432018 PHILOSOPHY OF PHYSICS (Spring 2002)

Lecture 3: The Quantum Description of Reality

Reading: Preliminary reading: Sklar, pp. 164-72.

In Lecture 2, we considered some of the experiments which lead physicists to question the adequacy of the 'classical' view of the world. In this lecture, we shall consider some more experiments, namely the double-slit experiment and the Stern-Gerlach experiment, which start to reveal just how mysterious the 'quantum world' really is. Along the way, we shall start to familiarise ourselves with some aspects of the formalism of this theory and see the problems associated with some of the initial, pseudo-classical, interpretations of it. In the next lecture, we shall use some of this material to gain a deeper insight into the formalism of QM.

1 The double-slit experiment

To initiate our study of the double-slit experiment, we shall study the outcome of performing the experiment with bullets (i.e. classical particles) and water waves (i.e. classical waves).[1] We shall then compare and contrast this with what we observe when we perform the experiment with electrons (i.e. quantum objects) and try and see how we should account for the results. (A more detailed explanation of these experiments is available in Feynman *et al* [1], pp. 1-1 to 1-5.) In particular, this will lead us to our first explanations of what is happening in QM and we shall consider why these initial attempts at an explanation are inadequate.

1.1 Bullets

Imagine the experimental set-up illustrated in Figure 1 (a). We have a machine-gun that fires bullets

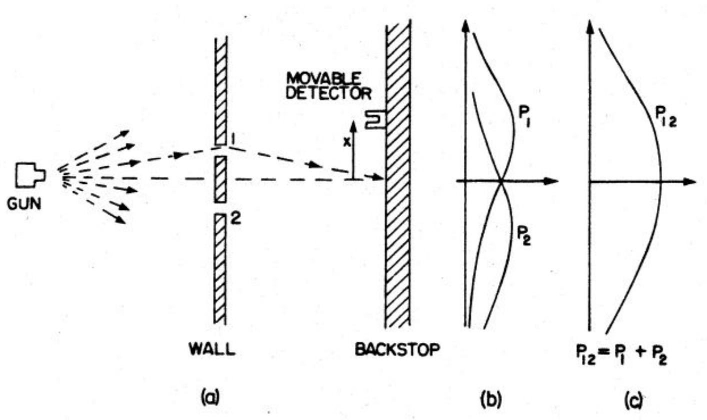

Figure 1: *The double-slit experiment for bullets. (Taken from [1], p. 1-2.)*

in a random direction over a fairly large angular spread. These bullets travel in straight-lines until they reach a wall which has two holes (labelled 1 and 2) in it, at which point: some of them pass right through a hole; some scatter off the sides of a hole[2] and others miss the holes completely and just become embedded in the wall. Beyond the wall there is a movable 'bullet detector' which can be moved up and down (in the x-direction) and we use this to measure the number of bullets that

[1] We use water waves since they are an unashamedly classical phenomena. Perhaps everyone would be happier if we used light 'waves' since their behaviour in a double-slit experiment can be described in a classical manner using Maxwell's [classical] theory of electromagnetism. However, light is really a 'quantum object' — i.e. we sometimes want to talk about photons instead of light waves — and so we don't use it here to avoid any possible classical/quantum confusion.

[2] The scattering makes the bullets change direction, but after they have been scattered in this way, they continue to travel in a straight-line.

reach the backstop for a given value of x (where $x = 0$ is located on the line that goes through the line of symmetry of the apparatus).

With this device in the appropriate place, we can experimentally answer the question: 'What is the probability that a bullet which passes through the holes in the wall will arrive at the backstop at a distance x from the centre?' Note that here, we talk about *probability* for only one reason, namely that the gun fires bullets in random directions and so we cannot predict the direction in which a given bullet will travel. So here, the probability is *epistemic*, it's a measure of our ignorance about some aspect of the gun. Thus, by 'probability', we mean the chance that a bullet will arrive at the detector and, by placing the detector at the appropriate value of x, this can be measured in one of two ways, namely:

1. count the number of bullets that hit the detector in a given time interval and then divide this by the *total* number of bullets that reached the backdrop during this time interval.

2. assuming that the gun fires at a constant rate during the experiment, the probability is proportional to the number that reach the detector in some given time interval.

But, in order for this all to work nicely, we must assume that the experiment is somewhat *idealised*, that is, we must assume that the bullets are *indestructible*. In essence, this means that when we 'detect' bullets we are always detecting *whole* bullets of some fixed shape and size which is completely independent of the rate at which the gun fires. Let's face it, this isn't a great leap of the imagination since bullets rarely break and whatever rate the gun fires at, the bullets are all pretty much the same.

Now, we run the experiment and measure the probability as defined above as a function of x. The outcome of the experiment is sketched in Figure 1 (c) where the x-axis points upwards and is aligned with the corresponding values of x in Figure 1 (a) and the 'measured' probabilities are drawn on a scale that goes to the right. We call this probability *distribution* P_{12} since it tells us the probability that a bullet is detected at a given value of x if the bullets are free to go through either hole 1 or hole 2.

This result is pretty much what we expect. In particular, we expect that:

- P_{12} gets smaller as x gets larger (in magnitude) since the scattering which is required for a bullet to reach such large values of x is quite rare, and

- P_{12} is quite large around the values of x that correspond to 'line-of-sight' shots through either of the holes.

