

If you have read anything about quantum mechanics before, you probably expect that this book will begin by talking about particles. Most popular books on quantum mechanics talk about spooky behavior of particles that are here and not here, jump from one place to another without any cause, and so on.

There *is* something mysterious in quantum mechanics, but it actually has very little to do with the existence of particles. In talking of the philosophy of quantum mechanics, many people have focused on the topic of “wave-particle duality.” In this book, we will see that modern theory does not treat these two things on an equal footing. Waves are fundamental in quantum theory; particles are not. To use a big word, particles are “epiphenomena,” that is, theoretical constructs that are useful in some circumstances but that can be completely derived from other, more fundamental elements of the theory.

Saying something is not fundamental does not mean that it is a useless concept. For example, I could explain the sliding of different solids by invoking the concept of a friction force; the friction force is a very useful concept. But if we wanted, we could derive it from microscopic forces between atoms. Friction force is not a fundamental force of nature; it is a handy concept in many cases, such as when many atoms in two solids are sliding against each other. In other cases, if we insisted on the existence of friction force, we would run into strange conundrums. For example, if we wanted to talk about the friction force between single atoms, we would end up speaking nonsense. This often happens when a concept is pushed beyond the limits of its applicability.

Many scientists learned the particle picture of quantum mechanics in introductory physics in college and never questioned it afterwards. But anyone who has gone on to study the advanced theory of quantum mechanics taught in graduate school, known as *quantum field theory*, knows that particles arise as oscillations of the underlying fields in nature. This field theory is the most accurate theory we have for quantum mechanics; the Schrödinger single-particle equation model presented in introductory classes is known to be a simplified case of the more general field theory. Instead of discussing quantum mechanics in terms of that simplified model, let us go right to the heart of the more basic theory, the theory of fields.

1.1 Fields

Talking about fields sounds like science fiction, but it is not really so strange. The best way to imagine a field is to think about the physical system that was the historical basis




Figure 1.1

An example of a scalar field. The gray scale represents the numerical values.


of field theory: a fluid, like water. All of our mathematics for fields was developed in the early 1800s to describe the flow of fluids. This mathematical framework was then adopted by James Maxwell in the 1860s to describe the fields of electricity and magnetism, and then his concepts were eventually adapted by Paul Dirac and others in the 1920s and 1930s to give us our modern formulation of quantum mechanics. Along the way, most physics departments stopped teaching the theory of fluids (“hydrodynamics”) and now start with just Maxwell’s field theory. But thinking about water flow helps us to visualize fields.

Imagine a container filled with a fluid like water everywhere inside its volume. We can describe this in mathematical language by assigning numbers to every point inside the fluid. One example would be to write down a number for the density of the fluid at each point. The density of the fluid at all locations is an example of what is called a *scalar field* – at each point in space, there is just one number to write down, namely the density of the fluid at that point. Figure 1.1 illustrates this idea.

At every point in the fluid, we could also identify the direction and speed of its flow. To do this, we would need several numbers. One number could give us the speed, in miles per hour or meters per second, and another number or two could give us the direction. For example, we could give the angle of the flow relative to due east. If we allow up-and-down motion, we could also define the angle of ascent relative to the ground.

In this case, instead of just one number, we would have two or three numbers to assign to every location inside the fluid, to describe the flow speed and direction at each point. This set of numbers is collectively called a *vector*. But just like the case of the density, we could assign a numerical measurement to every point in space. The distribution of velocities of the fluid is what we call a *vector field*.

Figure 1.2 shows two ways of representing a vector field using pictures. In the first case, the direction of the flow is indicated by little arrows at each point in space. In the second, the direction is given by *field lines*. In both of these ways of illustrating fields, there is a lot of empty space between the arrows or lines, which is not labeled. One must imagine that, for a real field, every single point in space could have an arrow or line assigned to it. The illustrations draw the arrows or lines more sparsely just to make the picture easier to understand.

