, and so on." Simon was a distinguished political scientist, economist, and psychologist (among other things); in short, a brilliant polymath who probably deserves a chapter of his own in this book.

Simon proposed that the most important common attributes of complex systems are *hierarchy* and *near-decomposibility*. Simon lists a number of complex systems that are structured hierarchically—e.g., the body is composed of organs, which are in turn composed of cells, which are in turn composed of celluar subsystems, and so on. In a way, this notion is similar to fractals in the idea that there are self-similar patterns at all scales.

Near-decomposibility refers to the fact that, in hierarchical complex systems, there are many more strong interactions within a subsystem than between subsystems. As an example, each cell in a living organism has a metabolic network that consists of a huge number of interactions among substrates, many more than take place between two different cells.

Simon contends that evolution can design complex systems in nature only if they can be put together like building blocks—that is, only if they are hierachical and nearly decomposible; a cell can evolve and then become a building block for a higher-level organ, which itself can become a building block for an even higher-level organ, and so forth. Simon suggests that what the study of complex systems needs is "a theory of hierarchy."

Many others have explored the notion of hierarchy as a possible way to measure complexity. As one example, the evolutionary biologist Daniel McShea, who has long been trying to make sense of the notion that the complexity of organisms increases over evolutionary time, has proposed a hierarchy scale that can be used to measure the degree of hierarchy of biological organisms. McShea's scale is defined in terms of levels of *nestedness*: a higher-level entity contains as parts entities from the next lower level. McShea proposes the following biological example of nestedness:

> Level 1: *Prokaryotic* cells (the simplest cells, such as bacteria)
>
> Level 2: Aggregates of level 1 organisms, such as *eukaryotic* cells (more complex cells whose evolutionary ancestors originated from the fusion of prokaryotic cells)
>
> Level 3: Aggregates of level 2 organisms, namely all multicellular organisms
>
> Level 4: Aggregates of level 3 organisms, such as insect colonies and “colonial organisms” such as the Portuguese man o’ war.

Each level can be said to be more complex than the previous level, at least as far as nestedness goes. Of course, as McShea points out, nestedness only describes the *structure* of an organism, not any of its *functions*.

McShea used data both from fossils and modern organisms to show that the maximum hierarchy seen in organisms increases over evolutionary time. Thus this is one way in which complexity seems to have quantifiably increased with evolution, although measuring the degree of hierarchy in actual organisms can involve some subjectivity in determining what counts as a “part” or even a “level.”

There are many other measures of complexity that I don’t have space to cover here. Each of these measures captures something about our notion of complexity but all have both theoretical and practical limitations, and have so far rarely been useful for characterizing any real-world system. The diversity of measures that have been proposed indicates that the notions of complexity that we’re trying to get at have many different interacting dimensions and probably can’t be captured by a single measurement scale.