

THE DEVELOPMENT OF THE SPACE-TIME VIEW OF QUANTUM ELECTRODYNAMICS*

by

Richard P. Feynman

California Institute of Technology, Pasadena, California

Nobel Lecture, December 11, 1965.

We have a habit in writing articles published in scientific journals to make the work as finished as possible, to cover all the tracks, to not worry about the blind alleys or to describe how you had the wrong idea first, and so on. So there isn't any place to publish, in a dignified manner, what you actually did in order to get to do the work, although, there has been in these days, some interest in this kind of thing. Since winning the prize is a personal thing, I thought I could be excused in this particular situation, if I were to talk personally about my relationship to quantum electrodynamics, rather than to discuss the subject itself in a refined and finished fashion. Furthermore, since there are three people who have won the prize in physics, if they are all going to be talking about quantum electrodynamics itself, one might become bored with the subject. So, what I would like to tell you about today are the sequence of events, really the sequence of ideas, which occurred, and by which I finally came out the other end with an unsolved problem for which I ultimately received a prize.

I realize that a truly scientific paper would be of greater value, but such a paper I could publish in regular journals. So, I shall use this Nobel Lecture as an opportunity to do something of less value, but which I cannot do elsewhere. I ask your indulgence in another manner. I shall include details of anecdotes which are of no value either scientifically, nor for understanding the development of ideas. They are included only to make the lecture more entertaining.

I worked on this problem about eight years until the final publication in 1947. The beginning of the thing was at the Massachusetts Institute of Technology when I was an undergraduate student reading about the known physics, learning slowly about all these things that people were worrying about, and realizing ultimately that the fundamental problem of the day was that the quantum theory of electricity and magnetism was not completely satisfactory. This I gathered from books like those of Heitler and Dirac. I was inspired by the remarks in these books; not by the parts in which everything was proved and demonstrated carefully and calculated, because I couldn't understand those very well. At the young age what I could understand were

* This document is a revised version of Feynman's Lecture, with amendments made by Michael D. Godfrey and Michael A. Gottlieb (email: michaeldgodfrey@gmail.com and codelieb@caltech.edu).

the remarks about the fact that this doesn't make any sense, and the last sentence of the book of Dirac I can still remember, "It seems that some essentially new physical ideas are here needed." So, I had this as a challenge and an inspiration. I also had a personal feeling, that since they didn't get a satisfactory answer to the problem I wanted to solve, I don't have to pay a lot of attention to what they did do.

I did gather from my readings, however, that two things were the source of the difficulties with the quantum electrodynamical theories. The first was an infinite energy of interaction of the electron with itself. And this difficulty existed even in the classical theory. The other difficulty came from some infinities which had to do with the infinite number of degrees of freedom in the field. As I understood it at the time (as nearly as I can remember) this was simply the difficulty that if you quantized the harmonic oscillators of the field (say in a box) each oscillator has a ground state energy of $\frac{1}{2}\hbar\omega$ and there is an infinite number of modes in a box of ever increasing frequency ω , and therefore there is an infinite energy in the box. I now realize that that wasn't a completely correct statement of the central problem; it can be removed simply by changing the zero from which energy is measured. At any rate, I believed that the difficulty arose somehow from a combination of the electron acting on itself and the infinite number of degrees of freedom of the field.

Well, it seemed to me quite evident that the idea that a particle acts on itself, that the electrical force acts on the same particle that generates it, is not a necessary one—it is sort of a silly one, as a matter of fact. And, so I suggested to myself, that electrons cannot act on themselves; they can only act on other electrons. That means there is no field at all. You see, if all charges contribute to making a single common field, and if that common field acts back on all the charges, then each charge must act back on itself. Well, that was where the mistake was, there was no field. It was just that when you shook one charge, another would shake later. There was a direct interaction between charges, albeit with a delay. The law of force connecting the motion of one charge with another would just involve a delay. Shake this one, that one shakes later. The sun atom shakes; my eye electron shakes eight minutes later, because of a direct interaction across.

Now, this has the attractive feature that it solves both problems at once. First, I can say immediately, I don't let the electron act on itself, I just let this act on that, hence, no self-energy! Secondly, there is not an infinite number of degrees of freedom in the field. There is no field at all; or if you insist on thinking in terms of ideas like that of a field, this field is always completely determined by the action of the particles which produce it. You shake this particle, it shakes that one, but if you want to think in a field way, the field, if it's there, would be entirely determined by the matter which generates it, and therefore, the field does not have any *independent* degrees of freedom and the infinities from the degrees of freedom would then be removed. As a matter of fact, when we look out anywhere and see light, we can always "see" some matter as the source of the light. We don't just see light (except recently some radio reception has been found with no apparent material source).

You see then that my general plan was to first solve the classical problem, to get rid of the infinite self-energies in the classical theory, and to hope that when I made a quantum theory of it, everything would just be fine.

That was the beginning, and the idea seemed so obvious to me and so elegant that I fell deeply in love with it. And, like falling in love with a woman, it is only possible if you do not know much about her, so you cannot see her faults. The faults will become apparent later, but after the love is strong enough to hold you to her. So, I was held to this theory, in spite of all difficulties, by my youthful enthusiasm.

Then I went to graduate school and somewhere along the line I learned what was wrong with the idea that an electron does not act on itself. When you accelerate an electron it radiates energy and you have to do extra work to account for that energy. The extra force against which this work is done is called the force of radiation resistance. The origin of this extra force was identified in those days, following Lorentz, as the action of the electron itself. The first term of this action, of the electron on itself, gave a kind of inertia (not quite relativistically satisfactory). But that inertia-like term was infinite for a point-charge. Yet the next term in the sequence gave an energy loss rate, which for a point-charge agrees exactly with the rate you get by calculating how much energy is radiated. So, the force of radiation resistance, which is absolutely necessary for the conservation of energy, would disappear if I said that a charge could not act on itself.