**Figure 1.2**

Two ways to draw a vector field. (a) Arrows giving the strength and direction of the flow speed at many points in space. (b) The same field represented by field lines.

Thinking about a fluid like water helps us to understand what we think of as “real.” We have seen that there are several different ways of describing a water field. One way is as a set of numbers. (In most cases, we don’t actually have to write down all the numbers individually; instead, we can write down a mathematical formula that tells us how to calculate the numbers.) For the density field of the water, another way to describe the field is with a grayscale image like that shown in Figure 1.1. For the vector field of the water flow, other ways of describing it are with little arrows or with field lines, as in Figure 1.2. In each case, what is real is the water itself and its properties, specifically density and velocity. The sets of numbers or vectors, or the pictures, are not the fundamental reality; rather, these are ways that we describe the underlying reality. We could describe the field with different numbers. For example, instead of writing the speed of the fluid in miles per hour, we could write the speed in kilometers per second, or furlongs per fortnight.

It is a main contention of this book that this distinction doesn’t change when we switch to talking about quantum fields. Just as water is a “thing” that can be described by a mathematical formalism, so the fields of quantum mechanics are “things” that exist, whether or not we choose to write down mathematical descriptions or pictures for them. This is not always how physicists and philosophers talk, but, as we will see, if we accept the notion of the reality of water flow, then it is artificial to drop the notion of reality for other types of fields.

Electric and magnetic fields. In the early 1800s, the famous experimental physicist Michael Faraday made the intellectual leap that electricity and magnetism could be described by fields very similar to the flow of fluids; he actually drew pictures of electric and magnetic fields circulating around in space like water. Figure 1.3 shows how a magnetic field can be seen directly in the alignment of small iron filings. The somewhat disturbing implication of Faraday’s work is that electric and magnetic fields flow through all space, even through matter such as our bodies.




Figure 1.3

Magnetic field lines from a bar magnet (placed under the board) seen in the alignment of small iron filings. (Windell H. Oskay (www.evilmadscientist.com).)

The Scottish physicist James Clerk Maxwell later succeeded in writing down mathematical equations that accurately described this flow. Figure 1.4 gives these equations, but there is no need to understand these equations to see their main feature: they are very short and compact, written in just four lines. Maxwell's equations remain to this day one of the most impactful and elegant successes of physics. Although these equations are short, they fully describe a huge range of effects of electricity, magnetism, light and optics, X-ray radiation and gamma rays, infrared heat radiation, microwaves, and radio waves. They led directly to the technology of radio and television which revolutionized communications. They also needed no corrections when Einstein's theory of relativity came along. In fact, the accuracy of Maxwell's equations played a major role in inspiring Albert Einstein to come up with his theory of relativity; Einstein wanted to keep Maxwell's equations the same for all observers in the universe. In a way, Einstein was just mopping up the implications of what Maxwell wrote down. Not only that, Maxwell's field theory is not changed significantly by

$$\begin{aligned}
 \nabla \cdot \vec{E} &= \frac{\rho}{\epsilon_0} \\
 \nabla \cdot \vec{B} &= 0 \\
 \nabla \times \vec{E} &= -\frac{\partial \vec{B}}{\partial t} \\
 \nabla \times \vec{B} &= \mu_0 \epsilon_0 \frac{\partial \vec{E}}{\partial t} + \mu_0 \vec{J}
 \end{aligned}$$

Figure 1.4

Maxwell's equations for electric field \vec{E} and magnetic field \vec{B} , for a charge density ρ and current density \vec{J} . The values of ϵ_0 and μ_0 are universal constants that together give the speed of light, according to $c = 1/\sqrt{\epsilon_0 \mu_0}$.

quantum mechanics. In quantum theory, we learn that Maxwell's equations are correct for a special class of fields known as *coherent* fields.

Every physicist today learns Maxwell's equations as an example of a field theory. Just as we can assign a direction or density to water flow, these equations assign a direction and strengths of the electric and magnetic fields at every point in space. The electric and magnetic fields at every point in space (including inside your body) make up together what is called the *electromagnetic field*.

