CONTENTS

1

Reductionism and Its Discontents

My son is fascinated by the possibility of an infinite regress of "whys," an unending chain of answers spawning questions requiring new answers, which in turn lead to their own questions. No matter the starting point (Why does the moon rise over a particular hill? Why do worms fare poorly in the sun?), I am generally stumped after a short sequence of ever more incomplete answers. Fortunately, I am fairly sure that my son is not overly troubled by his father's profound ignorance. I suspect that at times he is not very deeply invested in the answers, but he always seems to love the idea that each answer contains within it the kernel of the next question. Perhaps this is somehow built into the mechanisms of the mind, since the endless "why?" and "how?" appears to be a universal feature of childhood. Even if this love of infinite regress is not actually part of our universal human experience (although I suspect it is), it certainly is part of our shared scientific culture, particularly in the physical sciences. The chain of whys is, at its heart, the meaning of reductionism in science.

My field of physics has been described as the progressive exploration of this chain of whys. For at least a hundred years, my colleagues have pursued the dream of a final Theory of Everything, one that admits no more fundamental challenge to what the universe is made of and how those elemental pieces interact to form everything we see and everything invisible to us. The theory of everything is meant to be the final stop in the chain of whys.

Reductionism is powerful and quite effective, and the allure of the final step in the chain is profound and tantalizing. But today, a number of scientists are questioning the idea that reductionism leads to a complete understanding of nature. Before we consider the insufficiency of the reductionist approach, let us marvel, at least briefly, at where it can take us.

Consider the act of pouring salt. From a sufficient distance, one might imagine that the stream of salt flows like a fluid. Why does it flow? To answer that, we look more closely. Quickly we note that the stream of salt is not a continuous body but rather is composed of small, typically cubic blocks, each of which is off-white and translucent. What we first saw as a stream is really a flock of particles flying through the air, bouncing and tumbling off one another. We have learned something. We might be able to make a movie of these blocks falling, tumbling, and colliding. You might even imagine that a stream of water or any atomic or molecular fluid is, in essence, similar—a set of much smaller particles moving through empty space and colliding with one another.

If we were to study the movie of flowing salt even more carefully, we would find that it does not behave like a typical fluid of atoms or molecules. As students of "granular materials" can explain, the key distinction appears to be that the collisions between salt particles do not conserve energy. Two salt particles on a collision course at high velocity rebound more slowly after their collision—some of their energy of motion is lost. Where did it go? To answer this question, we must look still more closely. During each collision between the grains of salt, some of that energy of motion is transferred to internal vibrations of each grain, sound waves that echo throughout the crystals and then decay into random motion, which we call heat. Pouring salt in effect transfers gravitational potential energy (the work you did to lift the salt shaker) into the kinetic energy of the flowing grains, and finally into heat. The pile of salt on your kitchen counter is very slightly warmer than the salt that remains in the salt shaker.

But why are there sound waves in each grain of salt? To answer that, we plunge down into the crystal to find that it is a quite orderly array of sodium and chlorine atoms arranged in a lattice.

This lattice is springy since the atoms repel when pushed closer together and attract when pulled apart. From this picture, we can develop a theory of sound waves in a crystal. My colleagues and I call them phonons.[1] But the chain of whys does not stop there. What determines the forces between the atoms? To answer that question, we need to explore the structure of

1. Actually, a phonon is a quantum of sound energy, in analogy to the quantum of light energy, the photon.

atoms and confront quantum mechanics. But we are still not finished. In our studies of the atoms, we learn that most of their properties can be understood from the complex interactions of their electrons and their nuclei. But why do electrons have the properties that they do? They have a particular charge and mass, but why? Today, we do not know. Our chain of whys has hit a roadblock, and we are left to assert that there are certain "fundamental constants" of nature, including those related to the electron. These fundamental constants are part of what we now call our Standard Model for the interaction of elementary particles. About the nuclei, we have much more to say, and the chain of whys extends a few more hard-won links. The electrical charge (and spatial distribution of electrical charge) and mass of the nucleus can be understood in terms of collections of their constituent particles: protons and neutrons. The properties of those particles can be understood in terms of elementary particles, quarks, which interact by exchanging gluons and photons. The properties of these particles we know, but we do not know the why. Here that particular branch of the chain of whys extends into the deep fog that may someday be penetrated by string theory or its descendants.

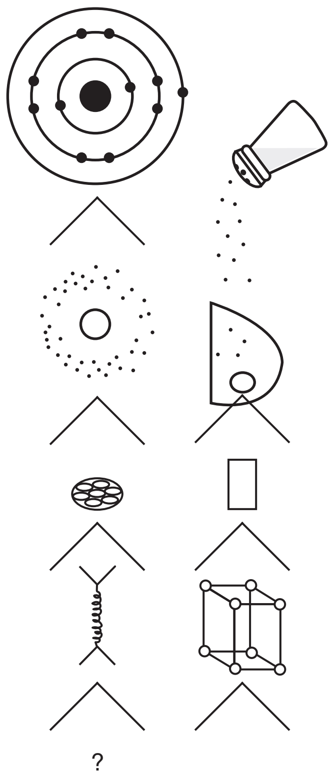

FIGURE 1.1 The chain of whys, from the salt flowing from a salt shaker through the atomic and nuclear structure of salt to the physics underlying fundamental processes in the universe.

Now let us look backward up the chain of whys to our small pile of salt, just to see how far we have come. There is no simple measure of that distance; no one number can fully report on the conceptual distinction between a flowing stream of salt, glinting in the light, and the underlying world of crystals, atoms, and elementary particles lying beneath. One imperfect but instructive measure is simply the change in scale. Through the mind's eye we have traveled across an immense chasm that separates what we experience directly with our senses to a nearly unimaginably small scale in the interior of the nucleus of a sodium atom, and finally down to the scale of a single proton.

