SASS
Wells Wulsin
22 September 2010
6013 words

QED for Beginners

My goal today is to explain quantum electrodynamics in a non-technical way. For those of you who have never encountered this theory before, I hope to give you a flavor of what QED is all about. Because after all, even if you do not need to perform QED calculations for your field of research, it still is the theory that governs how light and matter interact, and as I'll show, it can explain lots of familiar phenomena from our everyday experience. Others of you may be well-versed in QED already, whether you're an experimentalist or theorist. For you folks, I hope to show you how this theory can be presented to a general audience with minimal mathematical complications. After all, this theory is too beautiful and powerful to be restricted to only a small segment of highly-trained physicists.

This talk is based on the book QED: The Strange Theory of Light and Matter by Richard Feynman (Princeton University Press, 2006 edition). The book is short, just 152 pages. It is based on four lectures Feynman gave at UCLA in 1983. Each chapter of the book corresponds to one of those chapters. I don't know if a video of the UCLA lectures exists, but four similar lectures that Feynman gave in New Zealand are available online at http://vega.org.uk/video/subseries/8. I highly recommend these—they show how witty and sharp Feynman is on his feet.

Now, in a single seminar of under an hour I cannot hope to do justice to material that took Feynman four ninety-minute lectures to cover. So I have to move more quickly over some of the finer points, and I have left several sections out. Still, I hope to do justice to Feynman's original goal, which was to present QED accurately in an accessible way.

Feynman writes that to learn QED you have two options: either spend 7 years taking physics courses training or read his book. As an analogy to these two approaches, consider two ways of multiplying 123 times 456. In one case, someone says that you need to memorize a 9x9 table, and then follow a bunch of rules like, "Line up the two numbers on top of each other, then start with the right-most numbers, and look up the number in the table, and then put only the last digit of that number underneath, and carry the first digit…" This method corresponds to the 7 years of school. Alternatively, someone could explain to you, "Just find 123 jars, and into each jar put 456 pebbles. Now pour all the pebbles from all the jars into a big pile and count them out." This latter method is what Feynman is trying to do. He doesn't teach you all the tricks that make these calculations quicker and easier to perform, but he does give an accurate picture of how the theory works.

There are a few different formulations of quantum mechanics, and they each have different advantages, though they all give the same end result. Feynman here uses the path-integral formulation of quantum mechanics. For some problems this formulation is not very helpful for calculating the answer. For example, to find the energy levels of the hydrogen atom, the Schrodinger wave-function formulation makes things much easier. But

again, performing calculations are not the main goal here, and I think the conceptual advantages of the path-integral formulation are why Feynman chooses this approach.

Feynman's lectures were delivered at a level accessible to non-physicists. Since all of you are physicists, I'll move a bit more quickly than he did. In particular, I'll use diagrams that I've copied from his book. However, I would recommend that if you were actually presenting this material to a non-physics audience, there is a great advantage to drawing diagrams yourself on the chalkboard, because it gives your audience time to digest what is going on and to see your line of thinking. But for the sake of time, I'm going to use slides instead.

[1Introduction]

Now, let's get started. QED describes the behavior of light, electrons, and their interactions with each other. It does not explain gravity, or nuclear interactions, but it does explain everything else. Dynamite explosions, photosynthesis, springs, colors—all the phenomena of chemistry and biology are governed by the interactions of light and electrons, and QED lies at the root of all these phenomena. Now, the task of using QED to explain something like photosynthesis becomes prohibitively difficult, because there are so many electrons involved that the calculations are just too hard to carry out. But in principle it is possible, and whenever we look at systems that are simple enough that the QED effects can be calculated, the experimental measurements agree with the theory. In fact, its level of precision is higher than maybe any other theory. The prediction (in 1983) for the magnetic moment of the electron from QED theoretical calculations is 1.00115965246 +/- 20, while the experimental measurement of the same quantity is 1.00115965221 +/- 4. This level of accuracy corresponds to the measurement of the distance between New York and Los Angeles to a precision of the width of a human hair. This is my way of persuading you that QED actually works!

[2LightOffGlass]

[Fig1, p.15]

Let's begin with a simple experiment, the measurement of light. We use a photomultiplier tube to measure light, and it works in the following way: some light enters the detector, and strikes a metal plate, which knocks off some electrons. These electrons are then attracted to a second plate by a large voltage difference, and then they strike the second plate with greater energy, knocking off more electrons. These electrons are in turn attracted to a third plate by a voltage difference, and so on, over a series of plates, until a large number of electrons has built up that produce a current that can be amplified and then produce an audible click when passed through a speaker.