Understanding the nature of the electromagnetic field is one of the first major conceptual leaps that modern physics requires. Are electromagnetic fields “real”? In the case of water flow, the field is obviously “real” by any normal definition: we can look at the surface of a river and see the flow of the water. That is, we can see its velocity field directly. In the case of the electromagnetic field, the concept is a little more difficult. What we measure are forces; for example, the electric force needed to deflect a pointer in a meter – the same electric force you feel when you rub a balloon on your hair and stick it to a wall. These forces always act on objects. One could therefore argue that only the objects and forces are real, and the field in the empty space between them is just a mathematical trick for keeping track of the complicated forces between the objects. Some people do argue this way (e.g., Mead 2000). But the vast majority of scientists today view electromagnetic field as “real.” In fact, their sense of the reality of forces and fields is inverted: the field is the real thing, which flows throughout all space, and the forces we measure are just the specific effects of that field in particular times and places.

One reason for seeing the electromagnetic field as real is that electric and magnetic fields can carry energy and pressure through the vacuum of empty space. We all have a very immediate experience of this. The sun sends light radiation through the vacuum of outer space to us every day, and this is responsible for the warmth of the earth, and our life. A laser beam traveling through outer space is also an example of the field carrying energy and pressure. You might not think that light exerts pressure, but it does – a laser beam, and even a light bulb, pushes any object it hits, with a tiny pressure. Some scientists have proposed using vast sails in outer space to use the pressure from sunlight to sail around the solar system.

In the 1800s, it bothered people to think about how the vacuum of outer space could carry energy and pressure. Some argued that the particle picture is needed to understand this. In this view, light particles (photons) travel through outer space like little bullets. It is easy to imagine energy and momentum being carried by these little particles, and pressure coming from them when they hit something.

But the particle picture is not necessary for understanding how light carries energy and momentum through space. All the theory we need is in Maxwell's equations, which tell us exactly how much energy and pressure is carried by light waves through vacuum. The electromagnetic field carries these just as a water wave does, and Maxwell used the mathematics of fluid pressure to describe how electromagnetic waves do this.

Quantum fields. The next conceptual leap of modern physics is to envision the field that underlies everything we think of as normal matter. Just as there is an electromagnetic field that fills all of space everywhere, the fundamental theory of quantum mechanics says that there is also a “matter field” that extends everywhere in space. This field is described

by equations, based largely on the work of Paul Dirac, that are very similar in form to Maxwell's equations.

In classical electromagnetic field theory, the field at each point in space is determined by the electric and magnetic forces there. This has the advantage that we can equate the numbers of the field to physical measurements that we could have made, at least in principle. In quantum field theory, there is an additional conceptual hurdle, which is that the number value given to the field at each point doesn't even correspond to a force. We can only deduce the values of these numbers indirectly. For this reason, some physicists who accept that the electromagnetic field is real get off the boat at this point and say that the matter field is not real.

But in quantum field theory, the equations that describe the electromagnetic field and the equations that describe the matter field *have the same fundamental nature*.¹ We will look into this in more detail in the coming chapters; for those who are mathematically inclined, the basic math is given in Part III of this book. The general point is this, however: we have an ascending ladder of fields introduced in physics, from water currents and air flow, to electric and magnetic forces, to matter fields, which have exactly the same type of math. We are comfortable with calling the water currents "real," and most physicists are comfortable with calling electromagnetic fields "real." Why then make a cutoff and refuse to call quantum matter fields real, when the basic structure of the mathematics is the same, and the way the equations predict the results of experiments are the same? It is true that, just because some equations are formally similar, we need not take them as describing the same thing; science has many examples of equations that are similar but describe different things. But in the case of matter fields and classical fields, not only are the mathematical structures the same; the elements in the different fields are interchangeable – matter fields generate light and sound fields, and vice versa. (We will come back to this in Section 4.5.) And all classical fields can be deduced as special cases from quantum fields. There is simply no fundamental reason to draw a sharp line between the nature of physical reality of different types of fields.