This is not a mathematical book, but sometimes numbers matter; they matter the most when we address time or distance scales so far outside of our daily experience as to be inconceivable otherwise. And so, a few numbers are in order. A grain of salt might be about a few millimeters across, but the radius of one proton is about ten trillion times smaller. But this is simply too great a change to be meaningful to me. How can we conceptualize such a big number? To put it another way, imagine that our proton were the size of the original grain of salt, one millimeter across. To blow things up to this scale, the entire grain of salt would be about ten billion kilometers across. Again, we find ourselves at a loss for a sense of scale, but fortunately, our exploration of space has given us a few somewhat concrete notions of long distances. *Voyager One*, a space probe launched when I was in grade school, has been traveling outward toward interstellar space for most of my life. It is the fastest vehicle we have ever made. Right now, it is about ten billion kilometers from home. So, if we were to blow up each proton to the scale of an ordinary grain of salt, that single grain, which one can see easily on the tip of one's finger, would be so vast that, after traveling for forty years, the *Voyager* spacecraft would have just crossed one face of this massive single grain of salt. Looking back from one side of that stupendous grain to the other, the entire Earth would be reduced to a single pinprick of light. The original pile of salt would span the distance between the stars. We are truly prisoners of our senses, which are confined to a narrow range of accessible lengths and times. The chain of whys breaks the chains of our senses, freeing us to see the world that lies beneath.

> To see a world in a grain of sand
> And a heaven in a wild flower
> Hold infinity in the palm of your hand
> And eternity in an hour
>
> —WILLIAM BLAKE ("AUGURIES OF INNOCENCE")

But following the chain of whys is not entirely satisfying for at least a few reasons. The most obvious one is that our chain is actually a branching and rebranching tree. In order to tell the story as I did above, I pruned this tree down to a single branch. For example, why are the salt crystals typically cubic and what does that tell us about the underlying lattice of atoms making them up? Why are the crystals translucent and white? And what does that tell us about how light interacts with matter? These roads not taken end up leading us to some of the same endpoints as we reached earlier, those dealing with the Standard Model. And this is one of the marvels of reductionism. It appears that when we follow the tangled web of whys, we end up eventually on the same small set of threads. Those ends of the chain thus appear to be at the base of the entire network. In fact, all the paths we know seem to lead to the exact same threads. And this is the power of reductionism and the Standard Model. A large community of scholars is attempting to push along that chain deeper and deeper, but this is not our story here.

But you may wonder if being able to travel down the chain is the same as understanding the stream of salt. Another way to ask the question is: can one understand the details of pouring salt as a consequence of our knowledge of fundamental physics? In essence, we followed the chain down to the world of quarks, but can one follow it back up? Should one even want to do so? One reason we may not even want to try arises because nature appears to package up the details at one scale as a small set of numbers that are observable at the next scale. It is as if each link in the chain of whys opens a new set of actors, but their complex behavior informs the next level up through only a few constants.

Consider our sodium atom. To understand the behavior of the electrons in the sodium atom, we need to consider its nucleus. But, even very precise measures of the electronic structure of the atom rely on our knowing only charge and mass of the nucleus, two collective quantities selected out of all the countless ones describing its internal workings. All of the protons and neutrons in the nucleus, each of which is an internal sea of quarks and gluons, individually have a complex life of their own, and the interactions of these separate seas within the nucleus is, itself, complex. But from the point of view of the electrons of the atom, it is sufficient to know only the net charge and mass of the nucleus; the rest of the internal drama fades in importance. These remaining parameters may be hard to understand from the point of view of the Standard Model (and the mass certainly is), but nature appears to package all the complexity of the underlying scales into just a few numbers that might be considered "fundamental" at the next level up. To continue, the interaction between the atoms in the crystal can be described by a potential energy curve parameterized by, again, just a few numbers. Clearly, these numbers emerge from the complex interactions of the electrons and nuclei of the atoms. Once again, all the complexity at one scale blends into a few parameters necessary at the next one. When considering the sound waves in the crystal, the complexity of the potential energy of two atoms can be reduced by considering their interaction to be that of a (slightly nonlinear) spring. The lattice vibrations of the crystal are well understood in terms of a network of masses and springs. This seems rather prosaic. The sea of quarks and gluons is far below us, and very little of their complexity survives at the level of the crystal. Finally, at the level of the stream of salt, we might reasonably treat each grain as an inelastic solid, characterized again by just a couple of numbers: mass and a measure of how much energy is lost in a collision between grains. This final number, called the coefficient of restitution, is quite the Russian nesting doll. It contains features of lattice waves of salt crystals, which in turn depend on the interaction of their atoms, which in turn depends on the structure of those atoms, which in turn depends on the structure of their constituent parts. But the chain of whys is neatly packed up inside one number.

It is interesting to consider why our understanding of nature is packaged up so neatly. The answer appears to be that the various constituents we have discussed act at very different energy, length, and time scales. The work required to change the configuration of the electrons of an atom is about a million times smaller than the work required to change the configuration of the nucleons within its nucleus, which is about one hundred thousand times smaller than the atom itself. This is the reason that chemistry, acting on the scale of electronic configurations, cannot transmute the elements, which requires exchanging nucleons. In a rough sense, it is the underlying reason that chemists cannot turn lead into gold and why nuclear weapons are so much more devastating than chemical bombs. One can ignore intranuclear dynamics when considering the binding of two atoms chemically because those nucleons are under the influence of such large forces that mere electrostatic interactions with the electrons make no difference to them. Why this massive range of energy scales appears in nature is not immediately clear, but one of its implications is that our understanding of nature comes quite neatly packaged.

