

Probability Amplitudes

3-1 The laws for combining amplitudes

When Schrödinger first discovered the correct laws of quantum mechanics, he wrote an equation which described the amplitude to find a particle in various places. This equation was very similar to the equations that were already known to classical physicists—equations that they had used in describing the motion of air in a sound wave, the transmission of light, and so on. So most of the time at the beginning of quantum mechanics was spent in solving this equation. But at the same time an understanding was being developed, particularly by Born and Dirac, of the basically new physical ideas behind quantum mechanics. As quantum mechanics developed further, it turned out that there were a large number of things which were not directly encompassed in the Schrödinger equation—such as the spin of the electron, and various relativistic phenomena. Traditionally, all courses in quantum mechanics have begun in the same way, retracing the path followed in the historical development of the subject. One first learns a great deal about classical mechanics so that he will be able to understand how to solve the Schrödinger equation. Then he spends a long time working out various solutions. Only after a detailed study of this equation does he get to the “advanced” subject of the electron’s spin.

We had also originally considered that the right way to conclude these lectures on physics was to show how to solve the equations of classical physics in complicated situations—such as the description of sound waves in enclosed regions, modes of electromagnetic radiation in cylindrical cavities, and so on. That was the original plan for this course. However, we have decided to abandon that plan and to give instead an introduction to the quantum mechanics. We have come to the conclusion that what are usually called the advanced parts of quantum mechanics are, in fact, quite simple. The mathematics that is involved is particularly simple, involving simple algebraic operations and no differential equations or at most only very simple ones. The only problem is that we must jump the gap of no longer being able to describe the behavior *in detail* of particles in space. So this is what we are going to try to do: to tell you about what conventionally would be called the “advanced” parts of quantum mechanics. But they are, we assure you, by all odds the simplest parts—in a deep sense of the word—as well as the most basic parts. This is frankly a pedagogical experiment; it has never been done before, as far as we know.

In this subject we have, of course, the difficulty that the quantum mechanical behavior of things is quite strange. Nobody has an everyday experience to lean on to get a rough, intuitive idea of what will happen. So there are two ways of presenting the subject: We could either describe what can happen in a rather rough physical way, telling you more or less what happens without giving the precise laws of everything; or we could, on the other hand, give the precise laws in their abstract form. But, then because of the abstractions, you wouldn’t know what they were all about, physically. The latter method is unsatisfactory because it is completely abstract, and the first way leaves an uncomfortable feeling because one doesn’t know exactly what is true and what is false. We are not sure how to overcome this difficulty. You will notice, in fact, that Chapters 1 and 2 showed this problem. The first chapter was relatively precise; but the second chapter was a rough description of the characteristics of different phenomena. Here, we will try to find a happy medium between the two extremes.

3-1 The laws for combining amplitudes

3-2 The two-slit interference pattern

3-3 Scattering from a crystal

3-4 Identical particles




Fig 3-1. Interference experiment with electrons.

We will begin in this chapter by dealing with some general quantum mechanical ideas. Some of the statements will be quite precise, others only partially precise. It will be hard to tell you as we go along which is which, but by the time you have finished the rest of the book, you will understand in looking back which parts hold up and which parts were only explained roughly. The chapters which follow this one will not be so imprecise. In fact, one of the reasons we have tried carefully to be precise in the succeeding chapters is so that we can show you one of the most beautiful things about quantum mechanics—how much can be deduced from so little.

We begin by discussing again the superposition of *probability amplitudes*. As an example we will refer to the experiment described in Chapter 1, and shown again here in Fig. 3-1. There is a source s of particles, say electrons; then there is a wall with two slits in it; after the wall, there is a detector located at some position x . We ask for the probability that a particle will be found at x . Our *first general principle* in quantum mechanics is that the *probability* that a particle will arrive at x , when let out at the source s , can be represented quantitatively by the absolute square of a complex number called a *probability amplitude*—in this case, the “amplitude that a particle from s will arrive at x .” We will use such amplitudes so frequently that we will use a shorthand notation—invented by Dirac and generally used in quantum mechanics—to represent this idea. We write the probability amplitude this way:

$$\langle \text{Particle arrives at } x \mid \text{particle leaves } s \rangle. \quad (3.1)$$

In other words, the two brackets $\langle \rangle$ are a sign equivalent to “the amplitude that”; the expression at the *right* of the vertical line always gives the *starting* condition, and the one at the *left*, the *final* condition. Sometimes it will also be convenient to abbreviate still more and describe the initial and final conditions by single letters. For example, we may on occasion write the amplitude (3.1) as

$$\langle x \mid s \rangle. \quad (3.2)$$

We want to emphasize that such an amplitude is, of course, just a single number—a *complex* number.

