

Quantum Mechanics

David Tong

*Department of Applied Mathematics and Theoretical Physics,
Centre for Mathematical Sciences,
Wilberforce Road,
Cambridge, CB3 0BA, UK*

<http://www.damtp.cam.ac.uk/user/tong/quantum.html>

`d.tong@damtp.cam.ac.uk`

Recommended Books and Resources

There is no shortage of excellent textbooks on quantum mechanics. Two good introductory books, at a similar level to this course, are:

- David. J. Griffiths, “*Introduction to Quantum Mechanics*”
- Alasdair Rae, “*Quantum Mechanics*”

Both books excel in providing clear and straightforward explanations and if you’re struggling with the basics are good places to turn. They don’t however do anything sophisticated and won’t help much for later courses.

If you want a book to last a lifetime, rather than just one year, there are also many options. Here are some that I like.

- Shankar, “*Principles of Quantum Mechanics*”

A long and detailed book, but one where the author takes you by the hand and leads you gently through some challenging calculations.

- Landau and Lifshitz Volume 3 “*Quantum Mechanics: Non-Relativistic Theory*”

There’s no hand holding in Landau’s theoretical minimum, just encyclopedic information, explained without fuss. It is, however, remarkably readable.

- Steven Weinberg, “*Lectures on Quantum Mechanics*”

Weinberg is always worth listening to.

Contents

1	Introduction	2
1.1	The Wavefunction	3
1.1.1	Normalisation	6
1.1.2	Superposition	7
1.2	The Schrödinger Equation	9
1.2.1	The Hamiltonian	10
1.2.2	Conservation of Probability	11
1.2.3	The Collapse of the Wavefunction	13
1.3	The Double Slit Experiment	14
2	A Quantum Particle in One Dimension	19
2.1	The Free Particle	20
2.1.1	A Particle on a Circle	22
2.1.2	The Infinite Potential Well	24
2.1.3	The Gaussian Wavepacket	27
2.1.4	A First Look at Expectation Values	30
2.2	The Harmonic Oscillator	34
2.2.1	The Energy Spectrum	35
2.2.2	The Wavefunctions	37
2.3	Bound States	41
2.3.1	A Finite Potential Well	42
2.3.2	A Delta Function Potential	48
2.3.3	Some General Results	49
2.3.4	The Double Well Potential	52
2.4	Scattering	54
2.4.1	A Step Potential	56
2.4.2	Tunnelling	60
3	The Formalism of Quantum Mechanics	63
3.1	States	63
3.1.1	The Inner Product	64
3.2	Operators and Observables	66
3.2.1	Eigenfunctions and Eigenvalues	68
3.2.2	Hermitian Operators	71
3.2.3	Momentum Eigenstates and the Fourier Transform	75

3.3	Measurement	76
3.3.1	Expectation Values	79
3.4	Commutation Relations	81
3.4.1	The Heisenberg Uncertainty Principle	84
3.5	Interpretations of Quantum Mechanics	88
3.5.1	Hidden Variables	89
3.5.2	Copenhagen and Many Worlds	91
3.5.3	Shut Up and Calculate	93
4	A Quantum Particle in Three Dimensions	95
4.1	Angular Momentum	96
4.1.1	Angular Momentum Commutation Relations	96
4.1.2	The Eigenfunctions are Spherical Harmonics	99
4.2	Solving the 3d Schrödinger Equation	104
4.2.1	The 3d Harmonic Oscillator	105
4.2.2	Bound States in 3d	108
4.2.3	Why Can't the Wavefunction Diverge at the Origin?	110
4.3	The Hydrogen Atom	111
4.3.1	The Energy Spectrum	112
4.3.2	The Wavefunctions	116
4.4	A Pre-Quantum Quantum History	118
4.4.1	The Bohr Model	119
4.4.2	What About the Photon?	121
4.5	A First Look at Renormalisation	122
4.5.1	A Delta Function Potential in the Plane	124
4.5.2	Renormalisation in Quantum Mechanics	127

1 Introduction

Without wishing to overstate the case, the discovery of quantum mechanics is the single greatest achievement in the history of human civilisation.

Quantum mechanics is an outrageous departure from our classical, comforting, common sense view of the world. It is more baffling and disturbing than anything dreamt up by science fiction writers. And yet it is undoubtably the correct description of the universe we inhabit, providing insights into many aspects of the world around us.

Among the many successes of quantum mechanics are answers to old and very basic questions. Like why is matter stable? And why does the sun shine and why is it yellow? And why are all solids either conductors or insulators? But quantum mechanics also opens up vistas that we didn't previously know existed, from novel states of matter where the constituent particles become so entangled that they can be coaxed to perform seemingly impossible tasks, to the subatomic world of fluctuating ethereal fields, to a revised understanding of the meaning of information and what one can achieve with it.

Although quantum mechanics gives the right answers to all questions, it does so at a price. The answers are always statistical in nature. There are few certainties in the quantum world. Of course, there are few certainties in the classical world too, but we can always strive to eliminate them. In classical physics, knowledge is power: the more you know, the better off and less uncertain you are. This is because classical probabilities are always about our ignorance and any classical system that appears random does so because it is somehow difficult, but never impossible, for us to know what's going on inside.