But if we carefully examine that Russian nesting doll called the coefficient of restitution, we see that it contains something else quite remarkable and wholly distinct. Specifically, it contains an idea we have not yet confronted: a transition from the dynamical—following individual particles or their collections—to the statistical. Underlying the coefficient of restitution is the transfer of highly organized energy within the crystal (sound waves) into the disorganizing jiggling of the atoms (heat). Not only must we understand the sound waves launched into the crystal during a collision between grains, but we must understand how that energy eventually becomes heat. How long does a grain vibrate after one collision? If that time is long compared to the typical time between collisions, we have a very different sort of problem to consider, since grains in that case would have the ability to remember their history of collisions. To address this question, we have to be able to understand how, for example, a bell stops ringing and where the energy went. This is an entry point to a new set of deep and complex questions regarding thermodynamics, statistical physics, and the approach of systems to equilibrium. We will take this up later in some detail. At the moment, it is sufficient to realize that we have papered over a very complex set of questions that biology forces us to consider in some detail.

Finally, and most important, we realize that there is a continuing chain of whys extending away from the Standard Model. Assuming that all of the above is understood, can one understand the dynamics of the original stream of salt flowing from the shaker and onto the table? Can we predict the shape of the stream, and the shape of the resulting pile of salt on the table? What we are really asking is whether we can predict where each and every grain of salt will end up and by what trajectory it gets there. At first blush, this might seem to be a ridiculous question, somewhat akin to asking if one can predict whether a flipped coin ends up on its head or tail. In fact it is much worse, in that we are asking about millions of coins crashing into one another and the table! One of the great lessons of condensed matter and statistical physics is that there are well-defined questions of this type that can be meaningfully answered. In fact, many of those answers rely on the inherent randomness of the microscopic processes involved. Rather than recoil in the face of complexity, one can, by asking the right questions, revel in it.

Still, you may conclude that confronting our ignorance in something so quotidian as pouring salt is disappointing. Thinking about this problem may have convinced you that the amazing accomplishments of recent physics with their promise of a theory of everything in the offing will not, and cannot, finish the story. I think so. But, the flip side to admitting that reductionism is not the full story is acknowledging the possibility of other stories that science can tell. This is both hopeful and exciting. It means that the frontiers of human knowledge are closer than we might think. Frontiers of physics can be found by examining the properties of large collections of matter whose collective behavior emerges from the elementary interactions of its constituents, but whose collective properties cannot be simply predicted from those elementary interactions.

The physics of condensed matter was born to address the behavior of complex assemblies of atoms whose collective behaviors, from melting to superconductivity, emerge from the interactions of atoms, and which require a new language to describe and account for these collective phenomena.

Biological physics today is emerging from condensed matter physics as a way to take these basic notions that complexity engenders and that demand their own level of physical description and apply them in another venue. And what an elaborate and elegant world this is! The world of living matter is highly complex, with organization on many length scales. Furthermore, it is driven into nonequilibrium states. Life is a form of nonequilibrium or active matter in a way that the atoms thermally jiggling in your teacup are not. This point requires some detailed thought; I will return to it later. For now, let me point out that the living world exists as a peculiar type of transient organization of atoms, which continually consumes energy and matter and produces entropy. It is a collection of atoms that doesn't simply exist but also performs a function, even if that function is sometimes only to persist and reproduce. Sometimes, living matter does other things too. It even, on occasion, writes a book.

A well-read person with an interest in physics may have been led to believe that the current frontiers of physics lie far from his or her front door. They lurk in evacuated tunnels of colliders, like at CERN, where elementary particles are accelerated to nearly the speed of light and smashed to test our understanding of matter on ever-shorter length scales. They lie in the faint light of galaxies and the microwave echoes of the big bang, allowing us to explore the size and shape of our universe. Or the frontier sleeps in new quantum phases of matter that exist only at temperatures lower than those found in interstellar space. These frontiers are truly far from home, literally or at least metaphorically. In this research, my colleagues address profound questions that are far removed from our daily experience when they examine nature on scales much smaller than atoms or as large as the universe. These puzzles remain separated from our daily experience by vast ranges of length scale or energy.

But the frontiers of our knowledge also lie right at your fingertips. The boundaries of our ignorance remain deeply woven within our own bodies. You and I (and all living beings) are, in many ways, walking, talking, swimming, flying enigmas. The living world contains examples of the most complex organizations of matter known and is the site of intricate dynamical processes that we have yet to sufficiently describe, much less understand. The story of biological physics is one approach to address these questions from the point of view that the living world is shaped by conformity with natural laws and, as a consequence, can be understood at a fundamental level by the application of those physical laws. Physics informs biology and can be used to elucidate its most basic properties.