So, I learned in the interim when I went to graduate school the glaringly obvious fault of my own theory. But, I was still in love with the original theory, and was still thinking that with it lay the solution to the difficulties of quantum electrodynamics. So, I continued to try on and off to save it somehow. I must have some action develop on a given electron when I accelerate it to account for radiation resistance. But, if I let electrons only act on other electrons the only possible source for this action is another electron in the world. So, one day, when I was working for Professor Wheeler and could no longer solve the problem that he had given me, I thought about this again and I calculated the following: Suppose I have two charges—I shake the first charge, which I think of as a source and this makes the second one shake, but the second one shaking produces an effect back on the source. And so, I calculated how much that effect back on the first charge was, hoping it might add up to the force of radiation resistance. It didn't come out right, of course, but I went to Professor Wheeler and told him my ideas. He said,—yes, but the answer you get for the problem with the two charges that you just mentioned will, unfortunately, depend upon the charge and the mass of the second charge and will vary inversely as the square of the distance R , between the charges, while the force of radiation resistance depends on none of these things. I thought, surely, he had computed it himself, but now having become a professor, I know that one can be wise enough to see immediately what some graduate student takes several weeks to develop. He also pointed out something else that bothered me, that if we had a situation with many charges all around the original source at roughly uniform density and if we added the effect of all the surrounding charges the inverse R^2 would be compensated by the R^2 in the volume element and we would get a result proportional to the thickness of the layer, which would go to infinity. That is, one would have an infinite total effect back at the source. And, finally he said to me, and you forgot something else, when you accelerate the first charge, the second acts later, and then the reaction back here at the source would be still later. In other words, the action occurs at the wrong time. I

suddenly realized what a stupid fellow I am, for what I had described and calculated was just ordinary reflected light, not radiation reaction.

But, as I was stupid, so was Professor Wheeler that much more clever. For he then went on to give a lecture as though he had worked this all out before and was completely prepared, but he had not, he worked it out as he went along. First, he said, let us suppose that the return action by the charges in the absorber reaches the source by advanced waves as well as by the ordinary retarded waves of reflected light; so that the law of interaction acts backward in time, as well as forward in time. I was enough of a physicist at that time not to say, “Oh, no, how could that be?” For today all physicists know from studying Einstein and Bohr, that sometimes an idea which looks completely paradoxical at first, if analyzed to completion in all detail and in experimental situations, may, in fact, not be paradoxical. So, it did not bother me any more than it bothered Professor Wheeler to use advance waves for the back reaction—a solution of Maxwell’s equations, which previously had not been physically used.

Professor Wheeler used advanced waves to get the reaction back at the right time and then he suggested this: If there were lots of electrons in the absorber, there would be an index of refraction n , so, the retarded waves coming from the source would have their wavelengths slightly modified in going through the absorber. Now, if we shall assume that the advanced waves come back from the absorber without an index—why? I don’t know, let’s assume they come back without an index—then, there will be a gradual shifting in phase between the return and the original signal so that we would only have to figure that the contributions act as if they come from only a finite thickness, that of the first wave zone. (More specifically, up to that depth where the phase in the medium is shifted appreciably from what it would be in vacuum, a thickness proportional to $\lambda/(n - 1)$.) Now, the less the number of electrons in here, the less each contributes, but the thicker will be the layer that effectively contributes because with less electrons, the index differs less from 1. The higher the charges of these electrons, the more each contributes, but the thinner the effective layer, because the index would be higher. And when we estimated it, (calculated without being careful to keep the correct numerical factor) sure enough, it came out that the action back at the source was completely independent of the properties of the charges that were in the surrounding absorber. Further, it was of just the right character to represent radiation resistance, but we were unable to see if it was just exactly the right size. He sent me home with orders to figure out exactly how much advanced and how much retarded wave we need to get the thing to come out numerically right, and after that, figure out what happens to the advanced effects that you would expect if you put a test charge here close to the source? For if all charges generate advanced, as well as retarded effects, why would that test charge not be affected by the advanced waves from the source?

I found that you get the right answer if you use half-advanced and half-retarded as the field generated by each charge. That is, one is to use the solution of Maxwell’s equation which is symmetrical in time and that the reason we got no advanced effects at a point close to the source in spite of the fact that the source was producing an advanced field is this: Suppose the source is surrounded by a spherical absorbing wall

ten light seconds away, and that the test charge is one second to the right of the source. Then the source* is as much as eleven seconds away from some parts of the wall and only nine seconds away from other parts. The source acting at time $t = 0$ induces motions in the wall at time $t = +10$. Advanced effects from this can act on the test charge as early as eleven seconds earlier, or at $t = -1$. This is just at the time that the direct advanced waves from the source should reach the test charge, and it turns out the two effects are exactly equal and opposite and cancel out! At the later time $t = +1$ effects on the test charge from the source and from the walls are again equal, but this time are of the same sign and add to convert the half-retarded wave of the source to full retarded strength.

Thus, it became clear that there was the possibility that if we assume all actions are via half-advanced and half-retarded solutions of Maxwell's equations and that all sources are surrounded by material absorbing all the the light which is emitted, then we could account for radiation resistance as a direct action of the charges of the absorber acting back by advanced waves on the source.

Many months were devoted to checking all these points. I worked to show that everything is independent of the shape of the container, and so on, that the laws are exactly right, and that the advanced effects really cancel in every case. We always tried to increase the efficiency of our demonstrations, and to see with more and more clarity why it works. I won't bore you by going through the details of this. Because of our using advanced waves we also had many apparent paradoxes, which we gradually reduced one by one, and saw that there was in fact no logical difficulty with the theory. It was perfectly satisfactory.

We also found that we could reformulate this thing in another way, and that is by a principle of least action. Since my original plan was to describe everything directly in terms of particle motions, it was my desire to represent this new theory without saying anything about fields. It turned out that we found a form for an action directly involving the motions of the charges only, which upon variation would give the equations of motion of these charges. The expression for this action A is

$$A = \sum_i m_i \int \left(\dot{X}_\mu^i \dot{X}_\mu^i \right)^{\frac{1}{2}} d\alpha_i + \frac{1}{2} \sum_{\substack{ij \\ i \neq j}} e_i e_j \int \int \delta(I_{ij}^2) \dot{X}_\mu^i(\alpha_i) \dot{X}_\mu^j(\alpha_j) d\alpha_i d\alpha_j \quad (1)$$

where

$$I_{ij}^2 = [X_\mu^i(\alpha_i) - X_\mu^j(\alpha_j)][X_\mu^i(\alpha_i) - X_\mu^j(\alpha_j)]$$

and where $X_\mu^i(\alpha_i)$ is the four-vector position of the i^{th} particle as a function of some parameter α_i , and $\dot{X}_\mu^i(\alpha_i)$ is $dX_\mu^i(\alpha_i)/d\alpha_i$. The first term is the integral of proper time, the ordinary action of relativistic mechanics of free particles of mass m_i . (We sum in the usual way on the repeated index μ .) The second term represents the electrical interaction of the charges. It is summed over each pair of charges (the factor $\frac{1}{2}$ is to count each pair once, the term $i = j$ is omitted to avoid self-action).