This is not the way of the quantum world. The randomness is real, the unpredictability inherent. There is no way to game the system or gain some illicit knowledge of underlying quantum properties that would reduce the uncertainty. If, for example, a particle has probability $\frac{1}{2}$ to be in one place and probability $\frac{1}{2}$ to be elsewhere, this may well be because the particle really is in both places at the same time. Any attempt to eliminate the quantum certainty will simply shift it, like a bubble in wallpaper, elsewhere. As we will see, quantum particles are fragile objects and the very act of looking changes them, disturbing many of the other delicate properties that they possess.

The purpose of these lectures is to begin to get to grips with the quantum world and understand some of the strange behaviour that it exhibits. As should be clear already, this is something of a step into the unknown and our classical intuition will not be a great guide in the quantum realm. Fortunately we have a better guide because, for

reasons that we do not fully understand, the quantum world is described astonishingly well by the language of mathematics. Our strategy in these lectures, and later ones, will be to embrace this mathematical description and leverage it to build a new intuition for how the universe really works.

1.1 The Wavefunction

The difference between quantum and classical mechanics does not involve just a small tweak. Instead it is a root and branch overhaul of the entire framework.

This is manifest from the very beginning as we can see by comparing how we describe the *state* of a system in the two frameworks. The state is the information that tells us all we need to know about the system at a fixed time, with the idea that the laws of physics will then dictate how the state evolves at all later times. Throughout these lectures we will deal only with the dynamics of a single particle. For the present discussion, we'll think about a particle moving in \mathbf{R}^3 .

In the classical world, the state of the particle is determined by its position \mathbf{x} and its velocity $\mathbf{v} = \dot{\mathbf{x}}$. If you specify both bits of information at some time t_0 then you can use the equation of motion $\mathbf{F} = m\ddot{\mathbf{x}}$ to determine $\mathbf{x}(t)$ and $\mathbf{v}(t)$ for all time. Importantly, it's not enough to just know only, say, the position of the particle at $t = t_0$. You need both $\mathbf{x}(t_0)$ and $\dot{\mathbf{x}}(t_0)$. Mathematically, this is because the equation of motion is a second order differential equation and so you need to specify two integration constants to get a unique solution.

In the quantum world, the state of a particle is determined by its *wavefunction*. This is a complex valued function

$$\psi(\mathbf{x}, t)$$

As we will see, if you know the wavefunction at some time, say t_0 , then you have all the information that you need to determine the state at all other times.

The description in terms of the wavefunction is not a small amendment to classical mechanics. We've replaced the three position coordinates $\mathbf{x} \in \mathbf{R}^3$ with an infinite amount of information, a functions worth of information. Moreover, we haven't specified anything about the particle's velocity; that information must also be contained, in some manner, in the wavefunction $\psi(\mathbf{x}, t)$.

The wavefunction has a very simple interpretation. Or, more precisely, the mod-square of the wavefunction has a very simple interpretation. It tells us the probability that we will find a particle at a given position. The probability density P for a particle to sit at point \mathbf{x} is

$$P(\mathbf{x}, t) = |\psi(\mathbf{x}, t)|^2$$

This is known as the *Born rule*, after Max Born who first understood that quantum mechanics is a theory of probability rather than certainty.

From the probability density, you can compute actual probabilities by multiplying by a volume: the probability that the particle sits in some small volume dV centred around point \mathbf{x} is $P(\mathbf{x}, t)dV$.

In all other realms of science, probability arises because of our ignorance. If you throw a classical dice, and know its initial state with complete precision, then there is no doubt about what will happen. Once the dice leaves your hand, the probability that you will roll a six is either 0 or 1. If you were really good at solving differential equations, you could just figure out the answer and impress your friends. But, in most circumstances we don't have a good knowledge of the initial state of the dice and, besides, differential equations are hard. So we just give up, admit that we don't know what will happen, and say that the probability of rolling a six is $\frac{1}{6}$. Crucially, however, the introduction of probability is entirely due to our lack of knowledge; it is not an inherent property of the dice.

This is not the case in quantum mechanics. The state $\psi(\mathbf{x}, t)$ contains all the information about a particle and the probabilistic interpretation is not because of any failing on our part, but instead due to an inherent randomness in the quantum world. This is, to put it mildly, something of a departure from the Newtonian way of thinking.

The novelty in the wavefunction description of the particle might suggest other interpretations of the function $\psi(\mathbf{x}, t)$. You might, for example, wonder if perhaps we shouldn't think of a particle at all, but rather some fluid-like object, spread out over space. Such objects are commonplace in physics and are known as *fields*. The electric and magnetic fields are familiar examples. But that's *not* the right way to think about the wavefunction. This is because we never observe a fragmented particle, one that's lost its particle form and starts spreading jell-o like throughout space. Any measuring device that will allow you to determine whether a particle is in some designated region of space will return the answer yes or no. It won't tell you "well, the particle's a little bit here but mostly somewhere over there to the left". The wavefunction doesn't describe some real fluid; it is more nebulous. It is a wave only of probability.

