Path Integrals

Now, we just found out that if we have two slits in a screen, and we want to know the prob. that a photon will pass through the screen, we have to consider the two possible paths that the photon can take. But what if we have 3 slits, or 4, or 5?

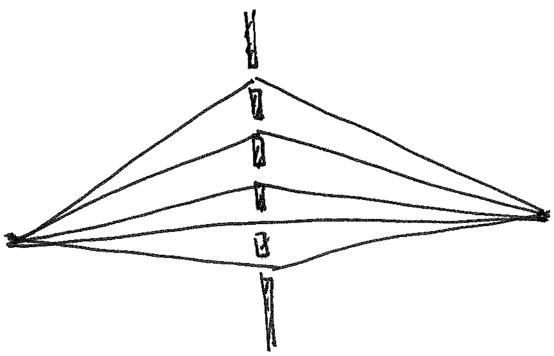

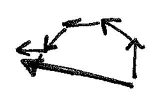

We have to add arrows, or amplitudes, for each individual path to find the total amplitude

Now what if we add another screen?

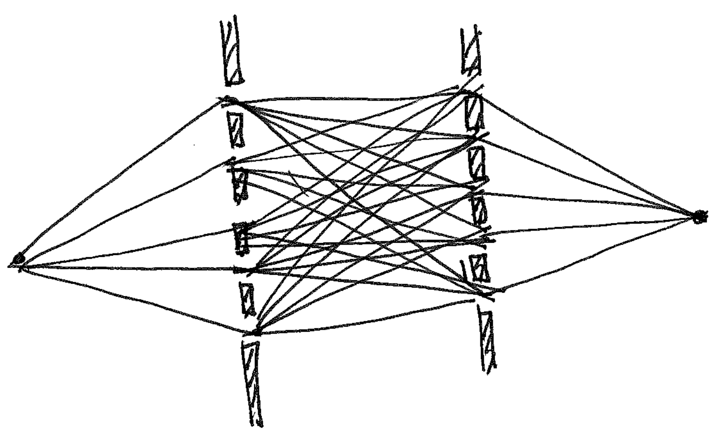

~~Now~~ We have to add

5 x 5 = 25 arrows!

Already, its getting complicated!

Now what if we add another screen, and another ...

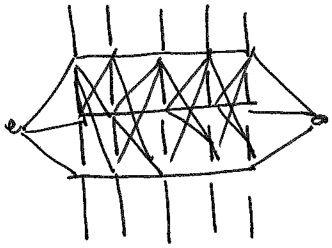

we keep increasing the number of arrows we have to add

If we add more and more screens, ~~and add more and more holes in each screen~~ we soon fill up the space between the emitter + detector. Now if we add more + more holes to each screen, what happens?

We can imagine that if we add more + more holes, the spacing between holes gets smaller, + the holes get closer + closer together. If we add enough holes (an infinite number, in fact), the screen disappears! ~~But if we add an infinite number of holes, we have to add~~ And if we do this to every screen, then each screen disappears! ~~But if we add an infinite number of holes to each screen, we have to add an infinite number of arrows, right? Yes.~~ And if all the screens disappear, we just have ... empty space! ~~Yes.~~

But wait. If we add an infinite number of holes + screens in order to get empty space, don't we have to add an infinite number of arrows to figure out how light travels through empty space?

Yes, that's exactly what we have to do. If we want to calculate the probability (and the most probable path taken) of a photon going from A to B through <u>empty space</u>, we have to add the arrows from an <u>infinite number of paths</u>, because all of those paths are possible

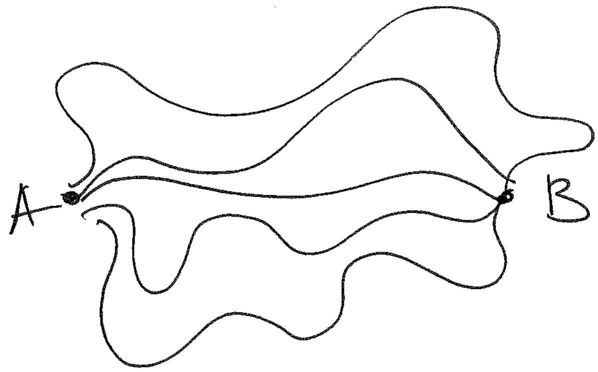

~~[illegible]~~

(The process of "adding up the arrows" or ~~[illegible]~~ accounting for all paths is called a path integral,– we'll talk about ~~this more in a~~ how to compute this in a minute)

So if light can travel on an ~~a~~ infinite number of paths, and all paths are possible, ~~how does light ever get from point A to point B?~~ why would we expect (and why do we see) that light travels in a straight line from A to B?