* It is apparent that Feynman meant *test charge* instead of *source*. This was noticed by Jonathan Hilairret.

The interaction is a double integral over a δ -function of the square of the space-time interval I^2 between two points on the paths. Thus, interaction occurs only when this interval vanishes, that is, along light cones.

The fact that the interaction is exactly one-half advanced and one-half retarded meant that we could write such a principle of least action, whereas interaction via retarded waves alone cannot be written in such a way.

So, all of classical electrodynamics was contained in this very simple form. It looked good, and therefore, it was undoubtedly true, at least to the beginner. It automatically gave half-advanced and half-retarded effects and it was without fields. By omitting the term in the sum when $i = j$, I omitted self-interaction and no longer had any infinite self-energy. This then was the hoped-for solution to the problem of ridding classical electrodynamics of the infinities.

It turns out, of course, that you can reinstate fields if you wish to, but you have to keep track of the field produced by each particle separately. This is because to find the right field to act on a given particle, you must exclude the field that it creates itself. A single universal field to which all contribute will not do. This idea had been suggested earlier by Frenkel and so we called these Frenkel fields. This theory which allowed only particles to act on each other was equivalent to Frenkel's fields using half-advanced and half-retarded solutions.

There were several suggestions for interesting modifications of electrodynamics. We discussed lots of them, but I shall report on only one. It was to replace this δ -function in the interaction by another function, say, $f(I_{ij}^2)$, which is not infinitely sharp. Instead of having the action occur only when the interval between the two charges is exactly zero, we would replace the δ -function of I^2 by a narrow peaked thing. Let's say that $f(Z)$ is large only near $Z = 0$ and has width of order a^2 . Interactions will now occur when $T^2 - R^2$ is roughly of order a^2 , where T is the time difference and R is the separation of the charges. This might look like it disagrees with experience, but if a is some small distance, like 10^{-13} cm, it says that the time delay T in action is roughly $\sqrt{R^2 \pm a^2}$ or approximately, if R is much larger than a , $T = R \pm a^2/2R$. This means that the deviation of time T from the ideal theoretical time R of Maxwell, gets smaller and smaller, the further the pieces are apart. Therefore, all theories involved in analyzing generators, motors, etc., in fact, all of the tests of electrodynamics that have been available since Maxwell's time, would be adequately satisfied if a were 10^{-13} cm. If R is of the order of a centimeter this deviation in T is only 10^{-26} seconds. So, it was possible, also, to change the theory in a simple manner and to still agree with all observations of classical electrodynamics. You have no clue of precisely what function to put in for f , but it was an interesting possibility to keep in mind when developing quantum electrodynamics.

It also occurred to us that if we did that (replace δ by f) we could reinstate the term $i = j$ in the sum because this would now represent, in a relativistically invariant fashion, a finite action of a charge on itself. In fact, it was possible to prove that if we did do such a thing, the main effect of the self-action (for not too rapid accelerations) would be to produce a modification of the mass. In fact, there need be no mass m_i term; all the mechanical mass could be electromagnetic self-action. So, if you would

like, we could also have another theory with a still simpler expression for the action A . In expression (1) only the second term is kept, the sum extended over all i and j , and some function f replaces δ . Such a simple form could represent all of classical electrodynamics, which aside from gravitation is essentially all of classical physics.

Although it may sound confusing, I am describing several different alternative theories at once. The important thing to note is that at this time we had all these in mind as different possibilities. There were several possible solutions of the difficulty of classical electrodynamics, any one of which might serve as a good starting point to the solution of the difficulties of quantum electrodynamics.

I would also like to emphasize that by this time I was becoming used to a physical point of view different from the more customary point of view. In the customary view, things are discussed as a function of time in very great detail. For example, you have the field at this moment, a differential equation gives you the field at the next moment and so on; a method, which I shall call the Hamiltonian method, the time differential method. We have, instead (in (1) say) a thing that describes the character of the path throughout all of space and time. The behavior of nature is determined by saying her whole space-time path has a certain character. For an action like (1) the equations obtained by variation of $X_{\mu}^i(\alpha_i)$ are no longer at all easy to get back into Hamiltonian form. If you wish to use as variables only the coordinates of particles, then you can talk about the property of the paths—but the path of one particle at a given time is affected by the path of another at a different time. If you try to describe, therefore, things differentially, telling what the present conditions of the particles are, and how these present conditions will affect the future, you see, it is impossible with particles alone, because something the particle did in the past is going to affect the future.

Therefore, you need a lot of bookkeeping variables to keep track of what the particle did in the past. These are called field variables. You will, also, have to tell what the field is at this present moment, if you are to be able to see later what is going to happen. From the overall space-time view of the least action principle, the field disappears as nothing but bookkeeping variables insisted on by the Hamiltonian method.

As a by-product of this same view, I received a telephone call one day at the graduate college at Princeton from Professor Wheeler, in which he said, “Feynman, I know why all electrons have the same charge and the same mass.” “Why?” “Because, they are all the same electron!” And, then he explained on the telephone, “suppose that the world lines which we were ordinarily considering before in time and space—instead of only going up in time were a tremendous knot, and then, when we cut through the knot, by the plane corresponding to a fixed time, we would see many, many world lines and that would represent many electrons, except for one thing. If in one section this is an ordinary electron world line, in the section in which it reversed itself and is coming back from the future we have the wrong sign to the proper time—to the proper four velocities—and that’s equivalent to changing the sign of the charge, and, therefore, that part of a path would act like a positron.” “But, Professor,” I said, “there aren’t as many positrons as electrons.” “Well, maybe they are hidden in the protons or something”, he said. I did not take the idea that all the electrons were the same one from him as seriously as I took the observation that positrons could

simply be represented as electrons going from the future to the past in a back section of their world lines. That, I stole!

To summarize, when I was done with this, as a physicist I had gained two things. One, I knew many different ways of formulating classical electrodynamics, with many different mathematical forms. I got to know how to express the subject every which way. Second, I had a point of view—the overall space-time point of view—and a disrespect for the Hamiltonian method of describing physics.

