Title	*020 Further Orbital Angular Momentum, Spectra of L2 and LZ*
Description	Twentieth lecture in Professor James Binney's Quantum Mechanics Lecture series given in Hilary Term 2010
Presenter(s)	James Binney
Recording	http://media.podcasts.ox.ac.uk/physics/quantum_mechanics/audio/quantum-mechanics20-medium-audio.mp3
Keywords	physics, quantum mechanics, mathematics, F342, 1
Part of series	*Quantum Mechanics*

Contributor Okay, so let's get going. So, yesterday we introduced these orbital angular momentum operators, the three of them. The three components of X cross P and divided by hbar to make it dimensionless. And we can make the total angular momentum operator L squared by squaring the individual components and adding. We established these results that the commentator of, for example, L X with Y is I times Z. So, similarly, if you do a commutation with L I with P J – I haven't written this down here – you find I times P K, times epsilon sorry.

So, for example, L X times P commuted with P Y would be I times P Z. And basically what we're learning here is that L commuted with the component of a vector operator gives you the third vector operator in that set.

And from that, you can go on very easily to show that the demonstration is exactly the same as what we already did with the angular momentum operator – the total angular momentum operator – that any of these angular momentum – these orbital angular momentum operators – commutes with scalars like X squared or P squared or X dot P.

Also, L squared commutes with any of these things, any of its components.

So we have a situation so far which is precisely analogous to the commutation relations to J comma – sorry, to the total angular momentum. Total. Or what we have so far called the angular momentum operators. So, for example, J I comma X J commutator is I sum over K epsilon I J K; X K Etc. Etc.

So, we have that J I comma J J is equal to I sum over K epsilon I J K J K which mirrors one that I didn't write up here and should have done. Which is that L I comma LJ commutator is equal to I sum over K epsilon I J K L K. Which is really a generalisation of this rule – I mean, an application of this rule – here to the vector L J being put in here, so L K appears over there. It just says that in fact L is a vector operator and similarly, this is a generalisation.

So, we have this exact analogy and the question obviously arises; are they the same operator? Is L the operator which – are the operators L X L Y and L Z the operators that generate rotations that we inferred had to exist just by thinking in the abstract about rotations at the end of last term.

And so the answer to that question is no. And here is the demonstration of it. Let's calculate J squared comma L I, and all those. Let's work out the commutator between the total angular momentum operator and one of these components. Then this is going to be – this is a commutator of a product – and I need to write this – this is a sum of squares, so I need to make that explicit, so I write that as the sum of over I, say, of J I squared comma – oops! No, no; I won't do because

I've got I busy; J L I. And now, each one of these is just a simple product so we use the rule for taking a commutator of a product so this is equal to the sum over J still of – sorry, that's a J; what we're summing over. I leave one of these standing idly by whilst the L operator commutes with the other one. And then I have the L operator commuting with the other one – the first one – while the other one stands idly by in his position.

Now, we know what the result of this is because this is a vector and we know we established, last term, what the commutator of J with the component of a vector is; it's going to be the third member. So this I can replace by an epsilon J I, say, K. So, this is going to be a sum over J and then there is going to be a sum over K coming up of JJ. There will be an I from the commutator. So, JJ – that's this one here, make sure that's clear – this is going to be epsilon J I K L K. So, that's the fundamental rule we established, that the commutator of this with any vector gave you the third component. And then here we will have the same thing; epsilon J I K times L K times J J in the back.

So, we can clean this up into the I times the sum over both J and K going from over X, Y and Z of epsilon J I K times J J L K plus L K J J.

So, this is the sort of thing which often vanishes because this is anti-symmetric. If you swap those two over, J and K, you get a change of sin. And if this was symmetric, we would get a change – oh, sorry; if this was symmetric in J and K – in other words, if you swap them over and you've got no change of sin; you've just got the same thing back – then we would have zero. But this is not symmetric. If you went to change the order of these – if you swapped the indices – so under a swap, what does this go to? It goes to J K L J plus L J JK, which is not the same as this.

So, this thing is not equal to zero but we know that J squared does commute with all of its components. So that sort of suggests that these are not the same and not the same thing.

So, let's try and understand what these orbital angular momentum operators and really understand their relationship to the things that generate rotations. We need to talk a little bit in abstract about rotation around – sorry; motion around a path. So, if we just have a translation – we've already studied translations last term – through one displacement, A; that's generated by the unitary operator which we called U of A which is E to the minus I A dot P on hbar. Alright? That is the operator which generates out of the state that we would have if our system was shoved along the bench by – or was at the different location – a distance – a vector displacement A away.