Now, what's interesting when we use these photomultipliers is that when we can turn down the intensity of the source of light low enough, rather than see a smooth amount of light coming with less energy, what we actually see is the light coming in chunks of the same size energy. These lumps of energy called photons are tiny, and we are normally bombarded by so many of them that light seems to be a smooth quantity. But in fact light is like raindrops, with the energy traveling in these little drops, and if light is a single color then all the drops have the same size. Whenever we have an instrument to detect light that is sensitive enough, it records light coming in discrete chunks, so we conclude that light is made of particles.

I should point out here that when I refer to light, I am not referring only to the wavelengths that the human eye can see, but rather to any part of the electromagnetic spectrum.

[Fig 2, p.17]

Now let's take a source a single color of light—let's say red—such as a laser pointer, and shine it at a piece of glass. Then let's place two detectors, one above the glass and one inside the glass. If we turn down the intensity of the light source low enough, then we only detect a single photon at a time, at either detector A or at B. If we wait until 100 photons have been detected, we find that on average 96 of them have been detected at B, while 4 have been detected by A. We say that for any given photon, there is a 4% chance that it is reflected back to A, and a 96% chance that it is transmitted to B. In fact, it turns out that this probabilistic description is all that we can say about what will happen to any given photon. This may strike you as unsatisfactory. Surely, nature must know whether a given photon will be reflected.

Maybe photons are shaped like footballs, and they tumble as they travel through space, and have to be oriented a certain way to pass through the glass. If that were the case then, you would think that you could put another sheet of glass above this first glass, to filter out the photons so they're all oriented the same way when they hit the final piece of glass. But it turns out that whatever filters we use, still 4% of the photons get reflected and 96% are transmitted.

[Fig 3, p.19]

Another explanation you might consider for why only some of the photons get reflected is that maybe there are "holes" in the glass and only certain "spots" that reflect light. This might be a good explanation, if not for the following experiment.

[Fig 4, p.20]

Now we put the second detector below the glass. You might think that now we would get 8 photons at detector A: 4 reflected from the top surface and 4 reflected from the bottom surface.

[Fig 5, p. 22]

But in fact we find that the number of photons reflected depends on the thickness of the piece of glass, as shown in this figure. With no glass at all, the number of reflected photons is 0, as we would expect, and as the glass gets thicker, the number reflected increases to 16%, but then drops back down to 0%, and the number reflected keeps increasing and decreasing as the thickness grows. Somehow, by putting a second surface just the right distance below, we can turn off the reflections from the front surface. And then putting the second surface at a slightly different distance causes the reflections to jump up to 16%. I think you'll agree that the "holes and spots" theory will have trouble explaining this effect. For many years, this effect was explained by a theory of waves, but that theory cannot explain why low light is always observed as particles.

So let's now turn to QED, which tells us that we can only calculate the probability for a photon to end up at point A or B. The way to calculate the probability for the photon to end up in one place is to draw an arrow for each possible path the photon can take, and the to add up the arrows, and square the final result. So let's see how that works.

[Fig 6, p.25]

In the first case, where we need only consider reflection from a single surface, the probability of reflection is .04, which we get by squaring an arrow of length 0.2.

[Fig 10-12, p.26-27]

Now, in the case where we need to consider reflection from two surfaces, we will have to find two arrows. Each one has the same length: 0.2. Now, to determine the direction of each arrow, imagine that we have a stopwatch with a single hand that can time a photon as it moves. The stopwatch hand turns very quickly—about 36,000 per inch for red light—and the direction that the hand ends up when the photon arrives at the photomultiplier is the direction we draw the arrow. Also, there is one other rule: when light bounces off the front of the glass, we reverse the direction of the arrow.

For a thin piece of glass, the photon bouncing off the back surface takes just a little bit more time than the one bouncing off the front surface. So when we add up the two arrows, they are nearly back to back, and the resultant arrow is relatively short. So there is a small probability, 0.25% in this case, that the photon is received at A.

[Fig 13-15, p. 30-31]

When the thickness of the glass changes, the relative position of the back arrow also changes, causing the resultant arrow to increase to 16%, and then decrease down to 0%. And as the thickness continues to increase, the resultant arrow goes back and forth between 0 and 0.4.

