CHAPTER 4
THE MARRIAGE OF RELATIVITY AND QUANTUM THEORY
Relativistic Quantum Mechanics

Almost all of us have a pretty good notion of what the word 'theory' means. Relativity and quantum mechanics, two of the other three ingredients of relativistic quantum field theory, have been introduced at some length in the previous chapter. What we have yet to talk about at any length is the idea of a 'field' so let's begin there.

I. FORCE FIELDS

As briefly noted before, our original descriptions of gravitation and electromagnetism – Newton's theory of gravitation (or Einstein's later 'general relativistic' formulation of the theory of gravitation) and Maxwell's theory of electromagnetism – predate the development of quantum mechanics, and are commonly thought of as 'classical' theories. In classical theories, the charge associated with the force in question generates a force field, or simply, a 'field', which can be represented by relatively straightforward functions of space and time(***1). Mass (gravitational charge) generates a gravitational field, electric charge generates electromagnetic fields, and so forth. Although any given charged object is localized, i.e., exists at a well-defined point in space at any given time, the field generated by the charge extends throughout space. Thus, in this 'classical' picture, the *fields* are responsible for a phenomenon of 'action at a distance' – the ability of some object carrying the charge associated with a certain force to influence another object, not in direct contact with the first object, which also carries some of the charge appropriate for the force in question. The earth is in orbit around the sun – not because the sun is reaching out a fiery arm and turning the earth about it, but because the considerable mass of the sun is generating a gravitational field throughout the solar system, through which field the massive earth moves, and whose resulting continual tug keeps the earth from flying off into interstellar space.

The fields fully represent the potential for a given charge to exert a force on any other appropriately charged object. With the field functions in hand, one knows precisely how the other object will be influenced by the force, regardless of where or when the object is placed in the field of the original charged object. In fact, if one knows the field functions, one can forget about the original charge that generated that field. The force that a second charge would feel, if placed at a certain moment at some point in space where the field generated by the first charge is appreciable, is simply given by the value of the field at that point in space and time, and the properties of the second charged object alone (the magnitude of the second object's charge and its motion through the field).

For those not completely uncomfortable with mathematics, the following concrete example will serve to illustrate this point. If an electrically charged object, of electrical charge magnitude q_2, is placed at rest a distance r from a stationary electric charge of magnitude q_1, the the mathematical expression for the electric force exerted on the charge q_2 by the charge q_1 is quite simple, as was determined by a series of careful experiments by the French physicist Charles Coulomb in 1785. The 'electrostatic' or 'Coulombic' force F_2 exerted on the second charged object is, of course, fixed (constant) in time, and depends on the spatial separation r of the two charges according to the 'inverse square law'

$$F_2 = k\frac{q_1 q_2}{r^2};$$

doubling the distance between the two charges reduces the strength of the force by $\frac{1}{2}^2 = \frac{1}{4}$, and so forth. The 'electrostatic force constant' k is just a number, determined by experiment, which expresses the overall strength of the electrical force between two charged objects – k is a fundamental constant which is completely situation-independent; no matter what the charged objects and their charges and separation are, k always has the same value (***2).

We introduce the concept of the electric field by a purely mathematical manipulation. We separate the expression into two factors

$$F_2 = (q_2) \times (k\frac{q_1}{r^2}) \qquad (1)$$

and say, by fiat, that the 'electrostatic' *field* E generated by a charge of magnitude q_1 sitting at rest is just the second of these two factors

$$E = k\frac{q_1}{r^2}. \qquad (2)$$

The electrostatic *force* F_2 exerted on the second charge by the first charge can then be interpreted as being due to the effect of the *field* E of the first charged object on the second charged object, and is obviously thus governed by the mathematical relationship

$$F = q \times E. \qquad (3)$$

Note that, once we know the electric field E at every point in space and time, we can simply unconcern ourselves about what it is that generated that field – a single stationary charge q_1, as in this example, or a multitude of q_1's moving around in some arbitrary and complicated motion relative to the charge q_2. The electrostatic *force* on the charge q_2 will always be related to the resulting electric *field* according to the simple expression (3). Of course, in the case of many charges moving chaotically about our bewildered charge q_2, the expression (2) which determines the value of the electric field E at any point in space and time will be tremendously more complicated.

The astute reader may well recognize this whole argument as a somewhat vacuous mathematical manipulation – who cares whether we express the force between two charged objects directly via expression (1), or break it up into two steps and make use of expressions (2) and (3)? In fact, the latter seems less economical, arbitrarily requiring the introduction of an unnecessary physical entity known as 'electric field'. It's cute, but why bother?

The answer lies in the Maxwell's fully unified theory of electromagnetism, for in this theory, it is possible for time-varying electric and magnetic fields, in combination, to support each other, leading to *electromagnetic* field disturbances which happily propagate endlessly through free space, *in the complete absence of any nearby electrically charged objects*. So, the fields themselves do seem to possess an independent physical meaning. Furthermore, Maxwell's theory not only predicts the possibility of such electromagnetic field disturbances, it also relates the speed at which this disturbance travels through space to the electrical force constant k and the corresponding magnetic force constant μ_0(***3). When this relation was evaluated with the known values of k and μ_0, it was found that such a disturbance would propagate through space with a speed of about 299,800,000 meters per second, or about 186,000 miles per hour.

This, to most, is a familiar number, for it is just the speed of light. Maxwell had accomplished nothing less than explaining the origin of light, the single most important mechanism for the transport of energy and information in the universe. This is why Maxwell's work is considered to be amongst the crowning achievements in long history of the development of our understanding of the physical universe. Without the introduction of the notion of a 'field', this work would not have been possible.

In this chapter, we will introduce the modern description of the interaction (forces) between charged objects, which in light of the discussion of the last few paragraphs, we now appreciate must necessarily be based on a description of the fields generated by these charged objects. To be truly modern, of course, this description must be consistent with the most basic tenets of modern physics – relativity and quantum mechanics – the essential points of which were brought forth in the previous chapter. Hence, the modern formulation of the mechanism by which charged objects influence one another bears the somewhat daunting name of 'relativistic quantum field theory', which is often shortened to 'quantum field theory', or even just simply 'field theory'. In its full presentation, quantum field theory is a subtle and intricate body of conjectures and calculational techniques, and is typically taught to graduate students in physics in their second year of graduate study. At its core, though, lie a relatively small number of accessible, and rather intriguing principles, which we will discuss at some length below.

Finally, it should be mentioned that field theory is not really a theory – it is something of even greater stature than that. Field theory is really a *framework*, on which all successful theories of fundamental interactions are expected to be based. Current fundamental theories, which bear names such as the 'electroweak theory', 'quantum electrodynamics', 'quantum chromodynamics', 'quantum gravity', 'string theory', 'supersymmetry', 'M-theory', and the like, whether supported by experimental evidence or merely conjectural, are all various different deployments of the principles of quantum field theory. Why it is that physicists have so much confidence in this field-theory framework will be discussed at the end of this chapter.

II. THE QUANTIZATION OF FIELDS

In 'conventional' quantum mechanics, the subject of the previous chapter, objects move in the presence of continuous fields, whose field strengths vary smoothly from point to point in space. These fields are incorporated into the Schroedinger Equation (the quantum mechanical wave equation which the all-encompassing wave function must satisfy) via their associated 'potential function' $V(x)$, which just gives, roughly, the amount of work the object must do against the force in moving to the position 'x', for any position 'x'. The need for the wavefunction to satisfy the Schroedinger Equation associated with the particular potential function $V(x)$, combined with the requirement that the wave function be well-behaved (not infinite) everywhere in space, leads to the 'quantization' of the motion of the object – the object can only find itself in one of a finite number of possible states, each of which corresponds to a specific value of the object's total energy.

Quantum field theory takes this notion one step further, asserting that the *fields* themselves are quantized. This view of fields, as we shall see, is radically different than the classical notion, introduced just above, and made use of in conventional quantum mechanics.

Recall that, in Maxwell's celebrated theory of electromagnetism, light is understood to be nothing more than a self-supporting disturbance in the electromagnetic field that carries energy from the original source of the disturbance (such as a far-away wiggling electric charge) to the observer. Recall also our discussion of 'wave-particle duality': in certain contexts (such as the delivery to the surface of a metal of the miniscule amount of energy needed to free an electron from the metal), light behaves as if it were composed of a very large number of very tiny particles, known as 'photons'. These tiny, yet indivisible, bits of light are prime candidates for the 'quanta' of the electromagnetic field. The electromagnetic field, instead of being a function which varies in some smooth but arbitrary way over space, becomes instead something entirely different – an assemblage of photons, each with a specific energy given by

$$E = hf$$

where h is Planck's constant, and f is the frequency (color) of the photon. The wave functions which describe the photons are indeed solutions to a quantum mechanical wave equation like the Schroedinger Equation – but not quite the Schroedinger Equation, since photons always travel at the speed of light, and so the wave equation must incorporate the tenets of Einstein's special relativity, which the Schroedinger Equation does not. (This is where the term 'relativistic' enters into the expression 'relativistic quantum field theory'). More on exactly what sort of an equation the photon's wave function must satisfy will follow below.

III. EXCHANGE FORCES

With the fields themselves quantized into discreet bundles which carry energy and behave like particles, the notion of how forces are transmitted becomes radically altered – but luckily, in a fairly intuitive way. Consider two ice skaters facing each other on a frozen pond. Although they are at rest, a small push to either would send that skater sliding across the pond, since there is essentially no friction between the skater's skate and the pond to impede the motion.

Now, picture one of the skaters throwing a ball towards the other skater. As she thrusts the ball forwards towards the other skater, she recoils against the action, and as the ball flies away, she begins to slide backward at some fixed speed. The harder she throws the ball, the quicker she recoils, but once the ball is thrown, her speed is set, and it doesn't change again.

Of course, in catching the ball, the second skater absorbs its momentum and energy, and goes sliding back in the other direction. The speed at which he recoils after catching the ball again depends on how hard the ball is thrown, but once the ball is caught, his speed is set and it doesn't change again.

So, the result of the exchange of the ball is simply that both skaters go sliding off in opposite directions – *just as if they had exerted some sort of force on each other*. From the point of view of quantum field theory, this is *exactly* how forces are conveyed. Replace the two skaters with two electrons, and the ball with a photon. We see that from this point of view, the repellent force between the two electrons is nothing more than the result of an exchange of one or more photons – the quantum of the electromagnetic field – between the two electrons.

