1
The Rules of the Game

We begin our journey of discovery on the ground floor, in the familiar world of things we can see and hold in our hands. A world where we know our way around and can make careful preparations for our descent. For example, we'll need to understand how scientists discover physical laws by studying patterns in nature. It's also important to see that when we search for answers to the big questions, the laws we discover, however simple, can have a huge impact on our society. Sometimes the patterns are perfectly clear but the answers are hidden at a level invisible to the naked eye. How on earth can you see something a million times smaller than the smallest thing visible to the naked eye (spoiler: with a particle accelerator) and how do you learn your way around that strange world? We'll begin our adventure on solid ground before we start digging.

The Starting Point

Over the past centuries, we have uncovered many of nature's secrets. But how do you actually do that? An apple can't tell you why it falls when you let go, and you can't ask the sky why it's blue. Instead, to learn how nature works and unearth its secrets, you have to study the universe very systematically. How does nature behave under various conditions, ordinary or extreme? What phenomena do we observe, in all their details?

Instead of just sitting around waiting for nature to offer you information, you can roll up your sleeves and create a variety of conditions in a controlled way. That's what experiments are: a way of asking nature to answer the questions we have about its behavior. Experiments produce a wealth of facts that we can document and categorize. And they often show us that however complex the phenomena of nature may seem at first, they can usually be traced back to a small number of deceptively simple underlying principles. To find your way to these deeper rules and laws that nature strictly obeys, you have to be able to recognize patterns in the facts that you've collected. Children use this strategy almost instinctively when trying to understand the world around them. How will my parents react if I draw on the wall with my felt-tip pen? What will happen if I suddenly scream while we're waiting in line at the supermarket? And will it really hurt if I stick my finger in the flame?

Although parents may react in all sorts of ways (as I can tell you from experience), nature responds according to unbreakable, ironclad principles. These regularities, the laws of nature, give us information about how things work, and they're universally applicable. As soon as you've figured out how nature behaves, you can predict how it will behave in the future and in other situations. That's how, over the years, we've gradually learned more about the world.

Making scientific progress is not as easy as it may sometimes seem in retrospect. By definition, scientists are almost always in unknown territory, driven by an irresistible urge to find answers and not even knowing for certain that answers exist. Real scientific progress comes in spurts. We usually move forward in small steps, but every now and then comes one of those rare moments when we suddenly make a great leap. At a moment like that, we recognize the underlying mechanism and hit on a more fundamental set of laws. After the initial euphoria, we take our first careful steps into this new reality, this new world. And when we do, we find time after time that we observe new phenomena there. This kind of breakthrough could be the result of some genius's new insight, but it could just as easily be a chance discovery, or the outcome of a new experimental technique that allows us to study nature at a very different level. One good example is the invention of the microscope, a new technique that revealed a host of living mechanisms concealed in something as simple as a drop of water and so brought to light a hidden world.

This invention was a crucial step forward for medical science. And this type of discovery, which unearths a deeper level, takes us from the *how* to the *why*, as we learn to understand the strange phenomena we previously observed in terms of the new laws we've just discovered. But that's not all it does for us. It also often gives us the tools to make connections between phenomena that up to that moment had seemed completely unrelated.

Before we embark on our journey of discovery into the abstract world of elementary particles, I'd like to give a few examples of how we've persuaded nature to spill its secrets. Sometimes it was easy, sometimes hard, but each example shows that the drive to understand phenomena has both irreversibly transformed our understanding of nature and produced the knowledge on which our modern society is based. For the rest of this book, our guiding theme will be the gradual discovery of the world of elementary particles. But each time we look at an insight or discovery, however abstract and fundamental it may be, I will also try to show its practical applications, which have become indispensable parts of our everyday lives. We will see that fundamental research not only brings deeper insight into the workings of nature, but also has a profound long-term influence on the economy and society.

Going on a journey of discovery means being the first to enter new terrain, so you're almost certain to run into unexpected problems that don't yet have a solution. For example, no matter how well you can build strong, sturdy bridges, if you want to know what's on the far side of the ocean, you really have no choice but to build a boat. Other times, it's not hard work that leads to the next step, but a clever idea. To find out what's inside a walled-off area, you could chip away at the wall with a hammer and chisel for years, but the smart thing to do is make a ladder. This all seems deceptively simple, because we already know how to solve these particular problems, but imagine being the first person ever to come up with these ideas.

In short, scientists are real adventurers. They may not become millionaires, but imagine the triumph and everlasting fame of being the first to reach the summit of Mount Everest or set foot on the moon, or, in my own field, of being the first to identify the fundamental building blocks of nature, figure out why no antimatter can be found anywhere in the universe, or learn whether empty space is truly empty or actually full of a mysterious substance that gives all particles their mass.

There's a great analogy I once heard from a colleague that sheds light on how challenging it can be for a scientist to identify patterns and construct a theory. Imagine you're an alien who has just landed on earth. Chances are almost everything you see on this planet surprises you, but you decide to be systematic and start with something simple. So you ask yourself, "What are the rules of the popular game soccer (or "football") played in every country on this planet?" It's a clear question and may seem easy to answer. But there's a catch: you can watch as many games as you like, but you can't talk to anyone about it or read anything on the subject. All you can do is watch. Give that a minute to sink in, and then ask yourself how long it would take for you to come up with a complete list of the rules.