¹ The only significant difference is that what we typically think of as matter fields are *fermion* fields, while electromagnetic fields are *boson* fields. The distinction between these two will be discussed in Section 2.4; in the math, the only difference is a change of a single sign from positive to negative. Classical fields such as water velocity fields (described by the *Navier–Stokes* equations) and electromagnetic fields (described by Maxwell's equations) are now understood to be special cases of the more general formalism for quantum boson fields, known as *coherent states*, discussed mathematically in Section 12.5.

A commonly repeated statement is that the mathematical structure of quantum fields involves *operators*, while classical fields do not, and that this gives an essential difference between the two (the mathematics of operators are discussed in Section 11.1). This is actually a misunderstanding based on different uses of the term "field" in the literature. Classical and quantum fields can both be acted on by mathematical operators, and in both cases, the operators by themselves do not carry any information about actual physical states. When the operators have a spatial dependence, one can talk of an "operator field," but this is not a physical entity for either classical or quantum fields. For the mathematical details of this, see the discussion of spatial field operators in Section 12.2.3.

Others have argued (e.g., Griffiths 2003) that the fact that matter waves involve complex numbers means that they are fundamentally different from electromagnetic waves. But as shown mathematically in Section 13.1.2, complex numbers are just a useful bookkeeping device to keep track of two degrees of freedom of a fermion field. Section 9.1 also addresses this view.

The perspective shared by most physicists familiar with quantum theory is that matter fields really do exist. All the classical fields we are familiar with, such as the electric fields that raise our hair in the presence of static electricity, and currents in water, can be derived from underlying quantum fields. The quantum fields are more fundamental, or one might say “more real,” than electric forces and water currents.

Quantum field theory is not some alternative version of quantum mechanics. It is the most generally accepted, basic theory of quantum mechanics as understood by experts. And it makes no sharp distinction between light waves, electron waves, and water waves.

1.2 Waves

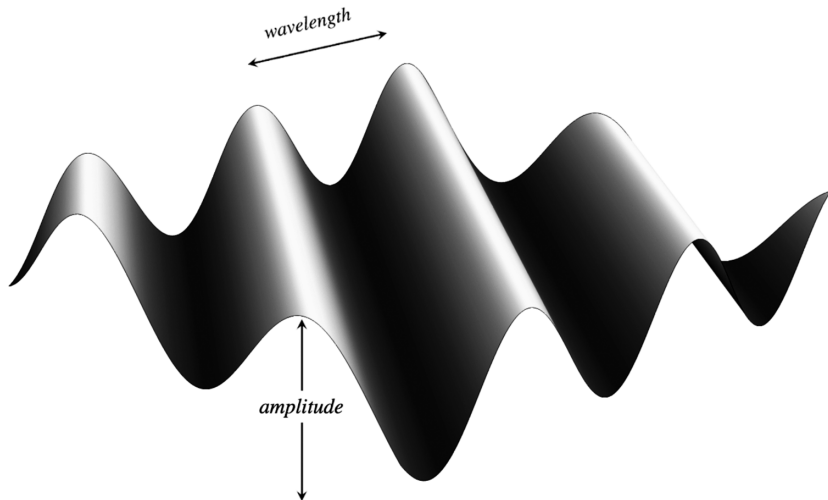
The previous section introduced the concept of a *field*. The next important concept is a *wave*. Just as we defined a field with the concrete example of a fluid, we can use a very tangible example to define a wave, which is just what its name sounds like: a water wave. If you go to the ocean and see waves, or if you make waves in a swimming pool, you know what a wave is. The important thing to keep in mind about waves is that energy and pressure are carried by waves from one place to another even though the water itself goes nowhere. Think about it: if someone were sitting on the opposite side of a swimming pool, you could push the water, creating a wave, and the wave would move to that person and hit them. Yet the water level in the pool would stay the same. You put energy and pressure into the water, and the water carried that energy and momentum somewhere else, even while the water itself stayed in the same place.

