


I. Experimental Evidence for Quantum Mechanics

Quantum mechanics (QM) is a *tool* that was invented to *explain* experimental results. It is nothing more and nothing less than that. The utility of QM is therefore based entirely upon its ability to predict and explain experimental results, and by this measure it is a phenomenal success. There has yet to be an experiment of any type that violates the basic principles of QM. Thus, to begin with, we should discuss some of the experimental results that illustrate key principles of QM. Since this is a chemistry course, we will slant our perspective towards *chemically relevant* experiments, but similar effects can be found in any situation where the systems are small enough and the temperature is low enough.


a. Polarization of Light

Light waves can be polarized in any direction perpendicular to the direction of motion of the wave. So, for example, if we have a laser propagating in the \bar{z} direction, the light beam can be polarized either along \bar{x} or \bar{y} . In this sense, light can be thought of as a transverse wave (i.e. one whose oscillations are perpendicular to the direction of propagation) and the two polarization directions can be thought of as: These two polarization components can be separated using a




polarization filter. Typically, the filter consists of a crystal composed of rows of aligned molecules. Then, light whose polarization is not aligned with these rows will not pass through the crystal; meanwhile, light whose polarization is aligned with the crystal axis will be able to pass through the gaps between the rows.

Light from typical sources (such as a lamp or the sun) is not polarized; if you pass it through a polarization filter, some of the light passes through, and some does not. We will depict this simple experiment by:



Where the round circle represents a polarization filter, and the vertical lines indicate that it is a polarization filter in the \bar{x} direction. The polarization filter performs a simple **measurement**; it tells us how much of the light is polarized in a given direction.

This measurement is, however, very boring. It gets interesting when we start to consider multiple polarization measurements being applied to one laser beam. For example, if the first filter is \bar{x} while the second filter is \bar{y} , we get no light transmitted:




To put it another way, the first filter *measures* the polarization of the light and tells us that a certain part of the wave is \bar{x} -polarized. Then, the second filter *measures* how much of the resulting \bar{x} -polarized beam is actually \bar{y} -polarized. The obvious result of this experiment is that **none** of the \bar{x} -polarized light is simultaneously \bar{y} -polarized. This makes sense from a physical perspective (none of the \bar{x} -polarized waves would fit between the \bar{y} -oriented slits) and also from

a logical point of view (it is impossible to have the light polarized in two mutually exclusive directions).


On the other hand, if we perform an experiment where the first filter is \bar{x} while the second filter is aligned at a 45° angle to \bar{x} (along a direction \bar{x}') we get do get some transmission:

Expt. 2



In fact, if we have very good filters, we can get 50% of the \bar{x} -polarized light to pass through the \bar{x}' filter. We can rationalize this, as well, because \bar{x}' is half-way between \bar{x} and \bar{y} , so having half of the \bar{x} -polarized beam pass through the \bar{x}' filter makes sense.


Now, we come to the key experiment. Let's take the beam of light produced in Expt. 2 and measure its polarization in the \bar{y} direction:




We find again that 50% of the light passes through the final filter. This may seem benign, but notice that if we remove the \bar{x}' filter, we recover Expt. 1, where **none** of the light passed through the final filter. So, basically, the \bar{x}' filter takes in a beam of light that is 0% polarized in the \bar{y} direction, **removes** some of the light, and the resulting beam is then 50% \bar{y} -polarized! What is going on here?

Could the Filters Be Rotating the Polarization?

In class, it was pointed out that there is one obvious explanation of these experiments: perhaps the polarization filters do not *filter* the light, but instead *rotate* the polarization. This is a credible interpretation of the experiment, and is actually the explanation given in most freshman physics texts. However, we should note two things that make this a very strange rotation. First, different polarizations of light are rotated in *different directions*. For example, if we pass \bar{x} -polarized light through the \bar{x}' filter, the light polarization is 'rotated' $+45^\circ$ to \bar{x}' . On the other hand if we pass \bar{y} -polarized light through the same filter, it again comes out \bar{x}' , a net rotation of -45° . If we want to denote this with a picture, we might note that a normal rotation acts on the two independent polarizations like this:



while the 'rotation' of polarizations looks quite different:




where the double arrow on the right indicates that the two independent polarizations on the left get 'collapsed' into one polarization on the right. The second weird feature of this 'rotation' is that we always get *less* light out after the rotation than we put in, and often the difference is significant. This contradicts our notion that with perfect optics we should be able to get perfect transmission.

