

## 2. Wavefunctions

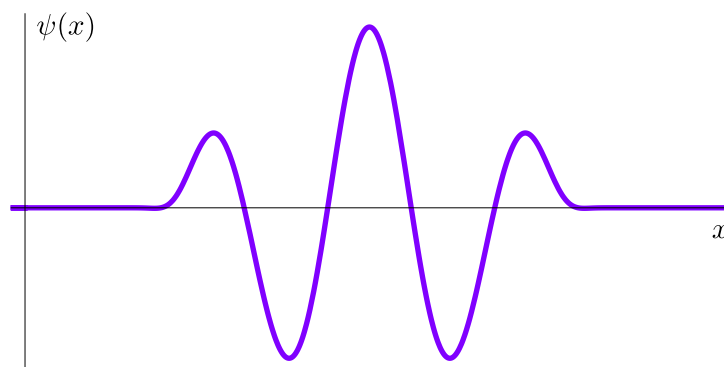
Copyright ©2015–2016, Daniel V. Schroeder

To create a precise theory of the wave properties of particles and of measurement probabilities, we introduce the concept of a *wavefunction*: a function of space that encodes the current state of a system.

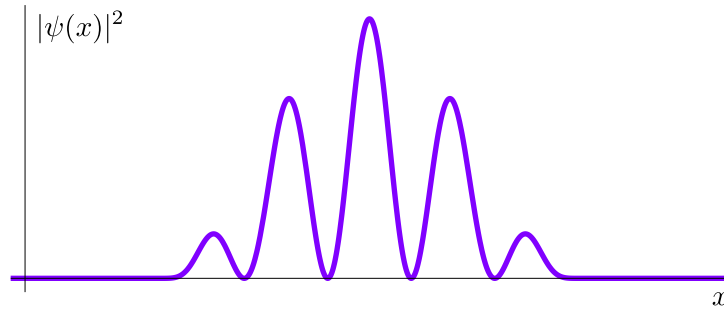
For now, we’ll assume that the “system” consists of a single particle living in a one-dimensional universe. (We’ll generalize to more complicated systems in a few weeks.) Then, if this were a *classical* particle, the “state” of the system would consist of just two numbers: its position  $x$  and its momentum  $p_x$  (or velocity  $v_x$ , which you can easily calculate from the momentum). For a *quantum* particle, the state instead consists of the *wavefunction*  $\psi(x)$ , a whole infinity of numbers (one for each  $x$ ). Quantum states are vastly more complicated, and interesting, than classical states.

### An example wavefunction

For example, if we draw the  $x$  axis across the two-slit interference pattern illustrated in the previous lesson (and ignore the other two dimensions of space), then the wavefunction of each particle, just before it hits the detection screen, might look something like this:



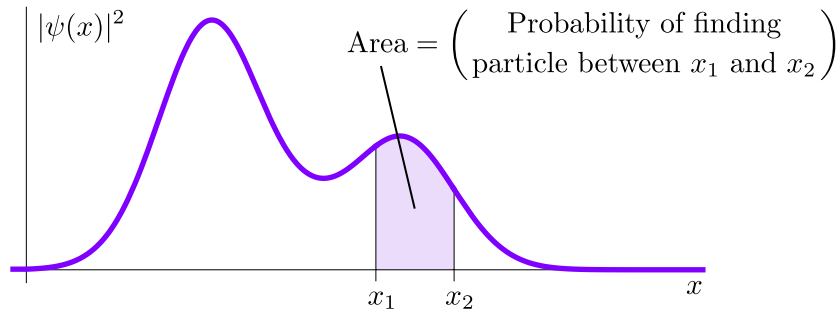
This wavefunction has five “bumps,” corresponding to the five bright lines in the interference pattern. The dark lines in the pattern are at the locations where the wavefunction is zero. More precisely, the brightness of the interference pattern is proportional to the *square* of the wavefunction, in analogy to the way the energy in a mechanical wave or an electromagnetic wave is proportional to the square of the wave amplitude. Here is a plot of the square of our five-bump wavefunction:



The height of this graph at any  $x$  is then proportional to the probability of finding the particle at  $x$ , when the subsequent interaction with the detection screen “measures” the particle’s position. After many such measurements are made on identically prepared particles, the five-line interference pattern emerges.

