



Introduction

Alix Maunier was very curious about physics and often asked me to explain things to her. I would do all right, just as I do with a group of students at Caltech that come to me for an hour on Thursdays, but eventually I'd fail at what is to me the most interesting part: We would always get hung up on the crazy ideas of quantum mechanics. I told her I couldn't explain these ideas in an hour or an evening—it would take a long time—but I promised her that someday I'd prepare a set of lectures on the subject.

I prepared some lectures, and I went to New Zealand to try them out—because New Zealand is far enough away that if they weren't successful, it would be all right! Well, the people in New Zealand thought they were okay, so I guess they're okay—at least for New Zealand! So here are the lectures I really prepared for Alix, but unfortunately I can't tell them to her directly, now.

What I'd like to talk about is a part of physics that is *known*, rather than a part that is unknown. People are always asking for the latest developments in the unification of this theory with that theory, and they don't give us a chance to tell them anything about one of the theories that we know pretty well. They always want to know things that we don't know. So, rather than confound you with a lot of half-cooked, partially analyzed theories, I would like to tell you about a subject that has been very thoroughly analyzed.

I love this area of physics and I think it's wonderful: it is called quantum electrodynamics, or QED for short.

My main purpose in these lectures is to describe as accurately as I can the strange theory of light and matter—or more specifically, the interaction of light and electrons. It's going to take a long time to explain all the things I want to. However, there are four lectures, so I'm going to take my time, and we will get everything all right.

Physics has a history of synthesizing many phenomena into a few theories. For instance, in the early days there were phenomena of motion and phenomena of heat; there were phenomena of sound, of light, and of gravity. But it was soon discovered, after Sir Isaac Newton explained the laws of motion, that some of these apparently different things were aspects of the same thing. For example, the phenomena of sound could be completely understood as the motion of atoms in the air. So sound was no longer considered something in addition to motion. It was also discovered that heat phenomena are easily understandable from the laws of motion. In this way, great globs of physics theory were synthesized into a simplified theory. The theory of gravitation, on the other hand, was not understandable from the laws of motion, and even today it stands isolated from the other theories. Gravitation is, so far, not understandable in terms of other phenomena.

After the synthesis of the phenomena of motion, sound, and heat, there was the discovery of a number of phenomena that we call electrical and magnetic. In 1873 these phenomena were synthesized with the phenomena of light and optics into a single theory by James Clerk Maxwell, who proposed that light is an electromagnetic wave. So at that stage, there were the laws of motion, the laws of electricity and magnetism, and the laws of gravity.

Around 1900 a theory was developed to explain what matter was. It was called the electron theory of matter, and

it said that there were little charged particles inside of atoms. This theory evolved gradually to include a heavy nucleus with electrons going around it.

Attempts to understand the motion of the electrons going around the nucleus by using mechanical laws—analogs to the way Newton used the laws of motion to figure out how the earth went around the sun—were a real failure: all kinds of predictions came out wrong. (Incidentally, the theory of relativity, which you all understand to be a great revolution in physics, was also developed at about that time. But compared to this discovery that Newton's laws of motion were quite wrong in atoms, the theory of relativity was only a minor modification.) Working out another system to replace Newton's laws took a long time because phenomena at the atomic level were quite strange. One had to lose one's common sense in order to perceive what was happening at the atomic level. Finally, in 1926, an "uncommon-sense" theory was developed to explain the "new type of behavior" of electrons in matter. It looked cockeyed, but in reality it was not: it was called the theory of quantum mechanics. The word "quantum" refers to this peculiar aspect of nature that goes against common sense. It is this aspect that I am going to tell you about.

The theory of quantum mechanics also explained all kinds of details, such as why an oxygen atom combines with two hydrogen atoms to make water, and so on. Quantum mechanics thus supplied the theory behind chemistry. So, fundamental theoretical chemistry is really physics.

Because the theory of quantum mechanics could explain all of chemistry and the various properties of substances, it was a tremendous success. But still there was the problem of the interaction of light and matter. That is, Maxwell's theory of electricity and magnetism had to be changed to be in accord with the new principles of quantum mechanics that had been developed. So a new theory, the quantum

theory of the interaction of light and matter, which is called by the horrible name “quantum electrodynamics,” was finally developed by a number of physicists in 1929.

But the theory was troubled. If you calculated something roughly, it would give a reasonable answer. But if you tried to compute it more accurately, you would find that the correction you thought was going to be small (the next term in a series, for example) was in fact very large—in fact, it was *infinitely*! So it turned out you couldn’t really compute *anything* beyond a certain accuracy.

By the way, what I have just outlined is what I call a “physicist’s history of physics,” which is never correct. What I am telling you is a sort of conventionalized myth-story that the physicists tell to their students, and those students tell to their students, and is not necessarily related to the actual historical development, which I do not really know!

At any rate, to continue with this “history,” Paul Dirac, using the theory of relativity, made a relativistic theory of the electron that did not completely take into account all the effects of the electron’s interaction with light. Dirac’s theory said that an electron had a magnetic moment—something like the force of a little magnet—that had a strength of exactly 1 in certain units. Then in about 1948 it was discovered in experiments that the actual number was closer to 1.00118 (with an uncertainty of about 3 on the last digit). It was known, of course, that electrons interact with light, so some small correction was expected. It was also expected that this correction would be understandable from the new theory of quantum electrodynamics. But when it was calculated, instead of 1.00118 the result was infinity—which is wrong, experimentally!