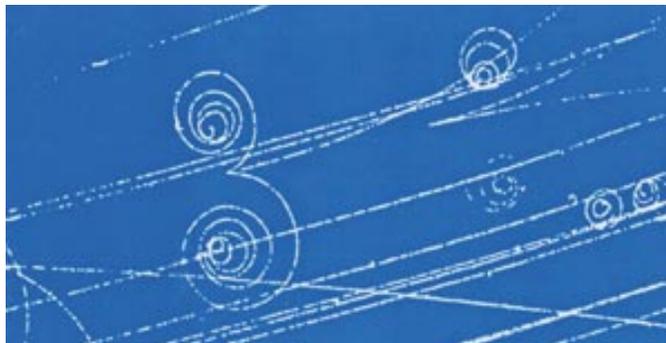


Figure 1. In any particular experiment, you only detect particles with very definite trajectories and no hint of the more nebulous underlying wavefunction.

This is illustrated in Figure 1 which shows the tracks left by electrons and positrons as they pass through a detector known as a bubble chamber. The fast travelling particles move in approximately straight lines, while those that are slower spiral in circles due to an applied magnetic field. The electrons bend one way, the positrons the other, giving the back-to-back spirals that you can see. For our purposes, the key point is that when the particles entered the detector, they were described by a wavefunction that was spread out over a large part of space. Yet, the particles don't appear as fluffy, insubstantial clouds of probability. Instead, they leave clear tracks, with a definite trajectory.

The introduction of probability at such a fundamental level means that we must abandon the idea of predicting, with certainty, what will happen in a given experiment. There is no way to say when and where the spirals will appear in the picture above. We can only compute the likelihood for this to happen. Clearly this is retrograde step from the Newtonian (strictly, Laplacian) dream of knowing everything that will happen, from the beginning to the end of time.

One might wonder if perhaps it's possible to restore this dream at some deeper level. Maybe the wavefunction $\psi(\mathbf{x}, t)$ is just some coarse-grained description of something finer and more subtle that's going on underneath. Maybe another layer of reality awaits discovery, one that allows us to bring back our comfortable deterministic, Newtonian way of thinking. There is good reason to think that this is not the case. Reality is, at heart, probabilistic. We'll mention some of the arguments for this in Section 3.5 where we touch upon some of the interpretations of quantum mechanics. A fuller description of why a deterministic world is incompatible with experiment will be given in the the lectures on [Topics in Quantum Mechanics](#) (see, in particular, the section on Quantum

Foundations). For now, you'll just have to accept that the world in which we live is not the one our forefathers grew up believing in. It is random and uncertain. It is also, as we will see, considerably more interesting because of it.

1.1.1 Normalisation

The probabilistic interpretation means that the wavefunction can't just be any old function. The particle has to be somewhere, and this translates to the requirement

$$\int d^3x P(\mathbf{x}, t) = \int d^3x |\psi(\mathbf{x}, t)|^2 = 1$$

Wavefunctions that have this property are said to be *normalised*.

In practice, this isn't such a big deal. Suppose that we have a wavefunction $\Psi(\mathbf{x}, t)$ that isn't normalised but instead obeys

$$\int d^3x |\Psi(\mathbf{x}, t)|^2 = N < \infty$$

If N is finite we say that the wavefunction is *normalisable*. Clearly a wavefunction is normalisable only if $\psi(\mathbf{x}, t) \rightarrow 0$ sufficiently fast as $|\mathbf{x}| \rightarrow \infty$. In this case, we can always construct a normalised wavefunction

$$\psi(\mathbf{x}, t) = \frac{1}{\sqrt{N}} \Psi(\mathbf{x}, t)$$

Quite often in these lectures, it will turn out to be useful to work with un-normalised wavefunctions Ψ and then remember to include the normalisation factor only at the end when computing probabilities.

From the discussion above, it should be clear that we will have little use for wavefunctions that cannot be normalised because

$$\int d^3x |\psi(\mathbf{x}, t)|^2 = \infty$$

These have no probabilistic interpretation and should be discarded. They do not describe quantum states. (I should warn you that this statement, while mostly true, comes with small and annoying caveat that will rear its head fairly soon. We will address this in Section [2.1](#).)

There is one other relation between wavefunctions that is important: two wavefunctions that differ by a constant, complex phase should actually be viewed as describing equivalent states

$$\psi(\mathbf{x}, t) \equiv e^{i\alpha}\psi(\mathbf{x}, t)$$

for any constant, real α . Note, in particular, that this doesn't change the probability distribution $P = |\psi|^2$. Nor, it will turn out, does it change anything else either. (The "anything else" argument is important here. As we'll see later, if you multiplied the wavefunction by a spatially varying phase $e^{i\alpha(\mathbf{x})}$ then it doesn't change the probability distribution $P = |\psi|^2$ but it does change other observable quantities and so multiplying by such a factor does not give back the same state.)

Combining the need for normalisation, together with the phase ambiguity, it is sometimes useful think of states as the collection of normalisable, complex functions with the equivalence relation

$$\psi(\mathbf{x}, t) \equiv \lambda\psi(\mathbf{x}, t)$$

for any complex $\lambda \in \mathbf{C}$ with $\lambda \neq 0$. The wavefunctions ψ and $\lambda\psi$ should be viewed as describing the same physical state.