However, we may be surprised that P_{12} takes its maximum value at $x = 0$, especially since bullets can only reach this point if they are scattered in a certain way by one of the two holes. But, we can see why this is the case if we repeat the experiment with hole 2 (hole 1) covered, yielding the probability distribution P_1 (P_2) sketched in Figure 1 (b).[3] In these experiments we find that P_1 and P_2 are both symmetric about their respective maxima and, importantly, we find that our three probability distributions are related by the simple formula:

$$P_{12} = P_1 + P_2,$$

i.e. the probabilities just add together. Thus, P_{12} peaks at $x = 0$ since P_1 and P_2 are both suitably large when $x = 0$.

The important lesson here is that the probability distribution due to both holes being open is just the sum of the probability distributions due to either one or the other hole being open. When this simple rule for combining probability distributions applies we will say that there is *no interference* and we will see why later.[4] This is pretty much how we would expect all classical particles to behave in such experiments: they behave as localised 'wholes' and their probability of arrival at a certain point shows no interference.

[3]Notice that, in this experiment, P_1 (P_2) takes on its largest values around the values of x that correspond to 'line-of-sight' shots through hole 1 (2).

[4]In the language of statistics, when such a rule applies we say that the events 'bullet passing through hole 1' and 'bullet passing through hole 2' are *statistically independent*.

1.2 Water waves

Imagine the experimental set-up illustrated in Figure 2 (a). This is very similar to the apparatus in

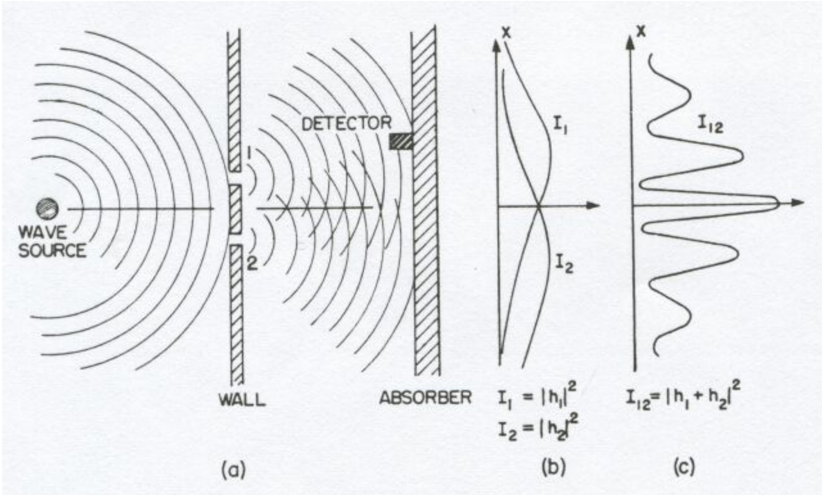

Figure 2: *The double-slit experiment for water waves. (Taken from [1], p. 1-3.)*

Figure 1 (a), but in this case, the apparatus is immersed in a shallow trough of water. And, instead of a gun we have a wave source, which moves 'up and down' (perpendicular to the plane of the paper) in a regular manner generating circular waves that radiate away from it uniformly in all directions as shown. To the right of this source, we have a wall which has two holes (labelled 1 and 2) in it and an absorber — both of which must be made of a material that stops the waves 'dead' (i.e. the waves can't be reflected by the wall or the absorber) unless they happen to pass through one of the two holes. The detector is now a device that measures the *intensity* of the wave a given point, i.e. the amount of energy per unit time that the wave is carrying from the source to the detector.[5]

Firstly, notice that the intensity can have *any* size. Unlike the 'whole bullets' we detected in the previous experiment, we can vary the height of the up and down motion of the wave source so that, at a given point, the intensity can be made as small or large as we please. To put it another way, in this case the intensity can take any non-negative *real* value whereas in the previous experiment we could only get certain non-negative *integer* values for the number of bullets detected.

Secondly, notice that when we run the experiment and measure the intensity at lots of different values of x, we get the intensity distribution I_{12} sketched in Figure 2 (c), and this is very different from what we saw in the previous experiment. Those of you who have done some physics will recognise this as the wave phenomenon known as *diffraction*, i.e. holes 1 and 2 act *like* sources of waves themselves and the waves from these two pseudo-sources *interfere* in such a way to produce the pattern sketched in Figure 2 (c).

In particular, notice that if I_1 (I_2) is the intensity distribution obtained when we run the experiment with hole 2 (hole 1) covered, we find that we obtain the intensity distributions sketched in Figure 2 (b). In these cases, there is only one pseudo-source of waves and so there is no interference since there are no other waves for them to interfere with.[6] Clearly then, I_{12} is *not* the sum of I_1 and I_2 — the former case includes interference effects which are absent from the latter. This is pretty much how we would expect all classical waves to behave in such experiments: they behave as non-localised wholes and their intensity distributions can't simply be added together due to interference effects.[7]

[5] In fact, in this experiment, the intensity is just proportional to the square of the *amplitude* of the wave at that point.

[6] In fact, these intensity distributions have a maximum when the distance between the pseudo-source and the absorber is minimised. Then, as this distance increases, the intensity decreases due to the fact that water waves lose energy as they travel — the further they travel, the less energy they have at the absorber.