Let's consider the following. Let's make a series of displacements; here's A one, here's A two, here's A three and so on, right. We're going to make a path by doing a series of displacements. The we have that U total is going to be the product of this displacement and this displacement, this displacement, this displacement. So it's going to be U A four operating on using the result of U A three, operating on the result of using U A two U A one. Each of these things is an exponential. So, this is E to the minus I A four dot P over hbar; E to the minus I A three dot P over hbar Etc.

And the operators occurring in these exponentials all commute with each other because momentum operators commute with all other momentum operators. They commute with themselves, obviously, and then P X commutes with P Y Etc. So, when we multiply these exponentials together, we have the usual magic of exponential functions; we don't have to worry that those are operators, so these are operators but we don't have to worry about that because everything commutes. So this can be written ad E to the minus I, the sum of these vectors summed over I dotted into P over H bar.

So the operative that generates you with the state you get as a result of all of these displacements is given by this; just a single exponential, it's an operator of exactly the original form where the displacement is simply the sum of them. So, the result of taking it this way and this way and this way and this way, we've just shown is the same as the result of just taking it in a straight line over the sum of the A I.

That's a general argument, and now we say special case. So, in particular, for a closed path – so if the path carries you all the way around, what does that mean? It means the sum of the A I zero – all the vectors add up to nothing – and that implies immediately that U total, because it's E to the nothing, is the identity transformation. Now, that might sound obvious, but you'll see in a minute that that's anything but an obvious result.

So, now let's specialise in our paths. Let our path be made up of the series of – so, let this be X; let this angle in here be delta alpha; let N be the point out of board. Then this displacement vector, this is going to be A equals delta alpha in cross X. Right, that's just ordinary classical geometry. That if I go like this and N is out of the board, then that's the displacement that we generate. So, what does that do? That means that the unitary operator that moves my system from here to here, U delta alpha, is going to be E to the minus I A, which is that; delta alpha N cross X dot P upon hbar. So that's just applying the standard stuff with the displacement vector of this form. This is delta A. Or I, whatever.

But what is that? This scalar triple product can be reordered; sort of standard vector algebra tells us that we can reorder this thing into N dot X cross P. So this is equal by vector algebra to I delta alpha N dot X cross P upon hbar. So that's the operator, okay? And there's the hbar that's here. But this is what we define to be L. So this is equal to E to the minus I delta alpha N dot L. We've now discovered what L does. L is the generator of rotations of movements around circles. So we conclude that L is the generator of translations around circles.

We can also get something interesting that's very important by combining these two things. So for a complete circle; we can multiply these things together. So, U for a circle is going to be the sum of this for all sorts of delta alphas, okay? E to the minus I; delta alpha; N dot L, which is going to be E to the minus I N dot alpha N dot L where this is the sum – sorry; it's going to be the product of these.

Right, if we make a series of transformations, each one by delta alpha, then that product of these exponentials can be written as the exponential of the sum of the arguments. The sum of the arguments, it always has N dot L as the operator, and the sum of the delta alphas we're going to call alpha. So, for a circle, this is going to be two pi. So, we're going to come to the conclusion that E to the minus two pi I N dot L is equal to the identity transformation because we've shown that U going all the way around in any closed path has to be the identity.

So, this thing has to be an identity for any vector, any unit vector N. Now, we'll need that result shortly.

So, what do we – come on! So, what is the general picture here? That we can now understand what the distinction is between orbital angular momentum and angular momentum. So, look at this L and then there's this arrow here. We're moving the point at the base of the arrow around a circle. We're just translating it. Let's just go back. We're just moving it around the circle but we're not doing anything to its internal structure; we're just translating it. So if it has a little arrow – it has an orientation as, for example, the direction of its spin – that's not going to change its direction. It'll just be carried round.

Now ask what J does. What J does is, it makes the system you would have had if you took what you got and you rotated it on a turntable around the origin. If you rotated it on a turntable around the origin, this is what happens; the orientation of the particle changes as well as its location. So, J is changing locations, L – sorry L is only changing locations; J is doing a complete job by putting the particle on a turntable and rotating it at the same time as translating it. That's the difference.