This picture is a far cry from the action-at-a-distance principle of classical field theories, discussed at the beginning of this chapter. On the other hand, once you know quantum field theory, with a little work you can show that for two charged objects sitting at rest, quantum field theory predicts that the force (and field) will have just the properties outlined in the beginning of the chapter, as if they were influencing each other via action at a distance. Quantum field theory doesn't really *invalidate* the classical theory of fields – it just substitutes a more refined notion of how those fields are generated, ultimately allowing for the extension of the 'field' interpretation of the interactions of matter into (quantum-mechanical) regimes for which the classical approach is inadequate.

Where do these photons come from that are exchanged between the two electrons as they repel each other? The answer is that they come from nowhere – they are created (thrown) and absorbed or annihilated (caught) for the sole purpose of transmitting the force. Before and after their brief passage between the electrons, they do not exist. But, as we discussed before, this is perfectly consistent with relativity, for it is no longer true that matter can neither be created nor destroyed. Instead, what is true is that *energy* (as long as mass is properly accounted for as energy according to the expression $E = mc^2$) can neither be created or destroyed; energy is conserved. So, if at some instant in time there are just two pieces of matter (two electrons), and at a slightly later instant there are three pieces of matter (two electrons, plus one photon – one piece of 'field' matter), everything is just fine, as long as the total *energy* of the system is unchanged.

A very interesting point arises if we change the situation slightly by throwing the two electrons at each other (rather than having them repel each other at rest). In this case, working out the energy balance, making sure that at every step of the way energy is conserved, we would find that while the exchanged photon is alive, it would have to have some amount of mass. The exact amount of mass required of the photon depends on the how hard the electrons are thrown at each other and how close they come before they repel, but in this case the photon always has to have *some* amount of mass.

But we know that photons, which compose light, don't weigh anything – how much does a light beam weigh? (Experimentally, we know that the mass of a photon is less than 10^{-51} kg, where a kg is about 2 pounds.) Here, we are saved by Heisenberg's uncertainty principle, discussed at some length in the preceding chapter. You may recall from this discussion that the product of the uncertainty in the energy and lifetime of a particle must be greater than $h/4\pi$, where again h is just Planck's constant. With a little algebraic manipulation, we can write this as

$$\Delta E = \frac{1}{\Delta t}\frac{h}{4\pi}.$$

Since the photon, as discussed above, lives for only a very short time Δt (while it is flying at the speed of light between the two electrons), it has a substantial uncertainty in its energy, and, since mass and energy are essentially the same, its mass. Thus, it is fully consistent with the tenets of quantum mechanics that a given particle, known experimentally to have certain mass (e.g., zero for the photon) can have a somewhat different mass if it lives only for the very brief period as it is being exchanged between two interacting objects. Technically, we say that such a short-lived particle, required by energy conservation to have a mass different than its known value, but perfectly allowed to according to quantum mechanics, is 'off mass-shell'(***4). Such particles are also referred to as 'virtual' particles.

IV. THE FOUR FORCES AGAIN

To reiterate, we've just seen that, from the point of view of quantum field theory, the electromagnetic force is described as the interaction between electrically charged objects that is caused by the exchange of photons – the quanta of the electromagnetic field. It's easy to extend this notion to the description of the other three forces. Quite simply, each force is generated by the exchange of the particle – the quantum – of the associated force field between objects that carry the charge appropriate for that force (electrical charge for the electromagnetic force, mass for the gravitational force, 'color charge' for the strong nuclear force, 'weak isopsin' for the weak nuclear force).

Thus, from the point of view of quantum field theory, the forces are delineated simply by the identity of the quantum (or quanta) exchanged between appropriately charged objects. To describe any given force, we need simply to state what the associated force-field quanta are, and describe their properties (for example, mass, charge, and 'spin', the latter of which will be described below). Developing the theory of any given force boils down to simply identifying the field quantum that gets exchanged when the force is play, and determining its properties.

In the case of the electromagnetic force, there is just a single field quantum (the photon), but in general there need not be just one. In fact, the weak nuclear force has associated with it not one but three separate field quanta. Two of these, the 'W bosons', are like each other in every way, except they have opposite *electrical* charge; the third, the 'Z^0 boson' (or simple 'Z boson') is electrically neutral, i.e., has no electrical charge. One can ask why it is that two of the three quanta of the *weak nuclear force* carry the charge associated with the *electromagnetic* force – for now, we'll simply say that this is merely one of the properties of the weak force field quanta that make the weak force what it is, and leave it at that. Of course, in the bigger picture, this just might be an important clue regarding the eventual unification of the theories of the weak and electromagnetic forces, but we'll have to leave this intriguing thread for later.

Quantum field theory is a great leap forward in our understanding of nature. Not only does it (as we'll see at the end of the chapter) provide a description of nature with phenomenal quantitative power, it also simultaneously simplifies and extends our description of the natural forces. For example, Maxwell's classical theory of electromagnetism, triumph that it was, was stated in terms of four rather intimidating interrelated (differential) equations, the development of which spanned a century of painstaking experimentation and brilliant theoretical strides. With quantum field theory, one need only identify and describe the electromagnetic field quantum (the photon), and the job is done. Not only that, but the resulting quantum field theory of electromagnetism is more broadly applicable, reproducing the classical theory when the energy associated with the electromagnetic interaction is large enough that quantum effects can be ignored, while providing the appropriate modification of that description in the case that such effects can not be ignored. Even more, the particular properties of the strong and weak nuclear forces (in particular, the fact that these forces are short-ranged, operating only over nuclear scales of 10^{-15} meters or less) make it impossible to deduce the corresponding, Maxwell-like classical field equations. The description of these two forces is only possible within the framework of quantum field theory. And of course, if we hold the 'unification' of the theories of the four forces into a single comprehensive theory as our ultimate goal, it is necessary that the initial individual theories of the forces be expressed within the same framework. Clearly, quantum field theory provides this framework, and is an essential first step towards this loftiest of goals.

Finally, it's a bit of an oversimplification to say that the statement of the quantum field theory associated with a given force boils down to simply identifying and describing the properties of the associated field quanta – there's another ingredient which must be included. This ingredient is the nature of the 'coupling' between the field quanta and the charged objects which are tossing these quanta back and forth – the details which govern the process of the creation of a given field quantum by the 'tossing' object and the absorption or annihilation of the quantum by the 'receiving' object. One aspect of the description of this 'coupling' is simply its overall strength – the stronger the coupling, the more quanta are exchanged in a given interaction (or the greater the likelihood that a quanta will be exchanged at all), and the stronger the force is. A good way to think of it is that the strength of the force's coupling to a given fundamental particle is related to the magnitude of the associated charge carried by the particle.

A second aspect of the description of this coupling between the field quantum and charged objects is referred to as its 'spacetime property'. Recall that we identified the weak nuclear force as being particularly interesting because of its tendency to irreverently violate some basic symmetry principles, such as the equitable treatment of matter and antimatter. It is the spacetime property of this coupling which determines, among other things, whether this (and other) symmetry principles are respected or violated, and if so, to what extent.

In practice, though, this is not as much of a complication as it may appear, because there are only a handful of different spacetime properties available to the coupling between field quanta and charged objects which are consistent with the tenets of Einstein's special relativity. In addition, once the field quantum is identified, its characteristics (specifically, its 'spin') restrict the possible spacetime properties of the coupling even further. So, once the field quanta are identified and their properties determined, it's usually fairly clear what the appropriate spacetime property of the coupling must be – and then the theory is essentially complete. A full list of the quanta associated with each of the four forces, and their essential properties, lies somewhat ahead (at the end of Chapter 5).

The one exception to this is in fact the aspect of the spacetime property of the weak nuclear coupling leading to the violation of matter/antimatter symmetry. In this case, while the representation of this property within the Standard Model is still rather simple, being specified by four numbers which much be determined by experiment, the determination of these numbers is a rather daunting experimental challenge, and is one of the driving forces behind a number of experiments that are either currently taking data or propose to begin so within the next ten years. The exciting prospect that matter/antimatter symmetry violation in the weak nuclear force may be found in the end *not* to be correctly described by the Standard Model with its four experimental inputs is an intriguing possibility that is on many people's minds. But we won't know until after these experiments have been run and analyzed.

V. FEYNMAN DIAGRAMS

One of the primary proponents of the development of quantum field theory, and its initial use in the description of the electromagnetic force, was the American physicist Richard Feynman(***5). Feynman's contributions play a central role in the formal, quantitative structure of quantum field theory, most of which lies well beyond the scope of this book. One of his contributions, however, provides a very straightforward way to determine and describe the possible interactions between objects subject to the influence of a given force. These descriptions, known as 'Feynman diagrams', will be of great use to us throughout the remainder of this book.

In order to understand the meaning of a given Feynman diagram (at the level appropriate for the discussion in this book), there is one mathematical hurdle to overcome: the description of the motion of a particle in terms of a graph of position versus time – such a graph is known as a 'spacetime plot' or 'spacetime diagram'. Take a look at Figure 4.1, which is just the coordinate axes of a graph, with the horizontal ('x') axis representing the position of the object, and the vertical ('t') axis representing the time at which the object has the given position x in space. By convention (it's an important convention, so please take note) distance increases from left to right, and time increases (elapses) from bottom to top. If you like, you can think of the x axis as a meterstick (or yardstick) with numbers increasing (as you would expect) to the right, while the t axis is represents the readings off of a ticking stopwatch.

Figure 4.2 shows how a particle 'at rest', i.e., with a speed of zero, would be represented on such a graph. If it has a speed of zero, then it's not going anywhere, so its position, represented by its location along the x axis, is fixed in time. You might expect such a particle to be represented by a dot on this graph, but remember – time waits for no man, or, more generally, for no object, so for this particle at rest, time marches inexorably forward. Thus, even though this particle is happily at rest, going nowhere in *space* (i.e., possessing an unchanging position along the space, or x axis), it is traveling through *time*, and so it is represented by a line parallel to the t axis – a line which has associated with it a constant value of the spatial position x, but an ever increasing value of the associated time t. Furthermore, this passage through time has a well-defined 'sense' – things always travel forward in time, we are told – which is represented by the arrow.

Now, take a glance at Figure 4.3a. The dotted lines show us that when the particle is at the (spacetime) point A, it has time T_1 and position X_1. At some later time T_2 (remember that time elapses from bottom to top), the particle is at the spacetime point B, and has a greater reading X_2 for its spatial position – as time elapses, the particle moves towards larger and larger readings on the meterstick. Thus, this graph represents a particle moving through space, from left to right, with some speed. The exact value of that speed, in meters per second, is not important for us to specify. The arrow on the line again just represents that fact that the particle is moving forward through time. Similarly, you should be able to convince yourself that the line in figure 4.3b represents a particle moving with some speed, but in this case going in the opposite direction – from right to left.