The water in this case is the fluid field we discussed in the previous section. The wave is an oscillation of that field. In the same way, the electromagnetic field and the quantum matter field can have wave oscillations.

Let us define some generic terms used for all waves. Table 1.1 gives a summary of these. The first important term is the *wavelength* of a wave, which is the distance between

Table 1.1 Standard wave terms.

Name	Usual symbol	Definition	Measurement
Wavelength	λ	Distance between the crests of a wave	Measured as a length
Frequency	f	Number of wave crests passing a point per second	“Hz” = number per second
Phase	ϕ	Degree of completion of the oscillation cycle of a wave	Ranges from 0 to 360°
Amplitude	ψ	Range of variation of whatever is oscillating	Different for each field

**Figure 1.5**

A typical wave, with the definition of two terms for describing the wave.

its crests, as illustrated in Figure 1.5. The next is the *frequency* of a wave, which is the number of wave crests moving past some location within a period of time. The frequency and wavelength of a wave are not independent of each other: if we pick one, then we can deduce the other from the properties of the field. A fast oscillation has high frequency (many wave crests per second) and short wavelength; a slow oscillation has low frequency and long wavelength. Frequency is measured in crests per second; instead of saying “crests per second,” scientists and engineers use the term *Hertz*, after the famous scientist of that name, and abbreviate this as *Hz*. The oscillating electric wave in your home electric circuits has a frequency of 60 Hz (60 crests per second), while radio waves have frequencies of millions of crests per second, that is, mega-Hertz (written in the metric system as MHz). Sound waves in the audible range have frequencies of tens of Hz up to a few thousand Hz, and ultrasound waves, used for seeing inside a human body, have frequencies around 1 MHz.

The electromagnetic field can oscillate just like the surface of a swimming pool. When the electromagnetic field oscillates, we detect it as light waves, radio waves, or microwaves, and so on. The only difference between these types of waves is the *frequency* of the waves. Radio waves have much longer wavelength than light waves but are otherwise just the same.

The quantum matter field can also oscillate. Matter waves can have short wavelength or long wavelength. In general, it takes more energy to create a wave with high frequency than it does to create a slow oscillation with low frequency.

In each type of wave, we call the range of oscillation the *amplitude* of the wave, as illustrated in Figure 1.5. In the case of water waves, it is easy to see the amplitude: as the water goes up and down, the amplitude is the height of the waves. It is not so easy to visualize the oscillation of electromagnetic waves or matter waves. But something is really oscillating. In the case of an electromagnetic wave, the strength and direction of the electric




Figure 1.6

The phase of a wave equated to the position of a clock hand. The angle θ is the *phase angle* corresponding to a point in the cycle.

force oscillates. If you had a small charged object, it would feel an oscillating force if an electromagnetic wave passed by. The force would change direction from left to right, or up to down, and then back the other way. The amplitude in this case would be measured as the maximum force felt by the charged object. In the case of a matter wave, nothing is moving left to right or up to down. But something is still oscillating.

One last general wave term is the *phase*. This concept, like many other scientific concepts, suffers from degradation in the popular media: “set phasors to stun.” Phase is actually very simple. Consider the left side of Figure 1.6, with a clock hand that can go around a circle. The phase of the clock hand is its position in the circle. This can be measured as an angle. For example, we could define 12:00 PM as angle zero, 3:00 PM as the 90° angle, 6:00 PM as 180° , and so on. The phase goes all the way around to 360° and then starts over again.

We can extend this definition of phase to any cyclic oscillation. (This is why we talk of the phases of the moon, in its cycle from new to full and back.) With a water wave, the wave goes up and down. We can define the highest point as phase of zero, and the lowest point, when it is halfway through its cycle, as phase of 180° ; when it gets back to the highest point, the phase reaches 360° and starts over. Figure 1.6 illustrates how the position in an oscillation can be equated to the position of a clock hand. The water is not really moving in a circle; it is moving up and down, but since it is moving through a periodic cycle, we can equate the different stages in its cycle as different phase angles around a circle.