Right, so we've got these operators and you want to know - obviously when you have operators, you need to know what their spectra are, what the allowed values of their eigen values are and we can immediately say what these are going to be. So we know that we can find a complete set of simultaneous eigen states of L squared of any one of its components – L Z; that's the same as when we were dealing with the angular momentum operators. And we obtained the eigen values of the angular momentum operators at J squared and J Z by using – by exploiting – the commutation relationships between the different components of J. That was the only thing we used. Right?

So, we got that the E values of J squared J Z were… They were J plus one – this is for J squared. J squared had eigen values of J J plus one where J was nought, a half, one, three half – curses! – three halves Etc. Right? And we had that J Z had eigen values M which lay between minus J and plus J. And we got these results only using that J I comma J J equals I epsilon I J K J K. If you look back at what we did, you'll find that that's the case, that nothing else went into this with these communtation relationships here.

For the N operators, we have the same commutation relations, right? So the L satisfy identical commutation relations. So the argument we had for J could be repeated line by line. There's no point in repeating it literally, but virtually we now repeat that argument line by line with every J replaced by an L. And we conclude that the E values can be. L squared has numbers like L plus one where L could equal nought, a half, one blah blah blah. And L Z has N lying between L and minus L.

But then we think, 'but...'; so these are the candidate – no other numbers are possible than those numbers. Alright? That's what that argument shows. Are all those numbers possible? No, for the following reason. But; E to the two pi I L Z has to be the identity operator. We've just shown that. That's this statement with N put equal to the unit vector in the Z direction. Actually, I need a minus sin. So we have that this thing has to be the identity operator. Consequently, if we use this operator on one of our states, which are going to be called L M, alright? This is the mutual eigen state of L squared and L Z - with eigen value L L plus one for L squared and M for L Z precisely by analogy with what we did with J.

Then this has to be simply L M. Because this is the identity operator. But what actually is this? So, we're doing an exponential of this operator. This operator looks at this and says, "That's my eigen function," and therefore it replaces itself with the eigen value, so this is E to the minus two M pi I times L M. Comparing this side with side we have E to the two pi I M is equal to the number one. And that implies that M is an integer. Because if it were a half integer, this would be E to a certain number of I pis which would be minus one. But we've show it cannot be minus one so it's an integer.

So, this is where there is a difference between what happens – between what L does and J does and why we have to keep track. So L is looking very like J, but it is not J because it is – well, we've seen physically that it's not because it merely translates you around a circle; it doesn't rotate you round a circle. And that has the consequence that M has to be an integer and therefore L has to be an integer.

So the eigen values of L squared are L plus one where L equals nought one, two, three, four, Etc. and therefore M is also going to be stuck on integers.

So, what we learn from this in some sense is that – let's just do this again. So, when we have translated our system all the way around a circle, it's come back to where it was, same orientation, everything the same. Right? And the result is that an identity operator is applied that's described as applying an identity operator to our state.

If on the other hand we translate J all the way round the system, you would think it came back to where – well, it comes back to the same place undoubtedly, and the little arrow comes back to the same orientation. Now, you would think you were back in the same place. But quantum mechanics is telling us the mathematical – well, it's really experimental physics that tells us – that it's not. I'm having to described the orientation of the spin of a particle using an arrow and I've already said this is a hazardous enterprise. Because in quantum mechanics, particles of giros and they is some sense have a spin, a point in a direction, but you do get into trouble as the Einstein Podolsky Rosen experiment shows, if you take this idea of pointing in a direction too seriously. And here we have another example of the risks of taking too seriously the idea that you can describe the spin of a particle by pointing in a direction.

Because when it has gone all the way around, the state vector has changed sin in the event that this is a half integer. And the Stern Gerlach experiment and other experiments – zillions of experiments – show that electrons and protons and so on do have half integer values. They're angular momentum comes in half integer amounts.

So when you take a real electron on one of these tours around the origin, its state is somehow different from the state it started from. It's difficult for us to understand that but that's what the mathematics combined with the experimental physics tells us.

Okay. The next item on the agenda is the eigen functions of L squared and L Z. So we already know what the eigen values are but what we now would like know is what do these states look like – these L M states look like – in the positional representation. So what we want is R theta phi L M. We'd like to find expressions for – we'd like wave functions which describe these things here.