1.1.2 Superposition

The set of wavefunctions form a *vector space*. This means that if $\psi_1(\mathbf{x}, t)$ and $\psi_2(\mathbf{x}, t)$ are both possible states of the system, then so too is

$$\psi_3(\mathbf{x}, t) = \alpha\psi_1(\mathbf{x}, t) + \beta\psi_2(\mathbf{x}, t)$$

for any $\alpha, \beta \in \mathbf{C}$. This is known as the *principle of superposition*.

The mathematics underlying this statement is simple. If ψ_1 and ψ_2 are states of the system then both must be normalisable. They may, indeed, be normalised but for now let's just assume that

$$\int d^3x |\psi_i(\mathbf{x}, t)|^2 = N_i < \infty \quad \text{for } i = 1, 2$$

Then ψ_3 is also a possible state of the system provided that it too is normalisable. But it is straightforward to show that it has this property. First we have

$$\int d^3x |\psi_3|^2 = \int d^3x |\alpha\psi_1 + \beta\psi_2|^2 \leq \int d^3x \left(|\alpha\psi_1| + |\beta\psi_2| \right)^2$$

where we've used the triangle inequality $|z_1 + z_2| \leq |z_1| + |z_2|$ for any $z_1, z_2 \in \mathbf{C}$. Continuing, we have

$$\int d^3x |\psi_3|^2 \leq \int d^3x \left(|\alpha\psi_1|^2 + |\beta\psi_2|^2 + 2|\alpha\psi_1||\beta\psi_2| \right) \leq \int d^3x \left(2|\alpha\psi_1|^2 + 2|\beta\psi_2|^2 \right)$$

where, now, in the last step we've used the fact that $(|z_1| - |z_2|)^2 \geq 0$ which, rearranging, gives $2|z_1||z_2| \leq |z_1|^2 + |z_2|^2$. So, finally, we get the result we wanted

$$\int d^3x |\psi_3|^2 \leq 2|\alpha|^2 N_1 + 2|\beta|^2 N_2 < \infty$$

We learn that ψ_3 is normalisable, and hence also an allowed state of the system.

The idea that functions form a vector space might be novel. There is a simple notational trick that should help convince you that it's not too far from things you've seen already. If we have some n -dimensional vector \vec{y} , then we often use subscript notation and write it as y_i with $i = 1, \dots, N$. We could equally well write it as $y(i)$ with $i = 1, \dots, N$. A function $y(x)$ is a similar kind of object, but now with a continuous label $x \in \mathbf{R}$ rather than the discrete label $i = 1, \dots, N$.

Of course, the big difference is that we're now dealing with an infinite dimensional vector space rather than a finite dimensional vector space, and I would be lying if I told you that this doesn't bring new issues to the table. Indeed, we've already met one of them above: we don't consider any old function $\psi(\mathbf{x})$ to be in our vector space, but rather only normalisable functions. There are many further subtleties in store but we will be best served by simply ignoring them at this point. We'll return to some issues in Section 3.

The principle of superposition has profound physical consequences. Suppose that you have a particle that, at some time t_0 , you know is localised somewhere near the point \mathbf{X} . For example, we could describe it by the Gaussian wavefunction

$$\psi(\mathbf{x}) = \frac{1}{\sqrt{N}} e^{-a(\mathbf{x}-\mathbf{X})^2}$$

for some choice of a that tells you how spread out the wavefunction is. Here N is a normalisation factor that won't concern us. (For what it's worth, we should take $N = (\pi/2a)^{3/2}$ if we want to ensure that ψ is normalised.) This is a state in which you might imagine that the particle still retains something of its classical particle nature, at least in the sense that the probability distribution is not spread out over long distances.

However, the superposition principle tells us that we should also entertain the idea that the particle can sit in a state

$$\psi(\mathbf{x}) = \frac{1}{\sqrt{N'}} \left(e^{-a(\mathbf{x}-\mathbf{X}_1)^2} + e^{-a(\mathbf{x}-\mathbf{X}_2)^2} \right)$$

for arbitrary positions \mathbf{X}_1 and \mathbf{X}_2 . But now the cat is really out of the bag! The interpretation of this state is that the particle has somehow split and now sits both near \mathbf{X}_1 and near \mathbf{X}_2 . Indeed, we'll shortly see clear experimental consequences of states like the one above, where elementary particles – which are, as far as we can tell, indivisible – are coaxed into travelling along two or more different paths simultaneously. Taken to the logical extreme, it is states like the one above that lead us to seemingly paradoxical situations with cats that are both alive and dead.