You'd probably figure out pretty fast that there are two teams of eleven players, that the whole game is played between the outermost white lines, that players switch sides after forty-five minutes, and that the basic objective of the game is to kick as many balls as possible into the opponent's goal. But who are those players, one on each side, who do a lot less running, dress differently from their teammates, and are allowed to use their hands? How long will it take you to realize what the two humans are doing who run back and forth along the sidelines with little flags, or to understand corner kicks, offside, substitutions, the mysterious extra time at the end of some matches, the strange lines on the playing field, penalty kicks, and so on? Imagine how difficult it would be to discover all the rules. Difficult, but not impossible—at least if you're really, really motivated and willing to invest enormous amounts of time. Scientists face exactly the same challenge, but this time, the playing field is the world around us. Nature doesn't just give away its secrets free of charge. Only by observing carefully, and designing experiments that ask nature the right questions, can we figure out what phenomena exist. That allows us to decipher nature's rulebook, bit by bit.

By the way, no one ever said the rules had to be logical. In fact, the laws of nature *don't* fit everyday logic—not one of them. Quantum mechanics and relativity, two of the most famous theories we'll encounter in this book, are both deeply weird. In a sense, you could compare them to the offside rule in football: absurd but true, and simply the way the game is played. As soon as you accept that it's a rule, it's only logical that certain goals count and others don't. Likewise, the bizarre principles that underlie theories like relativity and quantum mechanics are completely counterintuitive, but once you accept them, they do explain all the bizarre and complex phenomena that we see when we look at nature on the atomic scale. The theories are right. But logical? Nope.

By working out the consequences of these strange theories, we've learned how to apply them in useful ways in everyday life. Much of the current research and progress in nanotechnology and quantum computing are founded entirely on that peculiar theory known as quantum mechanics. Although many experimental results and observed phenomena are "logical" according to that theory (in other words, they're explained by the strange laws of quantum mechanics), there's not a scientist on earth who understands *why* the world obeys the quantum mechanical laws. For instance, how is it possible for a particle to be in two places at once or in an entangled state that would be unthinkable in our ordinary world? You can't be both pregnant and not pregnant at the same time, but in the quantum world, mixed states of that kind are completely normal.

The more successful a theory is, the more comfortable we get with the rules and laws of that strange world. At the same time, it's frustrating to have no explanation for the basic building blocks of your logical framework. In this case, we've gone beyond the *how* to the *why*: quantum mechanics. But no sooner do we take that step than the question shifts from the *how* of quantum mechanics to the *why* of quantum mechanics. In other words, the answer immediately gives rise to a new question. The poor scientists are always running after shifting goalposts. But their insatiable curiosity has led to huge gains for our society, because the insights and applications arising from their work have become foundation stones of our civilization. Still, it's important to realize that even the smartest scientists on earth will run out of explanations fairly quickly if you keep asking "why."

Over the past hundred years or so, elementary particle physicists have managed, step by step, to reach the innermost depths of the atomic nucleus. In the course of that adventure, we have taken a few great strides and arrived at fascinating insights—not only into the building blocks of all the matter in the universe, the elementary particles, but also into the fundamental forces of nature. There's every reason to take pride in that. Before we delve into the world of elementary particles, I'd like to show you a few simple laws of nature that we humans have learned to manipulate, and patterns that we've uncovered. These three examples will demonstrate that a few things we take for granted have no logical basis but nevertheless play a crucial role in everyday applications that arose from pure fundamental research. The first two relate to the questions of how electricity is generated and where inherited traits are encoded in the human body.

One of the greatest threats to our prosperity and way of life is a shortage of energy. We don't often pause to think about it, but our Western society is addicted to energy, and without electricity, it would come to a complete standstill in less than a day. Try to imagine a typical working day without electricity: no alarm clock, no lamps, no coffee machine, no car, no elevator, no ATMs, no television, no computer, no radio, no Internet, no telephone, and no dishwasher. You can clearly see why energy is such a hot topic of public and political debate. This book is not the place to review all the facets of this complex issue, such as the earth's finite supply of fossil fuels, the geopolitical interests at stake, carbon dioxide emissions, green energy, and the nuclear power debate. You could fill a library with expert studies of those subjects, and the debates are still in full swing. My job here, as a physicist, is to raise a question that I know doesn't play a central role in the discussion, but which I'd like everyone to be able to answer: *How do you make electricity?* For example, how do you turn a charcoal briquette, the kind you use in your barbecue grill, into an electrical current? We can do that thanks to a secret that nature revealed to us through one simple observation. The discovery of this seemingly straightforward law of nature changed our lives and our civilization fundamentally.

Some one hundred and fifty years ago, James Maxwell managed to capture all known facts about electricity and magnetism in four famous formulas that have carried his name ever since: the Maxwell equations. They show that magnetism

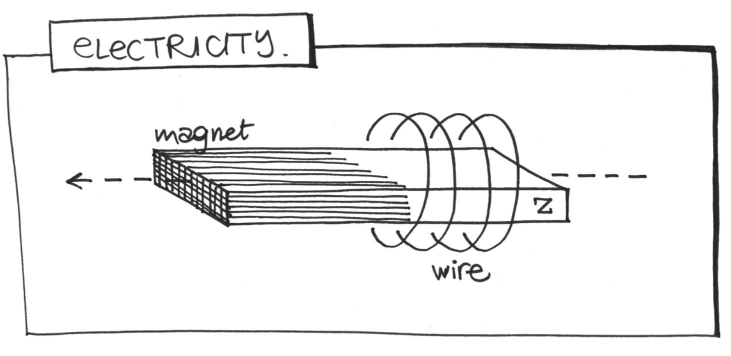

and electricity are intimately interrelated, and they describe electromagnetic phenomena that can be extremely complex—including a regularity discovered earlier by Michael Faraday that gives us a way of making electricity.

> **Observation (and law of nature):** *A current starts running through a coil of copper wire when the magnetic field in the center changes.*

That may not sound so exciting, but if you take a coil of copper wire and move a magnet through it, it really happens. The magnetic field is absent at first, becomes very strong as you move the magnet into the center of the coil, and then disappears once the magnet has been removed completely. This change creates an electrical current in the wire. That's the simplest way I can put it. But I couldn't make it much more complicated either. So there it is: the principle that allows us to generate electricity, whether in a bicycle lamp or in the most up-to-date nuclear power plant.