Well, this problem of how to calculate things in quantum electrodynamics was straightened out by Julian Schwinger, Sin-Itiro Tomonaga, and myself in about 1948. Schwinger was the first to calculate this correction using a new “shell

game”, his theoretical value was around 1.00116, which was close enough to the experimental number to show that we were on the right track. At last, we had a quantum theory of electricity and magnetism with which we could calculate! This is the theory that I am going to describe to you.

The theory of quantum electrodynamics has now lasted for more than fifty years, and has been tested more and more accurately over a wider and wider range of conditions. At the present time I can proudly say that there is *no significant difference* between experiment and theory!

Just to give you an idea of how the theory has been put through the wringer, I’ll give you some recent numbers: experiments have Dirac’s number at 1.00115965221 (with an uncertainty of about 4 in the last digit); the theory puts it at 1.00115965246 (with an uncertainty of about five times as much). To give you a feeling for the accuracy of these numbers, it comes out something like this: If you were to measure the distance from Los Angeles to New York to this accuracy, it would be exact to the thickness of a human hair. That’s how delicately quantum electrodynamics has, in the past fifty years, been checked—both theoretically and experimentally. By the way, I have chosen only one number to show you. There are other things in quantum electrodynamics that have been measured with comparable accuracy, which also agree very well. Things have been checked at distance scales that range from one hundred times the size of the earth down to one-hundredth the size of an atomic nucleus. These numbers are meant to intimidate you into believing that the theory is probably not too far off! Before we’re through, I’ll describe how these calculations are made.

I would like to again impress you with the vast range of phenomena that the theory of quantum electrodynamics describes: It’s easier to say it backwards: the theory de-

scribes *all* the phenomena of the physical world except the gravitational effect, the thing that holds you in your seats (actually, that's a combination of gravity and politeness, I think), and radioactive phenomena, which involve nuclei shifting in their energy levels. So if we leave out gravity and radioactivity (more properly, nuclear physics), what have we got left? Gasoline burning in automobiles, foam and bubbles, the hardness of salt or copper, the stiffness of steel. In fact, biologists are trying to interpret as much as they can about life in terms of chemistry, and as I already explained, the theory behind chemistry is quantum electrodynamics.

I must clarify something: When I say that all the phenomena of the physical world can be explained by this theory, we don't really know that. Most phenomena we are familiar with involve such *tremendous* numbers of electrons that it's hard for our poor minds to follow that complexity. In such situations, we can use the theory to figure roughly what ought to happen and that *is* what happens, roughly, in those circumstances. But if we arrange in the laboratory an experiment involving just a *few* electrons in *simple* circumstances, then we can calculate what might happen very accurately, and we can measure it very accurately, too. Whenever we do such experiments, the theory of quantum electrodynamics works very well.

We physicists are always checking to see if there is something the matter with the theory. That's the game, because if there *is* something the matter, it's interesting! But so far, we have found nothing wrong with the theory of quantum electrodynamics. It is, therefore, I would say, the jewel of physics—our proudest possession.

The theory of quantum electrodynamics is also the prototype for new theories that attempt to explain nuclear phenomena, the things that go on inside the nuclei of atoms. If one were to think of the physical world as a stage,

then the actors would be not only electrons, which are outside the nucleus in atoms, but also quarks and gluons and so forth—dozens of kinds of particles—inside the nucleus. And though these “actors” appear quite different from one another, they all act in a certain style—a strange and peculiar style—the “quantum” style. At the end, I'll tell you a little bit about the nuclear particles. In the meantime, I'm only going to tell you about photons—particles of light—and electrons, to keep it simple. Because it's the *way* they act that is important, and the way they act is very interesting.

So now you know what I'm going to talk about. The next question is, will you *understand* what I'm going to tell you? Everybody who comes to a scientific lecture knows they are not going to understand it, but maybe the lecturer has a nice, colored tie to look at. Not in this case! (Feynman is not wearing a tie.)

What I am going to tell you about is what we teach our physics students in the third or fourth year of graduate school—and you think I'm going to explain it to you so you can understand it? No, you're not going to be able to understand it. Why, then, am I going to bother you with all this? Why are you going to sit here all this time, when you won't be able to understand what I am going to say? It is my task to convince you *not* to turn away because you don't understand it. You see, my physics students don't understand it either. That is because *I* don't understand it. Nobody does.

I'd like to talk a little bit about understanding. When we have a lecture, there are many reasons why you might not understand the speaker. One is, his language is bad—he doesn't say what he means to say, or he says it upside down—and it's hard to understand. That's a rather trivial matter, and I'll try my best to avoid too much of my New York accent.

Another possibility, especially if the lecturer is a physicist, is that he uses ordinary words in a funny way. Physicists often use ordinary words such as “work” or “action” or “energy” or even, as you shall see, “light” for some technical purpose. Thus, when I talk about “work” in physics, I don’t mean the same thing as when I talk about “work” on the street. During this lecture I might use one of those words without noticing that it is being used in this unusual way. I’ll try my best to catch myself—that’s my job—but it is an error that is easy to make.

The next reason that you might think you do not understand what I am telling you is, while I am describing to you *how* Nature works, you won’t understand *why* Nature works that way. But you see, nobody understands that. I can’t explain why Nature behaves in this peculiar way.

Finally, there is this possibility: after I tell you something, you just can’t believe it. You can’t accept it. You don’t like it. A little screen comes down and you don’t listen anymore. I’m going to describe to you how Nature is—and if you don’t like it, that’s going to get in the way of your understanding it. It’s a problem that physicists have learned to deal with: They’ve learned to realize that whether they like a theory or they don’t like a theory is *not* the essential question. Rather, it is whether or not the theory gives predictions that agree with experiment. It is not a question of whether a theory is philosophically delightful, or easy to understand, or perfectly reasonable from the point of view of common sense. The theory of quantum electrodynamics describes Nature as absurd from the point of view of common sense. And it agrees fully with experiment. So I hope you can accept Nature as She is—absurd.

