Introduction

Introductory Remarks

This is probably your first real course in quantum mechanics. To be sure, it is understood that you have encountered an introduction to some of the basic concepts, phenomenology, history, and so forth, of quantum mechanics in a course on "modern physics". But this is presumably your first chance to get your hands dirty and see to some extent how the theory *really* works. Although you should have encountered already a survey of the key experimental underpinnings of quantum mechanics, it is worth setting the stage by giving some introductory remarks on what quantum mechanics is and why we use it. Of course, these remarks, at best, can serve only to whet your appetite.

Quantum mechanics – and, in particular, its extension to *quantum field theory* – is used to give a microscopic description of matter and its interactions at or below atomic length scales. It is the currently accepted description based upon extensive experimental investigation. Nevertheless, the *way* in which quantum mechanics describes nature is sufficiently different from successful descriptions of macroscopic phenomena (*e.g.,* from classical mechanics), that I feel obliged to justify the existence of the subject to you. The word "macroscopic" is an important clue to the need for quantum mechanics.

Consider a classical description of the structure and interactions of matter. Normally, one assumes that matter can be broken into "small" constituents, which we mathematically model as "particles", and which obey some set of physical laws, say, Newton's laws, and from which the behavior of matter can be explained and predicted. Of course, one can always imagine subdividing the previous constituents into a new set of "smaller" particles, *etc.* One logical possibility is that Newton's laws continue to hold at smaller and smaller length scales. If this is the case, then there is apparently no end to the classical reductionist approach to describing matter. In classical mechanics the explanation of phenomena involving "large" objects by simple laws describing the "small" objects is an infinite regression. Still, as I said, it is possible the universe could have been set up along these lines. ("Turtles all the way down!") But according to experiment it isn't.

The laws of nature do *not* seem amenable to an analysis along the lines just presented. There is a real distinction between macroscopic and microscopic phenomena. Macroscopic phenomena seem well-described by the laws of classical mechanics, while microscopic phenomena require a different set of laws, the laws of quantum mechanics being the current best set. In particular, at sufficiently small length scales it becomes clear that there are fundamental limits to our ability to make certain kinds of basic measurements – an ability that is taken for granted in classical mechanics. For example, the position and momentum of a particle are the basic observables of classical mechanics; the laws of classical mechanics naturally deal with the time evolution of these observables. In particular, by knowing the position and momentum at one time one can, using for example Newton's laws, deduce the trajectory of the particle in space for all time, hence the position and momentum (and any other observables) are known for all time. All this is quite familiar to you. It is tacitly assumed in classical mechanics that it is possible to determine with arbitrary accuracy the position and momentum of a "particle" at any given time. However, when describing microscopic systems, *e.g.,* atoms, it turns out that it simply is not true that, say, the electron has a sharply defined position and momentum at any given time. Consequently, the organization of the theory describing the atom in terms of position and momentum of its constituent particles is simply not appropriate. As you probably have heard, if the position of a particle is determined with great accuracy, the momentum has a wide variety of possible values, and vice versa. This is the celebrated "uncertainty principle", which we shall make more precise later. The uncertainty principle applies to a variety of observables—not just position and momentum; many of the "classical" attributes of particles cannot be determined with complete accuracy in the usual classical sense–*even in principle.* For this reason there is no compelling reason to expect that these attributes (such as position and momentum) are intrinsically part of a "particle"; something else may serve better to describe matter.

Evidently, there *is* a real distinction between large and small in the universe. One might say that "large" objects can have their classical observables (*e.g.,* position and momentum) determined at a given time with very, very good accuracy. "Small" objects are such that it is *impossible* to measure the position without disturbing the momentum and vice versa. In a real, operational sense, one must give up our classical notions of particles as "objects" that have a definite position, momentum, and so forth, at least in the classical sense. If position and momentum are denied their classical existence, then it follows that particle trajectories (defined by Newton's second law) are denied a microscopic existence. The usual *modus operandi* of classical mechanics is simply not available and a new mode of description of nature will be needed. Such a mode is provided by the (at first sight, slightly bizarre) laws of quantum mechanics.

An illustrative thought experiment

To illustrate the above remarks, I would like to briefly discuss some thought experiments (blatantly stolen from the *Feynman Lectures*), that will perhaps make the point better than my clumsy pseudo-philosophical arguments. These are variations on the famous "double slit experiments", which you may have encountered in a previous introduction to quantum mechanics. The "experiments" I will describe are over-simplifications of real experiments.

We shall consider a thought experiment in which we send a beam of particles through a barrier with two openings, the “slits”. We shall first indicate what would be seen in a macroscopic context, which will be pretty believable given our experience with everyday mechanical phenomena. We then discuss what would be seen in a microscopic context; the result is very surprising from the point of view of classical mechanics.

