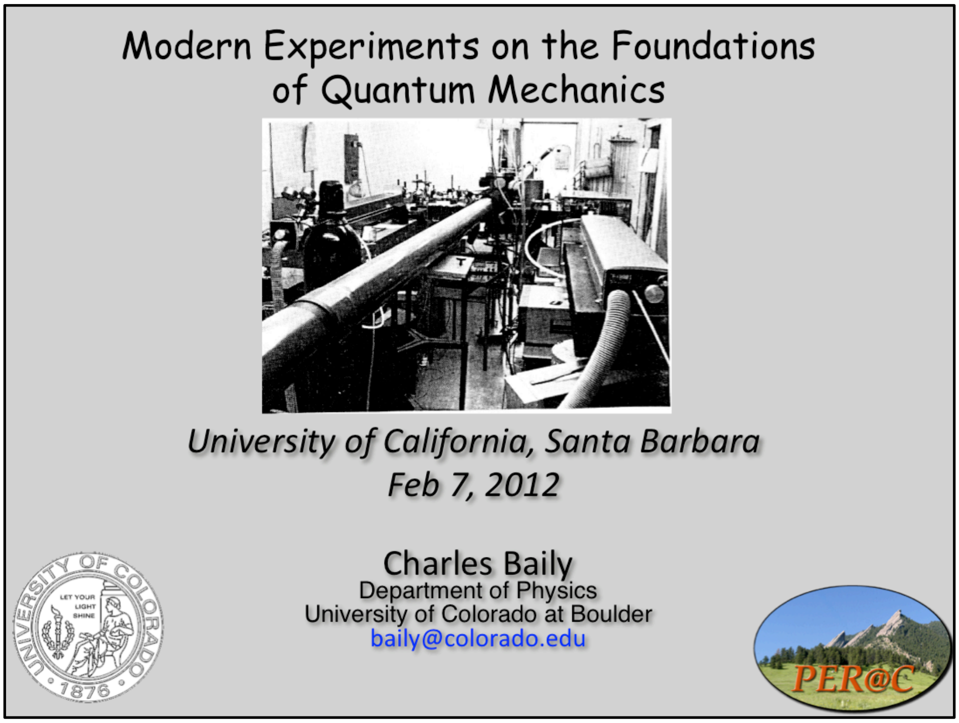

It may be true that at times the behavior of quantum objects (photons or electrons) may seem strange and mysterious to us, but the theory of quantum mechanics itself should not be a mystery. It has rules, and understanding quantum mechanics properly requires us to understand what those rules are, and when to apply them.

Quantum States

$$|\Psi_1\rangle = |E_1, p_1, \ldots\rangle \qquad |\Psi_2\rangle = |E_2, p_2, \ldots\rangle$$

$$|\Psi\rangle = a_1|\Psi_1\rangle + a_2|\Psi_2\rangle + \ldots$$

$$P_1 = |a_1|^2 \qquad P_2 = |a_2|^2 \quad \text{etc}\ldots$$

What does this superposition of states mean?

- Statistics for an ensemble of many identical systems?
- Actual state of each individual system?

We generally describe some quantum state with the letter psi, and in this talk we'll use a bracket notation, where we just place inside the bracket whatever information we have that distinguishes that state, such as its energy, momentum, etc…

Quantum mechanics is a probabilistic theory, meaning there may be different outcomes for measurements performed on identically prepared systems. When this is the case, we describe the quantum state as a linear **superposition** of the possible final states, weighted by their relative probabilities for occurring.

But what does it mean to construct a superposition state? Do these probabilities only have meaning when describing the statistics for a large ensemble of identical systems, or does it describe the true state of each individual system?

This is a question that physicists have debated since the advent of quantum mechanics, and for many decades, this question seemed more like a matter for philosophers, because physicists didn't have the ability to work with a single atom, or a single photon or electron – they could only work with large numbers of them at a time.

Double-Slit Experiment

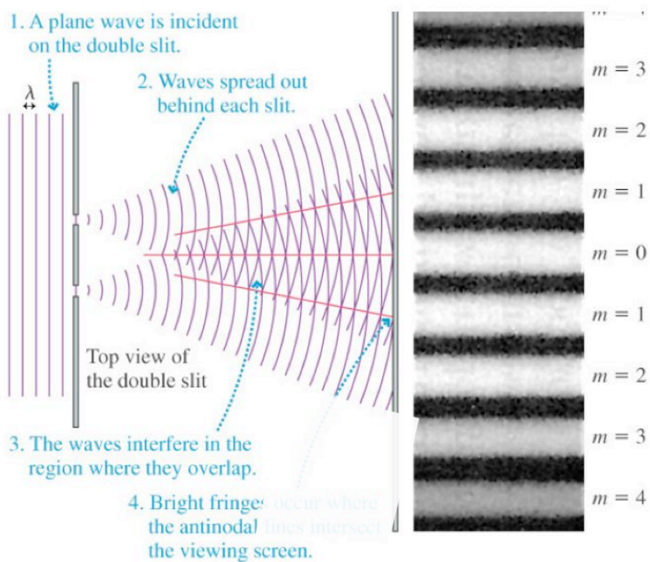

To illustrate the nature of the debate, let's consider a typical double-slit experiment with electrons. As it's usually described in introductory textbooks, we represent a beam of electrons with a wave function – where the wavelength is given by the deBroglie relation.

The beam of electrons is incident on the double-slit, and each slit acts as a coherent wave source. The waves interfere in the regions of overlap, and we see at the detecting screen a brightness where the waves constructively interfere, and dark areas where there is destructive interference. We call this an interference pattern.

But what's really interesting is that we can, instead of a beam of many electrons, send through one electron at a time. We see the electrons arrive at the detector one at a time, in what at first appears to be random locations. However, after many electrons have been detected, we see the same interference pattern as before.

Double-Slit Experiment with Single Electrons (1989)

http://www.hitachi.com/rd/research/em/doubleslit.html

A. Tonomura, J. Endo, T. Matsuda, T. Kawasaki and H. Ezawa, "Demonstration of Single-Electron Buildup of an Interference Pattern,“ *Amer. J. Phys.* **57**, 117 (1989).

this experiment **has** been performed – it was done by a group at the research and development division of Hitachi in 1989, and they created a movie we can watch so that we can see for ourselves that this isn't just a "thought experiment".

What we're seeing here is the arrival of individual electrons at the detecting screen. We see them arrive one at a time, so there can't be any question that essentially only one electron is passing through the apparatus at a time.

The actual length of the movie is about 20 minutes, but they've provided some time lapse for us, so that we can see in a relatively short time that, after many electrons have been detected, we indeed see the same kind of interference pattern as when we were shooting through many electrons at once.

Interpretation

Statistical: Each electron passes through one slit or the other, but we can't know which one without disrupting the interference pattern. (particle-like)

Quantum Wave: Each electron passes through both slits as a delocalized wave, interferes with itself, and then collapses to a point when interacting with the detector.

Agnostic: We can't know what each electron is doing between emission and detection.

How are we to interpret this experiment? For many years, different physicists have had different opinions about what's happening with each electron between emission and detection.

In what we'll call a "statistical interpretation" – it's decided that each electron (being a localized particle) passes through one slit or the other, but any time we try to determine which slit it passed through, we change the experiment and disrupt the interference pattern.

If we believe in a wave description of individual electrons, then we must conclude that each electron is passing through **both** slits, interfering with itself as a delocalized wave on its way to the detector, yet we still detect them at a single point in space when they arrive at the screen.

Another popular stance is to say that, in science, we can't talk about things that can't be observed or measured. And for many people, there didn't seem to be a clear answer to this kind of question, so why worry about it when we can apply the theory to calculate what we need?

This stance is embodied in a popular phrase, usually attributed to Richard Feynman. In fact, Feynman discusses this experiment in his famous lectures, where he's clear that, in the 1960's, this really was just a thought experiment

For the longest time, the closest people could come to performing this experiment was to pass an electron beam through a crystal lattice, which is in effect a "many-slit" experiment, and not really the one we've been describing.

Double-Slit Experiment with Electrons (2008)

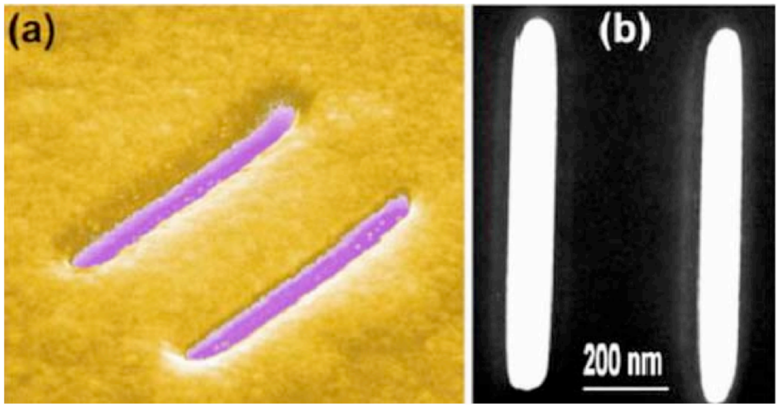

- SEM image of a nanometer scale double-slit system created using gold foil and a focused ion beam
- Slits are 83 nm wide and spaced 420 nm apart

"Nanofabrication and the realization of Feynman's two-slit experiment"
S. Frabboni, G. C. Gazzadi, and G. Pozzi, *App. Phys. Letters* **93**, 073108 (2008).

This experiment has indeed been done, in 2008. What we see here are Scanning Electron Microscope images of two slits created using a focused ion beam and gold foil. The slits are 83 nm wide, and spaced 420 nm apart.

