then why couldn't the new watch be a little bit better than the old one? Interestingly, Paley introduces chance into the argument when he says that the new watch is similar, but not identical. How similar? What determines what is the same and what is different in the offspring? What if there were millions of such watches, reproducing and exchanging information about their construction, passing every improvement to the next generation—would these watches not improve over time?

Before you object that watches don't have babies, let me remind you that I am simply following Paley's argument. Obviously, a watch that could reproduce itself would stop being a watch. A watch is an artifact made for an *external* purpose—to tell the time. Once the watch starts reproducing itself, it acquires an *internal* purpose—efficient reproduction. If we still had an external agent who selects watches for how well they tell time and only lets good timepieces reproduce, they would, over the generations, become better and better watches. But in the absence of an external agent, the watches would stop just being watches, because efficient reproduction would become their new raison d'être. After some time, they would radiate into many different machines—based on which could reproduce the best.

Now let's go one step further. Nothing works in the absence of energy. A watch needs to be wound up to work. A living organism must eat. Thus, reproducing watches would need to take in energy. They would need to find ways to beat competing watches in the race to sequester enough energy to reproduce. This would require them to find new ways of making a living. Biologists call these roles *niches*. Before long, we would not recognize our watches anymore. Few would still tell time. Their wheels would now be used for digestion or locomotion. The hands and the dial would be used for attracting a suitable partner with which to exchange information. Maybe, a glow-in-the-dark hour hand would drive the opposite sex wild. This is evolution—this is life.

Evolution

How do molecules evolve? Despite the histrionic debates in various American school boards, the mechanism of evolution is, as we saw in Chapter 1, quite obvious when contemplated with an open mind. This observation

prompted Darwin's supporter Thomas Huxley to lament his not having thought of it first. Molecules are subjected to the same natural selection that applies to the more macroscopic parts of an organism. As a matter of fact, the evolution of proteins is a good way to see how evolution works, because there is a direct relationship between the protein's amino acid sequence and the information encoding this sequence in DNA. Every molecular innovation—every new molecular machine that transports cargo a bit faster or makes fewer mistakes when transcribing DNA—will give an advantage to the organism it inhabits. Consequently, better molecular machinery will become more prevalent in a population. Or when conditions change, new machinery will emerge to deal with the changed conditions.

A famous example of the evolution of proteins was the 1975 discovery of a strain of *Flavobacterium* in a wastewater pond at a Japanese nylon factory. The bacteria in this pond had evolved to eat chemicals associated with nylon manufacturing—chemicals that do not exist in nature. On further investigation, researchers isolated three enzymes that had evolved inside these bacteria and that helped the bacteria break down nylon. None of these enzymes existed in *Flavobacterium* strains that were not raised in the nylon pond. How did the bacteria invent the new enzymes? Bacteria multiply very fast and exist in large numbers. Therefore, they can evolve very rapidly. In this case, the DNA replication machinery of a few *Flavobacterium* cells apparently made a mistake. The machinery read off a DNA sequence from the wrong starting point, leading to a so-called frame-shift mutation. It so happened that the resulting protein was helpful in breaking down nylon, which is helpful when you live in a pond full of the stuff. Is such serendipity really believable? Absolutely. Just consider that a human body contains 10^{14} (a hundred thousand billion) bacteria. The Japanese pond must have contained much more than that. A typical time for bacteria to multiply is twenty to sixty minutes. Assuming the slower time and assuming that the nylon factory was in production ten years, the bacteria would have gone through 87,600 generations of gazillions of bacteria. Considering this enormous number of bacteria and the many generations they pass through, a rather unlikely mutation now moves into the realm of definite possibilities. But we are not done yet: The first enzyme may not have been good at digesting nylon, but a bad nylon-digesting enzyme was

certainly better than none at all. Once the bad nylon-digesting enzyme spread through the bacterial population, it evolved and improved rapidly.

The glacial, step-by-step, and somewhat unpredictable process of evolution makes it difficult for people to believe that this mechanism could have led to the sophisticated machinery of our cells. But as the above example shows, sometimes evolution happens in a few years. Remarkably, the lion's share of the history of life (almost three-quarters of life's history, or three billion years) consisted of the evolution of single-celled organisms. Multicellular organisms only appeared in the last billion years. Why did it take so long for multicellular life to appear? When we look at the complicated machinery of our cells, an answer suggests itself: It took billions of years of evolution to turn the first primitive enzymes into our modern sophisticated cellular machinery. Multicellular organisms became possible only when a minimum degree of efficiency and sophistication was reached. This view of life's early evolution is supported by the observation that on a fundamental level, all multicellular animals (and all plants) are the same. Humbling as it may sound, at the nanometer scale little distinguishes a human from a fungus. The basic cellular toolkit is the same. The complexity of this kit justifies the length of time it took for it to develop. Once the toolkit was in place, evolution was free to create ever more amazing multicellular creatures, from octopi to redwood forests. In some sense, the real mystery of life lies at the molecular scale. This is where all the real work of evolution was done. The rest is icing on the cake.