Macroscopic version

Let us fire a steady stream of bullets at an indestructible partition with two holes just big enough to let the bullets go through. We assume that the gun firing the bullets sprays them with some angular spread so that it is possible that a given bullet could go through either hole. On the other side of the barrier there is a backstop that stops the bullets so that at any time we can have a look and see how many bullets went where. You can easily guess what will happen. Bullets will either pass through the holes, be “reflected” by the partition, or perhaps scatter from the edge of a hole either backward or forward. What is seen at the backstop? First of all, for sufficiently short time intervals, only one bullet impact is detected during the time interval. Thus the impacts come in discrete lumps, which confirms our understanding that bullets are localized in space at any given time.

To further quantify the results of this experiment we introduce (somewhat informally for now) the notion of probability. If we pick a spot on the backstop, after a fixed amount of time we can count how many bullets hit that spot. If we take the ratio of this number to the total number of bullets that hit the backstop in that time interval we can interpret the ratio (for a large enough total number of bullets) as the chance that the bullet was scattered to our chosen spot. This ratio is a real number between 0 and 1 and is the *probability* for scattering the bullet to the given location. If we graph the probability as a function of position we find a curve with a maximum centered between the slits and decreasing monotonically as we move away from the center. This is the *probability distribution* for finding a bullet in the backstop.

Finally, it is useful for later comparison to repeat the experiment with one of the slits closed. Then, of course, the bullets can only pass through the open slit and we obtain a probability distribution that is qualitatively similar to the one we just found, but now centered directly behind the open slit.

Let us call the probability distribution for each slit $P_1(x)$ and $P_2(x)$, where x locates positions from the center point between the slits. The distribution P_1 is measured with slit number 2 closed, *etc.* Let us denote the probability distribution obtained with both slits open by $P(x)$. Given your everyday experience with macroscopic objects, it will not surprise you that $P(x)$ is the sum of $P_1(x)$ and $P_2(x)$.

To summarize: the results of the experiment are that bullets arrive at the backstop in discrete “lumps”, and the probability $P(x)$ for finding a bullet at the location x is given by

$$P(x) = P_1(x) + P_2(x).$$

Of course, we interpret this result by noting that a bullet must have traveled through either slit 1 or slit 2, and the total probability is just the sum of the probability for these mutually exclusive processes.

Microscopic version

Let us now repeat this experiment using electrons and very small slits. Now we must use a more sophisticated means of detecting the electrons (*e.g.,* a Geiger counter), but let us not worry about the technicalities. To begin with, we note that the detector "clicks" in a discrete fashion, which is consistent with our view of the electron as a discrete, lump of matter. Next, let us try the experiment with only one slit open. As before, we find that the probability distributions P_1 or P_2 to be peaked at the x location of slit 1 and slit 2, respectively, with a monotonic decrease away from the peak. Now we try the experiment with both slits open. Here is where nature is surprising. Instead of a peaked distribution corresponding to the sum of P_1 and P_2 we find instead an oscillatory pattern within an envelope that is peaked about $x = 0$. This oscillatory pattern is exactly like that which occurs in the amplitudes of waves (*e.g.,* water or light) that are passed through a barrier with two slits. In the wave case we interpret the oscillating pattern in terms of *interference* of the waves scattered from the two slits.

It is possible to cook up a mathematical description of the electron experiment to match the interference pattern of waves if we assign a *complex* "amplitude" ψ_1 to the waves passing through the first slit and ψ_2 to those passing through the second slit. The total amplitude is then the sum:

$$\psi(x) = \psi_1(x) + \psi_2(x).$$

The probability $P(x)$ for finding an electron at detector location x will take the desired oscillatory pattern if we define

$$P(x) = |\psi_1(x) + \psi_2(x)|^2 = |\psi_1(x)|^2 + |\psi_2(x)|^2 + 2\Re[\psi_1^*(x)\psi_2(x)].$$

We can view the first two terms as representing the probability distributions coming from electrons that pass through slit 1 and slit 2 respectively. Indeed, if we cover up one of the slits, say slit 2, then we get $P(x) = |\psi_1(x)|^2$. However, because of the last "interference" term, we cannot ascribe the total probability distribution when both slits are open to the effects of each slit separately. In effect, this experiment prevents us from saying that the electrons always travel through slit 1 or slit 2. Our classical picture of electrons as just small "bullets" has not been supported by experiment.