And our strategy is basically that we're going to – so, the strategy is this - the detail is rather tedious - but our strategy is this; we're going to apply L plus to L L and get nothing. We're going to express this abstract – so this is L X plus – alright; I think we talked about this yesterday but I'm not absolutely certain - L X plus L Y. So this is the operator which will try and raise the M entry here to one larger, but that's not possible because it's already at its largest value, so it'll kill it.

So this is an operator equation. We will look at this operator equation in the position representation and they will become a first-order differential equation. We will solve it. It'll turn out to be dead simple to solve. So that will lead us to – we'll leave off the R because we're not really interested in R for the moment – it will lead us to this; L L. And then we will be able to say that theta phi L; L minus one is equal to theta phi – sorry, there's a horrible square root now – L; L plus one minus M; M minus one times L minus one – sorry, times L minus. So, we will having obtained this – this wave function – we will apply L minus to it in the positional representation, that will give us essentially the next one down. That'll lower this by one and by repeatedly doing this we'll be able to generate all of the wave functions associated with a particular total angular momentum quantum number L. That's the strategy.

To carry this out, what we need is expressions. We need expressions in the positional representation for these operators. Let's start by writing down what L Z is. What is L Z? It's one over hbar of X P Y minus Y P X in the positional representation. What is that? P Y in the positional representation is minus hbar D by D Y, so this becomes minus I X D by D Y minus Y D by D X. Now, to find out what that is, we would like to express everything in terms of theta and phi because we know angular momentum is to do with rotations. That's why we want theta and phi spherical polar co-ordinates.

And we could just use the chain rule to turn this into derivatives of theta or phi, but it's much easier to go – to use the chain rule; not going from this to D by D theta, D by D phi, but to go to find out what is D by D phi according to the chain rule? It's DX by D phi; D by D X; plus D Y by D phi; D by D Y; plus D Z by D phi; D by D Z. Alright? So, that's just the chain rule from calculus. Now we put in what these X, Y and Z are; we know that X in polar co-ordinates is R sin theta cos phi. We know that Y is R sin theta sin phi; and we know that Z is R cos theta, so this implies – so, take a derivative of this. With respect to phi, we have that the X by D phi is going to be minus R sin theta sin phi, which is the same as minus Y – minus is simply going to be minus Y. We have that. Similarly D Y by D phi; that becomes a cos sin and therefore this becomes X; and we have the Z by D phi is nothing at all.

So we take these results and stuff them back in here, and that tells me that D by D phi is equal to – that's a Y – sorry, that's a minus Y. So, let's write it down; minus Y; D by D X plus X; D by D

Y. What is the relationship to what we have up there? That is – okay. So this is precisely what's in that bracket for L Z, so this implies that L Z is minus I; D by D phi.

Of course, this is no surprise because L Z is the thing that rotates you around the Z axis, and everybody knows that the spherical polar co-ordinate is defined as the angle around the Z axis. So it's kind of obvious because it's got to be the case, but it's nice to see that the chain rule delivers the goods.

Now, the next bit is distinctly more tiresome. What we now have to do is – what we want to do is, as I said, express L plus and L minus – L X plus… you know, plus I LY – in terms of D B D thetas and D B D phis. You could go at this just by brute force but the algebra would be heavy, so – and the algebra's not going to be that wonderful now – here is, I think – this is the way to do it.

Right. So let's just calculate what D B D theta is using the chain rule. It's, obviously, D X by D theta; D by D X plus D Y by D theta; D B D Y plus D Z by D theta; D by D Z. Okay, we have expressions up there so D X by D theta is going to produce us an R cos(theta), cos phi Etc. This is going to produce us an R cos theta sin phi. So I've got a common factor of R cos theta, open a bracket and then we will be able to write down sin – sorry; cos phi D by D X plus cos phi – sorry; plus sin phi; D B D Y. And then finally we have D Z by D thingy which is going to be minus R because we're differentiating cos sin theta D B D Z.

Now let's take our expression for D by D phi and multiply it by cot(theta). So, we're going to write down cot(theta); D by D phi, just for the fun of it. So, in doing it, we're going to replace – well, let's take this one first. It's going to be R sin theta cos phi. When we multiply it by cot, which is cos over sin, the sin is going to go into a cos. So, this is going to be R cos theta brackets sin phi – excuse me, cos phi – D by D Y. So, this is here is R sin theta cos phi, which is X times cos over sin, which is cot.