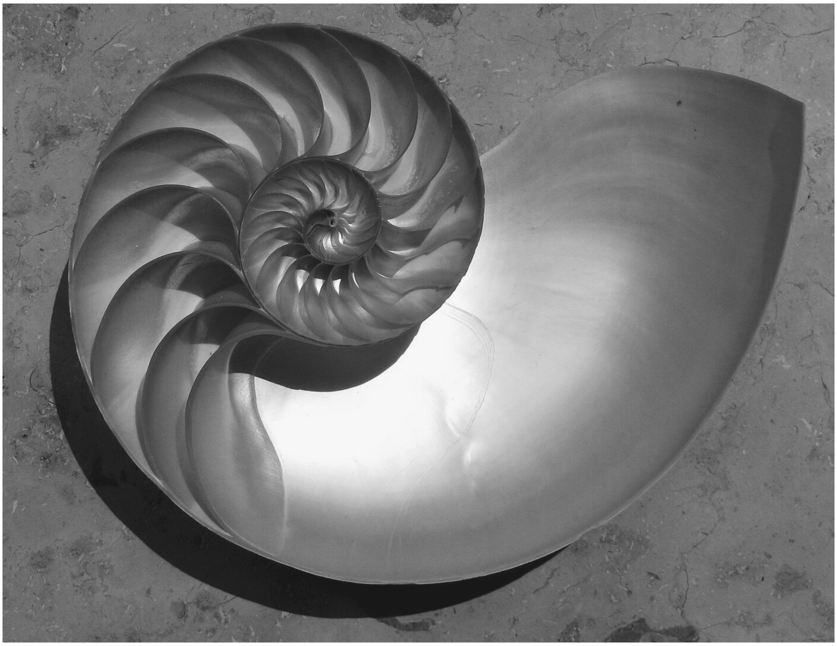

FIGURE 1.2 The chambered nautilus.

This is an idea as old as Galileo, and an unbroken thread can be traced following this idea from Galileo's day to ours. For instance, the nineteenth century saw brilliant insights, derived from mathematics and physics, into the structure of living things on length scales accessible to our unaided senses. One example is the elegant curve of the nautilus.

Due in large part to the fantastic advances of molecular biology in the past fifty years, we are now aware of complex structures of matter found in the living world on length scales ranging from single proteins up to collections of cells. Understanding the properties of these structures and their biological significance will surely involve the application of physics. As of old, physics will continue to inform biology.

But what I believe is emerging in our century is the understanding that the living world has much to teach us about physics. This is surely not the physics of super strings, quarks and gluons, or even the structure of atoms. It is not the physics that pertains to the large-scale structure of the universe, but rather a new type of condensed matter physics—the study of complex interacting systems of atoms, developed in the twentieth century and which led to our current understanding of the mechanical, optical, and electronic properties of solids and liquids.

One of the leading theorists in that field, Philip W. Anderson, wrote an essay titled "More Is Different." This manifesto proclaimed the intellectual independence of the field of condensed matter physics from what was considered to be more fundamental areas of physics research, and advanced the then revolutionary notion of emergent phenomena in nature. This is the idea that the behavior of billions of simple actors interacting by well-understood fundamental rules can develop complex and altogether surprising collective structures or dynamics. One concludes that these new structures and dynamical patterns are, in a real sense, not contained in any one of the elementary actors making them up. In addition, understanding these complex collective properties requires new ideas that operate at the scale of multitudes rather than at the level of the individual. Thus, the luster of gold is not contained in the nature of an individual gold atom, but in the sea of electrons moving through the lattice of gold atoms. The properties of ice are not built into individual water molecules, which remain unchanged in the transition from ice to water to steam, but rather in the properties of their organization on scales of trillions of molecules. The patterns of convection rolls in a teakettle brought to boil or in the plasma of the sun obey their own rules, which have a deeper commonality than the disparate nature of their constituents might suggest. More is truly different and condensed matter physics seeks its own principles to describe and predict the emergent properties of matter.

Anderson pointed out that the reductionist viewpoint does not imply what he called a "constructivist" one. In short, the ability to follow the chain of whys down to what we think of as the fundamental rules of nature does not imply that we can reasonably guess at the complexity of nature based solely on our understanding of those laws. A very fine understanding of the Standard Model does not appear to lead to a prediction for the viscosity of water, or the shape of the clouds, or the chance of rain tomorrow. Even if we have a theory of everything, we will still need a theory of everything else.

But what actually stops us from turning that theory of everything into a predictive theory of DNA replication or the flight of a curve ball? Significant insight into this question was achieved in the last century. From the time of Newton to the twentieth century, physicists were convinced that the laws of classical mechanics contained within them the future and that a sufficiently clever person (or, today, an electronic computer) could unfold that knowledge, in effect seeing the future with arbitrary precision. In the immediate aftermath of Newton's *Principia*, people believed that the rules by which the future unrolled from the present were known and that, with sufficient care, one versed in Newtonian mechanics could tell the future as one can tell the past. Like Anchises speaking to his son in the underworld, a Newtonian physicist could imagine that, with sufficient mathematical skill he or she could say, "Come, I will now explain what glory will pursue the children of Dardanus, what descendants await you of the Italian race." Today we know this is impossible, and we understand the reasons why.

Complex systems are typically chaotic and thus unpredictable in an essential way. Any arbitrarily small uncertainty in the precision with which we know the state of the system ends up being magnified by the subsequent dynamics. Eventually, one cannot use one's knowledge of that initial configuration of the system to make statements about its current state. Appreciating this point was hard won and fairly recent in our history. In the middle of the last century, John von Neumann was instrumental in building one of the world's first digital electronic computers. That computer was extremely good at doing a number of things, including calculating the details of the chain reaction occurring inside an exploding nuclear warhead (which provided at least part of the motivation for building those computers), but that computer and its descendants failed in predicting the weather more than a few days in advance. Von Neumann had hoped that, given a knowledge of fluid dynamics (which we have in principle), the computer could run a model of the Earth's weather forward in time and, in effect, predict the future, our future. This turned out to be impossible, because the atmosphere is fundamentally chaotic. Even knowing the rules in detail is not sufficient to apply them to the deep future, so no amount of measurement today will inform me about whether the Cubs will be rained out on a particular June day one year hence.