To answer this, we again have to ~~com~~ add the arrows, or amplitudes, for each possible path. Since it would take forever to do all of them, lets just consider a few:

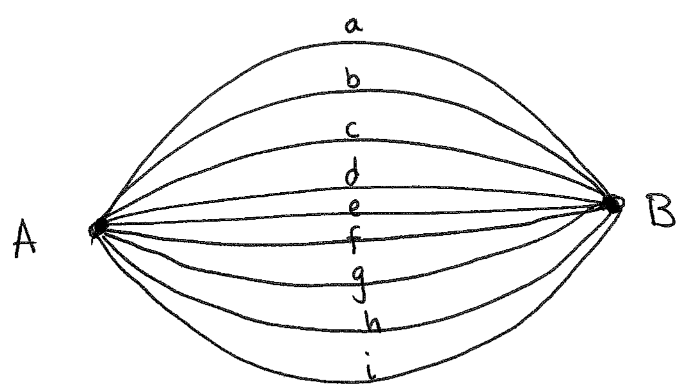

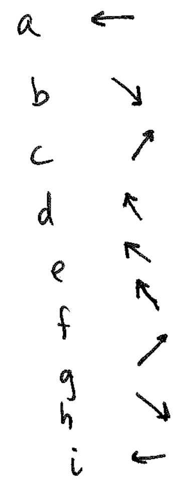

distance traveled changes a lot, so time changes a lot

distances traveled are approx equal, so time approx constant

When we add these all together :

these arrows add!

these arrows cancel each other out

We see that the majority of the contribution to the total amplitude is due to paths d, e, and f - the path close to the straight line joining A and B. If we do this calculation more rigorously, we find that the only sizeable contribution to the prob. amplitude is from the nearly-straight-line paths - all other paths cancel out!

~~Now here we must be careful.~~
~~We find that if light is going to travel from A to B, it will travel~~

Now, believe it or not, the way that we calculate probabilities is just by ... drawing arrows! That's right - we represent an "event" by an arrow. The probability that the event will occur is equal to the length of the arrow squared (from QED, Feynman)

To give you a better idea of what I mean, let me give you an example:

for our single slit, we said that 1 in 100 (or 1%) photons reached the detector. This corresponds to an arrow of length 0.1

Somehow, this ~~represents~~ represents our single-slit experiment

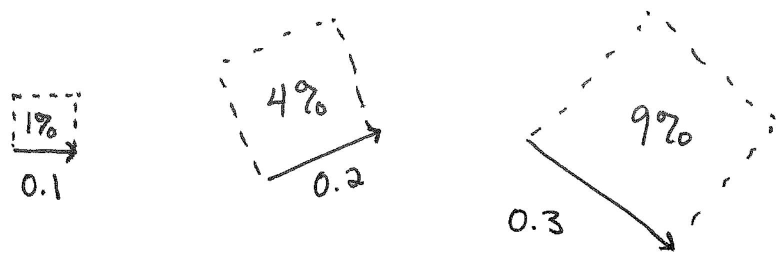

Now, these arrows can ~~in principle~~ point in any direction, and they can be any length ~~$~~between 0 and 1 (we can never have more than 100% probability). The size and direction of the arrow depends on the specific event (ex. ~~single slit~~ photon through slit A, through slit B, photon bounces of screen, etc.)

So how do we figure out what ~~length~~ direction the arrow should point for a specific event? How do we decide which arrows to draw?

