Serial: Counting the Quanta

As promised last time, we will introduce the Schrödinger equation in today's episode. We can start playing with it right away.

Great! Now I'm Lost!

However, before we present the mentioned equation, we must completely get away from classical (Newtonian) mechanics. We need to realize, that it's not just the equation governing the motion of bodies that changes but the quantities describing the motion themselves.

In classical mechanics, there are primarily two quantities: the position and velocity (momentum) of the particle at each point of the trajectory. If we know these two quantities for all particles, this is sufficient to describe their motion completely. The motion is thus described by time functions: $\vec{x}(t), \vec{v}(t)$. The former distinguishes whether the arrow is in the bow or has already been shot at the prey. The second then distinguishes a motionless arrow from a flying one (a problem that troubled Zeno in the ancient world so much that he called it a paradox).

In quantum mechanics, we replace these two quantities with a single function of all spatial coordinates and time, called the wave function. It is usually denoted by the Greek letter $\psi(\vec{x}, t)$. In addition to the classical position and momentum, this wave function can generally take on complex values. Its complexity is not surprising since it encodes information about both the particle's position and momentum. We say "encodes" on purpose because, in quantum mechanics, the precise information about where particles are or how fast they are moving completely disappears. Only probabilities can be decoded, such as the probability of a particle being at a given location.

This probability is given as the square of the absolute value of the wave function

$$P_{\text{particle}\in\vec{x}} \sim |\psi(\vec{x})|^2 \ .$$

We don't write equality here quite deliberately because considering the probability of occurrence at one particular point of space is mathematically tricky (the point occupies an infinitesimally small part of space, and therefore, according to mathematical logic, the probability of occurrence cannot be non-zero). We don't explicitly write the time dependence here for simplicity, but if our wave function evolves over time, this probability will also evolve over time. In technical terms, this quantity is called the probability density

$$\rho(\vec{x}) = |\psi(\vec{x})|^2 \ .$$

If we integrate this density over a region, we get the probability of a particle occurring in that region (similar to how mass is the integral of density over volume)

$$P_{\text{particle}\in(a,b)} = \int_a^b \rho(x) \ \mathrm{d}x \ .$$

This is a one-dimensional example, but it is similar in the multidimensional case. And if you have yet to be introduced to integrals, then know that this just means calculating the area

under some part of the graph. If the probability density were the same at all points in the region of space under consideration, the probability would be the product of the density and the volume.

A quite natural requirement is that we want the overall probability of the particle to be 1 (i.e., the particle is definitely somewhere)

$$P_{\text{Tot}} = \int_{-\infty}^{\infty} \rho(x) \, \mathrm{d}x = \int_{-\infty}^{\infty} |\psi(\vec{x})|^2 \, \mathrm{d}x = 1 \,,$$

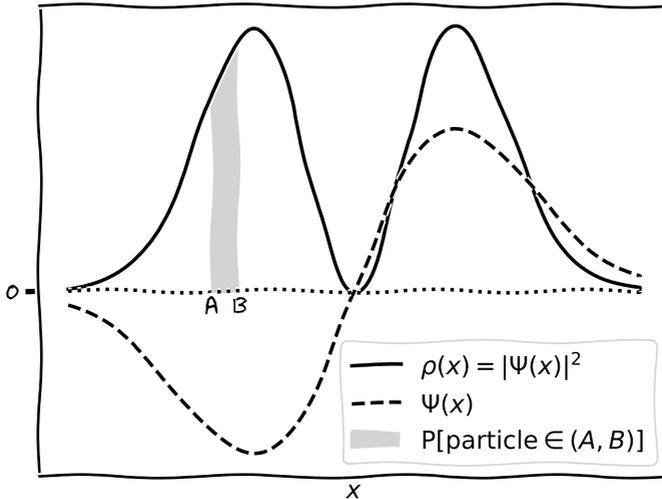which is also called the wave function normalization requirement.

Fig. 1: Wave function (dashed line) and probability density (solid line). The probability that a particle occurs somewhere between a and B is then the area of the grey-shaded region.

The momentum is then encoded by the so-called momentum operator (We'll say what this operator is in a bit, but we mark it with the circumflex symbol, also called "hat".), which is given by the derivative

$$\hat{p} = -i\hbar \frac{\mathrm{d}}{\mathrm{d}x} \,.$$

A derivative is a mathematical operation that replaces the functional value at each point of a function with its slope. Thus, for example, the derivative of a linear function (whose slope does not change) is a constant function. Momentum cannot be interpreted as simply as position, meaning that the momentum operator cannot be used to introduce a simple probability density for momentum. (And a complicated method is beyond the scope of this text.)