Does this unpredictability of complex systems destroy the chain of whys as we move up from simpler and more fundamental systems to more complex ones? The answer is rather nuanced because unpredictability closes some doors but actually opens other ones. For instance, it is a fool's errand to try to predict the trajectory of a particular molecule of oxygen in the room. In this, chaos closes one door. But it opens another. If I were to ask what fraction of the molecules in a cubic centimeter of air is made of oxygen, I can know that answer with a great deal of precision as long as I assert that the air is in equilibrium, in other words, as long as I can rely on molecular chaos to erase all the details of its past. The distinction between these two types of questions is that the former requires me to follow a particular molecule in its various collisions with other molecules, while the latter asks me about the statistical properties of a large number of such molecules. The beauty of the latter type of question is that it relies precisely on the veil of ignorance imposed by molecular chaos: the only reason I can make a definite statistical statement about a parcel of air is that its history is essentially irrelevant. The converse of the fact that I cannot follow the molecules in their courses, is that each parcel of air very quickly forgets its history. All of the dynamical variables telling me where each and every molecule is now and where it used to be become irrelevant in a sense that can be made very precise, at least when you ask the right questions. Understanding what questions one can ask and answer is really the story of equilibrium statistical physics and thermodynamics. This point allows us to assert that although the laws of nature at the microscopic level should allow us to (very nearly) see the future emerge deterministically from the past, this is not so for sufficiently complex systems. More about that later, but for now it is enough to observe that chaos prevents us from running the world of many interacting particles forward (by computer or otherwise) and blindly hoping for insight.

One of the insights of thermodynamics is that, out of all the trillions of variables that describe a small parcel of air (including the position and velocity of each molecule, and also perhaps the electronic and vibrational excitations of those molecules, and the states of all their nuclei, and so on and so forth), only a paltry few are necessary to follow—at least if one is willing to restrict oneself to the right set of questions. These variables are distinguished by the fact that they are history-independent. In this case, we know them to be pressure, temperature, volume, and mass; what is interesting about them is that they are either collective properties of all the molecules involved or properties of the statistical distribution of their positions and velocities.

This is another, and perhaps the most extreme, form of packaging of variables at one level into a much smaller set of variables at the next level. To move up to the next level in the chain of whys, we have jettisoned billions of microscopic variables and introduced a new and very small set of ones that had no meaning at the lower level of description. This process leaves the chain of whys intact, but breaks the notion of predictability. In other words, we assert that not only do we need new ways to think about the many atoms and molecules and new quantities to describe the phenomena we see, but also we must give up the notion that a sufficiently complex calculation or computation can meaningfully describe the world starting from the elementary rules of nature.

"More is different" means something precise. When Philip Anderson proclaimed that more is different (first in a speech at UC San Diego and later in an editorial in *Science* magazine), he meant something rather precise. He was, in effect, standing on a new hard-won hilltop after the battle was concluded, and he was in a position to comment on the new vistas obtained. In short, he was showing that condensed matter systems—like the salt crystal and more complex ones, too—obey a set of rules that is not evident in the elementary interactions of their constituents, but becomes clear in their collective properties in the limit of large systems. These concepts become mathematically precise in the limit of infinite systems, but complex matter composed of trillions and trillions of atoms are a good enough approximation to that infinity for these notions to be meaningful (and actually essential) for our everyday world.

The vista accessed by condensed matter physics allowed us to observe the phenomenon of spontaneously broken symmetry in action. In short, even though the underlying rules of the game—the fundamental laws of nature—may treat two or more situations identically in that those rules do not favor one situation over the other, we might nevertheless observe only one of these in nature. How can fundamental rules that do not discriminate between two outcomes lead to collective effects in which one of those outcomes is preferred?

Consider an example that occurs in condensed matter physics and may be the best-studied case, although there are many others. We might imagine the inside of a magnet as a lattice of atoms, each of which is a tiny bar magnet linked to the lattice, but free to flip so that the north pole of those tiny magnets points in one of two directions (think “north” or “south” or, in the preferred jargon of condensed matter physics, either “up” or “down”). The fundamental laws controlling the interaction of these magnets are such that there is no basic distinction between up and down. The only thing that matters to the magnets is how they are oriented relative to their neighbors.

This lattice of tiny (atomic) magnets admits a phase transition between two distinct phases. In one, the phase at high temperatures, the tiny magnets are found randomly distributed between ups and downs. The net magnetic field of the entire lattice of magnets cancels out so that there is no permanent magnetic field associated with the crystal. We call this the paramagnetic phase of the system, and it is associated with the magnets at sufficiently high temperatures where random thermal jostling disorders the orientations of the atomic magnets.

As one cools this magnet down, one reaches the Curie temperature (named after Pierre Curie, husband of Marie Curie), at which point something amazing happens. The strength of the interaction of the tiny magnets becomes strong enough that they do not flip up and down randomly and independently anymore; rather, larger and larger groups of nearby magnets

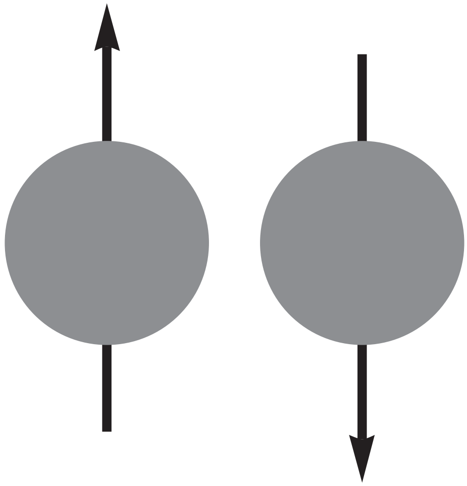

FIGURE 1.3 A lattice of magnets that can individually point in one of two opposite directions, "up" or "down." The up spins can be depicted by arrows pointing upward, or + signs, and the down spins by arrows pointing downward, or – signs.