In a matter wave, the phase of the matter field also oscillates in a cycle, so that its phase goes from 0 to 180° , 360° , and so on. The phase of a matter wave is not just a fictional concept. It can be measured in experiments, for example, in the current of superconductors and in small (“mesoscopic”) electronic circuits. One way this is done is with interference and tunneling experiments, which are described in Section 1.3.

1.3 Basic Wave Effects

All of these waves in fields, whether we are talking about water waves, sound waves (oscillations of the air), electromagnetic waves (light, radio, etc.), or matter waves, have certain

common properties. These properties all seem normal when we are thinking about typical “energy” waves like sound and light but can seem strange when we are talking about matter.

The first basic wave property is called *superposition*. This property means that two waves can pass through each other, and, while they are passing through each other, the behavior of the whole system is just the sum of the two waves.

This property is not absolute. It breaks down when there are “nonlinear” terms in the wave equations, and, in almost all real systems, there are some nonlinear terms that crop up in some cases. Many modern physics experiments study these nonlinear effects. But for the most part, we can assume that the fields are linear and therefore the principle of superposition is true. As discussed in Section 2.3, nonlinear effects are one reason why the notion of indivisible quantum particles breaks down in some cases.

Waves passing through each other seem normal for waves that we think of as pure energy waves. For example, it doesn't seem strange that a beam of light from one flashlight can pass through the beam of light from another flashlight, or that a water wave created by a person splashing on one side of a pool can pass through the wave created by another person on a different side of the pool. Our normal experience, however, is that solid matter does not pass through other matter. Matter waves passing through other matter waves seem very strange. It can happen, however. The reason why it does not happen more often has to do with the specific property of matter waves called *Pauli exclusion*, which we will discuss in Chapter 2.

Given the property of superposition of waves, a direct consequence is the possibility of *interference* of waves. Interference occurs when one wave cancels out another wave. Suppose that two water waves are passing through each other, as shown in Figure 1.7. At some points, the wave crest from one wave is moving up while the crest from the other wave is moving down. The two waves will cancel at those points, leaving the water motionless there. At other points, the two waves will add up, to give a wave higher than either. You may not have seen interference of water waves (though it is easy to demonstrate), but you have certainly seen it with light waves, and may have heard it with sound waves. If you have seen rings of different colors on an oil slick, as shown in Figure 1.8, you have seen interference





Figure 1.7

Computed interference pattern of two waves passing through each other.

**Figure 1.8**

Pictures of interference in the real world. (a) Fringes of light reflected from oil, caused by the interference of light from the top surface and the bottom surface of the oil, with slight thickness variations of the oil (Creative Commons Attribution – Share Alike 2.5 Generic). (b) Interference patterns of matter waves composed of whole atoms, measured at very low temperature. From Andrews 1997.

of light. What is happening is that the light waves bouncing off the top surface of the oil are interfering with light waves bouncing off the bottom surface, so that in some places the light of a certain color cancels out. Sound can do the same thing. There is now an industry selling active sound-reduction headphones. These generate sound exactly opposite of the sound coming in, canceling it out and creating silence.

Again, this doesn't seem so strange with light and sound waves, but very strange for matter. Imagine taking a solid object and canceling it out with another object, so that you have nothing. Yet experiments have been done that do essentially that: electronic current can be canceled out through interference. When electric current sent down one wire is sent into the same place as electric current sent down another wire, for certain experimental conditions, they can cancel out, giving no current. This is not science fiction; it is reality, and the




Figure 1.9

A vibrational resonance of a string produced by an oscillator at one end. The frequency is such that three half-wavelengths fit exactly into the available space. (Greg Severn, University of San Diego.)

basis of useful devices for measuring magnetic field, as discussed further in Sections 2.6 and 8.5. More recently, whole atoms have been canceled out, as shown in Figure 1.8(b), to make fringes similar to the ones you see on an oil slick.