Imagine that we have a stopwatch for each individual photon. When the photon leaves the laser, we start the watch. When the photon hits the detector, we stop the ~~d~~ watch. Whichever direction ~~we~~ the final watch hand is pointing is the direction of our arrow (in this example, ↘). As long as the photon is moving, the clock hand is rotating. In fact, the speed of the clock hand depends <u>only</u> on the frequency of light. For red light, the clock hand turns as many as 36000 times for each inch the photon moves!

So let's look more closely at our single slit
How many ways can the photon go to reach the detector?

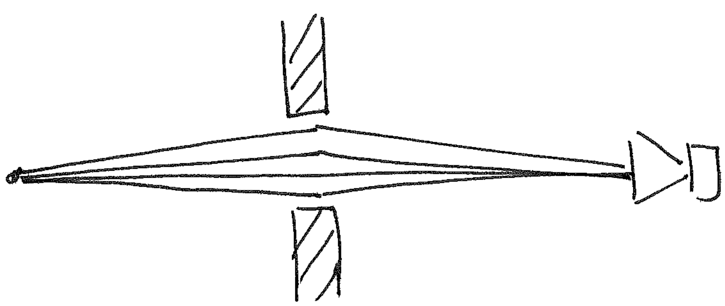

Let's just look at two of these paths, two next to each other:

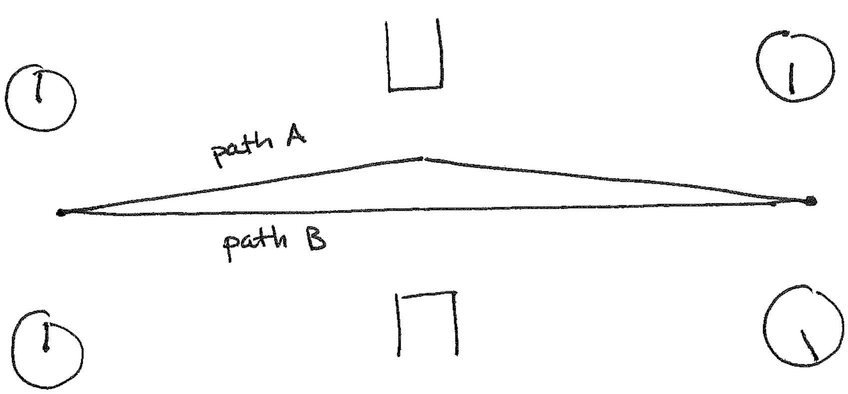

because path A is slightly longer than path B, its clock runs slightly longer.

Now we have 2 arrows, one for each path

path A   path B

Now here's the trick - if the photon can travel on path A <u>or</u> path B to get ~~the~~ to the detector, the total arrow representing the event ~~that~~ is the <u>sum</u> of the arrows for each path:

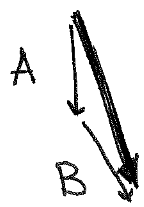

we "add" arrows by placing on arrow at the end of the other, head to tail. The sum is ~~the~~ an arrow going from the tail of A to the head of B.

Now we know that there are actually many more than two ways the photon can travel to the detector. Each way has its own arrow:

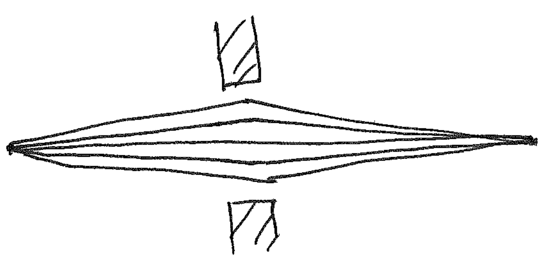

if we add these all together:

This arrow represents the total "amplitude" of the event (it includes all possible ways that the photon can travel).

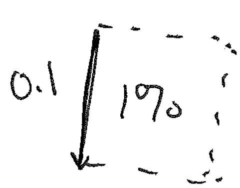

The square of this arrow is the total probability that a photon will reach the detector through a single slit - 1%!

Now, what happens if we place the slit above the detector?