### Computing probabilities

Because  $x$  is a continuous variable, the actual probability of finding a particle at any *exact* value of  $x$  is infinitesimal. To get around this awkwardness, we can instead ask for the probability of finding a particle *between* two values of  $x$ . We obtain this probability by calculating the corresponding area under the graph of  $|\psi|^2$ . For a generic wavefunction, you can visualize the calculation like this:



Written as an equation, the rule for calculating probabilities is to integrate:

$$\left( \begin{array}{c} \text{Probability of finding} \\ \text{particle between } x_1 \text{ and } x_2 \end{array} \right) = \int_{x_1}^{x_2} |\psi(x)|^2 dx. \quad (1)$$

The function  $|\psi(x)|^2$  is called the *probability density*, and I like to think of it as a function *whose purpose in life is to be integrated*. In order for the rule to work, however, we must impose the condition that the total probability of finding the particle *somewhere* equals exactly 100%:

$$\int_{-\infty}^{\infty} |\psi(x)|^2 dx = 1. \quad (2)$$

Any function  $\psi$  that satisfies this condition is said to be *normalized*.

## Computing averages

Sometimes, instead of calculating the probability of finding a particle at (or near) a particular location, we want to know its *average* position. For example, the average  $x$  for the five-bump wavefunction drawn above would be in the middle of the central bump, while the average  $x$  for the function drawn above equation 1 would be somewhere to the right of the highest peak but to the left of  $x_1$ . I'll use the symbol  $\langle x \rangle$  to denote the average value of  $x$ .

To figure out the general formula for  $\langle x \rangle$ , it's helpful to first imagine that the values of  $x$  are discrete, separated by intervals of  $dx$ , so that the probability of any particular  $x_i$  is  $|\psi(x_i)|^2 dx$ . Then the average  $x$  is just the sum of all the possible values  $x_i$ , weighted by their probabilities:

$$\langle x \rangle = \sum_i x_i |\psi(x_i)|^2 dx \longrightarrow \int_{-\infty}^{\infty} x |\psi(x)|^2 dx, \quad (3)$$

where in the last expression I've taken the continuum limit.

You can use a similar formula to calculate the average value of any *function* of  $x$ :

$$\langle f(x) \rangle = \int_{-\infty}^{\infty} f(x) |\psi(x)|^2 dx. \quad (4)$$

A common example would be to compute the average value of  $x^2$ , which lets you then calculate the *standard deviation* of  $x$ , denoted  $\sigma_x$ :

$$\sigma_x = \sqrt{\langle x^2 \rangle - \langle x \rangle^2}. \quad (5)$$

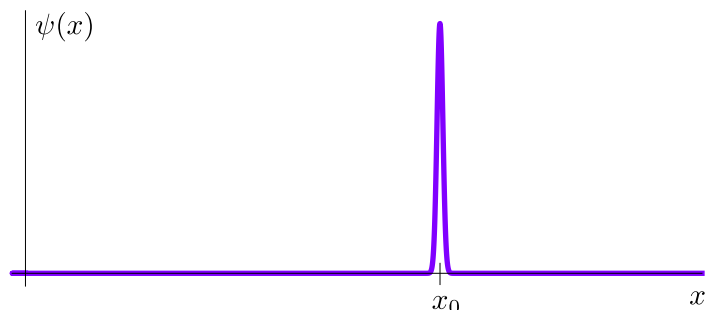
This is a rough measure of the “width” or “spread” of a wavefunction. (If you need a review of the concept of a standard deviation, or don't know how to derive the preceding formula, see any statistics textbook or Griffiths, Section 1.3.)

Although we commonly say that these average values apply to a single particle whose wavefunction is  $\psi$ , it's important to remember that if you actually *measure* the value of  $x$ , you'll get some random value that may or may not be near the average,  $\langle x \rangle$ . It's therefore impossible to determine  $\langle x \rangle$  (or  $\langle x^2 \rangle$  or  $\sigma_x$ ) with a single measurement. To determine these quantities in the laboratory you must first prepare *many* identical particles in the same wavefunction  $\psi$ , then measure all their positions and average the results.