[7] However, there is a relationship between these three intensity distributions. A discussion of this, and how interference phenomena come about, can be found in [1] pp. 1-3 to 1-4. Or, indeed, any good A-Level physics textbook.

1.3 Electrons

Imagine the experimental set-up illustrated in Figure 3 (a). This is very similar to our first exper-

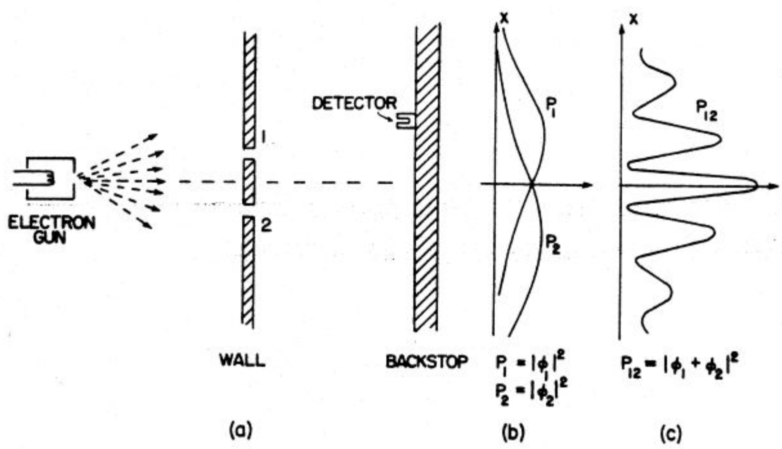

Figure 3: *The double-slit experiment for electrons. (Taken from [1], p. 1-4.)*

iment, but instead of a gun that fires bullets we have a gun that fires electrons. Moreover, instead of a bullet detector we have an electron detector, say a Geiger counter, that makes a clicking sound when an electron hits it.[8]

Now, when we do the experiment we only hear 'whole' clicks. That is, a click corresponds to the arrival of an electron at the detector and so, like bullets, electrons are 'whole' when they reach the detector. Indeed, we notice that:

- if we change the rate at which the electron gun fires out electrons, then we only observe a change in the rate at which we hear the clicks.
- if we have two detectors and fire one electron from the gun, we will only only hear one click from one of the detectors.
- if we run the experiment twice, then provided that we allow it to run for a sufficiently long period of time, we find that the number of electrons which we detect is [pretty much] the same.

And, as such, the electrons seem to be behaving like particles. Thus, like in the bullet case, we use the detector to measure the probability that an electron passing through the holes in the wall will arrive at the backstop at a distance x from the centre. But, unlike in the bullet case, we get the probability distribution given in Figure 3 (c), which is reminiscent of what we found with water waves. Strange indeed!

But, note that unlike the water wave case, this is a variation in the average rate at which the clicks are heard, i.e. we do *not* observe electrons spread out like waves, just 'whole' electrons. So, electrons seem to be behaving likes waves (they exhibit interference phenomena) and particles (they are detected as 'wholes').

Electrons interfere with one another?

When we detect electrons, we detect 'whole' electrons. So, despite the weird behaviour observed in Figure 3 (c), we may feel that we can still explain this phenomena. For example, since we only observe 'whole' electrons, they are like bullets and either go through hole 1 or hole 2. That is, we could claim that:

(A) Each electron *either* goes through hole 1 *or* it goes through hole 2.

[8]This is only a representation of the experiment that we would have to do for several reasons. But, for simplicity, we will not discuss these since such details are not essential.

If this is true, then every detected electron is either an electron that went through hole 1, or an electron that went through hole 2. However, if we repeat the experiment with hole 2 (hole 1) covered we get the probability distribution P_1 (P_2) sketched in Figure 3 (b). This seems reasonable, if electrons are like bullets then this is what we would expect. However, if (A) is true, then P_{12} should just be the sum of P_1 and P_2, just as we found with the bullets. But, this is clearly not the case — i.e. (A) is *false.* In fact, by analogy with our water wave experiment, we seem forced to conclude that there is interference of some kind taking place due to the presence of holes 1 *and* 2.

1.4 Another experiment with electrons

Imagine the experimental set-up illustrated in Figure 4(a). This is very similar to our last experiment,

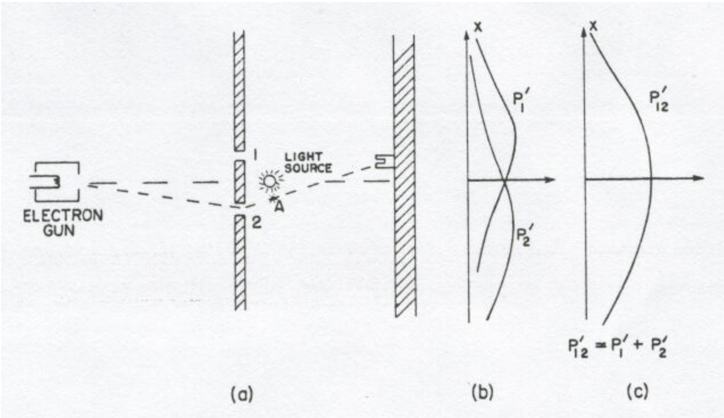

Figure 4: *The modified double-slit experiment for electrons. (Taken from [1], p. 1-7.)*

except now we have a very strong light source just behind the wall. The thinking is, that since electrons scatter light, if an electron goes through hole 2 we will see a little flash of light in the vicinity of the point marked A in Figure 4 (a). Similarly, if an electron goes through hole 1, we will see a flash of light near that hole. Presumably, if the electron can go through both holes we will see two flashes of light, one in the vicinity of each of the holes, at pretty much the same time.