1.2 The Schrödinger Equation

The wavefunction gives us a description of the state of the system. The next step is to understand how these states evolve in time. This is described by the *Schrödinger equation*,

$$i\hbar \frac{\partial \psi}{\partial t} = \hat{H} \psi \quad (1.1)$$

Here \hbar is *Planck's constant*, a fundamental constant of nature. Its value is

$$\hbar \approx 1.05 \times 10^{-34} \text{ Js}$$

In the quantum world, the energy unit of choice is the electronvolt, defined by the kinetic energy that an electron picks up when accelerated over a potential difference of 1 volt. It is related to the Joule as $1 \text{ eV} \approx 1.6 \times 10^{-19} \text{ J}$. In these units, we have

$$\hbar = 6.58 \times 10^{-16} \text{ eV s}$$

The numerical value of \hbar is important for specific applications. But, conceptually, the units that \hbar carries are more important. It has dimension of energy \times time. This is the same dimension as angular momentum $\mathbf{L} = m\dot{\mathbf{x}} \times \mathbf{x}$. It is also the same dimension as the action in classical mechanics, $S = \int dt \frac{1}{2}m\dot{\mathbf{x}}^2 - V$.

A slightly old-fashioned convention is to define Planck's constant to be $h = 2\pi\hbar$. This is bad practice. First, \hbar appears in equations much more often than the combination $2\pi\hbar$. But, more importantly, the quantity \hbar is used exclusively for Planck's constant and its appearance in any equation is jumping up and down and screaming at you that quantum mechanics is at play. Meanwhile, h can mean many things in physics, from “height” to “magnetic field density”. In these lectures we'll use only \hbar .

1.2.1 The Hamiltonian

The next object in the Schrödinger equation that needs explanation is \hat{H} . This is known as the *Hamiltonian*. Different choices of Hamiltonian describe different laws of physics. In these lectures, we will restrict ourselves to non-relativistic particles moving in three dimensions in the presence of a potential energy $V(\mathbf{x})$. For such a system, the Hamiltonian is given by

$$\hat{H} = -\frac{\hbar^2}{2m}\nabla^2 + V(\mathbf{x}) \quad (1.2)$$

This is an example of a *differential operator*. It takes a function $\psi(\mathbf{x})$ and gives you back a new function, but only after differentiating by the Laplacian

$$\nabla^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2}$$

For much of these lectures, we will denote operators with a hat, like \hat{H} . I should mention, however, that this notation is a little like training wheels on a bike and in future lecture courses hats will be discarded and you just have to remember what objects are operators.

A few comments are in order about the Schrödinger equation and Hamiltonian. First, the Schrödinger equation replaces the venerable $\mathbf{F} = m\mathbf{a}$ in classical mechanics. Importantly, the Schrödinger equation is first order in time, rather than second order. This, ultimately, is why the state of a quantum system is described by $\psi(\mathbf{x})$ and not both $\psi(\mathbf{x})$ and $\dot{\psi}(\mathbf{x})$.

Next, the “Hamiltonian” is not something entirely novel to quantum physics. It is actually a concept that arises in more sophisticated presentations of Newtonian physics. You can read about this approach in the lectures on [Classical Dynamics](#). In the classical world, the Hamiltonian is closely related to the energy of the system which, for us, is

$$E = \frac{1}{2m}\mathbf{p}^2 + V(\mathbf{x}) \quad (1.3)$$

where $\mathbf{p} = m\dot{\mathbf{x}}$ is the momentum of the particle. As these lectures proceed, we will also see the connection between the operator Hamiltonian \hat{H} and energy. For now, we can just note that there is clearly a similarity in structure between the equations (1.2) and (1.3). If we wanted to push this analogy further, we might postulate some relationship in which momentum is replaced by a derivative operator,

$$\mathbf{p} \rightarrow \pm i\hbar\nabla \quad ?$$

where it's not clear whether we should take $+$ or $-$ since the energy depends only on \mathbf{p}^2 . More importantly, it's not at all clear what a relationship of this kind would mean! How could the momentum of a particle – which, after all, is just a number — possibly be related to something as abstract as taking a derivative? As these lectures proceed, we'll see the connection between momentum and derivatives more clearly. (We'll also see that we should actually take the minus sign!)

It turns out that not all classical theories can be written using a Hamiltonian. Roughly speaking, only those theories that have conservation of energy can be formulated in this way. Importantly, the same is true of quantum mechanics. This means, in particular, that there is no Schrödinger equation for theories with friction. In many ways that's no big loss. The friction forces that dominate our world are not fundamental, but the result of interactions between many (say 10^{23}) atoms. There are no such friction forces in the atomic and subatomic worlds and, moreover, the formalism of quantum mechanics does not allow us to easily incorporate such forces.

(If, for some reason, you really do want to consider quantum friction then you're obliged to include the direct coupling to the 10^{23} other atoms and track what happens. There are interesting applications of this, not least an idea known as *decoherence*.)

One final comment: as we proceed to more advanced physical theories, whether relativistic quantum field theory or even more speculative endeavours like string theory, the Schrödinger equation (1.1) remains the correct description for the evolution of the quantum state. All that changes is the Hamiltonian (1.2), which is replaced by something more complicated, as too are the wavefunctions on which the Hamiltonian acts.

1.2.2 Conservation of Probability

We made a big deal about insisting that our wavefunctions are normalised so that they have a probability interpretation. But now we have the Schrödinger equation that tells us how the wavefunction evolves in time and we should make sure that this does not mess up a previously normalised wavefunction.