I would like to interrupt here to make a remark. The fact that electrodynamics can be written in so many ways—the differential equations of Maxwell, various minimum principles with fields, minimum principles without fields, all different kinds of ways, was something I knew, but I have never understood. It always seems odd to me that the fundamental laws of physics, when discovered, can appear in so many different forms that are not apparently identical at first, but, with a little mathematical fiddling you can show the relationship. An example of that is the Schrödinger equation and the Heisenberg formulation of quantum mechanics. I don't know why this is—it remains a mystery, but it was something I learned from experience. There is always another way to say the same thing that doesn't look at all like the way you said it before. I don't know what the reason for this is. I think it is somehow a representation of the simplicity of nature. A thing like the inverse square law is just right to be represented by the solution of Poisson's equation, which, therefore, is a very different way to say the same thing that doesn't look at all like the way you said it before. I don't know what it means, that nature chooses these curious forms, but maybe that is a way of defining simplicity. Perhaps a thing is simple if you can describe it fully in several different ways without immediately knowing that you are describing the same thing.

I was now convinced that since we had solved the problem of classical electrodynamics (and completely in accordance with my program from M.I.T., only direct interaction between particles, in a way that made fields unnecessary) that everything was definitely going to be all right. I was convinced that all I had to do was make a quantum theory analogous to the classical one and everything would be solved.

So, the problem is only to make a quantum theory which has as its classical analog this expression (1). Now, there is no unique way to make a quantum theory from classical mechanics, although all the textbooks make believe there is. What they would tell you to do, was find the momentum variables and replace them by $(\hbar/i)(\partial/\partial x)$, but I couldn't find a momentum variable, as there wasn't any.

The character of quantum mechanics of the day was to write things in the famous Hamiltonian way—in the form of a differential equation, which described how the wave function changes from instant to instant, and in terms of an operator, H . If the classical physics could be reduced to a Hamiltonian form, everything was all right. Now, least action does not imply a Hamiltonian form if the action is a function of anything more than positions and velocities at the same moment. If the action is of the form of the integral of a function, (usually called the Lagrangian) of the velocities and positions at the same time

$$S = \int L(\dot{x}, x) dt \tag{2}$$

then you can start with the Lagrangian and then create a Hamiltonian and work out the quantum mechanics, more or less uniquely. But this thing (1) involves the key variables, positions, at two different times and therefore, it was not obvious what to do to make the quantum-mechanical analog.

I tried—I would struggle in various ways. One of them was this: if I had harmonic oscillators interacting with a delay in time, I could work out what the normal modes were and guess that the quantum theory of the normal modes was the same as for simple oscillators and kind of work my way back in terms of the original variables. I succeeded in doing that, and I hoped then to generalize to other than a harmonic oscillator, but I learned to my regret something, which many people have learned. The harmonic oscillator is too simple; very often you can work out what it should do in quantum theory without getting much of a clue as to how to generalize your results to other systems.

So that didn't help me very much, but when I was struggling with this problem, I went to a beer party in the Nassau Tavern in Princeton. There was a gentleman, newly arrived from Europe (Herbert Jehle) who came and sat next to me. Europeans are much more serious than we are in America because they think that a good place to discuss intellectual matters is a beer party. So, he sat by me and asked, "what are you doing" and so on, and I said, "I'm drinking beer." Then I realized that he wanted to know what work I was doing and I told him I was struggling with this problem, and I simply turned to him and said, "listen, do you know any way of doing quantum mechanics, starting with action—where the action integral comes into the quantum mechanics?" "No," he said, "but Dirac has a paper in which the Lagrangian, at least, comes into quantum mechanics. I will show it to you tomorrow."

Next day we went to the Princeton Library, they have little rooms on the side to discuss things, and he showed me this paper. What Dirac said was the following: There is in quantum mechanics a very important quantity which carries the wave function from one time to another, besides the differential equation but equivalent to it, a kind of a kernel, which we might call $K(x', x)$, which carries the wave function $\psi(x)$ known at time t , to the wave function $\psi(x')$ at time, $t + \varepsilon$. Dirac points out that this function K was *analogous* to the quantity in classical mechanics that you would calculate if you took the exponential of $i\varepsilon$, multiplied by the Lagrangian $L(\dot{x}, x)$ imagining that these two positions x, x' corresponded to t and $t + \varepsilon$. In other words,

$$K(x', x) \text{ is analogous to } e^{i\varepsilon L(\frac{x'-x}{\varepsilon}, x)/\hbar}.$$

Professor Jehle showed me this, I read it, he explained it to me, and I said, "what does he mean, they are analogous; what does that mean, *analogous*? What is the use of that?" He said, "you Americans! You always want to find a use for everything!" I said, that I thought that Dirac must mean that they were equal. "No," he explained, "he doesn't mean they are equal." "Well," I said, "let's see what happens if we make them equal."

So I simply put them equal, taking the simplest example where the Lagrangian is $\frac{1}{2}M\dot{x}^2 - V(x)$ but soon found I had to put a constant of proportionality A in, suitably

adjusted. When I substituted $Ae^{i\varepsilon L/\hbar}$ for K to get

$$\psi(x', t + \varepsilon) = \int A \exp \left[\frac{i\varepsilon}{\hbar} L \left(\frac{x' - x}{\varepsilon}, x \right) \right] \psi(x, t) dx \quad (3)$$

and just calculated things out by Taylor series expansion, out came the Schrödinger equation. So, I turned to Professor Jehle, not really understanding, and said, “well, you see Professor Dirac meant that they were proportional.” Professor Jehle’s eyes were bugging out—he had taken out a little notebook and was rapidly copying it down from the blackboard, and said, “no, no, this is an important discovery. You Americans are always trying to find out how something can be used. That’s a good way to discover things!” So, I thought I was finding out what Dirac meant, but, as a matter of fact, had made the discovery that what Dirac thought was analogous, was, in fact, equal. I had then, at least, the connection between the Lagrangian and quantum mechanics, but still with wave functions and infinitesimal times.

It must have been a day or so later, when I was lying in bed thinking about these things, that I imagined what would happen if I wanted to calculate the wave function at a finite interval later.

