Principles of Quantum Mechanics

Quantum mechanics is the most successful quantitative theory ever produced. Not a single one of the untold thousands of experiments done to test it has ever found the basic principles to be in error, and the agreement can sometimes go to ten significant figures (as in some predictions of quantum electrodynamics). Quantum mechanics also underlies vast realms of science, from physics to chemistry to some aspects of biology (probably). All that we have studied so far in this course is, to some extent, an approximation to the fundamental quantum physics.

However, in studying quantum we run into unique difficulties. When one learns classical mechanics, it is possible to guide the intuition by thinking of personal experience with how objects fall, springs bounce, or whatever. Quantum mechanics, in contrast, is most important in a realm with which we have no direct experience. A table seems solid to us, not composed of fuzzy electron orbitals. As a result, we have to build up our intuition in a different way: we study the mathematics of quantum, knowing that this mathematics matches experiments incredibly well, then we gradually get a feel for how that math works in different situations. We therefore build up intuition from the math, rather than building it from real-world experience and then seeing how mathematics corresponds to that experience. Nonetheless, just as in more familiar classical physics, there are guiding principles and ways to do order of magnitude estimates. We'll talk about those, but first let's discuss briefly the historical motivation for quantum mechanics.

The nineteenth century produced the greatest triumphs for classical physics: the discovery of Neptune and Sirius B using Newton's laws, the development of electromagnetism, and the laws of thermodynamics. However, it also revealed inklings of puzzling experimental results that would ultimately lead to the "new physics" of quantum mechanics and relativity. One example has to do with the specific heats of gases, and was recognized by Maxwell. According to classical theory, every degree of freedom of a molecule has an energy $\frac{1}{2}kT$ associated with it; this includes motion in all three directions, and vibration and rotation in all possible directions for polyatomic molecules. In principle, each separate electron and nucleon would have their own degrees of freedom, but electrons and nucleons weren't known in the mid-1800s. The problem was that experimental results made it seem as if somehow at different temperatures some of these degrees of freedom were suppressed, a result that could not be explained by classical theory.

A better-known example has to do with blackbodies. A perfect absorber radiates with a spectrum that depends only on its temperature. However, the spectrum itself could not be derived from classical theory, which predicted an "ultraviolet catastrophe", i.e., that the intensity as a function of frequency should rise without bound as the frequency increased. Max Planck was able to derive the correct distribution with the assumption that light isn't continuous but instead comes in discrete packets, or photons. At the time, Planck considered this just a calculational device, but in 1905 Albert Einstein showed that a phenomenon known as the photoelectric effect could be explained if light really did exist as photons. This was the beginning of quantum mechanics, where "quantum" implies a discrete entity. As time went on, more phenomena became explainable under the assumption that nature was sometimes discrete rather than continuous. For example, it was shown (by Thomson and Rutherford, among others) that atoms are negative charges (electrons) around positive charges (nuclei, although originally it was thought to be just protons). Classical theory says that any charge that is accelerated will radiate, so the prediction would have been that atoms would decay in a burst of radiation within a matter of nanoseconds! It was also discovered that the spectra of hydrogen and other elements gave well-defined lines at specific frequencies, rather than a continuum. Both the lines and the stability of atoms were explained by Bohr with the ad hoc assumption that the electrons in atoms could only have certain specific values of angular momentum.

All of this came to a head in the 1920s. Louis de Broglie proposed (and it was later confirmed by experiment) that particles such as electrons act as waves in certain ways. The full theoretical structure for this was presented by Schrödinger and Heisenberg independently, with two different methods that were shown to be equivalent. The initial theory was nonrelativistic, but subsequent generalizations have been able to match experiments to incredible precision.

From the point of aesthetics or a warm fuzzy feeling of intuition, I have to admit that the cure may seem worse than the disease! However, it is important to remember that quantum *works*, and works better than anything else that has ever been proposed. We'll finish this lecture with a little bit about the philosophy of quantum mechanics, since that occupies the attention of a lot of people who think about quantum, but fundamentally this is a quantitatively successful theory.

Before getting into the basic axioms of quantum, let's address the single most important, and single most counterintuitive, aspect of quantum mechanics: the uncertainty principle. In classical physics we have no difficulty at all thinking of an exact, known position for a particle and simultaneously imagining an exact, known momentum for that same particle. Quantum mechanics says that this is impossible. Why? Ultimately, it is because measuring one changes the other by an amount that can't be known with complete precision. This statement, like some statements about thermodynamics (e.g., that you can't extract work from a system in thermodynamic equilibrium), is one that is not obvious, so it will help to consider several examples.