I’m going to have fun telling you about this absurdity, because I find it delightful. Please don’t turn yourself off because you can’t believe Nature is so strange. Just hear me all out, and I hope you’ll be as delighted as I am when we’re through.

How am I going to explain to you the things I don’t explain to my students until they are third-year graduate students? Let me explain it by analogy. The Maya Indians were interested in the rising and setting of Venus as a morning “star” and as an evening “star”—they were very interested in when it would appear. After some years of observation, they noted that five cycles of Venus were very nearly equal to eight of their “nominal years” of 365 days (they were aware that the true year of seasons was different and they made calculations of that also). To make calculations, the Maya had invented a system of bars and dots to represent numbers (including zero), and had rules by which to calculate and predict not only the risings and settings of Venus, but other celestial phenomena, such as lunar eclipses.

In those days, only a few Maya priests could do such elaborate calculations. Now, suppose we were to ask one of them how to do just one step in the process of predicting when Venus will next rise as a morning star—subtracting two numbers. And let’s assume that, unlike today, we had not gone to school and did not know how to subtract. How would the priest explain to us what subtraction is?

He could either teach us the numbers represented by the bars and dots and the rules for “subtracting” them, or he could tell us what he was really doing: “Suppose we want to subtract 236 from 584. First, count out 584 beans and put them in a pot. Then take out 236 beans and put them to one side. Finally, count the beans left in the pot. That number is the result of subtracting 236 from 584.”

You might say, “My Quetzalcoatl! What *tedium*—counting beans, putting them in, taking them out—what a job!”

To which the priest would reply, “That’s why we have the rules for the bars and dots. The rules are tricky, but they are a much more efficient way of getting the answer than by counting beans. The important thing is, it makes no difference as far as the *answer* is concerned: we can

predict the appearance of Venus by counting beans (which is slow, but easy to understand) or by using the tricky rules (which is much faster, but you must spend years in school to learn them)."

To understand *how* subtraction works—as long as you don't have to actually carry it out—is really not so difficult. That's my position: I'm going to explain to you what the physicists are *doing* when they are predicting how Nature will behave, but I'm not going to teach you any tricks so you can do it *efficiently*. You will discover that in order to make any reasonable predictions with this new scheme of quantum electrodynamics, you would have to make an awful lot of little arrows on a piece of paper. It takes seven years—four undergraduate and three graduate—to train our physics students to do that in a tricky, efficient way. That's where we are going to skip seven years of education in physics: By explaining quantum electrodynamics to you in terms of what we are *really doing*, I hope you will be able to understand it better than do some of the students!

Taking the example of the Maya one step further, we could ask the priest *why* five cycles of Venus nearly equal 2,920 days, or eight years. There would be all kinds of theories about *why*, such as, "20 is an important number in our counting system, and if you divide 2,920 by 20, you get 146, which is one more than a number that can be represented by the sum of two squares in two different ways," and so forth. But that theory would have nothing to do with Venus, really. In modern times, we have found that theories of this kind are not useful. So again, we are not going to deal with *why* Nature behaves in the peculiar way that She does; there are no good theories to explain that.

What I have done so far is to get you into the right mood to listen to me. Otherwise, we have no chance. So now we're off, ready to go!

We begin with light. When Newton started looking at light, the first thing he found was that white light is a mixture of colors. He separated white light with a prism into various colors, but when he put light of one color—red, for instance—through another prism, he found it could not be separated further. So Newton found that white light is a mixture of different colors, each of which is pure in the sense that it can't be separated further.

(In fact, a particular color of light can be split one more time in a different way, according to its so-called "polarization." This aspect of light is not vital to understanding the character of quantum electrodynamics, so for the sake of simplicity I will leave it out—at the expense of not giving you an absolutely complete description of the theory. This slight simplification will not remove, in any way, any real understanding of what I will be talking about. Still, I must be careful to mention all of the things I leave out.)

When I say "light" in these lectures, I don't mean simply the light we can see, from red to blue. It turns out that visible light is just a part of a long scale that's analogous to a musical scale in which there are notes higher than you can hear and other notes lower than you can hear. The scale of light can be described by numbers—called the frequency—and as the numbers get higher, the light goes from red to blue to violet to ultraviolet. We can't see ultraviolet light, but it can affect photographic plates. It's still light—only the number is different. (We shouldn't be so provincial: what we can detect directly with our own instrument, the eye, isn't the only thing in the world!) If we continue simply to change the number, we go out into X-rays, gamma rays, and so on. If we change the number in the other direction, we go from blue to red to infrared (heat) waves, then television waves, and radio waves. For me, all of that is "light." I'm going to use just red light for most of my examples, but the theory of quantum electro-

dynamics extends over the entire range I have described, and is the theory behind all these various phenomena.

Newton thought that light was made up of particles—he called them “corpuscles”—and he was right (but the reasoning that he used to come to that decision was erroneous). We know that light is made of particles because we can take a very sensitive instrument that makes clicks when light shines on it, and if the light gets dimmer, the clicks remain just as loud—there are just fewer of them. Thus light is something like raindrops—each little hump of light is called a photon—and if the light is all one color, all the “raindrops” are the same size.


The human eye is a very good instrument: it takes only about five or six photons to activate a nerve cell and send a message to the brain. If we were evolved a little further so we could see ten times more sensitively, we wouldn’t have to have this discussion—we would all have seen very dim light of one color as a series of intermittent little flashes of equal intensity.