But, when we perform the experiment we find that *every* click of the detector is preceded by a flash of light which is *either* near hole 1 *or* near hole 2. That is, we *never* observe a flash of light in the vicinity of *both* holes. Thus, in this case, we observe that (A) is true. Indeed, if we run this experiment with hole 2 (hole 1) covered we get the probability distribution P_1' (P_2') sketched in Figure 4 (b) which are similar to what we found for P_1 and P_2 in the previous electron experiment. But now, if we run the experiment with both holes open we find that we get the probability distribution P_{12}' given in Figure 4 (c), which looks suspiciously like what we found in the bullet experiment, and is in fact just the sum of P_1' and P_2'! Thus, in this experiment, where we use a light source to see which holes the electrons are going through, we find that:

$$P_{12}' = P_1' + P_2',$$

i.e. we have *no* interference. So, it would seem that the addition of the light source which allowed us to observe which hole the electrons were going through changed the outcome of the experiment!

1.5 Conclusions from our experiments

In summary then, we can see that:

- Bullets (i.e. classical particles) are localised and go through *either* hole 1 *or* hole 2 (i.e. (A) is true for them). In this case there is no interference in the probability distributions.
- Water waves (i.e. classical waves) are not localised and each wave has a part that goes through *both* holes (i.e. (A) is not true for them). In this case the parts of a wave that go through different holes interfere with each other.

- Electrons (i.e. quantum objects) are neither classical particles nor classical waves. In fact, depending on the experiment we perform we see that:
  - If we don't observe which hole the electrons go through, we find that the probability distribution indicates that there is interference. That is, the electrons seem to be acting a bit like waves (at least, before they are detected).
  - If we do observe which hole the electrons go through, we find that the probability distribution indicates that there is no interference. That is, the electrons seem to be acting a bit like particles.

  So, electrons seem to act like particles or waves depending on what experiments we perform on them. This is called *wave-particle duality* and it is a phenomenon which is observed with all quantum objects.[9]

For a further discussion of these experiments, especially the ones involving the electrons (including why some other 'plausible' explanations don't work), you are strongly encouraged to read [1] pp. 1-1 to 1-9.

2 The wavefunction and its initial interpretation

We saw in the last lecture that light, which at the end of the nineteenth century was considered to be an electromagnetic wave, could also act like a particle. In 1924, this prompted the physicist Louis de Broglie to conjecture that particles may, under circumstances, also act like waves. Clearly, this idea gains some plausibility when we consider the first of our electron experiments — the electrons appear to be interfering with each other, and so they do seem to be acting like waves in this experiment.

This idea lead Erwin Schrödinger to develop a mathematical theory that would explain this phenomena. This is the formalism of QM that is usually referred to as *wave mechanics*. Within this theory, each electron is described by a *wavefunction*, and the way the behaviour of the electron is governed by the Schrödinger equation. This is called *wave* mechanics since the wavefunction looks a bit like a function that describes a wave and the Schrödinger equation looks a bit like a wave equation.

The initial *interpretation* of this theory closely followed de Broglie's idea that in the quantum world, objects that we commonly thought of as particles could act like waves. That is, when the object acts like a wave it is described by the wavefunction which was thought to represent the object as a physically real non-localised wave — i.e. a *matter* wave. For example, the ϕ_1, ϕ_2 and ϕ_{12} in Figure 3 (b) and (c) are the wavefunctions of an electron and they should be treated in a way which is analogous to the treatment of the functions h_1, h_2 and h_{12} which describe the water waves in Figure 2 (b) and (c).

However, this proposed interpretation has some severe problems. Firstly, we have the intuitive problems:

- What, exactly, is acting like a wave? Is it the electron? Are we seriously saying that the electron is non-localised enough so that a single electron can act like a wave going through both holes and hence interfere with itself? If not, what?
- When we try and locate electrons (say using our detectors) we always find them concentrated in a *small physical region*. How can we reconcile this with the claim that an electron should be identified with a physically real *spread out wave*?

and, secondly, we have the problems for this interpretation which appear due to the formalism:

- We can only interpret the non-localised wavefunction (which describes an electron) as a wave in physical space and time if we have *one* electron. That is, if we want to describe the behaviour of *many* (say $n \geq 2$) electrons, the wavefunctions only look like waves in an abstract $3n$-dimensional space.[10]

[9] From a classical perspective, we would want to say that this is absurd, classical objects are *either* particles *or* waves, but this just goes to show how different the classical and quantum worlds are!

[10] In which, incidentally, the positions of all the electrons at a given time is treated as a single point.

- The wavefunction takes on complex values whereas we can only interpret real values physically. How can we make physical sense of a function that can take on complex values?[11]

Thus, even though this interpretation may sound plausible, we can't adopt it.