We have already seen in the discussion of Chapter 1 that when there are two ways for the particle to reach the detector, the resulting probability is not the sum of the two probabilities, but must be written as the absolute square of the sum of two amplitudes. We had that the probability that an electron arrives at the detector when both paths are open is

$$P_{12} = |\phi_1 + \phi_2|^2. \quad (3.3)$$




Fig. 3-2. A more complicated interference experiment.

We wish now to put this result in terms of our new notation. First, however, we want to state our *second general principle* of quantum mechanics: When a particle can reach a given state by two possible routes, the total amplitude for the process is the *sum of the amplitudes* for the two routes considered separately. In our new notation we write that

$$\langle x | s \rangle_{\text{both holes open}} = \langle x | s \rangle_{\text{through 1}} + \langle x | s \rangle_{\text{through 2}}. \quad (3.4)$$

Incidentally, we are going to suppose that the holes 1 and 2 are small enough that when we say an electron goes through the hole, we don't have to discuss which part of the hole. We could, of course, split each hole into pieces with a certain amplitude that the electron goes to the top of the hole and the bottom of the hole and so on. We will suppose that the hole is small enough so that we don't have to worry about this detail. That is part of the roughness involved; the matter can be made more precise, but we don't want to do so at this stage.

Now we want to write out in more detail what we can say about the amplitude for the process in which the electron reaches the detector at x by way of hole 1. We can do that by using our *third general principle*: When a particle goes by some particular route the amplitude for that route can be written as the *product* of the *amplitude* to go part way with the *amplitude* to go the rest of the way. For the setup of Fig. 3-1 the amplitude to go from s to x by way of hole 1 is equal to the amplitude to go from s to 1, multiplied by the amplitude to go from 1 to x .

$$\langle x | s \rangle_{\text{via 1}} = \langle x | 1 \rangle \langle 1 | s \rangle. \quad (3.5)$$

Again this result is not completely precise. We should also include a factor for the amplitude that the electron will get through the hole at 1; but in the present case it is a simple hole, and we will take this factor to be unity.

You will note that Eq. (3.5) appears to be written in reverse order. It is to be read from right to left: The electron goes from s to 1 and then from 1 to x . In summary, if events occur in succession—that is, if you can analyze one of the routes of the particle by saying it does this, then it does this, then it does that—the resultant amplitude for that route is calculated by multiplying in succession the amplitude for each of the successive events. Using this law we can rewrite Eq. (3.4) as

$$\langle x | s \rangle_{\text{both}} = \langle x | 1 \rangle \langle 1 | s \rangle + \langle x | 2 \rangle \langle 2 | s \rangle.$$

Now we wish to show that just using these principles we can calculate a much more complicated problem like the one shown in Fig. 3-2. Here we have two walls, one with two holes, 1 and 2, and another which has three holes, a , b , and c . Behind the second wall there is a detector at x , and we want to know the amplitude for a particle to arrive there. Well, one way you can find this is by calculating the superposition, or interference, of the waves that go through; but you can also do it by saying that there are six possible routes and superposing an amplitude for each. The electron can go through hole 1, then through hole a , and then to x ; or it could go through hole 1, then through hole b , and then to x ; and so on. According to our second principle, the amplitudes for alternative routes add, so we should

be able to write the amplitude from s to x as a sum of six separate amplitudes. On the other hand, using the third principle, each of these separate amplitudes can be written as a product of three amplitudes. For example, one of them is the amplitude for s to 1, times the amplitude for 1 to a , times the amplitude for a to x . Using our shorthand notation, we can write the complete amplitude to go from s to x as

$$\langle x | s \rangle = \langle x | a \rangle \langle a | 1 \rangle \langle 1 | s \rangle + \langle x | b \rangle \langle b | 1 \rangle \langle 1 | s \rangle + \cdots + \langle x | c \rangle \langle c | 2 \rangle \langle 2 | s \rangle.$$