But what we can look at (not entirely rigorously) is the validity of Heisenberg's uncertainty principle, i.e., the more precisely we know the position of a particle, the less we know about its momentum. If we imagine that we want to make the wave function as "flatten" as possible while keeping the constant area under the curve so that the uncertainty in position is small; we find that the momentum given by the slope of the curve will be enormous and will change a lot. Conversely, if we want to have small changes in slope, the wave function will "flatten" into space (but a wave function with a particular momentum value must be complex). This is the point of many scientific jokes, for example:

> Heisenberg gets pulled over by a policeman, who asks him:
> "Do you know how fast you were going?"
> "No, but I know exactly where I am."
> "You were doing 95 kmph in a 50."
> "Great! Now I'm lost!"

Now we can answer the fundamental question that was left over from the last part: Why doesn't the electron actually fall into the nucleus? The answer is simple because of the uncertainty principle! If the electron fell directly into the nucleus, it would have the lowest potential energy. However, it would be precisely localized so that the momentum uncertainty would grow beyond all limits. Because of this, the electron would then have infinite kinetic energy. Therefore, the electron prefers states where the uncertainty in position and momentum are balanced.

Finally, the Equation

Now that we have a reasonable description using the wave function $\psi(\vec{x})$ at each time, using the space and time function $\psi(\vec{x}, t)$, nothing prevents us from introducing the time-dependent Schrödinger equation

$$i\hbar\frac{\partial}{\partial t}\psi(\vec{x}, t) = \hat{H}\psi(\vec{x}, t). \tag{1}$$

This cannot be derived in any way from classical physics; Schrödinger himself guessed its shape from the fact that he wanted its solutions for a particle with no external potential to be plane waves. This was again a bit of a lucky guess by de Broglie, as you may remember from the last episode. Eventually, it became clear that this equation could describe the world around us. On the left side of this equation, we have the product of the imaginary unit, the reduced Planck constant, and the partial derivative of the wave function with respect to time. On the right-hand side, we have borrowed the Hamiltonian operator ($\hat{H}$) from theoretical mechanics to describe the system. However, for reasonable systems, it is nothing more than the sum of kinetic and potential energy. Its form in quantum (and without a hat, even in classical) mechanics is simple

$$\hat{H} = \frac{\hat{p}^2}{2m} + V(x) = -\frac{\hbar^2}{2m}\frac{\mathrm{d}^2}{\mathrm{d}x^2} + V(x),$$

in the case of 3D space, we only replace the second derivative with the Laplace operator. Note that this is not an ordinary physical quantity but an operator (hence the hat!). An operator is a mathematical object that takes a function and makes it into another function (a kind of "function on functions"). As we can see, an operator can be, for example, a derivative, multiplication by a constant, a third power, or any other operation we can do on functions.

Specifically, the Hamiltonian takes an arbitrary function (unlike the notation used for the functions, we do not write the function on which the operator acts in parentheses, but simply to the right of the operator), then firstly differentiates it twice, multiplies it by the factor $-\frac{\hbar^2}{2m}$ (this corresponds to the kinetic energy) and finally adds the product of the potential and the original wave function to the result (this corresponds to the potential energy).

The Schrödinger equation is much more complicated than the ones you usually encounter. It is an equation whose solution is not a number, but an entire function that satisfies the equation (such equations are called differential equations).

However, since solving this equation simultaneously for the spatial distribution and the time evolution would be disproportionately difficult, Schrödinger, in introducing it, also thought of finding solutions where the probability density does not change over time (in these states, a quantum system such as atom, for example, will exist). It's actually a similar idea to Bohr's model that only certain trajectories are allowed. Such states can be easily found by looking for solutions in the following form

$$\psi(x,t) = e^{-i\frac{E}{\hbar}t}\psi(x)\,,$$

where E is the total energy of the system. We have thus decomposed the position and time function $\psi(x,t)$ into the product of two functions, one of which depends only on time (the exponential part) and the other only on the position $\psi(x)$. Try for yourself that after substituting the wave function in this form into the time-dependent Schrödinger equation (1) the time dependence is eliminated. This gives an equation for only the spatial part of the wave function in the form