With this in hand, we are ready to draw our first Feynman diagram. Remember that these diagrams are Feynman's way of depicting the interaction of two object via a given force – just to be definite, we'll pick the electromagnetic force, and for the objects we'll pick two electrons, which carry the charge ('electric charge') appropriate for having them influence each other ('interact') via the electromagnetic force. Since this is supposed to be a 'Feynman' diagram, and we know how much Feynman loved his quantum field theory, this diagram had best represent the conveyance of the electromagnetic force from the point of view of field theory.

So, now take a look at Figure 4.4a. Using the interpretive skills developed in the last few paragraphs, what we see represented there is the following. At early times (towards the bottom of the plot), two electrons are approaching each other – one from the left moving towards the right, and the other from the right moving towards the left. At point 'A', the electron on the left emits (throws out) a photon (labelled 'γ') , the quantum of the electromagnetic field. This photon carries energy and momentum away from the electron that emits it, causing that electron to recoil and start moving back towards the left. The photon moves to the right until it is absorbed by the second electron at point 'B', at which point the second electron recoils back to the right. Of course, this process of photon exchange occurs over a *very* short distance – typically, the length of the photon's path will have atomic dimension (10^{-9} meters) or less – and the details of the photon exchange process are lost to the observer. What the observer would see is two electrons, originally moving towards each other, repelling each other and 'scattering' back in the direction they came from. Of course, what this diagram represents is the simplest way that two electrons can influence each other via the electromagnetic force, from the point of view of quantum field theory.

Figure 4.4b is very similar to 4.4a, except in this case the photon is emitted by the *right moving* electron at point 'B', and absorbed by the *left-moving* electron at point 'A'. Again, we have to step back and ask what the observer would see: two electrons, originally moving towards each other, repelling each other and 'scattering' back in the direction they came from. In other words, exactly what the observer would see in Figure 4.4a. The processes are identical, in the sense that there's no way for the observer to tell whether process 4.4a or 4.4b is what caused the electrons to scatter off of each other. In order to do that, the observer would have to cut into the diagram and detect the photon – but then the observer's apparatus has absorbed the photon, and it doesn't make it over to the other electron to play out the transmission of the force. In quantum mechanics, you may recall, there's no way to observe the system without disturbing it. So, if you *do* want to observe a given process, such as the scattering of two electrons via the electromagnetic force, you can only do it by measuring the system well before and then again well after the process has taken place, and the comparing the two measurements.

Since we can't tell in detail what went on during the interaction process, i.e., whether the electron on the left or the electron on the right threw out the ('virtual') photon, then there's no point in distinguishing between the two. The observable process is the combination of the two different indistinguishable and identical pieces – the one of 4.4a and 4.4b. Thus, we night as well just represent the process in the way shown in Figure 4.4c, which we know will always really mean the combination of 4.4a and 4.4b. Figure 4.4c is the true 'Feynman diagram' for this process.

The astute reader might be somewhat bothered by the fact that our Feynman diagram only represents motion in one dimension. We know, however, that things in general, and electrons in particular, move through three dimensional space, and not just back and forth on some line. Even if the electrons are approaching each other from opposite ends of a line segment, after they influence (scatter from) each other, they will usually be found going off in some completely different direction, which requires the other two dimensions of space to describe. So, the Feynman diagram 4.4c, you would say, can't really represent, in full generality, the process of the scattering of two electrons via the exchange of a single photon.

In fact, this is *not* the case. No matter how many dimensions you need to describe the motion of the electrons before and after the scatter, it's still fundamentally the same process – two electrons influencing each other via the exchange of a single photon. The purpose of Feynman diagrams is to categorize and describe the different possible *types* of interactions between a given set of interacting particles via a given force, and *not* to represent the exact motion (speed and direction) of the particles after the interaction. After all, this is quantum mechanics, and the uncertainty principle guarantees for us that the result *can not* be the same every time the scattering is performed. What characterizes the *type* of interaction taking place is simply the list of particles taking part in the interaction (two electrons and a photon in our case), and the pattern of interconnection of these in the Feynman digram (the photon connecting between the two electron trajectories in our case). Mathematicians would like us to say 'the geometry of the diagram doesn't matter, just the topology'. [Don't worry if you have no idea of what the mathematician might mean by this statement. In real life mathematicians are seldom, if ever, understood.]

On the other hand, the reason why Feynman diagrams are such useful representations of the interactions between particles is that the diagram contains all of the information that one can possibly know about the process that is represented (here, the electromagnetic interaction between two electrons conveyed, or 'mediated', by the exchange of a single photon). A person who knows quantum field theory can look at the Feynman diagram associated with a given process and, with enough paper and pencil lead (or pen ink if they're *really* experienced) turn this diagram into an explicit, quantitative calculation of everything that can possibly be predicted about the interaction.

To be specific, this calculation will reveal two things. The first, known as the 'total cross section', is just the probability that the two particles moving towards each other will actually exchange the photon and interact, rather than just passing by each other unheeded, like two ship passing in the night, if you will. The second, known as the 'differential cross section', is just the relative probability of any given final result of the scatter, given that the scatter really did occur. If one picks any possible energy and direction of the particles after the interaction has taken place, the differential cross section provides the probability that that particular result will be achieved, relative to the probability of any other possible result. What more could you ask for in a process like this, where two electrons influence each other for an exceedingly brief period of time, and then go on their merry scattered ways as if they had never met each other? (***6) The description is complete – it gives us all the possible a-priori (before-the-fact) knowledge admitted by the probabilistic nature of quantum mechanics.

VI. VERTICES AND THE MINIMAL INTERACTION

In this chapter, we have come to know the electromagnetic force as the force between electrically charged objects caused by the exchange of one or more photons between the charged objects. However, as we'll shortly see, the use of Feynman diagrams allows us to generalize (expand) our notion of what the electromagnetic 'force' really is, and in fact, they'll even incline us to drop the use of the expression 'electromagnetic force' in favor of the more general notion of the 'electromagnetic *interaction*'. In turn, this will force us to refine somewhat our view of what the electromagnetic interaction really is, at its most fundamental level. We will do this for more than reasons of mere philosophical interest; in fact, this refinement will be essential if we are to continue into the discussion of the developments beyond quantum field theory which led to the formulation of the Standard Model of particle physics.

Take another look at Figure 4.4c, the Feynman diagram depicting the mutual repulsion of two electrons via the exchange of a single ‘virtual’ photon. Question: is this the only possible type of interaction between these two electrons? Is there another possible way, within the paradigm of quantum field theory, that these two electrons can influence each other? If there is, then we ought to be able to depict this new process with another Feynman diagram. Figure 4.5 is just such a process – the interaction of two electrons via the exchange of *two* photons. From what we know about quantum field theory, there's no reason why this shouldn't happen – and it does! In the words of Feynman himself, "anything which *can* happen *does* happen". And, since good old-fashioned quantum mechanics lies at the heart of this picture, it's not that the force is exerted via the exchange of one photon some of the time, and two photons other times. Rather, any given time that the electrons repel each other and go flying back, there's a certain probability that it was due to the exchange of one virtual photon, and a certain other probability that it was due to the exchange of two virtual photons. We can never know (without disturbing the interaction) which it was, so it was both, every time. Of course, this progression doesn't stop with the exchange of one or two photons – any number of photons can be exchanged, and so the true electrostatic repulsion of the two electrons is effected by some sort of a combination of *all* of these (and other!) possibilities. More on this in a few pages.

But that's a bit of a digression. The real point is that, when you look at interactions between particles from the point of view of quantum field theory, there's always more than one – in fact an *infinite* – number of types of ways that any given process (such as the scattering of two electrons via the electromagnetic force) can take place, each one represented by its own unique Feynman diagram. If this is the case, though, then how can quantum field theory be of any use at all - there's an infinite number of difficult calculations to do every time you want to make a prediction about some process. The electrons can exchange *any number* of virtual photons – one, two, six, forty, whatever.

What saves us is the *coupling* between the photon and the electron. Look again at Figure 4.5. Every time an electron and photon meet in the Feynman digram, we have to remember to take into account the *coupling strength* between the electron and the photon (recall the discussion in the section entitled ‘The Four Forces Again’). This coupling is less than one – in fact, for the electromagnetic force, it's about 1/100(***7). So, for every intersection between and electron and photon we add to get the new diagram, we pay a price of about 100 in the probability that that diagram contributes to the overall process. Comparing Figures 4.4c and 4.5, we see that 4.5 has two extra electron-photon intersections, and so we would expect it to be roughly (1/100)*(1/100) = (1/10,000) as likely as figures 4.4c. So, it doesn't make a whole lot of difference if you include the more complicated diagrams involving two, three, and more exchanged photons. Only if you're interested in getting a *really* precise prediction would you be compelled to include these more complicated 'higher order' diagrams.

This discussion, it turns out, takes us somewhere quite interesting. If we take it seriously, and think for a moment, we might become inclined to start thinking of these *intersections*, or 'vertices', between electrons and photons – rather than the *exchange* of photons – as the defining element of the electromagnetic force. Note that every photon exchange involves two such electron-photon vertices, so perhaps the vertices are indeed more fundamentally characteristic of the electromagnetic force than is the notion of photon exchange.

If this is the case, we can make use of our new fundamental element – the vertex – to construct a new and very interesting Feynman diagram. Figure 4.6 shows this diagram. It begins, at early times towards the bottom of the diagram, with a real photon coming from the left, approaching an electron coming in from the right. (By 'real' we mean 'not virtual', i.e., a photon that was emitted from some sort of light source very long ago by quantum mechanical time scales, and so is *not* permitted by the uncertainty principle to have a mass other than its true mass of 0). At point A, the photon is absorbed by the electron at a photon-electron vertex – the vertex which, by our most current thinking, is the fundamental and characterizing happenstance of the electromagnetic force. Note that the vertex at A is identical to either of the vertices in Figure 4.4c – a wiggly line representing a photon terminates on a solid line representing an electron. Then, at point B. the digram employs another vertex, via which the electron spits out a photon (a *different* photon, which in general has a different energy as the initial one). Technically, in order to conserve energy at all times, the mass of the electron has to be slightly different than its true mass during the time interval between the absorption of the first photon and the emission of the second, but, just as for the virtual photon of 4.4c, that's fine because of the uncertainty principle. In this case, instead of having a virtual *photon*, we have a virtual *electron.*

Clearly, interpreting the fundamental nature of the electromagnetic force in terms of electron-photon vertices, rather than photon exchange, has allowed us to expand our list of processes governed by that force. From this more general point of view, the photon exchange of Figure 4.4c is just one possible mode of interaction afforded by the electromagnetic force. If we consider the more basic *vertices* as the fundamental description of electromagnetism. and play around a bit, we get the digram of Figure 4.6 – a whole new process, in which an electron beam can scatter a photon (light) beam by absorbing the photons and re-emitting them with different directions and energies.