## Localized wavefunctions

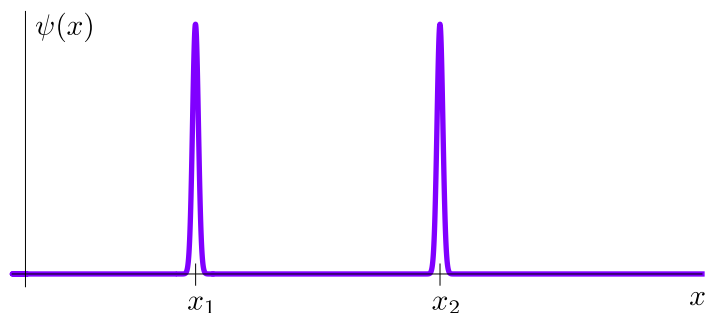
*Any* normalized function of  $x$  is an allowed wavefunction, which you could use to describe a quantum particle under suitable conditions. By a *function*, I mean any single-valued graph that you can draw on a set of axes—regardless of whether you can write down a formula for it.

Sometimes a quantum particle is localized within a very small region of space. To describe such a particle we would use a “spiky” wavefunction that’s zero everywhere except in that narrow region:



You could use a variety of formulas for such a function, but often the exact shape doesn’t matter, and it’s convenient to take the limit where the spike becomes infinitely narrow and infinitely tall to become a *Dirac delta function*, denoted  $\delta(x-x_0)$ . Borrowing a word from German, we say that a delta function is an *eigenfunction* (which could be translated “characteristic” or “particular” function) of position, meaning that it’s a function for which the particle’s position is precisely defined. A delta function is *not* a properly normalized wavefunction, however.

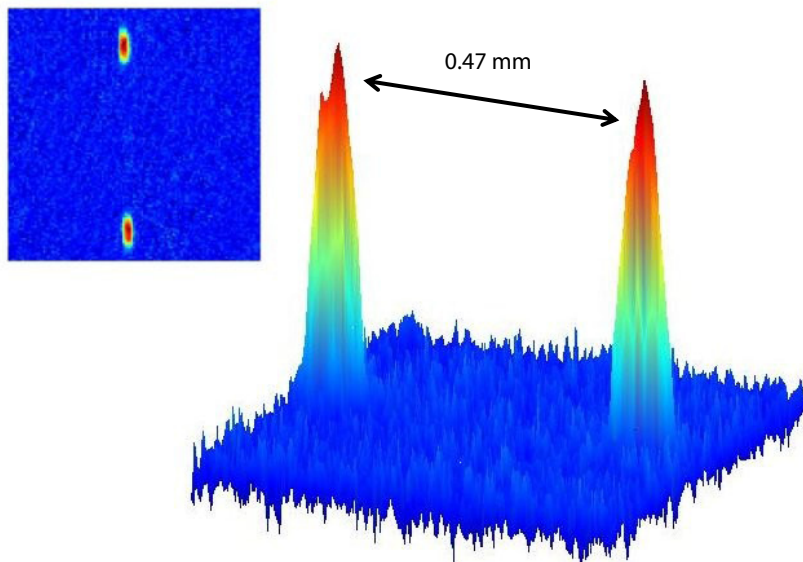
A wavefunction with a single narrow spike is in some ways like a classical particle. But an equally valid wavefunction could have *two* narrow spikes, separated by a gap:



This wavefunction has no classical counterpart, because it describes a particle that is (apparently) half in one location and half in another—and unambiguously *not* half way in-between the two locations!

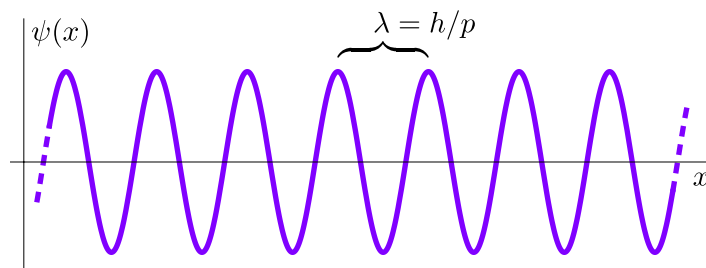
However, just because a wavefunction is mathematically *allowed* doesn’t mean it’s easy to physically *prepare* a particle to have that wavefunction. Putting a particle into a two-peak wavefunction is pretty easy if the particle is a photon (just aim it at a semi-reflective mirror), and becomes progressively more difficult for heavier particles. Here is a recent state-of-the-art image showing the probability density of a collection of about 10,000 identically prepared rubidium atoms, each of which was split between two distinct locations separated by about half a millimeter