Now that we've shown how QED works in this one example, let's write down the more general rules:

Grand Principle: The probability of an event is equal to the square of the length of an arrow called the "probability amplitude."

General Rule for drawing arrows if an event can happen in alternative ways: Draw an arrow for each way, and then add the arrows by hooking the head of one to the tail of the next. The final arrow is drawn from the tail of the first arrow to the head of the last one.

The technical term for the arrows is "probability amplitudes", and in practice physicists represent them as complex numbers rather than arrows, but the result is the same if you draw arrows and add them up.

Now I should mention at this point that I am leaving something out, which is the polarization of light—the fact that photons have a spin. This effect leads to some differences in the way that calculations are made, but it doesn't change the underlying idea of how QED works.

[Fig 18, p. 35]

Now, imagine that you spill some oil onto a puddle in the driveway. When you look at the puddle, you'll see a pattern of colors reflected back at you. The oil spreads out over the puddle, and acts like a piece of glass, able to reflect light. And since the film of oil has a different thickness in different places, for a given color of light, say red, there will be certain bands that appear black because no light is reflected and other bands where a maximum amount of red light is reflected. But for another color, say blue, the cycle from 0% to 16% repeats more quickly, since the stopwatch hand for blue light turns faster than for red light. So depending on the thickness of the film, in some places we see only red, or only blue, or a combination or red and blue which appears violet, or not much of either one which appears black. Since sunlight contains a mixture of many different colors, then the areas of the puddle which reflect each color overlap and produce the beautiful patterns that we see.

On the other hand, if you looked at the same puddle at night with the light of a sodium streetlamp, you would only see bands of yellow separated by black, since those

lights only emit light of one color. The phenomenon of producing colors by partial reflection of white light by two surfaces is called iridescence, and is how hummingbirds and peacocks get their brilliant colors.

[Fig 19, p.39]

Now let's apply the principles of QED to some other phenomena from everyday life. Let's begin with light reflecting from a mirror. The classical view says that light should reflect off a mirror at the same angle as its angle of incidence. So let's start with a source of low-intensity light of a single color at S, and let's put a detector at P, with a screen in between them so the photons cannot go straight across. Now, you might think that the only part of the mirror that matters is the section of mirror right in the middle, where the angle of incidence will equal the angle of reflection.

[Fig 20, p. 40]

But QED says that the arrows for all the different possible paths need to be added up to get the total arrow, which is then squared to get the probability. That means that we need to consider light that takes other paths. Now, you might say that's crazy, since the angle is wrong for those other paths, but as you'll see, we do need to consider them because the light really goes there!

[Fig 21, p. 40]

To simplify the problem, let's divide the mirror into short sections, which could of course be made arbitrarily small. Then a photon from S can get to P by reflecting off any one of those sections. So we just need to calculate the arrow for each of those possible paths. Now you might think that the arrow for the shortest path, when it hits the mirror in the middle, has to be longest, but that's not true. The arrow for each path has just about the same length. The only difference is the time that it takes the photon to travel each path—which means that the different arrows will be pointed in different ways based on the position of the hand of our imaginary stopwatch.

[Fig 23-24, p. 42-43]

So the length of the path the photon takes determines the direction the arrow points. We can plot the time of each path on this graph here, as a function of where the photon touches the mirror. At the center of the mirror, point G, the time taken is shortest, while farther away from the center, the time for the photon's journey gets larger. Now we just have to add up all the arrows from each of these paths to figure out the total arrow. At the beginning, the direction of the arrows are changing so much from point to point that we don't get very far, but then when we get close to the middle, the arrows are mostly pointed in the same direction, and the arrows actually move somewhere.

So now we see that the ends of the mirror don't matter so much: in those regions, all the arrows are spinning around so fast that the sum of the arrows just wanders around in pretty much the same place. It's the section of the mirror in the middle, where the arrows are mostly pointed in the same direction, that is mostly responsible for the length of the final arrow. The arrows in this middle section are mostly pointed in the same direction because the time of travel is nearly the same for all those points. It also happens to be where the time of travel is shortest. This is why we can get away with the approximation of saying that light only goes where the time is least—that's because other paths close to that path have nearly the same time, and so they add up to give an arrow of substantial length.