When you ride a bicycle, you generate the power for your own lamp. A hub dynamo is basically a very long, rolled-up copper wire, like the long iron wire coiled neatly around the hose of a vacuum cleaner. Inside that tube of coiled wire is a magnet connected to your wheel by a gear called a roller. As you pedal your wheel turns, and so does the magnet. And because of the law of nature we discovered, which says that changing the magnetic field in the tube of copper wire will make a current run through it, a current really does run through it. This current is then conducted through a thin metal wire into the lamp, making it warm up and start to glow. Ta-da: a bicycle light. And not to make the high-tech energy giants seem unimpressive, but a big coal-fired power plant works about the same way. There too, a magnet is rotated inside a coil of copper wire. The only difference is how you move the magnet. On your bicycle you do that by pedaling, while in a power plant the work is done, surprisingly enough, by a turbine. The turbine spins because of steam pushing hard against the blades, which are connected to the magnet by a gearbox. So how do we make steam? By heating up a large container of water. And how do we heat the water? By burning a heap of coal underneath it. It's that simple.

Of course, thousands of people work hard every day to make each step of this process as efficient as possible in power plants, and much more is involved than I describe here, but this is the basic principle. And a nuclear power plant works almost the same way. The only difference is how you heat the water. In a nuclear plant, this is done by the particles released when you split the nuclei of heavy atoms such as uranium. Wind energy uses the same principle: the spinning blades of the turbine turn a magnet inside a copper coil.

This very simple principle—generating current by changing a magnetic field—is fundamental to our economy and therefore to our prosperity. When Faraday first made these discoveries, no one could suspect how they would be applied. William Gladstone, who was in charge of the British Treasury, is said to have asked Faraday, "But, after all, what use is it?"—an understandable question from his point of view. It is the same question that scientists are still asked today whenever we apply for research funding. Unfortunately, we can no longer get away with Faraday's famous reply: "Why, sir, there is every probability that you will soon be able to tax it!" Back then, people were fairly content with candlelight, and it probably seemed more sensible to give money to the candle industry to find a more efficient production method or design a better wick. But in retrospect, we can see that we would never have discovered the light bulb that way.

Even though a lot of scientific research leads absolutely nowhere, this is a striking example of how true innovation can't always be planned in advance. The game-changing discoveries often come from unexpected places. That's an important message for politicians and for society as a whole: alongside innovation for industry, we need to create enough opportunities for free, unrestricted fundamental research. Applications are sure to follow.

For physicists, it's often both intriguing and frustrating to follow public debates about electric cars and hydrogen-fueled vehicles. We're sometimes startled to realize that, despite all the well-meant efforts of scientists and manufacturers, politicians and policymakers are unaware of the most basic scientific facts. The discussion often focuses on distant future scenarios, but a little more technical savvy among politicians could have a greater impact than all the scientific reports in the world. True, a Tesla doesn't give off carbon dioxide as you drive, but even so, it's striking how rarely people think to ask where the energy in the battery comes from. That battery is charged, via an electrical socket, with power generated by a coal-fired station that produced lots of carbon dioxide. On top of that, the battery is made of heavy metals and very aggressive acids, not exactly "green" technology. Of course it's a good idea to reduce our consumption of fossil fuels, and of course a large power plant is more efficient than a thousand separate automobile engines. And of course you can also charge the Tesla battery with solar energy. Even so, the popular belief that electric cars are squeaky clean is an exaggeration.

A similar debate is in progress about the hydrogen economy. The idea is to generate power by mixing hydrogen and oxygen, stored in separate tanks, and then burning them as fuel. That's certainly an ultra-clean process; it produces nothing but energy and water. You might think it's the ultimate in clean fuel for your car, but again, there's a catch. Where do you think you find the pure hydrogen and pure oxygen? Well, you start with water, which is made up of oxygen and hydrogen atoms. To separate the two, you need energy. It's exactly the same process as combining them for fuel, but in the opposite direction. And where do you find the energy for doing that? That's right—usually a large coal-fired power station or a nuclear plant.

Or from a wind turbine, of course. There's no denying that we can also use green energy to separate hydrogen and oxygen from water. But my main point here is that hydrogen is an energy *carrier*, rather than an energy *source.* It does offer certain advantages: no carbon dioxide or soot is released in the area where the car or bus is driven, the city center, and it's a fantastic way of storing excess energy from turbines, solar cells, and power plants for short-term use. But it's not *the* solution to the energy problem.

Not only physicists, but scientists in many fields, are in search of explanations for undeniable patterns that they see but don't understand. Fundamental research has not only given us electricity but also led to major discoveries that are now central to medical science. When you look at nature systematically, you always find a treasure trove of data. Sometimes, once you've collected enough information, you can recognize patterns and find your way to deeper insights. The technique of reductionism—zooming in on the basic building blocks of which things are made—is not reserved exclusively for physicists. One good example of the progress we've made by probing ever deeper into the building blocks of cells is the discovery of DNA and the coding of genetic information. This discovery stands head and shoulders above all others in its impact on medical science.

Biologists and farmers have known for a very long time that animals and other organisms pass on traits to their offspring. The best-known example is probably the eye color of parents and children. If a child's parents both have brown eyes, the chance that the child will have brown, green, or blue eyes is 75 percent, 19 percent, or 6 percent, respectively. There are many tables of this kind for different traits, from the color of a cat's fur, to the resistance of crops to certain diseases, to the ability of plants to adapt to saline soil or other extreme conditions. In the case of eye color, the only serious problem arises when we need to explain why a child has brown eyes even though both parents have blue eyes (since the probability of that is zero). But heritable traits do have a tremendous influence on our food supply. In the agricultural sector and the food industry, knowledge about heritable traits is used daily in an effort to pass on desirable traits, such as resistance to disease, high milk productivity for cows, or the adaptability of rice plants to extremely dry or wet climates. Selective breeding of plants and animals over many generations can make desirable traits more widespread in the population, using the laws of nature for our benefit.