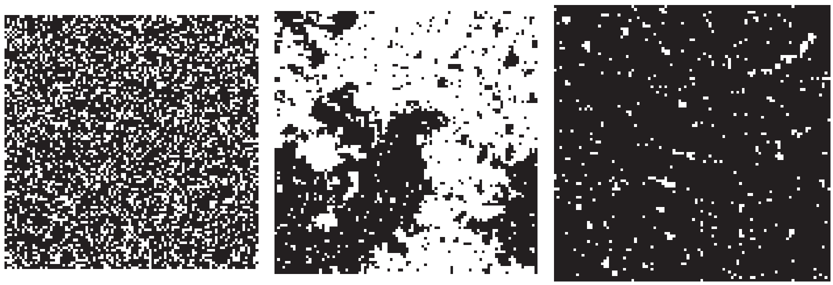

FIGURE 1.4 Characteristic configurations of the spins in a large lattice, black where they point up and white where they point down. *Left*, above the Curie temperature, where they are disordered and point randomly up or down. *Right*, below the Curie temperature, where one type of spin (up, in this case) dominates. *Center*, at the Curie temperature, where the number of up and down spins is roughly equal, but the spins have segregated into large up or down domains.

change direction together. The random collection of up or down magnets spontaneously self-organizes into larger and larger islands of regions where the magnets are all pointed the same way. Finally, as one passes through the Curie temperature, each region picks one direction for the magnets in it, either up or down. We now have a ferromagnet.

We also have spontaneously broken a symmetry. The rules of magnets do not care if one were to flip all the magnets from up to down. Another way to say this is that you cannot change the rules of nature simply by standing on your head to redefine which way is up! The magnet, however, has made an irrevocable decision about up or down when it was cooled through the Curie temperature from its disordered paramagnetic phase to its ordered ferromagnetic one. Today, we say that the magnet broke the symmetry between up and down spontaneously in making that one decision.

This simple model allows us to see precisely how many-body systems can be more complex than the underlying rules that make them up. The rules for the interactions of magnets are reasonably simple (but not necessary to review here) and complete. But the rules of the magnet in its ordered state require one extra rule telling us if north pole up or north pole down is the preferred state of the system. This basic picture can occur over and over again in more and more complex guises and serves as a mathematically exact way to see how complexity can create its own rules. In this process we also see how the packaging takes place at least in a restricted sense. Note that at no time did we have to discuss exactly how strongly the tiny atomic magnets interact with one another. In fact, all of those interactions (which in the end rely on some complex atomic physics) end up producing just one number, the Curie temperature. For iron, that temperature is about one thousand degrees Celsius, but that one temperature packages up a lot of detailed information about how the tiny iron atom magnets interact with one another. That information relies on the details of the tiny magnetic properties of individual electrons . . . and so on. We can open the Russian nesting dolls if we choose!

As an aside, consider the situation of tiny scientists living inside our magnet. Presumably, at some point these creatures learn the rules of magnets and discover that these rules make no distinction between up and down. Their fellow creatures look around the universe (their magnet) and see that it clearly makes a distinction between up and down. A clever creature might then postulate that they live in a broken-symmetry state of the universe. This is essentially what our colleagues in high-energy physics do today when they observe that the fundamental rules (the Standard Model) admit symmetries that the universe evidently does not enjoy. Strictly by analogy, we assume that the universe is in a phase (has ordered) in a way that has broken its symmetries.

We will return to this rather abstract notion of broken symmetry and our discussion of an enlarged meaning of "more is different" in later chapters. Understanding these ideas more fully requires us to pull them from the level of abstraction and apply them to particular instances as we consider a number of biological systems.

Before we go on, though, it is useful to observe that not only do statistical properties of a magnet obey the "more is different" principle but also the dynamical properties of complex systems obey rules of their own that emerge from the chaos of microscopic interactions. Consider how different air and water are. Certainly, if you are apt to confuse them, I would not want to go diving with you! But the rules by which air and water flow are essentially the same. The Navier-Stokes equation was known long before we had a good microscopic understanding of their constituents, or air and water, or any other fluid. This equation describes the flow of all sorts of fluids, from water to wine, from honey to alcohol, and from air to liquid nitrogen. Of course, fluids are fluids and we recognize them as such specifically because we expect there to be certain commonalities in their behavior. But it is interesting to pause for a moment to appreciate how remarkable this is. After all, each of these fluids is composed of very different objects, mixtures of small molecules or larger ones. If one could be shrunk to the size of these molecules, one could easily imagine that colliding sugar molecules in honey at room temperature are really quite different from diatomic nitrogen molecules at minus two hundred degrees Celsius. However, the Navier-Stokes equation describes the flow of each essentially incompressible fluid by just two numbers: the mass density (the mass per unit volume of the liquid) and the viscosity, a number describing how hard it is to shear the fluid. Honey behaves differently from water in degree, not kind. Its viscosity is a few thousand times higher. That is all.

We now have an example of a new type of Russian nesting dolls for dynamics. All the interesting molecular degrees of freedom, and the things that make honey sweet and alcohol intoxicating, are still there but do not affect the properties of the flow of these fluids outside of the control over two numbers. Physicists have given a great deal of thought to these ideas and learned how to identify how many numbers survive the transition from one level to the next and how these numbers emerge from the microscopic details. Ideas like conservation laws (the conservation of mass demands that the next level know about the mass density) and broken symmetry play a dominant role in the survival or emergence of new parameters at the next level and determine how these depend on properties of the lower ones. But, once again, the packaging is remarkably effective.