The fossil ancestors of our molecular machines are, for the most part, gone forever. Proteins do not keep for over three billion years, and bacteria with primitive machinery would have been eaten a long time ago. Even the so-called archaea microorganism, which have been found to be significantly different from bacteria, are not really archaic. In some sense, bacteria and archaea are more evolved than we are. After all, they had a lot more time, and they reproduced much faster. With this in mind, is there *anything* that can be done to determine how molecular machines may have evolved?

Trying to figure out the exact evolutionary steps leading to the ribosome or a kinesin molecular motor is like trying to solve a crime hundreds of years after it happened. Who was Jack the Ripper? It is impossible to tell. The trail has gone cold. Yet, using the few reports and other scant ev-

idence that remains, we can make some plausible arguments about what kind of person he may have been. In much the same way, when it comes to the evolution of molecular machines, we have to look at the few remaining clues and try to come up with a plausible story. In this case, *plausible* means that the story matches the evidence and is in accordance with known physics and chemistry. Once a plausible story has been hypothesized, parts of it can be tested in the laboratory. If it passes these tests as well, we end up with a *likely* story, but we will never get a proven story. Jack will remain at large.

An instructive example is the evolution of the ribosome. In Chapter 7, I suggested that RNA is believed to be a more ancient molecule than either DNA or proteins. In the chicken-and-egg problem of what came first—DNA, which encodes protein, or proteins, which are needed to read the DNA—the answer is clearly neither. RNA contains information and can catalyze reactions. It is a kind of egg on feet, which can lay its own eggs. No chicken needed.

RNA's ability to catalyze reactions is a fairly recent discovery. In 1982, Tom Cech and coworkers at the University of Colorado–Boulder discovered that a certain RNA strand in a bacterium was able to splice parts of itself and reconnect the RNA strand, without any protein-based enzymes. This was the first indication that RNA could act as an enzyme. It took another ten years before Harry Noller and his group at the University of California–Santa Cruz demonstrated that the RNA in the ribosome also has catalytic properties. Over the years, it became clearer that *all* the hard work in the ribosome is done by RNA. When researchers removed the protein components, the RNA was still able to process messenger RNA and produce an amino acid chain, although at some loss to efficiency and fidelity.

The finding that the ribosome needed its RNA, but not its proteins, suggested that catalytic RNA may have been the basic constituent of early life. In a paper in 2010, researchers from the Weizman Institute of Science in Rehovot, Israel, and the European Molecular Biology Laboratory in Heidelberg, Germany, suggested that the RNA pocket where the amino acid chain is assembled is universal in all ribosomes and may constitute the original ribosome precursor. If RNA really was first, and it could catalyze its own evolution through splicing and reshaping, it may have eventually hit on a structure that could produce proteins. After that, proteins that

assisted the primitive organism by being better catalysts than the RNA could have formed. This first protein-producing RNA may not have been able to create well-controlled protein products, but less control would have also led to more mutations and possibly faster evolution. Eventually, the combined forces of RNA and proteins invented DNA, and the modern cell was on its way. It is a likely story.

Another way to consider the evolution of molecular machines is to look at family relationships. In the previous chapter, we mentioned that kinesin and myosin share many similarities in their motor domain. This suggests that they may have evolved from a common ancestral protein. Myosin and kinesin share certain loops in their switch domains, which are associated with shape changes upon binding of ATP. Intriguingly, they also share these structures with so-called G proteins, which are not machines, but molecular switches. Molecular switches communicate chemical signals from the outside of the cell to the inside. In Chapter 6, we speculated that molecular motors may have evolved from enzymes that could change their shape when they bound a control molecule. This allosteric effect is what makes molecular switches work.

The details of how G proteins connect to kinesin and myosin is lost in the fog of billions of years of evolution. Nevertheless, that motor proteins would have evolved from molecular switches is very plausible, and the relationship with G proteins confirms this idea. The relationship between kinesin, myosin, and G proteins shows that in evolution, similar parts in different molecules can often serve different purposes. The evolution of a sophisticated machine like kinesin does not require that each part be invented from scratch or that all parts come into existence simultaneously. When it comes to evolution, almost anything goes.