Of course, to confirm that this picture is correct, all we need to do is go to the lab and discover such a process. In fact, as is often the case, the historical sequence of events was the exact opposite. Such an effect had already been discovered and explained by the American physicist Arthur Compton in the 1923 (Nobel Prize, 1927), well before the development of quantum field theory (recall that it was Compton who coined the term 'photon' for Einstein's electromagnetic field quantum). So, apparently, it *is* the more basic and general electron-photon vertex which we should look to as the fundamental component of the electromagnetic force, and not photon exchange, which is only one of several modes of interaction which can be constructed from electron-photon vertices. This most basic, irreducible action of electromagnetism is known as the 'minimal interaction'. It is not too great of an oversimplification to state that the job of anyone developing a theory to describe any given fundamental force (be it a description of any of the individual forces, or a theory of their unification, or whatever) is to make a list of all such 'minimal interactions' associated with the force, and to delineate the strength and spacetime properties of those interactions.

For electromagnetism, there is only one such minimal interaction – the electron-photon vertex (or, more generally, the fermion-photon vertex; the concept of a fermion will be introduced in the next section). It has a coupling strength equal to about 1/100, and the spacetime property of a 'vector'. The spacetime property designation of 'vector' certainly has no meaning to most readers of this book, it's just thrown in for completeness' sake. In any regard, that's it – that's basically all there is to our quantum field theory of electromagnetism, known as 'quantum electrodynamics', for which development Feynman and two others (Schwinger of the United States and Tomonaga of Japan) won the 1965 Nobel Prize. Well, that's not quite all, for as construed just above, the theory doesn't quite work, and in fixing it up, we'll enjoy yet another fascinating leap into the world of the counterintuitive. This discussion will have to wait a few pages, though.

Note that no one ever talks about the electric (or electromagnetic) force between electrons and light. And for good reason – there isn't any! Light doesn't get bent by charged objects. Take a key full of a painfully large amount of static charge (easily produced on a dry day on a plastic playground slide) and hold it up to the light - from the sun, a light bulb, or a laser beam. The light will not bend towards or away from the key. Electrically charged objects to not exert a force on light. But they do *interact* with it, by absorbing and re-emitting it in a different color (with a different energy). On the other hand, two statically charges objects will exert a force on each other. In both cases – light and charge or charge and charge – we understand the fundamental workings of the interaction to be describable in terms of a single fundamental component – the electron-photon vertex. Only in one of the two cases is a force exerted, but in both cases there is an *interaction*. Thus, in quantum field theory, we speak most correctly in terms of theories of fundamental *interactions*; 'forces' (as we usually construe them in terms of the action-at-a-distance interplay between two appropriately charged objects) are just one facet of the more general notion of an 'interaction'.

In a nutshell: in the mid-1800's, the introduction of the concept of the 'force field' liberated our thinking in a way that eventually led to the explanation of the phenomenon of light. In the mid-1900's, the *quantization* of these very same fields led to a further liberalization of our thinking, allowing the very concept of 'force' itself to be generalized to the notion of 'interaction', and permitting the explanation and theoretical representation of a host of new behaviors of the fundamental constituents of the universe.

VII. RELATIVISTIC QUANTUM MECHANICS, ANTIMATTER, AND SPIN

At this point, we need to suffer through a brief digression regarding some of the 'relativistic' aspects of relativistic quantum field theory (which we have shortened to simply 'quantum field theory' in most of the discussion above). In fact, to engage in this digression we need to introduce the concept of *angular* momentum. So, we begin with some background to our digression, and then proceed to the digression forthwith.

To begin with, let's take another look at Planck's constant h. Recall that Planck's constant sets the scale for the most basic quantum mechanical behavior, including the wavelength of pieces of matter (de Broglie's relation: $\lambda = p/h$, where p is the object's momentum) and the uncertainty principle $\Delta p \Delta x \geq h/(4\pi)$. The 'unit' of Planck's constant – the measuring stick used to establish its (minute) size – is Joule-seconds, where the Joule is the unit of energy. The Joule-second, it turns out, is the unit of something called '*angular* momentum'.

To get a feel for the notion of *angular* momentum/energy, consider a carrousel at an amusement park spinning on its axis.(***8) Your intuition correctly tells you that the carrousel, during the ride, has a lot of energy of motion. If you were to try to stop the turning by planting your feet and grabbing one of the horses, the horse would most certainly win out. Yet this energy of motion (kinetic energy) is not in the form that we are used to. This kinetic energy is not associated with the motion of the carrousel through space – the carrousel is fixed at a point, and despite its great energy, poses no threat to anyone with enough wits not to walk into it. The energy derives not from the motion of the carrousel through the amusement park, but rather from the *spinning* of the carrousel about its axis.

Analogously to the carrousel, particles such as protons and electrons can also spin about their axes. Unlike the carrousel, though, fundamental particles can not be made to spin faster or slower, or have their spinning motion stopped. The energy of spinning – the 'angular momentum' – is a *fixed property* of the particle. Just like a particle's mass, the angular momentum associated with this spinning motion, known simply as 'spin', is a characteristic of the particle at hand. The direction of the axis about which the particle is spinning *can* be changed, but the energy of spinning about that axis, no matter which way it points, can not be changed.

We've seen that the behavior individual particles requires quantum mechanics for its understanding and description. Thus, we might expect that quantum mechanics, and in particular its attendant scale factor h, might have a role to play in the spin of fundamental particles, especially since the unit of h, as we've seen, is just that of angular momentum.

This is in fact the case: the amount of angular momentum possessed by an electron or a proton is just $(1/2)\hbar$, where for the sake of expediency we have defined the 'reduced Planck's constant' $\hbar = h/2\pi$. On the other hand, the amount of angular momentum possessed by an individual photon is just $\hbar$. We often forget about the $\hbar$ and just say that electrons and protons are 'spin-$\frac{1}{2}$', and that photons are 'spin-1'. Some other objects that we'll talk about later have even different values for their spin. 'Pions' have spin-0 (they never have any energy associated with spinning), and 'delta baryons' have spin-$\frac{3}{2}$ (angular momentum $3/2\hbar$). This notion of the particle 'spin' plays a central role in latter chapters, so you'll want to make sure you have a good grasp of the last couple paragraphs before moving on.

As a final point of nomenclature: the set of particles with 'half-integer' spin (with spin equal to something-halves $\hbar$) are known as 'fermions', while particles with integer spin (with spin a pure multiple of $\hbar$) are known as 'bosons'. When in Chapter 5 we introduce the set of particles currently thought to be 'fundamental' (i.e., indivisible), we'll see that spin-$\frac{1}{2}$ fermions (such as electrons and quarks) are the components of what is conventionally thought of as 'matter', while the fundamental bosons (such as the photon) are the conveyers of force. So much for the background to the digression. Now for the digression itself.

Schroedinger's and Heisenberg's theory of quantum mechanics was a stunning success – as long as none of the particles in the system being studied were traveling so fast that Einstein's relativity needed to be taken into account. Since quantum mechanics was developed in the mid 1920's, 20 years after Einstein's original papers on special relativity, relativity was already part of the credo of physics when the quantum theory was starting to come together. So, almost immediately after the success of Schroedinger and Heisenberg, there was a growing dissatisfaction with the apparent inconsistencies between the world of the very small (quantum theory) and the very fast (special relativity). A number of young physicists turned their attention to this problem, and by the late 1920's, substantial progress had been made, along with the reputation of these youngsters(***9) as being the brightest of their generation. The imposition of relativity upon the new quantum theory did not come easily, but through these young physicists' brilliant re-interpretation of what seemed like fatal flaws in their relativistic quantum theory, a much deeper and richer understanding arose from this work than anyone had anticipated.

Recall our discussion of the Schroedinger equation. What we claimed is that this equation is nothing more than the quantum-mechanical statement of energy conservation, i.e., that the energy due to motion (kinetic energy) and the energy due to fighting against the field of whatever force you are considering (potential energy) is just equal to the total energy, the latter of which is a number that is fixed for all time (since energy is a conserved quantity) as long as the system under consideration is not disturbed from the outside. However, relativity tells us that the *mass* of the particle that we are trying to describe quantum mechanically had better also be included in the total energy, requiring the addition of another term in the wave equation.

To keep things from being too complicated, the proponents of the relativistic theory decided to begin with the description of 'free particles' – particles moving freely through space, *not* under the influence of some external force. So, the term in the Schroedinger equation associated with fighting back and forth against the force, the 'potential energy' or $V(x)$ term, could be forgotten about. It's important to note that this is no loss at all, for in the relativistic quantum *field* theory which grew out of this formulation of relativistic quantum mechanics, the forces are re-introduced via the minimal interaction and the exchange of (free) field quanta, as discussed at length above. So, free-particle relativistic wave equations are all that we'll ever need to consider.

So, what one ends up with is an equation which is essentially the quantum mechanical version of *relativistic* energy conservation (for a free particle): the sum of the kinetic energy and the energy associated with the particle's mass is just the total energy, which again must be conserved. (Looking at what we already know about quantum field theory, we see that if we create a virtual particle at a minimal interaction vertex, it's now OK, because the mass of this newly created particle is correctly taken account of in this *relativistic* formulation of quantum mechanics – the modeling of a force via the creation and absorption of virtual particles is an intrinsically relativistic point of view!) The resulting (differential) equation is known as the 'Klein-Gordon equation', and there's nothing to stop us from solving it to find the wavefunction of a particle moving with relativistic speed.

There is, however, something to make us very unhappy with the some of the wavefunctions we wind up with. For one, when we ask what value of energy is associated with the various possible solutions of the Klein-Gordon equation, we find that half of them have *negative* energy! In addition, recall that, in conventional quantum mechanics, when you square a particle's wavefunction (multiply it by itself at every point in space), you wind up with another function which just gives the probability of finding that particle at any given point in space. When one squares such a wavefunction which comes from s negative-energy solution to the Klein-Gordon equation, the probability function that one winds up with is negative.(***10) Negative energy solutions with negative probabilities – this does not sound healthy!