For our first example, let's consider a small ball. We'd like to measure its position. We do this by shining light on the ball, thus taking a snapshot of the ball and determining where it is. We know that we cannot measure the position more accurately than the wavelength λ

of the light. We also know that by shining the light, we are exerting radiation force on the ball, thus changing its momentum. Therefore, the best we can do is to shine just a single photon on the ball, to get its position while affecting the momentum as little as possible. However, we don't know whether the photon hits head-on or a glancing blow, therefore there will be some momentum imparted, in an unknown direction. The momentum of a photon of wavelength λ is $p = h/\lambda$, where $h = 6.63 \times 10^{-27}$ erg s is Planck's constant. Therefore, we can't know the new momentum to much greater than that precision. Of course, we could know the position and momentum to *worse* accuracy. Note that shorter wavelength λ will give you more information about the position of the particle, but less information about the new momentum of the particle. The product of the uncertainties for this problem therefore satisfies the inequality

$$\Delta x \Delta p \gtrsim \lambda(h/\lambda) = h \ . \tag{1}$$

Let's try another example, from the Feynman Lectures on Physics, volume 3. Suppose you want to know the position and momentum of a photon to high accuracy. You build up an apparatus which simply consists of a screen with a hole in it. Let's simplify to two dimensions, and say that the hole has size d and is oriented along the y axis. The screen is opaque but very thin. A light source an extremely long way away emits photons, so we know that the photons are all traveling in the x direction when they arrive at the screen. Behind the screen, at some distance, is a detector, oriented along the y axis.

With this setup, we know that any photon that makes it to the detector was localized to within a distance d. However, we also know that light suffers diffraction when it goes through a hole; the typical angle by which it is deflected is of order λ/d. If this is a small angle then the x momentum is still pretty well defined, but the y momentum (initially zero) becomes of order λ/d times the x momentum. This is the uncertainty in the momentum. Therefore, we have

$$\Delta x \Delta p \sim d\Delta p_y \sim d(h/\lambda)(\lambda/d) \sim h \ . \tag{2}$$

Once again, the product of the uncertainties is of order Planck's constant.

But wait! If you think hard about this problem, some difficulties may appear. For example, of the photons that get to the detector, we know that they passed through the hole (and thus were localized to within a distance d), but we also know that in order to get to the hole from a source a distance D away the angle relative to the x axis had to be less than d/D, which could in principle be arbitrarily small. Therefore, the momentum of the photon is known to order d/D, meaning that we can determine the position and momentum of the photon just before it got to the detector with essentially unlimited precision. Similarly, when the photon hits the detector, knowing where it hit tells us the angle it had to take from the hole, so once again we can figure out both the position and momentum to a level of precision that violates $\Delta x \Delta p \gtrsim h$.

What's the resolution? The point is that you have to be careful about what the uncertainty principle means. It does *not* forbid the possibility that *after the fact* you can establish the position and momentum something had to have to give the observed results. It restricts only predictions for the future: given your current measurements, what can you say about the position and momentum it will have some short time in the future? This is one of many points about quantum that people sometimes misunderstand. When one does the math carefully, the uncertainty principle says that

$$\Delta x \Delta p \geq \hbar/2 \tag{3}$$

where $\hbar \equiv h/2\pi$.

These two situations are examples of "thought experiments", meaning that we don't actually set up the experimental apparatus but we think about what would happen to get some insight. This is also used a lot in relativity theory. For decades, people have tried to think of ways around the uncertainty principle, but have found that every single real and imagined experiment, when analyzed carefully, leads to the same conclusions. You're welcome to think of experiments yourself, but be warned that if you think you've found a way around the uncertainty principle then you're missing something.

One thing that the uncertainty principle introduces that has caused people endless philosophical agonies is the concept that one cannot always predict with certainty what will happen when one measures a system. Instead, one can only determine the relative probabilities of different measurements.

Formally, this can be phrased in terms of a number of axioms for quantum mechanics. These axioms can be phrased in a mathematically rigorous way; for example, the textbook by Cohen-Tannoudji, Diu, and Laloë is a great introduction in every way. However, we'll just state the gist of a few of the axioms.