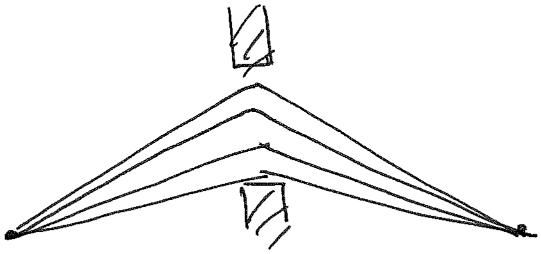

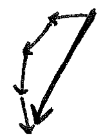

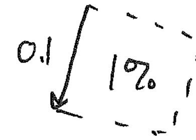

even though all the paths are longer, because they are close together, their amplitudes add ~~to give a 1% probability!~~ to make a sizable arrow!

as we "squeeze" the light, it doesn't necessarily travel in a straight line!

Now we are ready to look at the double slit:

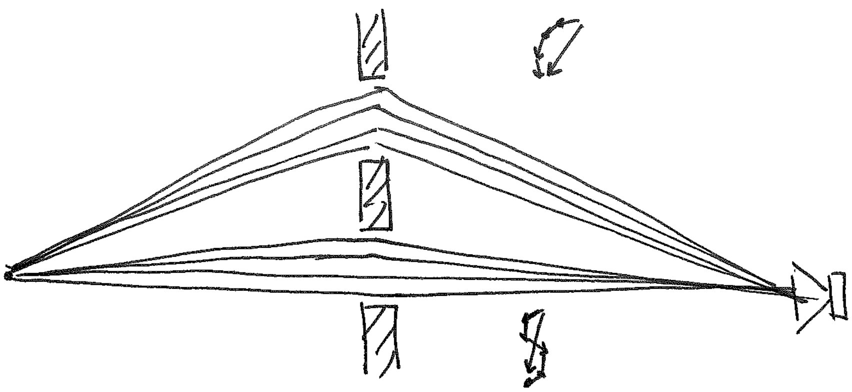

~~If we assume that the size of the slits is much smaller~~
We can redraw this as (averaging over the paths)

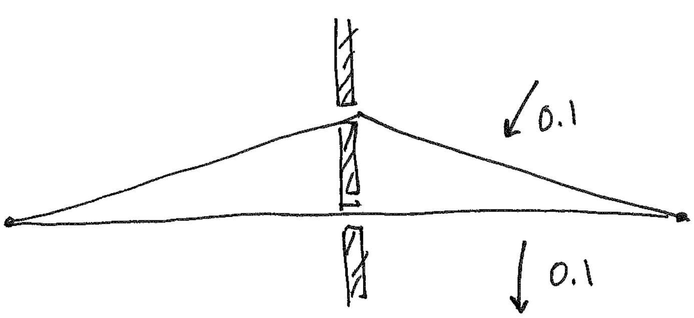

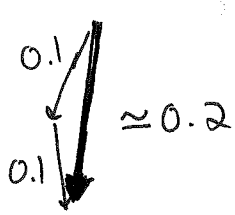

If the slits are very close together, the two arrows are almost aligned, and their sum is close to

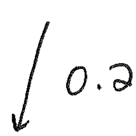

So, for two slits close together, the total prob. that a photon will reach the detector is

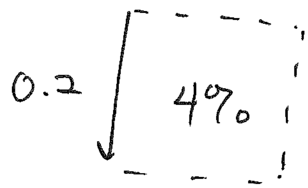

$(0.2)(0.2) = 0.04 = 4\%$

4 out of 100 photons reach the detector, as we said.

Now, what if we move the slits farther apart?

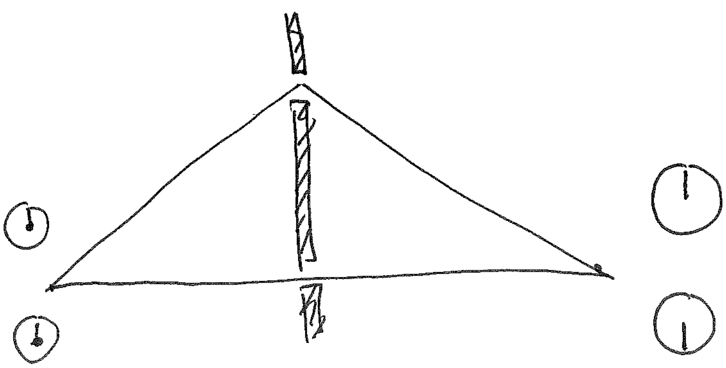

$\uparrow\downarrow = 0$

the arrows cancel each other out, and we have 0% probability that a photon will hit the detector – a photon will <u>never</u> reach the detector!