$$\hat{H}\psi(x) = E\psi(x)\,, \tag{2}$$

which is called the timeless Schrödinger equation. The advantage is that we've got rid of the need to solve the time evolution of the wave function, but we pay for it by the fact that our equation contains a second unknown – the energy corresponding to the wave function. We can only get rid of time this easily if the Hamiltonian does not depend explicitly on time. But this is actually not such an exceptional example; on the contrary, it usually happens that we are looking for the motion of a set of particles in some static potential. Even when we calculate the wave function of the entire molecule, we are satisfied with the timeless Schrödinger equation. Conversely, when we calculate how this molecule responds to the oscillating electric field of a powerful laser, we would need a time-dependent equation. (Although even in this case, it is quite often possible to find some tricks to get around this time dependence.)

The timeless Schrödinger equation has an interesting form from a mathematical point of view. We are looking for a function that remains the same after the operator is applied, except for the multiplicative constant. In mathematics, this process is called finding the eigenfunctions of the operator. The multiplicative constant for a given function is referred to as an eigenvalue. To us, eigenvalues are often as important, if not more, than the eigenfunctions themselves. Mathematicians have even borrowed physics terminology, and the set of all eigenvalues is called a spectrum. A trivial example is the "multiply by 7" operator. This operator multiplies a function by seven, so all functions are its eigenfunctions. The eigenvalue on the right-hand side is then, of course, seven. Another example might be the derivative operator. One of its eigenfunctions is e^x since the derivative of an exponential is an exponential. The eigenvalue, in this case, is one. (Another example is e^{2x} with an eigenvalue of 2.)

We can also verify that the Schrödinger equation (2) for a free particle in one-dimensional space (no external forces acting on it, the potential is zero $V(x) = 0$) with momentum p (note that this is a specific value, not an operator) holds for de Broglie waves in the following form

$$\psi(x) = e^{i\frac{px}{\hbar}} \, .$$

(These are indeed waves because if we decompose the complex exponential using Moivre's theorem, we get $\cos(\frac{px}{\hbar}) + i\sin(\frac{px}{\hbar})$ and see that both the real and imaginary parts are periodic, as we would expect from such a wave.)

If we apply the Hamiltonian on this function, we get

$$\hat{H}\psi(x) = \frac{\hat{p}^2}{2m}\psi(x) = -\frac{\hbar^2}{2m}\frac{\mathrm{d}^2 e^{i\frac{px}{\hbar}}}{\mathrm{d}x^2} = \frac{p^2}{2m}e^{i\frac{px}{\hbar}} \, .$$

Perhaps not surprisingly, in this case, the energy is given by the classical formula $E = p^2/2m$. Note that this equation thus has two solutions for any non-negative value of energy. We will see next that once we "confine" the particle with a potential in some finite space, this will not hold, and we will only have some specific "allowed" values of energy. (Remember the Bohr atom again!) It's also worth mentioning that, in this case, the particular solution is given by a single real parameter, the momentum p, which can take on both positive and negative values. Even if we apply the momentum operator on the wave function

$$\hat{p}\psi(x) = -i\hbar\frac{\mathrm{d}e^{i\frac{px}{\hbar}}}{\mathrm{d}x} = pe^{i\frac{px}{\hbar}} = p\psi(x) \, ,$$

we find that the de Broglie wave is an eigenfunction of the momentum operator. In quantum mechanics, this means that it has a specific momentum value. And to satisfy the uncertainty relation, the wave function is uniformly distributed along the real axis.

The simplest system with a potential we can study is a particle confined inside an infinitely deep well, also referred to as a particle in the box. In the one-dimensional case, we choose the interval of the "box" to be $(0, L)$ and the potential of this box to be

$$V(x) = \begin{cases} 0 & \text{if } x \in (0, L) \\ \infty & \text{if } x \in (-\infty, 0) \cup (L, \infty) \end{cases} \, .$$

In this case, we do not even need to solve the Schrödinger equation directly, but we can assume that we know the solution for each part of this equation and join them. When joining, we only need to ensure that the joined solutions have the same energy and that the resulting wave function is continuous.