And then that Y is going to be R sin theta sin phi, and this cot multiplication will turn the sin theta into a cos theta; I take it out and we will have a sin phi; here, D by D X. That's it. Right.

So, what we now do is – just again, for the fun of it – we take D B D theta and add on I times this. So, we look at D by D theta plus I times cot theta D by D phi. What happens then? We have R cos theta as common factors… Oops! Quick. Excuse me, yes that's right; that's fine. I'm looking at the wrong coefficients.

So what am I going to be doing? I'm going to be adding this and this. So we're going to add I times this to that, with a common factor of D B D X. We're going to have cos phi minus I sin phi, which is well known to be written as E to the minus I phi. That's how much of D B X we're going to have. And then we're going to have sin phi plus I cos phi and that is going to be plus I times E to the minus I phi D by D Y. We haven't quite finished, have we?

Right, so if you expand this out, it's going to give you an I cos phi, which we want – that's this one here times I – and it will contain a minus I sin phi times and I, which makes it a plus just plain sin phi, which is what we want there. Then finally we've got this in the back so we're – from here – so we have minus R sin theta D by D Z. Yes.

And this thing could be written R sin theta is this – no. No, no, no. Leave it alone.

So, what we do now is we multiply this by E to the I phi. So, we have E to the I phi brackets D by D theta, plus I cot(theta); D by D phi, which will turn out to be our bottom line; is equal to – this is going to be cancelled, because I'm multiplying through to E to the I phi. This will be cancelled. So, we will have R cos theta brackets; D by D X plus I; D by D Y. And then we will have minus R sin theta times E to the I phi, which I choose now to write as cos(phi) plus I; sin phi.

And now R sin theta cos(phi) is X and this, in the back, is going to be Y – R sin theta sin phi – so we'll be able to write the right hand side as R cos(theta) D by D X plus I; D by D Y minus – so this is going to be X plus I Y. Excuse me; this was times D B D Z; it always was. Where did this get lost? It was minus R sin theta D by D Z. Yes; it just got lost between this line and this line. Okay, Phew!

So, I now want to establish that that is actually L plus. So, in order to do that I write down – so what is L plus? L plus is L X plus I L Y. Don't need those brackets. Which is one upon hbar.

Okay. What's P X? It's Y P Z minus Z P Y plus I times L Y is Z P X minus X P Z. So, now I want to turn this into D B D Xs. So, this is minus I hbar D by D Z, so we're going to have – no, sorry; before we do that, let's gather things together, because we've got a P Z and a P Z. So let's just for the moment – right; this is one over hbar. We're going to have a common factor of Z. We're going to have I Z P X plus I P Y, I think. Right, this I and that I make the minus sin that we have there. Z is the common factor alright?

And then, what about the Z factors – the P Z factors? Well, we have a plus Y minus I X P Z. Now, we replace the Ps by minus hbar D by D X Etc. So, the hbars are going to go throughout. From here, we'll have a minus I times and I, which will give us a plus, so we'll have a Z; D by D X. Here we will have – we have two Is making a minus sin. We have another minus sin coming in from the minus I, so we end up with minus I; D by D Y. I'm rather worried about the sin there. `[[?? 0:44:00]]`. No, the two made a minus and I soaked up, sorry. These two Is made a minus; I brought in another minus from minus I hbar D by D E; the minus is cancelled, leaving me only with the I; and this is going to bring in a minus I.

If you propagate this minus I and sin here we get a minus sin there. So we get a minus X and this is going to plus I Y, because there's a minus I here and there's the minus sin and here's the I. D by D Z. And that, I hope, agrees with what we have written; it does because R cos(theta) is also known as Z. So we have Z D B X plus I D B D Y, and here we have a minus X plus I Y D B D Z.

So, we have established a very important result that L plus – in the positional representation – is the differential operator E to the I phi brackets; D B D theta – it's plus, isn't it? Plus I cotangent theta; D by D phi; close brackets.

There is an analogous calculation - which I'm not jolly well going to do - which tells us that L minus is equal minus – E to the minus I phi; D by D theta – is it minus? Yes; it is – minus I cot(theta); D by D phi. And if you try and do the two calculations together, with plus and minus sins, good luck to you; it's fantastically confusing. Right (Laughter). At least I got one of them right!

So, those will be our – well let's... No. No, let's leave it at that. Those are going to be our starting points for tomorrow deriving the eigen functions of these very important operators.