I would put one of these factors $e^{i\varepsilon L}$ in here, and that would give me the wave functions the next moment, $t + \varepsilon$, and then I could substitute that back into (3) to get another factor of $e^{i\varepsilon L}$ and give me the wave function the next moment, $t + 2\varepsilon$, and so on and so on. In that way I found myself thinking of a large number of integrals, one after the other in sequence. In the integrand was the product of the exponentials, which, of course, was the exponential of the sum of terms like εL . Now, L is the Lagrangian and ε is like the time interval dt , so that if you took a sum of such terms, that’s exactly like an integral. That’s like Riemann’s formula for the integral $\int L dt$, you just take the value at each point and add them together. We are to take the limit as $\varepsilon \rightarrow 0$, of course. Therefore, the connection between the wave function of one instant and the wave function of another instant a finite time later could be obtained by an infinite number of integrals, (because ε goes to zero, of course) of exponential (iS/\hbar) where S is the action expression (2). At last, I had succeeded in representing quantum mechanics directly in terms of the action S .

This led later on to the idea of the amplitude for a path; that for each possible way that the particle can go from one point to another in space-time, there’s an amplitude. That amplitude is e to the (i/\hbar) times the action for the path. Amplitudes from various paths superpose by addition. This then is another, a third, way of describing quantum mechanics, which looks quite different than that of Schrödinger or Heisenberg, but which is equivalent to them.

Now immediately after making a few checks on this thing, what I wanted to do, of course, was to substitute the action (1) for the other (2). The first trouble was that I could not get the thing to work with the relativistic case of spin one-half. However, although I could deal with the matter only non-relativistically, I could deal with the light or the photon interactions perfectly well by just putting the interaction terms of (1) into any action, replacing the mass terms by the non-relativistic $(M\dot{x}^2/2)dt$. When the action has a delay, as it now had, and involved more than one time, I had to lose

the idea of a wave function. That is, I could no longer describe the program as; given the amplitude for all positions at a certain time, compute the amplitude at another time. However, that didn't cause very much trouble. It just meant developing a new idea. Instead of wave functions we could talk about this: that if a source of a certain kind emits a particle, and a detector is there to receive it, we can give the amplitude that the source will emit and the detector receive. We do this without specifying the exact instant that the source emits or the exact instant that any detector receives, without trying to specify the state of anything at any particular time in between, but by just finding the amplitude for the complete experiment. And, then we could discuss how that amplitude would change if you had a scattering sample in between, as you rotated and changed angles, and so on, without really having any wave functions.

It was also possible to discover what the old concepts of energy and momentum would mean with this generalized action. And, so I believed that I had a quantum theory of classical electrodynamics—or rather of this new classical electrodynamics described by action (1). I made a number of checks. If I took the Frenkel field point of view, which you remember was more differential, I could convert it directly to quantum mechanics in a more conventional way. The only problem was how to specify in quantum mechanics the classical boundary conditions to use only half-advanced and half-retarded solutions. By some ingenuity in defining what that meant, I found that the quantum mechanics with Frenkel fields, plus a special boundary condition, gave me back this action (1) in the new form of quantum mechanics with a delay. So, various things indicated that there wasn't any doubt I had everything straightened out.

It was also easy to guess how to modify the electrodynamics, if anybody ever wanted to modify it. I just changed the delta to an f , just as I would for the classical case. So, it was very easy, a simple thing. To describe the old retarded theory without explicit mention of fields I would have to write probabilities, not just amplitudes. I would have to square my amplitudes and that would involve double path integrals in which there are two S 's and so forth. Yet, as I worked out many of these things and studied different forms and different boundary conditions, I got a kind of funny feeling that things weren't exactly right. I could not clearly identify the difficulty and in one of the short periods during which I imagined I had laid it to rest, I published a thesis and received my Ph.D.

During the war, I didn't have time to work on these things very extensively, but wandered about on buses and so forth, with little pieces of paper, and struggled to work on it and discovered indeed that there was something wrong, something terribly wrong. I found that if one generalized the action from the nice Lagrangian forms (2) to these forms (1) then the quantities which I defined as energy, and so on, would be complex. The energy values of stationary states wouldn't be real and probabilities of events wouldn't add up to 100%. That is, if you took the probability that this would happen and that would happen—everything you could think of would happen, it would not add up to one.

Another problem on which I struggled very hard, was to represent relativistic electrons with this new quantum mechanics. I wanted to do a unique and different

way—and not just by copying the operators of Dirac into some kind of an expression and using some kind of Dirac algebra instead of ordinary complex numbers. I was very much encouraged by the fact that in one space dimension I did find a way of giving an amplitude to every path by limiting myself to paths that only went back and forth at the speed of light. The amplitude was simply $(i\varepsilon)$ to a power equal to the number of velocity reversals, where I have divided the time into steps ε and I am allowed to reverse velocity only at such a time. This gives (as ε approaches zero) Dirac's equation in two dimensions—one dimension of space and one of time ($\hbar = m = e = 1$).

Dirac's wave function has four components in four dimensions, but in this case, it has only two components and this rule for the amplitude of a path automatically generates the need for two components. Because if this is the formula for the amplitude of a path, it will not do you any good to know the total amplitude of all paths which come into a given point, to find the amplitude to reach the next point. This is because for the next time, if it came in from the right, there is no new factor $i\varepsilon$ if it goes out to the right, whereas, if it came in from the left there was a new factor $i\varepsilon$. So, to continue this same information forward to the next moment, it was not sufficient information to know the total amplitude to arrive, but you had to know the amplitude to arrive from the right and the amplitude to arrive from the left independently. If you did, however, you could then compute both of those again independently and thus you had to carry two amplitudes to form a differential equation (first order in time).

And, so I dreamed that if I were clever, I would find a formula for the amplitude of a path that was beautiful and simple for three dimensions of space and one of time, which would be equivalent to the Dirac equation, and for which the four components, matrices, and all those other mathematical funny things would come out as a simple consequence—I have never succeeded in that either. But, I did want to mention some of the unsuccessful things on which I spent almost as much effort as on the things that did work.