Let us imagine two proteins A and B, encoded by certain genes in our DNA. Proteins A and B perform different functions. What if part of protein B could help make protein A work better, or what if the combination of parts from A and B were to create a new protein with a completely new function? No problem. Sometimes, whole protein sequences are translocated in our genome, either through copying errors or by viruses. This can lead to the combination of different proteins and the creation of an entirely new line of molecular machines. An example is the nylon-eating enzyme of *Flavobacterium*.

Ever since kinesin and myosin came into existence eons ago, they have evolved into many different forms, forming large superfamilies. We have seen an example for such a family tree in Chapter 7. The same is true for almost any molecular machine. Every protein is part of a family of related proteins whose jobs are often quite different, but which are clearly descendants of a common ancestor. Evolution never ends; it is ongoing. Once evolution discovers a new trick, such as a walking molecular motor, it soon creates many variants, all fulfilling specialized functions.

A common objection of creationists is that some biological structures are "irreducibly complex." What they mean is that a structure has many interdependent parts, so that if you remove just one, the whole thing could not work. For example, how could a car evolve? The engine could not evolve without already having a whole car in place. But the car could not evolve without an engine. All the parts of a car must be designed to fit together. No part can be left out. Thus, goes the argument, molecular machines must be *designed*, just as a car is designed. This is because (following their argument), a molecular machine is only functional when all the parts are in place. There can be no intermediate evolutionary steps. Every previous version of the molecular machine—without all the necessary parts in place—would have been utterly useless.

There are a number of problems with this superficially persuasive idea. First, as we have seen, structures are often put together from parts that previously served a completely different purpose. Take the car example. Clearly, different parts of the car can be developed independently of the whole car. An engine can drive a stationary machine. The Cardan shaft of today's automobiles, as we saw in Chapter 2, was invented for a water pump. Pistons come from air pumps. Gears from watches. And so on. There are countless examples of such versatility in evolution. The bones in our middle ear evolved from a jaw bone of an ancestral amphibian. The evolution of motor proteins from molecular switches, like G protein, is another example.

The second problem with the irreducibility argument is that incomplete structures are not as useless as one might think. Take the eye. Is an eye without a lens really useless? It sure beats no eye at all. Almost all intermediate stages of eyes exist in nature, from mere light-sensitive spots of some microorganisms to the sophisticated eyes of mammals. The same

applies to motor proteins. We have seen that two-legged motor proteins can be processive. However, one-legged motor proteins can work as well, although with much less efficiency and highly reduced processivity, using a pure Brownian ratchet mechanism. Indeed, as we have mentioned, there are one-legged kinesins—although the jury is still out as to whether they can pair to form a two-legged kinesin. Nevertheless, at least theoretically, there is no physical reason why such a one-legged kinesin would not work. It may not be as good as kinesin-1, but it's better than no kinesin at all.

Many biologists consider evolutionary changes of DNA the most important events in the history of life. One proponent of this view is evolutionary biologist Richard Dawkins, famous for his idea of the "selfish gene." While there is, of course, much to be said in favor of this view, biologists often underestimate the role of physical law. The DNA-centered view therefore emphasizes chance over necessity. This has been exploited by creationists who like to abuse the concept of chance in evolution to claim that evolution is random and that randomness alone could not have created the complexity of life.

If evolution were truly random, the probability of creating just one functional protein would be astronomically small. Calculating such probabilities is a common parlor game among creationists. But these probabilities are irrelevant. Evolution is *not* random: It is the collaboration between a random process (mutation) and a nonrandom, necessary process (selection). It is the result of the balance of chance and necessity. This is not unusual—*all of nature is the result of this balance*. If not, nature would be either a featureless structure that is the same everywhere (if necessity wins) or a random "mush" with no structure at all (if chance wins). The exquisite order and the amazing variety we see in nature at every level—from galaxies to molecules—is the result of the fruitful interaction of chance and necessity. What is the probability of Earth or a pebble? It's a meaningless question. Similarly, the question about the random assembly of a protein is also meaningless. Evolution is not random.

Another favorite question of creationists is "How did all the information get into DNA?" At first glance, questions about information in DNA seem legitimate. This is because the message in DNA has meaning. It encodes the structure of a protein or regulates the development of an organism. But is meaning the same as information? Information is measurable; meaning

is not. Creationists conflate these two terms to suit their own ends. In information theory, a message has more information the more random it is. As we mentioned before, a perfectly ordered message contains little information. AAAAAAA contains no information, while ACTTGATTC contains information. But does ACTTGATTC have meaning?