3 Born's interpretation of the wavefunction

Despite the fact that the initial interpretation of the wavefunction failed, the wavefunction and the description of its behaviour using the Schrödinger equation does provide a very good description of what happens in the quantum world. So, we want to keep the formalism and find a new interpretation. In 1926, Max Born suggested that:

> The *intensity*, $|\phi|^2$, of the wavefunction ϕ should be interpreted as a *probability*.

Thus, the wavefunction itself is a probability *representer* and not a physical wave in the world. So, in the above electron experiments we find that:

> The intensity of the wavefunction represents the *probability* with which a physical observable (such as the position of an electron) would be found to have a given value if an *appropriate measurement* were made (by, say, using the detector or a light source).

In particular, this gives us a way of reconciling the localised nature of particles and the non-localised nature of the wavefunction, i.e.

> The wavefunction does *not* represent an *actual* spread out particle, but it does allow us to calculate the probability of finding the particle in a given region of space.

which seems to tally with what we were actually trying to measure in our electron experiments.

However, there is a problem with this interpretation too — it doesn't account for the interference phenomena. To see why, notice that:

- Suppose that we the obtain an outcome, with a given probability, in two causally independent ways. For instance, let $P(O|A)$ and $P(O|B)$ be the probabilities of getting the outcome O given A and B (the two ways in which O can be obtained) respectively. By the rules of probability, since A and B are independent, the probability of O obtaining due to either A or B, i.e. $P(O|A \text{ or } B)$, is then just the sum of these two probabilities.

- So, suppose that we perform the first electron experiment and let O, S_1 and S_2 be the propositions 'the electron is detected at a given point', 'the electron went through hole 1' and 'the electron went through hole 2' respectively. According to the probability rule given above, this means that we have:
$$P(O|S_1 \text{ or } S_2) = P(O|S_1) + P(O|S_2).$$
(Cf. the way the probabilities worked in the bullet experiment.) But, this is *not* what happens! There is an interference phenomenon which is absent from this simple probability interpretation.

So, following Born, we do want the intensity of the wavefunction to represent a probability, but a simple probabilistic interpretation like the one above is not going to work. So the question is: If the wavefunction is a probability representer, and not a physical wave, how can such interference happen?

[11] Notice that we *do* sometimes describe physical waves using complex valued functions. But, when we do this, it is a mathematical convenience that can be avoided. In the case of wavefunctions, there is no way of avoiding the use of complex valued functions.

4 Heisenberg's wave mechanics

So far, we have a good formalism for making predictions about what happens in experiments involving quantum phenomena. We have wavefunctions which are governed by the Schrödinger equation and we interpret the intensity of the wavefunction as a probability (but not in the simplistic way shown above). But, in 1924, Werner Heisenberg developed an alternative formalism for QM called *matrix mechanics*. This formalism is equally good at making predictions and yet it says nothing about particles sometimes being waves. In brief, this alternative formalism discusses the evolution of a quantum system by describing how the *state* of the system changes over time. So, for example, within this formalism, unless we are explicitly studying the 'position states' of an electron (i.e. the states that tell us where it is) the [spatial] trajectory of an electron need not even be defined. In fact, it is easier to discuss QM using this formalism and so from now on we will tend to use this.[12] To get an idea of what happens in matrix mechanics we will start by considering another set of experiments. Then, in the next lecture, we will introduce the postulates of QM using this formalism and illustrate the use of these postulates and this formalism in the context of these experiments.

5 The Stern-Gerlach experiment

The Stern-Gerlach experiment, first performed in 1922, involved passing a stream of silver atoms through a non-uniform magnetic field. The details of the experiment don't really concern us, but a schematic illustration of the apparatus is provided in Figure 5 and to further simplify our discussion we will consider a stream of electrons rather than silver atoms. The surprising thing about this

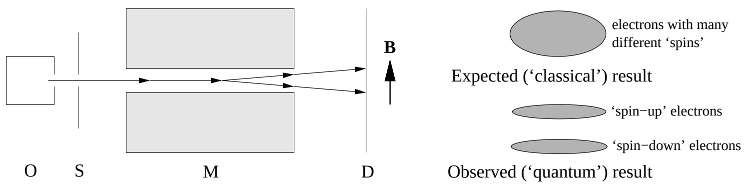

Figure 5: *On the left-hand side we have a schematic illustration of the Stern-Gerlach experiment. O (for 'oven') represents the source of the electrons, S is a screen with a slit in it, M is the magnet and D is the detecting device. The 'arrowed-line' represents the trajectory of the electrons. The vector* **B** *represents the direction of the magnetic field inside the magnet M. On the right-hand side we have a schematic illustration of the expected and observed results of the experiment.*

experiment is that Stern and Gerlach were *expecting* to find a continuous 'classical' distribution of electrons like the one in the top right-hand corner of Figure 5, but they actually *observed* a distribution like the one in the bottom right-hand corner of Figure 5. It turns out that the observed 'two part' distribution is indicative of a new quantum phenomena,[13] namely that electrons have an intrinsic magnetic moment or *spin* and this property of electrons is quantised. That is, the spin of an electron can only take on one of two distinct values which are normally referred to as 'spin-up' and 'spin-down' — see Figure 6.