This certainly seemed like a fatal flaw to Dirac, who nevertheless held the conviction that a meaningful relativistic wave equation ought to exist. Dirac carefully analyzed the failings of the Klein-Gordon equation to find the cause of the negative probability solutions, and was able to propose a wave equation with the same physical content (kinetic energy plus energy due to mass equals total energy for a free particle), but which avoided the plague of negative probability. This task was not particularly straightforward, for what Dirac ended up with was a 'matrix (differential) equation' – four separate but intertwined differential equations which needed to be simultaneously satisfied by any allowable wavefunction. Messy though this approach may seem, it immediately provided two wonderful dividends: the negative probabilities were gone, and the new equation, in its complexity, precisely incorporated the behavior of spin-$\frac{1}{2}$ particles, which had been delineated by Pauli a few years earlier.

With these successes in hand, Dirac was convinced that the final hurdle – the explanation of the 'negative energy' solutions – must be surmountable. His proposal was that the negative energy solutions could be interpreted as positive energy solutions for *antimatter* particles, which would behave in every way like their particle counterparts, but have opposite charge. this suggestion was unspeakably bold – no one had seen or even thought of antimatter at that time; it wasn't even yet the stuff of science fiction novels. Dirac dreamed it up out of the blue to satisfy his conviction that his candidate for a relativistic wave equation describing spin-$\frac{1}{2}$ particles had to be correct. Antimatter was simply the ingredient that he needed to make it all work out. So he proposed the existence of a whole new set of particles, with very specific properties, that nobody had ever dreamed of before.

Five years later, in 1932, Carl Anderson of the California Institute of Technology, in a balloon experiment, observed a particle with all the properties of an electron, except that it bent the 'wrong way' in the experiment's magnetic field. Anderson immediately knew what he had seen: the positively charged partner of the negatively charged electron; Dirac's antimatter *positron*. It was not long after this that both Dirac and Anderson found invitations to Stockholm in their mail (Dirac shared the 1933 Nobel with Schroedinger, while Anderson received his in 1936).

The place in quantum theory that Dirac carved out for antimatter – a necessary component for the self-consistency of the relativistic quantum theory – remains to this day, although the precise interpretation of the role of antimatter in the relativistic theory (i.e., the precise way in which it solves the negative energy problem) has undergone some evolution. The modern interpretation is due to none other than Feynman, and, the reader may thus infer, is the interpretation most appropriate for its application to quantum *field* theory. This interpretation is an essential step in our discussion, so here is Feynman's take on the issue, which dates from 1948. It's a curious one, but one that has become such an accepted part of the particle physicist's thought process that one tends to forget just how unusual it really is.

The wave function of a free particle that solves the (non-relativistic) Schroedinger equation can be split into two factors – one dictating how the wavefunction varies from point to point in space, and the other dictating how the wave function varies from instant to instant in time. The factor $\psi(t)$ relating to the time (t) dependence is quite simple:

$$\psi(t) = e^{-\alpha E t}$$

where E is as usual the particle's total energy (kinetic energy plus the energy due to the particle's mass), and α is just a fixed number whose value doesn't really concern us.(***11) Technically, the number $e = 2.718281828...$ is the 'base of natural logarithms', which really is of no concern to us. The only thing that we need to know is that the quantity $-\alpha E t$ in the superscript of this equation (the 'power' to which the constant e is raised) involves the product of the particle's energy E and the time t at which you want to know the value of the wavefunction ψ. If the energy E is positive (which it had better be!) the $-\alpha E t$ is negative (less than zero), simply because of the '$-$' sign.

For wavefunctions which solve *relativistic* wave equations, such as Dirac's wave equation for spin-$\frac{1}{2}$ particles, half of the solutions are of this form, while the other half are of the form

$$\psi(t) = e^{+\alpha E T} = e^{-\alpha(-E)t};$$

remember that when you multiply two '$-$' signs together you get a '$+$' sign, so that these two expressions for $\psi(t)$ are the same (the $-$ sign in front of the α multiplied by the $-$ sign in front of the E in the second expression give the $+$ sign in the first expression). So, looking at the rightmost expression in this equation, and comparing it to the equation just above, we see that this half of the solutions have *negative* energy, which as we mentioned above, was the cause of considerable consternation until Dirac was able to identify them (using a trick we won't go into here) as solutions describing the behavior of antiparticles. Feynman, on the other hand, looked at this function and saw something that was as apparently ridiculous as it was obvious: if we move the $-$ sign in front of the E so that it is instead in front of the t, then for these nettlesome solutions we can write

$$\psi(t) = e^{\alpha E t} = e^{-\alpha E(-t)}.$$

Comparing this to our original equation, we see that we are, in some sense, saved: the energy E is again positive, as our common sense requires. The price we have to pay for this repair job, though, seems quite steep: the fact that the wave function now depends on $-t$ rather than t means that this wavefunction describes a particle *moving backwards in time*. For some reason this didn't faze Feynman in the least. As boldly as Dirac predicted antimatter 5 years before it was discovered, Feynman made the statement that, as far as the quantum theory was concerned, antiparticles are simply particles traveling backwards in time.

As radical is this sounds, when you apply it to what we know about quantum field theory, everything falls neatly into place. Consider again the minimal interaction vertex of quantum electrodynamics (the quantum field theory of the electromagnetic force), shown in Figure 4.7a: an electron moving through space to the right emits a photon (eventually to be absorbed by another charged particle we presume), and in doing so recoils to the left. Now what if, at the vertex, the photon connects the incoming positive energy electron solution to an outgoing *positron* (anti-electron) solution rather than an outgoing electron solution. Remember that, as discussed in the previous paragraph, half of the solutions to the Dirac equation are antiparticle solutions, so this is fine as far as the relativistic quantum theory is concerned. As shown in Figure 4.7b, this particle, which we can think of as an electron traveling backwards in time, must recede it back in time from the electron-photon vertex, as indicated by the arrow (now we see why we bothered to carry those heretofore meaningless arrows around in our Feynman diagrams: when you change the direction of the arrow, you switch between descriptions of *particles*, going *forward* in time, and *antiparticles* – particles going *backward* in time).

Now, ask yourself the following question: have you ever observed a particle traveling backwards in time? I'll give you the benefit of the doubt, and assume the answer is no (let us know if otherwise). Thus, what an observer would see is *not* a particle receding backwards in time from the interaction vertex, but rather an antiparticle *positron* proceeding *forwards* in time towards an interaction at the vertex. And what does this positron appear to do in that interaction? Look at Figure 4.7b again. The interaction happens at a well-defined time, which is just the height of the vertex along the vertical (time) axis. Before this time, there are an electron and its antimatter counterpart, a positron, approaching each other. After this time, there is just a photon. The electron and positron have annihilated one another! This diagram represents nothing less than the annihilation of matter and antimatter – the conversion of mass (remember that the electron and positron have identical, non-zero mass) into a state of pure energy. It is the quantum field theoretical rendering of Einstein's most famous assertion: $E = mc^2$. And, last but certainly not least, it represents yet another class of interactions which are succinctly described by the single common denominator of quantum electrodynamics: the vertex of the electron-photon minimal interaction.

In dispensing at last with this lengthy but hopefully entertaining digression, two things need to be mentioned. First, Feynman's receding time interpretation of the 'negative energy' solutions is wonderfully general. Not only did it allow the Dirac equation to become the basis of the description of spin-$\frac{1}{2}$ particles in quantum field theory, it also resurrected other attempts to make quantum theory consistent with special relativity. For instance, Feynman's approach simultaneously solves both the negative energy and negative probability quandaries for the Klein-Gordon equation (see above; recall that these problems originally motivated Dirac's search for another approach and led eventually to the Dirac equation). This is a good thing, because the real world consists of more than just spin-$\frac{1}{2}$ particles. Profound as the Dirac equation may be, in order to develop a full, general relativistic quantum *field* theory of particle interactions, we need relativistic quantum mechanical descriptions of other types of particles. For example, for the electromagnetic interaction, the spin-1 photon must be described if we want to put electrons and photons together into minimal interaction vertices. With Feynman's approach in hand, it was soon recognized that the Klein-Gordon equation provides the basis for the description of spin-0 particles. From the comparison of the (spin-0) Klein-Gordon and (spin-$\frac{1}{2}$) Dirac equations, it was possible to deduce the proper form for the description of spin-1 particles (known as the 'Proca equation'), allowing for the description of photons and other spin-1 particles.

Finally, one can ask whether Feynman's receding-time approach is merely a 'formal' development, which allows for the incorporation of antimatter into the formalism of quantum field theory, or whether instead it is a *physical* discovery, reflecting some deeper truth about the differing relationship of matter and antimatter to the spacetime fabric. It seems that, as of yet, there is no definitive conclusion on this issue; perhaps the answer will suddenly become clear in the deeper context of some future leap forward in our understanding of the fundamental workings of nature.

So much for the digression; it's now time to get back to the final step – the final hurdle that had to be overcome – in the development of the quantum theory of fields.

VIII. RENORMALIZATION AND THE LIVING VACUUM

It's hoped that the discussion of this chapter has left the reader with the impression that the development of relativistic quantum field theory was a great leap forward, which both simplified and expanded our understanding of the fundamental mechanism which underlies the way in which objects in the universe influence one another. However, as alluded to just before our digression into the subject of the relativistic wave equation, the theory as presented so far, as inspired as it may appear, suffers the following shortcoming: it simply doesn't work. The problem lies in a set of interaction processes represented by a progression of Feynman diagrams which we have so far ignored.

Consider again the mutual repulsion of two electrons via the electromagnetic force. In its most basic form, this interaction is represented by the exchange of a single virtual photon, as shown in the spacetime plots of Figures 4.4a and 4.4b, which together are represented by the single Feynman diagram of Figure 4.4c. As mentioned towards the end of section VI, however, this interaction can also proceed via the exchange of any greater number of virtual photons (for example, the process involving the exchange of two virtual photons is shown in Figure 4.5). Recall that every time you add another minimal interaction vertex involving a photon and an electron to the Feynman diagram, the probability that the interaction takes place via the new diagram is about one percent of that represented by the original diagram. Since every time you add another virtual photon you have to add two new vertices to the diagram, the mistake you make by ignoring the possibility that more than one virtual photon can be exchanged is rather small.