(J. H. Burke, “Improvements and Applications of a Guided-Wave Bose-Einstein Condensate Interferometer,” PhD thesis, University of Virginia, 2010):



## Definite-momentum wavefunctions

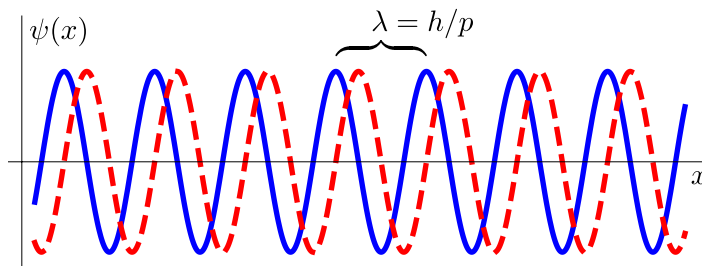
Completely different from a spiky, localized wavefunction would be one that describes a particle with well-defined *momentum*. For this purpose, according to the de Broglie relation, we should use a periodic function with wavelength  $\lambda = h/p$ . The simplest periodic function would be a sine or a cosine, which would look like this:



A long wavelength would correspond to a small momentum, and a short wavelength would correspond to a large momentum. A true sinusoidal function would repeat forever in both directions, and therefore wouldn't be normalizable, so we should remember in the back of our minds that the oscillations must eventually die out in both directions. As we'll see later, this means that no normalized wavefunction has a *perfectly* well-defined momentum; a perfectly well-defined momentum is an idealized but unphysical limiting case, just as perfectly well-defined position, described by a delta function, is an idealized limiting case.

But besides the issue of normalization, our sinusoidal wavefunction has two other awkward features. First, the square of this wavefunction alternates back and forth between some maximum value and zero, so it describes a particle whose probability density varies periodically with location. There's no reason a particle can't have such a wavefunction, but surely there must be a way to describe a particle with well-defined momentum whose probability density is more uniform. Second, our sinusoidal wavefunction doesn't seem to encode the *direction* of the particle's momentum; it would apparently look the same whether  $p_x$  is positive or negative (since  $\lambda$  can only be positive).

The way to fix both of these defects is to give the wavefunction a second component at each location in space—that is, to make it a *pair* of functions instead of just one. Then, to describe a particle with well-defined  $p_x$  and a uniform probability density, we use sinusoidal functions that are out of phase by a quarter cycle ( $\pi/2$ , or  $90^\circ$ ), so that one component is large in magnitude where the other is zero and vice-versa. And to distinguish left-moving from right-moving particles, we associate a quarter-cycle phase difference in one direction with positive  $p_x$ , and a quarter-cycle phase difference in the other direction with negative  $p_x$ . Here is a plot of both components, where I've used a solid line for the first component and a dashed line for second component:



By convention, we associate this particular phase relationship (second component a quarter-cycle to the *right* of the first component) with a *right*-moving wavefunction, that is, with *positive*  $p_x$ . To draw the analogous *left*-moving wavefunction, you would shift either component left or right by a half cycle, or, equivalently, flip either component upside-down.

## Complex notation

One way to write a formula for this two-component wavefunction would be to just list the two components in order:

$$\psi(x) = \{A \cos(kx - \phi), A \sin(kx - \phi)\}, \quad (6)$$

where  $A$  is the wave amplitude,  $\phi$  is a phase constant that determines where the cycles start, and  $k = 2\pi/\lambda$ . Notice that the two components share the same  $A$ ,  $\phi$ , and  $k$ . But we can write this formula much more concisely using *complex numbers*.