Due to the fact that the spin of an electron can only take two values it is particularly useful when investigating the bizarre behaviour of quantum systems. In particular, we note that in the Stern-Gerlach experiment the electrons which we classified as spin-up and spin-down moved 'up' or 'down' in the direction of the magnetic field described by the vector **B** in Figure 5. As such, an electron's spin is always found to be spin-up or spin-down in a *specified* direction. In particular, we will now consider using Stern-Gerlach-like apparatus to measure the spin of an electron in two mutually perpendicular directions which, for convenience, we label the x and y-directions.[14] So,

[12] In fact, wave mechanics and matrix mechanics are mathematically equivalent and so we are not actually questioning the validity of either formalism. It just so happens that there is less 'conceptual baggage' in the latter formalism, i.e. there is nothing in it which suggests that we should be talking about particles or waves.

[13] Which, incidentally, has no classical analogue.

[14] That is, one will have **B** pointing in the x-direction and the other will have **B** pointing in the y-direction. To see why we label the directions in this way, think of a Cartesian coordinate system.

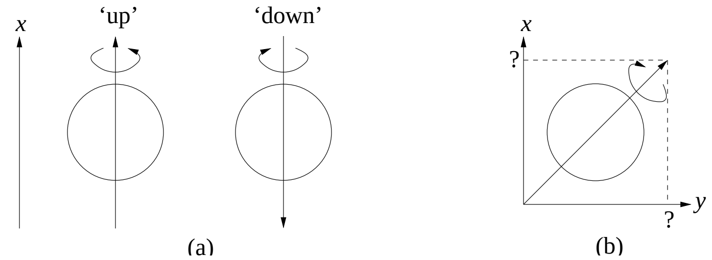

Figure 6: *One way of thinking about an electron's spin is to think of an electron as a 'little sphere'. In (a) we think of this 'little sphere' as rotating about an axis that points in the x-direction and take the state of anticlockwise (clockwise) rotation around this axis to correspond to what we have called 'spin-up' ('spin-down'). But, as seen in (b), this way of thinking about things is seriously flawed as such a 'little sphere' could be construed as rotating around an axis which is at an angle to the x and y-axis (say), i.e. we would get 'intermediate' values of spin in those directions (denoted by '?'s in the figure), and this never happens.*

using this phenomenon, let's consider another set of experiments that are indicative of the weirdness of quantum mechanics.

Stern-Gerlach experiments and states

We shall represent the Stern-Gerlach experiments that we are using schematically, as in Figure 7. In

Figure 7: *Schematically, we have electrons in some 'unknown' spin state entering 'in' to a Stern-Gerlach apparatus. (a) for an SGX-box, electrons leaving through the 'u' ('d') hole are in the 'up-x-spin' ('down-x-spin') state. (b) for an SGY-box, electrons leaving through the 'u' ('d') hole are in the 'up-y-spin' ('down-y-spin') state.*

particular, we shall utilise Stern-Gerlach apparatus that are aligned in two mutually perpendicular directions, let's call them the x and y-directions, so that we can see how two mutually perpendicular sets of spin states are related. For simplicity, we shall use a

- Stern-Gerlach apparatus with its **B**-field aligned in the x-direction, schematically indicated by the SGX-box in Figure 7 (a), to measure the spin of the incident electrons in the x-direction. Due to the nature of this apparatus, those electrons leaving through the 'u' ('d') hole are in the 'up-x-spin' ('down-x-spin') state.

- Stern-Gerlach apparatus with its **B**-field aligned in the y-direction, schematically indicated by the SGY-box in Figure 7 (b), to measure the spin of the incident electrons in the y-direction. Due to the nature of this apparatus, those electrons leaving through the 'u' ('d') hole are in the 'up-y-spin' ('down-y-spin') state.

Indeed, notice that after the electrons have left a Stern-Gerlach experiment *via* one of the holes, all we know about them is their spin state.

Stern-Gerlach experiments are 'repeatable'

To ensure that we are not wasting our time, we want to be sure that the results of our Stern-Gerlach experiments are 'good' or 'reliable'. In particular, we want to be sure that once we have ascertained which spin state an electron is in, another Stern-Gerlach experiment would agree with this result. To show that this is the case, consider the situation illustrated in Figure 8. Here, we pass some electrons in an 'unknown' spin state into an SGX-box and take those that come out of the 'u' hole, i.e. those

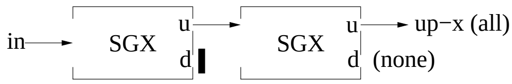

Figure 8: *Schematically, the fact that Stern-Gerlach experiments give us 'good' measurements can be seen by passing the 'up-x-spin' electrons from one SGX-box into another SGX-box. As one would expect from a 'good' measurement, the second SGX-box agrees that all of the electrons entering it are in the 'up-x-spin' state. Obviously, checking the result of an 'down-x-spin' measurement or a 'y-spin' measurement in this way would proceed along similar lines.*

that the Stern-Gerlach experiment assures us are in the up-x-spin state, and pass these (and only these) electrons through another SGX-box. Reassuringly, the second measurement on these electrons (due to the second SGX-box) sees them all coming out of the 'u' hole confirming that they are, as we initially expected, all in the up-x-spin state.

Indeed, as long as we don't 'tamper' with the electrons whilst they are between the two SGX-boxes (see below), this always works. Indeed, it also works if we look at the electrons which leave the first SGX-box by the 'd' hole, or if we consider the electrons that leave an SGY-box through the 'u' or 'd' holes.