To this end, let's look at how the probability density changes in time. We have $P = |\psi|^2$ and so,

$$\frac{\partial P}{\partial t} = \psi^* \frac{\partial \psi}{\partial t} + \frac{\partial \psi^*}{\partial t} \psi \quad (1.4)$$

The Schrödinger equation (1.1) and its complex conjugate give

$$\frac{\partial \psi}{\partial t} = -\frac{i}{\hbar} \left(-\frac{\hbar^2}{2m} \nabla^2 \psi + V \psi \right) \quad \text{and} \quad \frac{\partial \psi^*}{\partial t} = \frac{i}{\hbar} \left(-\frac{\hbar^2}{2m} \nabla^2 \psi^* + V \psi^* \right)$$

We see that the terms with the potential V cancel out when inserted in (1.4) and we're left with

$$\begin{aligned}\frac{\partial P}{\partial t} &= \frac{i\hbar}{2m} (\psi^* \nabla^2 \psi - \psi \nabla^2 \psi^*) \\ &= \frac{i\hbar}{2m} \nabla \cdot (\psi^* \nabla \psi - \psi \nabla \psi^*)\end{aligned}$$

We write this as

$$\frac{\partial P}{\partial t} + \nabla \cdot \mathbf{J} = 0 \quad (1.5)$$

where \mathbf{J} is known as the *probability current* and is given by

$$\mathbf{J} = -\frac{i\hbar}{2m} (\psi^* \nabla \psi - \psi \nabla \psi^*) \quad (1.6)$$

Equations of the form (1.5) have a special place in the laws of physics. They are known as *continuity equations* and arises whenever there is a conserved quantity. (For example, the same equation appears in [Electromagnetism](#) where it relates the charge density to the electric current and ensures the conservation of electric charge.)

To see that the continuity equation implies a conservation law, consider some region $V \subset \mathbf{R}^3$ with boundary S . We then integrate to get the probability P_V that the particle lies somewhere in this region.

$$P_V(t) = \int_V d^3x P(\mathbf{x}, t)$$

The continuity equation tells us that this probability changes as

$$\frac{\partial P_V}{\partial t} = - \int_V d^3x \nabla \cdot \mathbf{J} = - \int_S d\mathbf{S} \cdot \mathbf{J} \quad (1.7)$$

where, in the final step, we have used the divergence theorem. We see that that the probability that the particle lies in V can change, but only if there is a flow of probability through the surface S that bounds V . If we know for sure that $\mathbf{J} = 0$ everywhere on the surface S , then the probability that the particle is in the region V doesn't change: $\partial P_V / \partial t = 0$.

The importance of the continuity equation is that it doesn't just tell us that some quantity is conserved: it tells us that the object is conserved *locally*. In the present context, this quantity is probability. The probability density can't just vanish in one region of space, to reappear in some far flung region of the universe. Instead, the evolution of the probability density is local. If it does change in some region of space then it's because it has moved into a neighbouring region. The current \mathbf{J} tells you how it does this.

As a special case, if we take the region V to be all of space, so $V = \mathbf{R}^3$, then

$$P_{\text{anywhere}} = \int d^3x |\psi(\mathbf{x}, t)|^2$$

As we've already seen, $P_{\text{anywhere}} = 1$. The change in the probability is, using (1.7),

$$\frac{\partial P_{\text{anywhere}}}{\partial t} = - \int_{S_\infty^2} d\mathbf{S} \cdot \mathbf{J}$$

where S_∞^2 is the asymptotic 2-sphere at the infinity of \mathbf{R}^3 . Any normalised wavefunction must have $\psi \rightarrow 0$ as $|\mathbf{x}| \rightarrow \infty$ and, for all but the most pathological wavefunctions, this ensures that the current (1.6) also decays suitably quickly, $\mathbf{J} \rightarrow 0$ as $|\mathbf{x}| \rightarrow \infty$, so that

$$\frac{\partial P_{\text{anywhere}}}{\partial t} = 0$$

This is the statement that if the wavefunction is normalised at some time, then it remains normalised for all time.

1.2.3 The Collapse of the Wavefunction

There is one situation where the Schrödinger equation (1.1) does not describe the evolution of a quantum state. This occurs when we do a measurement.

As an example, suppose that at time t_0 we have a particle sitting in a superposition of two different positions far well-separated positions, \mathbf{X}_1 and \mathbf{X}_2 , something like

$$\psi(\mathbf{x}, t_0) = \frac{1}{\sqrt{N'}} \left(e^{-a(\mathbf{x}-\mathbf{X}_1)^2} + e^{-a(\mathbf{x}-\mathbf{X}_2)^2} \right)$$

Left to its own devices, this state will evolve through the Schrödinger equation. But suppose instead that we chose to do a measurement of the particle. For example, we can put a detector in the vicinity of \mathbf{X}_1 and see if it registers.

Suppose that the detector does indeed see the particle. Then we know for sure that the particle is sitting near \mathbf{X}_1 and that means that it can't possibly be elsewhere. The probability that it sits near \mathbf{X}_2 must then vanish! Correspondingly, after the measurement, the wavefunction can't have any support near \mathbf{X}_2 . Instead, at time t_+ immediately after the measurement the wavefunction must jump to something localised near \mathbf{X}_1 , like

$$\psi(\mathbf{x}, t_+) = \frac{1}{\sqrt{N}} e^{-a(\mathbf{x}-\mathbf{X}_1)^2}$$

This is known as the *collapse of the wavefunction*.