We can save writing by using the summation notation

$$\langle x | s \rangle = \sum_{\substack{i=1,2 \\ \alpha=a,b,c}} \langle x | \alpha \rangle \langle \alpha | i \rangle \langle i | s \rangle. \quad (3.6)$$

In order to make any calculations using these methods, it is, naturally, necessary to know the amplitude to get from one place to another. We will give a rough idea of a typical amplitude. It leaves out certain things like the polarization of light or the spin of the electron, but aside from such features it is quite accurate. We give it so that you can solve problems involving various combinations of slits. Suppose a particle with a definite energy is going in empty space from a location \mathbf{r}_1 to a location \mathbf{r}_2 . In other words, it is a free particle with no forces on it. Except for a numerical factor in front, the amplitude to go from \mathbf{r}_1 to \mathbf{r}_2 is

$$\langle \mathbf{r}_2 | \mathbf{r}_1 \rangle = \frac{e^{i\mathbf{p} \cdot \mathbf{r}_{12} / \hbar}}{r_{12}}, \quad (3.7)$$

where $\mathbf{r}_{12} = \mathbf{r}_2 - \mathbf{r}_1$, and \mathbf{p} is the momentum which is related to the energy E by the relativistic equation

$$p^2 c^2 = E^2 - (m_0 c^2)^2,$$

or the nonrelativistic equation

$$\frac{p^2}{2m} = \text{Kinetic energy}.$$

Equation (3.7) says in effect that the particle has wavelike properties, the amplitude propagating as a wave with a wave number equal to the momentum divided by \hbar .

In the most general case, the amplitude and the corresponding probability will also involve the time. For most of these initial discussions we will suppose that the source always emits the particles with a given energy so we will not need to worry about the time. But we could, in the general case, be interested in some other questions. Suppose that a particle is liberated at a certain place P at a certain time, and you would like to know the amplitude for it to arrive at some location, say \mathbf{r} , at some later time. This could be represented symbolically as the amplitude $\langle \mathbf{r}, t = t_1 | P, t = 0 \rangle$. Clearly, this will depend upon both \mathbf{r} and t . You will get different results if you put the detector in different places and measure at different times. This function of \mathbf{r} and t , in general, satisfies a differential equation which is a wave equation. For example, in a nonrelativistic case it is the Schrödinger equation. One has then a wave equation analogous to the equation for electromagnetic waves or waves of sound in a gas. However, it must be emphasized that the wave function that satisfies the equation is not like a real wave in space; one cannot picture any kind of reality to this wave as one does for a sound wave.

Although one may be tempted to think in terms of "particle waves" when dealing with one particle, it is not a good idea, for if there are, say, two particles, the amplitude to find one at \mathbf{r}_1 and the other at \mathbf{r}_2 is not a simple wave in three-dimensional space, but depends on the six space variables \mathbf{r}_1 and \mathbf{r}_2 . If we are, for example, dealing with two (or more) particles, we will need the following additional principle: Provided that the two particles do not interact, the amplitude that one particle will do one thing *and* the other one something else is the product of the two amplitudes that the two particles would do the two things separately. For example, if $\langle a | s_1 \rangle$ is the amplitude for particle 1 to go from s_1 to a , and $\langle b | s_2 \rangle$

is the amplitude for particle 2 to go from s_2 to b , the amplitude that *both* things will happen together is

$$\langle a | s_1 \rangle \langle b | s_2 \rangle.$$

There is one more point to emphasize. Suppose that we didn't know where the particles in Fig. 3-2 come from before arriving at holes 1 and 2 of the first wall. We can still make a prediction of what will happen beyond the wall (for example, the amplitude to arrive at x) provided that we are given two numbers: the amplitude to have arrived at 1 and the amplitude to have arrived at 2. In other words, because of the fact that the amplitude for successive events multiplies, as shown in Eq. (3.6), all you need to know to continue the analysis is two numbers—in this particular case $\langle 1 | s \rangle$ and $\langle 2 | s \rangle$. These two complex numbers are enough to predict all the future. That is what really makes quantum mechanics easy. It turns out that in later chapters we are going to do just such a thing when we specify a starting condition in terms of two (or a few) numbers. Of course, these numbers depend upon where the source is located and possibly other details about the apparatus, but given the two numbers, we do not need to know any more about such details.