The whole idea of DNA containing information is, in my opinion, one of the main culprits in maintaining the myth of creationism and intelligent design. First of all, without the genetic code and the entire machinery of transcription and translation, DNA contains neither information nor meaning. Worse, strictly speaking, DNA does not even encode proteins—at least not functional proteins. DNA only encodes the amino acid sequence of a protein. The functionality of a protein comes from its three-dimensional shape and the physical properties of various parts of the protein. This shape is the result of protein folding, which is the result of *physical forces* (sometimes helped by chaperonins) acting on a sequence of amino acids. Much of the information to shape a protein into its functional form is contained in the laws of physics and the action of these laws in space and in time. How would you quantify the information input life receives from physics and space-time? You can't. DNA only makes sense in the context of physical law, already-established order, and interactions with the environment. DNA does not tell us the final shape of an animal or a protein. Without context, DNA is meaningless.

An important statement of genetics is the *central dogma*: Information flows in only one direction, from DNA to RNA to proteins, never back from proteins to DNA. While the central dogma holds during replication, transcription, and translation, during the development of an organism, proteins control which parts of DNA are read at any stage of the development. There are feedback loops. The information to make a human being is therefore not encoded in DNA as in a blueprint. Although the word *blueprint* is often used for DNA, this is a misleading analogy. A much better analogy is *recipe*. To make a human being, DNA contains information to make proteins, which by their *physical interactions* with DNA, RNA, or other proteins, in the form of complicated regulatory feedback loops, shape the developing organism. This is similar to cooking a meal. A recipe does not contain a complete description of the result of cooking a meal; it just contains information about the ingredients (proteins) and the

timing of adding the ingredients (regulation). Then the physical interactions between the ingredients take care of making the meal.

Another way to put this is that organisms are emergent phenomena, emerging from complex interactions according to a specifically timed recipe. There is no way you could completely specify, in genetic code, every cell in a human being. How would you specify the trillions of connections in our brain? A few years ago, it came as a bit of a shock when the Human Genome Project revealed there are only about twenty-three thousand protein-encoding genes in the human genome. This is not much more than the number found in simple worms. I think the utter insufficiency of the information in DNA to specify an organism is one of the most powerful arguments for evolution. As argued before, life is a complex game played on the chessboard of physics and chemistry. I can think of no better analogy. Development of an organism needs information about proteins, but also needs space, time, physics, and complex feedback loops. None of these are encoded in DNA.

Evolution, like life, is also a game on the chessboard of time, space, and physics. The outcome of this game cannot be determined a priori. The game can create an enormous number of possible outcomes, and the role of evolution is to find some of these outcomes. By tweaking a protein here, or regulating a DNA sequence there, we see what effect this has; evolution has created a world inhabited by a limited set of all possible creatures that could theoretically exist. One of them happens to be us.

Creationists argue that humankind is the goal of Earth's history. If you start with this assumption, it would certainly be difficult to see how a playful process such as evolution could have necessarily ended up with us. Once you abandon this idea, however, and realize that evolution would have come up with *some* viable organisms, but not necessarily the same we encounter on our planet, it all starts making sense. Am I arguing for chance here? Chance is important, but I believe that life is inevitable and that myriad forms of life would have evolved in any case. As we have seen before, pure chance creates chaos; pure necessity, rigidity. Chance and necessity together become creation. What is created may be unpredictable, but creation itself is unavoidable.

This discussion is reminiscent of the differences between the views espoused by D'Arcy Thompson and Jacques Monod (Chapter 2). Monod be-

lieved that the existence of life is an incredible accident, the winnings of a cosmic lottery. It seems that Monod was too caught up in the DNA-centered view of life. Thompson, on the other hand, lived before we even knew about DNA. He emphasized necessity—he believed that all structures in living beings are the result of mathematical and physical laws. This is also clearly incorrect. Physical law by itself can make a rock, but without information, provided by evolution, we cannot make a living being. The views of Monod and Thompson can be combined to arrive at a fuller and more creationist-proof view of life: Information is important, but information comes from many sources—evolution, physics, chemistry, and the interaction of many complex entities in living cells.