However, now that we've understood how it is that antiparticles fit into the picture, we are prepared to discuss an entirely different way in which the basic diagram (Figure 4.4c) describing the repulsion of two electrons can be modified. This modification is represented in Figures 4.8a-c.

In these figures, we see that the understanding of the electron-electron repulsion presented by quantum field theory admits the following possibility: the photon exchanged between the two electrons can, via a standard minimal interaction vertex, instantaneously turn into an electron-positron pair (note that in 4.8a and 4.8b, one of the particles travels forwards in time, and so is an electron, and the other travels backwards in time, and so is a positron), reverting back to a photon again at a second vertex. This second photon is then absorbed by the second electron, completing the transmission of the force. As for Figure 4.4a and 4.4b, in Figure 4.8a the photon is emitted by the electron coming in from the left and absorbed by the electron coming in from the right, while in Figure 4.8b the roles are reversed. Also as for Figure 4.4, the two possibilities 4.8a and 4.8b are essentially the same process, so they are packaged into the single Feynman diagram 4.8c.

Again, we might question whether the process represented by the diagram of Figure 4.8 is consistent with the notion of mass/energy conservation, since the mass of a system consisting of an electron and a positron (both of which have a mass-energy of about 511,000 electron-volts) is decidedly *not* the same as the mass of the photon (zero) which produced them. As usual, though, we are rescued by Heisenberg's uncertainty principle: since the electron-positron pair exists only for a very brief period of time, its mass/energy is uncertain, and so mass/energy can still be conserved, even though the nominal mass of the electron-positron pair is different than the nominal mass of the photon. Just as the short-lived and unobserved photon exchanged between the electrons in Figure 4.4 is known as a 'virtual' photon, electron-positron pairs such as that of Figure 4.8, produced by an instantaneous fluctuation from and back to a photon via two minimal interaction vertices, are known as 'virtual' electron-positron pairs.

Comparing the Feynman diagrams of Figures 4.4c and 4.8c, we see that the latter diagram includes two extra vertices, and so again we would expect the probability of the more complicated, 'higher order' process of Figure 4.8 to be $(1/100)\times(1/100)$, or about one ten-thousandth, less likely than that of Figure 4.4. However, counting vertices is not the only thing that goes into the calculation of the interaction probability of a given Feynman diagram – it's merely a crude rule-of-thumb. There is a big difference between the higher order diagram of Figure 4.5 and that of 4.8c, which leads to a much different result for the interaction probability.

The essence of the difference is this: when the virtual photon fluctuates into the virtual electron-positron pair, its energy must divide itself up, part going to the electron and part to the positron. But there are a very large number – in fact, an *infinite* number – of different ways the energy can divide itself up between the positron and electron. The energy of one of the two can be pretty much anything, as long as the other has the correct amount of energy to compensate and make the total add up to the virtual photon's energy(***12). It turns out that when all of these different possibilities are taken into account, the interaction probability of 4.8c, instead of being ten thousand smaller than that of 4.4c, is *larger* – much larger. In fact, it's *infinite*.

This is clearly an unacceptable state of affairs. When this more complicated diagram is included (which it must be - recall Feynman's statement that 'anything that can happen must happen'), we find that our hallowed quantum field theory can not predict the relative probability of the electrons scattering off of each other into the various directions that you might mount a detector, since the probability of scattering into any given direction is predicted to be infinite. So, after all its great insights and advancements, and a number of predictions strikingly confirmed by experiment, are we to conclude that quantum field theory is dead and useless, having been pecked to death by these ephemeral matter-antimatter fluctuations?

Of course not, but it took some pretty serious soul-searching by Feynman and others to recognize and fix the problem. The approach that was eventually successful in overcoming this hurdle is known as 'renormalization'. We'll need to spend a bit of time here discussing the ideas underlying this fairly technical development, for later, the requirement that any useful model of fundamental interactions be 'renormalizable' will lead us to the particular class of models known as 'gauge theories'.

Feynman and others argued that (as we've seen above) quantum field theory is so profoundly successful on a number of fronts that it must be correct. So, when two electrons repulse each other, then it must be true that the description of this process includes the diagrams of both 4.4c and 4.8c (as well as 4.5, although as mentioned, these don't much matter). The way to incorporate these diagrams without having the theory fall apart, Feynman argued, is to step back for a moment and think very hard about exactly how the theoretical predictions of quantum field theory relate to physically observable quantities that can be measured in the lab.

Think again about the process represented by Figures 4.4 and 4.8 – what one observes in either case is just the scattering of one electron off of another. What goes on between the two electrons during the scatter is, according to the most basic tenets of quantum mechanics, off-limits to the prying eyes of experimentation – the intermediate particles can not be directly observed without disturbing the whole process so profoundly that it can no longer be interpreted as simply one electron scattering off another.

Just to be definite, let's say that the electron coming in from the left in Figure 4.8 is the 'projectile' in the scattering experiment, while the one coming in from the right is the 'target' (feel free to switch projectile and target if you wish – it will make no difference to the argument). Now consider Figure 4.9, which is identical to Figure 4.8c, except that there is now a dashed circle encompassing the target electron's electron-photon vertex, as well as the virtual electron-positron pair. There's nothing physical about the circle – it's just there to evoke an interesting interpretation of the process of Figure 4.8c. Figure 4.8c inclines us to think of the process as one in which the projectile emits a virtual photon which fluctuates momentarily into a virtual electron-positron pair, and then is absorbed (again as a virtual photon) by the target electron. Figure 4.9, however, suggests that we think of the process instead as being the exchange of a *single* virtual photon between a projectile and target electron, but where the target electron is not just the 'bare' electron represented by the single deflecting line, but rather the electron 'dressed up' by everything inside the circle – the bare electron itself plus the fleeting 'vacuum fluctuation' engendered by the virtual electron-positron pair and the other virtual photon. In this picture, it's not that the exchange mediating the repelling interaction between projectile and target can occur in two different ways (Figures 4.4c and 4.8c), but that the target electron itself can have two different forms – the 'bare' form of Figure 4.4c, or the 'dressed' form of Figure 4.9. And – since this is quantum mechanics and you can't ever tell what's going on between the two electrons, who is to say which is the correct point of view? They're both equally valid, and so let's pick the interpretation that serves us best in trying to connect our theoretical picture with the realities of the physical universe. This of course turns out to be the second interpretation – that of Figure 4.9 rather than Figure 4.8c.

We can now introduce the critical recognition behind the procedure of 'renormalization': *whenever* one makes a measurement involving an electron as a target, one is simultaneously measuring *both* of these possible manifestations of the target electron – the bare target electron of Figure 4.4, as well as the target electron 'dressed' by the fluctuation of the virtual photon into an electron-positron pair in the vacuum immediately surrounding the target electron. The bare electron of Figure 4.4 is *not* a physically observable particle; what *is* observable is the electron which is the combination of the bare electron of Figure 4.4 and the dressed electron of Figure 4.9. In fact, it's not just a matter of including the processes of Figures 4.4 and 4.9 – now that we know the rules for cobbling together electron-photon minimal interaction vertices, we can surround the target electron by an arbitrarily complex muck of virtual electron-positron pairs connected together by equally virtual photons. Figure 4.10 shows one of the infinite number of possible diagrams that we could draw. Indeed, the vacuum surrounding the target electron is veritably seething with these vacuum fluctuations, and if you want to talk about the *physical* electron – the one you will sense when you do experiments on electrons – you had better include in your considerations the full cadre of these vacuum fluctuations.

Thus, when you measure anything at all about the electron, such as say the strength of the electron's charge, which you must do by an experiment such as repelling another charged object from it, you are *not* measuring the strength of the charge of the bare electron, but rather are measuring the strength of the charge of the bare electron plus whatever effects are added in by this fool's gallery of virtual vacuum fluctuations. What we can then do is (essentially) to arbitrarily adjust, or *renormalize*, the charge of the bare electron so that, when you add in the effects of the seething 'vacuum' surrounding the electron (represented by diagrams like 4.10), the result gives precisely that appropriate for the *measured* electron charge.

Note that since the charge you measure by scattering off an electron is really that of the bare electron *and* its attendant cloud of vacuum fluctuations, and since you can *never* measure the bare charge directly (the dubious underworld of virtual particles will always be there), then you are free to make the bare charge whatever you darned well need to in order to make the 'effective' charge of the electron plus vacuum entourage agree with observational fact. In fact, the value that you need to chose for the bare electron charge, in order to compensate for the infinite probabilities imposed by the diagram of processes like those of Figure 4.9 and 4.10, is (predictably enough): infinity.

The charge of an electron, the field theorists tell us, is infinite. But they also tell us that no experiment you will ever mount will measure that infinity, so who cares if this seems ridiculous? It gets you back to a theory of quantum fields that makes reasonable predictions for any possible experiment that you could do to measure the electron's charge.

If it is the case that the adjustment of a *finite* number of physically arbitrary parameter(s) (such as the bare electron charge and mass) renders as finite all the possible Feynman diagrams associated with a quantum-field-theoretical description of a given force, then the theory is workable, and is a possible fundamental theory of that force. Such theories are called *'renormalizable'*. If this is not the case, then the theory is classified as 'unrenormalizable', and is *not* a candidate for the fundamental description of the force. As mentioned above, a certain subset of the possible implementations of quantum field theory, known as *'gauge theories'*, seem to have the property that they tend to be renormalizable. For this and other equally fundamental reasons (to be discussed in Chapter 8), gauge theories enjoy a very elevated position in the society of particle physics theories.

As mentioned above, the particular implementation of quantum field theory under scrutiny here is that of 'quantum electrodynamics' – the fundamental theory of electromagnetic interactions. In this implementation, it *is* true that all Feynman diagrams are rendered finite by the adjustment of a few arbitrary physical properties, such as the charge and mass of the bare electron. Quantum electrodynamics *is* renormalizable, and is our current candidate for the relativistic quantum theory of electromagnetic interactions. Fittingly, it turns out that quantum electrodynamics is itself a gauge theory; in fact, the simplest possible gauge theory that can be constructed. This last fact was not appreciated, of course, until the creation of gauge theory in the 1960's – one of the great and fascinating leaps forward in the development of particle physics that still lie ahead in our discussion.

Infinitely charged particles that can travel forward or backward in time, and a 'vacuum' that's as alive and shimmering as the air above a suburban parking lot on a hot summer's day – these are just a few aspects of the profoundly bizarre world-view that quantum field theory would ask us to take to heart. Just how well this unorthodox view of the world seems to work will be discussed shortly.