Fig. 3-3. An experiment to determine which hole the electron goes through.

3-2 The two-slit interference pattern

Now we would like to consider a matter which was discussed in some detail in Chapter 1. This time we will do it with the full glory of the amplitude idea to show you how it works out. We take the same experiment shown in Fig. 3-1, but now with the addition of a light source behind the two holes, as shown in Fig. 3-3. In Chapter 1, we discovered the following interesting result. If we looked behind slit 1 and saw a photon scattered from there, then the distribution obtained for the electrons at x in coincidence with these photons was the same as though slit 2 were closed. The total distribution for electrons that had been “seen” at either slit 1 or slit 2 was the sum of the separate distributions and was completely different from the distribution with the light turned off. This was true at least if we used light of short enough wavelength. If the wavelength was made longer so we could not be sure at which hole the scattering had occurred, the distribution became more like the one with the light turned off.

Let's examine what is happening by using our new notation and the principles of combining amplitudes. To simplify the writing, we can again let ϕ_1 stand for the amplitude that the electron will arrive at x by way of hole 1, that is,

$$\phi_1 = \langle x | 1 \rangle \langle 1 | s \rangle.$$

Similarly, we'll let ϕ_2 stand for the amplitude that the electron gets to the detector by way of hole 2:

$$\phi_2 = \langle x | 2 \rangle \langle 2 | s \rangle.$$

These are the amplitudes to go through the two holes and arrive at x if there is no light. Now if there is light, we ask ourselves the question: What is the amplitude for the process in which the electron starts at s and a photon is liberated by the

light source L , ending with the electron at x and a photon seen behind slit 1? Suppose that we observe the photon behind slit 1 by means of a detector D_1 , as shown in Fig. 3-3, and use a similar detector D_2 to count photons scattered behind hole 2. There will be an amplitude for a photon to arrive at D_1 and an electron at x , and also an amplitude for a photon to arrive at D_2 and an electron at x . Let's try to calculate them.

Although we don't have the correct mathematical formula for all the factors that go into this calculation, you will see the spirit of it in the following discussion. First, there is the amplitude $\langle 1 | s \rangle$ that an electron goes from the source to hole 1. Then we can suppose that there is a certain amplitude that while the electron is at hole 1 it scatters a photon into the detector D_1 . Let us represent this amplitude by a . Then there is the amplitude $\langle x | 1 \rangle$ that the electron goes from slit 1 to the electron detector at x . The amplitude that the electron goes from s to x via slit 1 and scatters a photon into D_1 is then

$$\langle x | 1 \rangle a \langle 1 | s \rangle.$$

Or, in our previous notation, it is just $a\phi_1$.

There is also some amplitude that an electron going through slit 2 will scatter a photon into counter D_1 . You say, "That's impossible; how can it scatter into counter D_1 if it is only looking at hole 1?" If the wavelength is long enough, there are diffraction effects, and it is certainly possible. If the apparatus is built well and if we use photons of short wavelength, then the amplitude that a photon will be scattered into detector 1, from an electron at 2 is very small. But to keep the discussion general we want to take into account that there is always some such amplitude, which we will call b . Then the amplitude that an electron goes via slit 2 and scatters a photon into D_1 is

$$\langle x | 2 \rangle b \langle 2 | s \rangle = b\phi_2.$$

The amplitude to find the electron at x and the photon in D_1 is the sum of two terms, one for each possible path for the electron. Each term is in turn made up of two factors: first, that the electron went through a hole, and second, that the photon is scattered by such an electron into detector 1; we have

$$\langle \text{electron at } x \mid \text{photon at } D_1 \mid \text{electron from } s \rangle = a\phi_1 + b\phi_2. \quad (3.8)$$

We can get a similar expression when the photon is found in the other detector D_2 . If we assume for simplicity that the system is symmetrical, then a is also the amplitude for a photon in D_2 when an electron passes through hole 2, and b is the amplitude for a photon in D_2 when the electron passes through hole 1. The corresponding total amplitude for a photon at D_2 and an electron at x is

$$\langle \text{electron at } x \mid \text{photon at } D_2 \mid \text{electron from } s \rangle = a\phi_2 + b\phi_1. \quad (3.9)$$