By convention, we associate the two components of  $\psi$  with the so-called “real” and “imaginary” parts of a single complex-valued function:

$$\psi(x) = A \cos(kx - \phi) + iA \sin(kx - \phi). \quad (7)$$

You may now recognize that  $\cos \theta + i \sin \theta$  is the *definition* of the complex exponential function,  $e^{i\theta}$ , so we can write simply

$$\psi(x) = A e^{i(kx - \phi)}. \quad (8)$$

Moreover, if we factor out  $e^{-i\phi}$  from the exponential, we recognize that this factor is just a complex constant, and so without loss of generality we can absorb it into the amplitude  $A$ :

$$\psi(x) = A e^{ikx}, \quad (9)$$

where  $A$  is now complex.

Again, this is the formula for a wavefunction with definite momentum, also called a momentum eigenfunction. Specifically, the momentum value corresponding to this function is

$$p_x = \frac{h}{\lambda} = \frac{hk}{2\pi} = \hbar k. \quad (10)$$

Moreover, as I said above, equation 9, with positive  $k$ , describes a particle with positive  $p_x$ . But if we allow  $k$  to be negative, then equations 9 and 10 actually work for negative  $p_x$  as well: the “imaginary” component of  $\psi$  is then a quarter cycle to the left of the “real” part, as desired. In essence, we have now implemented the de Broglie relation in the precise language of wavefunctions, by identifying the periodic function  $Ae^{ip_x x/\hbar}$  as the momentum eigenfunction corresponding to momentum value  $p_x$ .

Momentum eigenfunctions aren’t the only wavefunctions that are complex. Virtually all wavefunctions are complex, with two separate components that (in general) needn’t be related to each other in any particular way:

$$\psi(x) = \text{Re } \psi(x) + i \text{Im } \psi(x), \quad (11)$$

where  $\text{Re}$  and  $\text{Im}$  denote the “real” and “imaginary” parts, respectively. I started with examples of real-valued functions (for which  $\text{Im } \psi = 0$ ) only for simplicity. But I need to go back and explain more carefully how to square a complex-valued wavefunction, so you’ll know how to compute probabilities. The probability density,  $|\psi|^2$ , is actually the square modulus of  $\psi$ , defined as

$$|\psi|^2 = \psi^* \psi = (\text{Re } \psi - i \text{Im } \psi)(\text{Re } \psi + i \text{Im } \psi) = (\text{Re } \psi)^2 + (\text{Im } \psi)^2. \quad (12)$$

Here  $*$  indicates the complex conjugate, in which the “real” part is the same and the “imaginary” part has opposite sign. If you think of  $\psi$  (at a given  $x$ ) as a two-component vector, then  $|\psi|^2$  is just the squared magnitude of that vector (which is always a real, nonnegative number).

Incidentally, I've been putting scare quotes around the words “real” and “imaginary” because neither part of the wavefunction is any more real or imaginary than the other, in the everyday meanings of these words. Both contribute equally to the probability density. Whether wavefunctions are actually real physical entities or merely imaginary human constructs is another question entirely, which I'll mostly leave to the philosophers.

In any case, given an arbitrary complex-valued wavefunction  $\psi(x)$ , you can compute  $|\psi(x)|^2$  and then check whether it's normalized according to equation 2. If it is, then you can integrate it over any desired interval (according to equation 1) to obtain the probability of finding the particle within that interval.

There is a similar formula, which I'll give you later, for finding the probabilities of various outcomes of *momentum* measurements.

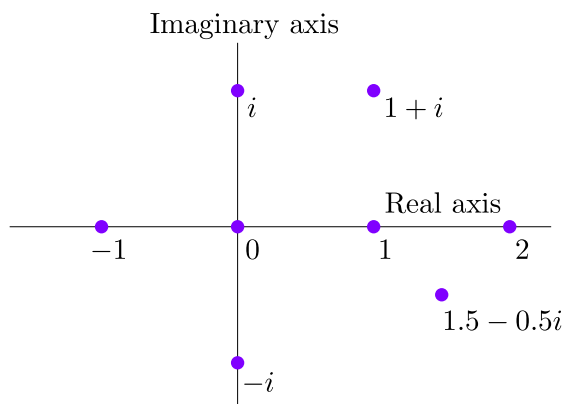
## The effect of a measurement

And now we come to the most important fact about wavefunctions and measurements. Immediately *after* you perform a measurement, the measured quantity (for example, position or momentum) will be well defined, and so the particle's wavefunction must be the corresponding eigenfunction. If you measured the particle's position and found it at  $x = x_0$ , then the particle's wavefunction is now the eigenfunction  $\delta(x - x_0)$  (at least approximately, and never mind the normalization). Or, if you measured the particle's momentum and found it to have  $p = p_0$ , then the particle's wavefunction is now the eigenfunction  $e^{ip_0x/\hbar}$  (at least approximately, and never mind the normalization). Either way, we say that the measurement process causes the wavefunction to *collapse*, in a random way, into whatever eigenfunction corresponds to the measurement outcome.