You might wonder how it is possible to detect a single photon. One instrument that can do this is called a photomultiplier, and I’ll describe briefly how it works: When a photon hits the metal plate A at the bottom (see Figure 1), it causes an electron to break loose from one of the atoms in the plate. The free electron is strongly attracted to plate B (which has a positive charge on it) and hits it with enough force to break loose three or four electrons. Each of the electrons knocked out of plate B is attracted to plate C (which is also charged), and their collision with plate C knocks loose even more electrons. This process is repeated ten or twelve times, until billions of electrons, enough to make a sizable electric current, hit the last plate, L. This current can be amplified by a regular amplifier and sent through a speaker to make audible clicks. Each time

a photon of a given color hits the photomultiplier, a click of uniform loudness is heard.

If you put a whole lot of photomultipliers around and let some very dim light shine in various directions, the light goes into one multiplier or another and makes a click of full intensity. It is all or nothing: if one photomultiplier

FIGURE 1. A photomultiplier can detect a single photon. When a photon strikes plate A, an electron is knocked loose and attracted to positively charged plate B, knocking more electrons loose. This process continues until billions of electrons strike the last plate, L, and produce an electric current, which is amplified by a regular amplifier. If a speaker is connected to the amplifier, clicks of uniform loudness are heard each time a photon of a given color hits plate A.



goes off at a given moment, none of the others goes off at the same moment (except in the rare instance that two photons happened to leave the light source at the same time). There is no splitting of light into “half particles” that go different places.

I want to emphasize that light comes in this form—particles. It is very important to know that light behaves like particles, especially for those of you who have gone to school, where you were probably told something about light behaving like waves. I’m telling you the way it *does* behave—like particles.

You might say that it’s just the photomultiplier that detects light as particles, but no, every instrument that has been designed to be sensitive enough to detect weak light has always ended up discovering the same thing: light is made of particles.

I am going to assume that you are familiar with the properties of light in everyday circumstances—things like, light goes in straight lines; it bends when it goes into water; when it is reflected from a surface like a mirror, the angle at which the light hits the surface is equal to the angle at which it leaves the surface; light can be separated into colors; you can see beautiful colors on a mud puddle when there is a little bit of oil on it; a lens focuses light, and so on. I am going to use these phenomena that you are familiar with in order to illustrate the truly strange behavior of light; I am going to explain these familiar phenomena in terms of the theory of quantum electrodynamics. I told you about the photomultiplier in order to illustrate an essential phenomenon that you may not have been familiar with—that light is made of particles—but by now, I hope you are familiar with that, too!

Now, I think you are all familiar with the phenomenon that light is partly reflected from some surfaces, such as water. Many are the romantic paintings of moonlight reflecting from a lake (and many are the times you got yourself in trouble *because* of moonlight reflecting from a lake!). When you look down into water you can see what's below the surface (especially in the daytime), but you can also see a reflection from the surface. Glass is another example: if you have a lamp on in the room and you're looking out through a window during the daytime, you can see things outside through the glass as well as a dim reflection of the lamp in the room. So light is partially reflected from the surface of glass.

Before I go on, I want you to be aware of a simplification I am going to make that I will correct later on: When I talk about the partial reflection of light by glass, I am going to pretend that the light is reflected by only the *surface* of the glass. In reality, a piece of glass is a terrible monster of complexity—huge numbers of electrons are jiggling about.

When a photon comes down, it interacts with electrons *throughout* the glass, not just on the surface. The photon and electrons do some kind of dance, the net result of which is the same as if the photon hit only the surface. So let me make that simplification for a while. Later on, I'll show you what actually happens inside the glass so you can understand why the result is the same.

Now I'd like to describe an experiment, and tell you its surprising results. In this experiment some photons of the same color—let's say, red light—are emitted from a light source (see Fig. 2) down toward a block of glass. A photomultiplier is placed at A, above the glass, to catch any

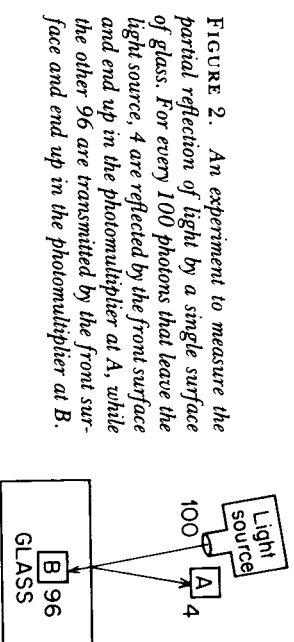


FIGURE 2. An experiment to measure the partial reflection of light by a single surface of glass. For every 100 photons that leave the light source, 4 are reflected by the front surface and end up in the photomultiplier at A, while the other 96 are transmitted by the front surface and end up in the photomultiplier at B.

photons that are reflected by the front surface. To measure how many photons get past the front surface, another photomultiplier is placed at B, inside the glass. Never mind the obvious difficulties of putting a photomultiplier inside a block of glass; what are the results of this experiment?

For every 100 photons that go straight down toward the glass at 90°, an average of 4 arrive at A and 96 arrive at B. So “partial reflection” in this case means that 4% of the photons are reflected by the front surface of the glass, while the other 96% are transmitted. *Already* we are in great difficulty: how can light be *partly* reflected? Each photon ends

up at A or B—how does the photon “make up its mind” whether it should go to A or B? (Audience laughs.) That may sound like a joke, but we can’t just laugh; we’re going to have to explain that in terms of a theory! Partial reflection is already a deep mystery, and it was a very difficult problem for Newton.

There are several possible theories that you could make up to account for the partial reflection of light by glass. One of them is that 96% of the surface of the glass is “holes” that let the light through, while the other 4% of the surface is covered by small “spots” of reflective material (see Fig. 3). Newton realized that this is not a possible explanation.¹ In just a moment we will encounter a strange feature of partial reflection that will drive you crazy if you try to stick to a theory of “holes and spots”—or to any other reasonable theory!

Another possible theory is that the photons have some kind of internal mechanism—“wheels” and “gears” inside that are turning in some way—so that when a photon is “aimed” just right, it goes through the glass, and when it’s not aimed right, it reflects. We can check this theory by trying to filter out the photons that are not aimed right by putting a few extra layers of glass between the source and the first layer of glass. After going through the filters, the photons reaching the glass should *all* be aimed right, and

¹ How did he know? Newton was a very great man; he wrote, “Because I can polish glass.” You might wonder, how the heck could he tell that because you can polish glass, it can’t be holes and spots? Newton polished his own lenses and mirrors, and he knew what he was doing with polishing: he was making scratches on the surface of a piece of glass with powders of increasing fineness. As the scratches become finer and finer, the surface of the glass changes its appearance from a dull grey (because the light is scattered by the large scratches), to a transparent clarity (because the extremely fine scratches let the light through). Thus he saw that it is impossible to accept the proposition that light can be affected by very small irregularities such as scratches or holes and spots; in fact, he found the contrary to be true. The finest scratches and therefore equally small spots do not affect the light. So the holes and spots theory is no good.

none of them should reflect. The trouble with that theory is, it doesn’t agree with experiment: even after going through many layers of glass, 4% of the photons reaching a given surface reflect off it.

Try as we might to invent a reasonable theory that can




FIGURE 3. One theory to explain partial reflection by a single surface involves a surface made up mainly of “holes” that let light through, with a few “spots” that reflect the light.

explain how a photon “makes up its mind” whether to go through glass or bounce back, it is impossible to predict which way a given photon will go. Philosophers have said that if the same circumstances don’t always produce the same results, predictions are impossible and science will collapse. Here is a circumstance—identical photons are always coming down in the same direction to the same piece of glass—that produces different results. We cannot predict whether a given photon will arrive at A or B. All we can predict is that out of 100 photons that come down, an average of 4 will be reflected by the front surface. Does this mean that physics, a science of great exactitude, has been reduced to calculating only the *probability* of an event, and not predicting exactly what will happen? Yes. That’s a retreat, but that’s the way it is: Nature permits us to calculate only probabilities. Yet science has not collapsed.

While partial reflection by a single surface is a deep mystery and a difficult problem, partial reflection by two or more surfaces is absolutely mind-boggling. Let me

show you why. We'll do a second experiment, in which we will measure the partial reflection of light by two surfaces. We replace the block of glass with a very thin sheet of glass—its two surfaces are exactly parallel to each other—and we place the photomultiplier below the sheet of glass, in line with the light source. This time, photons can reflect from either the front surface or the back surface to end up at A; all the others will end up at B (see Fig. 4). We might




FIGURE 4. An experiment to measure the partial reflection of light by two surfaces of glass. Photons can get to the photomultiplier at A by reflecting off either the front surface or the back surface of the sheet of glass; alternatively, they could go through both surfaces and end up hitting the photomultiplier at B. Depending on the thickness of the glass, 0 to 16 photons out of every 100 get to the photomultiplier at A. These results pose difficulties for any reasonable theory, including the one in Figure 3. It appears that partial reflection can be “turned off” or “amplified” by the presence of an additional surface.

expect the front surface to reflect 4% of the light and the back surface to reflect 4% of the remaining 96%, making a total of about 8%. So we should find that out of every 100 photons that leave the light source, about 8 arrive at A.

What actually happens under these carefully controlled experimental conditions is, the number of photons arriving at A is rarely 8 out of 100. With some sheets of glass, we consistently get a reading of 15 or 16 photons—twice our expected result! With other sheets of glass, we consistently get only 1 or 2 photons. Other sheets of glass have a partial reflection of 10%; some eliminate partial reflection altogether! What can account for these crazy results? After checking the various sheets of glass for quality and uniformity, we discover that they differ only slightly in their thickness.

To test the idea that the amount of light reflected by two surfaces depends on the thickness of the glass, let's do a series of experiments: Starting out with the thinnest possible layer of glass, we'll count how many photons hit the photomultiplier at A each time 100 photons leave the light source. Then we'll replace the layer of glass with a slightly thicker one and make new counts. After repeating this process a few dozen times, what are the results?

With the thinnest possible layer of glass, we find that the number of photons arriving at A is nearly always zero—sometimes it's 1. When we replace the thinnest layer with a slightly thicker one, we find that the amount of light reflected is higher—closer to the expected 8%. After a few more replacements the count of photons arriving at A increases past the 8% mark. As we continue to substitute still “thicker” layers of glass—we're up to about 5 millionths of an inch now—the amount of light reflected by the two surfaces reaches a maximum of 16%, and then goes down, through 8%, back to zero—if the layer of glass is just the right thickness, there is no reflection at all. (Do that with spots!)

With gradually thicker and thicker layers of glass, partial reflection again increases to 16% and returns to zero—a cycle that repeats itself again and again (see Fig. 5). Newton discovered these oscillations and did one experiment that could be correctly interpreted only if the oscillations continued for at least 34,000 cycles! Today, with lasers (which produce a very pure, monochromatic light), we can see this cycle still going strong after more than 100,000,000 repetitions—which corresponds to glass that is more than 50 meters thick. (We don't see this phenomenon every day because the light source is normally not monochromatic.)