To summarize the situation a few years after the war, I would say I had much experience with quantum electrodynamics, at least in the knowledge of many different ways of formulating it in terms of path integrals of actions and in other forms. One of the important by-products, for example, of much experience in these simple forms, was that it was easy to see how to combine together what was in those days called the longitudinal and transverse fields and, in general, to see clearly the relativistic invariance of the theory. Because of the need to do things differentially there had been, in the standard quantum electrodynamics, a complete split of the field into two parts, one of which is called the longitudinal part and the other mediated by the photons, or transverse waves. The longitudinal part was described by a Coulomb potential acting instantaneously in the Schrödinger equation, while the transverse part had an entirely different description in terms of quantization of the transverse waves. This separation depended upon the relativistic tilt of your axes in space-time. People moving at different velocities would separate the same field into longitudinal and transverse fields in a different way. Furthermore, the entire formulation of quantum mechanics insisting, as it did, on the wave function at a given time, was

hard to analyze relativistically. Somebody else in a different coordinate system would calculate the succession of events in terms of wave functions on differently cut slices of space-time, and with a different separation of longitudinal and transverse parts. The Hamiltonian theory did not look relativistically invariant, although, of course, it was. One of the great advantages of the overall point of view was that you could see the relativistic invariance right away—or as Schwinger would say—the covariance was manifest. I had the advantage, therefore, of having a manifestly covariant form for quantum electrodynamics with suggestions for modifications and so on. I had the disadvantage that if I took it too seriously—I mean, if I took it seriously at all in this form,—I got into trouble with these complex energies and the failure of probabilities adding to one and so on. I was unsuccessfully struggling with that.

Then Lamb did his experiment, measuring the separation of the ${}^2S_{\frac{1}{2}}$ and ${}^2P_{\frac{1}{2}}$ levels of hydrogen, finding it to be about 1000 megacycles of frequency difference. Professor Bethe, with whom I was then associated at Cornell, is a man who has this characteristic: If there's a good experimental number you've got to figure it out from theory. So, he forced the quantum electrodynamics of the day to give him an answer to the separation of these two levels. He pointed out that the self-energy of an electron itself is infinite, so that the calculated energy of a bound electron should also come out infinite. But, when you calculated the separation of the two energy levels in terms of the corrected mass instead of the old mass, it would turn out, he thought, that the theory would give convergent finite answers. He made an estimate of the splitting that way and found out that it was still divergent, but he guessed that was probably due to the fact that he used an unrelativistic theory of the matter. Assuming it would be convergent if relativistically treated, he estimated he would get about a thousand megacycles for the Lamb-shift, and thus, made the most important discovery in the history of the theory of quantum electrodynamics. He worked this out on the train from Ithaca, New York, to Schenectady and telephoned me excitedly from Schenectady to tell me the result, which I don't remember fully appreciating at the time.

Returning to Cornell, he gave a lecture on the subject, which I attended. He explained that it gets very confusing to figure out exactly which infinite term corresponds to what in trying to make the correction for the infinite change in mass. If there were any modifications whatever, he said, even though not physically correct, (that is not necessarily the way nature actually works) but any modification whatever at high frequencies, which would make this correction finite, then there would be no problem at all to figuring out how to keep track of everything. You just calculate the finite mass correction Δm to the electron mass m , substitute the numerical values of $m + \Delta m$ for m in the results for any other problem and all these ambiguities would be resolved. If, in addition, this method were relativistically invariant, then we would be absolutely sure how to do it without destroying relativistic invariance.

After the lecture, I went up to him and told him, "I can do that for you, I'll bring it in for you tomorrow." I guess I knew every way to modify quantum electrodynamics known to man, at the time. So, I went in next day, and explained what would correspond to the modification of the δ -function to f and asked him to explain to me how you calculate the self-energy of an electron, for instance, so we can figure out if

it's finite.

I want you to see an interesting point. I did not take the advice of Professor Jehle to find out how it was useful. I never used all that machinery which I had cooked up to solve a single relativistic problem. I hadn't even calculated the self-energy of an electron up to that moment, and was studying the difficulties with the conservation of probability, and so on, without actually doing anything, except discussing the general properties of the theory.

But now I went to Professor Bethe, who explained to me on the blackboard, as we worked together, how to calculate the self-energy of an electron. Up to that time when you did the integrals they had been logarithmically divergent. I told him how to make the relativistically invariant modifications that I thought would make everything all right. We set up the integral which then diverged at the sixth power of the frequency instead of logarithmically!

So, I went back to my room and worried about this thing and went around in circles trying to figure out what was wrong. Because I was sure physically everything had to come out finite, I couldn't understand how it came out infinite. I became more and more interested and finally realized I had to learn how to make a calculation. So, ultimately, I taught myself how to calculate the self-energy of an electron working my patient way through the terrible confusion of those days of negative energy states and holes and longitudinal contributions and so on. When I finally found out how to do it and did it with the modifications I wanted to suggest, it turned out that it was nicely convergent and finite, just as I had expected. Professor Bethe and I have never been able to discover what we did wrong on that blackboard two months before, but apparently we just went off somewhere and we have never been able to figure out where. It turned out that what I had proposed, if we had carried it out without making a mistake, would have been all right and would have given a finite correction. Anyway, it forced me to go back over all this and to convince myself physically that nothing can go wrong. At any rate, the correction to mass was now finite and proportional to $\ln(a)$ where a is the width of the function f which was substituted for δ . If you wanted an unmodified electrodynamics, you would have to take a equal to zero, getting an infinite mass correction. But, that wasn't the point. Keeping a finite, I simply followed the program outlined by Professor Bethe and showed how to calculate all the various things, the scatterings of electrons from atoms without radiation, the shifts of levels and so forth, calculating everything in terms of the experimental mass, and noting that the results as Bethe suggested, were not sensitive to a in this form and even had a definite limit as $a \rightarrow 0$.

The rest of my work was simply to improve the techniques then available for calculations, making diagrams to help analyze perturbation theory quicker. Most of this was first worked out by guessing—you see, I didn't have the relativistic theory of matter. For example, it seemed to me obvious that the velocities in non-relativistic formulas have to be replaced by Dirac's α matrices or in the more relativistic forms by the operators γ_μ . I just took my guesses from the forms that I had worked out using path integrals for non-relativistic matter, but relativistic light. It was easy to develop rules of what to substitute to get the relativistic case. I was very surprised to discover that it was not known at that time that every one of the formulas that had

been worked out so patiently by separating longitudinal and transverse waves could be obtained from the formula for the transverse waves alone, if instead of summing over only the two perpendicular polarization directions you would sum over all four possible directions of polarization. It was so obvious from the action (1) that I thought it was general knowledge and would do it all the time. I would get into arguments with people because I didn't realize they didn't know that; but, it turned out that all their patient work with the longitudinal waves was always equivalent to just extending the sum on the two transverse directions of polarization over all four directions. This was one of the amusing advantages of the method. In addition, I included diagrams for the various terms of the perturbation series, improved notations to be used, worked out easy ways to evaluate integrals which occurred in these problems, and so on, and made a kind of handbook on how to do quantum electrodynamics.