A particle (even a quantum one) cannot occur in regions with infinite potential, so the wave function there will be identically zero. This leaves only the region $x \in (0, L)$ to solve. But in that region, the particle moves freely, so the solution will be in the form of de Broglie waves. For each energy, we always have two waves differing only in the sign of the momentum (direction of propagation). So let us write the wave function for $x \in (0, L)$ as the sum of two waves of yet undetermined momentum with yet unknown coefficients A, B.

$$\begin{aligned} \psi(x) &= a\ e^{i\frac{px}{\hbar}} + Be^{-i\frac{px}{\hbar}} \\ &= a\ \left(\cos\left(\frac{px}{\hbar}\right) + i\sin\left(\frac{px}{\hbar}\right)\right) + B\left(\cos\left(\frac{px}{\hbar}\right) - i\sin\left(\frac{px}{\hbar}\right)\right), \end{aligned}$$

where we broke down complex exponentials using Euler's formula and the fact that sine is an odd function and cosine is an even function. If we want to satisfy the requirement of the wave function being continuous at $x = 0$ (i.e., we want $\psi(0) = 0$), we only need the sine components, which we ensure by choosing $B = -A$. Now we have to satisfy the requirement that such a modified wave function is continuous at $x = L$

$$\psi(x = L) = 2iA \sin\left(\frac{px}{\hbar}\right) = 0.$$

This cannot be achieved by choosing the coefficient A (we will determine this later to satisfy the normalization condition), but we must find a suitable value of the momentum p. By substituting in $\psi(L) = 0$, we can obtain the requirement from the properties of the sine function

$$\frac{pL}{\hbar} = k\pi,$$

where k is an arbitrary natural number. From this, we determine the admissible values of the momentum, and by plugging it into the original equation, we obtain the resulting wave function on the interval $(0, L)$

$$\psi(x) = \sqrt{\frac{2}{L}} \sin\left(\frac{k\pi}{L}x\right),$$

where we have also chosen the coefficient before the sine to satisfy the normalization condition – you can check for yourself that the following now holds

$$\int_0^L |\psi(x)|^2 \, \mathrm{d}x = 1$$

for all k.

The energies of the states are then obtained by operating the Hamiltonian operator on the wave function.

$$\hat{H}\psi(x) = \frac{k^2\pi^2\hbar^2}{2mL^2}\psi(x) \rightarrow E_k = \frac{k^2\pi^2\hbar^2}{2mL^2}.$$

The result is then a set of eigenstates described by a single parameter k, which is discrete, and its value can only be a natural number. The energy of these states is proportional to k^2.

Note that suddenly not all the energies are allowed. This is related to the particle being confined to a potential in a limited space. Thus, we impose boundary conditions on it. At the same time, like a de Broglie particle behaving as a wave in space, we can describe the confined particle as a standing wave, similar to a string. Even the motion equations are similar to the equations for our wave function. Again, for a string, they eventually lead to the fact that it cannot oscillate at any arbitrary frequency but has its particular harmonic frequencies. Isn't it a wonderful thing how universal physics is?

And if you feel that this is a completely artificial system, know that such a model is quite reasonable in many cases. It can model linear hydrocarbons with many conjugated bonds because electrons from such bonds can move relatively freely along the entire length of this conjugated system (motion in the well), whereas it is very difficult for them to leave the molecule (it is necessary to overcome a high potential).

Also, a simple potential well can be used to describe quantum dots, which are microscopic spheres of a semiconductor that have interesting optical properties, such as fluorescence. Moreover, these properties are not determined by the chemical composition of the sphere but by
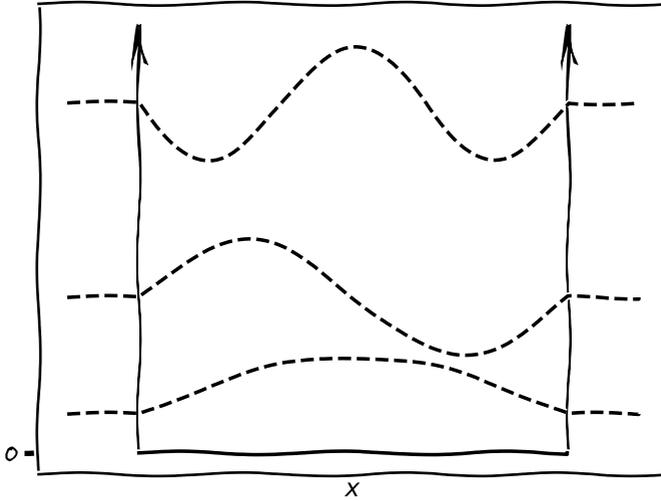
Fig. 2: Figure shows the lowest three states of an infinite potential well. The solid line shows the potential. The wave functions are shown in a way that is commonly used in quantum mechanics – the wave functions are shifted vertically from zero by the value corresponding to that state's energy. (Thus, in this figure, all wave functions outside the well are zero, not just constant as they might appear at first glance.)

its size, as this naturally limits the field of the electrons responsible for conductivity, forcing specific electronic states on them, as above.