There is something more than a little disconcerting about this aspect of quantum mechanics. Just a few pages ago, I told you that the probability in the wavefunction was not due to our ignorance, but instead to an inherent randomness of the system. Yet now we see that once we gain additional knowledge about the system – like where the particle actually is – the wavefunction necessarily changes. Which makes it sound like the wavefunction is describing some aspect of what we know after all!

The collapse of the wavefunction also demonstrates that our gain in knowledge is far from innocent. There is no way to learn something about a quantum system without perturbing that system, sometimes violently so. In the discussion above, it may appear that the collapse of the wavefunction only increases knowledge since we now have a better understanding of the particle’s position. However, the wavefunction contains other, more subtle information and this too has been disturbed by the collapse which, typically, will lead to an increase of uncertainty in other properties.

Furthermore, the collapse of the wavefunction doesn’t happen locally. Now the probability density really does just vanish from one region of space and re-emerge in the region where we detected the particle. You might worry that this will lead to contradiction with causality or with special relativity. It turns out that it doesn’t but we won’t get to fully understand why until later lectures. (See, for example, the section on Bell’s inequality in the lectures on [Topics in Quantum Mechanics](#).)

The fact that the wavefunction can evolve in two very different ways – the first through the smooth development of the Schrödinger equation, the second from the abrupt collapse after measurement – is one of the most unsettling aspects of quantum mechanics. It has led to much handwringing and philosophising and gnashing of teeth and mumbling of words like “ontic” and “epistemic”, none of which seems to be quite enough to calm the nerves. We’ll say more about these issues in Section 3.5 when we discuss the different interpretations of quantum mechanics.

1.3 The Double Slit Experiment

The double slit set-up is a simple but iconic experiment that vividly demonstrates some of the stranger aspects of quantum mechanics.

To set the scene, we first describe two experiments from classical physics. The set-up in both is the same: there is a wall, with two openings in it. These are the eponymous double slits. Behind the wall is a detector. You stand in front of the wall and throw things. The only difference between the two experiments lies in what you throw.

First, we throw particles. These are boring classical particles, like tennis balls. You throw randomly in the direction of the wall. Some simply bounce off the wall, while others get lucky and go through the slits. Then you measure the distribution of balls seen in the detector.

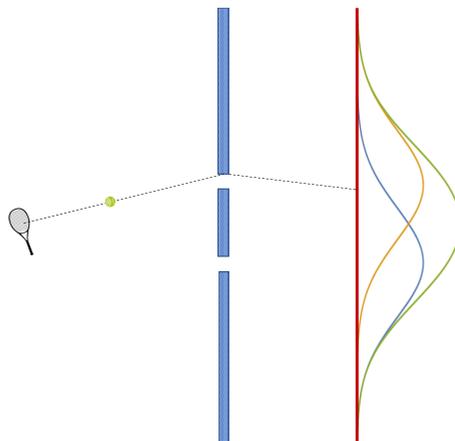


Figure 2. The distribution for classical particles.

In this case, things are nice and simple. Suppose that you close the bottom slit, keeping only the top slit open, and then measure the distribution of balls reaching the detector. The resulting distribution looks something like the yellow curve shown in Figure 2. Call this P_1 . Similarly, if you close the top slit, but keep the bottom slit open then you measure the same distribution just shifted downwards. This is the blue curve in the figure. Now open both slits. What you see is the sum of the previous distributions,

$$P_{12} = P_1 + P_2$$

This is the green curve in the figure. So far, so good.

Our second experiment is the same, but now with waves rather than particles. For our purposes, these can be water waves, sound waves or light waves. For any kind of wave, there is a new ingredient that dramatically changes the outcome: this is interference. Here we give a broad brush description of the physics.

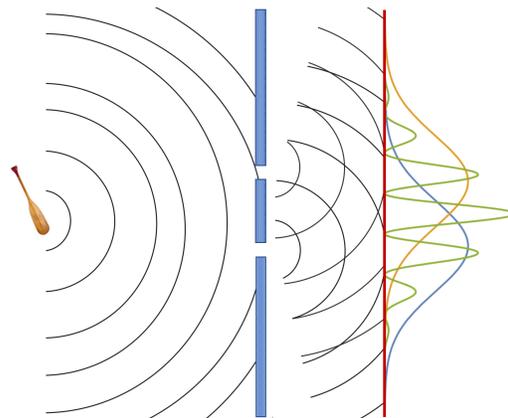


Figure 3. The distribution for classical waves.

First, we open one slit and not the other. The waves are diffracted at the slit and propagate out radially towards the detector where we can measure their intensity. The end result is very similar to that of a particle: a distribution spread out, peaked around the position of the first slit. This is shown in the yellow

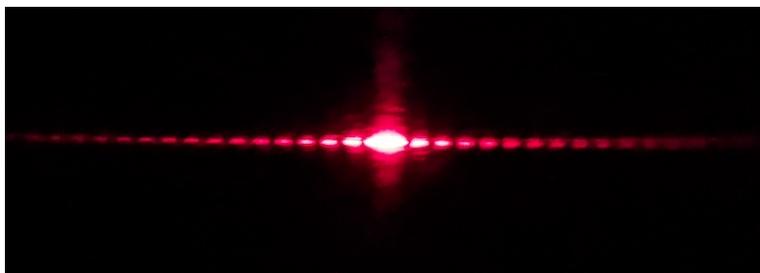


Figure 4. The interference of light waves through a double slit taken from [this website](#).

and blue curves in Figure 3, corresponding to keeping only the top, or only the bottom, slit open respectively.