Now we are finished. We can easily calculate the probability for various situations. Suppose that we want to know with what probability we get a count in D_1 and an electron at x . That will be the absolute square of the amplitude given in Eq. (3.8), namely, just $|a\phi_1 + b\phi_2|^2$. Let's look more carefully at this expression. First of all, if b is zero—which is the way we would like to design the apparatus—then the answer is simply $|\phi_1|^2$ diminished in total amplitude by the factor $|a|^2$. This is the probability distribution that you would get if there were only one hole—as shown in the graph of Fig. 3-4(a). On the other hand, if the wavelength is very long, the scattering behind hole 2 into D_1 may be just about the same as for hole 1. Although there may be some phases involved in a and b , we can ask about a simple case in which the two phases are equal. If a is practically equal to b , then the total probability becomes $|\phi_1 + \phi_2|^2$ multiplied by $|a|^2$, since the common factor a can be taken out. This, however, is just the probability




Fig. 3-4. The probability of counting an electron at x in coincidence with a photon at D in the experiment of Fig. 3-3: (a) for $b = 0$; (b) for $b = a$; (c) for $0 < b < a$.

light source L , ending with the electron at x and a photon seen behind slit 1? Suppose that we observe the photon behind slit 1 by means of a detector D_1 , as shown in Fig. 3-3, and use a similar detector D_2 to count photons scattered behind hole 2. There will be an amplitude for a photon to arrive at D_1 and an electron at x , and also an amplitude for a photon to arrive at D_2 and an electron at x . Let's try to calculate them.

Although we don't have the correct mathematical formula for all the factors that go into this calculation, you will see the spirit of it in the following discussion. First, there is the amplitude $\langle 1 | s \rangle$ that an electron goes from the source to hole 1. Then we can suppose that there is a certain amplitude that while the electron is at hole 1 it scatters a photon into the detector D_1 . Let us represent this amplitude by a . Then there is the amplitude $\langle x | 1 \rangle$ that the electron goes from slit 1 to the electron detector at x . The amplitude that the electron goes from s to x via slit 1 and scatters a photon into D_1 is then

$$\langle x | 1 \rangle a \langle 1 | s \rangle.$$

Or, in our previous notation, it is just $a\phi_1$.

There is also some amplitude that an electron going through slit 2 will scatter a photon into counter D_1 . You say, "That's impossible; how can it scatter into counter D_1 if it is only looking at hole 1?" If the wavelength is long enough, there are diffraction effects, and it is certainly possible. If the apparatus is built well and if we use photons of short wavelength, then the amplitude that a photon will be scattered into detector 1, from an electron at 2 is very small. But to keep the discussion general we want to take into account that there is always some such amplitude, which we will call b . Then the amplitude that an electron goes via slit 2 and scatters a photon into D_1 is

$$\langle x | 2 \rangle b \langle 2 | s \rangle = b\phi_2.$$

The amplitude to find the electron at x and the photon in D_1 is the sum of two terms, one for each possible path for the electron. Each term is in turn made up of two factors: first, that the electron went through a hole, and second, that the photon is scattered by such an electron into detector 1; we have

$$\langle \text{electron at } x \mid \text{photon at } D_1 \mid \text{electron from } s \rangle = a\phi_1 + b\phi_2. \quad (3.8)$$

We can get a similar expression when the photon is found in the other detector D_2 . If we assume for simplicity that the system is symmetrical, then a is also the amplitude for a photon in D_2 when an electron passes through hole 2, and b is the amplitude for a photon in D_2 when the electron passes through hole 1. The corresponding total amplitude for a photon at D_2 and an electron at x is

$$\langle \text{electron at } x \mid \text{photon at } D_2 \mid \text{electron from } s \rangle = a\phi_2 + b\phi_1. \quad (3.9)$$