But one step of importance that was physically new was involved with the negative energy sea of Dirac, which caused me so much logical difficulty. I got so confused that I remembered Wheeler's old idea about the positron being, maybe, the electron going backward in time. Therefore, in the time dependent perturbation theory that was usual for getting self-energy, I simply supposed that for a while we could go backward in the time, and looked at what terms I got by running the time variables backward. They were the same as the terms that other people got when they did the problem a more complicated way, using holes in the sea, except, possibly, for some signs. These I, at first determined empirically by inventing and trying some rules.

I have tried to explain that all the improvements of relativistic theory were at first more or less straightforward, semi-empirical shenanigans. Each time I would discover something, however, I would go back and I would check it so many ways, compare it to every problem that had been done previously in electrodynamics (and later, in weak coupling meson theory) to see if it would always agree, and so on, until I was absolutely convinced of the truth of the various rules and regulations that I concocted to simplify all the work.

During this time, people had been developing meson theory, a subject I had not studied in any detail. I became interested in the possible application of my methods to perturbation calculations in meson theory. But, what was meson theory? All I knew was that meson theory was something analogous to electrodynamics, except that particles corresponding to the photon had a mass. It was easy to guess the δ -function in (1), which was a solution of d'Alembertian equals zero, was to be changed to the corresponding solution of d'Alembertian equals m^2 . Next, there were different kinds of mesons—the one in closest analogy to photons, coupled via $\gamma_\mu\gamma_\mu$, are called vector mesons—there were also scalar mesons. Well, maybe that corresponds to putting unity in place of the γ_μ , perhaps what they called “pseudo vector coupling” and I would guess what that probably was. I didn't have the knowledge to understand the way these were defined in the conventional papers because they were expressed at that time in terms of creation and annihilation operators, and so on, which, I had not successfully learned. I remember that when someone had started to teach me about creation and annihilation operators, that this operator creates an electron, I said, “how do you create an electron? It disagrees with the conservation of charge,” and in that way, I blocked my mind from learning a very practical scheme of calculation.

Therefore, I had to find as many opportunities as possible to test whether I guessed right as to what the various theories were.

One day a dispute arose at a Physical Society meeting as to the correctness of a calculation by Slotnick of the interaction of an electron with a neutron using pseudo scalar theory with pseudo vector coupling and also pseudo scalar theory with pseudo scalar coupling. He had found that the answers were not the same. In fact, by one theory, the result was divergent, although convergent with the other. Some people believed that the two theories must give the same answer for the problem. This was a welcome opportunity to test my guesses as to whether I really did understand what these two couplings were. So, I went home, and during the evening I worked out the electron neutron scattering for the pseudo scalar and pseudo vector coupling, saw they were not equal and subtracted them, and worked out the difference in detail. The next day at the meeting, I saw Slotnick and said, "Slotnick, I worked it out last night, I wanted to see if I got the same answers you do. I got a different answer for each coupling—but, I would like to check in detail with you because I want to make sure of my methods." And, he said, "what do you mean you worked it out last night, it took me six months!" And, when we compared the answers he looked at mine and he asked, "what is that Q in there, that variable Q ?" (I had expressions like $(\tan^{-1} Q)/Q$ etc.) I said, "That's the momentum transferred by the electron, the electron deflected by different angles." "Oh," he said, "no, I only have the limiting value as Q approaches zero, the forward scattering." Well, it was easy enough to just substitute Q equals zero in my form and I then got the same answers as he did. But, it took him six months to do the case of zero momentum transfer, whereas, during one evening I had done the finite and arbitrary momentum transfer. That was a thrilling moment for me, like receiving the Nobel Prize, because that convinced me, at last, I did have some kind of method and technique and understood how to do something that other people did not know how to do. That was my moment of triumph in which I realized I really had succeeded in working out something worthwhile.

At this stage, I was urged to publish this because everybody said it looks like an easy way to make calculations, and wanted to know how to do it. I had to publish it missing two things; one was proof of every statement in a mathematically conventional sense. Often, even in a physicist's sense, I did not have a demonstration of how to get all of these rules and equations from conventional electrodynamics. But, I did know from experience, from fooling around, that everything was, in fact, equivalent to the regular electrodynamics and had partial proofs of many pieces, although, I never really sat down, like Euclid did for the geometers of Greece, and made sure that you could get it all from a single simple set of axioms. As a result, the work was criticized, I don't know whether favorably or unfavorably, and the "method" was called the "intuitive method." For those who do not realize it, however, I should like to emphasize that there is a lot of work involved in using this "intuitive method" successfully. Because no simple clear proof of the formula or idea presents itself, it is necessary to do an unusually great amount of checking and rechecking for consistency and correctness in terms of what is known, by comparing to other analogous examples, limiting cases, etc. In the face of the lack of direct mathematical demonstration, one must be careful and thorough to make sure of the point, and one should make a

perpetual attempt to demonstrate as much of the formula as possible. Nevertheless, a very great deal more truth can become known than can be proven.

It must be clearly understood that in all this work, I was representing the conventional electrodynamics with retarded interaction, and not my half-advanced and half-retarded theory corresponding to (1). I merely used (1) to guess at forms. And, one of the forms I guessed at corresponded to changing δ to a function f of width a^2 , so that I could calculate finite results for all of the problems. This brings me to the second thing that was missing when I published the paper, an unresolved difficulty. With δ replaced by f the calculations would give results which were not “unitary,” that is, for which the sum of the probabilities of all alternatives was not unity. The deviation from unity was very small, in practice, if a was very small. In the limit that I took a very tiny, it might not make any difference. And, so the process of the renormalization could be made, you could calculate everything in terms of the experimental mass and then take the limit and the apparent difficulty that the unitarity is violated temporarily seems to disappear. I was unable to demonstrate that, as a matter of fact, it does.