Another wave effect that will be extremely important in all of this book is *resonance*. Resonance means that some things have “natural” frequencies that they vibrate at. This is familiar to all musicians. For example, a tuning fork resonates at a certain frequency if you hit it, as does a stretched piano wire on a piano. Figure 1.9 shows an example of a resonant mode of a taut string.

An object can have more than one resonance. For example, if you blow into a bugle, you can get different notes by blowing differently; the different notes correspond to different resonances of the sound wave in the air in the bugle. In the same way, a tuning fork or a bell can have “overtones” (higher-frequency resonances) if you hit it hard enough.

A crucial property of wave resonances is that they can be sharply defined, with jumps in between. Blowing into a bugle, you get distinct notes, not a continuous smear of sound. Because resonances can be sharply defined, and the system can stay in a resonance for a long time, they are often called stable *states* of the system. There is always one lowest-frequency resonance, which corresponds to the longest wavelength; this is called the *ground state* of the system. Higher-energy states, which have shorter wavelength and higher frequency, are called *excited states*.

In quantum mechanics, not only sound waves and electromagnetic waves but also matter waves can have resonances. The famous “jumps” between electronic states of the atom, which we learn in chemistry class, correspond to transitions between resonances of the electron matter wave. These jumps are not instantaneous, however; we will return to discuss this in Chapter 3.

Resonances are not always sharply defined. If there is energy loss in the system, the resonances may be smeared out. For example, when you tune a radio, you are tuning it to resonate at different frequencies of radio waves. Each radio station sends out its signal at a different frequency; this is what the call numbers stand for (e.g., 104.5 MHz is 104.5 million wave crests per second). In a cheap radio, the resonances of the radio are not very




Figure 1.10

A typical resonance curve. The ratio of the center frequency to the width of the resonance is called its Q-factor.

distinct, and you can pick up more than one radio station at the same time. Figure 1.10 shows some of the characteristics of a resonance.

One more basic wave effect is *partial transmission*. You are quite familiar with this effect of waves: a person may be speaking in another room, and you hear the sound coming through the wall. The sound is quieter than it would be if you were in the same room as that person, because only part of the sound gets through the wall and the rest is reflected. The same thing happens with light: if light hits a glass surface, part of it goes through, and part of it is reflected, as shown in Figure 1.11. The same thing happens with the windows in your house.

Think how strange this would be if normal matter acted this way in daily life. Imagine a person hitting a wall, and one copy of the person passing through the wall while another copy of the person bounces off! We are used to matter being indivisible, going one way or




Figure 1.11

A laser beam being split into two by a glass surface. (C.-M. Zetterling, Royal Institute of Technology, Sweden.)

another but not both. We can break an object into two parts, but we never see one object that does two things simultaneously.

Yet this effect does happen with matter, also. When matter is partially transmitted through a barrier, it is known as *tunneling*. This has been seen in all kinds of experiments; for example, electrons in a superconductor can tunnel through barriers, and nuclear matter inside an atom can tunnel out, leading to the well-known effect of nuclear radiation. We will return to discuss this in Chapter 8.

1.4 The Return of the Ether

In the case of water waves, it is clear what is oscillating, or waving: the water. What is waving in the case of electromagnetic waves or matter waves? This question bothered many people in the 1800s. The answer, as understood in modern physics, is clear: the field is oscillating. Both the electromagnetic field and the quantum matter field, which exist everywhere in space, can oscillate in a measurable way.

This sounds a lot like saying that the “ether” oscillates. The concept of the *ether* is a very ancient concept, which scholars envisioned as a fluid filling all space. (From this term, we get the word “ethereal,” referring to something invisible and immaterial.) Many people who have read about philosophy of science have heard that the concept of the ether was put to death at the end of the 1800s, and now nobody believes in it. But actually, scientists have just substituted another word for the ether, and called it the “field”!