Before we enter into the discussion of its precise quantitative verification, we'll want to bring forth a few intriguing consequences of renormalization and the 'living vacuum' of quantum field theory. Although not immediately germane to our eventual introduction of the components of the Standard Model, they are nevertheless quite worthy of a few moments' digression.

The first of these is a process predicted directly by quantum field theory and its elevation of the minimal interaction vertex as the fundamental building block of electromagnetic interactions. This process truly flies in the face of everyday common sense, while simultaneously providing hard evidence for the existence of the sea of virtual electron-positron pairs that populate the vacuum.

By arranging the participants in the right way, we can join together two minimal interaction vertices to represent the process of Figure 4.11: a *real* photon, coming in from the left at the speed of light after being emitted from a distant source arbitrarily long ago, suddenly finds itself absorbed by an electron-positron vacuum fluctuation (point 'A' in the figure). Instead of fluctuating back into a photon, though, one of the two charged particles (in this case the electron) finds itself absorbing a second photon which just happens to be around (point 'B' in the figure) resulting after the dust clears in a *real* electron positron pair which go happily off on their own, both with the possibility of being sensed by a well-placed particle detector(***13).

In an experiment at the Stanford Linear Accelerator Center (SLAC) in 1995, directed by Adrian Melissinos of the University of Rochester, the leading edge of a powerful laser beam that had been Compton-scattered to very high energy by the SLAC electron beam (see the discussion of Compton scattering in Section VI above) was directed back against itself. Every once in awhile, one of the high energy scattered photons combined with a photon in the unscattered laser beam to form an electron-positron pair which was clearly identified in a detector surrounding the collision point of the two beams(***14).

This demonstrated experimentally, for the first time, that it is possible to 'conjure matter from light' as one particularly florid popular article on the experiment put it. Indeed, we all know that when two flashlight beams are shined against each other, nothing happens - they just pass right through each other without effect. On the other hand quantum field theory predicts – and experiment confirms – that if one beam is made energetic enough (by, say, bouncing it off of a very energetic electron beam), then every once in a long while two of the photons from the flashlight beams *will* interact, with the very concrete result that two matter particles (one electron and one positron) are produced.

In addition, although not for the first time, but perhaps more poignantly than ever before, this experiment showed that the living vacuum – this seething maggot's nest of virtual particle-antiparticle pairs – is very much alive and real. For there, making their presence known via highly characteristic flashes of light in the detector of the Melissinos experiment, were just those pairs of electrons and positrons, jarred into a concrete existence by photons from the two colliding light beams of SLAC Experiment E144.

But if that's not enough to convince you of the existence of the living vacuum, consider a prediction, dating from 1948, of the (very recently deceased) Dutch physicist Hendrick Casimir. Casimir recognized that, if all space were in fact seething with virtual electron positron pairs, then there ought to be a net interaction between the virtual particles in this not-so-vacuous vacuum and the free electric charges in an electrically neutral conductor(***15). Casimir thus predicted that two parallel conducting plates, electrically neutral and thus uninterested in each other according to standard electromagnetic theory, should be attracted to each other by the mutual effect of the electron-positron fluctuations in the interceeding vacuum. This 'Casimir effect' was finally measured in 1996 by Steven Lamoreaux of the University of Washington, confirming the strength of the attraction predicted by Casimir to within the 5% accuracy of the experiment, and providing yet another striking confirmation of the astoundingly unusual world of quantum field theory.

The final issue to be discussed before closing out this section and moving on to the issue of the precise quantitative confirmation of field theory is, if anything, even more interesting than the discussion of the Melissinos experiment or the Casimir effect. As mentioned some time back, when looking at something under a microscope, one can not resolve (bring into focus) features of the observed sample that are much smaller than the wavelength of the light which illuminates the sample (electron microscopy is so much more powerful because the wavelength of an electron is much shorter than that of a photon of the same energy). Recall that de Broglie's relation $\lambda = h/p$ tells us that wavelength *decreases* with energy. So, if we bounce particles of higher and and higher energy off of the sample, with correspondingly shorter and shorter wavelengths, we'll see smaller and smaller features of the sample.

Now consider the scattering of an electron off of another electron via electromagnetic repulsion, as we've discussed at length over the previous pages. As before, consider one electron to be the probe, like the electrons in an electron microscope, while the other electron is the target, like the sample in the electron microscope. As the momentum of the probe electron increases, its wavelength decreases accordingly. Thus, as the momentum of the probe increases, its ability to see the target electron for what it *really* is – a bare electron stripped of its cloud of virtual hangers-on – becomes better and better (but never so good, of course, that the target electron charge becomes infinite). Thus, if whole idea of renormalization is correct, as we go to higher and higher probe momentum, the effects of the virtual cloud should diminish, *and we should observe that the value of the electron charge increases*, i.e., becomes closer to its true, bare value (of infinity). This happens very slowly, but with modern particle accelerators, which achieve electron energies equivalent to the acceleration of the electron through 100,000,000,000 Volts (recall that an electron beam in a TV set has an energy of about 1,000 electron-Volts), the effect should be observable. In fact, current measurements show that, at these high energies, the electron's charge is about 2% greater than it is when measured via a low-energy table-top experiment – in perfect accord with the predictions of quantum electrodynamics.

This property – the dependence of the charge upon the momentum, or 'scale', of the interacting particles – is known as 'running', and *all* charges 'run'. In fact, if you look at the momentum dependence of the electromagnetic, weak, and strong force charges of the fundamental particles (exactly what the 'fundamental particles' are will be discussed in the next chapter), it is seen that they all appear to be 'running' towards the roughly the same value. We measure the variation of the charge from everyday scales out to interaction energies of about 100 billion electron-Volts (100 GeV); if we take the observed values and momentum dependence of the charges and use this information to extrapolate to higher momentum, the values all seem to coalesce at an interaction energy of about 10^{16} electron-volts(***16). (The measurement of the running of gravitational charge is, at this point, beyond the abilities of our experimental capabilities).

To physicists, this strongly suggests that the forces of nature are all in fact different facets of the same underlying 'grand unified' interaction, the single natural phenomenon that governs the way objects in the universe influence each other, and thus is fully responsible for absolutely every phenomenon that this exceedingly vast and rich universe puts forth – from the vast scale of the creation and evolution of the cosmos down to the most subtle intonations of human life. Although substantial strides have been made towards the formulation of this ultimate framework (the discussion of which lies ahead), it still remains an unattained goal. Particle physicists believe, however, that the next generation of experiments, due to be completed during the period between roughly 2005-2020, have the possibility to substantially further us towards that goal. Indeed, the discussion of the design of these near-future experiments, some of which are only in their earliest planning stages, is driven in large measure by this goal.

IX. THE GYROMAGNETIC RATIO OF THE ELECTRON

As we've seen above, our archetypical quantum field theory is quantum electrodynamics, the quantum field theory of the electromagnetic interaction. It is to this theory which we turn for the most exacting test of the tenets of quantum field theory.

Throughout this chapter, we have made incessant reference to the electrical repulsion of like-charged objects, in particular electrons. We have barely mentioned magnetism, but certainly our theory of electromagnetism should be expected to provide as thorough a description of magnetism as it does of the electrical force.

It is a property of electric charge (the sole type of charge associated with the full phenomenon of electromagnetism) that when it is in motion, it creates *magnetic* fields. In particle, when a charged objects orbits in a circle about some point in space, it 'generates a magnetic field', which is a fancy way to say that it acts like a magnet. More central to our discussion, though, is that, rather than orbiting, if a charged object spins about its axis, it also generates a magnetic field. Recall that electrons 'spin', i.e., have angular momentum (of magnitude $\frac{1}{2}\hbar$) associated with their continual and unyielding rotation about their axes, and thus each electron acts, among other things, as a tiny magnet. Everyday household magnets, such as the one that is holding your unpaid cable television bill to your refrigerator at this very moment, are magnetic by virtue of the collective magnetic effect of all the spinning electrons (the process of the 'magnetization' of a piece of iron is nothing more than getting some of the electrons' axes of rotation lined up and pointing in the same direction).

The ratio between the angular momentum ($\frac{1}{2}\hbar$ for the electron) associated with a charged particle's spin and the strength of the correspondingly generated magnetic field is known as the particle's 'gyromagnetic ratio', which is usually represented by the Greek letter μ. Recall that, in our digression above on the development of the relativistic wave equation(s) (see section VII), the Dirac equation is understood to provide the appropriate description for spin-$\frac{1}{2}$ particles, such as electrons. One of the nice things that follows from this understanding is a *prediction* for the value of these particles' gyromagnetic ratio. That prediction is:

$$\mu_{\text{pred}} = \frac{2q}{mc}$$

where q is the particle's electromagnetic charge, m is its mass, and c the speed of light. Since all of these quantities are known very precisely for the electron, this expression forms a very exacting prediction for the electron's gyromagnetic ratio μ_e.

Now, it's possible to directly measure the strength of the electron's magnetic field and thus its gyromagnetic ratio. If you place a magnet (such as a spinning electron) close to another magnet (such as the very precisely calibrated magnetic coil in the experiment we're about to discuss), the magnetic part of the electromagnetic force will try to align the first magnet in the direction of the field of the second magnet, just as a compass needle (which is just a little magnet) tries to align itself with the earth's magnetic field. However, if the first magnet (the electron) is placed in the field of the second magnet so it is not quite aligned, and then given a little push, it will begin to rotate ('precess') about the direction of the second magnet's magnetic field. It's exactly like a gyroscope, if you've ever had the pleasure of playing with one of those. If you get a gyroscope spinning and set it down on the table perfectly vertically (aligned with the earth *gravitational* field, which points up and down of course), it just stands there. However, if you set it down so it's a bit tilted, and give it a little push, it slow rotates (precesses) about the vertical.

For the case of the two magnets (the electron and the experiment's magnetic coil), the frequency of precession (number of seconds per rotation) depends in a very well known way on the strength of the two magnetic fields. So, if you measure this frequency and if you know the strength of one of the magnetic fields (the coil's), then you can calculate the strength of the other magnetic field (the electron's). This is a beautiful technique, since it's relatively easy to measure a frequency very accurately. You just need to trap the electron in the region of the coil's field and watch it go around and around for as long as you can bear. You just divide the number of rotations by the time elapsed, and that's your measured frequency. The longer you watch it, the more accurate your measurement is.