The honey in your tea is composed of quarks, but only a pedant would assert that those quarks make it sweet. Sweetness and viscosity and the ability to dissolve in water are higher-level abstractions that emerge on length scales, time scales, and energy scales so far from the world of quarks as to render them irrelevant to the discussion. The chain of whys may always descend to the quarks and beyond, but the packaging of nature demands new concepts at new levels of scale. Maybe this is disturbing in some primal sense, as it suggests that our understanding of the world at different scales will always require new language and that the theory of everything, when eventually found, will not provide all insight for all things. But this view, that reductionism does not imply constructionism, also seems eminently reasonable: If I sprain my knee, my doctor should not need to locate all my quarks and electrons to diagnose the problem. There is an appropriate scale of description for knees that does not admit quarks and gluons.

The air is rather thin at these ontological heights. Where does biological physics come into the story? The answer, I believe, is found in the continuing enterprise of condensed matter physics. The past century has led to a remarkably strong understanding of the collective properties of solids and liquids based on some of the ideas mentioned already. Along with this understanding came the realization that complex systems composed of many simple actors admit a simpler description of their behavior when viewed through the lens of the right variables, which relies on the inherent packaging of nature at finer scales into this smaller set of relevant variables, operating on the larger scale. “More is different” means that there is a different, but equally valuable, description of nature at the scale of complex interacting systems. This is the story of condensed matter physics in a nutshell, but it is only the jumping-off point for our new story of biological physics.

Is life different too?—I want to address a question related to Anderson’s. If more is different, and understanding that difference is the key to condensed matter physics, then is life different too? As a consequence of that difference, is there new physics to be learned from examining the living world, just as examining condensed matter in aggregate taught us fundamentally new ideas about how to conceptualize complex systems? Anderson’s essay recounted a battle won; it surveyed the new vista laid open by the hard work and new insights from a half century of condensed matter physics. It was a sort of victory lap. This book is deeply provisional. Here I propose that life might, in fact, be sufficiently different to drive my colleagues and me to find new physical principles that apply to complex systems maintained in a complex nonequilibrium state. I cannot take that victory lap; rather, I can share with you our thinking as to why there is a new vista to be seen, and one that will not be simply new biology or even just new physics, although that is exciting in itself. What I think might emerge from this work is something distinct from both biology and physics and, I suspect, quite surprising. Of course, I cannot organize and lay out the details of the surprise. I can, however, share with you why I think there is a new vista to be seen and, in the coming chapters, outline some of the things we have learned, using them as examples of the sorts of investigations that may take us to that new outlook.

A bacterium is a tiny object on the scale of things we normally manipulate. And I would venture that it is much more complex than any current human-designed object, although to be fair, it is hard to quantify and discuss meaningfully varying degrees of complexity. In spite of its tiny size and unusual complexity, it is at some level a massive collection of atoms of actually only a few types. Roughly speaking, one small bacterium contains about ten billion atoms. It is clearly the size of things where our rules of statistical and condensed matter should apply, but no one really expects this to be the case. What features of the bacterium make it fundamentally different from the things we understand quite well how to think about?

First, the bacterium has complex structures on many different levels of length scale ranging from the atomic up to its overall dimensions, roughly a millionth of a meter in length. These structures include various lipid membranes, bilayer sheets of molecules with a well-controlled thickness of tens of atoms (a few nanometers) but extending laterally for thousands of times this thickness. This membrane is dotted with complex proteins that act as gatekeepers, allowing water in or out of the cell and checking the identity of various ions before allowing them to pass. Our bacterium's more complex cousins (eukaryotes, like us) have a plethora of internal membranes in a bewildering variety of shapes and sizes, but this bacterium does not. Instead, inside this membrane is a soup of proteins, organized structures of a few thousand atoms each, that interact with one another to perform a type chemical computation—to run an algorithm deciding, for instance, whether to create new proteins to digest a particular class of sugar molecules, or whether to turn on a small rotary motor attached to a long tail in order to swim through the surrounding fluid. That tail also has a complex internal structure, based on an orderly arrangement of filamentous aggregates of other proteins.

The key point is that a simple crystal has but a few important length scales: the scale of the atoms at the small end and then perhaps the scale of the distance between defects in the packing of those atoms, and finally the scale of the entire crystal. Usually, the final scale of the entire crystal can for all intents and purposes be considered to be infinite. Everywhere inside the crystal is, in a sense, the same as everywhere else. The bacterium has a definite structure on scales ranging from atoms to those just visible in a light microscope.

This range of structural scales is perhaps unnerving for the physicist, but the real problem presented by the bacterium is that it manages to violate the nice packaging of information that we relied on in studying complex aggregates of matter in the nonliving world. There are at least two ways in which this happens. The first is due to the fact that the complex structures of the cell are not spaced out over a broad enough range of lengths and energies to keep the smaller Russian nesting dolls closed when looking at the larger ones. The reason that chemistry does not transmute elements is that the energy scale of the nucleus is at least a million times higher than that of the electrons that do the work of forming chemical bonds. The contents of the nucleus are locked away so tightly that only a very restricted bit of information about those contents (mostly just the total charge) can escape to influence chemistry. But in networks of filamentous protein aggregates acted on by molecular motors, the forces acting on the network on scales the size of the cell can influence the binding and activity of the motors at the scale of single (rather large) proteins. The fact of the matter is that the structures within the cell are not sufficiently separated in length or energy scales to render the packaging of smaller-scale features (e.g., molecular motor dynamics) impervious to influences from the larger-scale stresses and strains on the force-bearing network within the cell. Similarly, membrane tension can affect the opening and closing of channel proteins in the membrane. Biological structure violates the pristine separation of length scales on which a physical scientist has always relied.