Double-Slit Experiment with Electrons (2008)

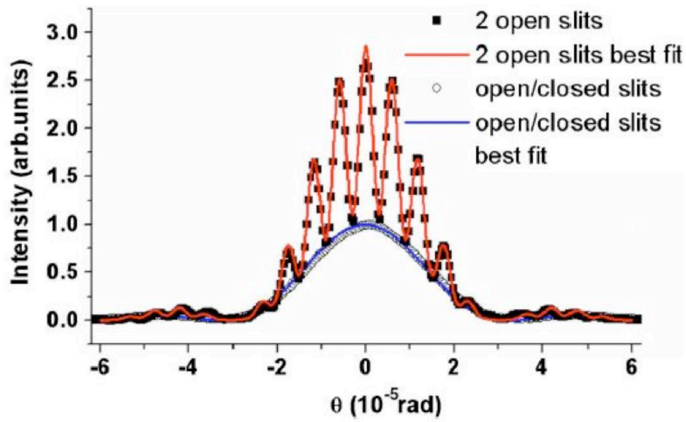

And when the experiment is performed with both slits open, we see exactly the kind of interference pattern we expect.

I've told you this story about electrons in order to illustrate the nature of the problem at hand, but we really haven't given definitive evidence regarding our original question about superposition states, so let's consider the same issue with respect to photons.

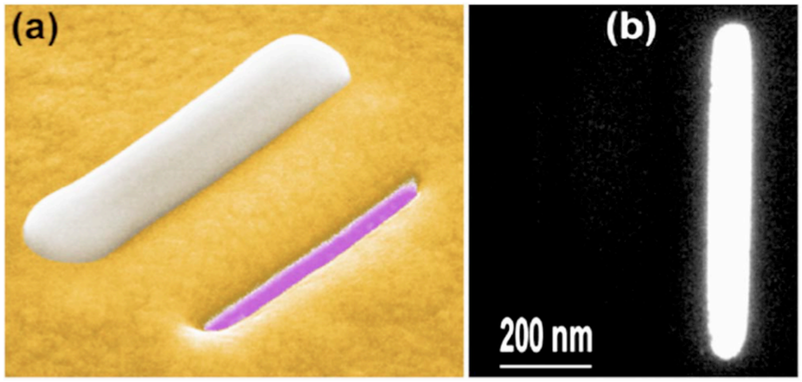

And just to double check that they're behaving as we've described, we can cover up one of the slits, and we see that there's an intensity maximum directly across from the open slit. No interference can be seen when there's just one path from source to detector, meaning there's no opportunity for electrons to interfere with themselves.

I've told you this story about electrons in order to illustrate the nature of the problem at hand, but we really haven't given definitive evidence regarding our original question about superposition states, so let's consider the same issue with respect to photons.

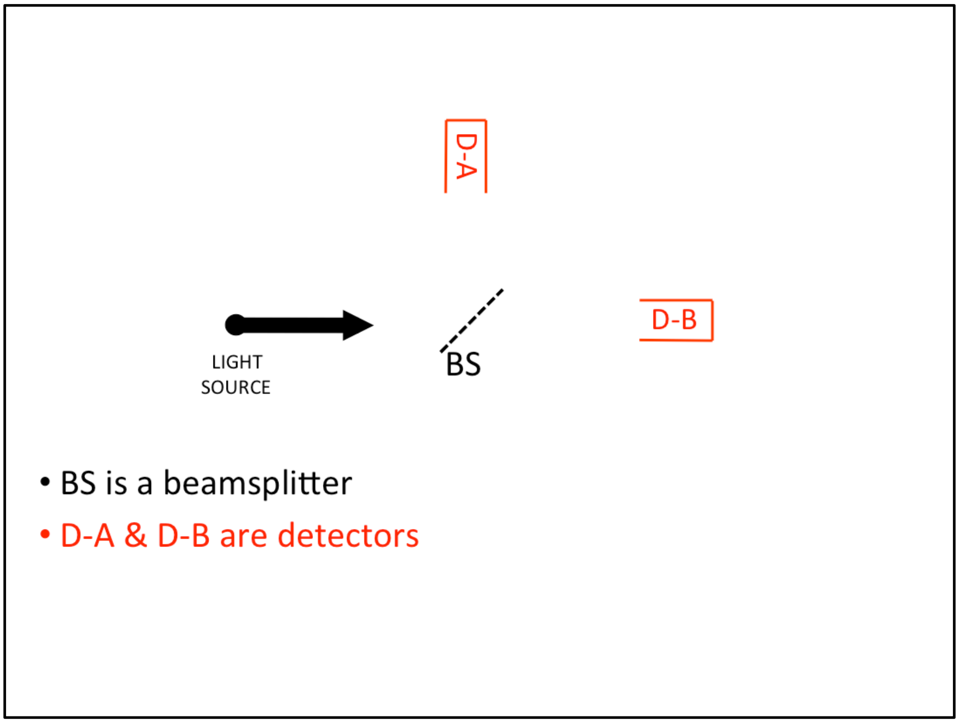

Quantum theory tells us that a typical light source can be described as containing a large number of individual photons, each carrying a specific amount of energy according to its wavelength

Let's consider this simple experimental setup. We have some kind of light source…

…and we'll aim the light source at a beamsplitter, which is really nothing more than a half-silvered mirror – a "two-way" mirror that both reflects and transmits light.

Behind the beamsplitter, we'll place two detectors in order to measure the intensity of the light that is transmitted and of the light that gets reflected.

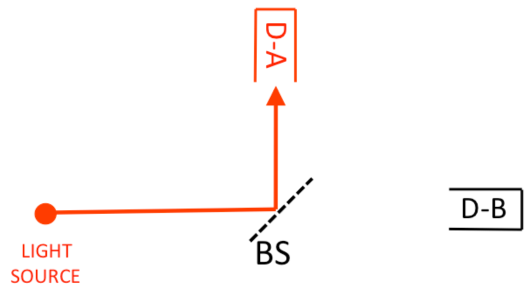

Two possibilities:

1. Photons are **reflected** at the beamsplitter

We could have photons reflected at the beamsplitter, which are then detected by Detector A…

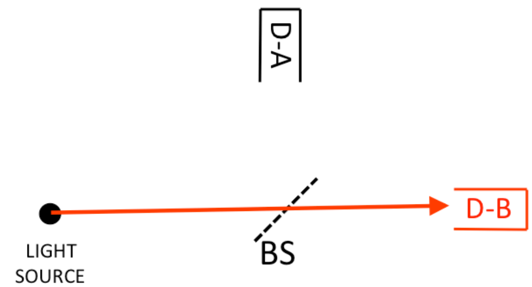

Two possibilities:

1. Photons are **reflected** at the beamsplitter
2. Photons are **transmitted** at the beamsplitter

…or they get transmitted by the beamsplitter, and are detected by Detector B

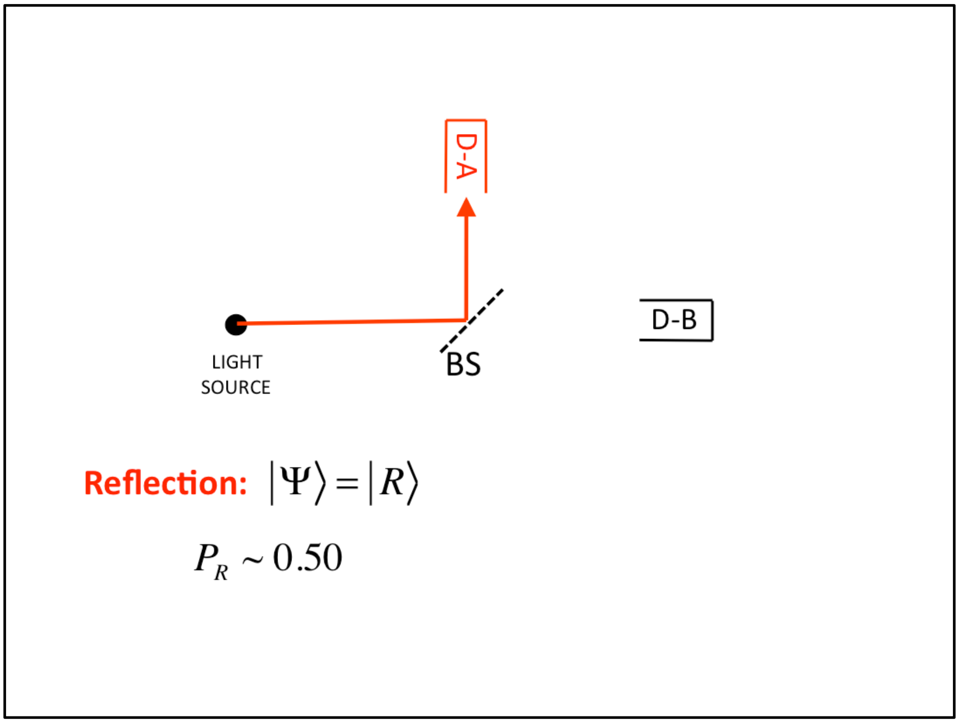

Let's represent the reflected state with an 'R', and when we look at the intensity of the light detected in D-A, we find that 50% of the light's intensity gets reflected at the beamsplitter.