[Fig 25, p.46]

Now you're probably thinking, that's just a waste of effort to bother with the ends of the mirror when we know that only the middle part of the mirror matters. But let me show you why we do need to consider all these different paths. Let me take a section of mirror from one of the ends, and zoom in on it. As expected, since the time is varying rapidly in this section, the arrows all point in different directions. And when you add them up, they spiral around in a circle and don't get anywhere, so they don't contribute to the total length of the final arrow.

[Fig 26, p. 47]

Now, let's say that I paint part of this section of mirror black, so that it can't reflect anymore. Now the arrows from the parts of the mirror that do reflect are all generally pointed in the same direction. According to QED, since these arrows do add up to a substantial final arrow, so we should get a strong reflection. And in fact we do! This kind of mirror is called a diffraction grating. By starting with a mirror where you don't expect a reflection, you can take bits of it away to create a reflection. Now, this grating was made for red light, but if you shine blue light on it, you can still get a reflection, but at a different angle. So when you shine white light on a CD or DVD, you see many different colors, because all the pits in the disc reflect different colors at different angles.

[Fig 28, p. 49]

An example of a diffraction grating found in nature is a salt crystal, made of an orderly lattice of ions—sodium and chlorine in the case of table salt. The crystals will reflect X-rays at a certain angle, based on the spacing between the ions, and so by measuring the angles at which the X-rays are reflected, we can figure out how far apart the ions are. So the lesson here is: we have to consider amplitudes from all the possible paths, not only paths that we think are going to be important.

[Fig 29, p. 50]

Now let's consider another example: light passing from air into water. So to calculate the probability of light being absorbed at detector D if it is emitted at source S, we consider all the ways that the light can go. And once again, the major contribution to the length of the final arrow comes from the paths close to where the light takes the least time—because each of these paths has an associated arrow pointed in nearly the same direction.

[Fig 31, p. 52]

This kind of analysis can also explain the effect of a mirage. One way to see the sky on a road is if there is a puddle of water on the road. But even without any water, it is possible to see the sky on the road, since light travels faster through hot hair than through cold air. Thus, the path of least time may bend close to the road and then up towards the observer.

[Fig 32, p. 53]

In the previous examples, I have always assumed that light travels in a straight line, but we can use QED to show where this general principle comes from. Let's look at all the ways that light from source S can travel to a photomultiplier at P. Here just a few possible paths are drawn in the figure. Notice that the straight-line path takes the least time, and the arrows for the nearby paths all point in the same general direction. The other, meandering paths, cover a much longer distance, and the arrows associated with them tend to point in different directions.

[Fig 33, p. 55]

Note that in order for the light to be visible, we need a substantial core of paths to be available—the shortest-time path all by itself is not enough. For example, let's look at what happens when light passes from source S to detector P. If there are enough paths available, then the light is detected at P, since the different paths each contribute arrows pointed in the same way that add together. If we put another detector over at Q, we won't get much light since the path lengths are quite different in time.

[Fig 34, p. 56]

Now, if we make the aperture smaller by moving the screens closer together, the number of paths to P is reduced so we see less light there. But at the same time the path-length difference decreases for the paths going to point Q, so we see more light over there. So we see that if you try to squeeze light too much, then it starts to spread out.

[Fig 35, p. 57]

Now, just as we were able to do a trick with our mirror to get light to reflect at funny angles, let's see how we can use a trick to get light to follow different paths. Consider light going from S to P, along the paths drawn in this figure. The dashed line here isn't a physical thing; it's just to guide the eye. Notice that the time to travel each of these paths is very different, so the arrows point in lots of different directions.

[Fig 36, p. 58]

Now let's trick the light into having the same time for each path. To do that, we make the light go slower on the paths that have a shorter distance—which we can accomplish by putting a medium like glass in the way of the light's path. If we calculate the amount of glass needed for each path, we may get a shape something like this, which gets the time for all of the paths to exactly match. Now when we add the arrows together, we get a very long final arrow. And this is exactly how a lens works.

[Fig 37, p. 60]

So to figure out the probability of an event that can happen multiple ways, you find the arrow for each way and add them up. Now let's consider the probability for a compound event—a process with multiple steps. For this, we find the arrow for each step, and we **multiply them**. To multiply arrows, we follow the rule that you add their angles and multiply their lengths. Take the example of light reflecting off glass to arrive at A, and then reflecting off another piece of glass to arrive at C. The first step, from S to A, is given by an arrow of length 0.2 pointed at 2 o'clock: that's a shrink of .2 and a turn to 2 o'clock. The second step, from A to C, is a shrink of 0.3, and a turn to 5 o'clock. Multiplying them together, we get a shrink of 0.2 times 0.3, which is 0.06, and a turn of 2 o'clock plus 5 o'clock, which is 7 o'clock.