- The physical state of a system can be described completely by a "wavefunction" $\psi(t)$ that can be a function of position and time. At a given time and place, the wavefunction is a complex number.
- The probability that a particle described by a wavefunction $\psi(t)$ is at some location is proportional to the squared modulus of ψ, or $\psi^*\psi = |\psi|^2$. For the probability to be well-behaved, this means that the wavefunction must be square-integrable over all space (i.e., integrating its square gives a finite number).
- Every measureable physical quantity is described by an "operator" that acts on the wavefunction. For example, an operator could be ∇, or x, or whatever. For any given operator, there are special functions (called "eigenfunctions") such that applying that operator to the function yields that function times some constant. For example, for the operator d/dx, e^{nx} is an eigenfunction, because $(d/dx)e^{nx} = ne^{nx}$. The constant thus produced (n in this case) is called an "eigenvalue".

- The only possible result of a measurement is an eigenvalue of the corresponding operator. This is what puts the "quantum" in quantum mechanics. The example of d/dx is in this case a bad one, because it can give a continuum of values, but when one considers the operator for the energy of an atom, you get only discrete eigenfunctions and eigenvalues. The measurement must of course give a real number! If the wavefunction is an eigenstate of the operator then the result of the measurement is certain (it's the eigenvalue of that eigenstate), but otherwise the result can't be predicted. However, the probability of the result of a measurement giving a particular eigenvalue is determined precisely.

There are additional axioms, but we'll leave those for later when we talk about the Schrödinger equation. Note that "eigenvalue" and "eigenvector" are used in matrix calculations as well, and in fact Heisenberg's matrix mechanics uses them directly.

In the next class we'll talk about how to use order of magnitude estimates in some quantum processes, and in the following one we'll revisit the axioms and use Schrödinger's equation to do some genuine calculations. However, I'd like to end this class with some comments about the philosophical implications of quantum mechanics.

Ever since quantum mechanics was introduced, philosophers (and more than a few physicists) have worried greatly about its implications. There is a lot of room for intelligent disagreement, but let me offer my perspectives. Fundamentally, I think that the purpose of a scientific theory is to explain and predict the results of experiments and observations. If a theory does that, it's successful. However, it's not satisfying to just apply equations, so people have inevitably looked for deeper intuition about theories. This leads to pictures about the world. Those pictures occasionally change drastically; think about absolute space and time in Newtonian theory being replaced by relativity, or determinism being replaced by uncertainty in quantum theory. What does *not* change, however, is the predictions. What I mean by that is that, in the domain of applicability of a theory, its predictions will always have the same accuracy that they always did, even if a new theory with a different world view comes along.

Let me give a great example from Isaac Asimov's short essay "The Relativity of Wrong". In ancient times, many people believed that overall the Earth is flat. That's not such a bad approximation over short distances; when you look at a map of the Maryland campus you don't worry about the overall curvature of the Earth. Many of the ancient Greeks argued that the Earth had to be a sphere, based on shadows on the Moon and other bits of evidence. That's a better approximation, but notice that if the sphere has a large radius then over a short distance it looks a lot like a plane. As measurements of the Earth got better it became clear that it bulges in the middle, and therefore is closer to an oblate spheroid than a sphere. However, here the deviations are very small indeed. Further refinement indicates that the northern hemisphere is different from the southern, making one think of the shape of a pear. But the point is that each approximation is better, and importantly, *the predictions are more and more accurate.* We do *not* have arbitrary freedom in models. It's not as if in a century we'll think that the Earth is a cube, and in a further hundred years we'll think that it has the shape of a donut!

Why bring this up? It's because regardless of the intuitive or philosophical picture we have about the microscopic world in a hundred years, we can be absolutely confident that the predictions of quantum mechanics will continue to be as valid as they are today, meaning that to the same accuracy that can be measured today, quantum mechanics will agree with experiments. Maybe more accurate experiments will point out discrepancies, but they won't be any larger errors than can be measured now. Thus, to an extent, the philosophical basis of quantum is a red herring that doesn't matter nearly as much as the predictions. Indeed, there are various different "models" for quantum mechanics that lead to the same equations, so they can't be distinguished from each other.

One final comment about this. The unpredictability of quantum mechanics strikes many non-scientists as a profound departure from classical physics. However, chaos in classical systems belies this. Even though in principle one might have imagined a perfectly determined clockwork universe, in reality even the tiniest 10^{-100} level error in one's initial conditions will be magnified to uncertainty in a ridiculously short time. Effectively, the classical universe is nondeterministic as well, so it's a bit silly to hold forth profoundly about how quantum mechanics allows free will or the like!