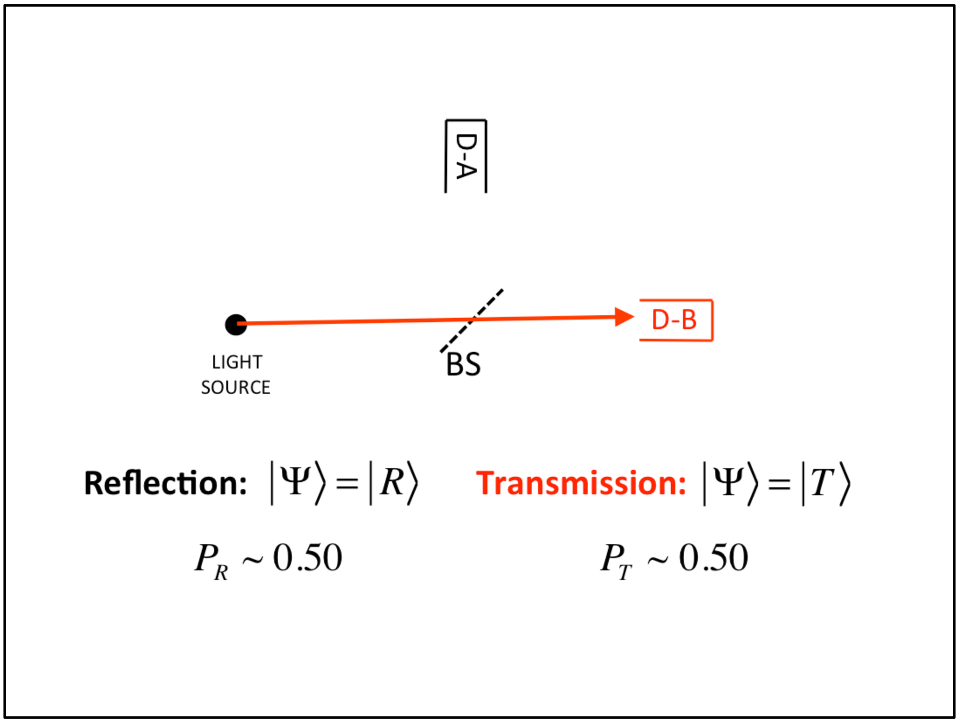

And similarly for the transmitted state, which we'll represent with a 'T'. The probability for transmission is also 50%. It must be, because the sum of the two probabilities must be unity – there are only two possibilities.

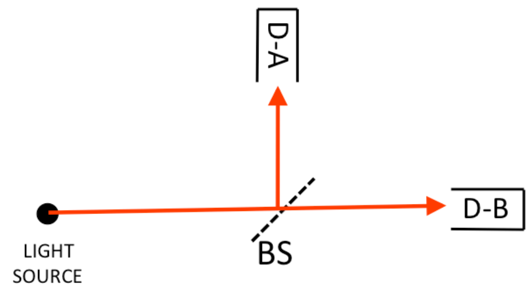

Reflection: $|\Psi\rangle = |R\rangle$ **Transmission:** $|\Psi\rangle = |T\rangle$

$$|\Psi\rangle = \frac{1}{\sqrt{2}}|R\rangle + \frac{1}{\sqrt{2}}|T\rangle$$

We can't know ahead of time whether any one photon will be reflected or transmitted, and quantum mechanics tells us that we must represent the ensemble of photons as a 50/50 superposition of being reflected or transmitted.

Interpretation

Statistical: Each photon is **either** reflected **or** transmitted at the beamsplitter (but not both). The superposition state represents our ignorance of its true state.

Quantum Wave: Each photon is **both** reflected **and** transmitted. The superposition state represents the true state of each photon after encountering the beamsplitter.

A statistical interpretation would say that each photon is either reflected or transmitted, and that the superposition state is a reflection of our ignorance of what each photon is going to do.

And just like with the electrons, in what we'll call a "quantum wave" interpretation, we'll say that each photon is both reflected and transmitted, and that the superposition state is the actual state of each photon.

If we're working with many, many photons at once, there's really no way to decide which explanation is correct, since they both predict the same kind of experimental outcome.

But what does it mean for a single photon to be **both** reflected and transmitted? Does that mean that half the photon energy goes one way and half the other way?

"The result of [the detection] must be either the whole photon or nothing at all. Thus the photon must change suddenly from being partly in one beam and partly in the other to being entirely in one of the beams."

P. A. M. Dirac, *The Principles of Quantum Mechanics* (1930)

Not according to Paul Dirac. Here's a quote from his famous textbook on quantum mechanics, which first appeared in 1930, long before any experiment could be performed with individual photons. He said:

"The result of [the detection] must be either the whole photon or nothing at all. Thus the photon must change suddenly from being partly in one beam and partly in the other to entirely in one of the beams."

Now, I'm not entirely certain how much he believed this to be literally true, but in the end it's not important. In his mind, physicists create models (or "pictures" in their heads) of what's going on behind the mathematics, and what was important was whether this kind of "picture" helped to make obvious the self-consistency of the fundamental laws of physics.

CALCIUM CASCADE PHOTON SOURCE

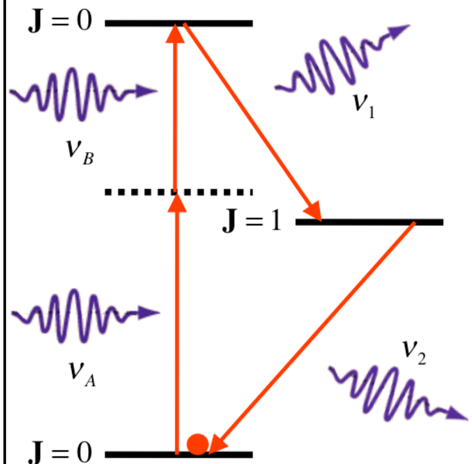

- Calcium atoms are excited by a two-photon absorption process (E_A = 3.05 eV) + (E_B = 2.13 eV)
- The excited state decays first by single photon emission (E_1 = 2.25 eV)
- The lifetime of the intermediate state is $\tau \sim 4.7$ ns
- High probability a second photon (E_2 = 2.93 eV) is emitted within $t = 2\tau$

Fast-forward 50 years, when Alain Aspect and his colleagues devised a way to actually perform this kind of experiment. What they realized was that two-photon radiative cascade from calcium-40 atoms, which they'd developed for a very different kind of photon experiment in 1981 would also be suitable for the kind of experiment we have in mind.

In this case, we're interested in some very specific energy levels for calcium 40 atoms. In its ground state, calcium 40 has two valence electrons outside a closed shell with their spins oppositely aligned, so that the total angular momentum (spin plus orbital) of this state is J = 0. The upper level of the cascade is also a J = 0 state, and the intermediate state is J = 1.

The calcium atoms are excited by a two-photon absorption process, delivered by two laser pulses, each tuned to match the energy between the relevant levels. From the upper level, the calcium atom first decays by emitting a single photon, which we'll refer to as nu-1.

The lifetime of the intermediate state is about τ = 5 ns, so there's a high probability that a second photon has been emitted within twice that time.

Why two-photon absorption? Well, one reason is that we can't get the valence

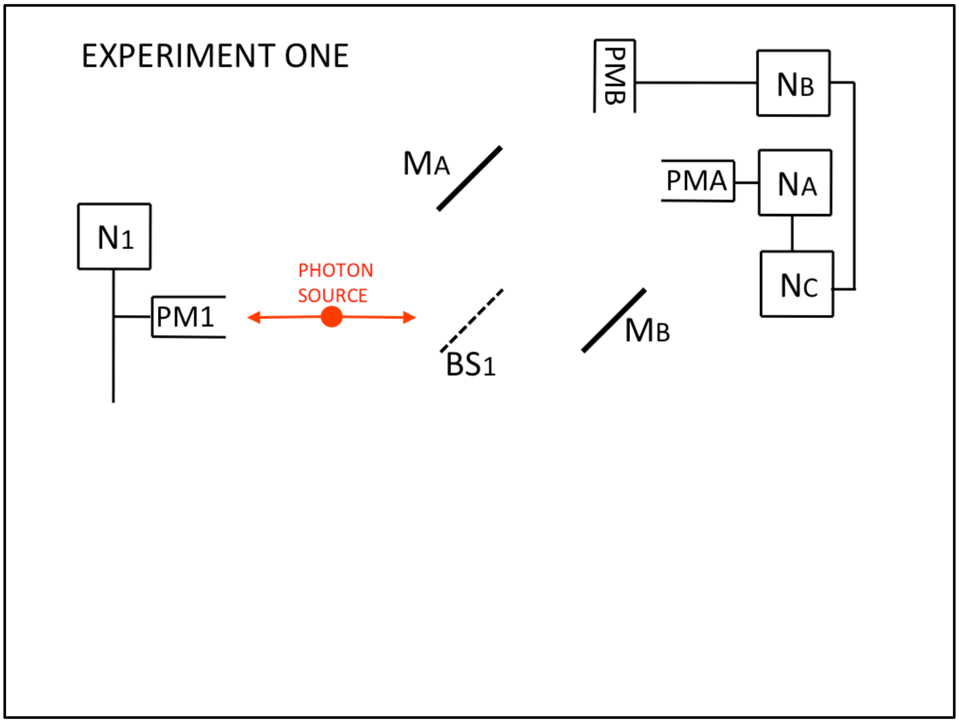

In these first two experiments, we're only going to be working with photons that were emitted back-to-back. Now, there's no reason they have to be emitted back-to-back, and in fact, most of them aren't, but if we collimate the photon source so that we're only working with ones that are, then we reduce the luminosity of the photon source even more.

This is a schematic diagram of the experimental setup. I'll explain in a moment what the various elements are…

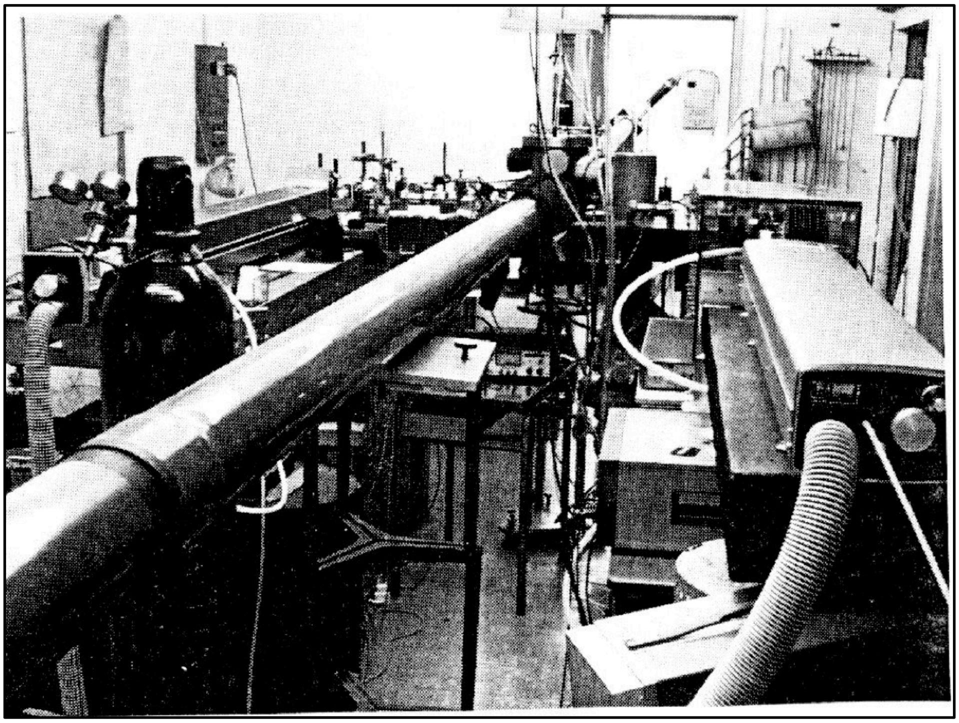

…but just to remind ourselves that the diagram is schematic, here's a picture of the actual experiment. There are evacuated tubes and other equipment, so all the empty space in the diagram isn't really empty – when we say there's just one path for a photon to get from here to there, that's it.