Let us take a short excursion into basic logic. A longstanding conclusion of pure logic is the principle that “nothing cannot exist.” This is a tautology: by using the verb “exist,” we imply that something exists which is doing the existing. Therefore, to exist is to *be something*. But if that something exists, then it is not “nothing.” Thus, to say that “nothing” exists is to engage in self-contradiction, to say that something exists, but it doesn’t.

When we talk about the vacuum of outer space, we mean that there is no *mass* there, but it is not “nothing.” There is always a *field* there; specifically, both the electromagnetic field and quantum matter field. These fields may have no waves happening there, but the fields still exist, just as in a pool with no waves, the water still exists.

So throughout all space, “something” exists, which we call a “field,” and that field can oscillate. We can call this the ether, if we want. Why then did the old concept of the ether get rejected?

Essentially, the old ether concept that got rejected was one that was pinned to a static concept of space and time. Einstein showed that our sense of time passing and our sense of distance depend on our point of view, and people could watch the same events and disagree about when and where they took place, even while agreeing that the same events happened.

The old concept of the ether was, in effect, nailed down to space, like the water in a lake. If you are in a boat and you go fast enough, you can measure your speed relative to the water, and you can even go faster than the waves you create. Every boater is familiar with this effect: as the boat speeds up, the waves stretch out to the rear. Waves only advance out in front of the boat if it is going slower than the waves do.

If light waves were like that, then if you were on a rocket going near the speed of light, and you shone a flashlight forward, you would see the light piling up at the front of the rocket, like water waves piling up at the front of a speeding boat. Einstein was specifically concerned to make sure that his theory did not allow this to happen, because he was certain that Maxwell's equations are universally true from every perspective in the universe. He based this on the intuition that in the vacuum of outer space, far away from any matter, there would be no way to know whether you were moving or not.

When the proper math of Einstein's relativity is taken into account, the result is that no matter how fast you go, light waves travel away from you at the same speed in all directions, and therefore you can't measure your speed relative to the field by looking at the light waves. The same is true for the quantum matter field: both the electromagnetic and quantum matter fields are *relativistic* fields. That means that the field equations are the same no matter how fast you go. (This incidentally implies that the speed of light which goes into the equations must always be the same, a premise that is often used in calculations of relativity.)

It is important to understand what relativity says and what it does not say. It does not say that anyone can believe anything he or she wants. It says that every observer everywhere can compute the numerical values of the fields using the same set of equations. If they are traveling at different speeds, they will come up with different numbers for the field values. But the laws of relativity give very definite rules for how to transform the field values for an observer at one time and place to the field values measured by an observer at another location, traveling at a different speed.

Thus, the field is always still there – it does not vanish for any observer. Every observer can detect a field filling all of space, and every observer can perceive that the field can oscillate, leading to waves. So it is very much like the ether of old. It is just not a nailed-down ether. It looks different from different perspectives, but it is always there. Physicists agree that the field is everywhere present, but some do not like to call the field a “medium,” because, by their definition, a medium is something pinned down, like the old version of the ether. But if we define a medium as simply whatever it is that is oscillating, then it is perfectly appropriate to say that the electromagnetic and quantum matter fields are the mediums for waves.

Not only is the vacuum field not “nothing,” but, in the modern view, the vacuum field is quite active. The remnants of the light emitted at the Big Bang still fill all of outer space, which means that there are electromagnetic waves bouncing around everywhere. The vacuum is filled with electromagnetic and matter fields, and these fields are as real and active as any physical entity we know.

References

- M. R. Andrews, C. G. Townsend, H.-J. Miesner, et al., “Observation of interference between two Bose condensates,” *Science* **275**, 637 (1997).
R. Griffiths, *Consistent Quantum Theory*, (Cambridge University Press, 2003).
C. Mead, *Collective Electrodynamics*, (MIT Press, 2000).