[Fig 40, p. 64]

So now we can look more carefully at a single reflection to see what is going on. We can break it into 3 steps. Step 1 is a turn of 5 o'clock, and a shrink of 1.0 (assuming no light is lost), Step 2 is a shrink of 0.2 and a turn to 6 o'clock, and Step 3 is a turn to 4 o'clock. The arrow for the whole process we get by multiplying the arrows for each of these steps together.

[Fig 47, p. 73]

Arrows are also multiplied for events when two or more things are happening simultaneously. For example, let's say we have two sources, X and Y, and two detectors A and B, and we want the probability for the following event: X and Y each lose a photon and A and B each receive a photon. Well, one way this can happen is shown here: a photon

goes from X to A and a second goes from Y to B. Each has an associated arrow, and we multiply those together to get the total arrow.

[Fig 48, p. 74]

But there's another way it can happen: the photon can go from X to B and another can go from Y to A. This process also has an associated arrow, and we have to add these two arrows together to get the total probability for the event. The relative directions of the arrows for the "parallel" and "crossed" paths can be changed by small changes in the relative position of A and B. This phenomenon, called the Hanbury-Brown-Twiss effect, has been used to distinguish between a single source and a double source of radio waves in deep space, when the sources are very close together.

[Fig 52, p. 86]

Up to now, I've mostly been talking about how photons move around, but QED describes the interactions of light with matter, so now I want to get to the matter bit. In QED, there are just three basic actions that can happen, and each one has an associated amplitude (arrow):

1) A photon goes from place to place.
2) An electron goes from place to place.
3) An electron emits or absorbs a photon.

Now the stage on which all of these things happen is spacetime, which is often drawn with the funny convention of putting time on the vertical axis, as shown in this spacetime drawing of a baseball.

[Fig 53, p. 87]

And here is a baseball bouncing against a wall.

[Fig 55-56, p. 88]

Now here is a photon going from spacetime point A to spacetime point B. The amplitude for this to happen is given by a formula, P(A to B), which only depends on the difference in space and in time between the two points. Now, as you might suspect, P(A to B) favors light traveling at the standard speed, c, but there are also contributions for light to travel at speeds other than c, which get smaller as the speed gets more different from c. For long distances, these contributions cancel out, but over short distances they can have important effects.

[Fig 57, p. 92]

The second action is for an electron to go from A to B, and the amplitude for this is given by another function E(A to B). There are many ways the electron can get from A to B, doing multi-leg journeys as shown in the figure, but E(A to B) includes the amplitudes for all of them.

[Fig 58, p. 92]

The third basic action is: an electron emits or absorbs a photon, as shown in this diagram. Note that electrons are represented by straight lines while photons are represented by wavy lines. This action is often called a junction or coupling. This coupling is just a number, let's call it j, which has a value of about -0.1: a shrink to one-tenth and half a turn. Now, I've left out some slight complications due to the polarization that I've left out, but otherwise, that's all there is to QED. Now let's put these actions together in different ways.

[Fig 59,60, p.93-4]

First, let's consider two electrons moving from points 1 and 2 to points 3 and 4. There are 2 straight-line paths to consider, as shown in the first diagram. These are each:
E(1 to 3)*E(2 to 4) and
E(1 to 4)*E(2 to 3). This is a good approximation for the amplitude of the event. But there is another way it could happen: they could exchange a photon, as shown in the second figure. Now we have amplitudes:
E(1 to 5)*j*E(5 to 3) * E(2 to 6)*j*E(6 to 4) * P(5 to 6) and
E(1 to 6)*j*E(6 to 4) * E(2 to 5)*j*E(5 to 3) * P(5 to 6).
But wait: positions 5 and 6 could be anywhere at all in space and time—yes, anywhere—and all of the arrows for all of those ways have to be added together. Yes, it's going to be a lot of work, which is why physicists have developed lots of tricks to make the calculations easier. But the rules are simple; it's just working them out that takes many years of training to master.

[Fig 61, p. 96]

Now the photon has an amplitude to go faster or slower than the speed of light, so the three diagrams here are all the same diagram. Rather than worry about which way the photon is actually going, we typically just say that a photon has been exchanged.