Now we are finished. We can easily calculate the probability for various situations. Suppose that we want to know with what probability we get a count in D_1 and an electron at x . That will be the absolute square of the amplitude given in Eq. (3.8), namely, just $|a\phi_1 + b\phi_2|^2$. Let's look more carefully at this expression. First of all, if b is zero—which is the way we would like to design the apparatus—then the answer is simply $|\phi_1|^2$ diminished in total amplitude by the factor $|a|^2$. This is the probability distribution that you would get if there were only one hole—as shown in the graph of Fig. 3-4(a). On the other hand, if the wavelength is very long, the scattering behind hole 2 into D_1 may be just about the same as for hole 1. Although there may be some phases involved in a and b , we can ask about a simple case in which the two phases are equal. If a is practically equal to b , then the total probability becomes $|\phi_1 + \phi_2|^2$ multiplied by $|a|^2$, since the common factor a can be taken out. This, however, is just the probability




Fig. 3-4. The probability of counting an electron at x in coincidence with a photon at D in the experiment of Fig. 3-3: (a) for $b = 0$; (b) for $b = a$; (c) for $0 < b < a$.

distribution we would have gotten without the photons at all. Therefore, in the case that the wavelength is very long—and the photon detection ineffective—you return to the original distribution curve which shows interference effects, as shown in Fig. 3-4(b). In the case that the detection is partially effective, there is an interference between a lot of ϕ_1 and a little of ϕ_2 , and you will get an intermediate distribution such as is sketched in Fig. 3-4(c). Needless to say, if we look for coincidence counts of photons at D_2 and electrons at x , we will get the same kinds of results. If you remember the discussion in Chapter 1, you will see that these results give a quantitative description of what was described there.

Now we would like to emphasize an important point so that you will avoid a common error. Suppose that you only want the amplitude that the electron arrives at x , *regardless* of whether the photon was counted at D_1 or D_2 . Should you add the amplitudes given in Eqs. (3.8) and (3.9)? No! You must *never add amplitudes for different and distinct final states*. Once the photon is accepted by one of the photon counters, we can always determine which alternative occurred if we want, without any further disturbance to the system. Each alternative has a probability completely independent of the other. To repeat, do not add amplitudes for different *final* conditions, where by “final” we mean at that moment the *probability* is desired—that is, when the experiment is “finished.” You do add the amplitudes for the different *indistinguishable* alternatives inside the experiment, before the complete process is finished. At the end of the process you may say that you “don’t want to look at the photon.” That’s your business, but you still do not add the amplitudes. Nature does not know what you are looking at, and she behaves the way she is going to behave whether you bother to take down the data or not. So here we must not add the amplitudes. We first square the amplitudes for all possible different final events and then sum. The correct result for an electron at x and a photon at either D_1 or D_2 is

$$\begin{aligned} & \left| \langle e \text{ at } x \mid e \text{ from } s \rangle \right|^2 + \left| \langle e \text{ at } x \mid e \text{ from } s \rangle \right|^2 \\ & \left| \langle \text{ph at } D_1 \mid \text{ph from } L \rangle \right|^2 + \left| \langle \text{ph at } D_2 \mid \text{ph from } L \rangle \right|^2 \\ & = |a\phi_1 + b\phi_2|^2 + |a\phi_2 + b\phi_1|^2. \end{aligned} \quad (3.10)$$

3-3 Scattering from a crystal

Our next example is a phenomenon in which we have to analyze the interference of probability amplitudes somewhat carefully. We look at the process of the scattering of neutrons from a crystal. Suppose we have a crystal which has a lot of atoms with nuclei at their centers, arranged in a periodic array, and a neutron beam that comes from far away. We can label the various nuclei in the crystal by an index i , where i runs over the integers 1, 2, 3, . . . N , with N equal to the total number of atoms. The problem is to calculate the probability of getting a neutron into a counter with the arrangement shown in Fig. 3-5. For any particular atom i , the amplitude that the neutron arrives at the counter C is the amplitude that the neutron gets from the source S to nucleus i , multiplied by the amplitude a that it gets scattered there, multiplied by the amplitude that it gets from i to the counter C . Let’s write that down:

$$\langle \text{neutron at } C \mid \text{neutron from } S \rangle_{\text{via } i} = \langle C \mid i \rangle a \langle i \mid S \rangle. \quad (3.11)$$

In writing this equation we have assumed that the scattering amplitude a is the same for all atoms. We have here a large number of apparently indistinguishable routes. They are indistinguishable because a low-energy neutron is scattered from a nucleus without knocking the atom out of its place in the crystal—no “record” is left of the scattering. According to the earlier discussion, the total amplitude for a neutron at C involves a sum of Eq. (3.11) over all the atoms:

$$\langle \text{neutron at } C \mid \text{neutron from } S \rangle = \sum_{i=1}^N \langle C \mid i \rangle a \langle i \mid S \rangle. \quad (3.12)$$




Fig. 3-5. Measuring the scattering of neutrons by a crystal.