So it turns out that our prediction of 8% is right as an overall average (since the actual amount varies in a regular pattern from zero to 16%), but it's exactly right only twice each cycle—like a stopped clock (which is right twice a day).

How can we explain this strange feature of partial reflection that depends on the thickness of the glass? How can the front surface reflect 4% of the light (as confirmed in our first experiment) when, by putting a second surface at just the right distance below, we can somehow “turn off” the




FIGURE 5. The results of an experiment carefully measuring the relationship between the thickness of a sheet of glass and partial reflection demonstrate a phenomenon called “interference.” As the thickness of the glass increases, partial reflection goes through a repeating cycle of zero to 16%, with no signs of dying out.

reflection? And by placing that second surface at a slightly different depth, we can “amplify” the reflection up to 16%! Can it be that the back surface exerts some kind of influence or effect on the ability of the front surface to reflect light? What if we put in a *third* surface?

With a third surface, or any number of subsequent surfaces, the amount of partial reflection is again changed. We find ourselves chasing down through surface after surface with this theory, wondering if we have finally reached the last surface. Does a photon have to do that in order to “decide” whether to reflect off the front surface?

Newton made some ingenious arguments concerning this problem,² but he realized, in the end, that he had not yet developed a satisfactory theory.

² It is very fortunate for us that Newton convinced himself that light is “corpuscles,” because we can see what a fresh and intelligent mind looking

For many years after Newton, partial reflection by two surfaces was happily explained by a theory of waves,³ but when experiments were made with very weak light hitting photomultipliers, the wave theory collapsed: as the light got dimmer and dimmer, the photomultipliers kept making

at this phenomenon of partial reflection by two or more surfaces has to go through to try to explain it. (Those who believed that light was waves never had to wrestle with it.) Newton argued as follows: Although light appears to be reflected from the first surface, it cannot be reflected from that surface. If it were, then how could light reflected from the first surface be captured again when the thickness is such that there was supposed to be no reflection at all? Then light must be reflected from the second surface. But to account for the fact that the thickness of the glass determines the amount of partial reflection, Newton proposed this idea: Light striking the first surface sets off a kind of wave or field that travels along with the light and predisposes it to reflect or not reflect off the second surface. He called this process “fits of easy reflection or easy transmission” that occur in cycles, depending on the thickness of the glass.

There are two difficulties with this idea: the first is the effect of additional surfaces—each new surface affects the reflection—which I described in the text. The other problem is that light certainly reflects off a lake, which doesn’t have a second surface, so light *must* be reflecting off the front surface. In the case of single surfaces, Newton said that light had a predisposition to reflect. Can we have a theory in which the light knows what kind of surface it is hitting, and whether it is the only surface?

Newton didn’t emphasize these difficulties with his theory of “fits of reflection and transmission,” even though it is clear that he knew his theory was not satisfactory. In Newton’s time, difficulties with a theory were dealt with briefly and glossed over—a different style from what we are used to in science today, where we point out the places where our own theory doesn’t fit the observations of experiment. I’m not trying to say anything against Newton; I just want to say something in favor of how we communicate with each other in science today.

³ This idea made use of the fact that waves can combine or cancel out, and the calculations based on this model matched the results of Newton’s experiments, as well as those done for hundreds of years afterwards. But when instruments were developed that were sensitive enough to detect a single photon, the wave theory predicted that the “clicks” of the photomultiplier would get softer and softer, whereas they stayed at full strength—they just occurred less and less often. No reasonable model could explain this fact, so there was a period for a while in which you had to be clever: You had to know which experiment you were analyzing in order to tell if light was waves or particles. This state of confusion was called the “wave-particle duality” of light, and it was jokingly said by someone that light was waves on Mondays, Wednesdays, and Fridays; it was particles on Tuesdays, Thursdays, and Saturdays, and on Sundays, we think about it! It is the purpose of these lectures to tell you how this puzzle was finally “resolved.”

full-sized clicks—there were just fewer of them. Light behaved as particles.

The situation today is, we haven't got a good model to explain partial reflection by two surfaces; we just calculate the probability that a particular photomultiplier will be hit by a photon reflected from a sheet of glass. I have chosen this calculation as our first example of the method provided by the theory of quantum electrodynamics. I am going to show you "how we count the beans"—what the physicists do to get the right answer. I am not going to explain how the photons actually "decide" whether to bounce back or go through; that is not known. (Probably the question has no meaning.) I will only show you how to calculate the correct *probability* that light will be reflected from glass of a given thickness, because that's the only thing physicists know how to do! What we do to get the answer to *this* problem is analogous to the things we have to do to get the answer to *every other* problem explained by quantum electrodynamics.

You will have to brace yourselves for this—not because it is difficult to understand, but because it is absolutely ridiculous: All we do is draw little arrows on a piece of paper—that's all!

Now, what does an arrow have to do with the chance that a particular event will happen? According to the rules of "how we count the beans," the probability of an event is equal to the square of the length of the arrow. For example, in our first experiment (when we were measuring partial reflection by the front surface only), the probability that a photon would arrive at the photomultiplier at A was 4%. That corresponds to an arrow whose length is 0.2, because 0.2 squared is 0.04 (see Fig. 6).

In our second experiment (when we were replacing thin sheets of glass with slightly thicker ones), photons bouncing off either the front surface or the back surface arrived at

A. How do we draw an arrow to represent this situation? The length of the arrow must range from zero to 0.4 to represent probabilities of zero to 16%, depending on the thickness of the glass (see Fig. 7).