So, even though there are more ways for the photon to travel to the detector, because the amplitudes that correspond to each path cancel (because the arrows cancel each other out), there is 0% probability that light will reach the detector!

(In reality, there is never exactly 0% probability, but the probability may be vanishingly small, like 1 in 100000000000000 ...)

Ok, now what if we move the detector?

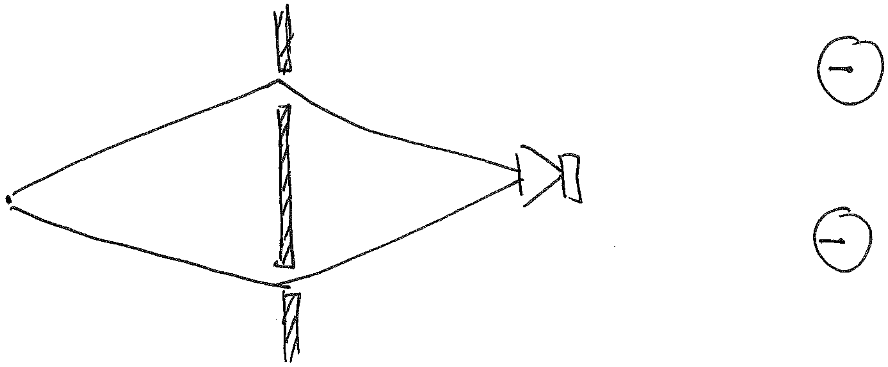

Now the two paths are the same length, so the clock hand moves the same amount in both cases, and the arrows are the same:

So, even for the same slit spacing, the amount of light we detect depends on where we place the detector

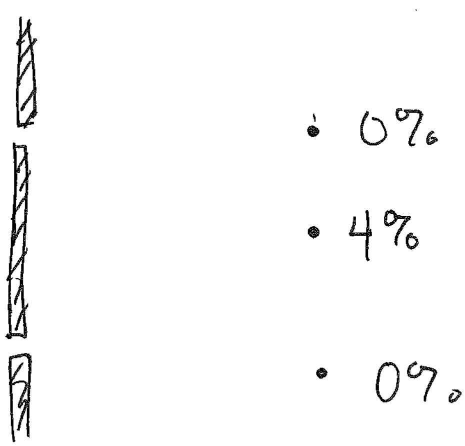

If we draw a line that represents the amount of light we get at each position, we begin to see something interesting

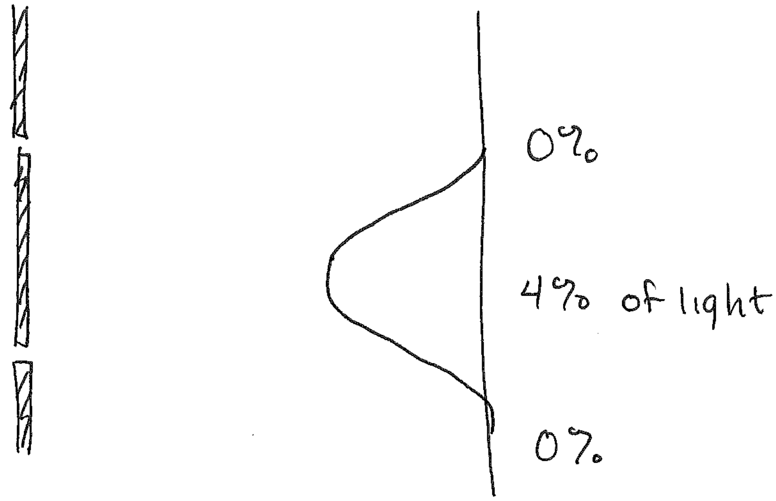

We see that there are bright regions (where more light hits the screen) and ~~less~~ dark regions (where no light hits the screen). Does this start to look like something weve seen?

Yes! Our diffraction pattern! The pattern we saw looks slightly more complicated:

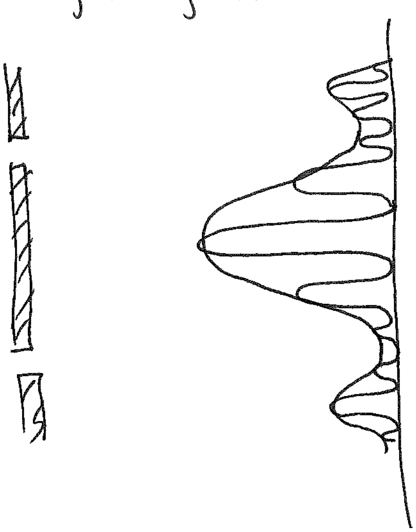

If we were to go through and calculate (add up arrows) the probability at each point on the screen (for every position of the detector) we would find this pattern: alternating regions of zero and non-zero probability, that correspond to alternating regions of bright and dark.

Now, when we shot individual photons at the same double slit, the best that we could do was calculate how probable it was that the photon would land somewhere (or where it was <u>most</u> probable to land). But, when we calculated the probability of seeing a photon <u>at each point on the screen</u> (probability distribution), we found that the probability distribution had a specific pattern that looked like the intensity pattern we saw on the screen. in our experiment

Comparing these two results, we can see that

<u>intensity pattern of the wave</u> $\longleftrightarrow$ <u>probability distribution of the particle</u>

So, even though we are measuring one particle at a time, the <u>particle</u> measurements (detecting individual photons) reveal something about the <u>wave</u> nature of light (intensity pattern). conversely, the wave measurements (intensity) tell us something about the particle nature (probability – where photons are likely to land).

And if we think about this a little more, it actually starts to make sense:

high intensity → lots of light.

If light is made of photons...

lots of light → lots of photons

and if the behavior of photons is governed by probabilities

lots of photons → high probability of photons landing there

From this, we conclude

high intensity → high probability

This wave and particle behavior go hand in hand - they are two sides of the same coin, and try as we might, we cannot remove one or the other. We are always stuck with both. Light ~~is~~ behaves like both a wave and a particle - this is a fundamental idea in quantum mechanics, and is called wave-particle duality

So when does light behave like what? Why do we sometimes see wave behavior, and other times see particle behavior?

As I mentioned earlier, <u>what</u> we observe depends on <u>how</u> we observe it.

More specifically, it depends on the size of our experiment relative to the size of the wavelength of the light. To explain this, let me give an example:

Say you're riding on a boat in the ocean. If your boat is really small compared to the ~~size~~ length of the wave, you really feel it rocking up + down:

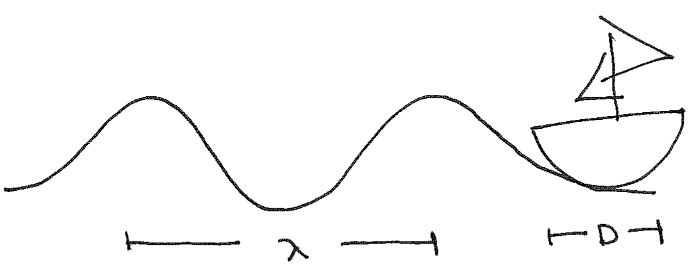

$D < \lambda$

acts like a wave

But, if ~~y~~ the length of the wave is small compared to the size of your boat, ... you feel like you crashed into something

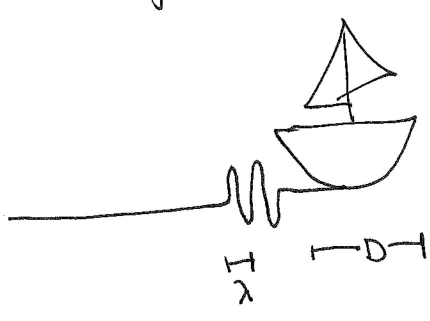

$D > \lambda$

acts like a particle

The water acts like a particle.

~~In this analogy, the wave is light and the size of the boat is roughly the size of the slit in our diff. expt.~~

Now while we're on the topic of measurement, there's one crazy behavior I'd like to talk about.

In our diff. expt, we saw that light behaves like a wave. — we saw this because we measured an intensity pattern on the screen. But what if we put detectors on the individual slits to see where the light is actually going? If light is a wave, we should see both detectors go off. But if light is a particle, we should see one detector or the other go off (unless the particle splits in two and goes through both).

If we send light in, we see that the two detectors never go off at the same time. Either one or the other goes off, but never both. We do not observe the wave passing through both slits, nor do we see particles splitting in two. But, more interestingly, we never see interference on the screen when we put detectors on the slits. So, in observing the particle behavior, we cannot observe the wave behavior!

Let me explain how we can calculate this:

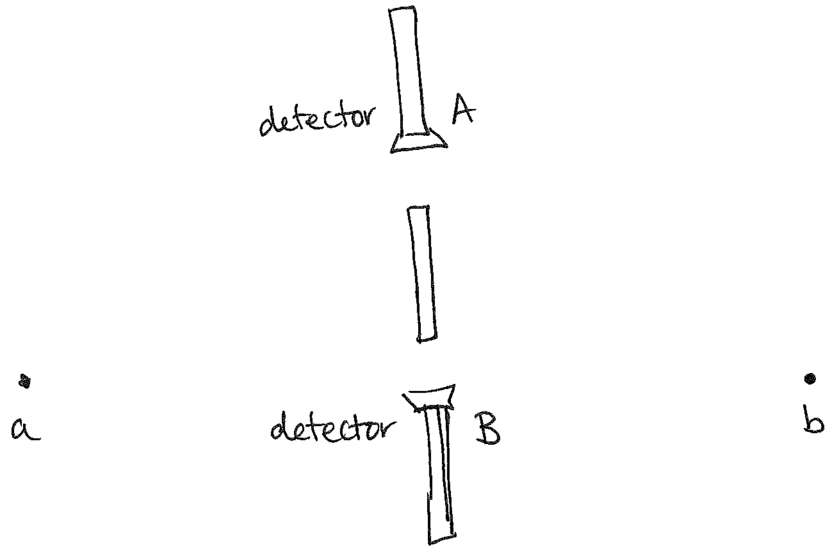

when we saw interference, we found that the prob. of detecting a photon varied from 1% - 4% depending on where we put the detector - this was due to interference.

Now, when we put detectors on the slits, what is the probability of detecting a photon?

Here we have to be a little careful. Before, we wanted to know the probability of a photon hitting "a." That was ~~what~~ what we call an "event" (event = photon hits "a"). The total probability was just the prob. of the event. Now we still want to know the prob. that a photon hits "a," but we have 2 events (can you guess what they are?)

event 1 = photon hits "A", photon hits "a"

event 2 = photon hits "B", photon hits "a"

Now, the probability that the photon hits "a" is the <u>sum</u> of the probabilities of the two events:

$$P_a = P_{\text{event } 1} + P_{\text{event } 2}$$

So what are $P_{event\ 1}$ and $P_{event\ 2}$?

→ They are just the probabilities that a single photon will pass through a single slit - 1%

$P_a = 1\% + 1\% = 2\%$ no matter where we put the detector "a".

So, if the probability is 2% everywhere, we never see interference! By putting ~~a~~ detectors on the slits, we have determined that light is behaving like a particle, ~~But in doing this, we have "removed" the wave nature. We cannot have it both ways!~~ and we know exactly which slit it is going through (we know it is not going through both). But in doing this, we have removed the wave nature! We cannot have it both ways!

To see interference, we cannot know which slit the photon~~light~~ passes through. We have to assume that the photon behaves like a wave and passes through both slits.

The bottom line is, the act of measurement physically alters the system, and thus alters the outcome. To see wave behavior, we cannot probe particle behavior. We have to assume that the particle, acting as a wave, passes through both slits simultaneously.