[Fig 62, p. 96]

The diagram with 2 photons will have 6 factors of E(A to B), 2 factors of P(A to B), and 4 factors of j. So length of these arrows are about 10,000 times smaller than the diagrams with no photons. So the more complicated diagrams contribute less and less to the final arrow, which is what allows us to ever make a calculation at all.

[Fig 63, p. 97]

Now let's consider the scattering of a photon off an electron. We can draw this in several ways, as shown here. Note that sometimes the electron emits a photon before it absorbs one, and sometimes it emits a photon, then rushes back in time to absorb another photon.

[Fig 64, p. 99]

This backwards-moving electron can hang around for a long time in the laboratory, and has the same properties as a normal electron, except that its charge is opposite. And we call it a positron.

[Fig 65-6, p. 100]

Now let's look at an electron in an atom. QED doesn't describe nuclear effects, but the forces between the electron and the proton of the nucleus are described by the exchange of photons. When a photon is scattered off an atom, the electron absorbs a photon, and then emits a new photon later. There is an amplitude associated with this scattering.

[Fig 67, p. 102]

Now let's return to our original scattering experiment. Really, light is not affected by surfaces. Rather, it is scattered by electrons in atoms all through the material. Now before, we imagined a stopwatch turning for each photon, but in P(A to B) there is no stopwatch. Where did that come from? Well, a monochromatic light source emits photons in a very predictable way: the amplitude rotates counter-clockwise as time moves forward. The rate for turning depends on the color of the light—a blue source turns almost twice as fast as a red source. But once the photon is emitted, there is no further turning of the arrow.

[Fig 68, p. 104]

So to simplify, let's divide the glass into 6 sections. Each of these sections will have a small arrow associated with it. These arrows slowly turn in a circle. So the front surface and back surface arrows were just a calculation trick—by adding them together (with the front arrow flipped around), we can come up with the total arrow from all the slices.

[Fig 69, p. 108]

The same analysis can be done for the light passing through the glass.

[Fig 70-71, p. 111]

The total amplitude for 2 photons to end up in the same place is the sum of two arrows: for each of the ways that the photons can get to that place. If we didn't think about interference, we would expect the probability to be twice as large on average, but instead it's 4 times as large. This effect is magnified by even more photons. This means that photons tend to get into the same state. The chance that an atom emits a photon is enhanced if some photons are already present. This is the phenomenon of stimulated emission which lasers take advantage of.

[Fig 72, p. 112]

The same effect would happen for electrons, except for the fact that electrons are polarized, which means that instead of adding the two amplitudes, they get subtracted—and they cancel out. This phenomenon—that electrons try not to end up in the same place—is known as the Pauli exclusion principle, and is responsible for giving all the elements there many varied properties. So all of chemistry comes from this basic principle of QED.

[Fig 73-75, p. 115-117]

The magnetic moment of the electron describes how the electron interacts with a photon in a magnetic field. The simplest interaction looks like the first figure, but there are higher-order diagrams that also must be considered. After adding all those up, we find that the experimental value agrees to a high precision with the theoretical value. This makes QED one of the best-tested theories in the history of science.

Outline

1. introduction
    a. agrees with experiment
2. light reflecting off glass
    a. light is quantized: photomultiplier 15
    b. detector in glass and above glass 17
    c. detector behind glass and above glass 20, 22
    d. soap bubble 34
3. rules for adding arrows
    a. 26, 28, 29, 30a, 30b, 31
4. Light bouncing off mirror 39, 40a, 40b, 42, 43,
5. DVD 46, 47, 48
    a. X-ray scattering 49
6. Light passing through water 50, 51,
    a. Mirage 52,
    b. Light going in straight line
7. Light passing through space
    a. Wide slit 55
    b. Narrow slit 56
    c. Lens 57, 58
8. Multiple steps 60, 63, 64, 66
9. Electron motion 79, 80, 83
10. Spacetime graphs 86, 87, 88, 92a
11. Electron emission of photon 92b, 94,
    a. Violation of time causality 96a, 97, 99
    b. Scattering of light by an atom 100a, 100b
12. Reinterpretation of light bouncing off glass 102, 104, 108
13. Two electrons propagating 93
    a. Pauli exclusion principle 111, 112
14. Magnetic moment of electron 115, 116, 117, 118
15. Loose ends
    a. Renormalization
    b. Nuclear effects: weak and strong forces
    c. Parameters
    d. Gravity