We start by considering the various ways that a photon




FIGURE 6. The strange feature of partial reflection by two surfaces has forced physicists away from making absolute predictions to merely calculating the probability of an event. Quantum electrodynamics provides a method for doing this—drawing little arrows on a piece of paper. The probability of an event is represented by the area of the square on an arrow. For example, an arrow representing a probability of 0.04 (4%) has a length of 0.2.

could get from the source to the photomultiplier at A. Since I am making this simplification that the light bounces off either the front surface or the back surface, there are two possible ways a photon could get to A. What we do in this case is to draw *two* arrows—one for each way the event can happen—and then combine them into a "final arrow"




FIGURE 7. Arrows representing probabilities from 0% to 16% have lengths of from 0 to 0.4.

whose square represents the probability of the event. If there had been three different ways the event could have happened, we would have drawn three separate arrows before combining them.

Now, let me show you how we combine arrows. Let's say we want to combine arrow x with arrow y (see Fig. 8). All




FIGURE 8. Arrows that represent each possible way an event could happen are drawn and then combined ("added") in the following manner: Attach the head of one arrow to the tail of another—without changing the direction of either one—and draw a "final arrow" from the tail of the first arrow to the head of the last one.

we have to do is put the head of x against the tail of y (without changing the direction of either one), and draw the final arrow from the tail of x to the head of y . That's all there is to it. We can combine any number of arrows in this manner (technically, it's called "adding arrows"). Each arrow tells you how far, and in what direction, to move in a dance. The final arrow tells you what *single* move to make to end up in the same place (see Fig. 9).

Now, what are the specific rules that determine the length and direction of each arrow that we combine in order to make the final arrow? In this particular case, we will be combining two arrows—one representing the reflection from the *front* surface of the glass, and the other representing the reflection from the *back* surface.

Let's take the length first. As we saw in the first experi-

ment (where we put the photomultiplier inside the glass), the front surface reflects about 4% of the photons that come down. That means the "front reflection" arrow has a length of 0.2. The back surface of the glass also reflects 4%, so the "back reflection" arrow's length is also 0.2.




FIGURE 9. Any number of arrows can be added in the manner described in Figure 8.

To determine the direction of each arrow, let's imagine that we have a stopwatch that can time a photon as it moves. This imaginary stopwatch has a single hand that turns around very, very rapidly. When a photon leaves the source, we start the stopwatch. As long as the photon moves, the stopwatch hand turns (about 36,000 times per inch for red light); when the photon ends up at the photomultiplier, we stop the watch. The hand ends up pointing in a certain direction. That is the direction we will draw the arrow.

We need one more rule in order to compute the answer correctly: When we are considering the path of a photon bouncing off the *front* surface of the glass, we reverse the direction of the arrow. In other words, whereas we draw the *back* reflection arrow pointing in the *same* direction as the stopwatch hand, we draw the *front* reflection arrow in the *opposite* direction.

Now, let's draw the arrows for the case of light reflecting from an extremely thin layer of glass. To draw the front

reflection arrow, we imagine a photon leaving the light source (the stopwatch hand starts turning), bouncing off the front surface, and arriving at A (the stopwatch hand stops). We draw a little arrow of length 0.2 in the direction opposite that of the stopwatch hand (see Fig. 10).




FIGURE 10. In an experiment measuring reflection by two surfaces, we can say that a single photon can arrive at A in two ways—via the front or back surface. An arrow of length 0.2 is drawn for each way, with its direction determined by the hand of a “stopwatch” that times the photon as it moves. The “front reflection” arrow is drawn in the direction opposite to that of the stopwatch hand when it stops turning.

To draw the back reflection arrow, we imagine a photon leaving the light source (the stopwatch hand starts turning), going through the front surface and bouncing off the back surface, and arriving at A (the stopwatch hand stops). This time, the stopwatch hand is pointing in almost the same direction, because a photon bouncing off the back surface of the glass takes only slightly longer to get to A—it goes through the extremely thin layer of glass twice. We now draw a little arrow of length 0.2 in the same direction that the stopwatch hand is pointing (see Fig. 11).

Now let’s combine the two arrows. Since they are both the same length but pointing in nearly opposite directions, the final arrow has a length of nearly zero, and its square is even closer to zero. Thus, the probability of light reflecting from an infinitesimally thin layer of glass is essentially zero (see Fig. 12).

FIGURE 11. A photon bouncing off the back surface of a thin layer of glass takes slightly longer to get to A. Thus, the stopwatch hand ends up in a slightly different direction than it did when it timed the front reflection photon. The “back reflection” arrow is drawn in the same direction as the stopwatch hand.





FIGURE 12. The final arrow, whose square represents the probability of reflection by an extremely thin layer of glass, is drawn by adding the front reflection arrow and the back reflection arrow. The result is nearly zero.

When we replace the thinnest layer of glass with a slightly thicker one, the photon bouncing off the back surface takes a little bit longer to get to A than in the first example; the stopwatch hand therefore turns a little bit more before it stops, and the back reflection arrow ends up in a slightly greater angle relative to the front reflection arrow. The final arrow is a little bit longer, and its square is correspondingly larger (see Fig. 13).

As another example, let’s look at the case where the glass is just thick enough that the stopwatch hand makes an extra half turn as it times a photon bouncing off the back surface. This time, the back reflection arrow ends up pointing in exactly the same direction as the front reflection arrow. When we combine the two arrows, we get a final arrow whose length is 0.4, and whose square is 0.16, representing a probability of 16% (see Fig. 14).

If we increase the thickness of the glass just enough so




FIGURE 13. The final arrow for a slightly thicker sheet of glass is a little longer, due to the greater relative angle between the front and back reflection arrows. This is because a photon bouncing off the back surface takes a little longer to reach A, compared to the previous example.