How are x and y-spin states related I — 'statistical independence'

Now, at this point you may be wondering why we are bothering to consider both the x and y-spin states of an electron. Well, the reason is that, *classically*, we may expect the spin of an electron in these two mutually perpendicular directions to be related in some way. Let's see if they are *statistically* related in some obvious manner. For example, we may expect up-x-spin electrons to yield more up-y-spin electrons and, such a relationship could reveal itself through a *correlation* between the number of up-x-spin electrons entering an SGY-box and the number of up-y-spin electrons leaving it.

So, to look for such a relationship, we take a set of electrons which are in the up-x-spin state, for example a set of electrons that have just left the 'u' hole of an SGX-box, and we pass these electrons through an SGY-box — see Figure 9. Now, if the x and y-spins of these electrons are

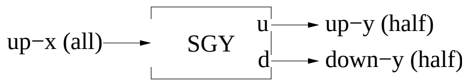

Figure 9: *Passing up-x-spin electrons through an SGY-box yields an equal number of up-y-spin and down-y-spin electrons. As such, there appears to be no correlation which allows us to decide whether an up-x-spin electron is more likely to leave the SGY-box as an up-y-spin or a down-y-spin electron.*

related, we would expect to see some correlation between the fact that all the incident electrons are spin-up in the x-direction and the results of the y-spin measurement performed by the SGY-box. But, on performing the experiment, we find that one half of the electrons go through the 'u' hole of the SGY-box and other half go through the 'd' hole of the SGY-box. That is, the fact that all of the incident electrons were spin-up in the x-direction doesn't seem to influence the outcome of the y-spin measurement. As such, since we observe no correlations, we infer that the x and y-spins are statistically independent, i.e. each up-x-spin electron is equally likely to emerge from an SGY-box as an up-y-spin or a down-y-spin electron.

How are x and y-spin states related II — 'disturbance'

However, you may still feel that there is some relationship between the x and y-spin states. For example, you might like to think that the absence of a correlation is just a coincidence along the lines of 'it just so happens that every set of up-x-spin electrons is just a half-half *mixture* of up-y-spin and down-y-spin electrons, and that's why we don't see any correlations'. If this was the case, the experiment in Figure 10 (a) would take only the up-x-spin electrons from the first SGX-box and then pass them through the SGY-box. As such, looking at the electrons which issued from the 'u' hole of this latter box we would have electrons which were *both* up-x-spin electrons (as they were before

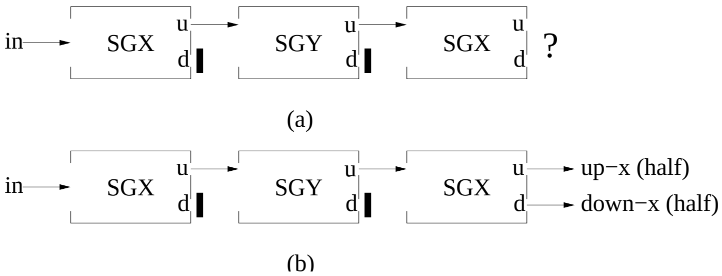

Figure 10: *In (a), what do you expect to happen? In (b), what actually happens. Clearly, passing an up-x-spin electron through an SGY-box makes it 'forget' that it initially had such a spin in the x-direction. That is, an SGY-box affects an electron's x-spin in such a way that further measurements by an SGX-box have a 50-50 chance of yielding an electron which is up-x-spin or down-x-spin.*

entering the SGY-box) and up-y-down electrons (as they are after leaving the SGY-box), i.e. the SGX and SGY boxes are just *sorting* these electrons out from the 'randomly oriented' set of electrons that were initially passed into the experiment. So, doing this experiment, the result we would expect from the second SGX-box (the '?') would be a stream of up-x-spin electrons (coming out of the 'u' hole) and *no* electrons coming out of the 'd' hole.

But, on running the experiment — see Figure 10 (b) — what we actually observe is half the electrons coming out of the 'u' hole and the other half coming out of the 'd' hole. That is, the up-x-spin electrons which entered into the SGY-box have emerged from this experiment as a 50-50 mixture of up-x-spin and down-x-spin electrons. So, there is still no correlation and furthermore, the measurement of the y-spin of these electrons has *altered* their x-spin. Thus, we say that a measurement of the y-spin of an electron *disturbs* its x-spin.

This may prompt you to ask two questions:

1. Could we build the SGY-box less crudely? That is, couldn't we build an SGY-box which doesn't affect the x-spin of the electrons?

2. For the SGY-box that we have been using (i.e. the crudely built one), what is it that determines precisely which electrons have their x-spin affected?

and there are two answers:

1. We *can* build SGY-boxes in different, less crude, ways. *But*, they all give results which are consistent with those that we have been discussing! What's important here is *not* that we can't build SGY-boxes that don't disturb the x-spin of an electron, but that however we construct an SGY-box we always get the *same* distribution of up-x-spin and down-x-spin at the end of such experiments. Thus, so long as the SGY-box fulfils the *definitional* requirements of such a piece of apparatus[15] it will always completely *randomise* the x-spins of the electrons that go through it! As such, it appears that the answer to this question is *no*, i.e. two apparently indistinguishable up-x-spin electrons can leave the experiment in different x-spin states!