There is a second way in which biology subverts the packaging of nature—through an active process. In fact, this sort of violation is, in some sense, written in the basic operating system of biology. Let us first recall a counter example from the inanimate world: if we were to study a fluid of honey flowing out of a pot, we would not need to concern ourselves with all of the molecular details of the honey, even if we were to concern ourselves with calculating its viscosity. For instance, the honey from orange blossoms and sage blossoms tastes different, so there must be some chemical difference, but I suspect there is no change in viscosity. This is due to the fact that taste relies on a very small concentration of impurities in the honey. The viscosity of that fluid is the result of an average over the dynamics of many, many molecules so the very small dynamical effects of these impurities are swamped by the averaging process. Biology, however, is built upon a method of amplifying very small chemical changes. Consider our bacterium again. A very small chemical change to one molecule out of the billions of molecules within it, its DNA, leads to very different chemical and physical properties, ones that literally spell the difference between life and death. To make one molecule matter, the cell has to work very hard to amplify the effect of a very small chemical change through a massively costly process of reading the DNA and producing proteins, some of which act back on the DNA. We will discuss this beautiful process at the heart of biology in more detail later, but for now we must confront the fact that the cell works very hard to amplify certain very small (atomic-scale) changes into having systemwide effects. The Russian nesting dolls of all sizes are running loose, and one of our best tools for organizing nature is actively broken by the living world.

Being actively broken brings up a whole new set of intriguing issues. One of our main tools for understanding the collective behavior of large numbers of atoms is equilibrium statistical physics and thermodynamics. These approaches are based on the fundamental principle that the system in question exchanges nothing, or only a very few things, with its surrounding environment. For instance, a cup of tea in thermal equilibrium at some temperature might be considered to be allowed to exchange energy in the form of heat with its environment in order to maintain its temperature. We have our tea coupled to a thermostat. It might also be allowed to exchange volume with the environment, to expand or contract inside the cup so that it remains at the same pressure as the surrounding environment. But that is all. Under these conditions, we know how the tea (or anything else) behaves once its system, on average, is not changing anymore. In such a state, we can, in principle, compute just about any statistical property of this many-atom system.

The living world remains out of equilibrium. The lifetime of a cell or a person is one long journey of maintaining a nonequilibrium state, characterized by continual input and output of mass and energy. Most important, it is characterized by the long-term increase in the total entropy of the universe. We will discuss this again in more detail, but for now we note that biology creates (or, perhaps more simply, *is*) a particular type of highly controlled nonequilibrium but essentially time-independent steady state. You and I are dynamical patterns of matter characterized by input of energy and mass, and the creation of entropy in the outside world. We are not our atoms, since those change on various time scales; we are really dynamical patterns written in atoms that persist, at least for a while.

There is a fairly long tradition of the theoretical study of nonequilibrium states of matter, from convection cells in a heated pot (or in a star) to crack propagation in a crystal, to driven magnetic flux lines in a superconductor. The equivalent of the basic principles by which we understand equilibrium systems have proven elusive, and the experiments on well-defined and persistent steady states out of equilibrium are typically difficult. After all, patience and isolation are generally sufficient to produce equilibrium states, but carefully and reproducibly driving systems out of equilibrium is more complex. Biology does this as a matter of course. It challenges us to understand how to think about nonequilibrium states of complex systems in a new way.

Finally, there is a less precisely defined but no less important distinction between the living and nonliving world. The internal workings of our bacterium have a purpose in a way that atoms do not. I do not mean a higher purpose of becoming complex, or interesting, or us. Rather, the mechanisms of the cell have been weeded out by evolution to get the job done of maintaining a particular nonequilibrium state that is capable, at least, of incorporating matter and energy from its surroundings and making copies of itself. Bacteria in a dish, when fed, grow. They *do* something. A dish of sodium atoms simply *is*. This is the point at which physics meets history, in that the structures we see in the cell are contingent on the evolutionary story of the creation and development of life.

Another way to think about this is that, given the building blocks of atoms, a reasonable person could postulate that when these building blocks occur at sufficiently high density and low temperature, atoms will form. That person should even be able to predict their properties. The same reasonable scientist, however, would be very hard pressed to postulate that atoms can then collect in larger structures to make hermit crabs, sycamore trees, and the Los Angeles Dodgers. In fact, one should well expect that if one were to start the universe over again, none of those interesting but historically contingent structures would come into being. Atoms, on the other hand, certainly would.

Because of this reasoning, the physicist is left in a bit of a quandary. Perhaps not all features of the living world are historically contingent, but clearly some of them are. It is not clear how we can tell the difference. Even more ambitiously, one might ask whether the evolutionary accretion of more and more complex structures and dynamics by biology—the transition from learning to metabolize fructose to learning the violin—follows its own understandable rules. Predicting the writing of the song "Brown-Eyed Girl" is certainly impossible, as is even predicting the formation of brown eyes, but maybe evolving the ability to detect light is built into the rules of forming living matter, at least in a place bathed in solar radiation. Of course, being able to address such questions has broad implications for understanding our place in the universe and perhaps whether the evolution of life is miraculous, statistically rare, or actually foreordained by having a collection of atoms at sufficient density and provided with an energy source.

In the rest of this book, I will address some of these ideas in more detail by examining a few systems of current study in biological physics. My goal is both to illustrate where new physics may be found and to acquaint you with some of the fascinating things my colleagues and I are learning about the living world from an application of new experimental techniques and theoretical models.

INDEX

Page numbers in italics indicate black-and-white illustrations; those in bold italics indicate color plates.