It is lucky that I did not wait to straighten out that point, for as far as I know, nobody has yet been able to resolve this question. Experience with meson theories with stronger couplings and with strongly coupled vector mesons, although not proving anything, convinces me that if the coupling were stronger, or if you went to a higher order (137th order of perturbation theory for electrodynamics), this difficulty would remain in the limit and there would be real trouble. That is, I believe there is really no satisfactory quantum electrodynamics, but I’m not sure. And, I believe, that one of the reasons for the slowness of present-day progress in understanding the strong interactions is that there isn’t any relativistic theoretical model, from which you can really calculate everything. Although it is usually said that the difficulty lies in the fact that strong interactions are too hard to calculate, I believe it is really because strong interactions in field theory have no solution, have no sense—they’re either infinite, or, if you try to modify them, the modification destroys the unitarity. I don’t think we have a completely satisfactory relativistic quantum-mechanical model, not even one that doesn’t agree with nature, but, at least, agrees with the logic that the sum of probability of all alternatives has to be 100%. Therefore, I think that the renormalization theory is simply a way to sweep the difficulties of the divergences of electrodynamics under the rug. I am, of course, not sure of that.

This completes the story of the development of the space-time view of quantum electrodynamics. I wonder if anything can be learned from it. I doubt it. It is most striking that most of the ideas developed in the course of this research were not ultimately used in the final result. For example, the half-advanced and half-retarded potential was not finally used, the action expression (1) was not used, the idea that charges do not act on themselves was abandoned. The path-integral formulation of quantum mechanics was useful for guessing at final expressions and at formulating the general theory of electrodynamics in new ways—although, strictly it was not absolutely necessary. The same goes for the idea of the positron being a backward moving electron, it was very convenient, but not strictly necessary for the theory because it is exactly equivalent to the negative-energy-sea point of view.

We are struck by the very large number of different physical viewpoints and widely different mathematical formulations that are all equivalent to one another. The method used here, of reasoning in physical terms, therefore, appears to be extremely inefficient. On looking back over the work, I can only feel a kind of regret for the enormous amount of physical reasoning and mathematical re-expression which ends by merely re-expressing what was previously known, although in a form which is much more efficient for the calculation of specific problems. Would it not have been much easier to simply work entirely in the mathematical framework to elaborate a more efficient expression? This would certainly seem to be the case, but it must be remarked that although the problem actually solved was only such a reformulation, the problem originally tackled was the (possibly still unsolved) problem of avoidance of the infinities of the usual theory. Therefore, a new theory was sought, not just a modification of the old. Although the quest was unsuccessful, we should look at the question of the value of physical ideas in developing a *new* theory.

Many different physical ideas can describe the same physical reality. Thus, classical electrodynamics can be described by a field view, or an action at a distance view, etc. Originally, Maxwell filled space with idler wheels, and Faraday with fields lines, but somehow the Maxwell equations themselves are pristine and independent of the elaboration of words attempting a physical description. The only true physical description is that describing the experimental meaning of the quantities in the equation—or better, the way the equations are to be used in describing experimental observations. This being the case perhaps the best way to proceed is to try to guess equations, and disregard physical models or descriptions. For example, McCullough guessed the correct equations for light propagation in a crystal long before his colleagues using elastic models could make head or tail of the phenomena, or again, Dirac obtained his equation for the description of the electron by an almost purely mathematical proposition. A simple physical view by which all the contents of this equation can be seen is still lacking.

Therefore, I think equation guessing might be the best method to proceed to obtain the laws for the part of physics which is presently unknown. Yet, when I was much younger, I tried this equation guessing and I have seen many students try this, but it is very easy to go off in wildly incorrect and impossible directions. I think the problem is not to find the *best* or most efficient method to proceed to a discovery, but to find any method at all. Physical reasoning does help some people to generate suggestions as to how the unknown may be related to the known. Theories of the known, which are described by different physical ideas may be equivalent in all their predictions and are hence scientifically indistinguishable. However, they are not psychologically identical when trying to move from that base into the unknown. For different views suggest different kinds of modifications which might be made and hence are not equivalent in the hypotheses one generates from them in ones attempt to understand what is not yet understood. I, therefore, think that a good theoretical physicist today might find it useful to have a wide range of physical viewpoints and mathematical expressions of the same theory (for example, of quantum electrodynamics) available to him. This may be asking too much of one man. Then new students should as a class have this. If every individual student follows the same current fash-

ion in expressing and thinking about electrodynamics or field theory, then the variety of hypotheses being generated to understand strong interactions, say, is limited. Perhaps rightly so, for possibly the chance is high that the truth lies in the fashionable direction. But, on the off-chance that it is in another direction—a direction obvious from an unfashionable view of field theory—who will find it? Only someone who has sacrificed himself by teaching himself quantum electrodynamics from a peculiar and unusual point of view, one that he may have to invent for himself. I say sacrificed himself because he most likely will get nothing from it, because the truth may lie in another direction, perhaps even the fashionable one.

But, if my own experience is any guide, the sacrifice is really not great because if the peculiar viewpoint taken is truly experimentally equivalent to the usual in the realm of the known there is always a range of applications and problems in this realm for which the special viewpoint gives one a special power and clarity of thought, which is valuable in itself. Furthermore, in the search for new laws, you always have the psychological excitement of feeling that possibly nobody has yet thought of the crazy possibility you are looking at right now.

So what happened to the old theory that I fell in love with as a youth? Well, I would say it's become an old lady, that has very little attractive left in her and the young today will not have their hearts pound when they look at her anymore. But, we can say the best we can for any old woman, that she has been a very good mother and she has given birth to some very good children. And, I thank the Swedish Academy of Sciences for complimenting one of them. Thank you.

Postscript

The records at Caltech indicate that the written version of Feynman's Nobel Lecture originates from a transcript of a lecture delivered by Feynman at Caltech some time after the Nobel ceremonies. Copies of the transcript were provided to the Nobel Foundation and to the editors of *Science* and *Physics Today*. The Nobel Foundation published the lecture in *Les Prix Nobel en 1965*, Norstedt, 1966, in *Nobel Lectures, Physics, 1963-1970*, Elsevier, 1972, and it appears in *The Selected Papers of Richard Feynman*, World Scientific Press, 2000.

This version of Feynman's Nobel Lecture was prepared to improve the readability of the text by correcting many small errors that appear in the previously published versions.

Michael D. Godfrey
Stanford University

Michael A. Gottlieb
California Institute of Technology

March 2008