Because we are adding amplitudes of scattering from atoms with different space positions, the amplitudes will have different phases giving the characteristic interference pattern that we have already analyzed in the case of the scattering of light from a grating.

The neutron intensity as a function of angle in such an experiment is indeed often found to show tremendous variations, with very sharp interference peaks and almost nothing in between—as shown in Fig. 3-6(a). However, for certain kinds of crystals it does not work this way, and there is—along with the interference peaks discussed above—a general background of scattering in all directions. We must try to understand the apparently mysterious reasons for this. Well, we have not considered one important property of the neutron. It has a spin of one-half, and so there are two conditions in which it can be: either spin “up” (say perpendicular to the page in Fig. 3-5) or spin “down.” If the nuclei of the crystal have no spin, the neutron spin doesn’t have any effect. But when the nuclei of the crystal also have a spin, say a spin of one-half, you will observe the background of smeared-out scattering described above. The explanation is as follows.

If the neutron has one direction of spin and the atomic nucleus has the same spin, then no change of spin can occur in the scattering process. If the neutron and atomic nucleus have opposite spin, then scattering can occur by two processes, one in which the spins are unchanged and another in which the spin directions are exchanged. This rule for no net change of the sum of the spins is analogous to our classical law of conservation of angular momentum. We can begin to understand the phenomenon if we assume that all the scattering nuclei are set up with spins in one direction. A neutron with the same spin will scatter with the expected sharp interference distribution. What about one with opposite spin? If it scatters without spin flip, then nothing is changed from the above; but if the two spins flip over in the scattering, we could, in principle, find out which nucleus had done the scattering, since it would be the only one with spin turned over. Well, if we can tell which atom did the scattering, what have the other atoms got to do with it? Nothing, of course. The scattering is exactly the same as that from a single atom.

To include this effect, the mathematical formulation of Eq. (3.12) must be modified since we haven’t described the states completely in that analysis. Let’s start with all neutrons from the source having spin up and all the nuclei of the crystal having spin down. First, we would like the amplitude that at the counter the spin of the neutron is up *and* all spins of the crystal are still down. This is not different from our previous discussion. We will let a be the amplitude to scatter with no flip or spin. The amplitude for scattering from the i th atom is, of course,

$$\langle C_{\text{up}}, \text{crystal all down} | S_{\text{up}}, \text{crystal all down} \rangle = \langle C | i \rangle a \langle i | S \rangle.$$

Since all the atomic spins are still down, the various alternatives (different values of i) cannot be distinguished. There is clearly no way to tell which atom did the scattering. For this process, all the amplitudes interfere.

We have another case, however, where the spin of the detected neutron is down although it started from S with spin up. In the crystal, one of the spins must be changed to the up direction—let’s say that of the k th atom. We will assume that there is the same scattering amplitude with spin flip for every atom, namely b . (In a real crystal there is the disagreeable possibility that the reversed spin moves to some other atom, but let’s take the case of a crystal for which this probability is very low.) The scattering amplitude is then

$$\langle C_{\text{down}}, \text{nucleus } k \text{ up} | S_{\text{up}}, \text{crystal all down} \rangle = \langle C | k \rangle b \langle k | S \rangle. \quad (3.13)$$

If we ask for the probability of finding the neutron spin down and the k th nucleus spin up, it is equal to the absolute square of this amplitude, which is simply $|b|^2$ times $|\langle C | k \rangle \langle k | S \rangle|^2$. The second factor is almost independent of location in the crystal, and all phases have disappeared in taking the absolute square. The