So we can *see* that nature follows certain patterns, and we could devote our entire lives to studying which traits are and aren't heritable. But the question we'd most like to answer, of course, is: *How* does it all work? The source of individual traits is apparently hidden away somewhere in the body. But *where*? Is it only in the egg cell and sperm, or in every one of the body's cells?

This question went unanswered until the 1960s, when new detection techniques enabled Francis Crick, James Watson, and Rosalind Franklin (the last of whom is often "forgotten" in the history books) to investigate structures much smaller than a human cell. They discovered the double helix structure in which genetic information is stored: DNA. Information about eye color and many other traits turned out to be kept in cell nuclei, and the language in which that information was written had an alphabet with only four letters, each corresponding to one nucleotide: C (cytosine), G (guanine), T (thymine), and A (adenine). Together, these four organic molecules encode all the traits and complex phenomena we observe in living things. What we learned through this revolutionary discovery is that, while our alphabet has twenty-six letters, only four are needed to record an individual's complete genetic code. If you were fluent in that language, you would know right away where a person's eye color is encoded, and you'd understand why some people are susceptible to certain diseases and others are not.

This was one of the most important scientific discoveries of all time and laid the groundwork for modern biomedical research and drug design. It yielded crucial insights into cell division and pointed to new questions: Where is each specific trait encoded? How does the cell copy the DNA when it divides? What is the effect of an "error" in a strand of DNA? How do you "read" the strand? What words are formed by the CTGA combinations? Can we find the source of cancer in the genes? And can we manipulate the genome to prevent disease? Although the basic idea of the alphabet of genetic material has been around for fifty years, we have not yet completely mastered the language. We find new combinations and patterns just about every week and only recently have succeeded in fully mapping large portions of the human genome.

But our knowledge is evolving at a dizzying pace. On the website of the U.S. National Human Genome Research Institute, you can find statistics showing that it cost a hundred million dollars to decode the entire genome in 2001. Today, it costs only a few thousand. And home kits can even be used to analyze part of your DNA from a sample of your saliva. Apart from medical scientists, physicists such as Cees Dekker in Delft, the Netherlands, are also participating in cutting-edge genetics research, and their work has led to the ability to read out long DNA strands efficiently. Once we can do that, the obvious step will be to start building our own DNA structures.

Like anything else, this development has a good side and a bad side. In recent years, advances in genetics have made the news almost every week. Sometimes it's a new genome that's been partly or fully decoded, or a genetic modification developed to track down or even repair the source of a disease. Everyone is in favor of new techniques for the early diagnosis of heritable disease, or new medicines tailored to individual genetics, but at the same time, these innovations give rise to many ethical debates. For example, do I want my health insurer to know that I have a high risk of cancer, and if so, what can we—or should we—do with that information? How will society handle the capability to detect many genetic disorders in unborn children, and do I have the right to decide for myself whether or not I want that information? While selective breeding of plants and animals for desired traits is pretty widely accepted, that's not the case with direct manipulation or synthesis of genetic material—in other words, genetic modification. The moral of this example is that simply compiling larger and larger books of tables for heritable traits would never by itself have led us to discover DNA; figuring out the pattern, the genetic alphabet, made the difference. Whether and how we use our new insights and techniques are matters of ongoing social debate.

Success stories about recognizing patterns, like the story of DNA, are an easy way for scientists to score points. But just as often, science is a frustrating tale of our inability to really understand the phenomena we study. Sometimes what we need is a brilliant insight, and sometimes we simply don't yet have enough information or knowledge to take the next step.

Unfortunately, you can't always find satisfactory answers to the questions you ask yourself about nature. If there's one thing we can count on, it's that the sun will rise again tomorrow. Right? It sounds so self-evident. But when Albert Einstein was young, he probably looked at the sun now and then and thought, "What keeps it burning?"

Now, I don't mean to pretend that I know what went on in Albert Einstein's head, but I know for certain he couldn't answer that question back then. What makes me so sure? Well, *no one* in the world knew the answer a century ago, because the scientific knowledge required to understand the answer, even in the most basic terms, was simply not yet available. Oddly enough, I never thought of that myself until about ten years ago, when I was preparing a talk for schoolchildren. It's a strange idea, because if no one knew what kept the sun burning, that means no one knew how long it had been burning and—not exactly an insignificant detail—how much longer it would go on burning. So what did Albert and his scientist friends believe? And what did the rest of humanity think? Why wasn't everyone on earth obsessed with the question?

Yes, of course, *now* we know that the sun keeps burning because energy is released whenever two hydrogen nuclei fuse into a helium nucleus in its core, thanks to the incredibly high temperatures there. But in the early twentieth century, the atomic nucleus had not yet been discovered. That didn't happen until thirty years later, when scientists pushed their machines to the outer limits of their capacity. That discovery, again, would inspire a variety of applications, as we'll see in

Chapter 2: not only nuclear energy and the atomic bomb, but also nuclear fusion, our great hope for a clean solution to the earth's energy problem.

Venturing into the World of the Very Small

Any expedition into unfamiliar territory demands the right equipment. If you want to go to the North Pole, you'll be better off with warm clothes, a pocketknife, and a dog sled than with a Hugo Boss suit, a cheese slicer, and a bicycle. And if you want to go to the moon, you'll need to build a rocket and find a spacesuit. Our branch of science is no different. To descend into the world of elementary particles and explore structures even smaller than the building blocks of DNA, we need one essential tool: the particle accelerator. This workhorse of particle physics is a complex piece of equipment that, like a Swiss pocketknife, can be used in various ways.