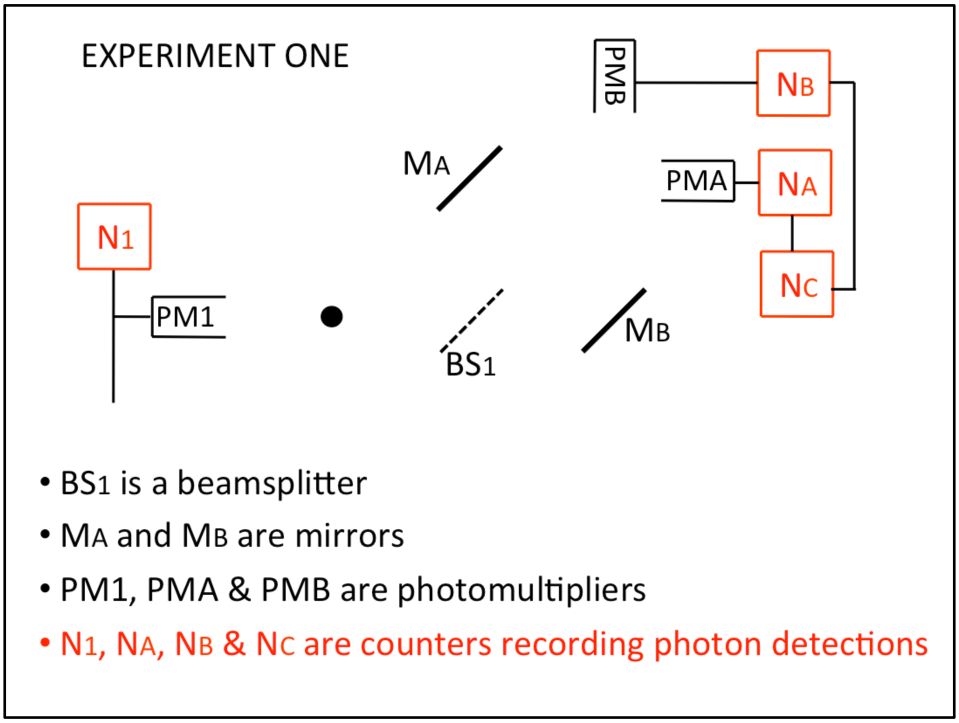

Like the earlier experiment, we're going to have a beamsplitter that we'll aim the photon source at.

Behind the beamsplitter, we'll place two mirrors that will reflect the photons that are reflected and transmitted…

…we'll have photomultipliers set up to detect photons…

…and counters that will record every time a photomultiplier detects a photon.

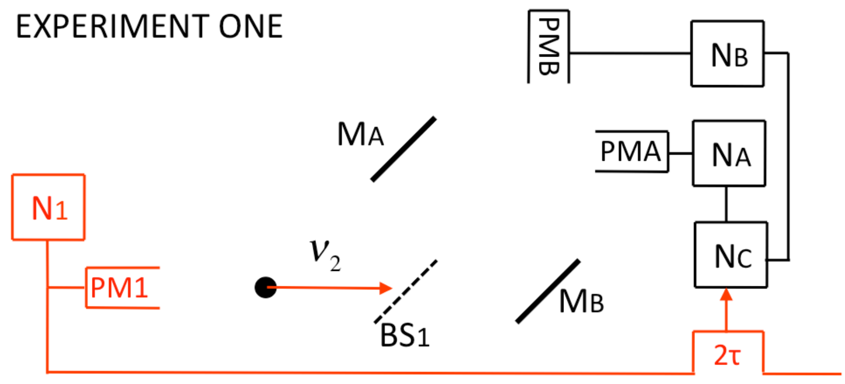

• Detection of first photon (ν_1) is counted by N1

• A signal is sent to tell the other counters (NA, NB & NC) to expect a second photon (ν_2) within a time 2τ

So, this is how the experiment goes. The first photon emitted in the calcium cascade gets detected by photomultiplier 1.

When that first photon gets detected, the electronics sends a signal to the other counters to expect the second photon – the "gate" is open for twice the lifetime of the intermediate state (10 ns), so there's a high probability the second photon has been emitted in that time. After 10 ns, the counters shut down until the next time a trigger photon is detected.

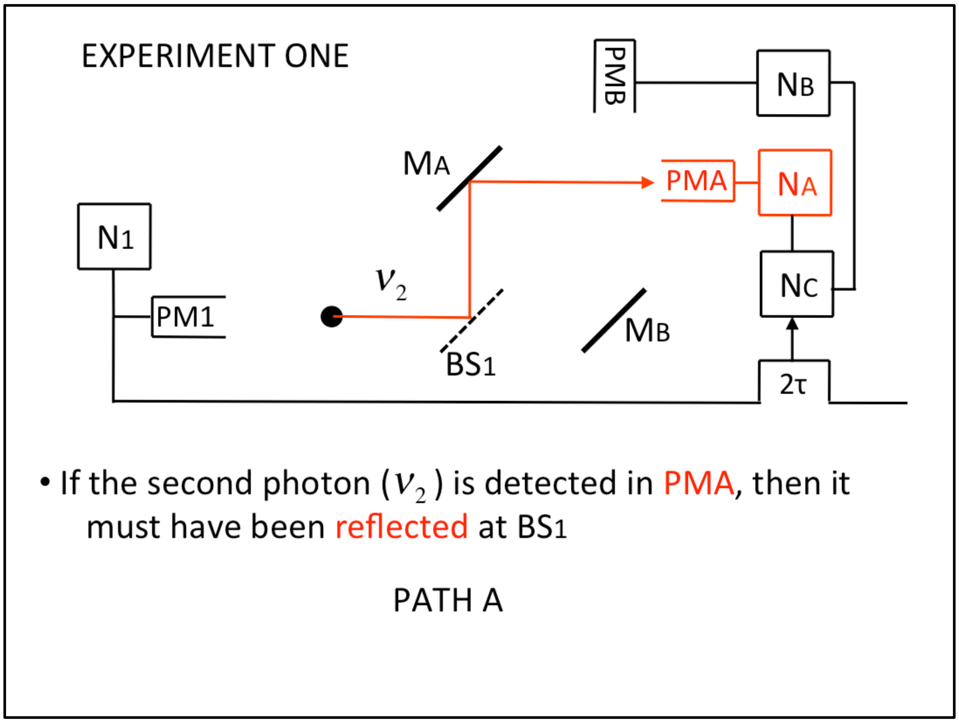

If a photon is detected in photomultiplier A, then it must have been reflected by the beamsplitter, because there's no way for it to have gotten there by any other path. We'll call this Path A.

EXPERIMENT ONE

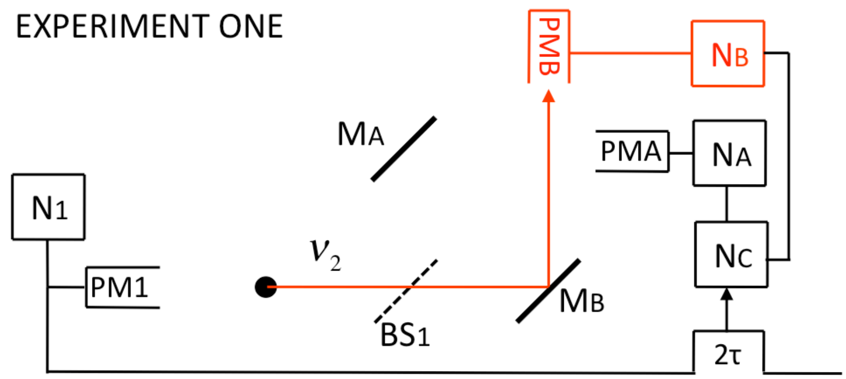

• If the second photon (ν_2) is detected in PMB, then it must have been transmitted at BS1

PATH B

And similarly, if a photon is detected in photomultiplier B, then it must have been transmitted by the beamsplitter, because there's no other way for it to get there. We'll call this Path B.

EXPERIMENT ONE

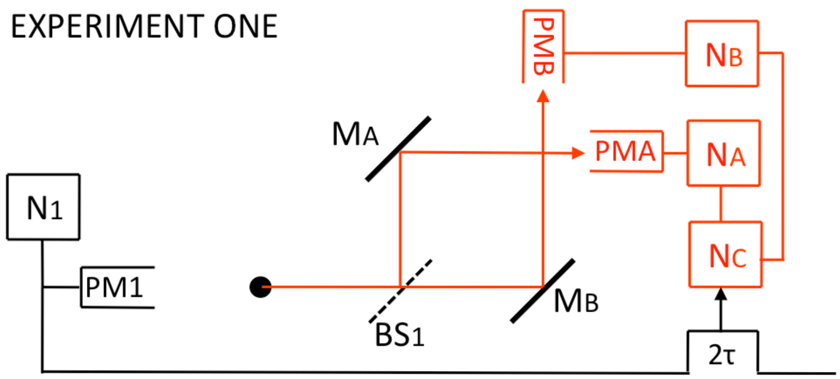

- If both PMA and PMB fire within t = 2τ, then the coincidence counter (NC) is triggered

Many photons…

…or classical wave-like behavior at the beamsplitter

If both photomultipliers fire during the time the gate is open, then the coincidence counter records this.