FIGURE 14. When the layer of glass is just thick enough to allow the stopwatch hand timing the back reflecting photon to make an extra half turn, the front and back reflection arrows end up pointing in the same direction, resulting in a final arrow of length 0.4, which represents a probability of 16%.

that the stopwatch hand timing the back surface path makes an extra full turn, our two arrows end up pointing in opposite directions again, and the final arrow will be zero (see Fig. 15). This situation occurs over and over, whenever the thickness of the glass is just enough to let the stopwatch hand timing the back surface reflection make another full turn.




FIGURE 15. When the sheet of glass is just the right thickness to allow the stopwatch hand timing the back reflecting photon to make one or more extra full turns, the final arrow is again zero, and there is no reflection at all.

If the thickness of the glass is just enough to let the stopwatch hand timing the back surface reflection make an extra $\frac{1}{4}$ or $\frac{3}{4}$ of a turn, the two arrows will end up at right angles. The final arrow in this case is the hypotenuse of a right triangle, and according to Pythagoras, the square on the hypotenuse is equal to the sum of the squares on the other two sides. Here is the value that's right "twice a day"— $4\% + 4\%$ makes 8% (see Fig. 16).

Notice that as we gradually increase the thickness of the glass, the front reflection arrow always points in the same direction, whereas the back reflection arrow gradually changes its direction. The change in the relative direction of the two arrows makes the final arrow go through a re-




FIGURE 16. When the front and back reflection arrows are at right angles to each other, the final arrow is the hypotenuse of a right triangle. Thus its square is the sum of the other two squares— $8^2\%$.




FIGURE 17. As thin sheets of glass are replaced by slightly thicker ones, the stopwatch hand timing a photon reflecting off the back surface turns slightly more, and the relative angle between the front and back reflection arrows changes. This causes the final arrow to change in length, and its square to change in size from 0 to 16% back to 0, over and over.

pearing cycle of length zero to 0.4; thus the *square* on the final arrow goes through the repeating cycle of zero to 16% that we observed in our experiments (see Fig. 17).

I have just shown you how this strange feature of partial reflection can be accurately calculated by drawing some damned little arrows on a piece of paper. The technical word for these arrows is “probability amplitudes,” and I feel more dignified when I say we are “computing the probability amplitude for an event.” I prefer, though, to be more honest, and say that we are trying to find the arrow whose square represents the probability of something happening.

Before I finish this first lecture, I would like to tell you about the colors you see on soap bubbles. Or better, if your car leaks oil into a mud puddle, when you look at the brownish oil in that dirty mud puddle, you see beautiful colors on the surface. The thin film of oil floating on the mud puddle is something like a very thin sheet of glass—it reflects light of one color from zero to a maximum, depending on its thickness. If we shine pure red light on the film of oil, we see splotches of red light separated by narrow bands of black (where there’s no reflection) because the oil film’s thickness is not exactly uniform. If we shine pure blue light on the oil film, we see splotches of blue light separated by narrow bands of black. If we shine both red and blue light onto the oil, we see areas that have just the right thickness to strongly reflect only red light, other areas of the right thickness to reflect only blue light; still other areas have a thickness that strongly reflects both red and blue light (which our eyes see as violet), while other areas have the exact thickness to cancel out all reflection, and appear black.

To understand this better, we need to know that the cycle of zero to 16% partial reflection by two surfaces repeats more quickly for blue light than for red light. Thus at

certain thicknesses, one or the other or both colors are strongly reflected; while at other thicknesses, reflection of both colors is cancelled out (see Fig. 18). The cycles of reflection repeat at different rates because the stopwatch hand turns around faster when it times a blue photon than




FIGURE 18. As the thickness of a layer increases, the two surfaces produce a partial reflection of monochromatic light whose probability fluctuates in a cycle from 0% to 16%. Since the speed of the imaginary stopwatch hand is different for different colors of light, the cycle repeats itself at different rates. Thus when two colors such as pure red and pure blue are aimed at the layer, a given thickness will reflect only red, only blue, both red and blue in different proportions (which produce various hues of violet), or neither color (black). If the layer is of varying thicknesses, such as a drop of oil spreading out on a mud puddle, all of the combinations will occur. In sunlight, which consists of all colors, all sorts of combinations occur, which produce lots of colors.

it does when timing a red photon. In fact, that's the only difference between a red photon and a blue photon (or a photon of any other color, including radio waves, X-rays, and so on)—the speed of the stopwatch hand.

When we shine red and blue light on a film of oil, patterns of red, blue, and violet appear, separated by borders of black. When sunlight, which contains red, yellow, green,

and blue light, shines on a mud puddle with oil on it, the areas that strongly reflect each of those colors overlap and produce all kinds of combinations which our eyes see as different colors. As the oil film spreads out and moves over the surface of the water, changing its thickness in various locations, the patterns of color constantly change. (If, on the other hand, you were to look at the same mud puddle at night with one of those sodium streetlights shining on it, you would see only yellowish bands separated by black—because those particular streetlights emit light of only one color.)

This phenomenon of colors produced by the partial reflection of white light by two surfaces is called iridescence, and can be found in many places. Perhaps you have wondered how the brilliant colors of hummingbirds and peacocks are produced. Now you know. How those brilliant colors evolved is also an interesting question. When we admire a peacock, we should give credit to the generations of lackluster females for being selective about their mates. (Man got into the act later and streamlined the selection process in peacocks.)

In the next lecture I will show you how this absurd process of combining little arrows computes the right answer for those other phenomena you are familiar with: light travels in straight lines; it reflects off a mirror at the same angle that it came in (“the angle of incidence is equal to the angle of reflection”); a lens focuses light, and so on. This new framework will describe everything you know about light.