Finally, we should mention that the structure of atomic nuclei was described using such a well, even before the nature of the forces acting between protons and neutrons was properly explained (and quite often even afterward because these forces are mathematically much more complex to calculate than the electromagnetic interaction between electrons. Even today, this is a problem at the computational limits of supercomputers).

In conclusion, we will quickly show (without inference) what happens if the potential outside the well is finite. This causes a particle that would previously have a wave function outside the well of zero to suddenly be able to escape, although the probability of this decreases exponentially with distance from the edge of the well. Otherwise, the solution procedure is similar to that of the potential well. We try to fit the wave function in such a way that it is continuous and joins smoothly. Mathematically, this means that we want continuity even in the first derivative, and leads to a set of algebraic equations that are not analytically solvable. (Feel free to try!) At the same time, it's probably not surprising that we can only fit a finite number of states into such a well. The exact number depends on the ratio of the depth to the width of the well, but mathematically it is guaranteed that there must always be at least one. Such states are also called bound states. The rest of the states have energy higher than the well's depth and

essentially resemble a free particle's wave function. (But unlike a free particle, they can, with some probability, bounce at the edge of the well and fly back.)
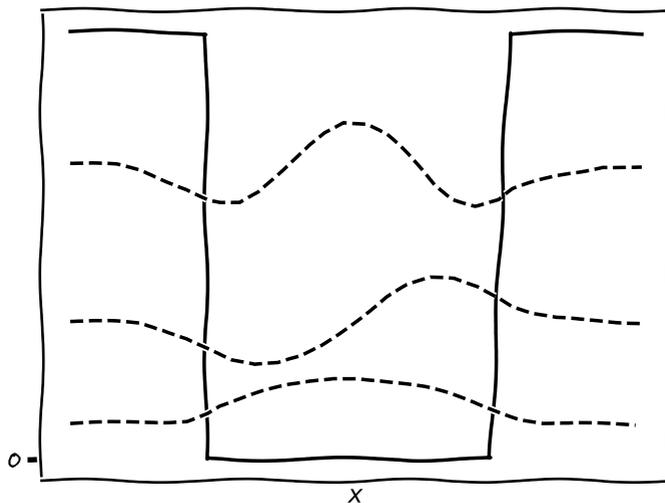

Fig. 3: Wave functions of the lowest states of the finite potential well.

What If There Are More Electrons?

When we solve the Schrödinger equation for multiple particles, the situation is much more difficult; we have one wave function for the whole system, but it's now a function of the positions of all the particles. To give you an idea, for a single particle in three-dimensional space, we have a differential equation for a function of three variables; for two particles, we have six, and so on. This generally deprives us of the chance to solve the problem on paper but also numerically using the computer because the standard approach of representing the function by its value at given grid points converts the problem into a giant system of classical linear equations.

But if we imagine dividing the space in each dimension into 10 points, which is still a very rough division, then for one particle in three-dimensional space, we need a thousand points, which is a piece of cake for a computer. For two particles, a million points, which an ordinary computer can still handle, but three particles would require a billion points, which takes a couple of gigabytes of memory just to store the data. Four can barely fit on a supercomputer, and for five, the data is on the order of petabytes, which is just about the current limits on RAM. Now think about a single atom of uranium with almost a hundred electrons!

Don't despair; we'll show in future episodes that it's not so hopeless if we get smarter about the problem. But today, we will end with a slight simplification that can often give a pretty good physical insight.

If we have a set of particles and neglect the interaction between them, each of them encounters the same potential. We then solve the Schrödinger equation separately for each particle so that all particles have the same eigenstates. Then we occupy the individual states successively with the particles of the lowest energy to satisfy the Pauli exclusion principle, which you've heard about in chemistry. If we have electrons, they come in pairs because they have a spin that can point in one of two directions. The total energy of this system of particles in this approximation is then given by the sum of the energies of the individual particles because they are non-interacting.

After all, the entire periodic table of elements is based on this idea. There, we successively occupy orbitals that are solutions of the Schrödinger equation for a single electron.