This could correspond to two situations. One is that we have more than one photon in the apparatus at a time – which is possible, but it should happen less and less as we turn down the intensity of our photon source more and more.

The other possibility is the reason we're doing the experiment in the first place. If the quantum waves corresponding to each photon behave in any way like classical waves, then we could expect that each wave gets coherently split at the beamsplitter, half the energy goes down one path, and half the energy down the other, and deposits energy coincidentally into both of the two photomultipliers.

Whether or not this is happening will tell us something about how the photons are behaving when they encounter the beamsplitter.

ANTI-CORRELATION PARAMETER (α)

Want some kind of measure of how often PMA & PMB are firing simultaneously (within t=2τ)

$$\alpha \equiv \frac{P_C}{P_A \cdot P_B}$$

- $P_A = \dfrac{N_A}{N_1}$ = probability for N_A to be triggered
- $P_B = \dfrac{N_B}{N_1}$ = probability for N_B to be triggered
- $P_C = \dfrac{N_C}{N_1}$ = probability for coincidence counter (N_C) to be triggered (PMA & PMB during t=2τ)

In order to measure how often the photomultipliers are being triggered at the same time, we define what we'll call an "anti-correlation" parameter – alpha.

P(A) is the probability for photomultiplier A to be triggered.

P(B) is the probability for photomultiplier B to be triggered.

P(C) is the probability that the coincidence counter is triggered.

ANTI-CORRELATION PARAMETER $\boxed{\alpha \equiv \frac{P_C}{P_A \cdot P_B}}$

- If NA & NB are being triggered randomly and independently, then $\alpha = 1$ $\quad P_C = P_A \cdot P_B$
- If NA & NB are being triggered separately (reflection **or** transmission) then $\alpha \geq 0$

  $P_C = 0$ when photons are detected by either PMA or PMB, but never both simultaneously
- If NA & NB are being triggered together (reflection **and** transmission) then $\alpha \geq 1$

  $P_C > P_A \cdot P_B$ means PMA & PMB are firing together more often than random (clustered)

What kind of values could alpha take on, and what do those values tell us?

If the two photomultipliers are being triggered randomly and independently, then alpha is going to equal one. The probability for both to be triggered is just the product of the probabilities for each to fire individually.

If the two photomultipliers are being triggered separately, which would correspond to photons being either reflected or transmitted (taking one path or the other but not both), then alpha should be greater than or equal to zero. It would be exactly zero if we never detect photons in both photomultipliers simultaneously.

If the two photomultipliers are firing together more often than random (clustered detections, as we'd expect for reflection **and** transmission of classical waves), then alpha should be greater than or equal to one. The probability for both to fire together is greater than the probability for both to fire as some random coincidence.

So, if we find that alpha is always less than one, we've ruled out the possibility that photons are behaving like classical waves at the beam splitter.

EXPERIMENT ONE

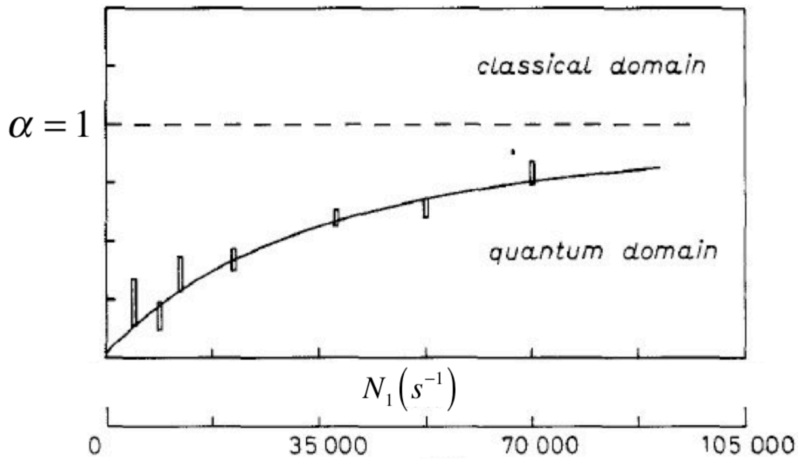

Photons take **either Path A or Path B**, but not both!!

Here are the results of the first experiment, where the horizontal axis represents the intensity of the photon source – the number of counts per second of the trigger photons. The bars represent plus or minus one standard deviation, and the line is what quantum mechanics predicts as we turn down the intensity – the prediction is based on several factors, such as detector efficiencies, the exponential decay of the intermediate state overlapping with the time of the gate, the angular correlation of the photon-pairs, and so forth.

Notice that we're getting down to fewer than ten thousand photons per second – considering how fast the photons are moving and the dimensions of the apparatus, we can be fairly confident that we have only one photon in the apparatus most of the time.

We have to extrapolate down to "single-photon" intensity because the apparatus has what's called a "dark rate" of about 300 photons per second. This means that even without our photon source turned on (lasers are off) we're still detecting thermal photons, and there's also the possibility for the photomultipliers themselves to fire from their own thermal energy without any source photon acting as a trigger. One of the ways we might fix this experiment up is to cool down the apparatus, but the results are fairly clear to be entirely in the "quantum" domain.

We interpret these results as meaning that the photons are always being either reflected or transmitted at the beamsplitter, but not both. This seems to support the notion that photons behave like localized particles, and the superposition state is a reflection of our ignorance of

EXPERIMENT TWO

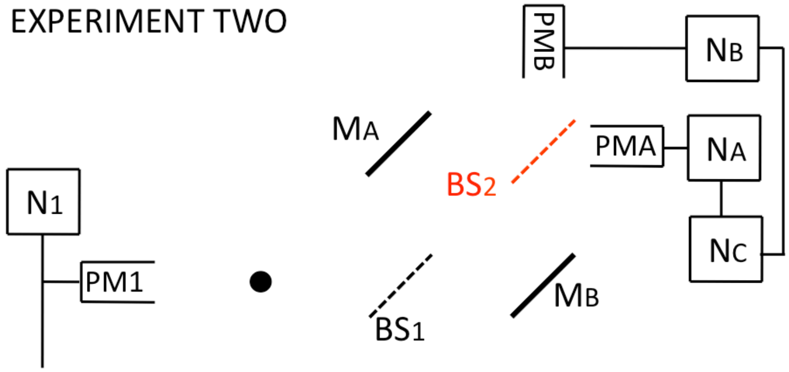

- Same setup as before, but now insert a second beamsplitter (BS2)

We can run the experiment a second time, but this time we'll insert another beamsplitter into the photon paths. This kind of setup is known as a Mach-Zehnder Interferometer.

EXPERIMENT TWO

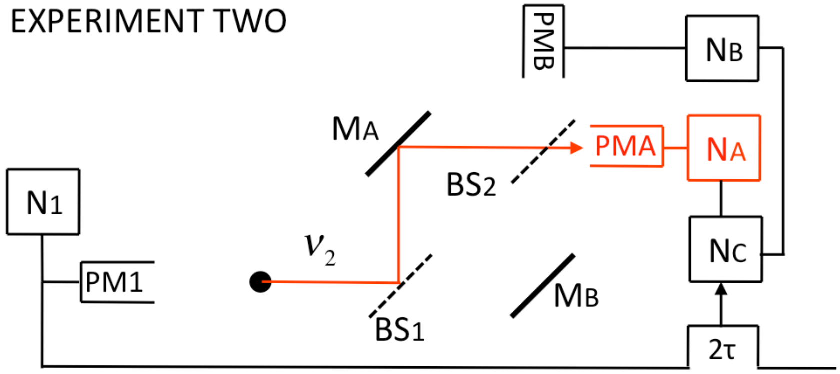

- If the second photon (ν_2) is detected in PMA, then it *could* have been reflected at BS1

This time, if a photon is detected in photomultiplier A, then it could have been reflected at the first beamsplitter and transmitted by the second…

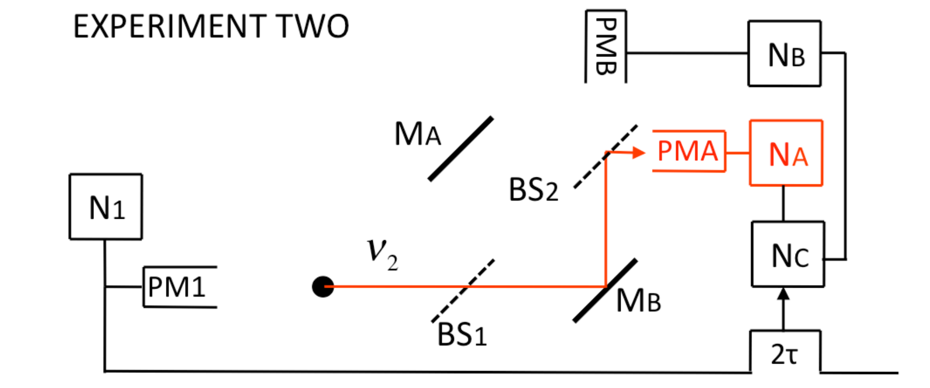

• If the second photon (ν_2) is detected in PMA, then it *could* have been reflected at BS1

OR

…it *could* have been transmitted at BS1

…or it could have been transmitted at the first beamsplitter and reflected at the second.

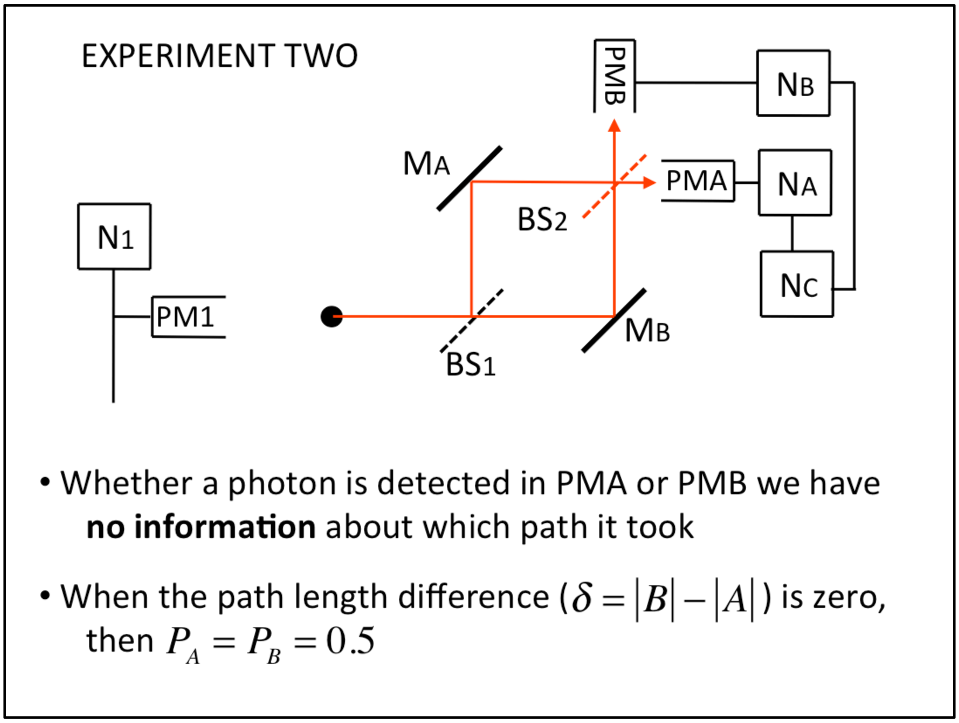

Either way, wherever the photon is detected, in PMA or PMB, we have no information about which path the photon took to get there. In other words, there are two paths any photon could take to get to a given detector.

When we run the experiment, we find that if the two path lengths (A & B) are the same, then there's an equal likelihood to detect the photon in either photomultiplier.

EXPERIMENT TWO $\boxed{\alpha = 0.18}$

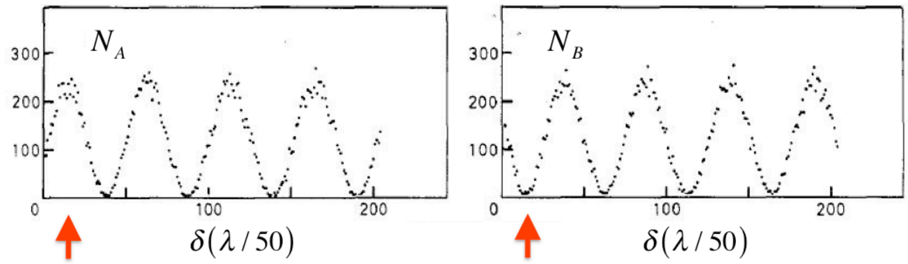

Slowly change the path length difference by moving just MB and we observe interference!

- For some path length differences $P_A \sim 1$ & $P_B \sim 0$
- For other path length differences $P_A \sim 0$ & $P_B \sim 1$

The behavior of every photon is somehow influenced by the relative lengths of both paths!!

However, if we change the path length difference by moving **just one** of the mirrors (Mirror B, for example), then we observe interference!

In physics speak, we say the probabilities for detection in either photomultiplier are oppositely modulated according to the path length difference.

The graph shows the detection rates in 15 second intervals for path length differences in units of lambda/50, where lambda is the wavelength of the photons.

For some path length differences, all of the photons are detected in photomultiplier A and none in photomultiplier B. For other path length differences, all of the photons are detected in photomultiplier B and none in photomultiplier A.

So, somehow the behavior of every photon is influenced by the lengths of both paths. Because otherwise, and here's the argument, how could the behavior of a photon that was supposed to have only gone along path A be affected by changes that are made to the other path B? Somehow, every photon "knows" something about both paths!

INTERPRETING EXPERIMENTS ONE & TWO

- Photons in **Experiment One** took **either** Path A **or** Path B when encountering BS1
  (reflection or transmission – particle-like behavior)
- Photons in **Experiment Two** took **both** Path A **and** Path B when encountering BS1
  (reflection and transmission – wave-like behavior)

When encountering BS1, how does a photon "know" whether we're conducting Experiment One or Two?

Can we "trick" the photon into behaving one way when it should be behaving in another?

Lets recap what's happened. In the first experiment we were forced to conclude that photons take only one path or the other – they were either reflected or transmitted at the first beamsplitter – the kind of behavior we associate with localized particles.

In the second experiment, we were forced to conclude that photons take both paths – they were both reflected and transmitted at the first beamsplitter – the kind of behavior we associate with classical waves.

But how could a photon, when it encounters the first beamsplitter, "know" whether or not the second beamsplitter is in place? And I put the word "know" in quotes because I don't like to anthropomorphize inanimate objects too much – the photon isn't conscious, but it's behavior seems to be determined by whether we're conducting experiment one or experiment two.

Is it possible to somehow "trick" the photon into behaving one way when it should be behaving in another? What if we started off performing experiment one, and then in the middle switch to experiment two?

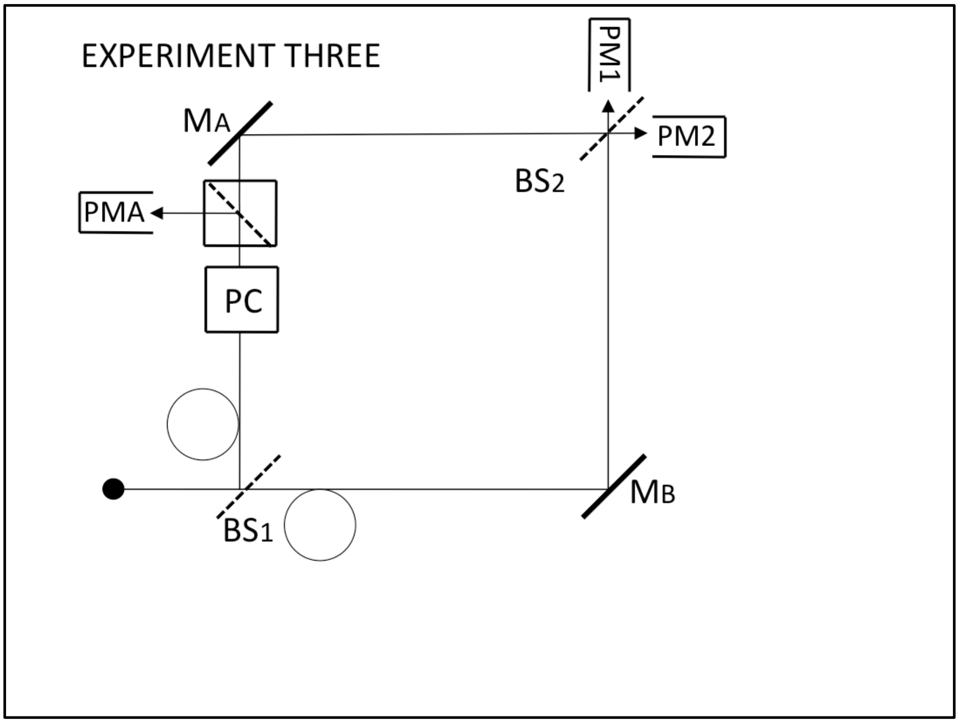

It would be impossible to physically insert or remove a beamsplitter in the time required, but this third experiment is equivalent to what we're proposing. This kind of experiment is called a "delayed choice" experiment, where "delayed choice" refers to the experimenter's choice to perform experiment one or experiment two **after** the photon has encountered the first beamsplitter.

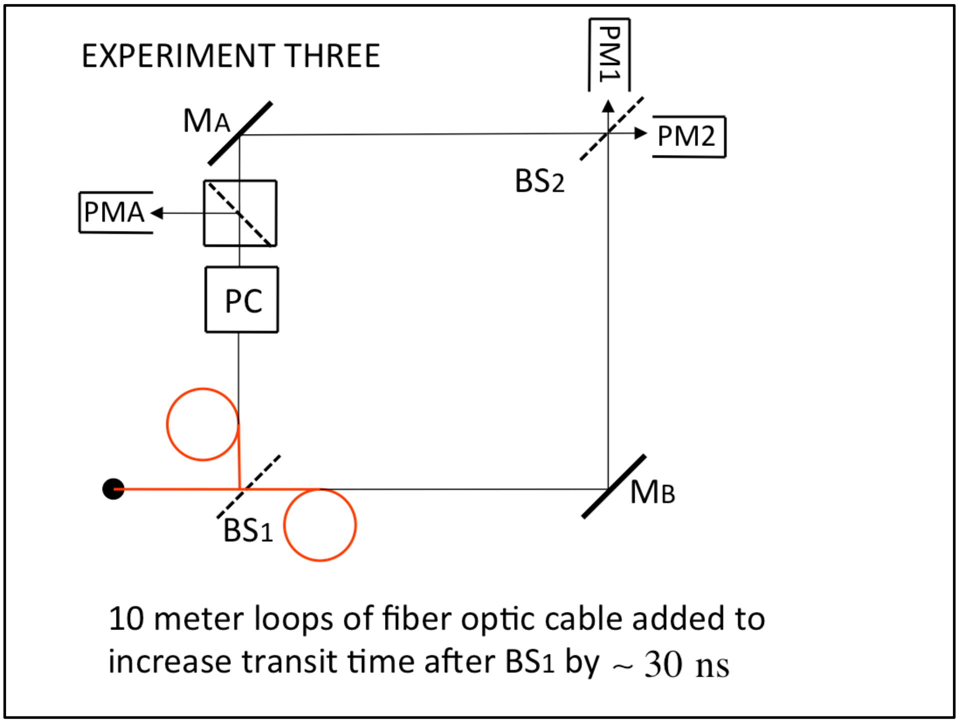

To do this, we insert 10 m fiber optic cables in order to increase the transit time after encountering the first beamsplitter by ~30 ns.

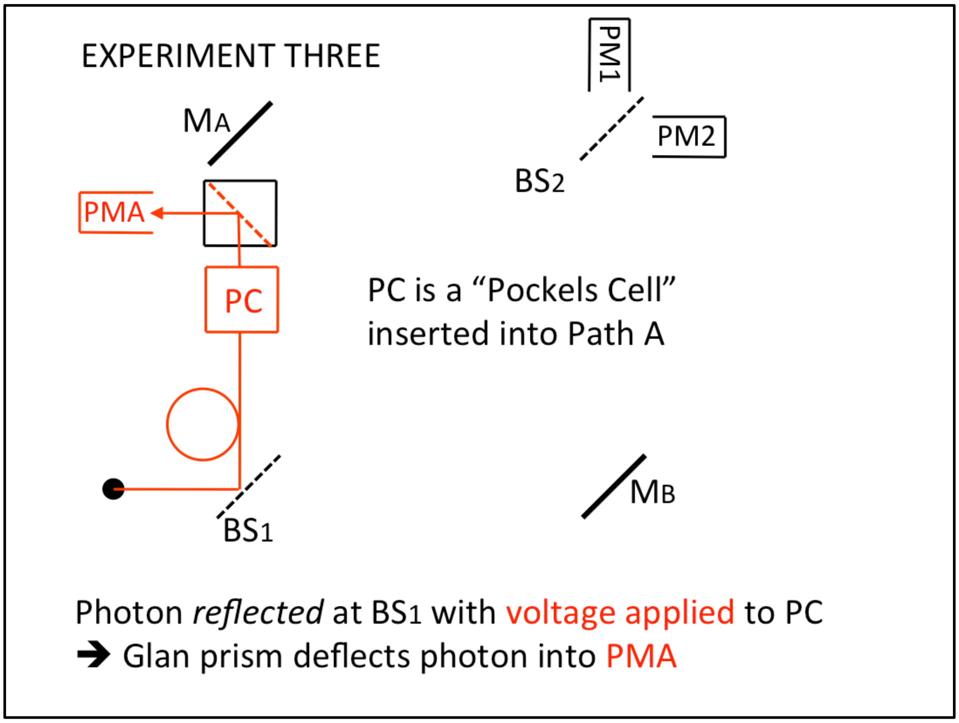

We insert into one of the paths a Pockels cell. If the photon is reflected at the first beamsplitter and a voltage is applied to the Pockels cell, it rotates the plane of polarization of the photon within 5 ns. A Glans prism will deflect a photon whose polarization has been altered into photomultiplier A...

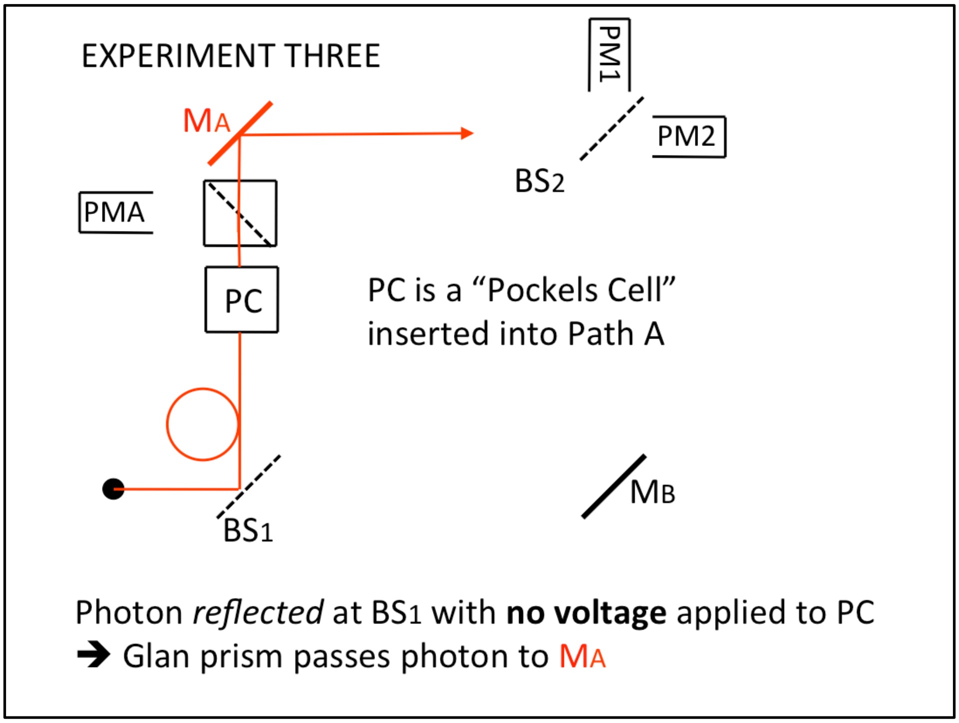

…but it will pass a photon whose polarization hasn't been rotated (voltage to Pockels cell turned off). We can't insert or remove a beamsplitter fast enough, but we can turn the voltage to the Pockels cell on and off at the required frequency.

EXPERIMENT THREE

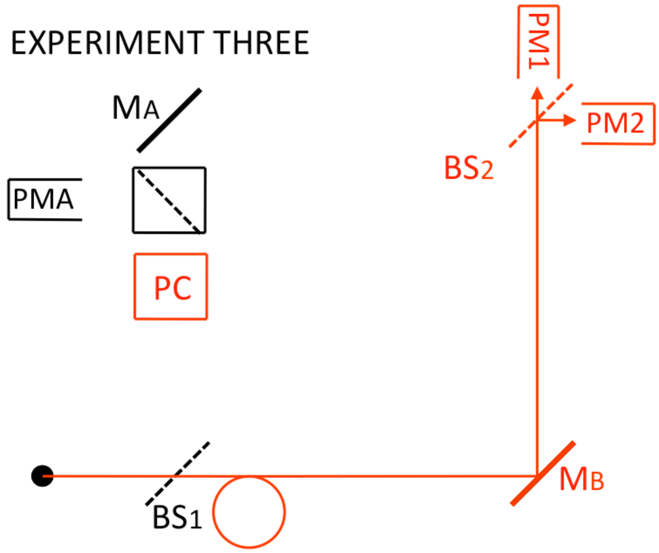

Photon *transmitted* at BS1 with voltage applied to PC

➔ $P_1 = P_2 = 0.5$ (no interference)

If the photon is transmitted at the first beamsplitter while a voltage is applied to the Pockels cell, then there's only one path to the second beamsplitter, and we observe no interference.

EXPERIMENT THREE

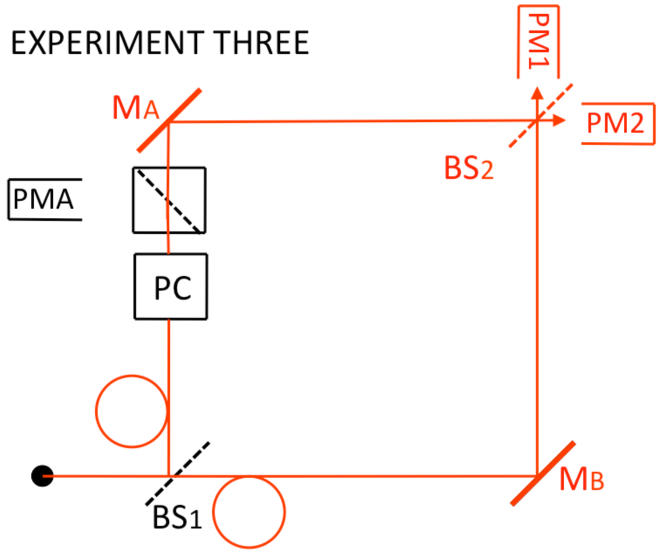

Both paths to BS2 possible with **no voltage** applied to PC
➔ $P_1 = 1 - P_2$ ($0 \le P_2 \le 1$ ➔ interference)

With no voltage applied to the Pockels cell, then both paths to the second beamsplitter are open, and we should observe interference, depending on the path length difference.

EXPERIMENT THREE

NO VOLTAGE applied to PC

- Two paths to BS_2 = Interference

VOLTAGE APPLIED to PC

- One path to BS_2 = No interference

"DELAYED-CHOICE" MODE

- Photon encounters BS_1 with **VOLTAGE APPLIED** (Experiment One)
- Photon *should* choose either one path (to PMA) **or** the other path (to BS_2)
- **VOLTAGE TURNED OFF** *after* photon encounters BS_1 (switch to Experiment Two)

To summarize, with no voltage applied to the Pockels cell, there are two paths to the second beamsplitter, and interference should be observed.

With a voltage applied, there is only one path to the second beamsplitter, and no interference should be observed, whatever the path length difference.

When we operate the experiment in "delayed choice" mode, we'll have a photon enter the apparatus and encounter the first beamsplitter with only one path available (voltage applied). Since we're conducting experiment one, that photon should be either reflected or transmitted at the first beamsplitter. If we turn off the voltage **after** the photon has encountered the first beamsplitter, then suddenly there will be two paths available to the second beamsplitter (switch to experiment two).

EXPERIMENT THREE

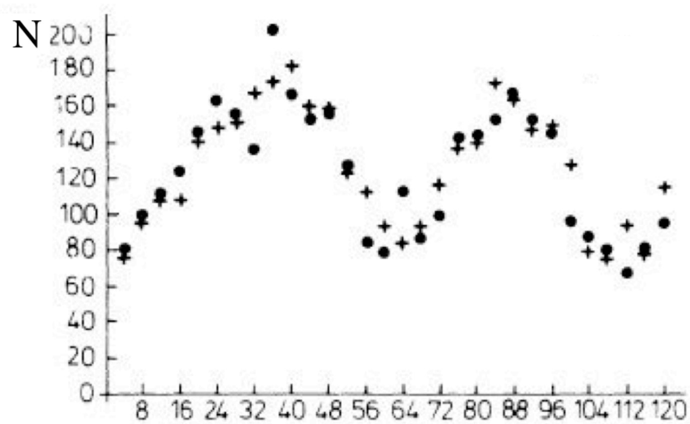

Dots: apparatus operating in “normal mode” – no voltage applied to PC

Crosses: apparatus operating in “delayed-choice mode”

The results of the experiment are shown here, where we again change the path length difference by moving just one of the mirrors. The dots represent the apparatus operating in “normal mode” – no voltage applied to the Pockels cell and two paths to the second beamsplitter. The crosses represent the apparatus operating in “delayed choice” mode. Both data sets indicate a clear interference pattern.

How are we to make sense of this, when the two experiments seem to indicate contradictory behavior of the photons at the first beamsplitter?

"The result of [the detection] must be either the whole photon or nothing at all. Thus the photon must change suddenly from being partly in one beam and partly in the other to being entirely in one of the beams."

P. A. M. Dirac, *The Principles of Quantum Mechanics* (1930)

Let's recall Dirac's explanation for the meaning of the superposition state, which is meant to find a way to make the behavior of the photons seem self-consistent.

According to Dirac, every time a photon encounters a beamsplitter, it is always both transmitted **and** reflected – the superposition of reflection and transmission represents the actual state of the photon – it takes both paths as a delocalized wave, and then randomly deposits all of its energy into a single point when it interacts with a detector.

Now, some people feel this notion of an instantaneous collapse of the wave function is itself so ridiculous that it couldn't be real, but it certainly provides a consistent way of understanding what would otherwise be inconsistent behavior of the photons when they encounter a beamsplitter.

http://phet.colorado.edu

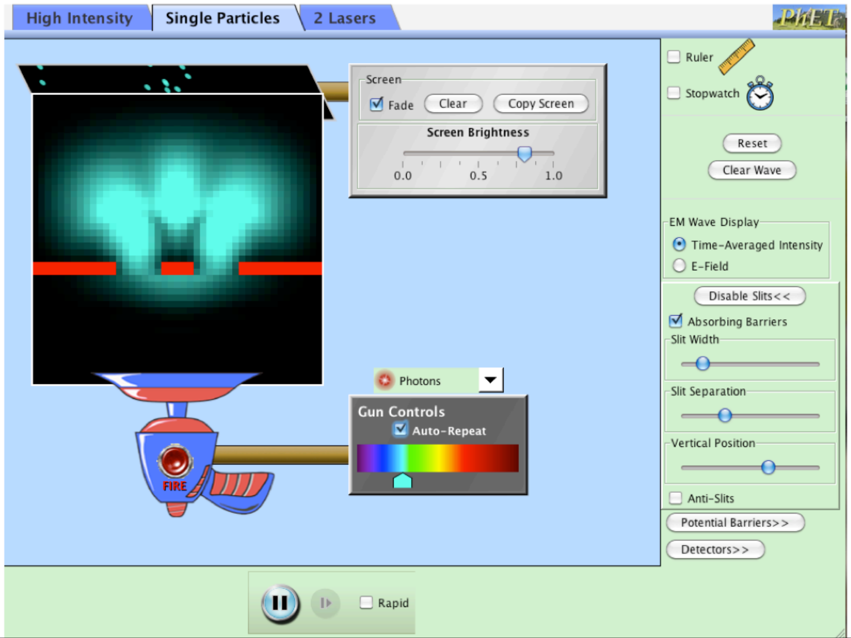

This kind of model is visually represented by a simulation designed by a team at the University of Colorado, led by Nobel-laureate Carl Wieman. Here we can see a double-slit experiment – the photon (or electron) is emitted by a "gun", propagates as a delocalized wave through both slits, interferes with itself, then collapses to a point when it interacts with the detecting screen.

Neutron Interference (1991)

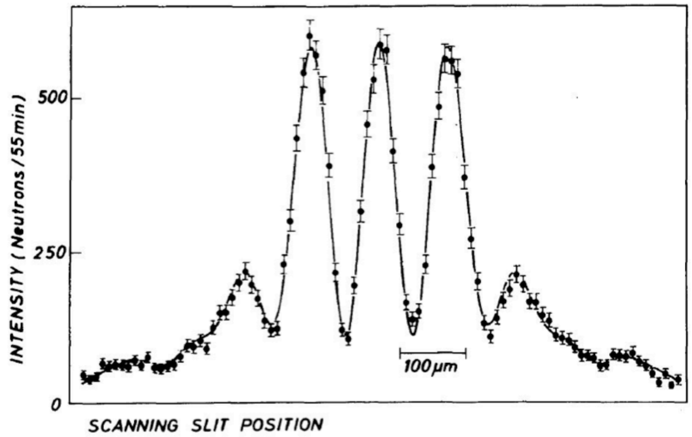

And this kind of behavior isn't limited to just elementary particles, electrons and photons. It's also seen in composite objects, like neutrons.

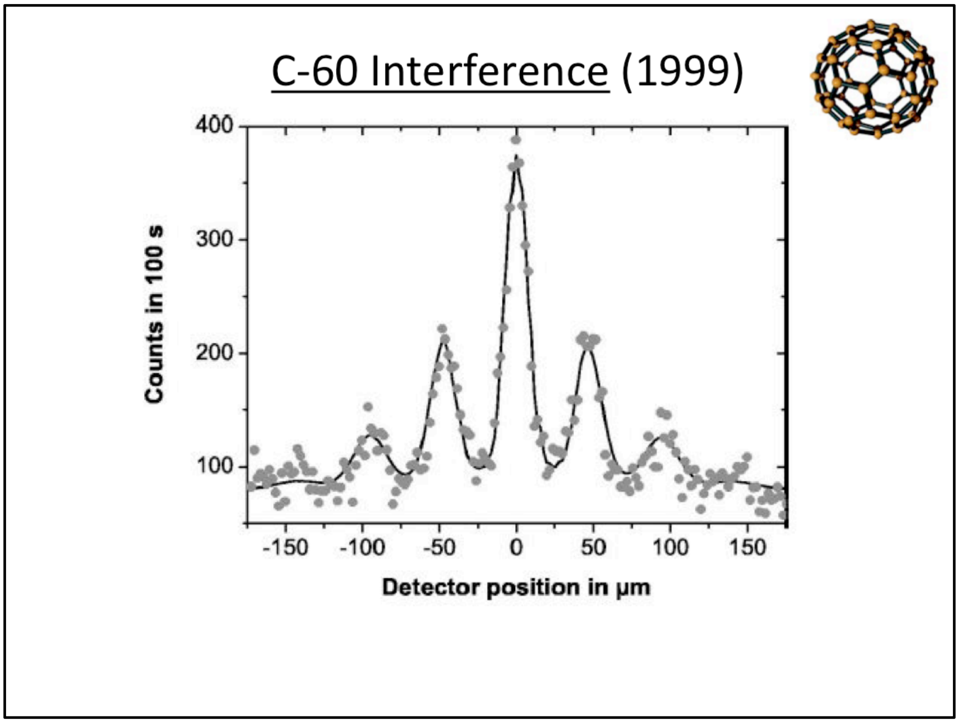

And not just for simple composite objects, also for “buckyballs” – carbon 60 molecules.

BEC Interference (1997)

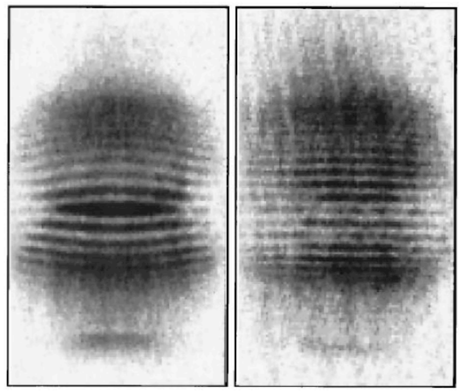

And now, even for macroscopic objects like Bose-Einstein condensates. Here, two separate BEC's that are contained in a magnetic trap, are suddenly allowed to expand outwards and into each other. A BEC is described by a single wave function, and we can see how they interfere with each other when they overlap.

http://tinyurl.com/BailyPER

Double-Slit Experiments:

Nanofabrication and the realization of Feynman's two-slit experiment,
S. Frabboni, G. C. Gazzadi, and G. Pozzi, *App. Phys. Letters* **93**, 073108 (2008).

Demonstration of single-electron buildup of an interference pattern,
A. Tonomura, J. Endo, T. Matsuda, T. Kawasaki & H. Exawa, *Am. J. Phys.* **57**, 117 (1989).

Single-Photon Experiments:

Experimental Evidence for a Photon Anticorrelation Effect on a Beam Splitter: A New Light on Single-Photon Interferences,
P. Grangier, G. Roger, and A. Aspect, *Europhysics Letters* **1**, 173 (1986).

Delayed-choice experiments in quantum interference,
T. Hellmuth, H. Walther, A. Zajonc and W. Schleich, *Phys. Rev. A* **35**, 2532 (1987).

The Reality of the Quantum World,
A. Shimony, *Scientific American* (January 1988, pp. 46-53).

More:

The Quantum Challenge: Modern Research on the Foundations of Quantum Mechanics,
G. Greenstein and A. J. Zajonc (Jones & Bartlett, Sudbury, MA, 2006).

None of these kinds of experiments were possible during the time when it was unclear as to whether quantum states could be used to describe single quanta. Now that they are possible, we are led to conclude that the wave nature of quantum objects is very real, and that superposition states reflect the actual states of each individual quanta, and not only to large ensembles of identical systems.