First of all, a particle accelerator is an extremely good microscope, with a power of magnification thousands of times greater than ordinary microscopes can ever achieve. Long before we had this tool, it became clear that there was a fundamental limit on the smallest objects we can see with a conventional microscope. But the particle accelerator finally allowed us to break through that seemingly impenetrable barrier into a much smaller world. We not only became familiar with the smallest distinct building blocks of each element, atoms, but also learned exactly how they were made up of even smaller building blocks. We also found out that on that scale, the laws of nature are fundamentally different than in our everyday world.

Encouraged by the success of this new technology, we built ever more powerful particle accelerators and discovered a second way of using them: as "nutcrackers." You see, a microscope is not always the right tool. If you want to know what's inside a walnut, for instance, then a microscope won't help. The microscope may show you the nutshell in incredible detail, but if you want to know what lies hidden within the shell, you'll have to break it open with a hammer or a nutcracker. And that's a perfect description of our second way of using a particle accelerator: by firing particles at an object at high velocity, we can actually crack open either those particles or their target. By studying the wreckage left by the collision, we can find out what was inside the material that we bombarded.

What will be most relevant to our journey in this book is the *third* way of using a particle accelerator: to create new matter. We've discovered—to our great surprise, I might add—that if you fire particles with extremely high energy at each other, they not only collide but also actually create *new* particles. This led to utter chaos at first, because we discovered hundreds of different kinds of tiny particles, but eventually the whole menagerie turned out to be a beautiful puzzle that could be assembled with a very limited set of building blocks. Those building blocks, and the ways they attract and repel each other, ultimately came together to form what we call the Standard Model: the outer limit of our knowledge of elementary particles, a magnificent framework that still stands tall and describes almost every phenomenon in that mini-world. We'll discuss this third function later on. But first, let's consider the particle accelerator in its role as a super strong microscope.

In everyday life, we use our eyes, nose, ears, mouth, and hands to perceive the world around us. Our eyes and nose, for example, are perfect instruments for distinguishing jam from butter in the morning. At the same time, we know that our eyes, as complex they are, fail us completely when we try to look at very small objects. Threading a needle is difficult enough;

forget about checking whether an ant has teeth or whether there are bacteria in a drop of water.

Long ago, we figured out how to combine lenses in an ingenious way to make a microscope, with which we can explore a smaller world. But even today's microscopes, which are far stronger than Antonie van Leeuwenhoek's original model, run up against a fundamental limit to their power. A microscope can never—*ever*—detect things smaller than a millionth of a meter. That's incredibly small, of course, and microscopes are perfect for looking at bacteria and cells, but they're just not the right tool for studying DNA or atoms. To break through that fundamental barrier, you need a clever trick. And we found one! Strangely enough, it required us to start looking at things without using our eyes.

Whenever we *see* something, it's because our eye has captured particles of light that have bounced off the object we're looking at. As you walk down the street, you can see the people around you because light from the sun bounces off them and straight into your eye. The retina in the back of your eye acts like a kind of digital camera with a whole lot of pixels, and your brain has learned to interpret the patterns and translate them into complex objects. That's how you can see the difference, at a glance, between a lamppost and a human being, or between stones and water. Our own built-in digital camera has two components: "rods" for measuring the light *level* and "cones" for seeing different *colors*. Together, these two types of photoreceptors in the human retina give us enough information to perceive the world around us.

But however sharp our eyes and ears may be, they're not perfect, and many worlds remain hidden from us. For example, there are pitches that our ears cannot detect but that dogs can hear perfectly. And besides good hearing, dogs have a phenomenal sense of smell. That's why they're used in harbors and airports to track down drugs and money hidden in suitcases. There are also worlds hidden from our vision—things that people can't see but that really are there. Our eye is not infinitely sensitive and can't see all colors.

We're all aware of our limits when it comes to the *intensity* of light. On a dark night, we humans can hardly see a thing, but we know that cats don't have that problem. Since cats have many more rods than cones, they can see even in very poor light. Their eyes work better because of their different structure. Cats see a whole world in the dark, a world hidden from humans because we simply aren't equipped to perceive it. But sometimes people do need to see in the dark—so, being the resourceful creatures that we are, we invented night vision devices. They intensify the few particles of light that do manage to reach us into a signal that we *can* see. Pretty smart!

Even when the sun is at its height, some worlds remain hidden from us simply because our eyes can't see every *color*. I'm not talking about color blindness, but about the limitations of the average human eye, which has cones that are sensitive only to the colors of the rainbow, the familiar range from red to violet. You see, each color of light corresponds to a distinct wavelength—meaning the length of the wave of light that flies toward you. The shortest waves that we see look violet, and the longest ones look red. Because of the shape of the rods in the human eye, they aren't at all sensitive to the colors corresponding to longer waves than red or shorter ones than blue. But those colors do exist. They're called infrared and ultraviolet, and plenty of animals can see them: an entire world hidden from us in bright daylight because our human eyes can simply not detect them.

A bee, for instance, can see ultraviolet light, those colors with a slightly shorter wavelength than violet, which is just barely in our visible range. Is that a useful skill? You bet. Along with the yellows and reds visible to us, some flowers display intense ultraviolet colors, which we can't see at all. So a bee flying over a field of grass can pick out the different flowers there effortlessly, while we struggle to make out a few vague hints of flowers in the green grass. At the opposite end of the spectrum are animals that can make out infrared colors, which we can't see with our eyes but can feel as warmth on our skin. Snakes use those colors to track down prey with ease.

We humans can invent special devices, like those clever night vision goggles, to use in scientific experiments. You might call them artificial eyes, ears, and noses that enable us to discover, observe, and explore worlds hidden from our senses. We must always remain aware that there's more around us than we can perceive with our senses and keep working hard on smart technology that makes those worlds visible to us in other ways.

A little more than four hundred years ago, a Dutch inventor of scientific instruments, Hans Lippershey, noticed that an ingenious combination of lenses made it possible to magnify a small object. The possible applications were endless. Galileo Galilei improved the telescope that allowed him to study the moon and the motions of the planets in great detail. But we also learned how to peer "into the depths" with microscopes, gaining access to another completely unknown realm. For example, Antonie van Leeuwenhoek discovered a world of wonders in something as simple as a drop of blood, becoming the father of modern microbiology. Although in the centuries since then we've steadily improved the design of the microscope, we knew that this avenue of exploration would one day come to an end, because of a fundamental limit to what you can see with light. Even using a microscope, we will *never* be able to see anything smaller than a millionth of a meter (about one hundredth of the thickness of a hair). That's not because we can't make better lenses, but because waves of light simply don't bounce off such small objects.

One thing we know about waves is that they bounce off only those objects larger than their own wavelength—a principle of physics that, for the time being, you'll just have to accept. Think of a marble rolling across the kitchen floor—it bounces when it hits the trash can, but rolls over a breadcrumb without any effect. To calculate the smallest object that light can bounce off, we need to know how large a light wave actually is. That varies a little by color—as I mentioned, red has a longer wavelength and blue a shorter one—but the light visible to us humans normally has a wavelength of a little bit less than a millionth of a meter. If you want to see something even smaller than that, you'll never succeed with a traditional microscope, not even the most powerful one in the world.

Fortunately, when faced with an insurmountable obstacle like this one, we don't have to throw in the towel. We just need a bright idea or a completely new approach. That's how fundamental science works: amazingly, no matter how big the problem is, a solution is eventually discovered that breaks through the barrier and expands our horizons. This problem was no exception. But we did have to look very far beyond traditional techniques.

Besides looking with your eyes, there are plenty of other ways to figure out an object's shape. If you close your eyes, you can still easily feel the difference between a knife and a fork.

Scientists use a similar method to "feel out" objects, but instead of using our hands, we fire little bullets and watch how those bullets bounce off the object we're studying. The way they rebound, or *scatter*, gives us information about the shape of the object.

To picture how that works, imagine an object on the floor of your living room, about a meter away, closely surrounded by a curtain that hides it from you. Your goal is to find out what's behind the curtain, and your only equipment is a bag of one hundred little marbles. All you can do is bombard the object with marbles, rolling them across the floor and under the curtain. They'll bounce off the object and come ricocheting out again. By looking carefully at how the marbles bounce and sending them toward the object at different velocities, you can get an impression of what the object looks like. If there's nothing on the floor behind the curtain but bread crumbs, then the marbles will roll straight over them and reappear on the other side as if nothing had happened. But if there's a sheet of wood at a 45-degree angle, the marbles will rebound exactly to the left or to the right. Those are both straightforward examples. But if you have to find out whether there's a thin wooden partition or a thick sheet of iron behind the curtain, the job gets a little harder. And just imagine shooting marbles to figure out whether there's a stuffed Mickey Mouse or Donald Duck doll behind the curtain. That's more challenging—extremely difficult, in fact—but not impossible! What you would need are (1) your marbles, (2) some idea of how Donald and Mickey make marbles rebound differently, and (3) a way of keeping track of the angles at which the marbles bounce back. This technique for looking at small objects has a fairly long history among scientists, who nowadays use particle accelerators to make the little bullets that they fire at objects. (I'll go on using

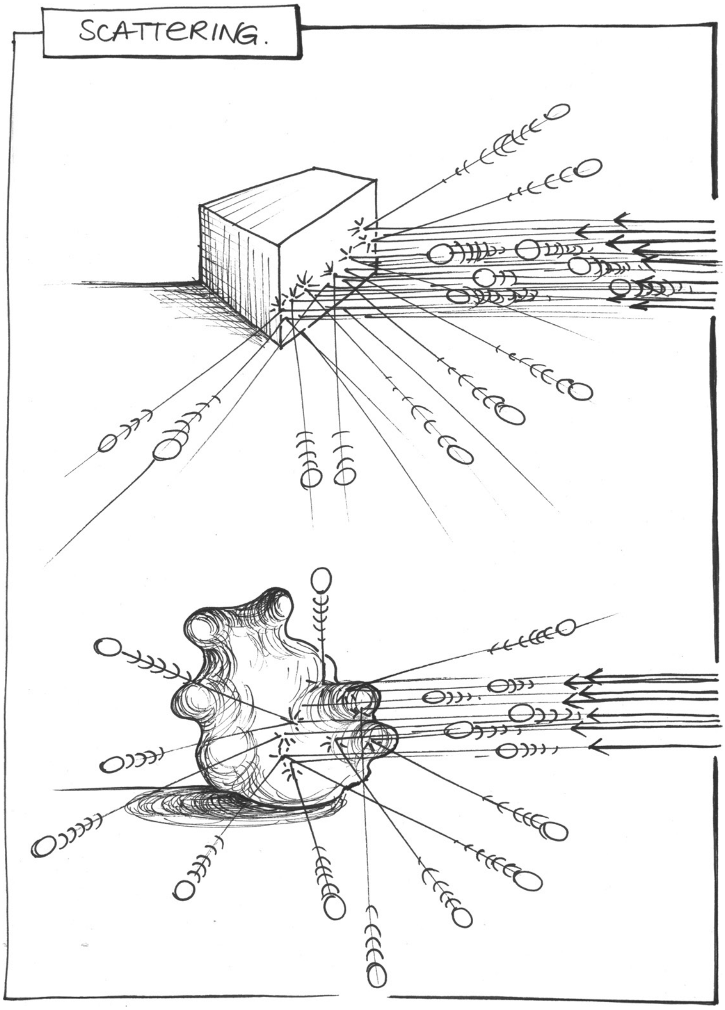
SCATTERING.

the word "bullet" here, but we're really talking about small particles.)

An electron microscope uses the very same method—rebounding particles—to create images of small objects like cells, details of an ant's eye, surfaces of a metal, or small structures in nanotechnology, so that we can study them. An electron microscope may use electrons, for instance, as projectiles to examine a surface or object. This technique has made it possible for scientists to go further into the micro-world than they can with conventional microscopes, ultimately to the point of unlocking the world of the atom.

There are three things we need to enter the world that remains hidden to our eyes and conventional microscopes:

1. Make tiny bullets and fire them—a particle accelerator.
2. Calculate how bullets are scattered by a particular shape—a theory.
3. Keep track of the scattered bullets—a detection device.

These are the three things with which my predecessors were equipped, and which we still use today in our search for ever smaller objects.

The particle-scattering trick is exactly what enabled the early twentieth-century scientist Ernest Rutherford to study the unique building blocks of each element, namely atoms, in detail. Later we'll see in more detail how he did that, but essentially, he fired small particles at a layer of gold atoms at high velocity and watched how they bounced. He had various ideas about the structure of atoms, but only one of them turned out to fit the measurements he and his assistants had made. They found that the atom was made up of a tiny, heavy, electrically charged particle, the nucleus, with a number of light electrons orbiting it. This new technique taught us that the world at its smallest scale had many secrets that would change science forever. That discovery set off an experimental gold rush, as scientists went searching for still more secrets of that mysterious world.

At this point I should mention that the traditional view of a particle as a hard little ball is not entirely correct. It's a bit more complicated than that. The rules of play for matter on that scale—quantum mechanics—state that particles also behave like waves. That sounds implausible, since the matter around us doesn't act like a wave at all. Yet experiment after experiment has shown conclusively that it is the case. Is that weird? Yes, very weird. It's good to realize that physicists think it's strange too. The only difference is that they're more used to the idea and have resigned themselves to the fact that not everything in the world of the atom is logical. One of those strange new rules states that the wavelength of a particle (and therefore its effective size) depends on its energy: the faster the particle moves, the greater its energy, and the smaller it effectively is. If we want the "bullets" we shoot from our particle accelerator to be sensitive to certain details of our object of study, then we need to use projectiles smaller than the structures we're investigating. By increasing the energy of the particles, we can make them smaller and smaller, so that they'll ricochet off even tinier objects and allow us to recognize even more details. The particle accelerator's job is to boost the energy of the particles as much as possible, so that they're as small as can be by the time we fire them at the object we're studying.

Like conventional microscopes, these "particle microscopes" have come a long way since the early days and are now used daily in many branches of science to reveal the surfaces of cells or materials. Any photos you see of nanostructures, the eyes of an ant, red blood cells, or cancer cells, for instance, were made with this technology. So the world's most powerful microscope is whatever particle accelerator can boost particles to the highest energy. Right now, that's the proton accelerator at CERN in Geneva: the Large Hadron Collider. The particles used in Geneva are at such high energy that their wavelength is approximately 10^{-20} meter, about one millionth the size of an atomic nucleus.

Although we often think of the particle accelerator as scientific equipment, it plays a larger role in our society. The best-known use of a particle accelerator, outside of science, used to be in an old-fashioned television, the kind with a screen that bulges outward. That screen is lit by a beam of electrons fired at high velocity, and magnets make the beam slither back and forth, hitting every point on the screen to create an image. The rise of LCD televisions has made that example awfully dated, but fortunately, particle accelerators also have a broad range of applications in health care and industry.

The ability of particle accelerators to make small objects visible has led to major scientific breakthroughs, because we can now produce detailed images of objects as small as blood cells or parts of ordinary cells. And for the past twenty years, manufacturers have been racing to build ever smaller structures, especially on computer chips. After a new idea or technique has been invented, it's important to examine the results (and any potential problems) in detail. That's why almost all high-tech companies use electron microscopes to produce images of the metal surfaces or nanostructures that they've made. The applications in health care are even more fascinating, because they're less familiar, so it's worthwhile to look more closely at a few of them. Even though these particle accelerators are, of course, not nearly as powerful as CERN's Large Hadron Collider, new developments in acceleration and detection technology in the scientific community have also benefited industry and hospitals.

When I'm giving a talk and I say that every modern hospital has a couple of particle accelerators, people look at me in astonishment, because few people realize that to make an ordinary X-ray or treat tumors, you just can't do without a particle accelerator. Even though, sad to say, many people have first-hand experience of X-ray images or radiation therapy for tumors, hardly anyone knows how those technologies work. Patients have other things to worry about, and when you're lying on that table, you don't ask exactly how it works. But because these are applications of particle accelerators, I'd like to tell you more about them.

The first well-known medical magic trick is looking straight through your skin to see your bones. Making an X-ray image is a lot like bouncing small particles off objects with a particle accelerator. The doctor fires something at you—in this case, X-rays (which are really just waves of light)—and examines whether they *did* or *didn't* pass through your body. Those rays have enough energy to go straight through soft skin and flesh, but *not* enough to penetrate hard bones. As soon as they hit bone or another hard material, they come to a stop. So in this case the main question is not the angle at which X-rays rebound, but whether or not they are blocked on their way through your body. Even though you can't see the rays that pass through the body with your naked eye, you can make them visible in a different way. It turns out that some materials turn dark when hit by X-rays (like the film in an old-fashioned camera, which turns dark when exposed to visible light). If you put a plate of this type of material behind the patient being scanned, the result is a detailed photo showing the exact locations of hard objects in the body. Of course, it also shows you where the hard parts are missing—for instance, if a bone is broken. You might compare it to the silhouette on a wall after a passing truck splashes mud all over you. That too is a negative image, showing where you stopped the mud from traveling onward. A brilliant idea, but where do those X-rays come from? That's where the particle accelerator comes into it. It whips up electrons to a high velocity and then fires them at a metal plate, releasing X-rays, which are aimed at the area the doctor wants to photograph. So without particle accelerators, there could be no X-ray images.