It is worth trying to see what happens if we attempt to track the electrons to see through which slit they pass. We can imagine shining a light over each slit and, for a sufficiently small intensity of incident electrons, "see" the electron pass through a slit and then record its location with our detector. The result is that the interference pattern disappears and we get

$$P(x) = |\psi_1(x)|^2 + |\psi_2(x)|^2.$$

This result might be explained by supposing that by shining the lights we have affected in some noticeable way the movement of the electrons. If the electrons were "classical" particles, like bullets, we could simply make our light weak enough so that we disturb the particles by an arbitrarily small amount in the process of seeing which slit they went through. For classical particles this is ok and is compatible with the observed probability distribution. With electrons, however, no matter how weak we make the measuring disturbance we get the classical distribution; but if we do not measure which slit the particles went through we get the new, oscillatory distribution. Thus it is problematic to suppose that the electron actually passes through just one of the slits in the original experiment. It is simply not a good description of microscopic particles, such as electrons, to have them behave as just some small classical bullets. A new, profoundly different way of describing (indeed, defining) microscopic particles is needed. Such a description is provided by quantum mechanics.

Observables and States

Typically, a theoretical description of a physical system involves a mathematical representation of (at least) two basic notions: *observables* and *states.* An observable is something you can measure about a system. For example, a particle's observables include position, momentum, energy, angular momentum, *etc.* The state tells you what you will find when you make a measurement of an observable. Usually the value of a suitable number of observables will uniquely determine the outcome of all other observables. So a measurement of these observables will suffice to determine the state. For example, if we know that a harmonic oscillator has zero energy, we know that its position will be at an equilibrium point, its momentum will vanish, *etc.* More generally, once the values of position and momentum of a particle are known all other observables are determined.* We know that the state of a system has changed when at least one observable has changed its value. We shall spend the bulk of this class seeing precisely how quantum mechanics represents states and observables.

In quantum mechanics, knowing the state of a system is equivalent to knowing the probability distribution for measurements of all observables. Sometimes this notion of

* This is the meaning of the term "particle".

state yields relatively simple information: "the energy of a hydrogen atom in its ground state is -13.6 eV". Here it is tacit that the probability for finding this result is one and that all other energies have probability zero. Not all states have such straightforward interpretations. For example, one can have a state in which "the probability of finding the particle on each side of the barrier is 1/2". And it turns out that "if a particle is in a state of statistically certain momentum it has equal probability to be almost anywhere".

The introduction of probability distributions as the essence of a state is unsettling to a classical mechanic who is used to asking: "If the initial position and momentum are given, where will the particle be in 10 seconds?" To be sure, all such classical mechanics questions can be rephrased in terms of probability distributions. For example, the previous question could be rephrased as: "If the system has the given coordinates and momentum with probability one at the initial time, what is the probability distribution for position after 10 seconds have passed?" In classical mechanics the probability distributions are quite simple, *e.g.,* the probability for the particle to be at its classical location at any given time (as specified by Newton's second law) is one, and zero otherwise. This is why it is cumbersome to use probabilities to do classical mechanics, and we usually don't do that.* In quantum mechanics it is impossible to have a state in which position and momentum of a particle are known with probability unity. Moreover, quantum mechanics will not, in general, lead to simple probability distributions (1 or 0) for observables upon time evolution even if the observables had statistically certain values at some initial time. For such reasons the probability description of states and observables is mandatory.

If you are experienced with statistical mechanics, then viewing the state of a system using probability distributions is probably familiar and reasonable. It is natural to ask: Is quantum mechanics just some kind of statistical mechanics? It seems the answer is "no" for a couple of reasons. Firstly, as briefly touched on earlier, probabilities in quantum mechanics are combined by adding probability "amplitudes" (wave functions) then taking the modulus-squared; this intermediate step allows for interference phenomena that cannot be obtained using probabilities as one does in statistical mechanics. Secondly, in statistical mechanics we are accustomed to thinking that the value of an observable for a system in a given state is part of the "reality" of that system even if we choose to settle for a statistical description for an ensemble of many such systems. So far as anyone knows, nature is described by quantum mechanics, and quantum mechanics uses a probabilistic notion of "state" as the complete description of any single system so that the "reality" of observable features of a system is now a much more subtle thing. Hopefully we will have time to discuss some of this in detail after we have acquired the necessary tools.

Has all of this introductory discussion left you feeling a little lost? That's OK; if it all

* Or course, when the system is at a non-zero temperature it is very useful to work with probabilities, as in statistical mechanics.

were clear to you then you wouldn't need this course!

We shall now turn to a more formal development of quantum mechanics, in which we will terminate the vague, qualitative discussions of the theory and try to figure out how to "do" it, quantitatively speaking. For quite a while we will for simplicity just consider a single particle in a universe which is one-dimensional. Later we will amend this oversimplification.