So we are left with one of two 'weird' interpretations of the experiment: either we somehow increase the \bar{y} -polarization when we measure the \bar{x}' polarization, or else the filters generate some kind of strange 'rotation' of the polarization.

Quantum Interpretation of Polarization Experiments

This experiment illustrates two key concepts of quantum mechanics. The first is the idea that the order in which we perform experiments matters. Consider the following set-up:



This differs from Expt. 3 only in the order in which the filters are applied (i.e. $\bar{y} \bar{x}' \bar{x}$ versus $\bar{x}' \bar{y} \bar{x}$), but in the one case, we get light transmittance and the other we do not. If the order of two experimental observations does not change the result, the two observations are said to **commute**. The surprise that QM brings is that, in general, experimental observations **do not commute** with one another. In this case, for example, the action of applying the \bar{y} filter does not commute with the use of an \bar{x}' filter.

In describing these polarization experiments, it is often useful to use a shorthand notation. Represent the polarization filter in the \bar{x} direction by \hat{P}_x , and similarly for $\hat{P}_{x'}$ and \hat{P}_y . Further, a string of these symbols placed next to one another can be interpreted as sequential measurements read off, by convention, from right to left. Thus

$$\hat{P}_{x'} \hat{P}_x$$

translates into “first apply an \bar{x} -filter and then apply an \bar{x}' filter”. Using this shorthand, we can denote the inequivalence of Expts. 3 and 3a by the symbolic equation:

$$\hat{P}_y \hat{P}_{x'} \hat{P}_x \neq \hat{P}_{x'} \hat{P}_y \hat{P}_x$$


or, if we strip off the leading factor of \hat{P}_x on both sides:

$$\hat{P}_y \hat{P}_{x'} \neq \hat{P}_{x'} \hat{P}_y$$

A consequence of non-commuting observables is the **uncertainty principle**. If two observations do not commute, it is impossible to predict exactly what the result of each observation will be; by measuring one, you automatically disturb the state of the other, and so one cannot know both values at the same time. In the polarization experiments, for example, we established that the \bar{y} filter and the \bar{x}' filter do not commute. This means it is impossible for us to know the degree of polarization of the light in both directions simultaneously. This may seem trivial in the present case (after all, how could light be polarized in two different directions at once?) but we will see that the same effect appears again and again.

b. Single Molecule Fluorescence

Suppose we want to measure the properties of individual sodium atoms. To do this, we can begin with sodium vapor at high pressure, and allow the gas to expand through a small nozzle to an area of lower pressure (supersonic expansion). The expanding gas can be collimated by placing a screen a distance away from the nozzle. This results in a very cold, very dilute beam of atoms.




Now, sodium has an important excited state (the 'D line' state) and we are interested in measuring its **lifetime**. That is, we are interested in determining how long it takes an excited sodium atom to emit a photon:



We know from experience that fluorescence typically occurs on the nanosecond time scale, so our measurements need to be **fast**. In order to accomplish this task, a team of graduate students working round the clock for three years designs a tuneable ultrafast laser that operates in the region of spectral interest (in this case, about **589** nm)

with picosecond time resolution. They also design a very accurate detector with similar time resolution in the same spectral range. Using these two devices, we can measure the lifetime of the sodium 'D line' state by hitting the molecular beam with a laser pulse at one point along the beamline at time t_1 and placing a detector further down the line to register any sodium fluorescence at a later time t_2 :

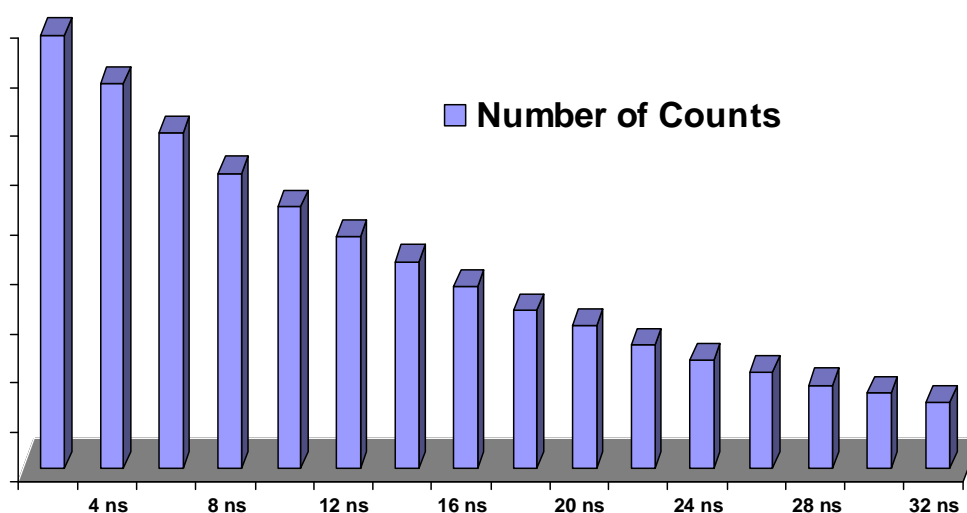


The lifetime is then determined by the time delay between excitation and emission ($t_2 - t_1$). Now, even a very accurate detector misses many photons; further, very often, the pulse will not succeed in exciting a sodium atom, because sometimes there won't be any atoms in the path of the laser beam (recall that the molecular beam is very dilute). Hence, we will have to try this many, many times before we get a successful result. But when we do, we can be very confident that the count came from the fluorescence decay (assuming we are careful to isolate our experiment from outside light sources).

So, we do this experiment several thousand times (we can fire the laser every microsecond, so this doesn't take as long as you might think) and eventually we register a successful count at the detector, and determine that the lifetime is 12.554(2) ns, where the uncertainty results from our picosecond time resolution in the excitation and detection. Just to check our results, we run the experiment again. This time we register a lifetime of 8.492(2) ns. If we run the experiment a third time, we register a lifetime of 22.100(2) ns. Again, we have been very careful to ensure that we have ps time resolution, and the initial states of the sodium atoms are **absolutely identical** in the repeated experiments. What could be going on here?

Quantum Interpretation of Fluorescence Experiments

These experiments illustrate another tenet of QM; all observations are by their nature **probabilistic**: one cannot in general *predict* the precise outcome of an experiment. The best one can do is predict the *probability* that one outcome will be observed as compared to another. In fluorescence experiments above, we found that decay times around 10-20 ns are very common, while a decay time of, say 2 hours is very uncommon. If we repeat the experiment many, many times and make a histogram plot of the number of counts that appear after a given elapsed time (t_2-t_1), we obtain a probability distribution




that looks like:


This tells us how likely it is that the molecule will fluoresce within a given period of time. The shocking thing is that this is absolutely the best one can do in terms of predicting the fluorescence time accurately. This limitation is closely related to the **uncertainty principle** discussed above. The best one can hope for in practice is a very narrow probability distribution (or, equivalently, a small uncertainty). It is important to recognize that this is not due to some weakness of our experiment, nor to any inhomogeneity of our sample. Instead, this arises from a very deep limitation on what we can know about a quantum system. If we are interested in knowing whether the system is in a given **state**, α , the best one can hope for is a way to predict the probability, P_α , that we will find the system in that state. If we perform the measurement of (t_2-t_1) once, we cannot predict the result; but if we perform the experiment many, many times and keep track of the individual times, we will **always** obtain the distribution above. In this sense our experiments are reproducible.

c. Nanoscale Electrical Conduction


The experiments described in this section are inspired by some work done recently in Bob Westervelt's group at Harvard (<http://meso.deas.harvard.edu/>). The details of the experiments are really quite challenging, but the general set-up is quite intuitive. They study electron motion in a thin layer of semiconductor sandwiched between two insulating layers – a so-called 2 dimensional electron gas (2DEG).




By attaching a metallic lead to the outer oxide layer, they create a non-invasive but localized source of electrons within the semiconductor:



They can then use lithographic techniques to deposit negatively charged dopants on the surface. Since electrons are also negatively charged and like charges repel one another, these deposited features act like huge repulsive barriers within the semiconductor layer. A simple configuration might consist of a 'wall' with a small gap in it:




Now, we're interested in being able to tell how the electrons emitted by the source behave after they pass through the gap. This is a very difficult question, since it is difficult to conceive of a microscope small enough to 'see' an individual electron. The elegance of their experiment is that they can measure the electron's location with an atomic force microscope (AFM) *even though the AFM tip is much larger than an electron!* So the final set up looks like this:




We're going to look at the results one gets with a bunch of different barrier configurations. In this case, it is a lot of work to re-draw the above picture every time, so we will resort to the following simplified cartoon:

Expt.1




In this case, we expect a rather simple result: we should detect electrons at the detector only in the region of space that was not blocked by the barrier. This is essentially right. In practice the observed electron distribution is slightly smeared out (due, in part, to the uncertainty principle):

Expt.1




Now, we consider a second experiment where we put **two** small holes in the barrier:

Expt.2




In this case, one might expect the distribution of electrons at the detector to look something like:

Expt.2




However, this is not what is seen in an experiment. Instead, the distribution of the electrons that arrive at the detector after passing through the two slits is:


Expt.2



This is somewhat surprising, but it is similar to something we have seen before: the **interference** of two waves. If we add up two equal waves where one is shifted in phase by an amount δ relative to the other, we find that the waves add in a non-intuitive way:

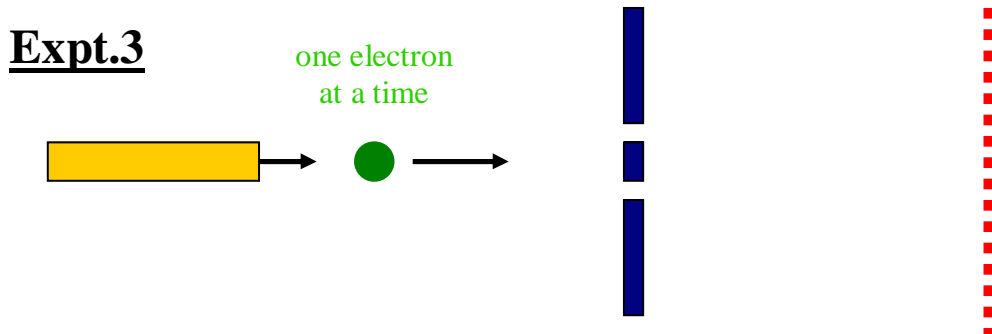


We see that for this relatively small value of δ the amplitude of the resulting wave is much larger than either wave by itself: **constructive** interference. On the other hand, if we have $\delta = \pi$ we get completely **destructive** interference:




Thus, the somewhat surprising conclusion of this experiment is that electrons can behave as waves. This is an illustration of the **wave-particle duality**: things we normally think of as particles (like electrons) sometimes behave like waves, while things we normally think of as waves (like light) sometimes behave as particles.

So, then the question arises: what is the electron really doing when it encounters the barrier? When is it a particle and when is it a wave? To address this, let us slow down the experiment above. Instead of a continuous stream of electrons, let's make the source emit at most one electron at a time:




Now, according to the **probabilistic** nature of QM, we do not expect to be able to predict exactly where this one electron will hit the detector. However, we *can* predict the distribution of electrons; that is, the result expected if one performs the experiment many, many times and tabulates the results. If we do this in the “slowed down” experiment, we find:



Thus, the form of the distribution does not depend on how many electrons we send at once. This may not seem surprising; if the electron leaves the source as a wave and then passes through the two slits as a wave, we expect the interference pattern to show up through the interference of different parts of this wave.

There is, however, one huge wrench that QM throws in this interpretation. Suppose we put two detectors in the two holes in the barrier:

Expt. 4

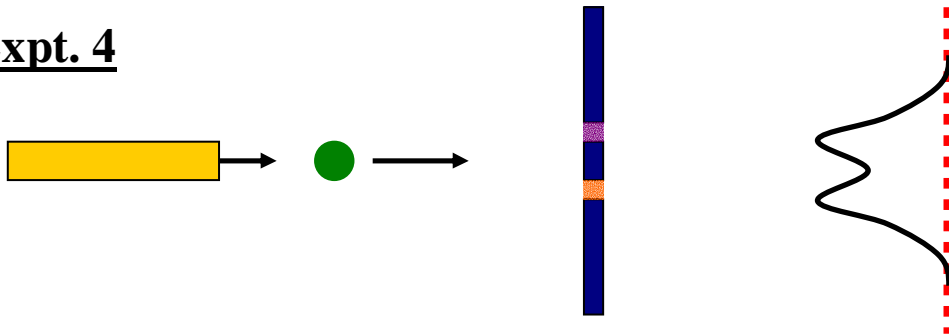


These detectors will tell us if the electron goes through the upper hole, the lower hole or both. This is actually very difficult to accomplish within the Westervelt experiments. However, other experiments have been constructed that perform essentially these measurements. For our purposes, it is enough that it can be done. The detectors do not deflect the electron any significant amount; they merely tell us which path(s) it takes by flashing (purple for the upper slit, orange for the lower). So, we can perform this experiment many times and we find:

	upper	lower	both
Run 1	X		
Run 2		X	
Run 3		X	
Run 4		X	
Run 5	X		
Run 6		X	
.....			

If we add it up, we find the purple light turns on 50% of the time, and the orange light turns on 50% of the time, but they are never on simultaneously. The conclusion of this is clear: the electron always passes through one slit or the other, but never through both at the same time. This is consistent with our understanding of the behavior of a particle: it always exists in a definite place, and never in two places at once. However, this poses a rather serious difficulty for Expt. 3: each electron is passing through one slit or the other, but in order to generate an interference pattern, it must somehow “know” that the other slit is there. This conundrum is complicated by the fact that the transmitted particle distribution in Expt. 4 is different from that in Expt. 3:

Expt. 4



Thus, the interference pattern disappears when we detect which slit the electron passed through!

Quantum Interpretation of Two Slit Experiment

First, we address the point of why the interference pattern is present in Expt. 3 but absent in Expt. 4. As was mentioned above, the detectors are sufficiently non-invasive that we are confident they do not disturb the path of the electron. The reason the results are different is that observing which slit the particle goes through **does not commute** with the observation of its position at the detector. Hence, no matter how delicate the detectors at the slits are, they will influence the outcome of the second experiment. This is very similar to our experiments inserting an \bar{x}' polarization filter drastically alters the degree of \bar{y} polarization in the light.

The other difficulty is understanding how the electron can generate an interference pattern in Expt. 3 when it is apparently only passing through one of the slits. The trick here is that we don't actually know that it doesn't pass through both slits in Expt. 3; Expt. 4 tells us with certainty that the electron is not passing through both slits at once, but in making this measurement, we affect the outcome of the experiment! Thus, in practical experiments, there is never a contradiction: one cannot know that it passed through one slit or the other and simultaneously observe an interference pattern. We can describe Expt. 3 and Expt. 4 individually; the challenge is to describe them together in a consistent manner.

QM accomplishes this by inventing the concept of a **superposition** of states. A **superposition** of states exists in any system that has multiple physically accessible states (e.g. particle goes through upper slit and particle goes through lower slit) and no measurement has

been made to tell us which one is actually realized. When a measurement is made, the **superposition** collapses to one of the observable states. Thus, in Expt. 3, the electron is said to be in a **superposition** of two states while it is between the source and the detector; one state goes through the upper slit, and the other goes through the lower slit. The two states can **interfere** with one another, and this gives rise to the **interference** pattern. In Expt. 4 we destroy the superposition by making a measurement. The measurement always gives us one of the two superposed states, but never both. After the **superposition** collapses to one state, the interference between states is destroyed.

This is popularly recounted in the Schrödinger's cat paradox. Suppose you have a sick cat and you put it in a box for a long time, and do not look at for a long time. Then according to quantum mechanics, it exists in a **superposition** of two states: "cat is dead" and "cat is alive". Once you open the box, however, it will either be dead or alive – the **superposition** will have collapsed to a single state. This is a bit of an unrealistic experiment, but it does illustrate one key point: Schrödinger did not like cats. This also illustrates the point that QM assumes no knowledge of anything that is not experimentally observable, because observation has the potential to change the outcome of later experiments. In between experiments our uncertainty about how the system evolves factors very heavily in how we make predictions about future experiments.