2. Currently, there are no [observable] correlations between the up-x-spin state of an electron and the states of an SGY-box which will allow us to determine which of the initially up-x-spin electrons will have their x-spin changed and those which won't! So, this question seems to have no answer…

Out of desperation, you may ask a third question, namely:

3. Couldn't we avoid this problem by building a device that measures *both* the x and y-spin of electrons?

to which one would get the reply:

[15]That is, so long as it is a device which we can *reliably* use to determine the y-spin of an electron.

3. How could we build such a device? Surely it would need to contain both an SGX and an SGY-box, but the former would affect the output from the latter, and *vice versa*, leading to *unreliable* information about the x and y-spins of the electrons. So, since we have no conceivable way of constructing such a device, it seems to be fundamentally beyond our means!

This last point is actually an example of the quantum mechanical *uncertainty principle*,[16] i.e. measurable physical properties like x and y-spin are said to be *incompatible* with one another since measurements of one will (as far as we know) *always* disrupt the other.

How are x and y-spin states related III — 'superpositions'

We end this discussion of Stern-Gerlach experiments and the strange nature of electron spin by discussing the most bizarre set of experiments. In Figure 10 we took the up-x-spin electrons from an SGX-box and then passed them through an SGY-box. We then took the up-y-spin electrons from this SGY-box and observed what happened when we put them through another SGX-box. Notice that, since we *only* took the up-y-spin electrons from the SGY-box we made a measurement at this point, i.e. prior to passing these electrons through the second SGX-box we knew that they were all in the up-y-spin state. And, at the end of the experiment we observed that half of these electrons were in an up-x-spin state and the other half were in a down-x-spin state.

Now, let's consider what would happen if we passed *both* the up and down-y-spin electrons from the SGY-box through the second SGX-box. In fact, as illustrated in Figure 11 (a), let's count the number of electrons that come out of the SGY-box in the up and down-y-spin state. If we do this, we

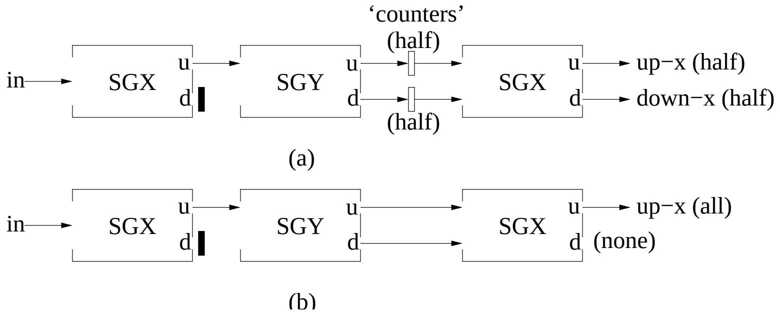

Figure 11: *The effect of measuring the spin. In (a) we count the number of electrons that leave the SGY-box (and, as such, gain information about their y-spin). In (b) we don't. Clearly, making such a measurement has an effect on the electrons since these two experiments give different results.*

find (unsurprisingly given the experiment in Figure 9) that half of the electrons are found in each of these two states. And, passing these 'counted' electrons through the second SGX-box we find (again, unsurprisingly given the results of the experiment in Figure 10) that half of the electrons leave the experiment in the up-x-spin state and the other half are in the down-x-spin state.

However, what happens if we run this experiment again, but this time we don't count (or observe in any other way) the electrons that leave the SGY-box? That is, we know that the electrons leaving the first SGX-box are all up-x-spin electrons (since we don't allow the down-x-spin electrons to leave this apparatus), but we don't make any attempt to observe the state of the electrons leaving the SGY-box. This experiment — illustrated in Figure 11 (b) — has a rather surprising outcome, namely that (*unlike* in the previous experiment) we now observe that *all* of the electrons leave this experiment in the up-x-spin state. That is, in this experiment, the SGY-box seems to have had *no* effect on the x-spin of the electrons entering into it!

Consequently, as hinted at above, it is *not* the SGY-box itself that affects the x-spin of the electrons, but the fact that up until now we have always been *observing the results* of passing electrons through an SGY-box. Indeed, compare this with the different results we got when we ran the electron double-slit experiment *with* a way of observing which hole the electrons went through[17] and *without* such an observational aid. The fact that making or not making a *measurement* can have such an

[16] Something which we have taken great pains to avoid in this course!

[17] Recall that we did this with a 'strong light source' positioned between the holes and the 'backstop'.

extreme effect on the outcome of an experiment is, perhaps, the key to trying to understand the weirdness of QM.[18]

References

[1] R. P. Feynman, R. B. Leighton and M. Sands, *The Feynman Lectures on Physics, Volume III* (Addison Wesley, 1965).

James Ward (e-mail: `j.m.ward@lse.ac.uk`)

[18]Note: although this all revolves around the term 'observation', it seems clear that this is observation *qua* 'making a measurement using some physical apparatus' and not observation *qua* 'coming to know the outcome of a measurement'. That is, these effects do *not* seem to be epistemic, they appear to be real physical effects that would occur even in the absence of a conscious entity who would actually 'observe' the results of such measurements. (Although, as we shall see, some interpretations of QM disagree with this!)