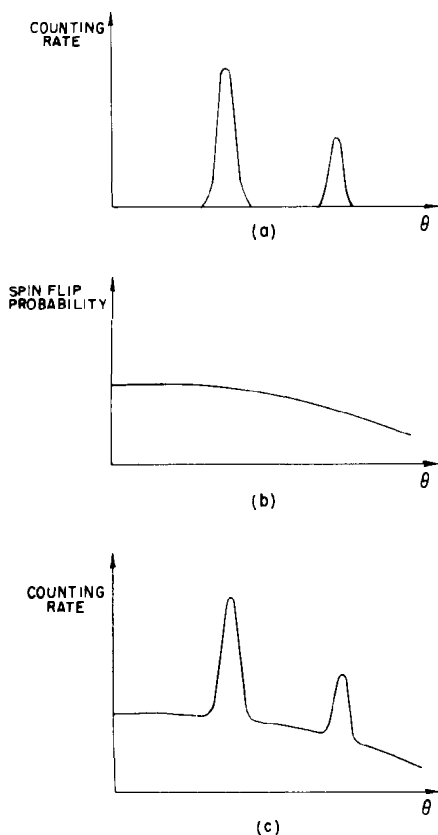


Fig. 3-6. The neutron counting rate as a function of angle: (a) for spin zero nuclei; (b) the probability of scattering with spin flip; (c) the observed counting rate with a spin one-half nucleus.

probability of scattering from *any nucleus* in the crystal with spin flip is now

$$|b|^2 \sum_{k=1}^N |\langle C | k \rangle \langle k | S \rangle|^2,$$

which will show a smooth distribution as in Fig. 3-6(b).

You may argue, "I don't care which atom is up." Perhaps you don't, but nature knows; and the probability is, in fact, what we gave above—there isn't any interference. On the other hand, if we ask for the probability that the spin is up at the detector and all the atoms still have spin down, then we must take the absolute square of

$$\sum_{i=1}^N \langle C | i \rangle a \langle i | S \rangle.$$

Since the terms in this sum have phases, they do interfere, and we get a sharp interference pattern. If we do an experiment in which we don't observe the spin of the detected neutron, then both kinds of events can occur; and the separate probabilities add. The total probability (or counting rate) as a function of angle then looks like the graph in Fig. 3-6(c).

Let's review the physics of this experiment. If you could, *in principle*, distinguish the alternative *final* states (even though you do not bother to do so), the total, final probability is obtained by calculating the *probability* for each state (not the amplitude) and then adding them together. If you *cannot* distinguish the final states *even in principle*, then the probability amplitudes must be summed before taking the absolute square to find the actual probability. The thing you should notice particularly is that if you were to try to represent the neutron by a wave alone, you would get the same kind of distribution for the scattering of a down-spinning neutron as for an up-spinning neutron. You would have to say that the "wave" would come from all the different atoms and interfere just as for the up-spinning one with the same wavelength. But we know that is not the way it works. So as we stated earlier, we must be careful not to attribute too much reality to the waves in space. They are useful for certain problems but not for all.

3-4 Identical particles

The next experiment we will describe is one which shows one of the beautiful consequences of quantum mechanics. It again involves a physical situation in which a thing can happen in two *indistinguishable* ways, so that there is an interference of amplitudes—as is *always* true in such circumstances. We are going to discuss the scattering, at relatively low energy, of nuclei on other nuclei. We start by thinking of α -particles (which, as you know, are helium nuclei) bombarding, say, oxygen. To make it easier for us to analyze the reaction, we will look at it in the center-of-mass system, in which the oxygen nucleus and the α -particle have their velocities in opposite directions before the collision and again in exactly opposite directions after the collision. See Fig. 3-7(a). (The magnitudes of the velocities are, of course, different, since the masses are different.) We will also suppose that there is conservation of energy and that the collision energy is low enough that neither particle is broken up or left in an excited state. The reason that the two particles deflect each other is, of course, that each particle carries a positive charge and, classically speaking, there is an electrical repulsion as they go by. The scattering will happen at different angles with different probabilities, and we would like to discuss something about the angle dependence of such scatterings. (It is possible, of course, to calculate this thing classically, and it is one of the most remarkable accidents of quantum mechanics that the answer to this problem comes out the same as it does classically. This is a curious point because it happens for no other force except the inverse square law—so it is indeed an accident.)

The probability of scattering in different directions can be measured by an experiment as shown in Fig. 3-7(a). The counter at position 1 could be designed to detect only α -particles; the counter at position 2 could be designed to detect