Ratchets

The interaction of chance and necessity in evolution is mirrored by the interaction of chance (as molecular storm) and necessity (structure and physical laws) in the functioning of molecular machines. The second law of thermodynamics predicts that everything moves toward bland uniformity. Yet we have seen that the emergence of the bewildering complexity around us does not violate the second law, as long as we pay the free-energy cost. Still, to arrive at this complexity, we need some kind of free-energy-fueled mechanism. In our cells, directed motion, "purposeful" activity, is created by the action of molecular ratchets—molecular machines, enzymes, and motors, which by degrading free energy and due to their asymmetric structures, can rectify the random motions of the molecular storm to create order. Evolution is also a ratchet: It rectifies the random input from mutations into the creation of an ever larger number of possible creatures. This rectification is achieved by natural selection. Thus there is a pleasing analogy between evolution and its products, our molecular machines.

There is also a more direct connection between the molecular storm and evolution. As we saw from Delbrück's green pamphlet, which inspired Schrödinger to write his book *What Is Life?,* thermal motion is the main contributor to mutations. Even more to the point, replication and DNA repair are performed by molecular machines, which are subject to the molecular storm and therefore sometimes, although rarely, make mistakes. These mistakes supply fodder for the ongoing evolution of life on the

planet. Interestingly, evolution strives to minimize mutations. The extremely high fidelity of replication (one base-pair mistake in ten billion base pairs) shows that there is, paradoxically, an evolutionary advantage to not evolve. Evolution is rarely radical. The low error rate ensures that it is a gradual process. Nevertheless, small differences can sometimes have a large impact on the final result, because DNA encodes a recipe, rather than a blueprint. The same is true in cooking. For example, consider leaving the baking powder out of your cake!

The observation that evolution acts like a ratchet also discredits the probability arguments of the creationists. Evolution builds improbability step by step, mutation by mutation, selection by selection. The question "What is the probability of creating a kinesin by randomly combining amino acids?" is irrelevant to how evolution works. Kinesin did not spring into existence fully formed; nor was it a goal of evolution. It is simply something evolution stumbled upon, as it ratcheted up more and more complexity, one small change at a time.

There Is No Other Way . . .

Looking at molecular machines has made me realize that evolution is the *only* way these machines could have come to exist. As we have seen, life exploits all aspects of the physical world to the fullest: time and space, random thermal motion, the chemistry of carbon, chemical bonding, the properties of water. Designed machines are different. They are often based on a limited set of physical properties and are designed to resist any extraneous influences. The tendency of molecular machines to *use* chaos, rather than resist it, provides a strong case for evolution. Why? If life started by itself, without a miracle, then life had to start at the molecular scale. The molecular scale has always been dominated by the molecular storm. The ability of life to somehow incorporate thermal randomness as an integral part of how it works—as opposed to giving in to the chaos—shows that life is a bottom-up process. It is not designed from the top down. A top-down design would have avoided the complications of thermal motion by making the fundamental entities of life larger, so they could resist the molecular storm more easily. This is what machines designed by humans do—until recently, as nanotechnologists have learned from life's nanobots to create tiny machines of their own.

Molecular machines' exquisite adaptation to their molecular environment is also a strong argument for evolution. Evolution is tinkering—the gradual improvement and better adaptation of biological structures. The history of life has been long, and evolution had ample time to create these amazing physics-exploiting machines that run our bodies. To achieve such near perfection, you need a process that designs dynamically. A onetime design is not enough. Conditions change over time, and our molecular machines need to remain adaptable. An external designer would do best if the designer used evolution to do the work. Adaptation is assisted by the fact that physical laws provide the missing ingredient. For example, many structures in our cells are made through self-assembly processes, which are the result of physical forces (vesicles, collagen, etc.). Evolution does not overdesign: It designs just enough to take advantage of physical laws. If physics does the work for you, then why bother designing what is already designed?

We also ought to consider the commonality of the molecular apparatus in the cells of every living being. Many molecules and cellular processes in an *E. coli* bacterium, a yeast, a bluebird, a begonia, or a human are almost identical. This strongly suggests common ancestry. At the same time, looking at the differences between organisms, we see how various molecular machines have been adapted to fulfill specialized functions peculiar to each species.

Yet, the best argument in favor of evolution and against a static-design view may be that any designer would have to work hard to keep organisms from evolving. This is what I alluded to in the beginning of the chapter. Paley's reproducing watch would evolve. As we have seen, mutations happen. Some of these may impart an advantage to its bearer. Such an advantage would tend to spread through a population. How could you stop it? And why should anybody want to? If I were the all-powerful being in charge of the world, I wouldn't bother. Why not sit back, relax, and enjoy what wonderful things evolution can create for you?