A traditional X-ray shows your bones in considerable detail, but sometimes we will want a sharper focus. One possibility is simply to fire more rays, but that does pose certain risks. There's a good reason that the professionals leave the room during an X-ray in the hospital or at the dentist's office: the radiation is hazardous. As it passes through your body, it can do serious cell damage and possibly even cause cancer. To limit that risk, we all accept that the image won't be quite as sharp as it theoretically could be. But there's a second way of obtaining a sharper image, namely by making the photographic plate more sensitive. This is like increasing the number of pixels in a digital camera. Scientists, like camera designers, are constantly working to improve their detection techniques. They've been very successful, and there are various initiatives to introduce the new technology in hospitals. After all, whenever we make our detection equipment more sensitive, we can take a better X-ray with the same amount of radiation—or we can take an X-ray with the same resolution while using less radiation. This doesn't make much difference if you're only taking one X-ray, but a CT (computed tomography) scan is the equivalent of two hundred X-rays. In that case, better technology means much less radiation and lower health risks for patients.

A good example of a joint initiative by science and industry—one that focuses on medical applications, detection techniques, and read-out chips—is the MediPix project, which has opened up possibilities for the color X-ray. The X-ray images we've discussed so far are always black and white, showing whether or not rays with a certain energy passed through the body. One disadvantage of black and white is that you can't see whether the rays were stopped by a bone five centimeters thick or a thin, two-millimeter iron plate. But suppose that instead of a single energy level, we could work with a whole spectrum of rays? That would give us more information, because we could see how much energy the particles need to pass through each obstacle. This principle has already been proven by researchers like Enrico Schioppa in his doctoral research at Nikhef, the Dutch National Institute for Subatomic Physics in Amsterdam. Dutch particle physicists perform a large range of experiments, as well as think about how they can contribute to new technologies. For a research and development group like Enrico's, seeing their research applied in an actual hospital would be a huge breakthrough. This requires a phenomenal amount of number-crunching, but fortunately, the Science Park in Amsterdam, where Nikhef is located, is also home to the Centrum Wiskunde en Informatica (CWI), the renowned Dutch research institute for mathematics and computer science. Together, these institutes are working on a test setup to show the world that a full-scale color X-ray system is possible.

Unfortunately, almost everyone knows somebody with cancer who needs radiation therapy. That type of treatment uses high-energy radiation, like X-rays but with much more energy. It turns one of the main drawbacks of that type of radiation, namely cell damage, into an advantage, by using it to destroy cancer cells. “Irradiation” means nothing more than being bombarded with energetic light rays that destroy cells. The particle accelerator comes into play because it is what generates the radiation. It’s a lot like producing X-rays: ordinary particles are first accelerated to a high-energy state and then fired at a plate, releasing radiation that can be focused on the patient. Without a particle accelerator, there’s no radiation. That’s why every major hospital in the Netherlands has a fairly powerful particle accelerator, known as a cyclotron.

The good thing about radiation therapy is that, in the place where the tumor is located, it does destroy the DNA in the cancer cells. But it’s no secret that on its way to the tumor, the radiation can also destroy or damage a lot of healthy cells, possibly even turning them into cancer cells. We’d like to keep that to an absolute minimum, for obvious reasons. Since our many experiments over the past decades have taught us how the particles lose energy in living tissue, we know that a great deal of damage is done in the area around the tumor. Doctors try to limit the damage to healthy cells by aiming particles at the patient from many different angles, so that the area right around the tumor receives the most radiation. But sadly, there is no way to completely prevent damage to healthy cells. That’s a tragic fact for a patient with a tumor close to one eye, or near the spinal cord, but radiation therapy is also hazardous to children, because they have many cells that will go on dividing in the future.

One recent development that tackles these problems is called *hadron therapy*. It’s a new kind of radiation treatment in which the entire dose is delivered to one place, so that it does the least possible damage to healthy tissue. Thanks to particle physics, we know how radiation moves through the body, losing more and more energy along the way. It’s like a bicycle braking in the sand, which slows down fairly quickly and at an even pace, eventually coming to a stop. But we've also seen that there's another way. There are particles (protons) that pass through materials with almost no trouble at all. They slow down gradually but release their energy abruptly, all at once, when they are almost at a standstill. That sounds like just what we need: radiation therapy with protons instead of the usual photons, so that we destroy only the tumor instead of damaging all the healthy cells around it. If you know precisely how deep under the skin the tumor is found, then you can work out how fast to shoot the protons into the body so that they stop moving in that exact spot and do their damage there. And of course, you need a particle accelerator that gives you protons with exactly the right energy. Fortunately, that just happens to be the core business of particle physicists!

Right now, hadron therapy is still very expensive and available only in a couple of places. I'm no economist, so I can't calculate all the costs and benefits (and I don't want to get tangled up in a moral and ethical debate about the value of a human life, or of one extra year of it). But this is an amazing new possibility. In the years ahead, scientists will be looking to see what other innovations we can offer and how we can work together with doctors to achieve new successes.

In recent centuries, the drive to explain the things we see around us has led to a treasure trove of knowledge. Science keeps developing new techniques to overcome obstacles and delve further into the secrets of nature. In my own branch of physics, the particle accelerator proved to be the key to exploring the world of elementary particles. It allowed us to enter the world of the atom, study the atomic nucleus, and finally, with the discovery of the illustrious Higgs boson, travel to the deepest level we know of today.