These measurements confirmed, more or less, Dirac's prediction of the size of the electron's gyromagnetic ratio. However, modern experimental technique allows for the trapping of electrons for a very long time, so modern experiments are *very, very* precise. Within all of this precision, a small deviation from Dirac's prediction emerged. If we let g_{meas} represent the ratio of the observed and predicted gyromagnetic ratio of the electron, i.e.,

$$g_{\mathrm{meas}} = \frac{\mu_{\mathrm{obs}}}{\mu_{\mathrm{pred}}},$$

then the result of the most accurate modern experiment is

$$g_{\text{meas}} = 1.00115965219$$

to within an uncertainty of about 1 in the very last digit.

The fact that g is not precisely 1 means that Dirac's prediction is not quite right. It's not wrong by much – just a little over one tenth of a percent. But the experiments are very accurate, and so that small discrepancy is very significant, being about *eight orders of magnitude* (factors of 10) larger than the uncertainty on its measurement.

But – and this is the critical point – Dirac's relativistic wave equation is *not* our quantum theory of the electromagnetic interaction, it is merely one (albeit very important) ingredient in that theory; specifically, the ingredient that allows us to determine the wave functions of the spin-1 $over 2$ particles (electrons) that enter into the Feynman diagrams of the full theory. So quantum field theory tells us that Dirac's prediction for the electron's gyromagnetic ratio is just a start, and to *really* predict it, we need to use the full theory and really do our homework. We need to calculate all the Feynman diagrams associated with the interaction of the quantum of the magnetic field (which is of course just the photon again) with the electron.

Of course, as we've seen, when we include the quagmire of virtual particles continually crawling in and out of the vacuum, there are many diagrams we need to calculate. However, as usual, the relevance of each possible diagram decreases with increasing number of vertices (as long as we used the appropriately *renormalized* electron properties, of course), so we can ignore the really complicated ones. However, the experimental result from above is so accurate that we need to include some of these more complicated diagrams if the accuracy of our theoretical prediction is to match that of the experiment – it turns out that we need to include all Feynman diagrams with 7 or fewer minimal interaction vertices. These diagrams are shown in Figure 4.12. (XXX - will need to get reproduction permission here).

The calculation of these diagrams took several *years*, but in the end, it was seen that Dirac's original prediction did indeed need to be modified. The value g_{pred} of the predicted modification was

$$g_{\text{pred}} = 1.0011596522.$$

Comparing this predicted modification to the measured difference

$$g_{\text{meas}} = 1.00115965219$$

from above shows an agreement to mind-boggling precision. *Quantum field theory predicts the value of the electron's gyromagnetic ratio to better than one part in* 10^{10}, or about three parts in one-hundred billion!! This certainly establishes quantum field theory in general (and quantum electrodynamics in particular) as the most quantitatively successful theoretical framework known to man. Quantum field theory – this collection of innovative, counterintuitive, irreverent, and sometimes almost perverse, but inarguably profound ideas – has received the loftiest possible imprimatur that could ever be afforded to a physical theory – precise, quantitative confirmation to a fraction of a fraction of a fraction (if that) of a gnat's ass. It *must* be right.

****

(***1) The modern notion of the 'force field' was first introduced by the British Physicist Michael Faraday in the 1830's in order to interpret his revolutionary experiments on the interrelation between the electric and magnetic forces, which laid the foundation on which much of Maxwell's later work synthesizing the two forces into a single unified electromagnetism was based. The story of Faraday's rise from humble beginnings as a laboratory apprentice in London's Royal Institution to the receipt of persistent offers to assume its leadership (all politely refused) and recognition as one of the great experimental physicists of his day is a particularly inspiring one.

(***2) Technically, $k \simeq 8.99 \times 10^9$ Newton-meter^2 per Coulomb. Later, when we discuss the *gauge theory* of physical forces, there will be only two things, one of which is the value of this overall strength or 'coupling constant', which need to be determined by experiment. The other will be the mathematical 'symmetry group' underlying the force in question.

(***3) The relation Maxwell derived is simple; if we denote by c the speed with which the disturbance travels though space, then

$$c = \sqrt{\frac{4\pi k}{\mu_0}}$$

where $\pi \simeq 3.14159$ is as usual the ratio between the circumference and diameter of a perfect circle, and $\mu_0 = 4\pi \times 10^{-7}$ Tesla-meters per Amp is the experimentally measured strength of the magnetic force.

(***4) We also use this term, tongue-in-cheek, to describe colleagues who continually put forth conjectures that don't make any sense to us.

(***5) In the popular culture of the United States, Feynman was relatively well known for his collections of humorous anecdotes (both scientific and non-scientific), as well as his brinkman-like demonstration in a congressional hearing that the thermal properties of the material used to make the O-rings was culpable in the mid-flight explosion of the Challenger space shuttle – a hypothesis later adopted as the most likely explanation for the disaster. Within the culture of professional physicists, Feynman is considered to be one of the most pronounced and free-thinking physicists since the great icons of the early twentieth century.

(***6) Technically, there is one more thing you could ask for: the quantum mechanical *phase* (see Chapter 3) of the scattered wavefunction. This is necessary to know if you have to combine the results of the calculated process (say, single photon exchange) with the result of the calculation of another process (say, double photon exchange) which has the same input and output particles, so that you can't tell from any given trial which of the two processes was responsible for the scatter. This is really a technicality.

(***7) This coupling goes by another name – the 'fine structure constant'. In the early days of atomic physics (before quantum field theory revealed the true meaning of the fine structure constant to be the strength of the coupling between the electron and photon), it was thought to have a value so close to being precisely 1/137 that numerologists started to establish cultish associations with the number 137. However, quantum field theory tells us that the fine structure constant actually depends upon the energy of the virtual photon, and is only very close to 1/137 at 'low energy' – i.e., the energies available to experimenters in those early days of atomic physics. So, there's nothing magical about 1/137. Quantum field theory predicts, and experiments confirm, that at virtual photon energies equivalent to about 100 proton masses (remember that mass and energy are related via $E = mc^2$), the value of the fine structure constant is about 1/128.

(***8) The board walk in Santa Cruz, California, has a particularly nice old-fashioned carrousel which turns to the music of 1900's vintage mechanical calliope. It's well worth the visit if you ever find yourself in town.

(***9) The English Paul A. M. Dirac and the Austrian Wolfgang Pauli are perhaps the most prominent of this crowd, both of whom were in their late 20's when this work was done. Soon after this work, Dirac was awarded the Lucasian Chair of Mathematics at at Cambridge, perhaps the most prestigious academic position on the planet, having seated such greats as Isaac Newton and the renowned astrophysicist Stephen Hawking.

(***10) Don't be discouraged if you've been told that the square of any number is a positive number – this is true, but only for *real* numbers. In general, solutions to any quantum mechanical wave function, relativistic or not, are *complex*. This means that the solutions contain factors of i, the number which when squared (multiplied by itself) somehow gives -1.

(***11) In fact, $\alpha = i/\hbar$ where $i = \sqrt{-1}$. Raising e to an imaginary number (i.e., one containing a factor of i) produces a function which oscillates up and down, like a wave, only in time rather than in space. So, such functions are very appropriate for wavefunctions.

(***12) The extremely wary reader may notice an apparent logical inconsistency here. In this discussion, we are requiring energy to be conserved absolutely, in that the sum of the energies of the electron and positron add up precisely to the energy of the virtual photon from which they derive. In the paragraph just above, however, we relied on the uncertainty principle in a way that may appear to release us from the requirements of energy conservation, so that the fact that real (non-virtual) photons have zero mass does not restrict the virtual photon from temporarily fluctuating into an electron-positron pair. These two statements are not inconsistent – the reconciliation of these two apparently incompatible statements lies in the recognition that, uncertainty principle or not, we *always* require energy to be conserved – it's just that, in doing so, the mass of a short-lived state (photon, electron-positron pair, etc.) can temporarily be different from what it would be for the corresponding state with indefinite lifetime whose mass you would measure in the lab.

(***13) Note that a photon in isolation can't 'fall apart' by getting absorbed by an electron-positron vacuum fluctuation, resulting in the spontaneous conversion of photon into an electron-positron pair. Since both states (the initial photon and the final electron-positron) live arbitrarily long, there is *no* uncertainty on either's mass/energy, and so the differing mass of the initial isolated photon and final electron- positron pair will prohibit the reaction from happening. The photon which is absorbed by an electron-positron vacuum fluctuation will always quickly fluctuate back into an isolated photon. The second photon of Figure 4.11 is absolutely essential if the electron and positron are to be made to stick around in the final state.

(***14) Note that a laser beam, while very intense, is composed of photons which are themselves rather ordinary. Recall that the energy of any photon, whether in a laser beam or coming from a child's nightlite, is just given by Planck's constant time the color (or frequency f) of the photon: $E = hf$.

So, the individual photons in a laser, no matter how intense the laser may be, are no more energetic than the photons in the nightlite. What makes lasers so interesting and powerful is that there are an unspeakably large *number* of photons concentrated in a small region of space and time, and, what's more, all with synchronized phase ('coherent'). On the other hand, once the laser beam is bounced off a high energy electron beam, you have a situation in which *both* the density *and* the energy of the individual photons is tremendous, and interesting things can happen (such as the creation of matter from light). As a final note, it should be pointed out that, in the Melissinos experiment, even after bouncing off the SLAC electron beam, the high energy photons still required simultaneous interaction with four or more ordinary energy photons from the beam in order to produce an electron-positron pair. So, Figure 4.11 is a slight oversimplification of the pair production process in the Melissinos experiment, but not one that changes the qualitative conclusions we are drawing from the experimental result.

(***15) An electricaly neutral, or uncharged, conductor (such as a chunk of copper) is just a bunch of neutral atoms connected together in some sort of regular pattern. What makes the conductor a conductor is that the outermost (most weakly bound) electrons in the atoms become free to wander at will about the conductor. Thus, the conductor contains a large number of free electrons (of order one per atom) even though the conductor as a whole is eletrically neutral. These, of course, are the electrons responsible for the 'conduction' of electricity when the conductor is hammered into the shape of a wire and connected between the two terminals of a battery.

(***16) In fact, for any given particle, the charge strengths associated with the three different forces (electromagnetic, weak, and strong) do not *quite* coalesce at the same energy unless additional physical laws, beyond those of the Standard Model of particle physics, are introduced. A set of physical laws which seem to quite naturally lead to the fine adjustments necessary to make these three separate extrapolations precisely coalesce are those provided by 'supersymmetry', a discussion of which can be found in Brian Greene's book *The Elegant Universe*. This is an intriguing hint that supersymmetry may in fact be the next great leap forward in our fundamental understanding of the universe.