The real difference comes when we open both slits. The waves pass through along two paths and, when they recombine on the other side, add constructively or destructively. The result is the distinctive interference pattern shown in the green curve. Crucially, at some positions the intensity when both slits are open can be less than when any single slit is open. This happens, for example, at positions where the two waves cancel themselves out completely so the intensity on the screen vanishes. An interference pattern of this kind can be seen vividly using lasers, as shown in Figure 4.

With this classical intuition in place, we now turn to the quantum world. We will perform the double slit experiment one last time, now with quantum particles, say electrons. For appropriately sized slits, the result is a combination of the two classical experiments described above. First, when we detect an electron it appears just like a classical particle would appear. That means that each time we play the game, the electron appears as a point on the detector, rather than something continuous and smeared out like a wave. However, after playing the game many times, the probability distribution seen on the screen agrees with that of a classical wave, exhibiting an interference pattern. An example is shown in Figure 5.

Perhaps the most striking aspect of this result is that there are certain places on the screen where fewer electrons are detected when both slits are open than when any one slit is open. There's no way to explain this from a classical worldview where the particle follows a trajectory that goes *either* through the first slit *or* through the second. It must be that each particle somehow knows about the presence of both slits.

We can understand this behaviour from what little we know about quantum mechanics. First, when one slit is open, the particle is described by a wavefunction $\psi_1(\mathbf{x})$ and

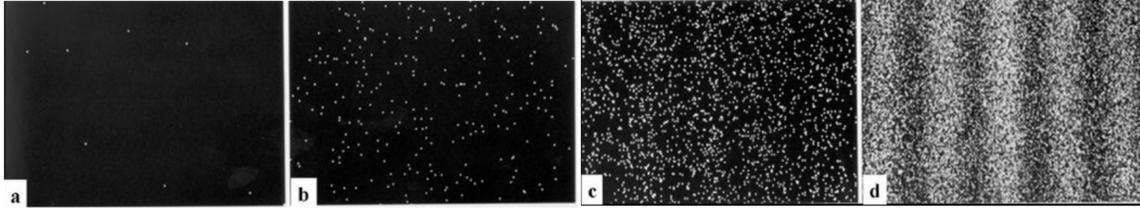


Figure 5. The double slit experiment, performed by Hitachi. The results show the build up the interference pattern from 8 electrons, to 270 electrons, to 2000 electrons and, finally, to 160,000 electrons where the interference fringes are clearly visible.

the probability distribution that we see on the screen is

$$P_1(\mathbf{x}) = |\psi_1(\mathbf{x})|^2$$

Similarly, with the second slit open and the first closed, we have $P_2(\mathbf{x}) = |\psi_2(\mathbf{x})|^2$. But when both slits are open, the principle of superposition tells us that it is the wavefunctions that add, not the probabilities. This means that we have $\psi_{12}(\mathbf{x}) = \psi_1(\mathbf{x}) + \psi_2(\mathbf{x})$ and so the new probability distribution is

$$P_{12}(\mathbf{x}) = |\psi_1(\mathbf{x}) + \psi_2(\mathbf{x})|^2 = P_1 + P_2 + \psi_1^*(\mathbf{x})\psi_2(\mathbf{x}) + \psi_2^*(\mathbf{x})\psi_1(\mathbf{x})$$

We see that the cross-terms mean that $P_{12} \neq P_1 + P_2$. These cross-terms are responsible for the interference pattern.

The conclusion that each electron necessarily passes through both slits is, to say the least, a little surprising. One could try to find a clever way to evade it, perhaps by setting up some ingenious apparatus near one of the slits that will tell us for sure which slit the particle travelled through. It turns out that any attempt to do this will necessarily disturb the state of the particle. For example, if you know for sure that the particle passed through slit 1 (or, equivalently, that it did not pass through slit 2) then the wavefunction collapses to

$$\psi_{12}(\mathbf{x}) \longrightarrow \psi_1(\mathbf{x})$$

and the interference pattern disappears. There are general principles – notably the Heisenberg uncertainty relation that we will meet in Section 3 – that say that any apparatus that can determine through which slit the particle passed will necessarily also disturb the state of particle enough to destroy the interference fringes.

Quantum mysteries are often phrased as the question of whether an electron is a particle or a wave. The answer that is usually given is: both. A much better answer is: neither. Electrons, like all quantum objects, exist in a regime where our everyday experience is of little use. The idea that such objects can be described in terms of things that we can find on the beach is, to put it mildly, a little immature. A quantum particle is, instead, described by a wavefunction. The right way to proceed is not to search for some classical metaphor that will capture some aspect of the quantum particle, but instead to better understand the rules that govern these wavefunctions.