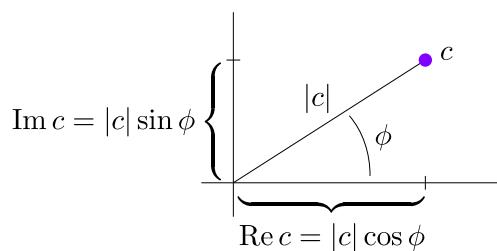
But if a measurement causes the wavefunction to “collapse,” then what exactly constitutes a measurement? Amazingly, quantum mechanics provides no clean answer to this question. Fortunately, though, we won't need a clean answer in order to apply quantum mechanics to laboratory experiments. The messy answer is that a measurement always involves some sort of interaction between the quantum system you're studying and the apparatus you're using to study it. Usually this apparatus is much larger than the system, so you can think of it as a classical (that is, non-quantum) object. Fundamentally, of course, quantum mechanics should apply to everything in your laboratory, including your measurement apparatus. For that matter, quantum mechanics should also apply to you and to everything else in the universe. How, then, should we think about quantum measurements fundamentally? Does the universe have a wavefunction and if so, does it ever “collapse”? These questions are fun to contemplate, but again I will leave them mostly to the philosophers.

## Visualizing complex functions

It's often useful to visualize a complex number as a vector or point in a plane. We usually plot the real part along the horizontal axis and the imaginary part along the vertical axis. Here I've plotted a few examples:



Sometimes, though, it's easier to describe a point  $(x, y)$  using polar coordinates,  $(r, \phi)$ . For a complex number  $c$ , the radial coordinate is just the modulus  $|c|$ , while the angular coordinate  $\phi$  is  $\tan^{-1}[(\text{Im } c)/(\text{Re } c)]$ :

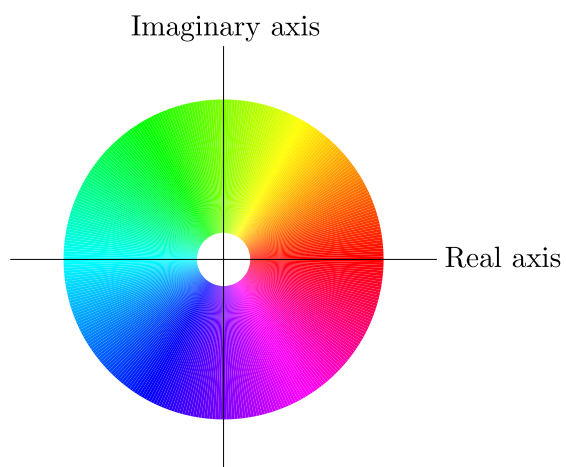


From the diagram you can now see that

$$c = |c|(\cos \phi + i \sin \phi) = |c| e^{i\phi}, \quad (13)$$

a handy expression that's explicitly in terms of the polar coordinates.

One way to visualize a complex-valued function  $\psi(x)$  is to simply plot separate graphs of the real and imaginary parts, as I did above for a momentum eigenfunction  $Ae^{ikx}$ . But another method, which is becoming common in this age of color computer graphics, is to plot a single graph of either  $|\psi|$  or  $|\psi|^2$ , and then fill in the area under the graph with a color that depends on the phase  $\phi(x)$ . We assign colors (or more technically, hues) to the phases starting with red for  $\phi = 0$  (pure real), then blending to yellow, green, cyan, blue, magenta, and back to red, as we go counter-clockwise around the complex plane:



Here is an example of an arbitrary wavefunction plotted with  $|\psi|^2$  on the vertical axis and phase as color, with a plot of the real and imaginary parts underneath for comparison:

